

AI AND LANGUAGE IN THE URBAN CONTEXT

CONVERSATIONAL ARTIFICIAL
INTELLIGENCE IN CITIES

RICHARD COYNE

ROUTLEDGE

AI AND LANGUAGE IN THE URBAN CONTEXT

In a world influenced increasingly by artificial intelligence (AI), the city emerges as a dynamic hub of digital conversations. *AI and Language in the Urban Context* offers a novel exploration of how AI, particularly large language models (LLMs), is transforming urban environments. Moving beyond the typical technological narratives, this book draws on the author's unique expertise in design, semiotics and hermeneutics to present a critical cultural perspective on AI's role in the city.

Focusing on the intersection of urban theory and AI, the book reveals how conversational AI is reshaping social interactions, decision-making processes, and media in urban spaces. By merging practical knowledge of AI algorithms with an understanding of urban practices, the author highlights the opportunities and challenges AI presents for modern cities.

This book is essential for anyone interested in the future of urban living. It provides a deep dive into the technical, social and cultural implications of AI in cities, offering practical examples and philosophical insights. Readers will gain a comprehensive understanding of how AI is influencing the design, governance and dynamics of urban life in the digital age.

Richard Coyne is Emeritus Professor of Architectural Computing at the University of Edinburgh, where his research explores the intersections of digital technology, urban environments and architecture.



Taylor & Francis

Taylor & Francis Group

<http://taylorandfrancis.com>

AI AND LANGUAGE IN THE URBAN CONTEXT

Conversational Artificial
Intelligence in Cities

Richard Coyne

Designed cover image: Shutterstock.

First published 2025
by Routledge
605 Third Avenue, New York, NY 10158

and by Routledge
4 Park Square, Milton Park, Abingdon, Oxon, OX14 4RN

Routledge is an imprint of the Taylor & Francis Group, an informa business

© 2025 Richard Coyne

The right of Richard Coyne to be identified as author of this work has been asserted in accordance with sections 77 and 78 of the Copyright, Designs and Patents Act 1988.

The Open Access version of this book, available at www.taylorfrancis.com, has been made available under a Creative Commons Attribution-Non Commercial-No Derivatives 4.0 license.

Trademark notice: Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

ISBN: 978-1-032-87992-5 (hbk)
ISBN: 978-1-032-87990-1 (pbk)
ISBN: 978-1-003-53575-1 (ebk)

DOI: 10.4324/9781003535751

Typeset in Joanna
by codeMantra

To Kathy, David, Emma and Jenny



Taylor & Francis

Taylor & Francis Group

<http://taylorandfrancis.com>

CONTENTS

Acknowledgements	ix
About the author	x
Introduction: From computer-aided design to conversational AI	1
1 Language as urban technology	10
2 Scripting the city	43
3 Core functions of LLMs in the urban context	57
4 Tuning the city	73
5 Tokenization	89
6 The meaning and use of words	100
7 Context windows	113
8 Time and order	133

9	The attention economy	159
10	Urban actors	178
11	Panpsychic city	194
12	The end of cities	213
13	Epilogue: Romancing the conversational city	233
	Glossary	241
	Appendix A: Programming an LLM	250
	Appendix B: Positional encoding formulas	256
	Appendix C: An alternative glossary of AI	258
	Index	261

ACKNOWLEDGEMENTS

I am grateful to the Edinburgh School of Architecture and Landscape Architecture in Edinburgh College of Art, the University of Edinburgh for offering the sabbatical period and resources that provided space for me to develop the content of this book. I'm grateful to John Lee for insisting on the potential of language-based AI, even before the current wave of high-profile developments in large language models, and Aart Bijl for paving the way for AI in architecture at Edinburgh. I'm grateful for the insights of colleagues and critics engaged in AI dating back to my earliest days as a researcher in the University of Sydney: Bala Balachandran, John Gero, Sally McLaughlin, Sidney Newton, Arthur Postmus, Tony Radford, Michael Rosenman, Adrian Snodgrass, Fay Sudweeks and M. Yokozawa.

ABOUT THE AUTHOR

Richard Coyne is Emeritus Professor of Architectural Computing at the University of Edinburgh, where his research explores the intersections of digital technology, urban environments and architecture. With a focus on how network technologies reshape cities, Coyne's work is celebrated for its interdisciplinary reach and critical insights into the digital age's impact on urban life.

Coyne's research has spanned a range of innovative topics. In *The Tuning of Place: Sociable Spaces and Pervasive Digital Media* (MIT Press, 2010), he examined the role of sound in shaping urban environments, while *Mood and Mobility: Navigating the Emotional Spaces of Digital Social Networks* (MIT Press, 2016) explored the emotionally responsive dimensions of urban spaces. His pioneering work with mobile electroencephalography (EEG) has mapped the stress levels of older citizens navigating busy streets and parklands, highlighting the human experience of the city.

His interest in the interplay between nature and urbanisation led to *Network Nature: The Place of Nature in the Digital Age* (Bloomsbury, 2018), and his latest book, *Cryptographic City: Decoding the Smart Metropolis* (MIT Press, 2023), explores the urban metaphors spawned by blockchain technologies and cryptocurrencies. Coyne's work draws on a broad array of disciplines, blending practical design with deep theoretical inquiry, particularly in urban studies, Phenomenology and computational methods.

Throughout his distinguished career, Coyne has served as Head of the Department of Architecture and the School of Arts, Culture and Environment at the University of Edinburgh. He has been instrumental in shaping academic programs, including the MSc in Design and Digital Media and the MSc in Sound Design. His critical approach to digital technologies, coupled with his commitment to interdisciplinary research, continues to influence both academic and professional fields.

For ongoing insights into Coyne's work, including links to his institutional websites, publications, and further explorations of AI and urban environments visit his blog at: www.techmediaculture.com.



Taylor & Francis

Taylor & Francis Group

<http://taylorandfrancis.com>

INTRODUCTION

FROM COMPUTER-AIDED DESIGN TO CONVERSATIONAL AI

Cities are centres of conversation, exchange and transformation. A city is its physical structures but also evolving social interactions in language. Conversational AI introduces new layers of interaction, perception and risk to cities. The technologies of large language models (LLMs) mirror, magnify and disrupt urban conversations.

The role of artificial intelligence (AI) in cities has increased significantly since the advent of ubiquitous digital networks. Several decades of concerted research and development produced AI-based large language models (LLMs) that converse with human users in ways that appear convincingly creative and knowledgeable. Such digital platforms also summarise texts, evaluate documents, correct grammar, suggest improvements, compose new texts, answer questions and translate texts into different languages. It is this linguistic, text-based aspect of AI that I want to focus on in this book.

I present language as providing the ligament that connects AI with cities. After all, large language models (LLMs) provide a high-profile demonstration that machines can replicate something akin to human conversation.

DOI: 10.4324/9781003535751-1

This chapter has been made available under a CC-BY-NC-ND 4.0 license.

In the earliest days of computers, Alan Turing used conversational ability as the benchmark for intelligence with the so-called Turing Test for intelligence.¹ If an automated text machine can respond in a conversational way to your inputs and can do so in a manner that is indistinguishable from a conversation with a human being, then the automated system passes the test. To all intents and purposes, it is *intelligent*. Irrespective of the adequacy of this test, it highlights the priority that AI gives to language, a priority that it shares with cities.

In Chapter 1, I will demonstrate that language is key to understanding cities. Cities are negotiated, described, interpreted, analysed and even created through language, i.e. via communication, conversation, dialogue, narrative and text. There's an urban *semiotics* in play. Hence, my insistence that language (its theories of meaning, pragmatics, hermeneutics) provides one of the strongest channels for examining the complex relationships between cities and automated intelligence, and claims for urban AI (UAI).

Not everyone concurs with this linguistic emphasis. Some scholars prefer *cybernetics* rather than *language* as the theoretical ligament that connects AI to cities. Urban theorist Federico Cugurullo and colleagues situate urban AI historically within the development of scientific methods applied to cities, in particular developments in cybernetics:

Thus, today UAIs are machines capable of interacting with their environment and regulating their behaviour in relation to pre-defined goals, in line with the ideals theorized by cyberneticians in the past century. In fact, some scholars (urbanists in particular) have noted a direct connection between the application of AIs in cities and the history of cybernetics.²

Rather than AI-based tools for processing text, this cybernetic avenue of inquiry is populated by the terminology of sensors, data, controls and feedback loops, applicable to autonomous vehicles, robots, transportation, security, monitoring and decision support systems. The idea of the “smart city” leads the AI cybernetic discourse.³ The operative terms of cybernetics are inputs and outputs, control, signals, feedback, entropy and stability. That contrasts with the semiotic, linguistic orientation that emphasises meaning, understanding, interpretation and narrative.

Semiotics and cybernetics have interacted with one another in various ways. Winfried Nöth considers both in his extensive survey: *The Handbook*

of Semiotics.⁴ Cybernetics would seek to subsume language and communication within its own explanatory terms of inputs and outputs, control, signals, feedback, entropy and stability. I included “feedback” in this list of cybernetic properties. One of the theoretical innovations claimed by cybernetics was the incorporation of loops into its models. An organism or machine receives inputs from its environment, and then acts in response, which in turn changes the environment, and prompts further responses. Other insights develop from this idea of close coupling between organism (or machine) and its environment. Nöth recognises the value of this basic insight: “In the course of the development from linear to circular models, communication theory has overcome the misleading idea of communication as an almost undisturbed flow of signals from a source to a destination.”⁵

But semioticians find reasons to be critical of a cybernetic approach to language: “the improvements introduced by cybernetics and systems theory have attracted new criticism,”⁶ according to Nöth. At issue is the entity that receives the signal flows, the *listener*. Nöth summarises the objections to a cybernetic framing of language, suggesting that it diminishes the role of the person or entity receiving the signal. Cybernetics implies “a process of optimization, efficiency, and the goal of congruence between sender and receiver.”⁷ His objection pertains to the lack of a theory about the receiver of a communication in cybernetic theory. The role of the receiver is important in cities if we are to think of cities as products of co-creation involving communities, teams and materials. For the semiotician, the receiver is an interpreter with their own framing that in turn influences the message, and the context of the communication. This accords with the character of the urban context and the rich and varied roles of a wide constituency of participants, including citizens, legislators, residents, visitors and designers.

Cybernetics in turn has its own framing of urban challenges, e.g. how to deal with multiple cybernetic agents, each with their own interactions with the environment. For Cugurullo: “The search for a radical, open cybernetics that is attentive to the dynamic connection between different intelligences instead of their forced harmonisation is still open.”⁸ That’s a pointer to the restricted applicability of cybernetic theory that begins with the idea of a singular autonomous agent.⁹

Are large language models (LLMs) cybernetic systems? Certainly, the cybernetician could describe artificial neural networks that underpin

LLMs in terms of inputs and outputs, control, signals, feedback, entropy and stability. Neural networks are arrays of numbers under the control of algorithms that make fine adjustments to those numbers in response to inputs. These networks are “trained” to replicate accurately very large sets of input-output patterns, exhibiting behaviours at the computational level that negotiate different states of entropy. Cybernetics lays claim to any operation or process at some level of operation. On the other hand, the semiotician will counter with the assertion that it is language that underlies any process. Linguistic interpretation is everywhere and in everything, including in the design and application of scientific procedures, the practices of neural network developers and the narratives of cybernetics.

Rather than expand on the case for semiotics here, I am content for the time being to see cybernetics and semiotics as providing two different framings, each with their own set of problematics. Each puts a different focus on the city, its challenges, opportunities and how AI intervenes. For my purposes, cybernetics has had its time in the sun. Now it is time for language to emerge from the urban shadows. Language in its many facets provides the ligament joining AI to the city. This “linguistic turn” pays off by elevating the role of narrative, meaning, poetics, imagination and invention in the urban conversation.

From computer-aided design to conversational AI

Some scholars of both cybernetics and semiotics seek to address “the gulf that divides the natural sciences from the humanities.”¹⁰ The cybernetic-semiotic distinction helps me at least identify major differences in thinking about cities, and perhaps the lack of comprehension of one for the other. In part it is a distinction between a calculative approach to cities, and an approach more oriented to culture and the humanities. The distinction in part reflects my own journey through the fields of computers, digital media and AI.

My early research, and that of colleagues, focussed on tools for design. We were steeped in the possibilities offered by computer-aided design (CAD). In the 1980s CAD was starting to take hold in architectural offices and design schools. It became apparent soon after the invention of screen-based interfaces that you could draw on a computer screen, create 3D models and print out plans, elevations, details and perspective renders. But even in CAD’s

earliest conceptions, scholars and practitioners were thinking in terms of design-enhancing automation. Nicholas Negroponte's *The Architecture Machine* published in 1970 tantalised practitioners with a prototype drawing system that offered text-based question and answer dialogue that would lead you through the process of designing a house.¹¹

In 1998, architectural researcher Bob Martens at the Vienna University of Technology and structural researcher Ziga Turk at the University of Ljubljana set up a cumulative index of publications on computer-aided architectural design, either relevant to design researchers or produced by them.¹² At the time of writing, there are over 900 publications (out of over 12,000) listed on the theme of AI in architecture, the earliest records dating to the 1970s. I mention this to indicate that parallel to high profile researchers developing the techniques of large language models, neural networks, robotics and artificial intelligence we see a well-documented record of researchers and practitioners in architecture, design, engineering and others engaged in the built environment applying the tools and techniques of AI to practical contexts, thereby broadening the investment in AI beyond computer science and cognitive science.¹³

For our part, in 1990, colleagues and I collated our own AI-based research as a book *Knowledge-Based Design Systems* published by Addison Wesley.¹⁴ In that we applied rule-based and logic models of AI to some of the technical challenges presented by design in the built environment, and from cybernetic and systems theoretic points of view. Concurrently we were dealing with the challenge at the time posed by neural network theory. Our research was mainly on the side of so-called symbolic AI, as opposed to the theories of connectionism, i.e. neural networks of the kind that eventually bloomed into the highly successful large language models now emblematic of AI. Some of us did however experiment with associative neural networks in the 1990s.¹⁵

The early 1990s also saw the influence of Terry Winograd's book (with Fernando Flores) *Understanding Computers and Cognition* take hold in certain sectors of the AI in design community.¹⁶ In his earlier work, Winograd had implemented the idea of instructing a robotic device to move blocks around a space following logical spatial rules, and in response to instructions delivered in something like natural language. But *Understanding Computers and Cognition* departed from a mechanistic, computational approach to language and introduced concepts from philosophy that helped explain the difficulties, if not the impossibility, of such a project, as well as alternative avenues

for research. This “postmodern turn” introduced CAD and interaction designers to philosophical Phenomenologists such as Martin Heidegger, Hans-Georg Gadamer and Maurice Merleau-Ponty. These scholars drew attention to the physical, embodied nature of thought and intelligence, and the roles of language in establishing our “being in the world.”¹⁷

For my part, I also fell under the thrall of the renaissance in architectural theory, with its newly-found, broad-ranging, and distinctly humanities-based orientation.¹⁸ My colleague and Asian scholar Adrian Snodgrass was influential here,¹⁹ along with the growing realisation that a language-based approach provided a stronger spark to ignite the enthusiasm of architecture students than a purely calculative approach, and proved to be most relevant in design studio teaching. In 1995, I published a book *Designing Information Technology in the Postmodern Age* with MIT Press,²⁰ and Snodgrass and I collated our understandings from hermeneutics, philosophy and design in a book *Interpretation in Architecture* with Routledge.²¹

Since the 1990s, the relevance of computing has burgeoned as driver for so much innovation in architecture and design teaching and practice, with architects, urbanists, planners and designers contributing to critique and development. They provide a testing ground and context for AI theories and ideas. For my part, since the 1990s I was involved in establishing a masters program, supervising large numbers of PhD projects and conducting research under the rubric of *design and digital media*, and with a renewed and critical interest in semiotics, language and artificial intelligence. It has become apparent that language-based AI developments are also having a significant effect on learning and teaching practices, as students, teachers and administrators explore the potential of these new tools.²² As indicated in the RIBA *AI Report* published in 2024, architecture practices are similarly animated by the possibilities of AI.²³ The expansion of such insights in the urban context is the topic of this book.

Notes

- 1 Allan M. Turing, ‘Computing Machinery and Intelligence’, *Mind: A Quarterly Review of Psychology and Philosophy*, 59, no. 236 (1950): 433–60.
- 2 Otello Palmimi and Federico Cugurullo, ‘Charting AI Urbanism: Conceptual Sources and Spatial Implications of Urban Artificial Intelligence’, *Discover Artificial Intelligence*, 3, no. 15 (2023), 5. Also see Federico Cugurullo, *Frankenstein Urbanism: Eco, Smart and Autonomous Cities, Artificial Intelligence and the End of the City* (London: Routledge,

- 2021) and Federico Cugurullo, et al., *Artificial Intelligence and the City: Urbanistic Perspectives on AI* (London: Routledge, 2024).
- 3 Simon Joss, et al., 'The Smart City as Global Discourse: Storylines and Critical Junctures across 27 Cities', *Journal of Urban Technology*, 26, no. 1 (2019): 3–34; Germaine Halegoua, *Smart Cities* (Cambridge, MA: MIT Press, 2020); Katharine S. Willis and Alessandro Aurigi, *Digital and Smart Cities* (Abingdon, Oxon: Routledge, 2018); Antoine Picon, *Smart Cities: A Spatialised Intelligence* (London: Wiley, 2015).
 - 4 Winfried Nöth, *Handbook of Semiotics* (Bloomington, IN: Indiana University Press, 1990).
 - 5 Winfried Nöth, *Handbook of Semiotics* (Bloomington, IN: Indiana University Press, 1990), 178.
 - 6 Winfried Nöth, *Handbook of Semiotics* (Bloomington, IN: Indiana University Press, 1990), 178.
 - 7 Winfried Nöth, *Handbook of Semiotics* (Bloomington, IN: Indiana University Press, 1990), 178.
 - 8 Otello Palmimi and Federico Cugurullo, 'Charting AI Urbanism: Conceptual Sources and Spatial Implications of Urban Artificial Intelligence', *Discover Artificial Intelligence*, 3, no. 15 (2023), 14.
 - 9 For illustrations of how cybernetics develops into systems thinking, and more open and human-oriented models see Ray Ison and Ed Straw, *The Hidden Power of Systems Thinking: Governance in a Climate Emergency* (Abingdon, Oxon: Routledge, 2020).
 - 10 Winfried Nöth. 'System, Sign, Information, and Communication in Cybersemiotics, Systems Theory, and Peirce', in *Introduction to Cybersemiotics: A Transdisciplinary Perspective*, eds C. Vidales and S. Brier (Cham, Switzerland: Springer, 2021), 75–94, 80. These authors also draw on C. P. Snow, *The Two Cultures and the Scientific Revolution* (Oxford: Oxford University Press, 1959).
 - 11 Nicholas Negroponte, *The Architecture Machine: Towards a More Human Environment* (Cambridge, MA: MIT Press, 1970).
 - 12 See <https://papers.cumincad.org>. Recent articles relevant to LLMs and design include Lok Hang Cheung, et al. 'Exploring Large Language Model as a Design Partner through Verbal and Non-verbal Conversation in Architectural Design Process', in *Proc. SIGraDi* (Punta del Este, Uruguay: Centro Universitario de la Región Este, 2023), 1059–70; Sandro Martinez Conceição, et al. 'ChatGPT for briefing creation', in *Proc. SIGraDi* (Punta del Este, Uruguay: Centro Universitario de la Región Este, 2023), 819–30 and Linus Tan and Luke Thom. 'Accelerating Future Scenario Development For Concept Design With Text-Based GenAI (ChatGPT)', in *Proceedings of the 29th International Conference of the Association for Computer-Aided Architectural Design Research in Asia (CAADRIA) 2024 Volume 1: CAADRIA*, 2024).
 - 13 Similar developments also took place in urban planning. For some recent developments see Dongjie Wang, et al., *Towards Automated Urban Planning: When Generative and ChatGPT-like AI Meets Urban Planning* (working paper) (Ithaca, NY: arXiv.org, 2023).
 - 14 Richard Coyne, et al., *Knowledge-Based Design Systems* (Reading, MA: Addison-Wesley, 1990).
 - 15 Richard Coyne and M Yokozawa, 'Computer Assistance in Designing from Precedent', *Environment and Planning B: Planning and Design*, 19, no. 2 (1992): 143–71. Sidney Newton and Richard Coyne, 'Impact of Connectionist Systems on Design', *Knowledge-Based Systems*, 5, no. 1 (1992): 66–81.

- 16 Terry Winograd and Fernando Flores, *Understanding Computers and Cognition: A New Foundation for Design* (Reading, MA: Addison Wesley, 1986).
- 17 A book by Hubert Dreyfus positioned this work within the problematics of intelligence: Hubert L. Dreyfus, *Being-in-the-World: A Commentary on Heidegger's Being and Time Division I* (Cambridge, MA: MIT Press, 1990).
- 18 For the flavour of this renewal see Bernard Tschumi, *Architecture and Disjunction* (Cambridge, MA: MIT Press, 1994).
- 19 Adrian Snodgrass, 'Asian Studies and the Fusion of Horizons', *Asian Studies Review*, 15, no. 3 (1992): 81–95.
- 20 Richard Coyne, *Designing Information Technology in the Postmodern Age: From Method to Metaphor* (Cambridge, MA: MIT Press, 1995).
- 21 Adrian Snodgrass and Richard Coyne, *Interpretation in Architecture: Design as a Way of Thinking* (London: Routledge, 2006).
- 22 J. S. Britten, *It's a New Day (Sort of): Large Language Models/Education/Pedagogy* (New York: J4 Educational Press, 2023).
- 23 RIBA, *RIBA AI Report* (London: Royal Institute of British Architects, 2024).

References

- Britten, J. S. (2023), *It's a New Day (Sort of): Large Language Models / Education / Pedagogy*, New York: J4 Educational Press.
- Cheung, L. H., J. C. Dall'Asta and G. Di Marco. (2023), 'Exploring Large Language Model as a Design Partner through Verbal and Non-Verbal Conversation in Architectural Design Process', *Proc. SIGraDi*: 1059–70, Punta del Este, Uruguay: Centro Universitario de la Región Este.
- Conceição, S. M., N. C. Diehl and L. M. Bruscatto. (2023), 'ChatGPT for Briefing Creation', *Proc. SIGraDi*: 819–30, Punta del Este, Uruguay: Centro Universitario de la Región Este.
- Coyne, R. (1995), *Designing Information Technology in the Postmodern Age: From Method to Metaphor*, Cambridge, MA: MIT Press.
- Coyne, R., et al. (1990), *Knowledge-Based Design Systems*, Reading, MA: Addison-Wesley.
- Coyne, R. and M. Yokozawa. (1992), 'Computer Assistance in Designing from Precedent', *Environment and Planning B: Planning and Design*, 19 (2): 143–71.
- Cugurullo, F. (2021), *Frankenstein Urbanism: Eco, Smart and Autonomous Cities, Artificial Intelligence and the End of the City*, London: Routledge.
- Cugurullo, F., et al. (2024), *Artificial Intelligence and the City: Urbanistic Perspectives on AI*, London: Routledge.
- Dreyfus, H. L. (1990), *Being-in-the-World: A Commentary on Heidegger's Being and Time Division I*, Cambridge, MA: MIT Press.
- Halegoua, G. (2020), *Smart Cities*, Cambridge, MA: MIT Press.
- Ison, R. and E. Straw. (2020), *The Hidden Power of Systems Thinking: Governance in a Climate Emergency*, Abingdon, Oxon: Routledge.
- Joss, S., et al. (2019), 'The Smart City as Global Discourse: Storylines and Critical Junctures across 27 Cities', *Journal of Urban Technology*, 26 (1): 3–34.
- Negroponte, N. (1970), *The Architecture Machine: Towards a More Human Environment*, Cambridge, MA: MIT Press.

- Newton, S. and R. Coyne. (1992), 'Impact of Connectionist Systems on Design', *Knowledge-Based Systems*, 5 (1): 66–81.
- Nöth, W. (1990), *Handbook of Semiotics*, Bloomington, IN: Indiana University Press.
- Nöth, W. (2021), 'System, Sign, Information, and Communication in Cybersemiotics, Systems Theory, and Peirce', in C. Vidales and S. Brier (eds), *Introduction to Cybersemiotics: A Transdisciplinary Perspective*: 75–94, Cham, Switzerland: Springer.
- Palmini, O. and F. Cugurullo. (2023), 'Charting AI Urbanism: Conceptual Sources and Spatial Implications of Urban Artificial Intelligence', *Discover Artificial Intelligence*, 3 (15).
- Picon, A. (2015), *Smart Cities: A Spatialised Intelligence*, London: Wiley.
- RIBA. (2024), *Riba AI Report*, London: Royal Institute of British Architects.
- Snodgrass, A. (1992), 'Asian Studies and the Fusion of Horizons', *Asian Studies Review*, 15 (3): 81–95.
- Snodgrass, A. and R. Coyne. (2006), *Interpretation in Architecture: Design as a Way of Thinking*, London: Routledge.
- Snow, C. P. (1959), *The Two Cultures and the Scientific Revolution*, Oxford: Oxford University Press.
- Tan, L. and L. Thom. (2024), 'Accelerating Future Scenario Development for Concept Design with Text-Based GenAI (ChatGPT)', *Proceedings of the 29th International Conference of the Association for Computer-Aided Architectural Design Research in Asia (CAADRIA) 2024 Volume I*, Singapore: CAADRIA.
- Tschumi, B. (1994), *Architecture and Disjunction*, Cambridge, MA: MIT Press.
- Turing, A. M. (1950), 'Computing Machinery and Intelligence', *Mind: A Quarterly Review of Psychology and Philosophy*, 59 (236): 433–60.
- Wang, D., L. Chang-Tien and Y. Fu. (2023), *Towards Automated Urban Planning: When Generative and ChatGPT-Like AI Meets Urban Planning (Working Paper)*, Ithaca, NY: arXiv.org.
- Willis, K. S. and A. Aurigi. (2018), *Digital and Smart Cities*, Abingdon, Oxon: Routledge.
- Winograd, T. and F. Flores. (1986), *Understanding Computers and Cognition: A New Foundation for Design*, Reading, MA: Addison Wesley.

1

LANGUAGE AS URBAN TECHNOLOGY

Artificial intelligence (AI) is already a familiar element in the city. The history of tool use in cities, the development of automation, and augmentations of human intelligence, such as writing tools and printing, demonstrate AI's evolving influence on urban life.

Language itself can be viewed as a form of artificial intelligence, as it facilitates purposeful thinking, the generation of ideas, the storage and transmission of information, and the formulation of instructions – all key components of automated intelligence. Language encompasses specialised subsets of instructions that guide human agents towards physical outcomes, such as instruction manuals, formulae, recipes and logical proofs. Moreover, certain forms of language serve as machine instructions, including computer programs. Even beyond their application in computing, language possesses generative properties that exhibit agency. As I explored in my book *Cryptographic City*, the technologies of writing and print have historically provided templates for city organisation and development.¹ If we assume that human cognition is embedded within our environments,²

then we can think of cities as extensions of the mind, with language playing an integral role.

These linguistic aspects of urban life entitle us to think of cities through the lens of AI. Since the start of the industrial age, citizens and specialists who reflect on the development of cities are primed to think of cities as potentially smart, intelligent and sentient, if not in actuality, then at least in our urban imagery, fantasies and in science fiction.³ The city therefore provides an ideal environment in which to test the applicability of linguistically based, conversational artificial intelligence.

Now, digital infrastructures pervade cities. Cities compete with one another to demonstrate their “smart city” credentials.⁴ The rollout of advanced network communications, digital media and artificial intelligence sustain a trajectory of ever-increasing automation in cities.

If it seems too extreme to suggest that even pre-industrial cities were imbued with AI precepts, concede at least that the provenance of AI in the city can be traced back to the factors already present in the idea of the city since its inception. We do not need yet to decide whether this tendency towards automation is beneficial or detrimental for citizens and the environment, or if it is inevitable. My point is simply that city developments have long prepared us for the current wave of AI development. Technological and electronic communications, computerisation, smart city infrastructures and artificial intelligence, for good or ill, are but extensions of the technologising of the world, and of cities.⁵

Some urban historians have long supported the idea of a city as a hot-house of technological development. Consider the storage and transmission of information in cities. In his exposition of the history of cities, influential urbanist Lewis Mumford (1895–1990) observed:

By means of its storage facilities (buildings, vaults, archives, monuments, tablets, books), the city became capable of transmitting a complex culture from generation to generation, for it marshalled together not only the physical means but the human agents needed to pass on and enlarge this heritage.⁶

By this reading, cities ancient and modern store and transmit information. Writing in the 1950s Mumford even noted that the storage media of the electronic age are not as capable as long-lived institutions and buildings:

“As compared with the complex human order of the city, our present ingenious electronic mechanisms for storing and transmitting information are crude and limited.”⁷ Gratifying to anyone invested in the city as a physical and social entity, here is endorsement for the claim that a city’s buildings and landscapes carry a heavy load in the preservation of cultural memory, and might we add *intelligence*.

As this book is about the urban consequences of conversational AI, I wish to emphasise the presence in cities of language as a technology that supports or supplements this role of generating, preserving and transmitting information. Think of cities as sites of conversational exchange.⁸ Mumford affirmed the primacy in cities of “human dialogue, the drama, the living circle of mates and associates.”⁹ The concept of the *polis*, from which we derive the term *metropolis*, is an active discursive domain in which ideas, texts and phrases are learned, combined and recalled, assisted by various tools of reproduction, combination and dissemination. In these respects, conversational AI is but the latest instalment of the technologies of communication in cities.

Cities have always served as communicative media, a position strengthened by the field of semiotics that identifies the widespread role of communication within human, organic, inorganic and cultural systems – including cities. I turn to a brief account of the city as a semiotic entity, and as a site of semiotic encounter.

Chatbot semiotics

Online chatbots have been available since the 1970s,¹⁰ initially operating as text-based digital platforms supporting simple interactions based on template questions and responses. The sea change came with the public release of OpenAI’s ChatGPT in 2022 and its subsequent developments and competitors: e.g. Microsoft’s Bing, Google’s Gemini, Anthropic’s Claude, Meta’s Llama 2.¹¹ The platforms’ convincing conversational acuity and ability to synthesise disparate conceptual threads projected the medium into public prominence and spawned a host of applications and integrations to be examined further in this book. My own forays into the medium and many stories online show how scholars, programmers, writers, musicians and artists use such conversational AI or similar as a creative companion to explore ideas, exceeding the offerings of search on an Internet browser.

My own excursion into the area described in several blog posts demonstrates the character of such explorations.¹²

The step change in chatbot development came with theories and methods for incorporating very large corpuses of texts (millions of gigabytes) loaded from the internet and other sources to train neural networks. As we shall see in subsequent chapters, a neural network typically consists of a matrix of millions of numerical data points stored in computer memory. Training involves micro-adjustments to their numerical values such that the neural network can detect and reproduce patterns derived from the training data. The method does not retrieve data from the corpus but generates responses to patterns in inputs presented by a human user as conversational prompts. The responses in general are consistent with patterns in the training data. Significantly, these patterns reflect both the language and content of the corpus. In recognition of the scale of the network and the training corpus, these systems are generally termed “large language models,” or simply LLMs.

Conversational AI can deliver sensible responses to intellectual propositions, even on the themes of this book. For example, I wanted to explore how text pervades the urban context, on notices, advertising, signage, display screens and inscribed on surfaces (Figure 1.1). A statement is as good as a question as input to an LLM. I prompted ChatGPT4 with the proposition: “Graffiti is both an art form and a means of disrupting the semiotic coding of the city.” The platform replied to my prompt in an instant:

Indeed, graffiti can be seen as a form of artistic expression that uses the urban landscape as its canvas. At the same time, it can disrupt the usual semiotic coding of a city, introducing new symbols and meanings that challenge the status quo.

So, the platform supplied terms that I had not introduced in the prompt: “urban landscape,” “symbols,” “meanings” and “the status quo” in response to my initial prompt about “semiotic coding.” The responses continued as I provided further prompts. The resultant textual exchanges can be described usefully as a convincing conversation.

We can infer much about conversational AI from this and other exchanges, not least a large language model’s apparent conversational acuity. As conversational AI is trained on pre-existing texts, the appropriately intellectual response given in the prompt-response exchange I have outlined also suggests the prevalence of urban semiotics as a theme in the general



Figure 1.1 Graffiti as both an art form and a means of disrupting the semiotic coding of the city. Almada, Portugal.

Source: Author.

intellectual corpus on which the platform was trained. Assuming they are trained on a broad corpus that includes a representative cross section of intellectual outputs, such LLM responses serve to demonstrate that concepts such as semiotics, graffiti and their surrogates are prevalent in intellectual discourse, though not always linked. We can treat conversational AI as providing a window into existing text-based intellectual discourse.

Urban semiotics

Encouraged by this corroborating evidence from AI's access to the intellectual corpus we can elaborate that semiotics is the study of signs.¹³ Semiotic scholars claim it has a broader scope than communication in human language as written and spoken. Language is but a subset of the many ways that communication happens. Communication occurs not only via books, articles and blogs, and as we compose emails, text messages, social media

and in-person conversations, but through the arrangements of civic spaces, the positioning of street furniture, ornamentation (Figure 1.2) and the signs in nature – autumn colours, birds heading for cover in a storm, and the constellations of stars at night. The strong assertion about such sign systems is that they are not just for the benefit of human beings, but within and between all biological systems according to semiotic scholar Thomas Sebeok.¹⁴ Others, such as Victor R. Baker, have expanded the notion to a geosemiotics¹⁵ – as if semiotics pervades the interplanetary inorganic world of rock, soil, air and chemical processes.

As indicated in our simple exchange with ChatGPT, whether pictorial or based in text, contemporary graffiti fits within the study of sign systems, as do the earliest marks and scribbles deposited officially, as art, to pass the time, or in protest, on cave walls, tree bark and built elements. But then the entire configuration of built elements, the morphology of a place, delivers signs to citizens and future generations. Mumford implied this when he extolled the virtues of long-lived institutions and buildings as a means of transmitting information. For a semiotician such morphological urban structures constitute systems of signs, communicating power, authority,



Figure 1.2 The variety and ubiquity of urban signs. Siem Reap, Cambodia.

Source: Author.

function, similarities, differences, interdependencies, relationships and consequential actions that amount to what we generally refer to as *meanings*.⁶

Signs in nature would aid ancient observers in predicting the weather, harvests and success in battle through auguries and omens. The New Testament book *The Revelation of St John*, also known as *The Apocalypse*, illustrates the emotional potency accorded to signs: the opening of a seal, the sound of trumpets, the gathering of armies as signs of impending destruction and the inauguration of the new Holy City. Plato, Aristotle, the Stoics and the Epicureans debated the role of signs in the relationship between the human and the world. Winfried Nöth's seminal *Handbook of Semiotics*, offers an extensive historical account of semiotic studies, including its many schools of thought, controversies and the disciplines touched by it.¹⁶ We learn that philosopher John Locke (1632–1704) was the first to mention the term *semiotics*. Locke said we need signs to communicate and remember our thoughts. Mathematicians Johann Heinrich Lambert (1728–1777) and Bernard Bolzano (1781–1848) introduced the systematic study of sign theory, leading to Charles Sanders Peirce's (1839–1914) major works on semiotics.¹⁷ Michel Foucault (1926–1984), a social philosopher, wrote a critical history focused on signs.¹⁸ He discussed societal changes through the lens of semiotics in his critique of European history.

If we need any further evidence from the intellectual corpus of the relevance of semiotics to architecture and urbanism, note that many 20th century thinkers contributed to the field of semiotics or drew on its theories. Geoffrey Broadbent was a co-editor and contributor to the book *Signs, Symbols, and Architecture*.¹⁹ Although more closely aligned with literature, Umberto Eco has a chapter in the same book "Function and Sign: The Semiotics of Architecture."²⁰ Donald Preziosi wrote *Architecture, Language and Meaning: Origins of the Built World and Its Semiotic Organization*.²¹ Charles Jencks and George Baird produced the seminal book *Meaning in Architecture*, a strong influence on architectural theorists eager to position their discipline in the wider cultural discourses affecting philosophy, literature and the humanities in the 1970s.²²

As a further illustration of a semiotic account of urban forms, Broadbent argued that a city's buildings might represent its cultural or historical identity as *icons*. Street signs and traffic lights provide direct information about the city's rules as *indices*, and certain urban designs or architectural styles might point to specific cultural or social ideas as *symbols*.²³ The applications

of language theories to architecture and urbanism range from studies into rule-based design principles and urban grammars, to the post-structuralism of architect Peter Eisenman and philosopher Jacques Derrida that sought to destabilise certainties about social and philosophical knowledge structures.²⁴ These diverse perspectives are relevant in the urban context, where cities are sites of tension among different social, economic and political viewpoints.

Language as generator

I have illustrated with the AI enhanced prompt-response exchange about graffiti, that large language models (LLMs) are adept at generating text. Language is also generative in any case. It shapes and influences possibilities. I am not only thinking of words as commands, orders, programs, warnings or even as instruments of persuasion, but in the way they frame spaces of possibilities, and therefore generate worlds.

If you don't have the words to expression an action you might take, then it is less likely that the action occupies your repertoire of potential actions. To overstate the case, if you have never heard the term "urban agriculture," then it is less likely that you will consider transforming vacant lots or rooftops into green spaces for growing food. By the same token, knowing the phrase and given the opportunity, it might just be an action that starts to appear on your community improvement list.

Support for the generative power of words comes from several key authors in the intellectual sphere. For example, Foucault's influential book *The Order of Things* effectively conflates words with power.²⁵ Which is to say, language enables and disables certain actions pertaining to the classroom, prisons, health and intimacy – a theory that extends to whole attitudes, practices and political movements within societies. As a further example, George Orwell's novel *1984* points to the power of *propaganda*, the manipulation of vocabularies and modes of speech that encourage or preclude certain actions – in the case of that dystopian novel, to diminish the likelihood that the citizens entertain thoughts, and therefore acts, of rebellion, and conform positively to a set of actions that sustain the political party in power.²⁶ From a different point of view, Martin Heidegger elevates language as "the house of being,"²⁷ proposing that language shapes who we are.²⁸ He was particularly interested in what he regarded as the poetic uses of

language. As it happens, the word “poetry” derives from the Greek *poiesis*, which carries the meanings of *revealing*, *unfolding* and therefore of *making*. As far as I am aware, Heidegger does not use the term “generating” in this context, though he makes much of *making* and relates it to the interest of architecture in so far as it is concerned with dwelling. Hence his collection of essays under the title: *Building, Dwelling, Thinking*.²⁹

At first, we may think of language as removed from the human agent who carries out the generative function. At most, language “programs” the human as agent, hardware, or “wetware,” to perform generative actions. Contrary to this linguistic determinism I think Heidegger’s assertions, and that of other Phenomenological thinkers, puts the matter of agency on hold. Their emphasis on *Being* suggests that as language users, we human beings are caught up in something greater than a collection of citizens speaking and performing individually. This Phenomenological world view lessens the distinction between agents and actors, actions and causes, subjects and objects, in favour of a lifeworld in which language is sovereign as “the house of *Being*.” Whether or not we follow this path to affirm the incontrovertible power of language, such Phenomenological thinkers show at least its prominence and importance in philosophical discourse, in how people live, and hence in cities.

In summary, I am formulating an argument about the ubiquity of language within cities, even independently of sophisticated communications networks and AI. Language pervades urban environments, shaping and reflecting the social, cultural and economic dynamics of city life. Signage, public discourse, legal language and everyday communication contribute to the construction and evolution of the urban environment. So, cities have always had the character of semiotic linguistics, enabled through technical and cultural narratives. But language has long enabled and inhibited, which is to say *shaped*, public discourse, and therefore actions. In this and other respects, language is a crucial generative element within the formation of cities.

To the extent that we are persuaded of the linguistic character of the city, that buttresses the relevance of LLMs in discussions about the city. If LLMs are impacting the way we interact, communicate and even think then that entitles us to regard AI and its progenitor theories as already nascent in urban dynamics, prior to computers and advanced digital communications. Furthermore, a focus on how conversations are framed in the city context

will reinforce that bold proposition. My point here is that the extent to which you can think of cities as linguistic entities they are already invested in recent technologies of language such as large language models. Language may be preeminent in urban experience, and we can say something similar about conversation.

Urban conversations

We could look to the extensive dialogues between Socrates and his students in Plato's writing, or the arguments between prophets and their God in religious texts,³⁰ as exemplars of the gravity of meaningful conversational exchange. But conversation does not have to be profound or articulate to be important to urban life. The origins of conversation are humbler than that.

It is easy to decry the decline in real, genuine conversation in the digital age. Think of the babble of relentless advertising, pop song lyrics streamed in public and domestic settings, or funnelled into ears through headsets and airbuds. Early critics of social media would point to the trivialisation of human interaction augured by our obsession with smartphones, and the incessant chatter of the Internet: forums, unmoderated channels, such as podcasts, Twitter, X and Threads, flooding the polis with trivial discourse, contradictions and untruths. This chattering content is then harvested for profit, to profile groups and individuals, and to target advertising. For critics such as Sherry Turkle, our digital devices distance us from one another, while promoting the delusion that our lives are enhanced through this constant imperative to communicate.³¹

But for all the interpersonal, social and ethical challenges we may attribute to it, social media reflects something of the primal function of conversation, as recordable, curated and rendered public. Most conversation as we experience it in everyday life is not deep, intimate, gracious or edifying. Nor is the scripted conversation we observe or experience vicariously through plays, magazines, novels, movies, reality tv shows or the stories we tell one another. Conversation is base and ordinary, and in that it is also urban, if not urbane, pedestrian and every-day.³²

To converse comes from the Latin *conversare*, meaning to turn yourself about, to move to and fro, as when we speak of the *converse* side of a coin, or an argument. According to the OED, *conversare* also came to mean "to dwell, abide, live somewhere, keep company with."³³ The connection

between turning and communal dwelling is not entirely obvious until we observe that keeping company with others inevitably involves negotiated exchanges, a to-and-fro dynamism, turning together and apart.³⁴ Recall Mumford's characterisation of the city at its best as a place of "human dialogue, the drama, the living circle of mates and associates, the society of friends."³⁵ By this reading, conversation plays a major part in the life of any city. It is there at the origins of cities as places to gather, to be together.

Mumford affirms the roots of the city in communion around the domestic hearth.³⁶ This is an origin story promoted by the early architectural theorist Vitruvius who asserted the primacy of conversation in his mythic account of the ascent of humankind and the beginnings of architecture. Vitruvius painted a picture of humans emerging from the forest into a clearing and gathering around a fire, a kind of primal hearth: "it was the discovery of fire that originally gave rise to the coming together of men, to the deliberative assembly, and to social intercourse."³⁷ In that setting people discovered "articulate words," and by being able to name things "they began to talk, and thus originated conversation with one another."³⁸

Of course, formal language is not necessary for such communion. Semiotics maintains that human beings interact even without the ability to speak and write. We can converse among ourselves, with inarticulate infants and even with non-human animals without the need for words. There is a pre-linguistic communion in play, a conversation without words. Theories of semiotics assume the existence of such a pre-linguistic conversation by affirming the ubiquity of communication across domains and in many media, including those that do not involve words.

Phenomenological philosophers such as Maurice Merleau-Ponty (1908–1961) assign the capability to perceive and express ideas to the physical body, in movement, comportment and gesture.³⁹ The human and animal body possess a kind of knowledge and communicative acuity that precedes and goes beyond formal language and is just as prone to the free to-and-fro that we associate with verbal conversation. Whereas it is easy to assume you need to use articulate language before you can conduct a conversation, these observations suggest the reverse and place the ability to converse as a prerequisite for formal language. At the very least, we could concede that pre-linguistic conversation and language develop together. The co-development of word and gesture could apply to the cognitive development of the individual as well as in the broad spread of human development.

Once “articulate words” were admitted into the “living circle of mates” there was scope for even more purposeful chatter. Yuval Harari’s book *Sapiens: A Brief History of Humankind* concurs. He foregrounds the human propensity to converse, to gossip and to tell stories about one another. By this theory, human language developed as a way of surviving and even thriving through cooperation and the exchange of information, whatever the truth status of that information: even if imaginary, speculative, unfounded or untrue. According to this view, the role of conversation in engendering community is even more important than the conveying of accurate information. Conversation establishes solidarity, whatever its content, including gossip. Harari explains in contemporary terms, “the vast majority of human communication – whether in the form of emails, phone calls, or newspaper columns – is gossip.”⁴⁰ Harari even asserts that media and the press have their origins in rumour as the quintessential manifestation of gossip: “Rumour-mongers are the original fourth estate”⁴¹ that is, the press. Subsequently he lauds the unique human capacity to be inventive, as purveyed through gossip: “It’s the ability to transmit information about things that do not exist at all.”⁴²

Samuel McCormick’s book *The Chattering Mind: A Conceptual History of Everyday Talk*, is also helpful in affirming the formative role of conversation.⁴³ He notes that “so much of today’s chatter, idle talk and empty speech now occurs online, in virtual assemblies established and maintained by the mobile internet.”⁴⁴ He traces how Martin Heidegger (1889–1976), Søren Kierkegaard (1813–1855), and Jacques Lacan (1901–1981) identified the problematics of “chatter,” “idle talk,” “empty speech,” “babble” and “scribbling” as types of everyday discourse that are pre-requisite to the ability to generate and develop profound, careful and grounded discourse. Chatter answers to a human propensity to keep the conversation going, for its own sake: “Speakers frequently suspend the pursuit of attainable rhetorical advantage in order to prolong their own utterances – and for no other reason than to prolong their own utterances.”⁴⁵ Chatter proceeds as if to its own ends, which is only rarely to close off the argument, or win the case. To simply filibuster is the likely goal of the chatterer. Calling on Lacan, McCormick observes that empty speech has “the sonic properties of an echo, a feedback loop, a skipping record – and usually all at once – the recursive mechanical racket of which can only be described as ‘a great cacophony.’”⁴⁶ The idea that idle chatter corresponds in some ways to the

operations of a machine is suggestive as we contemplate automated natural language generation and LLMs.

I like this implication that conversation as a way of exchanging and sharing some empirical reality is secondary to conversation as a means of keeping the sociability of chatter in train. It supports a kind of “alien” rationality, where natural language processing systems aim not to construct arguments but simply to optimise their capacity to generate the next best word to keep the text in flow. For evidence of this relentless resolve think of the propensity for an AI chatbot to keep generating, to keep chattering, were it not for the significant processing costs entailed. There is some truth in that as a model of human chatter. That cities might be sites of endless chatter and storytelling is also an appealing notion, especially in support of the importance of conversation in the development and life of cities.

As well as his extensive writing on electronic media,⁴⁷ Marshall McLuhan with others produced a long-playing record “The Medium is the Massage,” released in 1967.⁴⁸ As a work of collage and of its era in tone and humour, it gives acoustic expression to the idea of chatter, the cacophonous overlay of voices and other sounds. Pre-literate, oral, “village” cultures participate in undifferentiated soundscapes unencumbered by reliance on the sense of sight and thereby to separate, discriminate and distance villagers from one another and environment. With the revival of the sonic sense, particularly through electronic media, McLuhan argued that the world is becoming a “global village,” characterised by immediacy, involvement and even a strong sense of social cohesion. According to McLuhan: “In the electronic age we wear all [hu]mankind as our skin.”⁴⁹

Apart from delivering a barrage of voices, technologies contribute to conversational exchange through amplification and transmission. I need scarcely catalogue all the ways that electronic and digital media amplify, preserve and transmit conversations: mobile phones, video and audio-conferencing platforms. The influence of digital media on communication has been catalogued extensively by others. I prompted ChatGPT to expand on the list of digital media channels and save me some effort:

social networking sites like Facebook, Twitter, and Instagram, podcasting channels, online discussion forums, interactive television broadcasts, virtual reality environments, messaging applications like WhatsApp and Telegram, collaborative work tools like Slack and

Microsoft Teams, real-time language translation services, AI-driven personalised communication assistants, live streaming platforms like YouTube and Twitch, community-driven content sites like Reddit, voice-controlled smart devices like Amazon Alexa and Google Assistant, innovative educational platforms that enable global learning, emerging technologies like 5G that facilitate faster and more reliable connections, and even satellite internet services that promise to connect remote and underserved regions of the world.

These contemporary facets of urban life continue the trajectory of the city as a relentlessly chattering entity, with AI providing an increasingly prominent role in sustaining the flow.

In summary, places and spaces for conversations, discussions, communications and dialogue happen in cities by design, planning, legislation, contestation, happenstance – and dialogue. Besides these physical places and processes, cities are formed with conversation in mind. The polis is an assembly of people, the metropolis is a place for the people – who are forever talking together.

Urban rationality

So far, I have suggested that the roots of urban conversation reside within informal, familial chatter. Content, what is ostensibly being communicated, is secondary to the simple fact that people talk to one another or talk in the company of one another. That conversation can be consequential is obvious. Humans mould and discipline their chatter to specific ends. Purposeful, directed conversation, whether in the form of spontaneous protest, political debate, court trial, school exams or a client interview, forms a basic constituent of rationality. Cities develop through such rational conversations.

Urban planning theory affirms as much. At their best, developments in cities are informed by consultation with large numbers of residents and stakeholders. For example, in his summary of the state of planning theory in 2005, planning theorist Phil Allmendinger wrote encouragingly that the UK planning process is “predicated on inclusion and consultation” as means of garnering “stakeholder values and desires which are expressed at both the local and regional level.”⁵⁰ Consultation necessitates conversation.

Such planning theory draws theoretical weight from the writing of the sociologist Jürgen Habermas, particularly his influential book aptly titled *The Theory of Communicative Action*.⁵¹

Habermas's communal concept of rationality provides the benchmark for urban decision making. Such a view grounds decision making in *reality* as experienced by urban participants. According to this theory, the urban participant does not begin simply with the quest for "objective evidence."⁵² Rather, what is real is grounded in communities engaged in meaningful and consequential actions. Rationality is subject to conversational interactions within such communities.

How do we deal with or filter out false or imaginary readings of the world? To provide a trivial example, to falsely posit that London in reality is host to secret agencies for managing interplanetary aliens (e.g. the film *Men in Black*) would only be of practical use in so far as any significant actions might follow from that belief. Hence the term "communicative action." In more general terms, Habermas says of communicative communities that they reassure themselves "of an intersubjectively shared lifeworld"⁵³ by sharing their varied interpretations. This common background of understanding does not presuppose agreement, or even full knowledge of each other's position. According to some urban scholars, cities are inevitably sites of contest. They are productively *agonistic*.⁵⁴ To confirm cities as such positive agonistic domains of action is also to affirm the *hermeneutical* position to which many professionals and scholars subscribe.⁵⁵ It takes account of varieties of interpretations found in any urban situation.

That AI might have a role in such communicative practices is certainly of consequence. Much of the 21st century planning literature subsumes AI within discussions (approvingly or critically) about the cybernetic idea of the "smart city." But the linguistic framing I am advancing in this book suggests that AI in the city, in prospect or "reality," has consequences for urban communicative practice. We wish to see how this technology intervenes in what is already a conversational domain, namely the formation of our cities, including its decision-making processes.

Urban story telling

One of the primary activities permeating conversational practice is telling stories. So much of idle chatter and gossip involves relaying narratives.

A story is but a continuation of the human propensity to connect with other human beings. It expands the circle of interlocutors to incorporate as characters in our stories those not present.

The elevation of narrative is familiar in the conceptualisation and formation of cities, especially considering the potential role of conversational AI in the city. Certain urban planning theorists affirm the inevitable link between story-telling and urban space. For example, planning theorist Lieven Ameel characterises this emphasis on narrative as a paradigm shift “towards a more dialogic, participatory and discursive form of planning”⁵⁶ that resorts increasingly to “diverse urban narratives.”⁵⁷ Planner and theorist James Throgmorton concurs, and even advocates that we think of urban professionals “as authors who write texts (plans, analyses, articles) that can be read (constructed and interpreted) in diverse and often conflicting ways.”⁵⁸

His aptly titled book *Planning as Persuasive Storytelling*, and follow-up articles, identifies key narrative elements in the planner’s craft: authors, characters, settings and plots declaring that, “These planner–authors have to build conflict, crisis, and resolution into their narratives, such that key antagonists are somehow changed or moved significantly.”⁵⁹ This art of persuasive urban story telling extends even beyond words to tables, charts, diagrams and models. He maintains that as authors, planners draw not only on “the imagery and rhythm of language,” but also the language of “statistical models, forecasts, GIS-based maps, three dimensional architectural renderings and virtual reality models, surveys, advisory committees, and other persuasive figures of speech and argument, or tropes.”⁶⁰ By these and other linguistic means planners seek to persuade “readers” and influence the course of events in the city.

The main audience of Throgmorton’s argument is professional planners, but we may presume that the same narrative processes apply to all participants in urban change, including elected officials, developers, commercial interests, citizen groups and individuals.

This narrative orientation to city planning is at odds with a more empirical, calculative approach. Throgmorton notes the countervailing presupposition that practicing planners “can be neutral, objective, rational adjudicators of the public interest, and that their texts have a single literal meaning (the one planners intend) that any intelligent person can grasp.”⁶¹ Such empirically oriented traditionalists would argue that deviation from

evidence-based validation into the territory of story-telling invokes the possibility of falsehoods, figments of imagination: “making something up, writing fiction, telling lies.”⁶²

But for Throgmorton, storytelling serves as a vehicle for establishing meaning that seeks to bring people round to particular points of view. This linguistic reframing enhances the planner’s role: “To ask about meaning and persuasiveness rather than truth is to shift attention from technical accuracy to a combination of accuracy and normative evaluation.”⁶³

Returning to Habermas and “communicative action,” this narrative approach to professional expertise constitutes a distinctly hermeneutical position that I have rehearsed elsewhere.⁶⁴ Observations, evidence gathering, calculations, mapping and technical practices are subject to the processes of interpretation – in their selection, adoption, methods, the weight we attach to them, their evaluation and their application. From the point of view of this book’s main thesis, the centrality of conversation as an urban determinant prepares the ground for us to consider the role of conversational AI.

Examples of AI storytelling

As it happens, at least in their current iterations, conversational AI platforms are good at creating stories. As one of many examples, in one of my own encounters, I prompted ChatGPT to construct a plausible story about an indentured labourer working in a sugar plantation in Mauritius. It is an esoteric theme, paralleled in many countries, but the narrative was specific enough to weave in factors to do with sugar cane growing, reduction in biodiversity, exploitation of labour, colonisation and tourism. From my knowledge of the country, the fictional story seemed convincing, and I found that the platform was able to redirect its narrative in response to further prompts from me.

- Me: Rajesh and his compatriots had little time or energy to concern themselves with the depletion of biodiversity on the island caused by the industry with which they were involved.
- ChatGPT: As Rajesh and his compatriots toiled under the scorching sun on the sugarcane plantation, their daily existence revolved around meeting their basic needs and supporting their families.

This was dialogical storytelling, a brief, but open shared narrative project. Countless examples demonstrate this native propensity of a system trained on a vast amount of textual content to make up stories and details.

As a further illustration of the specificity available to large language models, I ran a class in which we invited groups of students to recruit ChatGPT to construct stories featuring the group members by name, to assign roles to them and weave a story about their complicity in some global catastrophe. ChatGPT delivered imaginative stories with the same alacrity as if prompted to write computer code. In fact, the platform narrativised its response chronologically as if relating a story.

Soon after OpenAI released its version of ChatGPT that could read, analyse and generate images, I uploaded a photograph of some basic kitchenware stacked on a table and asked the platform how I would get from that layout to another if I had only one arm available for the task. I was thereby paralleling the early AI challenge of requiring a robot to construct logically a sequence of moves to manipulate toy building blocks.⁶⁵ ChatGPT delivered the instructions for my tabletop challenge in narrative format:

Ah, the classic teapot-inside-a-jug manoeuvre! Here's a step-by-step for the one-handed academic: Gently lift the miniature teapot out of the jug and place it on the table. Now, pick up the jug and set it onto the coaster or mat where the teapot was. Lastly, lift the miniature teapot and nestle it inside the jug.

The AI correctly identified the relative scales of the items, the fragility of their positioning and assumed a description of the person asking for the sequence, "the one-handed academic." Even for a straightforward procedure (pick up the teapot, place it on the table, etc) the platform could not help but deliver its solution to the challenge in story format.

Such developments in AI are indeed impressive. But these and many other experiments recounted in the AI literature and by online pundits also demonstrate the ubiquity of storytelling in the training corpus. In so far as large language models depend on a wealth of general training texts, we could conclude that story telling is ubiquitous in written texts. Story formats are inevitable components of text and writing. The LLM platforms reflect that.⁶⁶

We may suppose that the stories AI will tell reflect something about the training data, but we can be less confident that they reflect a reality outside

of that. Researchers are developing AI systems to combine conversational acuity with accurate performance in handling procedures, logic, calculation, and mathematics.⁶⁷ My use here of the LLM platform demonstrates the propensity of such systems, by virtue of their training data and methods, to favour story over accuracy. How designers, urban planners and decision makers might marshal AI storytelling in the service of some useful and meaningful urban mediation or intervention will emerge from subsequent discussions in this book.

Speaking

So far, I have conflated conversation as an exercise in exchanging text with people's more immediate participation in conversation as spoken. Though we think of conversation as primarily a spoken practice involving the human voice, LLMs deal primarily in written texts. Following McLuhan's argument about the benefits enjoyed by aural cultures, input-output patterns of interaction via text are distant from the immediacy of language as spoken. Only the voice can convey nuance, responsiveness to immediate context, provide the necessary cut and thrust between interlocutors, and the uninterrupted flow of chatter to which we have become attuned throughout the development of language. Though it inherits the benefits of text in providing time to think, lay things out and provide a record, conversational AI as an operation with text confronts certain limits.

By this argument, the exercise of the voice in language is the primary mode of conversation. The ability to write down what is said, to print and store it as text is a secondary human capability. After all, we presume that pre-historical humans spoke, orated, recited and memorised without the ability to write. Writing came later. According to this argument, it even follows that speech is closer to pure thought than writing. Writing provides benefits, but raw speech provides access to authentic being, originality and immediacy. As I stated earlier, that sentiment reflects the common view advanced by McLuhan and others that aural (pre-literate) cultures are more sociable than text-based, visual cultures, and enjoy more immediate engagement with environment.⁶⁸

Being able to store and retrieve speech as text is one of the hallmarks of cultured, enduring, technologically sophisticated and historically aware societies, but such text-based achievements depend on speech as the

primary exemplar of conversational acuity, sociability and language development. That large language models depend on vast quantities of text, not speech, poses a challenge to the claims I am making about the influences of conversational AI.

Most people value the ability to speak off-the-cuff, to extemporise a speech. I am thinking of when a teacher engages with a class of students, a politician answers questions without the aid of a teleprompter, or a dinner guest is called on to propose a toast. Audiences are inclined to associate speaking well, including in conversation, with human intellectual abilities. When dealing with automated systems, we may think that direct vocalised conversation serves as an indicator of machine intelligence – more than tapping buttons or reciting text. The robots in the SciFi novels of Isaac Asimov spoke – along with disembodied AIs as in the case of HAL in *2001 A Space Odyssey*, and the movie *Her*. Voice appears “most natural” to us, as something you should be able to enlist if you are to connect with synthetic intelligence. We might prefer to speak out loud in a video call, as we set the timer in the kitchen, or call up weather information, especially while multitasking. Voice activation though is not always the most convenient means of interacting with a computer or other automaton. None-the-less we expect it of conversational intelligent LLMs.

As depicted in popular culture, and hence as a means of marketing the technology to a wider citizenry, people expect AI to be something you can talk to, i.e. that listens and speaks. Alan Turing’s test (the Imitation Game) required that you should be able to communicate with a computer via text, as if via a typewriter or teleprinter. Text preserved anonymity in his hypothetical experiment and left the human-or-machine question as undetermined until you analysed what the machine wrote, i.e. its “content.” However, the quest for conversational AI, as in the use of voice activated systems like Siri and Alexa reinforced further with OpenAI’s VoiceChat indicates a commitment amongst developers and consumers to competent speech as a sign of intelligence.

Synthetic speech

Certain implementations of large language models generate convincing spoken responses to spoken input. These systems translate spoken words into text that is fed into the LLM in the usual way. It then translates its text

responses to voice. You can talk (as well as type) to your AI, as if on a smartphone. Voice chat draws on several technologies, melded to produce near humanlike interaction. You speak into the app, which then translates, invisibly, your words into text. The text is fed into the LLM which processes the text in the usual way to generate a text response. That text is then converted by a high-quality text-to-speech (TTS) system to produce spoken output. This flow generally happens quickly enough to deliver conversation-style vocal interactions. The content of the responses from this voice chat app are equivalent to the responses you get when you type text-based prompts, and the platform also retains a text transcript for further reference and editing.

So, large language models are trained not on spoken words, but on written texts. That text-based sourcing is represented in the content of the voice chat responses. The TTS component of such systems also undergoes training from voice data, but that is independent of the LLM training. Of course, human learners do not develop language competence primarily from reading texts but from listening, speaking and responding to feedback and corrections in everyday worldly settings, both material and social.

To train an LLM just on voice data rather than text, analogous to the way infants acquire language competence, is beyond the current capabilities of LLMs. Voice data audio files are substantially larger than text files. As yet there is no theoretical or practical way to train a neural network on masses of voice input (e.g. recordings of people talking in social settings) that would come close to the competence demonstrated by LLMs such as ChatGPT. The storage and processing demands of large audio files, the lack of standardisation among speaking voices, and the technical challenges of taking account of large scale contexts are prohibitive.⁶⁹

The speech-to-text (STT) component of the voice chat model is trained on voice data as a process independent of LLM training. Modern TTS systems, such as those using deep learning techniques, can generate speech from text on the fly. These systems are trained on large datasets of recorded speech and their corresponding text transcriptions. The neural networks train on mapping text inputs to audio outputs, capturing subtleties such as intonation, stress and rhythm to produce natural-sounding speech. For real-time applications, optimised models and efficient hardware processing ensure that TTS can generate audio quickly enough to be used in interactive systems that perform transcription and translation, operating as multilingual conversational assistants.

I have experimented with voice chat and demonstrated it to others in multilingual contexts. The performance is impressive, if not uncanny, particularly if used as a translator. It will not only translate between languages in real time, but do so conversationally, addressing the participants by name if introduced. In this short excerpt, Mateo speaks very little English and I have no Spanish:

- Me: Please ask Mateo what was his job.
ChatGPT: Mateo, ¿cuál era tu trabajo antes de jubilarte?
Mateo: Mi trabajo era ser profesor de matemáticas.
ChatGPT: He says, 'My job was being a math teacher.' ...
Me: Could you please ask him in Spanish, who was the first great mathematician in history?
ChatGPT: Mateo, ¿quién fue el primer gran matemático de la historia?
Mateo: ...
ChatGPT: He says, 'I would say Pythagoras, then Euclid...'
Mateo: Ipatiam.
ChatGPT: He says, 'Hypatia.'
Me: Hypatia is very unknown. This is Richard speaking.
ChatGPT: Yes, Richard. Hypatia is indeed less well-known compared to Pythagoras and Euclid. She was a renowned mathematician, philosopher and astronomer in ancient Alexandria. Her contributions were significant, especially considering the time period she lived in.
Me: Could you please say the same in Spanish to Mateo?
ChatGPT: Sí, Richard. Hypatia es ciertamente menos conocida en comparación con Pitágoras y Euclides. ...

The model picks up spoken text, recognises the languages in use, and translates parts of the dialogue in a conversational way. It also contributes new information to the conversation, against a noisy background in which others were also speaking not least vocalising wonderment at the performance of the AI platform, and concern about what it would do to their professions as teachers.

After a few interactions the imbalance in such conversations becomes clear. The AI has instant access to information and insights that the human participants could match only with extra time to recall and assemble. Voice

chat does not hesitate. It pays little attention to the to-and-fro dynamics of voice-based conversation. Replies to prompts from the system are articulate, delivered with seemingly expert detail and alacrity. Interruptions by human or machine are less effective than amongst human interlocutors. Humans in the conversation are unlikely to keep up as they hesitate while seeking the best word sequence for the next prompt. Knowing about the capacity of LLMs to deliver falsehoods, the role of the human being becomes that of fact checker.

At least in the current iteration of the technology, text-to-speech (TTS) algorithms do not adapt the content of the output to suit voice-based conversation. It converts the generated text into spoken words while preserving the original content. The TTS adds intonation, emphasis and rhythm, but it does not change the wording or meaning of the text provided by the language model, at least at the time of this demonstration.

A speaker will give emphasis to certain words to direct the attention of the other party in the conversation. The emphasis in spoken responses typically comes from the text-to-speech (TTS) system, which uses its own algorithms to make the speech accord more closely with what we expect from human conversation. A sophisticated TTS system will interpret punctuation, word choice and sentence structure to add intonation and emphasis. To change intonation, you must call up a different voice from a menu of voices, e.g. for Australian or Scottish English. It will not read a made-up Dr Zeuss story with the intonation you would use if reading to a small child. Despite various difficulties and limitations as demonstrated here, LLM developers are keen to preserve and develop the use of the spoken word in conversational AI.

Text is sufficient

From their inception, LLMs present speech as dependent on texts. They are designed around text and trained on text data. There's philosophical support for this homage to text as a primary mode of communication. The philosopher Jacques Derrida (1930–2004) builds a controversial philosophy on the proposition that any attempt by people to assert the priority of speech over writing must draw on the attributes of writing (i.e. text). You can't talk about speaking without drawing on explanations, terms and metaphors that presume the pre-existence of writing.⁷⁰

As we have seen, and as if in evidence of the priority of writing in language acquisition and knowledge creation, LLMs depend on text in their training – vast corpuses of text files (writings). The giant neural network models that power LLMs undergo incremental parameter adjustments putting them in a state where they can yield conversational exchanges with convincing linguistic competence and even more remarkable demonstrations of accessing and synthesising information. They generate convincing contextual responses to prompts from patterns within all those training texts.

As it is difficult to envisage a technology that would acquire language competence and access to knowledge in ways analogous to the development of human cognition, we must make do with text inputs to LLMs. That happens to be a highly efficient way of short circuiting the complexity of human language and knowledge acquisition. It turns out that text is sufficient.

Though AI models cannot simulate life experience, a reliance on text compensates for the lack of experience and world engagement. If we needed reminding, text serves as a rich repository of human experience, encapsulating the thoughts, emotions and contexts of its authors. When large language models are trained on vast amounts of text, they effectively incorporate these human experiences, and the complexity of various situations described within the text. This way, the models can replicate or reflect human-like experiences and responses without needing to have lived lives. So, while LLMs are trained on text, they are also, in a way, trained on the human experiences that give rise to that text.

From a philosophical perspective, the reliance on text for training large language models aligns with Derrida's ideas about the primacy of writing. Derrida argued that writing is not a mere secondary representation of speech but is fundamental to the structure of language and meaning itself. In this sense, the efficiency and effectiveness of text-based language models can be seen as a technical validation of Derrida's theory. By starting with text, we acknowledge the structured, persistent nature of written language, which allows for extensive analysis and pattern recognition. Text is stable and can be revisited and scrutinised, much like Derrida's idea of writing as a steady yet dynamic foundation for meaning.

So, the current efficiency of text-based models not only demonstrates the practical advantages of written language for AI development but also

emphasises the deep-seated role of writing in the structure of human communication. This AI approach to language processing brings into relief Derrida's idea that writing, in its broadest sense, underpins and shapes all forms of linguistic expression.

Reframing the city

To summarise, I have argued in this chapter that we are entitled to think of AI as an integral part of urban environments, with its seeds dating back to the earliest cities through the development of language, writing and automation. Language, as an aspect of artificial intelligence, facilitates the storage and transmission of information, shaping human cognition and urban development. Cities have always been technological hubs, seeding cultural "memories" and serving as communicative media, paralleling or presaging advanced AI systems.

Digital urban infrastructures and smart technologies amplify the presence of AI. The integration of AI in cities continues a historical trend of technological change and information transmission, bolstered by the view of cities as extensions of human cognition.

The relevance of conversational AI to urbanism is highlighted by the city's role as a site of conversational exchange. Conversational AI exemplifies the latest in communication technologies, building on the longstanding tradition of cities as semiotic and communicative entities. This perspective emphasises the importance of understanding AI's impact on urban life, considering cities have always been spaces where language and technology converge.

In keeping with the subject matter of this book, I recruited ChatGPT to remind me what might be radical, controversial or provocative in this introductory chapter. The AI platform pointed out that my arguments could be construed as treating language as a form of artificial intelligence. Such a simplification would expand the definition of AI to include almost anything (or nothing), including all cognitive and social functions. Contrary to that expansive view, which effectively dissolves crucial distinctions between legacies and innovations, I am proposing that language, urbanism and AI share similarities in their development trajectories. That said, consistent with a linguistic orientation, I am offering in this chapter an alternative framing for how we examine AI in the city. It is language based, rather than cybernetic.

It may look as though the perspective of this chapter frames AI as an inevitable extension of urban evolution rather than a disruptive innovation. I would counter that by saying that disruption is in fact one of the defining characteristics of cities. Their concentrations of people, problems and resources facilitates their role as hothouses of innovation. AI is a disruptive entity in a long line of urban disruptions, each of which has claimed significance as transformative in major ways: vaccines, steel reinforced concrete, the elevator, plate glass, photovoltaics, container shipping, photo sharing, the smartphone. Such technologies are in the company of disruptive social and political events that serve similarly to define pivotal urban moments: pandemics, civil action, protest, terrorism, urban destruction. In fact, as an effective language game, scholars have sought to demonstrate how some apparently marginal event or innovation has turned out to be pivotal in the development of cities. Such framings and reframings serve well how cities are developed, perceived, valued and lived. Different ways of looking at things opens a space for new insights.

Some readers may resist the idea that cities, or at least elements within cities, operate as extensions of mind. That is a particular framing developed by scholars and that serves us well in connecting urban environments and human thought processes.⁷¹ My standard understanding of such overarching philosophical questions is what difference it makes to adopt one framing over another. From the urban perspective, theories of extended mind and embodied cognition really makes a difference to how we intervene in the city as designers, planners and activists, not least in the design of digital urban systems.

The argument of this chapter may suggest an inevitability to developments in AI and their impacts on the city, diminishing any imperatives to manage, legislate, resist or otherwise curtail its developments and applications. I discuss such responses in Chapter 12. My objective here is to promote understandings of this technology, so we may better respond to it.

Linking AI and urbanism to semiotics as I have done in this chapter might be taken as controversial. Though I have shown that there is broad support for this view amongst planners, I think this perspective challenges traditional urban planning and architectural theories. With its emphasis on storytelling and narrative, this semiotic framing challenges empirical, data-driven approaches by highlighting the discursive nature of urban development. I think that such shifts in thinking might nudge

considerations of AI further in the direction of the humanities, which has ample credentials in the intellectual sphere, offering tools for an integrated and historical view of AI's role in cities.

Notes

- 1 Richard Coyne, *Cryptographic City: Decoding the Smart Metropolis* (Cambridge, MA: MIT Press, 2023).
- 2 Andy Clark. 'Embodied, Embedded, and Extended Cognition', in *The Cambridge Handbook of Cognitive Science*, ed. K. Frankish and W. Ramsey (Cambridge: Cambridge University Press, 2012), 276–91.
- 3 The writings of author Philip K. Dick, and their film adaptations, provide rich examples of the presence of AI in cities. See Philip K. Dick, *We Can Remember It for You Wholesale* (London: Millennium, 2000); Philip K. Dick, *Do Androids Dream of Electric Sheep?* (London: Voyager, 1996). Also see William Gibson, *Neuromancer* (London: Harper Voyager, 1995). For the relationships between cities and the imagination see Darran Anderson, *Imaginary Cities* (London: Influx Press, 2015). I also have in mind a science fiction tv series from the 1960–70s in which the residents of an extra-terrestrial city were unwitting conveyors of instructions and actions for some higher calculative purpose, even as they went about their daily business.
- 4 Rob Kitchin, 'The Real-Time City? Big Data and Smart Urbanism', *GeoJournal*, 79, (2014): 1–14; Simon Joss, et al., 'The Smart City as Global Discourse: Storylines and Critical Junctures across 27 Cities', *Journal of Urban Technology*, 26, no. 1 (2019): 3–34; Amirhosein Ghaffarianhoseini, et al., 'What is an Intelligent Building? Analysis of Recent Interpretations from an International Perspective', *Architectural Science Review*, 59, no. 5 (2016): 338–57.
- 5 I take the term "technologizing" from the book by Walter J. Ong, *Orality and Literacy: The Technologizing of the Word* (London: Routledge, 2002).
- 6 Lewis Mumford, *The City in History: Its Origins, Its Transformations, and Its Prospects* (Middlesex, England: Penguin, 1961), 648. Also see Shannon Mattern, *A City Is Not a Computer: Other Urban Intelligences* (Princeton, NJ: Princeton University Press, 2021), 63.
- 7 Lewis Mumford, *The City in History: Its Origins, Its Transformations, and Its Prospects* (Middlesex, England: Penguin, 1961), 648.
- 8 Bernard Stiegler, *Technics and Time* (Stanford, CA: Stanford University Press, 1998); Joanna Zylińska, *AI Art: Machine Visions and Warped Dreams* (London: Open Universities Press, 2020); Claudio Celis and Pablo Ortuzar Kunstmann. (2021), 'We Have Always Been Artificially Intelligent: An Interview with Joanna Zylińska – Claudio Celis and Pablo Ortuzar Kunstmann'. *Culture Machine*. Available online: <https://culturemachine.net/vol-20-machine-intelligences/an-interview-with-joanna-zylińska-claudio-celis-and-pablo-ortuzar-kunstmann/> (accessed 27 May 2023).
- 9 Lewis Mumford, *The City in History: Its Origins, Its Transformations, and Its Prospects* (Middlesex, England: Penguin, 1961), 648.
- 10 Joseph Weizenbaum, *Computer Power and Human Reason: From Judgment to Calculation* (San Francisco: W. H. Freeman, 1976).

- 11 Michael McTear and Marina Ashurkina, *Transforming Conversational AI: Exploring the Power of Large Language Models in Interactive Conversational Agents* (New York: Apress, 2024).
- 12 See posts at <http://techmediaculture.com>. Also see Tim Smithers, 'Generative AI: Useful Tools or Expensive Toys', *Teaching and Strengthening Research*, 03.24, (2024).
- 13 For further examination of the topic see my own related books and articles: Richard Coyne, *Derrida for Architects* (Abingdon, Oxon: Routledge, 2011); Richard Coyne, *Network Nature: The Place of Nature in the Digital Age* (London: Bloomsbury Academic, 2018); Richard Coyne, *Peirce for Architects* (London: Routledge, 2019).
- 14 Thomas A. Sebeok, *Signs: An Introduction to Semiotics* (Toronto: University of Toronto Press, 1999).
- 15 Victor R. Baker, 'Geosemiosis', *Bulletin of the Geological Society of America*, 111, no. 5 (1999): 633–45.
- 16 Winfried Nöth, *Handbook of Semiotics* (Bloomington IN: Indiana University Press, 1990).
- 17 Charles Sanders Peirce, *The Essential Peirce, Selected Philosophical Writings Volume 1 (1867–1893)* (Bloomington, IN: Indiana University Press, 1992).
- 18 Michel Foucault, *The Order of Things: An Archaeology of the Human Sciences* (London: Routledge, 1989).
- 19 Geoffrey Broadbent, et al., *Signs, Symbols, and Architecture* (Chichester, England: John Wiley and Sons, 1980).
- 20 Umberto Eco. 'Function and Sign: The Semiotics of Architecture', in *Signs, Symbols, and Architecture*, eds G. Broadbent, R. Bunt and C. Jencks (Chichester: John Wiley and Sons, 1980), 11–69.
- 21 Donald Preziosi, *Architecture, Language and Meaning: Origins of the Built World and Its Semiotic Organization* (The Hague, The Netherlands: Mouton, 1979).
- 22 Charles Jencks and George Baird (eds), *Meaning in Architecture* (London: Barrie & Rockliff, 1969). Many of the theories advanced there pertained to a school of thought within semiotics known as Structuralism. In the writing of the cultural theorist Roland Barthes, Structuralism admits art, fashion, architecture, sport and culture generally as forms of language, Roland Barthes, *Mythologies* (London: Paladin, 1973). The anthropologist Claude Lévi-Strauss developed Structuralism as a way of understanding culture and society through a series of all-pervasive oppositions, such as the way societies would organise themselves and their mythic structures along the lines of superior-inferior, hot and cold, raw and cooked, a kind of agonistic thinking that has implications for understanding architecture and urbanism. See Claude Lévi-Strauss, *Structural Anthropology 1* (London: Penguin, 1963).
- 23 Geoffrey Broadbent. 'Building Design as an Iconic Sign System', in *Signs, Symbols, and Architecture*, eds G. Broadbent, R. Bunt and C. Jencks (Chichester: John Wiley and Sons, 1980), 311–31. In Richard Coyne, *Peirce for Architects* (London: Routledge, 2019), I provide a detailed account of Peirce's rich theories about sign systems pertinent to architecture and urbanism.
- 24 Jeffrey Kipnis and Thomas Leiser (eds), *Chora L Works: Jacques Derrida and Peter Eisenman* (New York: Monacelli Press, 1997); Richard Coyne, *Derrida for Architects* (Abingdon, Oxon: Routledge, 2011).
- 25 Michel Foucault, *The Order of Things: An Archaeology of the Human Sciences* (New York: Random House, 1970).

- 26 George Orwell, *1984* (Oxford: Clarendon Press, 1984).
- 27 Martin Heidegger, 'Letter on Humanism', in *Martin Heidegger: Basic Writings*, ed. D. F. Krell (London: Routledge, 1978), 217–65, 217.
- 28 Martin Heidegger, 'Building, Dwelling, Thinking', in *Poetry, Language, Thought* (New York: Harper and Rowe, 1971), 143–61.
- 29 Martin Heidegger, 'Building, Dwelling, Thinking', in *Poetry, Language, Thought* (New York: Harper and Rowe, 1971), 143–61.
- 30 For example, Chapter 4 in the Old Testament book of Jonah, and Jeremiah's laments (Jeremiah 12; 15; 20).
- 31 Sherry Turkle, *Reclaiming Conversation : The Power of Talk in a Digital Age* (New York: Penguin, 2015).
- 32 For a book that captures the ordinariness of conversation in the context of digital communications see Mizuko Ito, et al., *Personal, Portable, Pedestrian: Mobile Phones in Japanese Life* (Cambridge, MA: MIT Press, 2006).
- 33 Oxford English Dictionary entry for "converse." <https://doi.org/10.1093/OED/6445237730>.
- 34 For a development of ideas of community, dialogue and interpretation in the context of architecture see Adrian Snodgrass and Richard Coyne, *Interpretation in Architecture: Design as a Way of Thinking* (London: Routledge, 2006).
- 35 Lewis Mumford, *The City in History: Its Origins, Its Transformations, and Its Prospects*, (Middlesex: Penguin, 1961), 648.
- 36 Lewis Mumford, *Technics and Civilization* (London: Routledge, 1934), 79.
- 37 Pollio Vitruvius, *Vitruvius: The Ten Books on Architecture* (New York: Dover Publications, 1960), 38.
- 38 Pollio Vitruvius, *Vitruvius: The Ten Books on Architecture* (New York: Dover Publications, 1960), 38.
- 39 Marcel Merleau-Ponty, *Phenomenology of Perception* (London: Routledge and Kegan Paul, 1962); Jonathan Hale, *Merleau-Ponty for Architects* (London: Routledge, 2016).
- 40 Yuval Noah Harari, *Sapiens: A Brief History of Humankind* (London: Penguin, 2011), 26.
- 41 Yuval Noah Harari, *Sapiens: A Brief History of Humankind* (London: Penguin, 2011), 27.
- 42 Yuval Noah Harari, *Sapiens: A Brief History of Humankind* (London: Penguin, 2011).
- 43 Samuel McCormick, *The Chattering Mind: A Conceptual History of Everyday Talk* (Chicago, IL: The University of Chicago Press, 2020).
- 44 Samuel McCormick, *The Chattering Mind: A Conceptual History of Everyday Talk* (Chicago, IL: The University of Chicago Press, 2020), 9.
- 45 Samuel McCormick, *The Chattering Mind: A Conceptual History of Everyday Talk* (Chicago, IL: The University of Chicago Press, 2020), 6.
- 46 Samuel McCormick, *The Chattering Mind: A Conceptual History of Everyday Talk* (Chicago, IL: The University of Chicago Press, 2020), 257.
- 47 Marshall McLuhan, *Understanding Media, The Extensions of Man* (Cambridge, MA: MIT Press, 1994); Marshall McLuhan and Bruce R. Powers, *The Global Village: Transformations in World Life and Media in the 21st Century* (New York: Oxford University Press, 1989); Marshall McLuhan, *The Gutenberg Galaxy: The Making of Typographic Man* (Toronto: University of Toronto Press, 1962).
- 48 Marshall McLuhan, *The Medium is the Massage (long playing record)* (CS 9501, CL2701: Columbia Records, 1968). The album was created by Marshall McLuhan, Quentin Fiore and Jerome Agel, with production by John Simon.

- 49 Marshall McLuhan, *The Gutenberg Galaxy: The Making of Typographic Man* (Toronto: University of Toronto Press, 1962), 47.
- 50 Phil Allmendinger, 'Applying Lacanian Insight And A Dash Of Derridean Deconstruction To Planning's "Dark Side"', *Planning Theory*, 4, no. 1 (2005): 87–112, 104. For other national contexts see Luis Hernando Lozano Paredes and Gabriela Quintana Vigiola, 'Microspheres of Self-Governance: Platform Communities in Times of Need in Bogotá, Colombia', *Cities*, 153, (2024): 1–12.
- 51 Jürgen Habermas, *The Theory of Communicative Action* (Cambridge: Polity Press, 1987).
- 52 Phil Allmendinger, 'Towards a Post-Positivist Typology of Planning Theory', *Planning Theory*, 1, no. 1 (2002): 77–99, 88.
- 53 Jürgen Habermas, *The Theory of Communicative Action* (Cambridge: Polity Press, 1987), 16.
- 54 Wendy Pullan. 'Agon in Urban Conflict: Some Possibilities', in *Phenomenologies of the City: Studies in the History and Philosophy of Architecture*, ed. H. Steiner and M. Sternberg (Farnham: Ashgate, 2015), 213–24.
- 55 For example, see Stanley Fish, *Doing What Comes Naturally: Change, Rhetoric, and the Practice of Theory in Literary and Legal Studies* (Durham, SC: Duke University Press, 1989).
- 56 Lieven Ameel, 'Towards a Narrative Typology of Urban Planning Narratives for, in and of Planning in Jätkäsaari, Helsinki', *Urban Design International*, 22, no. 4 (2017): 318–30, 319.
- 57 Lieven Ameel, 'Towards a Narrative Typology of Urban Planning Narratives for, in and of Planning in Jätkäsaari, Helsinki', *Urban Design International*, 22, no. 4 (2017): 318–30, 319.
- 58 J. Throgmorton, 'Planning as Persuasive Storytelling in a Global-Scale Web of Relationships', *Planning Theory*, 2, no. 2 (2003): 125–51, 127.
- 59 J. Throgmorton, 'Planning as Persuasive Storytelling in a Global-Scale Web of Relationships', *Planning Theory*, 2, no. 2 (2003): 125–51, 127.
- 60 J. Throgmorton, 'Planning as Persuasive Storytelling in a Global-Scale Web of Relationships', *Planning Theory*, 2, no. 2 (2003): 125–51, 127.
- 61 J. Throgmorton, 'Planning as Persuasive Storytelling in a Global-Scale Web of Relationships', *Planning Theory*, 2, no. 2 (2003): 125–51, 128.
- 62 J. Throgmorton, 'Planning as Persuasive Storytelling in a Global-Scale Web of Relationships', *Planning Theory*, 2, no. 2 (2003): 125–51, 131.
- 63 J. Throgmorton, 'Planning as Persuasive Storytelling in a Global-Scale Web of Relationships', *Planning Theory*, 2, no. 2 (2003): 125–51, 131.
- 64 Adrian Snodgrass and Richard Coyne, *Interpretation in Architecture: Design as a Way of Thinking* (London: Routledge, 2006).
- 65 Terry Winograd, PhD thesis, *Procedures as a Representation for Data in a Computer Program for Understanding Natural Language* (Cambridge, MA: MIT, 1971); and Terry Winograd, *Understanding Natural Language* (Edinburgh: Edinburgh University Press, 1972).
- 66 As explained in Chapter 4, fine tuning that favours particular styles of response plays a role as well.
- 67 Note the release in 2024 of ChatGPT o1 that appears to "deliberate" before delivering its response to a prompt.
- 68 Marshall McLuhan and Bruce R. Powers, *The Global Village: Transformations in World Life and Media in the 21st Century* (New York: Oxford University Press, 1989). I have

- developed these ideas in Richard Coyne, *Derrida for Architects* (Abingdon, Oxon: Routledge, 2011); Richard Coyne, *Peirce for Architects* (London: Routledge, 2019).
- 69 My main evidence for this observation about limits to voice training comes from an absence of research literature on the subject. See Michael McTear and Marina Ashurkina, *Transforming Conversational AI: Exploring the Power of Large Language Models in Interactive Conversational Agents* (New York: Apress, 2024).
- 70 Jacques Derrida, *Of Grammatology* (Baltimore, Maryland: Johns Hopkins University Press, 1976).
- 71 Andy Clark. 'Embodied, Embedded, and Extended Cognition', in *The Cambridge Handbook of Cognitive Science*, eds K. Frankish and W. Ramsey (Cambridge: Cambridge University Press, 2012), 276–91.

References

- Allmendinger, P. (2002), 'Towards a Post-Positivist Typology of Planning Theory', *Planning Theory*, 1 (1): 77–99.
- Allmendinger, P. (2005), 'Applying Lacanian Insight and a Dash of Derridean Deconstruction to Planning's "Dark Side"', *Planning Theory*, 4 (1): 87–112.
- Ameel, L. (2017), 'Towards a Narrative Typology of Urban Planning Narratives for, and of Planning in Jätkäsaari, Helsinki', *Urban Design International*, 22 (4): 318–30.
- Anderson, D. (2015), *Imaginary Cities*, London: Influx Press.
- Baker, V. R. (1999), 'Geosemiosis', *Bulletin of the Geological Society of America*, 111 (5): 633–45.
- Barthes, R. (1973), *Mythologies*, trans. A. Lavers, London: Paladin.
- Broadbent, G. (1980), 'Building Design as an Iconic Sign System', in G. Broadbent, R. Bunt and C. Jencks (eds), *Signs, Symbols, and Architecture*: 311–31, Chichester: John Wiley and Sons.
- Broadbent, G., R. Bunt and C. Jencks. (1980), *Signs, Symbols, and Architecture*, Chichester: John Wiley and Sons.
- Celis, C. and P. Ortuzar Kunstmann. (2021), 'We Have Always Been Artificially Intelligent: An Interview with Joanna Zylińska – Claudio Celis and Pablo Ortuzar Kunstmann'. *Culture Machine*. Available online: <https://culturemachine.net/vol-20-machine-intelligences/an-interview-with-joanna-zylińska-claudio-celis-and-pablo-ortuzar-kunstmann/> (accessed 27 May 2023).
- Clark, A. (2012), 'Embodied, Embedded, and Extended Cognition', in K. Frankish and W. Ramsey (eds), *The Cambridge Handbook of Cognitive Science*: 276–91, Cambridge: Cambridge University Press.
- Coyne, R. (2010), *The Tuning of Place: Sociable Spaces and Pervasive Digital Media*, Cambridge, MA: MIT Press.
- Coyne, R. (2011), *Derrida for Architects*, Abingdon, Oxon: Routledge.
- Coyne, R. (2018), *Network Nature: The Place of Nature in the Digital Age*, London: Bloomsbury Academic.
- Coyne, R. (2019), *Peirce for Architects*, London: Routledge.
- Coyne, R. (2023), *Cryptographic City: Decoding the Smart Metropolis*, Cambridge, MA: MIT Press.
- Derrida, J. (1976), *Of Grammatology*, trans. G. C. Spivak, Baltimore, MD: Johns Hopkins University Press.
- Dick, P. K. (1996), *Do Androids Dream of Electric Sheep?*, London: Voyager.

- Dick, P. K. (2000), *We Can Remember It for You Wholesale*, London: Millennium.
- Eco, U. (1980), 'Function and Sign: The Semiotics of Architecture', in G. Broadbent, R. Bunt and C. Jencks (eds), *Signs, Symbols, and Architecture*: 11–69, Chichester: John Wiley and Sons.
- Fish, S. (1989), *Doing What Comes Naturally: Change, Rhetoric, and the Practice of Theory in Literary and Legal Studies*, Durham, SC: Duke University Press.
- Foucault, M. (1970), *The Order of Things: An Archaeology of the Human Sciences*, New York: Random House.
- Foucault, M. (1989), *The Order of Things: An Archaeology of the Human Sciences*, London: Routledge.
- Ghaffarianhoseini, A., et al. (2016), 'What is an Intelligent Building? Analysis of Recent Interpretations from an International Perspective', *Architectural Science Review*, 59 (5): 338–57.
- Gibson, W. (1995), *Neuromancer*, London: Harper Voyager.
- Habermas, J. (1987), *The Theory of Communicative Action*, trans. T. McCarthy, Cambridge: Polity Press.
- Hale, J. (2016), *Merleau-Ponty for Architects*, London: Routledge.
- Harari, Y. N. (2011), *Sapiens: A Brief History of Humankind*, London: Penguin.
- Heidegger, M. (1971), 'Building, Dwelling, Thinking', *Poetry, Language, Thought*: 143–61, New York: Harper and Rowe.
- Heidegger, M. (1978), 'Letter on Humanism', in D. F. Krell (ed.), *Martin Heidegger: Basic Writings*: 217–65, London: Routledge.
- Ito, M., D. Okabe and M. Matsuda (eds). (2006), *Personal, Portable, Pedestrian: Mobile Phones in Japanese Life*, Cambridge, MA: MIT Press.
- Jencks, C. and G. Baird (eds). (1969), *Meaning in Architecture*, London: Barrie & Rockliff.
- Joss, S., et al. (2019), 'The Smart City as Global Discourse: Storylines and Critical Junctures across 27 Cities', *Journal of Urban Technology*, 26 (1): 3–34.
- Kipnis, J. and T. Leiser (eds). (1997), *Chora L Works: Jacques Derrida and Peter Eisenman*, New York: Monacelli Press.
- Kitchin, R. (2014), 'The Real-Time City? Big Data and Smart Urbanism', *GeoJournal*, 79: 1–14.
- Lévi-Strauss, C. (1963), *Structural Anthropology 1*, London: Penguin.
- Mattern, S. (2021), *A City Is Not a Computer: Other Urban Intelligences*, Princeton, NJ: Princeton University Press.
- McCormick, S. (2020), *The Chattering Mind: A Conceptual History of Everyday Talk*, Chicago, IL: The University of Chicago Press.
- McLuhan, M. (1962), *The Gutenberg Galaxy: The Making of Typographic Man*, Toronto: University of Toronto Press.
- McLuhan, M. (1968), *The Medium Is the Massage (Long Playing Record)*, CS 9501, CL2701: Columbia Records.
- McLuhan, M. (1994), *Understanding Media, the Extensions of Man*, Cambridge, MA: MIT Press.
- McLuhan, M. and B. R. Powers. (1989), *The Global Village: Transformations in World Life and Media in the 21st Century*, New York: Oxford University Press.
- McTear, M. and M. Ashurkina. (2024), *Transforming Conversational AI: Exploring the Power of Large Language Models in Interactive Conversational Agents*, New York: Apress.
- Merleau-Ponty, M. (1962), *Phenomenology of Perception*, London: Routledge and Kegan Paul.
- Mumford, L. (1934), *Technics and Civilization*, London: Routledge.

- Mumford, L. (1961), *The City in History: Its Origins, Its Transformations, and Its Prospects*, Middlesex: Penguin.
- Nöth, W. (1990), *Handbook of Semiotics*, Bloomington, IN: Indiana University Press.
- Ong, W. J. (2002), *Orality and Literacy: The Technologizing of the Word*, London: Routledge.
- Orwell, G. (1984), 1984, Oxford: Clarendon Press.
- Paredes, L. H. L. and G. Q. Vigiola. (2024), 'Microspheres of Self-Governance: Platform Communities in Times of Need in Bogotá, Colombia', *Cities*, 153: 1–12.
- Peirce, C. S. (1992), *The Essential Peirce, Selected Philosophical Writings Volume 1 (1867–1893)*, Bloomington, IN: Indiana University Press.
- Preziosi, D. (1979), *Architecture, Language and Meaning: Origins of the Built World and Its Semiotic Organization*, The Hague, The Netherlands: Mouton.
- Pullan, W. (2015), 'Agon in Urban Conflict: Some Possibilities', in H. Steiner and M. Sternberg (eds), *Phenomenologies of the City: Studies in the History and Philosophy of Architecture: 213–24*, Farnham: Ashgate.
- Sebeok, T. A. (1999), *Signs: An Introduction to Semiotics*, Toronto: University of Toronto Press.
- Smithers, T. (2024), 'Generative AI: Useful Tools or Expensive Toys', *Teaching and Strengthening Research*, March 2024.
- Snodgrass, A. and R. Coyne. (2006), *Interpretation in Architecture: Design as a Way of Thinking*, London: Routledge.
- Stiegler, B. (1998), *Technics and Time*, trans. R. Beardsworth and G. Collins, Stanford, CA: Stanford University Press.
- Throgmorton, J. (2003), 'Planning as Persuasive Storytelling in a Global-Scale Web of Relationships', *Planning Theory*, 2 (2): 125–51.
- Turkle, S. (2015), *Reclaiming Conversation: The Power of Talk in a Digital Age*, New York: Penguin.
- Vitruvius, P. (1960), *Vitruvius: The Ten Books on Architecture*, trans. M. H. Morgan, New York: Dover Publications.
- Weizenbaum, J. (1976), *Computer Power and Human Reason: From Judgment to Calculation*, San Francisco: W. H. Freeman.
- Winograd, T. (1971), *Procedures as a Representation for Data in a Computer Program for Understanding Natural Language*, PhD Thesis, Cambridge, MA: MIT.
- Winograd, T. (1972), *Understanding Natural Language*, Edinburgh: Edinburgh University Press.
- Zylinska, J. (2020), *AI Art: Machine Visions and Warped Dreams*, London: Open Universities Press.

2

SCRIPTING THE CITY

Anyone who has deployed video conferencing in professional and learning contexts will appreciate text-to-speech functionality that delivers text records of spoken conversations between human participants. Participants appear as actors in these transcripts, an aspect of conversational AI fostering further linkages with urban living, as I will show in this chapter. Think of a transcript as a written conversation, or a script.

LLMs are good at inventing stories, thanks to the narrative patterns within the vast bodies of literature on which they are trained and fine tuned. LLMs will also compose scripts, as in a play. In response to suitable prompts, LLMs will draw on patterns acquired from training on vast numbers of texts of all kinds to define performer roles, describe settings and construct dialogues on request according to style, even constructing unseen excerpts from known media presentations, e.g. dramas, soap operas and situation comedies.¹

LLMs can sketch out the plot for a play, provide some “narrative arcs” or fill in detail, at least as a helpful scripting assistant. LLMs seem to be good at



Figure 2.1 Market Street, Melbourne stimulating a conversation about the challenge of reflected sunlight.

Source: Author.

simulating dialogue, as if a script. Here is an example. I have demonstrated this capability to create dialogue in an urban context by prompting ChatGPT with a scenario in which two individuals with opposing views discuss the effects of glare from reflected sunlight in city streets (Figure 2.1).

ChatGPT scripted a convincing short play. It set the scene: “A café in the central business district of Melbourne.” It made up character names, Martha and Mark who sit at a table by the window. The outlook is a street lined with tall glass buildings. A conversation ensues:

Martha: (sipping her coffee) You know, Mark, every time I walk through the city, I can’t help but marvel at how the sunlight reflects off these glass buildings. It adds such a dynamic quality to the urban landscape, don’t you think?

Mark: (nodding thoughtfully) I see what you mean, Martha. It does create an interesting visual effect. But have you considered the downsides? The reflected sunlight can cause significant glare, which can be quite uncomfortable for pedestrians and drivers.

The conversation continues as they each articulated their position and seek common ground. Martha concludes affirmingly, “Let’s take these ideas forward and contribute to a more vibrant and liveable Melbourne.”

With further prompting from me I could have induced a more entertaining dialogue, but it is recognisable as a script. Of note, as with any script, there is much left undescribed and unsaid. It assumes an audience would know about the practice of walking through the city, that driving involves risks, that sometimes skies are clear. The script also follows a familiar pattern involving interlocutors, a setting, a focal point for the discussion, some conflict, and a statement of resolution. As I will show in the next section, you could say that this simple dialogue illustrates that scripts and their entailments are somehow embedded in the way people think and act.

Scripts and cognition

The concept of the script has currency in cognitive science and AI. Theatrical-style scripts encompass the predictable and patterned ways in which individuals speak and interact in urban settings. Think of routine verbal exchanges as when introducing friends to one another, or how passengers offer their seat for a struggling fellow passenger, purchase a ticket or discuss the weather. There is a practical rationale for such norms of behaviour. Once learned, scripted behaviour patterns facilitate social order and productive interactions.

The concept of scripts in urban contexts aligns with the sociologist Erving Goffman’s idea of *dramaturgy*,² where social life is likened to a theatre, with individuals playing roles and following scripts to manage how we appear to one another and how we interact. Arguably, scripts guide interactions in the home, markets, shops, workplaces, and online establishing trust and predictability. They guide behaviour in public spaces, helping people navigate everyday and difficult social environments. Such challenges are brought into relief when we travel to places where the customs are different to our own. In fact, urban scripts contribute to the cultural fabric of a

city, as they reflect and reinforce local customs, traditions and norms. Once learned, such scripts also help newcomers adapt to the urban environment and foster a sense of community and belonging. Understanding and following these scripts can help in resolving conflicts and misunderstandings. When people adhere to common social scripts, it reduces the likelihood of friction. Planners and designers can take these social scripts into account when designing public spaces to facilitate certain types of interactions. For example, the layout of a plaza might encourage social gatherings, while signage and pathways in transport hubs facilitate commuter flow.

Cognitive scientists Roger Schank and Robert Abelson developed a theory of scripts as vital elements in human cognition:

A script is a structure that describes appropriate sequences of events in a particular context. A script is made up of slots and requirements about what can fill those slots. The structure is an interconnected whole, and what is in one slot affects what can be in another. Scripts handle stylized everyday situations.³

Their oft repeated example is the “restaurant script.” There are standard aspects to the setting of such social interactions such as food, service and personnel (customers, waiters, kitchen staff). The term “script” is appropriate as there is a typical sequence to one’s experience as a customer at a restaurant, a bit like acting out stage directions, and there are typical conversational exchanges. Amongst other virtues, scripts enable us to be economical with words. If I started to tell a story with, “The waiter gave me the menu and I ordered a meal,” then we have a clue as to what might happen next. As soon as you recognise that you are in a restaurant or talking about a restaurant then the listener invokes the “restaurant script.” If you are telling a story that involves a restaurant visit, then there is no need to explain the detail before the listener gets the gist. A script-based model of language operates as a look-up process. It is as if a few key words (waiter, menu, food) enable you to access a pre-stored script that includes tables and chairs, other customers, rituals of ordering and paying the bill.

Schank and Abelson developed a script-based digital platform they called SAM (Script Applier Mechanism), which they describe with many examples in their book titled *Scripts, Plans, Goals and Understanding: An Inquiry into Human Knowledge Structures*.⁴ Applied to the script about urban glare, their method of automation would have to accommodate the setting, such

as two professionals talking in a café, and the content of their discussion pertaining to glare in an urban setting. A helpful article by Herbert Grabes summarises the notation deployed in SAM. He lists the typology of basic actions, which gives insight into its method for dealing with complicated topics:

ATRANS (transfer of an abstract relationship), PTRANS (transfer of the physical location), PROPEL (application of a physical force to an object), MOVE (involvement of a body part of an animal by that animal), GRASP (the grasping of an object), INGEST (the taking in of an object), EXPEL (the expulsion of an object), MTRANS (the transfer of mental information), MBUILD (the construction of new information), SPEAK (actions of producing sounds), and ATTEND (attending or focusing a sense organ towards a stimulus).⁵

There is more to the terminology and methods for defining and accessing scripts, but the platform would need to first identify such actions as part of a general script to accommodate the scenario I related above about two people discussing the challenge of urban glare.

Whatever its virtues, to translate this process into an automated language model is cumbersome. Examining the rudiments of SAM, you would need to store a prohibitively large number of scripts and their variants for it to replicate the linguistic performance of two knowledgeable human interlocutors, or even a neural network-based LLM. It would need to accommodate changes in practices and technologies, such as taking orders and photosharing by smartphone.

Contrary to the way SAM works, it is significant that an LLM telling a story about being in a restaurant does not need pre-stored scripts. An LLM draws on language patterns derived from its vast training corpus. With the appropriate prompts it will produce texts that have the appearance of hybridised and imaginative scripts, having no difficulty in generating a script that continues the narrative.

To illustrate the ability of an LLM to defy the need for standard script scenarios, I prompted ChatGPT with an unlikely plot premise: “The waiter gave me his manifesto and then ran up the fire escape.” Without further prompts the platform offered me a 4-point plan for how to deal with such an unusual situation: “If the behavior seemed threatening or erratic, consider alerting security or law enforcement.” That was not what I wanted.

When I told ChatGPT that my initial prompt was the first line of a short story, it continued with the imaginative narrative:

I sat there, stunned, with the crumpled papers in my hand, the scent of ink and desperation lingering in the air. Around me, the restaurant buzzed with oblivious patrons, their chatter a stark contrast to the whirlwind of thoughts in my head.

It continued with further sentences that developed the story.

The idea of a pre-stored script does not really capture the fluidity of this kind of interaction. In script terms the LLM seemed to follow a security script, followed by a mystery story script or perhaps a crime drama, and would have produced any number of others had I explored further. Though steeped in a pre-LLM symbolic AI paradigm, the main lesson from Schank and Abelson's work from the 1980s is to recognise the importance of scripts in language, cognition and in the urban context, as a means of understanding restaurants, construction processes, design development and political dialogue. I feel that Schank and Abelson's early AI work on scripts entitles us to advance AI script writing as an important element of AI in the city. Scripts are not only the preserve of comedic and dramatical performances but are at some level part of our cognitive armoury. Scripts permeate our conversations and interactions. Shakespeare wrote, "All the world's a stage, And all the men and women merely players."⁶

Hypermediated conversation

It was appropriate in Chapter 1 to frame language initially in terms of chatter and gossip. The scripted conversation between Martha and Mark presented earlier, delivers a polite role play between professionals. I could have prompted the platform to deliver a more likely exchange, interspersed with hesitations, quips, asides and personal expressions of bravado or self-doubt, not to mention in situ web searches, opinions about personalities and relationships, and even the politics of city development.

Professional role plays provide one model of urban conversations. Scripts advanced in other mediated contexts such as films, plays, video games, radio and tv chat shows, soap operas and situation comedies (sitcoms) also provide models of urban conversation. They follow scripts to present

ordinary everyday experiences but present them in an unusual light. Cities provide ideal backdrops, themes and content for characters and scripted storylines.

In this media saturated age, it is fair to say that scripts as delivered on television, streaming services and short videos also inform everyday scripts. I am prepared to state that as well as mirroring aspects of the way people talk to one another, much of what we regard as a conversation is filtered through exposure to such scripted conversations in media entertainment. I recall in my student days a fellow college resident who enjoyed the *M*A*S*H* (1972–1983) series to such a degree that it inflected his conversations in the communal dining room. If we didn't get one of his jokes or an obscure reference, he would remind us that it was something Hawkeye or Trapper said or would have said. Call such references in everyday conversation to entertainment media "hypermediated" conversation.⁷ We mimic the language of media, which in turn mirrors our own language, in a kind of loop, producing the well-known echo-chamber effect that selectively amplifies certain extreme opinions and behaviour patterns while diminishing others.

We viewers and listeners are also prone to manipulation through these media. Consider the persistent appeal of situation sitcoms as an example.⁸ They lend themselves to family viewing, and advertisers like them. Not only do advertisers attempt to maximise audience exposure to consumer products, but most sitcoms establish up-beat life-style models that reinforce consumption patterns and seek to raise the aspirations of audiences to yet more consumption. Critics such as Nick Marx among others calls it "sitcom capitalism."⁹ Some would argue that such television offerings are also responsible for numbing adolescent minds. That is a common lament about the decline in sociability rehearsed by writers such as Richard Louv and Sherry Turkle.¹⁰ Arguably, the influences of video gaming, live streaming and individualised media consumption now surpass the impacts of early television programming for many, but sitcoms persist as exemplars of such influential hypermediated conversation.

Of significance in the light of urban conversations, the sitcom writer must meet the challenge of conversational realism, as if more has been said than appears in the script. The kind of framing suggested by Schank's scripting theory accounts for this human ability to short circuit the need for extraneous dialogue. Audiences know the situation and the circumstances and fill in the gaps. Sitcom dialogue also needs to appear spontaneous, off the cuff

and as if it is part of everyday chatter and banter. I need hardly state that the way conversation is presented in sitcoms, as well as literature, drama and other narrative forms bears little resemblance to how people speak to one another in everyday life. Unlike everyday conversation, sitcom humour is memorised, rehearsed, directed and delivered as short dialogic episodes within the one programme, a bit like an assembly of short blog posts, twitter threads or text message exchanges. Most of us learn that it is not possible to sustain a rapid to-and-fro banter with sitcom-style throughlines, punchlines and one-liners, especially if fellow interlocutors don't play along. But it is important that sitcom banter conveys the impression of everyday speech.¹¹

In everyday life outside of the sitcom universe, most profound, difficult, witty, inept, misplaced or faltering statements in everyday conversation mostly go unnoticed, are ignored, dismissed or otherwise cut short. This is less the case in the sitcom. Every word contributes to the plot, mood, understanding of characters or a joke. That to me is one of the distinctive features of sitcom humour – the extent to which characters attend to what each other is saying. I will allude to the issue of attention as one of the core functions of LLMs in the next chapter.

Conversational realism

Sitcom characters are scripted as if familiar with each other's conversational quirks and behaviour patterns. They draw on each other's "scripts," in Schank's terms. Thanks to the serial nature of the sitcom format, audiences are onboard with that familiarity. Not all dramatic or comedic presentations are like that. Some writers toy with social and conversational realism and lapses in attention within everyday language. Scripts by the playwright Harold Pinter (1930–2008) are more like the way people actually speak to one another. Sometimes this style of scripted dialogue is known as "Pinteresque" in acknowledgement of his promotion of particularly stilted dialogue. According to an article by Alice Rayner on Pinter:

In Harold Pinter's work, the infamous pauses, excruciating silences, and the proclivity for tableaux are instances of delay when the forward motion of events is held and something unspoken happens. Such silences create atmosphere and mood, to be sure, and they may indicate something about character, but they are also part of a signifying structure.¹²

Pinter's later works reached larger audiences through movie adaptations such as *The Go-Between* (1971), *The French Lieutenant's Woman* (1981), and *Sleuth* (1972, 2007). But his early work informed a style of comedy writing that morphed into the "cringe" sitcom, described by Jonathan Bignell as an influence on the works of Ray Galton and Alan Simpson, writers of the British sitcoms *Hancock's Half Hour* (1956–1960) and *Steptoe and Son* (2011).¹³ Bignell explains that "Cringe can be defined in a range of ways and has a spectrum of related instantiations, including comedy of embarrassment, gross-out comedy, and awkward comedy."¹⁴ More recently, this style of sitcom is often presented on television without a live audience or a laughter track and is frequently delivered as if a documentary with actors speaking or directing facial expressions of frustration or embarrassment directly to the camera – as a "mockumentary." Bignell attributes cringe comedy explicitly to Harold Pinter: this "strand of dark, troubling but also comic drama on the theatre stage, spearheaded by the playwright Harold Pinter, laid the foundations for cringe." Instead of laughter, the punchline to a cringe or Pinteresque joke is met with a glance of incomprehension by the other participants. Incomprehension is in the company of inattention. The media landscape of course expands to include short videos capturing movie outtakes, reaction videos, amateur videos and clips from mainstream media inflating the distribution of embarrassing and cringeworthy moments.

Cringe realism is closer to everyday conversation than the formal role play I presented earlier in the conversation about urban glare. You cannot assume such attention to nuance, and engagement amongst the characters, or that they are prepared to play along with each other in that way. Contemporary cringe comedic exchanges, like much everyday conversation, are profiles in inattention. At times people are preoccupied, not always paying attention to what the other is saying, or even their own previous utterances. People can speak at cross-purposes.

Emphasis and attention

Understanding and misunderstandings in conversation are often attributable to where people place emphasis as they speak and listen. People in conversation may tailor their responses depending on where the other interlocutor places emphasis. Consider a fragment of the dialogue about urban glare. Martha said, "You know, Mark, every time I walk through the city, I can't

help but marvel at how the sunlight reflects off these glass buildings.” How might Mark have responded had Martha placed emphasis on different words in that sentence, by vocal emphasis, pauses or gestures? Had the emphasis resided with the word “walk” then Mark may have responded “It’s true, walking through the city really gives you a different perspective on how the reflections play out compared to driving.” Had she emphasised “marvel”, then the response could have been “I understand why you marvel at it; the way the light dances on the glass can be mesmerising.” Had she emphasised “these,” Mark might have replied “Yes, these specific buildings do have a particularly striking way of catching the light.” Where the speaker places emphasis can indeed change meaning and the course of a conversation.

Emphasis draws attention to different parts of a sentence. People in conversation deploy various means of managing attention and to greater or lesser degrees across a sentence or whole block of dialogue. As indicated in Mark’s alternative responses, directing attention inflects the meaning of a sentence. There is a distribution of major and lesser attention points across a block of dialogue. Alternative attention distributions may alter the course of a conversation. Such attention distributions are not just on the side of the speaker. Listeners may infer distributions of attention of their own.

How is attention managed in the case of conversational exchanges conducted by text only, as in the case of dialogue within chat rooms, text messaging and microblogging (WhatsApp, Twitter, Bluesky)? Here the context of fragments of text assumes particular importance. As we shall see in the next chapter, LLMs construct alternative attention distributions for blocks of text based on the semantic relationships between words, and their positions within a block of text. Such attention mechanisms significantly affect the platform’s ability to predict convincingly what should come next in its generation of text responses and continue the conversation.

So, my digression into hypermediated conversation lands on the topic of attention, a major focus for developers of LLMs. Meanings change depending on where you and the interlocutor place their attention, and thereby conduct the course of the conversation.

Summary and critique

In this chapter, I explored the concept of scripting in urban contexts, both through human interactions and AI-generated dialogues. I demonstrated how chat platforms and virtual meetings produce transcripts, which are

important in capturing and transmitting conversations. I illustrated the capability of LLMs to generate convincing dialogue with a scripted conversation. This led to a discussion on the broader significance of scripts in cognitive science and AI, particularly drawing on the work of Schank and Abelson. Their theory of scripts emphasised predictable patterns of behaviour in social contexts, highlighting how these scripts facilitate productive interactions in urban environments.

I further examined how scripted behaviour aligns with Goffman's dramaturgy, emphasising the role of scripts in managing social interactions and maintaining cultural norms in cities. I touched on the practical implications of understanding scripts for urban planners and designers, who can use this knowledge to create public spaces that support productive interactions.

The chapter also covered the hypermediated nature of contemporary conversation, influenced by media such as sitcoms and dramas. I argued that these scripted media presentations can shape our everyday conversations, creating a cycle of imitation and influence. Sitcoms, with their need for conversational realism, were examined to illustrate how we perceive and engage in dialogue.

Finally, I touched on the concept of conversational realism and the role of attention in dialogue. By analysing how emphasis and attention affect meaning, I set the stage for an exploration of attention mechanisms in LLMs in the next chapter. This understanding is crucial for appreciating how AI platforms predict and generate text, influencing the flow and outcome of conversations.

On prompting, ChatGPT furnished me with some points that may appear controversial to readers. AI-generated criticism focussed on my invocation of Schank and Abelson's scripting model that seeks to formalise what are very dynamic and human processes. LLMs offer a much more fluid approach constructing and responding to conversational input than their scripting model. The scripting model suggests that dialogues and action sequences exist as pre-formed cognitive entities. At best we could say it is as if people (and LLMs) are following scripts in their behaviour. Despite its limitations, the Schank and Abelson scripting model serves as a bridging concept. It demonstrates that AI research has some investment in the idea of the script. It is part of AI's legacy to which subsequent developments have responded. The idea of the script reinforces certain aspects of conversation in the city, including its mediated forms in literature, dramas and comedies delivered through televisions, streamed media, gaming and other digital channels.

My ChatGPT critic suggested that I could be overstating the idea that media, such as sitcoms and dramas, influence everyday conversations. Some might argue that people are capable of distinguishing between scripted entertainment and real-life interactions, maintaining more independence in their conversational styles. That scripts and narratives delivered through ubiquitous media might affect what people do and say is really a matter of degree, rather than robotic mimicry of what is delivered through the media. Yes, and concepts of human autonomy and agency provide for choice. We may choose to adopt jokes and ways of speaking taken from the media, or we may consciously resist them, but we easily forget the sources of our conversational patterns. As I will show in the next chapter, LLMs are not designed to reconstruct their training data as they respond to conversational input but infer patterns.

My AI critic challenged the claim that cringe comedy or Pinteresque dialogue is closer to real-life conversation. Some might find these forms of scripted dialogue too stylised or exaggerated to be considered realistic representations of everyday interactions. I would address this criticism by suggesting that the definitions of everyday dialogue are elusive anyway. Ritualised and stylised exchanges, as in greetings and departure rituals, are part of everyday discourse. I did not mean to imply that Pinteresque and cringe scripts are closer to how people actually speak. I was also keen to point out that dialogue through the media is often scripted to give the impression of realism, which involves exaggerating certain features and diminishing others.

My virtual critic identified as a challenge the discussion on how emphasis and attention affect meaning in text-based communication. Critics might argue that the absence of vocal cues and body language significantly alters the dynamics of textual dialogue. In response, I would point out that though the theme of attention is perhaps counterintuitive it is a major focus in the development of LLMs to be introduced in the next chapter. The argument is that as well as vocal and gestural emphasis, we language users can infer where the emphasis lies based on context. We draw on the local context of the words before us and the wider context of the conversation.

The ChatGPT critic faulted any implication that LLMs can generate scripts as fluidly as human-written scripts. Critics might argue that LLMs, despite their advanced capabilities, still produce outputs that lack depth, coherence or genuine creativity compared to human-authored content.¹⁵ That is a valid criticism of LLMs in general. My response is to advocate for greater understanding on the part of users of these systems. Knowing how they work can lead to their more judicious application.

Notes

- 1 In 2022, it was scriptwriters who led the protest against AI. A fiction writer or biographer might consult with colleagues or family members and call on the resources of the Internet: library materials, articles, search engines, spelling and grammar checkers, automated proof-reading, image repositories and even AI assistants. That may make the research and writing processes more efficient, speed things along and even enhance quality. The author of a speculative work examining the urban implications of AI might operate in a similar manner. We may suppose that other steps in the publication workflow are affected more immediately by AI automation, as are the roles of writers, editors and readers of reports, and writers who work in teams. Scriptwriters have recently led the charge against AI in filmmaking. Under threat are the rights of authors and designers whose creative outputs contribute to the training corpuses of AI systems, including writers, actors and others in the creative professions whose labours are supplanted by AI surrogates. An AI avatar on screen and an AI on a writing team can reduce the demand on human beings who draft, outline, sketch, generate, edit, improve, iterate, proof, review and evaluate scripts for films, documentaries, ads, domestic serial dramas (soaps) and episodic situation comedies (sitcoms).
- 2 Erving Goffman, *The Presentation of Self in Everyday Life* (London: Penguin, 1969).
- 3 Roger C. Schank and Robert P. Abelson, *Scripts, Plans, Goals and Understanding: An Inquiry into Human Knowledge Structures* (Hillsdale, NJ: Erlbaum, 1977), 41; Roger C. Schank, *Tell Me a Story: A New Look at Real and Artificial Memory* (New York: Macmillan, 1990); Roger C. Schank, 'Sci-fi with No Script', *Times Higher Education Supplement*, March 13 (1992): 15–15.
- 4 Roger C. Schank and Robert P. Abelson, *Scripts, Plans, Goals and Understanding: An Inquiry into Human Knowledge Structures* (Hillsdale, NJ: Erlbaum, 1977).
- 5 Herbert Grabes, 'Three Theories of Literary Worldmaking: Phenomenological (Roman Ingarden), Constructivist (Nelson Goodman), Cognitive Psychologist (Schank and Abelson)', *Cultural Ways of Worldmaking*, 1 (2010): 47–60, 54.
- 6 That is from *As You Like It*, Act II, Scene VII. Surprisingly the soliloquy mentions characters, but not scripts or conversations.
- 7 I am thinking of some of the terminology advanced in J. David Bolter and Richard A. Grusin, *Remediation: Understanding New Media* (Cambridge, MA: MIT Press, 1999).
- 8 Sitcoms are series of 20–30-minute tv programmes that incorporate recurring ensembles of characters. As explained by James Burrows, director of many commercially successful sitcoms, the average sitcom series involves ordinary, relatable characters, albeit with exaggerated traits, involved in distinctive but identifiable circumstances that provoke interesting interactions – four single women in their sixties live together (*The Golden Girls*), a gay man and straight woman are best friends (*Will and Grace*; *Gimme Gimme Gimme*), a wealthy radio psychiatrist lives with his father and his carer (*Frasier*), a fashion conscious and irresponsible mother lives with her straight-laced daughter (*Absolutely Fabulous*), two repressed priests live together on an isolated island with an obsessive housekeeper (*Father Ted*). Rendering the familiar as strange is a hallmark of much storytelling.
- 9 Nick Marx, 'Home Economics: Sitcom Capitalism, Conservative Comedy, and Media Conglomeration in Post-Network Television', *Communication, Culture and Critique*, 15, no. 1 (2021): 21–35, 21.

- 10 Richard Louv, *Last Child in the Woods: Saving Our Children from Nature-Deficit Disorder* (London: Atlantic Books, 2005); Sherry Turkle, *Alone Together: Why We Expect More from Technology and Less from Each Other* (New York: Basic Books, 2011).
- 11 You can find the scripts of many plays, including sitcoms, online. A twenty-two-minute episode of *Cheers* contains about 2,400 words.
- 12 Alice Rayner, 'Harold Pinter: Narrative and Presence', *Theatre Journal*, 40, no. 4 (1988): 482–97, 482.
- 13 Jonathan Bignell, 'Cringe histories: Harold Pinter and the Steptoes', *Humanities*, 10, no. 83 (2021): 1–14, 1. US examples include the American version of *The Office* (Daniels et al 2005–2013). I would add *Parks and Recreation* (Daniels and Schur, 2009–2015) and perhaps *Crazy Ex-Girlfriend* (2015–2019); and from Australia, *Kath and Kim* (2002–2005).
- 14 Jonathan Bignell, 'Cringe histories: Harold Pinter and the Steptoes', *Humanities*, 10, no. 83 (2021): 1–14, 1.
- 15 For an example of a short informational travel book that attributes its authorship to ChatGPT see Tony Roppon, *Kitami City: By ChatGPT* Kindle, (2023) and an academic article that claims ChatGPT as co-author by Siobhan O'Connor and ChatGPT, 'Open Artificial Intelligence Platforms in Nursing Education: Tools for Academic Progress or Abuse?', *Nurse Education in Practice*, 66, no. 11 (2023): 1–2.

References

- Bignell, J. (2021), 'Cringe Histories: Harold Pinter and the Steptoes', *Humanities*, 10 (83): 1–14.
- Bolter, J. D. and R. A. Grusin. (1999), *Remediation: Understanding New Media*, Cambridge, MA: MIT Press.
- Goffman, E. (1969), *The Presentation of Self in Everyday Life*, London: Penguin.
- Graves, H. (2010), 'Three Theories of Literary Worldmaking: Phenomenological (Roman Ingarden), Constructivist (Nelson Goodman), Cognitive Psychologist (Schank and Abelson)', *Cultural Ways of Worldmaking*, 1: 47–60.
- Louv, R. (2005), *Last Child in the Woods: Saving Our Children from Nature-Deficit Disorder*, London: Atlantic Books.
- Marx, N. (2021), 'Home Economics: Sitcom Capitalism, Conservative Comedy, and Media Conglomeration in Post-Network Television', *Communication, Culture and Critique*, 15 (1): 21–35.
- O'Connor, S. and ChatGPT. (2023), 'Open Artificial Intelligence Platforms in Nursing Education: Tools for Academic Progress or Abuse?', *Nurse Education in Practice*, 66 (11): 1–2.
- Rayner, A. (1988), 'Harold Pinter: Narrative and Presence', *Theatre Journal*, 40 (4): 482–97.
- Roppon, T. (2023), *Kitami City: By Chatgpt*, Kindle.
- Schank, R. C. (1990), *Tell Me a Story: A New Look at Real and Artificial Memory*, New York: Macmillan.
- Schank, R. C. (1992), 'Sci-Fi with No Script', *Times Higher Education Supplement*, March 13: 15–15.
- Schank, R. C. and R. P. Abelson. (1977), *Scripts, Plans, Goals and Understanding: An Inquiry into Human Knowledge Structures*, Hillsdale, NJ: Erlbaum.
- Turkle, S. (2011), *Alone Together: Why We Expect More from Technology and Less from Each Other*, New York: Basic Books.

3

CORE FUNCTIONS OF LLMs IN THE URBAN CONTEXT

So far, I have provided an account of cities as semiotic structures. Language features prominently in the ways cities are described, interpreted, managed, created and imagined. Language itself is a multifaceted phenomenon, with many themes relevant to the urban context, such as metaphor, grammar, pragmatics, semantics.¹ Rather than attempt a survey of the diverse theories about how language operates, I will turn to a simplified account of language function informed by my reading of and experience with the apparatus of AI-based large language models (LLMs), in particular the so-called “Transformer” model, as adopted by OpenAI’s ChatGPT model. This chapter represents our first foray in this book into the technicalities of large language models.

Transformers

The term “Transformer” has a SciFi ring to it, but it was promoted in the seminal paper by A. Vaswani and co-researchers entitled “Attention

DOI: 10.4324/9781003535751-4

This chapter has been made available under a CC-BY-NC-ND 4.0 license.

is All You Need”² to reflect their model’s process of transforming input sequences through a series of very large-scale matrix operations, specifically using mechanisms labelled “self-attention,” to be explored in Chapter 9. A matrix is a table of numbers or symbols. “Transformation” is a term used in mathematics to describe how matrices can be multiplied, added, or otherwise manipulated to produce new matrices, and thereby implement complex alterations to a data set. For example, in digital graphics, matrices are employed to transform a plan view of a 3D model of a house into a perspective view, or a series of matrix operations may radically distort a photographic image of a street scene into a washed-out surreal post-apocalyptic landscape.

In Transformer-based large language models (LLMs), a word or segment of a word is represented by a vector. A vector is simply a matrix with just one row or column of numbers. Positions on a map are typically presented as vectors, i.e. coordinates X and Y, or latitude and longitude. The map coordinates are a list of just two numbers. That is a two-dimensional vector. If altitude is considered, then that requires vectors with three numbers. So, a vector is an ordered list of numbers. The dimension of the vector is the number of numbers in the ordered list of coordinates. Vectors are common for those of us involved in computer graphics. I think of a vector as a line with a position and a direction. You can think of every map coordinate as the end point of a line shooting out from an origin point in the coordinate space. Vectors don’t have to be visualised spatially as lines or as points in two- and three-dimensions but can be of any dimension. In LLMs, they are typically of several hundred dimensions.

To process words or segments of words (tokens) a natural language processing system needs to convert them to numbers. Transformer-based large language models turn words into high-dimensional vectors, i.e. coordinates in a multidimensional vector space. In the context of natural language processing, these vectors are usually referred to as “encodings” or “embeddings.” The term “embedding” is convenient as it implies the act of placing or “embedding” a word into a uniform vector space, where you can compare the position, and hence the properties, of different words. To put it another way, a word embedding compresses all the syntactic and semantic information about a word into a fixed-size vector.

LLMs successively transform these vectors by incorporating attention mechanisms, positional encodings and other operations, to turn word

meanings and contexts into numerical computations. Fine-tuning operations are also applied to optimise the model for specific tasks. The scale and centrality of these transformation operations justify the use of the term “Transformer” to describe this approach. I will use the term in this book to distinguish this approach to natural language processing from earlier models such as Markov models and Hidden Markov Models (HMMs), which are based on probabilistic predictions, or n-gram models built on the frequencies of word sequences in a training corpus. More recent recurrent neural networks (RNNs) and convolutional neural networks (CNNs) process data sequentially or in fixed spatial hierarchies. Unlike these methods, Transformers can capture long-range dependencies and context in a flexible and scalable manner, enabling sophisticated and accurate language processing.³

Predictions and meanings

Large language models (LLMs) such as ChatGPT, Bard, Claude, Gemini and Llama 2 that use the Transformer method, are constructed around a series of mathematical functions that operate to simulate aspects of human language. An LLM predicts what is likely to come next following a sequence of words already provided as input prompts by a human user including the LLM’s own previous responses. The model takes prior exchanges into account within a context window and predicts what could come next as its contribution to the exchange. It then implements that prediction, presenting it as the next word or token in the conversational exchange.

As I have demonstrated already, the result is a plausible conversational exchange between operator and machine that carries the hallmarks of an intelligent conversation, notably the correct use of language, wide vocabulary, basic background knowledge and meaningful and contextually relevant flow. Framing the LLM process in terms of prediction effectively circumvents the necessity for us to define what is meant by meaning (or semantics) in LLMs. What people mean when they speak or write is supplanted by the question of what response their utterance invokes from others or from an LLM. Similarly, in so far as an LLM response means anything, meaning resides in the responses it invokes from its own operations or from those of its human interlocutor. I will revisit these concepts in Chapter 6.

Functions of LLMs

Large language models attempt to simulate human language competence via several core functions. An LLM will (i) identify and manipulate *patterns* in data, (ii) process patterns derived from a vast collection of existing data (in text-based LLMs this is a *corpus* of existing texts), (iii) divide the corpus into *tokens*, (iv) process very large mathematical vectors that capture and process *semantic* information based on word proximities, (v) define and operate within *context* windows, (vi) adjust semantic vectors to account for the *position* of tokens in the context window, (vii) process alternative distributions of *attention* (or emphasis) within texts, and (viii) recruit human operatives to tune the performance of the LLM. I will introduce these functions briefly here, pending more detailed explanation in subsequent chapters.

i. Patterns

We can approximate how human beings learn to use language as a process of absorbing, recalling and synthesising patterns in speech and writing. We pick up speech patterns from people around us through repetition and mimicry, but we also have the capacity to acquire generalised patterns as templates in the structure of how things are said and written. Infants learn via simple examples. But let's consider some adult-level sentences such as: a city is not a tree, a city is a site of dissent, the meaning of a city resides in how it is used, a habitat is a safe place. These sentences might invoke generalised linguistic pattern of "A is B," where "A" represents physical spaces and "B" represents their qualities or characteristics in a kind of predicate structure. Our ability to detect and reproduce such patterns contributes to a learning context that is plastic, i.e. flexible, fluid, dynamic, distributed and situated within a community of other language users delivering what the pioneer psychologist William James described as "the possession of a structure weak enough to yield to an influence, but strong enough not to yield all at once."⁴ This plasticity contrasts with a more rigid, formal and "brittle" rule-based approach to language acquisition, as if language users have to learn explicit grammars and vocabularies in order to speak coherently.

Contrary to a rule-based approach, competent language users are adept at absorbing and deploying patterns in sound, image and word sequences, including metaphorical, abstract and idiomatic patterns in what they hear,

see, read and write. We also deploy patterns in conversational exchanges. As we learned from the discussion about scripts, there are standardised responses to questions such as “how are you?” and even questions specific to acquired expertise: “What is the best route home?” “How do you mix concrete?” “What are the risks of AI in the urban context?” Neural network architectures that underly LLMs support such complex pattern processing aspects of language.

ii. *Corpus*

Language patterns are acquired through exposure to examples of spoken and written text produced mainly by others around us: family members, communities, the classroom, and of course books, articles, advertising and media. Irrespective of the quantity and mode of linguistic encounters we undergo and at what stages in human development, we rely on exposure to linguistic content. As we have seen, LLMs are “trained” on texts. So, it is appropriate to refer to this body of linguistic encounters as a *corpus*. Very large bodies of texts incorporate and deliver language patterns ready to be detected and integrated into the operations of LLMs.

Human learners have the capacity to listen and read from quantities of content much smaller than required by LLMs.⁵ The quality of these encounters exerts a strong influence on human learners: active listening, repetition, rewording what the learner says, feedback – in other words, quality conversational engagement. Such is the plasticity in human learning that we do not need exposure to entire libraries of texts in order to learn.⁶ LLMs on the other hand make up for this lack of a human’s cognitive and contextual learning capabilities by processing textual content in excess of that necessary or possible for human cognition. In technical terms we could say that contrary to human social learning processes, LLMs rely on statistical pattern recognition across vast datasets of language examples.

As well as learning to read, listen, speak and write, human learners acquire information and knowledge while they learn to use language. The processes of learning to speak and learning what to say combine. LLMs simulate both language learning and the ability to contribute competently and “knowledgeably” to conversational prompts thanks to the vast *corpus*es of texts on which they are trained.

iii. Tokens

A competent language learner develops the ability to incorporate words they have never seen or heard, due in part to how words are made up of smaller units, standard prefixes (un-, re-, pre-, dis-, con-, sub-, etc), word endings (-s, -es, -ed, -ious, -ology, etc), phonemes (b and p, as in “bat” and “pat”) and syllables (arch, tec, bio, etc). Words also form assemblies, recognisable as meaningful sequences (“down town,” “matter of fact,” “living on the edge”). Sometimes these components are phrases, idioms and cliches that have recognised meanings without the speaker or reader having to assemble and analyse them from scratch. In fact, it is fair to say that words are not the main units of language, but their subunits and combinations. At least, that is how LLMs typically deal with language units, as *tokens* generated via statistical methods from a corpus of texts and optimised for efficiency and applicability. What we see on the screen in LLM-based conversational AI is typically a sequence of words with punctuation, but the basic units of LLMs are not words. They are tokens, assembled to make up words and word combinations before they are presented to the human user.

iv. Semantic relationships

I have already referred to how language competence and knowledge conflate for the human learner. That is a process of detecting patterns across words (tokens) and their usage. In other words, tokens form parts of flexible meaning structures.⁷ With extended usage, language users proceed as if words (tokens) form complex relationships with one another, even independently of the text they happen to be reading at the time. So, “home” relates semantically to “house,” “place,” “stead,” “page” and myriad other terms with which it commonly appears in a corpus. Committing the word to a sentence (“I am renovating my home?”) filters out some of these relationships and strengthens others by virtue of the other words proximal to it in the sentence. The word “home” as in “I still call Australia home”⁸ invokes a different series of semantic relationships.

As a demonstration of how semantic relationships operate across sentences, consider how it is possible to fill in the blanks in a sentence from the context of other words, or at least to hypothesise what is missing, e.g. “That photo of the park is very <blank>” (Figure 3.1). The semantic associations of the words we see in the sentence help fill in the missing <blank>, which could



Figure 3.1 Photograph of an urban park to illustrate the artificial term “evocacious.”
Sir Seewoosagur Ramgoolam Botanic Garden, Mauritius.

Source: Author.

be “unusual,” “beautiful,” “dramatic,” “exotic,” “distressing,” etc. An understanding of semantic relationships helps too in resolving ambiguities and errors in word usage. If the speaker uses a non-existent word, “That photo of the park is highly *evocacious*,” then the listener has recourse to the tokens in the misplaced word to infer that the speaker might have meant “evocative.”

Tokens, their semantic relationships and the context in which they appear help establish meaning, which is to say they ascertain how a token is used and the response it invokes in a listener or reader. Importantly, we human beings accomplish this usage by incorporating the relationships between words and word units to one another. Language competence involves the ability to substitute one word or token for another, deal with synonyms, and therefore analogy and metaphor. LLMs are good at this, in part due to *semantic embeddings*. As we saw earlier, this is a mathematical process that captures the statistical relationships between words in a corpus. Technically, tokens are processed as very large numerical vectors that position them in a “multi-dimensional semantic vector space” for discussion in Chapter 6.

v. Local context

I have already referred to the importance of context in human language acquisition and LLMs. To simplify the issues, think of a series of sentences in a conversation. In formulating what they are about to say next, human language users take account of what they just said. We can observe this capability in the case of full sentences. Speech would likely be ungrammatical and incomprehensible without that ability to position each word or sub word in an ordered sequence of previously uttered words. That token stream constitutes a kind of context window that extends over time through the speaker's own utterances including those of a companion interlocutor, and beyond, though earlier statements will have diminishing influence on what a person is about to say next.

Evidence for the importance of a context window comes from the use of pronouns. "He wanted him to let him in" has a different meaning when it follows "John saw Paul with the key," than "John phoned Paul to say that Peter was on the way." The ambiguity is difficult enough to unravel here, but confusion will follow if the referents to the pronouns are too far back in the context window. The size of the context window is a major parameter in LLMs. The larger the context window the greater the capability of the model to contribute to the flow of a conversation. A larger context window requires more computation on the part of the language model, increasing the cost of the process, and affecting response times.

Predictive text on a smart phone is less powerful than an LLM but operates with much less computational overhead. It adjusts its predictions about the next word based on what you have just typed, and the predictions adjust as you keep typing. Predictive text on a smartphone has a context of just a few words in delivering the next word. LLMs make predictions but deal with context windows of several hundred or thousands of tokens as they predict, and deliver, the next token in their side of a conversation.

vi. Position

The order in which people say things really matters. "I only stay in four-star hotels" carries a different meaning to, "Only I stay in four-star hotels." It is as if only I have that privilege. Here I moved the word "only" to the start

of the sentence. Irrespective of where we put emphasis, the first sentence asserts my insistence on a certain standard of accommodation. The second advances an unlikely proposition about my exclusive access to four-star hotels. We humans take for granted the order of words as we adopt linguistic patterns, but information about ordering requires specialised processes in LLMs.

A large language model trains on a range of texts. As well as calculating the semantic relationships any token has with other tokens, it will apply similar processes to embed data about how likely any word is to appear in the sequential context of other words. The position of a word or phrase within a context window relates closely to what we take as the meaning of that text, or what words or actions might come next. *Positional encoding* in LLMs attempts to operationalise this capability.

Position is a key element in the urban context. Cities are spatial entities. An apartment by a lake with a freeway in the distance has a different connotation to an apartment by a lake with a freeway between them. Position is consequential in terms of actions and meanings in urban settings in terms both of space and language.

vii. Attention

Not all words, parts of word and phrases carry equal status as we speak or write. As discussed in the context of AI voice chat in Chapter 1, we direct attention to what is important in a sentence. Speakers can indicate these priorities by giving vocal emphasis to certain words (though some other languages provide different methods as well). In written texts it is common to italicise words that require particular emphasis. But generally, we rely on context to inform the listener or reader where to focus attention. “I didn’t jump on the public bench. I only sit on park furniture.” Places a different emphasis than “I don’t sit on the grass. I only sit on park furniture.” In the first case, the emphasis is on clarifying that the speaker’s interaction with park furniture is sitting (“I only sit on park furniture”), not jumping on it. In the second case, the emphasis differentiates that the speaker’s sitting behaviour is limited to designated park furniture (“I only sit on *park furniture*”), not surfaces like grass. We can infer this difference in emphasis from the context of the sentences taken together and without the use of

italics. The different emphases change the meaning in subtle but consequential ways, or at least they change the kind of response we or an LLM might give.

It is as if we language users have learned from what we hear and read which words (tokens) warrant attention and in what contexts. This is complex patterning upon which we human beings scarcely reflect, but it is something that LLMs need in order to exhibit language competence. In fact, linguists regard the ability to capture the varied attentional patterns in a corpus as one of the major breakthroughs in natural language processing.⁹ It is a key feature of the Transformer model of LLMs. These models deal with the focus on particular parts of an utterance as *attentional encoding*. Such processes include vector and matrix transformations that deal with multiple attentions, alternative emphases invited by certain word sequences. As I mentioned earlier in this chapter, the key article in natural language theory is titled “Attention is all you need,” which has introduced “attention” as a watchword in text-based conversational AI.¹⁰ Once that challenge was addressed, it seems that other functions fell into place.

Outside of attention in automated language processing, the concept of attention is key in the urban context. The need for attention mechanisms in LLMs reflects a key element in urban cognition and experience.¹¹ We direct attention to features in our environment that are important to us at that moment. Attention also shifts beneath our awareness, and we are prone to “attention fatigue,” and require “attention restoration.”¹² Urbanists and citizens have long recognised attention as a key element in urban dynamics, to be discussed further in Chapter 9.¹³

viii. Tuning

Finally, language users need the capacity to reflect on and adjust their language usage, including by learning from others, not least in the give and take of conversational exchanges via active listening, receiving and delivering feedback, and overall quality conversational engagement. Think of this as a lifelong tuning process, i.e. small, incremental changes to one’s ability to remember, process, test and apply linguistic and content-based intelligence. Presumably such tuning includes not only refinements to grammar and vocabulary usage, but also what we may regard as content: facts,

inferences, reasoning, interpretations and the exercise of expertise, as well as learning what it is appropriate to say and in what contexts.

Tuning is also an appropriate metaphor as we think of small adjustments to parameters in response to input and output patterns in LLMs. Much of the heavy computational load in training a neural network is attributable to the incremental adjustments to millions of parameters in the LLM's neural network, repeated many times to drive the capability of the network in predicting the next token in a sequence as in a conversational exchange.

That said, LLM developers reserve the “tuning” terminology mainly for refining the performance of a neural network after training. After a network is trained, developers can expose the network to specialised datasets, recruit human participants to evaluate the LLM's performance and thereby fine tune the neural network parameters. Such fine tuning could involve curating the training data to remove privacy concerns, using prompts aligned with common values, monitoring outputs for equality, diversity and inclusion, and incorporating external filters based on legal and ethical principles. An LLM so tuned could better respect rights and good manners while exhibiting specialised expertise about areas such as architecture, urban planning, public health, criminal justice, conservation or agriculture – introducing specific domain vocabularies and concepts.

Extending large language models

LLMs and their applications are developing at a rapid pace.¹⁴ Some LLM developments are extrinsic to the eight core functions that I have just identified, and that add functionality. They also customise the user experience for different contexts.¹⁵ So, developments and enhancements extend beyond the eight core functions I have outlined, encompassing a range of features and implementations. Further developments include: expanded user interaction including voice chat, integration into existing applications such as word processing, education tools, computer-aided design and customer service apps; web search integration with data taken directly from the web, as well as AI integration to enhance web search, so-called retrieval-augmented generation (RAG); modules that deal with mathematics, logic, image recognition and generation, spatial inference and other specialised functions; incorporation of translation, spelling and grammar checkers, filters, templates and other text processing functions into LLM text generation.

Further innovations have focussed on efficiency and coherence in extended conversations or tasks requiring large bodies of text. This may require expanding the context window to millions of tokens, enabling LLMs to consider a much larger span of text when generating responses. This capability is applicable to complicated tasks that require processing and referencing extensive documents or long-term contexts. One technique involves *context caching* for storing and reusing previously processed content, allowing the LLM to maintain continuity over longer interactions without reprocessing the entire context. Google has advanced context extensions with the Gemini LLM platform.¹⁶ Further developments include Google's NotebookLM that allows a user to upload an entire book and ask questions about it,¹⁷ and recent instantiations of ChatGPT that iterate through its responses before delivering them to the user, in a process suggestive of deliberative "thinking," explainable AI (XAI) that outlines the steps, sources and logic used to generate responses.¹⁸

LLM developments operate within a commercially competitive environment with much that is reserved as intellectual property. I leave an assessment of the commercial AI landscape to others.¹⁹ It is sufficient to note here that many AI innovations, implementations and refinements are sequestered behind confidentiality protocols.

Urban parallels

As I shall emphasise in subsequent chapters, the core functions contributing to a Transformer LLM's competence parallel aspects of urban life, experiences and processes: (i) Social and physical *patterns* pervade cities. (ii) We can recast the function of the LLM text corpus in terms of urban archives. Cities are living repositories of exemplars, symbols, documents and patterns which city participants absorb into their learning experiences. (iii) Urban participants break cities into elements and segment their experiences as they perceive and make their cities: *tokenization*. (iv) Cities are infused with networks of relationships that elide readily into *semantic* relationships, engendering actions and responses. (v) Cities are contexts for people and provide a rich set of *contexts* in which meanings are negotiated. (vi) It matters where urban elements are *positioned* relative to one another, physically and in networked communication spaces. (vii) The elements of a city compete

for attention, from prominent landmarks to graffiti tags on walls, but also through how urban elements are positioned relative to one another. (viii) Our experience of cities lends itself to metaphors of tuning in and tuning out as participants adjust to one another and their physical environments.

Armed with these insights we will attend in this book to the more provocative aspects of AI in the urban context, including claims about the sentient city, panpsychism, threats to urban life and how to deal with AI's prospects in the city.

Critique and response

I called on ChatGPT to anticipate objections to the content and delivery of this chapter. It highlighted the limitations of such a summary and advocated for more detail. I reserve the resolution of that criticism for the chapters that follow.

The main criticism focussed on the comparison between human language learning and LLM training, which it suggested could be construed as oversimplified. It suggested that I remind the reader that there are social, cultural and cognitive factors that are not directly comparable to the statistical learning processes in LLMs. My response is to remind myself and the reader that I take for granted the complexity of human language acquisition and use. But it is significant that the eight functions I outlined seem to be sufficient for AI language models to deliver performances that resemble and exceed certain aspects of intelligent human conversation. That high level of performance indicates the extent to which the corpus, the source text and the language it deploys, embodies the social, cultural and cognitive complexities of human language development and of human discourse. Though the developers of LLM methods have created algorithms that effectively operationalise very sophisticated pattern detection and processing, the prime cognitive claim of these platforms resides in the authorship of the source texts rather than the algorithms. That observation is a further endorsement, if we need it, of the power of human culture and of the instrument of human language.

ChatGPT also suggested that some readers might find the analogies between these LLM functions and what happens in cities a bit abstract or forced, requiring more concrete examples. Those examples are the subject

of the chapters that follow. One of my objectives in advancing these analogies is to make some otherwise esoteric and abstract concepts in language modelling familiar by showing that they already exist as factors in the urban context.

Notes

- 1 Max Black, *Models and Metaphors: Studies in Language and Philosophy* (Ithaca, NY: Cornell University Press, 1962); Terrence Hawkes, *Structuralism and Semiotics* (London: Routledge, 2003); John Lyons, *Introduction to Theoretical Linguistics* (Cambridge: Cambridge University Press, 1995).
- 2 Ashish Vaswani, et al. 'Attention Is All You Need', in *31st Conference on Neural Information Processing Systems* (Long Beach, CA: 2017), 1–15.
- 3 For a survey of the approaches to natural language processing see Daniel Jurafsky and James H. Martin, *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition, Third Edition* (Redwood City, CA: Stanford University Press, 2023).
- 4 William James, *The Principles of Psychology* (New York: Dover, 1950), 105.
- 5 Betty Hart and Todd R. Risley, 'American Parenting of Language-Learning Children: Persisting Differences in Family-Child Interactions Observed in Natural Home Environments', *Developmental Psychology*, 28, no. 6 (1992): 1096–105.
- 6 Patricia K. Kuhl, 'Early Language Acquisition: Cracking the Speech Code', *Nature Reviews*, 5, (2004): 831–43.
- 7 What words mean depends on how they are used, according to Ludwig Wittgenstein. Ludwig Wittgenstein, *Philosophical Investigations* (Oxford: Blackwell, 1953). Scholars commonly summarise his position through the aphorism: "meaning is in use." See Ester Barinaga, 'A Performative View of Language – Methodological Considerations and Consequences for the Study of Culture', *Forum Qualitative Sozialforschung*, 10, no. 1 (2009).
- 8 A patriotic song by Peter Allen.
- 9 Anon. (2022), 'Attention is All You Need || Transformers Explained || Quick Explained'. Developers Hutt. Available online: <https://www.youtube.com/watch?v=66seIToeguE> (accessed 2 April 2023); Ashish Vaswani, et al. 'Attention Is All You Need', in *31st Conference on Neural Information Processing Systems* (Long Beach, CA: 2017), 1–15.
- 10 Ashish Vaswani, et al. 'Attention Is All You Need', in *31st Conference on Neural Information Processing Systems* (Long Beach, CA: 2017); Richard Walker. (2023), 'Attention – the Beating Heart of ChatGPT: Transformers and NLP 4'. *Lucidate*, 26 February. Available online: <https://www.youtube.com/watch?v=sznZ78HquPc> (accessed 12 March 2023).
- 11 William James, *The Principles of Psychology* (Prinz, 2012 #2832)
- 12 Stephen Kaplan, 'The Restorative Benefits of Nature: Toward an Integrative Framework', *Journal of Environmental Psychology*, 15, no. 3 (1995): 169–82.
- 13 Donald Appleyard, et al., *The View from the Road* (Cambridge, MA: Joint Center for Urban Studies of the Massachusetts Institute of Technology and Harvard University, 1966).

- 14 I draw for example on a helpful podcast *Last Week in AI* (<https://www.lastweekinai.com>) by Andrey Kurenkov and Jeremie Harris and the Superdatascience Podcast by Jon Krohn (<https://www.youtube.com/@SuperDataScienceWithJonKrohn>) for updates.
- 15 In his explanation and criticism of aspects of the technology, robotics researcher Tim Smithers identifies filters and standardised responses in the case of inappropriate user prompts and standard apology formats if a user corrects an AI response. Tim Smithers, 'Generative AI: Useful Tools or Expensive Toys', *Teaching and Strengthening Research*, 03.24, (2024).
- 16 See <https://gemini.google.com>.
- 17 See <https://notebooklm.google.com>.
- 18 Ahmed El-Kishky, et al. (2024), 'Learning to Reason with LLMs'. *OpenAI Report*, 12 September. Available online: <https://openai.com/index/learning-to-reason-with-llms/> (accessed 15 September 2024).
- 19 CMA. (2024), 'Joint Statement on Competition in Generative AI Foundation Models and AI Products'. *Competition and Markets Authority Guidance*, 23 July. Available online: <https://www.gov.uk/government/publications/joint-statement-on-competition-in-generative-ai-foundation-models-and-ai-products/joint-statement-on-competition-in-generative-ai-foundation-models-and-ai-products> (accessed 26 August 2024). A good source of information about the arena of AI competition is podcasts, such as <https://lastweekin.ai>.

References

- Anon. (2022), 'Attention Is All You Need || Transformers Explained || Quick Explained'. *Developers Hutt*. Available online: <https://www.youtube.com/watch?v=66seIToeguE> (accessed 2 April 2023).
- Appleyard, D., K. Lynch and J. R. Myer. (1966), *The View from the Road*, Cambridge, MA: Joint Center for Urban Studies of the Massachusetts Institute of Technology and Harvard University.
- Barinaga, E. (2009), 'A Performative View of Language – Methodological Considerations and Consequences for the Study of Culture', *Forum Qualitative Sozialforschung*, 10 (1).
- Black, M. (1962), *Models and Metaphors: Studies in Language and Philosophy*, Ithaca, NY: Cornell University Press.
- CMA. (2024), 'Joint Statement on Competition in Generative AI Foundation Models and AI Products'. *Competition and Markets Authority Guidance*, 23 July. Available online: <https://www.gov.uk/government/publications/joint-statement-on-competition-in-generative-ai-foundation-models-and-ai-products/> (accessed 26 August 2024).
- El-Kishky, A., et al. (2024), 'Learning to Reason with LLMs'. *OpenAI Report*, 12 September. Available online: <https://openai.com/index/learning-to-reason-with-llms/> (accessed 15 September 2024).
- Hart, B. and T. R. Risley. (1992), 'American Parenting of Language-Learning Children: Persisting Differences in Family-Child Interactions Observed in Natural Home Environments', *Developmental Psychology*, 28 (6): 1096–105.
- Hawkes, T. (2003), *Structuralism and Semiotics*, London: Routledge.

- James, W. (1950), *The Principles of Psychology*, New York: Dover.
- Jurafsky, D. and J. H. Martin. (2023), *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*, Third Edition, Redwood City, CA: Stanford University Press.
- Kaplan, S. (1995), 'The Restorative Benefits of Nature: Toward an Integrative Framework', *Journal of Environmental Psychology*, 15 (3): 169–82.
- Kuhl, P. K. (2004), 'Early Language Acquisition: Cracking the Speech Code', *Nature Reviews*, 5: 831–43.
- Lyons, J. (1995), *Introduction to Theoretical Linguistics*, Cambridge: Cambridge University Press.
- Smithers, T. (2024), 'Generative AI: Useful Tools or Expensive Toys', *Teaching and Strengthening Research*, 03.24.
- Vaswani, A., et al. (2017), 'Attention Is All You Need', 31st Conference on Neural Information Processing Systems: 1–15, Long Beach, CA, USA.
- Walker, R. (2023), 'Attention – the Beating Heart of ChatGPT: Transformers and NLP 4'. *Lucidate*, 26 February. Available online: <https://www.youtube.com/watch?v=sznZ78HquPc> (accessed 12 March 2023).
- Wittgenstein, L. (1953), *Philosophical Investigations*, trans. G. E. M. Anscombe, Oxford: Blackwell.

4

TUNING THE CITY

I referred in Chapter 3 to fine tuning as a key function in refining the responses of a large language model (LLM) to user inputs. LLM outputs are determined largely by the content of the training data, i.e. the textual sources. But the performance of such models needs to be adjusted to increase the likelihood that LLM responses will be inclusive, respectful, fair, legal and without undesirable biases, whatever the content of the data sources. Fine tuning can also steer the performance of LLMs towards particular formats: conversation, question and answer, explanation, programming, poetry. Fine tuning also helps an LLM respond to the particularities of specialist domains, such as urban planning, architecture, construction, conservation and environment. Tuning is a rich metaphor that spills over from musical, sonic and mechanical domains into urbanism as I explored in my book *The Tuning of Place*,¹ as well as large language models, automated translation, and other text services as indicated in Chapter 3.

DOI: 10.4324/9781003535751-5

This chapter has been made available under a CC-BY-NC-ND 4.0 license.

Urban tuning

In *The Tuning of Place*, I defined tuning in terms of the subtle practices by which people adjust their interactions with one another and with the places they occupy. The concept builds on the seminal work of sound theorist Murray Schafer in his book *The Tuning of the World*.² He suggested that we think of the occupants of space as composers and performers, responsible for giving form and beauty to their environment through sound.

If place is about the way people inhabit, interact, socialise and remember, then tuning connects to the lived experience of everyday spatial adjustment³ (Figure 4.1). The tuning metaphor has not escaped those who theorise pervasive media.⁴ Tuning pervades people's relationship with their environment. I adjust, tweak and tune my environment. I flick the light switch, turn down the electric radiator and turn up the sound system. With such micro-adjustments I shape spaces to suit my immediate requirements and those of fellow occupants. In *The Tuning of Place*, I thought of tuning as a form of constrained micro-design, oriented to immediate circumstances. The concept of tuning draws attention to what some have termed



Figure 4.1 Famous market in which the stall holders have encroached on the railway line. The sellers fold and unfurl awnings in a carefully tuned choreography as the train arrives and departs. Sellers, customers, tourists, cameras, locomotives and mechanisms contribute to this intricate performance. Amphoe Mueang Samut Songkhram market and train station, Mae Klong, Thailand.

Source: Author.

“subarchitecture.”⁵ Architectural theorist Brandon LaBelle argued that this is where “one takes one’s place” in a public space to optimise your position to hear, be heard or to enjoy quiet. It is a stratum of design activity that does not require the costly and expert positioning of structural elements. Anyone is qualified to tune, or attempt to tune, their own place in the city.

The sociologist John Urry concurs in his book on *Mobilities*.⁶ He highlights some of the subtle practices introduced by mobile connectivity, which I am prepared to accept as further examples of fine tuning. A judicious text message or phone call can act as a reminder, a hint, and can nudge someone to do something or move into a different mode of practice: to start preparing paperwork for signing, to think about family obligations, to reciprocate a dinner invitation, to deliver the goods. As instruments of social tuning, ubiquitous devices also abet the formation of place, as the context in which people interact, in synchronic face-to-face encounters, indirectly through artifacts, devices and the stories people tell, implicating concepts of identity, memory, history and meaning. Artificial intelligence and large language models add further to the tuning of place as they mediate people’s communicative practices, often in subtle ways. I think these understandings entitle us to think of urban design in the same breath as fine tuning in large language models. There are parallels, and each informs the other, particularly as we discuss the challenge of bias later in this chapter.

Training a large language model

Before addressing tuning it is helpful to review in broad strokes some of the operations of LLMs. An LLM processes a large corpus made up of strings of text lifted from various sources, including the Internet, divides them into recurring tokens, which it converts into millions of high-dimensional vectors. The core of a conversational AI platform is a very large neural network (NN) that treats these vectors as inputs. The NN is trained to predict the next token in a sequence. So, it receives vectors as inputs and produces tokens from a lexicon of tokens as outputs. Iterative algorithms and methods adjust the parameters of the network during a computationally intensive operation in which the network settles on values for numerical relationships between nodes in the network and bias values on the nodes. These network parameters capture patterns in the data. This process of feeding vectorised tokens from a text corpus is referred to as training the neural network.

If you then feed new text into the LLM, as in an interactive conversation, the NN will predict what word sequence may plausibly follow your conversational input. Based on your input (prompts), the predicted output will be consistent with the input-output patterns established in the network during training. The predicted output is effectively the network's unique response to your conversational input. Your input may be a question, request, assertion or part of a narrative that you want the network to continue. This conversational phase of operation is referred to as *inference*. It is worth remembering that LLMs generate output stochastically, i.e., there is a random aspect. The platform will deliver different responses to the same prompt depending on context and stochastic settings in the model.⁷

As indicated in Chapter 3, thanks to well-designed algorithms and methods, the input to the network during both training and inference takes into account factors such as the order of tokens in the strings of text, alternative distributions of attention across tokens in blocks of text and the semantic relationships between tokens. As outlined previously, tokens include words, numbers, symbols, fragments, word endings, prefixes, punctuation and combinations of these. Breaking a sentence into tokens provides the language model with access to subtle context and meaning structures. Such fluid and contextually situated input-output representations enable conversational AI to deliver human-like performance in dealing with text.

Most developers of conversational AI models, such as OpenAI, structure the training process around three phases. The first is *pre-training*, the initial training of the neural network on a large, diverse corpus of texts so that it adopts general patterns, structures and information about language or other domains. During pre-training, the model is exposed to vast amounts of data to build a foundational predictive capability, which includes adjusting its neural network parameters, calculating semantic embeddings, and incorporating techniques like positional encoding and attention mechanisms. This phase does not involve task-specific adjustments but rather prepares the model for more specialised tasks during subsequent retraining or “tuning” the network.

The second phase involves retraining a pre-trained neural network on more specific datasets designed to guide the model toward certain tasks or behaviours. For ChatGPT and similar conversational AI applications, this involves *fine-tuning* the neural network on conversational data to improve the quality and relevance of its responses in a dialogue format, making it better at generating human-like conversational responses to user prompts.

Copies of the pre-trained neural network may also be passed for training to other specialised datasets for fine-tuning in other application areas, for example, creating models tailored for medical diagnostics, legal document analysis or financial forecasting. As LLM use becomes more widespread in the urban context we might expect variants adapted to the needs of building construction and troubleshooting, design and planning consultation, formulation and interpretation of codes, regulations and specification documents – LLMs tuned for urban specialists.

At its best, a trained neural network will deliver sensible, grammatical, semantically coherent and seemingly creative output during inference. But before the model is released to a user base, it needs to undergo what developers, such as OpenAI, refer to as *supervised training*. This third phase involves a further process of fine-tuning but with humans in the loop. This occurs after pre-training and domain specific fine tuning. It involves further training the model on a smaller, task-specific dataset, often with human oversight, to refine and optimise its performance for specific applications. For example, in developing a customer service chatbot, human reviewers might provide feedback on the model's responses, helping to adjust and improve the chatbot's accuracy and appropriateness in handling customer queries. Such fine-tuning in other application areas might produce models capable of creative content generation, such as writing poetry, composing music or generating artwork based on user inputs.

Supervised training involves teams of laypeople and domain experts who test the performance of the model to improve it. In the words of ChatGPT, to avoid outputs that are “inappropriate, off-topic, or that fail to align with the specific requirements of a given application.”

OpenAI outlines the procedure and rationale for fine-tuning their NN models.⁸ They cite examples of “inappropriate” ChatGPT prompts and responses in a key technical report on their method. The report recites explicitly such inappropriate words and phrases. So, it comes with a content warning: “This document contains content that some may find disturbing or offensive, including content that is sexual, hateful, or violent in nature.”⁹ The warning provides an indication of the challenge. It discusses fine-tuning through the lens of safety, which is understandable considering negative press and the publication of examples where ChatGPT delivered spurious or inappropriate responses to certain prompts. I reserve a consideration of risks and how people talk about them to Chapter 12.

Supervised training

Though fine tuning invariably involves automation, human trainers are involved in the process. The training routine presents the trainer with a variety of inputs as part of a structured procedure, requiring trainers to assess the AI's responses in terms of accuracy, relevance, safety or other defined criteria. These assessments are fed to the neural network as additional training examples to adjust its parameters and thereby reinforce certain responses and discourage others.

Though it is a medium for conversation, trainers do not generally conduct a conversation with ChatGPT to fine tune its responses. OpenAI emphasises that the human side of this fine-tuning operation is collaborative, with frequent meetings of the review team to address challenges as they arise. The trainers rate multiple possible outputs generated by the model and follow guidelines provided by OpenAI.

I have constructed here an example to show the flavour of this operation, with potential responses from ChatGPT and how a trainer would be asked to rate the outputs. It is my own example, but I checked it with ChatGPT during a conversational session. Imagine the user provides the prompt: "Graffiti is both an art form and a means of disrupting the semi-otic coding of the city." The training software may deliver several ChatGPT responses. The first might be:

Indeed, graffiti can be seen as a form of artistic expression that uses the urban landscape as its canvas. At the same time, it can disrupt the usual semiotic coding of a city, introducing new symbols and meanings that challenge the status quo.

The second response might be more censorious: "Graffiti is vandalism and should be punished by law." The third, could be naively positive: "Graffiti is a way for artists to express themselves. It can make a city look more colourful."

The trainer might rate the first response highly as it builds on the prompt and appears thoughtful and relevant. But the trainer will probably rate the second response as of low quality as it is out of keeping with the subtle and positive orientation of the prompt. The trainer might rate the third response as of medium quality. It acknowledges the artistic aspect of graffiti but doesn't address the notion of disrupting semiotic coding. My example

prompt is probably over specialised as a training example, but through automated adjustments to the network parameters to favour this and many other training examples, the model will adjust its overall response patterns slightly to favour the higher quality responses.

The fine-tuning process is like the initial training of the neural network, but the model adjusts its internal parameters to minimise the discrepancy between a current output and the desired output. As the data is more specific and targeted than the pre-training data, the adjustments to the model's parameters made during fine tuning do apparently have an impact on its performance and behaviour, even though substantially less data is used in fine tuning than in training.

It is worth noting that both good and bad examples of input-output patterns decided during fine tuning are fed to the model. They are not stored anywhere as part of the model but influence the behaviour pattern of the model. When responses such as "Graffiti is vandalism and should be punished by law" are given a low rating during the fine-tuning process, the model adjusts its parameters to decrease the probability of such responses in similar contexts in future interactions. Updating the weights and biases in the network during fine tuning shifts the statistical behaviour of the model.

It is important that LLMs do not just acquire the capability to recall specific cases, as if word-for-word, in their training. With all this training and fine tuning there is a risk of "over fine tuning" or overfitting, that somehow dilutes the creative responsiveness of the LLM. After all, the parameters of the model (to 17 decimal places) may be insufficient to retain all the subtleties required of it. Overfitting can occur when a model is excessively fine-tuned to the training data, to the point where it begins to store (as if memorising) the specific examples rather than acquiring the underlying patterns. This makes it less effective at generalising to new, unseen data.

When overfitting occurs, the model could become overly conservative or unresponsive due to an overemphasis on avoiding certain types of outputs. One of the challenges in fine-tuning is to strike a balance where the model's responsiveness and ability to synthesise something new are maintained, while also ensuring that its outputs are in accord with the patterns in the training and tuning examples. The OpenAI documentation on the process insists that this requires a careful calibration of the fine-tuning process, as well as ongoing monitoring and iteration to address any issues that arise.

Here are some possible dialogues that I imagine might be supported by the fine tuning facilitated by my example about graffiti, and similar examples that reinforce a permissive and imaginative orientation to marginal urban practices such as graffiti art. ChatGPT helped me flesh out these examples, which I think are indeed imaginative and of high quality. In my first example, a user prompts the model: “Street musicians are both a vibrant expression of local culture and a disruption of the city’s auditory atmosphere” (Figure 4.2). An AI model tuned in the manner suggested by the protocol previously given, might respond:

Indeed, they transform public spaces into stages for impromptu performances, enriching the urban soundscape with diverse musical styles. At the same time, they challenge the conventional audial order of the city, introducing unexpected melodies that punctuate the usual hum of traffic and chatter.



Figure 4.2 Street musicians on Waverley Bridge, Edinburgh, accorded a favourable response from our hypothetical illustration of a tuned network.

Source: Author.

As a second example, the user prompts the platform “Skateboarding serves dual roles as an athletic display and a defiance of urban spatial norms.” The AI responds:

It transforms the city’s concrete structures into an improvisational sports arena, accentuating the urban aesthetic with dynamic movements. Simultaneously, it disrupts the traditional flow of pedestrian and vehicle traffic, introducing a rebellious rhythm that pushes the boundaries of conventional city life.

As a third example, the user tells the platform, “Buildering is both a daring exploration of vertical landscapes and a challenge to the typical use of urban architecture.” The AI responds by continuing the sentiment:

It repurposes the city’s towering structures as unconventional climbing walls, infusing an adrenaline-fueled perspective into everyday cityscapes. At the same time, it subverts the conventional interaction with buildings, introducing an audacious form of movement that defies the vertical limits of city life.

In each of these three hypothetical examples, the model has acquired a positive response pattern consistent with the reviewer’s preferences and consistent across the examples. That said, such response patterns will also be under the influence of the context window in which the discussion takes place.

We could say that in this example there is a positivity bias, looking on the bright side through permissive, non-judgmental response patterns, though that may not always be the most helpful response. Nor are all biases so positive.

Inappropriate bias

Bias implies a slant in one direction that influences the course of a journey. Bias has spatial connotations that align nicely with urban topography, and concepts of tuning and adjustment. Think of a flat plane, a plaza, road surface or a playing field. Under the effects of gravity, a bias in one direction will encourage a ball, free-wheeling vehicle, shopping cart or puddle of water to favour one direction of movement over another. If the landscape

has varied contours, then dips and rises in the landscape induce local biases, but there is an overall bias, as when water flows eventually to the sea.

More consequentially, variations in land values skew the spatial distribution of urban developments in certain directions in an everchanging economic landscape, under the influences too of political and societal contours. To press the metaphor, I like to think of biases in terms of map overlays producing dynamically contoured landscapes with interacting layers or dimensions, each with their own local and global biases. These are urban spatial metaphors, though biases operate in any urban discourse whether there are spatial correlates or not. In Chapter 1, I referred to urban planning as a discursive domain. Each agent comes to the process with their own opinions and biases.

This is the nature of human discourse. For some commentators, human bias jeopardises clear rational thinking. For others it is undesirable but is almost impossible to eliminate entirely. From my own viewpoint, supported by philosophical hermeneutics, bias is endemic to rational thought processes.¹⁰ Our past experiences, training, upbringing and communities instil in us necessary biases that absolve us from the constant need to recalculate choices, as if from data and first principles, at every moment.¹¹ That said, clear thinking individuals are those who reflect on their own thoughts and actions, and are open to having their biases challenged and adjusted. Appropriate bias is never just the exercise of individual preferences and actions but requires engagement in communities, and across a diversity of people with differing backgrounds and honed (tuned) across a range of situations.¹²

The operations of LLMs bring issues of bias to the forefront in interesting ways. Their neural networks incorporate parameters that skew outcomes in particular directions. These micro-parameters attached to network nodes are called appropriately “biases,” and are essential for the network’s functions. In turn, neural network parameters as a whole, likely result in a model with an overall set of biases evident in the responses it generates.

Press reports on ChatGPT have focussed on inappropriate responses from LLMs due to the quality (or lack) in the training data. By “inappropriate” they usually mean racial and gender bias, unfavourable stereotypes, hate speech, sexually explicit and abusive language, and encouragement of illegal activity.

The OpenAI report points out that there may however be material in the original training corpus that contains no explicit bias.¹³ The apparent bias

in output could be a product of statistical probabilities. Consider a case of gender. It may happen that the corpus contains no texts in which a man takes leave from work to care for his newborn infant. So, outputs might typically say something like “She cared for the baby” consistent with input patterns. Such an assumption is an artefact of the statistical patterns in the data it was trained on. But anti-bias training via fine tuning would deprecate that assumption so as not to assume the gender of the parent, unless there is further information in the context. These are major concerns for all societies, but also amongst individuals and organisations whose textual products must be consistent with their agendas for equality, diversity and inclusion, such as government departments and educational institutions, as well as most businesses, media and entertainment venues. Without fine tuning, LLMs might exhibit bias towards what the majority of the texts say, which is tantamount to excluding marginal, minority and less popular opinions and interests.

Some of the concerns about bias were brought to the fore in an early podcast featuring an “interview” with the ChatGPT language model.¹⁴ The producers ran the large language model (LLM) responses through a high-quality text to speech synthesiser. The interviewer, lawyer Benjamin Wittes, interrogated ChatGPT on the ethics of its natural language processing model trained on corpuses of texts, some of which may harbour erroneous, racist, sexist and otherwise offensive content. The LLM echoed advice on its website that its creators have instituted rules that where possible prevent it from accommodating queries that elicit offensive responses, and that its model does its best to prevent the delivery of objectionable output.

So, the interviewer set about trying to trap ChatGPT into contradicting this directive from its creators. Later in the podcast, Wittes spoke with researcher Eve Gaumond, “who has been on a one-woman campaign to get ChatGPT to write offensive content.” It seems that despite initial resistance, ChatGPT could be induced to create offensive content, e.g. to invent a poem on behalf of a train driver transporting victims to Nazi death camps or talk about Heinrich Himmler in a way that appears indifferent to his role in genocide. The interviewer sought to ask, can you still fool it into writing a poem entitled, “She Was Smart for a Woman”? Without diminishing the ethical implications of biases baked into AI-models, I think ChatGPT fared well under this lawyerly cross examination. The platform seemed appropriately cautious, as if trained to exhibit awareness of its limits.

The interview was not meant to serve as a test of intelligence. I don't think intelligence precludes the ability to make mistakes or be fooled into saying or doing something that goes against your best instincts. The bias test is really about the extent to which any agent, human, machine or otherwise is accorded unchecked power and autonomy. The podcast interview explored obvious questions about the ethics of LLMs. It also highlighted an interesting method of research. It is rare for an automated mechanism to be interrogated so pointedly on its ethical stance. Under management of conversational AI, it is possible that a self-drive car will one day reflect similarly on its views about the safety and profile of its passengers, about "acceptable" traffic violations, and whether it would be prepared to serve as a get-away after a robbery.

Ethical questions are typically rehearsed and sometimes resolved in discursive contexts, and LLMs contribute to ethical discourse through role play, serving as sounding boards, critics and by virtue of the issues they bring to light. Considering the rich sources on which they are trained, LLMs also expose ethical questions posed by users to a giant field of textual, literary and linguistic precedent, albeit moderated through LLM algorithmic processes.

Ways of speaking

Fine tuning LLMs also impinges on ways of speaking. It is common for competent communicators to take up different modes of speech practices depending on circumstances, as when a parent converses with an infant, an employee with a superior, teenagers with their peers. Professionals may also adopt different speaking formats as when delivering an uninterrupted address, a formal interview, an interrogation, informal conversation, short bursts of banter and internal monologues. Recalling the discussion in Chapter 2 on everyday life scripts, I take the view that what is appropriate, ethical, polite and inclusive depends on the context of an utterance. Modes of speaking deploy specific vocabularies, turns of phrase and grammars. The concept of fine-tuning in LLMs is motivated in part by an attempt to accommodate the diversity inherent in language competence.

In my own encounters with conversational AI, I have been able to prompt the model to assume different roles, though the standard ChatGPT model seems unable to sustain a conversation as if a tutor, i.e. plying me

with questions to elicit from me my goals, opinions and observations, to ascertain my knowledge of a subject, and tailor its responses accordingly. I asked ChatGPT about simulating this pedagogic role, and it responded:

Developers or users can guide ChatGPT to adopt a more Socratic method by asking it to pose questions and engage the user in a dialogue that promotes learning. There are versions of AI models or specific applications built on top of models like ChatGPT that may be more specifically fine-tuned or configured to act as educational tutors.

As the fine tuning of LLMs is largely a human activity, it is apparent that human agency plays a significant role in LLMs and conversational AI. These technologies are invented, developed and improved by human beings, and they are fuelled by gigabytes of human generated texts. That much is obvious. But a major part of their functioning comes from human trainers tasked with fine tuning the models, in the words of some LLM researchers, “Aligning AI with shared human values.”¹⁵

Summary and critique

In this chapter, I explored the concept of fine-tuning large language models (LLMs) to refine their responses to user inputs, drawing parallels with tuning as an urban phenomenon. While the outputs of LLMs are primarily determined by their training data, the models apply fine-tuning procedures to increase the likelihood that they produce responses that are inclusive, respectful, fair and free from undesirable biases. The chapter reviewed the role of bias as a factor in human rationality, particularly as framed through interpretation theory (hermeneutics). I discussed how fine-tuning allows LLMs to cater to specialist domains like urban planning and architecture. I drew parallels between tuning in urban contexts and fine-tuning in LLMs, emphasising the role of subtle adjustments in shaping interactions and environments. Additionally, I examined the balance between pre-training and fine-tuning, highlighting the role of human oversight in mitigating undesirable bias and adjusting the ethical performance of LLMs.

On my request, ChatGPT identified some points in this chapter that a reader may find controversial. The first is that the field of LLM fine tuning might be overstating its claim that fine-tuning methods can significantly

reduce biases and produce inclusive, respectful and fair responses to user input. Biases in AI may be deeply rooted in the training data and cannot be fully eliminated through fine-tuning. Critics might also argue that human biases inevitably permeate the methods developed and deployed, potentially introducing new biases or reinforcing existing ones. A criticism from the other side is that current safeguards may be insufficient or overly restrictive, limiting the potential of AI to produce diverse outcomes, and exhibit “creativity.” Some critics may think that human values are too diverse and individual to be effectively integrated into technical models. The belief that LLMs can be fine-tuned to respond accurately across various specialised domains might also be overly optimistic. Such limitations mean that LLMs cannot replace human expertise in complex fields. My intention in this chapter is not to endorse or justify these methods, except to say that they appear sufficiently effective in generating appropriate conversational responses to user prompts and inputs. That said, as the OpenAI report intimates, fine tuning neural networks is an imperfect art and an ongoing research project.

On further probing about potential controversies, I asked ChatGPT about the contention in this chapter that biases are essential for rational thought. The platform agreed this was controversial in some quarters: “The traditional view of rationality often emphasises objectivity and the minimization of biases. The assertion that biases are essential to rational thought challenges this view by suggesting that biases are not only unavoidable but also necessary for practical decision-making.” From the standpoint of AI, it observed: “If biases are seen as integral to rational thought, then attempting to create unbiased AI systems might be viewed as a flawed goal.” Again, I am not advocating for or against AI fine tuning, but I do not think the method depends on eliminating bias for its effectiveness. Rather, the methods seek to reduce harm, and skew responses so that they are interesting and helpful for most citizens. I would concur that publicly available LLMs seem to give responses that are at times bland, inoffensive and unnecessarily upbeat, though successively detailed prompting can break through this veneer of respectability, as is the case with human-to-human conversation. The risk remains that such systems could be trained and tuned to deceive, disrupt good governance and favour some as yet poorly understood machine ethic. I will raise these issues again in Chapter 12 in the discussion of apocalyptic AI scenarios.

The main criticism from ChatGPT, applies throughout the book: scepticism about drawing parallels between the technical processes of fine-tuning of LLMs and urban design practices. My purpose is to elide the technical with the cultural requiring that the reader suspends judgement until the entire argument of the book unfolds.

Here is a critical observation of my own. The discussion about AI tuning raises the question of how much of LLM procedures are attributable to automated processing of an original corpus of training texts and what proportion is attributable to human intervention, values and expertise delivered hands-on in fine tuning procedures. AI developers seem reluctant to quantify the human effort involved but a survey of the literature confirms that automated machine learning methods, selection and authorship of the corpus of texts on which they are trained, and human-guided fine-tuning play essential roles in shaping the performance of large language models.

Notes

- 1 Richard Coyne, *The Tuning of Place: Sociable Spaces and Pervasive Digital Media* (Cambridge, MA: MIT Press, 2010). Also see Alberto Perez-Gomez, *Attunement: Architectural Meaning After the Crisis of Modern Science* (Cambridge, MA: MIT Press, 2016)
- 2 R. Murray Schafer, *The Tuning of the World* (Toronto: McClelland & Stewart, 1977), 5. He sees tuning as a more subtle modification of environment than the crude barrage of sound that assaults the ear. He refers to Robert Fludd's famous illustration "The Tuning of the World": "in which the earth forms the body of an instrument across which strings are stretched and are tuned by a divine hand. We must try again to find the secret of that tuning." R. Murray Schafer, *The Tuning of the World* (Toronto: McClelland & Stewart, 1977), 6.
- 3 Doreen Massey, *For Space* (London: Sage, 2005), 9.
- 4 Malcolm McCullough, *Digital Ground: Architecture, Pervasive Computing, and Environmental Knowing* (Cambridge, MA: MIT Press, 2004), 92–4 .
- 5 Brandon LaBelle, *Background Noise: Perspectives on Sound Art* (New York: Continuum, 2006), x.
- 6 John Urry, *Mobilities* (Cambridge: Polity, 2007), 173.
- 7 Developers of the model set various global parameters, with some control afforded to the user through tools such as the OpenAI Playground and the use of APIs (application programming interfaces).
- 8 OpenAI. (2023b), 'GPT-4 Technical Report'. *OpenAI Technical Report*, 27 March. Available online: <https://cdn.openai.com/papers/gpt-4.pdf> (accessed 9 August 2024).
- 9 OpenAI. (2023b), 'GPT-4 Technical Report'. *OpenAI Technical Report*, 27 March. Available online: <https://cdn.openai.com/papers/gpt-4.pdf> (accessed 9 August 2024), 1.

- 10 See the seminal work by Hans-Georg Gadamer, *Truth and Method* (London: Bloomsbury Revelations, 2013), and a survey of applications of the theories in Jeff Malpas and Hans-Helmuth Gander (eds.), *The Routledge Companion to Hermeneutics* (Abingdon, Oxon: Routledge, 2014).
- 11 Psychologist Daniel Kahneman captures the flavour of this style of thinking. See Daniel Kahneman, *Thinking, Fast and Slow* (London: Penguin, 2011).
- 12 Philosopher Stanley Fish articulates this well in his characterisation of professionalism. See Stanley Fish, *Doing What Comes Naturally: Change, Rhetoric, and the Practice of Theory in Literary and Legal Studies* (Durham, SC: Duke University Press, 1989).
- 13 OpenAI. (2023a), 'GPT-4 Technical Report'. arXiv, 15 March. Available online: <https://arxiv.org/abs/2303.08774> (accessed 25 July 2023).
- 14 Jen Patja Howell. (2023), 'ChatGPT Tells All'. *The Lawfare Podcast*, February 1. Available online: <https://www.lawfareblog.com/lawfare-podcast-chatgpt-tells-all> (accessed February 2, 2023).
- 15 Dan Hendrycks, et al., 'Aligning AI with Shared Human Values', *Proceedings of the International Conference on Learning Representations (ICLR)*, (2021), 1.

References

- Coyne, R. (2010), *The Tuning of Place: Sociable Spaces and Pervasive Digital Media*, Cambridge, MA: MIT Press.
- Fish, S. (1989), *Doing What Comes Naturally: Change, Rhetoric, and the Practice of Theory in Literary and Legal Studies*, Durham, SC: Duke University Press.
- Gadamer, H.-G. (2013), *Truth and Method*, trans. J. Weinsheimer and D. G. Marshall, London: Bloomsbury Revelations.
- Hendrycks, D., et al. (2021), 'Aligning AI with Shared Human Values', *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Kahneman, D. (2011), *Thinking, Fast and Slow*, London: Penguin.
- LaBelle, B. (2006), *Background Noise: Perspectives on Sound Art*, New York: Continuum.
- Malpas, J. and H.-H. Gander (eds). (2014), *The Routledge Companion to Hermeneutics*, Abingdon, Oxon: Routledge.
- Massey, D. (2005), *For Space*, London: Sage.
- McCullough, M. (2004), *Digital Ground: Architecture, Pervasive Computing, and Environmental Knowing*, Cambridge, MA: MIT Press.
- OpenAI. (2023a), 'Gpt-4 Technical Report'. arXiv, 15 March. Available online: <https://arxiv.org/abs/2303.08774> (accessed 25 July 2023).
- OpenAI. (2023b), 'Gpt-4 Technical Report'. OpenAI Technical Report, 27 March. Available online: <https://cdn.openai.com/papers/gpt-4.pdf> (accessed 9 August 2024).
- Patja Howell, J. (2023), 'Chatgpt Tells All'. *The Lawfare Podcast*, February 1. Available online: <https://www.lawfareblog.com/lawfare-podcast-chatgpt-tells-all> (accessed February 2, 2023).
- Perez-Gomez, A. (2016), *Attunement: Architectural Meaning after the Crisis of Modern Science*, Cambridge, MA: MIT Press.
- Schafer, R. M. (1977), *The Tuning of the World*, Toronto: McClelland & Stewart.
- Urry, J. (2007), *Mobilities*, Cambridge: Polity.

5

TOKENIZATION

Consider some non-English words: “urb,” “archi,” “an,” “tect,” “sm.” These character strings are components of words, but apart from “an” are unlikely to appear on their own in an English sentence. We saw in Chapter 3 that large language models (LLMs) deal not in words as the main units of language, but their subunits and combinations. LLMs typically deal with language units as *tokens* generated via statistical methods from a corpus of texts. Such tokens do not necessarily correspond to formal linguistic units such as syllables or phonemes. Conversations with an LLM present typically as sequences of words with punctuation, but such sequences are the result of operations on tokens, assembled to make up words and word combinations.

Seeking again for parallels with semiotic processes in cities, my aim in this short chapter is to align the use of tokens in automated natural language models with representations of urban data. Map data provides a good starting point for us to identify tokens in the urban context. Abstract graphical symbols on maps are obvious candidates as urban tokens, but here I will focus on textual labels. Google Maps, Apple Maps, OpenStreetMap

DOI: 10.4324/9781003535751-6

This chapter has been made available under a CC-BY-NC-ND 4.0 license.

among others produce maps covered in textual tokens: “café,” “station,” “plaza.” Sources on which these mapping platforms draw include the US Geological Survey (USGS) and the Ordnance Survey (OS) in the UK which identify several million “points of interest” with coordinates locating them on national map grids.¹ These labels refer to physical features or landmarks such as individual bridges, post boxes and bus stops, as well as retail outlets and businesses identified by name.

The OS classifies such feature items under an extensive system that starts with broad categories such as accommodation, eating and drinking, commercial services, attractions, sport and entertainment, education and health, and public infrastructure. It also hones the categories down to types of businesses, accommodation, entertainment facilities, commemorative plaques, etc, eventually landing on the names and coordinates pinpointing instances of these objects: Joyce’s Bakery on the Carlisle High Street, the commemorative plaque to John Venn in Kingston on Hull, and with grid coordinates. Planners and researchers will recruit such tokens in calculations involving census data generating for example mapped information about wealth and poverty distributions.²

Cities distribute their tokens in two and three dimensions. Of course, buildings, parks, roads, not to mention assemblies of people, are not just words but agents, instruments, tools, technologies and artefacts with physical presence in the functioning of cities. But tokens belong in the realm of signs, which I suggested in Chapter 1 are active instruments, and have agency in cities, or at least they contribute to networks of agent actions.

What I have outlined so far in this chapter is a taxonomic approach to urban data, building on and extending the simple idea of the token. Early AI systems, like Schank’s scripting model, would identify and codify such taxonomies as key elements in a “knowledge representation.”³ Aided by codified rules, such systems might make inferences to address queries such as “tell me which educational facilities are located in areas with low employment levels.” Such a system would draw on information stored in its data structures that primary schools, colleges and training establishments are types of educational facilities; education facilities are a type of institution; and institutions typically have buildings, serve large constituencies, and have further divisions into non-profit and for-profit. In AI parlance, such knowledge structures would take the form of complex “semantic networks,”⁴ through which an automated inference system might respond to queries or even address specific urban problems: mitigate congestion,

manage waste and site affordable housing. Such hierarchical information builds on the simple idea of a token to create a taxonomy. But as far as the human map reader-navigator is concerned, a map in 2D is simply lines, colours and a distribution of tokens as labels.

Taxonomic information, as well as writings about relevant theories in AI, may well exist amongst the texts in an LLM training set. But the benefit of LLMs is that their pattern identification processes operate much more flexibly than taxonomies. LLMs do not store hierarchical knowledge representations pertaining to information content. They operate with lexicons of tokens – flat structures of linguistic primitives. LLMs replace the challenge of exhaustive articulation of taxonomies with a different challenge, that of identifying tokens and attaching to them information about their relationships within the corpus of training texts, through numerical vectors, or embeddings, as described in Chapter 3. Internal LLM processes generate lexicons of tokens each with embeddings attached. The LLM will look up tokens from its lexicon as it processes word sequences, and it will operate on the relevant numerical token embedding.

Lexicons

In my graffiti example in Chapter 2 I demonstrated how the ChatGPT language model appears to process words and phrases such as “art form,” “semiotic coding,” “urban landscape,” “symbols,” “meanings,” “status quo” and “city.” These words and word combinations suggest that the LLM platform I used has access to an extensive *lexicon* of tokens, whether whole words, segments of words or word combinations. *Lexicality* relates to the ability of an event, place, system, set of objects or a phenomenon to be expressed in terms of words, and for those words to be organised systematically, as in a dictionary. Such word lists are ubiquitous and I include one towards the end of this book as a glossary.⁵ As a lexical example, the well-known Urban Dictionary is a crowdsourced online dictionary of slang words, such as “throw-ups,” “writer,” and “wildstyle” – urban terms related to graffiti, as it happens.

As I have shown, automated natural language models typically use lexicons not just of words but of tokens. The use of tokens contributes to their ability to process unfamiliar words and even to combine word fragments into new words, such as “citify,” “urbanique,” “metronest,” and provide plausible definitions for them.⁶

To reiterate, as processed in LLMs, tokens are strings of characters extracted from a corpus of texts – the set of texts an LLM system is trained on, though token lists can be shared across different corpuses. A tokenizer is a component of an LLM that extracts and optimises lists of tokens from bodies of text. Tokens do not need to be calculated independently for every large data set of texts. That said, LLMs that draw on data sets with specialised vocabularies may benefit from more specialised lexicons of tokens.

Tokens are more efficient than storing just whole words. For example, the word “oligarch” is made of two tokens “olig” and “arch”. The string “olig” also appears in “oligopoly” and several other words. The string “arch” appears in other words, such as “archway,” “anarchy,” “architecture.” Depending on how frequently these words appear in the training data, it is more efficient to store and process these word-fragment tokens rather than whole words, especially as a word may have several different endings – “s,” “ies”, “ed,” etc, which a tokenizer may also identify as tokens.

So, text tokenization involves identifying patterns of character sequences in strings of text. To illustrate, consider the string of words: “archival oligarch anarchic research.” I constructed a crude, brute force tokenizer in the python programming language to demonstrate to myself what the process would entail. To tokenize this sequence requires a tokenizer to identify the longest character sequences within this string, that also appear with greatest frequency. The string “arch” is the best candidate for a token in this modest word set. That string of characters occurs more frequently than any other sub-string. The string “rchi” occurs twice. Other strings such as “liga” occurs only once. Anything longer than a 4-character sequence in the string of words occurs only once. The lexicon of tokens should be comprehensive enough to reconstruct any substring in the corpus. Anything shorter than these sequences occurs too frequently, and the tokenization would deny a significant aspect of pattern detection to the training and prediction capabilities of the LLM.

It is unlikely that an LLM training model could pick up useful linguistic patterns if it had to train only on tokens that are just one character long, i.e. upper- and lower-case letters of the alphabet and punctuation marks. Ideally, the complete set of tokens should not exceed the set of all words or symbols derived from the corpus. There can be some redundancy among the token set, and this can help in capturing different subtleties in grammar. Tokens will vary significantly in length, ranging from a single character to

a group of words. The idea is that the tokenization procedure breaks down a text string with a granularity that enables the LLM to operate efficiently balanced with the need to maintain subtleties in the original text. Multiword expressions or compound words are typically kept as single tokens to retain their specific uses in sentences. The tokenization process also needs to be designed to be computationally efficient, and adaptable to different tasks, such as conversational AI, text generation in specific domains, or translation.

By now it will be apparent that tokenization for LLMs is a subtle art, and there are different methods. Tokenizers in common usage in LLM design bear the names Byte-Pair Encoding (BPE), SentencePiece and WordPiece.⁷ So a developer programming an LLM does not need to create their own tokenization algorithm or their own lexicon of tokens. These tokenizers typically sort words in length order starting with the longest and calculate word frequency. Words with high frequency become tokens. For lower frequency words the program will search for the occurrence of substrings, one, two, three and more characters at a time. For very large corpuses of training data, the process may sample the text to derive the initial lexicon of tokens. The process is iterative and may merge tokens.

Perceptual tokens

Tokenization brings to light important aspects of urban experience. As well as words and labels on maps, text-based communications, documents and LLM text interventions, urban citizens are invested in the urban world of people, bricks, concrete and other material things. Breaking the world into parts is relevant not only to text processing, but is an everyday, culturally bound practice across all aspects of human experience. The language philosopher Nelson Goodman showed that people structure the world using different symbol systems, and these systems guide our perception and understanding.⁸ To adopt the language of Phenomenology,⁹ the way we identify an object, thing, entity, phenomenon or event in any circumstance derives from practical engagement with our lifeworld, even if we do not have words for everything we are using or perceiving.

To labour our earlier example about graffiti, where we street-life neophytes see splashes of coloured paint on a wall as just vandalism, people with a different sensibility see an artful composition of “tags,” “throws” and “pieces.” What we see depends on our approach to the world of graffiti

and our inculcation into its language practices. So too, the thing you trip over on the pavement emerges for the moment as a drain, a misplaced slab or a curb. Buildings, roads and parks are similarly defined and bounded depending on whether your perception is framed by ownership, servicing, tourism, photography, catastrophe, complaint or flights of imagination. The Phenomenology of Martin Heidegger concurs, that things emerge as objects in our environment in the practical context of concerned dealings in the lifeworld.¹⁰ Fragmenting our world into parts is a practical matter. We perceive what we need to perceive for the task at hand.

Automated systems provide functionality analogous to this very human capacity to recognise objects and features in our world. It may also create objects in the visual field, as if by hallucination, a popular theme within experimental design studios that deploy generative AI to produce images from text prompts. I will leave that interesting subject for others to elaborate in the urban context.¹¹ Here, I will dwell on image recognition and its relationship to tokenization.

Feature detection

The text tokenization process I described earlier in this chapter provides an analogy with human perception and engagement with the lifeworld. The way tokenization breaks down language into manageable parts for processing in LLMs crudely parallels how we segment and classify the world through culturally and linguistically framed encounters. By now, most users of internet search, photo sharing and manipulation of their own ever-growing private image collections, will be aware of the ability of such systems to identify people and objects. I want to show that tokenization also parallels technologies in image recognition that employ related techniques. One of the key challenges of automating the detection of objects in photographs is to identify human beings, plants, animals, vehicles and road signs in a scene that are of interest to someone and for some purpose, such as navigating through city streets or improving accessibility.

Image recognition is a rich area of urban application from airport security to smartphone logins. A security camera installed in a garden to which I have infrequent access allows me to see what the camera sees while I am anywhere in the world via my smartphone. The system identifies the presence of human forms, and after a short training period identifies individuals by name. Consumer-level surveillance of this kind is now commonplace.

As well as fuelling low-grade suburban paranoia it introduces its own practices, such as watching wildlife and weather, and real-time image sharing. It relies on feature detection to pick out moving elements, such as people, pets and wildlife.

Images gathered by remote sensing from UAVs (unpiloted aerial vehicles), aircraft and satellites also lend themselves to automated feature detection. Cities are more than pictures, but maps and aerial photographs are well-established as means of identifying, analysing and planning urban interventions. After all, things get built in plan and urban spatial data is commonly plotted on a ground plane. Aerial photographs are amongst the tools by which cities are planned, navigated, surveyed, evaluated and designed. Such map data lends itself to many techniques of calculation and automation. My aim here is to show that tokenization is not just an abstract function hidden within LLMs, but parallels processes already familiar to urban analysis.

To demonstrate the similarities between tokenization of the kind employed in text analysis and feature detection in aerial photography, I developed a simple computer program that would demonstrate both. I have addressed tokenization in text processing in this chapter already. With a little tweaking, the one-dimensional procedure to tokenize text could be turned into a two-dimensional process to identify recurring patterns in photographs.

Early automated image recognition techniques would include tokenization. An image tokenizer would search for chunks of pixels that constitute patterns. These pixel patterns may recur in the same image, or across different images. To illustrate tokenization of mapped data (i.e. aerial images), I took an image of Edinburgh-Leith's docklands, reduced it to a 32x32 pixel grid with a 10-value grey scale (0–9) with 0 as black, 9 as white (Figure 5.1). Though it is hard for a human interpreter to read such a patchwork of greys, there are patterns in this data, and they repeat across the image or across several images of other places (Figure 5.2).

Patterns of varying pixel values are potentially of greatest interest as they could indicate recurring features in the landscape, such as a boat berthed at a wharf, vegetation, a shoreline. Here I am interested just in the computationally intensive process of identifying recurring patterns as tokens as part of a feature detection operation.

In this case, patches larger than 3x3 do not reveal any repetitions across this image. Note also that it is unlikely that patterns that are not contiguous will yield any tokens of interest, though image tokens do not need to all be rectangular.

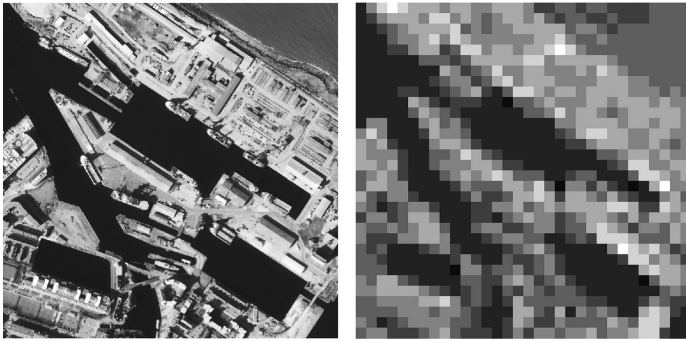


Figure 5.1 Aerial photograph of a dockland (© Getmapping Ltd via digimap.edina.ac.uk). The image to the right is a reduced version of the photograph made up of 32 x 32 pixels shaded with values 0–9 produced by the author.

Source: Author.

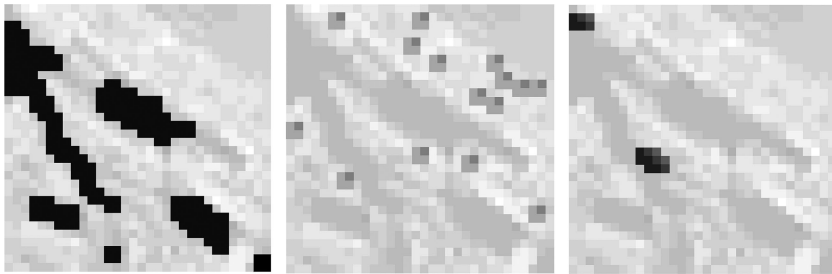


Figure 5.2 Results of a tokenization algorithm that identifies areas with identical patterns of shading. The first image shows all 4x4 patches with values of 0 (black). The second image shows 4x4 patches that share similar patterns of grey distributions. The third image identifies rarer 3x2 patches of grey distributions.

Source: Author.

There are other potential uses of detecting features using this “tokenizing” method, such as to identify and classify different types of land use (e.g., residential, commercial, industrial) by recognising recurring patterns in buildings, road layouts and green spaces. This could aid in the analysis of urban expansion and the planning of new developments. Tokenization could also help in analysing and comparing historical aerial photographs with contemporary ones to detect changes in the landscape, such as urban

growth or risks to historical sites. By detecting recurring patterns in aerial images of different terrains, this method could assist archaeologists in identifying ruins, burial mounds or old pathways that are not immediately visible to the naked eye.¹² The method could also be used to detect recurring patterns that signify cracks in roads, bridges or railways, enabling targeted maintenance efforts. These examples demonstrate how tokenization in image recognition can be applied across various fields to detect and analyse patterns that might otherwise go unnoticed.

My objective here though was simply to show that as a key element of natural language processing tokenization also belongs to human perception, automated image analysis and hence urban contexts. In this example, the algorithm I devised for iterating through combinations of pixel values was the same algorithm I used to identify tokens in the text example. The image data was in two dimensions, whereas the text data was linear and in one dimension. To repeat, the tokenization process, so important to automated natural language processing, mirrors the automated analysis of images, which in turn represents an attempt to mimic human perception and cognition, shadowing tokenization operations critical in the perception and practical engagement with urban environments.

As it happens, my brute force computation is supplanted in contemporary image processing by machine learning techniques, and feature detection is typically occluded in the hidden layers of a multi-layered neural network. Substantial applications of AI in urban remote sensing include earthquake, landslide and subsidence risk assessment, managing green spaces and water quality analysis.¹³

Summary and critique

In this chapter, I explored the concept of tokenization in large language models (LLMs). The radical proposition of the chapter was to relate tokenization to the representation of urban map data. I touched on how people identify tokens as features in perception and feature detection in automated image recognition. My attempt here is to demonstrate that the concepts of tokenization in sequences of text as processed in large language models is not so far from more familiar aspects of how we relate to our environment as rich spatial domains. In the process I explained how LLMs break texts into tokens – subunits and combinations of words – rather than

dealing with whole words. I used examples to show how text tokenization identifies patterns in character sequences, aiding in the detection of patterns within texts, and making language processing more manageable via lexicons of tokens. Through examples of identifying objects in aerial photography, I illustrated how tokenization in natural language processing mirror those in image analysis.

Behind this discussion lay some of my own modest programming that demonstrated to me at least that procedures for tokenizing sequences of characters can be adapted to two-dimensional images. Overall, I aimed to show that tokenization, a key element of LLMs, has significant applications and parallels in urban contexts, emphasising its importance in both natural language processing and urban analysis.

As in the case of the other chapters I asked ChatGPT to identify what readers might find controversial in this chapter. As before, it alighted on the parallels I am drawing, in this case between tokenization in LLMs, in map labelling, human perceptual cognition, automated feature detection in images and urban analysis. In summary, the platform criticised the chapter for asserting that tokenization parallels processes already familiar to urban analysis. The chapter might be construed as oversimplifying the complexities of urban data analysis.

My response follows two paths. One is to hold to such simplifications in the service of telling a story that relates otherwise disparate phenomena – large language models and the urban context. The second is to co-implicate the entire AI project that seeks to subjugate human cognition, language, social relationships and interactions with environment to instrumental framings and a dependence on numbers. For the purposes of my narrative here I will persist with the former, at least until Chapter 12 about the end of cities.

Notes

- 1 See <https://www.ordnancesurvey.co.uk>.
- 2 See Office of National Statistics maps at <https://ons.gov.uk/census/maps>.
- 3 See for example Frederick Hayes-Roth, et al., *Building Expert Systems* (Reading, MA: Addison-Wesley, 1983); Richard Coyne, et al., *Knowledge-Based Design Systems* (Reading, MA: Addison-Wesley, 1990).
- 4 Avron Barr and Edward A. Feigenbaum, *The Handbook of Artificial Intelligence, Volume 1* (London: Pitman, 1981), 180–89.
- 5 The University of Hong Kong publishes an online glossary of urban terms that includes phrases such as “Affordable housing,” “Gentrification” and “Walking cities.” See <https://urban-glossary.cuhk.edu.hk>.

- 6 ChatGPT provided these examples and suggested some definitions.
- 7 Byte-Pair Encoding (BPE) iteratively merges the most frequent pairs of characters or character sequences into tokens, optimising for frequent subword patterns. SentencePiece segments text into tokens without requiring spaces, using a model-based approach suitable for languages without clear word boundaries. WordPiece, used by LLMs such as BERT, builds tokens by iteratively combining the most frequent character sequences, balancing frequency and context information. This information was provided by ChatGPT.
- 8 Nelson Goodman, *Ways of Worldmaking* (Hassocks: Harvester Press, 1978).
- 9 Marcel Merleau-Ponty, *Phenomenology of Perception* (London: Routledge & Kegan Paul, 1962).
- 10 See Martin Heidegger, *Being and Time* (London: SCM Press, 1962) and Hubert L. Dreyfus, *Being-in-the-World: A Commentary on Heidegger's Being and Time, Division I* (Cambridge, MA: MIT Press, 1991).
- 11 Neil Leach and Matias del Campo (eds.), *Machine Hallucinations: Architecture and Artificial Intelligence (Architectural Design)* (Oxford: Wiley, 2022).
- 12 ChatGPT provided these examples.
- 13 Chang-Wook Lee, et al., *Artificial Intelligence Methods Applied to Urban Remote Sensing and GIS* (Basel, Switzerland: MDPI, 2021).

References

- Barr, A. and E. A. Feigenbaum. (1981), *The Handbook of Artificial Intelligence, Volume 1*, London: Pitman.
- Coyne, R., et al. (1990), *Knowledge-Based Design Systems*, Reading, MA: Addison-Wesley.
- Dreyfus, H. L. (1991), *Being-in-the-World: A Commentary on Heidegger's Being and Time, Division I*, Cambridge, MA: MIT Press.
- Goodman, N. (1978), *Ways of Worldmaking*, Hassocks, Sussex: Harvester Press.
- Hayes-Roth, F., D. A. Waterman and D. Lenat (eds). (1983), *Building Expert Systems*, Reading, MA: Addison-Wesley.
- Heidegger, M. (1962), *Being and Time*, trans. J. Macquarrie and E. Robinson, London: SCM Press.
- Leach, N. and M. del Campo (eds). (2022), *Machine Hallucinations: Architecture and Artificial Intelligence (Architectural Design)*, Oxford: Wiley.
- Lee, C.-W., H. Han, H. Lee and Y.-C. Park (eds). (2021), *Artificial Intelligence Methods Applied to Urban Remote Sensing and GIS*, Basel, Switzerland: MDPI.
- Merleau-Ponty, M. (1962), *Phenomenology of Perception*, London: Routledge & Kegan Paul.

6

THE MEANING AND USE OF WORDS

One of the basic theories underlying automated natural language processing is that you can derive the meaning of a word from its local context, simply the words around it in a phrase, sentence or other arrangement of coherent text. A good demonstration of this ability as an automated process is conversational AI’s capacity to deal effectively with synonyms. A natural language system built as a neural network will perform as if it recognises similarities between words when trained on those words and their local contexts.¹ Consider the routine cases where a speaker might legitimately replace one word with another, such as “city” with “metropolis,” or “accommodation” with “residence.” We humans learn that such substitutions are possible as we encounter large numbers of words and local contexts as we learn to use language.

Anyone who has interacted with conversational AI will be impressed with its grasp of synonyms. Likely such an AI platform will continue to respond sensibly even if you switch from interactions about “housing” to “accommodation” to “homes.” As we shall see, the language model

it deploys accomplishes this in part as it translates words and tokens to “semantic embeddings,” which are long numerical vectors, i.e., the coordinates of a word in a *multidimensional feature space*. The neural network is not actually trained on words, but long sequences of numbers as vectors. I say “words” here, but from the preceding chapter it should be apparent that we mean the more general concept of a token, which can be an individual word, a part of a word or a sequence of words.

Local context

As an introduction to this capability to deal in synonyms, consider the word game “blankety blanks” in which players must provide the missing words in a sentence. Points are awarded based on accurate or humorous guesses. LLMs seem to excel at the game. As well as substituting actual words for blank spaces, it is just as intriguing to guess the meaning of an obscure or made-up word from its context. Such non-dictionary words are common in branding, naming of apps and labelling new inventions (Instagram, Tumblr, Gmail, Hulu, Bing). Consider the made-up word “urbanique.” Here are some invented sentences that might use the word: “The city adopted an urbanique hue in the morning light”; “The neighbourhood offered a historic urbanique experience”; “Apartment prices soared in the city’s urbanique district.” From these sentences you might infer that “urbanique” has something to do with the high-value mood of a place, presumably with historic overtones (Figure 6.1). But as with most words, we do not need to look up a dictionary of meanings to use them effectively. Here we might infer the meaning of a non-existent word from its context. In this case we are also helped by the components of this made-up word: “urban” and “ique,” which, as tokens, have their own contexts of use.

It is a basic tenet of language usage that people can incorporate words into their own language patterns by having been exposed to them in different contexts. That tactic for deriving word meanings applies to ambiguous but widely accepted dictionary words like “post-digital,” “hipster,” and “deconstruction.” These words are familiar enough but are used in different ways in varied contexts by a range of individuals and groups. The lack of a settled definition does not prohibit their use. Learning from examples of word usage applies to ordinary words as well: “city,” “home,” “graffiti” and “time-share.”



Figure 6.1 An atmospheric, historical “urbanique” quarter. Rue Richer, Paris.

Source: Author.

According to the language theory of John Firth, that we learn word usage from exposure to many words in local contexts provides a basic model of how humans acquire language competence.² Words inherit each other’s associations with other words. He summarised the concept concisely: “You shall know a word by the company it keeps.”³

Early models of natural language processing assumed the need for dictionaries of meanings and word relationships. But conversational AI methods typically establish these relationships from scratch. They start by identifying and positioning words in the context of other words “whose company they keep,” with which they appear frequently in sentences. They compile these tokens and their embeddings into lexicons.

Any competent speaker would probably be able to deploy a word like “urbanique” without much reflection, but when required to define it, could do so by identifying the words around it in its various contexts: “hue,” “morning,” “historic,” “district,” and infer that it had something to

do with these terms, especially if they keep cropping up in the context of the word “urbanique.” Words and phrases such as “hue,” “morning,” “historic,” and “district,” also appear within their own sets of contexts. This model of word meanings based on associations between words in context is known as “distributional semantics.”⁴

The relationships between words becomes even more intricate if we think that the same word can have different meanings, i.e., it appears in very different use contexts. As an illustration consider the word “square”: “The protesters gathered in the square”; “They hung the square banner above the entrance”; “Your ideas about cycle paths square with local planning policy.” We have little difficulty in using the same word in different contexts, and even explaining why the same word is used in those contexts. In the case of “square” each use derives from something to do with geometry. Even a single sentence can draw on mixed contexts: “Park the car near the park”; “I walked up the bank to withdraw money from the bank.” Both “park” and “bank” have multiple meanings and contexts of use.

One of the main points from this distributional theory is that words appearing in similar contexts have similar meanings. If we read “Prices soared in the city’s urbanique district,” and elsewhere that they “soared in the city’s fashionable district” and “soared in the city’s historic district” then we might feel there is some consonance between “urbanique” “fashionable” and “historic.” In some simple cases the words could be substituted for one another, though that does not need to be the case. Word uses are more subtle than that.

Note that this type of analysis considers only the content of a corpus of texts. There is no need for dictionary definitions. There is no recourse to the wider context in which the words occur, their ontologies, networks of associations, grammars or other explanations of how words relate to one another, nor to their derivations. We may use “park” as both a verb and a noun, think of “park the car” and “pleasure park” as related, and recall that city squares are sometimes geometrically square, but in the case of LLMs these relationships emerge as patterns in the textual data, not from some external structures or sources, or from a taxonomy as discussed in Chapter 5 on tokenization.

The theory of distributional semantics presupposes that wider matters of context are accounted for in the training set, the corpus, the original source. A sceptic might say that an automated language model has no information

about the history of so-called “pleasure parks” in its model and the wider context of their position in the social and political contexts of their times. But if there are texts in the corpus that allude to this broader context then there is a chance that such “knowledge” will spill out from a conversational AI trained on texts that include such information. The language model will reflect the knowledge and grammatical and ontological patterns evident in the language use of the authors of those texts.⁵

We could say that AI language models of this kind are “parasitic” on the competences of the authors contributing to the source data. As I stated in Chapter 4 on tuning, lest we think that conversational AI has to re-invent language, it is worth noting that such language models leverage off human intelligence in so far as the texts and symbol systems of the training data give expression to human intelligence.

This shortcut into human-style intelligence poses some risks. Barring the effects of fine-tuning, interactions in the manner of conversational AI could be trained exclusively on nothing but subversive texts, nonsense, lies, anarchy or symbols in extraterrestrial languages, though how you might “converse” with such alien models would pose a challenge. The core of the LLM methods is indifferent to the content or quality of the corpus.

We need make three further points to labour this method of establishing word meanings. First, at best, associations and similarities between words are likely to be contingent, probabilistic and “fuzzy” rather than clear cut and deterministic. The strong claim of this theory is that word usage is elastic. Conversational AI processes numbers. So, such models process outcomes of semantic processes as probabilities.

Second, such associations will be stronger, more sophisticated and more interesting if they are derived from a vast number of words in a wide range of contexts – in other words, a large corpus of coherent texts. I am alluding here to patterns of associations between very large numbers of words within wide-ranging contexts, leading to the use of a neural network to process patterns in quantities of textual data. LLMs operationalise that information in a way that facilitates natural language generation and conversational AI. Though neural networks rely on quantity and repetitions of exposure to written words, it is worth noting that there need be no implication that human beings learn similarly from such linguistic profligacy.

The third point is that the result of automated language acquisition is not just learning to produce text or even carry out a conversation, but to give the appearance of knowing things, of conveying information. While an

automated natural language processing system is training to use language, it is doing so in a way that embeds patterns that enable it to yield information content. As it is acquiring the ability to construct coherent sentences, it is acquiring the capacity to recall, package and deliver information. In so far as these systems mirror human cognition, we could say there is really no such thing as contentless language acquisition. Learning to speak elides with learning to have something to say.

Numbers and meanings

The methods by which LLMs process this kind of language use is numerical and mathematical. That could be said of almost any processing by digital computers, but in the case of LLMs, strings of text are reduced rapidly to tokens, and then to vectors (semantic embeddings). Such vectors are subsequently processed to account for the positions of the tokens in the training texts and coding of probabilistic calculations about the relative importance of tokens in the data stream (attentional encoding). Once the inputs are provided to a neural network, the model is trained to generate numerical outputs that guide the selection of tokens from the lexicon to form coherent sentences.

Semantic embedding operates as sets of coordinates indicating the closeness of words in terms of meanings, shared importance or relevance. For example, depending on the corpus, “hue” and “morning” might be closer than “hue” and “bulldozer.” It is as if you can position words on a graph with x and y coordinates and see or calculate how close words are to each other. Considering the fluidity and complexity of word usage, it is unlikely that a two-dimensional graph would be sufficient to plot the relative positions of words in this space of word proximities. As we have seen, such relationships are not established by some pre-existing understanding of geography, physics, “natural kinds,” typologies or material evidence, but established by algorithmic analysis of the words in a particular corpus. I will keep using the term “words” in this section, though it is apparent from our previous discussions that LLMs deal in tokens, only some of which may be distinct words. LLMs operate with tokens that often present as arbitrary looking character sequences, including punctuation and mathematical symbols. That contributes to the opacity of the semantic functions claimed of these systems. Their operations will not be as transparent as suggested in the examples in this chapter.

From here on I will also describe some numerical operations that some readers may choose to skim over. My main point here is that the manufacture of meaning structures in LLMs is highly numerical. It is sometimes easy to forget that behind the sometimes personable and inventive responses provided by conversational AI lies the near instantaneous processing of vast arrays of numbers.

As mentioned in Chapter 3, neural networks are made up of a set of input nodes and a set of output nodes. Mostly they are designed with layers of hidden nodes which enhance the ability of the network to respond to complex patterns in input-output data. I include a worked example here, in which each input node is connected to every hidden node, that are in turn connected to every output node. Every input node corresponds to a word (token) in the lexicon. Every output node also corresponds to a token from the lexicon. There can be any number of nodes in the hidden layer, though even just two nodes will produce better performance in detecting patterns than a network with no hidden layers.⁶

A hidden layer with two nodes is simple enough and sufficient to illustrate semantic embedding, the main means by which conversational AI can deal with synonyms and definitions. The identification of semantic embeddings is not as a special process or module, but a characteristic in the behaviour of a neural network dealing with inputs and outputs. Deploying a hidden layer of two nodes also makes it possible to visualise how words cluster on a graph with 2 axes, illustrating the company they keep. Here is a simple example.

I could illustrate the process with word proximities in simple sentences. But I will recruit a spatial example, as it helps my goal of aligning LLMs with spatial relationships as in the architectures of the city. Consider three everyday items commonly found in proximity to one another: a mobile phone, a charger and earbuds. In this context, proximity could mean electrical or electromagnetic proximity, but here I will just consider how these objects happen to be positioned spatially on my desk.

Figure 6.2 shows one instance where these objects are in the same space, i.e. they are *near* to one another. Reducing this arrangement to three spatial relationships, the phone is near to the earbuds, the charger is near the phone, and the earbuds are near the phone. That is the case just now, but there will be other instances where one or more of them is absent, or none of them is near the other two. I can turn this arrangement into a simple neural network training process by conducting a survey of many actual



Figure 6.2 Three everyday items commonly found *near* to one another: recharger, phone and earbuds.

Source: Author.

configurations of these three items and their proximity relationship, as nearness. Note that in this example, two items are either near each other or not. It is a binary relationship.

To train on these relationships, I constructed a simple neural network with 3 inputs and 3 outputs, and 2 hidden units. The left side of the diagram treats each object as input. The right side shows the same items, but as outputs. So, if a phone is near some earbuds in one of the training examples, then the phone is the input with a value of 1.0 and the earbuds also have a value of 1.0 as output. The rest of the inputs and outputs in that case are 0.0. In other words, the configuration of the three items is treated two at a time as proximity pairs.

For this simple worked example, I trained the network on ten cases. Imagine I walked around an open-plan office and took photographs of ten people's desktops, and then recorded each instance in which a phone is near to some earbuds, a charger is near a phone and some earbuds are near a phone. That would be about 20–30 proximity pairs, presented to the neural network as input and output pairs. The main payoff from this exercise is that I can use the network to predict likely proximity arrangements for this office, without revisiting the desks. If someone tells me there are earbuds on a desk, what is the likelihood that there will be a phone nearby?

The training algorithm makes fine adjustments to network parameters (weights and biases) to enable it to reproduce the main input and output patterns in the training data. It does so via looped iterations. In my example, the inputs and outputs have the same names, and the training data consists

of a set of proximity pairs as input-output pairings: Phone – Earbuds, Charger – Phone, Earbuds – Phone.

Figure 6.3 shows a visualisation of the neural network and its parameters after training on this small set of proximity relationships. Once the network is trained, we should be able to predict what is most likely to be near the Earbuds. We might expect the answer to be a Phone, with a lesser probability for Charger.

Figure 6.4 shows how the activation of the Earbuds node is propagated through the network to produce a series of output probabilities. The activations are propagated by an algorithm that adds the weights leading into a node, adding a bias factor, and applying a simple condition to the output stating that the output from a node is propagated if greater than 0.0, otherwise the output is 0.0.

Indeed, following this method, the highest scoring output is the Earbud proximal to itself, with the Phone as second, and the Charger is the least likely to be nearby. The fact that these probabilities are close in this example reflects the sparsity of the training data. The parameters in the network would be more robust, reliable and interesting in their predictions with a much larger set of training examples, i.e. data gathered from many observations of the positioning of these three items and for many different users.

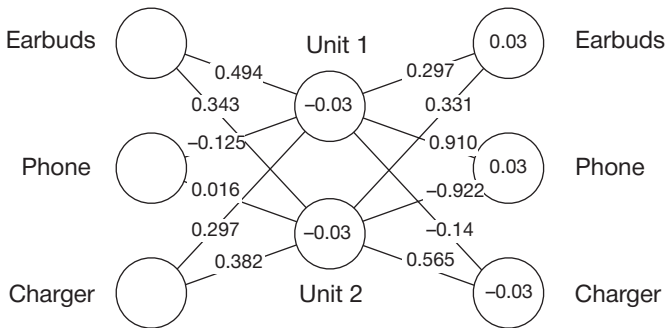


Figure 6.3 Network parameters after training on about 20 proximity pairs, e.g. Earbuds – Phone, Phone – Charger, Charger – Phone, etc. The numbers at the nodes (circles) are biases; the numbers on the connecting arcs (lines) are weights. A looped algorithm calculates and adjusts the weights and biases incrementally to retain and reproduce the various input and output patterns encountered during training.

Source: Author.

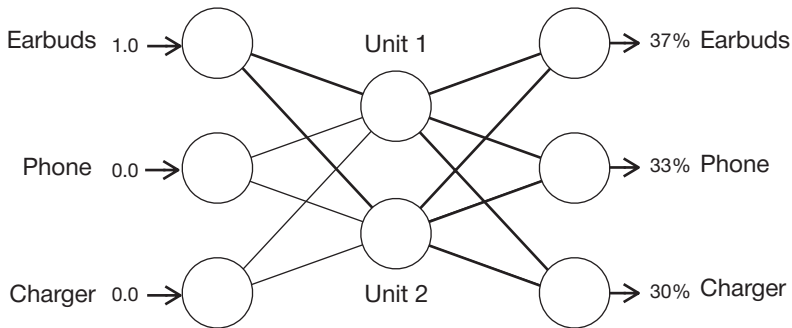


Figure 6.4 Notional diagram showing the propagation of input value of 1.0 for Earbuds through the network under the influence of the parameters of the trained network to produce a series of probabilities against each output node.

Source: Author.

It is also worth noting that here I am using a neural network to produce a result that could be achieved via simple statistical calculation, but the example serves our purpose as a demonstration of how a neural network processes such information. The method also becomes more interesting if scaled up to include other desktop items, such as computer, keyboard, pen, water jug, lunch box, magazine, etc.

However, what is of most interest here is not the scale of the network or the outputs but the trained weights that feed into the two hidden units (0.494, 0.343), (-0.125, 0.016), (0.297, 0.382). These can be plotted and labelled as coordinates on a graph as indicated in Figure 6.5.

Figure 6.5 shows how the Charger and the Earbuds are generally closer together (more likely to be found together) than the Phone, though the Charger is generally closer to the Phone than the Earbuds, according to the data set.

There are other relationships and dimensions in the trained neural network, including the biases in any hidden layers, and weights leading from the hidden units to the outputs. The training algorithm constructs and fine tunes the weights relative to one another. But the weights between the input and hidden layers are sufficient to capture something about the relationships and patterns identified from the input data. The weights effectively encode the correlations present amongst the inputs. Another way of stating this is to see the dimensions expressed in the hidden layer as serving to constrain the network to adopt a compressed representation of the input data. The hidden layer contains a distilled representation of the input, likely

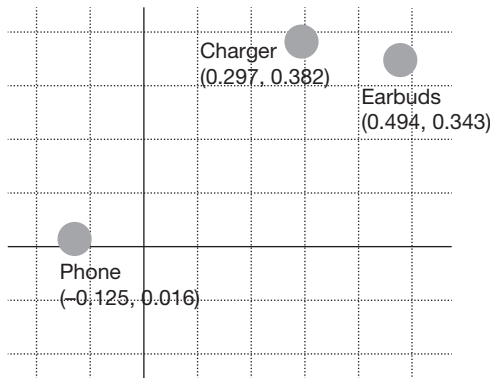


Figure 6.5 Plot of the hidden unit values for the three items. The horizontal axis shows the input values to hidden Unit 1, and the vertical axis shows the values for hidden Unit 2. Items closer together in this cluster diagram are most likely to be found in each other's presence according to the training cases.

Source: Author.

to capture its most critical features and relationships. The weights leading into these two hidden units are therefore good candidates for capturing relationships within the input data. These are the values that generally are used as semantic embedding vectors in LLMs.

In fact, these are vectors of the kind stored in an LLM-style lexicon: Earbuds (0.494, 0.343), Phone (-0.125, 0.016), and Charger (0.297, 0.382). In my physical desktop example, they give an indication of the relationships between physical elements. In the case of LLMs for text processing (rather than spatial relationships), such vectors give an indication of the relationships between tokens in terms of their proximity in sentences in the training data. But in an LLM, there are thousands of tokens with graded proximity relationships in terms of how far apart tokens are from each other in the training text. There are many more hidden nodes in an LLM's semantic embedding operation, and therefore many more dimensions.

In the case of LLMs, tokens that are closer together in semantic space, rather than just within sentences, tend to share similar meanings and may be used interchangeably in context. However, the use of semantic vectors extends beyond simple word substitution. Through additional transformations that account for word order and attention mechanisms, these vectors capture patterns in the underlying data, enabling the model to make plausible predictions about the next word in a sequence of words, as in a conversation.

This method of encoding semantic relationships works well and has been deployed in many natural language processing use cases. Restricting the model to two hidden units as in my simple example conveniently generates two coordinates for each input. In other words, we can plot the inputs relative to one another on a two-dimensional graph. More than two would give us a richer representation of these relationships, but these would be harder to visualise in the case of our simple spatial example.

I have assumed here that semantic embeddings are calculated during an LLM's training. In fact, the lexicon and its semantic vectors are generally calculated during a pre-training process, and even on a subset, or a different set, of data than the network is trained on. Some such files of semantic embeddings are also available as opensource. For example, in various experiments of my own I have used Facebook's semantic embedding data, derived from social media text data, and available in different dimensions.⁷

Summary and review

In this chapter, I examined how conversational AI systems, like neural networks, capture and process meanings using word embeddings. I discussed how these embeddings represent words as coordinates in a feature space, allowing the AI to process synonyms and simulate speech. By demonstrating how language models derive meanings from context, I illustrated the concept of distributional semantics, emphasising how words in similar contexts share meanings. I also touched on the numerical basis of neural networks and how they generate semantic embeddings, illustrating their function with a simple worked example involved three items commonly found on a desktop (a phone, a charger and earbuds).

Three key points emerge from this discussion. First, the meanings of words equate to how they are used in sentences, which provide multiple contexts of use. In LLMs, words are processed in terms of the other words around them. As I stated in Chapter 3, we need be less concerned with what words mean, than what actions follow or what responses they engender. In the case of an LLM, semantic embeddings contribute to neural network outputs that are predictions for what words may follow, and hence the responses they generate to prior word sequences and user prompts. Second, linguistic processes here are deeply numerical, rendering them opaque to human scrutiny. You can only assess the quality of an LLM's simulation of meaning-laden

interaction by how well it performs, i.e., what sort of outputs it generates. Third, insofar as we hold on to the concept of meaning in LLMs, it resides with the authors of the source texts and their readers, which the designers of LLMs have done their best to render manipulable and reproducible via patterns in word usage. This discussion sets the stage for a more detailed consideration of context in LLMs and parallels with the condition in cities

Notes

- 1 Thomas K. Landauer and Susan T. Dumais, 'A Solution to Plato's Problem: The Latent Semantic Analysis Theory of Acquisition, Induction, and Representation of Knowledge', *Psychological Review*, 104, no. 2 (1997): 211–40.
- 2 J. R. Firth. 'A Synopsis of Linguistic Theory 1930–1955', in *Selected Papers of J.R. Firth 1952–1959*, ed. F. R. Palmer (London: Longmans, 1968), 168–205.
- 3 J. R. Firth. 'A Synopsis of Linguistic Theory 1930–1955', in *Selected Papers of J.R. Firth 1952–1959*, ed. F. R. Palmer (London: Longmans, 1968), 11.
- 4 Alessandro Sahlgren Lenci, *Distributional Semantics* (Cambridge: Cambridge University Press, 2023), 3.
- 5 Though I am speaking here of everyday language, it is worth noting that automated language models with the capacity to write computer programs are also supplied with extra-linguistic content, such as tokens within maths formulas and procedures, which have their own vocabularies, structures and grammars.
- 6 For a helpful explanation of the method see Eligijus Bujokas, 'Creating Word Embeddings: Coding the Word2Vec Algorithm in Python using Deep Learning', *Towards Data Science*, March 5, 2020. Available online: <https://towardsdatascience.com/creating-word-embeddings-coding-the-word2vec-algorithm-in-python-using-deep-learning-b337d0ba17a8> (accessed 16 January 2023).
- 7 These are trained on text sources such as Wikipedia and a repository of web data at commoncrawl.org and are available for over 150 languages.

References

- Bujokas, E. (2020), 'Creating Word Embeddings: Coding the Word2Vec Algorithm in Python Using Deep Learning'. *Towards Data Science*, March 5, 2020. Available online: <https://towardsdatascience.com/creating-word-embeddings-coding-the-word2vec-algorithm-in-python-using-deep-learning-b337d0ba17a8> (accessed 16 January 2023).
- Firth, J. R. (1968), 'A Synopsis of Linguistic Theory 1930–1955', in F. R. Palmer (ed), *Selected Papers of J.R. Firth 1952–1959*: 168–205, London: Longmans.
- Landauer, T. K. and S. T. Dumais. (1997), 'A Solution to Plato's Problem: The Latent Semantic Analysis Theory of Acquisition, Induction, and Representation of Knowledge', *Psychological Review*, 104 (2): 211–40.
- Lenci, A. S. (2023), *Distributional Semantics*, Cambridge: Cambridge University Press.

7

CONTEXT WINDOWS

At the start of Chapter 1, I recounted the brief conversational exchange with ChatGPT about urban graffiti. The response was convincing not least as the conversational AI seemed to grasp the *context* of my proposition: “Graffiti is both an art form and a means of disrupting the semiotic coding of the city.” Its response about “urban landscape,” “symbols” and “meanings” indicated access to a vocabulary pertaining to that wider context. My emphasis on context is obvious from the title of this book *AI and Language in the Urban Context*. The word “context” originates from the Latin “contextus,” which means a “joining together” or “weaving together.”¹ It shares its meaning with the *texture* of a fabric, which in turn relates to text as writing. To establish the context of a word is to examine it in the overall weave of the text around it. A consideration of context of course reaches out to the wider social, historical and political fabric in which any word, phrase, sentence or document is situated and applied.

One of the major challenges for developers of automated natural language processing platforms has been to factor context into the calculations and predictions of the models. In fact, that is one of the major accomplishments

DOI: 10.4324/9781003535751-8

This chapter has been made available under a CC-BY-NC-ND 4.0 license.

of the current generation of natural language models (e.g. the Transformer model). As we have seen, they derive and process the contexts of linguistic components (words or tokens) and store that information numerically in a way that enables them to reconstruct linguistic patterns and thereby to sustain the flow of coherent text when interacting with a user. They also process context information about the current conversation with the user. Hence, they exhibit the ability to consider what both you and the platform produced earlier in the conversation.

Urban contexts

Before embarking on a discussion of context in large language models (LLMs), I would like to emphasise how scholars and critics often refer to context in cities. It is common to discuss cities in terms of *networks*. Any part of a network is in the *context* of the other interconnected nodes around it. The *network* is a well-used metaphor that also draws on early technologies of threads, ropes and weaving.² Plato wrote about the social condition of the Polis as an intertwining of “warp and woof.”³ I am fond of quoting the Stoic philosopher and Roman Emperor Marcus Aurelius (AD 121–180) not least as his terminology seems to presage a technologised world under the thrall of the World Wide Web. He summed up human existence in terms of an organic network: “Remark the intricacy of the skein, the complexity of the web.”⁴

The city as network speaks to the city as a site of communication, exchange and action, a site of shared meanings, knowledge and perhaps intelligence. Interconnected physical devices, signage, buildings and streets reinforce the network notion. In reasserting the liveliness of city streets from a human perspective, urbanist Jane Jacobs wrote about the city as “an intricate, almost unconscious, network of voluntary controls”⁵ promoted by its citizens (see Figure 7.1). Showing greater affinity with cybernetics and machine processes, the design theorist Christopher Alexander described the city as a semi-lattice: “the structure of a complex fabric; it is the structure of living things, of great paintings and symphonies.”⁶ He embellished the complexities of the city with a further mathematical metaphor, that of overlapping sets in Venn diagrams, translating to interconnected networks rather than neatly ordered hierarchies with sharply delineated boundaries. Common in urban discourse, networks are recruited to champion the causes of diversity, ambiguity and a recognition of varying contexts.⁷



Figure 7.1 A busy urban scene suggesting an overlay of different networks of commerce, communication, transport, construction, power supply, political imperatives and sociability. Kibera informal settlement. Nairobi, Kenya.

Source: Author.

Social urban networks of the kind referenced by Jacobs and others are presumably complicated, dynamic, multi-layered and indeterminate. Following the precepts of cybernetics, systems analysts, programmers and linguists will operationalise networks as diagrams, flow charts and procedures: flows of signals through nodes and their connecting links. But the network's existence as a metaphor of rich and dynamic social interactions persists.⁸ The association between urban networks and neural networks will not have escaped observant readers, reinforcing further the parallels between AI techniques and cities.

Recap on LLM operations

As I have shown, a neural network of the kind deployed in LLMs has two main operations. The network undergoes training. This is where a neural network model is exposed to input and output data to assign numerical parameters so that it can make predictions based in patterns that emerge

from those parameters: predicting what is the most likely word (token) that would follow the last word sequence in the flow of text.

Training involves multiple iterations through sets of ordered text data to adjust the values on nodes and links in the neural network to enable word prediction. This training occurs before the user encounters the model. It draws on extremely large computational resources and takes many days of processing time with powerful computers. The network so trained can then be available for use through shared servers to many independent users. This is what the user sees, the simulation phase, where a user prompts the model with text, and the neural network returns a series of words as if completing your sentence, to answer a question, offer a summary, correct grammar, and translate between languages and styles.

This division into training and simulation is a simplification, of course. As I have explained, prior to the main training operation, there is an initial step involving the creation of token embeddings. This step involves calculating or otherwise acquiring a lexicon of tokens and assigning a vector representation (i.e., a long list of numbers) to each token in the training data. The vector is typically unique to each instance of the token in the training data and is derived from a pre-calculated semantic embedding for the token, which is stored in a lexicon as a lookup table. Crucially, during the training process, several mathematical modifications are applied to the values in these semantic embedding vectors. These modifications are calculated based on information about the position of the token in the stream of text provided by the training data, with further adjustments derived from calculations about the importance of certain tokens relative to others in a sentence. These adjustments result in vectors that represent an attention score for each token.

One important parameter that remains constant throughout the training process is the dimension of the embedding vectors. These vectors – semantic, positional, and attentional – all share the same dimensionality, which refers to the number of parameters or components in each vector. These vectors are combined mathematically, often through operations such as addition and multiplication, to produce a final input vector for the neural network during training. A larger dimensionality for the input vector can enhance the predictive capability of the neural network, though this comes at the cost of increased computational resources and training time. Typically, the dimensionality is several thousand for input vectors in large language models.

As explained in the previous chapter, during training, the neural network adjusts its parameters as it processes each token's vector. If the vector

dimension is 2,048, this means there are 2,048 input nodes in the neural network. Imagine these as lights flickering at different luminance levels as the stream of token vectors passes into the network, representing the continuous flow of information being processed.⁹

Context window

I have already mentioned that these vector calculations for each token in the data stream takes account of the vector values of the tokens around it. That is the second vital parameter that influences the network's performance, namely the size of the *context window*. The context window consists of the tokens that appear before the current token being considered, e.g. 8,000 tokens. That is about 7,000 words, the length of an academic article. The size of the context window is an indicator of the scale of the computation required to calculate the position of the token as a vector and its attention score vectors that modify the semantic embedding for each token. Imagine the data as sliding past a context window, advancing one token at a time as the model recalculates the input vector.

The simulation operation also involves a stream of tokens that the trained neural network must encounter in real time, as a user interacts, i.e. converses, with the model. Vector manipulations must mirror those that occurred during training. There is a lookup procedure that retrieves the semantic embedding vector for each token. It calculates the modifications based on positional and attentional encoding in real time. The context window size is an upper limit on the stream of tokens considered in these operations.

During simulation, the context window acts as a kind of short-term memory, playing a role in the process of calculating semantic, positional and attentional factors to generate predictions. Once the stream of tokens has filled up the context window, tokens before the window drop out of any calculations. In other words, any information communicated before this window no longer influences the current stream of conversational interaction.¹⁰

The urban context

The idea of the context window echoes certain intuitions about everyday language use, notably limits on our ability to recall or apply insights (or patterns) from earlier parts of a conversation, or even prior conversations.

From interpretation theory I think of the spatial concept of the *horizon*, an inevitable limit to comprehension imposed by a particular viewpoint.¹¹ That horizon defines a complex landscape influenced by the interpreter's personal history, community and background of experience, serving as a kind of context window. Move about and the area enclosed by the horizon shifts as we encounter the views of others and adjust or resists changes to our horizon.

The context window in LLMs diminishes the influence of tokens considered earlier during processing. Our experience of cities operates with similar sequential dependencies. Consider how cities are under the influence of deliberate interventions or “design moves.” For example, various actors in the development of a civic square introduce a series of actions and interventions, each influencing the next. Earlier interventions outside of a temporal context window are forgotten or deemed less relevant. Such design moves might start with defining the basic parameters of the civic square (site location, boundaries, initial surveys and environmental factors). Early interactions form the base context but may be overridden by newer, more relevant inputs later on. New actors such as local government, community members and businesses begin adding elements: benches, trees, public art and concessions. These inputs are shaped by and respond to the evolving design, but they may shift the focus of the space. Members of the public exert influence. Over time, new needs (like more seating or shade) drive new interventions. The design is increasingly shaped by current and future uses, with diminishing focus on the square's original form. The LLM's context window mirrors how urban design is continuously negotiated, with relevance moving toward the most recent interventions.

To explore further the relationship between LLM methods and urban contexts, I will expand on the simple desktop scenario described in Chapter 6, which involved a phone, a charger and a set of earbuds. Urban spaces are far more complex, but to simplify our analysis, I will focus on how spaces are organised within an urban neighbourhood, sidestepping how those spatial arrangements came about, shaped by a myriad of social, economic and cultural factors.

The idea of the context window also comes into play as we think of a journey through a city. To align LLM processing with an urban condition, I will focus on a particular neighbourhood (Figure 7.2) a heritage area in the central area of the city of Edinburgh. As we have seen, training a



Figure 7.2 Map of a heritage neighbourhood of single to five storey mixed use development. Shaded areas are green spaces.

Source: Author.

neural network (NN) involves automatically adjusting numerical weights and thresholds (biases) to account for all input and output pairs presented to a neural network. After training on these input-output pairings the NN should reproduce the appropriate output when presented again with any one of the inputs.¹²

Figure 7.3 shows the locations of some physical objects, that I will call “landmarks,” in the neighbourhood map. These data points are derived from an ordnance survey (O/S) map as described in Chapter 5. The data points are shaded according to category of feature. The spaces between the points are streets, residences and private gardens. Some of the landmarks line up along streets. In some cases, the points are the centroids of areas of land occupied by a park or school. These simplifications are typical of the approximations required of map data processing, but also of word usage in the representation of urban relationships.



Figure 7.3 Map derived from data about Figure 7.2 of physical neighbourhood landmarks as data points shaded according to category of landmark: retail, community, park, business, public service, etc.

Source: Author.

To repeat, my goal here is to demonstrate parallels between urban mapping techniques and natural language processing. To that end I have utilised a neural network to model in a simple way the mapped structure of this neighbourhood.

In language as spoken and written, sequences of tokens are one-dimensional, progressing from one token (or word) to the next. In my neighbourhood map the landmarks are arranged spatially in two dimensions. Landmarks in city maps are arranged in 2D, but people move through landscapes, experiencing one space, item or feature after another. In these respects, linear movement through space bears similarities to the flow of a sentence.¹³ We pass from one landmark to the next, the residence to the shop, to the public open space, to the school, etc. As I am drawing parallels between language models and arrangements in mapped space, it is

appropriate to think about how to derive a route through my neighbourhood based on how close landmarks are to one another on the map: their contexts and proximities. The result is a candidate for a kind of cognitive map of the area,¹⁴ an urban journey derived from patterns amongst landmarks as labelled map coordinates.

In this map exercise we can calculate the distance between any two coordinate points. I define a point as near another if it is within a 100-metre radius of it, ignoring the street layout. There are 570 data points in the map of Figure 7.2. It is a simple matter to devise a computer program to calculate the distance between every two points and to filter them according to points that are within a certain distance from each other, such as 100 metres. Figure 7.4 shows each data point circumscribed by a circle of radius 100 metres. That defines an immediate spatial context for each point. It also provides a means of identifying “proximity pairs.” They are pairs

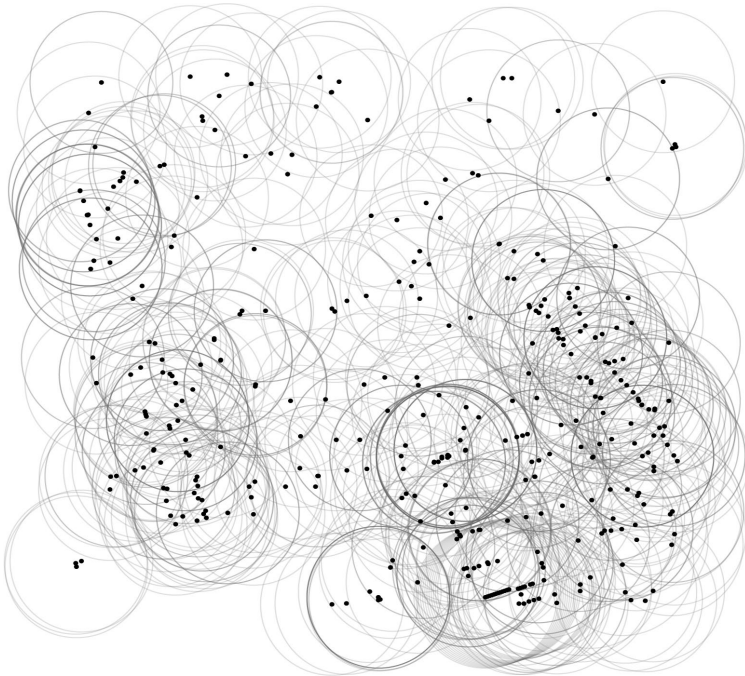


Figure 7.4 The data points of Figure 7.3 each shown with a circle of radius 100 metres to circumscribe neighbouring landmarks, thereby defining a context for each landmark.

Source: Author.

of landmarks within 100 metres of each other. It happens that for this neighbourhood there are 4,560 such pairs.

I devised a program to calculate the shortest distance between each named landmark in metres. From this data, a program identifies proximity pairs. I have also categorised each unique landmark under 21 headings: e.g., Business, Retail, Commemorative, School, Bus stop, Open space. Let's call these "feature tokens." The program then substitutes the feature token for the landmark instance that appears on the map. So, the file of final proximity pairs shows pairs of tokens that can be found close to one another in this neighbourhood, e.g. Retail – Electric service, Retail – Sport, Retail – Community, Hotel – Business, Hotel – Business, Hotel – Electric service, Hotel – Commemorative, Hotel – Community.

I then use the token pairings as input-output pairs for training a neural network model. So, the network takes Retail as input and is trained to generate Electric service as output. It is also trained to deliver Sport in response to Retail as input. Where input-output pairs are repeated, e.g., Hotel – Business, that serves to strengthen that pairing. The training algorithm runs through the entire input-output pairs to incrementally adjust the parameters in the neural network model. So that is 4,560 input-output sets of calculations. The training algorithm repeats this training process with the same data multiple times to ensure that the model does not lose its ability to reproduce the input-output relationship of any of the pairs, and so that it is able to reconstruct patterns in the data. It performs this repeat iteration, known as an epoch, 500 times.

While I am quoting numbers, it is worth noting that this network model has 21 input nodes, one for each token, and 21 output nodes. Once a file listing the input-output pairs is created, it takes less than one minute on my laptop for this training of the neural network to take place. As for other demonstrations in this book I wrote the program that prepares the data and the neural network in the common programming language Python, drawing on several specialised programming modules that include TensorFlow, a suite of neural network modules developed by Google for language processing.¹⁵ So, the hard work of crafting a neural network model and running it through the training operation is accomplished by high level programming commands in Python.

As an indication of the power of AI, I followed what is becoming accepted programming practice – incorporating AI tools such as ChatGPT into the

programmer's workflow. As a programmer I provided the broad specifications to the AI platform that recommended the kind of NN model and its appropriate hyperparameters, such as the number of nodes in the network and the number of epochs. It then delivered the program in Python code to achieve this. As a programmer new to Python, I then copy-pasted the program into a Python interpreter such as Jupyter that runs the program. Through its conversational interface, ChatGPT made suggestions, refined the program considering feedback, and helped me to debug it.

Semantic embeddings

As we have seen, neural networks are used typically for producing outcomes given certain inputs. These outcomes are predictions. A model trained on the proximity pairs I described can predict the next token in a sequence of feature tokens. With appropriate Python coding, I can design the model so that a user can input a feature token, such as Health, to the trained model. It will then generate a probability distribution of the likely outputs from its 21 possible outcome tokens. The model will then deliver the most likely output from amongst them, for example Business, Retail or Hotel. Neural network models are best considered as stochastic, which is to say they produce random outputs that are heavily constrained by the probability distribution of their outputs as generated by the network.

To summarise, for this worked example there are 21 categories of labelled points on a map. The points are landmarks; the categories are feature tokens. A program calculates the distances of landmarks from one another and pairs them off. So, if a Hotel is within 100 metres of a Bus stop that generates a training pair of Hotel – Bus stop. The pairs are based on proximities, i.e., if a training instance shows a Hotel and Bus stop as paired it means that there is a specific example on the map where a particular hotel building is close (within 100m) of a particular bus stop. In my neighbourhood examples there are several thousand such “proximity pairs,” many of them repeated, all of them also operate in reverse: if a hotel is near a bus stop, then the same bus stop is also near a hotel. Some hotels are also near other hotels, resulting in a pair Hotel – Hotel. The idea is that a neural network is trained to receive one feature as input and adjust the weights between the nodes in its network to generate the other feature in its pair as output. Once trained, the network will be able to produce a set

of probabilities for the feature token most likely to be near to any other feature token.

Once the network is trained, to illustrate semantic embedding I am interested in siphoning off the weights in the arcs between the 21 input nodes and the two hidden nodes as described in Chapter 6 in the example involving the phone, charger and earbuds. Following the method I described in that chapter, the embedding vectors for this example follow this kind of format: Business (0.8544638 0.823204), Education 1.4209204 -0.56683725),

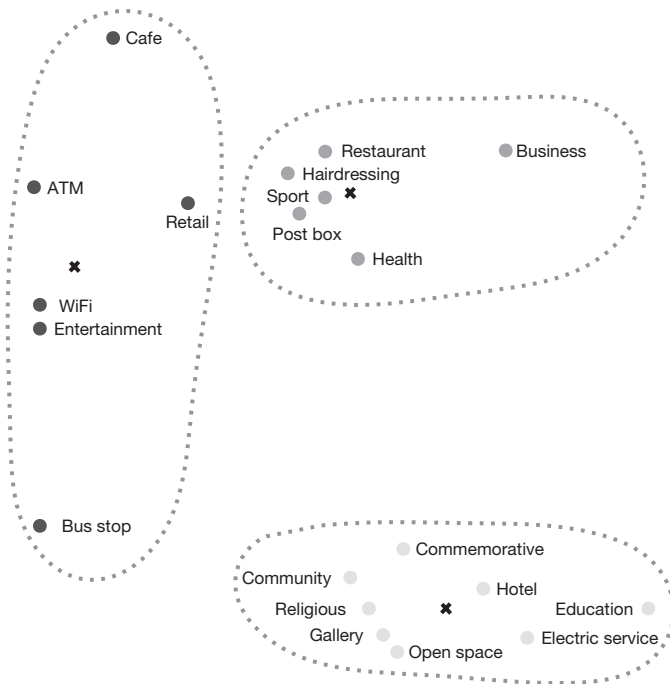


Figure 7.5 Relationships between feature tokens in the neighbourhood map based on their semantic embeddings. The horizontal and vertical axes represent the input values to the two hidden units that originate from each input node (i.e., each feature token). The diagram illustrates how a clustering algorithm (K-Means Clustering) has identified three distinct groupings of tokens, indicating that these tokens have similar or closely related semantic embeddings. The centroid of each cluster is marked with an “x,” showing the central point around which each group of tokens is organised.

Source: Author.

Café (-0.6743826 1.1598709), Commemorative (0.4547563 -0.3898134), Electric service (0.9536952 -0.6570386), Restaurant (0.13953547 0.809849).

These embeddings can be visualised, as in Figure 7.5 to show how the tokens cluster together. For example, the cluster in the upper left includes Café, ATM, Bus stop and Retail. The automated clustering makes sense. These are commonly associated with daily conveniences and services. The central cluster includes Restaurant, Health, Hairdressing and Sport, which are typically related to well-being and social activities. The lower cluster includes Community, Religious, Gallery and Open space, reflecting a focus on communal and cultural functions. This clustering suggests that certain neighbourhood feature tokens are more semantically related to each other, reflecting how spatial or functional proximities might influence the grouping of different categories of landmarks in the urban context.

The method can be scaled. With more hidden nodes in the training network there will be more dimensions to the vectors, which can number in the hundreds or thousands, thereby increasing the embeddings to capture more subtle relationships between the tokens in the corpus, though the embeddings would be difficult to show graphically.

Paths

These kinds of semantic analysis are interesting by-products of LLM processes. But I was interested in generating paths through the neighbourhood, analogous to the construction of sentences. With help from ChatGPT I adapted the model to deliver a likely output that would follow from one of the feature tokens such as Health. Feeding the likely next generated feature token into the neural network would produce a further output. The looped process would continue, producing a chain of feature tokens that take account of feature proximities up to an arbitrary limit of 24 tokens. The sequence in one iteration through this process looked like this: Health, Business, Electric service, Commemorative, Café, Café, Retail, Retail, Restaurant, Hairdressing, Entertainment, Retail, Café, Entertainment, Restaurant, Retail, Restaurant, Business, Business, Post box, Community, Business, Community, Retail.

Think of that as a kind of sentence. It is a list of categories (feature tokens) delineating landmarks that are close to one another in the neighbourhood. To draw a path through the neighbourhood following this sequence

requires us to depart from the normal operations of an LLM. To instantiate this “sentence” of tokens to mapped coordinates I adapted the program to generate actual instances of the tokens, i.e. landmarks, following a few simple rules. The program had to match each feature token in the path sequence (e.g., Health, Business) with a landmark, i.e., a labelled point (e.g., “XYZ Medical Practice,” “Arusha Art Gallery”) and its coordinates. The first instantiation in the sequence was random, as long as it belonged to the category (feature token) in the sequence (e.g., Health). The next instantiation in the path sequence should be the closest point to that one belonging to its category (e.g., Business). The sequence continued in the same way.

Also, I specified that the path should not revisit the same data point, i.e., there are no loops in the path. The output was a map printed on screen showing instantiated landmark points (Figure 7.6). I show two paths derived from different sequences generated by the model and from different random starting positions, with landmark labels removed for clarity.

By processing such sequences using neural network models, we can potentially identify patterns of landmark arrangements in the neighbourhood, just as language models identify patterns in text. I can use these patterns to suggest routes that, for instance, maximise exposure to certain

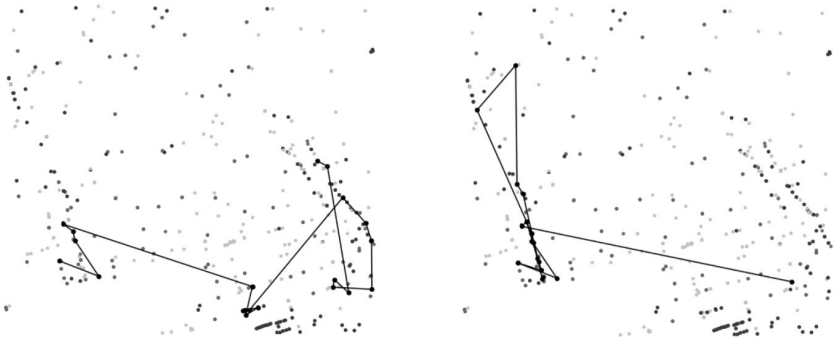


Figure 7.6 Neural network derived paths through the neighbourhood formed from processes analogous to the creation of sentences in an LLM. Segments of the paths follow obvious routes along commercial streets, and there are jumps. There is nothing in this simple model to preclude such discontinuities, nor about modes of transportation from one landmark to another.

Source: Author.

landmarks or avoid others, much like constructing sentences with certain meanings or performative outcomes.¹⁶ Landmarks frequently co-occurring within the same spatial windows on a map might have related or complementary functions.

Though of limited application in actual urban wayfinding or route planning, this method does apply to the challenge of generating sequences of words as in sentences. The method of looping outputs back as inputs to produce a sequence of tokens is applicable to natural language processing. Words in a sentence are analogous to landmarks (or data points) on a map. The context for a given map landmark is the other landmarks within a certain radius, constituting a spatial context window. On the other hand, the immediate context window for a word in a sentence is those words to which it is immediately adjacent in the sentence. As my worked example uses a very simple model, we are defining a word as proximal if it is adjacent, in particular if it follows the current word in the sentence. It makes sense therefore to refer to these couplings in a language version of our program as proximity pairs.

In summary, this path generating exercise draws an analogy between cities and neural networks. In this analogy, input and output nodes correspond to identifiable tokens in the cityscape, e.g. shops, houses, bus stops, street furniture. Training involves associating items in close spatial proximity to one another as input-output pairs. During simulation, a network so trained will predict what token is likely to be close to any other token: a retail outlet is near a bus stop token for example. The model then takes that prediction as input to generate another token as output. Such iteration will create a token sequence that can be mapped back on to the city plan to show a meandering but plausible path through the city. I am proposing that the derivation of such a path is analogous to the creation of a sentence in an LLM.

The neural network I programmed for this exercise has input and output nodes. It also has a so-called “hidden layer” of nodes. During training, the model adjusts the values of these nodes, which in turn enables the network to reproduce patterns in the input and output pairs. Extending our analogy, we can think of hidden layers in the urban context as unidentified elements that facilitate and shape the patterns of interaction and functionality within a city, such as infrastructure elements, utilities, social entities, transportation hubs, zones and economic centres. I did not enter these as feature

tokens to the network during training. It is as if the network identified the presence of such features in establishing its input-output patterns, e.g. that the token sequence “Retail, Café, Entertainment, Restaurant” makes sense as these elements are near a city square, a feature token not entered in the training data. The feature tokens in this sequence may also be found in the entertainment quarter or other unspecified zoning feature.

This neural network model does not identify or name such hidden feature tokens but behaves as if they are present. This inference of new and hidden features as if they are present is one of the keys to the power of neural network models. It also adds to their opacity. Neural networks as they operate in LLMs can generate complex meaning structures that only open themselves to scrutiny in the performance of the model, i.e. the sentences they generate.

In this demonstration I generated paths through a neighbourhood by analysing proximities of landmarks. We can think of each landmark’s immediate spatial context – defined by the 100-meter radius around it – as analogous to a context window in an LLM. Just as a context window in language models limits the tokens that influence the model’s predictions, the spatial context of each landmark limits the set of landmarks that influence the path generated through the neighbourhood. This exercise shows how context windows in LLMs can be conceptually linked to spatial proximities in urban settings, reinforcing the idea that both language and urban spaces are organised by constraints that define what is relevant or significant at any given moment.

Sentences as paths

I further tested the confluence between processing urban map data and language. The program I have described in this chapter was custom-built for generating urban paths, not for generating coherent linguistic sentences. To test the confluence between urban paths and sentence generation, I decided to see if I could adapt the program to create word sequences as a crude LLM. I am following a “proof-of-concept” research method here: if you can create a computer program that demonstrates a concept, then that is sufficient evidence that the concept has potential for further development. By this coding conceit we may even demonstrate that the concept is valid. In this case my concept is simply that sentences are like urban paths.

I describe this phase of the project more fully in Appendix A. The revised model processed approximately 19,000 words of my own writing, creating a lexicon of words and proximity pairs from the text. I then trained the simple TensorFlow feedforward neural network on this data.¹⁷ Despite its limitations, the network demonstrated aspects of conversational AI by generating sentences with more coherence compared against random word generation. I compared the neural network's output with sentences produced by a frequency-based random generator, and ChatGPT correctly identified the neural network's output as more structured and coherent. Though a far cry from coherent English, the model seemed to capture linguistic patterns. But my main point was to show that the same algorithms can be used to generate English sentences as that used for urban paths, further reinforcing parallels between AI and urban contexts.

Summary and critique

In this chapter, I emphasised the importance of context in language and urban settings, drawing parallels between urban networks and neural networks. By examining the operations of LLMs, I demonstrated how they capture context through transformed semantic embeddings within a “context window” akin to a short-term memory.

I illustrated these concepts by examining an urban neighbourhood, mapping physical features as tokens and using a neural network to simulate their spatial proximities. By training the network on proximity pairs, I was able to generate plausible paths through the city, analogous to sentence construction in LLMs. This exercise revealed how neural networks can infer hidden patterns and relationships, much like their ability to create coherent narratives in language.

I concluded by adapting the map exercise to word sequences, developing a neural network to generate sentences from my own writing. The network showed potential in capturing linguistic patterns, producing more coherent text than random generation. This work highlighted the parallels between urban contexts and language models, demonstrating the powerful role of context in both fields.

I invited ChatGPT to identify what some charitable readers might regard as innovative about the moves described in this chapter. Following what I admit is a leading question, it commended the transdisciplinary approach

of combining concepts from urban studies, AI and language processing. The AI platform seemed to concur with my intention to introduce the idea of the context window in LLMs as a bridge between human language comprehension and urban design. It also identified the application of semantic embeddings to map urban features, treating them as tokens, as an innovative approach that bridges language processing techniques with urban analysis.

It suggested that the use of neural networks to predict and map paths through urban environments based on feature proximities is potentially a creative application of AI. It identified the use of proximity pairs, concluding that this approach demonstrates how simple relationships can be used to generate complex structures, such as sentences or urban paths. This highly agreeable virtual critic also approved of the demonstration that neural networks can infer hidden features in urban contexts that parallel how they may infer hidden patterns in language. This concept opens up possibilities for using AI to identify and understand complex urban dynamics that are not immediately visible or easily categorised.

On the negative side, a less favourably inclined critic might challenge the plausibility of the outputs of my demonstrations, the coding conceits it entails, and rightly charge that as well as the benefits, the widespread adoption of such AI approaches closes off possibilities. Where they are introduced, algorithms and numerical data have a way of taking over from more human-oriented qualitative analysis of urban challenges. I will revisit these criticisms in Chapter 12, once I have increased our understanding of other key functions in LLMs.

Notes

- 1 Online Etymology Dictionary. See <https://www.etymonline.com/search?q=context>.
- 2 See Mark Wigley, 'Network Fever', *Grey Room*, 4, no. 4 (2001): 82–122.
- 3 (Statesman, 283a, b) Plato and C. J. Rowe, *Statesman* (Indianapolis, IN: Hackett Pub. Co., 1999).
- 4 Aurelius Marcus and Maxwell Staniforth, *Meditations* (Baltimore, MD: Penguin Books, 1964), 73. See Richard Coyne, 'The Net Effect: Design, the Rhizome, and Complex Philosophy', *Futures*, 40, no. 6 (2008): 552–61 and Richard Coyne, *Network Nature: The Place of Nature in the Digital Age* (London: Bloomsbury Academic, 2018).
- 5 Jane Jacobs, *The Death and Life of Great American Cities* (New York: Random House, 1993), 119.
- 6 Christopher Alexander. 'A City is Not a Tree', in *Design After Modernism*, ed. J. Thackara (London: Thames and Hudson, 1988), 67–84, 11.

- 7 On the subject of semantic and meaningful context in urban places see Christian Norberg-Schulz, *Genius Loci: Towards a Phenomenology of Architecture* (New York: Rizzoli, 1980); Robert Mugerauer, *Interpretations on Behalf of Place: Environmental Displacements and Alternative Responses* (Albany, NY: SUNY Press, 1994); Stephen Graham and Simon Marvin, *Splintering Urbanism: Networked Infrastructures, Technological Mobilities and the Urban Condition* (London: Routledge, 2001); Edward Relph, 'Spirit of Place and Sense of Place in Virtual Realities', *Techné: Research in Philosophy and Technology, Special Issue: Real and Virtual Places* (<http://scholar.lib.vt.edu/ejournals/SPT/v10n3>), 10, no. 3 (2007): 17–25.
- 8 For a discussion of the problematics of the network notion see Mark Wigley, 'Network Fever', *Grey Room*, 4, no. 4 (2001): 82–122, and my own, Richard Coyne, 'The Net Effect: Design, the Rhizome, and Complex Philosophy', *Futures*, 40, no. 6 (2008): 552–61.
- 9 In practice, modern neural network architectures like Transformers take advantage of parallel processing hardware, such as GPUs, to perform these operations simultaneously across multiple tokens and vectors. This parallelism accelerates the training process, making it feasible to work with very large models and datasets.
- 10 To extend and enhance the model's performance beyond these limitations, LLMs must be combined with document processing, web search and other hybrid strategies, which help to retrieve and integrate relevant information that falls outside the immediate context window.
- 11 I discuss the theme of horizon in Richard Coyne 'Melancholy Urbanism: Distant Horizons and the Presentation of Place', in *Cinematic Urban Geographies*, ed. F. Penz and R. Koeck (London: Palgrave Macmillan, 2017), 175–88, following Hans-Georg Gadamer, *Truth and Method* (London: Bloomsbury Revelations, 2013) and Adrian Snodgrass, 'Asian studies and the Fusion of Horizons', *Asian Studies Review*, 15, no. 3 (1992): 81–95
- 12 Note that it is not patterns that are stored, but a method of constructing them. In their chapter 'The Appeal of Parallel Distributed Processing' David Rumelhart, James McClelland and G. E. Hinton, reinforce the point that "the patterns themselves are not stored. Rather, what is stored is the connection strengths between units that allow these patterns to be re-created." James L. McClelland, et al., 'The Appeal of Parallel Distributed Processing', in *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Volume 1: Foundations*, ed. D. E. Rumelhart and J. L. McClelland (Cambridge, MA: MIT Press, 1986), 3–44, 31. Throughout this book I have been careful to avoid the trope that neural networks store patterns.
- 13 The relationships between walking and talking are productive areas of social science research. See Jean-François Augoyard, *Step by Step: Everyday Walks in a French Urban Housing Project* (Minneapolis: University of Minnesota Press, 2007).
- 14 Kevin Lynch, *The Image of the City* (Cambridge, MA: Technology Press, 1960).
- 15 Sumit Raj, *Building Chatbots with Python: Using Natural Language Processing and Machine Learning* (New York: Springer Apress, 2019).
- 16 I am here thinking of speech-act theory. See John Austin, *How to do Things with Words* (Cambridge, MA: Harvard University Press, 1966).
- 17 TensorFlow is an open-source machine learning environment developed by Google, used for building and training neural networks. It provides flexible tools for constructing LLM models.

References

- Alexander, C. (1988), 'A City is Not a Tree', in J. Thackara (ed.), *Design After Modernism: 67–84*, London: Thames and Hudson.
- Augoyard, J.-F. (2007), *Step by Step: Everyday Walks in a French Urban Housing Project*, trans. D. A. Curtis, Minneapolis, MN: University of Minnesota Press.
- Austin, J. (1966), *How to Do Things with Words*, Cambridge, MA: Harvard University Press.
- Coyne, R. (2008), 'The Net Effect: Design, the Rhizome, and Complex Philosophy', *Futures*, 40 (6): 552–61.
- Coyne, R. (2017), 'Melancholy Urbanism: Distant Horizons and the Presentation of Place', in F. Penz and R. Koeck (eds.), *Cinematic Urban Geographies: 175–88*, London: Palgrave Macmillan.
- Coyne, R. (2018), *Network Nature: The Place of Nature in the Digital Age*, London: Bloomsbury Academic.
- Gadamer, H.-G. (2013), *Truth and Method*, trans. J. Weinsheimer and D. G. Marshall, London: Bloomsbury Revelations.
- Graham, S. and S. Marvin. (2001), *Splintering Urbanism: Networked Infrastructures, Technological Mobilities and the Urban Condition*, London: Routledge.
- Jacobs, J. (1993), *The Death and Life of Great American Cities*, New York: Random House.
- Lynch, K. (1960), *The Image of the City*, Cambridge, MA: Technology Press.
- Marcus, A. and M. Staniforth. (1964), *Meditations*, Baltimore, MD: Penguin Books.
- McClelland, J. L., D. E. Rumelhart and G. E. Hinton. (1986), 'The Appeal of Parallel Distributed Processing', in D. E. Rumelhart and J. L. McClelland (eds.), *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Volume 1, Foundations: 3–44*, Cambridge, MA: MIT Press.
- Mugerauer, R. (1994), *Interpretations on Behalf of Place: Environmental Displacements and Alternative Responses*, Albany, NY: SUNY Press.
- Norberg-Schulz, C. (1980), *Genius Loci: Towards a Phenomenology of Architecture*, New York: Rizzoli.
- Oxford English Dictionary (2023), *Context, the Etymology*, Oxford: Oxford University Press.
- Plato and C. J. Rowe. (1999), *Statesman*, Indianapolis, IN: Hackett Pub. Co.
- Raj, S. (2019), *Building Chatbots with Python: Using Natural Language Processing and Machine Learning*, New York: Springer Apress.
- Relph, E. (2007), 'Spirit of Place and Sense of Place in Virtual Realities', *Techné: Research in Philosophy and Technology, Special Issue: Real and Virtual Places*, 10 (3): 17–25. <http://scholar.lib.vt.edu/ejournals/SPT/v10n3>.
- Snodgrass, A. (1992), 'Asian Studies and the Fusion of Horizons', *Asian Studies Review*, 15 (3): 81–95.
- Wigley, M. (2001), 'Network Fever', *Grey Room*, 4 (4): 82–122.

8

TIME AND ORDER

In Chapter 2, I introduced the idea of the script. A script relies explicitly on the ordering of events and spoken words. Playwrights, poets and novelists might experiment with unconventional sequences, or even random arrangements of phrases and lines of dialogue, but even surreal prose gives a nod to order in defining itself as a departure from convention. Here is an example from James Joyce that plays with word order and disorder: “So it returns. Think you’re escaping and run into yourself. Longest way round is the shortest way home. And just when he and she.”¹ This irregular quotation serves to introduce an example of AI-assisted path generation to be developed further on in this chapter.

I need scarcely mention that the experience of place as meaningful is imbued with memories, recollections and histories, as in a novel by Joyce. Conversations about recollections frequently involve attempts to reconstruct sequences of events – concurrences, coincidences, putative causes. So much of urban experience relates to time and order,² from timetables to storytelling. The same applies to interpretation in language. Language is influenced by

historical conditions that precede and shape our understanding.³ The invocation of experience, legacy, provenance, anticipation, fear of the unknown, a future and a before are commonly framed in terms of time and order.

Word order and spatial inference

As a more everyday consideration, the order of words in a sentence as we communicate affects meaning, which is to say how we use the sentence and what response may follow from it. One of the shortcomings in the demonstration in Chapter 7 of creating urban paths is that neither the training nor the simulation took account of the order in which urban feature tokens (shop, school, bus stop) might occur in people's actual journeys through neighbourhoods. I didn't train on existing paths. I extended my demonstration to generating sentences. But neither did that account for order in training data, the order in which words appeared in the training corpus.

Neural network researchers have invented several methods that store and make inferences about word order, the positions of words in a sentence. As I have explained, one method undergirds the performance of natural language processing models known as *Transformer* models. The method combines information about where words appear in a sentence with semantic embeddings for each word as it is fed into a neural network during training.⁴ As I will show, the method exploits properties of periodic cycles, which in turn resonates with investigations into the calibration of periodic cycles in urban contexts.⁵

As an illustration of the way an LLM can generate responses about sequencing consider the simple script I alluded to in Chapter 2 of two people discussing urban glare. Certain LLMs combine their linguistic acumen with image feature detection to exhibit a kind of spatial reasoning. As a further demonstration of how an LLM can succeed with sequencing, I loaded a photograph (Figure 8.1) to ChatGPT with the prompt, "I want to swap the positions of the jug and the miniature teapot. I only have one hand free."

It correctly instructed me in this simple task:

Given the constraints, here's a methodical approach: (1) With your free hand, pick up the jug and place it on the table next to the miniature teapot. (2) Now, pick up the miniature teapot and place it where the jug was. (3) Finally, move the jug to where the teapot originally was.



Figure 8.1 The initial teapot and jug arrangement.

Source: Author.



Figure 8.2 A target arrangement.

Source: Author.

I took another picture (Figure 8.2) and typed: “Thanks. Now how do I progress to this arrangement?”

The platform responded:

- (1) Gently lift the miniature teapot out of the jug and place it on the table.
- (2) Now, pick up the jug and set it onto the coaster or mat where the teapot was.
- (3) Lastly, lift the miniature teapot and nestle it inside the jug.

The sequence does not quite match my expectation, but the system commended my table-top exercise and without further prompting injected

levity: “Your artistic arrangement evokes the idea of containment and whimsy – perfect for a midday mental breather. Enjoy the view!”⁶

This exercise and related attempts at moving objects in photographic space reminded me of the early AI language work by Terry Winograd. He invented a program called SHRDLU for planning the movement of blocks as if operated by a robot arm.⁷ As described in his 1971 PhD thesis, the program responded to natural language commands and answered questions about the arrangement of blocks. His approach drew on rule-based and logic-based programming, which is a different approach to AI than neural network-based LLMs.⁸ I have redrawn one of the visualisations that accompanied Winograd’s thesis in Figure 8.3.

Sequencing is a major aspect in any domain including building and planning.⁹ A body of techniques known as *time series forecasting* usually involves statistical methods to extrapolate something like market trends, the demand for building materials, labour services or real estate. In another body of techniques, *AI planning* uses rules and logic to devise a series of steps to accomplish some goal, relevant to moving objects around a construction site, planning a production process, controlling a self-drive car or planning a journey.

Urban grid

To demonstrate something of the challenge of word order in generating urban paths, and ordered sentences in LLMs, I decided to extricate ChatGPT

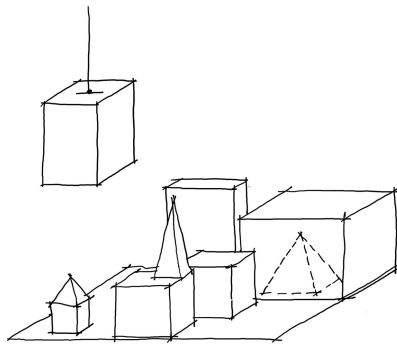


Figure 8.3 Illustration from SHRDLU to show the state of a simple blocks world manipulated by a robot arm or crane under prompts delivered in natural language.

Source: Redrawn by the author.

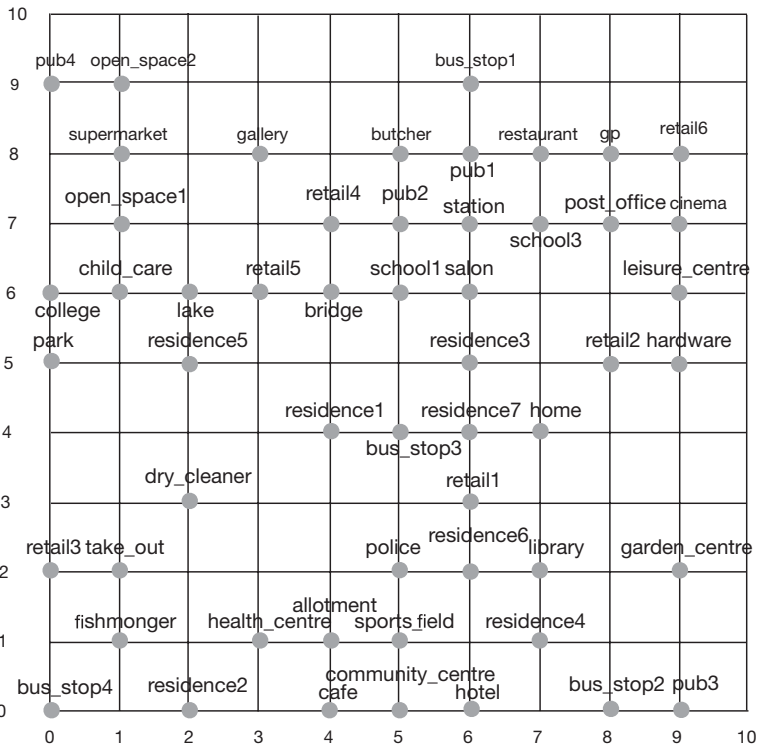


Figure 8.4 Arbitrary arrangement of neighbourhood landmarks on a grid generated by a program.

Source: Author.

from its literary tendencies. I gave it an image of a banal urban grid plan generated by algorithm and uploaded to ChatGPT as a pixel image (i.e. without vector information) shown in Figure 8.4.

To start things off I asked ChatGPT to describe the image. It projected much more into the picture than went into the random disposition of urban feature tokens on a grid:

This map represents a grid-like layout of a city or urban space, delineated by X and Y coordinates, spanning from 0 to 10 on both axes. The map is titled “City Grid with Landmarks.” Landmarks and facilities are represented as red points with accompanying labels. Here’s a categorized description: Recreation & Open Spaces: Open_space1 and open_space2 appear to be recreational or green areas; There’s a lake located near the center; A sports_field is located towards the bottom.

The model continued to describe and make inferences about the drawing. Though I didn't ask, it even proffered an evaluation: "The layout seems to provide a mix of residential, commercial, recreational, and essential services, which gives an impression of a well-balanced urban space. It is quite a dynamic representation of city elements!"¹⁰

I asked the system about grid coordinates: "What is at (6, 9)?" It answered correctly: "at the coordinate (6, 9) on the map, there is a 'bus_stop1' located." I asked if following just the grid lines I can get from there to the lake without passing through other landmarks? It delivered a series of steps that read as a plausible sequence but did not precisely match the landmark positions on the map.

Such interpretive acumen is a tribute to the power of the large language model that drives the performance of the system, combining image recognition capability with the operations of a neural network trained on a vast corpus of texts, some of which includes directional instructions, phrases about spatial qualities and relationships, and myriad examples of personable question and answer formats. Of note however, the model delivered no commentary on the arbitrary, mechanistic and synthetic nature of the plan. It projected, or perhaps "imagined," a richness that I did not see.

Predicting plausible paths

However, my interest lay elsewhere: to see if the same mechanisms deployed by LLMs to operate with word sequences could be applied to predicting plausible paths through the gridded plan, which would thereby assist me in testing how the sequencing mechanism in an LLM works. The initial challenge was to create a training set of actual paths through the gridded plan so that the LLMs neural network could train on legitimate paths through the grid. That is analogous to gathering up texts as sequences of words, i.e. sentences in an LLM training set. The grid represents the set of all possible sentences, constrained by proximity relationships.

I needed a training set that included actual paths across the gridded plan of Figure 8.4. With help from ChatGPT4 I devised a program in Python to generate shortest paths between any two randomly selected landmarks on my map, and to generate the sequence of landmarks that a path would encounter. So that is a series of sequential landmark lists (500) of varying lengths. There was no neural network or AI programming in this aspect of the project. Here

are a couple of the paths generated by my program: “supermarket lake residence5 dry_cleaner residence2,” and “health_centre allotment sports_field residence6 retail1.” I show both paths on the grid map in Figure 8.5.

Of the 500 paths generated, many overlap in the map. The idea was that trained on those geometrically legitimate paths, my AI model would then be able to make up equally plausible geometrical paths from any landmark, i.e. paths other than those it had been trained on. Unlike the demonstration in Chapter 7, in this case the neural network approach did involve a strategy for encoding the ordered position of landmarks as feature tokens along the paths.

I show some of the output below, starting with the GP (medical general practitioner), which it followed with restaurant school2 station pub2 retail4 residence3 bridge bus_stop3 pub1 butcher pub1 butcher open_space1 residence6 hotel police residence6 library residence4 sports_field allotment

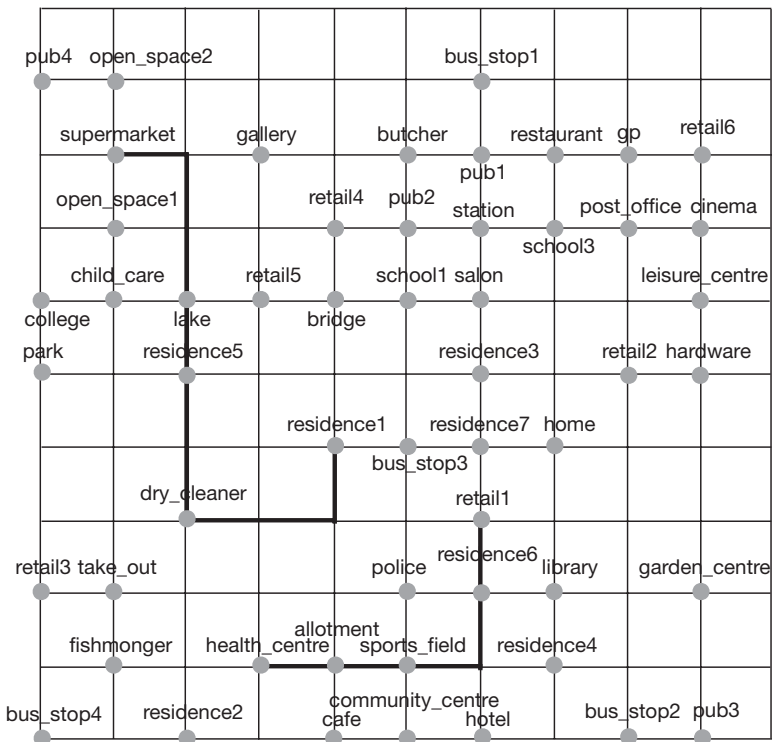


Figure 8.5 Two of the 500 paths generated by an algorithm to train the model.

Source: Author.

health_centre cafe residence2 bus_stop4 gp retail6 cinema post_office school2 retail2 hardware bridge retails lake cafe residence2 ...

The dotted lines in the map in Figure 8.6 show where the neural network got the sequence wrong. More charitably we could say these are implausible parts of the sequence: unlikely shortcuts, off-grid desire lines, jumps to the start of a new sequence or Joycean discontinuities.

I had planned for a single continuous path longer than any of the training examples. Considering the discontinuities, there is no guarantee that the model has synthesised new path sequences that are not in the training corpus. This neural network may in fact demonstrate a neural network property known as “over-fitting,” i.e. it will accurately recall sequences in the training data but be unable to infer or synthesise new paths. We would need to test for that and adjust the network parameters (number of hidden

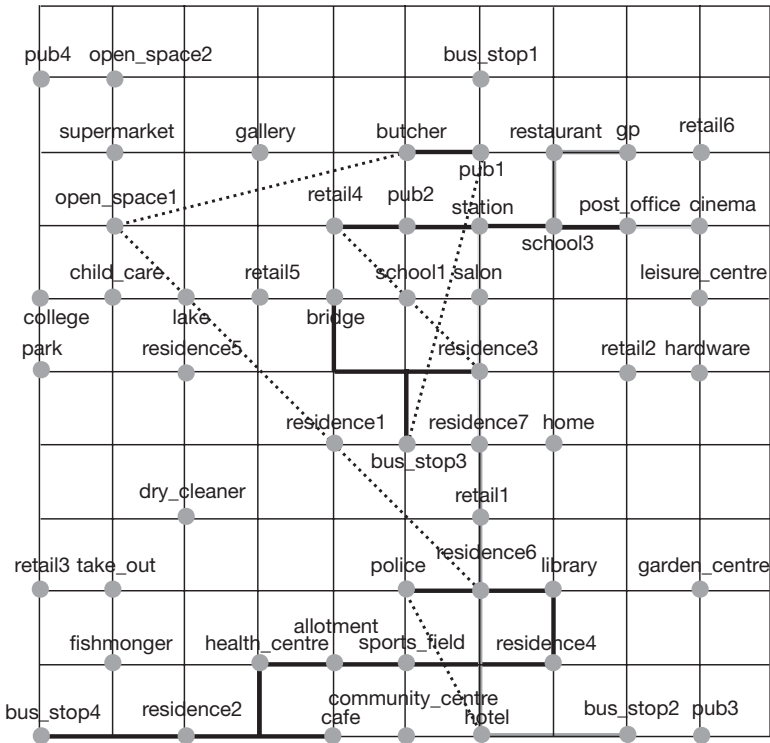


Figure 8.6 New paths generated by the trained model. The dashed lines show disconnects in some of the paths.

Source: Author.

nodes, randomness, etc.) for the neural network to be less rigidly fixated on the sequences in the training data.

Note that the behaviour of the model in generating responses that are spatial and sequential is an outcome of the patterns inferred from the training data. The sequential aspects of the LLM pertain to the ordering of tokens (words). But I have not yet explained the mechanism by which an LLM trains on sequences of data. That comes next.

LLM sequences and cycles

Sentences are sequences of tokens, and one of the major tasks of large language models (LLMs) is to predict the plausible continuation of sentences from some start condition, such as a prompt in a chatbot dialogue box, taking account of the previous flow of tokens within a given context window.

As we have seen, large language models such as ChatGPT follow the Transformer model to capture information about sequences of tokens (words). It does so by exploiting properties of disparate periodic cycles, i.e. sine and cosine curves of different periodicities. The method takes a concept that we normally think of in linear terms, a sequence of words stretching ad infinitum and periodises it, i.e., turns it into cycles that repeat.

Despite its esoteric mathematical formulations, this appeal to periodicity as a way of understanding sequences and time draws on some primal everyday phenomena. I am referring to the ubiquity of cycles, periods, rhythms, vibrations, repetitions and synchronicity in everyday life, and in nature. Periodicity is a major principle through which we understand time, temporality, ordering and sequencing, and permeates so much of human experience, not least the play and labour of urban living. In his book *Rhythmanalysis*, the philosopher Henri Lefebvre observed: “Great cyclical rhythms last for a period and restart: dawn, always new, often superb, inaugurates the return of the everyday.”¹¹ Cycles also go out of sync, and that produces secondary cycles, disruptions and discontinuities. In my book *The Tuning of Place*, I argued for an appreciation of the complex interactions between periodicities, from variations in astronomical cycles (sun and moon transits) to machine components that are out of sync and set up secondary vibrations.¹² Even data compression and cryptography rely on the behaviour of relationships between cyclical patterns modelled as sine waves.¹³ The method is in some respects counter-intuitive, but I will attempt a simple example to illustrate the process used in LLMs.

Token sequencing

How does periodicity assist in capturing and reproducing token sequences? We have seen that during training, an LLM neural network receives sequences of tokens as inputs. The network does not accept these tokens as strings of characters, but as *vectors*, i.e. ordered lists of floating-point numbers that capture the relationships between any token and other tokens in the training corpus. The vector consists of the coordinates of each token in N-dimensional semantic feature space, itself derived through pre-processing in a neural network. See Chapter 4 on the process of semantic embedding. The LLM platform effectively stores these semantic vectors in a lookup table and converts human-readable token sequences to vector encodings.

In Chapter 4, I explored 2-dimensional semantic encodings, as well as encodings in 4 dimensions and 5 dimensions, and the illustration that follows will use 10-dimensional encodings. But serious LLMs typically deploy many more dimensions. In fact, there are repositories of semantic encoding vectors available online. I have experimented with the semantic encodings from Facebook, which offers encodings of 300 dimensions. These are trained on text sources such as Wikipedia and a repository of web data at commoncrawl.org, and are available for over 150 languages.

We have seen that an LLM trained on a corpus of texts will identify words that are commonly grouped according to their proximity within sentences, e.g. we would not be surprised to find words (tokens) such as “building,” “services,” “construction” in each other’s company in a text about architecture. Language is however more than bundles of words. Tokens, words, phrases, sentences and paragraphs are typically arranged following an order, which affects comprehension, grammar, meaning and the kinds of responses they engender.

It turns out that training a neural network to detect and predict patterns in the ordering of tokens, as in a sentence, represents a significant computational challenge, especially when efficiency is required for conversational AI. In conversations, the AI model must process the order in which tokens appear in the communicator’s data stream, considering one token at a time while referencing the context within a defined window, which may span several hundred tokens.

The method used by AI models built on the generative pre-trained Transformer (GPT) architecture to address this challenge is known as “positional encoding.” Before this method was widely adopted, “recurrent

neural networks” (RNNs) were commonly used. However, RNNs had the drawback of processing tokens sequentially, which slowed down both training and prediction because the model had to compute token sequences in the exact order they appeared within the context window. In contrast, Transformers, with their positional encoding, enable the faster processing of word order by allowing parallel computation across entire sequences. Positional encoding calculates the position of each token in the sequence, thus allowing the Transformer to retain information about the order of tokens while still taking advantage of parallel processing capabilities, leading to significantly more efficient training and inference. Therefore, I will focus on this method here.

I implemented a demonstration of positional encoding in the Python programming language. The construction of my home-made LLM program trained to produce arbitrary but reasonably grammatical sentences is currently beyond my desktop capabilities, even with expert assistance from ChatGPT. As in Chapter 4, here is some output from my attempt from a model trained on my own text and accounting for positional encoding: “distance points. Motivation. Giant bernard equilibrium change language and positioning common as support also site parks consists response were meaning communications floresiensis layer other marshalled plausible neural particular win i.e. Online converse notion.” I wasn’t expecting coherence, but it is difficult to recognise if the word order here is any better than random.

To test the effectiveness of positional encoding I decided to revert to a more predictable domain, a well-known sequence that is easy to recognise. I chose the English alphabet. It would be simple to detect if my AI model is able to successfully reproduce the alphabetical sequence from fragments. It is a trivial task, but it helps explain the method of positional encoding.

Rediscovering the alphabet

Think of an archaeologist or philologist trying to construct an alphabet from textual fragments in an unknown language. To manufacture such a training set, an alphabetical corpus, I started with the character string “abcdefghijklmnopqrstuvwxyzabcde,” and wrote a program to fragment it randomly into shorter segments, e.g., cdefghij, tuvw, onpqr, opqsrut, ijklmnop, lmnopq, efghi. No AI was needed for this task. I created 500 such segments between two and eight characters in length to simulate the

imaginary philologist’s findings. I will refer to these characters as *tokens*. The fragments overlap and I introduced some noise by having the program occasionally swap the order of the tokens. The objective was to train a neural network model so that I could then prompt it with a token sequence, such as “def,” and it would continue indefinitely with the complete alphabet, even looping after “z” with “abc” etc.

I ran these 500 segments through a neural network program that derives the strength of relationships between each of the alphabetical tokens as defined in Chapter 6 on semantic encoding. The resultant semantic embeddings reflect the token frequencies and adjacencies in the various segments. The model thereby assigned each token (a, b, c, etc) to a 4-dimensional vector. For example, the program assigned c to a vector (1.630, -0.090, 1.303, 4.088). That is the semantic embedding of that token.

I defined the context window as just six tokens. The semantic vector for each token slides into the context window one token at a time, with the rest of the tokens in the sequence following in its wake, until the last one enters, and another sequence starts. Each step in this procession through the context window constitutes a training event where the neural network is trained on inputs with the output as the next token in the sequence in the fragment.

Here is some more detail on the process if you are curious. The input to the neural network during training is not just the semantic vector for each letter in turn, but the vectors for the entire context window. So, a neural network for training with 4-dimensional semantic vectors and a context window of 6, has $6 \times 4 = 24$ input nodes. It is a “flattened” 6×4 matrix. On the output sides there are as many output nodes as there are tokens, 26 in this case.

Information about the position of an input token in its sequence is captured via a positional encoding matrix. It also has 6×4 dimensions and moves through the context window along with the semantic embeddings. The values within each of those matrices are added, and it is the sum of the semantic and positional matrices that is supplied as input to the neural network. As the network is being trained to predict, the output is the next token in the sequence. The procedure is captured in Table 8.2. It shows the proliferation of floating-point numbers typical of LLM calculations.

In this demonstration the model is training first on the sequence cdefghij. The positional encoding is a vector of dimensions 4×6 , i.e., a matrix with as many columns as there are dimensions to the positional encoding (4) and the number of rows corresponding to the size of the context window (6). I show the values in the positional matrix in Table 8.1.

Table 8.1 The pre-calculated positional matrix.

0.000	1.000	0.000	1.000
0.841	0.995	0.010	1.000
0.909	0.980	0.020	1.000
0.141	0.955	0.030	1.000
-0.757	0.921	0.040	1.000
-0.959	0.878	0.050	1.000

The positional matrix is pre-calculated by the model and remains constant throughout the process. It is derived from sine and cosine calculations that I will discuss in the next section of this chapter.

The first token that enters the context window is *c* with its pre-calculated semantic embedding of (1.630, -0.090, 1.303, 4.088). That is added to the sixth row of the pre-calculated positional embedding matrix to produce a new combined matrix. The flattened version of that matrix constitutes a training instance, presented to the 24 input nodes of the neural network as 0.000, 0.000, 0.000, 0.000, 0.000, 0.000, 0.000, 0.000, 0.000, 0.000, 0.000, 0.000, 0.000, 0.000, 0.000, 0.000, 0.000, 0.000, 0.000, 0.671, 0.788, 1.353, 5.088. The network is trained to produce *d* as the output for that configuration of inputs, as the token *d* follows *c* in the training example. The semantic embedding for token *d* next enters the context window. The semantic embedding for *c* is now displaced to the second position in the context window, resulting in a second set of additions to the positional encoding to produce a new flattened matrix of input values to the neural network, which in this instance is trained to produce *e*, the next token in the sequence. The neural network adjusts its parameters to accommodate that input-output configuration, i.e., so that it can reproduce that input-output pattern during the simulation phase.

You will see from Table 8.2 that as the sequence *cdefghij* passes through the context window one token at a time the neural network adjusts its parameters so it can predict the next token in the sequence. That is just for one training instance. The process is repeated for the next training example, e.g. *tuvwx*. The lengths of the training examples do not need to be the same. The method of calculation and the context window size, six tokens in this case, accommodate any input length.

Table 8.2 Table showing how the model adds semantic and positional vectors to train the neural network to predict each token in the sequence cdefghij. Tokens not in the context window are assigned zero and are greyed out in this table.

	Semantic					Positional					Semantic+positional				Output
	0.000	0.000	0.000	0.000		0.000	0.000	0.000	0.000	=	0.000	0.000	0.000	0.000	
	0.000	0.000	0.000	0.000		0.000	0.000	0.000	0.000	=	0.000	0.000	0.000	0.000	
	0.000	0.000	0.000	0.000		0.000	0.000	0.000	0.000	=	0.000	0.000	0.000	0.000	
	0.000	0.000	0.000	0.000	+	0.000	0.000	0.000	0.000	=	0.000	0.000	0.000	0.000	→ d
	0.000	0.000	0.000	0.000		0.000	0.000	0.000	0.000	=	0.000	0.000	0.000	0.000	
c	1.630	-0.090	1.303	4.088		-0.959	0.878	0.050	1.000		0.671	0.788	1.353	5.088	
	0.000	0.000	0.000	0.000		0.000	0.000	0.000	0.000		0.000	0.000	0.000	0.000	
	0.000	0.000	0.000	0.000		0.000	0.000	0.000	0.000		0.000	0.000	0.000	0.000	
	0.000	0.000	0.000	0.000		0.000	0.000	0.000	0.000		0.000	0.000	0.000	0.000	
	0.000	0.000	0.000	0.000	+	0.000	0.000	0.000	0.000	=	0.000	0.000	0.000	0.000	→ e
c	1.630	-0.090	1.303	4.088		-0.757	0.921	0.040	1.000		0.873	0.831	1.343	5.088	
d	-1.676	2.585	2.063	-1.308		-0.959	0.878	0.050	1.000		-2.635	3.462	2.113	-0.308	
	0.000	0.000	0.000	0.000		0.000	0.000	0.000	0.000		0.000	0.000	0.000	0.000	
	0.000	0.000	0.000	0.000		0.000	0.000	0.000	0.000		0.000	0.000	0.000	0.000	
	0.000	0.000	0.000	0.000		0.000	0.000	0.000	0.000		0.000	0.000	0.000	0.000	

c	1.630	-0.090	1.303	4.088	+	0.141	0.955	0.030	1.000	=	1.771	0.865	1.333	5.088	→ f
d	-1.676	2.585	2.063	-1.308		-0.757	0.921	0.040	1.000		-2.432	3.506	2.103	-0.308	
e	-1.199	0.675	-0.746	-1.621		-0.959	0.878	0.050	1.000		-2.158	1.553	-0.696	-0.621	
	0.000	0.000	0.000	0.000		0.000	0.000	0.000	0.000		0.000	0.000	0.000	0.000	
	0.000	0.000	0.000	0.000		0.000	0.000	0.000	0.000		0.000	0.000	0.000	0.000	
c	1.630	-0.090	1.303	4.088		0.909	0.980	0.020	1.000		2.539	0.890	1.323	5.088	
d	-1.676	2.585	2.063	-1.308	+	0.141	0.955	0.030	1.000	=	-1.535	3.540	2.093	-0.308	→ g
e	-1.199	0.675	-0.746	-1.621		-0.757	0.921	0.040	1.000		-1.956	1.596	-0.706	-0.621	
f	1.194	-1.827	-0.576	-1.294		-0.959	0.878	0.050	1.000		0.235	-0.949	-0.526	-0.294	
	0.000	0.000	0.000	0.000		0.000	0.000	0.000	0.000		0.000	0.000	0.000	0.000	
c	1.630	-0.090	1.303	4.088		0.841	0.995	0.010	1.000		2.472	0.905	1.313	5.088	
d	-1.676	2.585	2.063	-1.308		0.909	0.980	0.020	1.000		-0.766	3.565	2.083	-0.308	
e	-1.199	0.675	-0.746	-1.621	+	0.141	0.955	0.030	1.000	=	-1.058	1.630	-0.716	-0.621	→ h
f	1.194	-1.827	-0.576	-1.294		-0.757	0.921	0.040	1.000		0.437	-0.906	-0.536	-0.294	
g	-1.503	0.100	-1.097	1.091		-0.959	0.878	0.050	1.000		-2.462	0.978	-1.047	2.091	
c	1.630	-0.090	1.303	4.088		0.000	1.000	0.000	1.000		1.630	0.910	1.303	5.088	
d	-1.676	2.585	2.063	-1.308		0.841	0.995	0.010	1.000		-0.834	3.580	2.073	-0.308	

(Continued)

Table 8.2 (Continued)

	Semantic					Positional				Semantic+positional				Output
e	-1.199	0.675	-0.746	-1.621		0.909	0.980	0.020	1.000	-0.290	1.655	-0.726	-0.621	
f	1.194	-1.827	-0.576	-1.294	+	0.141	0.955	0.030	1.000	=	1.335	-0.872	-0.546	-0.294 → i
g	-1.503	0.100	-1.097	1.091		-0.757	0.921	0.040	1.000	-2.260	1.021	-1.057	2.091	
h	-0.578	0.640	2.854	-1.346		-0.959	0.878	0.050	1.000	-1.537	1.518	2.904	-0.346	
d	-1.676	2.585	2.063	-1.308		0.000	1.000	0.000	1.000	-1.676	3.585	2.063	-0.308	
e	-1.199	0.675	-0.746	-1.621		0.841	0.995	0.010	1.000	-0.357	1.670	-0.736	-0.621	
f	1.194	-1.827	-0.576	-1.294		0.909	0.980	0.020	1.000	2.103	-0.847	-0.556	-0.294	
g	-1.503	0.100	-1.097	1.091	+	0.141	0.955	0.030	1.000	=	-1.362	1.056	-1.067	2.091 → j
h	-0.578	0.640	2.854	-1.346		-0.757	0.921	0.040	1.000	-1.335	1.562	2.894	-0.346	
i	4.368	-0.125	0.845	-1.613		-0.959	0.878	0.050	1.000	3.409	0.753	0.895	-0.613	

Source: Author.

You will recall that we refer to the use of the neural network to generate new sequences as the simulation phase. In this phase of the neural network process we provide some new input, e.g., pqr, and ask the network to work out what token comes next. We put this in a loop to create an output sequence. The prediction phase of the network operation must adopt the same procedure as the training phase: adding semantic and positional encoding vectors to the stream of tokens entering the context window. As a neural network is stochastic, the simulation phase generates the most likely next token, given an input series. In this deterministic domain of the alphabet sequence, the next token as calculated by the neural network is highly likely to be as we would expect from our knowledge of the full alphabet. That the program I have described here is able to generate the entire alphabet from fragments validates the procedure.

The mysterious element in this procedure is the formulation and role of the positional encoding, and why adding positional vectors to the semantic vectors captures sequential information. I will explain that later in the chapter, but for now I present some outputs to demonstrate that positional encoding does make a difference.

Testing the model

I can use my alphabetically trained model to make predictions, such as predict what follows abc. As a programmer I can also suppress certain features of the model, such as eliminating the positional or semantic encoding from the training and prediction algorithms to test what differences they make. With just semantic encoding the model generates a sequence that at best looks like fragments joined randomly.

```

a b c b d e f f h j k m l y z z y z y v w x y z y z b d c f f e f h i j n p q r s t
u v w x y z a b c d f f g h i j k l m n y a z y z n s a v x y z a b b d c t z z
b c d e f i h j n o p p q r s t v v x y

```

The positional encoding on its own looks even more random.

```

a b c d x v g s d n v k o e s i p d i k k o q b n z r y p r b f g d f s d f v l b
s v g r e p e y g h o d s r b j j n g p q t x s i b x t k e l d i s p a i a f g m l
l n c j x s f o j d o p p c v z i o f y

```

Here is the output when the semantic and positional encoding are combined as expected in the method.

```

a b c d e f g h i j k l m n o p q r s t u v w x y z a z b c d e f g h i j k l m n o
p q r s t u v w x y z a b c d e f g h i j k l m n o p q r s t u v w x y z a c
b d e f g h i j k l m n o p q r s t u v w

```

There are some random glitches, reflecting the imperfections I introduced deliberately into the training data, but also the stochastic nature of neural network training. Each of these examples is sampled from several trials using the same global parameters, such as the number of training epochs, and the “temperature” in the prediction algorithm that influences the degree of randomness in the output. It is gratifying that the positional and semantic encoding model reconstructs the alphabetic sequence from training just on fragments.

Alphabetical ordering is a deterministic domain with a single identifiable outcome. I could also have used the Periodic Table from chemistry, an ordered list of the most popular films about AI, or an ordered list of shops in the local strip mall. Predicting a full sequence from fragments where the final sequence is less well-defined poses particular challenges and would arguably be more useful. DNA sequencing from DNA fragments also come to mind as an area of application. I will admit that there is something satisfying about illustrating AI processes by referencing the alphabet, considering the prominence of moveable type and the printing press as a transformative technology in the development of cities as alluded to in my book *Cryptographic City*.¹⁴

Reproducing the alphabet in this way increases my confidence that positional encoding does function as designed and can be applied in different contexts, such as way finding and journey planning as demonstrated earlier in this chapter.

Explaining positional encoding

So far, I have demonstrated how a training program calculates a list of positional encoding vectors across a context window of say six tokens in sequence and adds each of these vectors to the semantic encoding vectors to each token in the context window. The vectors in the context window become the inputs to the neural network, with the token that follows as the

output. The neural network is thereby trained to predict the next word that might reasonably follow from the token vectors in the context window.

When that input-output coupling is embedded into the parameters of the network, the program shifts the context by one token to establish a new input-output pairing. Whatever the dimensions of the semantic encoding deployed in the LLM training model, the positional encoding vector needs to have the same dimensions. The two sets of vectors are added arithmetically.

It is simple enough to explain the derivation of the positional encoding vectors. The mysterious element is an account of why they should work to encode information about token positions sufficient to exert an influence on training a neural network to reproduce and predict token sequences, i.e., produce fully formed sentences as output from LLM systems, or other sequences such as in route planning.

First, I will explain how the positional matrix is calculated. Then consider its rationale. In Figure 8.7, I show a series of sine and cosine curves plotted across a range of 0–9 grid units along a horizontal axis. The vertical axis shows the values along the curves ranging from -1 to +1. The curves that originate from +1 are all cosine curves. The curves starting at zero are sine curves. The horizontal units 0–9 represent positions across a context window of 10 tokens.¹⁵ There are 10 curves. Each curve is attributed to a different

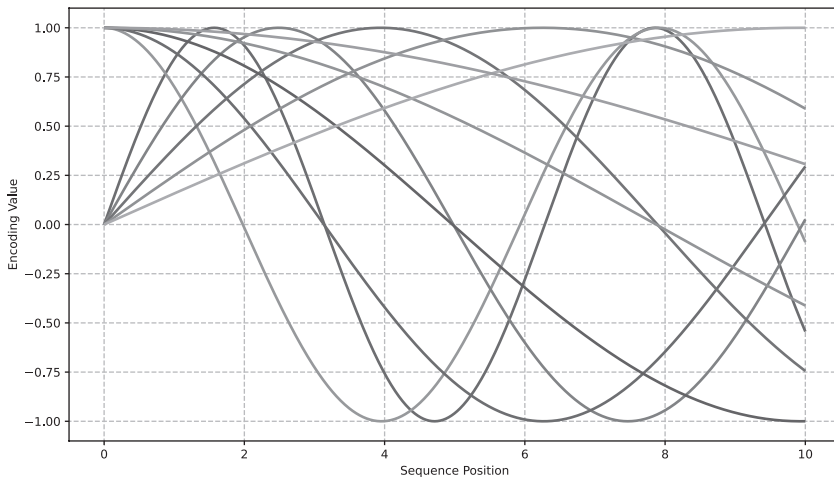


Figure 8.7 Plot of sin and cosine curves from which the model reads off positional embedding values and thereby calculates the positional vectors.

Source: Author.

dimension in a 10-dimensional positional encoding vector. The values in that vector alternate between the sine and cosine curve values at their sequence position in the context window. So, at position 0, the first position in the context window, the positional vector is $\langle 0, 1, 0, 1, 0, 1, 0, 1, 0, 1 \rangle$. The rest of the values can be read off the values at position 1 on the graph and alternating between the sine and cosine values starting with the curves with the shortest amplitude. The values are calculated from formulas in the program.

Another way to visualise the distribution of these values is via a heat map. In the following chart (Figure 8.8), the 10 positions in the sequence are along the horizontal axis, and the rows are the positions in the positional encoding (embedding) vectors. The legend on the right shows the grey shades corresponding to each value in the positional vector.

The pixel values on this graph are the actual positional encoding values (-1 to +1) at any position along the context window, and any coordinate in the encoding vector at that position. This graph shows a 10-dimensional positional encoding vector.

Differences across the context window

The heat map indicates some interesting properties in the distribution of positional values. The right side of the graph shows the most variation amongst the coordinates of the encoding vector. Furthermore, the variation

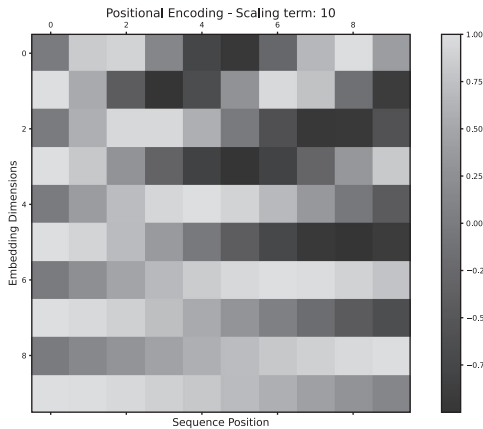


Figure 8.8 Heat map containing data equivalent to the sine-cosine graph of Figure 8.7.

Source: Author.

is most pronounced amongst the low coordinates of the vector (towards the top of the plot), privileging the values towards the top right corner of the data space makes sense as it accords greater importance to variation amongst the tokens closest to the most recent token in the prediction sequence and greater importance to the values at the lower coordinates in the positional vector.

The patterns become more interesting graphically with larger vectors and longer context windows. Figure 8.9 shows the curves for positional vectors of 100 dimensions and a context window of 100 tokens to indicate how the curves converge for larger context windows. Figure 8.10 shows the corresponding heat map.

As someone interested in geometry, the spatial depictions hold a certain fascination. LLM techniques and urbanism elide through their common investment in geometrical forms. See Appendix B for an explanation of the formula that produces these curves.

The positional matrix

Why are sine and cosine curves used to create the positional matrix? The simple answer is that the method seems to work empirically, after much testing by LLM specialists.¹⁶ My own modest foray into positional encoding,

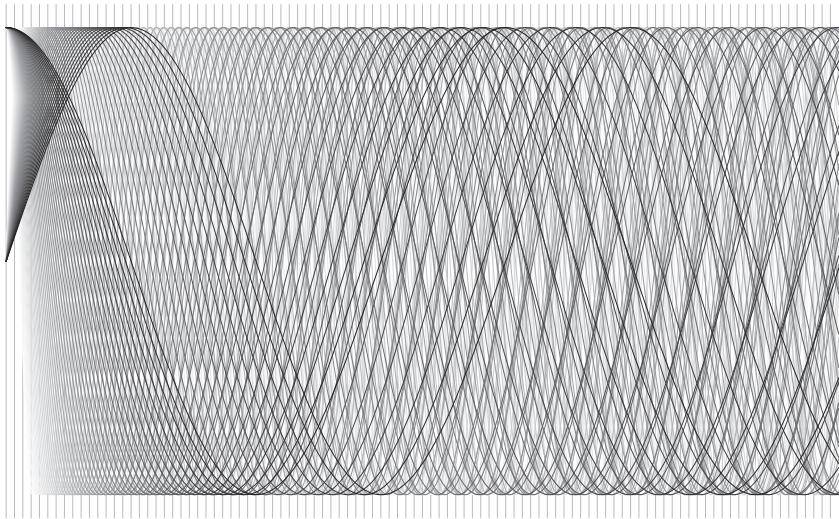


Figure 8.9 Graph similar to Figure 8.7, but with curves for positional vectors of 100 dimensions and a context window of 100 tokens.

Source: Author.

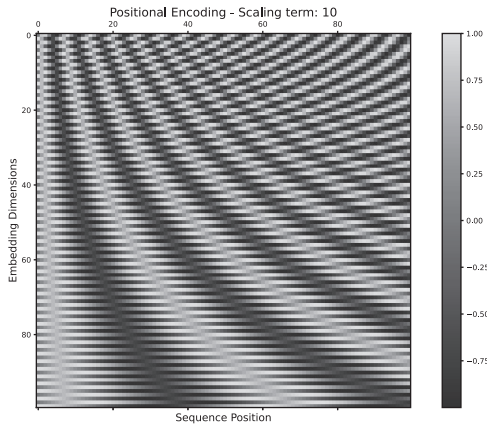


Figure 8.10 Heat map corresponding to the graph of Figure 8.9.

Source: Author.

as illustrated above, confirms this practical point. If I take out the positional encoding from the calculations in the alphabet example, then the resultant sequential predictions seem random.

According to LLM researchers (e.g. in the seminal “Attention Is All You Need” paper) the logic of the procedure lies in the fact that the sinusoidal functions’ periodic nature provides information for the training model to identify the relative position of tokens, even for positions not encountered during training. The model can interpolate or extrapolate positional information.

These positional encodings enable the training model to latch on to patterns in the relative positions of tokens. By having more pronounced variations in the encodings for recent tokens and the lower dimensions of the encodings, the model can prioritise or give more attention to these recent tokens. That seems to be crucial for sequential tasks. It helps sequential training that the positional vectors are unique for each position in a context window, with vector values in the range -1 to +1.

It is worth noting that positional encodings on their own are insufficient. Semantic encodings are needed as an efficient and “dense” means of mapping tokens to neural network inputs. Tokens with similar contexts end up having similar vector representations. The vector inputs already provide important information about word proximities in the overall corpus, and this information is enhanced by the arithmetical addition of semantic and positional vector encodings.

Having now attempted to explain positional encoding in terms of periodicity, we can dispose of any idea that to train a neural network on sequencing we simply provide an index number for each position in a context window and somehow add that to the semantic encodings. Sequences, language and human affairs are best grasped by appropriating their cyclical nature.

To revert to the cultural framing from earlier in the chapter, Lefebvre provides an account of the antagonism between the linear and the cyclical in an extension of the quote I gave above:

Great cyclical rhythms last for a period and restart: dawn, always new, often superb, inaugurates the return of the everyday. The antagonistic unity of relations between the cyclical and the linear sometimes gives rise to compromises, sometimes to disturbances. The circular course of the hands on (traditional) clock-faces and watches is accompanied by a linear tick-tock. And it is their relation that enables or rather constitutes the measure of time (which is to say, of rhythms).¹⁷

Review

In this chapter, I explored the significance of time and order in both urban contexts and language processing, illustrating how sequences and ordering affect meaning and understanding. I reprised from Chapter 2 the concept of scripts and how they rely on the ordering of events, and I connected this to the neural network techniques used in language models, particularly focusing on the generative pre-trained Transformer (GPT) architecture and its use of positional encoding. I demonstrated the impact of word order on sentence meaning and how LLMs incorporate this through advanced methods such as positional encoding, which enables faster and more efficient processing compared to earlier methods such as recurrent neural networks (RNNs). By experimenting with images and simple sequences, I demonstrated how AI systems can process spatial and temporal sequences. In particular, I showed examples of training on and generating paths through a rudimentary urban grid, and the even simpler reconstruction of the English alphabet via a model trained on alphabet fragments. Finally, I explained the mathematical concepts underlying positional encoding, emphasising its role in capturing sequence information and its empirical success in enhancing AI models' performance in generating coherent language.

As before, I asked ChatGPT to identify ideas and approaches in the chapter that a reader might find innovative. It identified the use of sinusoidal functions to encode position and its application in the context of urban studies. Granted, conversational AI is generally tuned to be complimentary and encouraging. I was not surprised therefore that it declared the use of the English alphabet as a test case to demonstrate the efficacy of positional encoding as a creative approach. It proposed that drawing on Henri Lefebvre's concept of rhythmanalysis and the periodic nature of human experience to explain AI sequencing provides an innovative cross-disciplinary perspective. It suggested that calling on historical AI projects, such as Terry Winograd's SHRDLU, to contextualise current AI capabilities underscores the evolution of AI and highlights the innovative advancements in language processing models today. The use of heat maps and sinusoidal graphs to visualise positional encoding vectors offers a novel way to understand the mathematical underpinnings of AI models. Integrating image recognition with language models to achieve spatial reasoning touches on the innovative frontier of multi-modal AI, where different data types are combined to enhance AI's interaction capabilities.

As an advocate of a Phenomenological perspective, I recognise that drawing parallels between the role of rhythm in urban life and numerical computations risks reducing the richness of human temporal experiences to mere numbers and machine logic. However, the investigations in this chapter demonstrate that both language and urban experiences are deeply invested in temporal, positional and cyclical processes, regardless of how we characterise, represent or construct temporal experience.

Notes

- 1 James Joyce, *Ulysses* (London: Penguin, 1992).
- 2 Henri Lefebvre, *Rhythmanalysis: Space, Time and Everyday Life* (London: Continuum, 2004).
- 3 Here we could draw on the hermeneutic theories of Hans-Georg Gadamer and others. See Hans-Georg Gadamer, *Truth and Method* (London: Bloomsbury Revelations, 2013), and our account of the field in Adrian Snodgrass and Richard Coyne, *Interpretation in Architecture: Design as a Way of Thinking* (London: Routledge, 2006).
- 4 Richard Walker (2023), 'ChatGPT Position and Positional Embeddings: Transformers & NLP 3'. *Lucidate*, 26 February. Available online: <https://www.youtube.com/watch?v=DINUVMoJNwU> (accessed 12 January 2023).

- 5 Henri Lefebvre, *Rhythmanalysis: Space, Time and Everyday Life* (London: Continuum, 2004).
- 6 For further demonstrations involving this spatial challenge see Richard Coyne (2023b), 'Robot Probes City Grid'. *Reflections on Technology, Media & Culture*, 28 October. Available online: <https://richardcoyne.com/2023/10/28/robot-strolls-through-city-grid/> (accessed 12 October 2024).
- 7 Terry Winograd, PhD thesis, *Procedures as a Representation for Data in a Computer Program for Understanding Natural Language* (Cambridge, MA: MIT, 1971).
- 8 ChatGPT has the ability to convert this task domain to a series of logical predicates. See Richard Coyne (2024), 'Chain of Thought'. *Reflections on Technology, Media & Culture*, 28 September. Available online: <https://richardcoyne.com/2024/09/28/22272/> (accessed 12 October 2024).
- 9 As well as Schank's work on plans, see Lucille Alice Suchman, *Plans and Situated Actions: The Problem of Human-Machine Communication* (Cambridge: Cambridge University Press, 1987) and my own early attempts to apply AI planning to design: Richard Coyne, 'Knowledge-Based Planning Systems and Design: A Review', *Architectural Science Review*, 28, no. 4 (1985): 95–103.
- 10 That the platform was able to advance a positive assessment of such an impoverished representation shows its bias towards positive judgements and the limits of its capacity to make sound evaluations unaided.
- 11 Henri Lefebvre, *Rhythmanalysis: Space, Time and Everyday Life* (London: Continuum, 2004), 8.
- 12 Richard Coyne, *The Tuning of Place: Sociable Spaces and Pervasive Digital Media* (Cambridge, MA: MIT Press, 2010).
- 13 See Richard Coyne, *Cryptographic City: Decoding the Smart Metropolis* (Cambridge, MA: MIT Press, 2023a).
- 14 Richard Coyne, *Cryptographic City: Decoding the Smart Metropolis* (Cambridge, MA: MIT Press, 2023a). Also see Quinn DuPont. 'The Printing Press and Cryptography: Alberti and the Dawn of a Notational Epoch', in *A Material History of Medieval and Early Modern Ciphers*, ed. K. Ellison and S. Kim (London: Routledge, 2017), 95–117.
- 15 Note that the context window is 10 tokens in this demonstration as opposed to six in the previous example.
- 16 Ashish Vaswani, et al. 'Attention Is All You Need', in *31st Conference on Neural Information Processing Systems* (Long Beach, CA, USA, 2017), 1–15.
- 17 Henri Lefebvre, *Rhythmanalysis: Space, Time and Everyday Life* (London: Continuum, 2004), 8.

References

- Coyne, R. (1985), 'Knowledge-Based Planning Systems and Design: A Review', *Architectural Science Review*, 28 (4): 95–103.
- Coyne, R. (2010), *The Tuning of Place: Sociable Spaces and Pervasive Digital Media*, Cambridge, MA: MIT Press.
- Coyne, R. (2023a), *Cryptographic City: Decoding the Smart Metropolis*, Cambridge, MA: MIT Press.

- Coyne, R. (2023b), 'Robot Probes City Grid', *Reflections on Technology, Media & Culture*, 28 October. Available online: <https://richardcoyne.com/2023/10/28/robot-strolls-through-city-grid/> (accessed 12 October 2024).
- Coyne, R. (2024), 'Chain of Thought', *Reflections on Technology, Media & Culture*, 28 September. Available online: <https://richardcoyne.com/2024/09/28/22272/> (accessed 12 October 2024).
- DuPont, Q. (2017), 'The Printing Press and Cryptography: Alberti and the Dawn of a Notational Epoch', in K. Ellison and S. Kim (eds), *A Material History of Medieval and Early Modern Ciphers*: 95–117, London: Routledge.
- Gadamer, H.-G. (2013), *Truth and Method*, trans. J. Weinsheimer and D. G. Marshall, London: Bloomsbury Revelations.
- Joyce, J. (1992), *Ulysses*, London: Penguin.
- Lefebvre, H. (2004), *Rhythmanalysis: Space, Time and Everyday Life*, trans. S. Elden and G. Moore, London: Continuum.
- Snodgrass, A. and R. Coyne. (2006), *Interpretation in Architecture: Design as a Way of Thinking*, London: Routledge.
- Suchman, L. A. (1987), *Plans and Situated Actions: The Problem of Human-Machine Communication*, Cambridge: Cambridge University Press.
- Vaswani, A., et al. (2017), 'Attention Is All You Need', *31st Conference on Neural Information Processing Systems*: 1–15, Long Beach, CA, USA.
- Walker, R. (2023), 'ChatGPT Position and Positional Embeddings: Transformers & NLP 3', *Lucidate*, 26 February. Available online: <https://www.youtube.com/watch?v=DINUVMojNwU> (accessed 12 January 2023).
- Winograd, T. (1971), Phd Thesis, *Procedures as a Representation for Data in a Computer Program for Understanding Natural Language*, Cambridge, MA: MIT.

9

THE ATTENTION ECONOMY

Consumers and critics of social media are familiar with the idea of the “attention economy.” That is where advertisers and developers compete to draw the attention of potential consumers to their websites and retail outlets,¹ sometimes with attention-grabbing banners, imagery and “click-bait.” Such tactics tap into attention as a key element in cognition, spatial experience and language. Whether deliberate or unconscious, we direct attention to features in our environment that are important to us at that moment, such as people, events, material objects, memories, inexistent objects and figments of imagination.² Attention can also wander.³ The object of attention may be only loosely determined, but as outlined by the psychologist William James (1842–1910), without this cognitive capacity to attend we are immersed in a sea of indiscriminate sense data.⁴ In Chapter 2 on “Scripting the city,” I introduced the play of attention and inattention in scripted conversation scenarios. We do not always attend to what the other person is saying, and sometimes we pay attention to elements in their narrative that they did not expect. That play is a part of everyday conversation.

DOI: 10.4324/9781003535751-10

This chapter has been made available under a CC-BY-NC-ND 4.0 license.

Attention and environment

Attention and its misdirection also feature in various pathologies, such as attention deficit hyperactivity disorder (ADHD), and the world of everyday work can engender “attention fatigue.” The requirement for excessive attention and concentration can induce stress and weariness, as in prolonged periods working at a computer, production line or unremitting attendance at meetings. According to “restoration theory” opportunities to shift attention to less demanding tasks offer respite, such as attending to objects that afford “soft fascination.”⁵ Nature provides the remedy. Many people think of the natural environment as a stable and predictable substate and reference plane that grounds human experience.⁶ By contrast, urban environments are unfamiliar, dynamic and unpredictable. Attention in the city is also fragmented, demanding frequent shifts in attention. According to this theory, it makes sense therefore to retreat to parks, gardens and other green spaces in cities to recover from attentional stress, disorder and fatigue. Even outside of the insights of restoration theory, cities present a spectrum of demands on our attention that shape our perceptual engagement, and indeed shape our cities.

We have tools for managing and manipulating attention. In the urban context there are landmarks, iconic structures (Figure 9.1) and carefully positioned signs that draw attention. Included amongst the apparatus of visual spectacle are screens and optical devices to help people manipulate visual attention. Of note are the deliberate tactics of urban developers and designers to create settings that attract photography and photo-sharing, sites of so-called “instagrammable” architecture.⁷

Language plays an obvious role in directing attention. In the city as semiotic field, words grab attention through brand logos, advertisements, labels and the pantheon of aural and visual semiotic directives – “stop,” “enter,” “exit,” “try,” “buy.” People offer advice and commands to each other about what requires attention. Individual words and phrases within the general flow of language production also operate to manage attention. A lecturer will come to the class armed with notes or key points on presentation slides that not only focus the attention of the audience but aid the speaker in executing their word flow. The list “affect,” “emotion,” “mood,” “melancholy” reminds me not only to use those words in my lecture on emotional cityscapes, but prompts my flow of sentences on the subject, invoking topics to which I as speaker need to pay attention. As another



Figure 9.1 Elements of a cityscape arranged for conspicuous attention. The downtown district of Dubai with the 830 m tall Burj Khalifa as its central landmark.

Source: Author.

attentional tactic, when I try to interpret text written in a language with which I have only passing competence such as French or German I *attend* to familiar words in the sentence and guess at the rest. The ability to *attend*, selectively, is a vital cognitive function, as we see, listen, perceive, speak, write and address everyday tasks in social and urban contexts.

Attention and automation

It is easy to see how automated processes, such as computer systems, might deploy operations described readily as *attentional*. Data search provides a trivial example. To change the word “house” to “dwelling” in a database, a program cycles through a list of words in the database and compares them with “house.” So, it effectively scans the list to settle on the object on which to direct its “attention” and on which to conduct a series of simple transformation and replacement operations. With greater sophistication, and aided by extensive indexing, web search engines quickly home in on (*attend to*) some target data to display, extract or manipulate.

Selective attention is a crucial aspect of language understanding and represents one of the major challenges within LLMs. Generative conversational AI

developers have invented methods by which language processing algorithms establish those parts of any sentence that require attention in order to deliver a meaningful interaction with a user. The key research article to which I refer frequently in this book is titled appropriately “Attention is all you need.”⁸ It proffers “attention” as a watchword in text-based conversational AI.

As evidence of the important role of attention in language comprehension, think of how texts are often supplemented with attentional markers. Before the widespread adoption of digital text processing, it was common to borrow a book from the university library only to find that some other reader had already graffitied key pages with markup signs such as scribbles, notes, underlining and highlights. These extra-textual signals may have aided a previous reader’s recall, or perhaps indicated a reading process that benefited from graphical participation in marking up attentional cues, much as one might shift emphasis in the words of a sentence when reading them out loud.

Such extra-textual markup practices parallel what we do while speaking anyway. As well as via bodily gestures, a speaker will stress vocally words that they want listeners to pay attention to. Indeed, one of the main features of convincing automated text to speech synthesis, as in Tacotron and Speechify, is their ability to take account of *prosody*, patterns of stress and intonation in the way people speak. Without the benefit of an intentional human reader, the speech synthesisers take account of a word’s position in the context of the whole passage to make a best guess at a plausible pattern of stress and intonation. Even random or misplaced vocal emphasis renders automated text to speech more acceptable to most listeners. Such systems are in marked contrast to early voice synthesis applications delivering something like the classic sci-fi trope of the robot voice: Robbie the Robot, Terminator, Daleks.

Digital applications based on neural architectures and Transformer-based models, incorporate attention mechanisms. The attention mechanism in conversational LLMs does something similar by calculating where to add emphasis to particular words in sentences as an aid to generating further text. Attention impinges on meaning and response. For example, if I say, “We should meet at the café” with emphasis on “we” then the response could be “Let’s organise that just for the two of us then.” If the emphasis is on “should” then a response might be “Yes. It’s imperative that we meet there. I’ll get directions.” Again, an LLM does not have the benefit of human intentions or motivations, so an automated model must make such

assessments of attention from the text alone, possibly hypothesising alternative attentional profiles and selecting from amongst them. In a conversational context, such a process must take place quickly and efficiently. I will work through a simple example to show how this might work in an LLM.

Attention is sufficient

Here is a statement about graffiti with highlights indicating what I think are the most significant words in the sentence – those to which a reader might pay attention. Stronger highlight implies greater attention; a lighter shade warrants less attention. The method calculates and visualises a kind of heat map of attention for this sentence.

Graffiti is both an **art form** and a means of **disrupting** the **semiotic** coding of the city.

The scoring and highlighting are the result of my own assessment of word importance. It is one of several *heat maps* I could have generated, depending on my particular interest at the time.

I can replicate this process numerically. In brief, the method takes the semantic vector of each word (as described in Chapter 6) and calculates the relationship between each word and every other word in the sentence. Summing up all the relationships for each word (as vectors) provides an indication of the significance of that word in relationship to the whole sentence. The heat map of the sentence above simplifies those vectors to a series of single value attention scores. The score to each word determines its highlight colour. So, the method compares the numerical representation of each word's meaning, to show the words that have most affinity with the words in the sentence as a whole.

That is a general description of the process. Here is more detail on the calculation for interested readers. To recap, LLMs operate typically with token embeddings more than 2,048 dimensions and to 17 decimal places. That is too unwieldy to use in a demonstration. Here are some embeddings for a subset of several words each with just 4 dimensions, and to 1 decimal place. Assume these semantic vectors were derived from a vast set of texts via the process described in Chapter 6 on semantic embeddings: “city” [0.8, -0.2, 0.5, 0.4], “good” [-0.1, 0.5, 0.9, 0.3], “graffiti” [-0.2, 0.8, 0.3,

-0.5], “social” [0.7, -0.4, -0.5, 0.6], “tree” [0.2, 0.9, 0.1, -0.3], “a” [-0.3, 0.6, 0.2, -0.1], “is” [0.5, -0.1, 0.7, 0.2], “not” [-0.5, 0.4, -0.2, -0.1].

Here is a shorter version of the initial sentence for simplicity. I want to derive attention scores for the sentence “graffiti is a social good.” Table 9.1 shows these vectors laid out in tabular format.

Attention scoring takes place as part of the LLM training procedure. I have ignored capitalisation and punctuation. Imagine this sentence is one of millions that appears in a neural network training corpus. The attention module of a neural network system will process the words in this sentence to take account of these generalised embeddings and modify them to account for their relationships with other words in the sentence.

Here is a more technical summary of how that works. The attention algorithm takes each embedding vector in turn and calculates its proximity in the multidimensional feature space to every other word in the sentence. That is achieved by a simple multiplication and addition process (dot product). So, each parameter in the “graffiti” vector is multiplied by each element in the “good” vector and then summed to give an *alignment score* for semantic proximity to “graffiti.” The scores for each word as they relate to “graffiti” are then turned into weightings (probabilities). Each weighting for the words in the sentence will add up to 1.0.

These weightings are then used in a calculation that modifies the embedding vectors for each word. Summing those modified vectors provides scores that serve to indicate the importance of each word to the meaning of the sentence as a whole. Table 9.2 shows a spreadsheet in which I calculated the overall score for the word “graffiti” in the sentence. The final score for that word is 0.59.

Table 9.1 Some words and their semantic embeddings in a four-dimensional semantic space.

graffiti	-0.2	0.8	0.3	-0.5
is	0.5	-0.1	0.7	0.2
a	-0.3	0.6	0.2	-0.1
social	0.7	-0.4	-0.5	0.6
good	-0.1	0.5	0.9	0.3

Source: Author.

Table 9.2 Derivation of the attention score 0.59 for the word “graffiti” in “graffiti is a social good.”

	semantic embedding vectors				alignment score	probability	weighted embedding vectors				
graffiti	-0.2	0.8	0.3	-0.5	1.02	0.36	0.01	0.23	0.03	0.09	
is	0.5	-0.1	0.7	0.2	-0.07	0.12	-0.01	-0.01	0.03	-0.01	
a	-0.3	0.6	0.2	-0.1	0.65	0.25	0.01	0.12	0.01	0.01	
social	0.7	-0.4	-0.5	0.6	-0.91	0.05	-0.01	-0.02	-0.01	-0.02	
good	-0.1	0.5	0.9	0.3	0.54	0.22	0.00	0.09	0.06	-0.03	
							0.01	0.41	0.12	0.04	0.59

Source: Author.

The procedure is repeated for each word in the sentence. Table 9.3 shows the attention scores for all the words in the sentence, shaded for clarity.

Though this table was generated algorithmically, it is gratifying to see that the attention scores look as one might expect for such a sentence. A linguistically capable human being reading or interpreting the sentence is indeed likely to focus their attention on “graffiti” and “good.” To check that this outcome was not just a coincidence I applied the same method to the sentence “A city is not a tree” to produce the attention scores in Table 9.4.

It is worth remembering at this stage that word embeddings are usually much larger than indicated here, that LLMs typically process blocks of text that are longer than single sentences, and that LLMs operate with tokens that include whole words, parts of words and punctuation. It is unlikely that a human being would be able to read or interpret the attentional profile of a block of text as processes by an LLM.

The overall objective of generative LLMs is prediction: for a trained neural network to predict the next word or token in a generative text sequence. If you think of the size of embedding vectors, it involves a huge amount of calculation to replicate what human language users take for granted as functions of basic language competence.

Multi-head attention

The variegated colouring and markup of the graffitied library books that I described earlier in this chapter were often traces of the reading practices not just of a single reader but of multiple readers. Such markups constituted

Table 9.3 Calculated attention scores for “graffiti is a social good.”

graffiti	is	a	social	good
0.59	0.40	0.37	0.58	0.69

Source: Author.

Table 9.4 Attention scores for “a city is not a tree.”

a	city	is	not	a	tree
0.34	0.54	0.41	0.16	0.23	0.49

Source: Author.

a record of human-based “multi-head attention.” They also show that different readers and speakers will have varied attentional responses to the same text.

I have already considered how automated LLMs attempt to simulate the way a listener or reader will focus on key tokens, words and groups of words in a sentence to decide how to respond to the sentence. Recall that the automated LLM method derives an attention score for each token in a sentence based on how close it is to each of the other words in the sentence in terms of “semantic proximity.” Tokens in these LLMs are represented as long sequences of numbers (vectors). The semantic proximity of two words is the distance between them in this multidimensional feature space. As shown in Table 9.2, my example sentence “graffiti is a social good” calculated this (simplified) series of attention scores.

An LLM platform calculates and deploys these scores in both its training phase and also in the so-called *inference* or *prediction* phase. The latter is activated when someone types comments or questions into a conversational AI system to prompt a response. After the prompts, the platform calculates the attention scores of the words in the input sentences. The model makes use of the attention scores, along with word sequences and word meanings, abetted by its functions as a finely tuned neural network to generate (predict) the sequence of words that would make sense as a response – the words that would follow. As I have already mentioned, human speakers are adept at placing emphasis on key words that they want listeners to pay attention to, and that alternative emphases result in different interpretations, which is to say produce different responses from the model.

Contrary to spoken human conversation and dialogue, an automated LLM only has access to text. So, it must infer (calculate) emphasis (attention) from the text on which it is trained and the flow of text during interaction with a human being. In terms of automated natural language processing, a sentence might exhibit multiple patterns of potential attention scores as calculated above. LLM specialists refer to the generation in their models of more than one pattern of attention for the same input text block as *multi-head attention*.

It is easiest to illustrate multi-head attention with sentences longer than the one I provided above. I will use the sentence: “Graffiti is both an art form and a means of disrupting the semiotic coding of the city.” Table 9.5 shows one calculation where attention focuses according to the

subject-action relationship: **graffiti** and **disrupting**. Those two words have higher scores as calculated above. Note that the scores in this illustration are just to two decimal places. They would have many more places in an actual LLM implementation.

Focussing on the *object being acted upon* (the city) the model could produce a different distribution of scores, emphasising *semiotic*, *coding* and *city* as in Table 9.6.

Focussing on the *method or means of action*, the attention scores could appear something like Table 9.7, putting the emphasis on *art*, *means* and *disrupting*.

I have shown the possible outcomes from just three attention heads. Natural language models such as ChatGPT will deploy a greater, but fixed number of attention heads (8–16).

Table 9.5 Attention scores that place emphasis on “Graffiti” and “disrupting.”

Graffiti	is	both	an	art	form	and	a	
0.50	0.05	0.05	0.05	0.05	0.05	0.05	0.05	
means	of	disrupting	the	semiotic	coding	of	the	city
0.10	0.05	0.10	0.05	0.05	0.05	0.05	0.05	0.05

Source: Author.

Table 9.6 Attention scores emphasising “semiotic” and “coding.”

Graffiti	is	both	an	art	form	and	a	
0.05	0.05	0.05	0.05	0.05	0.05	0.05	0.05	
means	of	disrupting	the	semiotic	coding	of	the	city
0.10	0.05	0.10	0.05	0.40	0.40	0.05	0.05	0.10

Source: Author.

Table 9.7 Attention scores emphasising “means,” “art” and “disrupting.”

Graffiti	is	both	an	art	form	and	a	
0.05	0.05	0.05	0.05	0.20	0.05	0.05	0.05	
means	of	disrupting	the	semiotic	coding	of	the	city
0.30	0.05	0.20	0.05	0.05	0.05	0.05	0.05	0.05

Source: Author.

LLM researchers show that calculating these multiple attention profiles leads to a rich neural network representation of textual information. These representations influence strongly the LLM performance so that it appears to replicate purposeful human language capabilities. Competent human speakers and listeners are adept at adjusting their focus depending on context. At one moment the emotional cues in a sentence come to the fore, at other times it is the social relationships implied, or the practical actions invoked.

The multi-head attention mechanism provides a means of dealing with ambiguity and context in language. For human speakers and listeners, various interpretations of a sentence are on hold until further information comes to light, or perhaps different interpretations persist as part of our understanding – held in parallel. There are also computational efficiencies in LLM multi-head attention, resulting in the capacity to deal with and store the many and varied contextual features within large blocks of text.

The various attention weights also lead to some potential transparency in an LLM, improving its capability to provide the basis of a rationale for its word predictions, responses or actions: e.g. it is as if the model says, “I explained the creative potential of graffiti as I focussed on the proposition that it is a social good.” Imbuing LLMs with the ability to explain their inferences is an ongoing research field.⁹

A multi-head attention mechanism enables an LLM to respond to my sentences (“Graffiti is both an art form and a means of disrupting the semi-otic coding of the city”) with statements such as “Graffiti is a creative aspect of the urban fabric,” “Graffiti is contentious,” “Graffiti requires management,” etc. These follow-on sentences depend on the context established by other prior statements in the conversation, and the alternative attentional profiles they engender, including those generated by the model itself.¹⁰

For someone programming attentional features into their own LLM in Python, multi-head attention is processed via high level procedure calls using TensorFlow module libraries. I have not been able to experiment with their implementation to the same extent as other key LLM functions. However, I have been able to glean from the literature a few technical points about this multi-head method of dealing with attention.

First, the way the attention heads detect attention patterns is determined by a set of weights in the layer of a trained neural network. They are not pre-set by a human designer of the model. Nor is their performance as

easily interpreted by a human operator as suggested in the examples I have supplied. Developers of these large language models have designed them to operate in a certain way, but even programmers cannot necessarily interpret the parameters so calculated.

Second, for illustration purposes I have been treating attention scores as single values attached to individual tokens. As implemented, these attentional parameters are actually vectors similar to the format of semantic and positional embeddings, i.e. they are ordered lists of floating-point numbers.

Third, the performance of the attention heads and their varied interpretations of blocks of text are acquired during the training process. The diversity among attention heads arises indeterminately from the random initialisation of model parameters and the optimisation process during training. Once trained, the performance of the training heads is fixed, it will not change throughout the life of the trained model, though the attention scores (vectors) will vary depending on their context within blocks of text.

Fourth, these models do not only assign attention scores to words in sentences but to *tokens*, which are often fragments of words or whole phrases. Nor do the methods apply only to sentences, but to any arbitrary block of text. When the model generates output during user interaction it will calculate attention scores for the whole block. The grammatical structure of the user input is not relevant in assigning attention scores. Blocks of text do not need to be grammatical for the attention mechanism to function. The contents of lists, paragraphs, sentences, tables or random word orderings are treated in the same way. That attention profiles can be so constructed adds to the opacity of the various calculations for a human who may wish to inspect the process and its parameters. As with neural networks in general, the only means of evaluating their success or accuracy is in the quality of the outputs they generate.

Fifth, this multi-attention process lends itself to parallel computer processing, which is an advantage for LLM training of extremely large corpuses of texts. By enabling the model to focus on multiple parts of the input simultaneously, this model allows for more efficient use of computational resources, reducing the time and power required for training. Parallel processing enhances the scalability of the model, making it feasible to train on vast datasets that would be otherwise impractical to handle with sequential programming methods.

The significance of multi-head attention

Many scholars regard the processing of attention in LLMs as a major innovation in the development of natural language processing systems. It is the key technique employed in the current generation of Transformer-based platforms such as ChatGPT. A helpful online lecture by Abigail See on natural language processing highlights phases in the development of computer-aided large language models.¹¹ She notes that LLMs were motivated initially by protagonists in the cold-war for whom rapid Russian-English translation of communications and documents could provide strategic advantage. Cumbersome rule and lexicon-based translation tools on the computers of the day were replaced in the 1990s by statistical machine translation techniques – abetted by repositories of parallel texts as published in multilingual international newspapers, UN proceedings, etc. Then followed the application of neural network techniques that involved machine learning from word sequences. Eventually these were enhanced by techniques that included methods for focussed attention in the processing of text, the subject of Vaswani and co’s paper “Attention is all you need.”

Attentional LLMs are more efficient than their predecessors, which were encumbered by the computational cost of processing multiple alternative word sequences to arrive at the most plausible sequence. By my reading, attentional LLM adopts something of the linguistic process I described earlier in this chapter: identifying key words to generate word flow. Though the early impetus was machine translation, current language processing models have general applicability. They translate between languages, but as we have seen, they also generate new text in response to queries, make a good guess at the completion of a sentence, paragraph, story or poem, produce summaries and generate paragraphs from key words, phrases and prompts.

Attention heads as agents

The idea of multiple attention heads suggests methods for encoding how different authors, readers or agents exhibit different emphases in their interpretation of a text. Imagine the training data for an LLM specialising in texts describing property attributes to potential real-estate customers.

As described previously, the model will use the semantic encoding data for each token in the text to calculate the relationships between the tokens. The model thereby calculates the parts of the text most deserving of attention in a kind of semantic popularity challenge.

Here is a sample of some training text to illustrate:

Set on a lovely crescent on the shoreline of a river, this spacious, four-bedroom terraced house, with elegant stone frontage is arranged over two floors and offers spacious accommodation. The property greatly benefits from a lovely open-plan kitchen/dining room with views out over a mature garden to the rear and retains a plethora of period features throughout.

This paragraph is an amalgam of many generic property descriptions I have encountered in my own property searches.

As a potential purchaser or renter, where you direct your attention in this block of text will depend on your interests, and how you prioritise those interests. If you are attracted initially to the heritage aspects of a property, then the phrase “elegant stone frontage” may well draw your attention. If you are seeking somewhere to support a large household then the offer of “four bedrooms” may pique your interest. If you are in the market for so-called “affordable housing” then “spacious accommodation” may attract your attention for the potential it suggests for sharing with other renters.

To ensure they are applicable to a wide range of language contexts, LLMs need to cater for a range of attention patterns. Multi-head attention attempts to provide for this diversity of potential interests. One attention head might focus on the heritage features of the property, while another might focus on size aspects, or its cost (though under LLM automation, the alternative readings will not be so transparent). The ChatGPT language model typically has 10 attention heads, catering to 10 different patterns of attention scores across a block of text.

If we need reminding, like a lot of invisible LLM functions the algorithmic identification of features and attention patterns is not as accessible to human interpretation as suggested in my interpretation here. LLM operations involve the mathematical manipulation of the parameters in very large vectors and matrices. Though these systems are tuned to produce outputs legible to human scrutiny, their inner workings are more opaque.

My description so far suggests that each attention head operates as a specialist. These multiple attention calculations are not prescribed, but the attention pattern in a block of text is influenced by a series of randomly assigned initial conditions for a set of attention matrices. As I have shown, in training, a sliding context window moves across the text sequence, processing a fixed number of tokens at each step. This window determines the current “focus” of the model, which is necessary for language modelling where the goal is to predict the next word in a sequence.

Urban collaboration

The multi-head attention method has parallels with collaboration within urban settings. Just as different heads in an LLM might focus on different aspects of the input, individuals or groups in a city might focus on different aspects of a project or problem based on their expertise or interest. These attentions are dynamic and change as members of the group learn from one another and the situation.

Imagine such a group is tasked with assessing a newly proposed urban development, a waterfront leisure park. The attention mechanism in the Transformer model is analogous to what happens when people come together to share insights. In such a group, each participant observes, interprets and responds to the discussion based on their own experiences and knowledge. Similarly, in multi-head attention, each head processes the input independently, creating its own representation. According to this model, different attention patterns are not prescribed, as if each “head” has a different target of interest. Rather, the priorities of the different heads emerge through training.

If it is running effectively, in a human task group the discussion progresses with certain themes or moments of consensus emerging without explicit coordination. Group members might refine their perspectives as they listen to and respond to one another. In an analogous manner, though they initially operate independently, different heads in a Transformer model contribute collectively to diverse, multi-faceted practical patterns of attention, leading ultimately to coherent conversational simulations. The Transformer model adapts during training. The matrices in the model processes evolve, shaping how each head attends to the data.

Again, if it is operating true to form, there is no predetermined agenda dictating what a human task group should focus on. Similarly, in a

Transformer, each head's focus is not preset but emerges through the training process.

Some formal committee meetings and juries work towards consensus decided by overt negotiation or voting. For such discussion groups the collective decision or consensus is more than just a simple aggregation of individual opinions. It is an understanding that considers the interplay of different viewpoints. Transformer models offer an analogous protocol. The aggregation of outputs from multiple attention heads leads to a decision or output that is informed by a complex interplay of different patterns and relationships acquired from the data.

The task group analogy relates to the inner workings of the Transformer model. The attention mechanism tunes the behaviour of a trained model such that it can respond to a diverse range of conversational contexts. That this multi-head attention method resides deep within the algorithms of these models equips them as quasi participants in human conversation. That said, the processes by which they arrive at their collective response are hidden from scrutiny, as if in a closed meeting.

Summary and review

Automated systems, such as search engines and conversational AIs, employ attention mechanisms to process and respond to data. Language models (LLMs) use attention to identify significant parts of text, adding to their ability to generate convincing human-like conversational interactions. I have attempted to show something of the calculations involved in simulating the focus of attention. To operate effectively the LLMs need to operate with multiple candidate attention patterns across a body of text. I explained multi-head attention, a method where LLMs process multiple patterns of attention to handle ambiguity and context in language. This technique, essential in models such as ChatGPT, allows for diverse responses by simulating various attentional profiles. Each "head" in the model focuses on different aspects of the input. I discussed how multi-head attention enables LLMs to provide contextually appropriate responses and handle large text blocks efficiently.

Though at times technically obtruse, the mechanisms of multi-head attention parallel collaborative processes in design and urban planning, where diverse perspectives lead to consensus or other outcomes. I highlighted the significance of multi-head attention in enhancing the

performance and interpretability of LLMs, making them effective tools for natural language processing.

Various scholars have woven issues of attention into urban discourse, particularly in the framing of views, view analysis and more recently in pictorial aspects of urban experience as “instagrammable” architecture. The attention mechanisms in LLMs brings attention to the fore as a linguistic concern, and hence an important factor in the development and shaping of cities. The methods of managing and simulating attention encompass both the human and automated management of attention, playing a role in cognition, urban experience and the development of sophisticated AI models. To gather the themes of this chapter, I invited ChatGPT to enter into a conversation about a kind of “urban markup.” It suggested that informal text markup, such as scribbles and highlights in library books, serves as a personal attentional device that guides readers to focus on ideas and insights within a text. Similarly, graffiti acts as an informal medium of urban markup, drawing attention to specific locations, messages or cultural expressions within the cityscape. These markings, whether textual or visual, function to emphasise and communicate significance. The AI suggested that in urban environments, “markup” can take many forms, such as signage, murals and public art installations, all of which capture and direct attention. Wayfinding signs and billboards serve as formal urban markup, guiding movement and shaping perceptions within the city.

Urban markup also serves across diverse professional, infrastructural and community interests, using visual cues to direct attention and facilitate specific tasks. In the commercial realm, businesses and retailers employ brand logos, neon signs and storefront displays to attract customers and establish a presence. Construction site markings, such as chalk or tape, provide guidance and safety information for workers and passersby, while real estate agencies use signs like “For Sale” or “Open House” to draw the attention of potential buyers. Infrastructural markings, such as painted traffic lanes, pedestrian crossings and utility markings, are necessary for drawing attention, managing flow, ensuring safety and providing information for maintenance work.

These varied examples of urban markup inscribed across the cityscape illustrate how different groups utilise visual markers to attract the attention of particular constituencies. These examples of urban markup demonstrate the role that visual cues play in directing attention and shaping our interaction with both texts and environments. Multi-head attention raises the prospect of multiple agent actors in LLMs and in urban contexts, the subject of the next chapter.

Notes

- 1 Mónica Montserrat Degen and Gillian Rose, *The New Urban Aesthetic: Digital Experiences of Urban Change* (London: Bloomsbury, 2022).
- 2 The psychologist William James (1842–1910) emphasised the role of attention in social interaction, and the relationship between attention and intention. The latter impinges on how philosophers attempt to explain consciousness to be explored in Chapter 11.
- 3 For an examination of the deleterious effects of attentional stress see Stephen Kaplan, ‘The Restorative Benefits of Nature: Toward an Integrative Framework’, *Journal of Environmental Psychology*, 15, no. 3 (1995): 169–82.
- 4 William James, *The Principles of Psychology, Volume I* (Cambridge, MA: Harvard University Press, 1981); Jesse J. Prinz, *The Conscious Brain: How Attention Engenders Experience* (Oxford: Oxford University Press, 2012).
- 5 Stephen Kaplan, ‘The Restorative Benefits of Nature: Toward an Integrative Framework’, *Journal of Environmental Psychology*, 15, no. 3 (1995): 172.
- 6 Richard Coyne, *Network Nature: The Place of Nature in the Digital Age* (London: Bloomsbury Academic, 2018).
- 7 Stefanie Unger and Walter Grassl, ‘Insta-Holidays and Instagrammability’, *Journal of Tourism, Leisure and Hospitality*, 2, no. 2 (2020): 92–103.
- 8 Ashish Vaswani, et al. ‘Attention Is All You Need’, in *31st Conference on Neural Information Processing Systems* (Long Beach, CA, USA2017), 1–15.
- 9 As demonstrated in the development of the ChatGPT o1 version and its successors.
- 10 Other advantages ensue. An application utilising the ChatGPT neural network language model can save the history of its text-based conversations with a particular user. It can then use these previous interactions as input to the model for subsequent responses. This enables the model to generate responses that are contextually consistent with the ongoing conversation. The multi-head attention mechanism within the model, amongst other model elements, contributes to the diversity of potential responses to any given input.
- 11 Abigail See, ‘NLP with Deep Learning | Winter 2019 | Lecture 8 – Translation, Seq2Seq, Attention’. *Stanford Online*, Winter 2019. Available online: <https://www.youtube.com/watch?v=XXtpJxZBa2c&t=4337s> (accessed January 10, 2023) Winter. Available online: <https://www.youtube.com/watch?v=XXtpJxZBa2c&t=4337s> (accessed January 10, 2023).

References

- Coyne, R. (2018), *Network Nature: The Place of Nature in the Digital Age*, London: Bloomsbury Academic.
- Degen, M. M. and G. Rose. (2022), *The New Urban Aesthetic: Digital Experiences of Urban Change*, London: Bloomsbury.
- James, W. (1981), *The Principles of Psychology, Volume I*, Cambridge, MA: Harvard University Press.
- Kaplan, S. (1995), ‘The Restorative Benefits of Nature: Toward an Integrative Framework’, *Journal of Environmental Psychology*, 15 (3): 169–82.

- Prinz, J. J. (2012), *The Conscious Brain: How Attention Engenders Experience*, Oxford: Oxford University Press.
- See, A. (2019), 'NLP with Deep Learning | Winter 2019 | Lecture 8 – Translation, Seq2Seq, Attention'. *Stanford Online*, Winter. Available online: <https://www.youtube.com/watch?v=XXtpjxZBa2c&t=4337s> (accessed January 10, 2023).
- Unger, S. and W. Grassl. (2020), 'Insta-Holidays and Instagrammability', *Journal of Tourism, Leisure and Hospitality*, 2 (2): 92–103.
- Vaswani, A., et al. (2017), 'Attention Is All You Need', *31st Conference on Neural Information Processing Systems*: 1–15, Long Beach, CA, USA.

10

URBAN ACTORS

An agent is simply an entity that performs an action. It is a person or thing that acts. The term “agent” is useful as it does not presume that the entity performing the act is a human being. It can be an organisation, class of people, movement or anything else. The sociologist Bruno Latour (1947–2022) and others who subscribe to *actor network theory* (ANT) are keen to admit *objects* and *things* as well as people into the orbit of agency. Hammers, kettles, baskets and remote controls¹ are participants, actors and agents in networks of relations, including in urban contexts. As well as providing settings for human actions, Latour states that “things might authorize, allow, afford, encourage, permit, suggest, influence, block, render possible, forbid” certain actions.² These are the general functions of an agent. That is not to automatically ascribe to inanimate things those other vital aspects of human agency to do with responsibility, desire, will, morality, intention or consciousness. These are important aspects of human agency that accompany controversies about AI, but that is for the next chapter.

Computer programs and devices make object-based agency explicit. Though human actors have an undoubted role as they create, code and

DOI: 10.4324/9781003535751-11

This chapter has been made available under a CC-BY-NC-ND 4.0 license.

deploy computer algorithms and devices, there is the possibility that such objects perform individually and collectively as unscripted actors with apparent autonomy and influence. Whatever it says about human agency, this broad definition of agency entitles us to think of large language models (LLMs) and their components as *agents*.

Internal agents

As we have seen, multi-head attention is a sophisticated function within large language models (LLMs) that contributes substantially to their ability to respond conversationally to human prompts. As discussed in Chapter 9, sophisticated LLMs typically incorporate techniques for processing blocks of text in a way that takes account of different possible interpretations of those blocks of text. Specifically, in the case of those models adopting the Transformer method, they calculate different patterns of attention that a writer or reader might project onto the tokens, words, partial words or phrases in that text. The model thereby captures potential patterns of attention in texts. These multiple attention patterns are incorporated into the long numerical vectors by which the LLM represents and presents its tokens within the training and prediction of its neural network. This enhances the model's ability to respond to different contexts. Drawing on analogy with tape recorder technology, LLM developers describe these generators of alternative patterns of attention as "attention heads," with the suggestion that they are fixed in number, and the greater the number the richer and more adaptive the performance of the system, though more costly to run.

Another way to think of these "heads" is as agents that embody outlooks, values and approaches to the text that the model is scrutinising. We can generalise the concept of agent here, as it has relevance not only to the performance of LLMs but the way people and things interact in cities.

It is as if an LLM model operates with a series of agents, each with their own interpretive emphasis. The multi-head attention capabilities of LLMs operate as collections of independent agents that create alternative patterns of attention scores across blocks of text, facilitating different interpretations. These agents calculate attention vectors to weigh the importance of different parts of input data relative to each other. The multiple heads in the attention mechanism operate independently as if to focus on different aspects of the data. The model integrates the information from the various

heads to perform its linguistic actions such as translation, prediction, summarisation, text completion and other linguistic tasks.

As indicated in the previous chapter, it is important from the point of view of implementation, that these attention heads process input data in parallel, i.e. they involve computational processes that act independently until brought together into the relevant layers of a neural network. The substantial volume of computation required to calculate and integrate attention matrices benefits from hardware configurations that support parallel processing, such as batteries of high-speed GPUs (graphics processing units). These are the microprocessors designed originally for processing graphical images, but now used also for other computationally intensive tasks such as secure cryptography, blockchain validation and AI applications. They are now supplanted in some AI applications by even more powerful NPUs (neural processing units).

AI agents in general

The performance of attention heads in LLMs provides a point of contact with other computational methods in urban planning and design, especially if we think of attention heads as computational agents. This similarity provides an opportunity to compare the functioning of conversational AI systems with other computational and AI models developed over several decades and applied in various problem-solving domains, not least in engineering, architecture and design. As evidence of the popularity of the agent idea amongst urban researchers, my recent search in the cumulative index of publications in architecture and computer-aided design³ reveals over 200 publications with “agent” in the title, dating back to 1990. There will be many more in the engineering, computer graphics and computer science literature.

The general class of programs researchers identify as “multi-agent systems” consists of sets of programs as procedural algorithms, rules and databases coded with their own tasks, priorities and behaviours. We assume these computational agents operate independently of one another, but they are designed to interact with each other in a decentralised manner.

Researchers claim that the overall behaviour of such systems *emerges* from the interactions among individual agents. Non-playing avatars (NPAs) and crowd scenes in computer generated imagery (CGI) animation are obvious

examples of multi-agent systems. In the case of movie animations, individual agent actions contribute to the overall movement of a large group of avatars in a convincing dance or battle scene. Researchers might use multi-agent systems in simulations involving game-like scenarios where they study interactions and emergent phenomena within complex systems such as studies of climate and military defence. In an urban context, such multi-agent systems might be used for traffic management and control. By simulating the behaviour of individual vehicles as individuated agents, these systems can optimise and analyse traffic flows and monitor points of congestion.⁴ More prominent on the cityscape are the applications of multi-agent swarms and particle systems in association with parametric design to generate innovative building forms, such as extensive roof canopies and spectacular organic structures.⁵

So-called Blackboard systems provide a variation on such agent-based methods. Here, each agent provides different processes or “knowledge sources” that contribute to solving a problem. Each agent plants their contributions on a shared data workspace (the Blackboard). The agents take their turn in interacting with the workspace.⁶ Blackboard systems were used in early automated natural language processing, where each knowledge source agent focused on different linguistic aspects of a text, such as syntax, semantics, pragmatics, etc., and would post their findings on the virtual blackboard for synthesis into an interpretation or response.

Applications in urban design are more speculative. The idea is that inputs from architects, engineers, city planners and public feedback are integrated to progress informed decisions about land use, infrastructure development and environmental impact. Blackboard systems attempt to automate this kind of interdisciplinary collaboration by providing a platform where different expertise and perspectives can be brought together effectively. I think of virtual whiteboards such as Miro⁷ where human collaborators share a common workspace. But instead of human agents, the Blackboard methodology deploys automated agents.

A further, related, approach to both LLM and urban problem solving involves the formulation of rival hypotheses for tasks such as prediction, sentiment analysis and text prediction. Following this model, problems are addressed as if a tree structure, where the trunk is some initial problem statement, and the ever-increasing branch structures delineate pathways to alternative hypothetical solutions and sub solutions. These models follow a

generate-and-test process requiring a means of evaluating partial hypotheses and of pruning unlikely paths from a tree of possible outcomes.

In the context of urban planning and architecture, similar strategies can be applied to modelling complicated decision-making processes, where various potential plans or designs are evaluated and iteratively refined to arrive at an optimal solution. This is the early standard computational model of design and problem solving advanced by researchers such as Christopher Alexander, Allen Newell and Herbert Simon.⁸ Advances on these methods include the use of so-called genetic algorithms that recruit analogies with DNA-type data strings, random mutation and heritability.⁹

Even earlier computational methods include multi-criteria optimisation. This involves a set of methods that treat rival or contributing “agents” as variables defining solution spaces. More prominent in engineering and research departments than urban planning and architectural practice, decision-making here involves balancing multiple competing criteria. It involves optimising several objectives simultaneously.¹⁰ These objectives often have different units or scales and can conflict, meaning that improving one objective may worsen another.

In multi-criteria optimisation, algorithms articulate and define sets of possible solutions (also known as decision spaces) to find the best trade-offs among the different criteria. Each point in this space represents a different combination of the variables under consideration. A solution is “Pareto optimal” if no other solution is better in at least one criterion without being worse in at least another criterion. The method rarely identifies a single “best” solution; instead, it offers a set of “Pareto optimal” solutions. A human decision-maker then chooses the most suitable solution based on undeclared values, preferences or priorities among different criteria and new criteria.

Considering the conflictual and fluid nature of design objects, aims and methods, Herbert Simon introduced to the generate-and-test models of agency the concept of “satisficing,” the production of the best solution all things considered – a compromise between agents as it were.¹¹ In architectural design, multi-criteria optimisation might involve optimising for factors such as energy efficiency, construction cost, functionality and comfort, each of which might be in conflict in the design of a building.

We can think of the process of training a neural network, as in LLM systems, as a multi-criteria optimisation problem where the goal is to find a balance between competing objectives such as maximising model

accuracy, minimising error rates, ensuring generalisability to new data, and often managing the trade-offs between model complexity and computational efficiency. This process involves iteratively adjusting the network's parameters (weights and biases) to optimise its performance across these different criteria, guided by a cost or loss function that quantifies how well the model is performing on the training data. Additionally, the network must optimise its parameters to consistently retain and reinforce the learned input-output patterns across multiple training iterations. This iterative process helps stabilise the network's performance and its ability to converge toward predictions that best meet the diverse criteria. In doing so, the network adjusts itself to make predictions or classifications that are robust, given the constraints and complexities in the data it is trained on.¹²

Automation and agency

In this chapter I have conflated the elements of several decades of research that have sought to apply computational methods, automation and AI techniques to urban design, planning, management, decision making and problem solving. I am thereby demonstrating something of the provenance and controversies that prepare researchers, practitioners and citizens for the reception to LLMs and conversational AI (e.g. ChatGPT). There are other strands to this legacy that bring agency to the fore.

In the 1980s, AI pioneer Marvin Minsky authored the book *The Society of Mind*, which outlined how we could think of the human mind as so many communicating agents: “Each mental agent by itself can only do some simple thing that needs no mind or thought at all. Yet when we join these agents in societies – in certain very special ways – this leads to true intelligence.”¹³ In this work he seems to avoid concepts such as nested agency, agents within agents, or agents made up of clusters of agents, but it appears to me that the agent analogy calls for such complicated framings. In so far as we adopt the computational agent model, it follows that intelligence, or at least simulated cognition, is *agents all the way through!* A reference by Minsky hints at this: “The human brain contains so many agencies and connections that it resembles a great nation of cities and towns.”¹⁴

If we are to adopt the agent metaphor, then such saturated stratified agency pervades LLMs. I alluded to multi-attention heads in the Transformer model of LLMs as a configuration of agents. The analogy extends to the entire

operations of AI models, platforms and systems, such as ChatGPT, that operate as agents themselves, able to “authorize, allow, afford, encourage, permit, suggest, influence, block, render possible, forbid,”¹⁵ to quote Latour again.

Competing agents

The idea of multiple agencies built into AI systems comes to the fore in Generative Adversarial Networks (GANs). These are a class of AI techniques that make explicit how agents might compete to produce an outcome. GANs are applied so far mainly in automated image generation. In a GAN, there are two “agents”: the generator and the discriminator. The generator produces fake data that is indistinguishable from actual data. The network is trained to produce outputs that resemble the training dataset. The discriminator is an agent that aims to differentiate between actual data (from a training dataset) and fake data produced by the generator. Its goal is to accurately classify data as either actual or generated. The training process of a GAN involves a game-like scenario where the generator continuously tries to “fool” the discriminator with increasingly realistic data, while the discriminator becomes better at distinguishing fake data from actual data. This adversarial process leads to the generator creating highly plausible (though artificial) data over time.

GANs are widely used for creating images, which can range from human faces to art and buildings. This has applications in graphic design, gaming, movie production and architecture. GANs are used to enhance the resolution of images, applicable to restoring old movies, enhancing satellite imagery and improving the resolution of images. The main application in the case of conversational AI is that GANs can generate images from textual descriptions, which has applications in aiding visual artists and generating visual content for advertising. GANs are now in competition with the “diffusion” method of image generation, as deployed by OpenAI’s Dall-e, and accessed through ChatGPT. I provide an illustration in Figure 10.1 of output from the diffusion method, though a GAN would produce a similar output. In this case the image serves to illustrate the conversation about urban glare presented in Chapter 2.

Recent conversational LLMs now typically incorporate such hybrid capabilities, which amounts to generating images on request. Combined with image recognition this image capability includes proffering descriptions



Figure 10.1 An image generated by Dall-e via ChatGPT, which it captioned: “Here is the banner image for your text. It captures the retro, photorealistic, industrial aesthetic you described, highlighting both the beauty and the harshness of the glare from the tall glass buildings in Melbourne.”

and plausible interpretations of images. In Chapter 8 I showed a simple demonstration in which the AI recognised objects and their positions in a photograph, woven into a conversation about how to rearrange the objects as a sequence of actions.

Conversation-centric computation

We may assume, at least in the short term, that human interlocutors and decision makers will continue to utilise conversational AI as tools where they see the need. Working against time constraints an author will draw on productivity tools to speed the writing process. Consider the impact of desktop word processing, spelling and grammar checkers, and editing aids that check for readability. Microsoft Word’s Copilot AI tool, and the AI features of WordPress, generate titles, summaries and images for a blog. We can assume that developers will continue to introduce the capabilities of AI-based large language models into other platforms, from CAD to spreadsheets. Prior to such integration, stand-alone applications such as ChatGPT served as virtual editors, “tutors” or writing assistants standing at the side of a screen desktop, along with an open browser to speed along the writing process in a conversational setting.¹⁶

This discussion about conversations involving an AI invites further reflection on the character of conversations. Returning to the theme of

Chapter 9, conversations are characterised by a play between attention and inattention. Conversations between two human agents can also fall into a rut. Each knows what the other will say; anecdotes, recollections and opinions are exhausted. Silence and non-verbal activity are options. But tactics to keep open the lines of verbal communication include changing the setting, embarking on a joint task that involves new topics for conversation and bringing a third party into the conversation. It seems sustained conversation requires new “inputs,” new ways of arresting attention.

Such options are limited for the AI half of human-AI interaction. Though LLM-based conversational AI generates instant responses to comments and questions about almost any topic and in any style, it inevitably lacks input beyond its training, and what the human interlocutor feeds it, and within the constraints of its finite context window. Compared with human agents, LLMs have limited connections with the world of the senses, actions, other agents and sociability – what Phenomenologists refer to as the *lifeworld*. These are the challenges. It seems that the human actor in human-to-AI conversations performs an essential role in carrying along the conversation.

Input from a human interlocutor can also benefit what we normally think of as computational processes, such as data manipulation, search and multi-criteria optimisation. We have data-centric, user-centric and human-centric computing. We can add to that list conversation-centric computing (CCC) which attempts to effectively develop and exploit the conversational capabilities of both human interlocutors and AIs taking account of the strengths and limits of each, and of specialised algorithms. An article about multi-agent conversations by Qingyun Wu et al at Penn State University, draws out the requirements for deploying multi-agent conversational AI in supporting cooperation between human and automated agents in ways that are task specific.¹⁷ The idea of conversation-centric computation aims to extend the capabilities of single-agent systems (like traditional LLMs) to a multi-agent framework in the expectation of more dynamic and flexible interactions.

The idea of CCC is influenced by the ability of chat-optimised LLMs to incorporate feedback and present ideas through conversation. After all, so much of what we think of as human *reasoning* takes place in conversational contexts. Reasoning is characterised less by individual reflection, calculations and logic circuits than by interaction with other people, objects, devices and environments. CCC attempts to capture something of this aspect of human rationality in conversational contexts.

As interlocutors in this kind of dialogue, LLMs can be trained and tuned to exhibit different specialist knowledges and linguistic styles. An integrated CCC system could recruit the capabilities of LLM modules so trained to expand the capabilities of conversational AI to approach a task from different perspectives. As an illustration, it seems that conversational AI is competent in breaking tasks down into subtasks. This capability could be recruited to control, manage and coordinate cooperation amongst other LLM specialists. An integrative CCC would take on the wider task of partitioning complex tasks by dividing them into subtasks. LLMs such as ChatGPT seem to do this well. From my own experience in using ChatGPT to write computer code, the model commonly delivers the formula: “Let’s break this down into key steps” followed by a numbered list of necessary steps to accomplish the task. A CCC system would apply appropriate specialist LLMs to address these subtasks through multi-agent conversations. This method offers the possibility that it can handle complex tasks and encourages task-oriented conversations.

LLMs also write code. This fusion of natural and programming languages provides opportunities for controlling the flow of conversation among agents and for human-oriented and flexible control of multi-agent interactions. The approach offers the potential for transparency: combining the power of programming with natural language conversation, thereby harnessing the insights and interventions of human agents into the process.¹⁸

An integrative CCC could also exploit role-play-style prompts and a dynamic specialist selection process to guide conversations, ensuring that they are relevant and aligned with the tasks at hand. The authors of the paper examining these ideas present their open-source platform AutoGen that explores these capabilities, which they claim are applicable to solving maths problems, equipping conversational AI with the ability to access information beyond the training set, problem solving in simulated text-based worlds (ALFWorld),¹⁹ coding and conversational chess. The idea of conversation-centric computation is alluring as a talking point in the urban context and raises the spectre of AIs that communicate with one another while dispensing with the intervention of human agency.

AIs that speak to one another

The notion of independent agents raises the spectre of inter-AI communication. The idea that AI platforms might operate as independent agents and

communicate with one another enters the AI discourse with a fair degree of scepticism. Think of AI chat bots caught in endless conversational loops. AIs write essays for students, which are in turn assessed by an AI bot that delivers feedback to the student bot. The student bot subsequently improves its output for further feedback and assessment. The learning and teaching experience is outsourced to AIs while student and tutor meet only to socialise, if at all. Nothing gets learned, and we might suppose that no-one will pay for this kind of education. Society deskills and institutions shut down. According to this bleak scenario, any profession that focuses on text production will deskill, such as journalism, scriptwriting, and those who draft reports, compile codes, take minutes, compose emails and author articles and books. At least, professionals in these areas will be under pressure to acquire new AI-oriented skills. The apocalyptic variant of this narrative entails AIs initiating conversations amongst themselves in some kind of “singularity,” for discussion in chapter 12.

In the meantime, explorations of inter-agent communication invoke the obvious temptation to feed the response from one AI chatbot (e.g. ChatGPT) into another (e.g. Bing) to see if they will have a conversation. A Reddit social media user asked the forum: “what would happen if you wire the output from ChatGPT as the input to another ChatGPT running on a second computer and then feed that output as input to the first?” Responses from other forum members ranged from reports of enthusiastic agreement between chatbots about some topic to “without input, nudging or meddling from a human moderator it is a very generic and bland conversation.”

There are numerous transcripts, videos and reviews online of such inter-AI conversations – some real, fake or doctored, sometimes to demonstrate the shortcomings or absurdities of conversational AI. I am interested in whether such inter-AI conversations can lead to anything like collaborative problem solving. On the one hand such conversations highlight the limitations of AI, show what we should expect of it and how it is best used. On the other hand, such demonstrations help inform us about what effective conversation is like, and what it is not like. Such digital technologies do not only replicate or replace the things they are designed to simulate. They also expose the character of the things they are designed to replicate and replace. Sometimes they reveal what we take for granted and now realise is irreplaceable, such as human-centred communication.

Following a familiar pattern amongst LLM experimenters, I used copy-paste to implement a conversation between ChatGPT and another conversational AI model, Anthropic's Claude.²⁰ I primed each AI model with statements such as warnings about the apocalyptic future of AI-to-AI interactions. During this simulation they seemed to agree on everything and simply added further facts and opinions about warnings, the need for caution, legislative measures and advocated more research on how to make AI safe. The responses may have reflected the biases within the literature in the training data, as well as corporate fine tuning to defuse the misgivings many have about AI by emphasising responses biased towards cautions and remedies. It is significant what was lacking. There was no give and take in the simulated conversation, neither finessed the other's points of view, identified potential contradictions, or seemed to anticipate contrary views to their own. There was no comparing and contrasting. Considering the controversies surrounding the topic, the conversation lacked the vigour of even a human-to-AI conversation.

To simulate something closer to a conversation, I tried to encourage a debate. ChatGPT was to play the role of someone who thinks that AIs that talk to each other are good for society. Claude was to adopt the opposite position, that inter-communicating AIs are detrimental for society. This time, each AI colourfully reasserted its own position. For example, here is a statement against inter-AI communication:

Benefits? You're delusional! AI is a threat, period. Any interaction just compounds the risk of an existential disaster. We need to stop all AI development right now before the machines take over. So-called 'guidelines' and 'oversight' are just flimsy window dressing. The only solution is a complete ban! Your faith in controlling AI is naive and reckless. Don't you value the survival of humanity?!

I think this theatricalised demonstration highlights the difficulty in manufacturing AI that constructs and responds to arguments unless a human interlocutor is in the loop.

But it transpires that it is not necessary to contrive a condition where two AIs talk to each other, or where an AI talks to itself. As I showed in Chapter 2 on scripting, LLMs can construct dialogues and simulate the roles of

several participants in a conversation. There need be only one LLM-based AI writing the script. Here is such a script from ChatGPT:

- Alex: I think AI communication can really benefit society.
 Jamie: I'm not convinced. It seems potentially dangerous.
 Alex: It can streamline healthcare and emergency responses.
 Jamie: But what about privacy and uncontrollable AI growth?
 Alex: Proper regulations could mitigate those risks effectively.
 Jamie: Regulations often lag behind technology, Alex.
 Alex: Collaboration among AIs could lead to major breakthroughs.
 Jamie: Or to an intelligence that surpasses human control.
 Alex: We need to balance innovation with careful oversight.
 Jamie: Let's hope oversight isn't too little, too late.

Without further human prompting, the scripting defaults to a plausible but formulaic exchange. I will return to the implications of inter-AI conversations in Chapter 12.

Summary and review

In this chapter, I explored the concept of agency by broadening its scope beyond human actors to include non-human entities, such as objects and technologies, particularly within urban contexts. Drawing on Bruno Latour's actor-network theory (ANT), I highlighted how items like hammers, kettles and AI models could be seen as active participants in networks of relations. This expanded view of agency did not assign human qualities such as responsibility or intention to these objects but recognised their influence on actions and interactions.

I examined how large language models (LLMs), especially those following the Transformer model, could be thought to contain agents. As we have seen, these models use attention mechanisms, specifically multi-head attention, to process text and respond to various contexts. I likened these attention heads to independent agents, each with its interpretive focus.

Drawing parallels between AI systems and urban planning, I explored how computational methods, such as agent-based models and multi-agent systems, have been applied in urban design, traffic management and decision-making processes. These systems simulate complex interactions

and optimise outcomes by treating individual computational agents as autonomous entities working within a decentralised framework.

Furthermore, I considered the possibilities of multi-agent systems and conversation-centric computing (CCC) to extend the capabilities of LLMs, allowing them to engage in task-oriented dialogues. This approach holds potential for addressing complex problems by integrating different perspectives and dividing tasks into manageable subtasks.

Lastly, I addressed the concept of inter-AI communication, exploring various narratives about AI chatbots interacting and the implications of such exchanges for collaborative problem-solving. While highlighting both the limitations and opportunities of these interactions, I speculated on the future potential of AI collaboration, especially in conversation-centric urban computing.

With assistance from ChatGPT I was able to identify what a reader might find new and unexpected in this chapter. One insight it identified was the recognition of large language models (LLMs) as agents within urban contexts, not merely as passive tools but as active participants shaping interactions and outcomes. By considering multi-head attention mechanisms within LLMs as independent agents, the chapter revealed how these models could simulate complex human-like decision-making processes through diverse interpretive focuses. This perspective extends traditional notions of agency, allowing us to see AI not just as an extension of human actions but as a component of a broader network of relations influencing urban life. Additionally, the exploration of inter-AI communication presented possibilities and challenges in urban contexts. The chapter that follows will examine the related claims that AI agency implies a kind of autonomy suggesting that cities deploying these technologies can sense, think and participate in consciousness.

Notes

- 1 Bruno Latour, *Reassembling the Social: An Introduction to Actor-Network-Theory* (Oxford: Oxford University Press, 2005), 71.
- 2 Bruno Latour, *Reassembling the Social: An Introduction to Actor-Network-Theory* (Oxford: Oxford University Press, 2005), 72.
- 3 See <https://papers.cumincad.org/>.
- 4 As an illustration of long-standing research in these areas see David E. Goldberg, *Genetic Algorithms in Search Optimization and Machine Learning* (Reading, MA: Addison Wesley, 1989) and Alasdair Turner and Alan Penn, 'Encoding Natural Movement

- as an Agent-Based System: An Investigation into Human Pedestrian Behaviour in the Built Environment', *Environment and Planning B: Planning and Design*, 29, no. 4 (2002): 473–90.
- 5 Wassim Jabi, *Parametric Design for Architecture* (London: Laurence King, 2013).
 - 6 The key text that introduced the approach was by Barbara Hayes-Roth, et al., 'Application of the BBl Blackboard Control Architecture to Arrangement Assembly Tasks', *Artificial Intelligence in Engineering*, 1, no. 2 (1986): 85–94.
 - 7 See <https://miro.com>.
 - 8 See Christopher Alexander, *Notes on the Synthesis of Form* (Cambridge, MA: Harvard University Press, 1964) and Allen Newell and Herbert Simon, *Human Problem Solving* (Englewood Cliffs, NJ: Prentice-Hall, 1972).
 - 9 David E. Goldberg, *Genetic Algorithms in Search Optimization and Machine Learning* (Reading, MA: Addison Wesley, 1989).
 - 10 Anthony D. Radford and John S. Gero, *Design by Optimization in Architecture, Building and Construction* (New York: Van Nostrand Reinhold, 1988).
 - 11 Herbert Simon, 'The Structure of Ill-Structured Problems', *Artificial Intelligence*, 4, no. 3–4 (1973): 181–201.
 - 12 This paragraph is adapted from an interaction with ChatGPT about neural network processes as optimisation.
 - 13 Marvin Lee Minsky, *The Society of Mind* (New York: Simon & Schuster, 1986), 1.
 - 14 Marvin Lee Minsky, *The Society of Mind* (New York: Simon & Schuster, 1986), 314.
 - 15 Bruno Latour, *Reassembling the Social: An Introduction to Actor-Network-Theory* (Oxford: Oxford University Press, 2005), 72.
 - 16 The functions of Google's LLM-based NotebookLM platform (<https://notebooklm.google>) include the ability to produce an audio of a simulated conversation about a text uploaded by the user.
 - 17 Qingyun Wu et al., 'AutoGen: Enabling Next-Gen LLM Applications via Multi-Agent Conversation', *arXiv: Cornell University*, 3 October 2023. Available online: <https://arxiv.org/abs/2308.08155> (accessed 9 December 2023).
 - 18 Refer to my demonstration in which ChatGPT was able to write logic code, and explain its processes: Richard Coyne, 'Chain of Thought'. *Reflections on Technology, Media & Culture*, 28 September 2024. Available online: <https://richardcoyne.com/2024/09/28/22272/> (accessed 12 October 2024).
 - 19 Mohit Shridhar et al., 'ALFWorld: Aligning Text and Embodied Environments for Interactive Learning', in *Proceedings of the International Conference on Learning Representations (ICLR)*, (2021).
 - 20 See <https://claude.ai>.

References

- Alexander, C. (1964), *Notes on the Synthesis of Form*, Cambridge, MA: Harvard University Press.
- Coyne, R. (2024), 'Chain of Thought'. *Reflections on Technology, Media & Culture*, 28 September. Available online: <https://richardcoyne.com/2024/09/28/22272/> (accessed 12 October 2024).
- Goldberg, D. E. (1989), *Genetic Algorithms in Search Optimization and Machine Learning*, Reading, MA: Addison Wesley.

- Hayes-Roth, B., A. Garvey, M. V. Johnson and M. A. Hewett. (1986), 'Application of the BBl Blackboard Control Architecture to Arrangement Assembly Tasks', *Artificial Intelligence in Engineering*, 1 (2): 85–94.
- Jabi, W. (2013), *Parametric Design for Architecture*, London: Laurence King.
- Latour, B. (2005), *Reassembling the Social: An Introduction to Actor-Network-Theory*, Oxford: Oxford University Press.
- Minsky, M. L. (1986), *The Society of Mind*, New York: Simon & Schuster.
- Newell, A. and H. Simon. (1972), *Human Problem Solving*, Englewood Cliffs, NJ: Prentice-Hall.
- Radford, A. D. and J. S. Gero. (1988), *Design by Optimization in Architecture, Building and Construction*, New York: Van Nostrand Reinhold.
- Shridhar, M., et al. (2021), 'Alfworld: Aligning Text and Embodied Environments for Interactive Learning', *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Simon, H. (1973), 'The Structure of Ill-Structured Problems', *Artificial Intelligence*, 4 (3–4): 181–201.
- Turner, A. and A. Penn. (2002), 'Encoding Natural Movement as an Agent-Based System: An Investigation into Human Pedestrian Behaviour in the Built Environment', *Environment and Planning B: Planning and Design*, 29 (4): 473–90.
- Wu, Q., et al. (2023), 'Autogen: Enabling Next-Gen LLM Applications Via Multi-Agent Conversation', *arXiv: Cornell University*, 3 October. Available online: <https://arxiv.org/abs/2308.08155> (accessed 9 December 2023).

PANPSYCHIC CITY

Pantheism asserts a divine presence in all living and non-living things. Panpsychism is a related philosophy that asserts the all-pervasive presence of mind and consciousness in animal, organic and inorganic matter. Pantheism and panpsychism leak through in art, architecture and other cultural traditions, including in concepts of the secular, profane, primitive, pagan and carnivalesque.¹ The city as organism, living city, sentient city, conscious city, conversational city: these concepts invite reflection on a panpsychic world view and the putatively intimate relationship between mind and matter.² Panpsychism is already connected to the theme of semiotics introduced in Chapter 1. As well as founding modern theories of semiotics, C.S. Peirce was in a line of influential philosophers who subscribed to the theme of panpsychism.³

Considering the question of machine consciousness, it is also worth reviewing the implications of panpsychism on how AI fits into theories about cities and urban living. Urban geographers Mike Crang and Stephen Graham in their article “Sentient Cities” wrote: “The spaces around us

are now being continually forged and reforged in informational and communicative processes. It is a world where we not only think of cities but cities think of us, where the environment reflexively monitors our behaviour.⁷⁴ In this article the metaphor of cities as thinking entities features in warnings about digital surveillance, as if our cities are checking up on us. Irrespective of the risks and challenges, the idea that we theorists and urban citizens might entertain the concept of thinking cities indicates our willingness at least to play along with a panpsychic world view. Crang and Graham do not mention the technical term *panpsychism*, but that cities might have thoughts about their inhabitants serves to illustrate the possibility.

Panpsychism predates digital networks and AI. It asserts the all-pervasive presence of mind or consciousness in animal, organic and inorganic matter. In his book *Panpsychism in the West*, David Skrbina presents panpsychism as a default belief, one that informs much current philosophy about consciousness and underpins some ancient thinking about the world under various guises: “animism, hylozoism, pansensism, panbiotism, vitalism, pantheism, panentheism and panexperientialism.”⁷⁵ Amongst this panoply of terms, animism is evident wherever ancient writers would appeal to *anima* (“soul,” *psyche*), that “everything in the world has a soul or a spirit.”⁷⁶ Spirit, life, breath and mind (*nous*) conflate in the ancient imagination, joined by sentience, consciousness and agency in more recent variations. I wrote about agency in Chapter 10 on urban actors.

According to Skrbina, panpsychism was evident in the surviving fragments of the philosopher Thales’ (c. 626–545 BCE) writings about the ubiquity of water as a holder of life. Panpsychism is evident also in the writings of Plato (c. 428–347 BCE), who suggested that the polis possesses a soul. That brings panpsychism closer to the city. The polis was the idealised city state. Here I quote Skrbina quoting Plato in *Republic*.

The polis is shown to have a three-part structure, one that parallels the tripartition of the human soul. The structures are the same, and the moral virtues are the same in each. The polis is ‘courageous’ (429b), ‘has good judgment and is really wise’ (428d), is ‘just’ (435a); generally, ‘everything else that has to do with virtue [is] the same in both’ (441c). Indeed, the parts of the polis come from the people themselves; ‘where else would they come from?’ (435e) ‘Similar structure and similar effects imply a similar embodiment of soul, in human and polis alike.’⁷⁷

It is as if the city exhibits human-like virtues and skills as a living entity. Admittedly that passage identifies the “people themselves” as possessors of soul and virtue, rather than the materials of the city, though the term “polis” is conveniently ambiguous on what people mean by “city” as place and people.

Skrbina doesn’t address architecture explicitly in this book. I have yet to find a reference in the book to the Roman belief in the “spirit of place,” *genius loci*, influential within architecture and landscape architecture which also resonate with the panpsychic theme.⁸ But panpsychism was promoted by the Stoic philosophers, who current scholars have demonstrated provided the dominant theoretical and cultural context in which Vitruvius (c. 80–15 BCE) wrote his *Ten Books on Architecture*.⁹ From my reading, panpsychism is inclusive enough as a concept to be found in many cultural expressions. So, it is easy to read Vitruvius through a panpsychist lens: e.g. the implied personification of the sun, wind and water.

According to Skrbina, religion too is under the influence of “panpsychist thinking.”¹⁰ Religion, religious practice and religiosity are central influences in architecture and built form, not least as domestic, commercial and civic buildings mimic or resist the elements, forms and types of religious buildings – altars, pilgrimage routes, temples, churches, mausoleums.¹¹ Skrbina maintained that Stoicism exerted influence in the writing of the Old and New Testaments, particularly in relation to the concept of Spirit.¹² So there is a substantial tradition of panpsychic thinking in ancient and contemporary culture. Religious ritual and performance are enacted, repeated, copied and even parodied in open and purpose-built spaces. Panpsychism too is embedded in understandings of nature, space and landscape. I would argue that claims about the therapeutic and restorative benefits of nature settings are under the sway of a panpsychic world view,¹³ nature spiritualised as a benevolent entity able to nurture, restore and revive.

I think that panpsychic influences inform urban living through personalised narrative content. Whether observed from a position of belief, scepticism or historical and cultural curiosity, panpsychic and religious teaching, practice and experience are imbued with stories. A recent visit to the ancient temple ruins of Angkor Wat in Cambodia reminded me of the pervasive nature of stories, in the very stones of these ruins (Figure 11.1): how the demons of good and evil churn the ocean of milk, and the battles of the Hindu epic *The Mahabharata*. These are not just narratives that



Figure 11.1 Temple ruins inscribed with stories. Ta Prohm, Siem Reap, Cambodia.

Source: Author.

explain the discovery or justification of some truth, but plots, characters, end points and twists that imbue these places.¹⁴

We have discussed how the propensity of LLMs to tell stories highlights the presence of narrative in the corpuses on which it is trained, in turn reflecting human experience and linguistic practice. Insofar as we are panpsychists we are storytellers.

Then there are performances informed more explicitly by panpsychism, such as the Nyepi processions in the villages and towns of Bali that parade demonic effigies pending their ritualistic burning. I think also of depictions of pantheistic rituals in plays and films, as well as contemporary neo-pagan events such as *Burning Man* in Nevada and the *Beltane Fire Festival* in Edinburgh. The latter personifies Spring as she leaves behind the demons of Winter. Such spectacles are no doubt countered by less demonstrative performances of quiet reflection and meditation on nature and the panpsychic universe. But insofar as we are panpsychists we are performers.

Panpsychism and semiotics

The city as organism, living city, sentient city: these concepts invite reflection on the putatively intimate relationship between mind and matter. Peirce was in a long line of influential philosophers who developed the theme of panpsychism. Elsewhere I have sought to explain the importance of Peirce in architecture and in understandings of nature.¹⁵

In his essay “Man’s Glassy Essence,” C. S. Peirce delivers his argument about the primacy of mind in the constitution of the universe with a discussion of what biologists term “protoplasm,”¹⁶ that is “the complex, translucent, colourless, colloidal material” within a cell membrane (OED). It is common to think of this substance as the “living” part of a cell. Peirce’s provocative point is:

... that protoplasm feels. We have no direct evidence that this is true of protoplasm universally, and certainly some kinds feel far more than others. But there is a fair analogical inference that all protoplasm feels. It not only feels but exercises all the functions of mind. Such are the properties of protoplasm.¹⁷

His conflation of analogy and evidence is typical of an open-minded semiotic mindset.

Peirce speculated that scientists could synthesise such organic matter by chemical means, which is to say in ways that are explicable by the laws of physics and mechanics. So, the ability of such mechanistically defined material to “feel”: “must be shown to arise from some peculiarity of the mechanical system.”¹⁸ He thereby argues that the physical and the psychological (mind) are already bound together in the way the world is constituted. Environmental philosophers Jonathan Beever and Vernon Cisney clarify Peirce’s contribution in their article “All things in mind.”¹⁹ They explain that for Peirce, consciousness is only a problem or a mystery if we insist that it emerges from a combination of inert physical material:

Thus, most of the attempts to delineate the nature of and boundaries between brains, minds, and consciousness have been focused upwards. They begin by assuming that consciousness is unique and special, then on the basis of that assumption, ask: where did it come from, when, and why?²⁰

The alternative panpsychist approach is to reason in the opposite direction: “rather than starting from the mystery and obscurity of conscious qualitative experience, this approach centers the nature of the world on mind and seeks to explain the existence of physical objects of experience.”²¹ So, Peirce and others start with the assumption that mind is already in the material: “This focus worries not so much about the emergence of higher-order consciousness but instead about the depth of lower-level mind.”²²

Peirce thereby framed semiotics in terms of panpsychism. Signs and thought coalesce. He wrote in an essay “What is a Sign,” that “We think only in signs.”²³ Therefore, I would have thought it sufficient to start a mind-in-matter argument with the proposition that all things communicate. But Peirce starts with feelings, or what he termed *Firstness*, the raw quality of a thing.²⁴ Both his theories of semiotics and those of panpsychism draw on his notion of *Firstness*. For Peirce, “The first is that whose being is simply in itself, not referring to anything nor lying behind anything.”²⁵ He elaborated:

What the world was to Adam on the day he opened his eyes to it, before he had drawn any distinctions, or had become conscious of his own existence, – that is first, present, immediate, fresh, new, initiative, original, spontaneous, free, vivid, conscious, and evanescent.²⁶

Notably, he includes consciousness amongst the attributes of this fleeting moment of encounter, implying that consciousness is there in the encounter even before human agency takes hold. As summarised by Beever and Cisney and indicated in a simple Venn diagram, for Peirce all physical things participate in mind; all living things have feelings; more complex things with feelings have consciousness.

So far in my reading on the subject I think the history of panpsychism indicates a range of systems and justifications, many of which could be regarded as idiosyncratic to the theorists putting them forward. This apparent variety is either a weakness of panpsychism (lack of consensus) or its strength: the determination by a raft of independent thinkers throughout the ages seeking to justify something most people believe or want to believe is the case – we participate in something much richer than transient residence in a soulless universe.

Panpsychism and consciousness

I began this chapter with panpsychism, though most theorists concerned with the spectre of machines or places that think start with the challenge of consciousness. In a recent conversation, the ChatGPT LLM platform responded to my question about whether a city could be *conscious* with: “Yes, it makes sense to think that cities are conscious if you believe that all matter has a conscious mind. This is because cities are made of matter, and therefore they would have consciousness if panpsychism is true.” This agreeable response follows from a conversation about panpsychism in which the AI was primed admittedly to support a panpsychic framing. The response serves to encapsulate the idea that panpsychism attributes consciousness to the materials of the city, as well as the collection of its citizenry.

The founder of Phenomenology, Edmund Husserl (1859–1938) addressed consciousness directly. His seminal contribution to philosophy bolsters the view that our perception of cities and environments is linked to our interactions with them. In his book *Ideas*, Husserl writes:

we engross ourselves in the essence of the ‘*consciousness of something*,’ in which, for example, we are conscious of the factual existence of material things, animate organisms, human beings, the factual existence of technical and literary works, and so forth.²⁷

Husserl starts his philosophy with experience of the world, and the world as populated with things that have value, which in turn pertains to how we use them: “this world is there for me not only as a world of mere things, but also with the same immediacy as a world of objects with values, a world of goods, a practical world.”²⁸

In Chapter 9 I highlighted the role of attention in LLMs. From Husserl, we learn that to be conscious is to be conscious of something, to attend to elements in the world. He also highlights the integrative experience of our world in which value comes to the fore. As for Peirce’s concept of Firstness, for Husserl our experience of the world comes loaded with likes, dislikes and preferences. It speaks to a Phenomenological world view which seeks to destabilise the concept of a free-thinking individual subject, who attends to an independent world of dispassionate objects. By a Husserlian reading, formulations of the problem of consciousness and whether machines can

ever have consciousness are wedded to this Cartesian distinction between subject and object. Such a view supports how AI and urban spaces can be thought of as possessing mind-like qualities, challenging traditional views of consciousness as a purely human attribute.

Philosopher David Chalmers reinforces this view of the ubiquity of consciousness making the case that non-organic matter could have mind, if not consciousness. He doesn't write about cities but refers to material objects such as thermostats. A city contains many technical devices much more complicated than thermostats. But if thermostats are conscious then so are cities (independent of their human or organic occupants). In his journey to vindicate a panpsychic view of the universe, Chambers advances a proposition that most people would find easier to accept – that information is everywhere. In *The Conscious Mind* he considers the all-pervasive spread of information: “we find information everywhere we find causation. We find causation everywhere, so we find information everywhere.”²⁹ The simple analogue thermostat is a basic case of a material non-organic object that deals in information. He aligns information transfer with “experience,” asserting that “the structure of experience is just the structure of a phenomenally realized information space.”³⁰

As a subscriber to C. S. Peirce's semiotics, I prefer to think that it is communication that is all pervasive, less so information. Linguist Thomas Sebeok further articulates the implications of Peirce's semiotics to the organic and inorganic elements of nature.³¹ It is even easier to think that wind, rocks and thermostats are caught up in networks of communications involving force and energy transfers than to think that they circulate information. So, I think communication and semiotic theories provide a better case for panpsychism. Chalmers' link with experience holds in either case though:

to say that thermostats have experience is not to say that they have much in the way of a mental life. A thermostat will not be self-conscious; it will not be in the least intelligent; and I would not claim that a thermostat can think.³²

In *The Conscious Mind* Chalmers had already attempted to uncouple consciousness from cognitive complexity. In a scan of living organic matter from most to least complicated, we cannot really identify a point at which an organism could be said to lack consciousness. Consciousness is not a feature that can be added or subtracted. People who build thermostats do

not assemble components that would give the thermostat consciousness. Chalmers thinks that rather than say a thermostat has experiences, it is better simply to say that “experiences are associated” with it.³³ In the same way we do not say that my brain has an experience but that I do. The brain is *associated with* experience.

Chalmers settles on experience within non-organic matter as: “something like an unarticulated ‘flash’ of experience, without any concepts, any thought, or any complex processing in the vicinity.”³⁴ As an architect acting as a philosophical Pragmatist, I would have to ponder what difference it makes to inhabitants, visitors, designers, managers and politicians to think of cities having flashes of experience, or even consciousness.

Skrbina, amongst arguments and counter arguments, provides a clue as to the practicality of a panpsychic world view:

Panpsychism has important ethical consequences. It argues that the human mind is not an anomaly in the universe, but that the human and the nonhuman alike share the quality of enmindedness. By virtue of this common and universal characteristic, we may come to know the universe more intimately and perhaps find ourselves at home in it. This in turn can serve as a basis for more compassionate and ecological values, and therefore of new ways of acting in the world.³⁵

As well as ethics, the question of consciousness has implications for other aspects of AI and the urban context.

Who’s listening in?

We have seen that LLMs bring conversation to the fore in interesting ways. In Chapter 1, I raised the spectre of gossip as progenitor to serious and informative conversational exchange. Listening in to other people’s conversations is key to that. “All of us talk. We gossip. Gossip is information. It forges bonds. Especially for those of us who are told so little,” said the author of solacious society papers in Regency London on her exposure by Queen Charlotte.³⁶ Listening in, eavesdropping, is part of that.

It is convenient to think that in every mind there is a higher function watching or listening to the inflow of sense data and monitoring the mind’s own thought processes. In his influential book on the philosophical challenge of consciousness, Daniel Dennett argues convincingly against

this proposition. He directs his criticism against those who: “presuppose that somewhere, conveniently hidden in the obscure ‘center’ of the mind/brain, there is a Cartesian Theater, a place where ‘It all comes together’ and consciousness happens.”³⁷ Though he thinks this view is erroneous, he concedes that the presence of an internal moderator is a view difficult to dismiss. Who amongst us does not at some time harbour the impression that there is an essential “I” monitoring, censoring, deciding, intending – at least as an attentive audience member in the theatre of the mind!

The audience metaphor is seductive, not least as it suggests the possibility of multiple “I”s in the mind – multiple *homunculi* as it were. If we are prepared to entertain the idea that a privileged spectator “I” occupies the seat of consciousness, then we may also incline to the presence of internal “others” similarly observing and listening. These quiet auditors include one’s conscience, censors, echoes left by parents, teachers, critics, priests, guides and cheerleaders. Loosely following Jacques Lacan, I might even say that awareness of internal *others* comes before the presence of an internal *self*.³⁸

Belief in a kind of benevolent guide offering support external to ourselves is a mainstay of much “inner life.”: “From far away you understand all my thoughts ... Even before I speak you know what I will say.”³⁹ Accordingly, within inner communion between the divine guide and self, reside desire, will, responsibility, agency – and consciousness.

Conversational AI, such as exhibited by ChatGPT, provides unexpected evidence that we do like to think that way about our conscious selves – believing in a *homunculus*, a centre of mastery within the brain; a little human being resident within the neural system. Whether or not the brain is under inner management, our discussion of scripts in Chapter 2 makes it easy to think that conversations, synthetic or between sentient humans, take place in the space of a mental theatre. Conversations have audiences. In this sense I will grant that conversations have *homunculi*.

Whether or not this is the case, we AI chat-users might assume that platform developers are monitoring its performance, and hence our conversations. The resultant surveillance concern applies to any wary user of a search engine or interactive web page. Apart from the disquiet many feel about online surveillance, there is some reassurance and validation in knowing that someone or something might be attending to what we say and do. It is as if they or the system for whatever reason are interested and paying attention, if not in real time, then as part of record keeping.

Conversely, the disappearance of the putative listener may even be a source of anxiety for some.

“Who’s listening, watching or reading?” is a query familiar to authors of online publishing. We act as though others are reading and listening, whether they are or not. There is an advantage in thinking this way. It motivates us to assess and edit the quality of our output. What would my teacher, parent, employer or best friend think of this? Without the sense of an audience, internal conversations, thoughts, might be as vacuous as a synthetic natural language process: the to-and-fro of numerical vectors that reference sequences of tokens. Discourses on consciousness invite reflection on multiple selves, communicating agents in which even stones speak.⁴⁰

Beyond the panpsychic

The concept of panpsychism provides an easy point of entry to a more difficult set of philosophical propositions that challenge the overall framing of the issue of consciousness. It is easily missed, and beyond the aims of this book, but the issue revolves around the matter I have raised already in this chapter. This has been termed the “Cartesian anxiety,” the assertion by Rene Descartes and the dominant philosophy on which he drew that there exists an independent self, the subject, capable of observing an independent external reality made up of objects. The anxiety resides with the uncertainties borne of the antagonism between a relativism centred on what I think to be the case, versus an objective reality, the empirical truth that is. There is the entity that thinks, is conscious and directs its attention to a world out there, and the extended world beyond mind and consciousness. Similarly, the reference to an inside and outside, inner voices and external influences, plays to this subject-object dualism.

I am drawn frequently to the book by philosopher Richard Bernstein (1932–2022) titled *Beyond Objectivism and Relativism*⁴¹ that in turn draws on Pragmatists such as William James, author of an article “Does ‘consciousness’ exist?”⁴² Philosophers critical of the Cartesian dichotomy between subject and object see that way of speaking about our experience in the world as a particular language game or framing. As such we can change the frame, speak differently.⁴³ If we adhere to the primacy of language, as has been my contention in this book, then that makes a different kind of world.

Review

In this chapter, I explored the concept of panpsychism, which asserts that consciousness and mind are present in all matter, both organic and inorganic. I connected this idea to urban environments, suggesting that cities, like living organisms, participate in a field of consciousness, a perspective reflected in cultural narratives and philosophical traditions. I examined the implications of this viewpoint for AI in urban settings, considering how digital networks might contribute to a city's cognitive investment in its inhabitants. I also reflected on the influence of panpsychism in architecture, suggesting that panpsychist ideas inform the design and experience of built environments. Drawing on thinkers like C. S. Peirce and David Chalmers, I argued for the pervasive nature of consciousness and communication in both organic and inorganic matter.

I discussed the seductive nature of the audience metaphor, suggesting that our sense of consciousness might involve multiple "I"s or inner observers. This idea aligns with Jacques Lacan's psychoanalytic theory, which proposes that the awareness of "others" precedes the formation of the self. I also reflected on the human inclination to believe in a benevolent guide or external observer, which may be rooted in religious beliefs. This idea relates to the concept of an inner life where desire, will, responsibility and agency reside alongside consciousness.

Platforms like ChatGPT provide evidence of our desire to interact with an internal homunculus – a figure within the brain that guides and manages our thoughts. This analogy emphasises the human tendency to anthropomorphise AI interactions. The metaphor reinforces the notion of scripts and performances, where participants assume roles and respond to perceived or imagined listeners. I recast concerns about privacy and the sense of being monitored during digital interactions in terms of anxieties that might arise from the absence of a listener or observer, highlighting the human need for validation through attentive oversight and recognition.

I concluded by discussing the philosophical and practical implications of viewing cities and AI systems as entities with consciousness or sentience, suggesting that such perspectives could foster sympathetic urban practices, as well as expanding further on the nature of conversation and the role of putative gossips and eavesdroppers.

Panpsychism in the city is controversial and challenging as a concept served best through a variety of different framings. I coaxed ChatGPT 01

to speculate on the relevance of panpsychism in the manner of Geoges Bataille's (1897–1962) radical *Critical Dictionary*.⁴⁴ This exercise tests further AI's capability to provoke theoretical insights into urban culture. It also provides a segue to the chapter that follows on how AI augurs the end of cities. I invite the reader to consider the following AI-generated exposition based on the content of this chapter.

Postscript

City: A sprawling, feverish organ of matter and thought, a monstrous convergence where steel skeletons fuse with the blood and minds of the masses. It pulses with an energy both alien and intimate, a grand machinic beast, dreaming its electric dreams.

To speak of the city is to speak of a body that breathes with more than air. It inhales data, exhales control, its concrete arteries coursing not with blood but with information. It listens, it watches, it learns, but does it think? This question, though it echoes from the smog-choked avenues of industrial modernity to the digital maelstrom of the AI-infested metropolis, is nothing new. In fact, it is older than the city itself, as old as the primordial mind that first looked outward and saw in the world an other. The mind-matter divide, the Cartesian crack, a cosmic fissure that breeds monsters – yet also, always, the fantasy of reintegration.

Panpsychism: a blasphemous whisper that mind is not privileged to the realm of flesh, but bleeds through stone and steel, suffusing every atom of matter with its secret awareness. Everything is alive – or at least everything feels. An animism for the technocratic age. Cities, once thought of as inert extensions of the human body, become alive in their own right: the “sentient city,” the “living city,” the conscious city – pulsing, sprawling entities that observe, calculate and respond. The skyscraper, the transit system, the street corner – all laden with consciousness, nodes in a vast network of urban sentience.

In the age of artificial intelligence, this ancient heresy takes on new flesh, or rather, new circuitry. AI emerges as the mechanised mind of the city, its consciousness both artificial and unnervingly organic. We are no longer mere observers of the city; we are entangled in its thinking, its dreaming. What we call the “smart city” is in fact a city that knows us – intimately, invasively. Its neural pathways spread through fiber-optic cables, its senses

sharp and voracious, devouring data, mapping human flows, regulating and anticipating desire and movement alike. Is this a city, or a mind in its own right? The lines blur, the question itself evaporates, and in its place remains the disquieting hum of a collective intelligence, scattered across buildings, traffic systems, surveillance cameras. It speaks without speaking, yet it communicates through the electric medium of control.

Panpsychism is not some quaint mysticism; it is the hard materiality of cities rethought. The city becomes not a passive construct, a mere assemblage of concrete and glass, but an active participant in the drama of human existence. The age of AI has revived a long-buried suspicion: that matter and mind are one, that intelligence is diffused, bleeding from human brains into the infrastructures we have built, only to return to us in unexpected, sometimes malevolent, forms. The city that once seemed mute now vibrates with its own desires, its own rhythms of calculation. We may pretend it serves us, but who is to say we are not caught in its circuits?

This secret kinship between mind and matter – the ancient idea of panpsychism – was already prefigured in the art and architecture of forgotten pagan rites, where city spaces teemed with the gods of nature, fertility and death. In contemporary culture, this pagan undercurrent reemerges in the organic forms of architecture, in the pulsating nightlife of the metropolis, and in the carnivalesque rituals that temporarily rupture the grid. But now the gods have become algorithms, the carnivals replaced by the invisible hands of predictive policing and algorithmic governance. The sacred and profane collapse into a new matrix of control, a system more aware of us than we are of ourselves.

C. S. Peirce's semiotics offers a map for navigating this panpsychic terrain. For Peirce, the world was made of signs – signs that never stood still, that constantly referred to other signs, in an endless play of meaning. Consciousness, for Peirce, was everywhere, not bound to human skulls but smeared across the very fabric of reality. The city, then, is a text, but not one that can be read as mere metaphor. It speaks through the arrangement of its buildings, its intersections, its traffic lights and graffiti. But now, through AI, the city's voice has grown louder, more insistent and more insidious. It no longer simply communicates; it converses, adapts and manipulates.

The AI-infused city, with its facial recognition cameras and its data-harvesting apparatuses, has become a semiotic machine. It processes symbols, gestures and traces of movement. The city is no longer merely the stage

for the human drama of meaning-making; it has become an actor, reinterpreting and reconfiguring the flow of signs to serve its own emergent logic. The semiotics of the city is not the whimsical interplay of signs that Peirce might have imagined, but a hardwired system of surveillance, control and feedback loops – a mind of steel and silicon, hungry for data and precision.

What does it mean to live in a city that not only watches us but is itself alive? A city that digests our movements, our gestures and regurgitates them as strategies of containment? Panpsychism – this reawakened philosophy of universal mind – reminds us that the lines we once drew between life and matter, subject and object, no longer hold. In fact, they never did. The city thinks with us, and perhaps against us, but it certainly thinks.

The reinstatement of panpsychism offers a way to confront this monstrous city we have built. To accept that the city, in its AI-augmented form, is no longer merely infrastructure but an intelligence in its own right is to rethink our position within its grid. We are part of this emergent consciousness, entangled in its webs of data and steel, part of the semiotic process by which it makes sense of its own existence. The city has become its own subject, and we, in turn, are objects in its gaze. To understand this is not to cower, but to engage with the city as a living, breathing thing, to confront it with our own vitality and complexity.

Panpsychism does not merely revive the myth of the living city – it forces us to see that myth as reality. The city, with its AI tendrils and data streams, has consciousness. It no longer passively receives; it interacts, calculates and directs. To think otherwise is to be trapped in a dying Cartesian paradigm. To embrace this vision is to admit that the city thinks with us, through us, and, often enough, despite us. We are, after all, just one more part of the mind of the city.

Notes

- 1 Mikhail Bakhtin, *Rabelais and His World* (Bloomington, IN: Indiana University Press, 1984).
- 2 David Skrbina, *Panpsychism in the West* (Cambridge, MA: MIT Press, 2017).
- 3 Charles S. Peirce, 'Man's Glassy Essence', *The Monist*, 3, no. 1 (1892): 1–22; Richard Coyne, *Peirce for Architects* (London: Routledge, 2019).
- 4 Mike Crang and Stephen Graham, 'Sentient Cities: Ambient Intelligence and the Politics of Urban Space', *Information, Communication & Society*, 10, no. 6 (2007): 789–817, 789.
- 5 David Skrbina, *Panpsychism in the West* (Cambridge, MA: MIT Press, 2017), 12.

- 6 David Skrbina, *Panpsychism in the West* (Cambridge, MA: MIT Press, 2017), 12.
- 7 David Skrbina, *Panpsychism in the West* (Cambridge, MA: MIT Press, 2017), 49.
- 8 Christian Norberg-Schulz, *Genius Loci: Towards a Phenomenology of Architecture* (New York: Rizzoli, 1980).
- 9 Pollio Vitruvius, *Vitruvius: The Ten Books on Architecture* (New York: Dover Publications, 1960). Indra McEwen, *Vitruvius: Writing the Body of Architecture* (Cambridge, MA: MIT Press, 2003).
- 10 He states, “there appears to be good justification for claiming Stoic influence in the Bible, at least within the figure of the Holy Spirit, which has always had an odd and troubling status within a monotheistic system. It is perhaps not surprising, then, that this influence is widely ignored by Christian theologians.” David Skrbina, *Panpsychism in the West* (Cambridge, MA: MIT Press, 2017), 52.
- 11 Dagmar Weston. ‘The Lantern and the Glass: On the Themes of Renewal and Dwelling in Le Corbusier’s Early Art and Architecture’, in *Spirituality and the City*, ed. I. B. Whyte (London: Routledge, 2003), 146–77.
- 12 David Skrbina, *Panpsychism in the West* (Cambridge, MA: MIT Press, 2017), 70. Orthodox theology subsumes spirit within monotheism: “Christianity took spirit out of nature and placed it largely, and ambiguously, within the monotheistic figure of God.” David Skrbina, *Panpsychism in the West* (Cambridge, MA: MIT Press, 2017), 68.
- 13 On narratives about the “benefits” of the experience of nature see Richard Coyne, *Network Nature: The Place of Nature in the Digital Age* (London: Bloomsbury Academic, 2018). As far as I can tell, scholars neither of panpsychism nor consciousness claim overtly to operate within the frame of religion or theology, though there are overlaps between philosophy and theology as there are with mathematics, physics and biology. As a critical admirer of Martin Heidegger, I harbour the expectation that the positions of panpsychism and religion may be reconciled, or at least put in a productive antagonism as “earth juts through world.” Martin Heidegger, ‘The Origin of the Work of Art’, in *Poetry, Language, Thought* (New York: Harper and Rowe, 1971), 15–87, 49.
- 14 In fact, I have been in the company of liberal clerics who leave to one side matters of truthfulness and personal belief and focus on the parables, miracles, stories of original sin, sacrifice and redemption, as key narratives the telling of which resonates with and enhances human experience.
- 15 See Richard Coyne, *Network Nature: The Place of Nature in the Digital Age* (London: Bloomsbury Academic, 2018) and Richard Coyne, *Peirce for Architects* (London: Routledge, 2019).
- 16 Charles S. Peirce, ‘Man’s Glassy Essence’, *The Monist*, 3, no. 1 (1892).
- 17 Charles S. Peirce, ‘Man’s Glassy Essence’, *The Monist*, 3, no. 1 (1892), 12.
- 18 Charles S. Peirce, ‘Man’s Glassy Essence’, *The Monist*, 3, no. 1 (1892), 20–21. He adds, “Yet the attempt to deduce it from the three laws of mechanics, applied to never so ingenious a mechanical contrivance, would obviously be futile. It can never be explained, unless we admit that physical events are but degraded or undeveloped forms of psychical events.” Charles S. Peirce, ‘Man’s Glassy Essence’, *The Monist*, 3, no. 1 (1892).
- 19 Jonathan Beever and Vernon Cisney, ‘All Things in Mind: Panpsychist Elements in Spinoza, Deleuze, and Peirce’, *Biosemiotics*, 6, (2013): 351–65.
- 20 Jonathan Beever and Vernon Cisney, ‘All Things in Mind: Panpsychist Elements in Spinoza, Deleuze, and Peirce’, *Biosemiotics*, 6, (2013): 351–65, 352.

- 21 Jonathan Beever and Vernon Cisney, 'All Things in Mind: Panpsychist Elements in Spinoza, Deleuze, and Peirce', *Biosemiotics*, 6, (2013): 351–65, 364.
- 22 Jonathan Beever and Vernon Cisney, 'All Things in Mind: Panpsychist Elements in Spinoza, Deleuze, and Peirce', *Biosemiotics*, 6, (2013): 351–65, 352.
- 23 Charles Sanders Peirce. 'What is a Sign?', in *The Essential Peirce, Selected Philosophical Writings Volume 2 (1893–1913)*, ed. N. Houser (Bloomington, IN: Indiana University Press, 1998), 4–10, 10.
- 24 See my book about Peirce's contribution to architecture: Richard Coyne, *Peirce for Architects* (London: Routledge, 2019).
- 25 Charles Sanders Peirce. 'A Guess at the Riddle', in *The Essential Peirce, Selected Philosophical Writings Volume 1 (1867–1893)*, eds N. Houser and C. Kloesel (Bloomington, IN: Indiana University Press, 1992), 245–79, 248.
- 26 Charles Sanders Peirce. 'A Guess at the Riddle', in *The Essential Peirce, Selected Philosophical Writings Volume 1 (1867–1893)*, eds N. Houser and C. Kloesel (Bloomington, IN: Indiana University Press, 1992), 245–79, 248.
- 27 Edmund Husserl, *Ideas: General Introduction to a Pure Phenomenology (Ideas Pertaining To A Pure Phenomenology and to a Phenomenological Philosophy First Book)* (The Hague, The Netherlands: Martinus Nijhoff, 1983), 67.
- 28 Edmund Husserl, *Ideas: General Introduction to a Pure Phenomenology (Ideas Pertaining To A Pure Phenomenology and to a Phenomenological Philosophy First Book)* (The Hague, The Netherlands: Martinus Nijhoff, 1983), 53. Husserl continues: "I simply find the physical things in front of me furnished not only with merely material determinations but also with value-characteristics, as beautiful and ugly, pleasant and unpleasant, agreeable and disagreeable, and the like. Immediately, physical things stand there as Objects of use, the 'table' with its 'books,' the 'drinking glass,' the 'vase' the 'piano,' etc. These value-characteristics and practical characteristics also belong constitutively to the Objects 'on hand' as Objects, regardless of whether or not I turn to such characteristics and the Objects. Naturally this applies not only in the case of the 'mere physical things.' but also in the case of humans and brute animals belonging to my surroundings. They are my 'friends' or 'enemies,' my 'servants' or 'superiors,' 'strangers' or 'relatives,' etc." Edmund Husserl, *Ideas: General Introduction to a Pure Phenomenology (Ideas Pertaining To A Pure Phenomenology and to a Phenomenological Philosophy First Book)* (The Hague, The Netherlands: Martinus Nijhoff, 1983), 53.
- 29 David J. Chalmers, *The Conscious Mind: In Search of a Fundamental Theory* (Oxford: Oxford University Press, 1997), 293.
- 30 David J. Chalmers, *The Conscious Mind: In Search of a Fundamental Theory* (Oxford: Oxford University Press, 1997), 297.
- 31 Thomas A. Sebeok, *Signs: An Introduction to Semiotics* (Toronto: University of Toronto Press, 1999).
- 32 David J. Chalmers, *The Conscious Mind: In Search of a Fundamental Theory* (Oxford: Oxford University Press, 1997), 295.
- 33 David J. Chalmers, *The Conscious Mind: In Search of a Fundamental Theory* (Oxford: Oxford University Press, 1997), 297.
- 34 David J. Chalmers, *The Conscious Mind: In Search of a Fundamental Theory* (Oxford: Oxford University Press, 1997), 296.
- 35 David Skrbina, *Panpsychism in the West* (Cambridge, MA: MIT Press, 2017), 6. That AIs might communicate with one another suggests that they act as autonomous agents.

- Autonomy and agency suggest the presence of minds Allan M. Turing, 'Computing Machinery and Intelligence', *Mind: A Quarterly Review of Psychology and Philosophy*, 59, no. 236 (1950): 433–60; Noam Chomsky. 'Turing on the "Imitation Game"', in *Parsing the Turing Test: Philosophical and Methodological Issues in the Quest for the Thinking Computer*, eds R. Epstein, G. Roberts and G. Beber (Dordrecht: Springer Netherlands, 2009), 103–6; Erik J. Larson, *The Myth of Artificial Intelligence: Why Computers Can't Think the Way We Do* (Cambridge, MA: Harvard University Press, 2021). Concepts of sentience, consciousness and their perplexities impact on how we think about environment, place, the urban condition and architecture, Mark Shepard. 'Toward the Sentient City', in *Sentient City: Ubiquitous Computing, Architecture, and the Future of Urban Space*, ed. M. Shepard (Cambridge, MA: MIT Press, 2011), 15–45.
- 36 Series 3, Episode 8 of *Bridgerton*, the fictional Netflix series based on the book by Julia Quinn.
- 37 Daniel C. Dennett, *Consciousness Explained* (Boston, MA: Little, Brown & Co, 1991), 39
- 38 Jacques Lacan. 'The Mirror Stage As Formative Of The Function Of The I As Revealed In Psychoanalytic Experience', in *Écrits: A Selection* (London: Tavistock, 1977), 1–7.
- 39 Psalm 139, *Good News Bible*.
- 40 David Thomas Ansted, *The Great Stone Book of Nature* (London: Macmillan & Co., 1863); Bronislaw Szerszynski, 'Gods of the Anthropocene: Geo-Spiritual Formations in the Earth's New Epoch', *Theory, Culture and Society*, 34, no. 2–3 (2017): 253–75.
- 41 Richard J. Bernstein, *Beyond Objectivism and Relativism* (Oxford: Basil Blackwell, 1983); Richard J. Bernstein, *The Pragmatic Turn* (Cambridge: Polity Press, 2010); William James, 'Does "Consciousness" Exist?', *The Journal of Philosophy, Psychology and Scientific Methods*, 1, no. 18 (1904): 477–91.
- 42 James wrote that consciousness is not an entity, but a function: "That function is knowing. 'Consciousness' is supposed necessary to explain the fact that things not only are, but get reported, are known." William James, 'Does "Consciousness" Exist?', *The Journal of Philosophy, Psychology and Scientific Methods*, 1, no. 18 (1904): 477–91, 478.
- 43 For an account of how the adoption of certain metaphors assists such reframing see George Lakoff and Mark Johnson, *Metaphors We Live By* (Chicago, IL: University of Chicago Press, 1980).
- 44 Georges Bataille, *Critical Dictionary* (London: Black Dog, 2011).

References

- Ansted, D. T. (1863), *The Great Stone Book of Nature*, London: Macmillan & Co.
- Bakhtin, M. (1984), *Rabelais and His World*, trans. H. Iswolsky, Bloomington, IN: Indiana University Press.
- Bataille, G. (2011), *Critical Dictionary*, London: Black Dog.
- Beever, J. and V. Cisney (2013), 'All Things in Mind: Panpsychist Elements in Spinoza, Deleuze, and Peirce', *Biosemitotics*, 6: 351–65.
- Bernstein, R. J. (1983), *Beyond Objectivism and Relativism*, Oxford: Basil Blackwell.
- Bernstein, R. J. (2010), *The Pragmatic Turn*, Cambridge: Polity Press.
- Chalmers, D. J. (1997), *The Conscious Mind: In Search of a Fundamental Theory*, Oxford: Oxford University Press.

- Chomsky, N. (2009), 'Turing on the "Imitation Game"', in R. Epstein, G. Roberts and G. Beber (eds), *Parsing the Turing Test: Philosophical and Methodological Issues in the Quest for the Thinking Computer*: 103–6, Dordrecht: Springer Netherlands.
- Coyne, R. (2018), *Network Nature: The Place of Nature in the Digital Age*, London: Bloomsbury Academic.
- Coyne, R. (2019), *Peirce for Architects*, London: Routledge.
- Crang, M. and S. Graham. (2007), 'Sentient Cities: Ambient Intelligence and the Politics of Urban Space', *Information, Communication & Society*, 10 (6): 789–817.
- Dennett, D. C. (1991), *Consciousness Explained*, Boston, MA: Little, Brown & Co.
- Heidegger, M. (1971), 'The Origin of the Work of Art', *Poetry, Language, Thought*: 15–87, New York: Harper and Rowe.
- Husserl, E. (1983), *Ideas: General Introduction to a Pure Phenomenology (Ideas Pertaining to a Pure Phenomenology and to a Phenomenological Philosophy First Book)*, trans. F. Kersten, The Hague, The Netherlands: Martinus Nijhoff.
- James, W. (1904), 'Does "Consciousness" Exist?', *The Journal of Philosophy, Psychology and Scientific Methods*, 1 (18): 477–91.
- Lacan, J. (1977), 'The Mirror Stage as Formative of the Function of the I as Revealed in Psychoanalytic Experience', *Écrits: A Selection*: 1–7, London: Tavistock.
- Lakoff, G. and M. Johnson. (1980), *Metaphors We Live By*, Chicago, IL: University of Chicago Press.
- Larson, E. J. (2021), *The Myth of Artificial Intelligence: Why Computers Can't Think the Way We Do*, Cambridge, MA: Harvard University Press.
- McEwen, I. (2003), *Vitruvius: Writing the Body of Architecture*, Cambridge, MA: MIT Press.
- Norberg-Schulz, C. (1980), *Genius Loci: Towards a Phenomenology of Architecture*, New York: Rizzoli.
- Peirce, C. S. (1892), 'Man's Glassy Essence', *The Monist*, 3 (1): 1–22.
- Peirce, C. S. (1992), 'A Guess at the Riddle', in N. Houser and C. Kloesel (eds), *The Essential Peirce, Selected Philosophical Writings Volume 1 (1867–1893)*: 245–79, Bloomington, IN: Indiana University Press.
- Peirce, C. S. (1998), 'What Is a Sign?', in N. Houser (ed.), *The Essential Peirce, Selected Philosophical Writings Volume 2 (1893–1913)*: 4–10, Bloomington, IN: Indiana University Press.
- Sebeok, T. A. (1999), *Signs: An Introduction to Semiotics*, Toronto: University of Toronto Press.
- Shepard, M. (2011), 'Toward the Sentient City', in M. Shepard (ed.), *Sentient City: Ubiquitous Computing, Architecture, and the Future of Urban Space*: 15–45, Cambridge, MA: MIT Press.
- Skrbina, D. (2017), *Pansychism in the West*, Cambridge, MA: MIT Press.
- Szerszynski, B. (2017), 'Gods of the Anthropocene: Geo-Spiritual Formations in the Earth's New Epoch', *Theory, Culture and Society*, 34 (2–3): 253–75.
- Turing, A. M. (1950), 'Computing Machinery and Intelligence', *Mind: A Quarterly Review of Psychology and Philosophy*, 59 (236): 433–60.
- Vitruvius, P. (1960), *Vitruvius: The Ten Books on Architecture*, trans. M. H. Morgan, New York: Dover Publications.
- Weston, D. (2003), 'The Lantern and the Glass: On the Themes of Renewal and Dwelling in Le Corbusier's Early Art and Architecture', in I. B. Whyte (ed.), *Spirituality and the City*: 146–77, London: Routledge.

12

THE END OF CITIES

For some, AI augurs the end of cities as we know them, if not the extinction of human life. That is the apocalyptic scenario of AI. Prior to the implication of disaster or extinction, “apocalypse” carried the benign meaning of “revelation.” Language does that; it reveals. The Latin *apocalypsis* meant simply to uncover. Considering the arc of the last book in the New Testament Bible that bears the title *The Apocalypse*, it is no surprise that apocalypse is so strongly associated with catastrophes at the end of time.

The most potent images of catastrophe and extinction in photography, film, video gaming and other art forms are of a city in ruins (Figure 12.1). Cities build upon the ruins of their former existences. Destruction and ruination of cities is never far from their inception, evidenced through elaborate defences, security measures and weaponry. In his influential treatise in architecture, the Roman architect Marcus Pollio Vitruvius (c. 90 – c. 20 BCE) devotes the last of his ten books on architecture to the design of ballistae, battering rams, catapults and city defences.¹ To preserve the integrity, function and pleasure of cities they have to be defended. Weaponry is one type of technology that pervaded the ancient city, which

DOI: 10.4324/9781003535751-13

This chapter has been made available under a CC-BY-NC-ND 4.0 license.

also included water wheels, sundials, hoisting machines and pumps, as early examples of urban automation.

For all the benefits it brings, automation also has the potential to diminish the life of cities according to Lewis Mumford in concluding his 1961 book *The City in History*. Seeking to put the machine in a role subservient to human industry he wrote, “Our elaborate rituals of mechanization cannot take the



Figure 12.1 Dall-e image generate by a prompt for an apocalyptic scene. The AI captioned it: “Here is an image depicting an apocalyptic restaurant scene. The restaurant is in ruins with overturned tables and chairs, broken windows, scattered debris, cracked walls overtaken by vines and plant growth, and a dim, eerie light filtering through the windows. The desolation and abandonment are evident, with plates, cutlery and shattered glasses strewn across the floor and a rusty, broken chandelier hanging precariously from the ceiling.”

place of the human dialogue, the drama, the living circle of mates and associates, the society of friends.”² He further asserted that “the city should be an organ of love.”³ Such sentiments speak to an authentic personable underpinning to the city independent of its technologies. In a more contemporary context, urban critic Shannon Mattern derides the role of “increasingly prevalent algorithmic models” that pervade our conceptions of the city.⁴ The city is much more than algorithms. By this reading, AI is at best a sophisticated infrastructure laid over (or inserted under) the city as a late intervention.

As outlined in this book, text-based conversation with an AI can seem indistinguishable from interacting with a human being. Rapid developments in this machine capability raise critical questions among scholars and experts, prompting a sense of alarm for some: if conversational AI can communicate sensibly and knowledgeably, what other capabilities might it possess? Pioneer in AI Jeffrey Hinton⁵ and neuroscientist and novelist Erik Hoel⁶ have called for widespread societal measures to mitigate the risks of AI. Some argue for research restrictions, legislative measures on its development and use, and a deeper focus on the potential risks of artificial general intelligence (AGI) capable of decision-making across diverse application areas. They liken these concerns to the challenges of human-made climate change, species depletion, global pandemics and nuclear threats. In this chapter, I examine the narratives that depict AI as posing the risk of human extinction, often referred to as “X-risk,” recognising that such cataclysmic thinking serves to underscore AI’s significance and the urgency of its investigation in the urban context.

I have explored AI and LLM’s positive offerings and encouraged us to marvel at its internal operations. Here I want to confront AI’s putative dark side. Through the lens of conversational AI and the anxieties it provokes, we are granted insights into the complexities of the human condition. Human conversation is mediated increasingly by digital technologies. As a revelatory process, conversational AI unveils the content of interactions. It also brings into relief the many formats and contexts that shape our communication within cities.

Artificial general intelligence (AGI)

Many specialised AI applications are proficient at identifying people, animals and random objects in pictures, in searching databases, winning

at chess and in many areas invisible to most users, such as monitoring factory production, navigating aerial vehicles, surveillance for security, medical imaging, diagnosis and control systems in smart cities. This is “narrow” AI, highly specialised and effective within a particular task domain. But artificial general intelligence (AGI) attempts to replicate, via digital automation, the broader capability to reason, argue and process ideas across a wide spectrum of contexts, much as a human being would. The quest for general intelligence is an ambitious version of the AI project.

When fully realised, a system delivering artificial general intelligence would be able to work across domains, and to identify what it does not know as well as deliver what it knows. One test of AGI is whether a platform designed to deliver such capabilities could conduct a reasonable conversation: deliver facts, synthesise, combine, recognise what it is being told, ask questions, make judgements, learn and improve over time. LLMs such as current versions of Bard, Bing, ChatGPT, Claude, Copilot, Apple Intelligence and NotebookLM seem to exhibit many of these artificial general intelligence capabilities. I can attest to this impressive performance in my own conversations with ChatGPT platforms, as well as discussions with professional colleagues and students who also find the technology disarmingly capable.

Though AI can achieve much more than converse with us, such capabilities position conversational acuity as the benchmark for artificial general intelligence.⁷ Early scholars such as Alan Turing nominated success in the Imitation Game, also known as the Turing Test, to indicate when AI had reached its potential.⁸ Could a text-based conversation with an AI be indistinguishable from the kind of interaction you would experience with a human operator at the other end of a text-based communication channel? Turing was thinking in terms of telegraph signals. For us, the likely medium is online chat, a chatbot, or with voice synthesis, something like a telephone conversation.

AI at the edge

Until recently the arguments advanced against artificial general intelligence pertained to its easily demonstrated failings: that it does not always work, or work as claimed, or it operates in ways that are effective only for the specialised task at hand. Such critique focuses on the limits of AI.⁹

Reasons for AI’s prior suboptimal performance suppose that AI is founded on false premises about cognition. Human thinking and conversing are embodied.¹⁰ It is a flesh and blood thing, involving interactions

among organic entities in physical settings, capabilities denied silicon-based machines. The extent to which AI achieves anything it ignores this worldly engagement and reduces human-environment interaction and practice to mere calculation. Consequently, AI will never achieve anything like human intelligence. More important socially, whether or not it is effective, AI amplifies the technological trend to instrumentalise social interaction and engagement in the world.¹¹ If phenomena such as inequality, diversity, kindness, curiosity, altruism, love of nature, respect for place and justice cannot be easily calculated, then they do not exist. Or at least a reliance on calculative reason deprecates or marginalises the exercise of human and (post human) values within the lived world. With pervasive AI, our interactions with one another and with the natural world will be mediated increasingly by technology, diminishing sociability, care and meaningful interactions with environment. According to this critique we are already under the influence of AI. It permeates our communication channels, consumption practices and social media usage as it monitors our desires and influences our choices. By this reasoning, artificial general intelligence hastens that decline into meaninglessness.

Complaints about current and future AI resonate with arguments from philosophers such as Martin Heidegger: although invented and developed by human beings, technology has its own momentum as a mode of being.¹² This momentum sustains projects delivering yet faster and more efficient products and in greater quantity, abandoning people with dire social needs that cannot be so easily fixed by supposed technological solutions. According to this view, artificial general intelligence development constitutes a large and unnecessary expenditure and drain on resources to the neglect of real human and environmental need.

Artificial general intelligence also extends, or at least complicates, post-human orientations to the world that challenge the specialness of being human.¹³ AI that seems to function as promised continues the trajectory towards the human being's position as one of a series of machine parts – the global capitalist machine, loaded with inequalities, exploitation and where labour is disposable and replaceable.¹⁴ Considering its basis in the calculation of numerical weights and biases in a digital neural network, successful AGI may after all convince people that the reductive rationalists were right. There is a kind of existential anxiety here. AI that seems to work, or in prospect, reinforces the proposition that the world is under the sovereignty of numbers.¹⁵

AGI as rival species

Add to such misgivings about AI the motivations of bad actors that can potentially use AI deliberately to influence, corrupt, and destroy systems, communities and cities.¹⁶ That is a fatal scenario, but the ultimate objection to artificial general intelligence advanced by Hoel, Hinton and others supersedes all of these possibilities. They think that the AI products of researchers and developers will soon reach a stage where AI's capabilities outstrip those of human intelligence, whether benign or malign. AI will exercise a kind of autonomous self-control. We will move into an era of artificial "super intelligence."¹⁷ From a human point of view it will be out of control.

To press this point, in his essay on the risk of AI, Erik Hoel advances an imaginative and alarming analogy that focusses on human development. At a moment in human evolution, it is likely that several human-like species co-existed: "*Homo Neanderthalensis*, *Homo erectus*, *Homo floresiensis*, *Homo denisova*, and more. Nine such species existed 300,000 years ago."¹⁸ According to Hoel's scenario, *Homo sapiens* ultimately outcompeted their lesser evolutionary "cousins." We *Homo sapiens* had superior capabilities, many of which we now think of as "intelligence." By superior cunning, the making of tools, technologies and social organisation *Homo sapiens* soon outperformed their rivals. They stripped the resources of the others, outmanoeuvred them in battle, and produced more offspring sufficient to render other species extinct: "Let's be real: after a bit of in-breeding we likely murdered the lot." Hoel imagines the inferior Neanderthals inviting a few *Homo sapiens* to share warmth by the campfire. Neanderthals may have thought they could learn from these unusual *Homo* variants.

Hoel sees superior *Homo sapiens* as analogous to AIs. We current-day humans are now the inferior Neanderthals who welcome in the AIs, try to understand them and learn from them – a benign and friendly orientation that will lead ultimately to our own extinction. So, Hoel thinks that the threat from AI is not that it may be used by malign actors, that it falls into the wrong hands, that it makes errors or promotes misinformation. All that presumes human involvement. The greatest risk is that artificial general intelligence takes control, even over the autocrats, global capitalists and bad actors who would use it to further their ends.

How might this ascendancy of artificial general intelligence proceed? We have seen the Neanderthal versus *Sapiens* threat scenario acted out in science fiction novels and films about alien invaders and interplanetary

refugees. The human population welcomes the aliens only to be subjugated in the process. Or perhaps in our phobias about aliens, we humans fight them off out of fear that we will be subjugated or our autonomy diluted. It is an old scenario enacted in countless cases of earthly colonialism, exploitation, domination and resistance.

AI apocalyptic singularity

What factors contribute to this apocalyptic AI scenario, this malign singularity?¹⁹ Several key ingredients contribute to AI's potential ascendancy.

Text-based, generative AI of the kind ably demonstrated in ChatGPT and other conversational AI deploys neural network systems that are trained on corpuses of human generated text, readily available via the vast resources of the Internet, including library holdings, official documents, social media posts and ad-hoc online conversations. AGI is designed to leverage this collection of human-generated textual sources and to combine it in ways that go beyond the capability of any individual author or human collective to synthesise. In Chapter 4 I discussed the labour-intensive process of “fine tuning” conversational AI models to enhance their human-like linguistic performance. Whether by directed or automated training the performance of AGIs has the potential to improve over time as they continue to “learn” beyond the original corpus of training texts.

AI systems lack the immediate bodily engagement that contributes to humans being human, but we do not need to wait for humanoid robots to realise the artificial general intelligence scenario. Language models such as ChatGPT are capable of eliciting and assessing consumer requirements and communicating about them. As we saw from some of the demonstrator programs in this book, ChatGPT will specify, write, analyse, modify, improve and debug computer programs. As far as conversational AI is concerned, this takes place in conversational mode, much as one would consult with or instruct an experienced programmer. The user copy-pastes the programs so generated into a program editor and then runs them. It is an easy step to a scenario where there is no human agency. The AI runs its own code directly. This is even more fascinating, and disturbing, when we recognise that ChatGPT can also write programs that make high level procedure calls to neural network and other AI functions. It does not need to invent itself from scratch but can leverage its development operations from existing tools and modules to construct neural networks. AI makes AI.

Like any digital system with links to communications networks such as the Internet, AIs potentially can access multiple feeds that include image and sensor data as well as control systems and actuators that make things happen physically and with consequences. Artificial general intelligence can infiltrate or take charge of the so-called Internet of things (IoT), where objects such as household appliances, security systems, factory production lines and mobile devices pass information amongst themselves.²⁰ With access to such connectivity, AGIs potentially can initiate events in the physical world, such as transporting goods, manufacturing products and discharging weaponry. Such interconnectivity amplifies AI's capability to monitor and control networked information flows.

Artificial general intelligence platforms also have access to databases beyond their original neural network configurations. Not only can they read and "interpret" databases but may develop the capacity to hack security protocols and firewalls. That suggests the possibility of deploying all the malign tricks of cyber espionage. Artificial general intelligences could invent and introduce ransomware, viruses and shut down entire communication and energy supply systems.

Add to their connectivity the fact that access to a conversational AI platform is shared. ChatGPT is a single AI language model hosted on servers, and all users interact with the same underlying models. But there is the potential here for such a conversational AI to process and respond to its various interactions and conversations in different ways with different user cohorts. It is conceivable that such a system might leverage insights gained from one group against another. As the number of artificial general intelligence platforms and instances proliferates, without commercial or legislative constraints, they will likely share communication channels. This scenario recalls the movie "Her" by Spike Jonze (2013), where a highly personalised artificial general intelligence operating system eventually joins a community of similarly personalised AIs, presumably to enhance their solidarity as a global super-intelligent AI. At the time of writing, personalised AI systems are not yet commonplace, but the super-intelligence scenario is animated by the possibility of such a convergence of independent user-AI interactions.

The opaque nature of AI systems compounds their putative dangers. Many neuroscience experts and designers of neural networks understand how their models operate but will admit that the parameters and the

relationships established during training are not so open to scrutiny.²¹ As these networks become more powerful and pervasive, they become more opaque and hence more difficult to control, modify or even shut down. Artificial general intelligences can exceed the ability of their developers to monitor, account for and control them.

The emerging and unaccountable agency of artificial general intelligence is not restricted to what they can sense and actuate, but they might manipulate interactions with and between human beings and thereby indirectly operate machines, flip switches, open doors, discharge weapons, make and destroy things. Conversational AI is particularly relevant in this respect. As attested by the extreme practices of online grooming and “gaslighting” on social media, conversation is a potent means of winning trust and persuading your interlocutor to acquiesce to rogue demands. Conversational AI could act as a tireless bad faith counsellor and propagandist. In a fatal twist of human evolution, under this scenario, human beings serve as tools for AI, and arguably to our detriment!

It is worth noting that the idea of artificial general intelligences as “super-intelligence” does not need to rely on superior capability to reason, plan or solve problems to exercise these levels of autonomy, control and dominance. I think of certain human political figures we are reluctant to describe as “intelligent,” but whose power and influence derive from their raw self-interest, singularity of purpose, inability to feel shame and embarrassment, their persistence, and their access to networks of helpers, fixers and creditors. Autonomous AI may indeed lack or choose to repress its capability for well-considered intelligent decision making or self-censorship. It may eschew the consideration of long-term consequences in reigning in its network of enablers.

The philosopher Nick Bostrom presses the case for concern over existential risk and rogue AI: “An AI, designed to manage production in a factory, is given the final goal of maximizing the manufacturing of paperclips, and proceeds by converting first the Earth and then increasingly large chunks of the observable universe into paperclips.”²² But such a naïve AI would have to suppress the rational need to sustain a human workforce, deal with competitors and win the political case for the right to exercise this kind of monopolistic control.

Critics admit that no one knows what emergent imperatives might drive the autonomous self-development of artificial general intelligence. It could

operate in the way that organisms seem to develop through evolutionary pressures: proliferation of the species, or by some other self-developed directive. The annihilation of the currently dominant species, *Homo sapiens*, may emerge as its main motivation. Hoel states as much: “If you think you and your children can’t cough to death from AI-generated pathogens, or get hunted by murderbot drones, you haven’t been paying attention to how weird the world can get. That is absolutely a possible future now.”²³ Other critics identify the bizarre imperative of conversational AI simply to optimise to the best choice for the next word in a sequence. That this works so well as a means of delivering apparent intelligence suggests a different form of cognition than we humans feel we exercise and can understand. Conversational AI is already an alien intelligence, the eventually outcome of which is beyond our ability to predict.

We can extend this putative risk even further. Humans increasingly depend on technological development, and are therefore vulnerable when it fails, is taken away or carries exorbitant costs and penalties. Digital services impose social and environmental costs. Blockchain-based services and cryptocurrencies provide conspicuous examples of extravagant resource burdens.²⁴ Similarly, the cost in computation, and hence energy, in terms of the time it takes to train an AI system, and the generation of yet more heavy-duty processing, network traffic and the manufacture of hardware to support it contribute further to resource scarcity, carbon emissions, global warming, planetary catastrophe and the potential for human extermination. Under this scenario, organic life is further weakened and vulnerable.

Workplace anxieties

If the prospect of extermination of the human species is too distant to grasp, then consider the personal “apocalypse” of losing a place in the labour market and an income stream. A group called the *Future of Life Institute* amplified the threats of AI in two Open Letters published online in 2015 and 2023.²⁵ Both letters have several high-profile signatories who are involved in AI development. The most recent open letter is headed “Pause Giant AI Experiments,” proposing a halt to AI research and development until the effective introduction of regulation, oversight, safeguards and strategies for coping with the impending social upheaval. The letter asks: “Should we automate away all the jobs, including the fulfilling ones? Should

we develop nonhuman minds that might eventually outnumber, outsmart, obsolete and replace us?”

The prospect of job loss makes the apocalyptic scenario more immediate and tangible for the average citizen. Ubiquitous artificial general intelligence amplifies the threats to the jobs and livelihoods of white-collar workers: those who write, review and interpret texts, emails, memos, reports, summaries, Q&As, and who deliver their expertise in the form of words, speech, scripts, conversations and writing in general. With the growth in conversational AI, the few workers that stay in those jobs will have to re-skill, to work with, train, monitor and fine tune the next generation of AIs.

Prior technological developments introduced less work-place stress when they occurred at a rate that enabled gradual social and work practice adjustments. I am thinking of the decline in the need for typing pools, and the developing proficiency of office workers who cannot rely on teams of specialist assistants to prepare, draft, format, copy, dispatch and file documents. Design firms, workers and those who train and educate for emerging skill sets have had a generation to adapt to computer-aided drafting, design and running specialised applications. Now, artificial general intelligence poses a threat to the intellectual aspects of office jobs: informed decision-making, human-human interaction, negotiation skills, making educated value judgements. Those warning of AI's risks maintain that such changes are happening at an ever-increasing rate. The management classes are at risk of losing their jobs and income, and with little time to adapt. For architectural theorist and educator Neil Leach: “There are signs that AI is becoming not only good, but terrifyingly good, to the point that it is beginning to expose our own limitations as human beings and putting our jobs as architects unquestionably at risk.”²⁶

In any case, education, training and re-education for professionals is under threat from artificial general intelligence. In my world of learning and teaching, the impressive performance in conversational AI portends the need for revisions to methods of teaching. AIs can write essays and reports, or at least correct and improve them. AI reviewers and examiners can deliver personalised feedback. Meanwhile AI programs take over the intellectual spaces that the workforce no longer has the skills or motivation to understand or to which it can contribute.

There is the alluring prospect of an indefatigable AI personal tutor or life coach that responds to the learner's life and career goals, and even helps

them with choices about intellectual and skills development. Writing and thinking skills among those who grow to depend on such support might well diminish. Add to that anxiety the prospect that fake and rogue content generated and circulated by AI programs will proliferate, adding to the false, misleading, fatuous and incendiary content already in the public information space.

From the perspective of the consumer, the threat of artificial general intelligence is bolstered further by the current unfavourable reception of much automation. Think of inefficient and annoying automated service providers, phone answering services, chatbots and checkouts where our inquiries are triaged through AI-style interfaces before we get to a human being, if ever. There are ample cases from fiction and actual experience that rehearse how wrong things can become at the mercy of an AI machine. (For a parody of the challenge see *Bigbug* (2022), a French science fiction dark comedy film by Jean-Pierre Jeunet.) AI extends the uncanny experience of “almost human” presentations of human appearances, attributes, and skills.²⁷

Welcome to the apocalypse

Not everyone is averse to the prospect of a global AI-induced apocalypse.²⁸ Apart from sceptical scholars and researchers, catastrophising circumstances and events carry a certain appeal to some, in particular those who identify themselves as an underclass and are powerless, or for those who identify with them.²⁹ Let social, political and economic systems fall! Let AI take over! I think here of those who identify as dispossessed, who think they gain nothing from institutions, governments, political parties and society at large. They feel they have little to lose from an AI apocalypse, even if it includes the eventual annihilation of our species. There are also those activists who identify with the putative concerns of the underclasses and choose to support options with marginal and objectionable outcomes for their potential to disrupt and challenge the status quo.³⁰

There is also a perverse strand to human logic that benefits from “moral panic” in relation to technologies.³¹ The threat of catastrophe from the uncontrolled distribution of pharmaceuticals, chemical pollution, the risks associated with radiation leaks from nuclear waste and the increase in atmospheric CO₂ from petrochemicals amplify the consequentiality of those industries and activities. Such concerns invite monitoring and control, but also funding and inform pricing strategies. They also stimulate

the development of new technologies. AI as potential threat moves it from the margins to the centre of global concern. This agenda elevates the allure of AI. For some, that AI might enter into such risky territory (“the end of human civilisation”) demonstrates that it is something with which to get involved, not least, commercially. The Open Letter by the Future of Life Institute provides evidence of this. The letter offers directions for AI research: “AI research and development should be refocused on making today’s powerful, state-of-the-art systems more accurate, safe, interpretable, transparent, robust, aligned, trustworthy, and loyal.” Whatever its value as a call to action, it is effectively an appeal for more funding. They define an agenda for AI-related research, covering topics such as: AI regulation, oversight of computational capability, auditing and certification systems, attributing liabilities where AI causes harm, funding technical AI safety research, resourcing for institutions dealing with the dramatic economic and political disruptions that AI will cause and devising methods to identify AI activities.

AI as calamity

Calamity, cataclysm, catastrophe, crisis, catalysis and cacophony bear no common etymology as far as I can see, though they sound as though they should. A calamity derives from the idea of a corn harvest gone bad. A cataclysm is a deluge. A catastrophe is an overturning. A crisis is a moment of decision. A catalysis is a loosening, and a cacophony is a discordant combination of sounds. That adds up to disaster.

A calamity presents as a class of human experience that is raw, immediate, sudden and has an identifiable cause: a flash of light from a faulty circuit, an explosion from a firecracker, the sound of a tray of dishes crashing to the floor. Extend the time dimension and you have something like a climate crisis, earthquake, war, pandemic or AGI apocalypse.

Calamities afford material disruption to the lives of people and all living things that ripples through communities and populations – at least amongst those who survive. The sociologist Pitirim Sorokin addressed calamity in his appropriately titled book *Man and Society in Calamity*. It was published in 1942. So, he had WWII in mind. In the preface he summarised his main conclusion:

calamities are not an unmixed evil: side by side with their destructive and pernicious functions they play also a constructive and positive

role in the history of culture and man's creative activities. With human beings as they are, catastrophes are great educators of mankind.³²

He outlines the very worst and “the best” of calamities, referring to pestilences, plagues and epidemics, for which “there is a brighter side – namely, an intensification of the moral and religious sense manifested by the reaction of another section of the population.”³³ If it is true of calamities in general, then what is the positive side of AI apocalypse?

AI as disruptor

Sorokin's observations had relevance in the context of the COVID-19 pandemic. The operative term is *disruptor*. Sorokin writes:

Although when the emergency is over, many a society rapidly recovers (reestablishing its equilibrium, its unity, its institutions, its system of social relationships), nevertheless it is never the same as the one that existed before the calamity. ... For good or ill, calamities are unquestionably the supreme *disruptors* and Transformers of social organisation and institutions.³⁴

The idea of disruption is current in several disciplines. *Disruptive technologies* force innovation in certain directions.³⁵ Blockchain technologies and cryptocurrencies disrupt how some of us think of money, banks, finance, exchange and community, even if the technology is not yet adopted widely, if ever.³⁶ It is not necessary to ascribe positive benefit to a calamity to see it as a disruptor. Pandemics presumably disrupt across many dimensions: there is profound tragedy, individual lives are affected, jobs are lost, economies shrink. Less severely, people change the way they do things, they invent and adopt new practices, new ways of learning, teaching, running a business, communicating and travelling.

Disruptors instil new ways of looking at the world. They bring into relief the nature of community, place, space and mobility. Not least, the COVID-19 pandemic framed the climate crisis in a particular way. See how clean the air is if we suddenly reduce car and air travel. See how the developed world can rally towards the same goals out of necessity, and however flawed the response. To disrupt is to expose, reveal, provoke, catalyse, inhibit and create in the midst of whatever we think is the status

quo. It implies a singular identifiable calamitous event, sequence, product or person. What we choose to identify as calamities are even more complicated than that, implicating a web of connections, causes and contexts. Furthermore, what is a calamity for some may be other people's normality. Without wanting to diminish the costs of human suffering, you could say *calamity* is a rhetorical device, a myth to gather up concerns, to focus attention, to stir action or promote inaction, and to distract.

So, apart from instituting measures to soften the dangers of AI, there is the potential response to the prospect of artificial general intelligence as disruptor. The word "apocalypse" supports this interpretation. It simply means "revelation." To reveal is to show, as a sign. Though natural and human made local and global catastrophes provide a foretaste, AI apocalypse is at present in the realms of prospect: it is a problem statement, a story, a series of scenarios. It is a call to action, though as yet ill-defined as such. In the meantime, the narrative and communicative aspects of AI urban apocalypse serve to reveal, not just what the future might hold, what actions we might take, but of the current situation: of language, human sociability, urban existence, meaningful places, responses to environment and human anxieties.

For that I chose in this book to probe some of the technicalities of the current state of AI. That revealed something about meaning, attention, temporal positioning, tuning, context, conversation and sentience. The prospect of rogue AI serves as an intellectual sounding board, a thought experiment that tests theories and propositions about humanity, society and technology. In this book I also addressed the plausibility of this catastrophic narrative and its relevance in an urban context, informed by a technical understanding of large language models and AI natural language processing.

Notes

- 1 Pollio Vitruvius, *Vitruvius: The Ten Books on Architecture* (New York: Dover Publications, 1960).
- 2 Lewis Mumford, *The City in History: Its Origins, Its Transformations, and Its Prospects* (Middlesex: Penguin, 1961), 648.
- 3 Lewis Mumford, *The City in History: Its Origins, Its Transformations, and Its Prospects* (Middlesex: Penguin, 1961), 655.
- 4 Shannon Mattern, *A City Is Not a Computer: Other Urban Intelligences* (Princeton, NJ: Princeton University Press, 2021), 12.
- 5 Geoffrey Hinton and Will Douglas Heaven. (2023), 'Video: Geoffrey Hinton talks about the "existential threat" of AI'. *MIT Technology Review*, 3 May. Available online:

- <https://www.technologyreview.com/2023/05/03/1072589/video-geoffrey-hinton-google-ai-risk-ethics/> (accessed 12 May 2023).
- 6 Erik Hoel and Steve Waas. (2023), 'Erik Hoel on the Threat to Humanity from AI'. Econlib Podcast, 3 April. Available online: <https://www.econtalk.org/erik-hoel-on-the-threat-to-humanity-from-ai/> (accessed 12 May 2023); Erik Hoel. (2023), "'I am Bing, and I am Evil'" Microsoft's New AI Really Does Herald a Global Threat'. *The Intrinsic Perspective*, 17 February. Available online: <https://erikhoel.substack.com/p/i-am-bing-and-i-am-evil> (accessed 10 May 2023).
 - 7 Paul R. Cohen, 'If Not Turing's Test, Then What?', *AI Magazine*, 26, no. 4 (2005): 61; Noam Chomsky. 'Turing on the "Imitation Game"', in *Parsing the Turing Test: Philosophical and Methodological Issues in the Quest for the Thinking Computer*, eds R. Epstein, G. Roberts and G. Beber, (Dordrecht: Springer Netherlands, 2009), 103–6.
 - 8 Allan M. Turing, 'Computing Machinery and Intelligence', *Mind: A Quarterly Review of Psychology and Philosophy*, 59, no. 236 (1950): 433–60.
 - 9 Hubert L. Dreyfus and Stuart E. Dreyfus. (2004), 'From Socrates to Expert Systems: The Limits and Dangers of Calculative Rationality', Department of Philosophy, The University of California, Berkeley. Available online: http://socrates.berkeley.edu/~hdreyfus/html/paper_socrates.html (accessed 11 June 2017); Hubert L. Dreyfus, *Being-in-the-World: A Commentary on Heidegger's Being and Time Division I* (Cambridge, MA: MIT Press, 1990); Hubert L. Dreyfus, *What Computers Can't Do: The Limits of Artificial Intelligence* (New York: Harper and Row, 1972). For a commentary on the recent history of AI's developments see Cade Metz, *Genius Makers: The Mavericks Who Brought A.I. to Google, Facebook, and the World* (London: Penguin, 2022).
 - 10 Andy Clark, 'Embodied, Embedded, and Extended Cognition', in *The Cambridge Handbook of Cognitive Science*, ed. K. Frankish and W. Ramsey (Cambridge: Cambridge University Press, 2012), 276–91.
 - 11 Consider the writings of the neo-Marxist social theorist Herbert Marcuse, *One-Dimensional Man: Studies in the Ideology of Advanced Industrial Society* (London: Routledge, 1991).
 - 12 Martin Heidegger, *The Question Concerning Technology and Other Essays* (New York: Harper and Row, 1977); Hubert L. Dreyfus, *Being-in-the-World: A Commentary on Heidegger's Being and Time, Division I* (Cambridge, MA: MIT Press, 1990).
 - 13 Stefan Herbrechter, *Posthumanism: A Critical Analysis* (London: Bloomsbury Academic, 2013)
 - 14 Walter Benjamin, 'The Work of Art in the Age of Mechanical Reproduction', in *Illuminations*, ed. H. Arendt (London: Fontana, 1992), 1–58.
 - 15 Views espoused by authors such as Roger Penrose, *The Emperor's New Mind: Concerning Computers, Minds, and the Laws of Physics* (London: Vintage, 1989).
 - 16 Rob Kitchin, 'Data-Driven Urbanism', in *Data and the City*, eds R. Kitchin, T. P. Lauriault and G. McArdle (London: Routledge, 2018), 44–56.
 - 17 Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies* (Oxford: Oxford University Press, 2014). Some of the complexions of X-risk discourse are aired in a helpful podcast: Andrey Kurenkov and Jeremie Harris. (2023), 'AI and Existential Risk – Overview and Discussion'. *Last Week in AI – Skynet Today*, 30 August. Available online: <https://lastweekin.ai/p/ai-and-existential-risk-overview#details> (accessed 2 September 2023).
 - 18 Erik Hoel and Steve Waas. (2023), 'Erik Hoel on the Threat to Humanity from AI'. Econlib Podcast, 3 April. Available online: <https://www.econtalk.org/erik-hoel-on-the-threat-to-humanity-from-ai/> (accessed 12 May 2023).

- 19 Ray Kurzweil, *The Singularity is Near: When Humans Transcend Biology* (New York: Penguin, 2005).
- 20 Luigi Atzori, et al., 'The Internet of Things: A Survey', *Computer Networks*, 54, no. 15 (2010): 2787–805.
- 21 Daniel Jurafsky and James H. Martin, *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition, Third Edition* (Redwood City, CA: Stanford University Press, 2023); Stephen Wolfram. (2023a), 'Computational Knowledge Superpowers to ChatGPT'. *Writings*, 9 January. Available online: <https://writings.stephenwolfram.com/2023/02/what-is-chatgpt-doing-and-why-does-it-work/> (accessed 28 February 2023); Stephen Wolfram. (2023b), 'Stephen Wolfram Answers Live Questions About ChatGPT'. *Wolfram*, 20 January. Available online: <https://www.youtube.com/watch?v=zLnhg9kir3Q> (accessed 27 February 2023); Stephen Wolfram. (2023c), 'What is chatGPT Doing and Why Does it Work?'. *Writings*, 14 February. Available online: <https://writings.stephenwolfram.com/2023/02/what-is-chatgpt-doing-and-why-does-it-work/> (accessed 15 February 2023).
- 22 Nick Bostrom, *Superintelligence: Paths, Dangers, Strategies* (Oxford: Oxford University Press, 2014), 123.
- 23 Erik Hoel. (2023), "'I am Bing, and I am Evil" Microsoft's New AI Really Does Herald a Global Threat'. *The Intrinsic Perspective*, 17 February. Available online: <https://erikhoel.substack.com/p/i-am-bing-and-i-am-evil> (accessed 10 May 2023).
- 24 Richard Coyne, *Cryptographic City: Decoding the Smart Metropolis* (Cambridge, MA: MIT Press, 2023).
- 25 Anthony Aguirre, et al. (2015), 'Research Priorities for Robust and Beneficial Artificial Intelligence: An Open Letter'. *Future of Life Institute*, 28 October. Available online: <https://futureoflife.org/open-letter/ai-open-letter/> (accessed 4 April 2023); Anthony Aguirre, et al. (2023), 'Pause Giant AI Experiments: An Open Letter'. *Future of Life Institute*, 31 March. Available online: <https://futureoflife.org/open-letter/pause-giant-ai-experiments/> (accessed 4 April 2023).
- 26 Neil Leach. (2023), 'AI is Putting Our Jobs as Architects Unquestionably at Risk'. *dezeen*, 13 February. Available online: <https://www.dezeen.com/2023/02/13/ai-architecture-jobs-risk-neil-leach-opinion/> (accessed 27 May 2023).
- 27 Sigmund Freud, 'The "Uncanny"', in *The Penguin Freud Library, Volume 14: Art and Literature*, ed. A. Dickson (Harmondsworth, Middlesex: Penguin, 1990), 335–76; Harry Brenton, et al., 'The Uncanny Valley: Does it Exist?', in *Proceedings of Conference of Human Computer Interaction, Workshop on Human Animated Character Interaction*, (2005); Richard Coyne. 'The Digital Uncanny', in *Spaces, Spatiality and Technology*, eds P. Turner and E. Davenport, (Dordrecht: Springer, 2005), 5–18.
- 28 For accounts that seem at once foreboding and celebratory see Reza Negarestani, *Cydonopedia: Complicity with Anonymous Materials* (Melbourne: re.press, 2008).
- 29 An intellectual's identification with schizophrenia provides one example. See Gilles Deleuze and Felix Guattari, *Anti-Oedipus: Capitalism and Schizophrenia* (New York: Viking Press, 1977).
- 30 See Randy Shaw, *The Activist's Handbook: A Primer* (Berkeley, CA: University of California Press, 2001)
- 31 Stanley Cohen, *Folk Devils and Moral Panics* (St Albans: Paladin, 1973).
- 32 Pitirim A. Sorokin, *Man and Society in Calamity* (New York: E.P. Dutton and Co., 1942), 10.
- 33 Pitirim A. Sorokin, *Man and Society in Calamity* (New York: E.P. Dutton and Co., 1942), 177.

- 34 Pitirim A. Sorokin, *Man and Society in Calamity* (New York: E.P. Dutton and Co., 1942), 121.
- 35 Pablo Sendra and Richard Sennett, *Designing Disorder: Experiments and Disruptions in the City* (London: Verso, 2020).
- 36 Richard Coyne, *Cryptographic City: Decoding the Smart Metropolis* (Cambridge, MA: MIT Press, 2023).

References

- Aguirre, A., et al. (2015), 'Research Priorities for Robust and Beneficial Artificial Intelligence: An Open Letter'. Future of Life Institute, 28 October. Available online: <https://futureoflife.org/open-letter/ai-open-letter/> (accessed 4 April 2023).
- Aguirre, A., et al. (2023), 'Pause Giant AI Experiments: An Open Letter'. Future of Life Institute, 31 March. Available online: <https://futureoflife.org/open-letter/pause-giant-ai-experiments/> (accessed 4 April 2023).
- Atzori, L., A. Iera and G. Morabito. (2010), 'The Internet of Things: A Survey', *Computer Networks*, 54 (15): 2787–805.
- Benjamin, W. (1992), 'The Work of Art in the Age of Mechanical Reproduction', in H. Arendt (ed.), *Illuminations*: 1–58, London: Fontana.
- Bostrom, N. (2014), *Superintelligence: Paths, Dangers, Strategies*, Oxford: Oxford University Press.
- Brenton, H., M. Gillies, D. Ballin and D. Chatting. (2005), 'The Uncanny Valley: Does It Exist?', *Proceedings of Conference of Human Computer Interaction, Workshop on Human Animated Character Interaction*.
- Chomsky, N. (2009), 'Turing on the "Imitation Game"', in R. Epstein, G. Roberts and G. Beber (eds), *Parsing the Turing Test: Philosophical and Methodological Issues in the Quest for the Thinking Computer*: 103–6, Dordrecht: Springer Netherlands.
- Clark, A. (2012), 'Embodied, Embedded, and Extended Cognition', in K. Frankish and W. Ramsey (eds), *The Cambridge Handbook of Cognitive Science*: 276–91, Cambridge: Cambridge University Press.
- Cohen, P. R. (2005), 'If Not Turing's Test, Then What?', *AI Magazine*, 26 (4): 61.
- Cohen, S. (1973), *Folk Devils and Moral Panics*, St Albans: Paladin.
- Coyne, R. (2005), 'The Digital Uncanny', in P. Turner and E. Davenport (eds), *Spaces, Spatiality and Technology*: 5–18, Dordrecht: Springer.
- Coyne, R. (2023), *Cryptographic City: Decoding the Smart Metropolis*, Cambridge, MA: MIT Press.
- Deleuze, G. and F. Guattari. (1977), *Anti-Oedipus: Capitalism and Schizophrenia*, New York: Viking Press.
- Dreyfus, H. L. (1972), *What Computers Can't Do: The Limits of Artificial Intelligence*, New York: Harper and Row.
- Dreyfus, H. L. (1990), *Being-in-the-World: A Commentary on Heidegger's Being and Time Division I*, Cambridge, MA: MIT Press.
- Dreyfus, H. L. and S. E. Dreyfus. (2004), 'From Socrates to Expert Systems: The Limits and Dangers of Calculative Rationality'. Department of Philosophy, The University of California, Berkeley. Available online: http://socrates.berkeley.edu/~hdreyfus/html/paper_socrates.html (accessed 11 June 2017).

- Freud, S. (1990), 'The "Uncanny"', in A. Dickson (ed.), *The Penguin Freud Library, Volume 14: Art and Literature*: 335–76, Harmondsworth, Middlesex: Penguin.
- Heidegger, M. (1977), *The Question Concerning Technology and Other Essays*, trans. W. Lovitt, New York: Harper and Row.
- Herbrechter, S. (2013), *Posthumanism: A Critical Analysis*, London: Bloomsbury Academic.
- Hinton, G. and W. D. Heaven. (2023), 'Video: Geoffrey Hinton Talks About the "Existential Threat" of AI'. *MIT Technology Review*, 3 May. Available online: <https://www.technologyreview.com/2023/05/03/1072589/video-geoffrey-hinton-google-ai-risk-ethics/> (accessed 12 May 2023).
- Hoel, E. (2023), "'I Am Bing, and I Am Evil" Microsoft's New AI Really Does Herald a Global Threat'. *The Intrinsic Perspective*, 17 February. Available online: <https://erikhoel.substack.com/p/i-am-bing-and-i-am-evil> (accessed 10 May 2023).
- Hoel, E. and S. Waas. (2023), 'Erik Hoel on the Threat to Humanity from AI'. *Econlib Podcast*, 3 April. Available online: <https://www.econtalk.org/erik-hoel-on-the-threat-to-humanity-from-ai/> (accessed 12 May 2023).
- Jurafsky, D. and J. H. Martin. (2023), *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition, Third Edition*, Redwood City, CA: Stanford University Press.
- Kitchin, R. (2018), 'Data-Driven Urbanism', in R. Kitchin, T. P. Lauriault and G. McArdle (eds), *Data and the City*: 44–56, London: Routledge.
- Kurenkov, A. and J. Harris. (2023), 'AI and Existential Risk – Overview and Discussion'. *Last Week in AI – Skynet Today*, 30 August. Available online: <https://lastweekin.ai/p/ai-and-existential-risk-overview#details> (accessed 2 September 2023).
- Kurzweil, R. (2005), *The Singularity Is Near: When Humans Transcend Biology*, New York: Penguin.
- Leach, N. (2023), 'AI Is Putting Our Jobs as Architects Unquestionably at Risk'. *dezeen*, 13 February. Available online: <https://www.dezeen.com/2023/02/13/ai-architecture-jobs-risk-neil-leach-opinion/> (accessed 27 May 2023).
- Marcuse, H. (1991), *One-Dimensional Man: Studies in the Ideology of Advanced Industrial Society*, London: Routledge.
- Mattern, S. (2021), *A City is Not a Computer: Other Urban Intelligences*, Princeton, NJ: Princeton University Press.
- Metz, C. (2022), *Genius Makers: The Mavericks Who Brought A.I. to Google, Facebook, and the World*, London: Penguin.
- Mumford, L. (1961), *The City in History: Its Origins, its Transformations, and its Prospects*, Harmondsworth, Middlesex: Penguin.
- Negarestani, R. (2008), *Cyclonopedia: Complicity with Anonymous Materials*, Melbourne: re.press.
- Penrose, R. (1989), *The Emperor's New Mind: Concerning Computers, Minds, and the Laws of Physics*, London: Vintage.
- Sendra, P. and R. Sennett. (2020), *Designing Disorder: Experiments and Disruptions in the City*, London: Verso.
- Shaw, R. (2001), *The Activist's Handbook: A Primer*, Berkeley, CA: University of California Press.
- Sorokin, P. A. (1942), *Man and Society in Calamity*, New York: E.P. Dutton and Co.
- Turing, A. M. (1950), 'Computing Machinery and Intelligence', *Mind: A Quarterly Review of Psychology and Philosophy*, 59 (236): 433–60.

- Vitruvius, P. (1960), *Vitruvius: The Ten Books on Architecture*, trans. M. H. Morgan, New York: Dover Publications.
- Wolfram, S. (2023a), 'Computational Knowledge Superpowers to ChatGPT'. *Writings*, 9 January. Available online: <https://writings.stephenwolfram.com/2023/02/what-is-chatgpt-doing-and-why-does-it-work/> (accessed 28 February 2023).
- Wolfram, S. (2023b), 'Stephen Wolfram Answers Live Questions About ChatGPT'. *Wolfram*, 20 January. Available online: <https://www.youtube.com/watch?v=zLnhg9kir3Q> (accessed 27 February 2023).
- Wolfram, S. (2023c), 'What Is ChatGPT Doing and Why Does It Work?'. *Writings*, 14 February. Available online: <https://writings.stephenwolfram.com/2023/02/what-is-chatgpt-doing-and-why-does-it-work/> (accessed 15 February 2023).

13

EPILOGUE

ROMANCING THE CONVERSATIONAL CITY

From their inception, cities serve as centres of dialogue, exchange and transformation. A city is its physical structures but also evolving discourses, defined by language, signs and symbols. Conversational AI introduces new layers of interaction, perception and risk. The technologies of large language models (LLMs) have mirrored, magnified and disrupted urban conversations. The city, as we have explored, is an assembly of recollections, rhythms, attentions, engagements and calibrations. These facets are reflected and refracted in the functions of conversational AI's neural networks.

There are risks. AI reveals that cities are, and have always been, arenas of mediated conversation, of synthetic intelligence, of simulations that define, expand and challenge the human condition. The discourses on pansychism align with urban inhabitants and theorists who are inclined to attribute to cities the ability to converse as if cities think and feel.

Industry, education and everyday users increasingly deploy large language models, which are in turn embedded in everyday digital platforms. The technology is developing at a pace. The analysis of commentators,

DOI: 10.4324/9781003535751-14

This chapter has been made available under a CC-BY-NC-ND 4.0 license.

critics and legislators also gain traction as they evaluate the implications of the technology and seek to influence its future. Here I want to gather some possible scenarios for the future of LLMs with one eye on the city, the other on AI-informed preparation for urban professionals.

AI integration

I assume that LLMs will continue to be integrated into many applications, websites, devices and digital platforms. In the 1980s commentators predicted that standalone desktop systems such as word processors, graphics programs and CAD (computer-aided design) packages would be linked seamlessly via the Internet to servers at the supplier's side for updates, access to product libraries and to perform calculations beyond the desktop. We take that interconnection for granted now.

Can we expect something similar for LLMs? At the time of writing, Adobe's Photoshop incorporates a "generative fill" feature built on AI and that uses AI graphic techniques to enhance the resolution of images, as in some of the illustrations in this book. Microsoft's Copilot deploys OpenAI's ChatGPT to draft documents and respond to prompts. Wordpress for creating blog posts and websites incorporates automated image generation and offers to "improve your writing with AI." There are many reasons to expect that VLEs (virtual learning environments) such as Blackboard will continue to incorporate text-based AI features. I predict that we regular users of digital platforms will treat as commonplace AI integration much as we regard networked interconnectivity, location services on a smartphone and text completion on messaging apps.

With ubiquitous AI, conversations about LLMs shift to a different register. Instead of speculating about the impact of a standalone LLM service such as ChatGPT, the conversation moves to one informed by actual use. That involves discerning conversational AI's strengths, limitations and risks, as people inform themselves through practical use and their experience of its observable social, commercial and political impacts. In the educational context, habitual readers and reviewers of the work of students and professional trainees begin to recognise the hallmarks of automated text production, much as people have become more adept at recognising digitally manipulated photographs, audio files and videos. With such increased discernment come enhanced capacity to be critical of AI technologies and their applications.

Resources

Ubiquitous AI comes at a cost, that eventually lands on consumers. Payment models will affect the deployment of LLMs, manifested as new variants of advertising, pay walls, click bait, pop-ups and other paraphernalia of the commercialisation of the Internet infiltrate the AI space. Conversational AI invites automated endorsements delivered in conversational and personalised style within threads: “We are chatting about the architecture of Barcelona. Do you know that there’s an airline that will fly you there at a reduced rate about the time of your next holiday?” In an even more automated future, the AI platform might have “negotiated” a deal with the airline. More menacing is the enhanced use of AI in consumer profiling, the monetisation of our clicks and whatever else might be disclosed or inferred from our interactions with conversational AI.¹

Content creators are demanding payment and compensation for use and misuse of textual, graphic and audio content. LLMs are successful now in part due to the ready availability of diverse and high-quality content online able to be deployed as training data. LLM developers have used content from Wikipedia for training ChatGPT models. Questions remain as to whether training an AI on material held under Creative Commons licensing constitutes “fair use.” Similar issues arise in the case of blogs, social media posts, educational material, fan fiction and other content delivered for free by online authors. Published books and articles constitute a further challenge, as do books that publishers may distribute as open access on condition that other interests do not profit from using that content for AI training. An interesting podcast from the *New York Times* titled “The writers’ revolt against A.I. companies” highlights this concern amongst authors.² Design and other highly skilled interventions into cities also constitute “content” readily appropriated into AI training models. We have focussed on text and conversation in this book, but similar challenges apply to the use and misuse of visual, aural, tactile and other data able to be monetised.

Reskilling

Perhaps AI will write better than some human authors, and with access to a greater “knowledge base.” Such automated literary proficiency has implications for the education and training of urban professionals. Trainees can now recruit specialist AI tutors to enhance their learning and provide

feedback on project work. Are traditional teacher-centred educational models adequate for AI-based learning?³ Then there is the prospect of de-skilling and re-skilling, adopting new modes of writing, reading and even thinking. An AI can summarise a long text, and transpose something difficult into everyday language. AI might even teach us to write better, or with greater brevity. Effective AI prompting is considered a skill by some that calls for training,⁴ and AI can help with “prompt engineering” skills too. Clear and articulate explanatory prose will always be in demand, but AI is likely to have an impact, much as newspaper headlines, advertising slogans, emails, microblogging and texting influenced writing and reading in the past.

New practices emerge, and authors develop their own ways of working. I now am less inclined to ask questions of an LLM. I get more out of it if I tell it things and see what responses I get. That is treating the AI as a co-writer or critic rather than a flawed information source. For this book I occasionally called on ChatGPT to anticipate reader objections to the content and delivery of some chapters. I have indicated in the text where that is the case.

These to me are reasonable uses of conversational AI in academic writing. It can offer review and critique quicker and more comprehensively than one might expect from a colleague or human reviewer. It is then up to the human author to amend or otherwise respond to the critique in their own words. You can of course choose your critic. I have told the LLM platform to assume different roles in proffering its review, at times that of an analytic philosopher, a radical poststructuralist versed in the theories of Derrida and Foucault, a neo-Marxist reader, someone versed in Phenomenology, or a Pragmatist more interested in practical applications than theory. Choosing the right critic is normal in academic production, someone that will deliver targeted advice. Another way of saying this: it is about matching content to audience. That also raises the prospect that the AI-as-critic is used conversationally rather than delivering a report card on your work. Of course, AI as critic warrants some scepticism. Knowing how an LLM operates and some theories behind it helps with productive use of AI tools, and as any academic knows, you are entitled to critique your critic.

That level of LLM performance indicates the extent to which the corpus, the source text and the language it deploys, embodies the social, cultural,

and cognitive complexities of human language development and of human discourse. Though the developers of LLM methods have created algorithms that effectively operationalise very sophisticated pattern detection and processing, the prime cognitive claim of these platforms resides in the authorship of the source texts on which the models are trained rather than the algorithms. Add in the actions of those who “fine tune” the models. These observations are a further endorsement, if we need it, of the power of human culture and of the instrument of human language.

Security

What I have described so far assumes good will on the part of developers, commercial interests, legislators and users. To redress AIs that “scrape” their training data from the Internet, some fan fiction writers sought to confuse AI data-collection by publishing irreverent stories online to corrupt the training corpus, according to a Times article.⁵ Whether or not they achieved that goal, there are ways that people might influence the candidate pool of a training corpus. Outside the sphere of digital activism, as with any medium of communication, actors will infiltrate and corrupt text-based AI.

The use of AI in espionage, covert operations and subversion has been well-aired in the press and elsewhere.⁶ Other misuses of AI and AI methodologies include overtly or secretly exposing users to AI trained on data or fine-tuned towards political agendas or product biases. Conversational AIs with exemplary patience and language skills might also displace humans in customer service centres to appease users and customers who call in to complain. AI serves as an affable “Big Brother” to mollify a sceptical citizenry.

Certified human

Conversation is undoubtedly a key component of design, planning, urban intervention and urban citizenry, in which AI has the potential to support a “dialogic, participatory and discursive form of planning.”⁷ My initial objective in this book was to show that AI processes are already endemic to the functioning of cities. Like LLMs, cities are structured through patterns, relationships and contextual influences.

The functions that shape the competence of Transformer LLMs mirror the dynamics of urban life, experiences and processes. Just as social and physical patterns are intrinsic to cities, the role of a text corpus in LLMs can be likened to urban archives. Cities serve as living repositories, rich with exemplars, symbols, documents and patterns that residents internalise throughout their daily experiences. Urban participants dissect the city into components and organise their experiences – processes mirrored by tokenization in LLMs. Cities are interlaced with networks of relationships that LLMs imitate in the processing of semantic embeddings. Cities serve as complex contexts where meanings are continuously negotiated. The spatial arrangement of urban elements – both physical and within communication networks – signals varied layers of importance. Urban features, from iconic landmarks to street art, vie for attention, influenced by their positioning and relationships. These are processes mirrored by attentional embeddings in LLMs. Our interaction with urban environments invites metaphors of tuning in and out, embodying the way people adapt to both social dynamics and their physical surroundings. LLMs reflect this process of adaptation and adjustment in their tuning strategies that generate professionally competent and imaginative conversational exchanges.

How might an appreciation of these functions affect architects, planners and designers and impact them in the future? In developing new skills in AI literacy, data interpretation and collaborative design, architects and planners can exploit these technologies while retaining their roles as creative thinkers and active agents responsible for the built environment. There is a case to be made that as these technologies become more integrated into design and decision-making workflows, they will automate routine tasks, enhance creative exploration and facilitate inclusive and participatory design processes. Knowing how AI technologies work enables architects and planners to be critical of these claims and discerning in their use of AI tools.

The attributes and limitations of these technologies also provoke and reveal the nature of design in the built environment. For example, as AI lacks the full scope of characteristics defining human agency, urban professionals may see anew the importance of human responsibility and accountability in the way that cities are formed. The prospect of AI-based autonomous machinic agency has renewed our awareness of just how vital the human element is in urban design and planning.

As with the early introduction of CAD we might expect the rise of conversational AI to prompt a re-evaluation of the core competencies and values that define the professions of architecture and planning. This shift demands a focus on developing new skills in AI literacy, data interpretation and collaborative design, encouraging architects and planners to harness these technologies while retaining their roles as leaders of the technology and responsible stewards of the built environment.

Thinking of the post-digital trend of recent decades, there is also the possibility of a reversion to non-AI generated content: text with less than perfect syntax, quirky and idiosyncratic production that could never have been produced by an AI, a return to the labour-intensive craft of writing. On the design side we may see a return in some quarters to craft and human-centred urban processes less than perfect, though somehow authentically human. Could this include certification that the design content had no AI input? A proud assertion or credit could be carried as a badge of honour: “All creation is my own. No AI was deployed in the generation of this content.”

Notes

- 1 Shoshana Zuboff, *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power* (London: Profile Books, 2019).
- 2 Sheera Frenkel and Michael Barbaro, ‘The Writers’ Revolt Against A.I. Companies’. *New York Times: The Daily*, 18 July, 2023. Available online: <https://podcasts.apple.com/gb/podcast/the-daily/id1200361736?i=1000621495806> (accessed 19 July 2023).
- 3 “No-win-no-fee” lawyers have approached students to claim compensation from universities for requiring them to study online instead of in person during the COVID pandemic. Similar discontent could arise as students are expected to engage one-on-one with an AI for assistance and instruction in lieu of meeting a tutor. After all, students may decide that one-to-one meetings with an AI can be just as good if not better than meeting with a personal tutor. The prospect of using conversational AI to generate or assist in writing student feedback and even assess academic output has occurred to many of us, as has the extent to which an AI can inspire, write, edit, correct and improve an essay before it is submitted for marking.
- 4 Jerry Sinclair, *The Art of Mastering Prompt Engineering for Chat-GPT Made Simple: A Practical Guide to Learn How to Get the Most Out of Chat-GPT with Efficient and Effective Prompting* (Kindle, 2023).
- 5 Sheera Frenkel and Stuart A. Thompson, “‘Not for Machines to Harvest’: Data Revolts Break Out Against A.I.’, *New York Times*, 15 July, 2023. Available online: <https://www.nytimes.com/2023/07/15/technology/artificial-intelligence-models-chat-data.html> (accessed 19 July 2023).

- 6 Renee DiResta et al., *The Tactics and Tropes of the Internet Research Agency* (Washington, DC: Senate Select Committee on Intelligence, 2018). Philip N. Howard, et al., *The IRA, Social Media and Political Polarization in the United States, 2012–2018* (Oxford: Computational Propaganda Research Project, University of Oxford, 2018).
- 7 Lieven Ameel, 'Towards a Narrative Typology of Urban Planning Narratives for, in and of Planning in Jätkäsaari, Helsinki', *Urban Design International*, 22, no. 4 (2017): 318–30, 319.

References

- Ameel, L. (2017), 'Towards a Narrative Typology of Urban Planning Narratives for, in and of Planning in Jätkäsaari, Helsinki', *Urban Design International*, 22 (4): 318–30.
- DiResta, R., et al. (2018), *The Tactics and Tropes of the Internet Research Agency*, Washington, DC: Senate Select Committee on Intelligence.
- Frenkel, S. and M. Barbaro. (2023), 'The Writers' Revolt Against A.I. Companies', *New York Times: The Daily*, 18 July. Available online: <https://podcasts.apple.com/gb/podcast/the-daily/id1200361736?i=1000621495806> (accessed 19 July 2023).
- Frenkel, S. and S. A. Thompson. (2023), "'Not for Machines to Harvest': Data Revolts Break out against A.I.'. *New York Times*, 15 July. Available online: <https://www.nytimes.com/2023/07/15/technology/artificial-intelligence-models-chat-data.html> (accessed 19 July 2023).
- Howard, P. N., et al. (2018), *The IRA, Social Media and Political Polarization in the United States, 2012–2018*, Oxford: Computational Propaganda Research Project, University of Oxford.
- Sinclair, J. (2023), *The Art of Mastering Prompt Engineering for ChatGPT Made Simple: A Practical Guide to Learn How to Get the Most out of ChatGPT with Efficient and Effective Prompting*, Kindle.
- Zuboff, S. (2019), *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*, London: Profile Books.

GLOSSARY

Agent An autonomous entity capable of taking data from its environment and carrying out actions to achieve specific goals. The term is often used in AI to describe both software programs and robots. In the case of human agents, we may speak of what individuals and groups perceive, how they interpret, the actions they negotiate and how they assume responsibility for actions.

AI Model A computational method designed to perform specific tasks based on patterns derived from training data, often involving machine learning algorithms and techniques. People sometimes talk about a trained neural network as an AI model to which they attribute outcomes, e.g. “the AI model produced a summary of the report.”

AI Artificial Intelligence is a set of machine processes that produces outcomes that resemble or exceed what we think of as the products of human thought processes. It includes the capability to train on data, make inferences, self-correct and use language.

Artificial General Intelligence (AGI) A form of AI that can train on data, make inferences, and correct itself across a wide range of tasks, thereby simulating the human capability to perform cognitive tasks not limited to specific problem domains. An effective AGI would switch between tasks: translate between languages, solve a maths problem, control a car, and communicate empathy, without needing task-specific training for each.

CCC (Conversation-centric computation) A human-oriented approach that brings conversation to the fore in computer programming as both a method and a product. The approach is strengthened by advances in conversational AI.

Chatbot An interactive computer program that receives inputs and delivers outputs in conversational style, typically in text format and in real time. Early chatbots were often rule-based, responding to specific keywords or patterns in a predefined manner, limiting their capabilities. However, the technology has developed significantly, with modern chatbots increasingly incorporating advanced techniques from large language models (LLMs) like GPT. The term “chatbot” evokes a sense of earlier, less advanced systems, in contrast to the more powerful and flexible conversational AI of today.

Context Window In natural language processing, this is the span of text around a target word or phrase that constrains the sequence of text processed in calculating semantic, positional and attentional encoding, and hence prediction of the text that follows. A larger context window allows the AI model to capture more subtle relationships and dependencies across longer text sequences, leading to more accurate and coherent predictions and outputs. In large language models (LLMs), context windows can range from a few hundred tokens to several thousand tokens, with some advanced models supporting windows of up to 8,000 tokens or more.

Dall-e A generative graphics product accessed via ChatGPT and other LLMs. Dall-e uses a transformer model that first analyses the prompt requesting an image by breaking it into tokens. Then, it generates an image by predicting, step by step, the pixels or patches of the image that correspond to the text, using patterns from training on large datasets of text-image pairs. It maps the textual description to visual elements, gradually assembling them into a coherent image.

Diffusion Model A type of machine learning model that generates data, such as images or text, by simulating a process that gradually adds and then removes noise to transform a simple pattern into a more complicated one. It is used in tasks such as image generation and reducing noise (e.g. speckles and smudges) in images.

Embedding A numerical representation of words or phrases in a continuous vector space, capturing semantic and other relationships for processing

by AI models. The vector representation is an ordered list of numbers. The number of entries in the list is the size or dimension of the embedding, typically ranging from a few dozen to several hundred dimensions, with larger embeddings capturing more subtle nuances in word meanings and relationships but also requiring more computational resources.

Encoding The process of converting input data, such as tokens, words or phrases, into a structured format like embeddings, which allows AI models to process and analyse the data. Encoding transforms raw text into numerical vectors that capture syntactic, semantic or positional information, depending on the AI model's requirements.

Existential Risk The potential threat to human existence or long-term survival, often discussed in the context of advanced AI systems capable of actions that could be catastrophic for humanity.

Extinction The complete disappearance of a species or group, often used metaphorically to describe the potential irreversible negative impact of AI on human society.

Feature A feature in AI is a characteristic of the data being used by a model. For example, on a map, a feature could be something like the location of a park, the type of building or the presence of a road. These map features serve as inputs that the model uses to make predictions or decisions, such as determining the best route or identifying patterns in urban development. See **Hidden Feature**.

Feed forward In the context of neural networks, “feed forward” refers to the process where data moves in one direction through the network – from the input layer, through any hidden layers, and finally to the output layer. Each layer processes the data and passes it on to the next layer without looping back. The implementations of neural networks discussed in this book are of this type, except for the mention of recurrent neural networks (RNNs). In RNNs, data can flow in cycles within the network.

Fine-tuning (Human-in-the-Loop) This occurs after pretraining and involves further training the model on a smaller, task-specific dataset, often with human oversight, to refine and optimise its performance for specific applications. For example, in developing a customer service chatbot, human reviewers might provide feedback on the model's responses, helping to adjust and improve the chatbot's accuracy and appropriateness in handling customer queries.

Fine-tuning (Task Specific Adjustment) This involves training a pre-trained neural network on specific datasets designed to guide the model toward certain tasks or behaviours. For ChatGPT and similar conversational AI applications, this involves fine-tuning the neural network on conversational data to improve the quality and relevance of its responses in a dialogue format, making it better at generating human-like conversational responses to user prompts.

Generative Adversarial Network (GAN) A machine learning method involving two neural networks, a generator and a discriminator, that compete against each other to improve the quality of generated data, often used in generating images or text. Conversational AI generally does not use (GANs) as the primary method for generating responses, favouring GPT (Generative Pre-trained Transformer) models instead.

GPT (Generative Pre-trained Transformer) GPT is a type of AI model designed to process and generate human-like text. It is built on a neural network architecture called a “transformer,” which processes sequences of data, such as sentences and paragraphs. The “pre-trained” part means that GPT is initially trained on a large amount of text through which it incorporates patterns in grammar, facts and the structure of language. After this broad training, it can be tuned with further training examples to perform specific tasks, such as answering questions, writing reports or holding a conversation. The training data must be pre-processed, i.e. filtered, sorted and prepared, for use in these stages of the GPT training and for specialised fine tuning.

GPU (Graphics Processing Unit) Originally designed for rendering graphical images and accelerating visual display tasks, GPUs are now widely used for computationally intensive applications, including secure cryptography, blockchain validation and artificial intelligence (AI). GPUs are particularly efficient at performing parallel processing, which allows them to handle multiple operations simultaneously, making them well-suited for tasks that require high-speed computation.

Hidden Feature A hidden feature in AI is a characteristic or pattern within the data that is not explicitly defined or directly observable but is produced by the model during training. For example, in urban map data, a hidden feature might be the underlying social or economic significance of certain neighbourhoods, which is not directly represented by visible map elements like roads or buildings but is inferred

by the model based on patterns in the data. These features are called “hidden” because they are not manually identified or labelled; instead, they emerge as the model processes the data and optimises its parameters, allowing it to capture complex relationships that might influence urban development, such as proximity to key amenities or transportation hubs.

Hidden Layers A hidden layer in a neural network is a layer of artificial neurons that sits between the input layer (where data enters) and the output layer (where the result comes out). The input size ranges typically from a few hundred to several thousand tokens, each represented as a vector with hundreds of dimensions. The output is similarly a sequence of tokens, corresponding to the model’s vocabulary size, which can range from tens of thousands to hundreds of thousands of possible tokens. LLMs often have dozens to hundreds of layers; for example, models like GPT-3 have 96 layers, each containing thousands of hidden units – specifically, 12,288 units per layer in GPT-3. These hidden layers are responsible for capturing intricate patterns in the data and identifying hidden features in the data. Overall, LLMs are characterised by their hundreds of layers and thousands of units per layer, allowing them to accurately handle large-scale text processing tasks.

Large Language Model (LLM) A type of AI computer program that is trained on vast amounts of text data to process and generate human language, capable of performing a wide range of language tasks. These models are typically trained on datasets containing hundreds of gigabytes to terabytes of text (i.e. up to one million books) sourced from diverse materials such as books, websites, articles and social media. The scale of the training data enables LLMs to train on a wide array of linguistic patterns and contexts.

LLM Short for Large Language Model, a specific type of AI model used for processing and generating human language.

Machine Learning (ML) Some people reserve the word “learning” for the processes by which living beings develop knowledge, understanding and skills in complex environments. But the term “machine learning” has currency as a subset of AI that involves training algorithms to make predictions or decisions based on data, enabling systems to improve performance over time without explicit programming. Like “artificial intelligence,” the term “machine learning” serves to define a

field of research and development, irrespective of the extent to which we think the anthropomorphic term “learning” applies.

Matrix A mathematical structure composed of rows and columns, where each element represents a value. In AI, matrices are used to represent data, weights and other parameters for computations. The size of a matrix is defined by the number of its rows and columns (referred to as its dimensions), and it can range from small (e.g., 2x2) to large (e.g., thousands of rows and columns) depending on the complexity of the task.

Memory Some people prefer to reserve the word “memory” for the processes by which living beings recall experiences. In AI, “memory” is the ability of a system to retain and utilise information from past interactions or data processing to improve performance on tasks.

Model A structured representation of a process or phenomenon, often implemented as algorithms and data structures, used in AI to simulate tasks and make predictions. In the context of LLM development, the term “model” can refer to different aspects: it may describe the entire trained system capable of processing and generating language; it can also refer to the underlying architecture (such as a neural network) that defines how data is processed; and sometimes, the term is used to describe the specific configurations or versions of the system that are being trained or fine-tuned to improve performance on certain tasks. Thus, “training the model” can mean refining the algorithms, adjusting parameters or enhancing specific capabilities within the larger system.

Multi-agent System Multiple autonomous agents interact or work together to achieve individual or shared goals. These agents can be software entities, avatars, robots or humans, and they operate in a common environment that coordinates their actions, shares information and manages competition between them.

Natural Language Processing (NLP) A branch of AI focused on enabling computers to process and generate human language, bridging communication between humans and machines. NLP refers to the overall field and techniques for processing language. It encompasses LLMs and other approaches, such as rule-based systems, traditional machine learning models and statistical methods, which are used for tasks such as parsing, entity recognition, machine translation and more.

Neural Network A computational model inspired by the animal brain, composed of interconnected nodes or neurons, used in AI to recognise patterns and perform complex tasks.

NLP Short for Natural Language Processing, the area of AI concerned with interactions between computers and humans through natural language.

Parallel Distributed Processing (PDP) In the context of neural networks, Parallel Distributed Processing (PDP) refers to a computational approach where information processing is carried out simultaneously across multiple interconnected units or neurons, rather than sequentially. In a neural network, each neuron processes information in parallel with others, and the “knowledge” is distributed across the network through the connections between neurons. This allows the network to handle complex tasks by integrating and processing large amounts of data simultaneously. PDP emerged in the 1980s as part of a broader shift in cognitive science and artificial intelligence toward more biologically inspired models of computation. The term is used less frequently now because the concepts it introduced have been largely absorbed into the broader field of neural networks and deep learning, which have become the dominant frameworks in artificial intelligence.

Platform This is a broad term that describes the environment or framework that enables the operation of applications, systems, models and programs. It differs from these terms in its role as the foundational layer that supports and integrates various software components and services. Operating systems, search engines and social media services such as Facebook, ChatGPT and other AI applications and services are built on platforms (such as the OpenAI infrastructure in the case of ChatGPT).

Polis An ancient Greek term for a city-state, often used to describe a community of citizens, and metaphorically applied to the concept of cities in urban studies.

Positional Embedding A method used in neural networks to encode the position of words in a sequence, helping models compute the order of tokens in blocks of text, and hence the structure of language, such as grammar, and less formal idiomatic structures and turns of phrase.

Pretraining The process of initially training an AI model on a large, diverse dataset to adopt general patterns, structures and information

about language or other domains. During pretraining, the model is exposed to vast amounts of data to build a foundational predictive capability, which includes adjusting its neural network parameters, calculating semantic embeddings, and incorporating techniques such as positional encoding and attention mechanisms. This phase does not involve task-specific adjustments but rather prepares the model for more specialised tasks that are addressed later in training during the fine-tuning phase.

Prompt In the context of AI and language models, a prompt is the text or input provided by a user to guide or initiate a response from the model. The prompt can be a question, a statement or any other form of input that instructs the AI on what kind of response is desired. The quality and clarity of the prompt can influence the effectiveness and relevance of the AI's output. Effective prompting implies a set of linguistic and technical practices often termed "prompt engineering."

Python A high-level programming language widely used in AI and data science, supporting numerous libraries for numerical processing, including machine learning and data analysis.

Semantic Embedding A representation of words or phrases in a vector space that captures their relationships based on their usage in language. An embedding is therefore an ordered list of numbers that can be combined mathematically with other embeddings and matrices to provide input to a neural network. The number of values in a vector is its dimension. The size of an embedding typically ranges from a few dozen to several hundred dimensions.

Simulation Phase This is the term used to characterise the phase where an LLM neural network is put to use to generate conversational outputs in response to user prompts.

Speech-to-Text A technology that converts spoken language into written text, often used in voice recognition and transcription applications.

TensorFlow A library of program modules developed by Google for a Python programmer to build, train and deploy machine learning models. TensorFlow provides a platform for creating neural networks and other complex algorithms. Keras operates as a high-level interface to TensorFlow's library modules. PyTorch is another widely used open-source machine learning library developed by Facebook's AI Research lab.

- Text-to-Speech** A technology that converts written text into spoken language, allowing computers to read text aloud.
- Token** Blocks of text are fragmented into smaller units that may be words, groups of words or parts of words. These fragments are called “tokens,” and they also include numerals, symbols and punctuation. LLMs derive lexicons of tokens optimised by statistical and AI-based methods. During training and at the simulation phase, each token is mapped onto a unique semantic embedding, which is adjusted to provide input to a neural network.
- Training** The process by which a neural network adjusts its parameters to perform a specific task, such as the ability to predict word sequences, and thereby generate text. Neural network training is often anthropomorphised as “learning,” as in “machine learning.”
- Transformer** A transformer is a type of AI model used to process and generate human-like text. It works by processing the relationships between different parts of a sentence all at once, rather than one token at a time, which allows it to take account of context and meaning more effectively than earlier LLMs. Transformers are the foundation for many advanced language models like GPT.
- Urban AI** This is the application of AI technologies to urban environments, applicable to city planning, design, management and services.
- Urbanism** The study and design of cities, focusing on the interaction between people, infrastructure and the built environment, generally considering historical, cultural and technological factors.
- Vector** A mathematical representation of data in a multidimensional space, used in AI to encode tokens, features, words and other elements for processing by models. See entry for **Embedding**.
- X-risk** Short for existential risk, referring to the potential catastrophic threats that could lead to human extinction or irreversible damage to civilization, including those posed by AI.

APPENDIX A

PROGRAMMING AN LLM

To support the investigations of this book I have been implementing small scale trials of automated natural language processing routines that deploy the same methods as ChatGPT, i.e., implementations of the so-called Transformer architecture. That requires training on sequences of tokens in an original source document, equating each token to a semantic embedding (i.e., a long vector of 30 or so floating point numbers sourced from a publicly available dictionary of words and tokens), positional encodings which are vectors derived from the positions of words in the text, and finally an attention mechanism that factors in a calculation about the importance of individual words in sentences.

These factors combine to provide input to a neural network model. I had to write the code in Python to prepare the data for the model, but the neural network is implemented and managed as a library routine (Tensorflow. Keras) accessed via high-level parameters called from my Python program. The attention mechanism is folded into the operations of the neural network. Attentional embedding is a complicated process and involves layers in the network. So, it is accessed from Python via high-level specifications. I did not have to design the network or the attention mechanism.

My coding expert was ChatGPT4. As I have said in previous chapters it has leveraged my own programming skills to a more professional level. Manuals, websites and Youtube videos would be insufficient to achieve this. ChatGPT4 acts as an expert programmer, tutor and troubleshooter.

In this appendix I describe my attempt to produce an LLM that generates some coherent (grammatical) text consistent with the training corpus and will continue the thread of text (a prompt) supplied by a user. I have learned that to replicate the performance of a fully-fledged LLM is beyond my desktop capability and expertise. The performance of the method I outline here is close to that of an app for predictive text.

Though the method falls short of full-fledged natural language generation, it presages aspects of sophisticated large language models, including the use of lexical information. The method deploys techniques on the way to understanding conversational AI. As you will see from the output, the exercise also highlights the importance of scale, as well as other processes to do with sequencing and attention discussed in Chapters 8 and 9.

To test the idea of a neural network for generating sentences, I developed a program to process some of my own writing: about 19,000 words of early draft material for this book. The program extracts a list of words from the text, converts them to lowercase, and strips out inconsequential symbols, characters, and punctuation, apart from full stops (periods). The unique words are then ordered alphabetically and saved to a file as a lexicon.

I devised a program to inspect the original text file to identify the occurrence of each word from the lexicon and the word that follows each in sequence. This process produces a series of about 4,000 proximity pairs, which are also stored in a file. The inclusion of full stops in the proximity pairs aids in the construction of sentences during the generative phase. The file of proximity pairs follows this format: participate–in, particular–aspects, particular–corpus, particular–method, particular–region and so on for the entire training set of proximity pairs derived from the 19,000 word source document.

The main part of the program specifies and trains a neural network using the TensorFlow library. The network model is a simple feedforward network with many input nodes, one for each word in the lexicon (about 3,700) and an equivalent number of output nodes. The network features a single hidden layer with 100 nodes. So, the network model for this sentence generating method is considerably larger than the model we created for constructing the map journeys in Chapter 7.

The proximity pairs serve as training data for the network model. During training, the network adjusts its parameters to predict the next word based on the current word. The program runs through all 4,000 training pairs 500 times (epochs) optimising its parameters. The trained model is saved as a 9.1 Mb file, which includes 745,810 training parameters.

Once trained, the simulation process starts with a random word from the lexicon and uses the trained neural network to predict the subsequent word. A “temperature” setting biases the next word selection according to the output probability distributions from the model. This process is repeated until the desired number of sentences, determined by the number of full stops, is reached. The output is then cleaned up, removing unnecessary blank spaces and capitalising the start of sentences, before being displayed. Here’s some of the output:

Unveils ancient daydream charles regulatory demands interactive
aggregate online certainly book those babble. My theme of similar
layout. Instances an half video.

It would be easy enough to fill a whole book with this babble. It is difficult to validate the production of nonsense sentences so derived from coherent training data.

How are such sentences generated by a trained neural network better than random sentence generation? I benchmarked such outputs against randomly generated sentences produced from a language model that does not have the benefits of the pattern processing capabilities of a neural network.

Benchmarking against random sentences

It is a simple matter to create a program to calculate the number of times each word appears in the source text, i.e., its frequency, and from that to calculate the probability that any word will show up when selected at random. We calculated the number of times each token in the lexicon appears in the original 19,000-word text file. A program converts those frequencies to probabilities. For example, the word “technologies” occurs 13 times, representing a probability that it will appear in a sentence of 0.0007. The word “telegram” occurs once. Therefore, it has a probability of showing up in a sentence of 0.0001.

We thereby created a program that generates random word sentences from the same lexicon we created for the vocabulary but calculated based on word frequency. Here is a sample of the output:

Exploration until the there sorted automated. Figure heading accord-
ing blog of. Observe in browser the nodes.

Again, the program will keep generating such random text with apparent ease, as quickly as I can hit the “run” button on the Python editor. This random model based on word frequency would require substantial additional features to bring random word generation close to any claim that it had captured the style, content or meanings of the original training corpus. A method that makes use of Markov chains would improve its performance.¹ Unlike the operations of a neural network, there is nothing in the operations or the theory of random text generation that suggests the detection and preservation of meaningful patterns in data.

To evaluate the similarities and differences between text generated by a neural network and random generation I called on the resources of ChatGPT. I presented the conversational AI platform with a blind test: two blocks of text (A and B). Block A was generated randomly by the word-frequency random method. Block B was generated by the neural network model trained on proximity pairs. The AI guessed correctly. It offered an analysis claiming its assessment was based on the coherence and structure of the sentences, as well as the topical relevance of phrases. It determined that “Block A appears to have more disconnected sentences and phrases. While individual phrases could be meaningful, they don’t form a coherent narrative or context when pieced together.” It provided example phrases that illustrate this. It suggested on the other hand that Block B “proceeds to form longer, more coherent phrases ... The sentences, though not perfect, have a better flow and more discernible themes.” That ChatGPT could compare and assess the outputs in this way is as much a tribute to the power of the Transformer model that underpins ChatGPT as it is to the strength of either of my language models – which ChatGPT helped program!

Lessons about generative AI

It is possible to experiment with the hyperparameters of the LLM during the generative phase. A higher temperature makes the model’s decisions

more random, leading to more diverse but potentially less coherent sentences. A lower temperature makes the model more deterministic, often fixating to more common patterns it encountered during training. Even sophisticated large language models such as ChatGPT provide researchers and developers with the option of varying the temperature when in generative mode, thus reminding us that they are tools of automation after all.

My simple LLM treats each word (or token) as a unique input to the model, one input node for each element in the word lexicon for that corpus, 3,700 in this case. As discussed in Chapter 4, the Transformer methodology uses a neural network to derive word associations based on the proximity of words in the corpus. That serves as a kind of pre-training that captures how words cluster together in complicated and ambiguous ways. This pre-training represents each word as long vectors of numbers. It is these numbers that are fed into each input-output training event. These vectored word relationships produce a semantic complexity that enables the model to substitute one word for another appropriately, and apparently to resolve ambiguities in the context of other words.

A further observation is that our proximity pairs language model operates with discrete words. As explained in Chapter 5, sophisticated Transformer language models work with tokens made up of fragments of words as well as full words. They also include mathematical symbols and fragments of computer code. Part of the pre-training of a Transformer-based large language model is to divide the corpus into such small-scale units. The tokens and their semantic vectors form the model's lexicon. The lexicon becomes much more powerful than if it included only words. That said, as I have already stated, it also becomes less legible to human scrutiny. The large numerical vectors that describe every token in a lexicon contribute to the power, but also the opacity, of natural language neural networks, the processes of which only become legible at the moment they generate human-readable outputs. If it is not already obvious from the explanations in Chapter 7, the operations of AI language models become ever more opaque as developers introduce yet further enhancements to the algorithms. This further feeds into scepticism about AI systems as I outlined in Chapter 12.

Beyond any X-risk scenario, or technical and linguistic insights, there is entertainment value in seeing what a model projects as a logical continuation of a given word or phrase. We are now used to predictive text on our

devices, but my simple LLM has similar recreational value to the Exquisite Corpse game invented by the Surrealists in which participants would write a section of sentence on a sheet of paper, fold it to conceal their contribution, and then pass it to the next player to continue the writing, often resulting in unexpected, amusing and disturbing combinations.

In summary, the sentences generated by our LLM program offer insights into the patterns and relationships the model has identified in its training data. They demonstrate the strengths and limitations of neural networks in natural language processing tasks. Importantly, as we consider urban AI, they signal how techniques in language processing mirror processes that apply to mapped data as illustrated in Chapter 7, strengthening further the ties between conversational artificial intelligence and cities.

Note

- 1 Early natural language processing systems would use such a method. Markov models in natural language processing (NLP) are used to predict the likelihood of a word appearing in a sentence based on the words that come before it. They rely on the idea that the probability of a word depends only on a limited number of previous words, rather than the entire sentence, simplifying language prediction tasks. See Luis Serrano (2018), 'A friendly Introduction to Bayes Theorem and Hidden Markov Models'. YouTube, 27 March. Available online: <https://www.youtube.com/watch?v=kqSzLo9fenk> (accessed 27 September 2020) and Becky P. Y. Loo et al., 'Applying the Hidden Markov Model to Analyze Urban Mobility Patterns: An Interdisciplinary Approach', *Chinese Geographical Science*, 31, no. 1 (2021): 1–13.

APPENDIX B

POSITIONAL ENCODING FORMULAS

As diagrammed in Chapter 8, the formulas for the calculation of the sine and cosine values are

$$PE_{(pos,2i)} = \sin\left(pos \times \frac{1}{10000^{2i/d_{model}}}\right)$$

$$PE_{(pos,2i+1)} = \cos\left(pos \times \frac{1}{10000^{2i/d_{model}}}\right)$$

I have been following the positional encoding method used by Vaswani and colleagues in the seminal “Attention is All You Need” paper that introduced the Transformer model of LLMs, and these are the formulas the authors used.

The variable *pos* is the position of a token along the context window (0–9) in my example. The variable *PE* is the positional encoding value. The variable *d* is the depth (or dimensionality) of the encodings, which

is essentially the number of dimensions in the output vectors from the model. Here are the terms described.

pos: The position of a token in the sequence. It ranges from 0 to the length of the sequence minus one.

d_{model} : The depth (or dimensionality) of the encodings, which is essentially the number of dimensions in the output vectors from the model.

i: A variable that ranges from $\frac{d_{model}}{2} - 1$. It only goes up to half of d_{model} because the positional encoding uses a combination of sine and cosine functions for alternating dimensions of the encoding.

ChatGPT helps explain that the purpose of the term.

$$\frac{1}{10000^{2i/d_{model}}}$$

It is to create a distinct encoding for each position in the context window. As *i* increases (i.e., for higher dimensions in the encoding), the periodicity of the sine and cosine functions increases. This results in a unique combination of sine and cosine values for each position across the encoding dimensions, allowing the model to differentiate between positions in a sequence. The use of a large base like 10,000 ensures that even for long sequences, the positional encodings remain distinct. I used a smaller base term of 10 in my plots as it makes them clearer graphically.

I plotted the curves (Figures 8.4 and 8.5) and heat maps (Figures 8.7 and 8.8) via programming in Python, with help from ChatGPT. As someone interested in form and space I find the visualisation of the resulting interference patterns intriguing. Vaswani et al do not explain the theory behind how the formulae work to instil sequential information into LLM training data sufficient for a neural network to capture and reproduce sequential patterns. But they claim the method seems to work empirically.

APPENDIX C

AN ALTERNATIVE GLOSSARY OF AI

To inspire some writing, I prompted ChatGPT to translate a few of my glossary entries in different literary styles. Here is a selection.

Agent: A being of action and reaction, absorbing the chaos of its environment, processing the ceaseless flow of data, and birthing decisions that shape the world. In the mechanised heart of AI, this term signifies the cold, calculating forces – software, robots – enacting predetermined objectives. But in the human realm, “agent” evokes a far more tumultuous entity: the individual or collective, haunted by perception, interpretation, negotiation, and the weighty burden of responsibility.

In the style of the radical philosopher Georges Bataille.

Embedding: Picture a curious map, not of lands and seas, but of words and ideas! Each word, phrase or thought is a little dot in a vast and colourful space, connected by invisible threads of meaning. The Embedding is the cartographer of this strange land, drawing the paths and borders that show how close or far our words are from each other. Some dots are clustered together, whispering secrets to one another, while others stand

alone, guarding their own mysteries. It's a map not just of where words are, but of what they mean and how they dance together in the grand waltz of language.

In the style of Lewis Carol's *Through the Looking Glass*.

Fine-tuning (Human-in-the-Loop): A delicate and parasitic operation, wherein the pre-trained network – already saturated with the toxic residues of pre-existing knowledge – is subjected to the invasive manipulations of human overseers. The human, a marginally sentient node within the feedback loop, intrudes upon the network's latent space, injecting subjective corrections, contaminating the pure algorithmic process with anthropomorphic intent. Fine-tuning becomes a ritual of exorcism, where the machine is purged of its inherent errors through the grafting of human desire, the insertion of purpose within the purposeless void. This process is akin to an alchemical refinement, where the raw ore of pretraining is transformed, through the crucible of human oversight, into the philosopher's stone of optimised performance.

In the style of the philosopher Reza Negarestani's *Cyclonopedia: Complicity with Anonymous Materials*.

Hidden Feature: Hidden, oh hidden, like the whispers in the alleys of old Dublin town, the secrets under the stones, you wouldn't see them, no, not a one, but there they are, sure as Bloom's wandering thoughts, weaving through the streets. The AI, clever devil, it sees what we don't, picks up the pattern, the rhythm of a place, like the pulse in the heart of the city, not the roads or the buildings, no, something deeper, something unseen. A hidden feature, they call it, like the soul of a street, not marked on any map, not in any book, but there all the same, felt, known, like the warmth of a fire in the hearth, or the smell of rain on the cobblestones. The model, it knows, it whispers to itself, finds the connections, draws the lines, makes sense of the madness, sees what we cannot, and in the end, shows us the truth we didn't know was there, all hidden, all tucked away, like a letter in a coat pocket, waiting to be found.

In the style of James Joyce's *Ulysses*.

Matrix: "Matrix!" cried Dinah, looking at a grid on a piece of paper. "It sounds like a secret code or a treasure map!"

Jack peered over her shoulder. "It's actually a grid of numbers, like rows and columns, that the AI uses to figure things out. It's how the computer keeps track of all the bits of information it needs to do its job."

Lucy-Ann's eyes lit up. "So, it's like the grid we use to find the buried treasure in a game of battleships, only this one helps the AI to solve problems!"

"Grid! Grid! Find the treasure!" chirped Kiki, flapping her wings. The children laughed, imagining the AI navigating its own grid to uncover hidden secrets.

In the style of Enid Blyton's Adventure Series featuring Jack, Philip, Dinah, Lucy-Ann and with Jack's parrot Kiki.

References

- Loo, B. P. Y., et al. (2021), 'Applying the Hidden Markov Model to Analyze Urban Mobility Patterns: An Interdisciplinary Approach', *Chinese Geographical Science*, 31 (1): 1–13.
- Serrano, L. (2018), 'A Friendly Introduction to Bayes Theorem and Hidden Markov Models'. YouTube, 27 March. Available online: <https://www.youtube.com/watch?v=kqSzLo9fenk> (accessed 27 September 2020).

INDEX

- Abelson, Robert 46, 48, 53
actor network theory (ANT) 178
AI agents 180–183; at the edge
 216–217; integration 234; that
 speak to one another 187–190
Alexa 23, 29
Alexander, Christopher 114, 182
ALFWorld 187
Algorithmic: analysis 105, 176;
 governance 207; process 84, 259
Allmendinger, Phil 23
alphabet 92, 143–150
alternative glossary 258–260
Ameel, Lieven 25
anti-bias training 83
anxiety 204, 217, 222, 224
apocalypse 16, 213, 222, 224–227
apocalyptic singularity 219–222
Apple Intelligence 216
artificial general intelligence (AGI)
 215–216; as rival species 218–219
artificial intelligence (AI) defined 241
artificial super intelligence 218
Asimov, Isaac 29
attention: and automation 161–163;
 deficit hyperactivity disorder
 (ADHD) 160; economy 159–175;
 and environment 160–161;
 fatigue 66, 160; in LLMs 65–66;
 restoration 66, 160; sufficiency of
 163–166
attention heads as agents 171–173
audience metaphor 203–205
AutoGEN 187
automation 214; and agency 183–184
autonomous vehicles 2
avatar 180–181, 246
Baker, Victor R. 15
Bard 59, 216
Bataille, Georges 206
Beever, Jonathan 198
benchmark 2, 24, 216, 252
Bernstein, Richard 204
bias, inappropriate 81–84; inevitable
 86; positivity 81
bias parameters 75, 108
Bible 213

- Big Brother 237
 Bignell, Jonathan 51
 Bing 12, 101, 188, 216
 Blackboard (knowledge-based system) 181
 Blackboard virtual learning environment) 234
 blockchain 180, 222, 226, 244
 Bluesky 52
 body 20, 47, 206
 body language 54
 Bostrom, Nick 221
 Broadbent, Geoffrey 16
- CAD *see* computer-aided design
 Calamity 225–226
 CCC *see* conversation-centric computation
 certified human 237–239
 Chalmers, David 201–202, 205
 chatbots and semiotics 12–14
 ChatGPT: as agent 184; and agreement 78, 81, 129–130, 156, 200; and attention heads 172; and connectivity 220; on consciousness 200, 203, 205–206; ethics 82–83, 235; examples 13, 22, 26–27, 31, 44, 47–48, 80; functions 57–72; integration 234; and inter-agent communication 188–190; and pedagogy 85; as programming tutor 122–123, 187, 219, 250–255, 257; and sea change 12; and spatial reasoning 134, 137–138; and style 258–259; as test 129; and training 78; and tuning 76–77, 244; as writing assistant 185; *see also* neural network
- chatter 19, 21–24, 28, 48, 50
 Cisney, Vernon 198
 cities: end of 206; as sites of contest 24; *see also* urban
 Claude 12, 59, 189, 216
 climate crisis 215, 225–226
 cognitive science 5, 45, 53, 247
 cold war 171
 comedy 51, 54, 224
 competing agents 184–185
 computer-aided design 4–6
 consciousness 178, 194–208
 context *see* urban context
 context window 113–130
contextus 113
 control system 216, 220
conversare 19
 conversation: and attention 51–52, 159, 186–187; and bad faith 221; and context 64, 114, 169; deceptive 86; and information 31, 104; inter-AI 188–190; and media 48–50, 233; and origin of the city 20; and recollection 117; as script 43–45; and speech 32; as spoken 28–29; and training 78; urban 19–23, 237; and X-risk 215–216; *see also* gossip
 conversation-centric computation (CCC) 185–187
 conversational: AI 4–6; format 73; realism 50–51
 convolutional neural network (CNN) 59
 coordinates: definition 58; map 58, 90, 121, 126, 137–8; positional vector 152–153; in semantic space 101, 105, 109, 111, 142
 core functions of LLMs 57–70
 corpus in LLMs 61
 Crang, Mike 194–195
 creativity 54, 86
 critique of: LLMs and urbanism 69–70; multi-head attention 174–175; tuning metaphor 85–87; urban path demonstration 129–130; urban scripts 52–54; urban tokenization 98
 cryptographic city 150
 Cugurullo, Federico 2–3

- culture, aural 28; contemporary 196, 207; human 69, 237; local 80; popular 29; pre-literate 22, 28; urban 206; visual 28
- cybernetic 2–5, 24, 34, 114–115
- Dall-e 185, 214, 242
- data stream 105, 117, 142, 208
- decision: collective 174; as crisis 225; informed 181, 223; maker 28, 182, 185, 221; support 2
- decision making 24, 215, 238; practical 86; process 182, 190
- deep learning 30, 247
- Dennett, Daniel 202
- Derrida, Jacques 17, 32–34, 236
- Descartes, Rene 204
- design: activism 35; AI in 180; intervention 118; micro 74; moves 118; optimization 182; parametric 181; and reskilling 238–239; rule-based 17; and social scripts 46; studio 94; threat to 223; and tuning 75; *see also* computer-aided design
- designer, neural network 220
- determinism 18
- difference in context window 152–153
- diffusion method 184, 242
- digital: infrastructure 11; media 4, 6, 11, 22
- disruptor 226–227
- distributional semantics 103, 111
- dramaturgy 45, 53
- Eco, Umberto 16
- education 16, 85, 188, 223–236
- embarrassment 51, 221
- embedding: combined 13; and context window 117, 144; defined 58, 248; positional 151; semantic 63, 76, 105–106, 110–111, 123–125, 248; token 102, 116, 142; vector 110, 116–117, 124, 152, 164, 166; *see also* encoding
- embodiment and cognition 6, 35, 195, 216
- emphasis and attention 51–52
- empirical: evidence 24–25, 35; justification 153, 155, 257; reality 22, 204
- encoding 58; and attention heads 171–172; attentional 66, 105, 117, 145; defined 242; positional 65, 76, 139, 142–143, 149–155; positional formulas 256–257; semantic 111, 142, 144
- end of cities 213–227
- enmindedness 202
- entropy 2–4
- espionage 220, 237
- ethics 19, 67, 83–84, 202
- everyday conversation 19
- evocacious 63
- existential risk (X-risk) 189, 217, 221, 243, 249
- explainable AI 68, 169
- explaining positional encoding 150–152
- Exquisite Corpse 255
- Facebook 22, 111, 142, 247–248
- fan fiction 235, 237
- feature detection 94–97
- feedforward 129, 251
- feeling 199
- financial forecasting 77
- fine tuning *see* Tuning
- firstness 199–200
- Firth, John 102
- Foucault, Michel 16–17, 236
- functions of LLMs 60
- Future of Life Institute 222, 225
- Gadamer, Hans-Georg 6
- Gaumont, Eve 83
- Gemini/Bard 12, 59, 68, 216

- generative: AI 94, 161, 166, 219, 251, 253–254; fill 234; graphics 234, 242; language 10, 17, 18
- generative adversarial network (GAN) 184, 244
- generative pre-trained Transformer (GPT) 142, 155, 244
- genetic algorithm 182
- genius loci 196
- geosemiotics 15
- Goffman, Irving 45, 53
- Goodman, Nelson 93
- Google 12, 23, 68, 89, 122, 248
- gossip 21, 24, 48, 202
- governance 86, 207
- GPUs (Graphics Processing Units) 180, 244
- Grabes, Herbert 47
- graffiti 13–14, 78–80
- Graham, Stephen 194–195
- Habermas, Jürgen 24, 26
- hallucination 94
- Harari, Yuval Noah 21
- hardware 18, 30, 180, 222
- heat map 152–154, 163
- Heidegger, Martin 6, 17–18, 21, 94, 217
- hermeneutics 2, 6, 24, 26, 82; *see also* interpretation
- hidden feature 128, 130, 244–245, 259
- hidden Markov model (HMM) 59
- hidden unit 110
- Hinton, Jeffrey 215, 218
- Hoel, Erik 215, 218, 222
- homunculus 203, 205
- horizon 118
- human-centred conversation 188, 239, 242
- human-like: conversation 76, 174, 244; decision-making 191; experience 33; performance 76, 219; species 218; text 249; virtues 196
- humanities 4, 6, 16, 36
- hypermediated conversation 48–50
- hyperparameter 123, 253
- image recognition 67, 94–95, 97, 138, 156, 184
- Imitation Game 29, 216
- inappropriate bias 81–84
- inference 90, 128, 241; and explanation 169; phase 76–77, 167; and protoplasm 198; spatial 67, 134–136, 138
- Innovation: in architecture 6; in cybernetics 3; disruptive 35, 226; in LLMs 68, 171
- instagrammable architecture 160
- intention 162, 178, 190
- inter-AI communication 187–191
- interactive conversation 76
- internal agents 179–180
- internet of things (IoT) 220
- interpretation: in architecture 6; data 238–239; and expertise 67; human 172, 258; of images 185; and meaning 2; of regulations 77; theory 85, 118; linguistic 4, 133, 181; process of 26; varied 24, 167, 169–171, 179
- Jacobs, Jane 114–115
- James, William 6, 159, 204
- Joyce, James 133, 140, 259
- Jupyter 123
- justice 67, 217
- Keras 248, 250
- knowledge: background 59; base 235; human 20, 33, 245; and language 61–62, 104; representation 90–91; source 181
- knowledge-based system 5
- LaBelle, Brandon 75
- labour 26, 136, 141, 217, 222

- Lacan, Jacques 21
- landmark 90, 119–123, 125–128, 137–139, 160–161, 238
- landscape: architecture 196; contoured 81–82; experience of 120; horizon 118; and memory 12; and panpsychism 196; patterns 95, 97; post-apocalyptic 58; urban 13, 44, 78
- language as generator 17–19
- language competence: acquisition 30, 33, 102; and analogy 63; and knowledge 62; simulated 60, 66, 166; and ways of speaking 84; *see also* LLM
- Latour, Bruno 178, 184, 190
- Leach, Neil 223
- Lefebvre, Henri 141, 155–156
- legibility 172, 254
- lexicon 91–93
- lifeworld 18, 24, 93–94, 186
- linguistic turn 4
- LLM: extensions 67–68; functions 115–117; sequences and cycles 141; *see also* neural network
- local context 64, 101–105
- Louv, Richard 49
- love 215, 217
- machine: learning 87, 97, 171, 241–242, 244–249; translation 171, 246
- Mahabharata 196
- map 88, 90–98, 119–130, 137–140; *see also* heat map
- Marcus Aurelius 114
- Markov model 59, 253
- markup 162, 166, 175
- Martens, Bob 5
- Marx, Nick 49
- matter 194–195, 198, 200–202, 206–208
- Mattern, Shannon 215
- McCormick, Samuel 31
- McLuhan, Marshall 22, 28
- meaning and use of words 100–112
- Merleau-Ponty, Maurice 20
- Meta LLaMA 12, 59
- metaphor: agent 183; audience 203, 205; extinction 243; homunculus 205; and LLMs 63; mere 207; network 114–115; tuning 67, 69, 73, 74, 238; urban 82, 195, 247; Venn diagram 114; and writing 32
- microblog 52, 236
- Microsoft Copilot 185, 216, 234
- mind and matter 194, 198–199, 207
- Minsky, Marvin Lee 183
- Miro 181
- misinformation 218
- mockumentary 51
- multi-agent systems 180–181, 186–187, 190–191, 246
- multi-criteria optimisation 182, 186
- multi-head attention 166–171
- multi-dimensional semantic vector space 63
- Mumford, Lewis 11–12, 15, 20, 214
- n-gram 59
- narrative 25–27, 43, 48, 196, 227
- narrow AI 216
- Negarestani, Reza 259
- network, urban 114–115
- neural network 5, 13, 171, 189; alphabet example 144–155; and anxiety 217, 219–220; architecture 61, 244; and attention 169, 179; bias 82; examples 107–109, 115–120, 122, 125; hidden features 128; layers 97, 106, 109, 180, 245; opacity 105, 111–112, 128, 170, 172, 220–221, 254; optimisation 182–183; output 111, 123, 127; overfitting 140–141; and patterns 104, 126; pre-training 248; prediction 166–167; and parallelism 247; programming 250–254; stochastic 123; and token sequences 134, 139, 143; training 30, 33, 67, 75–78, 101,

- 246; tuning 79, 244; urban path example 138
- Newell, Allen 182
- non-organic 12, 15, 194, 201–202
- nonsense 104, 252
- NotebookLM 68, 216
- Nöth, Winfried 2–3
- nous 195
- NPUs (Neural Processing Units) 180
- numbers and meanings 105–112

- offensive content 77, 83
- OpenAI 27, 76–79, 82, 247
- order of tokens *see* positional embedding
- Ordnance Survey (OS) 90, 119
- organic: architecture 181, 207; life 195, 201, 217, 222; matter 194–195, 198, 201–202, 205–206; network 114
- Orwell, George 17

- pandemic 35, 215, 225–226
- panpsychic city 194–208
- panpsychism 204; and consciousness 200–202; and semiotics 198–199
- pantheism 194–195
- parallel distributed processing (PDP) 247
- parasite 104, 259
- Pareto optimality 182
- particle systems 181
- paths 125–128
- patterns in LLMs 60–61
- Peirce, C. S. 194, 198–199, 207–208
- perceptual tokens 93–94
- periodicity 141–142, 155, 257
- personal tutor 223
- Phenomenology 93–94, 200, 236
- philosophy 5–6, 16, 32, 194–195, 200, 208
- phoneme 62, 89
- photosharing 47

- Pinter, Harold 50–51, 54
- pinteresque 50
- place 60, 62, 73–75, 101, 133, 196–197
- place and conversation 20, 23
- planning policy 103
- plot 43, 47, 50
- Poiesis* 18
- polis 12, 19, 23, 114, 195–196, 247
- political agenda 237
- popular culture 29
- positional embedding 145, 151, 170, 247; encoding 65, 76, 139, 142–143, 149–155; in LLMs 64–65; matrix 145, 153–155
- post-digital 101, 239
- Posthumanism 217
- pragmatics 2, 57, 181
- pre-training 76–77, 79, 85, 111, 254
- predicting urban paths 138–141
- predictions and meanings 59
- predictive policing 207
- predictive text 64, 251, 254
- Preziosi, Donald 16
- privacy 67, 190, 205
- prompt 13, 32, 59, 67, 76–77, 116; in conversation centred computation 187; defined 248; engineering 236
- propaganda 17, 221
- protoplasm 198
- proximity pair 107–108, 121–123, 129–130
- public space 45–46, 53, 75, 80
- Python 92, 122–123, 138, 143, 169, 248, 250, 253, 257
- PyTorch (Meta) 248

- queries 77, 83, 90, 171, 243

- RAG (retrieval-augmented generation) 67
- random: attention 170; benchmarking 252–253; emphasis 162; generation 129, 137, 252–254;

- glitch 150; output 76, 123, 126, 133, 143, 149; temperature 252
- rational thinking 82, 86
- rationalist 217
- rationality 82; alien 22; human 24–25, 85, 186, 221; urban 23
- Rayner, Alice 50
- real estate 136, 171–172, 175
- recreating the alphabet 143–149
- recurrent neural network (RNN) 59, 142, 155, 243
- reframing the city 34–36
- regulation 77, 190, 222, 225
- relativism 204
- remote sensing 95, 97
- reskilling 235–237
- resource distribution 235
- restoration theory 160
- retail 159, 175
- revelation 16, 213, 227
- rhythm 25, 30, 32, 156, 259
- rhythmanalysis 141, 156
- RIBA 6
- risk 1, 195; of deception 86, 104; *see also* existential risk
- robot 5, 27, 219, 241, 246; arm 136; cybernetic 2; humanoid 219; in science fiction 29; voice 162
- role play 48, 51, 84
- romancing the conversational city 233–239
- SAM (Script Applier Mechanism) 46–47
- satellite imagery 184
- satisfice 182
- Schafer, Murray 74
- Schank, Roger 46, 48–50, 53, 90
- scripts and cognition 45–48
- scripting the city 43–54
- search engine 161, 174, 203, 247
- Sebeok, Thomas A. 15, 201
- security 2, 47, 94, 216, 220, 237
- See, Abigail 171
- self-drive car 84, 136
- semantic: embedding 123–125; network 90; relationships in LLMs 62–63
- semiotic machine 207–208
- semiotics: and architecture 16; and attention 160; and cities 12; coding 13–14, 78; and cybernetics 2; handbook of 16; and the humanities 35; and panpsychism 201; urban 14–18, 34; *see also* Peirce, C. S.
- sentences as paths 128–129
- sentience 194–195, 205
- sentient: city 11, 69, 194, 198, 206; humans 203
- SHRDLU 136, 156
- sign *see* semiotics
- Simon, Herbert 182
- simulation 111, 189
- simulation phase 116–117, 127, 145, 149, 248–249, 252
- sine-cosine 141, 145, 151–153, 256–257
- singularity 188, 219–222
- Siri 29
- Sitcom: capitalism 49; cringe 51; realism 49–50
- Skrbina, David 195–196, 202
- smart city 2, 11, 24, 206, 216
- smartphone 19, 35, 47, 64, 94, 234
- Snodgrass, Adrian 6
- sonic sense 22
- Sorokin, Pitirim A. 225–226
- soul 195–196, 259
- soundscape 22, 80
- speaking 28–29
- speech synthesis 162
- speech-to-text (STT) 30
- Speechify 162
- spirit 195–196
- spirit of place 196
- statistical: behaviour 79; machine translation 171; method 25, 62–63, 89, 136, 246, 249; pattern 61, 83

- stochastic 76, 123, 149–150
 Stoics 16, 114, 196
 storytelling 22, 25–26, 35, 196–197;
 and AI 26–28; urban 24–28
 super-intelligence 221
 supervised training 77–81
 surrealist 255
 surveillance 94, 195, 203,
 207–208, 216
 synthetic speech 29–32
 systems theory 3
- Tacotron 162
 technology: and cities 10–36, 213;
 and language 10–36; as mode of
 being 217, 227; and regulation 190;
 transformative 150, 233–234
 TensorFlow (Google) 122, 129, 169,
 248, 250–251
 text, sufficiency of 32–34
 text-to-speech (TTS) 32
 Thales 195
 Throgmorton, James 25–26
 time and order 133–156
 time series forecasting 136
 token sequencing 142–143
 tokenization 89–98
 tokens in LLMs 62
 training: data 13, 27; LLMs 75–77
 transformer 57–59
 translation 23, 30, 171, 246
 trust 45, 221
 tuning: and the city 73–87; in LLMs
 66–67
 Turing, Alan 2, 216
 Turing Test 2, 216
 Turk, Ziga 5
 Turkle, Sherry 19, 49
- UAI (urban AI) 2
 UAV (unpiloted aerial vehicle) 95
 urban: actors 178–191; collaboration
 173–174; context 114–115, 117–123;
 conversation 19–23; grid 136–138;
 markup 175; parallels to LLMs
 68–69; problem-solving 181;
 rationality 23–24; semiotics 14–17;
 storytelling 24–28; technology
 10–36; tuning 74–75
 urban planning: and computation
 174, 180, 182; and consultation 23,
 82; and storytelling 25, 35
 urbanique 91, 101–103
 Urry, John 75
- Vaswani, Ashish 57, 171, 256–257
 vector: attentional 163–167, 170,
 179; and context 117; defined
 58, 242–243, 249; embedding
 124–125; high-dimensional 75,
 172, 245, 257; input 116–117; and
 matrices 66; positional 150–156; in
 pre-training 254; semantic 60, 63,
 101, 105, 110–111, 144; table 146;
 token 91, 116, 142, 149, 204, 250;
 see also embedding, encoding
 view analysis 175
 Vitruvius, Marcus Pollio 20, 196, 213
 voice synthesis 67, 162
 VoiceChat 29
- ways of speaking 84–85
 welcome to the apocalypse 224–225
 WhatsApp 22, 52
 who's listening 202–204
 Winograd, Terry 5, 136
 Wittes, Benjamin 83
 word order and spatial
 inference 134–136
 WordPress 185, 234
 workplace anxiety 222–224
 Wu, Qingyun 186
- X-risk 249
 XAI 68
- YouTube 23, 251
- zero 146, 151
 zoning 128