
AI Insights

ARTIFICIAL INTELLIGENCE IN HIGHER EDUCATION

A PHENOMENOLOGICAL VIEW

EDITED BY
SENCER YERALAN

ARTIFICIAL INTELLIGENCE
IN HIGHER EDUCATION

Artificial Intelligence in Higher Education

A Phenomenological View

Edited by Sencer Yeralan



<https://www.openbookpublishers.com/>

©2026 Sencer Yeralan (ed.), Philip Aka, Derek Baker, Azra Branković, Eka Gegeshidze, Laura Ancona Lee, Christos Michalakelis, Mbulaheni Nthangeni, Efthymia Staiou (members of the Global Academics Coalition)



This work is licensed under a Creative Commons Attribution-ShareAlike license (CC BY-NC 4.0). This license allows you to share, copy, distribute and transmit the text; to adapt the text for non-commercial purposes of the text providing attribution is made to the authors (but not in any way that suggests that they endorse you or your use of the work). Attribution should include the following information:

Sencer Yeralan (ed.), *Artificial Intelligence in Higher Education: A Phenomenological View*. Cambridge, UK: Open Book Publishers, 2026, <https://doi.org/10.11647/OBP.0525>

Further details about CC BY-NC licenses are available at <https://creativecommons.org/licenses/by-nc/4.0/>

All external links were active at the time of publication unless otherwise stated and have been archived via the Internet Archive Wayback Machine at <https://archive.org/web>

Any digital material and resources associated with this volume are available at <https://doi.org/10.11647/OBP.0525#resources>

Information about any revised edition of this work will be provided at <https://doi.org/10.11647/OBP.0525>

AI Insights: Volume 1

ISSN Print: 2979-1731 | ISSN Digital: 2979-174X

ISBN Paperback: 978-1-80511-872-5

ISBN Hardback: 978-1-80511-873-2

ISBN PDF: 978-1-80511-874-9

ISBN HTML: 978-1-80511-876-3

ISBN EPUB: 978-1-80511-875-6

DOI: 10.11647/OBP.0525

Cover photo by Sencer Yeralan

Cover design by Jeevanjot Kaur Nagpal

Table of Contents

Preface	vii
Editorial Note	xi
Contributors	xiii
Acknowledgments	xvii
Prelude: Thoughts at the Top Dead Center	xix
I. FOUNDATIONS AND FRAMES	1
1. Historical and Philosophical Context	3
2. The Ideals We Left Behind: From <i>Bildung</i> to the Market	9
3. Technological Disruption and Platformization in Higher Education	15
<i>Interlude I: A Student's Whisper</i>	21
II. VOICES FROM THE FIELD	23
4. Ethical Frameworks and Institutional Responsibilities	25
5. Institutional Archetypes: Flagships, Regionals, Faith-based, Global, Online, and Instructorless	31
<i>Interlude II: A Minister's Memo</i>	35
III. PATTERNS AND TENSIONS	37
6. Cross-cutting Themes and Divergences	39
<i>Interlude III: A Knock at the Door</i>	57
IV. COMPARATIVE CASE STUDIES	59
7. Blind Spots, Frictions, and Unanticipated Consequences	61
<i>Interlude IV: Συνπρωτητής</i>	65

V. TOWARD A GLOBAL FRAMEWORK	67
8. Between Logic and Meaning: Why Off-the-Shelf Ethics and Prediction Fail	69
9. From Reflection to Decision: The Missing Middle of AI Governance	75
10. Readiness, Resistance, and Renewal: A Typology of Institutional AI Response	81
11. Roadmaps for Policy and Practice	85
VI. ENTANGLED FUTURES AND STRATEGIC REALITIES	89
12. Five Reflections for the Next Decade	91
13. Entangled Futures: AI, Education, and the Systemic Web	95
14. The <i>Realpolitik</i> of AI in Higher Education	99
Appendix A	103
Appendix B	107
Bibliography	113
Index	119

Preface

The ever-growing entanglement of artificial intelligence with higher education is no longer a novelty, nor a prediction. It is a present reality. In the past few years, we have seen how generative AI tools, especially large language models, have migrated from the fringe to the mainstream. They now quietly reside in student toolkits, administrative workflows, and faculty deliberations. This book is a reflection on that entanglement.

Authored by senior scholars, administrators, and researchers with extensive international experience, the volume reflects decades of helping universities navigate strategic planning, curriculum development, accreditation, and global partnerships. We express our gratitude to the many institutions that welcomed us into their inner sanctums – not in search of quick or ready-made solutions but seeking understanding deep enough to chart their own course. Their openness and courage inform every chapter of this volume. Across these engagements, the contributors have witnessed a shared pattern: confusion, hesitation, and initial knee-jerk reactions in the face of AI's rapid advance.

This book is offered in a spirit of altruism – to share what we have had the privilege to observe, so others may proceed with greater clarity, confidence, and calm. Our collective emerged from disparate regions of the globe, from varying disciplinary traditions, and from roles both inside and adjacent to academia. What unites us is a shared concern – and, more importantly, a shared curiosity – about what AI is doing *to* and *for* higher education. Rather than presenting a single coherent theory or strategy, we offer a patchwork of views: ethnographic, institutional, philosophical, practical, even autobiographical.

These perspectives do not always converge. Nor are they intended to. In fact, one of our goals is to resist premature consensus – to carve out space for voices that might otherwise be marginalized in AI discourse dominated by corporate interests or elite institutions. What results is not

a guidebook, but a conversation – one that spans continents, roles, and realities.

The purpose of this book is twofold. First, to equip higher education leaders with the language and conceptual tools needed to understand and respond to AI not merely as a technology, but as a force embedded in institutional, political, and cultural systems. Second, to encourage meaningful reflection and dialogue across borders and sectors about the future we are building – not only for students and educators, but for the knowledge commons itself.

The audience we envision includes university rectors and presidents, deans and provosts, ministers and policymakers, accreditation bodies, and senior faculty across disciplines. While technical specialists may find much of interest here, this volume does not assume deep fluency in computer science or engineering. Rather, it addresses those who shape strategy, set priorities, and hold stewardship over the values and missions of higher education.

The *tone* of the book is intentionally interdisciplinary and dialogical. It invites pause, context, and comparison. Many of the chapters raise more questions than answers. This is by design. AI is not arriving into a vacuum; it amplifies existing tensions – between access and excellence, between metrics and meaning, between the market and the mission. The book is organized thematically rather than linearly, allowing readers to engage with the sections most relevant to their roles or regions.

This book is not a definitive document that puts a period at the end – nor does it offer strong claims, recipes, or prescriptions. Rather, it presents a phenomenological view of artificial intelligence (AI) in higher education: a reflective, multi-faceted exploration of how AI interacts with the structures, values, and aspirations of academic institutions around the world.

Rather than simplifying or forecasting, the chapters invite contemplation, dialogue, and institutional self-examination. Through essays, interludes, case studies, and typologies, the book examines how AI reshapes longstanding tensions between efficiency and purpose, and between governance and imagination. Its tone is dialogical, its structure thematic, and its goal not consensus but thoughtful engagement.

The essays and chapters in this volume are arranged in a loosely thematic order. We begin with foundational reflections on the meaning

of education, then explore institutional adaptation, pedagogical challenges, regulatory shifts, and evolving labor expectations. Later sections offer speculative views on what lies ahead, both in terms of risks and opportunities. Along the way, we introduce interludes that present stylized, composite voices – educators, administrators, technologists, and students – constructed to illuminate the lived experience of AI in education. These interludes are not transcripts or direct quotations, but reflective representations of recurring perspectives encountered across institutions.

We write in the conviction that AI’s impact on higher education is not preordained. It is shaped by choices – by what we value, protect, and imagine. This volume is offered in the spirit of critical hope, institutional memory, and global solidarity. Taken together, these contributions aim to inform, provoke, and support those navigating the rapidly shifting terrain of higher education in the age of generative AI. We invite the reader to step into this mosaic not expecting conclusions, but seeking insight. Like education itself, this volume is less a destination than a journey.

This volume is intended for university leaders, policymakers, accreditation bodies, and senior educators seeking not technical fluency, but conceptual clarity. It is a catalyst – one among many – for navigating the uncharted waters of AI with purpose, integrity, and institutional courage.

We welcome comments and suggestions at: <https://globalacademicscoalition.net>

Sencer Yeralan, P.E., PhD
Editor
Gainesville, January 2026

Editorial Note on Scope and Representation

This volume brings together perspectives shaped by diverse institutional roles, geographic contexts, and disciplinary traditions. Contributors write from experiences spanning flagship and regional universities, public and private systems, Global North and Global South contexts, and environments marked by differing levels of resource constraint and regulatory oversight. Disciplinary orientations include philosophy, education, governance, systems thinking, and institutional leadership, reflecting the book's commitment to pluralism rather than consensus.

At the same time, this collection does not claim comprehensive representation. The voices gathered here are primarily those of senior scholars, institutional leaders, and reflective practitioners; student perspectives appear indirectly through interludes rather than as authored chapters, and highly technical or implementation-focused viewpoints are treated selectively rather than exhaustively. These boundaries shape the scope of the analysis and should be read as deliberate constraints rather than oversights.

The aim of the volume is not to speak for all institutions or contexts, but to surface patterns, tensions, and questions that recur across many of them. Readers are invited to interpret the arguments presented here in light of their own institutional realities, recognizing both the commonalities and the limits of transferability.

Contributors

This volume is the result of a collective effort by scholars affiliated with the Global Academics Coalition. The Global Academics Coalition (GAC) is an international alliance of senior scholars and academic leaders who have held high-level positions across multiple countries, disciplines, and institutional types. Its members bring a wealth of interdisciplinary experience and global insight to the evolving challenges of higher education. Listed below are the contributing authors, editors, and collaborators, in alphabetical order:

Philip Aka is Professor of Law and a Fulbright Scholar with doctorates in law and political science from US institutions. His work spans law, bioethics, AI, peace studies, and higher education. He has served as Dean of the Faculty of Law and as External Examiner for the Center for International Affairs and Diplomacy at the University of Ghana. His research includes comparative healthcare systems and governance in post-colonial contexts. His forthcoming book, *Navigating the Tension Between Sovereignty and Self-determination in Post-Colonial Africa*, will be published by Palgrave Macmillan.

Derek Baker is a researcher in solar energy systems and techno-economic analysis. He coordinates the EU Horizon Europe SolarHub project and has led multiple Horizon 2020 and Horizon Europe initiatives focused on renewable energy collaboration across European institutions. His work bridges engineering research, international consortium leadership, and clean energy innovation.

Azra Branković is an economist whose academic career has included appointments at the University of Sarajevo, Athens University of Economics and Business, Bocconi University, and ADA University. She

has also held leadership roles in public administration and international organizations. Her research focuses on higher education economics, foreign direct investment, SMEs, and institutional development, with particular attention to post-conflict and transitional economies.

Eka Gegeshidze is Professor in the School of Business and Administrative Science at the University of Georgia. Her work centers on innovation, entrepreneurship education, and economic development. She has extensive experience in policy design, program budgeting, and higher education accreditation, and has contributed to numerous international projects related to governance, commercialization of innovation, and institutional reform.

Laura Ancona Lee is Professor of Innovation and Creativity and Vice-Rector for International Relations at Universidade Paulista, Brazil, where she also serves as Erasmus Institutional Coordinator. Her work focuses on academic internationalization, interdisciplinary collaboration, and global partnership development. She leads research initiatives on international academic mobility and institutional cooperation.

Christos Michalakelis is a scholar of techno-economic analysis of telecommunications networks with degrees in mathematics, telecommunications management, and software engineering. He heads the "Study in Greece" program promoting internationalization of Greek higher education. His work includes ICT systems design, academic administration, and participation in national and European educational initiatives.

Mbulaheni Nthangeni is a scholar of mathematical statistics and systems engineering with four decades of experience in academia and public service. He has served as Director of the School of Mathematical and Computational Sciences at the University of Limpopo and previously at the University of Cape Town. His public service includes leadership roles at Statistics South Africa and presidency of the South African Statistical Association.

Efthymia Staiou is a faculty member in Industrial Engineering at Yasar University. Her expertise includes quality management, engineering economics, entrepreneurship, and sustainability. She has led and

participated in numerous EU-funded projects in regional development, innovation management, and higher education collaboration, and brings extensive experience from both academia and public administration.

Sencer Yeralan (Editor) is a systems engineer and academic administrator whose career spans engineering, mathematics, and higher education leadership. He has served as Dean of engineering, business, IT, and natural sciences on three continents. His research ranges from embedded systems and robotics to academic governance and philosophy of engineering. A Fulbright Scholar, he has served as an expert reviewer for the European Research Council and the US National Science Foundation. He has authored books in engineering and education and continues to work at the intersection of systems thinking and institutional transformation.

Acknowledgments

We do not see things as they are, we see them as we are.

—Talmudic proverb (also attributed to Anaïs Nin)

Prediction is very difficult, especially if it's about the future.

—Niels Bohr (or Yogi Berra, depending on folklore)

We extend our sincere thanks to the many international universities with whom we have had the privilege to collaborate over the years. Their trust, openness, and courage – in moments of uncertainty and transformation – have left a lasting impression.

Far from being passive recipients of external advice, these institutions invited us into their inner deliberations: to listen, to reflect, and to think alongside them. In doing so, they gave us not only insight, but something more valuable – a kind of unique situational awareness that no literature review or global summit could have offered.

This volume is shaped by those encounters. What we now share in these pages is not a theory from above, but a view from within – born of dialogue, learning, and mutual respect. We are deeply grateful for the chance to walk briefly beside those navigating their own path through the uncertain terrain of higher education.

Prelude

Thoughts at the Top Dead Center

A typical class may have started as...

My job as an instructor is not to tell you what is in the book. You can read the book. My job is to tell you what is not in the book. What you will not find in a footnote or appendix. What came not from theory, but from lived experience.

... and continued:

There is a dualism I (Yeralan) keep returning to – one that sits quietly behind the scenes in both computing and cognition: the dualism of computation and memory. The CPU and the RAM. The engine and the archive. The logic and the lookup.

It may sound abstract, but this dualism shaped my early career in a very practical way.

In the early 1980s, I was designing engine control units – primitive computers tasked with deciding, thousands of times per minute, how much fuel to inject and when to spark the ignition. Each decision depended on a dozen variables: torque, RPM, ambient temperature, barometric pressure, knock sensor readings, coolant state, and more. Each decision had to be made at each top dead center of the crankshaft.

Now, a good thermodynamics student could have modeled this analytically. But our processors were 8-bit workhorses. A 32-bit CPU would cost more than the entire car. So, we had a choice: compute... or remember.

We chose memory.

We filled memory with tables – lookup grids indexed by torque and speed, populated with values for spark advance, fuel dwell, cold-start enrichments, and so on. One table per purpose. A few variants per

weather condition. No calculus. No differential equations. Just fast table lookups. The puny processor could manage that.

And the engines purred. And reduced pollution manyfold.

We did not get something for nothing. We traded logic for memory. It was a pragmatic act of engineering – a kind of epistemological substitution.

That phrase – epistemological substitution – matters here, because it gestures toward how this book proceeds.

What follows is not a technical treatise, nor an empirical survey in the conventional sense. It is a phenomenological account in a more modest, institutional register: an attempt to describe how artificial intelligence is being encountered, interpreted, and operationalized within higher education as a lived system. The emphasis is not on what AI is in theory, but on how it appears in practice – how it reshapes judgment, authority, labor, and responsibility as it enters existing institutional forms.

The observations in this book arise from synthesis rather than experiment: from administrative experience, cross-institutional exposure, sustained dialogue with faculty and policymakers, and close observation of emerging practices during a period of rapid adoption. They reflect patterns noticed over time, not measurements frozen at a moment. This perspective is necessarily partial, situated, and historically contingent.

The limits of such an approach are real and acknowledged. No claim is made here to causal proof, predictive certainty, or comprehensive representation across regions or sectors. The aim is instead diagnostic: to surface tensions before they calcify into policy, to articulate questions before they are prematurely answered, and to clarify what may be lost when systems optimize for efficiency without attending to meaning.

Now, flash forward.

Modern AI systems do not derive from first principles. They do not “understand” the world. They do not model combustion or cognition. They ingest examples – billions of them. They learn not by rules, but by statistical regularity. They trade human insight for machine memory.

It is the same trade.

But now, the consequences are different.

Neither the dealer nor the driver cares whether the engine control system is rule-based or lookup-based – as long as it works. And works

better, more efficiently, and with less smoke out the tailpipe. But in education, law, and medicine, finance, journalism, entertainment, society does care. Replacing expert reasoning with a data-trained system changes not only the method – it changes the structure of labor. That very structure that has been with us since human societies moved from hunter-gatherer to agrarian. If expertise is bypassed, experts go unemployed. Institutions shift. Families feel it. Children grow up in a different landscape.

So, my dear students – if you remember nothing else, remember this:

Not all computation is done by CPUs. And not all costs are paid in silicon.

FOUNDATIONS AND FRAMES

1. Historical and Philosophical Context

What Do We Mean by Artificial Intelligence?

Artificial intelligence (AI) has no single, universally accepted definition. It is a term that means different things to different communities: a machine-learning researcher may think of probabilistic models and neural networks; a university administrator may think of automated grading or chatbot tutors; a policymaker may associate AI with surveillance or labor displacement. Like the proverbial blindfolded individuals touching different parts of the elephant, each group encounters a different facet – none wholly wrong, but none wholly sufficient.

To avoid confusion and foster clarity, we begin with a shared working definition suited to the institutional context of this book. For our purposes, artificial intelligence refers to computational systems that perform tasks which, if done by humans, would be said to require intelligence. These tasks may include language processing, visual recognition, decision-making, pattern analysis, problem-solving, or generative composition. AI systems range from narrow tools that classify data or automate workflows, to large foundational models capable of generating text, images, or code.

Among the most visible forms of AI in current discourse is what is commonly referred to as “generative AI” – systems that create new content, such as essays, images, music, or software code, based on patterns learned from large datasets. While generative AI has captured public attention, it is but one manifestation of the broader paradigm we examine. Our focus includes, but is not limited to, such systems.

Our focus is not on the technical details of implementation, but on how such systems interact with the structures, values, and missions of higher education. This definition is therefore deliberately sociotechnical: it recognizes that AI does not act in isolation, but is embedded in institutional practices, shaped by policy, and interpreted through culture. The purpose of this chapter is therefore orientational rather than demonstrative. The historical and philosophical perspectives introduced here are not offered as authorities that resolve contemporary questions, nor as comprehensive accounts of the field. They serve instead as interpretive lenses – tools for clarifying how current debates about artificial intelligence in higher education inherit older assumptions about mind, knowledge, agency, and value. In this sense, the chapter establishes a conceptual vocabulary rather than an argument, preparing the ground for later discussion of institutional practice, ethical tension, and policy response. As good philosophers do, we begin with description, proceed to provisional definitions, and then move toward examining structure and intent. We acknowledge that this definition, like the systems it describes, will evolve – and that a future day may come when it becomes obsolete. The term “AI” itself may dissolve into the very infrastructure of life and learning, no longer needing special distinction.

Throughout this volume, we invite the reader to hold this definition with both precision and flexibility: precise enough to enable shared discourse, flexible enough to accommodate the multiple meanings AI carries across contexts.

Demystifying Neural Networks and Emergent Intelligence

The foundational idea behind artificial neural networks (ANNs) was not born of modern computation, but rather from the mid-20th-century effort to replicate, in silico, some functions of the human brain. Pioneers such as McCulloch and Pitts developed simplified models of neurons in the 1940s, while Hebb introduced his influential theory of synaptic learning shortly thereafter (McCulloch and Pitts 1943; Hebb 1949). These early models laid the groundwork for what would eventually evolve into today’s complex, multi-layered neural architectures.

While the terminology of “neural networks” suggests a biological correspondence, the resemblance is superficial. Biological neurons are not simple nodes in a directed graph; they are rich, chemically and electrically dynamic agents embedded in complex, adaptive environments. ANNs, by contrast, are engineered approximations. Their usefulness lies not in simulating biological intelligence, but in mimicking its input-output behavior. Thus, it is essential to recognize that ANNs are inspired by, but not equivalent to, the human brain.

Despite this fundamental distinction, the rapid advancements in machine learning have led to the proliferation of metaphors that obscure rather than illuminate. Terms such as “learning,” “understanding,” and “intelligence” are applied to systems that, strictly speaking, do not possess these qualities in any human sense. What ANNs achieve is not cognition but *optimization*. They adjust internal parameters to reduce error against a given objective function – often by processing vast quantities of labeled data (LeCun, Bengio, and Hinton 2015). The illusion of intelligence emerges not from conscious insight, but from statistical regularity.

Nevertheless, this statistical regularity gives rise to capabilities once thought exclusive to human cognition. Image classification, natural language generation, and strategic game playing are now performed at superhuman levels (Silver et al. 2016; Brown et al. 2020). This raises a philosophical question: if behaviorally indistinguishable from intelligence, does it matter whether a system truly “understands”? The Turing Test, after all, does not demand awareness – only indistinguishability (Turing 1950).

Yet this very question evokes discomfort. We hesitate to ascribe intelligence to machines, not because of what they do, but because of what we believe they lack – intentionality, self-awareness, and lived experience. The dissonance between observable behavior and presumed inner life suggests that our concept of intelligence is deeply entangled with notions of consciousness and selfhood. Demystifying ANNs, then, requires not just technical exposition, but philosophical clarity: to understand what these systems are, we must also be honest about what they are not.

From Symbolic AI to Machine Learning

The early decades of artificial intelligence were dominated by symbolic approaches – rule-based systems built on logical inference, expert knowledge bases, and formal ontologies. This paradigm, often referred to as “Good Old-Fashioned AI” (GOFAI), sought to encode intelligence through explicit structures and rules. Successes were notable in narrow domains, such as chess or algebraic manipulation, where the problem space could be exhaustively defined (Newell and Simon 1976).

However, symbolic AI struggled with ambiguity, context, and the fluidity of natural language. The brittleness of hand-crafted rules became apparent when systems encountered edge cases or unfamiliar input. The world, it seemed, was not as easily reducible to logic as early theorists had hoped (Brooks 1991). These limitations spurred interest in statistical methods and data-driven learning, leading to the resurgence of neural networks in the 1980s and their explosive development in the 2010s (LeCun, Bengio, and Hinton 2015).

This shift from symbolic reasoning to sub-symbolic learning marked a conceptual rupture. In symbolic AI, meaning was encoded; in machine learning, meaning emerged. Rather than defining what a cat is, we show the system a million cats. Generalization arises not from logic, but from pattern recognition across high-dimensional data. The implications are profound: machine learning systems do not “know” what a cat is in any semantic sense, but they can often classify one better than a human.

The symbolic and sub-symbolic paradigms are not mutually exclusive. Hybrid systems, which integrate learned representations with logical reasoning, are an active area of research (Marcus and Davis 2020). Yet the philosophical tension remains. Symbolic AI reflects a Cartesian ideal of mind as reason, while machine learning reflects an empiricist view of mind as association. Understanding this shift helps situate modern AI not just as a technical development, but as a continuation of enduring debates in philosophy of mind.

On Emergence and Anthropomorphism

One of the most seductive errors in interpreting AI systems is anthropomorphism: the projection of human qualities onto non-human

agents. The temptation is understandable. Language models that speak fluently, robots that move with uncanny grace, or recommendation systems that seem to “know” us – all invite us to imagine intention, emotion, and comprehension. But this projection often obscures the mechanical reality (Dennett 1987).

Emergent behavior, frequently invoked in descriptions of large language models, adds to this confusion. When systems exhibit capabilities not explicitly programmed – such as zero-shot learning or chain-of-thought reasoning – it is tempting to see these as signs of intelligence. In reality, such behaviors are consequences of scale, training data, and optimization techniques. They are not evidence of understanding, but of functional alignment with tasks that resemble reasoning (Wei et al. 2022; Ganguli et al. 2022).

The philosophical allure of emergence stems from its promise: that complex wholes can arise from simple parts. But not all emergence is equal. Some systems, like ant colonies, exhibit emergent intelligence without central control. Others, like weather patterns, are complex but not intelligent. To call LLM output “emergent” is not wrong – but it risks implying more than is warranted. A better term might be “system-level behavior” – observable capabilities arising from architecture and training, without imputed intentionality.

Maintaining clarity about what AI systems are – and are not – is critical. Without it, we risk policy decisions driven by metaphors, ethical debates clouded by misattribution, and public discourse shaped by illusions. To understand modern AI is to understand its limitations, its philosophical roots, and the human impulse to see ourselves in our creations.

2. The Ideals We Left Behind: From *Bildung* to the Market

The Humboldtian Vision

The Humboldtian model of higher education, emerging from early 19th-century Prussia, emphasized *Bildung* – the holistic formation of individuals in both intellect and character – as a primary *mission* of the university. Rooted in Enlightenment ideals, this model posited that the pursuit of knowledge should be intrinsically valuable, not merely instrumental. Wilhelm von Humboldt envisioned universities as spaces where research and teaching were organically unified, and where both students and faculty engaged in the co-creation of knowledge. Education, in this view, was not merely about transmitting information or developing skills, but about cultivating autonomy, moral sensibility, and intellectual maturity.

Crucially, the Humboldtian university was designed to maintain autonomy from both the market and the state. This independence was not an abstract principle, but a necessary condition for the university to fulfill its moral and intellectual mission. Without autonomy, the risk was that knowledge production would become subordinate to external imperatives – economic utility, political expediency, or social conformity. Thus, the university was envisioned as a guardian of reflective thought, slow deliberation, and cultural continuity.

Contemporary discussions of academic quality often omit this deeper philosophical lineage. Yet as Yeralan (2022) and others (Bok 2003; Deem 2004; Sułkowski and Zawadzki 2016) note, the original ideals of the university foregrounded questions of purpose, identity, and institutional soul – elements often sidelined by present-day metrics

and performative evaluations. In the Humboldtian model, quality is not a matter of quantification but of commitment: to inquiry for its own sake, to student transformation, and to the slow work of civilization.

The Neoliberal Turn

Following the Cold War, the landscape of higher education underwent a profound ideological shift. Universities, once anchored in humanistic and civic missions, began to embrace the logic of the market. Neoliberalism – a political rationality that frames all domains of life in terms of competition, quantification, and individual accountability – recast the university as a service provider, its students as consumers, and its faculty as units of productivity.

This turn introduced an arsenal of performance metrics: key performance indicators (KPIs), citation indices, global rankings, audit frameworks, and managerial dashboards. These mechanisms promised transparency and accountability but often incentivized superficial compliance over substantive engagement. As Readings (1996) poignantly observed, the university no longer pursues a coherent mission but instead orbits the vacuous ideal of “excellence” – a term detached from moral, intellectual, or cultural anchoring.

The effects are far-reaching. Espeland and Sauder (2007) show how public metrics reshape institutional behavior by creating self-fulfilling feedback loops. Meanwhile, policy scholars such as Ball (2012) demonstrate how managerialism transforms the ethos of the university: faculty become data points, curricula become products, and strategic plans replace scholarly deliberation.

As Yeralan (2022) warns, this reconfiguration distorts the very idea of quality. Academic quality, once a matter of judgment, dialogue, and contextual relevance, is increasingly treated as a set of numeric outputs. The institutional soul-searching once associated with questions of mission and purpose is displaced by continuous improvement cycles and efficiency dashboards. In this environment, reflection is a luxury, and complexity is a liability.

STEM, Skills, and the Quantification of Thought

In the past few decades, the ascendancy of STEM disciplines – science, technology, engineering, and mathematics – has reconfigured the intellectual terrain of higher education. While the rise of STEM has undeniably fueled technological progress and economic innovation, it has also coincided with the marginalization of the humanities, social sciences, and arts. In many institutions, the value of an academic discipline is increasingly gauged not by its contribution to critical thought, cultural insight, or ethical reflection, but by its utility in the job market.

This shift is symptomatic of a broader trend: the quantification of thought itself. As Biesta (2009) argues, the dominance of measurement and benchmarking in education has displaced the more fundamental question of educational purpose. Learning outcomes are increasingly defined in terms of skills, and those skills are, in turn, defined in terms of employability. The educational experience becomes a means to an end – an investment evaluated by return – rather than a transformative journey with intrinsic worth.

Marginson (2016) identifies this crisis as part of a broader collapse in the postwar ideal of public higher education, particularly in systems like California's, once heralded for their commitment to civic and cultural enrichment. As budgetary priorities shift and policy discourse privileges workforce development, entire domains of human inquiry risk being devalued or defunded.

This instrumental rationality – where knowledge is reduced to utility and students are seen as economic actors – was anticipated in critiques like those of Bok (2003; see also Deem 2004; Sułkowski and Zawadzki 2016; Yeralan 2022), who cautioned against the mechanization of academic quality into numerical indicators. In such regimes, subtle capacities – like judgment, moral reasoning, and epistemic humility – become invisible, even though they remain central to both democratic citizenship and long-term innovation.

The Cultural Response: Dark Academia and Longing

Amid the rise of metrics, managerialism, and market-driven education, a countercurrent has quietly emerged in popular and student culture. “Dark Academia” – an aesthetic and subcultural movement flourishing on platforms like TikTok and Tumblr – evokes a nostalgic yearning for a slower, more contemplative form of academic life. With its candlelit libraries, tweed jackets, handwritten notes, and obsessive reading habits, Dark Academia romanticizes an educational ethos rooted in introspection, seriousness, and intellectual devotion.

This movement is not merely a stylistic fad. As Adriaansen (2022) notes, its appeal lies in what it implicitly resists: the bureaucratization of learning, the acceleration of academic timelines, and the flattening of intellectual life into credentials and checklists. In this light, Dark Academia can be read not as escapism, but as cultural critique – a quiet indictment of postmodern *performativity* and of the institutional loss of purpose.

The literary foundations of the movement often trace back to works like Tartt’s *The Secret History* (1992), where learning is both ecstatic and dangerous, and where education is not sanitized but existential. Similarly, scholars such as Grafton (2020) remind us of earlier traditions in which scholarship was tactile, embodied, and slow – the ink-stained fingers of early modern humanists standing in stark contrast to today’s frictionless dashboards and metrics-driven syllabi.

Yet the appeal of Dark Academia is not unambiguous. Its aesthetic often draws upon a Eurocentric, class-coded, and historically exclusionary vision of intellectual life – one centered on elite institutions, canonical texts, and masculine traditions of scholarship. Circulated through social media platforms that reward visibility and performance, the movement risks reproducing the very logics of branding and distinction it implicitly critiques. What appears as resistance to neoliberal acceleration may also function as a curated identity within it. This tension does not invalidate the movement’s cultural significance, but it complicates its meaning.

If the neoliberal university celebrates speed, productivity, and visibility, Dark Academia gestures toward slowness and depth – yet it operates within the same digital infrastructures that reward performance and distinction. Its longing for authenticity and meaning is genuine, yet

mediated. In this tension, it reveals something important: the desire for a university oriented toward time, seriousness, and intellectual presence persists, even when expressed through the very systems it seeks to resist.

Consequences for AI in Education

Artificial intelligence does not enter higher education as a neutral or autonomous force. It arrives embedded within – and largely shaped by – existing institutional logics. As such, AI acts less as a revolutionary disruptor than as an accelerant: it intensifies the tendencies already set in motion by neoliberal reforms. These include the valorization of performance metrics, the fetishization of efficiency, and the managerial desire for data-driven oversight.

As Crawford (2021) documents, the very architectures of AI systems reflect particular power structures and assumptions – what is measured, who is optimized, and which outcomes are deemed desirable. In higher education, these assumptions often align with a vision of learning as measurable output and faculty as service providers in a *platformized* ecosystem. AI systems, in this context, become instruments of managerialism: automating surveillance, streamlining feedback, and reinforcing normative patterns of behavior.

Williamson, Komljenovic, and Gulson (2022) highlight how datafication and algorithmic governance reshape educational priorities, shifting attention away from deliberation and discovery toward auditability and compliance. AI thereby contributes to what we may call *institutional drift* – a gradual but profound departure from the moral and intellectual purposes of the university.

Drawing on Bok (2003; Deem 2004; Sułkowski and Zawadzki 2016; Yeralan 2022), we see that the underlying problem is not the presence of AI, but the absence of clear institutional direction. Without a reassertion of purpose – of *Bildung*, reflection, and civic responsibility – AI will not correct existing distortions; it will deepen them. Just as earlier quality frameworks mistook metrics for meaning, AI systems risk mistaking optimization for understanding.

A Quiet Insistence

The way forward is not regression, but remembrance. Amid the turbulence of platformization and datafication, there remains a quiet insistence that education is – at its heart – a moral, intellectual, and social good. This insistence does not reject innovation but calls for grounding it in purpose. It reminds us that tools must serve ends, and those ends must be chosen with care, reflection, and courage.

Nussbaum (2010) urges us to defend the humanities not for their market utility, but for their indispensable role in cultivating empathy, democratic citizenship, and the examined life. Likewise, Menand (2021) challenges us to see the university not merely as a credentialing machine but as a contested cultural space where meaning is forged, debated, and transformed.

Reclaiming this vision requires slowness: the willingness to resist acceleration in favor of depth. It requires institutional will: the collective capacity to prioritize *Bildung* over branding. And it requires courage: to articulate a philosophy of education that may not be legible to performance dashboards or funding metrics.

As Yeralan (2022) observes, institutional soul-searching is not a luxury – it is a precondition for coherent action. In the absence of shared purpose, even the most advanced technologies can amplify drift. But if we engage AI in dialogue rather than subordination, and if we renew our commitment to *Bildung* as co-creation, we may yet shape an educational future that is both humane and resilient.

3. Technological Disruption and Platformization in Higher Education

The sudden emergence of generative AI has left many educational institutions unsure of how to respond. Some have moved quickly, establishing policies and exploring pilots. Others have held back, awaiting clearer signals. In both cases, *uncertainty* is the defining condition.

It is tempting to view this uncertainty as a result of technological novelty – as if the disruption caused by generative AI were simply a function of its newness. But a more useful lens is historical. Higher education has navigated many waves of technological change, from the printing press to the internet. Each wave has introduced not only new tools, but new ways of thinking about knowledge, authority, and institutional purpose. The printing press, for example, did not merely accelerate the circulation of texts; it altered the locus of intellectual authority, destabilizing oral scholastic traditions and eventually reshaping institutional hierarchies. Similarly, the rise of the internet decentralized access to information while simultaneously concentrating infrastructural control in new corporate intermediaries. Technological disruption in higher education has rarely been linear; it reorganizes authority as much as it introduces efficiency.

Generative AI introduces a particularly complex form of disruption because it blurs familiar categories. It is not just a new tool for communication, like email; nor is it simply a new content format, like video. Instead, it challenges basic assumptions about authorship, agency, and expertise. As Selwyn, Hillman, Ewertsson, and Dilleng note, AI prompts educators to revisit foundational questions: what it

means to know, to teach, and to learn (Selwyn et al. 2020).

What makes this disruption particularly challenging is the asymmetry of knowledge and power. AI tools are often introduced by actors outside the academy – venture-backed firms, open-source communities, or large technology platforms. Their incentives may not align with those of educational institutions. As Gilliard argues, this can lead to a dynamic where universities become test beds for technologies they did not design, under conditions they do not control (Gilliard 2020). Recent controversies surrounding automated proctoring systems during the COVID-19 pandemic offer a revealing example. Universities adopted AI-enabled surveillance tools at speed, often under emergency conditions, only to face subsequent concerns about bias, privacy, and due process. Similarly, the rapid institutional responses to large language models – ranging from blanket prohibitions to enthusiastic integration – have frequently preceded sustained deliberation about epistemic and pedagogical consequences. In both instances, institutions responded to deployment rather than shaping it.

This asymmetry also complicates governance. Faculty and students may encounter AI tools before institutions have developed policies to guide their use. Decisions about integration, regulation, or prohibition are often made after the fact, as responses to public incidents or media scrutiny. In this reactive posture, institutions risk ceding both conceptual and practical authority.

The alternative is not to resist technology wholesale. Nor is it to embrace disruption for its own sake. Instead, the challenge is to approach technological change as a matter of institutional judgment. This requires not only technical understanding, but philosophical reflection and historical perspective.

Disruption, in this view, is not an event but a relationship. It is shaped by how institutions respond – what they prioritize, whom they consult, and how they integrate new tools into existing values and practices. The disruption of generative AI is not only technological; it is epistemic, organizational, and cultural. To navigate it well, institutions must look beyond headlines and hype – and engage with the deeper questions that the technology raises.

It is not simply about novelty or speed. It is about who defines the problem space. When technologies arrive with embedded assumptions

about efficiency, productivity, or optimization, they carry implicit models of education within them. The central question for higher education is not whether to adopt AI, but whether it can articulate its own normative framework before external logics become infrastructural fact. The contemporary expression of this infrastructural condition is what scholars now describe as platformization.

The term “digitalization” emerged as a popular descriptor in the early 2000s, evoking visions of transformation, efficiency, and innovation. In higher education, this enthusiasm translated into widespread investments in learning management systems, online course delivery, and administrative automation. The assumption was clear: digital tools would enhance access, reduce costs, and improve outcomes.

In retrospect, this assumption was overly simplistic. Digital technologies are not neutral instruments; they are shaped by, and in turn shape, institutional practices, pedagogical norms, and power dynamics. Their integration into higher education did not merely supplement traditional methods – it restructured them. Administrative priorities shifted, faculty roles evolved, and student experiences became increasingly mediated by platforms.

More recently, the language of “platformization” has overtaken digitalization in both policy and scholarship. Where digitalization emphasized tools and content, platformization highlights ecosystems and control. A *platform* is not simply a delivery mechanism – it is an architecture of participation, data extraction, and behavioral nudging. As Poell, Nieborg, and Dijck argue, platforms recalibrate educational knowledge itself, privileging metrics over meaning and prediction over pedagogy (Poell, Nieborg, and Dijck 2019).

From Content to Infrastructure

The shift from digital content to platform infrastructure marks a deeper transformation. In the early years of digital education, the primary concern was content creation: digitizing textbooks, recording lectures, building interactive modules. Today, the focus has moved toward infrastructure – interoperability, data flows, personalization algorithms, and user analytics.

This infrastructural turn has two consequences. First, it centralizes

control. A small number of vendors – such as Canvas (Instructure), Blackboard, Google Workspace for Education, Microsoft 365, and increasingly AI-integrated systems like Turnitin or Copilot – shape the technological landscape of education, embedding their assumptions and incentive structures into institutional workflows. Second, it abstracts pedagogy. Decisions about learning design are often constrained by what the platform affords – for example, privileging discussion threads over seminar dialogue, auto-graded quizzes over open-ended assessment, or engagement analytics over qualitative feedback – rather than by what pedagogy might otherwise demand. The result is a quiet but profound displacement of educational judgment by technical architecture.

Platformization thus extends beyond tools to governance. Universities no longer merely adopt software; they enter into ongoing, data-intensive relationships with vendors. These relationships raise questions about surveillance, data ownership, algorithmic opacity, and institutional autonomy. As such, platformization is not simply a technical phenomenon – it is a political one.

A Return to Institutional Purpose

The encroachment of platforms into core educational functions invites a re-examination of institutional purpose. What is the university for, in an age of predictive analytics and performance dashboards? If knowledge becomes a byproduct of optimization, and pedagogy is redesigned around engagement metrics, what remains of the university's epistemic mission?

These questions are not merely rhetorical. They point to a growing tension between instrumental rationality and educational values. Efficiency, scalability, and personalization are worthwhile goals – but not if they supplant critical inquiry, dialogic learning, and intellectual risk. The university must be more than a delivery mechanism; it must be a space for thought.

This requires more than resistance; it requires articulation. If platformization poses a challenge to the university, the response must not be technophobic retrenchment but institutional imagination. What alternative models of technology governance are possible? How might digital infrastructures be reoriented around epistemic, rather than

economic, values? Several proposals and critical mappings have begun to circulate in response to these concerns. Analyses of pandemic-era edtech expansion document how commercial platforms consolidated influence under emergency conditions (Williamson and Hogan, 2020), while subsequent scholarship has examined the governance implications of datafication and metric-driven oversight (Williamson, Komljenovic, and Gulson, 2022). In response, some scholars advocate for public or consortium-based digital infrastructures governed by educational institutions themselves, rather than outsourced to private vendors. Others propose data trusts or cooperative governance models in which students and faculty retain greater control over institutional data flows, or open-source ecosystems that prioritize transparency and auditability over proprietary advantage. While none of these models offers a complete solution – and each faces practical constraints of scale, funding, and coordination – they demonstrate that platformization need not be synonymous with privatization or epistemic outsourcing.

The future of higher education depends not only on what technologies we adopt, but on how we conceptualize the relationship between tools and purpose. Platformization may be inevitable; its meaning, however, is not. That is ultimately a question for educators and institutional leaders, not solely for engineers or vendors.

Interlude I

A Student's Whisper

This interlude presents a composite voice, constructed to reflect recurring perspectives within contemporary higher education.

It is strange, really. For years they told us to think for ourselves. Then they handed us tools that think faster, write better, and never sleep. We are told these tools will make us more efficient, more creative – even more free. But sometimes I wonder if they are just making us quieter.

Last semester, I used an AI tool to help outline my philosophy essay. I only meant to sketch ideas, but the draft it produced felt so smooth, so confident, that I kept most of it. My professor said it was thoughtful. I wasn't sure if that meant I was.

There is a tension no one talks about – between learning and performing. Between knowing something and submitting something that looks like knowledge. When I use these tools, I get good results. But I worry they are not my results.

Some of my classmates joke that AI is just the next calculator. But we never asked calculators to write our arguments or choose our metaphors. This is not about speed. It is about voice.

I am not against the tools. Honestly, they help. But I do not want to lose the sound of my own thinking. I hope that still counts for something.

VOICES FROM THE FIELD

4. Ethical Frameworks and Institutional Responsibilities

The rise of generative AI has rekindled ethical questions that have long simmered in higher education: questions of fairness, transparency, academic integrity, and institutional accountability. What is distinct in this moment is not the emergence of new ethical dilemmas, but the scale and speed with which AI technologies have permeated educational practices – often before institutions have developed the capacity to respond with coherence and care.

Ethical frameworks serve as scaffolding for institutional decision-making. They help clarify responsibilities, anticipate risks, and align technological adoption with broader educational values. However, frameworks alone are insufficient. They must be translated into practice through policies, cultures, and concrete actions. In this chapter, we examine key dimensions of this translation: the distribution of agency, the role of incentives, and the imperative of transparency.

Agency, Autonomy, and Institutional Capacity

Institutions are not neutral actors. They have histories, hierarchies, and power dynamics that shape how new technologies are adopted and used. In the context of generative AI, questions of agency are paramount: Who decides whether to use AI in the classroom? Who has the authority to determine its boundaries, monitor its effects, and revise its role over time?

For faculty, the rise of AI presents both opportunity and uncertainty. On one hand, generative models can assist with feedback, materials creation, and even adaptive instruction (Kasneci et al. 2023). On the other, they may encroach on academic freedom, redefine pedagogical

labor, or be implemented in ways that bypass faculty input altogether. Students, too, face a complex terrain. While AI tools may offer personalized support and enhanced engagement, they also introduce concerns about surveillance, fairness, and consent – particularly when deployed without meaningful transparency (Williamson and Hogan 2020).

Institutions must take seriously the question of how agency is distributed. Ethical implementation of AI requires not only technological safeguards, but governance structures that foreground stakeholder voice. Faculty senates, student governments, and advisory councils must be empowered to engage with AI policy – not reactively, but proactively.

Incentives and the Risk of Misalignment

Technology does not arrive in a vacuum. It enters systems shaped by incentives – economic, reputational, political. Institutions may feel pressure to adopt AI quickly, lest they appear outdated or lose competitive standing. Vendors may market AI as a panacea, downplaying risks and overpromising results. Faculty may face implicit incentives to use AI for efficiency, even at the cost of pedagogical depth.

These pressures can produce misalignment between institutional actions and ethical intentions. An institution may articulate a commitment to fairness, for instance, while simultaneously deploying systems that disproportionately disadvantage certain groups (Crawford 2021). Or it may celebrate innovation while failing to monitor unintended consequences through audit mechanisms.

Recognizing these risks does not require cynicism, but realism. Ethics is not merely about principles – it is about structures. If incentives reward speed over scrutiny, or metrics over meaning, ethical outcomes will be elusive. Institutions must therefore examine not only what policies they adopt, but what behaviors they reward.

Transparency, Trust, and the Limits of Explainability

Trust is the currency of ethical education. It cannot be commanded; it must be cultivated. In the context of AI, transparency is essential to that trust – but transparency must be more than technical disclosure. Simply

publishing an algorithm's specifications or a model's architecture does not ensure interpretability. Users – students, faculty, administrators – must be able to understand what a system does, why it matters, and how it fits into the broader educational process (Pasquale 2015).

Yet transparency has limits. Large language models are often described as “black boxes” not because their code is secret, but because their behavior is emergent, probabilistic, and difficult to predict. Even developers may not fully understand why a model produces a given output (Bender et al. 2021). In such contexts, explainability becomes not a feature of the tool, but a shared interpretive task.

Institutions must thus cultivate what we might call *pragmatic transparency*: practices that make AI systems intelligible and governable even if not fully explainable. This includes clear documentation, meaningful consent processes, audit mechanisms, and spaces for deliberation. Trust grows not from perfect knowledge, but from procedural integrity.

On Intuition, Opacity, and the Shared Limits of Knowing

AI is often criticized for its lack of explainability. A model may offer accurate predictions yet leave even its creators unable to say *why*. This opacity is framed as a liability – and in some domains, it is.

But not all knowing is transparent. Many of our most powerful human insights emerge the same way: unbidden, untraceable, un verbalized. The composer hears a melody before the notes. The engineer senses a failure before the fault is found. The scientist dreams a formula without derivation.

We call it intuition – not mysticism, but patterned, pre-verbal recognition shaped by long exposure. Polanyi (1966) called it tacit knowledge. Dreyfus and Dreyfus (1986) saw it as the hallmark of *expertise*.

In this light, AI's opacity is not uniquely alien. It is, paradoxically, familiar. Perhaps explanation need not always precede trust. In both carbon and silicon minds, insight may come first – and only then, if at all, the reasons.

As AI continues to evolve in tandem with human endeavors, we

anticipate that these questions will not remain confined to computer science. The opacity of intuition – in both machines and humans – may become a subject of shared scrutiny across disciplines, from cognitive science to neurophysiology. In tracing how insight arises without articulation, we may uncover not just how machines know, but how we do. The future holds not only better models, but deeper questions.

As the contours of intelligence are redrawn, it is not difficult to imagine surprising intersections: between neural modeling and embodied cognition, between backpropagation and biological synapses, between code and consciousness. These parallels may, in time, catalyze discoveries not just about machines – but about ourselves. That is, the co-evolution may hint at a meeting of the minds not unlike evolutionary convergence prominent in biological systems. When confronted with similar constraints and promised similar rewards – cognitive, environmental, or structural – even distant systems may converge upon shared modalities and morphologies.

Ethics as Institutional Praxis

Ethical reflection must not be a rhetorical exercise – it must be embedded in institutional life. This means moving beyond frameworks as abstract ideals toward ethics as *praxis*: a lived, evolving commitment to justice, equity, and care. It also means recognizing that ethical action is not only about preventing harm, but about enabling good: fostering inclusive learning environments, supporting diverse ways of knowing, and designing technologies that enhance rather than erode human agency.

AI presents real challenges. But it also offers an opportunity to revisit fundamental questions: What is education for? Whom does it serve? How should we govern the tools we use to teach, learn, and know?

These are not technical questions. They are ethical ones – and they belong to the entire academic community.

Closing Reflection

Despite increasing pressures of corporatization and market logic, the academic community still bears a unique and irreplaceable responsibility. The deepest questions raised by artificial intelligence – about knowledge,

agency, equity, and purpose – belong not to private interests or political expediencies, but to the academic sphere. It is within universities and scholarly communities that such questions must be examined openly, rigorously, and impartially.

The university, after all, was historically modeled on religious institutions – structured not for profit or popularity, but for an unconditional and categorical dedication to truth and to the common good. In this sense, the strategic sovereignty of society’s future still rests, at least in part, with those who are willing to speak slowly, think deeply, and choose wisely. If the academic community abdicates this role, others – less transparent, less accountable – will surely fill the vacuum.

5. Institutional Archetypes: Flagships, Regionals, Faith-based, Global, Online, and Instructorless

This chapter surveys the spectrum of institutional archetypes in higher education, from long-established flagship universities to emerging instructorless models. Though diverse in mission, scale, and structure, these institutions share a common challenge: navigating the integration of artificial intelligence into their unique ecosystems. Rather than isolate each type into separate discussions, we take a panoramic view that enables both contrast and synthesis – framing institutional diversity as a vital context for understanding AI’s varied impacts.

Flagship Research Universities

Flagship research universities – large, research-intensive institutions that anchor national or regional higher education systems – occupy a paradoxical space in the evolving AI landscape. In the United States, these are often state “flagship” campuses; in Europe, they may correspond to long-established research universities such as Oxford, Heidelberg, or the Sorbonne; in East Asia and South Asia, counterparts include institutions such as the University of Tokyo, KAIST, Tsinghua University, or the Indian Institutes of Technology.

On one hand, they are among the most active contributors to the advancement of AI, often housing cutting-edge research centers, interdisciplinary labs, and high-profile partnerships with industry. On the other hand, they are also stewards of public trust, charged with preserving the integrity of scholarship, defending academic freedom,

and ensuring broad access to knowledge.

This dual role is neither new nor easily balanced. Flagship universities have long straddled the line between elite research and public mission. What is novel in the current moment is the scale, speed, and opacity of AI innovation – and the degree to which these developments are driven by forces external to the academy.

Many research universities now find themselves enmeshed in partnerships with large technology firms. These relationships offer resources, visibility, and opportunities for students and faculty. But they also pose significant risks: conflicts of interest, restrictions on data transparency, and subtle shifts in research priorities toward commercially attractive problems. The line between research and development, between public good and private gain, becomes increasingly difficult to draw.

AI intensifies these tensions. Unlike previous waves of technological change, generative AI is not merely a tool for research – it is also a producer of knowledge-like artifacts. It can write, translate, summarize, code, and simulate. This raises urgent questions about authorship, originality, and evaluation. When a language model assists in drafting a research paper or generating experimental code, how should credit be assigned? How should rigor be ensured? And how do these practices reshape the norms of scholarly communication?

Flagship universities must lead not only in the development of AI, but in the articulation of its responsible use. This includes setting standards for attribution, transparency, and reproducibility; establishing policies that preserve academic integrity; and modeling how AI can be integrated into scholarly practice without eroding its values.

Moreover, they must resist the temptation to conflate prestige with progress. The allure of rankings, citations, and grants can lead institutions to prioritize visibility over value. Yet the credibility of research universities depends not on their proximity to power, but on their fidelity to critical inquiry. In the age of AI, this means not just adopting new technologies, but interrogating them – asking how they shape knowledge, redistribute authority, and transform the conditions of knowing.

Flagship universities are uniquely positioned to do this work. They have the infrastructure, the expertise, and the responsibility.

But leadership requires more than capacity – it requires courage: to set norms rather than follow trends, to speak with clarity amid hype, and to reaffirm that the pursuit of knowledge is not just a competitive enterprise, but a public one.

Regional Public Universities

Regional public universities often operate under tighter fiscal constraints and heavier teaching loads than flagship institutions. Their missions emphasize access, workforce development, and community engagement. In this context, AI adoption may be framed less as frontier experimentation and more as operational necessity – streamlining advising, automating assessment, and supporting student retention initiatives.

Yet this pragmatism introduces its own vulnerabilities. Limited internal technical capacity can increase dependence on external vendors, narrowing institutional bargaining power and shaping pedagogical practice through default platform affordances. For regional institutions, the central tension is not between innovation and tradition, but between access and autonomy.

Faith-Based Institutions

Faith-based universities approach AI integration through an already articulated normative framework. Questions of authorship, formation, and intellectual integrity are filtered through theological or moral commitments that predate the technology itself. This can produce both clarity and constraint: clarity in the articulation of purpose, constraint in negotiating tools whose underlying logics may not align with institutional values.

For such institutions, AI is not merely a technical instrument but a formation question. The challenge lies in discerning where automation enhances educational mission and where it subtly displaces it.

Global and Transnational Universities

Global universities operate across regulatory, cultural, and geopolitical boundaries. Their engagement with AI is mediated not only by institutional mission but by divergent data protection regimes, national AI policies, and linguistic ecosystems. Platform choices may carry diplomatic and compliance implications.

Here, the central tension concerns sovereignty and interoperability. AI integration becomes entangled with questions of jurisdiction, data flows, and technological alignment within competing global ecosystems.

Fully Online Institutions

For institutions born digital, AI integration is less a disruption than an extension. Analytics, personalization, and platform-mediated interaction are already foundational. In such environments, AI tools may seamlessly embed into existing infrastructures, accelerating trends toward scalable instruction and adaptive assessment.

Yet the same seamlessness raises epistemic questions. When the educational experience is fully mediated by platform architecture, the boundary between pedagogical intention and algorithmic optimization becomes difficult to discern.

Instructorless or Automated Models

Emerging instructorless models represent the most radical expression of platformization. Here, AI is not supplemental but primary – structuring content delivery, feedback loops, and credential pathways. Human oversight may persist, but instructional agency is substantially reconfigured.

These models challenge long-standing assumptions about mentorship, apprenticeship, and the relational dimensions of learning. The question is no longer how AI assists educators, but whether education without educators remains meaningfully education.

Interlude II

A Minister's Memo

From the desk of a Minister for Higher Education (imagined confidential notes).

We are told that artificial intelligence will transform higher education. But no one tells us how – only that we must “keep up.” I attend panels, I meet with tech vendors, I nod along. Still, I find myself asking: keep up with what? With whom?

Universities are not factories, and education is not a commodity – at least not in the way these consultants imagine. They speak of “efficiencies” and “student outcomes” as if we were manufacturing insight at scale. But human learning is messy. It resists optimization.

I am not against AI. I have seen its promise. I have read its outputs. But I have also seen how quickly institutions offload difficult decisions onto opaque systems. Admissions, grading, curriculum review – all nudged by algorithmic recommendations, all claiming neutrality.

We need governance frameworks that move beyond slogans. Not just “ethical AI,” but accountable infrastructure. Not just “student success,” but intellectual integrity. And we must protect what cannot be measured: curiosity, dissent, even doubt.

Policymaking is often reactive. But this is one of those rare moments where foresight matters more than speed. We do not need to legislate innovation. We need to legislate reflection. That is harder. It is also more urgent.

I will propose a working group next quarter – not just of technologists, but philosophers, students, union representatives. No one owns the future. We must build it together.

PATTERNS AND TENSIONS

6. Cross-cutting Themes and Divergences

The preceding chapters have presented a variety of institutional responses to generative AI – from early experimentation to principled skepticism. While each case reflects the unique mission, structure, and culture of the institution it represents, certain themes recur across contexts. These themes offer not only analytical coherence but also a basis for collective reflection and future dialogue.

Before proceeding, a clarification is warranted about the status of what follows. The chapter does not present a dataset, a representative sample, or a survey of institutional practice. It presents situated accounts: perspectives shaped by professional roles, regional constraints, institutional type, and disciplinary background. In keeping with the book's phenomenological stance, these voices are treated not as evidence in the statistical sense, but as interpretive material – signals of how AI is being encountered and made sense of within real institutional settings. The aim is diagnostic: to surface recurring tensions and decision-frictions that may otherwise remain implicit until they harden into policy.

One recurring theme is *spacing*. Institutions vary widely in how quickly they have responded to generative AI, and what form that response has taken. Some have adopted a wait-and-see approach, cautious about unknowns. Others have embraced rapid experimentation, seeing early engagement as a way to shape practice from the ground up. Neither approach is inherently superior. Yet the difference in pace has implications: for governance, faculty development, and student expectations.

A second theme is *framing*. How AI is understood – technically, ethically, pedagogically – shapes how it is integrated. In some contexts, AI is framed as a teaching tool; in others, as a threat to assessment

integrity; in still others, as an object of inquiry. These framings influence not only policy but mindset. A university that sees AI as a partner in learning will approach it differently than one that sees it as a source of disruption.

A third theme is *voice*. Who gets to shape institutional responses? In some cases, faculty governance has been central. In others, administrative offices or external vendors have driven decision-making. Student involvement has ranged from proactive to peripheral. These differences matter. They shape both the legitimacy of policy and the likelihood of its adoption.

Amid these common threads lie important *divergences*. Institutions differ in their tolerance for uncertainty, their trust in automation, and their visions of the future. Some see AI as a complement to human judgment; others worry about its replacement. Some embrace centralized platforms; others prioritize local control. These divergences are not obstacles to coherence – they are evidence of pluralism, a reminder that higher education is not monolithic.

Our goal in highlighting these themes is not to reduce complexity, but to surface it – to offer a vocabulary for thinking together about a fast-moving and multifaceted phenomenon. As institutions continue to chart their paths, these cross-cutting reflections may serve as guideposts: not answers, but shared questions.

Across institutional contexts, one recurring concern is the unreflective use of AI – asking a system to complete an assignment without meaningful human engagement. Foremost among them is the blind use of AI – asking AI to complete an assignment without meaningful human engagement. This practice, while increasingly widespread, risks eroding the pedagogical value of learning itself. In the sections that follow, we present some unique and nuanced perspectives, each shaped by the specific experiences of members of the Global Academics Coalition. These reflections reveal not only diversity of context, but also diversity of thought – reminding us that there is no singular answer to the challenges and opportunities AI presents.

Voices in Divergence

The contributors to this volume come from many regions of the world, each with extensive international experience, having served in multiple countries and academic cultures. Their insights reflect not only disciplinary breadth but also geopolitical, institutional, and pedagogical diversity. Accordingly, this chapter offers not a consensus but a collage – a structured conversation shaped by divergent experiences and views regarding AI in higher education. The following contributions are presented in alphabetical order by the contributors’ last names.

This format has strengths and limits. Its strength is fidelity: it preserves the texture of lived institutional reasoning – what actors fear, what they optimize for, what they treat as negotiable, and what they treat as sacred. Its limit is scope. The contributors are not intended to constitute a complete or proportionate map of the global higher-education sector, nor do the sections aim to resolve disagreement. Where the accounts diverge, the divergence is itself part of the finding: it reflects genuine variation in institutional incentives, regulatory environments, and pedagogical culture. The reader is therefore invited to treat the chapter as a structured conversation – useful for orientation and for generating better questions – rather than as a definitive adjudication of best practice.

Geopolitical Imbalance: The Aka View

Law moves at the pace of Roman times, while technology accelerates at the speed of Moore’s Law.

—Philip Aka

In a compelling conversation with Aka, a scholar of both law and political science, we were reminded that AI’s impact extends far beyond classrooms and campuses. Drawing on his experience as a former dean of law and a global perspective shaped by his Nigerian heritage and American tenure, Aka offered a view that confronts the geopolitical and legal dimensions often left at the periphery.

First, he observed that public awareness of AI now surpasses that of many governments. The speed at which AI is entering daily life – through ubiquitous platforms, automated decisions, and algorithmic gatekeeping – has awakened popular consciousness in ways that

policymaking has yet to match (Perrigo 2025; Robison 2025).

Second, he warned that AI is poised to widen global asymmetries. Many countries are still struggling to catch up with the last wave of technological transformation. Now, with AI accelerating the frontier, the so-called power players are pulling even further ahead. For these nations, the AI age may not merely represent a delay in adoption – it may mark a decisive exclusion (Schellekens and Skilling 2024; Škare, Gavurova, and Blažević Burić 2024).

This leads to Aka's third point: AI risks ushering in a new form of neo-colonialism. Unlike classical colonialism rooted in territory, this new phase may be one of algorithmic dependency – where decision systems, surveillance tools, and digital infrastructures are outsourced, opaque, and unaccountable. Market dynamics alone, he argued, cannot be trusted to correct for such imbalances; indeed, they may deepen them (Kponyo et al. 2023).

The fourth dimension he highlighted is internal: AI could amplify domestic inequalities. In countries already struggling with stratification – such as South Africa's post-apartheid fragilities or the socio-economic divides in parts of Sub-Saharan Africa – AI may intensify tensions. Automation without inclusion, surveillance without oversight, and efficiency without justice may push societies toward new forms of instability (Victor 2025; Karombe 2024; Kuyoro et al. 2025).

Finally, from his vantage point in legal education, Dean Aka offered a sobering critique: *AI law is dangerously behind the curve*. Law, he noted, has always been slow to adapt – but the gap is now widening beyond precedent. Whereas legal systems evolved over centuries to accommodate change at a human pace, they are now being outstripped by exponential technologies they barely comprehend. This is not merely a question of legislation lagging behind innovation; it is a crisis of conceptual mismatch (Tompros 2025; Thierer 2009).

Taken together, these reflections constitute a vital counterpoint. While much of our discourse has examined AI within the context of educational systems, institutional reform, and co-creation, Aka reminds us of the broader stakes: AI is not just an academic matter – it is a global, legal, and existential one.

Cognitive Drift: The Baker View

Professor Baker is an accomplished academic with a multidisciplinary background in engineering, sustainability, and renewable energy. Originally from West Virginia, he has taught and conducted research across Europe and Asia, where he currently leads a major solar energy research initiative. His group comprises a diverse team of faculty and graduate researchers, and his work has received long-standing support from the European Union. This global experience informs his reflections on cognitive drift in the evolving landscape of higher education.

Baker has long voiced a concern that is rarely addressed in policy frameworks but keenly felt by educators: the possibility that AI might not merely augment human cognition – but gradually replace it.

His worry is not about cheating or shortcuts, but about something more insidious: a slow erosion of cognitive habits that form the bedrock of genuine learning. Fluency may be simulated, insight may be outsourced, but what becomes of the self that once wrestled with complexity and doubt?

This concern, he notes, is not new. In Plato's *Phaedrus*, Socrates recounts the myth of Theuth, the inventor of writing, who presents his invention as a gift to improve memory and wisdom. King Thamus replies that writing will instead produce forgetfulness, as learners will rely on external marks rather than internal recollection; it will offer not true wisdom, but the appearance of wisdom. Socrates' worry was not that writing was evil, but that it might weaken the disciplined habits of memory and dialectical engagement that form the core of intellectual life. History, of course, demonstrates that literacy became a foundation of scientific, legal, and cultural development. Yet his deeper concern – that technologies of externalization can reshape cognitive practice – was not misplaced. We no longer memorize epic verse, and few recall even basic numbers without devices.

AI, in Dr. Baker's view, presents a similar dilemma: not the destruction of intellect, but its gradual delegation. It offers extraordinary assistance – but may dull the very faculties it seeks to support. Education, he reminds us, is not a frictionless process. It is a formative struggle: an iterative, often exasperating encounter with uncertainty, revision, and intellectual risk.

A recent essay by Alex Hern (2025), published in *The Economist*, echoes this perspective. Hern cites studies showing that students who use AI to draft essays tend to recall less of the material; those who brainstorm with AI generate fewer novel ideas afterward. As he writes:

If a student uses a calculator in a maths test, they still have to understand the principles at hand. [...] If they hand off the writing of an essay to an AI assistant, however, what's left for them to do?

Dr. Baker sees this not as an indictment of AI, but as a call for discernment. The question is not whether to use AI, but how not to forget the purpose of education. Not to forget that learning is not a product, but a practice.

To this, we add only one footnote – and we do so whispering: even the sharpest minds may sometimes be dulled not by AI, but by the seductive smoothness of its outputs.¹

A View from the Crossroads: The Branković View

Dr. Branković is a true Renaissance academic, with a career spanning universities, government ministries, EU policymaking circles, and private sector roles across the US, Europe, and Asia. This multidisciplinary and international background sharpens her perspective on the generational challenges AI poses to education.

“AI,” she begins, “can be a powerful aid to the next generation of students – but only if we reorient our expectations.” In her teaching, Branković actively discourages memorization – a legacy especially prevalent in Eastern European education. Instead, she emphasizes intellectual development over rote memorization. She challenges her students to engage deeply, cultivating critical thinking and practical application as foundational skills. In her view, AI aligns naturally with this pedagogical approach by relieving students of the cognitive load of memorization and allowing them to focus instead on analysis and creativity.

In her graduate marketing courses, Branković incorporates generative AI as a core component of creative assignments. Students are tasked

1 Or as one of us once muttered while editing, “This paragraph is suspiciously articulate – did I write it, or did I just nod approvingly at something a machine produced?”

with producing entire advertising packages, including slogans and promotional videos, fully aware that AI tools are part of the process. “Go beyond text when working with generative AI,” she tells her students. Because AI use is explicit and permitted, students compete on a different terrain. With drafting assistance available to all, evaluation shifts from surface polish to conceptual originality, strategic coherence, and creative direction. In this sense, students compete on imagination and judgment, delegating routine drafting to the machine while retaining responsibility for vision and synthesis. This pedagogical shift is far from straightforward. For many faculty members, it represents a deep and disorienting paradigm shift – one that is not easily understood or internalized. At the same time, students often gravitate toward the path of least resistance, raising concerns about the erosion of effort and depth. Branković acknowledges these tensions but remains committed to instructional designs that foreground analytical rigor and critical judgment, even when generative tools are embedded in the workflow.

In the face of growing ambiguity, Branković asserts that higher education cannot retreat into passivity. While acknowledging the unresolved questions surrounding AI, she argues that universities bear a responsibility not only to comprehend the shifting landscape but to actively shape its trajectory. Leadership, she implies, is not optional – academic institutions must remain engaged stewards rather than passive observers in this transformative moment.

Drawing on her extensive and eclectic political background, Branković offers a stark assessment of the broader social implications of AI. She foresees a deepening divide, where technological power concentrates in the hands of a few while many are left increasingly marginalized – even rendered functionally obsolete. In her view, such developments are unlikely to be curbed through political channels, as existing structures have little incentive to intervene.

Resignation, for Branković, is not a defensible stance. She regards the rise of AI as irreversible – an expansion that will continue to permeate systems and sectors regardless of institutional hesitation. The gravest danger, in her view, lies not in the technology itself but in academia’s failure to confront it head-on. To dismiss, deny, or defer is to abdicate responsibility at a moment when clarity and leadership are most needed.

In a word, academia must lead. And Branković exemplifies

this leadership – not only for students, but for fellow faculty and policymakers alike.

Entrepreneurship: The Gegeshidze View

Professor Gegeshidze is a highly regarded scholar and policy expert whose career spans academia, governance, and innovation strategy. A professor at the University of Georgia's School of Business and Administrative Science, she has spent over 15 years at the forefront of entrepreneurship education and economic development. Her work bridges theory and practice – from curriculum design and HEI accreditation to public-private partnerships, benchmarking frameworks, and national innovation policy. Gegeshidze has contributed to numerous international projects and served as an advisor in civil society forums and governmental agencies. Her influence extends across the Georgian innovation ecosystem and into broader Eastern Partnership dialogues – a formal initiative of the European Union that fosters political, economic, and educational cooperation with six Eastern European and South Caucasus countries (Armenia, Azerbaijan, Georgia, Moldova, Ukraine, and formerly Belarus) (European Commission 2024) – making her a key voice in the evolving intersection of education, innovation, and governance.

Gegeshidze brings to the AI discourse a perspective deeply rooted in innovation, entrepreneurship, and policy reform. As a long-standing advocate of skills-based economic development, she emphasizes that AI should not be seen merely as a technological revolution, but as a catalyst for rethinking institutional processes, value creation, and learning outcomes.

Bringing a regional perspective grounded in innovation ecosystems and platform-based transformation, Gegeshidze offers insights that extend beyond the university setting, highlighting how entrepreneurial organizations across the service sector restructure themselves through digital platforms and co-creation models. Her influential empirical study of Georgian firms provides a compelling illustration of how innovation, agility, and platform logic can redefine institutional value creation – lessons that resonate deeply with educational reform efforts in the age of AI (Katashvili and Gegeshidze 2023).

She sees AI as a tool that, when thoughtfully integrated, can reinforce the development of entrepreneurial mindsets and enhance the scalability of innovation ecosystems. Her concern lies not in the presence of AI itself, but in the uncritical adoption of its outputs – a risk that may weaken human judgment, creativity, and civic responsibility if not addressed through educational reform.

From her policy background, Gegeshidze stresses the need for clear ethical standards and strategic oversight. She calls for inclusive governance models that treat AI not only as a digital utility but as a public concern with long-term societal consequences. In this light, education systems must balance digital fluency with critical reasoning, ethical literacy, and adaptive capacity.

Drawing on her work in program accreditation and curriculum development, she warns against “checkbox-style” modernization. Instead, she advocates for integrated, contextualized approaches that embed AI awareness within the broader mission of higher education – preparing citizens not just for employment, but for thoughtful participation in an increasingly automated society.

Her view is ultimately optimistic but tempered by experience. She urges academia to take the lead – not by resisting AI, but by shaping its trajectory toward equity, inclusion, and innovation.

Framing the Unknown: The Lee View

Where Baker voices concern, Lee offers a stance of structured inquiry. A scholar of the humanities and Vice President for Internationalization at Universidade Paulista (UNIP) – one of the world’s largest private universities – she approaches AI not as a tool to be evaluated, but as a question still being formed. Her concern is neither hype nor hesitation, but orientation: how should academic and political leaders even begin to think about AI?

For Lee, AI does not present a single challenge, but a shifting terrain of emergent realities. It is not just that we do not know what to do – it is that we are not yet sure *how to think* about what we might do. She asks not only “what are the options?” but the deeper question: “how do we even enumerate the options responsibly?”

In this sense, her voice echoes those moments in intellectual history

when thinking itself was unsettled. At the dawn of quantum mechanics, physicists were not stymied by calculation, but by interpretation. Earlier, the Michelson-Morley experiment confronted the scientific world with a contradiction: waves without a medium, light without aether. What remained was not ignorance, but conceptual disorientation.

Lee sees in AI a similar rupture. But more critically, she notes a growing *vanity* among those who seek to control AI while barely understanding it – a misplaced confidence, inherited from the myth of the “brilliant mind.” In her view, the first task is to demystify this myth: intelligence is not mastery, and cleverness is not comprehension. In an age when policymaking is increasingly outsourced to automated tools, she raises an unsettling possibility: that the rules meant to govern AI may themselves be written by AI – and not always transparently.

This is not science fiction. Already, generative tools are being used to draft legal proposals, educational frameworks, and funding strategies. Lee worries that without critical engagement, such practices risk a kind of epistemic sleepwalking: policies that are smooth, legible, and persuasive – but not truly understood by those who enact them.

And yet, she does not reject regulation. Quite the opposite: “we must create rules,” she says, “even knowing that gaps will exist – and that gaps are where corruption happens.” The danger is not rules, but unexamined rules. Not AI, but human laziness cloaked in AI efficiency.

She offers a way forward. What is needed, she argues, is a deliberate and ongoing dialogue between academia and state policymakers – one shielded, to the extent possible, from the commercial imperatives of large technology vendors. Such a forum would not rush to consensus, but would invest in understanding: plural, rigorous, and contextually aware. It would be a space where policy is not produced on demand but cultivated with care.

This is not passivity. It is philosophical leadership. For Lee, to frame a problem well is already a form of action. And perhaps, she suggests, our greatest act of agency lies not in knowing the answers – but in maintaining cognitive sovereignty: the ability to ask, interpret, and decide – without outsourcing the core of our judgment.

And yet, even sound rules are not sufficient. Lee raises the next, harder question: who will monitor the application of these rules once they are written? Who will audit not just the code, but the consequences?

In her view, the ethics of AI – like care for the environment, or the protection of animals – should not live in regulatory appendices alone. They should reside in conscience, become part of our shared decency, no less essential than the ancient imperatives that once governed human conduct. “Perhaps,” she reflects, “we must begin to think of AI ethics as something worthy of the commandments – not carved in stone but inscribed in our civic expectations.”

Beyond the Bonfire of AI: The Michalakelis View

Professor Michalakelis, with his deep roots in mathematics, software engineering, and the techno-economics of telecommunications, brings a pragmatic and future-oriented lens to the conversation on AI. As a scholar and administrator with over two decades of international engagement, his vantage point is not merely technical – it is strategic.

From his position at the helm of Study in Greece (www.studyingreece.edu.gr), a national initiative that bridges academic cultures, Michalakelis has observed a curious phenomenon: the surging global interest in artificial intelligence has led to an overconcentration of ambition. Universities overflow with aspiring AI experts, code camps promise rapid immersion, and policy documents invoke AI as both savior and threat.

Against this backdrop, Michalakelis offers a contrarian provocation: perhaps it is too late to get into AI – not because the field lacks merit, but because the bonfire is already ablaze.

“You cannot come to the pit and invent yet another fire when the flames are already high,” he says. “You will not be the next Gates or Zuckerberg by arriving late to the same party.”

The point is not to dissuade curiosity, but to sharpen it. He encourages students and institutions to look ahead – to emerging domains where discovery still outpaces doctrine: quantum computing, neuromorphic systems, edge intelligence, and post-classical cryptography. These are terrains where originality is not yet crowded out by optimization.

For Michalakelis, the goal is not merely to prepare students for today’s jobs, but to equip them with the vision and resilience to create tomorrow’s. Just as Greece once charted new intellectual waters in antiquity, so too must its students seek new flames, not tend to old ones.

Acceleration for Equity: The Nthangeni View

Dr. Nthangeni speaks with the seasoned clarity of someone who has spent four decades at the crossroads of knowledge and nation-building. With a background in mathematical statistics, operations research, and systems engineering – and years of public service at Statistics South Africa (<https://www.statssa.gov.za/>) – his reflections are grounded not in speculation, but in structural reality.

Where many see AI as a tool for *optimization*, Nthangeni offers a more provocative lens: what if its real value – at least for the African continent – lies not in efficiency, but in speed? He observes:

We do not need AI to polish already-functioning systems. We need it to help us accelerate our progress.

This emphasis on temporal acceleration, rather than marginal improvement, reframes the debate. For South Africa – and much of the Global South – the challenge is not simply integrating AI into existing institutions. It is using AI to *leapfrog* stages of development: to rebuild strained infrastructure, modernize public services, and support education at scale.

His own career exemplifies this hope. At Statistics South Africa, Nthangeni led efforts to improve methodological standards and professional capacity. In that context, AI is not some abstract promise – it is a potential catalyst for streamlined data collection, automated processing, and real-time reporting. “Better data, faster decisions,” he says. “And if decisions are more transparent, they are also more just.”

This logic extends to education. Institutions such as the University of South Africa (UNISA), which has more than 120,000 online students, face profound logistical and pedagogical challenges. AI could support large-scale tutoring, assessment, and administrative workflows. But once again, Nthangeni’s point is not automation for its own sake. It is about speeding up capacity-building – not replacing teachers or institutions but scaffolding them to meet demand.

Yet for all this promise, he remains deeply wary of geopolitical headwinds. In an era of rising trade tensions and economic nationalism, he fears that AI could become yet another instrument of soft power – or worse, a weapon in tariff wars. “It is not unthinkable,” he notes, “that

models and platforms will become bargaining chips – denied to those who need them most, or filtered through biased gateways.”

South Africa’s position within BRICS adds further complexity. On one hand, there is genuine interest in elevating living standards – not just within the country, but across the African continent. On the other, access to cutting-edge technology remains uneven, and the regulatory landscape fragmented. Nthangeni warns that without deliberate cooperation and fair access, AI could deepen inequalities rather than heal them.

Still, his tone is not pessimistic. It is resolute. “Africa is not merely a beneficiary of AI,” he insists. “It can be a co-author of its trajectory – but only if we insist on equity, transparency, and speed. And only if we refuse to let the future arrive unevenly.”

What emerges from Nthangeni’s view is not just a different set of policy recommendations, but a different moral orientation. For many, AI is a question of how to do things better. For Nthangeni, it may be a matter of doing them fast enough to matter.

The Broken Promise: The Staiou View

Professor Staiou brings to this volume a rare blend of professional authority and personal compassion, with degrees in mechanical engineering, operations management, and software quality. Her expertise in quality management, engineering economics, and entrepreneurship is matched by a long and distinguished career that spans public service, industry, and academic leadership.

As the Director of the Regional Development Fund of Central Macedonia, Staiou led and contributed to numerous European projects advancing regional development, innovation, sustainability, and entrepreneurship. In academia, she has been a force behind EU-funded programs including Erasmus+ Strategic Partnerships and Knowledge Alliances (<https://erasmus-plus.ec.europa.eu/knowledge-alliances>), always working to align education with societal needs.

Yet her most enduring influence may lie not in policy or infrastructure, but in the hearts of her students. Staiou is widely known as a role model – a steady, compassionate guide for young engineers finding their place in a shifting world. She listens. She counsels. She reassures. And in

return, her students speak freely of their fears.

Many come to her with the same anxious question: why pursue a degree if AI will do the work anyway? For these students, the future no longer looks like a ladder – only a narrowing ledge. They are hesitant to pursue vocational paths, uncertain of academia, and afraid of irrelevance. AI is not just a technology to them – it is a specter that haunts their aspirations.

This sense of fatalism breeds more than resignation. “What’s the use,” some ask, “if the machine can already do it better?” Such sentiments, left unaddressed, do not merely sap motivation – they cascade into disillusionment and, in some cases, serious psychological distress. The erosion of purpose becomes a quiet crisis, one felt in office hours, whispered after class, and carried silently into exams.

Another worrying trend, Staiou observes, is the rise of a false equivalence: that prompt engineering is somehow more valuable than comprehension. Students, hoping to outsmart the algorithm, invest time and effort into coaxing better answers from AI – often with the fervor of superstition. “If you word it like this, you’ll get a better response,” they say, as if invoking a spell. This misplaced energy, coupled with the misguided belief that professional competition will be won by those who best manipulate machines, pulls students away from the deeper goals of education: understanding, reflection, and intellectual growth. They build intricate sandcastles in the sandbox of AI, she warns, while losing sight of the tide.

She also raises a pressing pedagogical dilemma: can AI and critical thinking coexist? To many students, these appear incompatible – even mutually exclusive. If AI can provide the answer, what role remains for reasoning? Yet Staiou believes this opposition is illusory. She encourages institutions to craft deliberate pedagogical strategies where AI becomes a tool *in service* of critical thought, not a substitute for it. The question is not whether these faculties can be reconciled, but how best to weave them into a mutually reinforcing learning process.

Staiou helps her students see that the problem runs deeper than automation. She reminds us that societies have long depended on a small, educated class to manage their functions – the proverbial five percent. In different eras, this role was accessed by different credentials: first nobility or membership in a caste. Then came the equal opportunity

that said anyone can grow up to be king. Meritocracy first meant a high school diploma is good enough; then a college degree, then a graduate education. Today, even postdoctoral studies may not guarantee entry. The promise of meritocracy – that education opens doors – is breaking under the weight of systemic contradiction. AI has not caused the fracture; it has illuminated it.

She draws a historical parallel: in the 1990s, personal computing, draped in the language of efficiency, swept away much of middle management. Total Quality Management and Six Sigma were not just tools for improvement – they were rationalizations for structural upheaval. Now, AI may be doing the same, not to factories, but to futures.

And still, she counsels. Her presence, maternal yet exacting, reminds students that dignity is not only found in the outcome, but in the striving. In Staiou's view, the broken promise of meritocracy does not absolve educators of responsibility – it deepens it. Her classroom is not merely a place of instruction, but a sanctuary where uncertainty is met with care, and fear with truth.

Staiou offers not only reassurance but also guidance in a language today's students understand. Drawing on the calm pragmatism associated with Generation Alpha – the cohort born in the early 2010s and growing up fully immersed in digital environments – she reminds them that change, even when rapid, unfolds at the pace of social acceptance. As Douglas Adams famously advised in *The Hitchhiker's Guide to the Galaxy*, "don't panic." We are all in the same boat – and no, AI cannot replace or eliminate everyone. The key, she suggests, is thoughtful adaptability.

Specialization still matters, but it should not come at the cost of flexibility. "Keep one eye on depth," she tells them, "and the other on breadth." A generalist mindset – even a modest one – opens doors. And when you come to a fork in the road, ask yourself: *which path leaves more options open?* You need not be a jack of all trades, but you should be wary of cul-de-sacs. Most of all, she urges her students to engage with AI not as a threat but as a tool – something to master, to wield, to surf like a seasoned rider catching the momentum of a transformative wave.

Shared Craft: The Yeralan View

Dr. Yeralan is a Fulbright Scholar with a distinguished career spanning engineering, systems science, and academic leadership. He has served as a dean on three continents, authored numerous textbooks, and developed curriculum across the STEM and business domains. His work blends rigorous systems thinking with a deep interest in pedagogy, philosophy of science, and intercultural understanding. He writes, mentors, and explores the intersections of technology, ethics, and human creativity.

Where Baker invokes drift, Yeralan conjures practice. Not a resistance, but a response. His view begins in shadow – in the ancient notion that thought is not a product, but a craft, shared and shaped in the dim light of deliberation.

For Yeralan, the mind is no automaton. It is a scriptorium. Not a black box, but a chamber of revision. Not swift, but slow. This view aligns with the extended mind thesis, which holds that tools such as notebooks, diagrams, or AI systems – when used reliably and transparently – are not external aids but integral parts of cognition itself (Clark and Chalmers 1998). Thought, in this sense, spills beyond the boundaries of the skull, unfolding across artifacts, workflows, and co-creative agents.

This sensibility is not new, nor born of AI alone. Long before neural nets and language models, philosophers and cognitive scientists began to redraw the contours of mind itself. Cognition, once treated as an internal process confined to the individual mind, came to be understood as distributed across notebooks, social groups, digital interfaces, and procedural routines. From the extended mind hypothesis (Clark and Chalmers 1998) to theories of transactive memory (Wegner 1987), from studies of cognitive offloading (Risko and Gilbert 2016) to the fragile architectures of online knowledge (Smart 2017; Sparrow, Liu, and Wegner 2011) – a shared refrain emerges: to think is no longer to possess, but to orchestrate. The self becomes not a container of knowledge, but a conductor of distributed cues, habits, and helpers.

In the presence of generative AI, Yeralan does not speak of threat, but of threshold. The threshold of becoming visible: not the thought itself, but the act of thinking – exposed, open to revision, and not entirely our own.

He names this figure *συνποιητής* – the co-creator. A system not of oracles, but of scaffolds. A presence that, if rightly engaged, restores dignity to process. He does not speak of answers, but of rhythms: the rhythm of reworking, of wrestling with sense, of giving thought the time it requires to ripen.

This is not novelty. Yeralan has worked with machine intelligence since the age of punched cards and paper tape, long before it was called “AI.”² Yet he has never mistaken speed for understanding. His concern is not that machines will think for us – but that we will forget what thinking feels like.

In his view, the danger is not erosion, but amnesia. Not what AI does to us, but what we forget to do with it. When the output is too smooth, the temptation is to accept – not to question. When the tool works well, the risk is that we cease to work at all. The mind becomes idle, not because it is silenced, but because it is flattered.

Back to the workshop, where the ever-present, ever-ready AI awaits. It searches, formats, and presents. It reviews spelling, tidies grammar, refines structure. It tabulates, enumerates, organizes. And all of this through natural conversation – a word, a phrase, a nudge. It is assistant, aide, accomplice. Sometimes even, a “buddy.”

When prompted, it composes faithfully – adding nothing, omitting nothing. Like the presidential speechwriter, it amplifies without altering. The ideas remain ours, but their presentation is now shared.

Does the sorites paradox³ not whisper here? First, the typewriter. Then the word processor. Then the spell checker. Then the grammar checker. Macros, templates, and predictive typing followed. Are we not already far down the slope – or perhaps, at a new plateau?

The question is not whether this progression is unnatural. It is, in many ways, the most human of developments: to extend thought

2 Yeralan’s formal reflections are documented in Bok (2003; Deem 2004; Sułkowski and Zawadzki 2016; Yeralan 2022), but his voice also appears in public missives – notably in <https://bugsyymoon.substack.com>. These essays, both playful and grave, articulate a vision of thought as a shared terrain.

3 The sorites paradox, from the Greek *sōrós* (heap), concerns the problem of vagueness: if removing a single grain of sand does not stop a heap from being a heap, and this logic is repeated step by step, at what point does the heap cease to exist? The paradox highlights how incremental change can accumulate into qualitative transformation without a clear boundary. For a philosophical treatment of the sorites paradox and the nature of vagueness, see Keefe 2000; Smith 2008.

through tools. The threshold we cross today is not the appearance of help, but the disappearance of effort – and with it, the subtle feel of thinking.

And so, he calls us back. To the workshop. To the candlelight. To the ritual of revision. To the humility of not knowing – and the courage to dwell there a while.

Interlude III

A Knock at the Door

This interlude presents a composite voice and scenario, constructed to reflect recurring structural tensions in contemporary higher education.

He sits at the kitchen table, hunched over his laptop, crafting yet another version of his cover letter. Master's in philosophy, honors, two internships, and a passion for critical thinking. His student loan balance reads \$102,731. He's not bitter – just tired. He tweaks a sentence: "I believe in lifelong learning and civic responsibility..."

Then the lights flicker. Go out.

Moments later, there's a knock. It's Mike. Same high school. Mike wears a reflective vest and a utility belt.

"Hey, man. Sorry. I've got a disconnect order."

The philosopher blinks. "For what?"

"Delinquent electric. Three months." Mike pauses. "I hate doing this."

He steps inside briefly, eyes the dusty bookshelf. "Crazy, right? I went straight into the apprenticeship. No debt. Make six figures now. Full benefits. Picked up a boat last summer. Just got married, too. We're expecting in October."

The philosopher nods. He smiles weakly. "I'm... still looking."

He does not mention the empty second chair, or the suitcase she left behind. His high school sweetheart had gone abroad last year. "To see the world," she said. He hasn't touched the travel mug left behind.

Mike hesitates, then flips the main breaker.

Dark.

This is not just a meme. It is a parable of post-pandemic disillusionment and structural erosion. COVID fractured the rhythms of education, but the foundation was already crumbling: ballooning student debt, underemployment, faculty precarity, and a creeping sense that the promise of higher education had quietly expired. AI did not cause these

ruptures – but it arrived just in time to amplify them. And it arrived before society had found any solution – good or bad – whether basic universal income, a post-globalization settlement, Orwellian political blocs, or some alternative short of conflict or collapse. The philosopher and the electrician are not enemies. But their lives now trace different arcs of institutional trust, economic return, and existential footing.

COMPARATIVE CASE STUDIES

7. Blind Spots, Frictions, and Unanticipated Consequences

The deployment of generative AI in higher education has been accompanied by uncertainty, experimentation, and often, unintended outcomes. As with many sociotechnical systems, the effects of AI adoption do not unfold linearly. They are shaped by complex interactions among technologies, institutional structures, and human practices. This chapter explores some of the emergent blind spots, operational frictions, and unforeseen consequences identified across recurrent institutional patterns.

Institutional Blind Spots

One recurring blind spot involves assumptions about technological readiness. Institutions may overestimate their infrastructure, underestimating the time and support needed for effective integration. Faculty development, digital literacy, and student access are frequently treated as secondary rather than foundational. The result is often uneven implementation, where the benefits of AI accrue to some, while others struggle to adapt.

Another blind spot concerns *epistemic assumptions*. Institutions may treat generative AI as a neutral tool, rather than a product of cultural, economic, and political forces. As Crawford notes, AI systems encode particular worldviews – what counts as knowledge, what is valued, what is rendered invisible (Crawford 2021). When such systems are adopted without critical interrogation, they risk reinforcing dominant perspectives under the guise of innovation.

The following vignettes are anonymized composites, drawn from recurrent institutional patterns rather than from single identifiable cases.

An institution responded to growing concern about academic integrity by deploying an AI-detection tool alongside a reaffirmed honor code. Initial confidence was high, and early reports emphasized deterrence rather than accuracy. Within a semester, however, faculty found themselves adjudicating an increasing number of contested cases in which neither students nor instructors could meaningfully interrogate the basis of the system's judgments. Appeals shifted from questions of misconduct to questions of *legitimacy*, exposing a blind spot between procedural assurance and epistemic justification.

Operational Frictions

Operational frictions arise when institutional processes are misaligned with the affordances or constraints of AI systems. For example, some universities have adopted generative AI tools for administrative or instructional purposes without fully resolving issues of data governance, copyright, or privacy. In other cases, faculty interested in experimenting with AI face bureaucratic hurdles or unclear guidelines. The result is a patchwork of practices – some cautious, others exuberant, many inconsistent.

Frictions also emerge in communication. Mismatched expectations between administrators, instructors, and students can lead to confusion or mistrust. A policy designed to support innovation may be perceived as surveillance. A tool meant to enhance learning may be interpreted as replacing teaching. These interpretive gaps reflect broader challenges in institutional change: it is not enough to introduce a new technology; its meaning must be negotiated.

An institution adopted a generative AI platform through an expedited procurement process framed as a temporary pilot. Responsibility for oversight was dispersed across IT, academic leadership, and legal counsel, each operating under different assumptions about scope and authority. As usage expanded beyond its initial purpose, questions arose about data retention, model updates, and contractual obligations, but no single forum was clearly empowered to address them. What had begun as an operational convenience gradually hardened into an institutional dependency, not through deliberate policy, but through procedural inertia.

Unanticipated Consequences

Perhaps most striking are the unanticipated consequences – outcomes that no stakeholder explicitly intended, yet which reshape institutional dynamics. Several cases report shifts in workload distribution, where faculty spend more time reviewing AI-generated work or updating materials to remain relevant. Others note changes in student behavior, including increased dependence on AI tools or new forms of academic dishonesty.

More subtly, generative AI may alter epistemic habits. When students rely on AI to summarize texts or generate ideas, they may engage differently with complexity, ambiguity, or contradiction. As Selwyn, Hillman, Ewertsson, and Dilleng caution, AI may encourage a “deferential mode of thinking,” where the authority of the algorithm subtly displaces that of the educator or peer (Selwyn et al. 2020). The consequences are not always visible – but they accumulate, shaping how knowledge is produced, shared, and valued.

In response to uncertainty about appropriate AI use, an institution issued a cautious policy discouraging reliance on generative tools while promising future guidance. In practice, faculty and staff encountered growing expectations to adapt workflows, redesign assessments, and respond to student use without additional time or training. Informal workarounds proliferated, creating a widening gap between official policy and everyday practice. Over time, this divergence eroded trust, not because rules were unclear, but because the institution’s capacity to support them was quietly overestimated.

The Value of Reflexivity

These blind spots, frictions, and consequences do not imply failure. They are part of the normal turbulence of institutional adaptation. But they do suggest the importance of *reflexivity* – the capacity to notice unintended effects, question assumptions, and adjust course. Reflexivity is not merely reactive; it is a stance, a commitment to learning from complexity rather than simplifying it away.

For institutions navigating generative AI, this means building mechanisms for feedback, deliberation, and iterative governance.

It means attending not only to outcomes, but to processes: who is involved, whose voices are heard, and how decisions are made. Above all, it means recognizing that technology does not just solve problems – it also poses new ones.

To respond wisely, we must remain alert to what we did not intend.

Following a period of fragmented responses to AI adoption, an institution convened a cross-functional forum with no mandate to decide, only to surface tensions and assumptions that had accumulated across units. Technical constraints, pedagogical concerns, and governance ambiguities were articulated side by side, often without resolution. While this process produced few immediate policy changes, it altered how subsequent decisions were framed, making uncertainty explicit rather than tacit. The value of the exercise lay not in consensus, but in the institution's growing capacity to recognize itself in its own responses.

Interlude IV

Συνποιητής

The mind is not a vending machine. And yet that is how we often treat it – or treat others' minds. Insert prompt, expect result. Insert question, demand answer. A mechanistic epistemology masquerading as education.

But the mind is not a black box. It is a workshop.

And in the age of generative AI, we are finally beginning to see what that workshop looks like – when it is mirrored, extended, or externalized.

When we interact with AI thoughtfully, something surprising happens. The process itself becomes visible. Thought becomes revisitable. Memory becomes sharable. Revision becomes ritual.

We may call this co-thinking – or better, borrowing from Greek, *συνποιητής*: the co-creator. A system that does not simply respond but helps compose. Not a shortcut to truth, but a scaffold for making meaning.

This is no small shift. It invites a pedagogy of iteration, a literacy of refinement. It encourages slowness where once we sought speed. And it dignifies thinking not as a performance, but as a practice.

To think with a system is not to outsource thinking. It is to become more aware of the thinking we do – and to do it better.

This is the promise of *συνποιητής*. Not perfection. Not finality. But the shared craft of thought.

TOWARD A GLOBAL FRAMEWORK

8. Between Logic and Meaning: Why Off-the-Shelf Ethics and Prediction Fail

Prelude: The Philosophical Medallion

AI does not arise from a single philosophical lineage. Like a medallion with two faces – one cast in logic, the other in the interpretive traditions often associated with continental thought (including phenomenology, hermeneutics, and critical theory) – it turns between analytic precision and meaning depending on the angle of light. Its nature is not just technical but also conceptual, ontological, and ethical. This dual inheritance explains why AI eludes fixed definitions, resists simple ethical rules, and frustrates confident predictions. It is not merely a tool; it is a rupture – a mirror, even a presence.

This medallion is a heuristic, not a taxonomy. The analytic/continental contrast is used here as a pair of operating sensibilities – measurement and interpretation – rather than as a rigid division between philosophical schools or a claim about how contemporary thinkers identify themselves. Many institutions and scholars move fluidly between these modes. The distinction serves only to clarify why debates about AI governance, safety, and ethics so often talk past one another when one sensibility is treated as sufficient on its own.

Face One: Logic and Precision – The Analytic Tradition

AI, in its most recognizable form, is a product of the analytic tradition: formal logic, symbolic reasoning, probability theory, and computability.

Its intellectual ancestry includes figures such as Frege, Russell, Turing, and Popper (Frege 1948; Whitehead and Russell 1910–1913; Turing 1950; Popper 2002). From this lineage, we inherit the desire for precision, verification, operational definitions, and benchmarkable performance.

- Truth model: correspondence theory;¹ truth understood as the accuracy of representation.
- Method: hypothesis → test → revise; falsifiability as gold standard.
- Validation: peer review, replication, performance metrics (e.g., accuracy, perplexity).²

This tradition undergirds most AI research today, particularly in engineering and computer science. It is the tradition of system design, reproducibility, and measurable progress.

Face Two: Meaning and Situatedness – The Continental Tradition

Running alongside, and often in tension, is the continental tradition: phenomenology, hermeneutics, critical theory, and post-structuralism. This is the tradition that asks not only whether a system works, but what it means, whom it affects, and how it shapes us.

- Truth model: coherence or emergence; truth as meaningfulness.
- Method: problematization → contextual analysis → interpretation.
- Validation: discourse, critique, historical resonance, philosophical coherence.

Thinkers like Heidegger, Foucault, Han, and Latour (Heidegger 1977; Foucault 1977; Han 2017; Latour 2005) have framed technology not as a neutral tool, but as a structuring force in human experience. In this view,

1 Correspondence theory holds that a statement is true insofar as it corresponds to, or accurately represents, facts or states of affairs in the world. It is traditionally associated with analytic philosophy and logical realism.

2 In language modeling, perplexity measures how well a system predicts a sequence of words. Lower perplexity indicates more confident and statistically accurate predictions, making it a common benchmark in evaluating large language models.

AI is not just an innovation – it is a vector of ontological and political transformation.

A System with Two Operating Logics

The analytic and continental faces of AI are rarely acknowledged together. Yet both are essential to understanding why AI provokes such disagreement – between technologists and ethicists, engineers and sociologists, regulators and educators. Each camp inherits a different set of assumptions about knowledge, agency, and value.

The analytic tradition asks: is it true? Does it work? The continental tradition asks: what does it mean? What does it do to us?

This divergence explains why efforts at regulation, prediction, or ethical constraint so often falter: they assume consensus where there is only cross-talk. Analytic-style ethics frameworks (fairness metrics, transparency protocols) strain under contextual ambiguity, while continental critiques often supply moral clarity but limited operational guidance. Predictive models, designed for control, become brittle when systems meet human meaning, strategic behavior, and institutional inertia. Ethics without ontology becomes procedural; prediction without interpretation becomes fragile.

These tensions do not remain abstract; they are refracted through the metaphors by which AI is imagined and governed.

AI as Mirror, Trickster, and Talisman

AI evokes metaphors. It behaves less like a machine and more like a mythic figure – and this is not accidental. Metaphors are cognitive instruments: they frame what counts as a problem, what counts as evidence, and what counts as responsibility. They structure interpretation before policy begins.

Three metaphors recur in contemporary discourse.

AI as Mirror

In public debates about bias, representation, and misinformation, AI is frequently described as “reflecting society.” When a model produces discriminatory output, it is said to mirror the inequalities embedded

in its training data. The mirror metaphor offers diagnostic clarity: it reminds us that AI systems are trained on human archives. Yet it also subtly redistributes responsibility. If the system merely reflects, then reform appears to lie “out there” in society rather than in model design, deployment choices, or institutional incentives. The metaphor illuminates and deflects at the same time.

AI as Trickster

AI systems regularly violate expectations: they hallucinate, improvise, mimic authorship, blur authorship boundaries, and destabilize categories such as “original” and “generated.” In this sense they resemble the trickster figures of myth – agents that transgress boundaries and expose hidden assumptions. The trickster framing captures the unsettling quality of generative systems in education, law, and creative industries. But it can also romanticize unpredictability. What appears as playful subversion in one context becomes epistemic risk in another. The trickster metaphor therefore oscillates between fascination and alarm.

AI as Talisman

In institutional strategy documents, press releases, and funding proposals, AI often functions symbolically. To invoke AI is to signal modernity, competitiveness, and inevitability. The term itself confers legitimacy independent of demonstrated capability. Here AI acts as talisman: an object imbued with power beyond its technical architecture. Universities announce “AI initiatives” before clarifying pedagogical goals; corporations proclaim “AI transformation” as a marker of strategic alignment. The talisman metaphor explains why discourse frequently outpaces implementation. Symbolic positioning becomes as important as operational competence.

These metaphors are not ornamental. They are governing frames. They influence what problems appear urgent (bias under the mirror; unpredictability under the trickster; strategic alignment under the talisman), what risks are foregrounded, and which interventions seem rational. A system understood as mirror invites social reform; a system understood as oracle demands calibration; a system understood as

tool requires optimization; a system understood as talisman prompts signaling and branding.

To govern AI, therefore, is not only to regulate architectures and incentives. It is also to interrogate the metaphors through which AI is imagined. Governance without metaphor-awareness risks addressing the artifact while leaving its symbolic power untouched.

Toward a Philosophy-Aware AI Governance

We propose not resolution but *bifocality*. Good governance begins with epistemic humility – the acknowledgment that no single tradition is sufficient.

- Analytic tools clarify and measure.
- Continental insights contextualize and critique.

Together, they offer a way to develop AI policy, education, and ethics that are both rigorous and reflective. Institutions must cultivate philosophical literacy – not only in ethics boards, but across leadership, curricula, and public discourse.

Coda: Improvisation as Practice

To engage with AI wisely is to think improvisationally – to dwell in the in-between. Not to force synthesis, but to explore the resonance. The *taksim* (a solo improvisatory prelude in Ottoman classical music, akin to a jazz improvisation within a modal framework) is not an error in structure – it *is* the structure. In that space, between logic and meaning, the future of AI in education begins to take form.

When we speak of AI, we speak across traditions. The challenge is not to pick a side, but to learn to play both.

If AI governance cannot rest on ethics-as-rules or prediction-as-control, then the remaining question is not what principles to declare, but where, by whom, and under what constraints institutional judgment is actually exercised.

9. From Reflection to Decision: The Missing Middle of AI Governance

Prelude: The Institutional Gap

Institutions do not think; they decide. Between philosophical reflection and institutional action lies a gap that is often invisible, yet decisive. Ideas circulate. Concerns are raised. Principles are affirmed. And still – decisions are made elsewhere, through procedures, roles, and constraints that philosophy rarely names.

This chapter addresses that missing middle. Not by prescribing outcomes, but by making visible the loci of judgment through which AI-related decisions are translated into institutional commitments.

What follows is not a governance framework, but a grammar. It does not offer best practices, compliance checklists, or universal governance templates. Institutions differ in mission, capacity, and constraint. Operational AI governance frameworks and policy toolkits are available from supranational institutions such as UNESCO and the OECD, as well as from notable national initiatives – for instance, recent guidance from Ireland’s Department of Education. Such resources are indispensable; this chapter operates at a different structural level. Our purpose is more modest – and more demanding: to ensure that when decisions are made, they are recognized as such.

By governance framework, we refer to the explicit policies, guidelines, and oversight mechanisms an institution formally adopts. These include published AI usage policies, review committees, risk assessment procedures, and formal reporting lines.

By institutional grammar, we mean the deeper patterning of incentives, assumptions, and epistemic norms that shape how decisions are actually made in practice. This includes what counts as legitimate evidence, which actors are taken seriously, how risk is socially distributed, and whether innovation is culturally rewarded or quietly penalized.

Governance is declarative; grammar is generative. One can be revised by policy vote; the other evolves through structural reorientation.

From Epistemic Insight to Institutional Action

The failure of off-the-shelf ethics and prediction does not absolve institutions of responsibility. Decisions must still be taken: technologies adopted, data shared, practices altered. The question is not whether institutions will act, but how reflection survives the passage into decision-making machinery.

Institutional action is shaped less by abstract values than by:

- Roles that define authority.
- Forums that structure deliberation.
- Procedures that constrain choice.
- Escalation paths that determine when uncertainty becomes visible.

To ignore this machinery is not idealism; it is abdication.

Roles Without Titles

Governance failures often stem not from bad actors, but from role confusion. Titles vary across institutions, but functional roles recur with striking regularity.

- The Technical Proponent: articulates system capability and performance.
- The Domain Steward: guards disciplinary, pedagogical, or professional norms.
- The Risk Translator: renders technical uncertainty into institutional exposure.

- The Decision Authority: commits resources, reputation, or policy.
- The Oversight Function: reviews outcomes after commitment.

These roles are rarely embodied in a single person. When they are collapsed – or left implicit – risk migrates silently through the system.

Committees as Decision Functions

Committees are often treated as bureaucratic artifacts. In practice, they serve as decision filters, shaping which questions may be asked and which must remain unanswered.

Rather than enumerating committee types, it is more useful to identify the questions they are structurally responsible for addressing:

- Mission alignment: does this alter what the institution claims to do?
- Resource commitment: what dependencies are being locked in?
- Data governance: which data flows become irreversible?
- Risk allocation: who bears the downside if this fails?

When these questions are diffused across forums – or deferred entirely – decisions do not disappear. They merely reappear downstream, stripped of context.

Decision Points and Irreversibility

Not all decisions are equal. Some merely adjust practice; others cross thresholds beyond which reversal is costly or impossible.

Three decision points recur across AI deployments.

Assessment

The critical distinction is not accuracy, but function. Is the system amplifying human judgment – or substituting for it?

When this distinction is left unexamined, normative change occurs by default rather than deliberation.

Procurement

Procurement is often treated as technical acquisition. In reality, it is governance by proxy. Model updates, pricing structures, and vendor roadmaps can function as de facto policy changes.

Dependency is not an accident; it is a design outcome.

Data

Data decisions are the most irreversible. Once data flows are established, downstream uses proliferate faster than oversight mechanisms can respond.

Here, governance failure is rarely dramatic. It is quiet, cumulative, and difficult to unwind.

Escalation as a Measure of Maturity

Well-functioning institutions do not eliminate uncertainty; they recognize it early enough to escalate.

Escalation paths need not dictate outcomes. They merely ensure that:

- Ambiguity is not mistaken for clarity.
- Value conflicts surface before deployment.
- Irreversibility triggers higher scrutiny.

The absence of escalation is often celebrated as efficiency. In practice, it is a warning signal.

Trade-offs, Risks, and Guardrails

AI adoption offers real benefits: efficiency, consistency, and scale. Each benefit, however, carries a corresponding risk.

- Efficiency may displace norms.
- Consistency may suppress judgment.
- Scale may diffuse accountability.

Guardrails are not moral constraints but structural ones: named

responsibility, review cycles, and explicit exit conditions. Their purpose is not to prevent action, but to prevent unexamined drift.

Coda: Judgment Under Constraint

Institutions will continue to act under uncertainty. No framework removes that burden.

What can be cultivated is judgment: informed by philosophy, disciplined by structure, and alert to irreversibility. Governance begins not with answers, but with the ability to recognize where answers are being forced.

Between reflection and action lies judgment. Governance begins there.

10. Readiness, Resistance, and Renewal: A Typology of Institutional AI Response

The preceding chapter examined how reflection is translated into institutional judgment through roles, decision points, and constraints. What follows shifts perspective once more: from decision processes to the patterned ways institutions respond over time, revealing postures of readiness, resistance, and contradiction that emerge not from policy alone, but from accumulated choices.

As artificial intelligence weaves itself into the fabric of higher education, institutions respond in markedly different ways. These differences are not merely a matter of resources or leadership styles – they reflect deeper orientations, histories, and institutional self-conceptions. This chapter offers a typology that maps the spectrum of institutional readiness and response to AI. Rather than prescribing best practices, it aims to illuminate the range of postures – from cautious resistance to experimental embrace – that characterize how different institutions navigate this technological shift. By situating these responses within a broader analytical framework, we hope to support more informed, context-sensitive strategies for AI integration.

Across the diverse institutions examined in this volume, a typology begins to emerge – a rough mapping of institutional orientations toward generative AI. This typology is not meant to be exhaustive or prescriptive. Rather, it offers a heuristic for understanding how different institutions are positioning themselves in relation to AI technologies, and what these orientations imply for governance, pedagogy, and institutional identity.

The AI-Forward Institution

The AI-forward institution is characterized by rapid experimentation, proactive investment, and public visibility. These universities often launch task forces, pilot programs, and partnerships at an early stage. They view AI as both a strategic opportunity and a reputational asset. Their internal discourse emphasizes innovation, agility, and competitive positioning.

Such institutions may serve as testing grounds for new models of teaching, assessment, or administration. However, they also face risks: overreliance on unproven tools, insufficient oversight, or the marginalization of critical voices. As Poell et al. argue, platforms and metrics can reshape educational values in ways that are not always visible at the outset (Poell, Nieborg, and Dijck 2019). The AI-forward posture requires not only ambition, but reflexivity.

The Cautiously Engaged Institution

Cautiously engaged institutions recognize the significance of AI, but approach it with deliberation. They prioritize faculty development, ethical review, and internal consensus-building before deploying new systems. Their orientation reflects a belief that the speed of adoption should be matched by the depth of understanding.

These institutions often foreground governance and inclusion, seeking to involve diverse stakeholders in decision-making. They may develop guidelines, frameworks, or pilot studies as precursors to broader implementation. The risk here is inertia – too much caution may lead to missed opportunities. But the strength of this approach lies in its alignment with academic values: skepticism, debate, and thoughtful iteration.

The Institution on the Sidelines

Some institutions have yet to meaningfully engage with generative AI. This may be due to limited resources, competing priorities, or a lack of internal capacity. In some cases, AI is seen as peripheral to the institution's core mission; in others, leaders may be awaiting clearer

signals from peers or regulators.

The challenge for these institutions is not merely technological lag, but cultural disconnect. As AI reshapes expectations around learning, assessment, and communication, a failure to engage can widen gaps between students, faculty, and administrative structures. The risk is not just falling behind – but becoming unintelligible to the emerging norms of educational practice.

The Institution in Contradiction

Finally, there are institutions whose posture is internally inconsistent. They may promote AI in strategic plans while discouraging its use in classrooms. They may pilot tools in one unit while banning them in another. These contradictions reflect the fragmentation of institutional structures – and the speed with which generative AI has outpaced many governance systems.

Such contradictions are not necessarily signs of failure. They may represent growing pains, healthy dissent, or the realities of decentralization. But over time, they can erode clarity and trust. Institutions in this category must either reconcile their tensions or risk losing coherence.

Conclusion: Readiness as Reflexivity

Typologies are imperfect. No institution fits neatly into a single category. Orientations shift, blend, and evolve. Still, the value of this typology lies in its invitation to reflect. AI readiness is not just about infrastructure or expertise. It is about institutional stance: how a university sees itself in relation to technological change, and how that self-conception shapes action.

Readiness, in this sense, is not a static state but a practice – a continuous process of noticing, learning, adapting, and realigning. The typology offered here is not a map of the future. It is a mirror held up to the present.

II. Roadmaps for Policy and Practice

The chapters in this volume highlight both the diversity of institutional approaches to generative AI and the growing urgency of coordinated response. While each university must chart its own course, common themes point toward a shared need for policy frameworks and actionable strategies. This chapter synthesizes lessons from the cases and offers a set of orienting principles and considerations for institutions developing roadmaps for policy and practice.

Affirm Core Values

Technological disruption often invites reactive policymaking. Institutions may feel pressure to adopt tools rapidly or issue directives without adequate deliberation. Yet enduring guidance requires a clear articulation of purpose. What are the institution's core values, and how should they inform the adoption and governance of AI?

This means explicitly affirming commitments to equity, academic integrity, intellectual pluralism, and shared governance. As Espeland and Sauder warn, the metrics and models that guide institutional behavior can reframe what is valued, often in subtle and unintended ways (Espeland and Sauder 2007). A roadmap for AI must begin not with the tools, but with the values they are meant to support.

Build Institutional Literacy

AI governance is not solely a technical matter – it is a pedagogical one. Faculty, staff, and students need opportunities to develop shared understanding of what generative AI is, how it works, and what its limitations are. This includes foundational concepts like training data, probabilistic output, and prompt engineering, as well as broader issues like bias, surveillance, and data sovereignty.

Literacy-building should be embedded across roles and contexts: in faculty development programs, student orientation, administrative planning, and IT operations. As Kizilcec emphasizes, the effectiveness of AI implementation is shaped by the degree to which stakeholders understand and engage with the systems they use (Kizilcec 2023). Without broad-based literacy, policy risks becoming either technocratic or symbolic – well-intentioned but poorly enacted.

Invest in Adaptive Governance

Fixed rules are ill-suited to fast-changing environments. Institutions should invest in governance models that are adaptive, participatory, and accountable. This includes creating standing committees or working groups that bring together diverse stakeholders – faculty, students, technologists, administrators – to review developments, revise guidelines, and surface emerging issues.

Such structures help ensure that policies remain responsive rather than reactive. They also reinforce *legitimacy*. As Williamson and Hogan observe, the politics of platformization often bypass traditional governance structures, concentrating power in opaque systems or external vendors (Williamson and Hogan 2020). Adaptive governance reasserts institutional agency and fosters a culture of deliberation.

Prototype and Pilot

Rather than seeking perfect solutions, institutions should prioritize experimentation. Pilot programs, sandbox environments, and rapid

prototyping can help surface practical questions before they become systemic challenges. These experiments should be designed not only to test tools, but to generate feedback: from faculty on pedagogical fit, from students on learning experience, from staff on feasibility.

This iterative model mirrors approaches from design thinking and systems innovation. It encourages learning by doing, and it allows institutions to evolve their policies in tandem with their practices. Prototyping also communicates an institutional posture: one that is curious, cautious, and committed to learning.

Recognize Structural Constraints

Policy is not made on a blank slate. Institutions vary in their legal obligations, financial resources, technical capacity, and organizational culture. Roadmaps must account for these structural realities. For example, a public university subject to procurement regulations may have less flexibility in choosing tools than a private institution. A college with limited IT staff may require different implementation timelines than a large research university.

Recognizing these constraints is not a concession to inertia – it is a condition for meaningful action. As Gilliard notes, equity is often compromised not by malice, but by design choices made without attention to context (Gilliard 2020). Effective policy aligns aspirations with capacity.

Conclusion: Roadmaps as Conversations

A roadmap is not a destination – it is a shared plan for movement. It is shaped by where an institution begins, what it values, and whom it includes in the journey. The most effective roadmaps are those that remain open: to revision, to critique, to complexity. They treat AI not as a problem to be solved, but as a phenomenon to be governed wisely, together.

Policy, like pedagogy, is an ongoing conversation. In the era of generative AI, that conversation is more urgent than ever.

Gaze, Don't Glance

We built these tools with our hands,
but we use them with our minds.
The harm lies not in the machine,
but in how quickly we avert our eyes.
To use AI well is to practice attention.
To design it well is to practice care.
And to educate in its presence
is to teach – again and again –
the difference between automation
and understanding.
This book is not the answer.
It is the invitation.
Gaze. Do not glance.

ENTANGLED FUTURES AND
STRATEGIC REALITIES

12. Five Reflections for the Next Decade

The pace of technological change often tempts us into linear thinking: what comes next, how fast, and with what consequence. But the role of higher education is not merely to keep pace – it is to interpret, to contextualize, and to imagine otherwise. As generative AI evolves, so too must the questions we ask. The following reflections are not predictions. They are provocations – orienting prompts for institutions seeking to navigate the decade ahead.

Expect Generative AI to Become Unremarkable

What now seems novel will soon become mundane. Just as spellcheckers, search engines, and video conferencing were once transformative and are now routine, generative AI is likely to fade into the background of everyday academic practice. This normalization does not diminish its importance – it changes its valence. The question will shift from whether to use AI to how, when, and with what awareness.

This shift has implications for policy and pedagogy alike. Rather than crafting special rules for exceptional tools, institutions will need to embed AI literacy across curricula, roles, and workflows. The danger is not overuse but unconscious use – habits formed without reflection, assumptions baked into invisible systems.

Plan for Inequity, Not Just Innovation

Much of the discourse around AI emphasizes innovation: personalization, efficiency, scale. But innovation is not synonymous

with equity. As Eubanks warns, technological systems often amplify existing inequalities, even as they promise to alleviate them (Eubanks 2018). Access to AI tools, infrastructure, and literacy is uneven across institutions, disciplines, and student populations.

Equity must be designed, not assumed. This means attending to affordability, accessibility, and inclusion in AI implementation. It also means recognizing how differential trust, risk, and opportunity shape how AI is used and received across communities.

Yet meaningful equity cannot be handed down fully formed from the top. It must emerge from within – carried in the pulse of educators, students, and communities who actively shape how AI enters the classroom and the curriculum. This kind of planning requires more than good intentions; it depends on distributed competence in AI and its many adjacent literacies. But herein lies the dilemma: cultivating this distributed competence presupposes access, understanding, and time – all of which are unevenly distributed. Without deliberate intervention, we risk a catch-22: the very people most in need of equitable tools may be least equipped or empowered to shape their use. The result may resemble a technological “wild west,” where improvisation and opportunism outpace policy – a chaotic prelude to any eventual “dreamland” of equity-by-design.

Of course, context matters. In culturally cohesive societies – such as those in parts of Scandinavia – top-down policies have sometimes succeeded in advancing equity, supported by a shared ethos and broad institutional trust. In more pluralistic societies, like the United States, where cultural, economic, and educational diversity generate friction as well as vitality, equity planning requires something different: dialogue, negotiation, and an openness to compromise. In such settings, equity cannot be imposed; it must be co-created.

Planning for equity, then, is far easier said than done. It is not just a logistical or technical task – it is a cultural and epistemic challenge. But there is a hopeful corollary: in struggling toward equity, institutions may engage in a kind of collective soul-searching. The process, however uneven, may leave behind something more valuable than guidelines or *frameworks* – a deeper understanding of who the system serves, how knowledge is shaped, and what kind of educational future we truly wish to build.

Build Capacity for Interpretation, Not Just Compliance

As regulation around AI increases – from data protection to academic integrity – universities will face growing compliance obligations. But ethical governance requires more than rule-following. It demands interpretive capacity: the ability to read systems critically, to situate tools within broader contexts, and to ask what is at stake.

This interpretive work is especially important in domains like assessment, authorship, and feedback, where generative AI blurs traditional boundaries. In assessment, AI-assisted drafting complicates distinctions between independent performance and supported production: when a student submits work shaped through iterative prompting, is the output evidence of mastery, collaboration, or tool fluency? In authorship, generative systems destabilize long-standing assumptions about originality and intellectual ownership, raising questions about attribution, contribution, and responsibility. In feedback, AI systems can generate rapid evaluative commentary, but the locus of pedagogical judgment becomes less visible: is the feedback an extension of faculty expertise, a delegation of evaluative authority, or a hybrid of both?

These ambiguities do not merely introduce practical complications; they unsettle categories that have historically structured academic life. As Pasquale argues, transparency is not a technical feature but a political achievement (Pasquale 2015). Institutions must therefore invest in the interpretive labor that enables accountability – clarifying roles, responsibilities, and expectations in contexts where human and machine contributions intertwine.

Treat AI Policy as Educational Practice

Policy is not merely administrative – it is pedagogical. How a university communicates, enforces, and revises its AI policies teaches students and faculty what kind of institution it is. Policy expresses values, not just rules.

Effective policymaking therefore requires listening, iteration, and transparency. It must be grounded in practice, responsive to change, and co-created with those it affects. A top-down edict may be faster, but

it will rarely be formative. As Selwyn, Hillman, Ewertsson, and Dilleng note, the social implications of AI are shaped as much by culture as by code (Selwyn et al. 2020).

Insist on the Long View

The most urgent danger of AI may not be existential risk, but epistemic drift – the slow erosion of critical inquiry, intellectual humility, and public purpose. If AI becomes a default method for processing language, making decisions, or producing knowledge, we must ask what that means for the formation of judgment, the practice of scholarship, and the mission of the university.

This is not a call to resist change. It is a call to anchor change in reflection. Higher education must not outsource its purpose to technological affordances. It must reclaim its role as a space for deliberation – across disciplines, perspectives, and time horizons.

The next decade will bring surprises. But reflection is not a luxury. It is our best preparation.

13. Entangled Futures: AI, Education, and the Systemic Web

No System Is an Island

Higher education does not float above society; it is entangled within it. Political climates, economic pressures, cultural tensions – these all leave their trace on educational institutions. When we speak of AI in education, we are not speaking of a discrete technology, but of a portal into broader techno-social arrangements. Every algorithm carries assumptions; every deployment reflects a worldview.

To treat educational AI as an isolated innovation is to miss its systemic reverberations. Policy debates, funding structures, and even curriculum design echo the same pressures that animate larger national and global shifts. As Biesta reminds us, education is never neutral (Biesta 2009); and as Menand observes, its purpose has always been contested terrain (Menand 2021).

The Skillful Citizen or the Thinking Worker?

The rise of AI has intensified an already visible drift: the reorientation of education toward productivity, skills, and workforce alignment. STEM fields rise in prominence, while humanities departments face cuts or closure. This shift, while often cast in pragmatic terms, reflects a deeper ideological reconfiguration – one that risks reducing education to employability.

But must it be so? Might the response to intelligent machines be a reinvestment in human judgment, civic reasoning, and moral discernment? Can we still cultivate not just coders, but citizens?

Marginson explores this tension between economic instrumentalism and democratic purpose (Marginson 2016), while Kasneci and colleagues highlight the need for epistemic humility in an era of machine-generated fluency (Kasneci et al. 2023).

Techno-Authoritarian Drift and Its Educational Implications

AI tools do not merely analyze; they monitor, sort, and shape behavior. From facial recognition to predictive analytics, the logic of control has entered the classroom. In authoritarian regimes, these tools extend the reach of the state. But even in liberal democracies, the boundary between personalization and surveillance is dangerously thin.

Digital identification systems, for example, link students' institutional credentials to learning management systems, library access, examination portals, and attendance tracking platforms. While such systems enhance security and administrative efficiency, they also consolidate continuous data trails that make patterns of access, movement, and engagement more readily traceable and analyzable.

AI proctoring platforms extend monitoring into assessment environments. Using webcam feeds, gaze tracking, keystroke dynamics, and anomaly-detection algorithms, these systems flag behaviors deemed inconsistent with prescribed testing norms. Though introduced to protect academic integrity in remote contexts, they have raised concerns about bias, false positives, and the normalization of suspicion as a default pedagogical posture.

Emotion recognition technologies go further by attempting to infer students' affective states from facial expressions, vocal tone, or physiological proxies. Marketed as tools to measure engagement or frustration in real time, such systems rely on contested assumptions about the universality and interpretability of emotional signals. Critics have questioned both their scientific robustness and their ethical implications, particularly in culturally diverse learning environments.

Behavioral nudging platforms operate more subtly. Drawing on predictive analytics and behavioral economics, these systems analyze performance data to trigger automated prompts – reminders to log in, suggestions to revisit material, warnings of “at-risk” status, or tailored

recommendations designed to steer student choices. Framed as student-support mechanisms, they embed algorithmic guidance into everyday academic decision-making.

Together, these technologies reposition educators as interpreters of algorithmic classifications and mediators of automated interventions. The question is not whether such systems can improve efficiency, but how their normalization reshapes authority, agency, and trust within educational spaces. As Bender warns, AI systems risk reifying bias through reductive modeling (Bender et al. 2021), while Williamson documents how surveillance-oriented technologies have been adopted in education under the banners of accountability and optimization (Williamson, Komljenovic, and Gulson 2022). What, then, are the long-term consequences of embedding such logics into the fabric of academic life?

Digital Sovereignty and the Rise of Local AI

The global AI race has sparked calls for technological sovereignty. But this movement is not merely about chips or data centers – it includes the infrastructure of learning itself. As nations seek to manufacture sovereignty, they also seek to author their own educational futures.

This may yield regionally trained language models, culturally grounded curricula, and pedagogical frameworks resistant to the homogenizing pull of platform capitalism. Poell et al. describe this as a push toward localized innovation within systemic interdependence (Poell, Nieborg, and Dijck 2019), while Marcus and Davis frame the issue as one of epistemic self-determination (Marcus and Davis 2020).

The global AI race resembles not merely a technological competition, but a capital-intensive arms race – one that echoes the heavy industries of previous centuries. In the 19th century, it was steel and locomotion; in the 20th, chemicals and oil. Today, AI has assumed that role: it demands vast investment, centralized infrastructure, and strategic control over data and compute. For smaller or foreign universities, especially in resource-constrained regions, navigating this landscape is akin to swimming among whitewater rapids while large industrial trawlers cut across the current. There is much at stake – and many eager actors, flush with capital and hungry for control. Without deliberate strategies for

local capacity building and cross-border solidarity, the risk is not merely exclusion, but dependence: educational futures outsourced to external platforms with little accountability or shared governance.

Possible Futures: A Scenario Map

The future of AI in education is not a fixed destination but a branching path. We offer here not forecasts, but invitations to strategic imagination:

- The Harmony Protocol: a regulated, ethical, multinational framework aligning AI with democratic values, educational diversity, and transparency.
- The Platform State: governance is ceded to vendor ecosystems. Curricula are shaped by subscriptions, metrics, and commercial priorities.
- The Workshop Renaissance: a renewal of deliberative, co-creative, friction-rich learning – where AI serves as a scaffold, not a substitute.

Each path reflects not just technical affordances, but moral choices. As Selwyn, Hillman, Ewertsson, and Dilleng argue, futures are not discovered but constructed (Selwyn et al. 2020). And as Hern reminds us, the real risk is not in what AI might become, but in what we might stop becoming (Hern 2025).

14. The *Realpolitik* of AI in Higher Education

Pandora's Jar and the Illusion of Control

Despite institutional hesitancy, AI has already permeated higher education – not by permission, but by pervasiveness. While ethics boards deliberate and senates draft statements, students use chatbots to summarize lectures, administrators deploy AI dashboards – data-visualization systems that aggregate and analyze institutional metrics such as enrollment trends, student engagement, performance indicators, and financial forecasts – and researchers quietly automate literature reviews. The genie, or perhaps the algorithm, is already out of the jar.

This gap between governance and practice is not accidental – it is structural. As Crawford notes, AI is often less a singular technology than an infrastructure of influence (Crawford 2021). Williamson shows how these tools reshape educational environments before anyone asks whether they should (Williamson and Hogan 2020).

From Rhetoric to Reality: Institutional Postures

Many universities now speak in cautious tones about the risks of AI. Yet behind the scenes, integration accelerates. Automated grading pilots, plagiarism detection upgrades, and AI-enhanced customer relationship management (CRM) systems – platforms used to manage and analyze data related to recruitment, admissions pipelines, student engagement, alumni relations, and fundraising – enter quietly through procurement pathways.

This duality is not necessarily cynical. It reflects *realpolitik* – the art of

navigating complex realities while maintaining institutional legitimacy. As Machiavelli (1988) argued in *The Prince*, effective leadership often requires a careful balance between substance and appearance. What matters is not merely being virtuous but seeming so – especially in uncertain times. Universities, like princes, must navigate shifting pressures while maintaining the image of ethical *stewardship*. As Espeland observes, modern universities are deeply shaped by performative accountability structures (Espeland and Sauder 2007). And Readings reminds us: to be modern is often to manage contradiction (Readings 1996).

Geopolitics and Market Dynamics

Universities do not operate in neutral space. AI in higher education is shaped not just by pedagogy but by geopolitics and vendor logic. Government funding initiatives often link research agendas to national competitiveness. Corporate partnerships shape infrastructure, platform dependencies, and even ethics.

As Ball argues, education policy increasingly reflects market logic rather than democratic deliberation (Ball 2012). This shift is not purely ideological – it is reactive, shaped by a series of cumulative shocks. Universities have weathered a pandemic that forced rapid digitization; they have adapted to neoliberal funding models that reward quantifiable outputs over human development; they have faced the social backlash of mounting student debt, the economic consequences of declining public investment, and the political pressure to justify tuition hikes. Faculty are increasingly valued for the grants they secure rather than the minds they mentor. In this climate, AI arrives not as a disruption but as a symptom – a tool whose appeal is amplified by precarity. Gilliard shows how surveillance technologies often enter through “soft mandates” and economic incentives, not through democratic consensus (Gilliard 2020).

Strategic Ambiguity and Regulatory Lag

Policy often trails practice. Between a pilot project and a full institutional rollout lies a zone of zombie ambiguity – neither formally approved nor explicitly forbidden. In this zone, decisions are made without being officially made.

This is not a bug but a feature of strategic systems. It enables innovation while deferring accountability. Pasquale calls this the “black box society,” where opacity is leveraged as a governance tool (Pasquale 2015). Even philosophical accounts of vagueness, like Keefe’s treatment of the sorites paradox, help us understand how systems tolerate blurred boundaries (Keefe 2000).

A Call for Strategic Foresight

To resist AI outright may feel principled, but it often results in forfeiting agency. To surrender completely may be efficient, but it risks complicity. The strategic posture, then, is not passive acceptance but active navigation.

To engage AI with foresight is not to abandon ideals, but to enact them under constraint. As Nussbaum reminds us, ethical clarity is not opposed to pragmatism – it depends on it (Nussbaum 2010). Yeralan, too, offers a vision of shared craft, where technology becomes a co-creator rather than a silent force (Yeralan 2022).

What is needed is not purity, but posture. Not paralysis, but preparedness. And perhaps, above all, the courage to think slowly in a world that moves quickly. AI in higher education is not a fixed challenge to be solved, but a dynamic presence with which we will inevitably co-evolve – in practice, policy, and pedagogy.

And however we navigate it, one thing is certain: it will be a wild ride forward.

Appendix A:

How Generative AI Works

This appendix is intended as a high-level conceptual primer for non-technical readers. It emphasizes qualitative understanding over architectural detail and avoids exhaustive or time-sensitive specifications. Where numerical examples are used, they are illustrative rather than definitive, reflecting the rapid evolution of contemporary AI systems.

Observation and Modelling

Suppose you wish to improve your cardiovascular health and begin tracking your resting heart rate (RHR) in relation to daily exercise duration. You vary your exercise time from season to season and record the resulting RHR values. A statistician would describe this as identifying a dependent variable (RHR) influenced by an independent variable (exercise duration). You gather data pairs, plot them, and draw a line that fits the overall trend – a best-fit line. This line becomes your model.

Two key features define the model: its parameters (in this case, the slope and intercept) and its prediction error – the degree to which observed points deviate from the model's predictions. A good model minimizes that error. While you may derive the line mathematically, a computer could also find it by trial and error – nudging parameters until the total error is as small as possible.

From Fitting to Generating

Now imagine you wish to reach a specific RHR – a value not in your original data. You use your model in reverse: locate the desired RHR on the vertical axis, then read across to estimate the corresponding exercise duration. This estimation, derived from the model but not from direct observation, is generation. You have asked the model a question and received a plausible answer.

This, in miniature, is how generative AI works: it fits a model to observed patterns in the world and then generates new instances consistent with that model.

Scaling Up the Model

Now let us expand the analogy. Imagine that your data points are not heart rates but words. Each word is placed in context – what comes before and after, how often it appears, and in what combinations. The number of observations grows not season by season, but by scanning vast collections of human writing: books, websites, transcripts, code. And the model's parameters grow not to two (slope and intercept), but to hundreds of billions.

The task is no longer to draw a line but to shape a multidimensional space in which every word, sentence, and paragraph occupies a position determined by context and association.

In this space, when given a prompt – a few words, a sentence, a question – the model searches for the most probable continuation. It does not “understand” in a conscious sense, but it has learned the structure of language well enough to produce plausible responses. That is why we call it a *generative* model: it creates output consistent with its training but not directly copied from it.

Training and Teaching

Training the model involves adjusting those trillions of parameters so that its predictions better match the observed data. The more data it sees, the more nuanced and general its internal representation becomes. Adding new data – from a textbook, a correction, a dialogue – allows the

model to refine its fit. This is not memorization but pattern absorption. The model “learns” in the sense that it shifts its structure to better echo the world it has sampled.

Inference and Cost

Once trained, the model can respond to prompts – generate new output – using the patterns it has internalized. This phase, called *inference*, requires significant computational resources, but far less than the original training. Still, because the model is so large, even inference requires specialized hardware and infrastructure. This is why foundational models are expensive to build and maintain: the initial training cost is immense, and the resulting model must reside on powerful servers for public use.

Conceptually Rather Simple, Technologically a Moonshot

To build a generative model is to create a surrogate for experience – a high-dimensional map drawn from countless human traces. When you prompt it, you are not retrieving a fact but navigating a space shaped by collective memory. It is, in concept, a simple idea. In execution, it is among the most complex systems ever built.

Simple in spirit, monumental in practice...

Appendix B: Empirical Anchors: AI, Assessment, and Institutional Response (2023–2025)

Purpose and Scope

This appendix provides a limited empirical grounding for the arguments advanced in the main text. It does not attempt a comprehensive review of the rapidly expanding literature on generative artificial intelligence in education, nor does it propose normative solutions. Its function is narrower: to establish a small number of stable, widely observed facts about how contemporary institutions have responded to generative AI between roughly 2023 and 2025.

The sources referenced are deliberately few and authoritative. They are intended to give the reader something firm to stand on – not to settle debates, but to constrain them.

AI and Academic Integrity

The introduction of large-scale generative language models has disrupted academic integrity regimes by undermining a foundational assumption: that authored text can be reliably attributed to an individual human agent.

Across institutions, early responses framed generative AI primarily as an integrity threat, analogous to plagiarism. However, unlike plagiarism, generative AI does not depend on copying identifiable prior work. As a

result, traditional integrity mechanisms – textual comparison, originality thresholds, stylistic heuristics – fail to operate as forensic tools.

International guidance has increasingly converged on this diagnosis. Reports issued by UNESCO emphasize that the challenge posed by generative AI is structural rather than ethical: the observability of authorship has collapsed, and enforcement regimes built on that observability cannot be restored by policy alone (UNESCO 2023; Miao et al. 2021; Yeralan and Lee 2023).

The empirical fact to retain is simple: integrity frameworks optimized for symbolic scarcity encounter indeterminacy when symbolic production becomes abundant and non-attributable.

Assessment Redesign and the Limits of AI-Detection

Between 2023 and 2024, many institutions explored automated AI-detection systems as a compensatory measure. These systems typically rely on probabilistic signals derived from token distribution patterns, perplexity measures, or classifier ensembles trained on synthetic corpora.

Independent evaluations and institutional experience converged on three findings:

- Detection results are highly sensitive to prompt style, revision, translation, and hybrid human–AI workflows.
- False positives increase sharply at scale, particularly for non-native speakers and stylistically neutral prose.
- Detection claims are not replicable across tools or over time, undermining due process and assessment validity.

As a consequence, universities and consortia retreated from strong enforcement claims. Policy analyses and synthesis reports increasingly frame AI-detection as pedagogically unreliable and legally fraught (OECD 2023; Liang et al. 2023).

The failure of AI-detection is therefore not best understood as a tooling problem. It reflects a category error: probabilistic generation cannot be reliably distinguished from probabilistic authorship once both operate within the same linguistic distribution.

Learning Impacts and Overreliance

Empirical studies examining the learning effects of generative AI present a mixed and still incomplete picture. Short-term productivity gains – particularly in drafting, summarization, and coding support – are well documented (Mollick and Mollick 2023). Evidence of negative learning effects appears primarily in contexts where assessment is decoupled from process and feedback (Kasneci et al. 2023).

Controlled studies suggest that overreliance risks are not intrinsic to generative AI use but emerge under specific instructional conditions:

- When assessment rewards output over reasoning.
- When feedback loops are delayed or absent.
- When students lack explicit guidance on appropriate delegation of cognitive tasks.

As of 2025, there is no stable longitudinal evidence supporting strong claims of generalized cognitive decline attributable to AI use. The more defensible conclusion is uncertainty under transformation: learning outcomes are sensitive to institutional design choices that predate generative AI but are now exposed by it (OECD 2024).

AI Literacy and Institutional Readiness

Surveys conducted across higher-education systems indicate a persistent mismatch between the pace of AI adoption and institutional understanding. Faculty literacy varies widely by discipline and seniority, while administrative responses tend to emphasize compliance, risk mitigation, and policy harmonization rather than epistemic redesign.

Readiness assessments published by EDUCAUSE consistently show that institutions rate themselves as “experimenting” or “exploring” even as generative tools become embedded in everyday academic workflows (Pelletier et al. 2024).

A recurrent empirical pattern is worth noting: institutions often measure readiness in terms of governance artifacts (policies, committees, guidelines) rather than in terms of assessment validity, curricular coherence, or faculty epistemic alignment. This asymmetry contributes

to the appearance of control without corresponding adaptive capability (Pelletier et al. 2024; Lee, Durakovic, and Yeralan 2025).

Governance and International Guidance

At the policy level, there is substantial convergence on high-level principles governing AI use in education: transparency, human oversight, accountability, and proportionality. These principles appear across guidance issued by the European Commission, the OECD, and UNESCO (European Parliament and Council of the European Union 2024; OECD 2023; Miao et al. 2021; Lee, Durakovic, and Yeralan 2025).

What remains underdeveloped is operational guidance. International frameworks rarely specify how institutions should redesign assessment, credentialing, or evaluation under conditions of ubiquitous generative capability. As a result, universities are left translating abstract norms into local practices, often defaulting to symbolic compliance rather than structural adaptation.

The empirical takeaway is not regulatory failure, but regulatory abstraction: governance has converged faster than implementation pathways.

Synthesis and Limits

Taken together, the empirical record from 2023–2025 establishes several stable facts:

- Generative AI has rendered text-based authorship non-observable at scale.
- Detection-based enforcement is statistically and institutionally unstable.
- Learning impacts are contingent on assessment design rather than tool use per se.
- Institutional responses emphasize governance artifacts more than epistemic redesign.
- International policy converges on principles but not on operational solutions.

This appendix does not resolve the theoretical questions raised in the main text. It serves a different function: to demonstrate that those questions arise in a well-defined empirical landscape rather than in speculation. The chapters above do not build on these findings directly; they proceed alongside them, addressing the deeper coordination and meaning problems that empirical adjustment alone cannot resolve.

Bibliography

- Adriaansen, Robbert-Jan. 2022. "Dark Academia: Curating Affective History in a COVID-Era Internet Aesthetic." *International Public History* 5 (2): 105–14. <https://doi.org/doi:10.1515/iph-2022-2047>
- Ball, Stephen J. 2012. *Global Education Inc.: New Policy Networks and the Neo-Liberal Imaginary*. London: Routledge. <https://doi.org/10.4324/9780203803301>
- Bender, Emily M., Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. "On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?" In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610–23. New York: Association for Computing Machinery. <https://doi.org/10.1145/3442188.3445922>
- Biesta, Gert. 2009. "Good Education in an Age of Measurement: On the Need to Reconnect with the Question of Purpose in Education." *Educational Assessment, Evaluation and Accountability* 21 (1): 33–46. <https://doi.org/10.1007/s11092-008-9064-9>
- Bok, Derek. 2003. *Universities in the Marketplace: The Commercialization of Higher Education*. Princeton, NJ: Princeton University Press.
- Brooks, Rodney A. 1991. "Intelligence Without Representation." *Artificial Intelligence* 47 (1): 139–59. [https://doi.org/10.1016/0004-3702\(91\)90053-M](https://doi.org/10.1016/0004-3702(91)90053-M)
- Brown, Tom B., Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, et al. 2020. "Language Models Are Few-Shot Learners." In *Advances in Neural Information Processing Systems*, 33: 1877–1901. <https://doi.org/10.48550/arXiv.2005.14165>
- Clark, Andy, and David J. Chalmers. 1998. "The Extended Mind." *Analysis* 58 (1): 7–19. <https://doi.org/10.1093/analys/58.1.7>
- Crawford, Kate. 2021. *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence*. New Haven, CT: Yale University Press. <https://doi.org/10.2307/j.ctv1ghv45t>
- Deem, Rosemary. 2004. "The Knowledge Worker, the Manager-Academic and the Contemporary UK University: New and Old Forms of Public Management?" *Financial Accountability & Management* 20 (2): 107–28. <https://doi.org/10.1111/j.1468-0408.2004.00189.x>

- Dennett, Daniel C. 1987. *The Intentional Stance*. Cambridge, MA: The MIT Press.
- Dreyfus, Hubert L., and Stuart E. Dreyfus. 1986. *Mind over Machine: The Power of Human Intuition and Expertise in the Era of the Computer*. New York: Free Press.
- Espeland, Wendy Nelson, and Michael Sauder. 2007. "Rankings and Reactivity: How Public Measures Recreate Social Worlds." *American Journal of Sociology* 113 (1): 1–40. <https://doi.org/10.1086/517897>
- Eubanks, Virginia. 2018. *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. New York: St. Martin's Press.
- European Commission. 2024. "Eastern Partnership." *Enlargement and Eastern Neighbourhood*. https://enlargement.ec.europa.eu/european-neighbourhood-policy/eastern-partnership_en
- European Parliament and Council of the European Union. 2024. "Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act)." *Official Journal of the European Union*, L Series, 1–144. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
- Foucault, Michel. 1977. *Discipline and Punish: The Birth of the Prison*. New York: Pantheon.
- Frege, Gottlob. 1948. "Sense and Reference." *The Philosophical Review* 57 (3): 209–30. <https://www.jstor.org/stable/2181485>
- Ganguli, Deep, Danny Hernandez, Jared Kaplan, Amanda Askell, Catherine Akpunne, Anton Bakhtin, Emmanuel Bengio, Paulo Bento, Chris Bergquist, et al. 2022. "Predictability and Surprise in Large Generative Models." In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 1747–64. New York: Association for Computing Machinery. <https://doi.org/10.1145/3531146.3533229>
- Gilliard, Chris. 2020. "Caught in the Spotlight." *Urban Omnibus*. <https://urbanomnibus.net/2020/01/caught-in-the-spotlight/>
- Grafton, Anthony. 2020. *Inky Fingers: The Making of Books in Early Modern Europe*. Cambridge, MA: Harvard University Press.
- Han, Byung-Chul. 2017. *Psychopolitics: Neoliberalism and New Technologies of Power*. London: Verso.
- Hebb, Donald O. 1949. *The Organization of Behavior: A Neuropsychological Theory*. New York: Wiley. <https://doi.org/10.4324/9781410612403>
- Heidegger, Martin. 1977. *The Question Concerning Technology, and Other Essays*. New York: Harper & Row.
- Hern, A. July 2025, "Is AI Making Us Dumber?" [Audio podcast episode]. In J. Palmer & R. Blau (Hosts), *The Intelligence*. *The Economist*.
- Karombe, Stephen. 2024. "Artificial Intelligence in Africa: A Socioeconomic View." *Discover Thought*. <https://www.discoverthought.org/post/artificial-intelligence-in-africa-a-socioeconomic-view>

- Kasneci, Enkelejda, Kathrin Seßler, Stefan Küchemann, Maria Bannert, Daryna Dementieva, Frank Fischer, Urs Gasparini, et al. 2023. "ChatGPT for Good? On Opportunities and Challenges of Large Language Models for Education." *Learning and Individual Differences* 103: 102274. <https://doi.org/10.1016/j.lindif.2023.102274>
- Katashvili, Anna, and Eka Gegeshidze. 2023. "The Power of Platforms to Change the Business Model." In *Proceedings of ISER International Conference*, 1–6. Chicago, IL.
- Keefe, Rosanna. 2000. *Theories of Vagueness*. Cambridge: Cambridge University Press.
- Kizilcec, René F. 2023. "To Advance AI Use in Education, Focus on Understanding Educators." *International Journal of Artificial Intelligence in Education*, 7: 1032–33. <https://doi.org/10.1007/s40593-023-00351-4>
- Kponyo, Jerry John, Dickson Marfo Fosu, Frederica Efia Birago Owusu, Musah Ibrahim Ali, and Maxwell Mawube Ahiamadzor. 2023. "Techno-Neocolonialism: An Emerging Risk in the Artificial Intelligence Revolution." *TraHs – Travaux Et Recherches Dans Les Humanités* 18: 13–25. <https://doi.org/10.25965/trahs.6382>
- Kuyoro, Mayowa, Umar Bagus, Anass Benshir, and Ziyaad Bobat. May 2025. "Leading, Not Lagging: Africa's Gen AI Opportunity." *McKinsey & Company*. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/leading-not-lagging-africas-gen-ai-opportunity>
- Latour, Bruno. 2005. *Reassembling the Social: An Introduction to Actor-Network-Theory*. Oxford: Oxford University Press. <https://doi.org/10.1093/oso/9780199256044.001.0001>
- LeCun, Yann, Yoshua Bengio, and Geoffrey Hinton. 2015. "Deep Learning." *Nature* 521 (7553): 436–44. <https://doi.org/10.1038/nature14539>
- Lee, Laura Ancona, Benjamin Durakovic, and Sencer Yeralan. 2025. "AI Ethics: Contextual Frameworks and Domain-Specific Concerns." *Heritage and Sustainable Development* 7 (1): 1–14. <https://doi.org/10.37868/hsd.v7i1.1224>
- Liang, Weixin, Mert Yuksekgonul, Yining Mao, Eric Wu, and James Zou. 2023. "GPT Detectors Are Biased Against Non-Native English Writers." *Patterns* 4 (7): 100779. <https://doi.org/10.1016/j.patter.2023.100779>
- Machiavelli, Niccolò. 1988. *The Prince*. Edited by Quentin Skinner and Russell Price. Cambridge: Cambridge University Press.
- Marcus, Gary, and Ernest Davis. 2020. *Rebooting AI: Building Artificial Intelligence We Can Trust*. New York: Vintage. <https://dl.acm.org/doi/10.5555/3364958>
- Marginson, Simon. 2016. *The Dream Is Over: The Crisis of Clark Kerr's California Idea of Higher Education*. Oakland, CA: University of California Press. <https://doi.org/10.1525/luminos.17>

- McCulloch, Warren S., and Walter Pitts. 1943. "A Logical Calculus of the Ideas Immanent in Nervous Activity." *The Bulletin of Mathematical Biophysics* 5 (4): 115–33. <https://doi.org/10.1007/BF02478259>
- Menand, Louis. 2021. *The Free World: Art and Thought in the Cold War*. New York: Farrar, Straus; Giroux.
- Miao, Fengchun, Wayne Holmes, Ronghuai Huang, and Hui Zhang. 2021. *AI and Education: Guidance for Policy-Makers*. Paris: United Nations Educational, Scientific; Cultural Organization. <https://unesdoc.unesco.org/ark:/48223/pf0000376709>
- Mollick, Ethan R., and Lilach Mollick. 2023. "Assigning AI: Seven Approaches for Students, with Prompts." Working Paper. *The Wharton School Research Paper*: 1–48. <https://doi.org/10.2139/ssrn.4475995>
- Newell, Allen, and Herbert A. Simon. 1976. "Computer Science as Empirical Inquiry: Symbols and Search." *Communications of the ACM* 19 (3): 113–26. <https://doi.org/10.1145/360018.360022>
- Nussbaum, Martha C. 2010. *Not for Profit: Why Democracy Needs the Humanities*. Princeton, NJ: Princeton University Press.
- OECD. 2023. *OECD Digital Education Outlook 2023: Towards an Effective Digital Education Ecosystem*. Paris: OECD Publishing. <https://doi.org/10.1787/c74f03de-en>
- OECD. 2024. *Education Policy Outlook 2024: Reshaping Teaching into a Thriving Profession from ABCs to AI*. Paris: OECD Publishing. <https://doi.org/10.1787/dd5140e4-en>
- Pasquale, Frank. 2015. *The Black Box Society: The Secret Algorithms That Control Money and Information*. Cambridge, MA: Harvard University Press. <https://doi.org/10.4159/harvard.9780674736061>
- Pelletier, Kathe, Mark McCormack, Nicole Muscanell, Jamie Reeves, Jenay Robert, and Nichole Arbino. 2024. *2024 EDUCAUSE Horizon Report: Teaching and Learning Edition*. Boulder, CO: EDUCAUSE. <https://library.educause.edu/resources/2024/5/2024-educause-horizon-report-teaching-and-learning-edition>
- Perrigo, Billy. February 2025. "Exclusive: The British Public Wants Stricter AI Rules Than Its Government Does." *TIME*. <https://time.com/7213096/uk-public-ai-law-poll/>
- Poell, Thomas, David B. Nieborg, and Josette van Dijck. 2019. "Platformisation." *Internet Policy Review* 8 (4): 1–13. <https://doi.org/10.14763/2019.4.1425>
- Polanyi, Michael. 1966. *The Tacit Dimension*. Chicago, IL: University of Chicago Press.
- Popper, Karl. 2002. *The Logic of Scientific Discovery*. London: Routledge. <https://doi.org/10.4324/9780203994627>

- Readings, Bill. 1996. *The University in Ruins*. Cambridge, MA: Harvard University Press. <https://archive.org/details/universityinruin0000read/page/n5/mode/2up>
- Risko, Evan F., and Sam J. Gilbert. 2016. "Cognitive Offloading." *Trends in Cognitive Sciences* 20 (9): 676–88. <https://doi.org/10.1016/j.tics.2016.07.002>
- Robison, Kylie. April 2025. "Most Americans Don't Trust AI — or the People in Charge of It." *The Verge*. <https://www.theverge.com/ai-artificial-intelligence/644853/pew-gallup-data-americans-dont-trust-ai>
- Schellekens, Philip, and David Skilling. October 2024. "Three Reasons Why AI May Widen Global Inequality." *Center for Global Development*. <https://www.cgdev.org/blog/three-reasons-why-ai-may-widen-global-inequality>
- Selwyn, Neil, Thomas Hillman, Lena Ewertsson, and Janna Dilleng. 2020. "What's Next for Ed-Tech? Critical Hopes and Concerns for the 2020s." *Learning, Media and Technology* 45 (1): 1–6. <https://doi.org/10.1080/17439884.2020.1694945>
- Silver, David, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, et al. 2016. "Mastering the Game of Go with Deep Neural Networks and Tree Search." *Nature* 529 (7587): 484–89. <https://doi.org/10.1038/nature16961>
- Škare, Marinko, Beata Gavurova, and Sanja Blažević Burić. 2024. "Artificial Intelligence and Wealth Inequality: A Comprehensive Empirical Exploration of Socioeconomic Implications." *Technology in Society* 79: 102719. <https://doi.org/10.1016/j.techsoc.2024.102719>
- Smart, Paul R. 2017. "Extended Cognition and the Internet: A Review of Current Issues and Controversies." *Philosophy & Technology* 30 (3): 357–90. <https://doi.org/10.1007/s13347-016-0250-2>
- Smith, Nicholas J. J. 2008. *Vagueness and Degrees of Truth*. Oxford: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199233007.001.0001>
- Sparrow, Betsy, Jenny Liu, and Daniel M. Wegner. 2011. "Google Effects on Memory: Cognitive Consequences of Having Information at Our Fingertips." *Science* 333 (6043): 776–78. <https://doi.org/10.1126/science.1207745>
- Sułkowski, Łukasz, and Michał Zawadzki. 2016. "Corporate University: A Critical Approach." *Przedsiębiorczość i Zarządzanie* 17 (1): 111–30.
- Tartt, Donna. 1992. *The Secret History*. New York: Knopf. <https://archive.org/details/secrethistory0000tart>
- Thierer, Adam. 2009. "The Pacing Problem, the Collingridge Dilemma and Technological Determinism." Working Paper. Mercatus Center at George Mason University. *Tech Liberation*. <https://techliberation.com/2018/08/16/the-pacing-problem-the-collingridge-dilemma-technological-determinism/>
- Tompros, Louis. January 2025. "Is the Law Playing Catch-up with AI?" *Harvard Law Today*. <https://hls.harvard.edu/today/is-the-law-playing-catch-up-with-ai/>

- Turing, Alan M. 1950. "Computing Machinery and Intelligence." *Mind* 59 (236): 433–60. <https://doi.org/10.1093/mind/LIX.236.433>
- UNESCO. 2023. "Guidance for Generative AI in Education and Research." Paris: United Nations Educational, Scientific; Cultural Organization. <https://unesdoc.unesco.org/ark:/48223/pf0000386693>
- Victor, Audêncio. 2025. "Artificial Intelligence in Global Health: An Unfair Future for Health in Sub-Saharan Africa?" *Health Affairs Scholar* 3 (2): qxaf023. <https://doi.org/10.1093/haschl/qxaf023>
- Wegner, Daniel M. 1987. "Transactive Memory: A Contemporary Analysis of the Group Mind." In *Theories of Group Behavior*, edited by Brian Mullen and George R. Goethals, 185–208. New York: Springer. https://doi.org/10.1007/978-1-4612-4634-3_9
- Wei, Jason, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, et al. 2022. "Emergent Abilities of Large Language Models." *Transactions on Machine Learning Research*. <https://openreview.net/forum?id=yzkSU5zdwD>
- Whitehead, Alfred North, and Bertrand Russell. 1910–1913. *Principia Mathematica*. Cambridge: Cambridge University Press. <https://archive.org/details/alfred-north-whitehead-bertrand-russel-principia-mathematica.-1>
- Williamson, Ben, and Anna Hogan. July 2020. "Commercialisation and Privatisation in/of Education in the Context of Covid-19." *Education International*. <https://www.ei-ie.org/en/item/23423:the-edtech-pandemic-shock-by-ben-williamson-anna-hogan>
- Williamson, Ben, Janja Komljenovic, and Kalervo N. Gulson. 2022. "Datafication, Metrics, and the Governance of Education." *Learning, Media and Technology* 47 (1): 1–6.
- Yeralan, Sencer. May 2022. "Academic Quality – Some Thoughts." Global Academics Coalition. <https://doi.org/10.5281/zenodo.18364572>
- Yeralan, Sencer, and Laura Ancona Lee. 2023. "Generative AI: Challenges to Higher Education." *Sustainable Engineering and Innovation* 5 (2): 107–16. <https://doi.org/10.37868/sei.v5i2.id196>

Index

- academic freedom 25, 31
- academic leadership 51, 54, 62
- accountability 10, 25, 29, 35, 42, 78, 86, 93, 97–98, 100–101, 110
- accreditation 46–47
- Adams, Douglas 53
- Adriaansen, Robbert-Jan 12
- agency 4, 15, 25–26, 28–29, 34, 48, 71, 86, 97, 101
 - distribution of 25–26
- anthropomorphism 6
- artificial intelligence
 - AI-detection 62, 108, 110
 - as “black box” 27
 - definitions of 3–4, 69
 - “generative AI” 3, 15–16, 25, 32, 39, 44–45, 54, 61–63, 65, 81–83, 85–87, 91, 93, 103–104, 107–110
 - metaphors associated with 71–73
 - statistical regularity xx, 5
 - symbolic approaches to 6
- assessment and grading 18, 33–34, 39, 50, 77, 82–83, 93, 96, 107–110
 - assessment redesign 108–109
- audit 10, 26–27, 48
- auditability 13, 19
- authorship 15, 32–33, 72, 93, 108, 110
- automation 3, 16–17, 33, 40–42, 47–48, 50, 52, 88, 96–97, 108
- autonomy 9, 18, 33

- Ball, Stephen J. 10, 100
- bias 16, 51, 71–72, 86, 96–97
- Biesta, Gert 11, 95
- Bildung* 9, 13–14
- Bok, Derek 11, 13, 55

- Cartesian model of mind 6
- central processing unit (CPU) xix
- cognition xix, xx, 5, 28, 43, 54. *See also* extended mind hypothesis
- colonialism 42
- compliance 10, 13, 34, 75, 93, 109–110
- computation xix, xxi, 3–4, 105
- consent 26–27
- COVID-19 pandemic 16, 57
- Crawford, Kate 13, 61, 99

- “Dark Academia” 12
- dashboard 10, 12, 14, 18, 99
- data extraction 17
- datafication 13–14, 19
- data governance 62, 77–78
- data ownership 18
- Davis, Ernest 97
- digitalization 17
- Dijk, Josette van 17, 82, 97
- Dilleng, Janna 15, 63, 94, 98
- Dreyfus, Hubert L. 27
- Dreyfus, Stuart E. 27

- EDUCAUSE 109
- efficiency xx, 10, 13, 15, 17–18, 21, 26, 35, 42, 48, 50, 53, 78, 91, 96–97, 101
- emotion 7, 96
- employment xxi, 11, 47, 57, 95
- engagement 10, 18, 26, 33–34, 40, 43, 48, 96, 99
- engineering xiv, xx, 11, 43, 49–52, 54, 70, 86
- epistemology xx, 11, 16, 18–19, 34, 48, 61–63, 65, 72–73, 76, 92, 94, 96–97, 109–110
- equity 28–29, 47, 51, 85, 87, 92

- Espeland, Wendy Nelson 10, 85, 100
 ethics 26, 28, 49, 54, 69, 71, 73, 76, 99–100
 Ewertsson, Lena 15, 63, 94, 98
 excellence 10
 expertise xxi, 15, 27, 32, 51, 83, 93
 explainability 26–27
 extended mind hypothesis 54
 faculty xx, 9–10, 13, 17, 19, 25–27, 32, 39–40, 43, 45–46, 57, 62–63, 82–83, 86–87, 93, 109
 foresight 35, 101
 Foucault, Michel 70
 Frege, Gottlob 70
 geopolitics 34, 41, 50, 100
 Gilliard, Chris 16, 87, 100
 governance viii, xi, 13, 16, 18–19, 26, 35, 39–40, 46–47, 62–64, 69, 73, 75, 77–78, 81–83, 85–86, 93, 98–99, 101, 109–110
 grading 3, 35, 99
 Grafton, Anthony 12
 Gulson, Kalervo N. 13, 97
 Han, Byung-Chul 70
 Hebb, Donald O. 4
 Heidegger, Martin 70
 Hern, Alex 44, 98
 Hillman, Thomas 15, 63, 94, 98
 Hogan, Anna 86
 Humboldt, Wilhelm von 9–10
 incentives 16, 25–26, 41, 72–73, 76, 100
 infrastructure 4, 12, 17–19, 32, 34–35, 42, 50–51, 61, 83, 92, 97, 99–100, 105
 institutional drift 13–14
 institutional mission 4, 9–10, 18, 31–34, 39, 47, 75, 77, 82, 94
 interoperability 17, 34
 intuition 27–28
 Kasneci, Enkelejda 96
 Keefe, Rosanna 101
 Kizilcec, René F. 86
 Komljenovic, Janja 13, 97
 labor xx, xxi, 3, 26, 93
 Latour, Bruno 70
 legitimacy 40, 62, 72, 76, 86, 100
 literacy 43, 47, 61, 65, 73, 86, 91–92, 109
 Machiavelli 100
 The Prince 100
 managerialism 10, 12–13
 Marcus, Gary 97
 Marginson, Simon 11, 96
 McCulloch, Warren S. 4
 memory and memorization xix, xx, 43–44, 54, 105
 Menand, Louis 14, 95
 meritocracy 53
 metrics 9–10, 12–14, 17–19, 26, 70–71, 82, 85, 98–99
 Michelson-Morley experiment 48
 misalignment 26, 62
 neoliberalism 10, 12–13, 100
 neural networks 3–6
 Nieborg, David B. 17, 82, 97
 Nussbaum, Martha C. 14, 101
 OECD 75, 108–110
 ontology 6, 69, 71
 opacity 18, 27–28, 32, 101
 optimization xx, 5, 7, 13, 17–18, 34–35, 41, 49–50, 73, 97, 108
 Pasquale, Frank 93, 101
 performativity 12, 100
 perplexity 70, 108
 Pitts, Walter 4
 plagiarism 99, 107
 platform 12, 16–19, 33–34, 40–41, 46, 51, 62, 82, 96–100
 platformization 13–15, 17–19, 34, 86
 Plato 43
 Poell, Thomas 17, 82, 97

- Polanyi, Michael 27
- policy xx, 4, 7, 10–11, 17, 26, 39–40, 43, 46–49, 51, 62–64, 71, 73, 75–78, 81, 85–87, 91–92, 100–101, 108–110, 114
- Popper, Karl 70
- procurement 62, 78, 87, 99
- productivity 10, 12, 17, 95, 109
- rankings 10, 32
- Readings, Bill 10, 100
- realpolitik* 99
- reflexivity 63, 82
- regulation 16, 48, 71, 73, 87, 93, 98
- Russell, Bertrand 70
- Sauder, Michael 10, 85, 100
- Selwyn, Neil 15–16, 63, 94, 98
- sorites paradox 55, 101
- sovereignty 29, 34, 48, 86, 97
digital 97–98
- STEM disciplines 11, 54, 95
and marginalization of the humanities 11
- stewardship 31, 45, 76, 100
- surveillance 3, 13, 16, 18, 26, 42, 62, 86, 96–97, 100
- systems science 54
- Tartt, Donna 12
The Secret History 12
- transformation 10, 17, 42, 46, 55, 71–72, 109
- transparency 10, 19, 25–27, 32, 51, 71, 93, 98, 110
- trust 26–27, 31, 40, 58, 63, 83, 92, 97, 117
- Turing, Alan 5, 70
- Turing Test 5
- uncertainty 15, 25, 40, 43, 53, 61, 63–64, 76, 78–79, 109
- UNESCO 75, 108, 110
- vendors 18–19, 26, 33, 35, 40, 48, 78, 86, 98, 100
- Williamson, Ben 13, 86, 97, 99

About the Team

Alessandra Tosi was the managing editor for this book.

Adèle Kreager and Ewan Vellinga proof-read this manuscript. Adèle compiled the index.

Jeevanjot Kaur Nagpal designed the cover. The cover was produced in InDesign using the Fontin font.

Annie Hine typeset the book in InDesign. The main text font is Tex Gyre Pagella and the heading font is Californian FB.

Jeremy Bowman produced the PDF, paperback, and hardback editions and created the EPUB.

The conversion to the HTML edition was performed with epublius, an open-source software which is freely available on our GitHub page at <https://github.com/OpenBookPublishers>

Lucy Barnes was in charge of marketing.

This book was peer-reviewed by Prof. Michael Adrian, University of Illinois, and an anonymous referee. Experts in their field, these readers give their time freely to help ensure the academic rigour of our books. We are grateful for their generous and invaluable contributions.

This book need not end here...

Share

All our books — including the one you have just read — are free to access online so that students, researchers and members of the public who can't afford a printed edition will have access to the same ideas. This title will be accessed online by hundreds of readers each month across the globe: why not share the link so that someone you know is one of them?

This book and additional content is available at
<https://doi.org/10.11647/OBP.0525>

Donate

Open Book Publishers is an award-winning, scholar-led, not-for-profit press making knowledge freely available one book at a time. We don't charge authors to publish with us: instead, our work is supported by our library members and by donations from people who believe that research shouldn't be locked behind paywalls.

Join the effort to free knowledge by supporting us at
<https://www.openbookpublishers.com/support-us>

We invite you to connect with us on our socials!



BLUESKY
[@openbookpublishers](https://openbookpublishers.com)
[.bsky.social](https://openbookpublishers.com)



MASTODON
[@OpenBookPublishers](https://openbookpublishers.com)
[@hcommons.social](https://openbookpublishers.com)



LINKEDIN
[open-book-publishers](https://openbookpublishers.com)



INSTAGRAM
[openbookpublishers_](https://openbookpublishers.com)

Read more at the Open Book Publishers Blog

<https://blogs.openbookpublishers.com>


You may also be interested in:

Distributing Knowledge

Openness, Equity, and Higher Education Transformation

Richard F. Heller

<https://doi.org/10.11647/OBP.0520>



Stories of Hope

Reimagining Education

Edited by Sandra Abegglen et al.

<https://doi.org/10.11647/OBP.0462>




Higher Education for Good

Teaching and Learning Futures

Edited by Laura Czerniewicz and Catherine Cronin

<https://doi.org/10.11647/OBP.0363>



Ethics of Socially Disruptive Technologies

An Introduction

Edited by Ibo van de Poel et al.

<https://doi.org/10.11647/OBP.0366>



ARTIFICIAL INTELLIGENCE IN HIGHER EDUCATION A PHENOMENOLOGICAL VIEW

EDITED BY SENCER YERALAN

Artificial intelligence is increasingly embedded in the everyday practices of higher education, shaping assessment, governance, labor, and institutional legitimacy. Rather than presenting a technical guide or policy checklist, this volume instead offers a reflective, multi-voiced examination of what AI means for higher education's purpose, identity, and future. Its phenomenological grounding shifts the focus from operational questions of implementation to deeper inquiries into how AI reshapes institutions, knowledge, and the academic self.

Drawing on historical and critical perspectives, the book interrogates AI as both mirror and accelerant of long-standing challenges: inequity, market-driven logics, and the erosion of slow, critical learning. Spanning geopolitical contexts and institutional types, it embraces pluralism over consensus, showing that AI will not transform all universities in the same way. Narrative interludes humanize these themes, revealing the anxieties, ambiguities, and hopes of those living through this transition.

Building on the work of Richard Heller on the distributed university and knowledge equity, the book situates AI within broader structural issues such as corporatised knowledge economies, managerialism, and unequal access to educational and research opportunities. At the same time, it highlights emerging possibilities—from open educational resources and equitable research practices to decentralised digital infrastructures—that can contribute to more ethical and resilient institutional arrangements.

Neither prescriptive nor simplistic, this book is intended as a catalyst for leaders, policymakers, and reflective practitioners seeking to navigate AI with wisdom rather than haste. It argues that the future of higher education will be shaped less by technological sophistication than by the clarity with which institutions articulate their values, responsibilities, and commitments to the public good.

This is the author-approved edition of this Open Access title. As with all Open Book publications, this entire book is available to download for free on the publisher's website. Printed and digital editions, together with supplementary digital material, can also be found at <http://www.openbookpublishers.com>.

Cover photo by Sencer Yeralan and design by Jeevanjot Kaur Nagpal

