


Theater as Data

Computational Journeys into Theater Research

Miguel Escobar Varela



THEATER AS DATA

Theater as Data

COMPUTATIONAL JOURNEYS
INTO THEATER RESEARCH

Miguel Escobar Varela

University of Michigan Press

ANN ARBOR

Copyright © 2021 by Miguel Escobar Varela
Some rights reserved



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. *Note to users:* A Creative Commons license is only valid when it is applied by the person or entity that holds rights to the licensed work. Works may contain components (e.g., photographs, illustrations, or quotations) to which the rightsholder in the work cannot apply the license. It is ultimately your responsibility to independently evaluate the copyright status of any work or component part of a work you use, in light of your intended use. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>

For questions or permissions, please contact um.press.perms@umich.edu

Published in the United States of America by the
University of Michigan Press
Manufactured in the United States of America
Printed on acid-free paper
First published August 2021

A CIP catalog record for this book is available from the British Library.

Library of Congress Cataloging-in-Publication data has been applied for.

ISBN 978-0-472-07479-2 (hardcover : alk. paper)
ISBN 978-0-472-05479-4 (paper : alk. paper)
ISBN 978-0-472-12863-1 (OA)

DOI: <https://doi.org/10.3998/mpub.11667458>

This open access version made available by the National University of Singapore.

Cover photo: From the performance *Pixel* by Adrien M & Claire B and CCN de Créteil et du Val-de-Marne / Compagnie Käfig—Mourad Merzouki, 2014. Photo © Raoul Lemerrier.

Acknowledgments

I would like to thank Yong Li Lan and the rest of the Asian Intercultural Digital Archives (AIDA) team at the National University of Singapore (NUS) for providing me with a conceptual home from which to embark on the study of theater as data, and for helping me understand the problems of turning theater into digital artefacts.

Paul Rae, Itty Abraham, Lonce Wyse, Eric Kerr, Maiya Murphy, Felipe Cervera, and the late John Richardson all read portions of this book and gave me invaluable feedback. At NUS, I would like to thank the Digital Cultures reading group at the Faculty of Arts and Social Sciences and the Quantitative Reasoning faculty members at the interdisciplinary University Scholars Programme for their conversations, feedback, and support. The University Scholars Programme and the Faculty of Arts and Social Sciences both provided financial support for making this book available through an open access and I am very grateful for this.

While this is not a book about the Indonesian performing arts, the traditions of Java do feature in some parts of this book. My understanding of these traditions has been shaped by many people, but I am especially indebted to Jan Mrázek, Eddy Pursubaryanto, Bima Slamet Raharja, Bernard Arps, Kathryn Emerson, Novi Marginingrum, Ki Catur Kuncoro, Dewi Ambarwati, Ki Aneng Kiswanto, Ki Ananto Wicaksono, Matthew Cohen, the late Ki Slamet Gundono, and the late Ki Ledjar Soebroto for bringing me closer to the performing arts of Java. I would also like to thank Imam Maskur for allowing me to use the Kluban data on *wayang kulit* performances. In Singapore I would like to thank The Necessary Stage, Centre 42, and *The Flying Inkpot* for letting me use their data.

The journeys of this book began in earnest when Gea Oswah Fatah Parikesit at the Department of Engineering Physics, Gadjah Mada University in Indonesia first invited me to work with him in a collaborative project

in 2014. It was then that I first saw how theater scholars can have productive collaborations with scientists. Thank you *mas* Gea, for your support, intellectual generosity, and for your willingness to engage with researchers across the disciplinary divide. I am also thankful to Andrew Schauf and Luis Hernández-Barraza, with whom I continued this collaborative journey and learned more about network science and biomechanics.

The global digital humanities community has accompanied and supported me in this methodological wayfaring. I am especially thankful to Clarisse Bardiot, Alex Gil, Isabel Galina, Sean Pue, Padmini Ray Murray, David Wrisley, Ernesto Priani, Christof Schöch, Radim Hladík, Frank Fischer, and Paul Spence. I'm also very happy to be part of the incipient DH community in Singapore, and I'm grateful for the conversations with, and support from, Kenneth Dean, Lim Beng Choo, Andrea Nanetti, Michael Stanley-Baker, Feng Yikang, Sayan Bhattacharyya, and the Digital Scholarship team at the NUS Libraries.

It has been an honor to work with the team at University of Michigan Press. Thank you to LeAnn Fields for believing in this project.

Last but not least, I wish to thank my family in near and far places. Thanks to my parents for being such inspiring and supportive role models on how to be academic researchers in today's world. Thanks to Dari for her useful feedback on interaction design and visualizations. My wife, Yingting, provided me with moral support throughout the writing of this book, and helped me better understand how scientists use data in their daily work. Sofia is too young to realize this, but she brought enormous joy to the process of writing the words that follow.

Contents

Introduction: In Pursuit of Theater's Digital Traces	I
--	---

Part 1: Pre-departure Reflections

1 Toward a More Nuanced Conversation on Methodology	23
2 The Roles of Statistics	41
3 The Roles of Visualizations	57

Part 2: Guided Tours

4 Words as Data	75
5 Relationships as Data	94
6 Motion as Data	116
7 Location as Data	141

Part 3: Ensuring the Journeys Continue

8 The Imperative of Open and Sustainable Data	163
9 The Roles of Software Programming	180

Appendix A: Data Biographies	189
Appendix B: Technical Glossary	193
References	195
Index	219

Digital materials related to this title can be found on the Fulcrum platform via the following citable URL: <https://doi.org/10.3998/mpub.11667458>

Introduction

In Pursuit of Theater's Digital Traces

I am on my way to the theater, and I can't stop generating data. When I use public transport, consult a map on my phone, and order a coffee at the theater foyer, I am leaving digital breadcrumbs of my activities behind me. The data points that have resulted from these activities will be aggregated with those generated by thousands of other people. Collectively, these datasets will shape the experience of riding buses, consulting maps, and ordering coffee for myself and others. From a technological perspective, it is doubtlessly exciting that data can be applied to such diverse aspects of life. But there are many complex ethical and political corollaries to this intense collection and use of data. Data is not only useful for optimizing bus routes and coffee pricing, it can also be used to increase social inequality (O'Neil 2016, 94) and to sway elections (Morgan 2018). We live in a world that is made *for* data and *from* data, as data shapes us in more ways than one. As Ribes and Jackson (2013) note, "we have entered into a symbiotic relationship with data—remaking our material, technological, geographical, organizational, and social worlds into the kind of environments in which data can flourish" (52).

I finish my coffee, switch off my phone, and enter the performance space. It seems like I have entered a data-free sanctuary and left the world of data and its ominous potentials behind. After all, the creative decisions that shape performances are most likely not driven by data. While data-driven content is becoming increasingly common in films and television (Suri and Singh 2018), and data plays a role in the marketing and selection of performances, I can't find any documented instance where the decisions about what happens on stage were shaped by data, not even in commercial productions. There are many performances which are *about* data, such as *Of All of the People in All of the World* by Stan's Café (2008),

where grains of rice are used to represent population statistics. However, creative decisions (lighting, casting, modes of interacting with the audience) are not *decided* by data. This is certainly not the case in the intimate, experimental productions and cultural performances to which the bulk of scholarly attention is devoted to. Likewise, most theater scholarship is focused on specific, thoroughly contextualized performance events. Theater and performance scholars deploy interpretive methods to study the performances that interest us. We don't usually make use of data—in the information science sense—in order to seek answers to our questions or to orient our discussions. That is, until recently.

With the advent of the digital humanities (henceforth DH), and the pioneering work of several researchers, using data is becoming more common in theater scholarship. This does not mean surrendering our methods to the nefarious potentials of data, but combining data and critical insight to offer fresh perspectives. For example, Holledge et al. (2016) have used a comprehensive dataset of *A Doll's House* performances to offer new insights into the global spread of Ibsen's masterpiece. Using network analysis, Trilcke et al. (2015) have identified previously unreported patterns in the interaction of characters in German drama across several centuries. Schöch (2017) has identified clear genre differences in the words used by classical French playwrights. Caplan (2017) has used network data to show that Yiddish theater performers are far more influential in popular entertainment than they are usually credited for. Wiesner and Stalnaker (2015) were able to estimate conceptual metaphors in dance from context-free movement patterns. And Miller (2017) has used Broadway data to show that the performance runs of musicals have consistently increased in length over time.

These are just some examples of work carried out at the intersection of theater and DH, an area with a long history. In an overview of this intersection, Spence (2013) outlines many projects that were decades in the making. In her overview of this area, Leonhardt (2014) describes projects that fit into seven categories: library and archival projects, projects on individual playwrights or actors, projects on theater architecture, projects on dance, larger network projects, applied DH in theater research and education/experiential DH projects in performance studies, and projects useful for or related to theater and performance research. Caplan (2015) distinguishes between four types of projects: digital archives and editions, digital theatrical environments, digital visualizations, and digital databases. And Bay-Cheng (2017) classifies existing projects according to three major types of research activities: collection, analysis, and dissemination.

Regardless of how one slices this interdisciplinary cake, one thing is certain: there is a great variety of projects carried out at the crossroads of theater and DH. This variety will feature in this book, but my main objective is not to chronicle this growing area. Rather, I focus on a more specific question: what does it mean to research theater in terms of data? I am interested primarily in studying theater as an event rather than as dramatic literature (an area that has been extensively studied elsewhere in DH). I believe, with Peggy Phelan (1993), that theater “constitutes itself through disappearance” (146). As performances blink in and out of existence, they leave traces behind. These include program booklets, play schedules, cast lists, critical reviews, and videos which can be transformed into digital data. Thus, my project is also aligned to what Manovich (2020) calls *cultural analytics*. While DH is a more expansive term that includes everything from the digitization of cultural heritage to the development of metadata standards, *cultural analytics* tends to be more narrowly focused on the computational analysis of large cultural datasets. This might well include the analysis of textual materials and historical data, but *cultural analytics* as practiced by Manovich and his lab since 2005 gives primacy to the analysis of contemporary, visual culture at a scale of millions of objects. In my projects, I also consider visual and other nontextual dimensions of theater, but some of my sources are historical rather than contemporary and my datasets are never as extensive as those used by Manovich.

What Is Data?

This book takes a narrow definition of data as “potential information.” This comes from information science, a field that assumes data is something that needs to be automatically processed to be useful (Pomerantz 2015). Processed data becomes information. Applied information becomes knowledge, and reasoned knowledge becomes wisdom. This hierarchy was famously inspired by T. S. Eliot’s play *The Rock* (1934):

Where is the wisdom we have lost in knowledge?
Where is the knowledge we have lost in information? (7)

This quote is often found in information science books. When theater scholars turn to other fields for definitions, we might find some comfort in being sent back to look at a theater play. A longer definition of data, as understood in this book, is *potential information in digital form that needs to be processed by a machine for it to become meaningful*. But the important question,

as Christine Borgman (2015) notes, is not “what are data?” but “when are data?” (17). The same digital object could be treated as information or as data. Take, for example, a theater script. If one close reads a theater play, one is not treating it as data, even if the text is stored in a digital format. But if one uses software to calculate the relative frequencies of a word in this text, then the same object is treated as data. The examples of data in this book are digital texts, videos, motion capture, geographical coordinates, and timestamps. In all cases, software is used to analyze and visualize this data. For example, the videos are not viewed by a human observer, but treated as sequences of two-dimensional images made of pixels. The distribution of pixels is used to make inferences about the ways actors move around the stage.

This narrow definition of data will help bring focus to the arguments in this book, but I will not ignore the many problems it poses. For example, when data is processed, it is often described as “flowing” from one form into another. However, as Jonathan Bollen (2017) notes, “if the metaphor applies at all, datasets have the fluidity of plastic; they flow when prodded, pushed, or pressed, when the heat is turned up and when placed under stress” (619). Data is never fully free of interpretation. As Lisa Gitelman (2013) writes: “Data are familiarly ‘collected,’ ‘entered,’ ‘compiled,’ ‘stored,’ ‘processed,’ ‘mined,’ and ‘interpreted’ [. . .] less obvious are the ways in which the final term in this sequence—interpretation—haunts its predecessors” (3). The history of pre-digital data further illustrates some of the problems associated with data. Daniel Rosenberg (2013) argues that data has often referred to aspects of an argument taken for granted for the sake of discussion or analysis. “The semantic function of data,” he tells us, has always been “specifically rhetorical” (8, original emphasis). In one of its earlier recorded uses, Euclid differentiates the given quantities of *data* from the *questia* of the quantities sought (Rosenberg 2013, 19). These views evolved during the seventeenth century, when the term became more common:

[In] philosophy and natural philosophy, just as in mathematics and theology, the term “data” functioned to identify that category of facts and principles that were, by agreement, beyond argument. In different contexts, such agreement might be based on a concept of self-evident truth, as in the case of biblical data, or on simple argumentative convenience as in the case of algebra, given $X=3$, and so forth. The term “data” itself implied no ontological claim. In mathematics, theology,

and every other realm in which the term was used, “data” was something given by the conventions of argument. Whether these conventions were factual, counter-factual, or arbitrary had no bearing on the status of givens as data. (Rosenberg 2013, 20)

The prophets of our own current techno-scientific era—scientists, managers, and entrepreneurs—often invite us to treat data as something that is beyond argument. A dangerous corollary of this attitude is the belief that data is always true or accurate. It is often assumed that data is the product of a semi-automatic, or at least systematic process. For example, data can be collected by scientific instruments or sensors. But for the sensors to work in a particular way, a series of assumptions about the phenomenon of interest needs to be made: “data need to be imagined as data to exist and function as such, and the imagination of data entails an interpretive base” (Gitelman 2013, 3 original emphasis). Or, as Borgman (2015) puts it “conceptualizing something as data is itself a scholarly act” (xviii). Data artist Jer Thorp suggests that data is never inseparable from a system where collection, computation, and representation inform each other. The following excerpt from his blog vividly captures the impact of choice on such a system:

Whenever you look at data—as a spreadsheet or database view or a visualization, you are looking at an artifact of such a system. What this diagram doesn’t capture is the immense branching of choice that happens at each step along the way. As you make each decision—to omit a row of data, or to implement a particular database structure or to use a specific color palette you are treading down a path through this wild, tall grass of possibility. It will be tempting to look back and see your trail as the only one that you could have taken, but in reality a slightly divergent you who’d made slightly divergent choices might have ended up somewhere altogether different. (Thorp 2017, n.p.)

In performance studies, Schechner (2013) famously distinguished between “as/is” performance, where certain events are generally thought to be performances whereas others (from political events to sporting competitions) could be thought of as performances (32). Following Gitelman, Borgman and Thorp, I argue that data never *is*—no natural category of data exists without preconditions. But many things can be considered as data, including theater performances. Perhaps all forms of theater

research require us to have an abstract model of what performance is, but conventional scholarly writing allows us to refine and revise definitions as we go along, through the conceptual process that Kwa (2011) describes as a “changing semantic network.” In contrast, data forces us to make definitions explicit. For example, let’s imagine a project that wants to show the average number of actors per performance throughout several decades for a given place. For this analysis, we will need an explicit definition of actor in order to produce the required dataset.

This definition doesn’t need to encompass all the essential qualities of an actor, but it must be narrow and implementable. In other words, it should consist of an actionable rule that enables the systematic creation of a dataset. For example, a research team might define actors, for the purpose of a project, as the people listed as cast members in program booklets. In some cases, this definition will be relatively uncontroversial. But other types of productions will be fraught with more intense disagreement. What about a forum theater performance where the audience members become *spect-actors*? Should each audience member be included as an actor in the dataset? Or only the resident company actors? Or should a separate category be created for *spect-actors*? All of these options are problematic, but building the dataset will force its makers to take a stance and decide on definitions of actor. It will be important to make choices that will enable the best inferences to be drawn from the data. For example, if we are interested in analyzing the careers of professional actors, we can leave the *spect-actors* out. Our data will not be a fully accurate model of reality, but it will be sufficient to find verifiable patterns, through the process that Craig and Greatley-Hirsch (2017) call “principled generalization” (7).

It is, however, also possible to use data in a way that takes this definitional problem head on. Instead of deciding at the outset whether *spect-actors* should be considered as actors, the goal of a project might be to frame this as a question. A data visualization might be used to explore the consequences of taking *spect-actors* as actors, showing how different solutions to this definitional problem will yield a completely different number of actors. An interactive implementation of this project might enable users to reach their own conclusions on the appropriate definition of actors. Perhaps the data points in the visualization can be linked to videos, where users can then see the way in which specific *spect-actors* interact with resident company actors. In this example, data is not used to answer a question, but to reframe the question of what it means to be an actor.

This example shows that there are two ways of working with data. One brackets off assumptions to find patterns, while the other uses data to problematize these assumptions. The former offers new answers, the latter poses the question in fresh ways. I call the first approach *data-driven* and the second, *data-assisted*. These two terms form the conceptual backbone of the present book.

Data-Driven and Data-Assisted Methodologies

Data-driven methodologies use computers to reason under formal constraints. Knowledge is conceived as a rational project, where logic and formal representations are used in the pursuit of replicable conclusions that disprove previous ideas. In contrast, data-assisted frameworks use computers to “imagine in different ways” (Harrell 2013, 79). Data opens speculative and subjective avenues for interpretation, which contribute to, but don’t necessarily displace, previous interpretations. Data-driven methodologies aim at producing consistent explanations that are “hard to vary,” as physicist David Detusch (2009) puts it. In turn, data-assisted methodologies produce a multiplicity of explanations.

In data-driven methodologies, we use data to answer questions. We create a formal representation of a question and automate a sequence of procedures to provide an answer. The criteria for evaluation are defined beforehand, and the answer is measured against these criteria. Researchers working in this mode know that framing a question in a narrow way leaves out many aspects of the object of interest, but they believe that there is something to be gained from this reduction. For example, they are able to identify features across a different scale (thousands of performances rather than a single instance). This also enables the possibility of countering our intuitions. This way of working with data brings evenness of attention, creating “an opportunity to be surprised: to back something other than the sentimental favorite and to reverse consensus views” (Craig and Greatley-Hirsch 2017, 8). Using data in this way, we keep our intuitions in check, and independently verify deeply held beliefs. While the answers aim to provide the best way to characterize a phenomenon, they are not final. They can be revised and disproved by further research. They also include probabilistic estimates of their accuracy in order to provide “measured uncertainty” (Craig and Greatley-Hirsch 2017, 3). These estimates can be revised by other researchers as more data becomes available.

In data-assisted methodologies, we use data to transform our view of a problem. The purpose of framing a theatrical event *as data* is not to offer a clear answer but to augment our capacity to think about such event. Data becomes *deformance*, to use the portmanteau of deformance and performance suggested by Lisa Samuels and Jerome McGann (1999) to characterize interpretive transformations in DH. Data, in other words, provides a good defamiliarization strategy. Data-assisted methodologies work well with theatrical problems that might be fundamentally unquantifiable. But in trying to quantify them, as in the spect-actors example above, we find new ways to address interesting aspects of these problems. According to Kramnick (2018), the objective of critical method in the humanities is not to give definite answers but to “keep the problem open, poke around its edges, ask whether it has been framed in the right way, resituate the conversation” (n.p.). This is also what data-assisted methodologies do.

The main distinction between these methodologies is the different criteria they require for evaluating conclusions as useful and valid. The gold standard of data-driven research is replicable, incremental knowledge. To assess data-driven claims, we should ask how likely the results are to be true and whether other independent sources of evidence corroborate these conclusions. In contrast, data-assisted frameworks ask to be judged for their generative capacity, their potential to trigger new questions and bring forth new perspectives. While in data-driven methodologies many different procedures and parameters should yield the same answer, in data-assisted frameworks the same approach should yield multiple answers.

Data-assisted methodologies are closer to a constructivist view of knowledge, where every piece of data and every conclusion is observer-dependent. Data-driven methodologies are closer to a realist epistemology. They seek to make intersubjectively verifiable claims, and conceive empirical research as a difficult, but worthy goal. Data-driven approaches consider the advantages of thinking in terms of aggregates, and data-assisted approaches demand close attention to individual instances. Table 1 summarizes the comparison made thus far between both approaches.

My term *data-assisted* is a nod to the DH community, where the term “computer-assisted criticism” is common (Siemens 2002; Rockwell and Sinclair 2016). Some of the distinctions between data-driven and data-assisted methods have parallels in the epistemological disagreements within the DH community. But this is not a book about DH as a field. It is a book about what it means to use data to study theater. My starting

point, though, is that epistemological discussions arise in any field that deals with data and that a book on data must be a book about methodology. It is disingenuous to think that we can describe data without talking about epistemology.

Thus, an epistemological vein runs through this book, as this is a meditation on how we come to know things about theater and the roles that data might play in this endeavor. For this reason, I have so far been characterizing data-driven and data-assisted approaches as *methodologies* rather than *methods*. A methodology is a framework for estimating the pertinence of given methods and for articulating criteria for evaluating results as useful, appropriate, or correct. A method is a protocol, a series of steps. The same method could be used within data-driven and data-assisted approaches. Take, for example, word frequencies. Within a data-driven perspective, one could use word frequencies to identify the most likely author of a dramatic text of unknown authorship. But word frequencies could also be used to radically transform the experience of reading a text. As Stephen Ramsay (2011) suggests, the font size of each word can be altered to represent the relative frequency of that word in the entire text. This data-assisted example doesn’t aim to provide a specific answer to a clear question, but to enable alternative readings via a computational defamiliarization strategy. Both the authorship attribution study and the defamiliarizing change in font sizes might use the same statistical methods for estimating word frequencies.

Statistics are at the heart of many data projects, regardless of whether they are conceived within a data-driven or a data-assisted framework. The presence of statistics does not in itself determine whether a project is

Table 1. Comparison of data-assisted and data-driven methodologies

Data-driven	Data-assisted
Reasoning under constraints (bracketing assumptions)	Problematizing assumptions
Answers that are hard to vary	Multiplicity of views
Productive reductionism	Deformance (productive distortion)
Replicability and falsifiability	Adding to a history of interpretation
Estimating the likelihood that an answer is correct	Resituating a question
Measured uncertainty	Fundamentally unquantifiable propositions
Realism	Constructivism
Aggregates	Individual cases
Disambiguation	Exploration of ambiguity

aligned with either methodology. Statistics are not always enlisted for the pursuit of objectivity, as they can also be used as defamiliarization strategies. There is, however, one key difference: while descriptive statistics are used in both approaches, inferential or explanatory statistics belong only to data-driven research. The different roles of statistics are described in more detail in chapter 2.

Visualizations are also central to many data projects, but the way in which they are used varies within both methodologies. Data-driven visualizations rely extensively on statistical graphs, using conventions borrowed from the scientific community. Data-assisted visualizations require the invention of new graphical conventions. Data-assisted visualizations aim to resituate conversations, and to communicate how ideological biases shape the data. Visualizations used in computational theater research can be static or interactive, regardless of which methodology is used. But interactivity is particularly useful for data-assisted visualizations that aim to provide context to each data point, as in the example above, where each spect-actors data point is linked to a video-recording. The different roles of visualizations are described in chapter 3.

When distinguishing between data-driven and data-assisted methodologies, it doesn't matter whether a project appears to rely extensively on statistics or not, or the extent to which it uses visualizations. The size of a dataset doesn't matter either: both methodologies can be applied to very small and very large datasets. What matters is *why* statistics and visualizations are used, and *how* they support the argument being made. The important thing, in other words, is which criteria for evaluation are enlisted in the communication of a research project.

There are many differences between data-driven and data-assisted methodologies, but they can also be used concurrently within the same research project, and both can contribute to a more comprehensive understanding of theater history and practice. As noted earlier, computational research exists in a spectrum from realism to constructivism. Data-driven research tends to be closer to realism and data-assisted to constructivism. Some researchers occupy radical positions at the opposite ends of this spectrum, affirming either that "we should strive for a science of culture" or that "data is always observer-dependent and socially constructed." But it is possible to carry out research from a more nuanced position in the spectrum. I believe, with the proponents of *critical realism*, that one can be a realist about some things and a constructivist about others. Critical realism is a metatheoretical perspective from the philosophy of science that

is becoming increasingly common in the study of the social world. While critical realism is not often invoked in the context of DH, I believe that it is consistent with the nuanced ways in which many researchers speak about their work. For example, many proponents of data-driven methods are quick to recognize that there is a certain subjectivity to their projects. The choice of a research topic, for instance, can’t ever be fully derived from objective first principles and always betrays a situated way of looking at the world. Likewise, many proponents of what I call data-assisted research note that the data used must conform to intersubjectively verifiable standards.

In contemporary humanities departments, calling someone a “positivist” is most certainly an insult. I have witnessed many data researchers defending themselves against accusations of positivistic allegiances by merely stating that no, their projects are not positivistic, but without offering further elaboration. I would like to suggest it is more useful to strive for a fine-grained vocabulary, such as the one that critical realism provides. Critical realism is given more attention in chapter 1, but its presuppositions run through the rest of the chapters in this book. For some projects, data-driven and data-assisted perspectives might coexist. But my goal is not to suggest that every research project must combine data-driven and data-assisted perspectives. Rather, my goal is to encourage any data project to aim towards critical realism. For a data-driven project, this means taking history and definitions seriously. For a data-assisted project, it means aiming to show how playful defamiliarizations and multiple perspectives can contribute to a fuller understanding of a complex, shared reality.

Some projects will be better suited for a data-driven perspective, and others for a data-assisted one. In my view, data-driven research is only useful for research questions that fulfill four conditions:

1. There is **low ambiguity** in the definitions of key concepts. Take, for example, the analysis of scene structure in theater. In some dramatic traditions, it is possible to establish consensus on what constitutes a scene. This is the case in Ibsen’s plays. But it would be harder to reach consensus in what constitutes scene boundaries in the work of other playwrights.
2. The area of interest can be thought of as **discrete features** rather than continuous processes. It is possible to analyze methodically the number of productions in which Andre Gregory has

been involved as a director. One could create a visualization of his creative output through time. But analyzing the aesthetics of one of his plays would require interpretive work more suited to data-assisted analysis, or other methods.

3. The data is **arithmetically malleable**. For example, one could measure the distance that two touring theater companies covered in a year. One could add these measurements and conclude that one company traveled twice as far as the second one. But one could not say that one company is twice as original as another, even in cases where the critical consensus is that one company is indeed more original.
4. Data can be combined into an **aggregate**. For certain research projects, it might make sense to take the entire choreographic oeuvre of Akram Khan as a unit. But it might not make sense to consider *migrant theater* as a unit. Perhaps what makes migrant theater interesting as a site of study is the diversity of individual approaches.

These four conditions are related, but they are not equivalent to each other. One could imagine a research area where there is high ambiguity, but researchers can still identify discrete, arithmetically malleable features. That is, many researchers identify discrete features but they don't agree on what they are. Let's take, for example, the plays attributed to Shakespeare. There is some disagreement on which plays he wrote, but plays are still discrete features that are arithmetically malleable. Different researchers can offer competing numbers of plays written by Shakespeare. The analysis of areas such as this, which only meet some of the four criteria, are better served by data-assisted methodologies, provided that discrete features can still be identified (otherwise there is nothing to be turned into data). Cases where none of the criteria can be identified are better served by other types of methodologies that don't use data, such as ethnography or practice-based research. This information is summarized in table 2.

I am sure that some people will disagree with my examples above. The-

Table 2. Criteria for different types of research

Criteria	Data-driven	Data-assisted	Other approaches
Low ambiguity	Yes	Maybe	No
Discrete features	Yes	Yes	No
Arithmetic malleability	Yes	Maybe	No
Aggregability	Yes	Maybe	No

ater and performance studies are, after all, built around *sophisticated disagreement* (Strine, Long, and Hopkins 1990). Perhaps there are researchers who consider the question of Shakespeare authorship settled, scholars of Ibsen who could point me to disputed scene boundaries in his plays, or dance scholars who believe that Akram Khan's work is too diverse to be considered as an aggregate. My proposed criteria are not meant to be taken as inherent properties of a research area, but as articulations of a researcher's perspective. Data-driven research is useful for someone who assumes that the four criteria are met for a given research project. Researchers who believe that such assumption would be misleading should better use a data-assisted methodology. Assuming that the four criteria are met doesn't mean sweeping problems under the rug, and low ambiguity doesn't mean no ambiguity. A researcher can still acknowledge disagreement in her field, while choosing to pursue a data-driven project. This acknowledgment can be communicated together with the results of the research. Or the consequences of this disagreement could be further analyzed through a data-assisted methodology.

In addition to the criteria described above, every data project requires good sources of data. The availability of data will always be connected to the existence of specific records, and this will limit the kinds of questions we can ask. The kinds of data available also signal the existence of specific histories of representation, and might raise many critical issues that should be addressed when using the data. Good data is also *thick data*. We need to know how and why the data was collected, and how it was processed. I use the term *thick data* following Tricia Wang (2016), who extends the anthropological concept of *thick description* to data. Thick data should also include thorough descriptions of our doubts and misgivings in the collection of data, a chronicle of the decisions made and a description of the roads not taken. *Data biographies* are useful for communicating these aspects of data and whenever possible, they should accompany a research project (more on this on chapter 8).

Computational Theater Research: Becoming Technical, Remaining Critical

In this book, I refer to data-driven and data-assisted methodologies under the umbrella term of computational theater research. *Computational* is a fraught term in DH, and some people consider it equivalent to what I call data-driven research. In the natural sciences, computational research relies

on algorithmic models and simulations rather than experimental and observational methods. My usage here differs slightly from these other cases.

I refer to computational research as that which requires computationally intensive processes to transform and analyze data. Some of the projects described in the book could theoretically be done by hand calculations. But this would be so time consuming that it would not enable the exploratory and iterative processes that characterize these research approaches. Thus, I can further refine my definition and describe computational theater research as *the exploratory and iterative use of computers to analyze digital theater data*. I choose the label *computational* over *digital* in order to distinguish the work described here from the study of digital performance (Dixon 2007; Bay-Cheng, Parker-Starbuck, and Saltz 2015; Chatzichristodoulou, Jefferies, and Zerihan 2009; Beyes, Leeker, and Schipper 2017; Causey 2006) and from other types of work at the confluence of DH and theater, such as digital archiving (Bardiot 2015; Carlin and Vaughan 2015; Sant 2017).

Computational theater research, as defined here, doesn't include simulations that enable users to recreate or devise performances in an interactive computer environment (Roberts-Smith et al. 2013; Delbridge and Tompkins 2009). While I give some attention to dramatic texts, I am primarily interested in the analysis of theater as an event, and I don't fully survey the extensive research on the computational analysis of dramatic literature. My project is conceptually similar to what Clarisse Bardiot (2017) terms *theater analytics*, which involves "algorithmic approaches and data visualization" (n.p.). However, her approach is somewhat closer to what I term data-driven research (Bardiot 2015, 2018) and, for me, computational theater research also includes data-assisted projects.

Another reason for choosing the label *computational* is to signal the importance of computer programming for both data-driven and data-assisted theater research. Some people believe that one must be a programmer to carry out DH work, a decidedly divisive attitude. I agree with the more nuanced position championed by Berry and Fagerjord (2017, 38–40). Learning to code is akin to learning to apply a given theoretical framework: knowing how to code equips a scholar with better understanding of the tools they use, even if the scholar is not a professional programmer. I also believe that being a programmer is not a simple binary matter, since there are different levels of expertise.

However, I do believe that those interested in computational theater research would benefit from some experience in programming (and ide-

ally some knowledge of statistics and visualization), regardless of whether they aim to pursue primarily data-driven or data-assisted work. There are epistemological and pragmatic reasons for this. One could perhaps say things about Japanese theater without speaking the language (see Leiter 2012), but having at least some familiarity with the language opens an entirely different realm of research avenues to the scholar of Japanese theater. The same is true for programming—although non-programmers can say many things about software tools, the nature of the questions that can be pursued changes when one is able to follow at least some basic aspects of the technical discussions in a field. From a pragmatic point of view, we also need theater scholars who can code or the future will not be sustainable (as I argue more fully in chapter 9). If we want to build research that is open, durable, and shareable, more of us need to get our hands dirty and do some programming. Otherwise, we are at the mercy of financial resources and tools developed by others. I argue that computational research does not need extensive financial resources, but it does require a growing community of programmer-scholars.

This, however, is not a technical handbook. I will describe all technical procedures in accessible terms, but interested readers can consult the appendices (and the web companion) where technical terms are described in more detail, and where the data used for the excursions can be downloaded by users interested in reproducing my results. For some of the projects described in the book, I have also created step-by-step Jupyter Notebooks, which can be downloaded from the University of Michigan’s Fulcrum website. These notebooks are aimed at users who want to learn more about the technical procedures described in this book. They include explanatory notes, and they are accessible to readers with no programming background who want to peek under the hood and get a sense of the technical devices behind some of the conclusions reported in this book. I don’t expect all readers of this book to become programmers, but I aim to make an argument about why they might consider dipping their toes into programming if they haven’t already done so. I also encourage readers with more programming experience to reuse my data and find new ways of analyzing it beyond what I have initially considered.

Finding a balance between data methods and critical awareness means that we need to be simultaneously makers and observers, cultural critics and data scientists, programmers and scholars. We need deep understanding of the scientific and technological principles of our tools, as much as we need a hermeneutic disposition that can unpack their ideolo-

gies, context, and ethical consequences. A great deal of scholarly attention from the humanities has focused on the critical analysis of data production and shown how data is always contingent on the financial, historical, and cultural contexts in which it is embedded. This type of critical awareness can also refine our data practices, enhancing the ways we collect and analyze data. Critical analysis and software programming can go hand in hand. As the writers of *Data Feminism* (D'Ignazio and Klein 2020) argue, it is possible to build on critical insights to improve our data methods. Many of the contributors to that volume describe themselves as both cultural critics and data scientists. We also need this combination of skills to create good scholarship at the intersection of theater studies and computational methodologies. This also means that we need to include not only data tables and code in our projects, but also thick descriptions of the ways in which the data was collected and analyzed. I believe that a responsible use of data requires us to critically scrutinize the contexts in which the data was produced. We need to acknowledge the ethical, social, and political complexities of the work we do, and the contexts in which we are located. In my own research, I try to demonstrate this by highlighting the cultural background of the performances I study and the conditions under which the data was collected. In my computational theater research projects, I ask narrow, highly specific questions concerning theatrical practices in Indonesia and Singapore. This book is not about my work, but it is informed by it.

This Book's Journeys: A Guided Tour with Excursions

This book is a meditation on methodology, not an analysis of any specific theater tradition or genre. I will report on many projects carried out by research teams around the world—this is primarily a book about other people, about my colleagues and predecessors who have worked at the intersection of theater and data in fields ranging from contemporary Portuguese choreography to adaptations of Shakespeare in Asia. The book is constructed as a guided tour of this incipient but diverse area, and those here represented might not always fully agree with the way I have characterized their work. But my hope is that my proposed vocabulary will lead to productive conversations and perhaps inspire alternative conceptual categories to describe computational research in theater and performance.

In this tour, I systematically apply my proposed vocabulary to characterize a wide range of projects from around the world as data-driven

and data-assisted theater research. I complement this guided tour with short excursions into my own work on Indonesian and Singaporean theater at the end of each chapter in part 2. These are *excursions* rather than fully fledged *case studies* that reference work that I have often published in more detailed formats elsewhere. The excursions don't aim to constitute a coherent analysis of Singaporean or Indonesian theater, but they serve to punctuate the more general guided tour I offer in the book. I don't include a comprehensive analysis of the many issues surrounding theater practice in the region where I live, but every time I embark on these excursions, I emphasize their histories and contexts. I do this to further a methodological tenet: data methodologies are wholly contingent on careful, critical adaptation to a given context.

Excursions into Singapore and Indonesia, two neighboring but fairly different countries, bring these contextual distinctions into sharp relief. Singapore, an island city-state with just over five and a half million inhabitants is the economic hub of Southeast Asia, and it has developed a vibrant contemporary theater culture over the past thirty years. Singapore-based theater companies collaborate extensively with foreign partners and Singaporean festivals are common stopovers for some of the most influential theater companies from around the world. The diversity of formats and languages found in Singaporean theater attests to the multicultural makeup of the city-state, and also to the material wealth that underwrites recent theatrical experimentation. Indonesia is Singapore's largest neighbor (with more than 260 million inhabitants, the fourth most populated country in the world) and it is home to a completely different theatrical environment. For the past twelve years, I have been primarily interested in the traditional performances of Java, which constitute only one of the cultural regions of this vast archipelagic nation. I have conducted data-driven and data-assisted work on traditional Javanese puppetry and dance. Some of these traditions are at least one thousand years old, but are still the focus of active experimentation, and are still performed dozens of times every day.

Some readers might think that these excursions are too idiosyncratic and too specific to offer generalizable lessons for theater and performance scholars working on other areas. For example, one of my projects on Javanese dance describes *character types*, a highly specific convention which will not speak directly to specialists of Russian ballet or Argentinean tango. But the methods I describe for the analysis of motion capture can be adapted to many other movement traditions. Any data proj-

ect could have been used as a springboard for epistemological reflections such as those found in this book. No example is perfect for a book with the methodological goals that animate the discussions that follow. A book like this one could have also been written by an expert on Arthur Miller or Chinese Opera. Their examples, constrained by context and interests, would also not generalize to most other contexts. But a reflection on how data changes theater can start from any example. The point is not that all examples are the same, but rather that any theater example is unique and culturally determined. One advantage of using relatively unknown examples from Singapore and Indonesia is that they will remind every scholar that their areas of study, too, are culturally specific.

Ultimately, my goal is to show that computational research has limited, but important, applications for theater studies. I will try to show that pioneering research in this field is scratching the surface of what is possible, with many tantalizing possibilities just within our reach. I do not want to convince every reader to use data and computational methodologies, but to contribute to the ongoing conversation around epistemology and methodologies in theater research. Methodological reflection has been at the forefront of theater studies for many years and these conversations stand to gain from paying some attention to the challenges posed by data and computers. While I don't expect every reader of this book to become a computational researcher, I do hope to provide both adherents and detractors of these new approaches with more nuanced vocabularies with which to carry on their conversations.

This book is divided into three sections: the first delves into the epistemological challenges that computational research poses for theater studies, the second is a guided tour of key areas of theater amenable to computational research, and the last considers the material and ethical implications of computational methodologies for theater research.

In the first part, chapter 1 argues for a more textured conversation on methodology. I continue the distinctions introduced here. I analyze the way methodology is usually conveyed in theater research and pedagogy. I then contrast these views with the way methodology is conceptualized in the natural sciences. I aim to draw cautionary lessons from the superficial application of scientific principles to the study of culture. My goal is to enact, by way of critical realism, a more nuanced epistemological attitude that doesn't fit into simplistic distinctions between the humanities and the sciences. Chapter 2 describes how both data-driven and data-assisted research can make use of statistics. This chapter is not a survey of statis-

tical procedures, but a description of how different statistical methods might lead to different types of bias, and what to do about these problems. Chapter 3 describes how data-driven and data-assisted research can use visualizations. This chapter characterizes different approaches to data visualization through the work of leading thinkers in statistical graphics, data journalism, and DH. I also argue that the conceptual lens of theater and performance research could be deployed to highlight the performative nature of certain data visualizations.

Part 2 is a guided tour of four key areas of theater research that can be modeled as data: words (chapter 3), motion (chapter 4), relationships (chapter 5), and locations (chapter 6). In the chapter on words, I don't delve into the digital analysis of dramatic texts but I focus on what digital text analysis can bring to the study of program booklets, theater reviews, interviews, and other textual sources—an area of research with enormous promise but few precedents. By *motion*, I refer to the movements of actors and objects on stage, and by *relationships* I refer to the links between fictional characters in a play (copresence in a scene), or between people in collaborative networks. In all of these areas, extensive work has been already carried out—but I argue that there is much more that could be done. Researchers in DH who work on theater and performance are not always well known in theater studies circles. For example, Schöch, Fischer, and Wiesner are seldom quoted in theater books and articles. In my survey, I bring work in theater studies in conversation with work from elsewhere in DH, as well as with perspectives from engineering, computer science, and statistics.

In each chapter, I first consider what data can mean for the area under consideration and describe a variety of methods for the analysis and visualization of this data. Then I describe how these methods can be used within data-driven and data-assisted methodologies. I maintain a firm distinction between methods and methodologies to show that the same procedures can often be used to achieve very different research objectives. My overview does not aim to introduce groundbreaking perspectives; I have the more modest goal of using a consistent vocabulary to describe the potential for data-driven and data-assisted theater research in each of these areas. As noted above, each chapter concludes with a short excursion into theatrical practices in Singapore or Indonesia. These are summaries of work that I have carried out, often together with my scientific collaborators Gea Oswah Fatah Parikesit, Andrew Schauf, and Luis Hernández-Barraza.

The third and last part consists of two chapters. Chapter 8 considers the conditions for the future of computational research in theater studies. Both data-driven and data-assisted methods require consistent ways to document, standardize, and preserve data and the software that can interpret it. Here, I tackle the challenges that sustainability poses for computational research, from making sure that resources are available in the future to thinking of ways to enable other research teams to reuse existing datasets and to reconstruct their original context. However, I also note that making sure every data point is fully conformant with standards for reusability and sustainability is futile for a variety of reasons, ranging from copyright restrictions to financial limitations. It is much better to emphasize which areas deserve priority and focus our attention accordingly.

Chapter 9 closes the book with a brief reflection on the ways computational research can further change the study of theater, and what is needed for that to happen. I return to an argument that I made earlier in this introduction. Doing sustainable and meaningful data-driven and data-assisted research into theater will benefit from a community of scholars who can do some programming. This is not just a technical skill, but a mode of thinking about the world. My joys and travails as a programmer have deeply shaped the journeys that follow.

PART 1

Pre-departure Reflections

Toward a More Nuanced Conversation on Methodology

In 2012, Su Wen-Chi was artist-in-residence at CERN, the largest scientific laboratory in the world. As a media artist and choreographer, she was very interested in understanding the ways that scientists work, and how this differs from artistic practice. In a workshop on “art and science” in Singapore several years later, she related her experiences and described, in minute ethnographic detail, many fascinating aspects about the lives and routines of scientists at CERN (Su 2019). She explained that when dealing with complex problems, scientists are assigned to two groups. If both group’s answers are identical, then the answer is considered correct. One participant in the workshop noted how this couldn’t be further removed from artistic practice. “If we gave a question to two groups of artists,” the participant observed, “and they both had the same answer, we wouldn’t think that the answer was correct—we would conclude that the question was wrong.” This observation wonderfully captures the differences between two ways of producing knowledge: one that values contextualized, unique responses, and one that goes to great lengths to validate and corroborate the answers it produces.

Replicability is at the core of a realist view of knowledge, traditionally closer to the sciences. Interpretation is the foundation of a constructivist approach to knowledge, conventionally closer to the humanities. Data-driven research places more emphasis on replicability and data-assisted research prioritizes interpretation. But thoughtful computational theater research requires a sophisticated balance. If data-driven research does not acknowledge the role of interpretation, it will have turned its back on the core tenets of theater research. Conversely, if data-assisted research has no replicable features, then how can data possibly enhance conventional theater research?

This chapter seeks to enable a nuanced conversation on methodology by pursuing three arguments. First, I explain why theater research is primarily interpretive, and why this entails some limitations. Then, I describe why replicability is fundamental in the sciences and explain why adapting scientific methodologies to the humanities has often been a misguided and controversial enterprise. Lastly, I present an overview of critical realism and suggest that this metatheoretical perspective can enable a more sophisticated description of the epistemological goals of computational theater research.

Theater Research and the Primacy of Interpretation

Theater scholars pursue knowledge through a variety of means, which include practice-as-research, ethnography, historical analysis, phenomenology, and using different theoretical perspectives for performance analysis. In spite of this diversity, most approaches rely on interpretive methodologies which are social, intuitive, situated, and rarely made explicit. Within interpretive methodologies, a method can't be followed to the letter, for it must first be transformed by each individual researcher. A method, in this perspective, is not a series of steps, but a mode of analysis and a perspective on the world. It is hard, for instance, to explicitly state what constitutes a good postcolonial analysis. Books and classes that teach this approach proceed by example and discussion. Researchers and educators often emphasize the *why* for postcolonial analysis rather than the *how*. There is ample space for disagreement on how to do a good interpretive analysis, and conclusions might differ when the method is carried out by a different researcher.

The same is true for methods such as ethnography. Clifford Geertz (1973) suggests that ethnography includes some explicit procedures such as “transcribing texts, taking genealogies, mapping fields” but that it is not these “techniques and received procedures that define the enterprise” (6). Ethnography is defined by thick descriptions, a notion Geertz took from Gilbert Ryle. Such descriptions constitute a layered analysis of actions that consider their situated meanings. Geertz’s example is “a wink.” A thin description would gloss a wink as the closing of an eyelid. But the same gesture can be understood as provocation, mockery, or even rehearsal of mockery—depending on where it is performed and how it is perceived. Thick description attends to the complicating factors of context, and describes not only actions, but the surrounding conditions and

histories that imbue them with specific meanings. Being able to unpack these layers of meaning is not a straightforward process. Geertz famously said that understanding a culture from the perspective of other people is more like “grasping a proverb, catching an allusion, seeing a joke or [. . .] reading a poem” (Geertz 1974, 70).

The bulk of theater research entails the same kind of epistemic operation: learning a method is more like grasping a proverb than it is like following an instruction booklet. In some cases, it is hard to distinguish method from theory, as the example of postcolonial analysis shows. For Soyini Madison, carrying out research means taking a “detour around a topic with theory” (Madison 2005, 4, original emphasis). This entails no discernible, reproducible steps:

[W]e see and feel theory as something like a collaborator with performance, a co-subject however uncomfortably removed from the stability of a subject/object relation. In embodied relation to performance, theory moves. It is less the primary figure in a new construction of performance than it is a reflexive participant in the poiesis of knowing, being, and acting that performance initiates. I would thus have to call performance-and-theory a project of interanimation: of discerning how many more vital possibilities (for performance, for theory, for the world) are wrought by the transactivity of performance and various ways of imagining it. (Madison 2005, 2)

Judging the quality of theater research is a task for experts who can appreciate the overall argument of a research piece, and who don’t judge its conclusions based on the appropriateness of specific methods for data collection and analysis. Writings in theater—and in the humanities more generally—are usually not structured according to inflexible, predefined formats. The opposite is true of the sciences, where a rigid system dictates the structure (literature review, data collection, results, discussion, future work). Rigor and intellectual merit are important in theater scholarship, but they are not often conveyed or judged according to their compliance to inflexible structures amenable to step-wise execution.

Theater scholarship follows conventions, but it is not easy to categorically ascertain what they are, and making them explicit is not a key concern of pedagogy and research practice. We do not generally characterize methods as sequences of incontrovertible, executable actions that should be strictly reproduced to ensure consistent results. A simplistic counterex-

ample is the way I prepare my coffee: I weigh and grind the beans, wait for a light to go off in the espresso machine, and press a button. This method produces consistent coffee, day in and day out, and this is why I like it. Many people would be appalled at the thought of doing performance analysis the way I prepare my coffee. Perhaps a reason for this is that a quality that is desirable in coffee—predictable consistency—is not desirable in performance analysis, as illustrated by Su Wen-Chi’s story.

In graduate school, I used to work as a tutor in a large introductory class to theater studies, where one of the student assignments was to carry out a performance analysis. Each of the tutors had to grade dozens of analyses and one of the most excruciating aspects of the task was the distinctly Sisyphean feeling one had when seemingly reading the same essay, over and over again. The essays were incredibly similar to one another and we took this to mean that they were poorly written. In our evaluation, we were looking for creativity and individual expression. Step-wise methods are generally not well suited for generating a multiplicity of creative and personal responses. They aim exactly at the opposite, at producing predictable results through processes that minimize personal expression. In the class, we tried to encourage a particular sensitivity in the students, enticing them to look at the world of theater in a thoughtful way. But as they listened to us, they were aiming to discern protocols: specific steps to write an essay that can achieve a reasonable grade. Understanding what we meant was particularly difficult for science students (as a general education class, a large portion of the students came from technical backgrounds). Part of the cognitive dissonance they experienced is due to the fact that reproducible methods (common in the sciences) aim to discover things about the world, but methods in the humanities seek something else.

In the *Routledge Introduction to Theater and Performance Studies*, performance analysis is described as “a cognitive process that should lead to the constitution of meaning” (Fischer-Lichte, Thomasius, and Arjomand 2014, 54). This wording captures the very essence of interpretive methods: they are processes for constituting meaning, and for developing a particular relationship to performance, rather than for discovering its fundamental, unchanging properties. The *Routledge Introduction* has a section on methodologies, comprising three chapters: performance analysis, theater historiography, and theorizing theater and performance. Performance analysis can only be used when the researcher has accessed the performances directly. The book insists that historiographic methods are needed

to analyze performances from the past, or those seen on video (even when they correspond to the very recent past). Performance analysis is then broken down into semiotics and phenomenology, as two distinct approaches that the authors encourage the readers to combine. A semiotic analysis could be conducted methodically (in the coffee-making sense). There are indeed guides the step-wise semiotic analysis of theater (Ubersfeld et al. 1999; Aston and Savona 1991; Pavis 1996). This systematic approach has now fallen out of fashion, but as a step-wise method, semiotic analysis could still be taught in an undergraduate-level introductory book. This is not what the *Routledge Introduction* does. Instead of describing the sequential and executable steps of semiotics, the readers are exposed to the philosophical underpinnings of semiotics. Granted, a strict semiotic analysis can be deeply unsatisfying. Although the steps are executable, they are not necessarily unambiguous. Perhaps the reason why the writers of introductory textbooks leave them out is that they feel some epistemological unease when dealing with the reductiveness of reproducible methods. I suggest there are several reasons for this unease, which I list below:

1. *A stance against positivism and philology:* Moving against positivism and philology has been an important strategy for theater studies as the discipline has found its way throughout the twentieth century (Jackson 2004).
2. *The legacy of poststructuralism:* This has endowed us both with a suspicion of method and a fascination with the vocabulary of method. Think, for example, of Derrida (1978) and Foucault (1998), both of whom wrote about method in an evocative way. As Bishop and Phillips note:

[I]n raising the question of method [. . .] one is concerned not with what might be called the “technical devices” of the theoretical disciplines (philosophy, sociology, psychology, political science, etc.) but with a more sustained and comprehensive understanding of the *conditions of possibility* for disciplines like these, which in their foundation and their horizons offer a kind of incomplete heritage for future knowledge. (Bishop and Phillips 2007, 267, original emphasis)

This passage signals a poststructuralist approach, where *conditions of possibility* are the central focus of the inquiry. However, in

the pages that follow I will argue that data requires a more sustained engagement with the “technical devices” of methods.

3. *An artistic ethos*: Many theater researchers are also artists or at least have an artistic sensitivity. They aim to transfer these qualities, or at least reflect them, in their academic writing. This is different in other fields that also deal with artistic practice. Take, for example, cultural policy, where a more dispassionate approach is preferred. In theater and performance studies, the ethos of artistic practice and a belief in individual expression permeate our own sense of what constitutes good and valuable research.
4. *A preference for implicit rules*: Theater scholars prefer the language of allusion and poetry (see the beautifully written passage from Soyini Madison above) and believe there is something prosaic in narrow definitions and step-wise sequences of goal-oriented actions.
5. *Comfort with ambiguity*: Theater scholars recognize the complexity of concepts such as *performance*, which have no clear definition. As mentioned in the introduction, most of us are comfortable in environments of sophisticated disagreement. Kagan (2009) argues that different disciplines attract different personality types, and that the sciences tend to attract people with lower tolerance for ambiguity.
6. *A principled stance against technocracy and the knowledge economy*: Like many humanities scholars, we believe that in our world, a superficial veneration of science and quantification has damaged many areas of life. Dwight Conquergood (2002) echoes common views when he advocates for the importance of seeking knowledge that is local, embodied, socially generated and implicit. Replicable methods are not well aligned to these objectives.

An excellent example of the principles outlined above can be seen in *Research Methods in Theatre and Performance* (Kershaw and Nicholson 2011), which is aimed at academics and graduate students, primarily those interested in practice-as-research. Their aim is described as a “challenge to outmoded perceptions that the terms ‘method’ and ‘methodology’ imply an attempt to capture, codify and categorize knowledge” (1). The authors constantly emphasize the unpredictable nature of theater, and explicitly articulate their discomfort with reproducible methods, asking “what are methods for but to ruin our experiments?” (15).

The books I’ve quoted thus far don’t constitute a full survey of the field. Areas such as cognitive approaches to performance warrant a more

nuanced take than what this overview permits, as they might not be fully described by the characteristics listed above. The books surveyed thus far are aimed at undergraduate, graduate, and academic researchers but they are also indicative of larger patterns. Fischer-Licthe, Arjomand, and Mosse provide an introduction to theater studies that is not dogmatic; Madison helps us think about theory in provocative and textured ways; and Kershaw and Nicholson provide unique insights into the relationship between artistic practice and academic research. Conventional ways of approaching methodology in theater are certainly useful, but there are three potential problems that demand attention.

First, *implicit definitions of methodology limit interdisciplinary exchange*. We want creative and situated responses, but it is very hard to explain this to people outside the humanities. Describing our methods as situated, social and intuitive does not diminish their value, it just describes them more adequately. I think that a more explicit characterization of interpretive methods is crucial for interdisciplinary research and for communicating our epistemological projects to society at large.

A second problem is that *some aspects of our scholarship would benefit from reproducible approaches*. Although a reductive strategy at first glance, this might help identify previously unsuspected features of theater, which can also aid in interpretive endeavors. This applies not only to computational research methods but also to standardized surveys such as the ones that are commonly used in the social sciences.

Third, *our lack of training in reproducible methods prevents us from critically assessing the flaws and strengths of relevant scientific research*. In reproducible research, creativity plays a role in devising new methods for collecting and analyzing data, or in identifying which method is relevant. But once the design has been made, the collection and analysis of data should proceed in a systematic, inflexible way. For these methods to be deemed rigorous, there is no scope for creativity in the execution of the methods, only in their design.

If you suspect you are sick and go to the doctor, you would expect interpretive flexibility on the part of a good doctor: empathy, creativity, and intuition. But you would not want the way your blood samples are collected and analyzed to be creative—you would want this to be conducted in a way that is as systematic and consistent as possible. Some aspects of medical analysis benefit from situated hermeneutics, while some require the rigorous consistency of impersonal, reproducible methods. In fact, diagnosis and treatment will work better the more these two areas are dif-

ferentiated. Might this also be true for theater scholarship? If we are to work with data, we need a better way to describe our methodologies and their objectives.

The Controversies of Scientific Research in the Humanities

The sciences traditionally favor a very different mode of evaluating evidence. When describing scientific practice, Richard Feynman (1985) famously said that “the first principle is that you must not fool yourself—and you are the easiest person to fool” (343). To avoid fooling ourselves, we need replication: “we’ve learned from experience that the truth will come out. Other experimenters will repeat your experiment and find out whether you were wrong or right. Nature’s phenomena will agree or they’ll disagree with your theory” (342).

Replication is fundamental in the sciences since it is the only way to ensure that explanations are likely to be true. Elsewhere, Feynman writes: “scientific knowledge is a body of statements of varying degrees of certainty—some most unsure, some nearly sure, but none absolutely certain” (Feynman 1955, 14). Absolute certainty is impossible, but the goal of science is to reduce the uncertainty of possible explanations. Science aims to achieve this through continued observation and experimentation. In another lecture, Feynman describes the process in the following way:

In general we look for a new law by the following process. First we guess it. Then we compute the consequences of the guess to see what would be implied if this law that we guessed is right. Then we compare the result of the computation to nature, with experiment or experience, compare it directly with observation, to see if it works. If it disagrees with experiment it is wrong. In that simple statement is the key to science. It does not make any difference how beautiful your guess is. It does not make any difference how smart you are, who made the guess, or what his name is—if it disagrees with experiment it is wrong. That is all there is to it. (Feynman 2017, 156)

Is a scientific study of culture possible? Here, I agree with Ted Underwood (2019b): “Questions that historians and literary critics used to debate are increasingly scooped up by quantitative disciplines [. . .] Instead of saying that the humanities are besieged and giving up ground, we could truthfully say that these disciplines are discovering new missions and new

ways to understand culture” (n.p.). But finding the right way to frame and conduct scientific research in the humanities is a difficult matter, as illustrated by the controversial work of Jonathan Gottschall. In what follows, I dedicate substantial attention to his work, since he presents a persuasive argument for a scientific approach to the humanities. But his work also provides a cautionary tale on the limitations of such approach. His attention is focused on literature but his arguments can be extended to other areas of the humanities.

Gottschall (2008) argues that the work of literary scholars is only of interest to themselves and that literary scholarship has remained irrelevant to the world at large. Whereas scientists might sometimes try to write for general audiences, this is rarely the case in literary studies. In part, this is because of the very way in which literary scholarship is construed. For Gottschall, the main problem is that literary scholars have failed to produce knowledge that stands the test of time as they “rarely succeed in accumulating more reliable and durable knowledge” (7). He identifies several reasons for this: a belief in masters who never proved anything, a circularity of theory-proof, and stasis in academic departments. He also identifies a strong ideological bias, as scholars remain convinced, *a priori*, of the need to fight for just causes and identify them at work in literature. He calls this last attitude the “liberationist paradigm” where scholars think they can set people free of the workings of power through the sheer force of literary criticism. As an antidote, Gottschall suggests that scholars should aim to narrow the space of possible explanation. This requires a change of theory, attitude, and method, as it would imply that literary scholars should submit their ideas to the test of others. Consider the following argument for embracing a scientific approach when studying literature:

[T]he praxis of the liberationist era has amounted to endlessly asking questions while despairing of more valid answers [. . .] The quintessence of the dominant paradigm is, then, constitutional and reflexive pessimism about the ability of humans to really know *anything* [. . .] But long before Derrida’s neon declarations, the scientists had been pretty comfortable with the idea that it is not possible for humans to know the truth of something in the sense of its ultimate reality [. . .] Popper’s concept of falsifiability, which has been a guiding philosophical principle of scientific investigation for more than a half century, is an attempt to grapple with the fact that it is not logically possible

to *prove* any scientific claim by experiment. But science's response to this realization was more reasonable and productive than that of the "great generation" of liberationist theorists [. . .] We can't know for certain what is true. Science makes no ultimate claims. But through a gradual process of rational thinking and falsifying tests, communities of scientists can show where the preponderance of evidence lies. This is the best that humans can do, and this is no small thing. (Gottschall 2008, 11)

Like him, I am also sometimes deeply dissatisfied by some aspects of literary scholarship, but I think that his analysis lacks nuance. While there are excesses in the "liberationist paradigm," there is also value in much literary criticism. The most important feature of good literary analysis is concerted attention to context and the careful exploration of concepts fraught with disagreement. Extrapolating this to theater, I don't think our objective will ever be to limit the possible space of explanation of what *performance* is. But Gottschall's insights are useful within limited aspects of scholarship. I do think that shrinking the possible space of explanation is useful and necessary within certain narrow inquiries, such as those that fulfill the criteria for data-driven research (see introduction). To better see the pitfalls of Gottschall's approach, let's consider the narrower case he makes for the application of scientific methods to literature. He advocates using evolutionary psychology (EP) to study literature. A substantial part of his research aims to show that certain narrative features are indicative of cognitive traits that were well suited to the evolution of early humans. To his credit, he is not merely applying EP concepts to literary analysis, but actually conducting scientific research with several collaborators. One problem of this approach, though, is that too many factors have shaped human evolution and it is impossible to accurately measure the impact of most of them. Scientists have robust and reliable genetic and climatological data which might explain certain aspects of human evolution. But we don't have sufficient data on the early uses of narrative. As Cameron (2011) points out in a response to Gottschall, EP research is often faulted for its use of circular logic: "why is this trait so common among humans now? Because it was adaptive for their ancestors. Why do we think it was adaptive? Because it is so common now" (62).

I think that Gottschall's larger epistemological argument is correct and directly relevant to data-driven research. But this approach only works on

narrowly focused questions, which can be answered in terms of the data we have. I dream of a world where we could explain the role art has played in evolution, or the ways in which it can make societies better. But these questions are very difficult to answer given the paucity of evidence and the complex interactions of the millions of factors that shape societies. A less satisfying, but more feasible approach, is to focus on small questions. Taking a cue from Gottschall's vocabulary, I suggest that data-driven methodologies should also aim to narrow the *space of possible questions*, not only the space of possible explanations. The things I've learned from my data-driven research are very limited in scope; for example, I found that the frequency of the word "audience" steadily decreased in the writings of a group of theater reviewers over a period of twenty years in Singapore, and that characters of Indian and Javanese origin are interconnected in vastly different ways in Javanese *wayang kulit* scene structures. These might seem as small conclusions, but they can be verified and disproved by other researchers. I wish to one day make more substantial discoveries on Indonesian and Singaporean theater, but the steps towards larger discoveries are incremental, small questions tightly wrapped around available data.

But leaving that aside, whether Gottschall's EP papers (or whether my own research) constitute solid science or not has no bearing on the validity of Gottschall's defense of science as a valid approach to study culture. Faulty research examples can't be used to condemn an approach. One could also find examples of superficial postcolonial perspectives and this would not in itself invalidate the whole enterprise of postcolonial studies. Gottschall's description of the scientific method and what it might do for the study of literature is sound and enticing, and relevant to the data-driven study of theater. Another line of critique, also espoused by Cameron, is that Gottschall's scientific analysis stands in fundamental opposition to literary study:

[C]ritics do not 'make progress' towards the 'true' interpretation of a Keats sonnet; they merely offer different readings. Unlike a new scientific theorem, moreover, a new reading of a text does not automatically supersede all previous interpretations. Rather it is of interest to the extent that it reveals additional meanings in a text, proposes hitherto unnoticed connections between texts, or foregrounds themes in texts which resonate with the concerns of a particular moment. (Cameron 2011, 67)

This is a very different kind of argument. The issue at stake here is not whether the scientific analysis of literature is possible, or whether Gottschall's own research conforms to the scientific standards he invokes. The question here is the purpose of literary study: is it about finding patterns in the production of literature or is it about generating situated readings? I think that the field is big enough to accommodate both scholarly criticism and a science of literature. The same argument can be extended to theater studies. Our discipline aims to enable situated conversations around the meaning of theater. But it is also a concerted effort to find verifiable patterns in the history and current practice of theater performances. And both projects benefit from the study of theater's digital traces. Data can be used to find replicable insights as well as to assist in interpretive, situated responses to performances.

Perhaps Gottschall's mistake is that he militates against all nonscientific perspectives in the study of culture and suggests that only science can prove useful. But this is not the only approach, not even when working with data. As Ramsay (2007) says: "We would do better to recognize that a scientific literary criticism would cease to be criticism." When analyzing the work of Virginia Woolf, Ramsay says that critics are not trying to solve Woolf: "they are trying to ensure that discussion of *The Waves* continues into further and further reaches of intellectual depth." Likewise, Rockwell and Sinclair (2016) have created tools for textual analytics which they describe as "interactive interpretive toys" (69), rather than as "microscopes revealing the inner structure" (103) of literature. Their goal is not to identify objectively verifiable patterns, but to "add to a history of interpretation" (103). Some research questions benefit from replicable approaches and others do not.

On the topic of replicability, I would like to also address the criticism against computational literary research brought forth by Nan Z. Da (2019) in a widely circulated—and widely controversial—paper. Da tried to replicate some DH papers and couldn't do so. Her analysis on topic modeling is particularly sharp. She shows that this procedure might be extremely sensitive to the parameters chosen: "When I randomly removed just 1 percent of the original sample, all the topics changed" (2019, 628). We certainly need more papers that seriously probe the foundations of computational research and that alert us to methodological blind spots such as this. There are different methods we can use to improve our conclusions and make them more robust to the problem of parametrization (a point I revisit in chapter 2). There are many useful things in Da's careful atten-

tion to computational errors. The problem is that she then moves on to say that such mistakes invalidate the computational study of literature as a whole. As Jannidis (2019) notes in his response to her paper, “there is no logical way to move from an error in a calculation of a researcher to a general statement about the fruitfulness of a research field in general” (n.p.). Da’s paper insists that, even if the errors in the calculations were solved, computation has nothing to add to the study of literature. This is similar to Gottschall’s argument, but in reverse. While Gottschall suggests that only a scientific study of culture is of interest, Da claims that science can add nothing to the field.

The arguments presented by Da and Gottschall are controversial for their totalizing conclusions. In both cases, though, we would be wise to recognize the very valid points they are making. Gottschall provides a vivid exhortation of what science can do, and Da’s paper should alert us to the importance of replication.

Critical Realism: Reconciling Different Epistemologies

We have seen that the natural sciences and the humanities have historically differed in their epistemic objectives. Data-driven and data-assisted methodologies are reconfiguring these faultlines. This has many implications for how we carry out research projects and also for how we might rethink teaching and evaluation practices. Speaking mostly about differences between contrasting approaches, as I have done so far in this chapter, is useful to orient ourselves in this changing landscape. But it has the unfortunate consequence of emphasizing divergence and opposition. Both interpretive and scientific approaches say things about reality in ways that seem at odds with each other. But believing that these two modes are fundamentally irreconcilable only makes sense in a zero-sum epistemology where we need to choose exclusive allegiances, and where we need to be either scientists or humanists, either realists or constructivists. Such opposition might be encouraged by disciplinary boundaries in academic institutions. In the humanities, we are usually taught to be suspicious of anything that seems remotely positivistic. In the sciences, we are taught to only trust that which can be measured and verified. But perhaps we can imagine a different kind of engagement with reality, where we recognize that some things can be measured whereas others cannot. Some areas benefit from systematic experimentation and rigorous replication practices, while others required the textured explorations of con-

tested terms that reflect subject-dependent experiences of the world. But these different components are sometimes part of the same phenomenon and keeping them separate reduces our understanding of complex issues.

Take, for example, climate change. The analysis of climatological patterns provides a good example of an area where a scientific attitude is necessary. Consistent, verifiable data has been central to one of the most consequential discoveries of our time. As scientists have learned from analyzing millions of data points, world temperatures are rising at alarming rates, and this rise is almost certainly due to human action (Bernstein et al. 2008). Reliable data will continue to be indispensable for tracking the impact of collective action to combat climate change. But climate change will also impact lives in ways that might not always be measurable. Imaginative and situated perspectives will be crucial if we want to understand how different communities make sense of climate challenges, and if we aim to devise creative and fair solutions to these problems.

The biggest challenges of our times will require us to reconcile a belief in solid scientific evidence with the recognition that scientific methods have been used to oppress and misrepresent people, and that they cannot answer every single question. We will need to accurately gauge the kinds of questions that science can answer, while complementing them with situated and interpretive modes of learning about the world. These are tall orders, and this book aims to make a small contribution to a minuscule part of the problem: to elucidate what computational research can or can't do for the study of theater.

In order to combine scientific and interpretive approaches, we can take a cue from critical realism (henceforth CR), a stance that rejects the extreme positions of both positivism and constructivism. The writings of Roy Bhaskar (1944–2014) are associated with CR, but here I refer mostly to later formulations, found in the work of Dave Elder-Vass, Ruth Groff, Paul Edwards, Frédéric Vandenberghe, and Berth Danermark. Rather than a theory or method, CR is a metatheoretical perspective. It is a relatively recent perspective and there are differences in the way CR is articulated by different theorists—what follows is a select summary of key ideas rather than a comprehensive overview.

From a CR perspective, the problem of positivism is that it denies the role that power, history, subjectivity, and discourse play in the construction of knowledge. Researchers in CR agree with constructivists on “the political nature of science and are equally skeptical of its truth claims, many of which simply represent the current orthodoxy within scientific

communities” (O’Mahoney and Vincent 2014, 5). In turn, the problem of constructivism is that it must reject any claims that the natural or social sciences provide a “better” understanding of the world. As long as someone holds a particular belief, then we must accept that is “their reality” and there is no way to ascertain that other theories might have better explanatory power. This conclusion is unreasonable in the natural sciences given the demonstrable advances of applied science. But it should also be untenable for the study of the social world:

We would hope discourses that girls are naturally bad at science, that Western cultures are superior to others [. . .] would not be accepted solely on the basis that some groups believe these statements to be true. (O’Mahoney and Vincent 2014, 6)

Following Bhaskar, critical realists suggest that both constructivism and positivism make similar mistakes, as their views are premised on an “epistemic fallacy” which conflates ontology (that which exists) and epistemology (that which we can know). Positivists reduce ontology to epistemology by claiming that only that which is observable exists. For constructivists, there is no ontology outside of epistemology, as the world only exists through our discourses about it.

In contrast to both perspectives, CR suggest that there is an external reality, but that our knowledge of that reality is necessarily subjective, contingent, and socially constructed. For CR researchers, this key distinction between ontology and epistemology leads to a “depth ontology” that distinguishes between three realms: the *empirical* (that which can be observed), the *actual* (that which exists), and the *real* (the causal mechanisms of reality) (O’Mahoney and Vincent 2014, 9). This means that a verifiable external reality exists, but we can only study it through uncertain processes, acknowledging that “different actors will define reality in different ways” (Edwards, Vincent, and O’Mahoney 2014, 321). Thus, CR-infused research aims to identify the causal mechanisms present in reality, while acknowledging that understandings of the world are socially constructed and subjectively experienced.

A CR perspective has clear implications for research; CR-infused research must be context-sensitive, as opposed to positivism. But it should also pursue the understanding of an objective reality, as opposed to constructivism; as a result, CR is committed to both truth and to thick explanations (which take politics, history, and subjective experiences into

account). As Vandenberghe (2013) notes, CR “enters the ‘science wars’ by fighting two fronts” (5). For CR researchers, “useful research is necessarily rich, ‘thick,’ and explanatory as opposed to the ‘thin’ descriptive approaches that positivism necessitates” (O’Mahoney and Vincent 2014, 4). But it also aims to pursue better explanations of an observer-independent reality.

What does this mean for computational theater research? I have argued that data-driven methodologies are closer to the sciences and data-assisted methodologies are closer to conventional humanities approaches. Through CR, we are invited to move beyond this dichotomy; to do science without abandoning the presuppositions of the humanities, and to study the humanities without abandoning the premises of science.

A humanities-aware science remains critical about the “facts” on which its conclusions rest. We are led by CR to analyze the influence of history and power on the construction of facts and categories:

Facts don’t speak for themselves [. . .], they are always categorized and schematized by one or another theory, philosophy or cosmology that is socio-historically determined, there is no observation that is not an interpretation and no interpretation that does not involve an imaginary representation of reality. (Vandenberghe 2013, 5)

Conversely, scientifically inclined humanities research aims to place historically and subjectively situated observations against a larger backdrop. This means piecing together explanations from different perspectives to elucidate larger, intersubjectively verifiable patterns. Consider, for example, what a CR-influenced ethnography would look like. From a CR perspective, ethnography can “provide a deeper understanding than subjectivism is capable of, one which is able to link the subjective understandings of individuals with the structural positions within which those individuals are located” (Rees and Gatenby 2014, 135). Extending this approach to data-assisted research means that *deformances* can still be understood as indicative of shared, objective structures of culture. However, this is a point with which many DH theorists would disagree, as some people carry out data-assisted research from a purely constructivist perspective. This is not something I want to disavow. The point I am making here is that it is possible to engage in CR-inspired, data-assisted research, in the same way that it is feasible to carry out CR-inspired ethnography, without wholly embracing positivism and without fully eschewing constructivism.

A possible criticism of CR is that it is only useful to study the natural world, and that there is no objective component in culture. To this, Elder-Vass (2012) responds that a purely subjectivist account of culture would be incoherent as “it would lack the means to explain how culture can acquire the shared quality that makes it culture” (39). Objective culture is a product of human agency, but nevertheless “exerts a causal influence of its own” once it is produced (41). Culture and beliefs have demonstrable agency and create an impact on an observer-independent reality: “the tooth fairy is not materially, but ideationally real: the discourse about it has real effects, for example, on the bedtime activities of children, even though it does not exist” (O’Mahoney and Vincent 2014, 7).

To bring these ideas back to theater, let’s consider the ontological status of performances. Performances belong to the realm of actuality, but any attempt at grasping their meaning or form is mediated by subjective experiences and socially constructed categories. Still, these categories have consequences. Let’s imagine an event that is classified as an “experimental mixed-media performance.” Each of these terms has contested meanings and many people might not agree with how well the label fits the specific event. However, if a company uses this label to describe their work, this will impact the way it is marketed and funded, and will likely change the composition of the audience that will come to see it. We can use data-driven approaches to understand how often this label is used, and data-assisted approaches to explore the limitations of this label.

Besides CR, there are other recent perspectives that challenge a binary distinction between realism and constructivism. Paul Rae (2018) offers a comprehensive survey of “new realisms” and what they mean for understanding the fraught connections between theater and reality (7–8). Perhaps other new realisms, as well as CR, can provide a foundation for carrying out data-driven and data-assisted research that has learned from both hermeneutics and science. Data-driven research should be aware of how politics, discourse, and subjectivity shape its methods and conclusions. Likewise, data-assisted research could contribute to the understanding of objective aspects of theater. Both perspectives could complement each other. One can use a data-assisted performative visualization to explore how power structures shape the categories of a data-driven project. Conversely, a data-driven analysis could uncover patterns within a data-assisted project.

But this does not mean that both approaches must always work in tandem. A data-driven project might address its limitations by reflecting on the impact of history and power structures in the way the data was defined,

collected, and processed. A data-assisted project can likewise indicate its contribution to the study of shared objective reality. Whether a research project requires primarily a data-driven perspective, a data-assisted one, or a combination of both depends entirely on the questions it seeks to explore. Data-assisted perspectives poke around the edges of a question, and data-assisted perspectives aim to find the best possible answer to a question under formal constraints (see the different criteria outlined in the introduction). Projects might also exist in a continuum. A project can be said to be driven by data inasmuch as it emphasizes replicability, and assisted by data to the extent it draws attention to situated interpretations. If we use the foundations of CR to describe the epistemological quests of computational theater research, we will be better able to justify why a particular approach works for a specific question, and how different kinds of questions can contribute to a fuller understanding of theater's history and current practice.

CHAPTER 2

The Roles of Statistics

In the majority of the projects reported in this book, at least some of the data is quantitative. For example, projects might report the number of words in a text, the number of partners in a collaboration network, the angular velocity of a dance movement, and the number of performances in a given place. But these quantities can be put to very different uses, as they can aid in both data-driven and data-assisted research. In data-driven research, numbers are used to find empirical patterns and convince others of the best possible description or explanation of a phenomenon. In data-assisted methodologies, numbers are deployed as deformances, in order to challenge the assumptions of a question and generate multiple interpretations that do not supersede each other as more correct or more fitting.

Often the same numerical methods can be used within very different epistemological frameworks. Term frequency, inverse document frequency (*Tf-idf*) is a statistic often used in text mining, and it could be adopted for the classification of documents in a textual collection. However, the same method could also be used for interpretive readings. In his study of Woolf's *The Waves*, Ramsay (2007, n.p.) uses *tf-idf* to give a sense of how certain characters' lexical patterns differ from each other's. As Ramsay notes, this statistic does not aim "to bring the results into closer conformity with 'reality,' but merely to render the weighting numbers more sensible to the analyst." Thus, quantitative transformations may enable new interpretive readings, and they can be enlisted as important allies in data-assisted research.

As Underwood (2019a) notes: "Quantitative models are no more objective than any other historical interpretation; they are just another way to grapple with the mystery of the human past, which doesn't become less complex or less perplexing as we back up to take a wider view" (xix). Piper (2018) echoes this view: "the literate and the numerate are not agons

engaged in a duel. They are two integral components of a more holistic understanding of human mentality” (5).

Numbers privilege some modes of representation over others. But the same is true for language, and scholars are well aware of this. As Katherine Bode (2012) notes, instead of abandoning language “scholars have sought to understand the ways it works and to challenge and critique the relations of power it perpetuates.” The same needs to be done with numbers, as we seek “to recognise them as a form of representation and, as such, to explore how they operate and the ways in which numbers accrue authenticity and authority” (12).

In this chapter, I answer Bode’s call to pay attention to the way numbers operate as I explore the ways they can be mobilized for data-assisted and data-driven research. I describe their potential for bias and insight by focusing on statistical methods, which are at the core of the projects I will describe later in the book. This chapter looks at a conventional distinction made in statistics, between descriptive and explanatory data analysis. But what follows is not a series of textbook descriptions. Rather, I offer an epistemological reflection on what these two types of statistical approaches reveal and obscure, and what they entail for computational theater research. I pay specific attention to the potential for bias in each, and the steps that can be taken to mitigate bias.

Within *descriptive analysis*, I refer to procedures that outline the shape of the data or which seek to identify outliers and patterns such as trends, correlations, similarity, and clusters. Within *explanatory analysis*, which are often called inferential statistics, I refer to methods that aim to explain causality in the phenomena under consideration. Descriptive analysis can be used in both data-driven and data-assisted projects. Explanatory statistical analysis, in contrast, belongs solely to the realm of data-driven research.

Descriptive Analysis

Some descriptive methods aim to characterize the general shape of the data. This includes calculating measures of central tendency (like the arithmetic mean) and measures of dispersion (such as the standard deviation), as well as conveying the distribution of the data (often via graphical means such as histograms). Descriptive analysis can also include ways of portraying the differences between groups (effect size) and estimating the relationship between variables (correlation).

These operations are often grouped together under exploratory data analysis (EDA), a concept first proposed by John Tukey (1977). Rather than a specific set of methods, EDA is an agnostic approach to data analysis that aims to identify the shape of the data without imposing assumptions about what it must be like. The NIST/SEMATECH e-Handbook of Statistical Methods (2003, n.p.) describes EDA as “an attitude/philosophy about how a data analysis should be carried out” which is used to achieve the following objectives: to uncover the underlying structure of a dataset, extract important variables, detect outliers and anomalies, test underlying assumptions, develop parsimonious models, and determine optimal factor settings for further analysis.

Statistician Allen Downey describes a typical data analysis routine as consisting of the following steps: (1) importing and cleaning the data, (2) single variable explorations such as distributions and summary statistics, (3) pair-wise explorations for possible relations between variables (correlations, linear fits), and (4) multivariate analysis, such as regression and control variables for more complex relationships. These exploratory procedures are then followed by inferential analysis for estimation and hypothesis testing (see the second part of this chapter). Visualizations are used throughout the entire process to aid in the analysis and to communicate the outcome of the research. Histograms, boxplots, and violinplots are particularly useful for identifying the shape of the distribution (I consider the role of visualizations in more detail in chapter 3).

According to Arnold and Tilton (2019), EDA permeates much of the digital humanities, even though it is not often named; it is an iterative, agnostic way of discovering the properties of the data. But the fact that it is agnostic in a technical sense (i.e., that it doesn't impose a model a priori) doesn't mean it will necessarily present an unbiased representation of the data: EDA is only as good as the data that is analyzed. And if there are biases in the way the data was collected, EDA will not necessarily reveal this. As Bode (2020) notes, no literary dataset is unbiased. The same is true for any theater dataset, or any dataset in the humanities for that matter. Bode is particularly concerned with the underrepresentation of women writers, which is often the result of systemic bias in many digital collections. Concerning word usage in a literary collection, EDA can offer an agnostic representation, but the systemic imbalances in the data collection can't be removed just by using EDA. Researchers will need to critically interrogate the ways the data was collected and pay attention to the history of the dataset. If undertaken from the perspective of critical

realism, EDA will need to be accompanied by a thorough analysis of the cultural and historical conditions under which the data was created.

Besides the problem described by Bode, which we could refer to as *sampling bias*, other features might also make descriptive data analysis less objective than it seems. Often, researchers don't have a direct way of measuring a phenomenon they are interested in, and turn to proxies. In some cases, the proxy might not be directly related to the phenomenon of interest and this can lead to *proxy bias*. In other words, the availability of seemingly neutral or complete datasets proves too tempting for researchers, and we might not always question the assumptions that went into the creation of the dataset.

Take for example a study of the evolution of novelty in films that uses IMDB *plot keywords* to measure innovation in films (Sreenivasan 2013) over the twentieth century. The author uses sound mathematical techniques to show a correlation between novelty and revenue: blockbusters tend to occupy the midsection of an innovation curve, with films that are either too innovative or not innovative enough trailing off at both ends of the curve. The main problem with this research project is the assumption that the number and variety of tags will indicate novelty. As any film buff knows, remakes will often be tagged with completely different tags than their original sources, and newer remakes will often have tags that are more varied and larger in number than the original films. For example, at the time of writing this book, *Abre los ojos* (dir. Amenábar 1997) has 140 tags, while the Hollywood remake *Vanilla Sky* (dir. Crowe 2001) has 176. The original *Ringu* (dir. Nakata 1999) has 154 keywords, while its remake *The Ring* (dir. Verbinski 2002) has 189; the original *Ghost in the Shell* (dir. Oshii 1996) has 222 keywords, and its remake (dir. Sanders 2017) has 332. It seems unlikely that these newer films are more innovative; perhaps the data just shows that newer films will tend to be tagged in more idiosyncratic ways (or that Hollywood productions are given more tags than films produced elsewhere). In other words, films have just become more *tagged*, with newer films being tagged in more granular ways. Likewise, when modeling theater as data, it is worth pausing to consider the extent to which available data directly measures the phenomenon of interest and the extent to which it is a convenient, but misleading proxy.

Besides EDA, other statistical methods can be used to identify patterns in the data, such as clusters and trends. The majority of the projects I describe in this book—and perhaps the majority of DH work—aims at identifying similarity (clusters) and change (trends). For example, in

chapter 4, I describe a project that identified changes in the words used by theater critics over a two-decade period. In chapter 6, I show ways to compare networks and track their changes over time. And in chapter 7, I look at clusters of performances in a given place, and on how that changes over time. These projects use methods from computational linguistics, network analysis, and geostatistics, as well as some general data mining methods for standardization and dimensionality reduction (reducing the number of variables in a large dataset). Elsewhere in DH, researchers sometimes use machine learning (ML) techniques rather than classical statistics to find clusters and trends. ML researchers distinguish between clusters and categories. Generally speaking, clusters are latent structures in a dataset discovered through unsupervised ML algorithms. In contrast, supervised ML approaches aim to assign objects to predefined categories, which are generally determined by human annotators. This book focuses on classical statistics rather than ML, and in what follows I refer to clusters as groups identified by classical statistical procedures, rather than by ML algorithms.

Clusters and trends are also descriptions of the data, but unlike EDA they require us to make assumptions about a dataset, and select parameters for what constitutes a trend or a cluster. A cluster is not an inherent property of the data as opposed to, say, the arithmetic mean. You can find the arithmetic mean of a dataset in an uncontroversial way without making any assumptions about what constitutes a meaningful relation between items. But a cluster requires a model of what is meaningful. As O’Neil (2016) notes, “models are opinions embedded in mathematics” (21).

When identifying trends and clusters, any method is susceptible to bias. In some cases, there are technical solutions and certain methods can be validated as more useful for given applications. But, as in the case of EDA, sometimes bias is ingrained in the way the data was collected. This is a particularly important problem for the identification of historical trends. Many DH projects aim at finding changes over time. But any attempt at analyzing historical data should consider *chronology bias*. This type of bias is more often discussed in the medical literature and refers to the errors in judgment one makes when comparing evidence from different historical periods (Feinstein 1971, 870). Medical researchers studying the incidence of a disease might conclude that a disease has become more prevalent over time, but the reason for an increase in reported cases might be that diagnostic methods improved. This is sometimes treated as a special kind of *confounding bias*, where an observed correlation is caused

by an agent not directly measured (in this case, this agent would be the improvement in diagnostic methods). A close parallel to this situation is common in DH. When using historical digitized texts, some words might appear to be more common today than in the past, but this might be the result of better optical character recognition (OCR) techniques for more recent texts. For theater, data might be more readily available for recent performances than for those in the past, and this might distort the perception of trends.

Another problem of historical analysis is that the meanings of concepts might change. An excellent example of this is offered by Rosenberg (2013) in his analysis of how the word “data” entered the English lexicon. In this study, he compared two different pattern analyses of the usage of the word “data,” which had been previously published online. He then closely read several early examples that formed the basis of such projects. His conclusion is that the patterns available in the published online pieces were misleading, since the relative usage of the word “data” needs to take a Latinized understanding of data common in the seventeenth century into account, and distinguish it from its modern English sense. Historical patterns offer tantalizing explanations for cultural change, but it is crucial that the data is examined in detail.

To find meaningful clusters in the data, one must acknowledge a pervasive problem: clusters will *always* emerge. This will be true regardless of the clustering method: there will always be some pattern to analyze. This is a mathematical necessity, a corollary of the Ramsey theory that states that a pattern is guaranteed to emerge given enough elements in a set (R. L. Graham, Rothschild, and Spencer 1990). Thus, researchers must resist the temptation to overinterpret the meaning of a cluster. One will always be able to find meaning in any clustering of the data, the way one can always find shapes in tea leaves. This is sometimes called *clustering illusion* (Bedeck et al. 2018), and in the scientific literature it is often linked to *harking*, or hypothesizing after the results are known (Kerr 1998).

For example, imagine a clustering method that shows that the words “chair” and “power” are part of the same cluster in a corpus of theater reviews. Based on this pattern, we might hypothesize that chairs are seen by the theater critics as symbols of power. But if we had chosen another clustering mechanism, it is entirely possible that “chair” would form a cluster with “contemporary,” and this will send us down an entirely different rabbit hole of possible associations and interpretations. We might assume that the first clustering method reveals semantic associations

between words. A few examples might convince us this is true. But we would need to independently calibrate such a method first. We would need to run tests on other datasets and confirm whether the method tends to reveal words that have semantic associations or not, otherwise the risk of deluding ourselves is very high. Alternatively, we could randomly select a subsample of sentences from the original dataset and closely inspect them to determine the extent to which our initial intuition about the results is true. We can also estimate the likelihood that the patterns that we see are due to chance. For this, probability values or p -values are often used. A p -value can be used to estimate the likelihood that the observed pattern is due to chance. This is different from the usage of p -values for hypothesis testing, which I describe in the next section of this chapter.

As noted in previous chapters, data-driven research aims at offering measured uncertainty. For this reason, it is important that some measure of uncertainty is offered. Instead or besides p -values, confidence intervals, effect sizes, and measures of accuracy (such as the F-score) can also be used (see the excursion in chapter 5 for an example of effect size reporting). But it is important to note that, in order to estimate the likelihood that a pattern is not random, we need a model that can tell us what would happen in this situation under random conditions. This means that we need to impose assumptions about what we would see in random situations and we need to critically inspect those assumptions. For example, if one assumes that the data is normally distributed, one would expect certain things to appear under random conditions. These would be very different if one were to assume an exponential distribution. The general audience books and technical papers by Nassim Taleb are good introductions to these problem (Taleb 2007; Cirillo and Taleb 2016).

Everything we choose to treat as data, and every procedure we perform on that data, is the result of specific perspectives on what is worthy of attention. This means that working with data can never be a fully objective endeavor. When confronted with this impossibility, we could throw our hands in the air and exclaim in despair, “all data approaches to the humanities are futile!” But we can also implement a more measured approach, guided by the principles of critical realism (see chapter 1). The fact that full objectivity is impossible doesn’t mean we can’t make progress towards consensus. If changing the definition of what counts as data in a given project leads to entirely different answers, then the onus is on the researchers to investigate what are the limitations of the data collection procedures used. If changing the settings of a given analytical

method yields different results, we need to figure out which are the best parameters. Finding foolproof configurations may well be impossible, but some parameters are better than others, and we can—and should—aim for improvement. This is why we need independent methods to verify data-driven results and to calibrate our data parameters. If our methods and data suggest the existence of a given trend, then we should find other approaches that might confirm or complicate these results. This can be achieved by fomenting a culture of replication where the goal is to collectively improve our methods rather than to dismiss someone else’s endeavors. Calibration methods, which use alternative approaches to model the same phenomena, or to manually verify a subset of computational results, should be central to our data research projects. We must also develop research environments that encourage the reporting of negative findings. Many computational projects will yield no conclusive evidence. Rather than overinterpreting vague patterns in the data, we need to be able to communicate research that fails to prove an initial hypothesis (this is something I attempt in the second excursion of chapter 6).

This short overview has described some potential sources of bias, but there are many more types of bias that might influence a computational result, and being on the lookout for those is part and parcel of computational research. However, seeking the best parameters and encouraging replication is not the only way in which we can proceed. In data-assisted research, patterns can be understood as generative devices. The “chair power” cluster in the example above might be used as a defamiliarization strategy to offer a fresh perspective on the corpus of theater reviews. Then, it doesn’t matter if the pattern used as a starting point is “accurate,” but the onus is on the researcher to use this defamiliarization to say something meaningful about the theater reviews under consideration.

Explanatory Analysis

The statistical methods surveyed thus far aim at uncovering structures and patterns in the data, and estimating the likelihood that they are not due to chance. But this is different from the explanatory statistical techniques that aim to test a hypothesis to establish causality. The gold standard, in this other scenario, is randomized controlled trials (RCTs). A common example is drug discovery, where a drug is administered to a group of individuals (the experimental group) and a placebo is given to another group (the control group). Cohort sorting should be randomized (i.e., any

participating individual has equal likelihood of being assigned to a given group) to try to minimize confounding factors (some underlying trait that would make an individual more likely to be assigned to a given group). This procedure was popularized in the early twentieth century when it was introduced to precision agriculture by Jerzy Neyman in 1923 (a translation of the original article is available in Splawa-Neyman 1990), but it is now routinely applied to all areas of science.

A common incarnation of RCTs is A|B testing, which is used in data-driven web design. Companies, such as Facebook, make design decisions based on experiments. Users around the world are randomly sorted into two groups, and they are exposed to almost identical versions of their products, which only differ by a single design element (perhaps a particular shade of blue in a button). Then a specific outcome is measured (perhaps the number of clicks on an advertisement), and if statistical analysis reveals a strategic advantage in one version of the design feature, it is then integrated into the product. Most of Facebook's interface design is based on such tests, and the company reportedly runs over a thousand such trials per day (Bakshy, Eckles, and Bernstein 2014).

However, RCTs are often impossible (or unethical) to carry out in the humanities and the social sciences. An alternative approach is identifying natural experiments, a strategy more common in social sciences such as economics and quantitative sociology. Natural experiments use events where people were "naturally" sorted into random groups, in conditions which approximate an RCT. In one famous example, researchers were interested in evaluating the impact in a person's life of attending an elite high school. In the United States, sometimes entry to elite high schools is determined by the grade in a standardized test. The cutoff point depends on the number of applicants in a year and is relatively arbitrary. Thus, it constitutes a good proxy for the random assignment of similar individuals into two groups. Researchers took two groups of people that had attained almost the same score in the test: those who barely made it and those who barely missed admission (differing by a single point). They assumed that both groups were roughly equal in terms of ability but one was admitted to an elite school and one wasn't, and this provided an ideal natural experiment. As the researchers found out, there was no significant difference in the academic and career achievements between groups in later life (Abdulkadiroğlu, Angrist, and Pathak 2014).

Whether experiments are controlled or natural, they are used in scientific research to identify causal explanations for phenomena. Causality

cannot be established outside the bounds of experimental conditions: researchers might be able to identify strong correlations between observations, but without experimental designs, it is impossible to assert underlying causes. A key tenet of the scientific method is that humans are not naturally good at distinguishing between causal relationships and spurious correlations. Take, for example, the graphs at the fabulously titled website <http://www.spuriouscorrelations.com>. One of my favorite instances shows a strong correlation between cheese consumption and death by sheet entanglement. This correlation is statistically significant, but that does not mean that cheese causes death by sheet entanglement (or the other way around). To establish such a causal relationship we would perhaps need to find a natural experiment (perhaps a region where cheese was banned) and see if it affected the number of fatal sheet entanglements as compared to a similar region where cheese was not banned.

As described above, to test drugs and medical procedures, scientists compare a group of animals or humans treated with a particular drug and one control group, where no such procedure was applied. The response of each individual won't be exactly the same but it will fall within a probability distribution. Effect sizes and *p*-values will be used to estimate the likelihood that the drug has an effect that is not due to chance. However, conventional hypothesis testing is the object of intense controversies and it is the result of specific histories. I will consider these controversies and histories briefly to show that statistics is not a unified field of practice but an area of intense epistemological debates. These discussions, in turn, can help us become more attuned to the affordances and limitations of statistics for computational theater research.

Hypothesis testing was developed in an earlier environment of less intense scientific activity. In the current landscape of frantic scientific research and thousands of scientific studies per day, the combination of certain social and mathematical conditions might mean that significance testing is inadequate to distinguish real results from false ones. The standard name for this statistical analysis is *p-value significance testing*. In the early twentieth century, when this procedure was developed, the word "significant" meant that something was signaled. As the century evolved, the word "significant" came to mean "important." As a result, sometimes a significant difference between samples seems to imply an important difference. However, when Fisher and others first used this term in the context of statistics, they merely wished to emphasize that there was an

identifiable difference, not necessarily an important or meaningful one (Salsburg 2002, 41). This has led to much confusion and unfortunate reports, both within academic circles and in the reporting of science to general audiences.

As statistics gained a more prominent role in all kinds of research, variations of the p -value significance testing came to dominate many areas of scientific inquiry, particularly in experimental research. Most experimental designs started using p -values to estimate the truth of a theory or the usefulness of a procedure. The most common variation of this method today is the Neyman-Pearson *hypothesis testing*, where a null hypothesis (no effect) is compared against a specific intervention (say, a drug). Although Neyman-Pearson hypothesis testing is now introduced in almost all statistical books and is enshrined at the top of scientific practice, its rise through the twentieth century was not unchallenged. The antipathy from Fisher to Neyman is well known (Salsburg 2002, 52–60), and the current primacy of a method that bears both their names would have baffled them. Neyman himself was not a blind proponent of simplistic hypothesis testing. He instead advocated using tests to distinguish between a family of distributions, but this more sophisticated approach did not make it into the statistics textbooks.

Hypothesis testing now runs rampant and is drawing criticism from many quarters. The abuse of the p -value has been exacerbated by the convention that there is a threshold for significance. This threshold is often set at 0.05, where results are considered significant when the p -value is lower than this threshold (in other words, when there is a 1/20 probability of getting a given result when the effect observed is not real, a ridiculously high threshold). This convention was set up in the early twentieth century, a time of less intense scientific practice. But now it has such a firm place in the practice of science that research groups would often waggle their data into significance. Some observers have also identified a practice termed *p-hacking*, where scientists try to coerce their data into p -value significance (Marasini, Quatto, and Ripamonti 2016; Farcomeni 2017). According to these observers, researchers should recognize that a p -value without context or other evidence provides limited information. For example, by itself a p -value higher than 0.05 offers only weak evidence that the proposed intervention has no actual effect (i.e., that the effect tested for, is false). Likewise, a relatively small p -value does not provide incontrovertible evidence for the correctness of a hypothesis; many other hypotheses may be equally or more consistent with the observed data. For these reasons,

data analysis should not end with the calculation of a p -value when other approaches are appropriate and feasible.

Some theorists argue that hypothesis testing as developed in the early twentieth century should be ditched altogether in favor of Bayesian analysis (Marasini, Quatto, and Ripamonti 2016, 320). Bayesian approaches (named after the seventeenth-century reverend Thomas Bayes), thinks of probabilities as ways to refine one's ideas about the world. Bayesian statistics require analysts to formally express their baseline beliefs (called a *priori odds*) before running an experiment. Calculating these odds effectively is a complex task that was all but impossible before the advent of the computer. Currently it is technically more feasible to run Bayesian statistical analyses, but there are several reasons why this is not yet a dominant trend. One is the lack of training, which is connected to the inherent difficulty of Bayesian statistical analysis. Another one is that frequentist statistics dominate journals, peer review, and industrial standards.

Beyond Description and Explanation?

This short detour into statistical history, shows that statistics is not a monolithic field, but a changing set of practices with many fierce and fascinating debates. Questions about the abuse of p -values and debates over Bayesian statistics are only scratching the surface. In the age of big data, some people have suggested that correlation is sufficient and that we don't need to understand causal mechanisms. A famous example is the provocative piece *The End of Theory: The Data Deluge Makes the Scientific Method Obsolete*, written by Chris Anderson (2008), former editor-in-chief of *Wired* magazine. This oft-cited essay stated that with enough data, the numbers speak for themselves. Anderson suggests that correlation supersedes causation, and science can advance without coherent models and unified theories. This point is also echoed by Mayer-Schönberger and Cukier's (2013) influential *Big Data: A Revolution that Will Transform How We Live, Work, and Think*.

However, many scientists have pushed back against this view. Using the aforementioned Ramsey theory and other mathematical explanations, Calude and Longo (2017) show that most correlations in big data will be spurious. Consider a paper on the dangers of relying only on correlation for medical diagnostics (Mullainathan and Obermeyer 2017). The authors use, by way of example, a project that applied ML to determine if patients arriving at the hospital were having a stroke. Using insurance claim data,

the ML algorithm found four unusual predictors of stroke: accidental injury, benign breast lump, colonoscopy, and sinusitis; in addition to more conventional predictors, such as cardiovascular disease. Subsequent qualitative analysis revealed the reason for the unusual predictors: they indicate that someone is a “heavy utilization patient”; in other words, someone who is likely to go to the hospital for relatively minor complications is more likely to have a stroke detected. Evaluating the ML results on their own, someone could conclude instead that sinusitis is a risk factor for stroke.

I here subscribe to the view that the only way to scientifically investigate causality is by devising experiments. In the humanities, natural or controlled experiments are both extremely rare. But if we were interested in scientifically studying causation in theater studies, we would need to devise experiments. Given the large number of possible factors that impact theater performances, it would be very hard to conduct such experiments. This doesn’t mean that we need to abandon the impulse to explain things. But perhaps we need to accept that a *scientific explanation* of theatrical phenomena is a very difficult goal. A scientific approach, of course, is not the only road we can take to arrive at explanations. We can also use reflection and intuition, as theater scholars have done for a long time, to explain theatrical phenomena. Even when working within a data-driven paradigm, it is important to recognize the extent to which our conclusions are grounded entirely on data and the extent to which we are extrapolating from the data, and speculating beyond what it actually reveals.

If we critically assess our data, and seek estimation and calibration, we can find accurate and objective patterns in the data. These observed patterns can be replicated and validated by other teams. But validating observed patterns is not the same as validating causal explanations. As I noted above, outside of experimental conditions, data can’t be used to establish causal explanations.

However, a data-driven methodology can help us estimate the correct description of a phenomenon. For example, we might be able to assert that the number of theater performances in a given city declined over time, within the bounds of a confidence interval. Explaining *why* the number of performances decreased is another matter. We might observe a strong correlation with declining population numbers in that city over the same period of time. This correlation might be very strong, but it won’t provide sufficient grounds for a causal explanation. We might be tempted to think that the decrease in the population caused the decrease in the number of

performances, but it might have very well been the other way around. The decrease in the number of performances might have driven people away from this city. Or perhaps both phenomena are the result of the same underlying cause, perhaps an economic crisis. An experiment might help settle the matter. A randomized controlled trial would be impractical and unethical (we would need to forcefully decrease the population in some cities and keep it stable in a comparable city to see what happens to the theater numbers). But perhaps natural situations could be found that approximate those conditions. Even then, cities are complex systems and it would be almost impossible to control for every possible confounding factor, and untangle true causality from a myriad of contributing causes.

To understand why performance decreased, perhaps the best method would be ethnographic: interviews and participant observation of different types of practices in the city. I am not suggesting that we abandon the pursuit of causal explanations. But merely that we recognize that data-driven methods can't carry us all the way. I should also stress that ethnography and data-assisted methodologies are very different. While ethnographers might refer to their notes as data, these notes don't fit the definition of data I use in this book. Also, ethnography cannot be carried out through computational procedures, as opposed to data-assisted methodologies which rely on them. To continue the example above, a data-assisted approach would consist of developing a performative visualization (see chapter 3) that shows, in a dynamic map, how the city loses its color and shape as the number of theater performances go down.

Institutional Bias and Statistics

This chapter has considered many sources of bias. But there is one more potential source that demands attention: the institutional conditions that arise in fields where statistics become prestigious. As theater moves into computational realms, it is important that we consider the challenges that statistics pose, and what we can do about them. Biostatistician John Ioannidis (2005) claims that research in many areas is only an "accurate measure of the prevailing bias" and that most published scientific research is false (700). He identifies six trends that decrease the likelihood that the findings are true:

1. The smaller the studies
2. The smaller the effect sizes

3. The greater the number and the lesser the selection of tested relationships
4. The greater the flexibility in designs, definitions, outcomes, and analytical modes
5. The greater the financial and other interests and prejudices
6. The hotter a scientific field (with more scientific teams involved)

We can find instances of all of these trends in DH. In particular, Ioannidis's trend 3, should give us pause. Digital methods make it almost trivial to run many iterations of the same procedure, tweaking parameters slightly to get different results. This is the very thing that makes digital research possible at a large scale. But the danger is that it becomes extremely easy to constantly adjust our methods until we get the results we want. An equally important problem is trend 4. If you can always reframe a problem in a slightly different way until it yields the results that you are after, you can easily trick yourself. This is why data-driven research requires agreed-upon definitions on what counts as data, if we want to reach intersubjectively verifiable conclusions. As noted in chapter 1, low ambiguity in the definition of data for a given research area is one of the four conditions for data-driven research. This is also why we need to find ways to calibrate our methods to avoid the clustering illusion and harking, as seen above. We should also recognize that DH is a "hot field" with significant financial and reputational incentives for finding new applications of digital methods (trends 5 and 6). But knowing this is a hot field is not a reason to stop our work, just to bring extra care to our methods. This is why being aware of the limitations of statistics is so important.

The only solution is to use statistical methods in a nuanced way and communicate our results with as much humbleness and skepticism as we can muster. On their own, the presence of statistics is not an unmistakable marker of objectivity. But under certain conditions, statistics can be used to help us avoid errors in our reasoning. We can turn again to Ioannidis (2014, 2) for suggestions to improve the likelihood that results are true:

- Large-scale collaborative research
- Adoption of a replication culture
- Registration (of studies, protocols, analysis codes, datasets, raw data, and results)
- Sharing (of data, protocols, materials, software, and other tools)

- Reproducibility practices
- Containment of conflicted sponsors and authors
- More appropriate statistical methods
- Standardization of definitions and analyses
- More stringent thresholds for claiming discoveries or “successes”
- Improvement of study design standards
- Improvements in peer review, reporting, and dissemination of research
- Better training in statistical literacy

These are all institutional suggestions. This means that, in order to generate better computational research we need to rethink the institutional environments that can aid in this project. Statistics have many roles to play in computational theater research: they can help us describe the data and find patterns. These patterns can be used as deformances and provocative defamiliarization strategies. Or, we can use models to estimate the likelihood that these patterns tell us something verifiable about the data. But there are also important limits to the usefulness of statistics, and we need to consider the many sources of bias in order to counter them. It will also be crucial that we report negative findings, rather than try to read too much into inconclusive data.

If critical realism informs our usage of statistics, we will be able to understand how cultural and historical conditions necessarily shape data collection. But we can still carefully use statistics for data-driven pattern discovery. We can also use statistical patterns as launching pads for useful data-assisted deformances, which can nonetheless contribute to a fuller understanding of theater history and practice.

The Roles of Visualizations

A data visualization is the graphical representation of data and its properties. In contemporary media environments, data and visualizations are almost inseparable from each other. Data could also be described solely by numbers and tables or conveyed through other means, such as aural *sonifications* or haptic *physicalizations* (Jansen et al. 2015; Moere 2008; Hogan 2015) but these are less commonly used.

Visualizations offer powerful insight into the data they represent. But visualizations can also be intentionally distorted to suit a particular narrative. Unpacking the ethical and political implications of visualizations is particularly urgent when data on national economies, public health, and environmental destruction are often communicated to specialized communities and to the general public mainly through visualizations. Alberto Cairo (2020), a leading thinker and practitioner of data visualization writes that “data visualization is a technology—or set of technologies—and, like artifacts such as the clock, the compass, the abacus, or the map, it transforms the way we see and relate to reality” (17). As Kennedy and Engebretsen (2020) note, visualizations are “cultural artifacts with distinct semiotic, aesthetic, and social affordances” (20). Data visualizations privilege certain modes of knowledge. For starters, they overemphasize the sense of sight, and pose a problem for inclusive design agendas. Understanding data visualizations requires a specific kind of literacy, which can lead to different types of inequality. As nation-states increasingly make decisions based on data, uneven levels of visualization literacy have troubling implications for democracies (Snaprud and Velazquez 2020). Visualizations prioritize correlation over causation, and this can entrench certain types of bias (Rettberg 2020, 41–42). In chapter 2, I discussed the problem of untangling correlation and causation from the point of view of statistics.

In this chapter, I will unpack the implications of data visualizations

for computational theater research against the background of theoretical discussions in statistical graphs, data journalism, and DH. Both data-driven and data-assisted research rely heavily on visualizations, but they use them in different ways.

Different Types of Data Visualization

To better understand how data-driven and data-assisted perspectives incorporate visualizations, we can refer to a parallel distinction between two communities that use data visualization: scientific and journalism communities. This distinction is not absolute, and the borders between these communities are porous (Kennedy and Engebretsen 2020). Still, there are some glaring differences in the ways both communities approach data visualization. The statistical graphs favored by the scientific community place a prime in the precise communication of quantitative information and include visual representations of statistical features (measures of dispersion and outliers). The scientific community tends to rely on standard graphical representations such as boxplots, scatterplots, and violinplots (some of these will be described later and used in chapters 4–7).

Statistical graphs are strongly associated with the work of Tukey (1977), as they are useful for EDA (as seen in chapter 3), but they have a longer history (Friendly and Denis 2001). Boxplots, which were popularized by Tukey were probably first used by Mary Eleanor Spear (1952), and scatterplots date back to the early eighteenth century (Friendly and Denis 2005). The work of Bertin (1983), Cleveland (1985; 1993) and Wilkinson (1999) has also been influential to the development of scientific visualization (Friendly and Denis 2001). Throughout their history, statistical graphs have generally aimed to be parsimonious. Tufte (1983), one of the most famous theorists in this community, says that graphs should aim for a low data-ink ratio. In a digital visualization, this means that every pixel should convey an aspect of the data, and any ornamental features should be removed. Another key feature of statistical graphs is that they enable the comparison across different facets of the data (for example, the treatment given to different samples or differences in populations). This is usually achieved through “small-multiples”: embedded plots that share the same scaling across the *x* and/or *y* axis to enable easy visual comparisons (for an example, see the pairwise plot in figure 5.7).

While the scientific community uses statistical graphs to communicate systematicity and precision, journalists often use visualizations for audi-

ence engagement. Visualizations in the media rarely indicate confidence intervals or measures of dispersion (such as standard deviations or interquartile ranges). Engagement and visual variety are very important, and visualizations in the media are often made by artists. Creativity is deployed to invent new modes of graphical display, rather than relying on established graphical forms such as scatterplots. A common approach is the usage of *infographics*, charts that explore a semantic link between the data and their graphical representation, often in ways which are humorous or aesthetically innovative. One of the most influential figures in the development of *infographics* is Nigel Holmes, who worked as a designer for *TIME* magazine in the 1970s. In a well-known example, he used bar charts to compare medical expenses across different countries, and each country's bar chart was represented as a sick patient lying in a hospital bed.

Tufte (1983) uncharitably described approaches such as this as “chartjunk,” as these types of *infographics* have a high data-ink ratio and the artistic distortions often make comparisons difficult. However, the design community has articulated different justifications for the use of *infographics*. Ichikawa (2016), writing from a user experience perspective, suggests that effective visualizations must “keep the viewer anchored” (n.p.). Visual variation and creativity are not superfluous, but part of an engaging experience that can enhance the comprehensibility and utility of a visualization aimed at the general audience. The work of Florence Nightingale is perhaps one of the most famous precursors of this approach. Her 1858 *coxcomb graph* aimed at showing the number of preventable deaths in the Crimean War, prompting officials into action. This is a landmark work of information graphics, but modern theorists of statistical graphs prefer bar charts, or other visualizations that enable straightforward comparisons (Gelman and Unwin 2013). For this reason, pie charts are commonly despised by the statistical graphs community.

Not all visualizations in the media are *infographics*. A growing trend in the past decade is the development of data-driven interactive stories, which have been popularized by the *New York Times*, the *Financial Times*, and *FiveThirtyEight*. The objective behind these stories is to enable users more direct access to the evidence. Many of these stories are both engaging and statistically sound, so they problematize any simplistic distinction between journalistic and scientific visualizations. There are other possible ways to distinguish among visualization communities (see Kennedy and Engebretsen 2020), but the lines between them are increasingly blurred. Still, it is important to identify the goals of different visualizations, even

if they can't be fully mapped to specific communities of practice. Engagement and statistical precision are distinct, if not mutually exclusive objectives. Data-driven theater research usually relies on the established visualization conventions developed by the scientific community. Data-assisted visualizations share some features with data-rich reporting in the media, but their objective can't be described merely on the basis of engaging a lay audience.

To better characterize the difference between data-driven and data-assisted visualizations, I will borrow Thudt et al.'s (2018) distinction between exploration and explanation in data visualizations. The aim of explanation is to communicate the author's view, while the aim of exploration is to enable users to find their own story in a dataset. Data-driven visualizations rely on explanatory facets and data-assisted visualizations rely on exploratory facets. As we have seen, the goal of data-assisted research is not to settle questions, but to challenge the assumptions of a question and find multiple possible answers. Likewise, the exploratory facets of a visualization "provide readers with the flexibility to ask a variety of questions of the data, to personalize their experience of the story based on their own interests, and to view the data from different perspectives" (Thudt et al. 2018, 64).

Data-driven visualizations rely on the graphical conventions of scientific visualizations to estimate the likelihood that an answer is correct given the data provided, or to indicate the most accurate characterization of the data, describing its structure, or drawing attention to trends and clusters (as seen in chapter 2). In contrast data-assisted visualizations aim to challenge the political, ethic and aesthetic conventions of data visualization tropes. Data-driven analysis can work within the constraints of established visualization conventions. But data-assisted analysis requires critical interventions to invent new forms suitable to its purposes, and it is to these interventions that I will now turn to. I want to emphasize, though, that both data-driven and data-assisted research might depend heavily on visualizations, or use none at all, as data can also be communicated through tables rather than visualizations. Later in this chapter, I will dedicate substantial attention to the role of interactivity in data-assisted visualizations. Both data-driven and data-assisted visualizations can make use of interactive features. But the goals of data-assisted research are particularly well served by interactive exploratory features.

Rethinking Visualizations from a Humanities Perspective

“Can graphical means assist the humanities in the project of interpretation?”

—Johanna Drucker

A data-assisted visualization doesn't take the assumptions of visualizations for granted, and aims to make its own assumptions transparent to viewers. The work of Johanna Drucker is a particularly useful conceptual guide for this endeavor. Drucker (2014) defines *graphesis* as “the study of visual epistemology” through a “dynamic, subjective process.” Its purpose is to “expose and describe the principles for structuring knowledge through graphical form,” and it thus seeks to create methods that are “generative and iterative, capable of producing new knowledge through the aesthetic provocation of graphical expressions.” In other words, the point of *graphesis* is to communicate humanities knowledge through visual means, understood as key rhetorical devices in their own right; they are “primary methods of analysis” that “create the data, not just display it” (n.p.).

In setting up a theory of *graphesis*, Drucker argues that graphical means of communication have long had a central role to play, but this role has often been dismissed in Western traditions of thought. The most familiar graphical forms, such as books and letters, have been dramatically overlooked. As Drucker notes, “basic codes for reading are graphically structured.” This includes conventions for text, footnotes, table of contents, and marginalia. She invites us to take these codes seriously and give them the attention they deserve: “A margin isn't an inert space, but a field of defining tension between text and page edge, and exerts a graphical force in relation to other elements on the page” (n.p.). If we give due consideration to these seemingly neutral elements, we can mobilize them for the reinvention of new graphical languages suited to the purposes of humanistic inquiry. She finds useful precedents for such a project in the work of several artists, such as Vasiliy Kandinsky, Paul Klee, Bauhaus artists László Moholy-Nagy and Josef Albers, El Lissitzky, Piet Zwart, Jan Tschichold, as well as filmmakers Dziga Vertov and Sergei Eisenstein. They all explored the rhetorical force of formal elements such as repetition, discontinuity and variation. One of Drucker's most fascinating examples is the work of Ernst Fraenkel, who presented his analysis of Stéphane Mallarmé's *Un Coup de Dés* through very idiosyncratic graphical arguments.

These are things that we can apply to the creation of new visualiza-

tions for data-assisted research. Drucker urges humanists to not reproduce the visual tropes of scientific knowledge, but to deploy visual codes to inscribe knowledge as provisional, situated, and observer-dependent. Visualizations should “bring the interpretive sensibilities of theoretical inquiry to bear on these assumptions while also acknowledging subjectivity as fundamental to the conception and expression of knowledge” (n.p.). A central preoccupation of Drucker’s own work is the graphical representation of time. She suggests that calendars, for instance, model temporal elements rather than represent a “natural” condition of time. In her own interventions, Drucker has been interested in modeling “the multi-linear (forking paths), heterogeneous (varied in density, rate, and scale), and discontinuous (broken, repetitive) temporalities that are part of human experience” (n.p.).

Writing elsewhere, Drucker (2011) offers a “call to imaginative action and intellectual engagement with the challenge of rethinking digital tools for visualization on basic principles of the humanities” (n.p.). To illustrate this, she asks us to imagine bar charts that illustrate population changes in different fictional nations, which divide people by gender. Visualizations such as these are common, and therein lie their problems. They are so naturalized that they effectively hide the interpretive basis on which they are built. Drucker takes apart the categories that underpin this visualization to show there is nothing natural about concepts such as “people,” “gender,” and “nation.” In collaboration with visual artist Xárene Eskandar, she offers an alternative mode of representation, where the bar chart is reimaged to show tension and indeterminacy in its constitutive categories. For example, blurred lines show gender ambiguity inside the bars, rather than binary distinctions. Floating points between the bars of two nations indicate the porosity of a political border and the transient state of migrant workers, revealing the instability of the concept of nation. The visualization could also indicate culturally specific notions of personhood. For this, Drucker imagines a given country where “women only register as individuals after coming of reproductive age, thus showing that quantity is an effect of cultural conditions, not a self-evident fact.” Even though the bar chart is based on fictional countries, it shows how the humanities can do more than merely apply existing visualization frameworks and develop graphical approaches that bring interpretation to the fore. Graphical elements are not discrete bounded entities but “conditional expressions of interpretative parameters” (n.p.). Drucker’s article includes other examples that should be required reading for anyone interested in developing data-assisted visualizations.

Drucker is a proponent of extreme constructivism, for whom interpretation is at the forefront of any knowledge enterprise in the humanities: “Nothing in intellectual life is self-evident or self-identical, nothing in cultural life is mere fact, and nothing in the phenomenal world gives rise to a record or representation except through constructed expressions” (n.p.). While working entirely from a constructivist perspective is a potential avenue for theater researchers interested in data, this is not the only possibility. In this book, I champion approaches that combine realism and constructivism, by recourse to critical realism (see chapter 1). While Drucker herself disavows any traces of observer-independent realism in humanities research, her ideas can still be incorporated into a critical realist approach to data visualization. Inspiration for this project can be found in feminist data visualization.

In D’Ignazio and Klein’s (2016) formulation, feminist data visualization “rejects neither the scientific process nor quantitative ways of knowing the world” but aims to “see how all knowledge is situated, how certain perspectives are excluded from the current knowledge regime” (n.p.). Feminist data visualization shows that it is possible to pursue evidence-based answers while remaining attentive to the problematic categories on which evidence is built. This means that the designers of visualizations should rethink binaries, consider edge cases that problematize categories, and legitimize embodiment and affect. It will also be crucial for such a project to trace each datapoint back to its source. D’Ignazio and Klein also stress the importance of using historically and culturally specific modes of representation and overturning hierarchies of knowledge transmission through participatory design practices. They ask: “What kinds of terminology, symbols, and cultural artifacts have meaning to end users, and how can we incorporate those into our designs?” (n.p.). Perhaps we can find an answer to their question in the usage of “cultural probes” (Gaver, Dunne, and Pacenti 1999), where culturally significant visual conventions are incorporated into the graphical language of a visualization. In chapter 6, I give an example of a *kayon plot* (figure 6.4), a visualization I developed for describing Javanese theater performances, which incorporates statistical information, has a low data-ink ratio, and is still premised on an interpretive, culturally situated view. This is a possible way in which critical realism can inform the development of visualizations that are statistically sound, but which also bring interpretation to the fore.

Another possible way of combining statistical information and culturally specific information is through Manovich’s (2011) notion of *direct visualization*. According to Manovich, data visualization is conventionally premised

on the representation of data through “graphical primitives such as points, straight lines, curves and simple geometric shapes to stand in for objects and relations between them” (36). Spatial variables such as position, size and shape represent key differences in the data and reveal patterns and relations. Direct visualization replaces the graphical primitives with actual media objects, such as images and video. These representations can be spatially organized according to quantitative measures (such as similarity in terms of color palette). This approach works specially well for images, and Manovich’s examples include visualizations of magazine covers.

Ethics of Representation

Whether theater researchers choose to work entirely within a constructivist perspective, or a critical realist one, a key tenet of data-assisted perspectives is that visualizations are never neutral. Specific visual configurations—colors, shapes, style—have rhetoric implications, and are argument-altering, persuasive statements. These choices also have ethical dimensions. As Hepworth and Church remind us, visualizations are not mere janitorial work, as important decisions made during the process of visualizing data can “critically shape historical narratives” (2019, n.p.). To highlight the ethical dimensions of data visualizations, they critique two projects that map lynchings in the United States. *Lynching in America* (<https://lynchinginamerica.eji.org/>) was made by Google for the Equal Justice Initiative, a mass incarceration not-for-profit organization. In an interactive choropleth map, the shade of each county indicates the number of African-Americans who were lynched between 1877 and 1950. Hepworth and Church, while sympathetic to the objectives of this project, show that the choice of visualization conventions implies that these murders were limited to a range of districts in the American South.

The authors contrast this with another project, *Map of White Supremacy Mob Violence* (<http://www.monroeworktoday.org/explore/>), which they describe as more nuanced and comprehensive. This second map focuses on the entire United States to depict lynchings and mob violence as a country-wide problem. For this they use a point visualization rather than a choropleth. State lines are not visible and each incident is represented as a single dot. Additional information is provided for each record, which includes Native American, Mexican, and Chinese victims, rather than only African Americans. In comparing the visual style and the possible interactions afforded by each site, Hepworth and Church demonstrate

that subtle design differences enact very different arguments. Building on these observations, the authors propose a framework of ethical visualization practices that consists of several steps: defining, reviewing, collecting, pruning, describing, surveying, and previsualizing. Each step is grouped into phases: pre-data collection (defining, reviewing); data collection and curation (collecting, pruning, describing); and data visualization and argumentation (surveying, previsualizing, visualizing, publishing). By showing the context and source of each data point, *Map of White Supremacy Mob Violence* resists the tendency of visualizations to show aggregates rather than individuals. Interactive data-assisted visualizations such as this one often try to link the representation of the dataset as a whole to the story behind each data point. In doing so, they illustrate the tension between each unique individual and the categories into which it is grouped. As seen above, this is also one of the guiding principles of feminist data visualization.

Thinking about categories and taking apart graphical conventions are not the only considerations to ensure an ethical approach to data visualization. As Diakopoulos (2018) notes, computational operations (such as algorithmic derivation, filtering, aggregation, and normalization) also change the ways the data is presented and can alter the conclusions reached. Algorithmic derivation means that data is not directly visualized as found in a dataset, but subject first to a computational transformation. Normalization is the mathematical operation of dividing a quantity by another in order to offer standardization. For example, instead of showing the raw counts of theater venues per city, the number of venues can be normalized by population counts (i.e., divided by the population counts). I give an example of this in figure 7.1, where I consider changes in the number of performances over time. Diakopoulos suggests that annotation and interactivity are ways in which the impact of computational choices can be communicated to the audience. Explaining how the data was obtained and transformed is paramount for ethical data visualizations (see Gray et al. 2016). For this reason, making the data and the methods available for replicability is also important from an ethical perspective, regardless of whether a visualization is part of a data-assisted or a data-driven project.

Performative Data Visualizations

Scheinfeldt (2012, n.p.) suggests that the humanities are becoming more performative (especially DH). His examples are web-advertised public

events. With the advent of social media, humanities work is done in public more and more: “Increasingly digital humanities work is being conceived as much as event as product or project.” He concludes that there are performative dimensions which are changing the rules of scholarly communication: “Performance is a different ball game than publication.” Building on this insight, Bay-Cheng (2017) suggests that history is also being increasingly staged, and that data has played a major role in this process. History is becoming more performative in the sense that it is increasingly performed for a public.

Data visualizations are often publicly available. But there is another way in which we can consider them as performative. These digital artifacts can be performative in the sense that they enact situated, subjective responses. Such performative character is less directly linked to the public character of a visualization and more closely connected to its dialogic, interpretive character. Ramsay (2011) builds explicitly on McGann and Samuels’ *deformance* and also takes inspiration from the Oulipo artists to describe how data can be used for playful and performative transformations. A larger “performative turn” is also discernible in recent thinking about visualizations. For example, Parry (2019) proposes the concept of *enactment* to critique visualizations: “Visualizations also act. They do things. They are visual-verbal-numerical enactments.” Visualizations, in Parry’s view, enact specific relations, histories, and politics. Bearing this in mind when analyzing a digital interface “attends to questions of who, under given social, cultural, and political conditions, is allowed to live, to speak, and to act—questions, in other words, of who and what gets to matter” (n.p.). These principles are not only analytical categories, as they can also be mobilized for specific design practices, what Parry terms “enactment-intensive data visualization,” and which include archival contestation, conditional revelation, and fluid interpretation.

Despite these rich discussions, references to performative writing as understood in performance studies are largely absent in DH. This is a missed opportunity, as the writings of people such as Peggy Phelan, Della Pollock, and Soyini Madison, well known to theater and performance scholars, are directly relevant to the project of reimagining performative visualizations. Della Pollock (1998), characterizes performative writing as evocative, metonymic, subjective, nervous, citational, and consequential. Finding ways to transpose these characteristics to visualizations provides fascinating design challenges and invitations to think differently about visualization. What is a nervous data visualization? I understand nervous-

ness here as a provisional, ever-shifting relationship to knowledge. For Madison (2006), a key characteristic of performative writing is the way it interpellates its imagined readers: “Performative writing emphasizes the relational dynamic between writer and reader in a spirit of caring about the dialogic and communicative quality of the connection” (n.p.). Performative visualizations likewise emphasize the relational dynamic between designer and user (or “subject,” as Drucker would have it). We can also gather data on how people interact with our visualizations and further contribute to this dialogue. The fact that digital visualizations can also be revised and are essentially open-ended makes them closer to performance in the sense that they are fundamentally unfinished and unfinishable. For Madison (2006), “performative writing is evocative because it is a braiding of poetry and reportage, imagination and actuality, critical analysis and literary pleasure” (n.p.). Performative visualizations, in turn, braid visual art, data analysis, data science, and data hermeneutics.

An example can be found in Manovich’s (2013) evocative data essay on Vertov’s films:

The presentation is an experiment. Normally an academic article consists from text with a small number of illustrations. Instead, this presentation is a portfolio of a large number of visualizations, with text serving as the commentary. The presentation also does not advance a single argument or a concept. Instead, I progressively “zoom” into cinema, exploring alternative ways to visualize media at different zoom levels, and noting interesting observations and discoveries. We can compare its genre to that of travel writing, where the organizing principle is the writer’s movements through space. (Manovich 2013, 45)

Borrowing Manovich’s felicitous metaphor, we can think of performative visualizations as more akin to travel writing than to the scientific exploration of landscape. Travel writing, like generative visualization, is aimed at suggesting more than what it contains, to signal longing and possibility, and perhaps to encourage future travelers to challenge the conventional wisdom of known landmarks.

Interactivity

The ethical, interpretive, and performative characteristics of data-assisted visualizations can be brought forth in either static or interactive visualiza-

tions. But interactivity adds specific affordances to data-assisted visualizations. The presence of interactivity per se is not a marker that a visualization is performative, or that the project is data-assisted. But there are three important objectives that can be enhanced by interactivity: multiple perspectives, multiple scales, and thick context.

A key tenet of data-assisted approaches is that they should allow multiple interpretive perspectives. In an interactive visualization, each of these perspectives can correspond to a specific pathway. For example, when users click on different buttons, the data can be reorganized into different interpretive categories. In interaction design, these pathways are sometimes explained with recourse to Aarseth's (1997) *ergodic narratives*. An ergodic narrative requires non-trivial effort on the part of the users to traverse a narrative. In the original formulation, any choice in an ergodic narrative forecloses other possible paths. A reader who makes a choice will not know what would have been the outcome of treading a different path. This is usually not the desired case in interactive data visualizations. Bateman et al. (2017) offer a refinement and describe interactive visualizations as "ergodic yet immutable" (108). Ergodic yet immutable visualizations can present contradictory points of view, aiming to emulate the sophisticated disagreement of theater studies, and the multiplicity of perspectives central to data-assisted research. As Hiippala (2020) notes, interactive data visualizations "require ergodic work both in the form of exploration and composition, a feature which separates [them] from static information graphics and non-dynamic data visualizations" (287). In a sense, the work of *designing* data visualizations is always ergodic, as designers must choose among different graphical perspectives. But interactive visualizations extend this ergodic experience to the users.

Moving between different scales is also important for data-assisted visualizations where attention to individual data points is as important as the consideration of data aggregates. Interactive systems are well suited to enable shifts between different levels of granularity. When thinking about ways to enable changes between scales, it might be useful to consider adding rich-prospect features to a visualization. For Ruecker, Radzikowska and Sinclair (2016) a *rich-prospect browser* enables users to see every item of a collection at once while hiding some of its details. Rich-prospect also enables users to see how different organizational criteria reveal varying connections within a collection (or dataset in this case). Rich-prospect browsers were originally developed for the visual display of digital cultural heritage and aim to provide "insight into how the collection was under-

stood by the people who made it” (Ruecker, Radzikowska, and Sinclair 2016, 176). This feature can be extended to interactive data visualizations to reveal how the data was understood by the people who collected or organized it, and highlight the situated perspective of these designers.

Lastly, interactivity can be used to provide thick context. Interactive visualizations enable users to retrieve the source of specific data points. For example, a network diagram might aggregate information on thousands of performers. By clicking on a node in the diagram, the user can find additional details about the performer represented by that node, as well as a summary of interpretive decisions on how the data was coded, or a note alerting the user to missing data (a similar example, for fictional characters is presented in chapter 5). For inspiration on how to merge visualization with contextual explanations, we can look at Bach et al.’s (2018) *narrative design patterns* for interactive visualizations, which are available at <http://napa-cards.net/>. A narrative design pattern is “a low-level narrative device that serves a specific intent” (Bach et al. 2018, 111). One of these patterns, *humans-behind-the-dots* is particularly useful for adding thick context. The idea behind this pattern is presenting individual stories in response to users’ clicks on data points. Although Bach et al. are writing from a data journalism perspective, this approach can be easily extended to data-assisted theater research.

In sum, interactivity can enable a combination of scales, bring together multiple perspectives, and add thick context. When thinking about the role of interactivity in visualizations, it is important to note that not every single element will be interactive, and that it will not be interactive to the same extent. According to Thudt et al. (2018), interactivity can be present in different degrees as users are only allowed to control certain aspects of a visualization: *view*, *focus*, and *sequence*. *View* means that the user can choose what is represented and how it is represented. A high degree of interactivity within this aspect would allow users to choose the parameters of visual encoding: color, size, and spatial placement. In a less interactive version, the users might select what visualizations they want to display side by side to see different aspects of the data. *Focus* refers to the segment of a visualization that the user wants to bring to the fore. When there is a large amount of data to choose from, users can decide where to look by filtering, selecting, zooming, or panning. Lastly, *sequence* refers to the extent to which users can make choices on the progression and temporal order of a visualization. Users can be guided through a predefined journey through *steppers* or *strollers* (in a recent trend called *scrolly-telling*). In this

case, interactivity enables users to choose the pace of an interactive story. But a higher degree of interactivity can be achieved when the users decide not just the pace but the next destination of a data journey.

When a visualization enables a high degree of interactivity for all of these aspects, it becomes more than just a visualization. Using the conceptual lens of theater studies, we can think of these as *intermedial essays*. The combination of scale, perspectives, and contexts requires not just interactivity but the co-presence of different medial forms: essays, videos, sounds, and graphs. Thinking of them as intermedial essays as opposed to data stories (a common term in data journalism) highlights their artefactuality and the media specificity of each of its components. As Mee (2018) notes when reflecting on the potential of web-based theater scholarship, “digital platforms allow our scholarship to embody our argument” (8). An intermedial essay can exist in between different formats, and this is similar to the way intermedial performances exist in between media (Chapple and Kattenbelt 2006; Bay-Cheng et al. 2010). Intermedial performances reflexively combine different media, and intermedial scholarship is well suited for the reflexive combination of data, textual arguments, and multimedia. For a performance to be intermedial, it must retain some aspects of theater, but also include some aspects of other media forms. Likewise, intermedial essays might still maintain the conventions of written essays, but also include videos and data visualizations.

A key aspect of intermedial essays will be self-reflexivity. Chapple and Kattenbelt (2006, 11) argue that “a self-conscious reflexivity that displays the devices of performance in performance” is an essential feature of intermedial performances. In other words, not all performances that use media are intermedial, but only those that integrate a certain level of self-reflexivity. Likewise, intermedial scholarship should include a level of reflexivity with regards to its own mechanisms, and make users aware of its own devices. A good example is the embedded “hermeneutic toys” by Rockwell and Sinclair (2016) discussed earlier in this book.

One of the most notable theater examples is Erin Mee’s *Hearing the Music of the Hemispheres*, which aims “to offer an alternative, performance-driven model for understanding spectatorship” by combining multimedia objects and arguments using the Scalar platform (2013, 149). Another influential set of intermedial essays is the Hemi Press’s *Gesture*, a series of “evocative digital works that combine multimedia and writing to make an original critical intervention in the fields of performance and politics”

(Hemispheric Institute 2020). A short excursion that prefigures some projects I will describe in part 2 is my intermedial essay “Wayang Kontemporer: Innovations in Javanese Wayang Kulit” (Escobar Varela 2015, available online at <http://cwa-web.org/dissertation/wayang-dis/>). A short video explanation of this project is available as video 3.1, in the web companion to the present book. *Wayang Kontemporer* includes videos, essays, and interactive visualizations. The different components are organized in two major “areas,” which can be described as *canvases* following Bateman’s (2017) semiotic vocabulary for interactive platforms. One canvas is dedicated to essays with embedded videos, and the other to interactive visualizations. As opposed to the majority of the visualizations I will describe later in this book, the visualizations in *Wayang Kontemporer* are entirely premised on interpretive methodologies (they are the result of fieldwork and interviews rather than computational transformations). In that project, I analyzed twenty-four performances, which I classified across five dimensions. The explanation for how each dimension was selected is described in a series of essays in the *essay canvas*. A series of radial charts can be loaded into the *visualization canvas*, and the user can overlay them on top of each other in order to make comparisons across the performances.

Since the users can choose which performances to compare, this interface affords an interactive *view* (in Thudt et al.’s terminology). The radial chart shows how each performance was classified along the aforementioned five dimensions. By placing the mouse over the visualization, the user can see a short explanation of why the performance was classified in such a way. Clicking on the visualization loads a longer essay that explains the interpretive decisions behind each visualization in greater detail. This is an example of how a visualization can provide thick context. The visualizations in *Wayang Kontemporer* constantly refer to the essays, and vice versa, but it is up to the users to follow such links. There are many possible pathways between the components of this intermedial essay, and the users can choose how to traverse those pathways. In intermedial essays, the distinction between visualization and interface collapses, and this poses major challenges for the digital sustainability of such projects (see chapter 8). But even simpler, static visualizations can achieve the interpretive and performative goals of data-assisted visualizations (think of Drucker’s bar charts).

This chapter has shown how data-assisted visualizations can be used to poke around the edges of a question, while data-driven visualizations

are used to offer the most plausible visual description of a dataset under a set of assumptions. These assumptions can in turn be interrogated through data-assisted means, visual or otherwise. The chapters in the next section show a variety of examples of how data visualization and data analysis are used to ask—and answer—a wide range of questions relevant to theater scholarship.

PART 2

Guided Tours

CHAPTER 4

Words as Data

The digital analysis of word patterns has the widest array of tools, methods, and theoretical perspectives of any area covered in this book. Techniques derived from computational linguistics and corpus analysis are often used in DH, and have been extensively applied to the study of dramatic texts (Craig and Greatley-Hirsch 2017). While I survey this type of work in some detail, my interest here is to argue that the same methods could be applied to study a variety of textual materials related to theater practice: advertisements, playbills, reviews, casting calls, production notes, academic articles, etc. This is currently an underexplored but promising area for computational theater research.

Text, as discussed in this chapter, denotes a collection of words rather than signaling more expansive definitions which could include diagrams and other visual items. Many researchers are interested in applying formal models to answer narrowly construed questions related to authorship and other similarly closed questions, whereas other researchers are drawn to the speculative potential of computers for the analysis of literary words (positions I have characterized as data-driven and data-assisted, respectively). Here I will describe four common methods—dimensionality reduction, time series analysis, measurement of linguistic differences, and topic modelling—and then explore how these methods can be used within data-driven and data-assisted methodologies.

Methods for the Analysis and Visualization of Texts

Principal component analysis (PCA) is the most common procedure within the family of dimensionality reduction techniques, and it is routinely used for authorship attribution, and computational stylometry (Binongo and Smith 1999; Eder 2015; Oakes 2014; 2017). When working with textual

data, each word in the corpus can be treated as a dimension. Each text can be thought of as a series of coordinates in this multi-dimensional space, where the value for each dimension corresponds to the number of times each word is present in that given text. With hundreds or thousands of dimensions it is hard to identify patterns in those numbers. PCA aims to make this easier by representing the words in a lower-dimensional space. These new dimensions, called principal components, aim to maintain most of the variance in the original multi-dimensional space. But having less dimensions (typically two) makes it easier to graph and analyze patterns in the data. For example, Schöch (2016) has used PCA to identify word usage clusters across genres in French drama. Craig and Greatley-Hirsch (2017) offer a more comprehensive definition of PCA than the one given here, and provide many applications to the study of dramatic texts.

Another common approach is the visualization of word usage changes over time. This approach rose to prominence with a paper by Michel et al. (2011) that used millions of datapoints from Google Books to identify trends in word usage patterns, and interpreted such trends as evidence of cultural change. This project relied extensively on line charts (like the one in figure 4.1) to visualize the change in word frequencies over time. Kulkarni et al. (2015) have developed a more robust set of techniques to identify statistically significant trends in word changes that introduces the kind of calibration checks I described in chapter 2. For example, they generated many alternative versions of a Wikipedia corpus where they artificially introduced word changes to verify that their proposed method could detect these shifts. In the excursion below, I describe a simpler statistical procedure for distinguishing word patterns over time and apply this to the study of a corpus of theater reviews.

A key interest of many digital literary scholars is developing measurements for comparing two texts. Some of these measures have been imported from corpus linguistics, such as the *type-token ratio*. To obtain this measure, we divide the total number of unique words by the total number of words in a text. The resulting number gives a sense of the lexical richness of a text; the bigger the number, the more varied the vocabulary. Many quantitative measurements have also been specifically developed for DH purposes, and the most influential is perhaps Burrows's (2002) *delta*. This and related measures were initially used for authorship identification, but are now commonly deployed to trace differences among sections of a text or among authors, genres, places, and times (Burrows 2007; Juola 2018) and they have also been extensively applied to the study of dramatic texts (Craig and Greatley-Hirsch 2017).

Topic models can be understood as groups of words that tend to occur together. The name might be misleading, as these “topics” might not represent semantic unities, or “belong to a common abstract theme (such as justice or biology)” (Schöch 2017, n.p.). Topic models have been applied extensively in digital literary criticism (Jockers 2013; Meeks and Weingart 2013; Piper 2018). Schöch’s (2017) example is particularly relevant to theater studies, as he constructed a topic model on 391 French plays from the classical age and the Enlightenment. Of the topics he found, some are thematic and others are related characters, recurring dramatic actions, or settings. Some of the findings confirm existing views while others hint at new hypotheses. With characteristic DH reserve, Schöch tempers the optimism of his results with a balanced explanation of the technical and interpretive limitations of this study. Although the corpus is of a decent size, and of the scale far beyond what a scholar could read in a short time, it’s still far from an ideal representation of the period in question. He was able to use this corpus, because of excellent efforts to encode French drama in free and open text encoding initiative (TEI) formats (more on which in a moment). The question of genre, of enormous importance to literary history, is also well-suited to the possibilities of topic modeling. Thus, here we see the happy confluence of available data in reusable formats, relevant methods, and questions pertaining to the specific intellectual history of French drama.

At this juncture, I want to bring attention to the importance of TEI, a series of guidelines for the digital markup of textual phenomena. Markup tags are divided into structural components (speaker parts, acts, and scenes), renditional aspects (font, size, hue, etc.), logical and semantic features (names, dates, addresses), and analytic features (notes and annotations). Consider Hamlet’s soliloquy, as encoded in TEI, taken from the Folger Digital Library (Mowat et al. n.d.):

```
<speaker xml:id="spk-1762">
  <w xml:id="fs-ham-0271840">HAMLET</w>
</speaker>
<l xml:id="ftln-1762" n="3.1.64">
  <w xml:id="fs-ham-0271850" n="3.1.64" lemma="to"
    ana="#acp-cs">To</w>
  <c> </c>
  <w xml:id="fs-ham-0271870" n="3.1.64" lemma="be"
    ana="#vvi">be</w>
  <c> </c>
```

```

<w xml:id="fs-ham-0271890" n="3.1.64" lemma="or"
  ana="#cc">or</w>
<c> </c>
<w xml:id="fs-ham-0271910" n="3.1.64" lemma="not"
  ana="#xx">not</w>
<c> </c>
<w xml:id="fs-ham-0271930" n="3.1.64" lemma="to"
  ana="#acp-cs">to</w>
<c> </c>
<w xml:id="fs-ham-0271950" n="3.1.64" lemma="be"
  ana="#vvi">be</w>
<pc xml:id="fs-ham-0271960" n="3.1.64">--</pc>
<w xml:id="fs-ham-0271970" n="3.1.64"
  lemma="that" ana="#d">that</w>
<c> </c>
<w xml:id="fs-ham-0271990" n="3.1.64" lemma="be"
  ana="#vvz">is</w>
<c> </c>
<w xml:id="fs-ham-0272010" n="3.1.64" lemma="the"
  ana="#d">the</w>
<c> </c>
<w xml:id="fs-ham-0272030" n="3.1.64"
  lemma="question" ana="#n1">question</w>
<pc xml:id="fs-ham-0272040" n="3.1.64">:</pc>
</l>

```

Hamlet is identified as the speaker, within the <speaker> . . . </speaker> tags. Each line is marked within the <l> . . . </l> tags. Individual words, nested within the line element, are marked within the <w> . . . </w> tags, which also provide information on the word's lemma, via the lemma attribute. This data can then be easily extracted and processed for a wide range of research purposes. As a result of large-scale digitization and encoding projects in literature, many researchers have access to textual data encoded in consistent TEI formats, and many software tools have been developed specifically to work with TEI-compliant files. Even then, TEI is not a panacea, as most researchers need to reformat the data to suit their specific needs and there are always omissions and inconsistencies in large textual collections. That being said, TEI provides an excellent data model and many analyses of dramatic literature build on TEI-compliant digital

text collections. But TEI can also be used to encode other types of materials relevant to theater research which are not necessarily playscripts. For example, *The Harry Watkins Diary: Digital Edition* is a TEI-encoded collection of diaries written by Harry Watkins (1825–1894), who regularly recorded the plays he saw between 1845 and 1860, and which is a digital companion to *A Player and a Gentleman: The Diary of Harry Watkins, Nineteenth-Century US American Actor* (Hughes and Stubbs 2018). Other textual materials relevant to theater research might not be encoded according to consistent data models, but many theater reviews and other texts are freely available online in machine-readable formats, and this is a feature that theater scholars could take more advantage of.

In the preceding overview, I have described some examples of dimensionality reduction, time series analysis, measures for linguistic comparison, and topic models. But it should be noted that there are many possible techniques within each of these approaches. Dimensionality reduction can be achieved through the rather traditional PCA method just described, but there are many other possibilities. Increasingly, researchers aim to arrive at dimensionality reduction through ML, and this approach doesn't require a predefined statistical model (Breiman 2001). Rather, inferences are made based on large amounts of data. There are many kinds of ML and these techniques can also be applied to time series analysis (Karsdorp et al. 2020) and the identification of linguistic clusters. Here, I just want to hint at this expansive area with many things to come—more extensive discussion of ML for textual analysis research can be found in Piper (2018) and Underwood (2019a).

The four major methods just surveyed, and combinations of them, can be applied to a wide range of projects and questions. One example is the analysis of text reuse: measures for comparing texts and for tracing change over time can be used to study how portions of a text have been copied, paraphrased, or cited in other texts. Sometimes text reuse is focused on how quotes (e.g., of sacred texts or theater plays) are propagated through time. Another major task of text reuse is alignment: the visual representation or description of parts of two or more texts that are identical (or at least equivalent). Alignment is particularly important for the study of translations, or the study of multiple versions of foundational texts. An example from theater is an analysis of thirty-seven different German translations of *Othello*, offered as a test case of the *Version Variation Visualization* project (Cheesman et al. 2011). Although text reuse is a growing area elsewhere in DH, it has yet to be more consistently applied

to the study of theater. It would be fascinating to see to what extent similar passages are repeated throughout theater scholarship, publicity materials, or theater criticism. Text reuse is relevant for questions that are less literary in nature, such as the study of adaptations for stage performance or patterns of textual transmission in oral traditions. Text reuse could be of interest to scholars working in theater traditions where improvisation based on textual formulas is important.

Procedures of increasing complexity are being developed for the computational analysis of texts. But there is still scope for relatively simpler methods. Take for example, *Style Inc.*, where Moretti (2009) analyzes seven-thousand titles of British literature. He looks both at the changes in the number of words per title and the types of words in the title to explain changes in the evolution of novels in the eighteenth and nineteenth centuries. Using simple statistics and visualizations, Moretti discovered that titles became progressively shorter, and that they also tended to become more abstract in nature. This kind of analysis can also prove inspirational for theater studies because of its source of data. Although Moretti's article is ostensibly about literature, it doesn't focus on the novels themselves, but on the registries of their titles. The availability of playbill records could constitute a similarly fertile territory for the kinds of analysis proposed by Moretti. There are several playbill repositories being compiled at the moment (Prince Lab for Digital Humanities n.d.; UPenn Libraries n.d.; NYPL/Zooniverse n.d.) and this might be an area of future growth. Of course, the kinds of questions that one could ask of such repositories are different from the ones warranted by a registry of book titles. In the case of novels, we can assume that each one is unique. In contrast, when it comes to theater playbills, one would expect that many titles are repeated. The reoccurrence of specific titles could be analyzed over time, as has been done for the study of the records of the Comédie Française Registers Project (York 2017). Compiling and arranging playbill data is a complex interpretive endeavor and not a straightforward process of transcription (Vareschi and Burkert 2017). While playbills provide a fascinating source of data, they are also highly cultural artifacts. Playbills make sense within certain production systems but not within others. They are commonly employed for commercial, experimental, school, and community theater productions in many parts of the world. But traditional theater performances (for example, in Southeast Asia) are rarely accompanied by playbills.

Data-Driven Text Analysis

A good example of a closed textual question, of the kind favored by data-driven analysis is: who is the author of a text with disputed or unknown authorship? This has been perhaps one of the most popular questions to be explored in digital literary analysis. The history of these questions goes back to the early 1950s, and they continue to be hotly debated in DH conferences today. An astounding variety of ML and classical statistical methods have been applied to study these questions, and all the methods seen above can be used for authorship attribution. This is not the space to review all possible approaches in full, as my objective here is just to show that the question of authorship is fundamentally scientific in the way it is pursued. There is a clear hypothesis, in response to explicitly formulated questions.

A favorite target of this analysis is Shakespeare. This fascination is closely linked to specific intellectual and cultural histories. From today's perspective, the mystery of Shakespeare's disputed authorship for certain texts is both unknown and important. Authorship in Shakespeare's time was collaborative and the existing records don't paint a definitive picture on authorial contributions. It is impossible to imagine that the kinds of questions stirred by Shakespeare's unconfirmed authorship would have much sway in the cultures and contexts not intrigued by individual authorship. Perhaps the question of authorship is so tantalizing because it's both not fully answered but still within our grasp—there is data that could, in principle, be processed to yield a satisfactory answer.

The data-driven analysis of text can be linked to *distant reading*, a term that was popularized by the Stanford Literary Lab. While distant reading sometimes relies on computational techniques, many influential distant reading projects rely on systematic reading by human annotators. The opposite of distant reading is not “reading,” but a literary history premised on seemingly haphazard anecdotes. Distant reading is mostly concerned with the systematic analysis of a wide sample of literary texts. In that sense it is less directly linked to the history of computational methods and closer to the tenets of the social sciences. This point is emphasized by Underwood's (2017) intellectual history of distant reading, which includes the work of Janice Radway (1991), whose analysis of romance novels is premised on conversations with a sample of readers. Underwood's own work also includes many instances of systematic reading in combination

with computational techniques (for example, Underwood 2017), which he also labels as distant reading. Tracing similar ideas between the social sciences and distant reading is not the only genealogy possible, and distant reading has also been linked to book history (Bode 2012). Distant reading is data-driven in the sense that it is premised on systematically gathered data (even if the methods of collection or analysis are not digital). Matthew Jockers's (2013) *macroanalysis* is more explicitly modeled on the quantitative analysis of digitally collected data. He doesn't dismiss the value of actual reading, but sees macroanalysis as an eminently quantitative approach that seeks to achieve a different objective: "we might think about interpretive close readings as corresponding to microeconomics, whereas quantitative distant reading corresponds to macroeconomics" (25). Macroanalysis, in contrast to macroreading (systematic reading), belongs only to the realm of computers. It can be used to explore a variety of questions, such as "the historical place of individual texts, authors, and genres in relation to a larger literary context" or the ways literary themes wax and wane over time (27). Both concepts—macroanalysis and distant reading—advocate for looking at literature from afar, and from considering a larger volume of texts that what is commonly used for conventional literary analysis. Macroanalysis and distant reading are also useful for the study of a range of theatrical texts, from dramatic literature to critical responses.

Data-Assisted Text Analysis

The potential for the speculative, interpretive, and situated analysis of words has been articulated in several influential theoretical perspectives. Stephen Ramsay (2003; 2007; 2011) has consistently argued that literary critics can reproduce their procedures on a computer, an approach he has named *algorithmic criticism*. This approach makes the steps of criticism, which are usually hidden, available for scrutiny and reproduction. Following Wittgenstein, he compares literary method to "a ladder that is discarded after one has used it to climb up" (Ramsay 2003, 171). Algorithmic criticism is the digital examination of this ladder, as it aims to reproduce on a computer the operations that a literary critic would carry out in their analysis. Unlike the purveyors of data-driven analysis, Ramsay is not interested in settling questions. Rather, his aim is to ensure that discussion of relevant literary works continues into grater depths.

Another influential data-assisted approach, which has also been

alluded to earlier, is Rockwell and Sinclair's (2016) *Hermeneutica*, which are defined as interactive, interpretive toys that can be embedded into digital essays. Rockwell and Sinclair are interested in how interactive visualizations can “add to a history of interpretation” rather than offer definitive answers to closed questions. Their web portal Voyant (<http://voyant-tools.org>) enables researchers to easily upload and visualize texts, creating interactive charts that can be embedded in other websites. By embedding a Voyant visualization in an intermedial essay, an author can exhort readers to explore multiple perspectives, multiple scales, and thick context (see chapter 2). An excellent example of these interactive affordances, included in Rockwell and Sinclair's companion website, is the analysis of two speeches on race, one by Barack Obama and one by Pastor Jeremiah Wright. Readers can find alternative ways to visualize the data, read the speeches in full and ask other questions of the same dataset used by the authors.

Voyant is a favorite of many researchers and is often used in the context of education (I, too, use it for my Intro to Digital Humanities course, as well as for the research project described in this chapter's excursion). The ease of use is one of the appeals: the platform automatically creates a set of visualizations and statistics. The visualizations and statistics are mostly inspired by corpus measurements, such as keywords in context (KWIC), collocate analysis, word clouds, and a time series visualization of terms in a corpus. The interface is customizable and enables users to select their own stop words (i.e., function words such as “the” that should be omitted). New features and tools are constantly being added, and at the time of writing this book, the portal has started offering maps and topic modeling. While the portal is excellent for the analysis of English-language texts, its assumptions about word boundaries delimited by spaces and consistent spelling make work difficult in languages where these characteristics are not found, such as in some South Asian languages (Battacharyya 2018). This is by no means a problem endemic to this platform, and digital text analysis is plagued by similar problems. Any tool is built on cultural assumptions—and this should be taken into account when applying it to any object of study.

Other platforms are more explicit on the assumptions on which they are built. While Voyant Tools relies mostly on automated procedures, CATMA (computer-assisted text markup and analysis; Meister et al. 2016), at <http://catma.de> enables users to manually or semi-automatically encode features of interest, rather than relying on its predefined algorithms. It

thus places greater emphasis on interpretation as a provisional and situated process and also enables collaboration, as different users can create their own usernames and collaborate in a project. This can highlight differences in interpretation across individuals, contributing to the multiplicity of perspectives central to data-assisted research. Both Voyant and CATMA enable generative and iterative explorations of texts, and can help scholars resituate questions of interest, rather than offer definitive answers.

As I have argued throughout this book, the same methods can be used within different methodologies. Data-assisted analysis might rely on delta scores, topic models, and time series as much as data-driven analysis, but it is premised on a different perspective, one that seeks to ask questions in new ways rather than to find better answers to previous questions. For example, the inquiries in *Literary Detective Work on the Computer* (Oakes 2014) are framed as closed questions on plagiarism detection, style, authorship, and decipherment. These are explicit questions that use straightforward methods, even when the answers might be inconclusive (as is the case in much scientific work). Is x the author of y ? Does genre matter more than personal style for the vocabulary choice of a given set of authors? What does a given set of graphic characters mean in an unknown language or secret code? These results require interpretation, but the quantitative methods used (for example, the aforementioned PCA to distinguish words that vary among authors) are aimed at trying to establish the answer to the question in empirical, replicable, and verifiable means.

This is different from the kinds of examples one finds in the previously mentioned *Hermeneutica* (2016). In the chapter “The Swallow Flies Swiftly Through: An Analysis of Humanist,” the authors look at the archive of *Humanist*, the online discussion group led by Willard McCarty that was central to the DH community at the time *Hermeneutica* was written. Rockwell and Sinclair were interested in analyzing how *Humanist* changed over the years. Several graphs showed changes in lexical choices over time (from 1987 to 2008), to indicate that the frequency of phrases such as “humanities computing,” “computing in the humanities,” and “digital humanities” had varied. The latter had been clearly rising in the years leading up to 2016. The authors also showed that certain disciplines, such as “Visual and Creative Arts” were in an upward trend, whereas “Classics” were becoming less often discussed. In their analysis, they used methods also espoused by Oakes, but the way the questions were posed is different. Rockwell and Sinclair were interested in how the discussion list had

changed. We find the same exploratory approach and *how* questions in their analysis of Hume's *Dialogues*, discursive changes in the *Game Studies* journal, and in the analysis of race in the two political speeches mentioned earlier. In all cases, the methodology is an exploratory analysis of a given corpus. In contrast, in the analyses carried out by Oakes, the methodology is less exploratory and less interested in helping answer *how* questions, and more focused in answering whether two things are similar or different: are two texts similar enough to warrant the conclusion of identical authorship? As Ramsay (2007) notes, there is disagreement in the closed questions posed by science, but the assumption is that there is a singular answer to a given problem. In contrast "literary criticism has no such assumption. In the humanities, the fecundity of any particular discussion is often judged precisely by the degree to which it offers ramified solutions to the problem at hand" (489).

Many researchers, though, weld data-driven inquiries together with data-assisted interpretation. Martin Paul Eve (2019), for example, developed a "computational microscope" to study a single novel: David Mitchell's *Cloud Atlas* (2004). There are many reasons this choice of novel is fascinating: differences in editorial intervention between a U.S. and a U.K. version, as well as the way the novel crosses genres and historical settings. Also, due to copyright restrictions, Eve had to retype the entire novel in its different versions, and this made him all the more aware of minute differences that he could then try to corroborate or expand through computational techniques.

Theatrical problems come in multiple guises, some of which will require closed questions to be settled computationally, whereas others will be better served by systems that aid interpretation. Let's imagine that we have a corpus of reviews written by theater critics in a given place (as explored in my excursion at the end of this chapter). Different questions could be asked about the same corpus. An example of a closed question could be: is there a difference in the way experimental and commercial productions are described? A question that could be better answered through an interpretive analysis would be: how has theater criticism in this place changed over time? In both cases, we might be using the same data, and some methods could be used to answer both types of questions. But the crucial difference is on how the question is framed, and how the answers are treated.

Two decades before the publication of this book, a framework for using corpus tools to study theater reviews was proposed by Roberts and

Woodman (1998). The article described the creation and analysis of a preliminary corpus of British theater reviews. Their database and software seem rudimentary from today's standpoint, but the authors identify an area of great promise which has sadly not been taken up by theater scholars, or by corpus linguists. They suggest, for example, that "world" is a more frequent term in theater reviews than in common English language usage. The prominence of "world" might suggest a "coherent directorial or design concept which is almost always, in the texts used for this project, implicitly naturalistic in orientation" (12). Whether this particular insight is of relevance or not, it seems clear that theater reviews constitute an extraordinary data source that could be easily analyzed through computational techniques. The work of Roberts and Woodman has lingered in silence for two decades, and I could not find other research teams following in their footsteps. But perhaps the wider interest in corpus analysis spurred by DH will mark a new beginning for the computational exploration of theater reviews.

In this chapter, we've seen corpus tools applied to study British theater reviews, topic models deployed to study genre in French drama, and text reuse techniques mobilized to identify variations in German versions of *Othello*. Given the variety of methods commonly applied to the study of textual phenomena, existing theater examples are merely scratching the surface of a wide range of possibilities. We can further mobilize this plethora of techniques to the study of dramatic texts but also extend their reach to the analysis of program booklets, tweets, and other media reactions to theater performances.

These approaches are particularly useful for texts do not lend themselves well to sequential reading because of their disjoint structure. Take, for example, marginalia. A series of articles in *Leviathan: A Journal of Melville Studies* (2008) used digital text analysis to analyze the marginalia written by Melville in the works of Shakespeare, Milton, and Homer. The marginalia are not beyond the scale where a single reader could easily read everything, but their disjoint structure means that it is hard to detect patterns in them. Digital text analysis revealed that Melville annotated Shakespeare's work much more than he did other writers, and that the passages annotated in Shakespeare tended to contain words with low frequency within Shakespeare's lexicon (Ohge et al. 2018). These two results can help draw conjectures about the reading practices, and perhaps the creative processes of Melville. New insight can similarly be drawn from the analysis of production notes or casting calls. It is perhaps in those texts

in the sidelines that the greatest promise for digital text analysis can be found for theater research. We can also apply digital textual methods to study performance scholarship itself, as Rachel Fensham (2019) has done for dance scholarship. As the most mature area of DH, more digital platforms, step-by-step books and easy to use software packages exist for the computational analysis of words than for any of the other area surveyed in this book. Our work is cut out for us.

Excursion: *The Flying Inkpot*

The Flying Inkpot was a volunteer-run website which published theater reviews from 1996 to 2015 in Singapore. The bulk of reviews dealt with theater and dance, but the project also included some reviews of poetry and other art forms in the early years. A branch project dedicated to classical music still exists (this is excluded from present consideration). All the reviews in the site were submitted on a voluntary basis by a core group of reviewers. The website itself was also maintained by volunteers, a remarkable feat for a project of this extent and longevity. This project was particularly relevant in Singapore, where an active theater review culture didn't previously exist within newspapers or other periodicals (but currently some newspapers have regular theater criticism columns). *The Flying Inkpot* was the brainchild of Matthew Lyon and Kenneth Kwok. Although the project is no longer active, an online archive maintained by Centre 42 provides access to the reviews for historical research (<https://inkpotreviews.com/>). The theater and dance section of this archive includes 1,154 reviews, written by 65 reviewers. The reviews include the time and date of the performance reviewed, the reviewer's name, a rating (in a scale of zero to five stars), and a commentary which is typically a few paragraphs long. The archive also includes shorter reviews, which are filed under "First Impressions" (these reviews were excluded from the current analysis). The shows reviewed include a mixture of local and touring productions. The reviews are extensive but, as they relied on the time and interest of reviewers, they don't constitute a comprehensive survey of Singapore-based theater. That being said, the 1,064,854 words written by this small army of reviewers are in themselves an important cultural object, and one that can be investigated by historians of Singapore-based theater in a number of ways. For the analysis below, I excluded the reviews from the first two years, since very few were written at that point in time.

Is there any identifiable change in the vocabulary usage of *The Flying*

Inkpot reviewers over time? To answer this question, I first uploaded the corpus to Voyant and selected the hundred most common words in the corpus, after removing articles (I used the default stopword list from Voyant). I then downloaded the trend data and further analyzed it in Python (see appendix A for details and visit the publisher’s website to download the data and code). I calculated the Mann-Kendall s statistic (Mann 1945; Kendall 1975; Gilbert 1987) on the trend data downloaded from Voyant. The Mann-Kendall s provides a statistical estimation on whether a trend is monotonic (i.e., consistently increasing or decreasing over time). The sign indicates the direction of the trend, an s of 103 and -103 are equally strong, but the former is increasing and the latter decreasing. The Mann-Kendall s has been previously used to detect changes in vocabulary trends on Twitter (Malakar et al. 2018). I calculated the Mann-Kendall s for the top 100 words, and then selected only the strongest trends, provided the p -value was lower than 0.05. The top five words with the clearest trends are shown in table 4.1 (upward) and 4.2 (downward). They are also graphed in figures 4.1 and 4.2, respectively.

I then downloaded the concordances for these words; that is, all the sentences where each of these words were used and annotated them by hand, pushing the analysis into the data-assisted realm. The annotation by hand aimed to be systematic, but as a situated and provisional process, it offers a different perspective on the monotonic trends and closely examines assumptions that might be hidden in the statistical analysis. Close

Table 4.1. Top five words with the strongest markers of an upward trend

Word	Mann-Kendall s	Counts in the corpus	p -value
feel	103	549	0.000112
makes	101	460	0.000152
great	87	510	0.001124
light	83	461	0.001897
work	81	1,470	0.002444

Table 4.2. Top five words with the strongest markers of a downward trend

Word	Mann-Kendall s	Counts in the corpus	p -value
woman	-87	595	0.001124
script	-75	845	0.005064
real	-71	510	0.008015
night	-61	520	0.023047
audience	-57	2,463	0.033909




Fig. 4.1. Top five words with the strongest markers of an upward trend.




Fig. 4.2. Top five words with the strongest markers of a downward trend.

annotation of each datapoint incorporates thick context and shifts attention from the aggregate to the singular sentence. This process revealed that, in the upward trend, “work” is most often used as a noun. A theater performance was increasingly referred to as a work, as opposed to a “show” or a “performance” (these other words were slightly more common in the early years of *The Flying Inkpot*, see figure 4.3). But in 2015, the last year, “performance” became again slightly more common than “work.” The trends of these other words were not as strong and were not captured by the statistical analysis reported above.

When manually disambiguating the multiple senses of “light” (“illumination,” “simple,” “funny”), its trend disappears. It is interesting to note that the verb “feel” (which also includes the conjugated forms “feels” and “felt,” but not derived nouns such as “feelings”) became a more common way for the reviewers to express their views. When it comes to the decreasing trends, “woman,” “script,” “real,” and “night” are slightly less common but they reveal a fascinating trend. One explanation is that preoccupation with scripted, nighttime theater that emulates reality dwindled over time. An alternative explanation is that, as the group of reviewers became more established and diverse, they started looking at different types of theater, and describing it with different words. In any case, identifying the reasons for this trend warrants additional research. I also found the decreased usage of “woman” surprising, since I previously had the impression that feminist readings of performances became more common over time. I suspect that, as feminist concerns became more common, a more nuanced vocabulary for the description of gendered experiences was developed and there was less need to directly use the word “woman.” However, I have not been able to prove this in terms of the data.

In the close reading of the downward trends, the most surprising insight came from reviewing instances of the word “audience” (which was overall more common than the other four words with strong downward trends). As I read each of the sentences where this word was used, I noticed that most of the times this did not indicate a description of the audience (i.e., “the audience laughed”) but rather rhetoric usages (“provoking the audience”), which constitute, in my opinion, an indirect way to phrase the reviewer’s own perspectives. I closely read all concordances for the word “audience,” and for each sentence indicated whether the usage was descriptive or rhetoric. This is a highly interpretive but systematic form of reading, that takes every sentence that uses the word “audience” as an individual case, and then groups them into categories. Other




Fig. 4.3. “Works” compared to “performance” and “show.”




Fig. 4.4. Changes in the relative frequency of all mentions of “audience,” and trends in purely rhetoric usages of “audience.”

scholars might disagree with my categories, or with the ways I classified individual sentences. For this reason, I am framing this part of my analysis as a data-assisted strategy, and highlighting the situated aspects of my interpretive process. The classification of sentences into descriptive and rhetoric usages is not necessarily an area of low ambiguity (see the criteria in the introduction). Therefore, interested readers are invited to download my data from the publisher's website and trace the interpretive move behind every data point. They can also reclassify my sentences according to other criteria and offer alternative interpretations of this trend.

Given my assumptions, the decrease in the usage of "audience" suggests that as the group of critics that coalesced around *The Flying Inkpot* refined their vocabulary and approach, references to the audience became less important as rhetorical devices. Perhaps this indicates that critics became more confident of their own voices, encouraged by the prominent place that *The Flying Inkpot* quickly earned among Singapore's theater circles. But proving this hypothesis would require additional research as one would need to identify other trends in the data that would confirm or disprove this hypothesis. For example, we could find other rhetorical devices that also indicate more confident authorial stances, or manually annotate a sample of reviews and assign them an "assertiveness" score. It is also important to note that the trend in this word usage does not constitute an argument about the importance of the audience for Singapore-based theater, but about how critics collectively chose to describe their impressions.

It could be argued that these changes in trends are not particularly special. One could imagine that, as the number of reviews increased, many terms became less frequent or more frequent. However, looking at general trends reveals that the opposite is true. The frequencies of many words remained stable across time, as critics came and went, and as the number of reviews waxed and waned. Other words have no discernible patterns. This is true for terms that have been the focus of intense academic attention in relation to Singaporean theater ("text," "stage," "politics," "state," "censorship," "queer," "postcolonial," "cultural," "Chinese," etc.). Close reading would be a better method to study how these concepts were used, how their usage responded to specific artistic events, and how they differ from the ways these terms are used in longer, academic pieces of writing. I want to emphasize that my results can't be taken as representative of Singapore-based theater as a whole. But they offer a new glimpse into the changes in the vocabulary of an influential group of critics. Given the importance that *The Flying Inkpot* had for practitioners over its almost two

decades of existence, these conclusions still reveal patterns that would be easy to miss if one just read all (or a portion) of the reviews.

The approach I demonstrate in this excursion shares some similarities with the culturomics project inaugurated by Michel et al. (2011), but that study used an entire corpus of digitized texts. It is perhaps closer then to the approach of Rockwell and Sinclair, who analyzed changes in the frequencies of words in the archives of *Humanist* and in the academic journal *Game Studies*. Although the identification of the trends in this excursion is driven by data, the results are then examined closely, and the conclusions are assisted, rather than purely driven, by data. This book advocates for the limited applicability of computational methodologies for studying theater performances, and related phenomena. The modest contribution of this excursion is showing that references to the “audience” in the reviews of *The Flying Inkpot* steadily decreased, and performances tended to be increasingly described as “works.” The scope of these observations is limited, but they can serve as a stepping stone for more comprehensive analysis of a key resource for the history of Singapore-based theater at the turn of the twenty-first century.

Code and data: Sample code for this chapter is available at <https://doi.org/10.3998/mpub.11667458.cmp.39>.
The data used can be downloaded from <https://doi.org/10.3998/mpub.11667458.cmp.26>, <https://doi.org/10.3998/mpub.11667458.cmp.27>, and <https://doi.org/10.3998/mpub.11667458.cmp.28>.

Relationships as Data

Theater depends on relationships. A theatrical production is impossible without the collaboration and co-presence of different people. By most definitions of theater, even the most minimal productions require at least a performer and one spectator. “As scholars of a collaborative art form, we are always dealing with relational data, for theater artists almost never work in isolation,” writes Caplan (2017). Even ritual puppet performances in Bali that don’t require a human audience are created by a collective of artists. In several dramatic traditions, characters and the relationships between them are also important. Relationships in these fictional and collaborative spaces can be modeled and analyzed as networks.

To see how a wide range of theater scripts and modes of collaboration can be represented and analyzed as networks, a short overview of network theory is needed. Following this overview, I will explain how network analysis can be used for data-driven and data-assisted theater research. This is complemented by an excursion into two of my own research projects: an analysis of collaborations in the Singaporean theater company The Necessary Stage (TNS) and an analysis of character co-presence in the fictional universe of traditional *wayang kulit* (shadow-puppet theater) in Indonesia. This chapter aims to show that the toolkit of network analysis can reveal hidden structures in relationships, which has great potential for theater research. Data-driven network analysis aims to show counterintuitive features of dramatic literature and artistic collaborations, while data-assisted network analysis enables a playful defamiliarization of theater history.

Methods for Network Analysis

In theater history, Jacob Levy Moreno is more commonly remembered as the father of sociodrama, but he was also the inventor of sociograms:

visual representations of connections among people. He was one of the first researchers to realize that networks could depict social relationships (Moreno 1960). However, the study of networks is older, as the first sketches of a theory of networks were proposed by Leonhard Euler in 1735, and they were inspired by the *Seven Bridges of Königsberg* in Saint Petersburg. This and many other key moments in the history of network theory are masterfully retold by physicist Albert-László Barabási in *Linked* (2002), a nontechnical introduction to the study of networks. In his book, Barabási (who is himself part of the history he tells), describes many ways in which networks can be analyzed, as well as several practical applications of the study of networks. His book blends history and context with mathematical theory, and it begins, like many an introductory network book, by describing the basic components of a network. Networks are made of two basic elements: nodes (things that are linked) and edges (the links between them). In the study of literary fiction, characters are often represented as nodes. But the edges can take a variety of meanings: they can represent blood ties between characters, words exchanged between them, or co-presence in a scene. Algee-Hewitt (2017) notes that, for the analysis of drama, the unit of the network “is not the node, but the edge: it measures not characters, but interactions” (752).

Many of the concepts in network analysis come from the sociological study of networks (often termed social network analysis or SNA). Quantitative analyses in this field often report measurements such as *degree* and *betweenness* (Knoke and Yang 2008). *Degree* indicates the extent to which a node is connected to other nodes. The higher the degree, the more connections a node has. For example, Haresh Sharma (the main playwright of TNS), is the node with the highest degree. A slightly different measurement is *betweenness*, which indicates the fraction of all shortest paths that pass through a given node. This gives a sense of the importance of a node to the network structure. For example, the character Karna in the wayang network has a very high betweenness (much higher than his degree). This means that although this character is not as well connected as others, he is often found in between many characters. Besides offering useful concepts such as this, SNA has also developed tools that are potentially useful for the study of theatrical networks. For example, there is abundant SNA literature on the construction of networks through interviews or surveys. These are called “egocentric networks” and they stand in opposition to “whole-networks.” Methods of SNA often estimate general properties of networks from just the subsection represented by egocentric

ones (Marsden 2005, 8–29). In computational social science, the purpose is estimating the range of variations in a network and determining how certain parameters such as gender or education level affect the formation of networks. A sociological study of theater networks could use this approach, but I am unaware of any study that uses such attempt.

In contrast, the study of networks in physics often focuses on general properties of the network structure, such as the distribution of degrees, the network diameter (the longest path between two nodes in the network), and its clustering coefficient (the average connection between two nodes). Many complex networks (from scientific collaborations to protein interactions) have three properties: their clustering coefficient is very high, the diameter is small, and their degree distributions are exponential (a few nodes have most of the connections, and most nodes have very few connections). The first two properties mean the network represents a “small world,” where even the least connected nodes are only a few nodes away from each other. The second property is often called “scale-free.” This lack of scale means that all values are expected since the distribution of the connections is very uneven (normal distributions, in contrast, cluster around an arithmetic mean). These properties have been identified in many networks and they are often reported as intellectual curiosities. But these properties can also be used to analyze the robustness of a network (say a phone network) in case of accidental failure or deliberate attacks. This has many practical applications, such as helping engineers devise resilient communication networks or help epidemiologists assess the risk of contagion within populations. Later, I will describe a project that seeks to understand the meaning of small worlds in theater networks.

Of all the areas covered in part 2 of this book, network analysis is the one with the most standard set of measurements. Networks, as mathematical models, have well-established quantitative properties that are straightforward to calculate, as seen above. This is not true to the same extent for texts, motion, images, and locations. As Algee-Hewitt (2017) notes, networks are “intuitive to grasp and yet mathematically complex” (752). Data-driven network analysis is well established in the sciences and social sciences. But what might one glean from network overviews in the humanities? In DH, networks have been applied to study relationships between artists (collaborations) and interactions between fictional characters. Both of these are directly applicable to the study of theater.

Network Analysis of Fictional Relationships and Artistic Production

Network analysis is an area where scientists have often applied quantitative measurements to datasets from the humanities. Researchers have studied, for example, the fictional networks of Greek mythology (Choi and Kim 2007) and Marvel superheroes (Alberich, Miro-Julia, and Roselló 2002). Some studies have also used fictional networks as test cases for information segmentation and retrieval (S. B. Park, Oh, and Jo 2012; G. M. Park et al. 2013). Some researchers have turned their attention to collaborative networks in the arts, and a network analysis of jazz musicians in the early twentieth century found clear evidence of racial segregation (Gleiser and Danon 2003). These papers all present fascinating insights from a network science perspective, but they don't try to contextualize their findings within humanities scholarship. Their priorities clearly lie elsewhere, and the lack of critical engagement is not a shortcoming in their fields. This is different in the examples that follow, which analyze fictional networks in drama, and collaborative networks in theater productions against more carefully considered historical and disciplinary contexts.

Moretti (2011) focused only on specific plays and his networks are therefore much smaller. He also built his networks by hand—the character nodes were not downloaded from a database or automatically extracted from a corpus. He reported some network measurements but did not dwell on their geometric and topological features. Still, his close analysis of networks reveals counterintuitive features of well-known texts. For example, his Hamlet network identifies a “region of death,” where all characters linked to both Claudius and Hamlet are killed, except for Osric and Horatio: “outside that region, no one dies in Hamlet [. . .] the tragedy is all there” (217). Moretti also found two separate components, sub-regions where each node is connected to every other node, which he interprets as the separate worlds of the court and the state (Moretti 2013, 228–30).

Algee-Hewitt (2017) takes this line of inquiry further, by analyzing network properties of 3,568 English dramatic texts written between 1550 and 1900 retrieved from the ProQuest Literature Online drama corpus. He developed his own code to identify the most probable recipient for each speech using a rule-based system. Speech was represented as directed edges that connect speakers to addressees. He then computed several

properties for the nodes. He presents a convincing case for taking eigenvector centrality (EC) and betweenness centrality (BC) as the most interesting measurements. EC is a measurement that computes the importance of a node by calculating its total number of connections, considering the relative importance of the nodes to which it is connected. As noted above, BC looks at characters that mediate between factions. Algee-Hewitt convincingly demonstrates that BC can be used to estimate protagonism. Rather than merely identifying a protagonist, the distribution of BC across characters in a play shows how protagonism is shared among characters, and how this changed over time, as plays were increasingly “less likely to feature a single central character (or cluster of characters) and more likely to depend on densely connected communities each featuring prominent characters in their own right: they are more likely to look like *A Midsummer Night’s Dream* than *Henry V*” (Algee-Hewitt 2017, 765). Similarly, BC shows what he calls “mediatedness.” Taking both measurements together, he identified plays where the protagonist is the mediator and those where mediatedness and protagonism are shared by different characters.

This line of analysis—tracking the history of drama through changes in network-theoretical measurements—has also been systematically explored by the Digital Literary Network Analysis (DLINA) group, led by Frank Fischer and Peer Trilcke. Over the span of several years, they have used different facets of network analysis to study a large corpus of German drama. First, they measured how some basic measurements, such as network density, changed over time (Trilcke et al. 2015). This made them realize that there are consistent markers of genre over the two centuries of data which they processed. The average density (the number of actual edges divided by the total number of possible edges) for tragedy remained constant, as was the case for comedy and libretto (a category where the authors include all musical theater).

In another project, these researchers wondered whether theater networks were small worlds (Trilcke et al. 2016). As seen earlier in this chapter, small worlds are common in nature and in social environments. The DLINA group looked for three markers of small worlds in their corpus: a high clustering coefficient, small average path lengths, and scale-free degree distributions. They found that only five plays (in a corpus of almost five-hundred plays) had small world properties, and this led them to pose questions of interest to the study of drama: what does it mean to have (or not have) small world properties? Although this is still an open question, this line of analysis sets this kind of projects apart from the ones

developed entirely by engineering and scientific teams. For engineers and physicists, it is merely interesting to note the presence or absence of certain quantitative markers in cultural data. But for theater researchers, the important questions are what these quantitative features mean for the history of our field.

On a subsequent project, the DLINA group focused on network dynamics (Fischer et al. 2017). Most of the other network approaches mentioned so far “spatialize” the flow of the drama into a single diagram, and thus remove the dimension of time. But one could also track the evolution of networks as they change. For this purpose, the DLINA group proposed a series of time-oriented measurements: *all-in index*, *central-character-entry index*, *final-scene-size measure*, and the *drama-change rate*. These measurements help identify outliers in drama history, such as extreme changes (“provocation”) or extreme uniformity (“boredom”). In the development of this project, they were inspired by IntNetViz (Xanthos et al. 2016), an interactive tool for visualizing drama as it unfolds.

Yet another of the DLINA group’s projects combined several quantitative measurements to automatically identify the protagonist in a drama. They looked at both network measurements (degree, closeness, betweenness, weighted degree, eigenvector) and word counts (words spoken, speech acts, frequency). They found that multidimensionality (i.e., combining different measurements) is the best solution. This project is reminiscent of work by another group, which tried to identify the most likely romantic couple in a dramatic text (Karsdorp et al. 2015) using a corpus of French drama. Also relying on a combination of network and other measures, they achieved very high accuracy, and were able to identify romantic couples in 80 percent of the cases. But, as they note, their algorithm is premised on the assumption that there is one romantic couple in each text. A harder problem would be identifying whether a romantic couple is to be found in the text (or how many such couples there are).

The DLINA group has announced plans to extend their work to English, French, and Russian dramatic texts and they have developed an impressive portal with detailed analysis, network visualizations, and an application programming interface (API), to automatically retrieve data queries (Trilcke and Fischer n.d.). The tools they propose are excellently matched to text-based drama, especially when such dramas are encoded in TEI formats that are easy to consume and reuse (see chapter 4). DLINA is leading the way by combining literary history and criticism with quantitative analysis. If more dramatic texts from around the world were readily

available for network analysis, this would open the door for comparative analysis across theatrical cultures.

So far, the projects surveyed have focused on dramatic texts. But another tantalizing possibility for theater history is the analysis of collaborations among theater artists. In an influential paper, Uzzi and Spiro (2005) analyzed the small world network of artists involved in Broadway musicals from 1945 to 1981. They found that the small world properties of the artists' networks affected their financial and artistic success. To estimate the latter, they used an index created by Suskin (1990), who manually assigned a numerical value to 315 productions.

AusStage (AusStage 2013) has amassed an impressive dataset on theater performances in Australia, and theater performances by Australian performers around the world. This type of dataset paves the way for new theater histories. As Caplan (2017) notes "data-driven theater history, at its best, can reveal previously invisible patterns about relationships among diverse groups of artists working across languages and cultures" (557). Caplan's own data consists of 290 Yiddish theater artists who worked on at least one Vilna Troupe production between 1915 and 1936. Using this data, she was able to identify ten previously distinct branches of the troupe, a fascinating and previously unreported finding. She suggests that this data can also be used to trace how gender influenced the formation of friendships and professional networks, or to estimate the influence of familiar ties—and of pedagogical figures—on collaborations.

Caplan wonders what collectives of scholars might be able to accomplish if they worked together to assemble large datasets on theater collaborations. In part, that potential has already been demonstrated by AusStage, and by projects that build on their data model, such as IbsenStage, which is a collectively assembled dataset of performances of Ibsen texts around the world. This comprehensive dataset, where the completeness is estimated at 60 percent, was the backbone of *A Global Doll's House* (Holledge et al. 2016). Using network visualizations together with other analytical techniques, Holledge and her coauthors showed that the global success of the play was linked to genealogies of transmission. One such visualization shows how a long chain connects productions from the late nineteenth century until the 1990s. In this directed network, edges indicate that at least one actor from a given production was involved in another more recent performance (more formally, this technique is called a unipartite projection). Thus, rather than showing co-presence, as other collaborative networks do, this visualization shows movement of actors across different

temporal layers of a production's history. Their work effectively demonstrates that it is not only ideas, but specific people, that ushered *A Doll's House* into global fame. Their network analysis also shows the decisive influence of women in the production and promotion of the play—rather than merely in the portrayal of the protagonist—another unexpected and important finding. Bardiot (2018) also used network analysis to study the collaborative networks of Merce Cunningham to show that after 1954, his working strategy brought a larger number of artists in collaboration with each other, shifting from a “star” to a “spiral” pattern of collaboration.

There are many other possibilities for network analysis in theater, and the more datasets are available for reuse, the more likely we will see network analysis in all areas of theater research. For example, one study modeled spectators' choices to attend a theater show in Belgium as networks, and tried to find the probability of people sticking to specific venues, estimating how much knowing another spectator influenced such decisions (Agneessens, Roose, and Waeye 2004). As theater researchers continue to apply network measurements in their work, they can draw inspiration from the work of Maximillian Schich, who has systematically used network analysis to study a wide range of topics central to art history. His work weds sophisticated quantitative analysis with thorough historical analysis (Schich, Lehmann, and Park 2008; Schich et al. 2014, 2017).

Data-Driven and Data-Assisted Network Analysis

The approaches reviewed so far vary in the level of mathematical formality used in the reporting of findings. However, most of them could be described as data-driven projects. This is often made explicit. Caplan (2017), for instance says that network analysis and other modes of data visualization “can offer an important corrective to our understanding of what is central and what is peripheral in theater history” (557). She thus urges theater scholars to engage in verifiable, consensus-driven studies of theater history premised on data. Holledge et al. categorically situate their work within a scientific paradigm. As described above, their interest is to study the global history of *A Doll's House*. They used network visualizations to identify the forces that led to the spread of the play across geographies and generations. The authors compare the analysis of the play's history to the scientific study of evolution. Using the zero-force evolutionary law proposed by philosopher Robert Brandon and paleobiologist Daniel McShea, they posit that cultural phenomena—such as *A Doll's House*—tend towards

“increasing diversity and complexity” until they find constraints. Natural selection is a constraint in biological systems, and “underlying political, social, economic, aesthetic, or technological forces” are constraints that shape the spread of theater plays (Hollidge et al. 2016, 18–19).

Some projects also use data-assisted network analysis to enable more situated and context-rich perspectives. Interactive visualizations, like Caplan’s digital portal for the Vilna Troupe network (<http://vilnatroupe.com/>) and Xanthos et al.’s IntNetViz (<https://github.com/maladesimaginaires/intnetviz>) encourage users to zoom into portions of the data and construct their own close analysis of texts and histories, in ways that depart from what is purely contained in the data. The interactive features enable multiple perspectives, different scales, and thick context (see chapter 3). When it comes to aesthetic provocations, the best example is *Brecht Beats Shakespeare! A Card-Game Introduction to the Network Analysis of European Drama* (Hechtel et al. 2018). This is a card game set where each card has a network visualization and measurements that correspond to different theater plays. The players try to win an opponent’s card with a higher value—but they must agree which network value to use for the game. The game relies on networks to produce a playful defamiliarization of theater history. It is worth noting that most of these data-assisted perspectives are based on data-driven projects. Caplan’s visualization is tied to a replicable analysis of the Vilna Troupe’s history. *Brecht Beats Shakespeare!* uses data from Fisher et al.’s data on European drama. But inspiration works on both directions: data-driven work can inform data-assisted projects and vice versa. Fisher et al.’s work, as noted earlier, also derived inspiration from the data-assisted IntNetViz for their analysis of network dynamics over time.

First Excursion: Collaboration Dynamics of The Necessary Stage

The first excursion in this chapter shows how network analysis can be used to study collaborative relationships among people involved in theater using data from the Singaporean theater company The Necessary Stage (TNS). I obtained the data as a spreadsheet (kindly provided by the theater company), which listed all the people involved in all their productions between 1987 to 2015. My research assistant Alyssa Chandra and I then went over all the records to standardize the names of the performers. Founded in 1987, TNS is one of the most active theater companies in Singapore with over four hundred productions to date. Through their outreach programs and sub-companies like Theater for Seniors (TFS) they have

built strong ties to different communities. The most prominent members are playwright Hareesh Sharma and director Alvin Tan, but the company has also served as an incubator for many other companies. Prominent artists such as Kok Heng Leun (Dramabox) and Chong Tze Chien (The Finger Players), have at some point been involved with TNS, and the company has actively sought to promote new playwrights and directors. This double focus—social involvement and talent incubation—gives the TNS network its unique structure. However, it is important to bear in mind that the company’s structure has changed over time and so has the meaning of “collaboration.” I will describe these two changes before analyzing some quantitative features of their collaborative network.

The people involved in the company have changed drastically over time, and only Tan has remained since the beginning, but the company has retained a distinctive mission. Artists and commentators describe the individual performances as part of a longer creative process: “whether you like their work or not, its significance outlives every single production they stage” (Birch 2004, 58). At the time of writing, TNS commissions an average of two performances per year, organizes an annual festival, have their own theater venue, and employ a full-time staff of around ten people. However, the company started out as a student collective and only eventually grew into the stable, professional institution it is today. Although they now are one of the most respected theater companies in Singapore, their political agenda and choice of artistic forms have created trouble for them in the past (K. P. Tan 2013).

According to Alvin Tan (2004, 253), in the beginning the artistic committee had to do administrative and production work to fight the “stardom syndrome” but groups were divided and some people left. This observation will be specifically important for the network analysis that follows. In 1992, only four people were working full time. One of the most significant structural changes took place in the year 2000, when they moved from their offices in the city center to the Marine Parade neighborhood in the eastern part of the city-state, a change that propelled a different kind of community engagement. In 2002, TNS established a lab to encourage more collaborations within group members. The people who have worked with TNS—the nodes in their impressive network—include famous actors, producers, academics in different fields, and members of parliament. These collaborations have taken many shapes, but they have retained a strong social and political objective. Alvin Tan (2004) asserts that “for TNS, collaboration is a methodology of resistance, resisting the

rationalized mindsets ingrained in cultures of both contemporary urban lifestyles and the production structure of the traditional Western model theater” (266). Artists affiliated with TNS often reflect on the meaning of collaborations, and the term appears extensively in the edited volumes they commissioned for their tenth and seventeenth anniversaries (Tan and Ng 2004; Krishan 1997).

Haresh Sharma is credited as the playwright of most TNS plays, but his writing is often undertaken in collaboration with the performers. According to Tan, the early phase of collaboration at TNS was mostly textual, where company members developed a working method based on improvisation, writing, revision, and rehearsal as distinct creative stages. Tan admits this way of working was still heavily text-centric and, after the first ten years of the company, he wanted to find more adventurous collaborative practices. A wider range of collaborative strategies can be observed in several productions. *Off Centre* (1993), a play about mental illness, required extensive research, interviews, improvisation, and writing. The research process was even more intensive in productions where the actual performers had personal experiences related to the subject of the plays. For example, *'Scuse Me While I Kiss the Sky* (1994), was based on the experiences of, and performed by, people who had attempted suicide. Similarly, *October* (1996), was also devised with, and performed by a group of elderly people. According to David Birch (2004, 61), their most collaborative work is *Completely With/Out Character* (1999), the result of working with the late Paddy Chew, the first Singaporean to publicly announce his HIV positive status. According to Wong (1997), these plays demonstrate that the actors “have not only become the subject of both play and interview, they have become more pro-active collaborators in determining the shape of the play as their experiences, improvisation work and interaction with one another are as much part of the process of play-making as they are events enacted for the stage” (195). This background should be considered when interpreting the network analysis below, where people are modeled as nodes and collaborations are modeled as undirected edges between every person who collaborated in a single play.

Before delving into the analysis of the TNS data, let's consider the first play of a fictional theater troupe by way of example. Let's say four people collaborate in this play, and there is one edge between each of them. There are four nodes and six undirected edges. Thus far, the network can be said to have a single *connected component*. If, for a second production, two of the original people and two new members are involved, then the shape of the

network will change. The number of nodes will increase from four to six and the edges will increase to eleven. But there will still be one single connected component (as the two people involved in both productions are linked to the other four who have not collaborated with each other). Now let's imagine a third production, also produced under the banner of the same troupe but with three entirely different people. Now the node number will have increased to nine, and the edge number to twelve. But we will observe two connected components, as there is no link between the makers of the first two productions and the group responsible for the third.

Figure 5.1 shows the staggering increase in the number of connected components within the collaborations of TNS. For this analysis, I omitted the names of the directors and playwrights (who are stable across many productions and would thus link the network into a giant connected component). This visualization shows that TNS was so consistent in its recruitment of new people that at some point twelve connected components are discernible. But notice also that the line goes up and down. The company often involves entirely new people, but then it draws from these distinct groups for new productions, bringing the number down. This ebb and flow illustrates the principles of collaborations seen in the history above: the radical pursuit of new collaboration strategies and the constant desire for social engagement by enlisting non-actors as performers. It is interesting to note that the last recorded number of components (11) remained stable since 2002 (when, as seen above, TNS established a lab to encourage more collaborations within group members).

We can further drill down into these results and separate the networks of people listed in the credits as cast from those listed as part of the production crew. Figure 5.2 shows that the number of components for cast members at one point soared to 25. In contrast, figure 5.3 shows the more closely knit world of the production crews, whose connected components have never exceeded 5. This could indicate there are fewer people working in production roles. But, in fact, as figure 5.4 indicates, there are—and have always been—more people working in the production side. The number of people (nodes) in this network increased at a far greater pace. The explanation is that the nature of the collaborations in the production side are different—figure 5.5 shows that edges grow at a greater pace in the production network.

These data-driven results add details to the history of the company, and help explain their position in the Singaporean theater landscape. However, this bird's eye overview should be read against a more situated




Fig. 5.1. The number of connected components over time.




Fig. 5.2. The number of connected components over time for cast members.




Fig. 5.3. The number of connected components over time for production crew members.




Fig. 5.4. The number of nodes over time separated by cast and production crew.




Fig. 5.5. The number of edges over time separated by cast and production crew.

reading of TNS's history, which shows that their understanding of collaborative work has changed drastically over time.

Second Excursion: Networks of Wayang Narratives

This excursion considers the relationships between characters in a fictional universe. There are different kinds of wayang theater in Indonesia and the most notable one is wayang kulit, where shadow puppets made of water buffalo hide are controlled by a single puppeteer in all-night performances. Other forms include *wayang wong* (where human performers imitate the leather puppets) and *wayang golek* (three-dimensional puppets made of wood). Most forms derive their narrative materials from the same sources, which include the epics *Mahabharata* and *Ramayana*. The former is often preferred for all-night puppet shows, where stories of feuding families, conspiracy, and betrayal are explored through extensive dialogue and philosophical explication. The *Ramayana* is more commonly used in dance, as its straightforward plot lends itself nicely to voiceless, visual drama (see chapter 6). Wayang theater dates at least to the ninth century CE, and the stories can be traced much farther back in South Asia (Escobar Varela 2017). In what follows, I focus on the Javanese versions of wayang kulit, and I refer to characters and stories by their Javanese spellings

(which might sound slightly odd to those more familiar with the South Asian versions of the stories). Following conventions in English-language scholarship on wayang I add an “s” to Javanese nouns to indicate plural forms, so I refer to *dhalangs* as the plural of *dhalang* (this, however, would not be correct in Javanese, which uses different grammatical conventions to indicate plural nominal forms).

The *Mahabharata* narrates the story of two groups of cousins, the Pandawas and the Korawas, who fight for the throne of Astina. The Pandawas are the rightful heirs, but the characters of both sides are often depicted in nuanced moral tones, offering the *dhalangs* ample opportunity for moral reflection. In Java, the *Mahabharata* is never performed in full and an entire all-night performance concentrates on a single episode. For example, a favorite *lakon* (story) is *Dewa Ruci*, the spiritual quest of Bima, the second Pandawa brother. This episode happens before the Great War (*Bharatayudha*) but does not fit into a specific chronology of the epic. Like most of the famous episodes, this is a Javanese addition to the South Asian versions. In fact, many key features of the *Mahabharata* as performed in Java are exclusively Javanese inventions. Another example of these additions is Semar, a hermaphrodite clown-servant of divine origin. Semar is another name for the Javanese god Ismaya, a remnant from an indigenous Javanese religion that predates the arrival of both Hinduism and Islam to Java. In the genealogy of wayang characters, Semar is said to be one of the brothers of Bhatara Guru (Shiva) and he appears in most stories as a clown-servant, offering advice to the protagonists of the story. Regardless of the episode, Semar and his sons Petruk, Bagong, and Gareng appear in the *gara-gara* scene, a comical interlude that takes place in the middle of all-night performances (a segment of the performances which itself often lasts a couple of hours). These characters are collectively referred to as *punokawan*, or clown-servants.

Working together with physicist Andrew Schauf, I built a network model of a set wayang stories (Schauf and Escobar Varela 2018). This model considered the numbers of shared scenes in which pairs of characters appear, giving weight values to each edge. Thus, an edge of weight twenty between two characters indicates that the corpus contains twenty different *adegan* (scenes) in which both characters appear. The resulting network can be explored through a data-assisted interactive visualization at the Digital Wayang Encyclopedia’s website at <https://villaorlado.github.io/wayangnetworks/html/canonical> (also available as video 5.1 in this book’s companion website).

In the interactive visualization, the size of a node corresponds to its degree (larger nodes have a higher number of edges). An edge (represented as a line) between two characters indicates that both appear together in at least one scene. The thickness of the edges represents weight of the degree. By clicking on a node, a user can access several network measurements for the corresponding character. These measurements are explained below:

- *Degree*: The number of nodes linked to the given node
- *Weighted degree*: The sum of the weights of all the node's links
- *Closeness centrality*: The inverse of the average length of the most direct paths between the given node and all other nodes in the network
- *Betweenness centrality*: As seen above, this is a measure of how often a node acts as a bridge between other nodes. A high betweenness centrality value indicates that the shortest paths between all pairs of nodes in the network often pass through the given node.
- *Eigenvector centrality*: As mentioned earlier, this is a measurement of the influence of the node in the network that considers the degrees of a node's neighbors. Nodes with high eigenvector centrality tend to be connected with neighbors who are themselves highly connected.

The online platform can be used to interactively explore many facets of the wayang characters and stories, as it enables situated, data-assisted interpretations. For example, a user can click on a given character's description and then see which other characters tend to appear in the same scene as this character. A user can also trace the different versions of this character's stories in South Asia and Java. The extensive notes on each character point add thick context to each data point. They also enable users to shift perspectives from aggregate categories to individual characters. The *exploratory features* (see chapter 3) of the visualization enable users to choose where to focus their attention, what to view and which sequence to follow.

This dataset can also be used to reach data-driven conclusions. We can combine the network-theoretical measurements with features that have long preoccupied scholars, such as the ways Indian-derived characters and characters of local Javanese invention interact in the stories as told today. When we considered the origin of the characters (India vs. Java), we

noted that almost half of the characters who appear in the epic are Javanese in origin. This seems counterintuitive; when I informally asked people familiar with wayang kulit to give an estimate of the percentage of wayang characters with Javanese origin, most (including several well-versed wayang experts) guessed that only about 20 to 30 percent of the characters were Javanese in origin. However, if we look at the weighted degrees of the characters, we note that the Indian characters tend to have significantly higher values than their Javanese counterparts. This means that although almost half of the characters are Javanese in origin, the Indian characters tend to appear more often—and tend to appear repeatedly alongside the same fellow characters—thus exhibiting stronger connections within the co-occurrence network. This greater prevalence, reflected in their weighted degrees, may account for the common perception that Javanese characters are fewer in number, even though Javanese and Indian characters appear in the same proportion. In some cases, scholars might disagree on whether a character is Javanese in origin, or a modified version of an Indian character. The Digital Wayang Encyclopedia includes a brief discussion on how these decisions were made, and the description of each character points to a variety of academic sources. However, readers might want to reclassify certain characters and run the analysis again. For this purpose, the data can be downloaded from this book's companion website and manually reclassified. The choice of stories might also entail bias, as we used a relatively small sample of stories. My hope is that more wayang stories and characters will be analyzed in the future following the procedures outlined here, in ways that extend or challenge our results.

My comparison of Javanese and Indian characters is summarized in table 5.1. I calculated the effect size of weighted degree difference between Indian and Javanese characters. Effect sizes are rarely reported in network analysis papers in DH, but they are common in other fields (Clemente et al. 2015). To estimate effect size, I used Cohen's (1988) d and obtained a value of $d = 0.505$. This means that there is half a standard deviation of difference between both groups' weighted degrees. In Cohen's original formulation, any effect bigger than 0.1 is considered significant, and any value bigger than 0.5 is considered a medium effect. However, more recent commentators (Lakens 2013) note that Cohen's original rules of thumb, which categorized effects as small, medium, and large, should just be taken as a rough estimate. The importance of an effect size needs to be contextualized within the field to which it is applied, and compared to other effect sizes reported in the literature. As more theater researchers

report effect sizes, it will become easier to gauge how large of an effect is 0.5 for the study of cultural transmission in theater traditions.

Figure 5.6 visualizes the pairwise comparison between network-theoretical measures of Javanese and Indian characters. The measurements considered are weighted degree, betweenness centrality, and eigenvector centrality. The usefulness of considering all three measurements together is that they all show marked differences between Indian and Javanese characters. In this pairwise plot, each measurement is paired against the other two, through six scatterplots. This visual presentation follows the “small multiples” convention often used in statistical graphs (see chapter 3). The variables in the scatterplots are determined by the labels of the shared x and y axis. Thus, the scatterplot center-top compares the betweenness centrality (x axis) against the weighted degree (y axis). The quadrant where a measurement would be compared against itself (such as the top left) shows instead a kernel density estimate (KDE) of the distributions of values for both Indian and Javanese characters. The KDE shows how many values (y axis) are present at a given measurement range (x axis). This KDE estimate is similar to the violinplots described in chapter 6 (except that here the kernels are rotated 90 degrees).

Both the scatterplots and the KDEs tell a consistent story: Indian and Javanese characters are very different in terms of their network measurements. For every measurement, the Indian characters are spread around all possible values, whereas the Javanese characters show two distinct peaks: the vast majority of characters are bunched together around lower values. But a few characters rank highly for each measurement. These are the aforementioned clown-servants, who appear in virtually every story and have an important role to play as mediators, as they bridge distinct factions. The KDE for betweenness centrality shows this dramatically important role, with three gray spikes at the far right of the quadrant. These spikes indicate that three of the clown-servants have the highest betweenness values of the entire dataset (this is also evident in the scatterplot). The key role that the clown-servants play in wayang is not a new insight, but it is strengthened by this data-driven investigation.

Table 5.1. Comparison of Indian and Javanese characters

	Percentage of characters	Average Weighted Degree
Indian Characters	52.4%	153.47
Javanese Characters	47.6%	73.36




Fig. 5.6. A pairwise matrix comparing the network-theoretical measurements of Javanese and Indian characters.

Elsewhere (Schauf and Escobar Varela 2018), we have generated random networks of the same size as this wayang network to estimate the likelihood that the values we see for each character are significant. In another paper, I have also considered differences between the network measurements of characters with a stable puppet representation (i.e., those which cannot be substituted other puppet), and *serambahan* or “wildcard” characters that can be represented by many puppets (Escobar Varela 2019).

The question of the interaction between Javanese and Indian characters can be answered with data, but the pertinence of the question itself

is open to other modes of critical analysis. Many dhalangs and Javanese scholars are interested in this issue, as can be seen from wayang dictionaries published in Indonesia that almost always indicate when Javanese characters are also found in the Indian versions of the stories. However, this line of inquiry can also be traced to the concerns of colonial-era Dutch scholars. In a previously quoted paper, I have also given a fuller account of how colonial scholarship has shaped this question and how it has changed in postcolonial Indonesia (Escobar Varela 2019). This data experiment should not be used to draw definite conclusions about wayang. As noted above, the dataset I used is relatively limited and the addition of more data will further refine the insights that network analysis can bring in to the analysis of Javanese wayang's history. However, the method of using pairplots to visualize differences across types of nodes in a network, and the practice of reporting effect sizes, can certainly be applied to other types of fictional and collaborative networks.

Conclusions

The networks in both excursions are complex systems with the same structure: a power-law distribution and small-world properties. This means that the same mathematical tools can be applied in both cases, but they will be mobilized to very different ends. The TNS network is used to analyze the history of a theater company. The wayang network reveals patterns in the ways characters of different origin interact with each other. The two excursions have described quantitative markers against the backdrops of history and culture that shaped these networks and their interpretations. Both examples look at things that make sense in their contexts. The analysis of collaborative relationships has been central to academic reflections on TNS's history, as is reflected by the edited collections that celebrate and critique their work. The analysis of the relationships between Indian and Javanese versions of wayang stories has long preoccupied scholars of Java.

Networks can model many things. Here, they model properties that are directly relevant to theater scholarship in Singapore and Indonesia. Besides these features, many other aspects of theater practice could be modeled as networks (international collaborations, funding sources, festival circuits, citations in theater scholarship, etc.). If analysis of these other areas repeats previous experiences, we might find the markers of complexity in most, if not all these networks. But the most promising research agenda is to zoom in on the networks, analyzing the way cultural

context determines their structures, and the way their quantitative properties tell the story of those contexts. The data currently available, and the explanatory power of standard measurements, places network analysis as the most promising area for the digital study of theater companies and the fictional universes of theater plays.

Code and data: Sample code for this chapter is available at

<https://doi.org/10.3998/mpub.11667458.cmp.40>.

The data used can be downloaded from

<https://doi.org/10.3998/mpub.11667458.cmp.29>,

<https://doi.org/10.3998/mpub.11667458.cmp.30>,

<https://doi.org/10.3998/mpub.11667458.cmp.31>,

<https://doi.org/10.3998/mpub.11667458.cmp.32>, and

<https://doi.org/10.3998/mpub.11667458.cmp.33>.

Motion as Data

The study of motion as data is a fascinating area, and one that holds vast promises for our field. As a generalization, motion is more important than language for the study of theater-as-event (where I am including dance and physical theater). Although there are many genres and traditions where language is absent or not central, it is difficult to imagine a theater performance with no motion. The only example I can think of is *While We Were Holding It Together* (2006, dir. Ivana Muller), where a group of actors remained in the same position for the entire performance, as their lines were projected behind them. But even then, the subtle eye twitches and mouth motions were essential to the meaning of the performance.

In this chapter, I look at how to use digital data to study motion, rather than at dance performances that incorporate digital technology (see Bleeker 2016). In this chapter, I use *motion* and *movement* interchangeably for the sake of variation. But my interest is analyzing the movement of actors and objects on stage. Objects don't generally move by themselves (unless they are powered by electrical or mechanical devices), and in this respect they are different from actors. But the movements of both actors and objects are computationally tractable and can be modeled as data. The analysis of motion is the hardest, but in some ways, most interesting area for the computational study of performance. It is also the area that most underscores the situated, constructed nature of data. For things such as geographical coordinates, data-collection devices are so embedded in everyday life that they can sometimes become transparent and we can forget their constructed nature. This is not possible for motion, where the methods of data collection are so invasive—and so artificial—that the cultural specificity and constructedness of the data will remain front and center.

For this reason, the structure of this chapter is somewhat different from other chapters in this section, as I first consider the possible ways motion can be understood as data. Then I continue the usual structure

and survey a range of methods for collecting and transforming this data and explain how these methods can be used within data-driven and data-assisted methodologies.

What Is Motion Data?

Collecting and processing motion data poses the hardest problem of all the areas discussed in this book. There is no standard format for representing motion (unlike networks for relationships, written words for text, and coordinates for maps). There is also no straightforward way to obtain motion data. Whereas it is relatively simple to obtain GPS coordinates or find the list of people working together in a theater company, obtaining motion data is a whole different issue. Options range from written records of choreographies that use dance terms (e.g., *demi-plié*), to notation systems (such as Labanotation), and to sophisticated motion capture systems. At the time of writing, all available systems are either inaccurate, cumbersome to use, or expensive. Since motion data is scarce and it comes in a variety of formats, there are no agreed-upon, widely used methods to analyze motion data that would fit the needs of computational theater research. For example, quantitative motion analysis is an active subfield in engineering, but the level of granularity and types of questions asked (e.g., calculating trajectories of pedestrians for autonomous vehicles) are not easily mapped to the study of motion in theater and dance. However, new software such as Open Pose (Cao et al. 2018) might change this soon, as will be discussed later in this chapter.

Motion can be represented in three ways, ordered here in increasing level of formalization:

1. *Concepts* can belong to a general-purpose analytical vocabulary such as Laban Movement Analysis (LMA) or be specific to a dance tradition (*grande jeté* in ballet, *tristik* in Javanese dance).
2. *Notations* are formal vocabularies that include syntactic rules and which are usually represented through a series of graphic symbols. The systems more commonly used today are Labanotation and Benesh Movement Notation (BMN).
3. *Numerical data* can be made of points on a 2D or 3D plane (usually joint positions, but also weight and muscle activity), or numbers that represent speed, effort, or some other magnitude for a given motion. Time might also be added to this data.

This is not the only possible classification of motion representation (for an alternative taxonomy see deLahunta and Jennet 2017, 65). The purpose of my classification is to explain what is required for the computational analysis of motion data. The more formal a language, the more amenable it is to automatic transformations, but the harder it is to obtain. The third level is required for the computational analysis of dance. Numerical data can be obtained directly (from sensors), automatically (from video) or converted from the first and second levels. The following overview describes these levels in order of progressive abstraction, which is not always equal to their chronological development.

Concepts

Most movement traditions that I am familiar with have a specialized vocabulary to describe its motions, and sometimes its aesthetic qualities as well. The oldest known description of dance is the *Nāṭya Śāstra*, attributed to Bharata Muni, which dates back to at least the second century CE, if not much earlier. This text describes *karanas*, the basic dancing blocks of classical dance which comprise posture, gait, and gestures (Pandya 2003). Many of the first notation systems to emerge in Europe use words and word abbreviations in order to record movement (Guest 1998). These systems were devised with the intention of documenting and teaching dance, and were not as explicitly prescriptive in nature as the *Nāṭya Śāstra*, but they were also premised on specific traditions of dance and social conventions. In spite of their claims to generalizable description, most systems were short lived since they could only function within the limited parameters of a specific dance practice.

Today, the most widely used conceptual system for the description of motion is Laban Movement Analysis (LMA), also known as Laban/Bartenieff Movement Analysis. Unlike Labanotation, which will be described below, LMA is not used for notation, but for analytical purposes. LMA is used in several scientific disciplines, for example to study human-computer interaction (Fdili Alaoui et al. 2017), or to establish the connection between movement and emotion in psychology (Tsachor and Shafir 2017).

Used by dance researchers to describe dance traditions from different parts of the world, LMA facilitates comparative approaches, but it requires imposing an external analytical system to describe a dance tradition. In other words, it constitutes an *etic*, as opposed to an *emic* approach to dance

theory (in the latter, the theory is derived directly from a community of practice). To generate emic descriptions that still enable comparisons, several dance anthropologists have developed comparative guides (Royce 1977). For example, Kurath's (1952) *Choreographic Questionnaire*, takes a formalist approach and includes questions on three sections: ground plan, body movement, and structure. Later in this chapter, I will describe projects that aim to map digital motion data to dance concepts.

Notation

The development of dance notation systems independent of natural language first emerged in the context of European traditions of dance. In a comprehensive historical survey of dance notation systems in the European tradition, Guest (1998) classifies notation systems according to the graphical notation medium: words and word abbreviations, track drawings, stick figures, music note systems, and abstract symbol systems. In the context of the present overview Guest's "words and word abbreviations" correspond to the previous category of "concepts." I include all other systems in the current category. Even though track drawings and abstract symbols look very different, they are similar from a data representation perspective since they require nontextual conventions of encoding. If one were to encode natural language conventions, one could make use of the many systems for the encoding, representation, and processing of text. If, in contrast, one wanted to represent any other system in a computer, a new format would need to be devised for such purpose.

As chronicled by Guest, before the twentieth century, the most influential systems were those devised by Raoul-Auger Feuillet, Arthur Saint-Léon, and Friedrich Albert Zorn. These systems are often referenced in historical documents but rarely used today, and I could not find any digital, machine-readable implementation of these systems. The situation is different for Labanotation, also known as Kinetography Laban, which was first developed by Rudolf Laban in 1928. The system is cumbersome to learn, but widely used today, and many dance traditions have been recorded using Labanotation. The Dance Notation Bureau has an extensive catalog of dances encoded in Labanotation, as do many libraries and academies.

There have been many efforts over the years to encode Labanotation in digital formats, and develop software for this purpose. Nakamura and Hachimura have written several papers describing an XML file for-

mat (similar to the one used for TEI in chapter 4) for Labanotation and a software with a graphical user interface for creating notation files interactively, the LabanEditor (Hachimura and Nakamura 2006; Hachimura 2006). The aim of this program is to generate visual representations of the dance automatically, which can be used for the digital reconstruction of dances. The authors have worked on Noh theater and other Japanese traditions with their systems. In their own estimation, the system works well for “simple motions,” but is as of yet insufficient to fully document complex dances (Hachimura 2006). Another program, GenLaban is aimed at generating Labanotation directly from motion capture data (Choensawat, Nakamura, and Hachimura 2015). Although these projects are described in great technical detail, the software programs themselves were not available for download and testing at the time of writing. Labanotator is a commercially available tool for the creation of Labanotation scores. The software is common but it does not enable users to export or import files in XML formats, which makes comparative analysis difficult. Gábor Misi’s Labanatory (2005) is a program that runs on AutoCAD in order to search for patterns across Labanotation scores. Misi’s approach does not convert the data to XML either, but uses the built-in capabilities of AutoCAD to find visual similarities in dance movements. Misi (1983) has used Labanotation to carry out formal analyses of male solo Transylvanian dances and developed methods for the algebraic representation and analysis of Labanotation (Misi 2008). Sadly, this approach has not been followed up by others. A widely used algebraic or XML representation would enable other teams to reuse data for comparison, retrieval, and analysis.

The other system which is widely used in dance today is Benesh Movement Notation (BMN). It does not have the same level of granularity as Labanotation, but it is easier to read as it is visually similar to staff notation for music. However, the development of digital projects based on BMN is a less active area, and I am not aware of any attempts to represent BMN in XML or through other digital formats, or to use BMN for quantitative or automatic formalist analyses. However, a program for BMN has been developed by the Royal Academy of Dance (Benesh Institute n.d.).

Labanotation is more common in the United States while BMN is more common in the United Kingdom. Both have been applied to dance traditions outside of Europe and North America, but they remain essentially bound to specific cultural assumptions about movement. El Raheb and Ioannidis (2014) report that different levels of familiarity with dance traditions produce different scores for the same dance.

Numerical Data

Numerical data can be acquired directly—from markers and sensors—or as the result of processing other kinds of data: words, notations, images, and videos. When generating numerical data, there is a trade-off between invasiveness and precision. Markers and sensors are expensive and they usually need to be attached to the bodies of the performers. This means that data can only be gathered in very specific spaces and conditions. Obtaining data from other sources is either computationally very difficult, or only yields coarse data.

Sensors are electronic components that send signals to a recording device. In biomechanics, sensors are most commonly used to track the electrical activity of muscles. This kind of data is rarely used for the analysis of dance, but is commonly deployed for the analysis of movement in sports and for the development of therapies for movement disorders (O'Donoghue 2010). Inertial measurement unit (IMU) sensors, like the ones commonly found in smartphones, can track the force, angular velocity, and direction of a movement. A single sensor can't provide enough data for motion capture, but it can be used to track some aspects of a given motion, such as its speed and direction. Cuykendall et al. (2016) have developed an art project that uses smartphones to gather this kind of data, which will be described later in this chapter, in the context of data-assisted dance research.

Some companies and research teams are developing motion capture technologies based on systems of IMU sensors. But a more common solution is to use markers instead. In contrast to a sensor, a marker does not have electronic circuitry. They only work in rooms fitted with special infrared cameras that can calculate the position of the markers in a three-dimensional space. Most marker systems come with their own software packages that can infer the position and angles of joints based on marker data. A system of forty-one markers is sufficient for most applications in biomechanics. Motion capture (mocap) for film and video games commonly requires more granular systems, which typically include full-body suits, masks, and gloves. Some dance and theater companies have also used very detailed mocap graphs in performance. Since my main concern is the utilization of mocap for the analysis of movement, I won't examine the extensive literature on the implications of mocap and computer graphics for contemporary performance practice (Kozel 2007; Chan et al. 2010).

The granularity and accuracy of the mocap data determines the kinds

of analysis possible, and the reliability of the results. Professional, full-body mocap can only be achieved in hi-tech production studios, where the costs are prohibitively expensive for most researchers and artists. There are many intermediate options, with portable camera systems and different configurations of markers. The most common setup in research contexts, both for the biomechanics of sports and the biomechanical analysis of dance, is the Vicon plug-in gate model (Duffell, Hope, and McGregor 2014). Hachimura and Nakamura used this commercially available optical-type motion capture system to record Japanese classical dance and Noh theater movements (Hachimura 2006; Hachimura and Nakamura 2006). One of their primary objectives was to use computer graphics to digitally reconstruct intangible heritage, for the purpose of preservation and education, and they were thus interested in the seamless conversion between XML, Labanotation, and mocap data. As discussed above, they developed systems for the interactive display of digital animations from Labanotation and for extracting Labanotation from mocap systems.

A way to circumvent the financial and technical restrictions of expensive mocap devices is to extract mocap from other sources of data. Nakamura's team has proposed systems for generating data from natural language commands, and from Labanotation. Another common approach is to obtain data from video. This was traditionally very hard to accomplish, but it is becoming easier with software packages such as Open Pose (Cao et al. 2018). In the first excursion below, I describe, in more detail, a simple technique that Gea O. F. Parikesit and I used to obtain motion data from wayang kulit videos. Our measurements are not very granular, since we can only estimate speed, but they are sufficiently accurate for very specific questions and for video data obtained under controlled environments, and they could easily be implemented by other research teams.

An intermediate approach between costly mocap systems and consumer-end video cameras is using commonly available motion cameras such as Microsoft Kinect and IpiSoft, which showed promising results in previous studies (Arulampalam, Pierrepont, and Kark 2015). Kinect is no longer in development, but motion cameras from Nintendo and PlayStation could perhaps be used for dance motion capture in the future. Cutting-edge algorithms can automatically estimate the number of persons in a scene, generate their corresponding 3D skeletons, and estimate their locations (Elhayek et al. 2018). Video data can also be used for gait identification, as a video surveillance system can identify people

based on their habitual patterns of walking, even when little data is available (Balazia and Sojka 2017). This has terrifying social implications for privacy, but it could be a boon for the analysis of dance. Still, these algorithms have yet to be applied to dance and physical theater data.

The advent of game-changing technologies might soon radically alter the types of research that are possible. As previously noted, Cao et al. (2018) have developed Open Pose, an open source software that uses a technique called Part Affinity Fields (PAFs) for real-time estimation of the position of dancers in a video with high accuracy, even when there is visual noise in the background of a video. Broadwell (2019) demonstrated in a workshop in the 2019 Digital Humanities conference that Open Pose could be used to obtain movement data from K-Pop dance videos and that this data could be potentially used to compare dance styles across a large number of videos.

Data-Driven Motion Analysis

Several scientific studies have used biomechanical methods to study dance, but they are mostly interested in preventing injuries in professional dancers or in using dance as a therapy for specific kinds of trauma or disease (Koutedakis, Owolabi, and Apostolos 2008; Luksys and Griskevicius 2016). Other studies have explored the intersection of biomechanics and more subjective experiences of dance. A review by Chang et al. (2016) looked at the correlation between the subjective perception of beauty and the biomechanical markers of skill, such as speed, vigor, and smoothness. The review found significant differences between expert and non-expert judgment, but no general trends in the correlation between aesthetic experience and biomechanical markers of skill. This is perhaps unsurprising to dance scholars who are familiar with the wide range of culturally specific ideals of dance (see Redding 2019, 62). For example, Javanese dance is generally slower than Balinese dance and there is little point in taking their different speeds as proxies for skill, appropriateness, or beauty. To determine whether a movement is beautiful, well-executed, or appropriate, one needs to be familiar with the specific dance culture where the motion is evaluated. A motion that is considered beautiful for one kind of dance would be considered ugly if performed in other dances. However, the authors of the review, writing from an engineering perspective, confidently hypothesize that the range of responses is due to the lack of clear goals in dance. This is different to sports (a much beloved field of biome-

chanical analysis) where there are usually clearly defined goals that lend themselves better to measurement and comparison.

Closer to the objectives of dance studies, Hachimura (2006) proposes a method for measuring similarities between dance gestures. The challenge, as he explains, is that two gestures perceived as identical by human observers might look very different in the data. Even when considering simple motions such as walking, “there are instances where you start with your right foot and instances where you start with your left foot” (Hachimura 2006, 59). There are also differences in terms of speed, direction, and position for different motions that could still be considered to belong to the same conceptual category of walking. Thus, he had to develop ways of measuring similarity that would be robust to these small variations. For this he used a dynamic programming (DP) algorithm which is often used in voice and handwriting recognition.

Hachimura also suggests that a minimum convex polyhedron can be calculated for different slices of time. An alluring insight hides under this technical term. To explain this, let me take a simpler concept first. A minimum convex polygon is a two-dimensional shape that encompasses a series of points. “Minimum” means that it is the smallest polygon in which no internal angle exceeds 180 degrees and which contains all the points. Minimum polygons are used in biological maps to show the extent of an animal population, or the area actively defended by an individual of that population. The minimum complex polyhedron is a three-dimensional version of this concept, or a solid in three dimensions with flat polygonal faces, straight edges, and sharp corners. For dance data, this is the minimum three-dimensional shape that encompasses all the joints of a dancer in a given gesture, as frozen in a moment in time. The total volume of this polyhedron can be calculated for any moment in a dance, and plotted in a graph as a function of time. This is useful for comparing two dances, and serves the purpose of retrieval based on similarity. Another interesting possibility, also mentioned by Hachimura, is that the principles of LMA, such as space, effort, and shape (which are extensively used by dance scholars, as discussed above) can be expressed in terms of polyhedra that change through time.

In the case of dance, it is also important to retrieve specific gestures of the feet or the hands, rather than whole body motions. Hachimura’s example is the *okuri* motion, which is important for Noh performances. Hachimura’s paper describes both methods with broad applicability and detailed examples that show engagement with specific dance traditions.

Another comprehensive set of dance methods has been proposed by Wiesner and her collaborators (2012). They created a database of movements, the ARTeFACT Movement Thesaurus (AMT) which encoded two hundred common dance movements into a semantic XML database with accompanying motion capture data. Subsequently, they used this data to automatically identify, annotate, and retrieve film data (Wiesner 2012). Using biomechanical analysis, they were able to achieve high accuracy in the automatic classification of these movements (Simpson, Wiesner, and Bennett 2014). Acknowledging that dance is more complex than individual steps, they have also developed systems to identify conceptual metaphors in dance (such as “conflict”). They have applied statistical procedures and borrowed principles from corpus linguistics to identify patterns in their biomechanical dataset and in verbal descriptions of movements, to identify features that correspond to conceptual metaphors.

As the previous projects show, the estimation of difference and similarity is central to the data-driven analysis of dance. MoComp is a tool specifically developed for the comparative visualization of mocap data (Malmstrom et al. 2016). In their interactive web-based visualizations, the authors used streamgraphs to represent changes in angles and angular velocities over time for a given joint. This project is explicitly premised on a scientific understanding of motion. Later, I will describe visualizations geared towards data-assisted analysis. We will see that many data-assisted projects are interested in the affective qualities of movement, but this is also important for some data-driven research projects. Li and Pasquier (2016) have looked at how different statistical and machine learning approaches can be used to identify affect in movement, using annotations made by performers and observers as ground truths (an example of the calibration practices discussed in chapter 2). Their project deals with affective dimensions but aims at consensus-seeking, intersubjectively verifiable observations. The projects discussed in the next section, in contrast, aim to enact affective responses through highly interpretive visualizations.

Data-Assisted Motion Analysis

In a more recent project, Wiesner and her collaborators (2016) worked with the creators of POEM, a mobile platform that can convert movement into whimsical poems, in order to provide a situated classification of movement through “embodied visual analytics” (Cuykendall et al. 2016;

Cuykendall, Soutar-Rau, and Schiphorst 2016). The collaboration between ARTeFACT and POEM builds on the strengths of both projects, combining the extensive data collection capabilities of POEM with the nuanced data model of ARTeFACT (Wiesner et al. 2016). This project helps validate and generalize the ARTeFACT data but also shows its potential for creative applications. Another example of performative visualization is found in Nakamura's work (2017), where she used lines of different colors to trace the movements of fingers, arms, and trunks of Balinese dancers as they moved through space. This provides a situated visual interpretation rather than replicable conclusions, but it also builds on a data-driven project, as is the case of the ARTeFACT/POEM collaboration.

Interactions between choreographers, engineers, and designers further expand the scope of data-assisted movement analysis. These collaborations have been thoroughly chronicled by deLahunta (deLahunta and Jenett 2017; deLahunta 2017). Of particular significance is *Synchronous Objects for One Flat Thing, Reproduced* by Ohio State University's Advanced Computing Center for the Arts and Design and the Department of Dance in collaboration with William Forsythe (Forsythe et al. 2009; Palazzi and Shaw 2009). This was perhaps the first collaboration between dancers and data scientists to achieve international notoriety in the dance world. In parallel with this project, William Forsythe and others started Motion Bank, with support from the German Federal Cultural Foundation and other sources. The project has generated reusable data and software packages. A video annotation tool named *Piecemaker* was developed by The Forsythe Company member David Kern between 2007 and 2013, and later redeveloped as *Piecemaker2* (PM2) and *PM2GO*.

Motion Bank features the work of several choreographers and, even though some tools are shared across the platform, each choreographer's portal attends to their specific ideas of dance practice. For example, the choreographies of Deborah Hay are not meant to be repeated exactly: "the movement may change, but the choreography itself does not change" (quoted in deLahunta 2017, n.p.). For Hay's work, the codification and transmission of inflexible dance scores makes no sense. Instead, six performers were invited to perform the same solo dance according to their different interpretations for Motion Bank. Each version was recorded and processed, and visualizations were invoked to explore these differences. One such visualization uses minimum convex polygons, the same mathematical objects I described earlier. Here, they are instances of what I call "performative visualizations" (see chapter 3), and they are used to

highlight differences between dance performers as much as to visually convey a key principle of Hay's choreographic philosophy. In the previous example, the same quantitative method was used for information-retrieval within a data-driven, scientific paradigm. The transferability of the same technique across methodologies emphasizes a key theme in this book: quantitative methods alone do not signal a specific attitude to knowledge. What matters is the way in which a method is contextualized within a methodology—the same method can be used to pursue consensus-seeking, replicable conclusions, or situated, performative hermeneutics.

Bermudez et al. (2011) also endeavor to communicate specific choreographic principles with data. They have developed interactive installations to explore the principles of *Double Skin/Double Mind*, the choreographic approach of Emio Greco | PC. Other platforms have also been developed for the creative visualization of dance data (see Carlson, Schiphorst, and Shaw 2011). But special attention should be paid to the work of Ribeiro, Kuffner dos Anjos, and Fernandes (2017), who have used 3D data derived from a Kinect device to create interpretive visualizations that represent the movement and improvisational principles of Portuguese dance creator João Fiadeiro's composition in real time (*Composição em tempo real* in Portuguese). Working together with him, the authors identified a subset of core concepts of Fiadeiro's method which were explored through several performative visualizations. One such principle is *suspension*, a moment in a choreography from where several options are available to a dancer. In the context of this dance approach, it is important to think about these moments as *possible futures*. A series of cubes juxtaposed in a 3D space were used to emphasize this perspective. In another visualization, they wanted to emphasize what happens when a performer leaves the stage and another enters. This is described as the *cycle of vitality* in Fiadeiro's work. This quality is performatively brought out in visualizations where the 3D representation of a dancer loses color as they leave the stage and another figure regains color as they enter the space and begin a new cycle of movement.

This last example highlights that all thinking about dance must be situated within specific cultural and analytical contexts. Ribeiro, Kuffner dos Anjos, and Fernandes's work is a felicitous combination of data visualization and contextual sensitivity and can serve as an example for projects in other dance traditions. Another striking performative visualization is EMVIZ (Subyen et al. 2011), which aims to produce visual representations of the Laban Basic-Efforts. They describe their project as an exam-

ple of artistic visualizations which combine “interpretive metaphoric mappings with aesthetic approaches in representing data from one domain to another” (Subyen et al. 2011, 121). They characterize artistic visualizations as data mappings that “are interpretive, subjective, and follow a different set of conventions than those governing information visualization or scientific visualization.” This echoes my descriptions of performative visualizations in chapter 3. Subyen et al’s system has been primarily used in the context of art galleries, where dancers wore sensors while performing a variety of movements, and real-time visualizations were projected behind them. It is important to highlight how EMVIZ, as well as the work of the teams led by Bermudez and Ribeiro differ from the MoCamp platform we encountered in a previous section of this chapter. All of these are computational transformations of dance data. But while MoCamp aims for the readability and systematicity that befit a scientific analysis, the projects we have encountered in this subsection enlist an interpretive sensibility for the production of data representations, which are in themselves art projects.

First Excursion: Calculating the Speed of Puppet Movement in a Wayang Video

In what follows, I describe two small research projects, which I carried out with my collaborators in order to study the implications of motion analysis for the theater traditions of Java. In the first, I collaborated with Gea O. F. Parikesit, a physicist who works in image and video processing. Here, I summarize our research, and show some new visualizations of the data we published earlier. We used Catur Kuncoro’s *Wayang Mitologi* (2012), a sixty-nine-minute video from the Contemporary Wayang Archive (CWA, available at <http://cwa-web.org/en/WayangMitologi>) and, calculated the speed of the video using *difference images*. A difference image is an image that results from subtracting one image from another. To get a better sense of our method, it is important to remember that a video is a sequence of digital images, each of which is made of pixels. In turn, each pixel contains information for three color channels (red, green, and blue), and each channel has a *grayvalue*, which indicates its brightness. Below is a short nontechnical summary of our methods. More comprehensive technical details can be found in our original publication (Escobar Varela and Parikesit 2017).

1. We converted the video to a series of static images.
2. We computed the difference image for every pair of subsequent images.
3. We calculated the total grayvalue of the pixels in each difference image where the value was above a certain threshold. The reason for including a threshold is that some of the differences are just noise. When recording a black image on a consumer-end video camera there will be small differences within black image frames. Thus, we used two black images in order to estimate the noise threshold and then counted every image with a gray value higher than such threshold—that is, every image that is a reasonable indicator that there was a change from one frame to the next.
4. We plotted the resulting values in a graph, with the number of non-identical pixels as a function of time. Figure 6.1 shows this in a slightly different form to the one in the initial publication.

A visual inspection of the patterns in non-identical pixels in difference images led us to believe that different scenes had different movement profiles. In the original paper, we explored these differences by manually segmenting the scenes and then using a statistical procedure to identify differences between them. Here, I use the same segmentation and descriptive statistics but present them through new visualizations. The main reason for offering new visualizations is that I have presented the original results several times after the initial publication, and in that intervening time I have thought of new ways to present the data, which I wish to explore here (the original images and analysis can be seen in the aforementioned paper). Part of my interest was to develop a more interpretive visualization, which I call the *kayon plot*, as I will now explain.

For background, let's consider that *Wayang Mitologi* is a contemporary version of the classical *wayang kulit* form described in chapter 5. There are many differences between classical and contemporary wayang, but the most important for present discussion is that contemporary wayang is usually shorter: a couple of hours rather than the seven to eight hours of the all-night, classical wayang (Escobar Varela 2015). *Wayang Mitologi* has different kinds of scenes: frame scenes, a comic interlude (or *gara-gara*, a fundamental part of a wayang performance), narrative scenes, and normal scenes. The comic interlude, narrative sections and the “normal scenes” are part of classical wayang conventions, while “frame scenes” are contemporary inventions. In the recording of *Wayang Mitologi*, the scenes




Fig. 6.1. Non-identical pixels in subsequent images as a function of time. Adapted from Escobar Varela and Parikesit (2017).

were segmented and labeled by hand. Figure 6.2 compares the mean and standard deviation of the number of non-identical pixels in subsequent images for each scene. A scene with a higher mean will be, on average, faster than one with a lower mean. A scene with a higher standard deviation will be less homogeneous than one with a lower one, as it will tend to include a wider mixture of slow and fast sequences. Figure 6.2 shows very clearly defined clusters. Narrative scenes stand out from the rest in terms of both mean and standard deviation, and can be seen in the upper right quadrant of the graph. Likewise, the two frame scenes have almost overlapping values near the bottom left. The normal scenes and the comic interlude occupy the midsection of the graph.

It is interesting to note that the comic scene is very similar to normal scenes. But it is also much longer than other scenes, and this information is not captured in the visualization in figure 6.2. Another missing aspect here is sequence: we don't know how scenes of different kinds follow one another. To include length and sequence information, I combined violinplots and bar charts (figure 6.3). The violinplots show the distribution of non-identical pixels within each scene, while the bars in the lower half show the scene's duration. A violinplot shows the distribution of values




Fig. 6.2. The standard deviation and mean of nonidentical pixels for each scene, grouped by scene type.




Fig. 6.3. Combined violinplots and bar charts. The violinplots show the distribution of non-identical pixels within each scene (the mean is indicated by a white circle within each plot). The bars show the length of each scene.

for a given observation. The violinplots in the upper half of this graph are cut off at 0 (since a difference image can't have a negative number of pixels). Thus, they look somewhat flatter and less like the musical namesakes of more conventional violinplots.

But like other violinplots, the width at each step of the y axis (e.g., 100 or 4,000) tells you how many such values were found within that range (this is similar to the KDE plots described in chapter 5, except that they are rotated 90 degrees). A wider contour at the bottom means that the scene in question had more lower numbers (i.e., slower parts). Longer shapes include higher numbers (i.e., faster parts). The white circle inside each shape represents the mean value. Narrative scenes (5 and 7) have both similar shapes. They are slender and long—this means that they have a wide distribution of values, but few instances of each value. If you go back to figure 6.2 you will see that these narrative scenes (represented there as diamonds) are in the upper-right section of the graph (another way of showing their high mean and high standard deviations). But figure 6.3 provides the additional information that scenes 5 and 7 are shorter than other scenes, as indicated by the bars at the bottom half of the graph.

When exploring this visualization, I was reminded of the shape of the *kayon*, a puppet with a distinctive shape with important functions in wayang kulit (figure 6.4). This puppet is used to indicate scene changes and forces of nature, and it is also used to open and close a performance (Arps 2016). The *kayon* has an important cultural significance in Java, and modified versions of its shape are used in the logos of many cultural institutions. I wanted to use this shape to visually present the information discussed thus far. Like all wayang puppets, this one is made of water buffalo hide and it has a protruding control stick (traditionally made of bone). Figure 6.5 represents each scene as a *kayon*. It is very similar to the previous graph (figure 6.3). The upper half is identical, but the y axis of the lower half is inverted, and the length of each scene is represented by a line rather than by a bar. These simple modifications make each shape somewhat reminiscent of a *kayon*. Let me offer a theoretical justification for the strangeness (and perhaps failures) of this graph. This plot doesn't look like most scientific visualizations and it is an attempt to answer Drucker's call for humanistic visualizations (see chapter 3). However, this visualization doesn't fully eschew the principles of statistical graphs either. It includes statistical information (mean and standard deviation), has a low ink ratio and it enables consistent comparisons (the principles advocated by Tufte). This visualization is a situated *deformance*, but is also statistically sound.




Fig. 6.4. A kayon puppet, from the front side.




Fig. 6.5. A kayon plot. Similar to figure 6.3 but the y-axis of the lower half is inverted, and the length of each scene is represented by a line rather than by a bar.

Even if other research teams don't develop their own performative visualizations, there is great potential in the usage of difference images for the analysis of theater video recordings. Difference image calculation is not as sophisticated as some of the other methods described above. However, its advantage is its very simplicity: its meaning is clear and it is easy to implement. This is also one of the methods proposed by Lev Manovich (2013) for the analysis of film. It must be noted, though, that the needs of quantitative motion analysis for theater are different from those of quantitative film analysis. The theater video recordings I use here are those developed for documentary purposes, and they constitute digital traces of a performance. In film, the materiality and aesthetics of the medium are central concerns. Most of the methods in *cinematics*—the quantitative analysis of film—aim to reveal formal aspects of film language, such as the number and length of cuts (Salt 1974; Tsivian 2005; Buckland 2008; Brodbeck 2011). Cuts are arguably one of the key aesthetic building blocks in film and patterns in cut usage are strong markers of authorial style. In contrast, the best kind of video for the digital analysis of motion in theater is one where there are no cuts, as the one we used here (since a cut would appear as a large number of pixels in a difference image).

If one is interested in the movement of actors or dancers, and wants to limit the influence of lighting, then the best video is one produced during a dry run. This, of course, depends heavily on context and traditions. Wayang kulit videos can be recorded from the side of the shadows, with minimal color and depth interference in the analysis of motion (it should be noted, though, that the majority of present-day spectators watch wayang from the front, which is the side where they can see the dhalang animating the puppets). For many theater forms and genres, an analysis of speed can give a quick overview of a recording, and can provide quantitative measurements for comparative analysis. The techniques used here could be adapted to seek answers that are pertinent to specific traditions of practice. Copyright still limits the kinds of usage that a video can be submitted to. This might likely change in the future if we implement systems for collecting video that enable future analytical uses without compromising present-day copyright. For example, a theater venue could record all the performances from a single camera position and only enable those videos to be used twenty or fifty years in the future (or to be only used for research rather than for distribution).

Second Excursion: Quantitative Analysis of Javanese Dance

For the research project summarize here, I worked with Luis Hernández-Barraza, a researcher in the field of bioengineering who works on the biomechanical analysis of sports. We borrowed some principles from sports biomechanics and tried to identify a question that would be of relevance to Javanese dance, and which could be answered with the tools at our disposal. We concentrated on *character types*, since these are of enormous importance to the pedagogy, analysis, and enjoyment of Javanese dance. As in many Asian theater traditions, characters in Javanese dances are defined along different types, each of which follows conventions of movement, speech, dress, and action. There are many dramas and dance-dramas in Java, but we focused on *Sendratari*, since this genre includes character types but no dialogue. *Sendratari* started in the 1960s as a touristic performance genre. It is still a dance that is mostly seen by tourists, both foreign and local. However, the *Sendratari Ramayana* has an important place in the city of Yogyakarta (where the present author lived for some years). Although some people use the moniker “touristic” to denote a lack of authenticity, touristic performances are authentic in their own right: they have their histories and politics and are presented for a specific audience to achieve a particular purpose (Bruner 2005). Different groups take turns performing the *Sendratari Ramayana* at the Prambanan temple in Yogyakarta, and there are many aspects of the performance that are open to creative improvisation by the different groups. But the rules governing character types are relatively stable, and they are amenable to biomechanical comparison. There are many qualitative descriptions of the differences between character types, and we used the terms and definitions proposed by the late R. M. Soedarsono (1983), a leading expert in Javanese dance and performance. While we think this list is authoritative, it should be noted that there is some disagreement among the dance community about the names of the character types in the *Sendratari Ramayana* and this is something that future research could investigate further. For our project, we focused on main character subtypes for male characters:

- *Impur* (refined-humble) and *kagok-kinantang* (refined-proud), two subtypes of the *alusan* (refined) type.
- *Kambeng* (strong-humble) and *kalang-kinantang* (strong-proud), two subtypes of the *gagahan* (strong) type.



Fig. 6.6. The theoretical location of all subtypes except Jatayu along the strong vs. refined, and proud vs. humble axes.

Figure 6.6 shows how these subtypes can be classified on the humble vs. proud, and refined vs. strong dimensions. We were also interested in the character of Jatayu, a mythological bird that has been only recently added to Javanese dance and which is not found in Seodarsono's comprehensive list of characters.

We had two questions, each of which was explored in a separate paper. First, we wanted to analyze whether there is a difference between the refined and the strong characters in terms of common biomechanical measurements. Second, we wanted to understand how the visualization of biomechanical data could contribute to an interpretive stylometry of dance. We invited a professional dancer to perform the same motion (standing up from a kneeling position) as it would befit the different character subtypes, and we used the Vicon motion capture system, which consists of eight infrared cameras and two force plates, to collect kinetic and kinematic data at a sample rate of 100 hertz, full technical details are available elsewhere (Hernandez-Barraza, Yeow, and Escobar Varela 2019). We used the Vicon Nexus 1.8.3 and Polygon 3.5 software for data collection

and processing. We distinguished between *kinetics* and *kinematics* following their technical distinction in biomechanics: kinematics refers to motion description whereas kinetics attempts to get at the cause of the motion. Thus, we obtained the following kinds of data:

1. Kinetic data: Ground reaction force (GRF) and joint moments
2. Kinematic data: joint angles and angular velocities

Our hypothesis was that refined (as opposed to strong) character subtypes would exhibit lower values for (1) range of motion (ROM), (2) angular velocities, (3) ground reaction force (GRF), and (4) joint moments. We used one-way ANOVA, followed by a Holm-Sidak posthoc test, to compare the peak vertical GRF, ROM, angular velocities, and joint moments of the character subtypes. All significance levels were set at $p = 0.05$. We found that we could only partially accept this hypothesis for the reasons summarized in table 6.1 (full analysis at Hernandez-Barraza, Yeow, and Escobar Varela 2019). The difference between the character subtypes is more nuanced than what we expected. While these results are not impressive, it is important to report negative findings, and resist the temptation to overinterpret the data. Many data-driven projects will find no evidence of the hypothesis they sought to prove. The road to replicable and incremental knowledge is built by negative results as much as by awe-inspiring breakthroughs.

For the second project, I created a website to explore the data interactively through different visualizations that are linked to animations of the dancer's skeleton. This project can be consulted at <https://villaorlado.github.io/dance/html/> (see video 6.1 in this website's online companion).

Whereas the results in table 6.1 are data-driven, the website enables an interactive, data-assisted analysis. In the website, users can load videos

Table 6.1. Comparisons of biomechanical markers between refined and strong characters

Measurement	Expectation	Result
ROM	Lower values for refined characters	Only true for right knee, and wrist
Angular velocities	Lower values for refined characters	Only true for left ankle, right knee, right shoulder, and right elbow
GRF	Lower values for refined characters	Only true for the anterior-posterior (AP) component of the right leg
Joint moments	Lower values for refined characters	Not true for any joint

that correspond to each character subtype. They can then choose any joint (e.g., left knee) and see how its angle changes along the x , y , or z axis. This change is displayed alongside the video, and it allows a defamiliarized way of looking at each character subtype. The page for each character subtype also includes photographs of the dance, and notes on the interpretation of this movement in the Javanese context. The website also allows users to compare the motion of the same joint for different character subtypes. For example, users might choose to compare the angles of the left knee for the *impur* and the *kambeng* character subtypes. Lastly, the ROM for the different character subtypes can be compared in interactive graphs. These visualizations add context to the data, as users can see how videos relate to the data points. The website also has several exploratory features that enable users to choose where to focus their attention, which data to compare and how to proceed in an “ergodic, yet immutable” sequence (see chapter 3).

The data used here is limited, but these visualizations are easy to implement and potentially very useful for the interpretive stylometric analysis of dance and of motion in theater. Whereas the scrutiny of character types only makes sense within certain traditions, combining animations and interactive visualizations is useful for the comparative analysis of a wide range of movement techniques, from *Butoh* to *Tango*. If movement data is extracted from video rather than generated from markers or sensors (as detailed earlier in this chapter), then the application of similar techniques to movement stylometry could grow at a faster pace. However, these tools will always be limited to very specific questions and there are many caveats to be considered when analyzing movement, as shall be explained in the remainder of this chapter.

Concluding Thoughts on the Application of Quantitative Stylometry to Dance

Any quantitative study of culture runs the risk of oversimplifying its object of study. Dance is in a constant state of change, and this is particularly true for the hectic centers of practice and scholarship in Java. Documentation might give a misleading impression of fixity in what is in reality a dynamic field. This impression might be created in good faith, as researchers engage in necessary simplifications in the pursuit of specific research questions. However, fixity might also help controversial political agendas. In the case of Java, fixity of cultural forms has been linked to both the Dutch colonial project and to the repressive Suharto dictator-

ship (Sears 1996). According to Dyah Larasati, the Suharto Regime (1966–1998) systematically prosecuted dancers, while at the same time trying to craft an unchanging, idealized version of Javanese dance (Larasati 2013). Larasati suggests that these phenomena should be read as the two sides of the same coin.

The history of dance is always political, as dancing bodies have often been contested sites. I believe that the specific things we tried to measure in the projects reported in this chapter are not contentious, but this doesn't exempt me and my collaborators from considering the unintended political ramifications of our work. The availability of technological tools should not be used to the detriment of the historical and ethnographic analysis of dance and other movement traditions. A data-driven analysis of motion which is grounded on critical realism should always consider the historical and political context of any performance tradition. We also shouldn't forget that dance is an embodied and affective practice, for both creators and spectators. When describing the complex, funny, and inter-sectional dances of Javanese dancer Didik Nini Thowok, Jan Mrázek notes that dance can't always be interpreted, labeled, or described: "laughter, tears or an upset stomach will be always a better reaction than an academic paper" (Mrázek 2005, 279).

Digital data might be an even less appropriate reaction. The dangers of oversimplification enhanced by technology are apparent to me when I speak about the excursions in this chapter to technical audiences. A common question is whether my data and results could be used to teach dance (or puppetry). My answer in the negative often earns me condescending smiles. In engineering disciplines, there is a strong emphasis on real world applications, actionable knowledge, and problem solving. In the projects I envision in this book the objective is not so much to design interventions in the world of performance, but rather to find different ways of analyzing culturally situated aspects of theater that are amenable to data and computation. My collaborators and I are often told by technical journal editors and conferences attendees that our research needs to be more practical and applied. Scientists are often more open than engineers to the pursuit of knowledge to gain understanding of a phenomenon, and they tend to be more interested in pursuing questions that might not have direct application (however, this sweeping distinction between pure and applied research must be taken with a pinch of salt, since it is context-dependent and changing rapidly).

I would be worried if the methods presented here were to be used

to assess dancers or other performers (for intake admissions to dance schools or for hiring into professional troupes). I would also dread the possibility that mocap data produced under very specific research conditions is one day considered the canonical or authentic version of a given dance. This is a real danger of computational theater scholarship that must be addressed in relation to all projects. It is important to document the conditions under which data was collected, and the assumptions that went into the collection of data. This is as important for data-assisted as it is for data-driven research.

Anthropologists of dance and scholars of physical theater are—and will continue to be—interested in things that motion capture and biomechanical analysis will never capture: the social contexts, meanings, and controversies of dance and physical theater. As Balme (2015, 165–66) notes, gender and gendered representations are key concerns in current dance scholarship. This is an area perhaps not well captured by motion data alone. The more enticing promise of quantitative motion analysis is that it can be used to understand dance history in more textured ways. Tracing similarities between movements in different historical contexts, or from neighboring geographic regions could be used to study the evolution and spread of dance patterns. As this chapter has laid out, such goals will require access to large datasets that are sufficiently standardized and shared across research teams. Although several researchers are working on related fields, the level of community work this would require is still far in the horizon. However, the range of projects chronicled in this chapter show the enormous promise for the computational analysis of dance, dance-theater, physical theater, or movement in theater more generally.

Code and data: Sample code for this chapter is available at

<https://doi.org/10.3998/mpub.11667458.cmp.41>.

The data used can be downloaded from

<https://doi.org/10.3998/mpub.11667458.cmp.34> and

<https://doi.org/10.3998/mpub.11667458.cmp.35>.

CHAPTER 7

Location as Data

Theater is bound in time and space more than many other artforms and more attention should be given to the geographic distribution of theater performances. This is an area where data is available in large quantities and computational methods are uniquely useful. Films and novels depend on specific, place-bound infrastructures: recording studios, book-signing events, awards ceremonies, etc. But the experience of reading and watching film is more distributed through time and place than experiencing theater events live. The visual arts have their own complex relationships to time and place, as artworks are created, sold, exhibited, and stolen in specific places. But the ontology of most visual artworks is less strongly tied to events that are etched in time and place. This chapter looks at computational methods for the analysis of theater events as spatial and spatiotemporal data. It draws heavily on GIS (geographical information systems) but considers other relevant geographical tools, including analytical methods for combining spatial and temporal analysis. I do not consider the 3D reconstruction of historical theater buildings, which is a separate area of research (Ioannides et al. 2016; Manzetti 2016).

I believe that spatiotemporal analysis is the area that holds the most promise for theater studies—in chapter 6 I said that the analysis of motion is potentially the most interesting, but the technologies are much harder to use and there are many more roadblocks. In contrast, for geospatial and geotemporal analysis, the needed data is more easily available, and there are many ways to use this data within different types of projects. As I have done in previous chapters, I will first describe some general methods for visualization and analysis before explaining how they can be used within data-driven and data-assisted methodologies for spatial theater research.

Spatial Visualizations

Perhaps the most common spatial visualizations are cartographies. These visual artifacts show the location of venues and events as points in a 2D map. Even moderate amounts of data can reveal interesting patterns when placed on a map and answer a host of interesting questions. Where are the contributors to a specific theater conference based? What are the national origins of the members of a circus troupe? What are the touring locations of commercial troupes and how do they differ from those of troupes that depend on festival circuits? Plotting points and regions on maps can constitute useful heuristics to identify inequalities or patterns that warrant more investigation.

In many cases cartographic visualizations are sufficient to trigger questions and suggest new interpretations of available data. However, as Jessop (2008) notes, GIS offers much more than just cartographic possibilities. And GIS might not necessarily be the best option. When dealing with historical data, for example, gazetteers (databases that include historical name changes) might be more useful for researchers in the humanities. Another common way to visualize geospatial data is to use choropleths. In these maps, areas are shaded according to different quantitative values, such as the number of performances in a given country or region. The first choropleth map was produced in either 1819 or 1826 by Pierre Charles François Dupin, to show rates of illiteracy in different regions of France (Friendly 2001). Since then, choropleths have been used for a variety of purposes but they are less common in the humanities than in the natural and social sciences. In the humanities, they are most commonly used in linguistics (Liao and Petzold 2010). One problem of choropleths is that large geographical areas will have larger shapes, which might suggest they are more important. Large but sparsely populated provinces will tend to take up a large space in the visualizations. A possible solution is to use cartograms, which distort the shape of geographical entities (countries, districts, etc.) according to a particular metric or statistic. The resulting maps show how much a given area contributes to the total for a given metric in a way that is not distorted by the land area of a province or country. They are also rare in DH but they have been used in the analysis of film distribution networks (Arrowsmith, Verhoeven, and Davidson 2014).

The difference between cartographies and spatial visualizations is not cut and dry. All maps are, in some sense, visualizations of geographical data that are the product of specific cultural and epistemological assump-

tions. Yet, it is useful to distinguish between maps that show events as points and lines and visualizations that include alternative modes of displaying geographical data. The reason for this distinction is that these other modes usually require some type of quantitative transformation, even if such transformations are not foregrounded (for example, the color scheme of a choropleth is the result of a statistical transformation of the data). Choropleths and cartograms require statistical transformations of numerical data, but these do not exhaust the possible uses of statistics for the analysis of locations.

Geostatistical Methods

A common purpose of statistical procedures, as described in earlier chapters, is to estimate the likelihood that an observed pattern is not a fluke, but that it reveals the true existence of a phenomenon. In geostatistics, this means that researchers are interested in showing that a geographical observation is due to an underlying cause, and is not random. Statistical methods can be used to test a given set of observations against the assumption of complete spatial randomness (CSR). Several statistical tests can be used to distinguish between spatial clustering and CSR, such as the Clark-Evans test (Ripley 1979) or a *chi-square* goodness of fit (Rogerson 1999). I am not aware of any usage of CSR measurements in computational theater research, but CSR has been applied to study linguistic data (Leino and Hyvönen 2008; Kretzschmar 2013) and tombstone locations (Streiter et al. 2012). To imagine a potential application for theater data, let's assume that we had data of people attending theater shows in a particular venue in a given city. When plotting the addresses of the theatergoers, we might see that they tend to cluster around certain neighborhoods. However, this pattern might well be random. To more rigorously validate the hypothesis that theatergoers tend to live in certain neighborhoods, a Clarke-Evans test could be used to distinguish spatial clustering from CSR in the addresses.

Another common object of geostatistical inquiry is to determine whether geographical proximity affects a given variable, a feature termed *spatial autocorrelation*. In social science, this often means determining whether a country's value for a certain metric (e.g., "indexes of democracy") is affected by the values for the same metric in the country's neighbors (Ward and Gleditsch 2019). The most common statistic used for this purpose is Moran's *I* (Moran 1948), which estimates the probability that a

value for a geographical unit is affected by the average value of neighboring units. In DH this statistic has been used to analyze linguistic variations across different places (Asnaghi, Speelman, and Geeraerts 2016). An example for theater would be to test whether the presence of certain kinds of festivals in a country are affected by the presence of such festivals in neighboring countries. The excursion at the end of this chapter offers a more detailed example of the usefulness of Moran's *I* for theater research.

Data-Driven Location Analysis

Geostatistical estimation methods fit squarely within a data-driven paradigm, but are uncommon in computational theater research. Cartographies and other spatial visualizations can be used both for data-driven and data-assisted location analysis.

An excellent example of data-driven cartographic work is the research by Holledge and her collaborators (2016). As seen in chapter 5, Holledge and her coauthors explicitly invoke a scientific paradigm in their work, and aim at answering a close question: what accounts for the global success of *Et dukkehjem*? Holledge and her collaborators use Ibsen's original Dano-Norwegian title, rather than the standard English version, *A Doll's House*, to refer to an entire body of work that comprises translations, adaptations, and productions in film and other media (2016, 4). In their analysis of the early global spread of the play before World War I, they show that women played a fundamental, and previously underrecognized role. While conventional narratives attribute the success to male backers of Ibsen, Holledge et al. convincingly demonstrate that women were the entrepreneurial force driving the spread of the play. Their success can be linked to international movements of women's emancipation, but also to the routes of colonization, migration, and specific biographies. Holledge et al. present this argument through detailed historical analysis and through a series of cartographic representations. They identified key performances in forty-six countries across five continents. To find these key performances, they listed all contributors in the previously mentioned IbsenStage database who were connected with performances in three or more countries before 1914, and then used this in conjunction with production records to identify twenty-two performer-entrepreneurs responsible for the spread of the play around the world (Holledge et al. 2016, 29). Their geographical visualizations show the locations of performances as dots, which are connected via straight lines that denote sequential travel patterns.

These routes disappeared with the war in 1914. International touring will reappear in the late twentieth century, but then it will be driven by different forces. Rather than land and sea routes, troupes will use air travel and subsidized international festivals (Holledge et al. 2016, 64–65). In another chapter of their book that discusses the role of the state in the success of the play, Holledge et al. present a map with all productions of *Et dukkehjem* imported to the Ibsen Stage Festival in Oslo from 1990 to the present day, next to a map of the international distribution of Norwegian government funds promoting *Et dukkehjem* as a global play text (98). These two maps are reverse mirrors of each other. Both use arrows, but their direction is reversed—from production places to Oslo in one instance, and from Norway to places receiving funding in the other. Taken together, these two maps show the close link between funding and international exposure. The cartographic displays in Holledge et al. were previously presented by Bollen and Holledge (2011) to an audience of cartographers. In that previous article, they argue that theater maps show the “importance of distributional flows through time, across geographical space, and between artists from production to production” (226). Although their cartographic work is primarily data-driven, the authors also note the importance of complementing visualizations with in-depth social, historical, and political analysis. Thus, their research occupies a nuanced position in the spectrum between data-driven and data-assisted work. Their work is data-driven inasmuch as it answers closed questions with verifiable data, but it is data-assisted when it zooms in and out of different scales, switching between interpretive historical analysis and the distant vision afforded by the maps.

As part of a project that examines the convergence of people, theater venues, and media, *Circuit: Mapping Theatre Performances in Victoria* (Tombe et al. 2017) used data from AusStage to visualize connections between touring productions among theater venues in Victoria, Australia. These connections were overlaid in interactive choropleths that also include statistics on the ethnic diversity of the cities where the venues are located.

Other types of geographical visualizations and geostatistical analyses are less common in theater research. Thus, I turn briefly to film studies in this overview to show projects that could provide inspiration for the data-driven analysis of theater locations. Arrowsmith, Verhoeven, and Davidson (2014) explore different modes of visualizing data from cinema venues. Besides choropleths, their work includes circos plots and cartograms. Circos plots are a kind of chord diagram which is sometimes also

used for network data (such as Caplan's interactive visualization of the Vilna Troupe, discussed in chapter 5). Circos plots show interrelations between data in a matrix, and that makes them popular choices for network analysis and for indicating movement across geographical locations. Arrowsmith et al. use them to show the influence of different distributors in the movement of Greek films across venues in Australia. They also used cartograms (see above) to display the different number of cinema venues in different Australian provinces from 1948 to 1971.

Quantification is sometimes used for creating choropleths and other geographic visualizations, but the explicit description of geostatistical estimation procedures is not common in DH. A notable exception is Verhoeven and Arrowsmith (2013), who use Markov chain analysis to test whether the distribution of Greek films from different distributors follows statistically discernible pathways through different exhibition venues in Australia. Markov chains are often used for analyzing the probability distributions of a sequence of events, such as the pathways that are of interest to Verhoeven and Arrowsmith. Markov chain analysis is a complex method which is sometimes used for other areas of DH, such as authorship attribution (Khmelev and Tweedie 2001) or social network analysis (Warren et al. 2016). This procedure could also be applied to theater research, for example to analyze the touring circuits of theater companies.

Data-Assisted Location Analysis

The spatial humanities have led to the conceptualization of many theoretical positions that support the usage of GIS and mapping within interpretive epistemologies: "little g" GIS, *deep mapping*, and *thick maps*. There is some overlap between these concepts, but taking them in turn will help clarify the possibilities for the data-assisted analysis of theater locations.

Bodenhamer (2013) sees a major challenge in the clash between the epistemology of cartographic GIS and the interest of historians. Like most humanists, historians are interested in what he terms "extractive scholarship": scholars constantly shift perspectives "in the pursuit of the fullest possible understanding of heritage and culture" (5). This means that, traditionally, narratives are used to construct and present arguments. Narratives enable the "interweaving of evidentiary threads, each of which can be qualified, highlighted or subdued through a variety of literary devices". In contrast, GIS privileges a world that values "authority, definition, and certainty over complexity, ambiguity, multiplicity, and contingency," the very

things that narrative enables and that historians value (7). Bodenhamer's suggestion is to construct "deep maps," GIS artifacts that are interwoven with critical commentary to enable the narrativization and contextualization of spatial data.

Klenotic (2011) has proposed "little g" GIS, as opposed to "big G" GIS, to signal that researchers need not use the full range of technical options available in GIS, but that there is value in "partial, self-taught, bottom-up applications of GIS" (59). Even basic, easy-to-use tools are powerful enough to trigger interesting questions, and visualize data in ways that problematize previous assumptions about spatial knowledge. Klenotic's approach is modeled after Knigge and Cope's (2016) notion of *grounded visualization*, which embraces messy, exploratory, and iterative visualization processes that are amenable to multiple perspectives. Klenotic suggests that applying the notions of grounded visualization to GIS means developing maps and visualizations with piecemeal iterations that require constant reanalysis of the data. Klenotic's view embraces a progressive and recursive ethos of praxis. "Deep" and "thick" maps introduce other theoretical dimensions. Thick mapping is the name of the project by Todd Presner and his collaborators. The name is a direct allusion to Geertz's *thick description*. This approach acknowledges the social construction of maps as unstable, culturally specific objects that "make claims and harbor ideals, hopes, desires, biases, prejudices, and violences" (Presner, Shepard, and Kawano 2014, 15). Considering maps as contingent invites makers of GIS devices to treat them as useful but transient representations that don't correspond neatly to an external reality. This enables an interpretive usage of GIS that tells complex, multilayered stories.

Deep maps also aim to produce multilayered descriptions of places, by integrating data and stories from different sources. Although this term activates similar references to thick mapping, the term reveals a different intellectual heritage. Deep maps show inherent contradictions by reflexively including different kinds of media, in ways that draw inspiration from critical geography (Bodenhamer, Corrigan, and Harris 2013). Deep maps aim to use cartographic and geospatial conventions while also drawing attention to the problematic, positivistic assumptions of GIS and related technologies.

There are some nuanced distinctions among the concepts just reviewed—"little g" GIS emphasizes process, thick maps highlight the constructedness of maps, and deep maps stress the multiplicity of sources—but all of them ultimately provide epistemological frameworks

for using GIS technology within data-assisted theater research. An example is the work of Bench and Elswit (2017), who are interested in using databases and maps to study the international tours of dance companies. They argue that tours are central to studying the global nature of dance, and that this endeavor requires a combination of historical research and digital methods. This combination triggers many fascinating questions and lines of inquiry, some of which I summarize below:

1. How do key choreographic works become canonical, even when other works were almost as commonly performed (for example, in relation to Pavlova's "Dying Swan," 583)?
2. Are aesthetic inclinations reflected in different modes of touring? For example, are the routes and repertoires of American and Russian modernisms substantially distinct from each other (587)?
3. What does the data on performers reveal? Even basic data such as nationality has the potential to become a window into historically contingent classifications. For example, in a 1941 American Caravan Tour, two German-born dancers were classified as stateless (582).
4. How do different geopolitical events affect touring? The most interesting examples are not necessarily big conflicts, but smaller scale wars such as the Ecuadorian-Peruvian war of 1941 which disrupted the travels of the American Ballet Caravan (590).

Bench and Elswit explore these avenues of inquiry and present them together with different cartographies of travel, where tours are represented as point-to-point lines, which are also offered through online visualizations.

On the surface, these visualizations might appear similar to the work of Holledge et al., but I believe that Bench and Elswit are more interested in using maps to find new ways of thinking about touring rather than to answer closed questions. The questions listed above are not entirely answerable in terms of the data, but they could not have been formulated without systematic, careful data collection. Bench and Elswit adamantly highlight the pitfalls of reducing the politics of touring to lines and points in a coordinate system, and seek instead "scholarly modalities" that move between stories and data in "dynamic spatial histories of movement" (575). They argue for scalability, moving between the close and the distant, the single data point and the aggregate. Movement across scales, as

I have argued in this book, is a key marker of data-assisted research. To achieve this movement when working with maps, Bench and Elswit stress that it is important to never use tools for their own sake, but to critically deploy them to rethink questions relevant to dance history. They rightfully note that a choropleth would be ill-suited to their needs, as they are more interested in travel as process rather than in the number of events in a given region. In chapter 3, I have also noted how choropleths can be misleading, by referencing two projects that map lynchings in the United States (Hepworth and Church 2019).

One fascinating perspective for interactive, data-assisted location analysis is to use maps as interfaces for multiform materials to enable alternative modes of scholarly communication. Interactivity is particularly important for bringing depth and thickness to maps. Maps as interfaces enable their makers to chart multiple, contradictory voices. In so doing, they critically reimagine common technologies in modes that transcend conventional mapping tropes. Maps as interfaces are prime examples of the intermedial scholarship I described in chapter 3.

While not specifically a theater project, *The Digital Literary Atlas of Ireland* (Travis and Breen 2017) is an excellent example of an interactive, deep map that deals with the lives of some playwrights. The project includes maps with geographical information for fourteen writers as well as several themed maps: *Emigration*, *House-Island and Provincial Town*, *Dublin Bricolage*, and *Northern Impressions*. The maps are also linked to pages with more detailed information about each of the writers. Although this project has comparatively little data, its strength is the depth and context of its resources, which can be used to examine different aspects of the writers' lifepaths. For example, Travis (2015) shows how the Atlas can be used to explore Beckett's connection to landscape, in ways that are deformative and ergodic (I have discussed both concepts in chapter 3). Designing digital platforms that enable ergodic experiences, and which are expressly defined as *deformances*, is particularly important for work on writers such as Beckett, who challenged many assumptions about narrative in his work. It would be a contradiction to impose a closed realist space as the sole mode to access Beckett's life and work. The Atlas aims to allow "bricolage by interactively juxtaposing different scales of time, space, text and image" (Travis 2015, 223). This approach helps resituate cartographic questions as provisional and subjective, and to reconsider maps as generative objects that trigger new interpretations.

Earlier I described the film venue visualizations of Arrowsmith, Ver-

hoeven, and Davidson (2015) as examples of data-driven work. However, they have also invented the *petal plots*, which address the unique challenges of their datasets and which are closer to the playful transformations required by data-assisted visualization. The dataset that gave rise to these diagrams showed multiple attributes of cinemas in Melbourne. The data included the lifespan of a cinema, transformations in seating capacity, and location. The petal plot is able to represent all these different dimensions at once. Each cinema is visualized as a curved line. Its location along a circular axis indicates its direction from the city center of Melbourne. The beginning and end points indicate the years in which the cinema began and ended operation. The curve of the line shows the change of cinema capacity—a convex curve indicates that a given cinema first increased and then decreased its seating capacity. These diagrams might not be readily applicable to other kinds of data. But the lesson here is that sometimes we should develop visualizations that are specific to our problems rather than merely copying visual tropes from scientific disciplines. This is one of the most effective examples I know of a performative visualization that echoes Drucker’s call for developing new humanistic visual tropes. Petal plots effectively deploy new visual metaphors but they also preserve the precision and readability of statistical graphs. I was directly inspired by their work for developing the *kayon plots* I describe in chapter 6. Davidson and collaborators convincingly show that the nuance of their historical data would have been lost had they employed other kinds of visualizations.

Excursion: The Geographies of Wayang Kulit

Wayang kulit performances (which were described first in chapter 5, and also mentioned in chapter 6) are very common in Java. My Facebook feed and several WhatsApp groups that I belong to are full of announcements of upcoming performances and discussion of recent performance events. It is surprising that such an old form continues to exist to this day and age. Wayang is at least one thousand years old, and a performance is full of allusions to a rich oral literature. Arps (2016) says that part of the pleasure of watching a wayang show is “distinctly philological,” as a wayang aficionado delights in catching references to old performances, or to different variations of the musical repertoire associated with specific artists (28). A traditional wayang show lasts all night (seven to eight hours). And even if people are free to come and go—and most spectators typically leave before the show concludes—wayang invites us to rethink temporal-

ity in ways that are not fully compatible with a modern concept of time, linked to ideas such as efficiency and opportunity costs. Jan Mrázek (2019) quotes his wayang kulit teacher saying that the nights lasted longer in the past (56).

Wayang, in other words, is not for the faint of heart. And it also does not seem particularly compatible with the demands of contemporary, pressed-for-time urban living. Wayang aficionados in Java often long for the past and are anxious about the future, and they often remark, in conversation, that wayang was more common in the past. Anxiety about the future of wayang is perhaps as old as the written record of the form itself. Thow (2018) has traced how *gamelan* music (which is an integral part of wayang) has been said to be disappearing and degrading for at least a century. In these nostalgic expressions, geography plays a role too. Those afflicted with wayang nostalgia say that people in the cities don't have time to enjoy wayang anymore, and that wayang is eminently a rural form (see Mrázek 2019, 98–107 for an excellent treatment of nostalgia and wayang). But is it really a rural form? This is where the data comes in.

Since the middle of 2015, Imam Maskur and his organization have been systematically collecting information on wayang kulit performances all over Java, Indonesia, which is then posted to the website <https://kluban.net> at the beginning of every month. For this he mobilizes a large group of volunteers, managers, and artists who mainly communicate over social media. Their main aim is to let people know which performances are taking place in the month to follow, but at the end of each month they also correct the records, as many performances get canceled and others are only confirmed a few days before the actual event. My research assistant Wejo Seno Nugroho ran a survey over one year in order to verify the accuracy of randomly selected performances and locations. We were surprised to conclude that the website is ~97% accurate and captures ~85% of all performances. However, this is just a preliminary assessment. The performance records are most likely biased toward self-reporting by slightly more famous artists, and very local performances in remote villages might go unreported. Thus, the results that follow should be taken with a pinch of salt. Since the data has only been collected since 2015, it is hard to draw historical comparisons. Previous data were collected in the 1960s and 1970s but they were based on surveys and it is hard to estimate how accurate they were. But even if the historical question is not accessible to us, we can use the data from the website to ask whether wayang is today primarily a rural form.




Fig. 7.1. Scatter plots with regression lines for average number of performances per month compared to population density, land area, and total population per regency.

The distinction between urban and rural areas in Java is not hard and fast, so I have refined the question: are performances more common in less densely populated areas? Truisms spoken by those longing for the past would have us believe so. But the data from Kluban shows that there is no correlation between population density and the number of performances, as shown in figure 7.1. As the unit of analysis, I used data on each regency (*kabupaten* in Indonesian, similar to a municipality). I used the census data from 2016, since this was the year for which data was available for all regencies from official Indonesian sources (Badan Pusat Statistik 2019). Therefore, I also used the data for performances in 2016 for this calculation. As I will show later, the number of performances has remained stable over the years under consideration, so it is reasonable to extrapolate from 2016 data.

As figure 7.1 shows, there is no correlation between the average number of performances in a regency and the regency's size, population density, or total population. So, there must be something else explaining why some regencies have more performances than others. Visualizing the distribution of performances in a map seems to suggest that areas with more performances tend to cluster together. Figure 7.2 is a choropleth of the regencies in Java colored according to the total number of performances. The colors are assigned based on the decile in which each regency is found. A decile is a division of the data into ten bins, where each bin holds an incremental 10 percent of the data. Thus, the regencies with the darkest hue are the ones in the first decile (from the first to the tenth percentile).

The map shows that the places with most performances are located at the peripheries of Central Java, in provinces that border West Java and East Java. As we saw earlier, Moran's *I* is the kind of statistical measure




Fig. 7.2. A choropleth map of Java, with districts colored by average number of performances per month (grouped into deciles).

that can be used to further characterize this type of geographical distributions. Moran's I tells us how much a variable (number of performances) is explained by its "spatial lag," or the number of performances in neighboring regions. Using the built in Moran's I function of the PySAL package (Rey and Anselin 2010) I obtained a Moran's I value of 0.173 with a p -value of 0.016 for my data. This means there is a reasonably strong spatial autocorrelation in the data. Understanding the meaning of Moran's I is clearer when it is presented in visual form. Figure 7.3 shows a scatter plot of each provinces' total number of performances (x axis) against the average value of its neighbors, also known as its spatial lag (y axis). Neighbors are defined here as those regions which border the regency in question. Using a nomenclature borrowed from chess, this is called *queen contiguity* (as opposed to *bishop contiguity*). All values are visualized in a scatterplot. The regression line indicates the correlation between the number of performances and its spatial lag, and Moran's I can be understood as the slope of this line (which is 0.173 in this case).

To calculate the p -value of Moran's I , or the likelihood that this result is merely the product of chance, the most common procedure is a *bootstrapping* method, where random permutations are generated of the same data (in this case 999 permutations were created). Then we can count how many of those permutations have a slope of 0.173 or larger, which in this case is less than 1.6% of all cases (more formally a calculation derived from permutations is called a *pseudo p-value*). If researchers aim to repli-




Fig. 7.3. A scatterplot of each regency's total number of performances (x axis) against the average value of its neighbors, also known as its spatial lag (y axis).

cate this result, they will get slightly different p -values, since the permutations will be different each time. But they should still be reasonably close to the numbers reported here. Figure 7.4 shows the distribution of these results as a shaded curve (a KDE, the same technique described in chapter 6). The dotted line shows the actual observed slope of 0.173. Although not huge, the p -value increases our confidence that the results we are witnessing are not a random effect, but that they reveal the presence of a real underlying phenomenon.

Thus, we can assert that regencies in Java are more likely to have more performances if their neighbors also have more performances. But this is still a very general observation. To delve deeper, we can see the areas where this effect is the most pronounced. For this we can calculate the local Moran's I , which gives us a sense of the "hot areas" where the spatial lag is the most pronounced. I used PySAL for this purpose again and found that the provinces with the highest local Moran's I are: Cilacap, Banyumas, Kebumen, and Ciamis. These regions might not mean much to the reader. And this might be true even if they are scholars of Javanese performing arts, as these are regions which are rarely the focus of academic attention. The bulk of local and foreign scholarship focuses on the Central Javanese court cities of Surakarta and Yogyakarta (where I have




Fig. 7.4. The results of 999 random permutations of the data. The shaded area shows the kernel density estimate of these results, and the vertical dashed line shows the actual observed Moran's I of 0.173.

also done most of my research). But the Kluban data shows that scholars could extend their attention to other places. I have decided, upon looking at these charts, to embark on further ethnographic research in these other regions. I do not yet have enough information to say why there are more performances there. But this data-driven analysis has shown a blind spot of current scholarship. This is similar to what Miller (2017) suggests with his analysis of underrepresented American dramas in theater sourcebooks. Looking at the actual data reveals a trend different from what scholarship alone suggests.

As noted earlier, the data was obtained from <https://kluban.net>. However, the data is not presented there in a machine-readable format, so I had to scrape it using a common Python library (as other procedures, this is described in more detail in appendix A, and the curated dataset is available for download in this book's companion site). I then wrote Python routines to check the data for accuracy and to detect spelling errors. For each performance, the website states the date, the name of the performer (*dhalang*, see chapter 6), and the location. Sometimes the location is very precise and includes a specific address, but this is only true for less than 40 percent of all records, so I had to resort to a slightly less granular category (regency) for comparisons. Other conclusions could be drawn if I

had access to the specific coordinates of the event for each of the 11,255 records. But here, as is often the case when working with data, there is a trade-off between accuracy and granularity. A researcher faces decisions and must aim to communicate them to others when the research is made public. One advantage of using regencies is that there are population statistics available for analysis. The number of performances per month per regency, as well as the population statistics can be easily traced to other sources, and they help move this project in the direction of replicable research (this also means that the conclusions reached here might be later challenged by other researchers).

Other types of information are absent or inconsistent in my wayang dataset. For example, most records don't state the reason for a performance (marriage, institutional celebration, village cleansing, etc.). This is something I would wish to know more about, but my dataset alone won't shed light into this issue. The data does include the name of the dhalangs, but this also carries problems. Many dhalangs have similar names, and name disambiguation is very tricky except in the case of the most famous performers. I identified between 4,000 and 5,000 possible distinct names in the dataset, but further disambiguation would require other kinds of research, beyond purely computational techniques as I would most certainly need to interview people directly. Thus, I have avoided looking at questions about who performs where, tantalizing as they are. But there is another variable in the data which is accurate and precise: the dates.

When I visualized the dates as a time series, I was struck by the regularity of their patterns. Figure 7.5 shows the total number of performances per month for 2016, 2017, and 2018. There are less performances in the fasting month of Ramadan (since only ritual performances are allowed then). This lunar month moves from one year to the next, but in the data graphed in the figure it fell in May and June, as can be seen from the dip in the performance numbers. The peaks are connected to the anniversary of the Indonesian Independence (August 17), the first month in the Islamic New Year (which also moves, but fell around September and October in the years when the data was collected) and the anniversary of the inscription of wayang into UNESCO's representative list of intangible heritage (the date itself is November 7, but the preceding week usually sees a surge in performances). This data can also be combined with geographical information. Even though the total number of monthly performances is constant over the years, some regencies see more performances than others in specific months. To visualize this, I created a choropleth map




Fig. 7.5. Performance totals per month across three years: 2016, 2017, and 2018.

for each month, and then joined the results into an animation that can be consulted in the book’s companion website (video 7.1).

This data-assisted animation, more performative and evocative than the static choropleth, shows Java as a “beating heart,” as the highest concentration of performances flows back and forth between the court cities of central Java and the eastern regencies (the ones in the local Moran hotspot mentioned earlier). This analysis of wayang kulit performance data is highly specific, but it shows a way of working with large amounts of location data, and how it can be analyzed through different kinds of visualizations and statistics. Similar research is certainly possible for many other theatrical forms around the world and I think it is likely that an explosion in this type of research will be seen in the near future.

Further Considerations: Time and Place

In this chapter I have spoken mostly about geospatial rather than geotemporal data, even if time features into many of the projects described earlier. All the pioneering projects in theater studies mentioned above accounted for time in one way or another. Bollen and Holledge (2011) superimposed networks into their cartographic displays to show the temporal sequence

in which performances followed one another. Bench and Elswit (2017) represented tours as point to point lines that can be animated in the online versions of their maps, to show the order in which the tours were carried out. This consistent attention to time is certainly linked to the nature of theater, and should attune us to the possibilities for geotemporal analysis in the future. Projects in other areas of DH can also provide inspiration. *Starting Out from 23.5°N* (Academia Sinica Center for Digital Cultures n.d.) explores the life and work of Taiwanese artist Chen Cheng-Po (陳澄波, 1895–1947) as he moved between China, Taiwan, and Japan in the first half of the twentieth century. This project combines maps-as-interfaces, interpretive resources, controlled-vocabularies, and interactive timelines to explore the life of an influential figure through deep maps. This and many other GIS projects have been affected by a change in the terms and conditions of the Google Maps API which took effect on June 11, 2018 (Google Developers n.d.), a time when I was doing research for the present book. This should raise the alarm bells of depending on commercial providers for critical DH architecture, a theme further explored in the final part of this book. There are other, more robust platforms that offer some degree of independence from commercial providers, such as Neatline (Nowwiskie et al. 2013). It is also possible to build custom interfaces with Leaflet (Open Source Community [2010] 2018) using only data from OpenStreetMap (OpenStreetMap Community 2004).

In this book, I have argued that data-driven and data-assisted methodologies can coexist to address different aspects of the same project. In the computational analysis of location data, we can seek statistical rigor to answer discrete questions, which can in turn be expanded and contextualized in complex narrative displays that don't take mapping tropes for granted, but that examine them as culturally situated objects. Location data can be prodded and deformed, in ways that situate observers front and center and which guide nuanced interpretive perspectives. But this data can, in turn, be used to test patterns. We can readily imagine complex hermeneutic circles, where the results of statistical patterns guide further situated analysis, and where *deformances* suggest new empirical investigations. These possibilities are contingent on the potential of our work to be accessible and reusable in the future. As hinted in the discussion above in relation to Google Maps, we need to pay attention not only to epistemological possibilities but also to the technical and institutional conditions that can limit or enable the future of computational theater research. This is the subject of the next, and last, part of this book.

Code and data: Sample code for this chapter is available at

<https://doi.org/10.3998/mpub.11667458.cmp.42>.

The data used can be downloaded from

<https://doi.org/10.3998/mpub.11667458.cmp.36>,

<https://doi.org/10.3998/mpub.11667458.cmp.37>, and

<https://doi.org/10.3998/mpub.11667458.cmp.38>.

PART 3

Ensuring the Journeys Continue

The Imperative of Open and Sustainable Data

The Australian Performing Arts Database (AusStage) is one of the most ambitious performing arts databases ever undertaken. It aims to capture all live theater performances for the entire history of Australia, as well as performances by Australian artists around the world (AusStage 2013). The project is remarkable also for its clear and robust data model and the fact that the entire dataset is openly available. Researchers can download and reuse the data as long as they provide proper acknowledgment. The progress of data-driven and data-assisted theater research will depend on the well-planned openness of projects such as this, and it will be underwritten by institutional infrastructures that can keep these kinds of projects available for future generations of scholars.

The Comédie Française Registers Project is also extraordinary in terms of its scope and openness, as it enables users to download and reuse data on the performances of the Comédie Française from 1680 until 1791 (Biet et al. 2015). There are many other carefully planned database projects around the world, such as IbsenStage (modeled explicitly in the footsteps of AusStage), Records of Early English Drama (Black et al. 2017), 19th Century Acts (Gonzalez et al. n.d.), the Digital Yiddish Theater Project (Baker et al. 2019), *Reseña Histórica del Teatro en México 2.0–2.1* (Historical Theater Reviews in Mexico 2.0–2.1) (Franco 2020), *Base de Datos de Comedias Mencionadas en la Documentación Teatral 1540–1700* (Database of plays mentioned in theater records: 1540–1700), (Ferrer Valls 2019), and the Cuban Theater Digital Archive (Manzor, Rimkus, and Ogihara 2013). Some projects integrate textual sources with other media, such as the Map of Early Modern London (Jenstad 2011), which combines maps and textual annotation. Other projects, such as the Hemispheric Institute's (2008) Digital Video Library, the Asian Shakespeare Intercultural Archive (Yong et al. 2015), and the Digital Dance

Archives (Fensham 2016) contain extensive video material, interactive visualizations, and copious scholarly annotations.

In many of these projects, data is bound up in other kinds of artifacts, as it is used for visualizations or it constitutes the backbone of archival projects (and it is not always readily available for download). Different projects also vary in their policies for sustainability and reusability, and in how clearly they display such information. A comprehensive survey of such policies is available elsewhere (Escobar Varela and Lee 2018). Projects have good reasons for sharing data in different ways. In some cases, this is limited by copyright restrictions or other institutional constraints. Sometimes sustainability policies are not necessarily explicitly stated. This is understandable, as many projects are new and struggling to find financial and technical support to ensure their subsistence. But it poses serious problems for the continued existence of computational theater research.

The bulk of this book deals with methodology and epistemology. Yet, the implementation of methods depends on material and institutional foundations, and it is to these that I now turn my attention to. I will consider three aspects of data: modeling, sharing, and preservation. I will dedicate some attention to each of them in isolation, but my larger aim is to show how these aspects are deeply interconnected. The way data is modeled limits and constraints how it is to be shared and preserved. Sharing without preservation is meaningless. And preserving things that won't be shared (at least eventually) would be preposterous.

Data Models and Metadata

Imagine an assiduous theater goer that keeps a notebook with reflections on every theater show she watches. The notes contain some factual information—name of the performers, time, maybe even ticket price—but these are interspersed with her thoughts on the performances. Now let's imagine that she wanted, after several years of note-taking, to turn this notebook into a digital database. For this, she would need to explicitly formalize the kinds of things she is capturing in her notebook. She thus designs a spreadsheet with the following column names: creative team, title, venue, and comments. Each performance's data will be entered in a separate row. This formalization is a *data model*, albeit a relatively simple one. To see where it falls short, let's imagine that she wanted to compare her data with someone else's. This other person included two categories—performers and director—to describe the information that the first the-

atergoer included within the single category of “creative team.” If they wanted to make their data more directly comparable, both people could settle on a shared model. But let’s say that both projects are too far along their development and that standardization is impractical. Each could continue with their own system, but establish ways to map one onto the other. For example, the theatergoers could note down that “creative team” in system A includes “director” and “performers” in system B, and that they are all instances of a larger category of data called “people,” which have properties such as dates of birth and names. In technical parlance, this means that the two systems would be *interoperable*. The systems will remain distinct but formal rules would map one to the other. These rules can also be said to be part of a more general data model.

In both situations—standardization and interoperability—one could decide to use an existing data model or to develop a new one. The latter strategy would add more nuance, but would require time and effort. More importantly, it would need to be adopted by other people who might not be easily persuaded to change to a bespoke, little-used system that is not theirs. Using an existing data model reduces cost and effort, but necessarily forgoes some nuance and granularity. In some of the excursions in this book, I have described Javanese wayang kulit performances. The dhalang is the creative leader of a wayang show, and he or she speaks all character parts, animates the puppets and cues in the musicians. One could describe the dhalang as a director or a puppeteer, but I think that both terms fail to capture the actual role of a dhalang in a wayang kulit show. To describe a wayang performance, I could use a theater specific data model, or I could just refer to a more general model such as Dublin Core (DC), a widely used, small set of vocabulary terms that can be deployed to describe resources. The core element set consists of 15 terms (description, format, identifier, language, publisher, relation, rights, source, subject, title, and type), and several dozen properties, classes, datatypes, and vocabulary encoding schemes. Using DC, I could specify the role of a dhalang as “creator.” This is indeed what I did in the Contemporary Wayang Archive (CWA). Table 8.1 shows an example of a wayang kulit performance described according to the DC elements in the CWA.

The elements modeled by DC are more formally called metadata elements. Metadata is data about data. In the introduction, I used Pomerantz’s definition of data as a potentially informative object. For him, metadata is “a statement about a potentially informative object” (Pomerantz 2015). A good data model includes some metadata. Without it, a model

would be much less useful. Continuing the example above, let's consider an entry in our imaginary theatergoer's diary for Robert Lepage's *Far Side of the Moon* (2002). The director's name (Robert Lepage) is a data point. But it does not belong to the same conceptual category as *Far Side of the Moon*, even though both are text strings. A human perusing the database (at least one familiar with conventions of play titles and Francophone naming conventions) would know which is the director and which is the title. But a machine would be unable to do this without some formal mechanism. Metadata provides such a mechanism, as it describes which data point belongs to which category. For these descriptions to be maximally useful, the relationships between different types of data statements need to be explicitly defined. One option is to use a formal language such as the Resource Description Framework (RDF). In the RDF, relationships between data and data categories are made through subject-predicate-object statements, known as triples. These statements indicate relationships between different data elements. An example of the previous record for Enthus Susmono's *Dewa Ruci* (2008) that uses RDF conventions is available in XML format at <http://cwa-web.org/en/metadata/DewaRuci.xml>. As seen in chapter 4, XML is the same language used to encode TEI files. The excerpt below gives a sense of how it is structured:

```
<rdf:RDF xmlns:rdf="http://www.w3.org/1999/02/22-
  rdf-syntax-ns#" xmlns:dc="http://purl.org/dc/
  elements/1.1/">
<rdf:Description rdf:about="http://dublincore.org/">
  <dc:title>Dewa Ruci</dc:title>
  <dc:identifier>http://cwa-web.org/en/DewaRuci</
  dc:identifier>
  <dc:date>2008</dc:date>
  <dc:publisher>Contemporary Wayang Archive</
  dc:publisher>
  <dc:creator>Enthus Susmono</dc:creator>
  <dc:language>id</dc:language>
</rdf:Description>
</rdf:RDF>
```

Metadata operates at different levels of granularity. In the CWA, the metadata describes a performance. But the performance itself is obviously not part of the archive, which includes only a video recording of the per-

Table 8.1. A DC metadata record for Enthus Susmono's *Dewa Ruci* (2008)

<i>Title</i>	<i>Dewa Ruci</i>
<i>Identifier</i>	http://cwa-web.org/en/DewaRuci
<i>Data</i>	2008
<i>Publisher</i>	Contemporary Wayang Archive
<i>Creator</i>	Enthus Susmono
<i>Language</i>	ID
<i>Description</i>	In this performance, Bima is on a spiritual quest to find the meaning of life. His teacher Durna tries to trick him by telling him the answer will be found at the bottom of the ocean. Bima dutifully follows, defeats a dragon that lives in the ocean and finds a miniature version of himself, Dewa Ruci, from whom he receives a lesson on the spiritual meaning of life.

formance. A more complex metadata standard could have been used, such as the CIDOC-CRM (the Conceptual Reference Model of the *Comité International pour la Documentation* [International Council for Documentation]). The CIDOC-CRM is often used to capture information on intangible cultural heritage, and has been proposed for documenting theater performances (Pendón Martínez and Bueno de la Fuente 2017). The CIDOC-CRM could be used to describe multiple instances of a performance and to indicate how multiple video recordings refer to them. But even this is not extremely detailed. We could imagine using an even more granular data model where the content of each frame of each video is described in detail. But there would need to be a category for each thing in the video. One can easily imagine how quickly this would run out of control—and Borges' (1946) short story “Del Rigor en la Ciencia” (On Exactitude in Science) comes to mind. The characters in this story are so obsessed with creating a map that fully represents reality—every bird, every leaf on a tree—that the resulting map ends up being so comprehensive that it overlaps with reality, and people don't know any more whether they are living in the map or in the reality it seeks to represent.

The limitations of any metadata model for the humanities are obvious. Calling the dhalang a creator misses the cultural-specific aspects of a dhalang's role in performance, as he or she is someone who is often reinterpreting oral traditions rather than making entirely new works. Here I am advocating for using metadata models for the sake of sustainability, but it is important to note that there are serious limitations with these systems. Brown and Simpson (2013) note that standards limit the ability to make

the nuanced statements that characterize humanities research. They use the example of a “writer” (Michael Field), who is not an actual person but the pen name used by the authorial collaboration between two women in the Victorian era (Katherine Harris Bradley and Edith Emma Cooper). The nuances of their collaborations can’t be captured by existing data models and, as Brown and Simpson note, this is not an extreme case, but a common example of the kind of careful attention to specificity that the humanities require.

When choosing data models and metadata standards, there is always a tradeoff. Flanders and Jannidis (2015) suggest different strategies for dealing with this tradeoff. When it comes to the purpose of a model, they distinguish between curation-driven and research-driven data modelers. The first group wants to achieve models that are as generalizable as possible, finding the most widely applicable common ground. The second group is interested in formalizing very specific research ideas for specific purposes and narrow domains. The CWA follows a curation-driven approach. My objective in relying on a widely used, though less granular standard, is to make it easy for people to find and cite these resources, as the DC records can be easily ingested by library systems. But I might one day want to develop a research-driven model, one that is aimed specifically at video analysis. I might then wish to manually annotate the video recordings and describe the actions that take place in the videos. The same action can be described in a wide variety of ways. Whichever model I end up developing and using will betray specific prejudices and preferences and will reflect my interpretation of what matters in a performance. As Jannidis and Flanders (2015) note, most researchers in DH know that models are social constructs rather than representations of objective reality, but this doesn’t mean all models are equally good. Thus, Jannidis and Flanders suggest that models can be assessed in terms of persuasiveness, intellectual elegance, or strategic value. Unlike natural objects, digital artifacts “are created with a purpose by identifiable agents and they have a history which is part of their identity” (235). We thus need to represent not only the history of the artifact, but also the history of the *ways* in which it has been described and contextualized—and there is always some degree of uncertainty in these histories.

Bollen (2017), who has been instrumental in the development of AusStage, analyzed twelve influential, large-scale theater databases and found that they tend to coalesce around five aspects: places, people, companies, performances, and works. Models that include these aspects are

fittingly described as persuasive, intellectually elegant, and strategic, using the terms proposed by Jannidis and Flanders. But Bollen is also quick to highlight cases where the proposed categories are too coarse or too value-laden to describe border cases, such as touring productions. Large-scale projects would do well to adopt the model proposed by Bollen and to pay attention to cases that test the limitations of such model. But what about smaller datasets? In many cases, we might not want to add metadata to each individual data point. For example, in my analysis of the geographic distribution of performances in Java described in chapter 7, I scrapped data from <http://www.kluban.net> and then put it into tables. The tables are consistently formatted but the data model is not formalized into a standard such as RDF, and I don't provide explicit metadata for each data point.

When gathering this data, I had very specific questions I wanted to answer, and these didn't require granular metadata or formal statements. I am interested in sharing this data with other people who might want to use them for their own analysis and I imagine their objectives would be one of the following: to verify my results, to use my data as an example in their teaching, to ask other questions of my data, or to combine it with other datasets. In each case, they will likely apply transformations to the data. So, rather than serve it via a semantic server which is hard to maintain, I will just offer the data for download in its entirety, in a simple and sustainable CSV format. I will add a description of how I obtained and transformed the data, but this will be a *data biography*, rather than a more formal, machine-actionable description. The data biography is a concept I borrow from Simon Eliot's (2002, 289) "biography of a data source," which I came to know by way of Bode (2012, 14). I will add metadata that is less granular, and which describes the entire dataset as a single object to facilitate citation, and the data will be maintained in this book publisher's website. This is the strategy followed by most recent DH books (Piper 2018; Eve 2019; Underwood 2019a; Mullaney et al. 2019).

Projects such as AusStage have many times more data than the example above, and their objective is to support a more varied range of research agendas. Thus, it makes sense for them to structure their data according to more formal data models. There is no one-size fits all solution. Perhaps the distinction introduced by Jannidis and Flanders, between research-driven and curation-driven data modelers, is best thought as a spectrum of possibilities. Each data team needs to think of what type of data model is practicable, and which best suits its needs. I agree with Toni Sant's (2014)

admonition that “performance scholars need to make use of best practices in information science more consistently,” but we also need to do what is technically and institutionally feasible. The ways in which data are to be shared are perhaps the most important considerations when choosing a data model, and it is to these that I now turn to.

Sharing, Reusing, Citing

Many projects require the creation of their own data. But a vibrant computational research environment is one where not everyone needs to produce all of their data, and research teams can reuse previously existing datasets. This is what literary studies, probably the most successful area of DH, have achieved through a shared model for literary data. This model is premised on the guidelines of the TEI that I discussed in chapter 4. These guidelines enable researchers to add structural markup to digital texts. For example, theater texts can be encoded in ways that specify the names of speakers, the content of their speeches, and the division of plays according to scenes and acts. This is the format that Trilcke and his collaborators have used to study the networks of German dramatic texts (as seen in chapter 5). Based on available TEI data, they could easily extract information on the interactions between speakers in a scene. They didn’t have to manually encode texts, as they were able to reuse extensive datasets that use the same format. The TEI format also includes metadata on each text, such as the author and year of publication. This can be easily extracted from the digitized texts and used to make comparisons across time. However, even standardized data needs to be transformed to suit specific research agendas. Even when their data was TEI-conformant, Trilcke and his collaborators spent a considerable amount of time cleaning their data (Trilcke et al. 2015). But their project would have been all but impossible without an open repository of texts in reasonably standard formats that they could tap into. The way in which network researchers use the TEI format to study theatrical texts is only one example of what this shared format enables. The widely used *Stylo* in R package (Eder, Rybicki, and Kestemont 2016) and *Voyant* can also ingest TEI-conformant texts.

Sharing data will be important for the future of computational theater research for several reasons: to allow others to verify our results, to enable other researchers to combine our data with their own datasets and ask new questions, and, equally important, for use in training courses. In ideal conditions, systems should be in place to enable other people

to always access all the data. There need to be clear policies about how the data is to be reused and the data must be described in ways that other people can understand. Lastly, the data must conform to formats that others can understand and reuse. The common names for these principles are: findability, accessibility, interoperability, and reusability, also known as the FAIR guiding principles for data management and stewardship (Wilkinson et al. 2016).

Christine Borgman (2009) wrote an often-cited “call to action” based on her keynote at the 2009 Digital Humanities conference. In it, she encouraged digital humanists to think critically about their data practices so that a robust infrastructure for data scholarship can be developed. Science, she says, offers inspiration—but it is important to design a scholarly information infrastructure that caters specifically to humanities concerns. For her, this infrastructure encompasses technology, services, practices, and policies. She identifies six factors for comparison between the humanities and the sciences, “selected for their implications for the future of digital scholarship in the humanities”: publication practices, data, research methods, collaboration, incentives, and learning (n.p.). Of particular relevance to the present discussion is her analysis of the reasons why people don’t share data: “(1) faculty get more rewards for publishing papers and books than for releasing data; (2) the effort of individuals to document their data for use by others is much greater than the effort required to document them only for use by themselves and their research team; (3) data and sources offer a competitive advantage and are essential to establishing the priority of claims; and (4) data are often viewed as one’s own intellectual property to be controlled.” The problem of copyright is indeed very complex in the humanities as researchers typically do not own the data they work with. But in some cases, research teams can make provisions to share some data in ways that don’t violate intellectual property laws.

A 2016 report on the state of Open Access in DH for the EU Digital Research Infrastructure for the Arts and Humanities (DARIAH) follows along similar lines and identifies several ways in which humanists can develop an ecosystem for data citation (Buddenbohm et al. 2016, henceforth “OA Report”). In their view, the data management plan (DMP), should be viewed as a key instrument for ensuring that data can be shared. The DMP is a formal document, proposed by the UK’s Digital Curation Center (Rusbridge et al. 2005), where a researcher describes how the data is to be used and shared and what metadata will be added (an example is

available at <https://dmponline.dcc.ac.uk/>). The DMP should be a “living” document where changes are made through time. The authors of the OA Report also cite eight principles for encouraging data reuse from the Joint Declaration of Data Citation Principles issued by the Data Citation Synthesis Group in 2014. These principles are: importance, credit and attribution (whenever claims are made on data, the data should be cited); unique identification; persistence; specificity and variability; interoperability; and flexibility. It is important that access to the datasets themselves (in CSV or another simple, machine-readable format) is always possible.

The OA Report highlights the key role that data citation might play in a scholarly ecosystem where data is openly shared and identify several things that prevent citation from happening. Some of these barriers are data-related, such as the diversity of formats, fragmentation, and missing organization. Other hindrances are technical, such as data that is not clearly separated from other digital artefacts and is, for example, interspersed with other aspects of a PDF file. There are also legal barriers such as privacy issues and copyright. A final category of barriers include things such as funding and its impact on sustainability, the difficulty in ascertaining provenance, as well as problems related to versioning and granularity. The authors of the OA Report identify the European Holocaust Research Infrastructure (<https://www.ehri-project.eu/>) as an excellent case study and mention several data repositories where individual researchers can deposit humanities data: Re3Data.org, Harvard Dataverse, DataDryad.org, Zenodo, Figshare, and Mendeley. At the time of writing, many teams working on computational humanities projects are also depositing their code and data at the Open Science Framework (<https://osf.io/>). The right incentives must be in place for people to share and cite data, and this must contribute to the prestige of academic careers if this is to be a sustainable model. Theater journals, conferences, and university departments have an enormous role to play in implementing a culture where these principles can guide academic practice.

Should All Digital Projects Be Preserved? A Multitiered Approach to Digital Sustainability

Many interactive visualizations require specific infrastructures to stay alive, and this is particularly difficult for the case of intermedial essays (see chapter 3). For both the preservation of data and the preservation of intermedial essays, a multitiered approach will come in handy. Sustain-

ability seems like a reasonable aim for all data and interfaces. But perhaps not every aspect of every single project should—or can—be maintained. While long-term preservation is a major challenge, it is also a choice. Not all projects might aim to endure for centuries and there is value in what I will call *bloom-and-fade* experiments in digital publishing, which can uncover new modes of scholarly expression without having to worry about the technical infrastructure required for long-term preservation. Performance’s essence has often been theorized in relation to its ephemerality—performance studies is thus uniquely positioned to conceptualize the contribution of digital projects that are only meant to exist temporarily. However, if the aim of a project is to become a data archive, then its creators must devise systems for supporting its continued existence into future centuries, as books and libraries have succeeded in doing. The designers of digital projects might opt for a multitiered approach, where certain digital objects are marked for long-term preservation and others are developed solely for their more immediate, ephemeral value. This section enables the strategic construction of this multitiered approach by bringing together conceptual insights from performance studies and best practices from digital curation and information science. I will place special emphasis on how projects with limited financial resources can achieve such goals.

In another short story, “Funes el memorioso” (Funes the memorious), Borges (1942) imagined a man who could remember everything. It soon transpires that this is a curse rather than a blessing. Being unable to forget things, Borges’ character was also unable to make abstractions and to isolate details from their contexts. To think, Borges concludes, is to forget. One of the challenges of digital scholarship today is that we need to constantly consider the preservation of our materials. But this is not entirely a new problem, as the long-term preservation of the cultural record has always been central to any endeavor in the humanities. The humanities could be defined as the intergenerational custodianship and commentary of the cultural record. This is true for areas as different from each other as architecture and literature. What to preserve, and how to preserve things are open questions. But the centrality of preservation is an assumption that underpins all endeavors of the humanities. Even scholars who analyze present phenomena do so in relation to the past. It is important to note that there is a cultural dimension to preservation. Not all intellectual traditions are interested in preserving things in the same way. There are different ways of keeping the past alive, as Western societies place a pre-

mium on the preservation of artifacts which are ideally kept in conditions that work against the passage of time (temperature-controlled, hermetic environments, for example). Museums and archives are very important in this context, and the intellectual traditions of the West use these artifacts to reconstruct specific histories. Questions about how and when things came to be the way they are today are considered central within these intellectual traditions. Indonesia, to give a contrasting example, is a place where keeping oral traditions alive is perhaps more important than the preservation of physical artifacts and the precise adjudication of provenance. Both attitudes, however, signal a preoccupation with the intergenerational preservation of cultural memory.

Memory institutions devoted to preservation have evolved over many centuries and many generations. As records become increasingly digital, the previous institutional structures for preservation need major overhauls. Within a paper-based model of scholarly publishing, scholars don't need to spend much time thinking about how their materials will be preserved. Or about which materials will be preserved. Well-developed infrastructures and procedures are in place. Journals, publishers, and libraries employ highly trained professionals and dedicate substantial financial resources to these purposes. Digital preservation still works better when preserving digital facsimiles of print materials. But interactive data visualizations and intermedial essays require interactive and multimedia platforms. Since there is no true-and-tested method of digital preservation for interactive online scholarship, many scholars have set up their own archives and portals. This is a key area of DH experimentation and, in many cases, initiatives are not driven by institutions but by individual researchers who must spend significant time and effort thinking about how to preserve their materials. One challenge is defining what to preserve—a significant problem in an era where generating data is much easier than storing it in sustainable formats. As Borges' short story suggests, perfect memory is a hindrance to thought. It is also impossible from a practical point of view. Bloom-and-fade projects enable a nimbler, more adaptive development strategy that eventually settles on the data, the data models, and the visualizations and interfaces worth keeping for posterity.

Modifying digital media is almost too easy. Its inherent instability is linked to a series of principles identified by Lev Manovich (2000): numerical representation, modularity, automation, variability, and transcoding. These are well known, so I will only describe them briefly to show how they make digital media easy to change and hard to preserve. Numerical rep-

resentation is the principle from which the others stem. In current computer architectures everything is represented by binary numbers (images, texts, videos, etc.) and it is very easy to apply mathematical transformations to these numbers. Any operation on a file (search, delete, change color) consists of a series of mathematical transformation applied to these series of numbers. The second characteristic, modularity, means that all new media files and systems are made of independent entities that can be easily modified. An image is made of many pixels and one can change the color of a single pixel without affecting the others. Likewise, a web page is made of many modular components, such as texts and images. It is very easy to change an image without affecting the integrity of the whole webpage. This is obviously impossible to do with print materials.

Automation, the third principle, means that any mathematical operation can be repeated in the same way with no intervention of the user. Changing thousands of images or reverting them back to their previous state is almost effortless. The fourth principle is variability. Digital media is so easy to change and replace, that it is difficult for media creators to decide when a product has reached a finished stage and digital files can “exist in different, potentially infinite versions” (Manovich 2000, 36). Transcoding, the last principle, is the translation of media from one format to another. It is easy to convert an image file to a lower resolution, or to export a video as a series of still images. Transcoding and variability mean that files can exist in an endless state of flux. These two last characteristics are of course underpinned by the numerical representation and modularity of digital files, as well as by the ease with which transformations can be automated. Manovich argues that the enormous creative potential of digital media can be explained by these five principles. But they also account for the difficulty of preserving digital media. The very instability that is so enticing for creative purposes is what makes preservation hard.

In addition to these characteristics of digital media, certain institutional aspects of DH work contribute to constant change (Escobar Varela 2016):

Funding: Funding for research projects around the world tends to focus on short cycles (3–5 years). Often, priority is given to new projects. For projects that are new versions of older ones, significant improvements need to be made. Usually funding is not given for keeping projects alive, but to revamp design features. These prob-

lems, and a series of potential solutions, are further explored in the Endings Project (The Endings Project Team 2019).

Teams change: Often the people responsible for a project change, bringing with them new expertise and ideas for new features in a project

Practice-based approach: Digital projects often begin without a clear sense of the final product, which means that teams often must feel their way forward, experimenting as they go along.

Constant user feedback: Obtaining user feedback on software products and web portals is not difficult. This is both a blessing and a curse. Feedback certainly makes products better, but it provides an endless supply of new ideas. Preservation requires keeping a product in a finished state for posterity and the essential variability of new media, coupled with the best practices of iterative design—which keep new users in mind—conspires against this goal.

This situation is compounded by the physical constraints of storage media (Bollacker 2010). As a rule of thumb, data written into smaller physical entities is easier to rewrite but it has a shorter lifespan. For example, it is easier to amend a message saved on a solid state drive than to change a message carved out in a stone surface. The stone does not lend itself easily to continuous rewriting, which makes each successive inscription cumbersome. The upside, though, is that stone inscriptions from previous centuries have endured to our day, often in harsh environmental conditions. Solid state drives can hold the same information as thousands of stone inscriptions, but they won't survive for more than a dozen years. Digital data requires constant migration to newer storage media, and managing this imposes manpower constraints.

To complicate matters further, it is not sufficient to maintain physical storage media. To read files from the past, it is crucial that we can run the software that can process those files, and this is not always a straightforward process. In other words, preservation won't happen automatically. People will need to design and enforce action plans to preserve digital data and infrastructure. Molloy (2014) suggests that theater makers are insufficiently familiar with the needs of digital data preservation and don't usually have a preservation plan in place. Initiatives such as the Long Now Foundation are trying to develop best practices (Bollacker 2010). I have thus far argued that it is impossible to save everything—constant change is easy and preservation is hard. But preservation is not always desirable, since it requires enormous costs in terms of money, infrastructure and

effort. As mentioned above, there is strategic value in doing faster, nimbler projects. I suggested that teams can better identify what their long-term goals are if they first work on bloom-and-fade projects that are not meant to last for a long time. When they have done so, they will be in a better position to preserve the key aspects of their work in the long term. In addition to this, computational research teams can implement the following suggestions:

Experiment widely: One of the best things that a DH team can do is carry out little experiments to figure out what they want to achieve in terms of preservation. Bollacker (2010) suggests that since no one knows what will be useful in the future, it is important to experiment with as many formats and combinations of formats as budget and conditions allow for:

Because we also can't predict the future to know the best data-representation choices, we try to do as nature does. We can copy our digital data into as many different media, formats, and encodings as possible and hope that some survive. This is our best shot at ensuring that at least some formats and some things will still be available in the future. (110)

Be informed: It is sometimes hard to fathom how volatile and fragile current storage media and infrastructures are. I'm always surprised when speaking to people who overestimate the durability of platforms, formats, and physical storage media. There are many resources that people can consult to understand better how media works. All the documents from the Long Now Foundation are useful reading for anyone interested in DH and digital preservation. *Minimal computing* also applies here (a concept I will revisit in chapter 9). It is important to choose minimal architectures since they tend to be easier to maintain, and more compliant with existing standards.

Arrange things into baskets: Based on small-scale experiments and on awareness of the technical and social constraints of long-term preservation, research teams can make decisions about what aspects of their projects should be marked for long-term preservation and which are meant to exist as relatively ephemeral digital projects. There are some aspects of a project that are essential and that must

be preserved even if everything else is lost. I think of this as the “preservation core” of a project. Having grown up in an earthquake-prone city, I learned to keep key documents in a sealed plastic bag by the door of my house. In case of an emergency, I could take these documents and run. The same principle applies to the design of a digital preservation core. A second basket will include digital artifacts that can be kept in case we have a bit more time to run away before the building (or the funding) comes crumbling down. This is similar to what the Digital Curation Centre (DCC) recommends (Higgins 2008). We often don’t want to believe that disaster will strike and when it does, we are caught off guard. I know of one project where the digital architecture collapsed overnight, when there was a mandatory update from the base system they were using. Since there was no multitiered approach to preservation it was very difficult for the team to recover functionality quickly. More importantly, this imperils the long-term preservation of this project. Data marked for long-term preservation should be stored according to explicit data models and ideally made openly accessible. Both preservation and openness depend on each other. Open sharing requires sustainability practices. But in turn, resources that are openly shared are more likely to be preserved. This is one of the core principles of LOCKSS or “Lots of Copies Keeps Stuff Safe” (Maniatis et al. 2005).

When planning projects, we also need to plan for our absence or retirement. This is of course hard for academics to do (as it is hard for everyone) but it is a must if digital projects are to endure into the future. For example, the key data of the CWA are the performance recordings and the translations and notes. They constitute the innermost preservation basket, or digital preservation core. For this reason, they are kept in the most standard possible files. The translations, transcripts and notes are stored in vanilla text files (UTF-8 encoding) and the videos in raw formats. On the next basket we have the documentation and source code on GitHub. Given that interactive websites are hard to preserve, it is important to take video snapshots of interactive features for future reference. This is what I have done with the interactive visualizations described in this book. The website companion includes videos that show how these visualizations are meant to work, in case the online versions of those visualizations stop working in the future due to any of the problems described earlier in this chapter.

This last part of this chapter might read as a prescriptive instruction manual, but its message underpins the entire project of *Theater as Data*. Earlier in the book, I discussed mostly epistemological and methodological concerns. We cannot, however, think about epistemology without paying attention to the material conditions on which data survives or perishes. Matters of digital preservation are more than just technical issues. They point to the essence of the humanities. If we want to have viable projects, careful attention to the conditions that encourage and limit preservation is essential. On a practical level, the issues discussed here will help us think about how to strategically construct a multitiered approach that values both ephemerality and long-term preservation. The double focus on ephemerality and long-term cultural transmission is, after all, a defining feature of theater.

The Roles of Software Programming

Programming is more than a technical skill, it is a mode of looking at the world. Engaging with programming has pedagogical, ethical, and institutional implications for computational theater research. In this short, concluding chapter, I will untangle some of these implications by suggesting that programming can best be thought as a range of creative and critical endeavors. I believe that programming will be useful to anyone attempting to work with theater data, whether they are interested in the data-driven quest to answer specific questions, or in the provisional and situated polyvocality of data-assisted research. My objective is not to say that those who are not programmers should refrain from doing computational research. Given the range of available software and online tools, many people can carry out computational theater research without knowing how to program. People who are not programmers have and will continue to make fundamental contributions to computational theater research. My objective in this chapter is not to defend programming knowledge as a precondition for working with data. Rather, I want to show how programming can help us think about theater in new ways, and to encourage more people to tinker with the art of programming.

Learning to program is not a binary choice. Programming skills are located along a continuum, and it is difficult to ascertain a single moment where someone becomes a programmer. But embarking on the road towards the acquisition of programming skills will make our work richer and deeper as it will expand our potential to decide how data is collected, processed, and presented. Both data-driven and data-assisted research can be carried out by deploying existing software, by tweaking existing packages, or by writing bespoke code. An interesting feature of online platforms, such as Voyant, is that they encourage playfulness and reward their users with the aesthetic joy of seeing texts transformed into

tables and visualizations which can then be tinkered with. A more direct version of this experience is available to users who write their own code. Several DH books take readers through step-by-step instructions of how to write code for text analysis (for example, Jockers 2014). An intermediate step between using platforms such as Voyant and writing code directly can be found in certain programming packages that have a graphical user interface (GUI) which can be run directly by novice users but also modified by users with more programming experience. For working with text, the most famous example of this approach is the *Stylometry in R* package (Eder, Rybicki, and Kestemont 2016). Many current projects depend on libraries built in the R and Python programming languages, which are often easy to use. The code I have used for projects reported in this book was mostly written in Python, and I made extensive use of libraries developed by other people for the analysis and visualization of data, as can be seen in the online companion to this book. These libraries are open source, which means others can reuse and extend them. In this way, programming, like theater, becomes a social activity that builds on collective effort and shared ideals.

Even a rudimentary knowledge of programming will enable us, theater scholars, to better understand our projects, and to be more attentive to latent assumptions and blind spots in our methodologies. The more we increase our programming proficiency, the more resilient we will be as a community that doesn't need to outsource decision making to other disciplines and to commercial companies. There is, of course, a danger of romanticizing self-sufficiency. No one can carry out the entire process on their own. You can't build a computer from scratch and write every piece of software that runs on it. But an incursion into the making of tools will sensitize us to things at stake, and to kinds of labor which are often invisible. A central theme running through this book is that every method is value-laden. When choosing a method, or a source of data, there is no value-free option, as every choice represents a set of preferences and prejudices. If we just choose "default" options in a software package or online portal, that means someone else has made the choice for us. If we want to carry out nuanced data-driven and data-assisted research, we need to be able to own our assumptions. This is easier to do if we rethink software programming as a mode of critical thinking.

For example, thinking as programmers will enable us to remain attentive to cultural specificity. If we can trace the assumptions of the algorithms we use, we will be better able to reflect on what gets misrepresented and

craft tools better suited to our objectives. Many named entity recognition algorithms are premised on the idea that personal names are made of a first name, an optional middle name, and a family name (expressed in that order). But Chinese names often begin with the family name, followed by a bisyllabic first name (where each syllable might be separated by a space when written in Roman characters or hyphenated). Mexican names (like mine) are based on Spanish naming conventions where a first name is followed by two family names (one from the father's and one from the mother's side); many Indonesians have only one name, or two first names but no hereditary family name. These are just some examples, but naming conventions vary extensively around the world. This is not often recognized by the designers of systems for name recognition which underpin, among other things, academic citation databases. If we can code, we can try to circumvent uneven practices in the work that we do. We might not be able to change the ways of big corporate systems, but we should be able to intervene in the systems that we need for data research in the humanities. We need open platforms that others can inspect and adapt, but we also need people who are able to do this. People who know how to program are more likely to demand rightful access to the source code. And open source code encourages people to program. Software programming will also better enable us to develop performative, data-assisted visualizations that challenge conventional visual tropes, and to craft intermedial essays in ways that are not predefined by other people, but that stay close to the topics and intellectual dispositions that interest us.

Learning to code expands a scholar's horizon in a way akin to learning another language. One could perhaps say insightful things about French theater without speaking French. But even a tentative incursion into the French language will enable a scholar to grasp subtle differences, ask better questions, and trace a more nuanced history of prejudices and intellectual positions. Code is also quickly becoming a *lingua franca*, and the imperative to learn to think as programmers has gained urgency in a world where software shapes so many aspects of life. As Kitchin and Dodge (2014, 1) note:

It is very difficult to avoid the effects—the work—of software in the world [. . .] because of the difference it makes to the constitution and practices of everyday life. Indeed, to varying degrees, software conditions our very existence. Living beyond the mediation of software means being apart from collective life.

These remarks have implications for pedagogy, and for what we expect of our students. Undergraduate students starting their theater studies degrees at the time of writing are digital natives, or so the conventional thinking goes. But are they really native to this world, if they are only users, rather than makers? If they can't rewrite the code that underpins their reality, they don't have full citizenship in this digital world. As Rushkoff (2011, 12) puts it, in equally provocative terms:

In the emerging, highly programmed landscape ahead, you will either create the software or you will be the software. It's really that simple: program, or be programmed. Choose the former, and you gain access to the control panel of civilization. Choose the latter, and it could be the last real choice you get to make.

This passage, which I often use in my introductory DH classes is also quoted by Ramsay (2012), who reflects on how to teach programming to humanities students. He makes a crucial point when he says that teaching students to program also teaches them to think, and that programming teaches humanities thinking as much as computational thinking. While there is a strong appeal in teaching humanities students to code in order to improve their chances in the job market, I agree with Ramsay's point that we need to teach programming skills to make sure that humanities thinking can continue into a digital age in our own terms, rather than merely on the conditions of dominant software ideologies.

When I teach my students to program, I also use another quote from Ramsay, in this case to set them at ease. Many feel inadequate when learning to program and they see these type of work as alien to their core interests. But via Ramsay (2014), I reassure them that the key ingredients of programming are curiosity and an open mind:

One thing is certain: Being good at mathematics in no way guarantees that one will be good at programming (or vice versa). My own (admittedly anecdotal) experience as a teacher suggests that being musical, enjoying games and puzzles, being a tinkerer, loving to cook, and being a good long-form writer are far better predictors of success. A very high tolerance for frustration helps as well. (n.p.)

There is a craft, an art even, to programming. And this point has been perhaps best captured in the words of Paul Graham (2003), who com-

pares hackers to artists. For him, as for many others, “hacking” refers to an approach to programming premised on trial-and-error rather than on engineering models:

The fact that hackers learn to hack by doing is another sign of how different hacking is from the sciences. Scientists don’t learn science by doing it, but by doing labs and problem sets. Scientists start out doing work that’s perfect, in the sense that they’re just trying to reproduce work someone else has already done for them. Eventually, they get to the point where they can do original work. Whereas hackers, from the start, are doing original work; it’s just very bad. So hackers start original, and get good, and scientists start good, and get original. (n.p.)

Programming is closer to the modes of working of theater makers and researchers than people conventionally assume. And even learning a little bit of programming, learning to playfully hack and adapt our tools can help us reimagine our relationship to our technologies and to our theater traditions. By learning to program—or rather, by discovering programming through creative tinkering—we can also choose to represent things differently, away from hegemonic paradigms of representation. Vikram Chandra, the renowned Indian novelist, worked for a long time as a software programmer, although this is a little known fact of his biography. Like me (and many others in DH), he did not receive formal training in software engineering but learned by endless creative tinkering. Chandra (2014) vividly describes the pleasure of experimentation as the driving force in his learning process:

The work of making software gave me a little jolt of joy each time a piece of code worked; when something wasn’t working, when the problem resisted and made me rotate the contours of the conundrum in my mind, the world fell away, my body vanished, time receded. And three or five hours later, when the pieces of the problem came together just so and clicked into a solution, I surfed a swelling wave of endorphins. [. . .] Even after you are long past your first “Hello, world!” there is an infinity of things to learn, you are still a child, and—if you aren’t burned out by software delivery deadlines and management-mandated all-nighters—coding is still play. You can slam this pleasure spike into your veins again and again, and you want more, and more, and more. (18–19)

What we need is this playfulness, this tinkering that characterizes hacking. As Mark Olson (2013) has suggested:

A hacker ethos is a way of feeling your way forward, through trial and error, up to and perhaps beyond the limits of your expertise, in order to make something, perhaps even something new. It is provisional, sometimes ludic, and involves a willingness to transgress boundaries, to practice where you don't belong. (238)

Besides the ludic pleasure that programming can bring, programming also frees up financial resources and enhances our imagination. These two related points are more cogently articulated by the proponents of *minimal computing* from GO::DH (Global Outlook Digital Humanities, a special interest group within the Association for Digital Humanities Organizations), from whom I have borrowed the concept:

We use “minimal computing” to refer to computing done under some set of significant constraints of hardware, software, education, network capacity, power, or other factors. Minimal computing includes both the maintenance, refurbishing, and use of machines to do DH work out of necessity along with the use of new streamlined computing hardware like the Raspberry Pi or the Arduino micro controller to do DH work by choice. This dichotomy of choice vs. necessity focuses attention on computing that is decidedly not high-performance. By operating at this intersection between choice and necessity minimal computing forces important concepts and practices within the DH community to the fore. (Minimal Computing Working Group 2015, n.p.)

The interface of choice and necessity is articulated in different ways by the most active members of the Working Group. Alex Gil (2015) suggests that scholars around the world (including librarians and students) need to ask themselves what counts as *sufficient*, and ask themselves what is their specific goal. He urges us to remember that the most important objective of our efforts should be the “renewal, dissemination and preservation of the scholarly record.” Doing this in a sophisticated, ethical, and sustainable way requires that we learn “to produce, disseminate and preserve digital scholarship ourselves, *without the help we can't get*, even as we fight to build the infrastructures we need at the intersection of and beyond our librar-

ies and schools” (n.p., original emphasis). The projects I advocate for and participate in are a response to this call for action. Like Gil, I also endeavor with my collaborators to do things ourselves, while also building long-lasting infrastructure. Jentery Sayers (2016) responds to Gil in another short position piece saying that minimal DH should be “not only what we need but what we want” (n.p.). Sayer’s crucial question is: “How might minimal computing increase our shared capacities to think or imagine, and not just our individual capacities to work or produce?” (n.p.). The question of minimal infrastructures is thus not purely practical, but carries with it an important disposition towards our code and machines.

Inspired by these views, I imagine data-driven and data-assisted theater projects built with the minimum possible financial investment and built in such a way that minimal conservation efforts are needed for their long-term sustainability. This is at odds with the way many DH projects are done, which require large financial investments, but allocate little resources (or thought) to long-term sustainability. An expensive project might be harder to maintain than one that requires modest investment, but this also depends on where the investment is made. This might seem counterintuitive to people who assume that keeping projects alive is very expensive. A good data project, in my view, is like a bicycle—it won’t be fast but it will be reliable, requiring much less maintenance than an expensive car. Even if you have the money to buy a car, it might be a smarter strategy to buy a bicycle or a series of them and distribute the money throughout the years, making sure the bike-enabled transportation is available in the future. Digital sustainability (see chapter 8) provides an important justification for the allocation of financial and other resources, including the substantial amount of time and effort that goes into the development of such projects. But sustainability should be a primary concern for intellectual reasons, as the traditions of the humanities—including theater and performance studies—are comparative and historical.

It is, however, difficult to talk about DH work without considering money (Liu 2012). The presence of large financial resources for DH funding has been a cause for both celebration and concern. In a time when funding for the humanities is decreasing across the Western world, DH has enabled academic positions, grants, and research centers to flourish. But this situation has also been assessed negatively by those who fear that the presence of large budgets has co-opted humanities interests into subservience to capital, technoscience, and managerism at universities (Berry and Fagerjord 2017, 10). Another reason why we should be vigilant

of the impact of financial resources is that they tend to be disproportionately appointed to performance traditions from rich countries and to well-known projects. My fear is that theater projects with no data materials will become even less well known. But it is precisely here where I see a shimmer of hope. Data scholarship has a greater potential to break through the gatekeepers of academia than other kinds of projects. Smartly planned data projects from around the world might disrupt knowledge distribution channels if resources are allocated in a way that ensures sustainability. My own work is mostly funded by generous research grants in Singapore, but I also work with cash-starved institutions in Indonesia and have developed several projects with minimal financial resources.

As Gil, Sayers, and other members of the Minimal Computing WG suggest, technical knowledge is required in order to ensure that minimal possible investment and sustainability are possible. Think of the bicycle again: if you know how to repair it yourself, the costs of maintenance will be much cheaper. Or, at least you will know when you are overpaying for a service. Financial constraints are conventionally seen as limitations for DH, especially in parts of the world where funding is harder to come by. But there is a way of doing sustainable, highly technical work with minimal investments: learning to code and relying on open infrastructures. This strikes a particular chord with my colleagues in Indonesia, where funding is scarce. However, this should be a goal even for well-funded projects in rich institutions, for the reasons discussed above. Funding is often seen as one of the greatest impediment, or enablers, of computational research. However, computational theater research can be carried out by people who understand the foundations of their tools, with the minimum possible financial investments and it can be built to last.

Even as we learn to code, we should still be able to engage in productive research dialogues. Good work will certainly come out of collaborations with colleagues in technical and scientific fields. But these collaborations will be all the more meaningful if all partners understand each other's tools and disciplinary histories. In carrying out research for projects reported in this book, I have been particularly lucky to collaborate with Gea O. F. Parikesit and Andrew Schauf. Both of them are physicists with an extraordinary awareness of the histories of the humanities and the complexities of Indonesian arts. When theater scholars collaborate with scientists and programmers, the communication will be more meaningful as both sides learn to understand different paradigms and approaches.

If we, as theater scholars want to engage with data, we need to learn

to craft our software for the variety of reasons I have argued above: to enhance our way of thinking critically about methodology, to create culturally sensitive work, to gain control of our tools, and to deepen our collaborations and to make work that lasts. Embedding programming into our academies will no doubt be tricky. But as Bench and Elswit (2017) remark, who better than theater scholars to think about new modes of scholarship that are collaborative, creative, and practice-based? The paradigm shift in theater studies that enabled the rise of practice-based research might provide the right model for computational work. We have already rethought the place of practice in the making of new knowledge, it is now time to consider the ways in which programming can further expand our methodological imaginations.

Appendix A

Data Biographies

All the code for the excursions in the book was implemented in Python 3.7.6, using the Pandas package, version 1.0.1. All images were generated with Matplotlib 3.1.3 (Hunter 2007) and Seaborn version 0.10.0.

Chapter 4

The theater reviews were downloaded from <https://inkpotreviews.com/archive.html> in January 2019. The reviews for each year were grouped into separate text files, which omitted the names of the reviewers and titles of the reviews. These files were uploaded to Voyant (<http://voyant-tools.org>), where I used built-in functions to calculate the trends over time and the concordances. I then downloaded these files for processing in Python.

4.1_trendOverTime.csv <https://doi.org/10.3998/mpub.11667458.cmp.26>

This file was downloaded from Voyant, and then analyzed in Python. To calculate the Mann-Kendall statistics I used the `pymannkendall` package version 1.4.1 (Hussain and Mahmud 2019) for Python.

4.2_audienceKWIC.csv <https://doi.org/10.3998/mpub.11667458.cmp.27>

This file is the concordance for sentences containing the KWIC concordance for the word audience, which I then manually categorized into descriptive and rhetoric mentions.

4.3_audienceTrends.csv <https://doi.org/10.3998/mpub.11667458.cmp.28>

This file includes the results of manual analysis of the audience concordance.

Chapter 5

The file with the list of cast members and roles was provided by The Necessary Stage (TNS) in January 2016. My research assistant Alysha Chandra and I verified all names and roles and standardized the spelling of names in the files for consistency. Based on this data, the network was built using NetworkX version 2.4 (Developers 2010). The number of nodes, edges, and components were calculated using the built-in functions of this library. The files below include data on the component size, as well as the number of nodes and edges over time.

5.1_TNScomponents.csv <https://doi.org/10.3998/mpub.11667458.cmp.29>

This file includes information on all people involved in the production.

5.2_TNScast.csv <https://doi.org/10.3998/mpub.11667458.cmp.30>

This file includes information on people labeled as cast members in the file provided by TNS.

5.3_TNSproductionCrew.csv <https://doi.org/10.3998/mpub.11667458.cmp.31>

This file includes information on people labeled as production crew members in the file provided by TNS.

The interactive visualizations for the Digital Wayang Encyclopedia are available at <https://villaorlado.github.io/wayangnetworks/html/> and in video 5.1 (<https://doi.org/10.3998/mpub.11667458.cmp.23>) in this book's web companion. The data was obtained from several published wayang encyclopedias (Hardjowirogo 1948; Sudibyoprano et al. 1991; Sudjarwo, Sumari, and Wiyono 2010; Purwadi 2013; Solichin et al. 2017). The data was verified by my research assistants Losheini Ravindran and Yosephine Novi Marginingrum.

5.4_wayangNodeInfo.csv <https://doi.org/10.3998/mpub.11667458.cmp.32>

This resource includes information on each character.

5.5_wayangNetwork.gephi <https://doi.org/10.3998/mpub.11667458.cmp.33>

This is the Gephi file of the full network.

Chapter 6

The original *Wayang Mitologi* (Catur Kuncoro, 2012) video recording is available at <http://cwa-web.org/en/WayangMitologi>

6.1_mitologiDifferenceImages.csv <https://doi.org/10.3998/mpub.11667458.cmp.34>

This resource includes the number of pixels in each difference image. The video was processed by Gea Oswah Fatah Parikesit using Scilab and the Scilab Image and Video Processing Toolbox (SIVP). In the video recording of *Wayang Mitologi*, the noise level is 10 greyvalues, so we only consider pixels with greyvalues higher than 10 as the non-static pixels. This was checked by recording two dark images and measuring the greyvalues of the supposedly dark pixels. The video is 69 minutes long. We obtained 1,000 images per minute, and then sampled one out of each 150 difference images, which corresponds to ~0.1 min. The file indicates how each of the 460 difference images corresponds to the 69,652 frame numbers and the 69 minutes.

This file also includes scene segmentation information, which I manually added.

For the Sendratari *Ramayana* dance, Luis Hernández-Barraza and I used a motion-capture system (Vicon MX, Oxford Metrics, Oxford, UK), that consists of seven infrared cameras to collect kinematic data at a sample rate of 100 hertz. Forty-one reflective markers (14 mm diameter) were attached to the dancer following the full body Plug-In-Gait Marker model (Vicon, Oxford Metrics, Oxford, UK) to facilitate capture of the dancer's motion. Two embedded force plates (AMTI, Watertown, Massachusetts, USA) were used to obtain GRF data at a sampling rate of 1,000 hertz. The force plates were synchronized to the motion capture system, and both were calibrated according to the manufacturer's recommendations before the dance data was recorded.

6.2_danceSubtypesAngleData.csv <https://doi.org/10.3998/mpub.11667458.cmp.35>

The material in this resource was obtained and processed by Luis Hernández-Barraza. The columns are named according to the character subtype, the joint and the plane (x, y, or z). Each row corresponds to the measurement of an angle (1 observation per ms).

Chapter 7

The data was obtained from <https://kluban.net>. I used the Python BeautifulSoup version 4.8.2 package to scrap the website and then format the results into a CSV file. The geographical regions were assigned via a rule based system. All cases where the rules could not sufficiently disambiguate the region, were manually resolved. I built a custom Python script to achieve both operations (rule-based and manual disambiguation). The population statistics were obtained from Badan Pusat Statistik (Indonesian Center for Statistics), and they correspond to 2016 estimates. The maps were generated with Geopandas version 0.8.1. Moran's I was calculated using PySAL version 1.14.4.

7.1_placeAndDate.csv <https://doi.org/10.3998/mpub.11667458.cmp.36>

This file includes the date and location of each performance. The date is meant to be used as a Python datetime object. The time information can be discarded, as only the dates are relevant. Time was assigned at 00:00 for consistency, but most performances begin at 21:00 and end around 04:00 the next day. The date in the file corresponds to the beginning of the performance.

7.2_performancesAndPopulation.csv <https://doi.org/10.3998/mpub.11667458.cmp.37>

This resource includes aggregate counts per regency as well as statistics on the size, population and population density of each regency.

7.3_performancesPerRegency.csv <https://doi.org/10.3998/mpub.11667458.cmp.38>

This resource includes aggregate counts per regency.

Appendix B

Technical Glossary

Betweenness centrality (network analysis): This is a measure of how often a node acts as a bridge between other nodes. A high betweenness centrality indicates that the shortest paths between all pairs of nodes in the network often pass through the given node.

Cartogram: A geographic visualization which distorts the shape of geographical entities (countries, districts, etc.) according to a particular metric or statistic. The resulting maps show how much a given area contributes to the total for a given metric in a way that is not proportional to the actual land area of a province or country.

Choropleth: A geographic visualization where areas (districts, counties, countries) are colored according to different quantitative values.

Circos plots: A type of chord diagrams that show interrelations between data in a matrix. Each data point is represented as a dot in the circle's circumference, which is connected by lines drawn across the area of the circle.

Closeness centrality (network analysis): The inverse of the average length of the most direct paths between the given node and all other nodes in the network.

Cohen's d : A measure of the difference between groups that can be used to estimate the magnitude of a phenomenon (effect size). Cohen's d is the difference between two means (one for each group) divided by the standard deviation for the data.

Complete spatial randomness (CSR): A distribution of points in a geographical location that conforms with a random model's prediction.

Degree (network analysis): In directed networks, the degree is the number of edges of a given node. In undirected networks, we can distin-

guish between in-degree (edges with the given node as target) and out-degree (edges with the given node as source).

Difference image: An image that results from subtracting one image from another. This is used to compare subsequent images extracted from videos. The number of pixels in an image difference can be used to estimate the amount of movement in a video.

Eigenvector centrality (network analysis): A measurement of the influence of the node in the network that considers the degrees of a node's neighbors. Nodes with high eigenvector centrality tend to be connected with neighbors who are themselves highly connected.

Kernel density estimation (KDE): A non-parametric (i.e., that doesn't impose a model) method for visually estimating the probability density function of a variable.

Mann-Kendall s : A measure for the strength of a monotonic trend (that is, one that increases or decreases consistently). A negative value means that the trend is decreasing, while a positive value means it is increasing.

Minimum convex polyhedron: A solid in three dimensions with flat polygonal faces, straight edges and sharp corners.

Moran's I : A score used to determine whether statistical proximity affects a given variable (spatial autocorrelation). Moran's I estimates the probability that a value for a geographical unit is affected by the average value of neighboring units.

Network density: The number of actual edges in a network divided by the maximum total number of edges.

Network dynamics: The study of networks as they change over time.

Weighted degree (network analysis): The sum of the weights of all the node's links.

References

- Aarseth, Espen J. 1997. *Cybertext: Perspectives on Ergodic Literature*. UK edition. Baltimore: Johns Hopkins University Press.
- Abdulkadiroğlu, Atila, Joshua Angrist, and Parag Pathak. 2014. "The Elite Illusion: Achievement Effects at Boston and New York Exam Schools." *Econometrica* 82 (1): 137–96.
- Academia Sinica Center for Digital Cultures. n.d. "Starting Out from 23.5°N: Chen Cheng-Po." Accessed December 18, 2018. http://chenchengpo.asdc.sinica.edu.tw/about_en
- Agneessens, Filip, Henk Roose, and Hans Waege. 2004. "Choices of Theatre Events: P* Models for Affiliation Networks with Attributes." *Metodoloski Zvezki* 1 (2): 419.
- Alberich, Ricardo, Joe Miro-Julia, and Francesc Rosselló. 2002. "Marvel Universe Looks Almost like a Real Social Network." *ArXiv Preprint Cond-Mat/0202174*.
- Algee-Hewitt, Mark. 2017. "Distributed Character: Quantitative Models of the English Stage, 1550–1900." *New Literary History* 48 (4): 751–82. <https://doi.org/10.1353/nlh.2017.0038>
- Angwin, Julia, and Jeff Larson. 2016. "Machine Bias." Text/html. ProPublica. May 23, 2016. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- Arnold, Taylor, and Lauren Tilton. 2019. "New Data? The Role of Statistics in DH." In *Debates in the Digital Humanities 2019*, by Matthew K. Gold and Lauren F. Klein. Debates in the Digital Humanities. Minneapolis: University of Minnesota Press. <https://dhdebates.gc.cuny.edu/read/untitled-f2acf72c-a469-49d8-be35-67f9ac1e3a60/section/a2a6a192-f04a-4082-afaa-97c76a75b21c#ch24>
- Arps, Bernard. 2016. *Tall Tree, Nest of the Wind: The Javanese Shadow-Play Dewa Ruci Performed by Ki Anom Soeroto: A Study in Performance Philology*. Singapore: NUS Press.
- Arrowsmith, Colin, Deb Verhoeven, and Alwyn Davidson. 2014. "Exhibiting the Exhibitors: Spatial Visualization for Heterogeneous Cinema Venue Data." *Cartographic Journal* 51 (4): 301–12. <https://doi.org/10.1179/1743277414Y.0000000096>
- Arulampalam, J., J. Pierrepont, and L. Kark. 2015. "Markerless Motion Capture: Validity of Microsoft Kinect Cameras and IPisof." *Gait & Posture*, 24th Annual Meeting of ESMAC 2015 Abstracts, 42 (September): S76. <https://doi.org/10.1016/j.gaitpost.2015.06.141>
- Ashworth, Peter D., and Man Cheung Chung, eds. 2006. *Phenomenology and Psychological Science: Historical and Philosophical Perspectives*. History, Philosophy, Psychology. New York: Springer.

- Asnaghi, Costanza, Dirk Speelman, and Dirk Geeraerts. 2016. "Geographical Patterns of Formality Variation in Written Standard California English." *Literary and Linguistic Computing* 31 (2): 244–63. <https://doi.org/10.1093/llc/fqu060>
- Aston, Elaine, and George Savona. 1991. *Theatre as Sign-System*. Routledge.
- AusStage. 2013. "AusStage." 2018 2013. <https://www.ausstage.edu.au/>
- Bach, Benjamin, Moritz Stefaner, Jeremy Boy, Steven Drucker, Lyn Bartram, Jo Wood, Paolo Ciuccarelli, Yuri Engelhardt, Ulrike Köppen, and Barbara Tversky. 2018. "Narrative Design Patterns for Data-Driven Storytelling." In *Data-Driven Storytelling*, edited by Nathalie Henry Riche, Christophe Hurter, Nicholas Diakopoulos, and Sheelagh Cappendale, 59–83. New York: A. K. Peters/CRC Press. <https://doi.org/10.1201/9781315281575-3>
- Badan Pusat Statistik. 2019. "Badan Pusat Statistik." 2019. <https://www.bps.go.id/>
- Baker, Zachary, Joel Berkowitz, Sonia Gollance, Debra Caplan, Barbara Henry, Faith Jones, C. Tova Markenson, et al. 2019. "Digital Yiddish Theatre Project." Text/html. Digital Yiddish Theatre Project. March 23. <https://yiddishstage.org/about>
- Bakshy, Eytan, Dean Eckles, and Michael S. Bernstein. 2014. "Designing and Deploying Online Field Experiments." In *Proceedings of the 23rd International Conference on World Wide Web*, 283–92. ACM.
- Balazia, M., and P. Sojka. 2017. "You Are How You Walk: Uncooperative MoCap Gait Identification for Video Surveillance with Incomplete and Noisy Data." In 2017 *IEEE International Joint Conference on Biometrics (IJCB)*, 208–15. <https://doi.org/10.1109/BTAS.2017.8272700>
- Balme, Christopher. 2015. *The Cambridge Introduction to Theatre Studies*. Cambridge: Cambridge University Press.
- Barabási, Albert-László. 2002. *Linked*. Cambridge: Perseus.
- Bardiot, Clarisse. 2015a. "Arts de La Scène et Big Data. Retracer et Analyser Le Processus de Création d'un Spectacle Grâce à La Visualisation de Données." In *Le Numérique à l'ère de l'Internet Des Objects: De l'hypertexte à l'hyperobjet*. Paris.
- Bardiot, Clarisse. 2015b. "ReCALL: An Environment for Notation/Annotation/Denotation." *Performance Research* 20 (6): 82–86. <https://doi.org/10.1080/13528165.2015.1111058>
- Bardiot, Clarisse. 2017. "Arts de La Scène et Culture Analytics." *Revue d'histoire de l'histoire Du Théâtre* 4.
- Bardiot, Clarisse. 2018. "Measuring Merce Cunningham: A Theatre Analytics Research." In *Bridges/Puentes*. Mexico City: ADHO.
- Bardzell, Jeffrey, and Shaowen Bardzell. 2013. "What Is Critical about Critical Design?" In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 3297–3306. ACM.
- Bateman, John A. 2017. "Triangulating Transmediality: A Multimodal Semiotic Framework Relating Media, Modes and Genres." *Discourse, Context & Media* 20 (December): 160–74. <https://doi.org/10.1016/j.dcm.2017.06.009>
- Battacharyya, Sayan. 2018. "Non-normative Data from The Global South and Epistemically Produced Invisibility in Computationally Mediated Inquiry—DH2018." In *Proceedings of the 2018 Digital Humanities Conference*. Mexico City. <https://dh2018.adho.org/en/non-normative-data-from-the-global-south-and-epistemically-produced-invisibility-in-computationally-mediated-inquiry/>
- Bay-Cheng, Sarah. 2017. "Digital Historiography and Performance." *Theatre Journal* 68 (4): 507–27. <https://doi.org/10.1353/tj.2016.0104>

- Bay-Cheng, Sarah, Robin Nelson, Andy Lavender, and Robin Nelson. 2010. *Mapping Intermediality in Performance*. Amsterdam: Amsterdam University Press.
- Bay-Cheng, Sarah, Jennifer Parker-Starbuck, and David Z. Saltz. 2015. *Performance and Media*. Ann Arbor: University of Michigan.
- Bedeck, Michael A., Alexander Nussbaumer, Luca Huszar, and Dietrich Albert. 2018. "Methods for Discovering Cognitive Biases in a Visual Analytics Environment." In *Cognitive Biases in Visualizations*, edited by Geoffrey Ellis, 61–73. Cham: Springer.
- Bench, Harmony, and Kate Elswit. 2017. "Mapping Movement on the Move: Dance Touring and Digital Methods." *Theatre Journal* 68 (4): 575–96. <https://doi.org/10.1353/tj.2016.0107>
- Benesh Institute. n.d. "Benesh Notation Editor—Benesh Notation Editor Help Centre." Accessed December 9, 2018. http://www.ok-edesign.com/Benesh/BNBNE_WhatIsBne.html
- Bermudez, Bertha, Scott Delahunta, Hoogenboom Marijke, Ziegler Chris, Frédéric Bevilacqua, Sarah Fdili Alaoui, and Barbara Meneses Gutierrez. 2011. "The Double Skin/Double Mind Interactive Installation."
- Bernstein, Lenny, Peter Bosch, Osvaldo Canziani, Zhenlin Chen, Renate Christ, Ogunlade Davidson, William Hare, Saleemul Huq, David Karoly, and Vladimir Kattsov. 2008. *Climate Change 2007: Synthesis Report: An Assessment of the Intergovernmental Panel on Climate Change*. IPCC.
- Berry, David M., and Anders Fagerjord. 2017. *Digital Humanities: Knowledge and Critique in a Digital Age*. Cambridge: Polity Press.
- Bertin, Jacques. 1983. *Semiology of Graphics; Diagrams Networks Maps*. Madison: University of Wisconsin Press.
- Best, Michael, and Janelle Jenstad. 2013. *The Internet Shakespeare Editions*. University of Victoria. <http://internetshakespeare.uvic.ca/>
- Beyes, Timon, Martina Leeker, and Imanuel Schipper, eds. 2017. *Performing the Digital: Performativity and Performance Studies in Digital Cultures*. Digital Society. Bielefeld: Transcript Verlag.
- Bhaskar, Roy, and Mervyn Hartwig. 2010. *The Formation of Critical Realism: A Personal Perspective*. Ontological Explorations. Routledge.
- Biet, Christian, Pierre Frantz, Florence Filippi, Georges Forestier, Sylvaine Guyot, Sara Harvey, Tiphaine Karsenti, Sophie Marchand, Jeffrey Ravel, and Agathe Sanjuan. 2015. "The Comédie-Française Registers Project—Comédie Française Registers Project." 2015. <https://www.cfregisters.org/en/>
- Binongo, J. N. G., and M. W. A. Smith. 1999. "The Application of Principal Component Analysis to Stylometry." *Literary and Linguistic Computing* 14 (4): 445–66. <https://doi.org/10.1093/lc/14.4.445>
- Birch, David. 2004. "Celebrating the Ordinary in Singapore in Extraordinary Ways: The Cultural Politics of The Necessary Stage's Collaborative Theatre." In *Ask Not: The Necessary Stage in Singapore Theatre*, edited by Chong Kee Tan and Tisa Ng, 266–90. Singapore: Times Editions.
- Bishop, Ryan, and John Phillips. 2007. "Of Method." *Theory, Culture & Society* 24 (7–8): 264–75. <https://doi.org/10.1177/0263276407084796>
- Black, Carolyn, John Craig, James Cummings, Matthew Davies, Alexandra Gillespie, Peter Greenfield, Diane Jakacki, et al. 2017. "REED Online." 2017. <https://reed.library.utoronto.ca/>

- Bleeker, Maaike. 2016. *Transmission in Motion: The Technologizing of Dance*. Oxon and New York: Routledge.
- Bode, Katherine. 2012. *Reading by Numbers: Recalibrating the Literary Field*. London and New York: Anthem Press. <http://site.ebrary.com/id/10595451>
- Bode, Katherine. 2020. "Why You Can't Model Away Bias." *Modern Language Quarterly* 81 (1).
- Bodenhamer, David J. 2013. "Beyond GIS: Geospatial Technologies and the Future of History." In *History and GIS: Epistemologies, Considerations and Reflections*, edited by Alexander von Lünen and Charles Travis, 1–13. Dordrecht: Springer Netherlands. https://doi.org/10.1007/978-94-007-5009-8_1
- Bodenhamer, David, John Corrigan, and Trevor M. Harris. 2013. "Deep Mapping and the Spatial Humanities." *International Journal of Humanities and Arts Computing* 7. <http://diginole.lib.fsu.edu/islandora/object/fsu%3A209941/>
- Bollacker, Kurt D. 2010. "Avoiding a Digital Dark Age." *American Scientist* 98 (2): 106–10.
- Bollen, Jonathan. 2017. "Data Models for Theatre Research: People, Places, and Performance." *Theatre Journal* 68 (4): 615–32. <https://doi.org/10.1353/tj.2016.0109>
- Bollen, Jonathan, and Julie Holledge. 2011. "Hidden Dramas: Cartographic Revelations in the World of Theatre Studies." *Cartographic Journal* 48 (4): 226–36. <https://doi.org/10.1179/1743277411Y.0000000026>
- Borgman, Christine L. 2009. "The Digital Future Is Now: A Call to Action for the Humanities." *Digital Humanities Quarterly* 3 (4). <http://www.digitalhumanities.org/dhq/vol/3/4/000077/000077.html>
- Borgman, Christine L. 2015. *Big Data, Little Data, No Data*. Cambridge: MIT Press.
- Bowker, Geoffrey C., and Susan Leigh Star. 1999. *Sorting Things Out*. Inside Technology. Cambridge: MIT Press.
- Breiman, Leo. 2001. "Statistical Modeling: The Two Cultures (with Comments and a Rejoinder by the Author)." *Statistical Science* 16 (3): 199–231.
- Broadwell, Peter M. 2019. "Automated Movement and Choreography Analysis of Video Data via Deep Learning Pose Detection." 2019. http://broadwell.github.io/DH2019_Movement_Choreography.html
- Broadwell, Peter M., and Timothy R. Tangherlini. 2017. "GhostScope: Conceptual Mapping of Supernatural Phenomena in a Large Folklore Corpus." In *Maths Meets Myths: Quantitative Approaches to Ancient Narratives*, edited by Ralph Kenna, Máirín MacCarron, and Pádraig MacCarron, 131–57. Understanding Complex Systems. Cham: Springer International. https://doi.org/10.1007/978-3-319-39445-9_8
- Brodbeck, Frederic. 2011. "CINEMETRICS—Film Data Visualization." 2011. <http://cinemetrics.fredericbrodbeck.de/>
- Brown, Susan, and John Simpson. 2013. "The Curious Identity of Michael Field and Its Implications for Humanities Research with the Semantic Web." 2013 *IEEE International Conference on Big Data*, 77–85. Silicon Valley, CA, October 6–9. <https://doi.org/10.1109/BigData.2013.6691674>
- Bruner, Edward M. 2005. *Culture on Tour: Ethnographies of Travel*. Chicago: University of Chicago Press.
- Buckland, Warren. 2008. "What Does the Statistical Style Analysis of Film Involve? A Review of Moving into Pictures. More on Film History, Style, and Analysis." *Literary and Linguistic Computing* 23 (2): 219–30. <https://doi.org/10.1093/lc/fqmq046>
- Buddenbohm, Stefan, Nathanael Cretin, Elly Dijk, Bertrand Gaiffe, Maaike De Jong,

- Nathalie Le Tellier-Becquart, and Jean-Luc Minel. 2016. "State of the Art Report on Open Access Publishing of Research Data in the Humanities." DARIAH.
- Bulger, Monica, Greg Taylor, and Ralph Schroeder. n.d. "Data-Driven Business Models: Challenges and Opportunities of Big Data," 74.
- Burrows, John. 2002. "'Delta': A Measure of Stylistic Difference and a Guide to Likely Authorship." *Literary and Linguistic Computing* 17 (3): 267–87. <https://doi.org/10.1093/llc/17.3.267>
- Burrows, John. 2007. "All the Way Through: Testing for Authorship in Different Frequency Strata." *Literary and Linguistic Computing* 22 (1): 27–47. <https://doi.org/10.1093/llc/fqj067>
- Busa, Roberto. 2005. "Index Thomisticus." <http://www.corpusthomicum.org/it/in dex.age>
- Cairo, Alberto. 2020. "Foreword." In *Data Visualization in Society*, edited by Martin Engelbrechtsen and Helen Kennedy, 17–18. Amsterdam: Amsterdam University Press. <https://doi.org/10.2307/j.ctvzgb8c7.6>
- Calude, Cristian S., and Giuseppe Longo. 2017. "The Deluge of Spurious Correlations in Big Data." *Foundations of Science* 22 (3): 595–612. <https://doi.org/10.1007/s10699-016-9489-4>
- Calvert, Thomas W. 1986. "Toward a Language for Human Movement." *Computers and the Humanities* 20 (1): 35–43. <https://doi.org/10.1007/BF02393462>
- Cameron, Deborah. 2011. "Evolution, Science and the Study of Literature: A Critical Response." *Language and Literature*, February. <https://doi.org/10.1177/0963947010391126>
- Cao, Zhe, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. 2018. "OpenPose: Realtime Multi-Person 2D Pose Estimation Using Part Affinity Fields." *ArXiv: 1812.08008 [Cs]*, December. <http://arxiv.org/abs/1812.08008>
- Caplan, Debra. 2015. "Notes from the Frontier: Digital Scholarship and the Future of Theatre Studies." *Theatre Journal* 67 (2): 347–59. <https://doi.org/10.1353/tj.2015.0059>
- Caplan, Debra. 2017. "Reassessing Obscurity: The Case for Big Data in Theatre History." *Theatre Journal* 68 (4): 555–73. <https://doi.org/10.1353/tj.2016.0106>
- Carlin, David, and Laurene Vaughan, eds. 2015. *Performing Digital: Multiple Perspectives on a Living Archive*. First ed. Farnham, Surrey, England; Burlington, VT: Routledge.
- Carlson, K., T. Schiphorst, and C. Shaw. 2011. "ActionPlot: A Visualization Tool for Contemporary Dance Analysis." In *Proceedings of the International Symposium on Computational Aesthetics in Graphics, Visualization, and Imaging*, 113–20. CAe '11. New York: ACM. <https://doi.org/10.1145/2030441.2030466>
- Casey, Matthew. 2006. *Theatre and Performance in Digital Culture*. Vol. 5. Routledge Advances in Theatre and Performance Studies. London: Routledge.
- Chan, Jacky C. P., Howard Leung, Jeff K. T. Tang, and Taku Komura. 2010. "A Virtual Reality Dance Training System Using Motion Capture Technology." *IEEE Transactions on Learning Technologies* 4 (2): 187–95.
- Chandra, Vikram. 2014. *Geek Sublime: The Beauty of Code, the Code of Beauty*. Minneapolis: Graywolf Press.
- Chang, Michael, Mark Halaki, Roger Adams, Stephen Cobley, Kwee-Yum Lee, and Nicholas O'Dwyer. 2016. "An Exploration of the Perception of Dance and Its Relation to Biomechanical Motion: A Systematic Review and Narrative Synthesis." *Jour-*

- nal of Dance Medicine & Science 20 (3): 127–36. <https://doi.org/10.12678/1089-313X.20.3.127>
- Chapple, Freda, and Chiel Kattenbelt. 2006. *Intermediality in Theatre and Performance*. Themes in Theatre: Collective Approaches to Theatre and Performance. Amsterdam: Rodopi.
- Chatzichristodoulou, Maria, Janis Jefferies, and Rachel Zerihan. 2009. *Interfaces of Performance*. Digital Research in the Arts and Humanities. Farnham and Burlington: Ashgate.
- Cheesman, Tom, Kevin Flanagan, Zhao Geng, and Sebastian Sadowski. 2011. "Version Variation Visualization." <http://www.delightedbeauty.org/vvv>
- Choensawat, Worawat, Minako Nakamura, and Kozaburo Hachimura. 2015. "GenLaban: A Tool for Generating Labanotation from Motion Capture Data." *Multimedia Tools and Applications* 74 (23): 10823–46. <https://doi.org/10.1007/s11042-014-2209-6>
- Choi, Yeon-Mu, and Hyun-Joo Kim. 2007. "A Directed Network of Greek and Roman Mythology." *Physica A: Statistical Mechanics and Its Applications* 382 (2): 665–71.
- Cirillo, Pasquale, and Nassim Nicholas Taleb. 2016. "Expected Shortfall Estimation for Apparently Infinite-Mean Models of Operational Risk." *Quantitative Finance* 16 (10): 1485–94. <https://doi.org/10.1080/14697688.2016.1162908>
- Clemente, Filipe Manuel, Fernando Manuel Lourenço Martins, Dimitris Kalamaras, P. Del Wong, and Rui Sousa Mendes. 2015. "General Network Analysis of National Soccer Teams in FIFA World Cup 2014." *International Journal of Performance Analysis in Sport* 15 (1): 80–96. <https://doi.org/10.1080/24748668.2015.11868778>
- Cleveland, William. 1985. *The Elements of Graphing Data*. Monterey, CA: Wadsworth Advanced Books and Software.
- Cleveland, William. 1993. *Visualizing Data*. Summit, NJ: Hobart Press.
- Cohen, Jacob. 1988. *Statistical Power Analysis for the Behavioral Sciences*. Hillsdale, NJ: Lawrence Erlbaum Associates.
- Conquergood, Dwight. 2002. "Performance Studies: Interventions and Radical Research." *TDR/The Drama Review* 46 (2): 145–56. <https://doi.org/10.1162/105420402320980550>
- Craig, Hugh, and Brett Greatley-Hirsch. 2017. *Style, Computers, and Early Modern Drama: Beyond Authorship*. Cambridge: Cambridge University Press.
- Craig, Hugh, and Arthur F. Kinney. 2009. *Shakespeare, Computers, and the Mystery of Authorship*. Cambridge: Cambridge University Press.
- Cuykendall, Shannon, Ethan Soutar-Rau, and Thecla Schiphorst. 2016. "POEME: A Poetry Engine Powered by Your Movement." In *Proceedings of the TEI '16: Tenth International Conference on Tangible, Embedded, and Embodied Interaction*, 635–40. TEI '16. New York: ACM. <https://doi.org/10.1145/2839462.2856339>
- Cuykendall, Shannon, Ethan Soutar-Rau, Thecla Schiphorst, and Steve DiPaola. 2016. "If Words Could Dance: Moving from Body to Data Through Kinesthetic Evaluation." In *Proceedings of the 2016 ACM Conference on Designing Interactive Systems*, 234–38. DIS '16. New York: ACM. <https://doi.org/10.1145/2901790.2901822>
- D'Ignazio, Catherine, and Lauren F. Klein. 2016. "Feminist Data Visualization." In *Workshop on Visualization for the Digital Humanities (VIS4DH)*. Baltimore, MD: IEEE.
- D'Ignazio, Catherine, and Lauren F. Klein. 2020. *Data Feminism*. Cambridge, MA: MIT Press.
- Da, Nan Z. 2019. "The Computational Case against Computational Literary Studies." *Critical Inquiry* 45 (3): 601–39. <https://doi.org/10.1086/702594>

- Diakopoulos, Nicholas. 2018. "Ethics in Data-Driven Visual Storytelling." In *Data-Driven Storytelling*, edited by Nathalie Henry Riche, Christophe Hurter, Nicholas Diakopoulos, and Sheelagh Carpendale, 59–83. New York: A. K. Peters/CRC Press. <https://doi.org/10.1201/9781315281575-10>
- Davidson, Alwyn, Deb Verhoeven, and Colin Arrowsmith. 2015. "Petal Diagrams: A New Technique for Mapping Historical Change in the Film Industry." *International Journal of Humanities and Arts Computing* 9 (2): 142–63. <https://doi.org/10.3366/ijhac.2015.0146>
- DeLahunta, Scott. 2017. "Dance Becoming Data: Part One Software for Dancers." *Computational Culture*, no. 6 (November). <http://computationalculture.net/dance-becoming-data-part-one-software-for-dancers/>
- DeLahunta, Scott, and Florian Jenett. 2017. "Making Digital Choreographic Objects Interrelate. A Focus on Coding Practices." In *Performing the Digital. Performativity and Performance Studies in Digital Cultures*, by Martina Leeker, Imanuel Schipper, and Timon Beyes. Bielefeld: transcript. <http://dx.doi.org/10.25969/mediarep/2079>
- Delbridge, Matthew, and Joanne Tompkins. 2009. "Using Virtual Reality Modelling in Cultural Management, Archiving and Research." In *EVA London 2009*, 260–69. London.
- Derrida, Jacques. 1978. *Writing and Difference*. Chicago: University of Chicago Press.
- Deutsch, David. 2009. *A New Way to Explain Explanation*. TEDGlobal 2009. https://www.ted.com/talks/david_deutsch_a_new_way_to_explain_explanation
- Developers, NetworkX. 2010. "NetworkX." Networkx. Lanl. Gov.
- Dixon, Steve. 2007. *Digital Performance: A History of New Media in Theater, Dance, Performance Art, and Installation*. Cambridge, MA: MIT Press.
- Downey, Allen. 2014. *Think Stats: Exploratory Data Analysis in Python*. Open Textbook Library. Needham: Green Tea Press.
- Drucker, Johanna. 2009. *Speculative Computing: Basic Principles and Essential Distinctions*. Chicago: University of Chicago Press. <http://chicago.universitypressscholarship.com/view/10.7208/chicago/9780226165097.001.0001/upso-9780226165073-chapter-2>
- Drucker, Johanna. 2011. "Humanities Approaches to Graphical Display." *Digital Humanities Quarterly* 5 (1).
- Drucker, Johanna. 2013. "Diagrammatic Writing." *New Formations: A Journal of Culture Theory/Politics* 78 (1): 83–101.
- Drucker, Johanna. 2014. "Graphesis." *Paj: The Journal of the Initiative for Digital Humanities, Media, and Culture* 2 (1).
- Duffell, Lynsey D., Natalie Hope, and Alison H. McGregor. 2014. "Comparison of Kinematic and Kinetic Parameters Calculated Using a Cluster-Based Model and Vicon's Plug-in Gait." *Proceedings of the Institution of Mechanical Engineers, Part H: Journal of Engineering in Medicine* 228 (2): 206–10.
- Dunne, Anthony, and Fiona Raby. 2013. *Speculative Everything: Design, Fiction, and Social Dreaming*. Cambridge, MA: MIT Press. <http://muse.jhu.edu/book/28148>
- Eder, Maciej. 2015. "Does Size Matter? Authorship Attribution, Small Samples, Big Problem." *Literary and Linguistic Computing* 30 (2): 167–82.
- Eder, Maciej, Jan Rybicki, and Mike Kestemont. 2016. "Stylometry with R: A Package for Computational Text Analysis." *R Journal* 8 (1): 107–21. <https://doi.org/10.1093/llc/fqt066>
- Edwards, Paul K., Steve Vincent, and Joe O'Mahoney. 2014. "Concluding Comments." In *Studying Organizations Using Critical Realism: A Practical Guide*, edited by Paul

- Edwards, Joe O'Mahoney, and Steve Vincent. Oxford: Oxford University Press. <http://www.oxfordscholarship.com/view/10.1093/acprof:oso/9780199665525.001.0001/acprof-9780199665525-chapter-17>
- Eijk, Gwen van. 2017. "Socioeconomic Marginality in Sentencing: The Built-in Bias in Risk Assessment Tools and the Reproduction of Social Inequality." *Punishment & Society* 19 (4): 463–81. <https://doi.org/10.1177/1462474516666282>
- El Raheb, Katerina, and Yannis E. Ioannidis. 2014. "From Dance Notation to Conceptual Models: A Multilayer Approach." In *MOCO*. <https://doi.org/10.1145/2617995.2618000>
- Elder-Vass, Dave. 2012. *The Reality of Social Construction*. Cambridge.
- Elhayek, Ahmed, Onorina Kovalenko, Pramod Murthy, Jameel Malik, and Didier Stricker. 2018. "Fully Automatic Multi-Person Human Motion Capture for VR Applications." In *Virtual Reality and Augmented Reality*, edited by Patrick Bourdot, Sue Cobb, Victoria Interrante, Hirokazu Kato, and Didier Stricker, 28–47. Lecture Notes in Computer Science. Springer International.
- Eliot, Simon. 2002. "Very Necessary but Not Quite Sufficient: A Personal View of Quantitative Analysis in Book History." *Book History* 5: 283–93. <https://www.jstor.org/stable/30228195>
- Eliot, T. S. 1934. *The Rock*. London: Faber and Faber.
- Endings Project Team, The. 2019. "The Endings Project." 2019. <https://projectendings.github.io/>
- Escobar Varela, Miguel. 2015. "Wayang Kontemporer: Innovations in Javanese Wayang Kulit." Singapore: National University of Singapore. <http://cwa-web.org/dissertation/wayang-dis/index.php>
- Escobar Varela, Miguel. 2016. "The Archive as Repertoire: Transience and Sustainability in Digital Archives." *DHQ: Digital Humanities Quarterly* 10 (4).
- Escobar Varela, Miguel. 2017. "From Copper-Plate Inscriptions to Interactive Websites: Documenting Javanese Wayang Theatre." In *Documenting Performance: The Context and Processes of Digital Curation and Archiving*, 203–14. London and New York: Bloomsbury Methuen Drama.
- Escobar Varela, Miguel. 2019. "Towards a Digital, Data-Driven Wayang Kulit Encyclopedia." *Indonesia and the Malay World* 47 (137): 23–46. <https://doi.org/10.1080/13639811.2019.1553382>
- Escobar Varela, Miguel, and Nala H. Lee. 2018. "Language Documentation: A Reference Point for Theatre and Performance Archives?" *International Journal of Performance Arts and Digital Media* 14 (1): 17–33. <https://doi.org/10.1080/14794713.2018.1453242>
- Escobar Varela, Miguel, and Gea Oswah Fatah Parikesit. 2017. "A Quantitative Close Analysis of a Theatre Video Recording." *Digital Scholarship in the Humanities* 32 (2): 276–83. <https://doi.org/10.1093/llc/fqv069>
- Eve, Martin Paul. 2019. *Close Reading with Computers: Textual Scholarship, Computational Formalism, and David Mitchell's Cloud Atlas*. Stanford: Stanford University Press.
- Farcomeni, Alessio. 2017. "Contribution to the Discussion of the Paper by Stefan Wellek: 'A Critical Evaluation of the Current p-Value Controversy.'" *Biometrical Journal* 59 (5): 880–81. <https://doi.org/10.1002/bimj.201700053>
- Fdili Alaoui, Sarah, Jules Françoise, Thecla Schiphorst, Karen Studd, and Frédéric Bevilacqua. 2017. "Seeing, Sensing and Recognizing Laban Movement Qualities." In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, 4009–20. ACM.

- Feinstein, Alvan R. 1971. "XI. Sources of 'Chronology Bias' in Cohort Statistics." *Clinical Pharmacology & Therapeutics* 12 (5): 864–79.
- Fensham, Rachel. 2016. "Searching Movement's History: Digital Dance Archives." In *Transmission in Motion: The Technologizing of Dance*, edited by Maaïke Bleeker, 94–103. Oxon and New York: Routledge.
- Fensham, Rachel. 2019. "Research Methods and Problems." In *The Bloomsbury Companion to Dance Studies*, edited by Sherril Dodds, 35. London: Bloomsbury.
- Ferrer Valls, Teresa. 2019. "Base de Datos de Comedias Mencionadas En La Documentación Teatral 1540–1700." 2019. <http://catcom.uv.es/consulta/>
- Feynman, Richard. 1955. "The Value of Science." *Engineering and Science* 19 (3): 13–15.
- Feynman, Richard. 1985. "Surely You're Joking, Mr. Feynman!": *Adventures of a Curious Character*. Edited by Ralph Leighton and Edward Hutchings. New York: W. W. Norton.
- Feynman, Richard. 2017. *The Character of Physical Law*. Edited by Alan Sleath. Cambridge, MA: MIT Press.
- Fischer, Frank, Matthias Göbel, Dario Kampkaspar, Christopher Kittel, and Peer Trilcke. 2017. "Network Dynamics, Plot Analysis: Approaching the Progressive Structuration of Literary Texts." In *Digital Humanities 2017 Conference Abstracts*. Montréal.
- Fischer, Frank, Peer Trilcke, Carsten Milling, and Daniil Skorinkin. 2018. "To Catch A Protagonist: Quantitative Dominance Relations In German-Language Drama (1730–1930)." In *Digital Humanities 2018 Conference Abstracts*. Mexico City.
- Fischer-Lichte, Erika, Ramona Thomasius, and Minou Arjomand. 2014. *The Routledge Introduction to Theatre and Performance Studies*. Routledge.
- Flanders, Julia, and Fotis Jannidis. 2015. "Data Modeling." In *A New Companion to Digital Humanities*, by Susan Schreibman, Ray Siemens, and John Unsworth, 229–37. John Wiley & Sons. <https://doi.org/10.1002/9781118680605.ch16>
- Forsythe, William, Maria Palazzi, Norah Zuniga Shaw, and Scott deLahunta. 2009. "Synchronous Objects for One Flat Thing, Reproduced." Website Installation or On Line Resource. Columbus, Ohio.
- Foucault, Michel, James D. Faubion, and Robert Hurley. 1998. *Aesthetics, Method, and Epistemology*. Vol. 2. New York: New Press. <http://www.yorku.ca/rajagopa/documents/Topic-conceptualthemes.docx>
- Franco, Israel. 2020. *Reseña Histórica Del Teatro En México 2.0–2.1. Sistema de Información de La Crítica Teatral (Historical Theater Reviews in Mexico 2.0–2.1, Theater Review Information System)*. Accessed May 21, 2020. <http://criticateatral2021.org/>
- Friendly, Michael, and Daniel J. Denis. 2001. "Milestones in the History of Thematic Cartography, Statistical Graphics, and Data Visualization." <http://datavis.ca/milestones/>
- Friendly, Michael, and Daniel Denis. 2005. "The Early Origins and Development of the Scatterplot." *Journal of the History of the Behavioral Sciences* 41 (2): 103–30. <https://doi.org/10.1002/jhbs.20078>
- Gaver, Bill, Tonny Dunne, and Elena Pacenti. 1999. "Cultural Probes." *Interactions* 6 (1). <https://interactions.acm.org/archive/view/jan.-feb.-1999/design-cultural-probes1>
- Geertz, Clifford. 1973. *The Interpretation of Cultures: Selected Essays*. New York: Basic Books.
- Geertz, Clifford. 1974. "'From the Native's Point of View': On the Nature of Anthropological Understanding." *Bulletin of the American Academy of Arts and Sciences* 28 (1): 26–45. <https://doi.org/10.2307/3822971>

- Gelman, Andrew, and Antony Unwin. 2013. "Infovis and Statistical Graphics: Different Goals, Different Looks." *Journal of Computational and Graphical Statistics* 22 (1): 2–28. <https://doi.org/10.1080/10618600.2012.761137>
- Gil, Alex. 2015. "The User, the Learner and the Machines We Make · Minimal Computing." 2015. <http://go-dh.github.io/mincomp/thoughts/2015/05/21/user-vs-learner/>
- Gilbert, Richard O. 1987. *Statistical Methods for Environmental Pollution Monitoring*. New York: Van Nostrand Reinhold.
- Gitelman, Lisa. 2013. "Raw Data" Is an Oxymoron. Infrastructures Series. Cambridge, MA: MIT Press.
- Gleiser, Pablo M, and Leon Danon. 2003. "Community Structure in Jazz." *Advances in Complex Systems* 6 (4): 565–73.
- Gonzalez, Anita, Clara McClenon, Katy Robinson, Allyson Mackay, Stacey Bishop, Amana Kazkazi, Justin Joque, et al. n.d. "19th Century Acts." Accessed March 23, 2019. <http://19thcenturyacts.com/>
- Google Developers. n.d. "Important Updates | Google Maps Platform." Accessed December 18, 2018. <https://developers.google.com/maps/billing/important-updates>
- Gorski, Philip S. 2013. "What Is Critical Realism? And Why Should You Care?" *Contemporary Sociology* 42 (5): 658–70. <https://doi.org/10.1177/0094306113499533>
- Gottschall, Jonathan. 2008. *Literature, Science, and a New Humanities*. First Edition. New York: Palgrave Macmillan.
- Graham, Paul. 2003. "Hackers and Painters." <http://www.paulgraham.com/hp.html>
- Graham, Ronald L., Bruce L. Rothschild, and Joel H. Spencer. 1990. *Ramsey Theory*. Wiley-Interscience Series in Discrete Mathematics and Optimization. Wiley.
- Gray, Jonathan, Liliana Bounegru, Stefania Milan, and Paolo Ciuccarelli. 2016. "Ways of Seeing Data: Toward a Critical Literacy for Data Visualizations as Research Objects and Research Devices." In *Innovative Methods in Media and Communication Research*, edited by Sebastian Kubitschko and Anne Kaun, 227–51. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-40700-5_12
- Groff, Ruth. 2004. *Critical Realism, Post-Positivism, and the Possibility of Knowledge*. Vol. 11. Routledge Studies in Critical Realism. London and New York: Routledge.
- Guest, Ann Hutchinson. 1998. *Choreo-Graphics: A Comparison of Dance Notation Systems from the Fifteenth Century to the Present*. Amsterdam: Psychology Press.
- Hachimura, Kozaburo. 2006. "Digital Archiving of Dancing." *Review of the National Center for Digitization* 8: 51–66.
- Hachimura, Kozaburo, and Minako Nakamura. 2006. "An XML Representation of Labnotation, LabanXML, and Its Implementation on the Notation Editor LabanEditor2." *Преглед НИЦД* 9: 47–51.
- Hajian, Sara, Francesco Bonchi, and Carlos Castillo. 2016. "Algorithmic Bias." In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2125–26. KDD '16. New York: ACM. <https://doi.org/10.1145/2939672.2945386>
- Hallam, Julia, and Les Roberts. 2013. *Locating the Moving Image: New Approaches to Film and Place*. Bloomington: Indiana University Press. <http://muse.jhu.edu/book/27089>
- Hardjowirogo. 1948. *Sejarah Wayang Purwa*. Jakarta: Balai Pustaka.
- Harrell, D. Fox. 2013. *Phantasmal Media: An Approach to Imagination, Computation, and Expression*. Cambridge, MA: MIT Press.

- Hecht, Angelika, Frank Fischer, Anika Schultz, Christopher Kittel, Elisa Beshero-Bondar, Steffen Martus, Peer Trilcke, Jana Wolf, Ingo Börner, and Daniil Skorinkin. 2018. "Brecht Beats Shakespeare! A Card-Game Intervention Revolving Around the Network Analysis of European Drama." In *Digital Humanities 2018 Conference Abstracts*, 595–96. Mexico City: ADHO.
- Hemispheric Institute. 2008. "Hemispheric Institute Digital Video Library." 2008. www.hemisphericinstitute.org/eng/hidvl/
- Hemispheric Institute. 2020. "Gesture—HemiPress." 2020. <https://hemi.press/gesture/>
- Hepworth, Katherine, and Christopher Church. 2019. "Racism in the Machine: Visualization Ethics in Digital Humanities Projects." *Digital Humanities Quarterly* 12 (4).
- Hernandez-Barraza, Luis, Chen-Hua Yeow, and Miguel Escobar Varela. 2019. "The Biomechanics of Character Types in Javanese Dance." *Journal of Dance Medicine & Science* 23 (3): 104–11.
- Higgins, Sarah. 2008. "The DCC Curation Lifecycle Model." *International Journal of Digital Curation* 3 (1).
- Hiippala, Tuomo. 2020. "A Multimodal Perspective on Data Visualization." In *Data Visualization in Society*, edited by Martin Engebretsen and Helen Kennedy, 277–94. Amsterdam: Amsterdam University Press. <https://doi.org/10.2307/j.ctvzgb8c7.23>
- Hirsch, Brett D., and Janelle Jenstad. 2016. "Beyond the Text: Digital Editions and Performance." *Shakespeare Bulletin* 34 (1): 107.
- Hogan, Trevor. 2015. "Tangible Data, a Phenomenology of Human-Data Relations." In *Proceedings of the Ninth International Conference on Tangible, Embedded, and Embodied Interaction*, 425–28. TEI '15. New York: ACM. <https://doi.org/10.1145/2677199.2691601>
- Holledge, Julie, Jonathan Bollen, Frode Helland, and Joanne Tompkins. 2016. *A Global Doll's House: Ibsen and Distant Visions*. Palgrave Studies in Performance and Technology. London: Palgrave Macmillan.
- Hughes, Amy E., and Naomi J. Stubbs. 2018. *A Player and a Gentleman: The Diary of Harry Watkins, Nineteenth-Century U.S. American Actor*. Ann Arbor: University of Michigan Press.
- Hunter, John D. 2007. "Matplotlib: A 2D Graphics Environment." *Computing in Science & Engineering* 9 (3): 90.
- Hussain, Md. Manjurul, and Ishtiaq Mahmud. 2019. "PyMannKendall: A Python Package for Non Parametric Mann Kendall Family of Trend Tests." *Journal of Open Source Software* 4 (39): 1556. <https://doi.org/10.21105/joss.01556>
- Ichikawa, Tomoko. 2016. "Visualization as Experience." *Digital Studies/Le Champ Numérique* 5 (3). <https://doi.org/10.16995/dscn.32>
- Ioannides, Marinos, Eleanor Fink, Antonia Moropoulou, Monika Hagedorn-Saupe, Antonella Fresa, Gunnar Liestøl, Vlatka Rajcic, and Pierre Grussenmeyer. 2016. *Digital Heritage. Progress in Cultural Heritage: Documentation, Preservation, and Protection: 6th International Conference, EuroMed 2016, Nicosia, Cyprus, October 31—November 5, 2016, Proceedings*. Cham, Switzerland: Springer.
- Ioannidis, John P. A. 2005. "Why Most Published Research Findings Are False." *PLoS Medicine* 2 (8). <https://doi.org/10.1371/journal.pmed.0020124>
- Ioannidis, John P. A. 2014. "How to Make More Published Research True." *PLOS Medicine* 11 (10): e1001747. <https://doi.org/10.1371/journal.pmed.1001747>
- Jackson, Shannon. 2004. *Professing Performance: Theatre in the Academy from Philology to Performativity*. Cambridge: Cambridge University Press.

- Jannidis, Fotis. 2019. "On the Perceived Complexity of Literature. A Response to Nan Z. Da." *Journal of Cultural Analytics Blog* (blog). June 17, 2019. <https://culturalanalytics.org/2019/06/on-the-perceived-complexity-of-literature-a-response-to-nan-z-da/>
- Jansen, Yvonne, Pierre Dragicevic, Petra Isenberg, Jason Alexander, Abhijit Karnik, Johan Kildal, Sriram Subramanian, and Kasper Hornbæk. 2015. "Opportunities and Challenges for Data Physicalization." In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 3227–36. CHI '15. New York: ACM. <https://doi.org/10.1145/2702123.2702180>
- Jenstad, Janelle. 2011. "Using Early Modern Maps in Literary Studies: Views and Caveats from London." In *GeoHumanities: Art, History, Text at the Edge of Place*, by Michael Dear, Jim Ketchum, Sarah Luria, and Doug Richardson. London and New York: Routledge. <https://doi.org/10.4324/9780203839270-22>
- Jessop, Martyn. 2008. "The Inhibition of Geographical Information in Digital Humanities Scholarship." *Literary and Linguistic Computing* 23 (1): 39–50. <https://doi.org/10.1093/lc/fqm041>
- Jockers, Matthew Lee. 2013. *Macroanalysis. Topics in the Digital Humanities*. Baltimore: University of Illinois Press.
- Jockers, Matthew Lee. 2014. *Text Analysis with R for Students of Literature. Quantitative Methods in the Humanities and Social Sciences*. Cham: Springer.
- Juola, Patrick. 2018. "Large-Scale Accuracy Benchmark Results for Juola's Authorship Verification Protocols." In *Digital Humanities 2018, DH 2018, Book of Abstracts, El Colegio de México, UNAM, and RedHD, Mexico City, Mexico, June 26–29, 2018*, 411. <https://dh2018.adho.org/en/large-scale-accuracy-benchmark-results-for-juolas-authorship-verification-protocols/>
- Kaepler, Adrienne L. 1978. "Dance in Anthropological Perspective." *Annual Review of Anthropology* 7: 31.
- Kagan, Jerome. 2009. *The Three Cultures*. Cambridge University Press.
- Karsdorp, F., M. Kestemont, C. Schöch, and A. P. J. van den Bosch. 2015. "The Love Equation: Computational Modeling of Romantic Relationships in French Classical Drama." In *Proceedings of the 6th Workshop on Computational Models of Narrative (CMN-2015)*, Atlanta, GA. <https://repository.ubn.ru.nl/handle/2066/142767>
- Karsdorp, Folgert, Enrique Manjavacas, Lauren Fonteyn, and Mike Kestemont. 2020. "Classifying Evolutionary Forces in Language Change Using Neural Networks." *Evolutionary Human Sciences* 2. <https://doi.org/10.1017/ehs.2020.52>
- Kealiinohomoku, Joann. 1974. "Dance Culture as a Microcosm of Holistic Culture." In *New Dimensions in Dance Research: Anthropology and Dance (The American Indians)*, 99–106. Tucson: University of Arizona, Committee on Research in Dance.
- Kendall, Maurice. 1975. *Multivariate Analysis*. London: Charles Griffin.
- Kennedy, Helen, and Martin Engerbretsen. 2020. "Introduction." In *Data Visualization in Society*, edited by Helen Kennedy and Martin Engerbretsen, 19–32. Amsterdam: Amsterdam University Press. <https://doi.org/10.2307/j.ctvzgb8c7.7>
- Kerr, Norbert L. 1998. "HARKing: Hypothesizing after the Results Are Known." *Personality and Social Psychology Review* 2 (3): 196–217.
- Kershaw, Baz, and Helen Nicholson. 2011. *Research Methods in Theatre and Performance*. Edinburgh: Edinburgh University Press.
- Khmelev, Dmitri V., and Fiona J. Tweedie. 2001. "Using Markov Chains for Identifica-

- tion of Writer.” *Literary and Linguistic Computing* 16 (3): 299–307. <https://doi.org/10.1093/lc/16.3.299>
- Kitchin, Rob, and Martin Dodge. 2014. *Code/Space: Software and Everyday Life*. Cambridge, MA: MIT Press.
- Klenotic, Jeffrey. 2011. “Putting Cinema History on the Map: Using GIS to Explore the Spatiality of Cinema.” In *Explorations in New Cinema History: Approaches and Case Studies*, by Richard Maltby, Daniel Biltereyst, and Philippe Meers, 58–84. Oxford.
- Knigge, LaDona, and Meghan Cope. 2016. “Grounded Visualization: Integrating the Analysis of Qualitative and Quantitative Data through Grounded Theory and Visualization.” *Environment and Planning A*, July. <https://doi.org/10.1068/a37327>
- Knoke, David, and Song Yang. 2008. *Social Network Analysis*. Vol. 154. Quantitative Applications in the Social Sciences. Los Angeles: SAGE.
- Koene, Ansgar. 2017. “Algorithmic Bias: Addressing Growing Concerns [Leading Edge].” *IEEE Technology and Society Magazine* 36 (2): 31–32. <https://doi.org/10.1109/MTS.2017.2697080>
- Koutedakis, Yiannis, Emmanuel O. Owolabi, and Margo Apostolos. 2008. “Dance Biomechanics: A Tool for Controlling Health, Fitness, and Training.” *Journal of Dance Medicine & Science: Official Publication of the International Association for Dance Medicine & Science* 12 (3): 83–90.
- Kozel, Susan. 2007. *Closer: Performance, Technologies, Phenomenology*. Cambridge, MA: MIT Press.
- Kramnick, Jonathan. 2018. “The Interdisciplinary Delusion.” *The Chronicle of Higher Education*, October 11, 2018. <https://www.chronicle.com/article/The-Interdisciplinary-Delusion/244772>
- Kretzschmar, W. A. 2013. “GIS for Language and Literary Study.” In *Literary Studies in the Digital Age: An Evolving Anthology*, edited by R. Siemens and K. Price. New York: Modern Language Association. <http://dx.doi.org/10.1632/ltsa.2013.0>
- Krishan, Sanjay, ed. 1997. *9 Lives, 10 Year of Singapore Theatre 1987–1997*. Singapore: First Printers.
- Kulkarni, Vivek, Rami Al-Rfou, Bryan Perozzi, and Steven Skiena. 2015. “Statistically Significant Detection of Linguistic Change.” In *Proceedings of the 24th International Conference on World Wide Web*, 625–35. International World Wide Web Conferences Steering Committee.
- Kurath, Gertrude P. 1952. “A Choreographic Questionnaire.” *Midwest Folklore* 2 (1): 53–55.
- Kwa, Chunglin. 2011. *Styles of Knowing: A New History of Science from Ancient Times to the Present*. University of Pittsburgh Press. <https://muse-jhu-edu.libproxy1.nus.edu.sg/book/1988>
- Lakens, Daniel. 2013. “Calculating and Reporting Effect Sizes to Facilitate Cumulative Science: A Practical Primer for t-Tests and ANOVAs.” *Frontiers in Psychology* 4. <https://doi.org/10.3389/fpsyg.2013.00863>
- Larasati, Rachmi Diyah. 2013. *The Dance That Makes You Vanish: Cultural Reconstruction in Post-Genocide Indonesia*. Difference Incorporated. Minneapolis: University of Minnesota Press.
- Larson, Jeff, and Julia Angwin. 2016. “How We Analyzed the COMPAS Recidivism Algorithm.” *Text/html*. ProPublica. May 23, 2016. <https://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm>

- Leino, Antti, and Saara Hyvönen. 2008. "Comparison of Component Models in Analysing the Distribution of Dialectal Features." *International Journal of Humanities and Arts Computing* 2 (1–2): 173–87. <https://doi.org/10.3366/E1753854809000378>
- Leiter, Samuel L. 2012. "Is the Onnagata Necessary?" *Asian Theatre Journal* 29 (1): 112–21. <https://doi.org/10.1353/atj.2012.0028>
- Leonhardt, Nic. 2014. "Digital Humanities and the Performing Arts: Building Communities, Creating Knowledge." In *Keynote Address, SIBMAS/TLA Conference, New York City*. Vol. 12.
- Li, William, and Philippe Pasquier. 2016. "Automatic Affect Classification of Human Motion Capture Sequences in the Valence-Arousal Model." In *Proceedings of the 3rd International Symposium on Movement and Computing*, 15:1–15:8. MOCO '16. New York: ACM. <https://doi.org/10.1145/2948910.2948936>
- Liao, Han-teng, and Thomas Petzold. 2010. "Analysing Geo-Linguistic Dynamics of the World Wide Web: The Use of Cartograms and Network Analysis to Understand Linguistic Development in Wikipedia." *Cultural Science* 3 (2).
- Lindinger, Christian, David Labbe, Philippe Pollien, Andreas Rytz, Marcel A. Juillerat, Chahan Yerezian, and Imre Blank. 2008. "When Machine Tastes Coffee: Instrumental Approach To Predict the Sensory Profile of Espresso Coffee." *Analytical Chemistry* 80 (5): 1574–81. <https://doi.org/10.1021/ac702196z>
- Liu, Alan. 2012. "The State of the Digital Humanities: A Report and a Critique." *Arts and Humanities in Higher Education* 11 (1–2): 8–41. <https://doi.org/10.1177/1474022211427364>
- Luksys, Donatas, and Julius Griskevicius. 2016. "Quantitative Assessment of Dance Therapy Influence On the Parkinson's Disease Patients' Lower Limb Biomechanics." *Mokslas* 8 (6): 583–86. <http://dx.doi.org.libproxy1.nus.edu.sg/10.3846/mla.2016.978>
- Madison, D. Soyini. 2005. *The SAGE Handbook of Performance Studies*. SAGE Publications.
- Malakar, Sourav, Saptarsi Goswami, and Amlan Chakrabarti. 2018. "An Online Trend Detection Strategy for Twitter Using Mann–Kendall Non-Parametric Test." In *Industry Interactive Innovations in Science, Engineering and Technology*, edited by Swapan Bhattacharyya, Sabyasachi Sen, Meghamala Dutta, Papun Biswas, and Himadri Chattopadhyay, 185–93. Singapore: Springer Singapore.
- Malmstrom, Carl, Yaying Zhang, Philippe Pasquier, Thecla Schiphorst, and Lyn Bartram. 2016. "MoComp: A Tool for Comparative Visualization Between Takes of Motion Capture Data." In *Proceedings of the 3rd International Symposium on Movement and Computing*, 11:1–11:8. MOCO '16. New York: ACM. <https://doi.org/10.1145/2948910.2948932>
- Maniatis, Petros, Mema Roussopoulos, Thomas J. Giuli, David SH Rosenthal, Mary Baker, and Mary Baker. 2005. "The LOCKSS Peer-to-Peer Digital Preservation System." *ACM Transactions on Computer Systems (TOCS)* 23 (1): 2–50.
- Mann, Henry B. 1945. "Nonparametric Tests against Trend." *Econometrica: Journal of the Econometric Society*, 245–59.
- Manovich, Lev. 2000. *The Language of New Media*. Leonardo. Cambridge, MA: MIT Press.
- Manovich, Lev. 2011. "What Is Visualisation?" *Visual Studies* 26 (1): 36–49. <https://doi.org/10.1080/1472586X.2011.548488>
- Manovich, Lev. 2013. "Visualizing Vertov." *Russian Journal of Communication* 5 (1): 44–55.
- Manovich, Lev. 2020. *Cultural Analytics*. Cambridge, MA: MIT Press.
- Manzetti, Maria Cristina. 2016. "3D Visibility Analysis as a Tool to Validate Ancient

- Theatre Reconstructions: The Case of the Large Roman Theatre of Gortyn." *Virtual Archaeology Review* 7 (15): 36–43. <https://doi.org/10.4995/var.2016.5922>
- Manzor, Lillian, Kyle Rimkus, and Mitsunori Ogihara. 2013. "Cuban Theater Digital Archive: A Multimodal Platform for Theater Documentation and Research." In *International Conference on Information Technologies for Performing Arts, Media Access, and Entertainment*, 138–50. Springer.
- Marasini, Donata, Piero Quatto, and Enrico Ripamonti. 2016. "The Use of P-Values in Applied Research: Interpretation and New Trends." *Statistica; Bologna* 76 (4): 315–25. <http://dx.doi.org.libproxy1.nus.edu.sg/10.6092/issn.1973-2201/6439>
- Marsden, Peter. 2005. "Recent Developments in Network Measurement." In *Models and Methods in Social Network Analysis*, edited by Peter J. Carrington, John Scott, and Stanley Wasserman. Vol. 28. *Structural Analysis in the Social Sciences*. Cambridge: Cambridge University Press.
- Mayer-Schönberger, Viktor, and Kenneth Cukier. 2013. *Big Data: A Revolution That Will Transform How We Live, Work, and Think*. Boston: Houghton Mifflin Harcourt.
- Mee, Erin B. 2013. "Hearing the Music of the Hemispheres." *TDR: The Drama Review* 57 (3): 148–50.
- Mee, Erin B. 2018. "Born-Digital Scholarship." *TDR: The Drama Review* 62 (3): 8–9.
- Meeks, Elijah, and Scott Weingart. 2013. "The Digital Humanities Contribution to Topic Modeling." *Journal of Digital Humanities*, April 9. <http://journalofdigitalhumanities.org/2-1/dh-contribution-to-topic-modeling/>
- Meister J. C, Petris, M, Gius, E, and Jacke, J. 2016. "CATMA 5.0 [Software for Text Annotation and Analysis]." <http://catma.de/>
- Michel, Jean-Baptiste, Yuan Kui Shen, Aviva P. Aiden, Adrian Veres, Matthew K. Gray, Joseph P. Pickett, Dale Hoiberg, et al. 2011. "Quantitative Analysis of Culture Using Millions of Digitized Books." *Science* 331 (6014): 176–82. <https://doi.org/10.1126/science.1199644>
- Miller, Derek. 2017. "Average Broadway." *Theatre Journal* 68 (4): 529–53. <https://doi.org/10.1353/tj.2016.0105>
- Minimal Computing Working Group. 2015. "About · Minimal Computing." 2015. <http://go-dh.github.io/mincomp/about/>
- Misi, Gábor. 1983. "Formal Methods in Form Analysis of Transylvanian Male Solo Dances." *Dance Studies* 7: 21–56.
- Misi, Gábor. 2005. "Labanatory." 2005. <http://www.labanatory.com/eng/software.html>
- Misi, Gábor. 2008. "An Algebraic Representation of Labanotation for Retrieval and Other Operations." In *Proceedings of the 25th Biennial Conference of ICKL, Mexico City*, 143–60.
- Moere, A. V. 2008. "Beyond the Tyranny of the Pixel: Exploring the Physicality of Information Visualization." In *2008 12th International Conference Information Visualisation*, 469–74. <https://doi.org/10.1109/IV.2008.84>
- Molloy, Laura. 2014. "Digital Curation Skills in the Performing Arts—an Investigation of Practitioner Awareness and Knowledge of Digital Object Management and Preservation." *International Journal of Performance Arts and Digital Media* 10 (1): 7–20. <https://doi.org/10.1080/14794713.2014.912496>
- Moran, Patrick A. P. 1948. "The Interpretation of Statistical Maps." *Journal of the Royal Statistical Society. Series B (Methodological)* 10 (2): 243–51.
- Moreno, J. L. 1960. *The Sociometry Reader*. Glencoe: Free Press.

- Moretti, Franco. 2000. "The Slaughterhouse of Literature." *Modern Language Quarterly* 61 (1): 207.
- Moretti, Franco. 2007. *Graphs, Maps, Trees*. Verso.
- Moretti, Franco. 2009. "Style, Inc.: Reflections on Seven Thousand Titles (British Novels, 1740–1850)." *Critical Inquiry* 36 (1): 134.
- Moretti, Franco. 2011. "Network Theory, Plot Analysis." *New Left Review*, no. 68 (March): 80.
- Morgan, Susan. 2018. "Fake News, Disinformation, Manipulation and Online Tactics to Undermine Democracy." *Journal of Cyber Policy* 3 (1): 39–43. <https://doi.org/10.1080/23738871.2018.1462395>
- Mowat, Barbara, Paul Werstine, Michael Poston, and Rebecca Niles, eds. n.d. *Hamlet*. Washington, DC: Folger Shakespeare Library. Accessed 13 April 2020. www.folgerdigitaltexts.org
- Mrázek, Jan. 2005. "Masks and Selves in Contemporary Java: The Dances of Didik Nini Thowok." *Journal of Southeast Asian Studies* 36 (2): 249–79. <https://doi.org/10.1017/S0022463405000160>
- Mrázek, Jan. 2019. *Wayang and Its Doubles: Javanese Puppet Theatre, Television and the Internet*. Singapore: NUS Press.
- Mullainathan, Sendhil, and Ziad Obermeyer. 2017. "Does Machine Learning Automate Moral Hazard and Error?" *American Economic Review* 107 (5): 476–80. <https://doi.org/10.1257/aer.p20171084>
- Mullaney, Thomas Shawn, Christian Henriot, Jeffrey Snyder-Reinke, David William McClure, and Glen Worthey. 2019. *The Chinese Deathscape: Grave Reform in Modern China*. Stanford: Stanford University Press.
- Nakamura, M. 2017. "The Postures and Movements of Balinese Dance." In *2017 International Conference on Culture and Computing (Culture and Computing)*, 63–64. <https://doi.org/10.1109/Culture.and.Computing.2017.37>
- NIST/SEMATECH. 2003. "NIST/SEMATECH e-Handbook of Statistical Methods." 2003. <https://www.itl.nist.gov/div898/handbook/>
- Nowiskie, Bethany, David McClure, Wayne Graham, Adam Soroka, Jeremy Boggs, and Eric Rochester. 2013. "Geo-Temporal Interpretation of Archival Collections with Neatline." *Literary and Linguistic Computing* 28 (4): 692–99. <https://doi.org/10.1093/llc/fqt043>
- NYPL/Zooniverse. n.d. "Ensemble@Yale." Accessed August 3, 2018. <http://ensemble.yale.edu/#/about>
- O'Donoghue, Peter. 2010. *Research Methods for Sports Performance Analysis*. Oxon and New York: Routledge.
- O'Mahoney, Joe, and Steve Vincent. 2014. "Critical Realism as an Empirical Project." In *Studying Organizations Using Critical Realism: A Practical Guide*, edited by Paul Edwards, Joe O'Mahoney, and Steve Vincent. Oxford: Oxford University Press. <http://www.oxfordscholarship.com/view/10.1093/acprof:oso/9780199665525.001.0001/acprof-9780199665525-chapter-1>
- O'Neil, Cathy. 2016. *Weapons of Math Destruction*. New York: Crown.
- Oakes, Michael P. 2014. *Literary Detective Work on the Computer*. Amsterdam and Philadelphia: John Benjamins Publishing.
- Oakes, Michael P. 2017. "Computer Stylometry of C. S. Lewis's *The Dark Tower* and Related Texts." *Digital Scholarship in the Humanities*. <https://doi.org/10.1093/llc/fqx043>

- Ohge, Christopher, Steven Olsen-Smith, Elisa Barney Smith, and Adam Brimhall. 2018. "At the Axis of Reality: Melville's Marginalia in The Dramatic Works of William Shakespeare." *Leviathan* 20 (2): 37–67.
- Olson, Mark J. V. 2013. "Hacking the Humanities: Twenty-First-Century Literacies and the 'Becoming-Other' of the Humanities." In *Humanities in the Twenty-First Century: Beyond Utility and Markets*, edited by Eleonora Belfiore and Anna Upchurch, 237–50. London: Palgrave Macmillan. https://doi.org/10.1057/9781137361356_13
- Open Source Community. (2010) 2018. *JavaScript Library for Mobile-Friendly Interactive Maps: Leaflet/Leaflet* (Version 1.3.4). JavaScript. Leaflet. <https://github.com/Leaflet/Leaflet>
- OpenStreetMap Community. 2004. "OpenStreetMap." <https://www.openstreetmap.org/>
- Palazzi, Maria, and Norah Zuniga Shaw. 2009. "Synchronous Objects for One Flat Thing, Reproduced." In *SIGGRAPH: International Conference and Exhibition on Computer Graphics and Interactive Techniques*, 2. New Orleans: ACM.
- Pandya, Ami. 2003. "Notation System in Indian Classical Dance—Bharatanatyam." PhD diss., India: Maharaja Sayajirao University of Baroda (India). <http://search.proquest.com/docview/1771297342/citation/1957BE1310BF4463PQ/1>
- Park, Gyeong-Mi, Sung-Hwan Kim, Hye-Ryeon Hwang, and Hwan-Gue Cho. 2013. "Complex System Analysis of Social Networks Extracted from Literary Fictions." *International Journal of Machine Learning and Computing* 3 (1): 107.
- Park, Seung-Bo, Kyeong-Jin Oh, and Geun-Sik Jo. 2012. "Social Network Analysis in a Movie Using Character-Net." *Multimedia Tools and Applications* 59 (2): 601–27.
- Parry, Kyle. 2019. "Reading for Enactment: A Performative Approach to Digital Scholarship and Data Visualization." In *Debates in the Digital Humanities 2019*, by Matthew K. Gold and Lauren F. Klein. *Debates in the Digital Humanities*. Minneapolis: University of Minnesota Press. <https://dhdebates.gc.cuny.edu/read/untitled-f2acf72c-a469-49d8-be35-67f9ac1e3a60/section/a2a6a192-f04a-4082-afaa-97c76a75b21c#ch24>
- Pavis, Patrice. 1996. *The Intercultural Performance Reader*. Routledge.
- Peer, Willie van, Jèmeljan Hakemulder, and Sonia Zyngier. 2007. *Muses and Measures*. Newcastle: Cambridge Scholars.
- Pendón Martínez, Alberto, and Gema Bueno de la Fuente. 2017. "Description Models for Documenting Performance." In *Documenting Performance: The Context and Processes of Digital Curation and Archiving*, by Toni Sant, 29–46. London and New York: Bloomsbury.
- Phelan, Peggy. 1993. *Unmarked: The Politics of Performance*. London and New York: Routledge.
- Piper, Andrew. 2018. *Enumerations: Data and Literary Study*. Chicago: University of Chicago Press.
- Pollock, Della. 1998. "Performing Writing." In *The Ends of Performance*, by Peggy Phelan and Jill Lane, 73–103. London and New York: New York University Press.
- Pomerantz, Jeffrey. 2015. *Metadata*. The MIT Press Essential Knowledge Series. Cambridge, MA: MIT Press.
- Presner, Todd Samuel, David Shepard, and Yoh Kawano. 2014. *HyperCities*. MetaLAB-projects. Cambridge, MA: Harvard University Press.
- Prince Lab for Digital Humanities. n.d. "English Playbills | Price Lab for Digital

- Humanities.” Accessed August 3, 2018. <https://pricelab.sas.upenn.edu/projects/english-playbills>
- Purwadi. 2013. *Mengenal Gambar Tokoh Wayang Purwa Dan Keterangannya*. Surakarta: Cendrawasih.
- Radway, Janice A. 1991. *Reading the Romance: Women, Patriarchy, and Popular Literature*. Revised edition. Chapel Hill: University of North Carolina Press.
- Rae, Paul. 2018. *Real Theatre: Essays in Experience*. Theatre and Performance Theory. Cambridge: Cambridge University Press. <https://doi.org/10.1017/9781316890752>
- Ramsay, Stephen. 2003. “Toward an Algorithmic Criticism.” *Literary and Linguistic Computing: Journal of the Association for Literary and Linguistic Computing* 18 (2): 167.
- Ramsay, Stephen. 2007. “Algorithmic Criticism.” In *A Companion to Digital Literary Studies*, by Raymond George Siemens and Susan Schreibman. Vol. 50. Blackwell Companions to Literature and Culture. Malden, MA: Blackwell.
- Ramsay, Stephen. 2011. *Reading Machines: Toward an Algorithmic Criticism*. Champaign: University of Illinois Press. <http://www.jstor.org/stable/10.5406/j.ctt1xcmmr>
- Ramsay, Stephen. 2012. “Programming with Humanists: Reflections on Raising an Army of Hacker-Scholars in the Digital Humanities.” In *Digital Humanities Pedagogy*, edited by Brett D. Hirsch, 3:227–40. Practices, Principles and Politics. Open Book Publishers. www.jstor.org/stable/j.ctt5vjtt3.14
- Ramsay, Stephen. 2014. “ENGL 4/878 FAQ.” 2014. http://jetson.unl.edu/syllabi/2014/fall/dh/html/01_overview.html
- Redding, Emma. 2019. “The Expanding Possibilities of Dance Science.” In *The Routledge Companion to Dance Studies*, by Helen Thomas and Stacey Prickett, 56–67. London: Routledge.
- Rees, Chris, and Mark Gatenby. 2014. “Critical Realism and Ethnography.” In *Studying Organizations Using Critical Realism: A Practical Guide*, edited by P. K. Edwards, Joe O’Mahoney, and Steve Vincent. First ed. Oxford: Oxford University Press.
- Rettberg, Jill Walker. 2020. “Ways of Knowing with Data Visualizations.” In *Data Visualization in Society*, edited by Martin Engebretsen and Helen Kennedy, 35–48. Amsterdam: Amsterdam University Press. <https://doi.org/10.2307/j.ctvzgb8c7.8>
- Rey, Sergio J., and Luc Anselin. 2010. “PySAL: A Python Library of Spatial Analytical Methods.” In *Handbook of Applied Spatial Analysis*, 175–93. New York: Springer.
- Ribeiro, Claudia, Rafael Kuffner dos Anjos, and Carla Fernandes. 2017. “Capturing and Documenting Creative Processes in Contemporary Dance.” In *Proceedings of the 4th International Conference on Movement Computing*, 7:1–7:7. MOCO ’17. New York: ACM. <https://doi.org/10.1145/3077981.3078041>
- Ribes, David, and Steven Jackson. 2013. “Data Bite Man: The Work of Sustaining a Long-Term Study.” In “Raw Data” Is an Oxymoron, by Lisa Gitelman, 147–66. *Infrastructures Series*. Cambridge, MA: MIT Press.
- Ripley, B. D. 1979. “Tests of ‘Randomness’ for Spatial Point Patterns.” *Journal of the Royal Statistical Society. Series B (Methodological)* 41 (3): 368–74.
- Roberts, David, and Lance Woodman. 1998. “A Corpus Linguistics Study of the Theatre Review: First Steps.” *Studies in Theatre Production* 18 (1): 6–28. <https://doi.org/10.1080/13575341.1998.10806987>
- Roberts, Les. 2015. “Navigating the ‘Archive City’: Digital Spatial Humanities and Archival Film Practice.” *Convergence* 21 (1): 100–115. <https://doi.org/10.1177/1354856514560310>

- Roberts, Les. n.d. "Liverpool—City in Film." Accessed December 14, 2018. <https://www.google.com/maps/d/viewer?mid=1f8v4bF2fBwk2Z5tjtmWq2QnQOlg>
- Roberts-Smith, Jennifer, Shawn DeSouza-Coelho, Teresa M. Dobson, Sandra Gabrielle, Omar Rodriguez-Arenas, Stan Ruecker, Stéfan Sinclair, et al. 2013. "Visualizing Theatrical Text: From Watching the Script to the Simulated Environment for Theatre (SET)." *Digital Humanities Quarterly* 7 (3).
- Rockwell, Geoffrey, and Stéfan Sinclair. 2016. *Hermeneutica: Computer-Assisted Interpretation in the Humanities*. Cambridge: MIT Press. <http://www.jstor.org/stable/j.ctt1c0gm6h>
- Rogerson, Peter A. 1999. "The Detection of Clusters Using a Spatial Version of the Chi-Square Goodness-of-Fit Statistic." *Geographical Analysis* 31 (1): 130–47. <https://doi.org/10.1111/gean.1999.31.1.130>
- Rosenberg, Daniel. 2013. "Data before the Fact." In "Raw Data" Is an Oxymoron, by Lisa Gitelman, 15–40. *Infrastructures Series*. Cambridge: MIT Press.
- Royce, Anya Peterson. 1977. *The Anthropology of Dance*. Bloomington: Indiana University Press.
- Rozik, Eli. 1999. "The Corporeality of the Actor's Body: The Boundaries of Theatre and the Limitations of Semiotic Methodology." *Theatre Research International* 24 (2): 198–211. <https://doi.org/10.1017/S0307883300020824>
- Ruecker, Stan, Milena Radzikowska, and Stefan Sinclair. 2016. *Visual Interface Design for Digital Cultural Heritage: A Guide to Rich-Prospect Browsing*. Farnham and Burlington: Ashgate.
- Rusbridge, Chris, Peter Burnhill, Seamus Ross, Peter Buneman, David Giarretta, Liz Lyon, and Malcolm Atkinson. 2005. "The Digital Curation Centre: A Vision for Digital Curation." In 2005 IEEE International Symposium on Mass Storage Systems and Technology, 31–41. Baltimore, MD: IEEE.
- Rushkoff, Douglas, and Leland Purvis. 2011. *Program or Be Programmed: Ten Commands for a Digital Age*. Berkeley, CA: Counterpoint.
- Salsburg, David. 2002. *The Lady Tasting Tea*. New York: W. H. Freeman / Owl Book.
- Salt, Barry. 1974. "Statistical Style Analysis of Motion Pictures." *Film Quarterly* 28 (1): 13–22. <https://doi.org/10.2307/1211438>
- Samuels, Lisa, and Jerome McGann. 1999. "Deformance and Interpretation." *New Literary History: A Journal of Theory and Interpretation* 30 (1): 25.
- Sant, Toni. 2014. "Interdisciplinary Approaches to Documenting Performance." *International Journal of Performance Arts and Digital Media* 10 (1): 3–6. <https://doi.org/10.1080/14794713.2014.912495>
- Sant, Toni. 2017. *Documenting Performance: The Context and Processes of Digital Curation and Archiving*. London and New York: Bloomsbury.
- Santos Unamuno, Enrique. 2017. "GIS and Telescopic Reading: Between Spatial and Digital Humanities." *Neohelicon: Acta Comparationis Litterarum Universarum* 44 (1): 65.
- Sayers, Jentry. 2016. "Minimal Definitions · Minimal Computing." 2016. <http://go-dh.github.io/mincomp/thoughts/2016/10/02/minimal-definitions/>
- Schauf, Andrew Johnathan, and Miguel Escobar Varela. 2018. "Searching for Hidden Bridges in Co-Occurrence Networks from Javanese Wayang Kulit." *Journal of Historical Network Research* 2 (1): 26–52. <https://doi.org/10.25517/jhnr.v2i1.42>
- Schechner, Richard. 2013. *Performance Studies*. Edited by Sara Brady. London and New York: Routledge.

- Scheinfeldt, Tom. 2012. "Game Change: Digital Technology and Performative Humanities." *Found History*. February 15, 2012. <http://foundhistory.org/2012/02/game-change-digital-technology-and-performative-humanities/>
- Schich, Maximilian, Christian Huemer, Piotr Adamczyk, Lev Manovich, and Yang-Yu Liu. 2017. "Network Dimensions in the Getty Provenance Index." *ArXiv Preprint ArXiv:1706.02804*.
- Schich, Maximilian, Sune Lehmann, and Juyong Park. 2008. "Dissecting the Canon: Visual Subject Co-Popularity Networks in Art Research." 5th European Conference on Complex Systems, September 3, Jerusalem. <https://archiv.ub.uni-heidelberg.de/artdok/711/>
- Schich, Maximilian, Chaoming Song, Yong-Yeol Ahn, Alexander Mirsky, Mauro Martino, Albert-László Barabási, and Dirk Helbing. 2014. "A Network Framework of Cultural History." *Science* 345 (6196): 558–62. <https://doi.org/10.1126/science.1240064>
- Schmicking, Daniel, and Shaun Gallagher, eds. 2010. *Handbook of Phenomenology and Cognitive Science*. Dordrecht: Springer Netherlands. <https://doi.org/10.1007/978-90-481-2646-0>
- Schöch, Christof. 2016. "Principal Component Analysis for Literary Genre Stylistics." Billet. *The Dragonfly's Gaze* (blog). 2016. <https://dragonfly.hypotheses.org/472>
- Schöch, Christof. 2017. "Topic Modeling Genre: An Exploration of French Classical and Enlightenment Drama." *Digital Humanities Quarterly* 11 (2).
- Sears, Laurie Jo. 1996. *Shadows of Empire: Colonial Discourse and Javanese Tales*. Durham: Duke University Press.
- Septi Rito Tombe. 2017. University of Melbourne Digital Studio, and Social Media and Performance Impact Survey. "Circuit: Mapping Theatre Performances in Victoria." <https://circuit.unimelb.edu.au/>
- Siemens, Raymond G. 2002. "A New Computer-Assisted Literary Criticism?" *Computers and the Humanities* 36 (3): 259–67.
- Simpson, Travis T., Susan Wiesner, and Bradford C. Bennett. 2014. "Dance Recognition System Using Lower Body Movement." *Journal of Applied Biomechanics* 30 (1): 147–53. <https://doi.org/10.1123/jab.2012-0248>
- Snaprud, Mikael, and Andrea Velazquez. 2020. "Accessibility of Data Visualizations:" In *Data Visualization in Society*, edited by Martin Engbreetsen and Helen Kennedy, 111–26. Amsterdam: Amsterdam University Press. <https://doi.org/10.2307/j.ctvzg8c7.13>
- Soedarsono. 1983. "Wayang Wong in the Yogyakarta Kraton: History, Ritual Aspects, Literary Aspects, and Characterization." Ann Arbor: University of Michigan.
- Solichin, Suyanto H., Sumari, Undung Wiyono, and Sri Purwanto. 2017. *Ensiklopedi Wayang Indonesia*. Jakarta: Mitra Sarana Edukasi.
- Spear, Mary Eleanor. 1952. *Charting Statistics*. New York: McGraw-Hill.
- Spence, Paul. 2013. "Teatro clásico y humanidades digitales: el cruce entre método, proceso y nuevas tecnologías." *Teatro de Palabras* 7 (9): 31.
- Splawa-Neyman, Jerzy. 1990. "On the Application of Probability Theory to Agricultural Experiments. Essay on Principles. Section 9." Translated by Dorota M. Dabrowska and T. P. Speed. *Statistical Science*, 465–72.
- Sreenivasan, Sameet. 2013. "Quantitative Analysis of the Evolution of Novelty in Cinema through Crowdsourced Keywords." *Scientific Reports* 3: 2758.

- Streiter, Oliver, Yoann Goudin, Chun (Jimmy) Huang, and Ann Meifang Lin. 2012. "Matching Digital Tombstone Documentation to Unearthed Census Data: Surveying Taiwan's Family Names, Ethnicities and Homelands." *International Journal of Humanities and Arts Computing* 6 (1-2): 57-70. <https://doi.org/10.3366/ijhac.2012.0038>
- Strine, Mary S., Beverly Whitaker Long, and Mary Frances Hopkins. 1990. "Research in Interpretation and Performance Studies: Trends, Issues, Priorities." *Speech Communication: Essays to Commemorate the 75th Anniversary of the Speech Communication Association*, 181-204.
- Su, Wen-Chi. 2019. "Dramaturgy and Technology." In *Dramaturgy and the Human Condition*. Singapore.
- Subyen, Pattarawut, Diego Maranan, Thecla Schiphorst, Philippe Pasquier, and Lyn Bartram. 2011. "EMVIZ: The Poetics of Movement Quality Visualization." In *Proceedings of the International Symposium on Computational Aesthetics in Graphics, Visualization, and Imaging*, 121-28. CAE '11. New York: ACM. <https://doi.org/10.1145/2030441.2030467>
- Sudibyoprono, R. Rio, Suwandono, Dhanisworo, and Mujiyono. 1991. *Ensiklopedi Wayang Purwa*. Jakarta: Balai Pustaka.
- Sudjarwo, Heru S., Sumari, and Undung Wiyono. 2010. *Rupa Dan Karakter Wayang Purwa*. Jakarta: Kakilangit Kencana.
- Suiter, Ted. 2013. "Why 'Hacking'?" In *Hacking the Academy: New Approaches to Scholarship and Teaching from Digital Humanities*, by Daniel J. Cohen and Tom Scheinfeldt. Ann Arbor: University of Michigan Press.
- Suri, M., and S. N. Singh. 2018. "The Role of Big Data in the Media and Entertainment Industry." In *2018 4th International Conference on Computational Intelligence Communication Technology (CICT)*, 1-5. <https://doi.org/10.1109/CICT.2018.8480281>
- Suskin, Steven. 1990. *Opening Night on Broadway: A Critical Quotebook of the Golden Era of the Musical Theatre, Oklahoma! (1943) to Fiddler on the Roof (1964)*. New York: Schirmer Trade Books.
- Taleb, Nassim. 2007. *The Black Swan: The Impact of the Highly Improbable*. London: Allen Lane.
- Tan, Alvin. 2004. "A Necessary Practice." In *Ask Not: The Necessary Stage in Singapore Theatre*, edited by Chong Kee Tan and Tisa Ng, 266-90. Singapore: Times Editions.
- Tan, Chong Kee, and Tisa Ng, eds. 2004. *Ask Not: The Necessary Stage in Singapore Theatre*. Singapore: Times Editions.
- Tan, Kenneth Paul. 2013. "Forum Theater in Singapore: Resistance, Containment, and Commodification in an Advanced Industrial Society." *Positions: East Asia Cultures Critique* 21 (2).
- TextGrid Consortium. 2006. "TextGrid: A Virtual Research Environment for the Humanities." 2014 2006. textgrid.de.
- Thorp, Jer. 2017. "You Say Data, I Say System." *Hacker Noon*, July 13. <https://hackernoon.com/you-say-data-i-say-system-54e84aa7a421>
- Thudt, Alice, Jagoda Walny, Theresia Gschwandtner, Jason Dykes, and John Stasko. 2018. "Exploration and Explanation in Data-Driven Storytelling." In *Data-Driven Storytelling*, edited by Nathalie Henry Riche, Christophe Hurter, Nicholas Diakopoulos, and Sheelagh Cpendale, 59-83. New York: A. K. Peters/CRC Press. <https://doi.org/10.1201/9781315281575-3>

- Törnberg, Anton. 2019. "Abstractions on Steroids: A Critical Realist Approach to Computer Simulations." *Journal for the Theory of Social Behaviour* 49 (1): 127–43. <https://doi.org/10.1111/jtsb.12194>
- Travis, Charles. 2015. "Acts of Perception: Samuel Becket, Time, Space and the Digital Literary Atlas of Ireland, 1922–1949." *International Journal of Humanities and Arts Computing* 9 (2): 219–41. <https://doi.org/10.3366/ijhac.2015.0150>
- Travis, Charles, and Richard Breen. 2017. "Digital Literary Atlas of Ireland." 2017. <https://uploads.knightlab.com/storymapjs/df287c1e4ed312fe68fb5e395ee7d57/digital-literary-atlas-of-ireland-time-map/index.html>
- Trilcke, Peer, and Frank Fischer. n.d. "Dracor.Org." Accessed September 5, 2018. <https://dracor.org/>
- Trilcke, Peer, Frank Fischer, Mathias Göbel, and Dario Kampkaspar. 2015. "Comedy vs. Tragedy: Network Values by Genre." *DLINA: Digital Linterary Analysis* (blog). <https://dlina.github.io/Network-Values-by-Genre/>
- Trilcke, Peer, Frank Fischer, Matthias Göbel, and Dario Kampkaspar. 2016. "Theatre Plays as 'Small Worlds'? Network Data on the History and Typology of German Drama, 1730–1930." In *Digital Humanities 2016 Conference Abstracts*. Krakow, Poland.
- Trilcke, Peer, Frank Fischer, and Dario Kampkaspar. 2015. "Digital Network Analysis of Dramatic Texts." In *Digital Humanities 2015 Conference Abstracts*. Sydney, Australia.
- Tsachor, Rachele P., and Tal Shafir. 2017. "A Somatic Movement Approach to Fostering Emotional Resiliency through Laban Movement Analysis." *Frontiers in Human Neuroscience* 11 (September). <https://doi.org/10.3389/fnhum.2017.00410>
- Tufte, Edward R. 1983. *The Visual Display of Quantitative Information*. Vol. 2. Cheshire, CT: Graphics Press.
- Tukey, John W. 1977. *Exploratory Data Analysis*. Vol. 2. Reading, MA: Addison-Wesley.
- Ubersfeld, Anne, Frank Collins, Paul Perron, and Patrick Debbeche. 1999. *Reading Theatre*. Toronto Studies in Semiotics. University of Toronto Press.
- UiO Universitetet i Oslo. n.d. "IbsenStage." Accessed September 5, 2018. <https://ibsenstage.hf.uio.no/>
- Ulmer, Gregory L. 1994. *Heuretics: The Logic of Invention*. Baltimore, MD: Johns Hopkins University Press.
- Underwood, Ted. 2017. "A Genealogy of Distant Reading." *DHQ: Digital Humanities Quarterly* 11 (2).
- Underwood, Ted. 2019a. *Distant Horizons: Digital Evidence and Literary Change*. Chicago: University of Chicago Press.
- Underwood, Ted. 2019b. "Dear Humanists: Fear Not the Digital Revolution." *Chronicle of Higher Education*, March 27. <https://www.chronicle.com/article/Dear-Humanists-Fear-Not-the/245987>
- UPenn Libraries. n.d. "19th Century Playbills, 1803–1939." Accessed August 3, 2018. http://dla.library.upenn.edu/dla/pacscl/detail.html?id=PACSCL_FLP_FLPTHCPLA_YBILL
- Uzzi, Brian, and Jarrett Spiro. 2005. "Collaboration and Creativity: The Small World Problem." *American Journal of Sociology* 111 (2): 447–504. <https://doi.org/10.1086/432782>
- Vandenberghé, Frédéric. 2013. *What's Critical About Critical Realism? Essays in Reconstructive Social Theory*. Routledge. <https://doi.org/10.4324/9780203798508>
- Vareschi, Mark, and Mattie Burkert. 2017. "Archives, Numbers, Meaning: The

- Eighteenth-Century Playbill at Scale." *Theatre Journal* 68 (4): 597–613. <https://doi.org/10.1353/tj.2016.0108>
- Verhoeven, Deb, and Colin Arrowsmith. 2013. "Mapping the Ill-Disciplined? Spatial Analyses and Historical Change in the Postwar Film Industry." In *Locating the Moving Image: New Approaches to Film and Place*, by Julia Hallam and Les Roberts, 106–29. Bloomington: Indiana University Press. <http://muse.jhu.edu/book/27089>
- Wang, Tricia. 2016. "Why Big Data Needs Thick Data." *Ethnography Matters* (blog). January 20, 2016. <https://medium.com/ethnography-matters/why-big-data-needs-thick-data-b4b3e75e3d7>
- Ward, Michael Don, and Kristian Skrede Gleditsch. 2019. *Spatial Regression Models*. Vol. 155. Quantitative Applications in the Social Sciences. Thousand Oaks: SAGE.
- Warren, Christopher N., Daniel Shore, Jessica Otis, Lawrence Wang, Mike Finegold, and Cosma Shalizi. 2016. "Six Degrees of Francis Bacon: A Statistical Method for Reconstructing Large Historical Social Networks." *Digital Humanities Quarterly* 10 (3).
- Thow, Xin Wei. 2018. "All the Good Musicians Are Dead': A Sense of Decline in Javanese Gamelan." MA thesis, National University of Singapore.
- Wiesner, Susan. 2012. "ARTeFACT Project Summary." Charlottesville: University of Virginia. <http://avillage.web.virginia.edu/ARTeFACT/default.asp>
- Wiesner, Susan, Bennett Bradford, and Rommie Stalnaker. 2012. "The ARTeFACT Movement Thesaurus: Toward an Open-Source Tool to Mine Movement-Derived Data." In *Digital Humanities 2012*. Hamburg, 413–14. See <https://dh-abstracts.library.cmu.edu/works/1503>
- Wiesner, Susan, Shannon Cuykendall, Ethan Soutar-Rau, Rommie L. Stalnaker, Thecla Schiphorst, and Karen Bradley. 2016. "Schriftanz: Written Dance/Movement Poems." In *Digital Humanities 2016: Conference Abstracts*, 402–4. Kraków: Jagiellonian University and Pedagogical University, Kraków.
- Wiesner, Susan, and Rommie L. Stalnaker. 2015. "Representing Conflict through Dance: Using Quantitative Methods to Study Choreographic Time, Stage Space, and the Body in Motion." In *With(out) Trace: Inter-Disciplinary Investigations into Time, Space and the Body*, by S. Dwyer, R. Franks, and R. Green. Oxford: Inter-Disciplinary Press.
- Wilkinson, Leland. 1999. *The Grammar of Graphics*. Berlin: Springer.
- Wilkinson, Mark D., Michel Dumontier, IJsbrand Jan Aalbersberg, Gabrielle Appleton, Myles Axton, Arie Baak, Niklas Blomberg, et al. 2016. "The FAIR Guiding Principles for Scientific Data Management and Stewardship." *Scientific Data* 3 (March): 160018. <https://doi.org/10.1038/sdata.2016.18>
- Wong, Audrey. 1997. "Collaboration in the Work of The Necessary Stage." In *9 Lives, 10 Year of Singapore Theatre 1987–1997*, edited by Sanjay Krishan, 266–90. Singapore: First Printers.
- Xanthos, Aris, Isaac Pante, Yannick Rochat, and Martin Grandjean. 2016a. "Visualising the Dynamics of Character Networks." In *Digital Humanities 2016: Conference Abstracts*, 417–19.
- Xanthos, Aris, Isaac Pante, Martin Grandjean, and Yannick Rochat. 2016b. "About IntNetViz." 2016. <https://maladesimaginaires.github.io/intnetviz/about.html>
- Yong, Li Lan, Eng Hui Alvin Lim, Ken Takiguchi, Chee Keng Lee, Hyon-u Lee, Ha-young Hwang, Michiko Suematsu, and Kaori Kobayashi. 2015. "Asian Shake-

- speare Intercultural Archive (A|S|I|A). National University of Singapore." English, Chinese, Japanese. Accessed 1. <http://a-s-i-a-web.org>
- York, Christopher. 2017. "Exploratory Data Analysis for the Digital Humanities: The Comédie-Française Registers Project Analytics Tool." *English Studies* 98 (5): 459–82. <https://doi.org/10.1080/0013838X.2017.1332024>
- Yuri Tsivian. 2005. "Cinemetrics—About." <http://www.cinemetrics.lv/index.php>
- Zahavi, Dan. 2013. *The Oxford Handbook of Contemporary Phenomenology*. Oxford: Oxford University Press.
- Zillner, Sonja, Margrit Gelautz, and Markus Kallinger. 2002. "'The Right Move'—A Concept for a Video-Based Choreography Tool." In. Graz, Austria. <http://eprints.cs.univie.ac.at/1148/>

Index

- accuracy, 5–7, 32, 36, 47–48, 53–54, 60, 99, 117, 121–23, 125, 151, 155–56
- actor, 100, 103–5, 116, 134. *See also* performer
- algorithmic criticism, 82. *see also* Ramsay, Stephen
- ambiguity, 9, 11–13, 27–28, 55, 62, 90, 92, 146, 156, 192
- annotation, 65, 77–78, 90, 125–26, 163–64
- artistic collaborations: as object of study, 17, 19, 81, 84, 94, 96–97, 100–114
- AusStage, 100, 145, 163, 168–69
- authorship attribution, 75–76, 81, 84–85, 146
- Bardiot, Clarisse, 14, 101
- Bay-Cheng, Sarah, 14, 66, 70
- Bayesian: statistics, 52
- Bench, Harmony, and Kate Elswitt, 148–49, 158, 188
- bias, 10, 19, 31, 42–48
institutional bias, 54–56
proxy bias, 44, 49
- biomechanics, 123–125, 135–40
- bloom-and-fade publishing, 173–77
- Borges, Jorge Luis, 167, 173–74
- Borgman, Christine, 4–5, 171
- boxplot, 43, 58
- calibration (of methods), 47–48, 53, 55, 76, 125, 191
- Caplan, Debra, 2, 94, 100–102, 146
- CATMA (computer-assisted text markup and analysis), 83–84
- causality / causal relationships, 37–39, 42, 48–54
- chance: probability and statistics, 47–50, 153
- character types: in Javanese dance, 135–38
- characters: dramatic, 77, 94–99, 108–14
- chartjunk, 59
- choropleth, 64, 142–43, 145–46, 149, 152, 153, 156–57
- CIDOC-CRM, 167. *See also* data model
- classification, 41, 92, 118, 125, 148
- clusters: statistics, 42, 44–48, 55, 60, 76, 79, 130, 143, 152
- Comédie Française Registers Project, 80, 163
- complete spatial randomness (CSR), 143
- consensus, 7, 11–12, 47, 101, 125, 127
- constructivism / constructivist, 8–10, 23, 35–39, 63–64
- copyright, 20, 85, 134, 164, 171–72
- corpus: linguistics, 48, 75–77, 83–88, 93, 97–99, 109, 125
- critical realism, 10–11, 18, 24, 35–40, 43–44, 56, 63, 139
- culturally specific / cultural specificity, 18, 62–63, 116, 123, 132, 139, 147, 181, 188, 208

- dance, 87, 108, 116–28, 135–140, 148–49, 163
- dance notation systems, 119–20. *See also* Labanotation
- data
- data biography; biography of a dataset, 13, 169
 - data collection, 25, 43, 47, 56, 65, 116, 126, 148
 - data management plan (DMP), 171–72
 - data model, 78–79, 100, 126, 163–70, 178 (*see also* CIDOC-CRM; Dublin Core)
 - data repository, 80, 170, 172 (*see also* digital archives)
- deep mapping, 146
- defamiliarization, 8–11, 48, 56, 94, 102, 138
- difference image, 128–29, 132, 134
- digital archives, 2, 84, 87, 93, 128, 163, 165–67. *See also* data repository
- digital curation, 65, 168–69, 171, 173, 178
- digital humanities (DH), 2–3, 8, 11, 13–14, 18–19, 34, 38, 43–46, 55, 58, 65–66, 75–79, 81, 83–84, 86–87, 96, 111, 144, 146, 158, 168–71, 174–75, 177, 181–87
- Digital Literary Network Analysis (DLINA), 98–99
- dimensionality reduction, 45, 75–76, 79. *See also* principal component analysis (PCA)
- director: theater, 12, 86, 103, 105, 164–66
- dispersion: statistics, 42, 58–59
- distant reading, 81–82
- A Doll's House*, 2, 100–102, 144–45. *See also* IbsenStage
- Drucker, Johanna, 61–63, 67, 71, 132, 150
- Dublin Core, 165–68, 168. *See also* data model
- effect size, 42, 50, 54, 74, 111–14
- Elswitt, Kate. *See* Bench, Harmony, and Kate Elswitt
- ephemerality, 173, 177, 179
- ergodic: platforms, 68, 138, 149
- ethics / ethical implications, 64–67
- ethnography, 12, 24, 38, 54
- exploratory data analysis (EDA), 43–45, 58
- feminist data visualization, 63, 65
- Feynman, Richard, 30
- Fiadeiro, João, 127
- film studies / cinema, 67, 134, 141–49
- Fischer, Frank. *See* Digital Literary Network Analysis (DLINA)
- The Flying Inkpot, 87–93
- funding (financial resources; money), 15–16, 20, 55, 114, 122, 145, 172–75, 178, 185–87
- genre: written drama and literature, 76–77, 82–86, 98
- geographical information systems (GIS), 141–48, 158
- Gottschall, Jonathan, 31–35
- graphical conventions, 10, 60, 65
- Hachimura, Kozaburo, 119–20, 122, 124
- hacking / hackers, 184–85
- Hamlet, 77–78, 97. *See also* Shakespeare
- Hay, Deborah, 97
- hermeneutics, 15, 29, 39, 67, 70, 127, 158
- IbsenStage, 100, 144–45, 163. *See also* *A Doll's House*
- Indonesia. *See* wayang, Sendratari
- interactivity, 10, 60, 65, 67–72, 149
- intermedial essays, 70–72, 83, 149, 172, 178, 182
- interoperability, 165, 171–72
- interpretive approach, 2, 5, 8, 12, 24, 26, 29, 34–36, 41, 62–63, 66–71, 77, 80, 82–83, 85, 90, 92, 125, 127–29, 136, 138, 145–47, 158
- iterative research, 14, 43, 61, 84, 147, 176

- Jannidis, Fotis, 35, 168–69
 Jockers, Matthew, 77, 82, 181, 206
- kernel density estimation (KDE), 112, 132, 154, 194
- Labanotation, 117–20, 122. *See also* dance
 literary scholarship / literary analysis, 30–34, 43, 75–77, 80–85, 95
- machine learning (ML), 45, 47, 49, 52, 79, 81, 125
 Mann-Kendall, 88, 101. *See also* trends
 Manovich, Lev, 3, 63–64, 67, 134, 174–75
 markers: motion capture, 121–22
 metadata, 164–71. *See also* data
 methods and methodology, 9–10
 minimal computing, 177, 185–87
 minimum convex polyhedron; minimum convex polygon, 124, 126
 mocap (motion capture), 17, 117, 120–22, 125, 136, 140
 models / modelling, 6, 14, 41, 43–48, 52, 56, 75, 79, 94, 96, 100–101, 104, 114, 116, 126, 164–70, 184
 Moran's I, 143–44, 152–55, 157
 Motion Bank, 126
 Mrázek, Jan, 139, 151
- natural experiments, 126
 negative findings, 48, 56, 137
 network theoretical measurements, 95, 98–99, 110, 112–13, 193
- open access, 171–72
 Open Science Framework (OSF), 172
 open source, 181–82
- p-value, 47, 50–52, 88, 153–54
 pairwise plot, 112–13
 parameters / parametrization, 8, 34, 45, 48, 55
- Parikesit, Gea Oswah Fatah, 122, 128–30, 187
 PCA. *See* principal component analysis (PCA)
 performance studies, 5, 13, 26, 28, 66, 173, 186. *See also* Schechner, Richard
 performative (data) visualization, 39, 54, 65–67, 126–28, 134, 150
 performer, 2, 69, 94, 100, 102, 104–5, 108, 121, 125–27, 140, 144, 148, 155–56, 164–65. *See also* actor
 permutations (bootstrapping method), 153–55
 Piper, Andrew, 41, 77, 79, 169
 playbills, 75, 80
 political aspects
 of dance, 138–39, 148
 of research, 1, 16, 27, 36, 57, 60, 66
 of theater, 102–3, 145
 positivism, 27, 36–38
 poststructuralist; poststructuralism, 27
 principal component analysis (PCA), 75–76, 79, 84. *See also* dimensional-ity reduction
 program booklets, 3, 6, 19, 25, 86
- Ramsay, Stephen, 9, 34, 41, 46, 66, 82, 85, 183
 Ramsey theory: mathematics, 46, 52
 randomized controlled trials; controlled trials, A|B testing, 48, 49, 54
 randomness; random conditions, 47–49, 54, 113, 143, 153–54
 replicability / reproducibility, 7–9, 15, 23–29, 28, 30, 34–35, 40, 48, 53, 55–56, 65, 82, 84, 102, 126–28, 137, 153, 156
 rich-prospect, 68
 Rockwell, Geoffrey, and Stéfan Sinclair, 8, 34, 70, 83–84, 93
- Schauf, Andrew, 19, 109, 113, 187
 Schechner, Richard, 5

- Schöch, Christof, 19, 76–77
- scholarship: theatre, 16, 25, 29, 30–32, 70, 72, 80, 87, 97, 114, 138, 140, 146, 149, 155, 171, 174, 185–88
- scientific: approaches to research, 10, 15, 18–19, 23–24, 29, 30–35, 46, 49–55, 58–62, 81, 84, 101, 118, 123, 125, 127–28, 144, 187
- semiotics, semiotic, 27, 57, 71
- Sendratari: Javanese dance, 135–38
- sensors: motion capture, 5, 118, 121, 128, 138
- Shakespeare, 12–13, 81, 86, 102, 163. *See also* Hamlet
- Sinclair, Stéfan, 68. *See also* Rockwell, Geoffrey, and Stéfan Sinclair
- Singapore, 16–19, 23, 33, 87–93, 102–8. *See also* The Flying Inkpot; The Necessary Stage (TNS)
- social network analysis (SNA), 95–96. *See also* network theoretical measurements
- spectator, 101, 134, 139, 150
- stylometry, 75, 136, 138–40, 181
- TEI, 77–79, 99, 120, 166, 170
- text reuse, 79–80, 86
- textual differences: measures of textual difference, 76
- The Necessary Stage (TNS), 94–95, 102–8
- thickness
thick context, 67–71, 83, 90, 102, 110
thick data, 13
thick description, 13, 16, 24, 147
- time series analysis, 75, 79, 83–84, 107–8, 156
- topic modelling, 34, 75, 77, 79, 83–84, 86
- training; education, 29, 52, 170, 183
- trend detection, 42, 44–46, 48, 60, 76, 84, 88–92, 155. *See also* Mann-Kendall
- Trilke, Peer. *See* Digital Literary Network Analysis (DLINA)
- uncertainty / doubt, 7, 9, 30, 47, 168
- Underwood, Ted, 30, 41, 81–82
- Verhoeven, Deb, 142, 145–46
- video, 27, 70–71, 109, 118, 121–23, 126, 128–29, 134, 137–38, 163–64, 166–68, 175, 178
- violinplot, 43, 58, 112, 130–32
- Voyant Tools, 83–84, 88, 170, 181. *See also* Rockwell, Geoffrey, and Stéfan Sinclair
- wayang, 33, 71, 94–95, 108–14, 122, 128–34, 150–57, 165–67
- Wiesner, Susan, 2, 19, 125–26
- word frequencies, 9, 76, 88–93
- XML, 77–78, 119–20, 122, 125, 166