

Artificial Intelligence and the Humanities: Toward a New Interdisciplinary Paradigm

Chen Jie

Zhejiang University, Hangzhou, Zhejiang, China.

ABSTRACT

The relationship between artificial intelligence (AI) and the humanities has grown more complex and generative over the past decade. While computing entered the humanities through relatively narrow tasks like textual indexing and corpus analysis, contemporary AI, particularly large language models and machine learning systems, now touches nearly every corner of humanistic inquiry, from literary analysis and historical research to philosophical ethics and cultural criticism. This paper examines how AI is reshaping the conditions of humanistic knowledge production and argues that the moment calls for a genuinely new interdisciplinary paradigm rather than the incorporation of AI as just another research tool. Drawing on established scholarship in digital humanities, philosophy of technology, and AI ethics, the paper traces the historical arc from early humanities computing to present-day applications, identifies recurring theoretical tensions, and proposes a framework for AI-humanities collaboration built around three principles: interpretive accountability, epistemic plurality, and critical reciprocity. The paper also addresses the ethical challenges that arise when computational systems interact with culturally sensitive materials and argues that humanists are well-positioned to interrogate the assumptions embedded in AI systems. Rather than treating AI as a threat to humanistic methods or as a neutral instrument, this paper advocates for an engaged, critically informed relationship between the two.

Keywords: Artificial intelligence, digital humanities, interdisciplinary research, humanities computing, AI ethics, large language models

INTRODUCTION

Something unusual is happening at the intersection of artificial intelligence and the humanities. It is not simply that humanists have adopted new tools (and that has been happening since at least the 1940s), when the Italian Jesuit priest Roberto Busa began working with IBM to create a machine-assisted concordance of the works of Thomas Aquinas (Busa, 1980). What is different now is the character of the AI systems involved and the scale at which they operate. Large language models can generate coherent scholarly prose, process vast archives of text, translate between languages with reasonable accuracy, and produce plausible interpretations of literary and historical materials. These capacities raise a different kind of question from the ones that earlier humanities computing had to face: not just whether machines can help us process texts faster, but what happens to interpretation when machines can interpret.

The humanities, understood broadly to include literary studies, history, philosophy, linguistics, art history, and cultural criticism, have always been organized around interpretation. The guiding assumption has been that understanding human cultural production requires human readers who bring historical consciousness, ethical sensitivity, and social situatedness to the act of reading. That assumption is now under pressure from a direction quite different from the earlier quantitative turn in the humanities. When Franco Moretti (2013) proposed distant reading as a method, he was still describing a process in which human scholars designed the research questions and interpreted the results that computational tools

returned. AI systems capable of generating interpretations on their own complicate that division of labor in ways that remain poorly understood.

This paper does not argue that AI will replace humanistic inquiry, nor that humanists should resist or ignore AI. Both of those responses would be intellectually premature and practically unhelpful. Instead, the paper argues that the current situation demands a new kind of interdisciplinary engagement, one in which humanists bring their distinctive skills in interpretation, critique, and ethical reasoning to bear on AI systems themselves, while also drawing on AI capabilities in ways that are theoretically informed and methodologically transparent.

The paper proceeds in several stages. First, it offers a compressed historical account of how computing entered the humanities and how the digital humanities emerged as a recognizable field. Second, it surveys the range of ways in which contemporary AI is being applied across humanities disciplines, with particular attention to large language models and machine learning techniques for text and image analysis. Third, it develops a theoretical framework for thinking about what genuine AI-humanities interdisciplinarity might look like, organized around three principles: interpretive accountability, epistemic plurality, and critical reciprocity. Fourth, it addresses the ethical challenges that arise in this domain, including questions of bias, cultural representation, and the politics of data. Finally, it sketches several pathways toward institutionalizing a new paradigm.

The argument draws on established scholarship in digital humanities, philosophy of technology, and AI ethics. It also draws on the recognition, noted by N. Katherine Hayles (1999) in a different context, that the boundary between human cognition and computational process is not fixed but produced, and that humanists are in a particularly good position to examine how that boundary is being drawn and with what consequences.

FROM HUMANITIES COMPUTING TO DIGITAL HUMANITIES

The history of computing in the humanities is longer than most people outside the field tend to assume. Busa's Index Thomisticus project, which began in 1949 and was completed only in 1980, represents the inaugural moment of what would later be called humanities computing (Busa, 1980; McCarty, 2005). Busa's collaboration with IBM established several patterns that would persist for decades: computing was understood as a tool for processing textual data at scales that exceeded human capacity, the humanist remained the interpreter of the results, and the collaboration required institutional support from outside the humanities.

Through the 1970s and 1980s, humanities computing developed through projects focused on corpus analysis, concordance creation, and textual editing. The Text Encoding Initiative, launched in 1987, formalized a set of protocols for encoding literary and historical texts in machine-readable form. The work was painstaking, technically demanding, and largely invisible to the broader humanities community. The assumption governing these projects was that the computer was a powerful but essentially passive instrument: it could store and sort and count, but the work of interpretation remained entirely human.

The emergence of the World Wide Web in the 1990s changed the landscape by making digital texts and tools widely accessible and by creating new possibilities for collaborative scholarly projects. Schreibman, Siemens, and Unsworth's (2004) *A Companion to Digital Humanities* represents something like the field's first major disciplinary statement. The volume mapped an emerging set of practices, tools, and communities, though it also reflected the relatively narrow scope of what the field was then attempting. The emphasis was still firmly on computing as a service to traditional humanistic aims rather than as a generator of new epistemological possibilities.

By the 2010s, digital humanities had become a recognizable, if contested, academic field. The Gold (2012) Debates in the Digital Humanities volume captures a moment of internal self-examination in which practitioners debated whether the field had a theoretical core, whether it was adequately addressing questions of race and power, and what its relationship to traditional humanities scholarship ought to be. Critics like Alan Liu (2012) argued that digital humanities had been conspicuously reluctant to engage with the cultural criticism that had animated literary and historical studies since the 1970s. That critique remains worth taking seriously.

These debates were productive, but the landscape has shifted again since then. The AI systems now available, built on deep learning, trained on billions of documents, and capable of generating sophisticated language outputs, are qualitatively different from the earlier tools that digital humanities relied on. They are not simply faster versions of concordance software. They embed interpretive tendencies, reproduce the biases of their training data, and produce outputs that can be difficult to interrogate or audit. Understanding how these systems work and what they do requires a kind of critical engagement that humanities scholars are well-positioned to offer, provided the institutional conditions exist to support it.

AI APPLICATIONS ACROSS HUMANITIES DISCIPLINES

The range of AI applications in the humanities has expanded rapidly over the past decade. The following discussion surveys the most significant developments organized by disciplinary area, and Table 1 provides a structured overview.

LITERARY STUDIES

In literary studies, the most influential AI application has been in the analysis of large text corpora using machine learning methods. Underwood (2019) demonstrated how computational models can identify patterns of literary change across centuries that would be impossible to detect through close reading of individual texts. His work on genre, prestige, and gender in literary history draws on a corpus of thousands of novels and produces findings that are counterintuitive and methodologically transparent. Jockers (2013) similarly used computational methods to map influence networks among nineteenth-century novels, tracing how thematic and stylistic features spread through literary culture.

More recently, large language models have been used in literary studies to assist with translation, to generate summaries of large corpora, and to examine stylistic features of texts. The interpretive status of these outputs remains an open question. When a language model generates a reading of a poem, that reading reflects statistical regularities in its training data rather than a situated act of reading by a historically located reader. The difference matters for how we understand what interpretation is and who or what can perform it, a question that is at once empirical and philosophical.

HISTORICAL RESEARCH

In historical research, AI applications include automated transcription of handwritten historical documents, classification of archival materials, and analysis of large-scale historical datasets. The Transkribus platform, developed through European research infrastructure projects, has made it possible to train automated transcription models on specific historical handwriting styles, dramatically reducing the labor required to make manuscript sources accessible to scholars. Historians have also used topic modeling and network analysis to identify thematic patterns and social connections in large documentary collections that exceed any individual's reading capacity.

Michel et al. (2011) demonstrated the potential of computational methods at a different scale, tracing patterns of word usage across millions of books spanning several centuries in a study published in *Science*. The authors described this approach as *culturomics*, and while the methodology attracted legitimate

criticism, particularly regarding the representativeness of the underlying corpus, and the study established that large-scale computational analysis of textual data could produce findings of genuine historical interest. The Google Books corpus on which it drew remains one of the largest collections of digitized text ever assembled.

PHILOSOPHY, LINGUISTICS, AND ART HISTORY

In philosophy and ethics, AI has become both a tool for research and a primary object of inquiry. The question Turing (1950) posed, namely whether a machine could exhibit behavior indistinguishable from intelligent human behavior, has taken on practical urgency as language models pass informal versions of his imitation game with some regularity. The development of AI systems capable of reasoning, translation, and creative production raises questions that belong squarely within the philosophical traditions of mind, language, and epistemology, and philosophers are beginning to engage with those questions in a sustained way.

In linguistics, AI tools are used for corpus analysis, automatic parsing, machine translation, and the computational modeling of language acquisition and change. These applications have accelerated research into phenomena that were previously difficult to study at scale, including long-term patterns of semantic shift and cross-linguistic variation. In art history and visual studies, AI image recognition tools are being used to examine large collections of images, identify stylistic patterns across periods and traditions, and assist in the attribution of works of uncertain provenance. Manovich (2013) outlined the theoretical implications of treating cultural objects as data, arguing that computational approaches to cultural analysis require their own conceptual frameworks rather than simply borrowing from existing humanistic methods, an argument that has become more rather than less relevant as the tools have grown more powerful.

Table 1. AI Applications Across Humanities Disciplines

Discipline	Primary AI Applications	Key Methods	Representative Works	Open Questions
Literary Studies	Corpus analysis, stylometry, LLM-assisted translation and interpretation	Topic modeling, neural LLMs, stylometry	Underwood (2019); Jockers (2013)	Interpretive status of machine-generated readings
Historical Research	Document transcription, archive classification, network analysis	OCR, ML classification, NLP	Michel et al. (2011); Schreibman et al. (2004)	Selection bias in digitized archives
Philosophy	Ethical analysis of AI systems, computational argument modeling	Conceptual analysis, argument mining	Turing (1950); Floridi et al. (2018)	Moral status of AI agents; limits of formal ethics
Linguistics	Corpus analysis, machine translation, language change modeling	NLP, statistical modeling, neural MT	Michel et al. (2011)	Training data representativeness across languages

Art History	Image classification, style analysis, provenance attribution	CNN, feature extraction, similarity analysis	Manovich (2013)	Correspondence of ML categories to art-historical concepts
-------------	--	--	-----------------	--

Note. LLM = large language model; ML = machine learning; NLP = natural language processing; OCR = optical character recognition; CNN = convolutional neural network; MT = machine translation.

TOWARD AN INTERDISCIPLINARY THEORETICAL FRAMEWORK

A genuinely interdisciplinary paradigm for AI and the humanities requires more than a list of tools and their applications. It requires a set of theoretical commitments that can guide practice, shape research questions, and provide criteria for evaluating outcomes. This paper proposes three such commitments: interpretive accountability, epistemic plurality, and critical reciprocity. These principles are presented schematically in Figure 1.

INTERPRETIVE ACCOUNTABILITY

Interpretive accountability refers to the requirement that researchers who use AI tools in humanistic inquiry be able to give an account of how those tools work, what assumptions are built into them, and how their outputs were produced. This is not a demand for technical expertise in machine learning across the board; the specialized knowledge required is substantial, and division of labor in collaborative research is inevitable. It is rather a demand that humanists not treat AI tools as black boxes whose outputs are simply accepted or rejected on the basis of whether they confirm prior intuitions.

Rockwell and Sinclair (2016) make a related point in their discussion of hermeneutical computing: the act of building or using computational tools is itself interpretive, and the decisions made in constructing those tools (what data to include, what categories to use, what algorithms to apply) reflect theoretical commitments that need to be made explicit. This principle applies with even greater force to large language models, which embed patterns absorbed from training data that is rarely transparent to the researchers using them. Interpretive accountability is not a burden imposed on humanists who use AI; it is a logical extension of the interpretive rigor that humanistic scholarship has always demanded.

EPISTEMIC PLURALITY

Epistemic plurality refers to the recognition that humanistic knowledge comes in forms that AI systems are not well-equipped to replicate or evaluate. Close reading, archival intuition, the interpretation of cultural contexts, the ability to recognize irony or register ambiguity, the capacity to understand why a particular historical moment matters: these are forms of knowledge that resist formalization. An interdisciplinary paradigm should not assume that computational analysis and humanistic interpretation are doing the same thing in different ways, or that one can substitute for the other.

Drucker's (2009) insistence that humanistic data is always *capta* (captured, constructed, interpreted) rather than simply given is a useful touchstone here. When literary texts are converted into datasets for machine learning, decisions have been made about how to tokenize them, what metadata to attach, which editions to use, and which texts to include at all. Those decisions are not neutral. They reflect existing hierarchies of scholarly attention and cultural value, and they shape the findings that computational analysis can produce. Acknowledging this is not an argument against computational methods; it is an argument for being clear-eyed about what they can and cannot tell us.

CRITICAL RECIPROCALITY

Critical reciprocity refers to the idea that the relationship between AI and the humanities should run in both directions. Humanists should not simply receive AI tools as instruments to be applied to their research questions. They should also bring humanistic methods of critique to bear on AI systems themselves, examining their assumptions, interrogating their training data, questioning their categories, and identifying the cultural values they embed. This is a form of scholarship that humanists are distinctively equipped to pursue, and it is scholarship with practical stakes.

Bender et al. (2021) examined the social risks of large language models trained on internet text, showing how such systems reproduce harmful stereotypes, generate convincing misinformation, and can be deployed in ways that concentrate communicative power in the hands of those who control the infrastructure. The humanistic tools developed for examining rhetoric, ideology, and cultural representation are directly applicable to this kind of critical examination. If the humanities do not rise to the challenge of producing rigorous AI criticism, that critical function will be left to disciplines with less developed frameworks for reading culture.

Figure 1. A Proposed Framework for AI-Humanities Interdisciplinary Research

Principle	Core Commitment	Practical Implications
Interpretive Accountability	Researchers must be able to explain the assumptions and mechanisms of any AI tool used in their work	Transparent methodological documentation; critical review of tool design choices; disclosure of training data limitations
Epistemic Plurality	Humanistic and computational knowledge are distinct in kind; neither substitutes for the other	Collaborative research designs that allocate interpretive authority appropriately; resistance to reductionism in either direction
Critical Reciprocity	The relationship between AI and the humanities runs both ways: humanists use AI tools and also scrutinize them as cultural artifacts	Development of AI criticism as a scholarly genre; humanist input into AI design; critical reading of AI outputs as part of standard research practice

Note. This framework proposes three principles that together constitute the basis for a new AI-humanities interdisciplinary paradigm. Each principle addresses a distinct dimension of the relationship between computational and humanistic inquiry and carries specific methodological and institutional implications.

ETHICAL CHALLENGES

The ethical challenges at the intersection of AI and the humanities are substantial and deserve sustained attention. Several stand out as particularly pressing for scholars working in this space.

BIAS AND REPRESENTATION IN TRAINING DATA

Large language models and machine learning systems are trained on data that reflects the historical patterns of cultural production and preservation. That means they tend to perform better on texts and images from dominant linguistic traditions, Western cultural contexts, and well-documented historical periods. They are often poorly calibrated for non-Western materials, minority language traditions, and the cultural production of historically marginalized communities. This is not an incidental flaw that will be

corrected over time as more data is added; it reflects structural inequalities in which texts have been written, preserved, digitized, and made accessible.

Crawford (2021) has documented the material and political dimensions of AI systems, showing how the infrastructures of data collection, annotation, and model training reproduce existing hierarchies of power in ways that are rarely visible to end users. For the humanities, which have spent decades developing critical frameworks for understanding how cultural materials have been selected, preserved, and interpreted in ways that serve the interests of dominant groups, these patterns are deeply familiar. The tools of cultural criticism are directly applicable to the analysis of AI systems, and humanists who develop this application are making a contribution that extends beyond their own discipline.

AUTHORSHIP, CREDIT, AND ATTRIBUTION

When AI systems are used to generate text, translate documents, or identify interpretive patterns in large corpora, the question of who should be credited for the resulting scholarship becomes complicated. Existing norms for scholarly attribution were developed in contexts where the labor of research and writing was performed by human beings who could be identified and held accountable for their claims. Those norms need to be rethought for contexts in which AI systems contribute substantially to research outputs, not just as tools that researchers operate, but as generators of content that may appear in published work.

This is not merely a technical question about citation practices. It connects to deeper questions about the nature of intellectual contribution, the social organization of credit in academic communities, and the legal frameworks governing copyright and intellectual property. Humanists who have thought carefully about the history of authorship, a concept that is historically contingent rather than naturally given, are well-positioned to contribute to these debates in ways that go beyond merely policing norms.

CULTURALLY SENSITIVE MATERIALS AND EPISTEMOLOGICAL RISKS

AI systems trained on large corpora of cultural materials may reproduce those materials in ways that violate the cultural norms of the communities from which they derive. Sacred texts, oral traditions preserved in archival recordings, historical records of trauma and violence, and materials created by communities with specific protocols governing their use can all appear in training datasets without the knowledge or consent of the relevant communities. Humanists who work with such materials have developed ethical protocols through decades of engagement with communities and archives. Those protocols need to inform the use of AI in humanities research.

There is also a subtler risk worth naming: AI systems produce outputs that are fluent and confident but that may lack the interpretive grounding that makes humanistic scholarship trustworthy. When a language model generates a reading of a historical document that sounds authoritative, it can be difficult for non-specialist readers to recognize that the reading may be statistically derived rather than interpretively grounded in historical knowledge and contextual understanding. The humanities have long been concerned with questions of interpretive authority and credibility; those concerns are directly relevant to evaluating AI-generated outputs in scholarly contexts.

Floridi et al. (2018) developed an ethical framework for AI that includes principles of beneficence, non-maleficence, autonomy, justice, and explicability. These principles are useful as general orientations, but they need to be supplemented by the more specific ethical frameworks that humanists have developed for working with cultural materials, marginalized communities, and contested historical records. The general principles of AI ethics and the specific ethical protocols of humanistic research are not alternatives; they need to be integrated.

PROPOSALS FOR A NEW INTERDISCIPLINARY PARADIGM

If the argument made in the preceding sections is correct, what follows from it practically? Several pathways toward institutionalizing a new paradigm present themselves, and they are complementary rather than sequential.

CURRICULAR CHANGE

Graduate programs in the humanities should develop training in AI literacy, meaning not programming per se, but a working understanding of how machine learning systems function, what kinds of data they require, how their outputs are generated, and what kinds of errors they characteristically produce. This literacy would allow humanists to engage critically with AI tools used in their research, to collaborate more effectively with computer scientists and data scientists, and to contribute to the growing field of AI criticism.

Similarly, computer science and data science programs that work on natural language processing and related AI applications should include training in humanistic methods of interpretation and critique. Hayles (1999) argued for a posthumanities that takes seriously the entanglement of human cognition and technological systems; that entanglement is now sufficiently specific and consequential to warrant being built into graduate training. This does not mean that every humanist needs to become a programmer or that every AI researcher needs to become a literary critic. It means that the institutions training the next generation of scholars in both domains need to prepare them for genuinely collaborative work across disciplinary lines.

INSTITUTIONAL AND FUNDING STRUCTURES

Research funding bodies should create explicit support for interdisciplinary projects that bring together AI researchers and humanists in genuinely bilateral collaboration, not projects in which one group provides tools and the other provides data, but projects in which the theoretical commitments of both inform the design of the research from the beginning. This kind of collaboration is more difficult to organize and evaluate than traditional single-discipline research, and funding structures need to reflect that difficulty rather than forcing interdisciplinary work into existing categories designed for disciplinary research.

METHODOLOGICAL NORMS AND TRANSPARENT REPORTING

The humanities need to develop norms for the transparent reporting of AI-assisted research. When computational methods are used in literary, historical, or philosophical research, the documentation of those methods should be sufficient for other researchers to evaluate the choices made and, where appropriate, to replicate the analysis. Underwood (2019) and Jockers (2013) both model this kind of methodological transparency, and their practice deserves to become standard across the field rather than an exemplary exception. Journals that publish digital humanities and AI-humanities work should develop reporting standards commensurate with those that exist in the quantitative social sciences.

Ramsay (2011) suggested that algorithmic criticism might be a genuinely new form of literary inquiry. Taking that suggestion seriously means developing not just new tools but new genres of scholarly writing in which the description of computational methods and the interpretation of results are integrated rather than separated. Such genres are still being developed, and their development is itself a scholarly project that deserves institutional recognition and support.

AI CRITICISM AS A HUMANISTIC SCHOLARLY GENRE

Humanists should develop AI criticism as a recognized genre of humanistic scholarship: systematic, theoretically grounded analysis of AI systems and their cultural implications. This genre builds on existing

traditions of technology criticism and cultural criticism while developing the specific expertise needed to examine how AI systems work and what they do. It is scholarship that the humanities are well-equipped to produce and that they have a particular responsibility to pursue, given the scale at which AI systems now interact with cultural materials and the speed at which those interactions are reshaping the conditions of humanistic knowledge production.

Such criticism would examine not just what AI systems produce but how they were built, whose labor made them possible, what data they were trained on and whose cultural materials that data represents, and what interests are served by their deployment. It would bring the analytic tools of rhetoric, ideology critique, discourse analysis, and cultural history to bear on systems that are shaping what can be said, what can be remembered, and what can be known. This is not a peripheral or supplementary function. It is one of the most consequential things humanists could be doing right now.

CONCLUSION

The relationship between artificial intelligence and the humanities is not new, but it has entered a qualitatively different phase. The AI systems now available are more powerful, more flexible, and more opaque than anything that earlier humanities computing had to reckon with. They raise questions that cannot be answered within any single discipline, and they create possibilities that cannot be realized without sustained collaboration across disciplinary lines.

This paper has argued for a new interdisciplinary paradigm built on three principles (interpretive accountability, epistemic plurality, and critical reciprocity) and has identified curricular, institutional, methodological, and critical pathways toward realizing that paradigm. The argument is not that AI will transform the humanities in ways that are uniformly beneficial, nor that it represents a threat to humanistic inquiry that needs to be resisted. It is that the current moment calls for an engaged, critical, theoretically grounded response of the kind that humanistic scholarship is distinctively capable of offering.

The humanities have something to gain from AI: access to new forms of analysis, new scales of inquiry, and new possibilities for making meaning from the vast accumulated record of human cultural production. AI systems have something to gain from humanistic engagement: the critical scrutiny, the interpretive frameworks, and the ethical sensibilities that humanists bring to their work. Realizing those mutual gains requires building the institutional conditions for genuine collaboration and developing the theoretical frameworks that can make collaboration productive rather than merely additive.

Turing (1950) imagined a future in which the question of whether machines can think would become as empirically tractable as any other question about behavior. Whether or not that future arrives, the question of what AI and the humanities can do together, and on what terms, is both tractable and urgent. Getting serious about answering it is the work of the paradigm this paper has tried to sketch.

REFERENCES

- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency* (pp. 610–623). ACM.
- Busa, R. (1980). The annals of humanities computing: The index Thomisticus. *Computers and the Humanities*, 14(2), 83–90.
- Crawford, K. (2021). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.
- Drucker, J. (2009). *SpecLab: Digital Aesthetics and Projects in Speculative Computing*. University of Chicago Press.

- Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). An ethical framework for a good AI society: Opportunities, risks, principles, and recommendations. *Minds and Machines*, 28(4), 689–707.
- Gold, M. K. (Ed.). (2012). *Debates in the digital humanities*. University of Minnesota Press.
- Hayles, N. K. (1999). *How we became posthuman: Virtual bodies in cybernetics, literature, and informatics*. University of Chicago Press.
- Jockers, M. L. (2013). *Macroanalysis: Digital methods and literary history*. University of Illinois Press.
- Liu, A. (2012). Where is cultural criticism in the digital humanities? In M. K. Gold (Ed.), *Debates in the digital humanities* (pp. 490–509). University of Minnesota Press.
- Manovich, L. (2013). *Software takes command*. Bloomsbury Academic.
- McCarty, W. (2005). *Humanities computing*. Palgrave Macmillan.
- Michel, J.-B., Shen, Y. K., Aiden, A. P., Veres, A., Gray, M. K., Pickett, J. P., Hoiberg, D., Clancy, D., Norvig, P., Orwant, J., Pinker, S., Nowak, M. A., & Aiden, E. L. (2011). Quantitative analysis of culture using millions of digitized books. *Science*, 331(6014), 176–182.
- Moretti, F. (2013). *Distant reading*. Verso.
- Ramsay, S. (2011). *Reading machines: Toward an algorithmic criticism*. University of Illinois Press.
- Rockwell, G., & Sinclair, S. (2016). *Hermeneutica: Computer-assisted interpretation in the humanities*. MIT Press.
- Schreibman, S., Siemens, R., & Unsworth, J. (Eds.). (2004). *A companion to digital humanities*. Blackwell.
- Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59(236), 433–460.
- Underwood, T. (2019). *Distant horizons: Digital humanities and the study of literature*. University of Chicago Press.