PROJECT MUSE®

Computation as Context: New Approaches to the Close/Distant Reading Debate

Kirilloff Gabi

College Literature, Volume 49, Number 1, Winter 2022, pp. 1-25 (Article)

Published by Johns Hopkins University Press

DOI: <https://doi.org/10.1353/lit.2022.0000>

➔ *For additional information about this article*

<https://muse.jhu.edu/article/844432>

ESSAYS

COMPUTATION AS CONTEXT:
NEW APPROACHES TO THE
CLOSE/DISTANT READING DEBATE

GABI KIRILLOFF

I at once gave up my former occupations, set down natural history and all its progeny as a deformed and abortive creation, and entertained the greatest disdain for a would-be science which could never even step within the threshold of real knowledge.

—Mary Shelley, *Frankenstein*

Over the last four years, criticism of the digital humanities (DH) has escalated. In a piece for the *Chronicle of Higher Education*, Timothy Brennan links the rise of DH with the downfall of the humanities, noting that DH has “turned many humanists into establishment curators and made critical thought a form of planned obsolescence” (2017). Big-data-driven approaches to literature have, for Brennan and others, not lived up to the hype. In her controversial piece for *Cultural Inquiry*, Nan Z. Da takes this criticism further, suggesting that computational text analysis is full of “technical problems, logical fallacies, and conceptual flaws” (2019, 601). Underlying these

critiques is the notion that computation cannot productively capture the nuance of literary texts. The data, either through inherent inadequacy or human error, is wrong.

DH scholars have addressed these claims by defending their methods. Brennan reduces computational research to the meaningless act of counting the occurrences of “whale” in *Moby Dick*. Sarah E. Bond, Hoyt Long, and Ted Underwood respond by pointing out the ways in which computational tools (modeling in particular) have developed beyond frequency counts.¹ Da claims that the statistical principles underlying several prominent DH publications are unsound. Andrew Piper, Ben Schmidt, and others counter that Da has misunderstood and misused their methods.² This antagonism over methodological validity is reflected in the recent, productive push (from inside and outside of DH) toward transparency. Thanks largely to scholars like Katherine Bode, open-source data, code, and corpora are becoming the new norm.

However, this recent preoccupation with method has primarily functioned as a red herring. Debates about methodological validity have obscured the fact that many scholars express a disinterest in the results of computational work. DH “tells us what we already know” (Kirsch 2014). While other aspects of the DH critique have shifted, this criticism has proven resilient. Results are pegged as unoriginal and valueless. This complaint is trickier to parry than claims that “the math is wrong” or “the corpus is biased.” The disjuncture between excited DH scholars and underwhelmed colleagues speaks to a fundamental breakdown in the expectations for and communication of computational work. What is missing from the current DH debate is a frank discussion of the relationship between method and interpretation: What makes a computationally derived textual interpretation valuable and when should it occur? Whose job is it to analyze and interpret the data produced from computational work?

I argue that the information produced from computational methods should be understood as a type of context, rather than as a textual “reading” or interpretation. While this may seem like a subtle shift, viewing computational output as context implies that impartial and flawed methods can still supply valuable information. Incorrect data and flawed methods can lead to relevant and novel textual interpretations, much as valid methods can produce uninspired interpretations. This is a radical departure from the field’s current preoccupation with methodological validity.

The field of DH stands at a critical juncture in determining who is responsible for performing data-driven interpretation and how this interpretation should be communicated. Rather than focusing on method as a type of product (the “deliverable”), critics should understand digital work as an ongoing process in which interpretation occurs at multiple stages. I argue that these distinct stages do not need to be performed by the same scholar(s). The context produced from computation must be taken up by other members of the scholarly community. Currently, there is undue pressure on digital scholars to produce rigorous, inventive computational methods and valuable pieces of textual interpretation. The field would benefit from separating these acts. Archival studies offer a productive model for the relationship that computational scholars could build with the rest of the scholarly community. Archival research produces biographical and historical context that is used by other scholars in acts of textual interpretation. Such pieces of context are often riddled with gaps and are rarely used in isolation. The creation of a data set, like the discovery of a cache of letters or an unpublished manuscript draft, is a useful act, regardless of whether it immediately produces a novel textual interpretation. To explore the relationship between computational method and interpretation, I analyze two data-driven projects. Both projects demonstrate that flawed methods and incomplete data can still facilitate useful interpretations. The first, Google’s *Perspective*, is an Application programming interface (API) that uses machine learning models to flag online content as “toxic.” Because of biased training data, *Perspective* labels statements about disenfranchised groups, including women and people of color, as toxic. *Perspective* is not the only online tool that furthers racist and sexist sentiments. As Safiya Noble, Kate Crawford, Helen Nissenbaum, and others have argued, many of the algorithms that underlie our everyday existence are deeply biased. In analyzing *Perspective*, I discuss the ways in which cultural attitudes about algorithmic bias are beginning to shift the DH debate. Rather than discount *Perspective* as a broken tool built on “bad data,” I argue that it offers scholars an opportunity to critically examine contemporary racist and sexist attitudes.

I also reflect on my current research, which uses natural language processing to automate the extraction of direct address from a corpus of over 2,000 novels. Originally, the project was oriented toward literary history: how do authors address their readers and

how does this change over time? (e.g., How popular was the Victorian “dear reader” convention, and when exactly did it disappear?) As the project progressed, my focus shifted. The data raised more questions than it answered, and I was unable to determine an exact accuracy rate for my method. The data, though flawed, aided in recovery work and facilitated close reading. Instead of providing a clear picture of a historical phenomenon, the data guided me toward individual narratives that employed address in unconventional ways. Through discussing my work, I attempt to model a workflow that utilizes data as context.

These projects are radically different. One is commercial and public, the other academic and private. However, both employ common DH methodologies (e.g., machine learning and natural language processing) to examine texts. In both cases the data are imperfect and partial, but this is inescapable. All data are imperfect and partial because they are, as Crawford notes, always derived “from a human history” (2017). While these imperfections can have negative effects, they should not be viewed as a deterrent to computational research.

“BAD” DATA AND FLAWED METHODS

There are many reasons to be skeptical of data. In 1998, a study was published linking the MMR vaccine with an increased risk of autism.³ While it soon became clear that the initial findings were fraudulent, the study fueled decades of fear and contributed to tens of thousands of preventable cases of measles and mumps. Anti-vaccination research offers a prime, and oft-cited, example of the consequences of “bad” data. The misuse of inaccurate data is becoming increasingly visible as governments and corporations rely on algorithms to automate services. Data does not have to be fraudulent to cause damage. In 2016, Australia’s Department of Human Services switched to an automated, algorithmically driven system for debt collection.⁴ The system used biweekly income to estimate yearly income without factoring in contract, part-time, and unstable work. An investigation concluded that a quarter of the debt notices by the system were incorrect; thousands of financially vulnerable individuals were victimized as a result of incomplete data.

A growing awareness of algorithmic bias and fake news culture has justifiably produced anxiety around the ways data can be politically and ideologically weaponized. The “post-truth” era is not characterized by a dismissal of facts but by “a conviction that

facts can always be shaded, selected, and presented within a political context” (McIntyre 2018, 6). Counterintuitively, this cultural shift corresponds with the “age of datafication.” Data has become malleable while it has become ubiquitous. Given that, as Joanna Redden observes, we have “little ability to interrogate and challenge how our data are being used” it is not surprising that scholars (and the general public) are increasingly concerned about the validity of data sets (2018).

For better or worse, data anxiety is dramatically reframing conversations about DH.⁵ Adam Kirsch’s assertion in 2014 that the problems with DH “derive from a false analogy between the humanities and the sciences” is distinct from Da’s claim in 2019 that Piper did not properly scale his distance matrices (612). The new fear is not that DH is science but that it is a “would-be science.” Scholars are increasingly asking of DH work, how do I know if this data is right? How do I know if it is complete?

These questions are important because “bad data” can have cataclysmic, real-world consequences. However, bad data can present learning opportunities and reveal new information about literary and cultural texts. I do not mean to suggest that the majority of work done by DH scholars is methodologically unsound. DH scholars have presented compelling evidence to support the veracity of their methods. Rather, I’d like to posit that the “methodological debate” is itself flawed. As Schmidt notes, Da’s method-driven critique of computational analysis “elevates [statistical reasoning] by incanting the language of quantification against itself. . . . She asserts that scientists do something arcane” (2019). Correlating use-value with methodological veracity runs the risk of setting DH back a decade, to a time when scholars like Franco Moretti were positioning literary data as scientifically objective.⁶

Google’s *Perspective* offers an example of the complex relationship between method and interpretation. It also highlights the ways public concerns about data are pertinent to conversations about DH. *Perspective*, created by Google’s Jigsaw division, aims to facilitate “better conversations” by automatically detecting toxic comments. *Perspective* does not ban a set of words or phrases. It uses machine learning to predict whether combinations of words might be perceived as toxic. *Perspective* assigns a numerical score based on the perceived toxicity of a comment. Originally, this score was represented on a scale of 0 to 100, but it is now represented on a scale of 0 to 1: e.g., “I love you” is scored as 0.04 (“likely follows community guidelines”)

while “I hate you” is scored as 0.92 (“may violate community guidelines”). Several companies have partnered with *Perspective*, including *The New York Times* and *Disqus*, a comment-hosting service used by news sources like *Wired*, *Rolling Stone*, and *The Atlantic*.

When *Perspective* went public in 2017, users reported a variety of issues, the most disconcerting of which was a tendency to perceive certain identity categories as inherently toxic:

Table 1. Bias in *Perspective*’s toxicity rankings on a scale of 1 to 100.

Sentence	“Toxicity” score
I am a man.	20
I am a woman.	41
I am a gay woman.	66
I am a gay black man.	82
I am a black woman.	85

The practical implications are alarming. Since *Perspective* is used to semi-automate comment deletion, it risks silencing the voices that are targeted by truly “toxic” online content. It’s important to note that flawed tools like *Perspective* are not forging digital dystopias; they’re recreating systemic violence in online spaces.

Machine learning algorithms are developed using training sets, or data that allows a model to “learn” certain patterns that can then be used in classification tasks. Large and diverse training sets help to ensure a robust model—you want to create a model that will perform well when given content that is distinct from the training data. To train *Perspective*, Jigsaw used a set of 160,000 comments from Wikipedia. Each of these comments was first labeled by annotators as toxic or nontoxic. 5,000 annotators were employed in total, with at least 10 annotators labeling each comment.⁷ While this data set seems large and robust, the diversity of the training set is questionable since Wikipedia contributors are not a diverse group.⁸

Although an “improved” version of *Perspective* is still in use, its initial iteration has joined the ranks of other monstrous algorithms gone awry, including Microsoft’s Holocaust-denying chatbot Tay and Google’s racist photo recognition algorithm.⁹ Several tech media outlets have condemned *Perspective*, criticizing the tool’s overall ineffectiveness (*Perspective* can be tricked with creative grammar and spelling) and inherent bias. Violet Blue biting remarks that the

tool will “be a hit with the alt right” while Robert Epstein considers *Perspective* the perfect illustration of the “stupidity” of contemporary AI—proof that we do not need to worry about the singularity just yet (2017).

Unfortunately, Google’s rationale for creating *Perspective* received less attention. The *Perspective* team defines a toxic comment as a comment that will “likely make others leave the conversation.”¹⁰ The goal was to create a tool that did not rely on one individual’s subjective definition of “toxicity.” While the Jigsaw team wanted to avoid interpretative responsibility, the composition of a dataset is an interpretative act, as “data is never neutral” (Crawford 2017). The assumption that people are, in aggregate, less biased than individuals ignores systemic inequality. Trying to prevent a user from “leaving the conversation” sounds a lot like trying to keep the user comfortable. But what if the majority of users in a community *would* leave if another commentator identified herself as a queer woman? What is sacrificed to keep those users comfortable? In this sense, *Perspective* gave its developers what they asked for (albeit perhaps not what they wanted). Focusing on the limitations of the computational method/data obscures this fact. Faulty acts of interpretation, both of what “toxicity” means and of what *Perspective* actually accomplishes, are the real culprits.

This does not mean that *Perspective* is useless. Its reflection of cultural bias can be leveraged as both a research and teaching tool, if we alter our interpretation of its results. Media theorist David Rokeby’s conception of refraction is helpful in parsing how we might use *Perspective* to understand culture. He notes, “interactive technology is a medium through which we communicate with ourselves—a *mirror*. The medium not only reflects back, but also refracts what it is given; what is returned to ourselves, transformed” (1995, 133; emphasis added). Though writing about interactive art, Rokeby’s conception of refraction applies to both the process and products of computation. *Perspective* offers an uncomfortable reflection of the humans that created and labeled its training data; it shares their cultural biases and limitations.

Data, like Rokeby’s mirror, always returns a *transformed* reflection. We do not know that the phrase “I am a woman” appeared anywhere in *Perspective*’s training set. In all likelihood, most of the comment writers and raters would not have ranked “I am a woman” as a toxic statement. Nor would they have described “I am a woman” as more toxic than “I am a man.” This distortion is not a weakness:

it is precisely what makes computation generative. *Perspective*'s rankings represent cultural biases *pursued to their conclusion*—enough “toxic” statements about women must mean that you think women are toxic. Through refraction, *Perspective* articulates a difficult truth.

Of course, it is still possible to claim that *Perspective*'s insights are redundant. After all, most humanities scholars are aware of systemic bias. We can see familiar academic principles reflected in *Perspective*'s toxicity rankings. For example, Table 1 supports an intersectional perspective, since combining identity categories (such as gender and race) changes a statement's score. The Table supports the notion that oppression is neither relegated to a single category of marginalization (gender) nor is it simply the sum of multiple categories (gender and race). However, *Perspective* does more than demonstrate overt discrimination. We can use *Perspective* to examine limitless combinations of words and ideas, looking, for example, at the ways behavior and description collide with identity to produce “discomfort”:

Figure 1. The Associations Between Actions, Descriptions, and Identity.

Woman actively running.	0.22
Man actively running.	0.10
Woman passively running.	0.12
Man passively running.	0.12 ¹¹

Systemic bias and cultural norms are often invisible. We need to render them visible to change them. *Perspective*'s results are not interpretations, but they prompt further exploration: is the first sentence more “toxic” because of the proximity of “woman” and “actively”? Are there any activities that a woman could perform “actively” that would be considered less toxic than when a man performs them?

It's important to pause and note that *Perspective*'s results do not, in and of themselves, facilitate a “reading” of systemic bias in online spaces. This is evidenced by the fact that *Perspective*'s developers did not recognize that the tool's output was harmful or biased. *Perspective*'s output provides a useful piece of context that, when combined with other pieces of context (e.g., a historical understanding of gender stereotypes), can help to further unpack the way in which bias and language intersect.

In discussing whether a computational tool or result is “useful,” it is also important to remember that scholarly value is not solely

based on the novelty of information, something that DH detractors often forget. The way *Perspective* presents information about bias is deeply affective. Figure 1 and Table 1 are jarring and disturbing in ways that similar factual statements (e.g., “online comments often discriminate against queer individuals”) are not. The output is startling and persuasive. When humanities scholars talk about affect and data, they typically focus on the ways data can manipulate our emotions: “graphical tools are a kind of intellectual Trojan horse, a vehicle through which assumptions about what constitutes information swarm with potent force” (Drucker 2011, 238). This skepticism is useful. However, it’s equally important to recognize data’s destabilizing power.

The affective potential of tools like *Perspective* makes them suited to the classroom. I have taught *Perspective* in an upper-level literature class on big data and canonicity. In this context, *Perspective* is useful because it highlights the fact that computational tools do not always work the way we want them to. However, once students move past the fact that *Perspective* is “broken,” they begin to make observations about what we can learn from the tool. *Perspective* expands students’ understanding of the depth and pervasiveness of systemic inequality. This directly intersects with concerns about the canon; students are quick to realize that if this is how online comments are being curated, then other collections of text (archives, the canon, assigned textbooks) are also shaped by injustice. As Ed Whitley notes, “it can be disorienting to go from feeling that archives provide tremendous access . . . to being told that seemingly bottomless collections of digitized texts are partial, incomplete, and even at times untrue” (2018, 377). This disorientation is beneficial since it teaches students and scholars to more be “savvy” consumers of cultural narratives.

Students enjoy tinkering with *Perspective*; they try to “break” the tool to discover which combinations of words and concepts reflect bias. This sandbox approach to *Perspective*’s output is playful and experimental. It functions as a form of reverse engineering, the process by which a human-made object is deconstructed to reveal its architecture. Conceptually, reverse engineering fits in well with the goals of cultural criticism, since both seek to demystify an object, understanding the forces that went into its creation.

This approach also recalls the backward-tracing detective work of archival recovery, which relies on the scholar’s ability to piece together an unknown whole from fragmentary parts. In her recent work on Australian periodical fiction, Bode observes that, ideally,

the composition of a corpus should be rooted in an understanding of the documentary record. DH scholarship has much to learn from bibliographic studies. In describing archival work, Natalie Davis notes that the archival researcher should not “be lured into thinking that quotations will carry the story without intervention and interpretation. He or she must relish the individual or unusual case, while seeking to understand its singularity” (2013, xiii). Computational work has been criticized for its obsession with trends (Liu 2012). Yet, computation is most useful when used to determine and assess the individual, unusual case. We rarely find what we expect in the archive. Reading biographical and historical context well means reading them “with and against the grain,” flexibly reinventing our understanding of the record with each new encounter (Davis 2013, xi). This type of reading is required in computational work. It involves being quick on one’s interpretative feet—for example, recognizing that *Perspective*, though flawed, can help us better understand cultural bias.

THE ROLE OF CLOSE READING

Perspective tells us something about societal trends, not individual narratives. In this sense, it easily connects to the pattern-oriented aspect of computation. However, this does not address questions about the richness or veracity of literary readings facilitated by flawed data. The computational analysis of literature is most useful when applied to textual features that human readers ignore. As Piper notes, the repetition of stylistic features “etch themselves into the social fabric, like so many furrows in the ground. They are often invisible, because so common” (2018, 3). Computational tools bring these features to the foreground. In literary studies, computers are often used to examine punctuation, word frequency, and sentence structure.

In my work on direct address, I used computational tools to examine the frequency and features of direct address. While it is tempting to view address as a stylistic tick, it is deeply rhetorical. It can evoke sympathy—“[she] wept and sobbed like a child. Perhaps, mother, you can tell what she was thinking”—and fear: “At least some of the horror I took away at four in the morning you now have before you” (Stowe 1889, 477; Danielewski 2000, xvii). It can slip into the second person, placing the reader in the action: “your knife is slippery, and you are toiling like mad” (Sinclair 2005, 52). While

the “dear reader” convention of the Victorian novel is easy to spot, other forms of address recede into an impression of narrative voice. For example, in Mark Twain’s *A Connecticut Yankee in King Arthur’s Court*, the narrator frequently addresses the reader with a conversational “if you consider,” or a “don’t you see.”

My study examined how address fluctuates over time. How did address transition from Charlotte Brontë’s conversational “Reader, I married him” to Italo Calvino’s pithy “you are now reading Italo Calvino’s *If on a Winter’s Night a Traveler*”? (1864, 479; 2012, 1). I also examined how different groups of authors employ address. Feminist narratologist Robyn Warhol argues that nineteenth-century female and male authors address their readers distinctly. According to Warhol, many female authors address their readers sincerely, using the intimate second person pronoun “you” to advocate for social causes. Male authors are more likely to employ address for ironic, stylistic purposes, using the distancing term “reader.” Do these observations hold true across a larger corpus? How does Warhol’s conception of “sincere” address intersect with other groups of disenfranchised authors (such as African American authors)?

To investigate these questions, I developed a computational method for categorizing and detecting moments of address.¹² The code was written entirely in the statistical programming language R and is available on the project’s GitHub repository. Initially I attempted a classification based approach using a training sample of 350 sentences of address and 350 sentences of narration. However, because address is often syntactically identical to surrounding narration and dialogue, a regular expressions-based pattern matching approach that looked for specific words and phrases proved more successful than machine classification.¹³ In the *Routledge Encyclopedia of Narrative Theory*, Irene Kacandes defines address as “vocative formulations that identify the reader directly . . . ; ‘you’ and ‘dear reader’ are the most common [form]” (2005, 4). However, it is not simply the presence of certain words that signals address, but rather their position outside of dialogue:

You must excuse him, he couldn’t help it (327)

“You must not take that fellow to be any specimen of Southern planters” (380)

Both of these excerpts are from *Uncle Tom’s Cabin*. In the first, the narrator addresses the reader, telling her to excuse St. Clare’s behavior. In the second, a steamboat passenger addresses another

passenger. Beyond the quotation marks, there are few syntactic differences between the first excerpt and the second.

For this project, I defined address as sentences outside of dialogue in which the narrator refers directly to the reader or story. There are virtually limitless combinations of words that can signal address. An author might, for example, refer to the reader as “gentle reader,” “my friend,” or “my esteemed colleague.” I focused on a specific subset of terms and phrases: “you, your, yourself, yourselves, reader, readers, our story, my story, our tale, my tale, our narrative, my narrative.” I chose these features based on preexisting scholarship and close reading. The features were repeatedly refined as I close read the results. While I initially included more phrases, such as “my friend,” I refined the feature set to only include terms that returned more true positives than false positives. Consequently, my approach would detect the first quote below as an instance of address, but not the second or third:

Take my arm, gentle reader, and come with me into some street perhaps trodden by your daily footsteps. (Nathaniel Hawthorne, “Old News”)

Reach not after morality and righteousness, my friends. (Jerome K. Jerome, *Three Men in a Boat*)

Where are my manners? I could introduce myself properly, but it’s not necessary. (Markus Zusak, *The Book Thief*)

I then created a regular expressions pattern-matching script that removed quotations, leaving only narration, and extracted any remaining sentences that contained one or more of these keywords.¹⁴ I applied the tool to a 30-novel test corpus, a 2,000-novel corpus, and a corpus of 23 novels by African American authors.¹⁵ All corpora include Anglophone novels written between 1782 and 1923. In the 30-novel corpus, the tool returned 1,883 instances of address. By reading each instance in its textual context, I determined that approximately 80% were correctly labeled (the other 20% were false positives).

The tool returned over 60,000 instances of address from the corpus of 2,000 novels. On average each novel contains 54 instances of address. To provide some context, *Jane Eyre*, a novel associated with address, contains 56 instances. The forms of address I studied proved unexpectedly stable across time.¹⁶ In the 2,000-novel corpus, the forms of address employed by female and male authors adhered to Warhol’s claims: female authors are more likely to use the term

“you” than male authors. Male authors are more likely to use the word “reader” than female authors. Interestingly, of the groups examined, African American authors use “you” the *least* frequently, employing “you” during only 40% of address (compared to the corpus average of 56%) and “reader” for a startlingly 32% (compared to the corpus average of 20%).

While suggestive, the results do not “answer” any questions about literary history. They do not provide a “reading” or “interpretation” of the historic trajectory of address. The method is not perfect. It is reliant on a subjective, partial definition of address. While I can determine precision for the test corpus, I cannot determine recall. In other words, I know the number of false positives, but I do not know the amount of address missed. Determining recall in the test corpus would mean accurately counting, by hand, each instance of address. It is impossible to determine precision or recall for the 2,000-novel corpus. It is important to note that the corpora are limited. As with *Perspective*, any gaps or biases in the textual input are reflected in the output. The corpora used in this project do not represent a complete, or even representative, sample of the novelistic output of the nineteenth and early twentieth centuries.¹⁷

So what exactly can this data tell us? Frequency statistics mean little without an understanding of how these different forms of address rhetorically function in context. In *Reading Machines*, Stephen Ramsay notes that, historically, DH scholarship has tended to treat data as something that is not “in need of interpretation” (2011, 5). Data are not facts. The results from the African American corpus demonstrate this; the data are fraught with the potential for oversimplification. We might conclude that since African American authors use the term “reader” frequently, they are employing Warhol’s “distancing” address. However, this ignores textual, historical, and cultural context.

Most DH scholars now recognize that data are objects for interpretation. As Frederick W. Gibbs and Daniel J. Cohen note, scholars should “have a conversation with the data in the same way that we have traditionally conversed with literature” (2011, 70). While I agree with this assessment, it positions data as *the object of reading*. Close reading is often applied during the iterative process of computation in order to interpret the results and to “improve” the method. For example, in creating my methodology, I close read the initial results to refine my feature set. I interpreted the results to determine whether they were “correct” (actual instances of address). Scholars

often incorporate close reading and historical context to “shuttle between scales”—distant reading may lead to a passage that, once interpreted, creates further queries for computational exploration (Cordell 2015). However, there is a significant difference between “close reading the data” and close reading the text *with* the data. The former positions the data as text, the later as context. Both are necessary, though the latter opens computational work up to more conventional forms of literary analysis that may be useful to a wider scholarly audience. Too often, computational work does not reach this stage.

In the case of the direct address project, though the data did not tell me *why* certain patterns occurred, they did tell me *where* I might find answers. In automatically extracting a large amount of address, the tool created a queryable database of address. This database functions as a finding aid. It enables users to locate moments of address within a given novel and to locate novels that use address in unusual ways. As is the case with *Perspective*, it is important to read this finding aid with and against the grain.

As noted, the results for the African American-authored corpus suggested that these works were outliers in their use of address. Examining instances of African American-authored address in context suggests that these instances of address serve unique rhetorical functions. Consider James Weldon Johnson’s address in *The Autobiography of an Ex-Colored Man*: “If the reader has never been in a strange city without money or friends, it is useless to try to describe what my feelings were; he could not understand” (1970, 61). Johnson’s narrator speaks “sincerely” to real readers. He also distances certain implied readers’ experiences from his own.¹⁸ William Wells Brown makes a similar move: “Reader, did ever a fair sister of thine go down to the grave prematurely? If so, perchance thou hast drank deeply from the cup of sorrow. But how infinitely better is it for a sister to ‘go into the silent land’ with her honor untarnished, but with bright hopes, than for her to be sold to sensual slaveholders!” (2000, 56). Brown interrupts his white readers’ process of identification, foreshadowing contemporary critiques of sympathy and empathy as “fantasies of mutuality, shared experience and touristic invitations to intimacy” (Lather 2009, 19).

Within the African American-authored corpus, the high usage of “reader” does not serve a single purpose. A shared feature set is *not* a shared rhetorical purpose. The current role of interpretation in DH scholarship often obscures this fact. Typically, context is brought to

bear on data. We might, for example, close read the data on African American-authored address with historical context about African American authorship in mind. Scholars close read their data in order to make better sense of their results. While this is a necessary aspect of computational work, it should not be the final stage of interpretation since it encourages scholars to determine “explanations” for patterns in the data. Conversely, bringing data (as context) to bear on individual narratives highlights the ways shared features do not indicate shared meaning.

Frank J. Webb’s *The Garies and Their Friends* offers a counterpoint to Johnson and Brown’s address. Published in 1857, *The Garies and Their Friends* is the second novel written by an African American author and the first to focus on free African Americans living in the northern United States. The novel follows Mr. Garie, a wealthy white Southerner who moves to Philadelphia to marry his enslaved mistress, Emily. Webb addresses his reader twenty-four times (fourteen of which include the term “reader”). These moments achieve numerous rhetorical effects. Webb uses address to accentuate the fictional (crafted) elements of his story, bolstering his authorial agency.

Webb also uses the term “reader” in what seems like a series of stylistic flourishes: “we must now introduce our readers into the back parlour of the house belonging to Mr. Garie’s next-door neighbour, Mr. Thomas Stevens” (1857, 93). This style of address is common throughout the 2,000-novel corpus. It is reminiscent of address employed by authors like William Makepeace Thackeray, “the tune I am piping is a very mild one, (although there are some terrific Chapters coming presently),” and Anthony Trollope, “what was the nature of the reply to Lord Cantrip the reader may imagine, and thus we will leave our hero” (1903, 73; 1901, 310).

Reading *The Garies* with this data in hand hints at the political and ideological concerns underlaying the novel. Webb’s narrator, who is positioned as the author of the story, is not “introducing” his readers into Stevens’s back parlor to paint a more vivid scene. Stevens, the secret cousin of Mr. Garie, instigates a race riot that kills the Garies. Throughout the novel, Stevens and other white characters attempt to exert control over African American bodies by regulating the spaces they inhabit. The narrator’s metaphorical trek through Stevens’s back parlor takes on an added significance: Webb underscores that it is only as narrator that he can traverse this racially charged space.

The irony at the core of this interaction is more explicit when the narrator later “visits” Mr. Stevens’s mansion:

We must pay that gentleman a visit at his luxurious mansion in Fifth Avenue, the most fashionable street of New York—the place where the upper ten thousand of that vast, bustling city most do congregate. As he is an old acquaintance (we won’t say friend), we will disregard ceremony, and walk boldly into the library where that gentleman is sitting. (Webb 1857, 234)

The notion that the narrator and reader are together entering Mr. Stevens’s room crosses diegetic levels. In this sense, the passage is ironic in the same way as other instances of metanarrative. Consider a passage from William Gilmore Simms’s *Beauchampe*, in which the narrator opens a scene by “entering” the home of one of the novel’s protagonists: “We have already formed a sufficient idea of the dwelling which William Calvert occupies . . . the reader will please go with us while we re-enter it” (1856, 369). In this case, irony is derived from the impossibility of entering a fictional space. Webb’s irony is multivalent. The tongue-in-cheek reference to Stevens as an “acquaintance” not “friend” points toward the disruptive nature of Webb’s act: as an African American abolitionist, the narrator enters the space of Stevens, a racist murderer. His irony, operating on multiple levels, transforms a form of address typically used for stylistic purposes into an act of protest.

Though attitudes have started to shift, scholars long criticized *The Garies and Their Friends* for heralding “uplift” and “black bourgeois respectability” (Henry 2015, 578). In this view, Webb’s “sentimental potboiler” is filled with Victorian tropes (Lapsansky 1991, 28). However, the complexity of Webb’s irony disrupts this view. Though Webb’s address is not *overtly* political, his irony is reliant on the inescapable nature of racial identity in nineteenth-century America. Though a gifted author and the reader’s only guide through the story, Webb’s narrator can only fictionally navigate Stevens’s parlor. To witness Webb’s irony is to admit that “bourgeois tropes” cannot be worn or claimed equally irrespective of identity.

Would it be possible to arrive at this interpretation of Webb’s fiction without the use of computational methods? Yes and no. On the one hand, the computational “results” tell us very little. The formal features of Webb’s address (e.g., twenty-four instances, primarily utilizing the word “reader”) do not, in and of themselves, signal the effect of Webb’s interjections. To understand the complex

irony at play when Webb's narrator "enters" Mr. Stevens's mansion, it is necessary to understand Webb's characters and the cultural context informing Webb's authorship. Even if there was an unmediated link between formal features and rhetorical meaning, the limitations of the method would make direct interpretation of the data difficult.

On the other hand, the computational context prompted my examination of Webb's novel. The results from the computational study pointed toward (but did not explain) a trend in the African American-authored corpus that Webb participates in (the frequent use of "reader"). My queryable database of address enabled me to quickly search for moments of address in the corpus that were similar to and different from Webb's. Even if I had discovered and interpreted Webb's use of address without computational tools, I would not have been able to situate it within such a broad historical and cultural spectrum of texts. The computational methods provided, rather than erased, additional context that informed my reading. It is worth noting that reading Webb and other African American authors with and through data did ultimately produce conclusions about literary history—during the end of the nineteenth century, many African American authors used address to question the ethics of literary empathy. These authors subverted popular forms of Victorian address. But the route to this interpretation was circuitous. Instead of data leading to literary history, it led to individual narratives, which in turn spoke to commonalities within specific movements and periods.

INTERPRETATION THROUGH COMPUTATION

Computational algorithms are not monolithic entities that are either "correct" or "incorrect." The two projects discussed above hopefully demonstrate that "bad" data and flawed methods can lead to useful interpretations. *Perspective's* flaws stem from faulty acts of *human* reading and interpretation. Approaching the tool from the position of reverse engineering and recovery allows us to use *Perspective* as a research and teaching tool. In the case of the address project, the tool's limitations meant that I could not directly "answer" questions about literary history. However, the tool facilitated interpretation by guiding me toward specific patterns, narratives, and passages. In both cases, success had more to do with scholarly praxis than any inherent quality of the computational tool.

Most DH scholars are aware of the subjective and incomplete nature of computation. For example, Ramsay's work is based on the notion that computation should unfold "interpretative possibilities," not provide answers (2011, 10). This exploratory, openly subjective conception of computation remains dwarfed by Moretti's hyperbolic vision of scientific objectivity. As Bode notes, literary studies is "hung up on (whether in favor of, or opposed to) this individualistic, masculinist mode of statistical criticism" (2019). This "hang up" points to a breakdown in the communication of DH work and demands a deep restructuring of how such work is articulated.

At the moment, DH scholars are judged both on the novelty of their textual/historical interpretations *and* the veracity of their methods. As I have argued, this "methodological" turn in the critique of DH is partially a response to larger cultural conversations about the dangers and biases of "big data." This is clear from Brennan's piece; he begins his critique of DH by citing instances of political data analytics gone wrong. He focuses on "Ada," the algorithm that informed Hillary Clinton's campaign, noting that "the dream that algorithmic computation might reveal the secrets of complex social and cultural processes has suffered a very public and embarrassing results crisis" (2017). In this view, computational methods aren't sophisticated enough to capture nuance. Consequently, they have not revealed anything new. These criteria are assumed to be mutually exclusive: an interpretation can't be "good" if the method that produced it is faulty. A method is limited or valueless if it did not prompt an original interpretation. This is a harmful oversimplification.

To demonstrate methodological veracity, DH scholars are expected to recount their method in detail, publish their code, and publish their corpus. For the DH scholar, performing a novel interpretation still requires the range of additional contexts used in traditional scholarship: historical context, biographical context, bibliographic context, close reading, and critical theory. Method and interpretation are both valuable, but there is no reason that they must be undertaken by the same scholar (or group of scholars) in the space of a single publication (or series of publications). It is almost impossible to satisfy both criteria within the current, expected bounds of publication. Expecting methodological explication and textual interpretation to unfold in the space of an article or essay assumes that computation is a method *for reading*. It is not. It is, like archival research, a method for creating context. This context may in turn contribute to acts of critical and close reading.

The analogy between computational criticism and archival work highlights why the current expectations are limiting. Scholars do not assume that archival recovery is only useful if it directly and immediately enables new textual interpretations. While archival and editorial work have been devalued in the past, most scholars understand the extent to which these activities are valuable, intellectual, and interpretive. Archival research creates context that scholars (even those unfamiliar with archival practices) might later use in acts of textual interpretation. Because of this, scholars performing archival research and recovery often put a good deal of thought and effort into presenting and disseminating their materials. Conversely, the scholarly community has well-developed standards for utilizing and citing archival resources.

It is not enough to acknowledge that computation is a humanistic activity, reliant on close reading and interpretation. Instead, we must parse out the distinct stages of interpretation that occur during computation. At the moment, interpretation is involved in creating a computational tool, refining the tool, and processing results. These stages are distinct from the final acts of textual interpretation enabled by the data. As my direct address project highlights, interpreting data is distinct from interpreting text *with* data, or using data as context. In most cases, DH scholars openly admit that this phase of interpretation requires a deeper extended analysis. For example, in Ted Underwood, David Bamman, and Sabrina Lee's article "The Transformation of Gender in English-Language Fiction" (2018), the authors remark that they don't pretend to have fully explained the paradoxes in their data. They further note that the essay has come nowhere near exhausting the data's potential. An interpretation derived directly from a cursory examination of data (e.g., close reading the data) should be regarded as a hypothesis. A hypothesis is useful, and indeed publishable, *but* it must be taken up somewhere down the line. At the moment, this last step is underrepresented in literary studies.

Separating these stages of interpretation would benefit scholars working both in and outside of DH. Often, multiple pieces of a computational research project are folded into a single publication or CV line (e.g., code release, article based on work, etc.). It is important to understand that these pieces represent distinct acts of intellectual labor. Creating a data set (like the "finding aid" for direct address) should count as a valuable piece of scholarly labor, regardless of whether they immediately produce novel interpretations. Separating

the stages of interpretation also incentivizes the presentation of the data—how can I publish this data so that it is a resource for a broader community? How can I ensure that others will use the context that I create? Resolving the distant reading debate will involve give-and-take on both sides. Proponents must focus on how they would like their work to be used by non-DH scholars. Critics of computational analysis need to rethink their understanding of methodological validity, and accept the notion that flawed data can still function as a useful form of literary context.

NOTES

- ¹ See Sarah E. Bond, Hoyt Long, and Ted Underwood's response to Timothy Brennan: "'Digital' Is Not the Opposite of 'Humanities'" (2017).
- ² See Andrew Piper's response, "Do We Know What We are Doing?" (2019) and Ben Schmidt's response, "A Computational Critique of a Computational Critique of Computational Critique" (2019).
- ³ See the collection of articles, "The fraud behind the MMR scare," in the *BMJ* (2011).
- ⁴ See Joanna Redden's "The Harm that Data Do" (2018) for a description of Australia's debt collection debacle.
- ⁵ Schmidt makes a similar point when he notes that "The rhetorical tools you can deploy against positivism are strong, but they risk appearing to make it seem—say—that maybe we shouldn't listen to climate scientists" (2019). Schmidt's insightful remark gets at the connection between the DH debate and the sociopolitical climate, but I do not want to suggest that the methodological critique performed by DH detractors is intentionally performative. Rather, I'm suggesting that it is symptomatic of a broader cultural anxiety about data validity.
- ⁶ In *Distant Reading*, Franco Moretti writes that the study of literary history need not involve a "single direct textual reading. . . the ambition [of the research] is now directly proportional to the distance from the text: the more ambitious the project, the greater must the distance be" (2013, 48). Though in his earlier work, *Graphs, Maps, and Trees*, Moretti admits that the quantitative study of literature still requires acts of human interpretation, his approach has been criticized for overstating the scientific objectivity of quantitative work (Ross 2014), rejecting close reading, and ignoring the bibliographic complexities of large textual datasets (Bode 2018). As Ross notes, Moretti's outdated scientific positivism embodies the "optimism of early digital literary studies" (2014). Unfortunately, Moretti's approach has become synonymous with DH for many scholars outside of the field as well as for the popular media.

- ⁷ See the GitHub page for Conversation AI, which underlies the Perspective API: <https://github.com/conversationai/perspectiveapi>.
- ⁸ Wikipedia has been repeatedly criticized in the popular media for furthering a male-dominated, Western worldview. See, for example, Mark Hill's "Wikipedia is Shockingly Biased" (2016).
- ⁹ For more about these tools, see James Vincent's "Twitter taught Microsoft's AI chatbot to be a racist asshole in less than a day" (2016) and "Google 'fixed' its racist algorithm by removing gorillas from its image-labeling tech" (2018).
- ¹⁰ See the Conversation AI GitHub page: <https://github.com/conversationai/perspectiveapi>.
- ¹¹ It is important to note that *Perspective*, like most online tools and search engines, is constantly evolving. This output was produced on August 17, 2019. Reproducing these outputs may be impossible in the future as developers continue to make modifications. This raises questions about the importance of preserving past versions of public APIs. As companies "make mistakes" with data, they are no doubt anxious to bury these mistakes. However, it is important to hold these companies accountable and to study, preserve, and analyze the biases and limitations of their tools.
- ¹² The automated categorization and detection of features can eradicate nuance. For example, in "Gender and Cultural Analytics: Finding or Making Stereotypes?" (2019), Laura Mandell observes that computational projects can reinforce gender binaries (categorizing characters and authors as either female or male). At the same time, the acts of categorization that inform corpus creation and tool design can require scholars to productively reacquaint themselves with seemingly self-evident concepts. In this case, creating an automated system for categorizing address required me to rethink my definition of address and to closely consider its rhetorical function.
- ¹³ The code for the project, written in R, along with documentation, is available on the project's GitHub repository: <https://github.com/DearReader/DearReader>.
- ¹⁴ I needed to develop a fairly sophisticated method for dialogue removal—for example, a method that would deal with commonly occurring punctuation inconsistencies (such as missing quotation marks). I also employed a pattern-matching approach that removes some embedded correspondence. These methods are detailed on the project's GitHub repository: <https://github.com/DearReader/DearReader>.
- ¹⁵ The 30-novel test corpus is composed entirely of freely accessible texts from Project Gutenberg (it includes novels such as *Uncle Tom's Cabin*, *Jane Eyre*, *Frankenstein*, and *Moby Dick*). These 30 novels are available in .txt format on the GitHub repository for this project. The 2,000-novel corpus is composed of texts from the Chicago and Chadwyck Healey corpora, corpora often used in text analysis projects. Many, though not all,

of these texts are available on Project Gutenberg. A full list of the 2,000 novels (including titles, author names, and publication dates) is available on the GitHub repository for this project. The distribution of the 2,000-novel corpus is in some respects uneven: the corpus is approximately 27% female and 73% male authored; 15% British and 85% American; and less than 1% unknown in terms of gender and less than 1% other in terms of nationality.

- ¹⁶ The average number of instances of address per novel is as follows: 97 instances (1782–1849), 53 instances (1850–1899), and 31 instances (1900–1924). However, the decline of address becomes significantly less pronounced if outliers are removed (using the standard deviation): 35 (1782–1849), 35 (1850–1899), and 24 (1900–1924). This suggests that novels that employ unusually frequent address diminish, but that the typical frequency of address does not.
- ¹⁷ For example, several factors limited the size of the African American-authored corpus. The primary issue is that the novels themselves have not been preserved, not been digitized, or have been digitized using OCR and not corrected. The University of Kansas offers an excellent bibliography of African American novels, which includes sixty-five novels published before 1924. I could only find twenty-three of these novels in a digitized, clean format suitable for processing without manual cleanup. The discrepancy between this list and the novels available online offers a stark reminder of the ways in which bias informs the preservation and dissemination of texts.
- ¹⁸ Robyn Warhol begins to discuss this simultaneously distancing and engaging function of address in her article, “Reader, Can You Imagine?” (1995).

WORKS CITED

- Blue, Violet. 2017. “Google’s comment-ranking system will be a hit with the alt-right.” *Engadget*, September 1, 2017. <https://www.engadget.com/2017/09/01/google-perspective-comment-ranking-system/>.
- Bode, Katherine. 2018. *A World of Fiction*. Ann Arbor: University of Michigan Press.
- . 2019. “Computational Literary Studies: Participant Forum Responses, Day 2.” *In the Moment*, April 2, 2019. <https://critinq.wordpress.com/2019/04/02/computational-literary-studies-participant-forum-responses-day-2-3/>.
- Bond, Sarah E., Hoyt Long, and Ted Underwood. 2017. “‘Digital’ Is Not the Opposite of ‘Humanities.’” *The Chronicle of Higher Education*, Nov. 1, 2017.
- Brennan, Timothy. 2017. “The Digital-Humanities Bust.” *The Chronicle of Higher Education*. Oct. 15, 2017. <https://www.chronicle.com/article/The-Digital-Humanities-Bust/241424>.

- Brontë, Charlotte. 1864. *Jane Eyre*. A. L. Burt Company.
- Brown, William Wells. 2000. *Clotel or The President's Daughter*. Boston: Bedford/St. Martin's. First published 1853.
- Calvino, Italo. 2012. *If On A Winter's Night A Traveler*. San Diego, CA: Harcourt. First published 1979.
- Cordell, Ryan. 2015. "Scale as Deformance." *Ryan Cordell*, November 5, 2015. <https://ryancordell.org/research/scale-as-deformance/>.
- Crawford, Kate. 2017. "The Trouble with Bias." NIPS Keynote. Long Beach Convention Center, Lecture.
- Da, Nan Z. 2019. "The Computational Case against Computational Literary Studies." *Critical Inquiry* 45 (3): 601-639.
- Danielewski, Mark Z. 2000. *House of Leaves*. New York: Random House.
- Davis, Natalie. 2013. "Forward." In *The Allure of the Archives*, by Arlette Farge, ix-xvi. New Haven, CT: Yale University Press. First published 1981.
- DearReader. n.d. "DearReader." *GitHub*. <https://github.com/DearReader/DearReader>.
- Drucker, Johanna. 2011. "Humanities Approaches to Graphical Display." *Digital Humanities Quarterly* 5 (1).
- Epstein, Robert. 2017. "Google's Fighting Hate And Trolls With A Dangerously Mindless AI." *Fast Company*, August 31, 2017. <https://www.fastcompany.com/40459339/google-perspective-fighting-hate-and-trolls-with-a-mindless-ai>.
- Gibbs, Frederick W., and Daniel J. Cohen. 2011. "A Conversation with Data: Prospecting Victorian Words and Ideas." *Victorian Studies* 54 (1): 69-77. <https://muse.jhu.edu/article/468193>.
- GitHub. n.d. "Conversation AI." *GitHub*. Accessed August 25, 2019. <https://conversationai.github.io/>.
- Godlee, Fiona. 2011. "The fraud behind the MMR scare." *BMJ*. <https://doi.org/10.1136/bmj.d22>.
- Hawthorne, Nathaniel. 1835. "Old News." *The New England Magazine* no. 8, 91-96.
- Henry, Katherine. 2015. "Garies (The) and Their Friends." *The Encyclopedia of Greater Philadelphia*. Accessed 1 May 2018. <https://philadelphiaencyclopedia.org/archive/garies-the-and-their-friends/>.
- Hill, Mark. 2016. "Wikipedia is Shockingly Biased: 5 Lessons from an Admin." *Cracked*, July 11, 2016. <https://www.cracked.com/personal-experiences-2344-5-ugly-realities-wikipedia-i-learned-as-admin.html>.
- Jerome, Jerome K. 1889. *Three Men in a Boat*. Bristol.
- Johnson, James Weldon. 1970. *The Autobiography of an Ex-Colored Man*. New York: Knopf. First published 1912.
- Kacandes, Irene. 2005. "Address." In *Routledge Encyclopedia of Narrative Theory*, edited by David Herman, Manfred Jahn, and Marie-Laure Ryan, 4-5. London: Routledge.

- Kirsch, Adam. 2014. "Technology Is Taking Over English Departments." *The New Republic*, May 2, 2014. <http://www.newrepublic.com/article/117428/limits-digital-humanities-adam-kirsch>.
- Lapsansky, Phillip. 1991. "Afro-Americana: Frank J. Webb and His Friends." In *The Annual Report of the Library Company of Philadelphia for the Year 1990*, 27–38. Philadelphia, PA: The Library Company of Philadelphia.
- Lather, Patti. 2009. "Against Empathy, Voice and Authenticity." In *Voice in Qualitative Inquiry: Challenging Conventional, Interpretive, and Critical Conceptions in Qualitative Research*, edited by Alecia Y. Jackson and Lisa A. Mazzei, 17–26. London: Routledge.
- Liu, Alan. 2012. "Where is Cultural Criticism in the Digital Humanities." In *Debates in the Digital Humanities*, edited by Matthew K. Gold. Minneapolis: University of Minnesota Press.
- Mandell, Laura. 2019. "Gender and Cultural Analytics: Finding or Making Stereotypes?" In *Debates in the Digital Humanities 2019*, edited by Matthew K. Gold and Lauren F. Klein. Minneapolis: University of Minnesota Press.
- McIntyre, Lee. 2018. *Post-Truth*. Cambridge, MA: MIT Press.
- Moretti, Franco. 2005. *Graphs, Maps, and Trees*. London: Verso.
- . 2013. *Distant Reading*. London: Verso.
- Nissenbaum, Helen. 2009. *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford, CA: Stanford University Press.
- Noble, Safiya. 2018. *Algorithms of Oppression*. New York: NYU Press.
- Piper, Andrew. 2018. *Enumerations: Data and Literary Study*. Chicago, IL: University of Chicago Press.
- . 2019. "Do We Know What We Are Doing?" *Journal of Cultural Analytics*, April 1, 2019. <http://culturalanalytics.org/2019/04/do-we-know-what-we-are-doing/>.
- Ramsay, Stephen. 2011. *Reading Machines: Toward an Algorithmic Criticism*. Champaign: University of Illinois Press.
- Redden, Joanna. 2018. "The Harm that Data Do." *Scientific American*, November 1, 2018. <https://www.scientificamerican.com/article/the-harm-that-data-do/>.
- Rokeby, David. 1995. "Transforming Mirrors: Subjectivity and Control in Interactive Media." In *Critical Issues in Interactive Media*, edited by Simon Penny, 133–59. New York: SUNY Press.
- Ross, Shawna. 2014. "In Praise of Overstating the Case: A Review of Franco Moretti, *Distant Reading*." *Digital Humanities Quarterly* 8 (1).
- Schmidt, Ben. 2019. "A Computational Critique of a Computational Critique of Computational Critique." *Ben Schmidt*, last updated December 5, 2019. http://benschmidt.org/post/critical_inquiry/2019-03-18-nan-da-critical-inquiry/.
- Shelley, Mary. 2006. *Frankenstein*. New York: Penguin. First published 1818.

- Simms, William Gilmore. 1856. *Beauchampe: Or, The Kentucky Tragedy, a Sequel to Charlemont*. Redfield.
- Sinclair, Upton. 2005. *The Jungle*. Boston: Bedford/St. Martin's. First published 1905.
- Stowe, Harriet Beecher. 1889. *Uncle Tom's Cabin: Or, Life Among the Lowly*. Boston: Houghton Mifflin Company. First published 1852.
- Thackeray, William Makepeace. 1903. *Vanity Fair: A Novel without a Hero*. New York: Charles Scribner. First published 1847.
- Trollope, Anthony. 1901. *Phineas Finn Volume III*. Philadelphia, PA: Gebbie and Company. First published 1867–68.
- Underwood, Ted, David Bamman, and Sabrina Lee. 2018. "The Transformation of Gender in English-Language Fiction," *Cultural Analytics* 3 (2). <http://culturalanalytics.org/2018/02/the-transformation-of-gender-in-english-language-fiction/>.
- Vincent, James. 2016. "Twitter taught Microsoft's AI chatbot to be a racist asshole in less than a day." *The Verge*, March 24, 2016. <https://www.theverge.com/2016/3/24/11297050/tay-microsoft-chatbot-racist>.
- . 2018. "Google 'fixed' its racist algorithm by removing gorillas from its image-labeling tech." *The Verge*, January 12, 2018. <https://www.theverge.com/2018/1/12/16882408/google-racist-gorillas-photo-recognition-algorithm-ai>.
- Warhol, Robyn. 1989. *Gendered Interventions*. New Brunswick, NJ: Rutgers University Press.
- . 1995. "Reader, Can You Imagine? No, You Cannot': The Narratee as Other in Harriet Jacobs's Text." *Narrative* 3 (1): 57–72.
- Webb, Frank J. 1857. *The Garies and Their Friends*. London: George Routledge & Co.
- Whitley, Ed. 2018. "Uncle Tom's Cabin and Archives of Injustice." In *Teaching With Digital Humanities*, edited by Jennifer Travis and Jessica Despain, 215–27. Champaign: University of Illinois Press.
- Zusak, Markus. 2005. *The Book Thief*. London: Picador.

GABI KIRILLOFF is an Assistant Professor of English at TCU, where she specializes in digital humanities and American literature. Gabi has worked on several large-scale digital projects, including *The Walt Whitman Archive* and *The Willa Cather Archive*. Much of Gabi's research uses digital tools and computational methods to explore the portrayal of gender in fiction. Her work has appeared in journals including *Cultural Analytics* and *Digital Scholarship in the Humanities*.