

UNIVERSITY OF CALIFORNIA

Los Angeles

Social Reading in the Digital Age

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy in English

by

Allison Hegel

2018

© Copyright by

Allison Hegel

2018

ABSTRACT OF THE DISSERTATION

Social Reading in the Digital Age

by

Allison Hegel

Doctor of Philosophy in English

University of California, Los Angeles, 2018

Professor Allison B. Carruth, Chair

The rise of online platforms for buying and discussing books such as Amazon and Goodreads opens up new possibilities for reception studies in the twenty-first century. These platforms allow readers unprecedented freedom to preview and talk to others about books, but they also exercise unprecedented control over which books readers buy and how readers respond to them. Online reading platforms rely on algorithms with implicit assumptions that at times imitate and at times differ from the conventions of literary scholarship. This dissertation interrogates those algorithms, using computational methods including machine learning and natural language processing to analyze hundreds of thousands of online book reviews in order to find moments when literary and technological perspectives on contemporary reading can inform each other. A focus on the algorithmic logic of bookselling allows this project to critique the ways companies sell and recommend books in the twenty-first century, while also making room

for improvements to these algorithms in both accuracy and theoretical sophistication. This dissertation forms the basis of a re-imagining of literary scholarship in the digital age that takes into account the online platforms that mediate so much of our modern literary consumption.

The dissertation of Allison Hegel is approved.

Ursula K. Heise

Todd S. Presner

Brian Kim Stefans

Allison B. Carruth, Committee Chair

University of California, Los Angeles

2018

Table of Contents

<i>Abstract</i>	<i>ii</i>
<i>List of Figures</i>	<i>vi</i>
<i>List of Tables</i>	<i>x</i>
<i>Acknowledgements</i>	<i>xi</i>
<i>Vita</i>	<i>xiii</i>
1 Introduction	1
2 The Genres of Amateur and Professional Reviews	36
3 The Genres of Goodreads	83
4 Modes of Plot Summary	144
Conclusion	192
Appendix	195
Bibliography	198

List of Figures

Figure 1.1: Model of social reading.	20
Figure 2.1: Navigation categories for book reviewing websites, both amateur (top row) and professional (bottom row).	37
Figure 2.2: Collocations appearing directly beside the terms “science fiction,” “scifi,” and “sf” in professional and amateur reviews.	46
Figure 2.3: Average portion of professional and amateur science fiction reviews dedicated to each topic.	54
Figure 2.4: Percent of amateur reviews discussing genre for four types of genre fiction.	54
Figure 2.5: Collocations appearing within 10 words of the terms “science fiction,” “scifi,” and “sf” in professional and amateur reviews.	56
Figure 2.6: For 20 popular science fiction books, the proportion of shelves Goodreads users place them on that are related to awards won (in blue) versus sentiments elicited (in orange), plotted over time.	59
Figure 2.7: Average portion of professional and amateur romance reviews dedicated to each topic.	64
Figure 2.8: Percent of professional and amateur reviews discussing plot for four types of genre fiction.	65
Figure 2.9: Most common words in reviews, broken down by professional vs. amateur reviewers and low-rated vs. high-rated reviews.	66
Figure 2.10: Average portion of professional and amateur mystery reviews dedicated to each topic.	71

Figure 2.11: The overall portion of all professional and amateur reviews that discussed the book's sentiment, broken down by the genre of the book under review.	74
Figure 2.12: The overall portion of all professional and amateur reviews that discussed the book's context, broken down by the genre of the book under review.	75
Figure 2.13: The overall portion of all professional and amateur reviews that discussed the book's characters, broken down by the genre of the book under review.	76
Figure 2.14: Differences in topic proportions between amateur and professional reviews.	78
Figure 2.15: Variation in topic proportions between amateur and professional reviews for three topics.	80
Figure 3.1: Goodreads homepage with elements related to genre boxed in red.	98
Figure 3.2: Example of a Goodreads recommendation based on genre from the homepage.	99
Figure 3.3: The Goodreads book page for <i>The Ultimate Hitchhiker's Guide to the Galaxy</i> , with its list of genres highlighted.	103
Figure 3.4: <i>Hitchhiker's Guide</i> genre list from book page (left), compared to a list of the most popular shelves for the book (right).	104
Figure 3.5: A visualization of the result of clustering Goodreads shelves based on the books they have in common.	109
Figure 3.6: Close-up of the scifi-fantasy cluster.	110
Figure 3.7: Close-up of the drama cluster.	111
Figure 3.8: Close-up of the horror cluster.	112
Figure 3.9: Number of reviews in dataset broken down by year.	117
Figure 3.10: Diagram of the prediction method.	120

Figure 3.11: Comparing the predictability of reviews from different types of clusters	123
Figure 3.12: Goodreads science fiction landing page (goodreads.com/genres/science-fiction).	125
Figure 3.13: Predictability of science fiction reviews on Goodreads.	126
Figure 3.14: Predictability of science fiction-related subgenre reviews on Goodreads.	129
Figure 3.15: Predictability of creepy reviews on Goodreads.	133
Figure 3.16: Comparison of the predictability of reviews on Goodreads, professional, and academic platforms.	136
Figure 3.17: Predictability of reviews appearing on the first, second, and third page of Goodreads results for a given book.	140
Figure 4.1: Diagram detailing this project's methods.	153
Figure 4.2: Breakdown of this project's corpus, which includes books that have plot summaries on at least three different platforms. The chart shows how many total summaries each platform had in the corpus.	156
Figure 4.3: Stanford CoreNLP dependency parse results for a sentence from the Wikipedia summary of Ruth Ozeki's <i>My Year of Meats</i> .	159
Figure 4.4: Total number of plot summaries and events from each social reading platform.	160
Figure 4.5: Average number of events in each summary, broken down by platform.	161
Figure 4.6: For each character in <i>No Country for Old Men</i> , the chart depicts the percent of all total character mentions attributed to that character as either subject or object.	166
Figure 4.7: For each platform, the box-and-whisker plot shows the distribution of the percent of total character mentions attributed to the characters on that platform.	167

Figure 4.8: The most frequent verbs in plot summaries of <i>Flowers for Algernon</i> by platform.	169
Figure 4.9: The most distinctive verbs in plot summaries of <i>Cat's Cradle</i> by platform.	171
Figure 4.10: The most distinctive verbs in plot summaries of <i>The Da Vinci Code</i> by platform.	174
Figure 4.11: Diagram of the “Rising Action” section of the plot diagram from 7 th -grade lesson plan (Jordan).	178
Figure 4.12: Chart showing the use of four different verbs over time in plot summaries of <i>The Poisonwood Bible</i> from seven platforms.	181
Figure 4.13: Chart showing the use of four different verbs over time in plot summaries of <i>Charlotte's Web</i> on four platforms.	182
Figure 4.14: The most frequent verbs in plot summaries of <i>Gone with the Wind</i> by platform.	184
Figure 4.15: Use of the verbs “hate” and “love” in Goodreads review plot summaries of books in this project's corpus.	188
Figure 4.16: Use of the verb “die” in Goodreads review plot summaries of books in this project's corpus.	189

List of Tables

Table 2.1: Topic guidelines and examples.	51
Table 2.2: Comparison of topics addressed in book review writing guides.	52
Table 2.3: Most predictive words for each topic.	53

Acknowledgements

The faculty, staff, and administration of UCLA's Department of English made this dissertation manageable, from my first weeks settling into a new city to the final frenzied hours of filing. Behind the scenes, the department fought for larger stipend packages and summer funding, organized career events and parties, and made the Humanities Building a cozy home away from home. Mike Lambert in particular was a steady source of wisdom and positivity who averted countless crises with his encyclopedic knowledge of the department, his vast network of contacts, and his willingness to entertain even the silliest questions. In addition, I am grateful to UCLA's Graduate Division, the Mellon Foundation, and UCHRI for funding that made my research possible.

I was also fortunate to receive support from outside my own department, including from UCLA's Center for Digital Humanities and the Stanford Literary Lab. Miriam Posner was the nerve center of the digital humanities program at UCLA, and her mentorship from my very first days of graduate school helped me find my community. Mark Algee-Hewitt welcomed me to Stanford like one of his own students and hosted some of the most productive and provocative events I've ever attended. Digital humanities projects operate at a scale that would have seemed overwhelming without the support of these generous scholars.

I couldn't have asked for a more supportive, accomplished, or inspiring dissertation committee. Ursula K. Heise sets the bar high in every conversation, whether it's a qualifying exam or a post-lecture chat about a Kim Stanley Robinson novel, motivating everyone around her to be the most thoughtful version of themselves. Her vocal support for alt-ac careers and the versatility of an English degree gave me the courage to pursue my current career path. Todd

Presner taught me that the humanities aren't solitary through his incredible ability to build coalitions among different disciplines. His courses are adventures, bringing together humanists, scientists, librarians, and students wielding our varied backgrounds and strengths to unravel 10-week mysteries. His vast knowledge of computational methods, practicality, and humor improved this dissertation immensely. Brian Kim Stefans was a sounding-board and co-conspirator for all of my wildest ideas, helping me to temper them, or—if they had legs—turn them up a notch.

My chair, Allison Carruth, supported me through the writing process with infinite patience and a contagious passion for the field. She has given me a model for the kind of writer, mentor, and colleague I want to become. I am inspired by her commitment to public outreach and her ability to make even the most daunting task seem manageable. Our weekly meetings fleshing out the details of this project remain my fondest memories of graduate school as we formulated hypotheses and tested ideas new to both of us.

And finally, I would never have finished this daunting project without the support of my closest friends, who kept me sane for the past five years. Kevin's stories made me laugh even at my most overwhelmed, reminding me that there was more to life than my research. Zach sees the world from a different angle, pointing out what everyone else missed and bringing it into focus, from the most sweeping literary theory to the smallest turns of phrase. Every good idea in the following pages spawned from our conversations. Matt has been my strongest source of support and my greatest advocate. His unfaltering conviction that I was equal to whatever task I took on, whether learning a new method or writing a new chapter, has expanded my small world and made me immeasurably excited for what's next.

Vita

Allison Hegel graduated from the University of Chicago in 2012 with a double major in English (with honors) and Fundamentals: Issues and Texts (with honors). She has presented work on genre and social reading at the MLA Annual Convention, the University of Pittsburgh, Stanford University, Berkeley University, the University of British Columbia, and UCLA. As a graduate student at UCLA, she served as the Digital Project Manager of the Cruikshank Digital Archive, General Editor of *The Programming Historian*, a Graduate Fellow at the Getty Summer Institute, and a Digital Project Scholar at UCLA Library Special Collections. Most recently, she was awarded an Insight Data Science Fellowship.

Chapter 1: Introduction

Now more than ever, literary scholars must bring their skills to bear on digitally networked reading. Researchers who are versed in reading's many cultures, economies, and conditions of reception know that it is never possible for a reading platform to be a 'passive conduit.' For reading has always been social, and reading's economies, cultures of sharing, and circuits of travel have never been passive. (Nakamura, "Words with Friends" 243)

Reading has always been social, but the platforms mediating our social reading have never been more active. These platforms have expanded on older, analog methods of cataloging readers through clear demographic and product categories by developing newly specific user segmentation schemas powered by algorithms and fueled by users' personal data. Online social reading platforms, which allow readers to write their own reviews and even create their own genres and affinity categories using social tagging systems, have taken over the role of predicting readers' preferences and reifying genre boundaries.

As these amateur review sites have grown more popular, they have begun to worry traditional publishers. Particularly in the film industry, the rise of amateur movie reviews paired with declining profits has caused the Hollywood establishment to notice—and attack—sites like Rotten Tomatoes, where users submit reviews and star ratings that have a tangible impact on new movies' success in the box office. The film industry's concern is the popularization of movie reviewing, a task once reserved for professional critics but now open to anyone with an internet connection: "Some filmmakers complain bitterly that Rotten Tomatoes casts too wide a critical net. The site says it works with some 3,000 critics worldwide, including bloggers and YouTube-based pundits. But should reviewers from Screen Junkies and Punch Drunk Critics really be treated as the equals of those from The Los Angeles Times and The New Yorker?" (Barnes). As this article frames it, the debate between amateur and professional reviewers is a question of

quality. And indeed, amateur reviews, which aren't curated or moderated, often exhibit a variety of biases: some sites are heavily skewed towards positive ratings (Hickey, "Ghostbusters"), many people rate movies before they even come out based on their political or cultural leanings (Hickey and Mehta), and men have even been shown to tank the ratings of movies meant for women (Hickey, "Men Are Sabotaging"). These same concerns arise in amateur book reviewing, as Goodreads eclipses traditional book reviewing platforms, and its reviews play a more significant role in readers' purchasing decisions.

In the shift from professional reviewers with mass audiences like Michiko Kakutani, the former chief book critic for *The New York Times*, to a distributed system of millions of amateur reviewers, audiences have become fragmented. Instead of critics like Kakutani dictating the success of a new book, algorithms have taken over the role of curator, recommending personalized content to more and more finely broken-down demographics. We see these increasingly specific consumer categories across the web, driven by technology companies who offer free services at the cost of users' personal information, which then helps to drive further demographic segmentation and targeting. Goodreads provides a platform to share and read literary opinions in exchange for your preferences and reading history, using genre fandom as an easy starting point for these categorizations. These platforms go further than publishers ever could in predicting and personalizing content to their customers. However, they work best when products are easily categorized, and work badly in the case of ambiguity, often causing public backlash. For example, take filmmaker Martin Scorsese's explanation for the negative popular reaction to the film *mother!*: "People seemed to be out for blood, simply because the film couldn't be easily defined or interpreted or reduced to a two-word description. Is it a horror movie, or a dark comedy, or a biblical allegory, or a cautionary fable about moral and

environmental devastation? Maybe a little of all of the above, but certainly not just any one of those neat categories” (Scorsese). Professional filmmakers and producers as well as book publishers work hard to ensure that their offerings fit neatly into clear market segments, lest they face the confusion and disappointment of their audience.

As companies double down on categorization to drive their marketing and algorithmic recommendations, the way we choose, evaluate, and interpret books has begun to shift. As literary critic Mark McGurl has noted, Amazon “transforms literary experience into customer experience” (“Everything and Less” 455). Companies like Amazon, with their well-known and obsessive focus on the customer, insist that they know their customers better than they know themselves, and can suggest the perfect product before customers even know they need it. But these marketing claims become self-fulfilling prophecies as Amazon’s algorithmic design goes beyond predicting what people want to shaping what people want. The evidence for this transformation lies in online reviews, where readers’ initial reactions to books transform under the influence of technology companies’ ideologies and priorities.

Parsing readers’ reactions has typically fallen within the domain of reader response criticism, as I will discuss below, but the field has only begun to scratch the surface of this new landscape of amateur criticism. We need to shift the scale of our analyses from individual examples and small-scale surveys to the types of large-scale algorithmic data mining these companies use themselves to fully understand the effect they’ve had on the reading practices of millions of users. Digital humanists have taken on the task of transforming literary methods into algorithms that can manage the vast increase in the quantity of available literary material. The Stanford Literary Lab alone has tried out methods that work at the scale of the word (Allison, Heuser, et al.), the sentence (Allison, Gemma, et al.), the paragraph (Algee-Hewitt, Heuser, et

al.), and the archive as a whole (Algee-Hewitt, Allison, et al.). While the field is still identifying and inventing best practices, digital humanities is best equipped with methods that can handle the vast quantity of twenty-first-century literary material.

With the rise of new online platforms that allow amateur book reviews to achieve the same prominence as professional reviews, amateur definitions of literary concepts like genre and plot have become a new source of literary insight. These reviews are a rich source of first-hand insight into modern readers' expectations for and reactions to books. The first two chapters of this dissertation examine how genre has changed in the twenty-first century due to new social reading platforms and marketing algorithms. In the first chapter, I analyze the language and context of book reviews to compare the ways amateur and professional reviewers discuss genre. I find that, for four popular types of genre fiction, professional reviewers have formal definitions of the genre while amateurs focus on more personal definitions, reinforcing the idea that these amateur platforms encourage their users to think of genre in personal terms. In this way, I argue, genre on these amateur platforms has transformed from a literary category to a consumer category, with its key purpose to differentiate readers into marketing demographics based on their identification with genre readership communities.

In the second chapter, I examine the effect of genre's shift from literary to consumer category on the most popular amateur social reading platform, Goodreads, to see what impact the site has had on millions of readers' ideas of genre. Using machine learning methods, I find that, while Goodreads' shelving system allows a diverse range of user-created, non-traditional genres, over the past decade, the Goodreads platform has in fact seen its users' depictions of genre grow less diverse. As users' conceptions of genre become more predictable, platforms can more easily recommend books and target advertisements, but readers become less likely to encounter

surprising examples of a genre or to have their notions of genre challenged. As amateur social reading platforms continue to thrive, their standardizing effect will only grow more effective, further shifting the meanings and functions of genre in the twenty-first century.

While genre remains one of the most important ways people evaluate a book, plot plays an equally important role. In the third chapter, I detail the diverse ways readers from different reading communities on social reading platforms summarize literary plots. Using natural language processing to extract events from plot summaries, I outline four modes of plot summary, each with its own style, audience, and goals. Social reading platforms use these modes to achieve different purposes, from creating a balanced portrayal of the plot to driving user engagement through suspense and emotion. I find that Goodreads in particular emphasizes emotionally-charged events through its algorithmic system for sorting reviews, bringing certain types of events to the forefront for its users. The design of social reading platforms shapes their users' conceptions of plot, an effect that is observable in the way these platforms encourage their users to summarize plot.

The rise of online platforms for buying and discussing books such as Amazon and Goodreads opens up new possibilities for reception studies in the twenty-first century. Online reading platforms rely on algorithms with implicit assumptions that at times imitate and at times differ from the conventions of literary scholarship. This dissertation interrogates those algorithms, using computational methods including machine learning and natural language processing to analyze hundreds of thousands of online book reviews in order to find moments when literary and technological perspectives on contemporary reading can inform each other. A focus on the algorithmic logic of bookselling allows this project to critique the ways companies sell and recommend books in the twenty-first century, while also making room for improvements

to these algorithms in both accuracy and theoretical sophistication. This dissertation forms the basis of a re-imagining of literary scholarship in the digital age that takes into account the social reading platforms that mediate modern literary consumption.

Social Reading in the Digital Age

There are several questions at stake in any study of social reading:¹ who reads? What do they read? How do they choose what to read? And how do the books they read position them within certain reading communities and hierarchies? In addressing such questions, critics' approaches to studying social reading have shifted in scale alongside the conventions of literary scholarship. While their studies come from a variety of critical perspectives—new historicism, formalism, and ethnography, among others—there are dominant trends in how they analyze reading practice in different time periods. As I will detail below, studies of nineteenth-century reading in the United States focus on the archive, where extensive knowledge of the historical context as well as contact with artifacts of reading like published book reviews, library records, diaries, and letters are the raw material critics use to piece together their depictions of American reading practice. However, these archival studies tend to elide groups whose reading took place outside of the traditional historical archive, such as working-class and African American readers. Studies of reading in the twentieth century, when close reading came to dominate literary criticism, begin to focus closely on the individual through ethnographic surveys and interviews,

¹ This dissertation conceives of the process of reading broadly, beginning before a reader even picks up the book to long after they put it down. Before reading, both the reader's process of book selection—often influenced by the opinions of friends or a community they belong to—and her preconceived notions about it are products of a reader's social context. After reading, readers often discuss their reaction with others either in person or on a social reading website. This broader conception of reading, which takes into account the social context of a reader's experience with a book, is what I will refer to as "social reading." I will elaborate on my definition of social reading in "Social Reading, Past and Present" below.

filling the gaps in the archive with first-hand accounts of previously overlooked reading practices. Today, as close reading methods encounter competing paradigms like Franco Moretti's "distant reading," the scale of analysis has widened yet again, for example via investigations of the educational and cultural institutions that dictate massive and public reading practices. In this vein, my project interrogates twenty-first-century social reading institutions and digital platforms like Goodreads in order to make sense of the reader's newfound position as both data to be mined for advertising purposes, and a reading public in a newly digital context.

Studies of social reading in nineteenth-century America tend to rely on small numbers of canonical texts to reach conclusions about the history of American reading more broadly. Richard H. Brodhead's *Cultures of Letters: Scenes of Reading and Writing in Nineteenth-Century America* offers detailed readings of Hawthorne, Melville, Stowe, Alcott, Jewett, and Chesnut, arguing that critics understand American literary history as a "history of the relation between literary writing and the changing meanings and places made for such work in American social history" (8). Brodhead's method, in each chapter, is to highlight a distinctive theme during the nineteenth century—whether debates about the canonization of literary taste (45), slavery and its relation to corporal punishment (47), or the emergence of public entertainment contrasted with a newly private idea of domestic life (53). He then tracks that theme through one or more canonical works of fiction, showing how the social context of that time period influenced the construction of those novels.² Another study of nineteenth-century reading by Sarah Wadsworth

² For example, Brodhead argues that we have to take into account the complex social valence of corporal punishment in an antebellum America conflicted over slavery in order to fully understand contemporary books like *The Scarlet Letter* and *Uncle Tom's Cabin*, which model contrasting theories of discipline. Brodhead complicates a reading of discipline in antebellum American novels by going beyond the idea that novels simply either promote or subvert corrective ideologies (47). Instead, Brodhead observes that novels in this antebellum moment use discipline so often because of the social context that placed discipline within the domestic sphere—the same place that novel-reading was privileged. Thus novelists of the time had to contend with the dominant specter of discipline governing family life.

explicitly states that it “is concerned less with individual historical readers than with large groups of readers joined through shared literary practices,” using case studies of Hawthorne, Alcott, Twain, and James to construct these practices (9). While there are some studies of nineteenth-century reading that take into account individual readers through letters and diaries (Hunter; Sicherman; Zboray and Zboray), marginalia (Jackson), and library records (Augst; Pawley; Todd), these studies still focus on readers’ reactions to canonical texts, or to the only slightly wider body of texts chosen by authorities such as the Book-of-the-Month Club.³ Noting the tendency to draw on canonical examples to develop accounts of national cultures in the nineteenth century, Lawrence Buell questions “the legitimacy of reading ‘the national’ through *N* number of putative touchstone narratives” (4). Even so, his own examination of the concept of the Great American Novel centers on case studies of canonical nineteenth-century novels, which he uses as the foundation for his larger historical claims.

Even when studies expand their view beyond the canon to look at readers’ responses to these books, they focus on professional reviews rather than responses from the amateur reading public (Baym; Machor, *Reading Fiction*; Cantor and Shuttleworth). The dominance of the professional reviewer wanes by the end of the nineteenth century, when “the influence of popular taste threatened to dwarf that of professional critics,” a trend that continues today (Wald and Elliott 39). In studies of nineteenth-century reading, however, the focus on professional reviews prevails. These reviews show the literary establishment’s cultured and calculated views of novels, carefully crafted for their reading audience. For example, Nina Baym analyzes professional reviews of novels in periodicals in the first half of the nineteenth century, citing

³ A possible reason for this focus on canonical and authoritative texts is the scarcity of materials outside of the canon. In Janice Radway’s study of the Book-of-the-Month Club in the mid-twentieth-century, she uses the lack of surviving evidence of readers’ responses to the club’s chosen books to justify her third section, which focuses on her own experiences with these books (*A Feeling for Books*).

reviewers who were typically well-educated (20) and male (21). Baym’s study frequently notes reviewers’ comments on public taste and their attempts to influence readers’ choices for the “better,” setting reviewers apart from the general public.⁴ We see this primacy of professional over amateur reviews in James L. Machor’s study of reading in antebellum America, where he focuses one of his four case studies on a non-canonical author, Caroline Chesebro, alongside chapters on Poe, Melville, and Sedgwick. Machor also looks at reviews of these authors, but chooses to read professional reviews from established magazines rather than examine artifacts that shed light on amateur readers, claiming that these reviews are better representative of writers’ conceptions of their own public reception (*Reading Fiction* x). In these studies, which use reviewers’ constructions of the popular reader in their reviews as a proxy for the reading public, we are limited to a second-hand view of nineteenth-century American readers.

These and other studies of nineteenth-century reading leave large portions of the reading public unexamined. For example, McHenry’s *Forgotten Readers* calls attention to the “historical invisibility of black readers” in the early nineteenth and twentieth centuries (4). Instead of studying literary experiences in classrooms, libraries, and other sites of reading that have systematically excluded African Americans, McHenry argues that we should expand our attention to include “nonacademic venues like churches, private homes, and beauty parlors,” which show a more diverse history of social reading (McHenry 10). In studying online social reading platforms, which typically have a low barrier for entry and a diverse user base,⁵ my

⁴ See Baym’s third chapter, which is an account of how reviewers negotiated the taste and demands of the popular reader (Baym 44–62).

⁵ While online social reading platforms are more diverse relative to face-to-face reading groups, there are still many inequalities in access to these online social reading platforms. Lisa Nakamura’s work argues that access to the Internet is more complicated than a binary “have” or “have not,” instead arguing for an understanding of how different races, genders, nationalities, and classes have varying degrees and types of access (*Digitizing Race* 15). Even the way that we measure access to the Internet obscures these differences, notably by considering Internet

project answers McHenry's call to study reading that takes place in unofficial and informal places. My project views non-traditional sites of social reading and discussion as areas of analysis that are in some ways more democratic and meaningful to readers whom previous discussions of social reading excluded. At the same time, however, I will also examine how these online platforms raise new barriers for entry and create new social hierarchies that limit users' reading experiences based on their gender, race, nationality, and class, as McHenry critiques of previous forms of social reading.

As opposed to studies of reading in nineteenth-century America, which use evidence from canonical texts and professional reviews, studies of the twentieth century tend to base their conclusions on individual readers' reactions, drawing influence from both new criticism and reader response theory. However, in replacing the canon with close attention to individual readers' habits and reactions to texts, these studies tend to build their analyses on evidence from only a few groups or individuals. Studies of twentieth-century reading use methods from reader response theory and ethnography to evaluate the experience of reading first-hand, but the individuals they choose to interview rarely represent a wide range of the American reading public. Even Janice Radway, whose interviews of women who read romance novels incorporate women into an otherwise male-dominated view of reading culture, nevertheless leaves open to further study "how other social variables like age, class location, education, and race intersect with gender" (*Reading the Romance* 9). Elizabeth Long's *Book Clubs* attempts to address this gap by adding race to the story of women's reading groups, but in less detail than McHenry's study of African American book clubs.

users (and Asian Americans in particular) to be passive *consumers* of online content rather than acknowledging the many ways in which they are also active *producers* of online content (177–78).

We can see the changing focus of twentieth-century reading studies when we compare Baym's reading of nineteenth-century professional reviews, which she took from 21 major American periodicals, to Amy Blair's reading of early twentieth-century professional reviews. Blair looks closely at the *Ladies Home Journal*, where reading advice columnist Hamilton Wright Mabie instructed casual readers in what and how to read so that they may be considered cultured. Blair credits Mabie with creating the idea of the middlebrow—which Radway and Joan Rubin examine as it develops in later decades—by placing popular realist novels alongside highbrow works in his columns (8). Mabie treads the line between highbrow and middlebrow by only recommending highbrow fiction—such as *The Rise of Silas Lapham*—that left room for romantic and sentimental “misreadings” and therefore ensured that readers would choose the proper books, even if it was for the wrong reasons (Blair 97–98). Blair's study details how early twentieth-century readers practiced “reading up,” or choosing reading material that promises to raise one's social status in a practice that often conflates the language of aesthetics with the language of economics (2–3). Yet this conclusion is based largely on the influence of one columnist, whose writing, while influential, tells us little about diverse readers' reactions to his columns and to the books he recommended.

As we enter the digital age, and the object of study shifts to postwar fiction, critics now focus on the systems and institutions that govern individual responses to literature, be they universities (McGurl, *The Program Era*), technologies (Chun, *Programmed Visions*), or governments (Foucault). McGurl's history of the American creative writing program, for example, focuses on “the actual institutions, technologies, and practices from which postwar fiction emerges” (*The Program Era* 31). New media theorists, most notably Wendy Hui Kyong Chun, similarly call attention to the neoliberal logic behind the technologies that run our

companies, governments, and social lives, “creating ‘informed’ individuals who can overcome the chaos of global capitalism by mapping their relation to the totality of the global capitalist system” (*Programmed Visions* 8). In this work, Radway’s intimate interviews with female readers give way to massive databases that gather more and more detailed personal information via digital surveillance, turning both books and readers into commodities. These theories are crucial to understanding corporations like Amazon and Goodreads, which allow unprecedented freedom for readers to preview and talk to others about books, but also exercise unprecedented control over which books readers see and buy, as well as how they respond to and communicate about them.

Book Reviews Online

In investigating how readers make sense of the books they read, literary scholars have studied not just the book, but also its context of author, reader, period, genre, and other factors that influence processes of interpretation. However, in this body of scholarship, one aspect of a book’s context has been neglected. For many readers in the era of networked computing and social media, the first step in reading a book isn’t turning to the title page, but instead turning to a web page where they can read reviews of that book to decide if it’s worth reading. These reviews become supplemental to the book itself as they influence readers’ choices of what to read, provide a context and ready-made evaluation of the book for readers before they read it, and finally record readers’ own interpretations of the book after they read it. Reading these reviews alongside the books they promote, critique, and otherwise evaluate gives us a more

complete picture of the way readers construct meaning from what they read, and how digital reading institutions and information technology both aid and interfere with this process.

While one might argue that amateur reviewers are unsophisticated and that advertising campaigns too often influence their reviews, it is precisely this democratization of authority and entanglement with marketing that best characterizes book publishing today, and which has, I will argue, profound effects on the way books inspire communities in the digital age. While twentieth-century studies often turn to book clubs to determine readers' responses to books, these twenty-first-century social reading websites are taking over the role of book clubs, boasting massively larger membership and a similarly increased impact on literary trends.⁶ And even though amateur reviews see little to no editing and their authors are not always trained or experienced in literary criticism, each review tells us about that reader's immediate and unpolished reaction to the book, giving us access to a range of responses we might never see if confined to professional reviews, which have standards for length, style, and literary value that amateur reviewers rarely heed. In the same way that a social scientist gives credence to every survey response, this project treats every review as a substantial answer to the question, "What did you think of this book?"—mediated through the oddly both personal and public medium of a particular online platform. By using digital humanities methods such as machine learning and natural language processing to analyze reviews at scale, there is no need to sift through reviews to examine only a select few deemed to be particularly representative or insightful. Instead, by analyzing every review, long or short, thoughtful or dismissive, we can gain a more comprehensive and nuanced understanding of the wide range of reactions readers have to the books they read.

⁶ While social reading websites' larger membership is easily measured (Goodreads alone had 20 million users as of 2013), the hypothesis that they have a greater impact on literary trends is one that my dissertation seeks to prove.

There are, however, some limits to the reviews I consider. My project focuses on English-language reviews and the American (.com) versions of the websites that host them. American companies own most social reading websites and target them at mostly American (though increasingly international) audiences.⁷ Another project might explore how other countries and cultures use these sites differently. Still, even with this limitation, it would be impossible for a human to read every one of these reviews. This project uses distant reading techniques to find patterns and outliers in the reviews, giving weight to each reader's individual voice but reaching conclusions at scale.

In studying how different communities of readers read and discuss books, I'm entering into a well-traveled field, but my raw material—hundreds of thousands of online book reviews—is unique. This project tracks how new digital tools have altered social reading practices in their shift from literary to monetary value and from face-to-face to online reading communities. I hope to restore content to the empty shells of books that appear on social reading website pages, and ultimately re-examine some of the assumptions that have become standard features of digital tools, from their reliance on personal data for accurate book recommendations to their overly simplistic depictions of genre and plot.

⁷ The most popular social reading websites are all founded and headquartered in the United States: Amazon (ranked the 6th most visited website internationally), Goodreads (rank 273), Scribd (rank 411), and LibraryThing (rank 15,910). These rankings are from Alexa, which, notably, is a subsidiary of Amazon. Companies such as Quantcast track the geographic breakdown of these sites' users, showing that they are overwhelmingly from the United States. For example, 49% of Goodreads users are from the United States, with the next most users coming from the United Kingdom (7%).

Social Reading, Past and Present

A review of the history of reading in America since the nineteenth century shows a trend of increasing access to books—from growing literacy to improved publishing technology—that continues in the modern practice of digital publication and consumption.⁸ Rather than reading writ large, the tradition modern digital reading practices build on is *social* reading: not the solitary nobleman in his study,⁹ but the laborers reading aloud in a public house, the students repeating lines of a poem, the women in a book club, and most recently, the digital natives and baby boomers on a social reading website. As is clear from these examples, social reading is not a new phenomenon, as much as social media companies claim to have revolutionized the way we read and interact. *The American Novel 1870-1940* details the many ways in which social reading played a role in entertainment, education, and civic life in nineteenth-century America, from the family home to the workplace (Wald and Elliott 38). More broadly, scholars have placed social reading at the center of early American constructions of a class-based hierarchy with elites at the top (Levine 255–56), and social reading continued to form the glue of American identity formation well into the nineteenth century (Machor, *Reading Fiction* 3). At the same time, however, social reading also disrupted any sense of American identity: Isabelle Lehuu sees social reading as a catalyst that exposes social differences and opposing values (9). Shifting from a national perspective on reading to a personal one, sociologists have studied book clubs and reviews to determine how social reading influences people’s book selections and self-positioning in a hierarchy of middlebrow versus elite culture (Radway, *A Feeling for Books*; Rubin). These

⁸ See Cavallo and Chartier's *A History of Reading in the West*, and Chartier's *The Order of Books*.

⁹ For the dominance of this image in historical depictions of reading, see Long's *Book Clubs* (8–11). For depictions of women reading, see Phegley and Badia's *Reading Women*.

earlier examples of social reading, with their analog contexts, are not far removed from how we conceive of social reading in the digital age. As Alan Liu argues, we replicate this dichotomy between solitary and social reading today in the debate between privacy and social media hyper-sharing (Liu, “The Big Bang of Online Reading” 278). While Liu emphasizes the similarities between online and analog reading, Ed Finn’s work on contemporary literary and authorial networks shifts the conversation from personal social media use to what he calls “the social lives of books,” or the digital traces of reading that millions of readers post on social reading websites by rating, reviewing, and buying books online (*The Social Lives of Books* 3). While earlier conceptions of social reading tend to construct top-down models of literary influence in which published authorities determine readers’ consumption, Finn’s definition recognizes readers as active and influential on digital social reading platforms:

Readers, then, are people who engage with books as social experiences, discussing them in group fora and engaging in other more or less public acts of literary exchange. This engagement does not begin and end with the interpretation of marks on a page, but rather starts with the complex series of social and economic transactions involved in selecting and purchasing a book and continues indefinitely as the reader continues consuming, discussing and contributing to literary culture. (Finn, *The Social Lives of Books* 9)

I follow Finn’s expanded conception of reading to include not just the physical act of looking at a page, but also how readers choose books in the first place, and what they do with those books after they read them. In accounting for the acts that precede and follow reading itself, my dissertation situates social reading in relation to the long history of reading, as well as in relation to modern social media and networked communication culture. Social reading, then, re-conceives of reading as an activity that is less and less solitary in the post-Internet age, enmeshed as it is in networks of relationships, technologies, production, and advertisement.

I should distinguish social reading from “sociable reading,” a term Cameron McLachlan uses to denote a type of reading that emphasizes enjoyment and shared community at the

expense of attention to the book's content and the difficult questions it raises—for example, a museum celebrating Katherine Mansfield that in numerous ways undermines the attention to detail, multifaceted personality, and distrust of historical artifacts apparent in her writing (70–72). Sociable reading, rather than social reading, describes what popular social reading platforms like Goodreads promote. Companies set up these platforms using standard database structures that are designed to hold any type of data, not just literary data. Amazon shows this interchangeability most clearly by using the same layout template for every type of product it sells, from books and music to groceries and jewelry. Goodreads' and Amazon's platforms place little emphasis on the content of the books they feature, instead providing elements like “helpfulness ratings” and “Customers Who Viewed This Item Also Viewed” that connect people in superficial ways and most importantly keep people involved in the community. By bringing user reviews into conversation with the content of books, I apprehend under what conditions and to what effects digital platforms built to promote sociable reading in fact facilitate deep social reading practices and communities. While online social reading platforms are typically designed to gloss over complexities of content and context, my project intends to show the diverse ways in which readers grapple with these complexities in their reviews.

Besides positioning books, readers, and readers' personal data as commodities, the digital age has also shifted the terms of our ever-present anxiety about machines. Studies of social reading in previous centuries overwhelmingly focus on the standardization of public education, popular reading material, and cultural value as the biggest contemporary threats to the reading public (Chartier; Lehuu; Machor, *Readers in History*; Machor, *Reading Fiction*). Today, however, personalization (from medical care to website profiles) has become the key threat to agency in the neoliberal age (Foucault; Chun, *Control and Freedom*). Social reading platforms

offer readers opportunities to share their own readings of a text, and receive personalized recommendations based on their own and others' reviews. As a result of social reading companies' investments in personalization, the publishing industry relies on consumer feedback to determine what to publish and how to market it, turning readers into content creators whose reviews feed the algorithms that shape other readers' choices and perceptions of books.¹⁰ While tech companies advertise personalization as a feature, both standardization and personalization limit readers' access to a diverse range of books in a time when more books are accessible than ever before.


Studying Social Reading: How do (and should) we study reading?

Studying social reading calls for a different technique than studying reading alone. Since social reading involves so many influences outside of just the reader, it will be useful to build a model of social reading to account for the interactions between each agent, including authors, readers, books, and book reviews. Each discipline that uses modeling uses it in widely varying ways, from biology's models of cell and chemical interaction, to climate science's models of weather systems, to linguistics' models of language structure and learning. While scientists and social scientists are more likely to use models than literary critics, models are a particularly useful tool for this project, which aims to study not individual agents (e.g., readers) but interactions between agents (how online social reading platforms influence readers and their reviews). Literary critics might view the goal of modeling reading as mechanical or reductive, but as Andrew Elfenbein argues, models are useful for literary critics because they don't

¹⁰ For more on readers as content creators in the digital age, see Cavallo and Chartier (27) and Córdón García et al. (4).

overlook the parts that are more routine, but still central to the act of reading (485). As long as literary scholars recognize that no model can ever include every possible influence, scholars can use them to better understand the complex and multifaceted social reading process.

Liu’s models of literary activity provide a starting point for my model of social reading. Below are two models Liu draws to highlight the democratic effect literary theory has had on the sites of analysis available to the study of reading. The model on the left is Liu’s “core circuit of literary activity: authors, publishers, readers (and interpreters) mediated by documents” (“From Reading to Social Computing”). On the right is what happened after the 1960s, when literary theory complicated the core circuit:



(Figs. 1 and 5, Liu, “From Reading to Social Computing”)

The model on the right adds many more inputs and outputs that influence the reading process, as well as acknowledging that these influences are bidirectional. This second model also includes non-traditional sites of social reading, as advocated by McHenry: instead of studying just canonical authors in major publishing cities, literary scholars now look to “the tavern, coffeehouse, Grub Street, and other scenes on which historical writing and reading were a hubbub of collective literary life and where there were no stable distinctions between primary and secondary players” (Liu, “From Reading to Social Computing”). In its wider reach, Liu

argues that literary theory’s expanded model is a more democratic and representative picture of the literary landscape, which has often ignored the marginal—whether marginal notes and ephemera, or marginalized people.

Liu’s models show us the change that literary theory has brought to the way we study reading, and also the deep complexity any model of reading in the digital age must represent. My own model, pictured below, does not include every possible influence on the reading process, but instead highlights the most important actors in my study of social reading online:

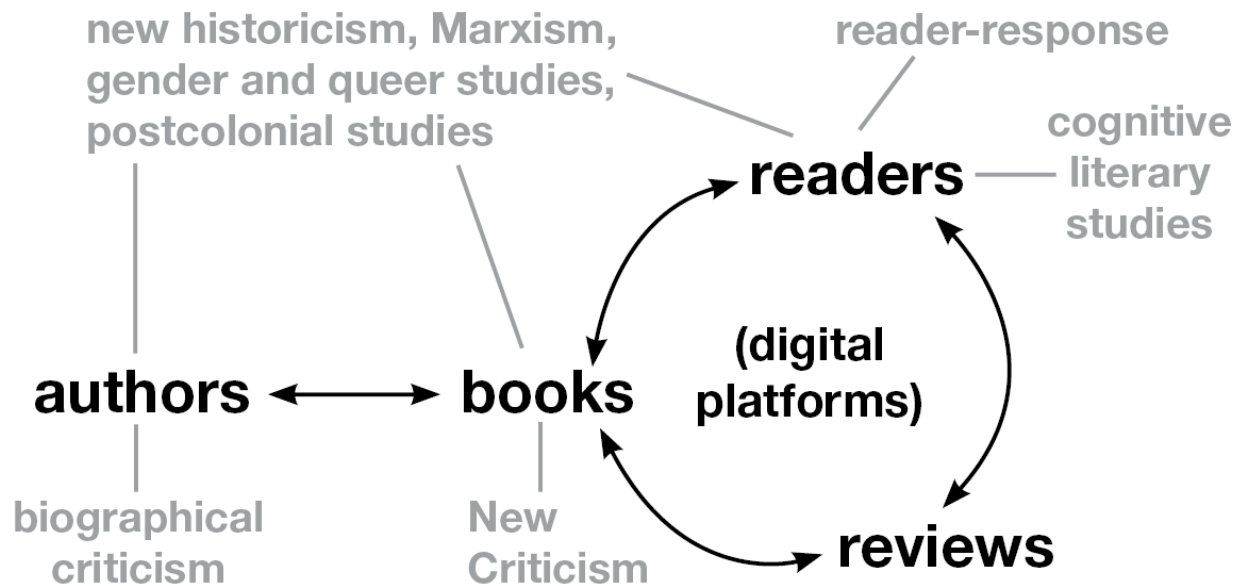


Figure 1.1: Model of social reading.

While Liu’s models have documents at the center, mine revolves around a cycle of interaction between books, readers, and reviews, mediated by the affordances and design of digital social reading platforms. In the following paragraphs, I will detail each aspect of the model and the literary critical fields that have examined them. As a whole, my model of social reading clarifies my dissertation project’s scope and main areas of analysis, accounting for how different schools of literary criticism have understood authors, books, and readers over the past century as well as

underlining the need to add online, amateur reviews to the current critical understanding of social reading.

When faced with the complex model of interacting parts in the social reading process, some literary critics have chosen to focus in detail on individual steps of the process. New Critics focus on the book, allowing formal elements like imagery, rhyme, and irony to become the primary variables in determining a book's meaning. By attending closely to the text, New Criticism brought needed attention to the text's details and structure. This is not to imply that the New Critics deny the importance of the author, the reader, or the book's social context. Cleanth Brooks clarifies the New Critical stance on whether the reader has a place in the evaluation of literary meaning: while he doesn't deny the importance of these external factors, he argues that "to put the meaning and valuation of a literary work at the mercy of any and every individual would reduce the study of literature to reader psychology and to the history of taste" (598). Brooks does not make explicit here his reasoning for placing literary meaning above psychological or historical meaning, but it remains that the New Critics view the reader as a distraction from the critical and scholarly work of finding meaning in literature. Similarly, W.K. Wimsatt and Monroe C. Beardsley set forth the New Critical stance on the author and the question of authorial intention, a red herring that only misleads readers away from the meaning of a literary work, since "[c]ritical inquiries are not settled by consulting the oracle" ("The Intentional Fallacy" 487). However, we should define New Critics not by such oppositions to reader-centered analysis, but by their positive declarations of what comprises literary meaning: "The poem belongs to the public. It is embodied in language, the peculiar possession of the public, and it is about the human being, an object of public knowledge" (Wimsatt and Beardsley, "The Intentional Fallacy" 470). In their privileging of the text above all else, the New Critics

intentionally set aside the influence of reader, author, and social context. While New Criticism has had a lasting impact on literary criticism and pedagogy, diverse schools of literary criticism since the mid-twentieth century start from the premise that a book's content and form alone cannot give you a full picture of its meaning.¹¹

Just as New Critics focus on the book, other schools of literary criticism have conducted a robust consideration of readers and their responses. Reader-response criticism both argued for the inclusion of the reader and conducted the first studies in this direction. At first, it looked much like New Criticism, where the reader emerged wholly from the text through close readings intended to decipher its intended audience. Wolfgang Iser's reader-response theory is explicit about this textual creation of readers' reactions to and experiences of books: "While an aesthetics of reception deals with real readers, whose reactions testify to certain historically conditioned experiences of literature, my own theory of aesthetic response focuses on how a piece of literature impacts on its implied readers and elicits a response" (Iser 57). Stanley Fish departs more sharply from New Criticism, arguing against the New Critical rejection of the reader known as the "affective fallacy" (Wimsatt and Beardsley, "The Affective Fallacy") in favor of "affective stylistics," which claims that "all poems (and novels and plays) were, in some sense, about their readers, and that therefore the experience of the reader, rather than the 'text itself,' was the proper object of analysis" (Fish 21). However, Fish still locates his method's legitimacy in its ability to reflect the full meanings of the text, with the reader merely a conduit for those meanings. As New Historicism gained prominence during the Cold War and after, scholars critiqued these previous reader-response studies as ahistorical, and attempted to incorporate into

¹¹ One prominent critique concerns New Criticism's restricted objects of study: "the trouble with close reading (in all of its incarnations, from the new criticism to deconstruction) is that it necessarily depends on an extremely small canon. This may have become an unconscious and invisible premiss by now, but it is an iron one nonetheless: you invest so much in individual texts *only* if you think that very few of them really matter. Otherwise, it doesn't make sense" (Moretti, "Conjectures on World Literature" 57).

their studies the historical context that influences readers' responses to a text (Machor, *Reading Fiction* xi).¹² However, historically-inflected reader-response studies still neglect individual readers in favor of tracking the broad historical trends that influence all readers in a given time period, and thus tend to generalize about such readers. Reacting to this continued absence of culturally-specific accounts of readers, several ethnographers have directly examined particular communities of readers in specific cultural and national contexts through surveys and observation as they read in book clubs or at home (Long, *Book Clubs*; Radway, *Reading the Romance*; Farr). These studies emphasize the importance of including a full account of individual readers' backgrounds, values, and personal interpretations in any attempt to understand a book's meaning.

But a reader's response isn't quite so simple to pin down. Cognitive literary studies further complicates our model of social reading by asserting that in order to understand how people read, we have to understand how their brains work while reading different kinds of texts in different conditions. Some scholars approach this problem by observing what happens to the brain in the moments when we empathize with or dislike a character (Zunshine). Others investigate how children learn to read, and study disorders like dyslexia as well as structural inequalities in language acquisition that cause children raised in poverty to have smaller vocabularies and narrower exposure to language (Wolf). Still others shift in scale from individuals to species, searching human brains for markers of social memory that have lasted for generations (László). Cognitive literary scholars are also interested in how reading online or on mobile devices changes how we read, from our attention span to how much we retain (Hayles, "How We Read"; Hayles, "Hyper and Deep Attention"). While it is often difficult to use the

¹² For more on historical reader-response, see Jauss's *Toward an Aesthetic of Reception*.

findings of cognitive studies to directly answer questions about literary mainstays like meaning, themes, and poetic language, these studies add nuance to our social reading model's account of the reader. Cognitive science adds another body of evidence to study, in addition to the reviews and diaries that serve as evidence in reception studies, helping literary scholars better understand the way readers make sense of books.

Biographical criticism brings yet another aspect of our social reading model into focus: the author. Biographical critics argue that the author's intentions and biography ought to be taken into account when determining the meaning of a book. Northrop Frye denounces this approach as mere "hero-worship" (24), and others offer detailed critiques of the futility of such a goal, calling it the "intentional fallacy" (Wimsatt and Beardsley, "The Intentional Fallacy"). But it is still difficult to argue that we should completely ignore the author, since biographical information can help readers better understand the text's context by illuminating ways in which the author's cultural and historical background influenced his writing, whether through religion, family circumstances, or political views. Rather, knowledge of an author's intention should not overshadow any other aspect of our model of social reading. Instead, literary critics should continue to add to and expand on each aspect of the process of social reading: the author is one factor, readers are another, and the more we know about them both, the more we can find out about a book's meaning.

Influencing each aspect of the social reading model discussed so far—book, reader, and author—is a web of social and historical contexts that literary critics must untangle in order to better understand social reading. The context of a single book is complicated enough: critics must account for the other books its readers have read and compared with it, the books that inspired and shaped the author, the books published in the same time period or geographic

region, the books categorized in the same genre, and so on. This expansive list of adjacencies and correspondences shows just how difficult a problem it is to construct an archive of related and relevant books for any literary study. Archival studies like the ones cited above on nineteenth-century American reading practices grapple with this question of the definition and limits of literary archives. Digital humanities scholars who focus on literature tend to question these studies when they examine only a limited collection of canonical texts and then extrapolate as if they had read the entire archive (Algee-Hewitt, Allison, et al.). Michael North's encyclopedic *Reading 1922* addresses this concern, painting a fuller picture of modernism that he claims is only possible through confronting an extensive and variegated archive of texts from 1922: as he says, "I read everything I could get my hands on" (vi). North's thorough archival work, despite its reliance on close rather than distant reading, is surprisingly in sync with digital humanities methods, as shown by a 2013 study confirming North's claims about the centrality of 1922 to the development of modernism that analyzed an even larger corpus, read programmatically rather than manually (So and Long). Digital humanities scholars take seriously the dream of the infinite archive touted by big data companies like Google, whose mission is "to organize the world's information and make it universally accessible and useful." What often goes unnoticed, however, is that their focus has shifted from the archival object itself to the object as represented by bytes of data within a database structure.¹³

Just as archival studies expands our model of social reading by adding context to books, a number of literary fields interrogate readers' social contexts. New historicism, Marxism, gender and queer studies, and postcolonial studies all argue that a book's meaning depends on its

¹³ Digital humanists methods like distant reading cause scholars to turn from individual texts to hundreds of texts, and in that way often lose sight of the intricacies of each text, instead representing texts as mere data to be mined and tabulated. My dissertation addresses this critique by performing a series of case studies which I will detail in the chapter breakdown below. For a fuller account of the various critiques of digital humanities methods, as well as how my project will address them, see the conclusion to this chapter.

placement within its social context. For example, in a series of distant reading studies that nevertheless show his Marxist roots, Franco Moretti uses publication and sales data to attempt to explain the success of certain books in a given time period to be a result of the relevant contemporary concerns that book addressed (*Graphs, Maps, Trees; Distant Reading; The Bourgeois*). The underlying assumption of these studies is that a book's sales are related to its ability to capture a spirit of the age, causing readers to buy it in droves.¹⁴ But many factors intervene to disrupt the relationship between a book's sales and the popularity of a book's message.¹⁵ There may be confounding variables, such as publicity campaigns, current events influencing what (and whether) people are reading, and readers' varying interpretations of the content, style, form, and themes of the books they read. Rather than using sales data as a proxy for these responses, my project uses reviews to gain direct insight into readers' responses to books, constructing a detailed picture of the reading communities that form around readers' reactions to books that can't be summed up in a dollar amount or a five-star rating.

This is of course only a rough sketch of the model of social reading as we might understand its history over the past century. My dissertation aims to complicate the model further by focusing on an aspect that is yet unexplored: user reviews, and what they teach us about the process of interpretation that every reader goes through and yet that is so difficult to model. In its focus on user reviews, however, my dissertation does not lose sight of the rest of the model, in

¹⁴ Moretti's argument has two steps: first, authors write books in accordance with a spirit of the age; second, readers speak with their wallets, only reading books that capture this spirit. An example of the first step is Moretti's claim that "[i]n the traumatic, fast-moving years between 1789 and 1815, human actions seem to have become indecipherable and threatening; to have—quite literally—lost their meaning. Restoring a 'sense of history' becomes one of the great symbolic tasks of the age: and a task uniquely suited for novelists, because it asks for enthralling stories (they must capture the explosive new rhythm of Modernity), but also well-organized ones (that rhythm must have a direction, and a shape)" (*Distant Reading* 20). Moretti articulates the second step of his argument in his strong view of readers' discernment in the books they read, citing, for example, detective fiction and the contemporary success of the clue as a plot device: "Readers discover that they like a certain device, and if a story doesn't seem to include it, they simply don't read it (and the story becomes extinct)" (*Graphs, Maps, Trees* 72).

¹⁵ For a detailed statement of this argument, see Radway (*Reading the Romance* 19).

which books, readers, authors, and social contexts both influence and are influenced by these reviews. Models allow researchers to make predictions, and then look for places where their results differ from what they expected. The working model of social reading articulated above serves as a starting point for my understanding of social reading, which I then update and complicate based on the results of my computational experiments. As detailed below, each chapter of my dissertation tests the above model of social reading by interrogating one aspect of the reading process, treating user reviews as the raw material for analysis. I ask a series of questions to test how well the model predicts the results we see in the user reviews, either confirming or contradicting our established ideas about genre and plot. In the next section, I elaborate on the specific methods I used, which each support my overall approach of testing and refining the literary model of social reading.

Methods

Machine reading tools and their applications to social reading

When you land on the webpage for any book on the social reading platform Goodreads, it offers much more information about the book, who has read it, and what they think of it than when you pull that book off a shelf in a library or bookstore, and even more information than when you locate it in a digital card catalog. Goodreads shows you all of the information you can find in a physical book (such as the title, author, and cover image) but it also offers an aggregated five-star rating, a list of genres the book belongs to, a list of similar books, and thousands of user reviews of that book. While this information is digestible for a reader interested in a single book, in order to understand the diverse ways this information shapes

readers and also draws from readers' own contributions to such digital platforms at the scale of 1.5 billion books (the number of books on Goodreads), we need the help of a computer.

Machine learning is a field of computer science that specializes in finding patterns and outliers in large datasets like those that inhere in the code and content of Goodreads, LibraryThing, Amazon, and other social reading platforms. For example, machine learning programs can classify news articles, detect faces in images, translate foreign languages, and forecast the weather. From these examples, it is clear that researchers can apply machine learning to many types of data, from text and images to meteorological data. Nevertheless, researchers rarely apply machine learning to literary fiction, perhaps because of fiction's difficult questions of meaning and interpretation, or its lack of business applications.¹⁶ The field's methods rest on a class of algorithms that allow computers to perform complex tasks repeatedly, improving their accuracy with each repetition. A typical machine learning project will start by constructing a model that researchers can then use to make predictions. This process of modeling, predicting, and testing the model is in line with my dissertation's aim to refine the literary model of social reading, as discussed in detail above. Digital humanists who study literature use machine learning to make sense of large amounts of textual data, from tens of thousands of eighteenth-century French encyclopedia articles to thousands of nineteenth-century British novels (Horton et al.; Jockers and Mimno). Fundamental to machine learning is the assumption that computers can find meaningful patterns, algorithmically, in large bodies of text.

My project uses machine learning to read book reviews, testing the model of social reading and how well it predicts answers to questions of genre and plot. As an example of how such a method works: instead of manually reading every review posted on a social reading

¹⁶ One notable exception is Jodie Archer's work using machine learning to delineate the characteristics of bestselling novels ("Reading the Bestseller").

website in order to develop a theory of genre, a researcher or team of researchers might teach a computer how to sort books into genres based not on the preset genre categories that social reading platforms apply via algorithm, but instead on the text of its reviews, allowing readers to determine a book's generic classification. However, teaching a computer is not like teaching a human. You don't tell it what concepts govern genre or recount the history of genre theory—instead, you give it a dataset with specific features (perhaps you mark each word's part of speech, or give each word a rating based on how positive, negative, or neutral it is) and it finds the procedure that works best on its own, with very little transparency. You often won't know which features it considered and with what weight, but you'll get an "accurate" result.

But how would a researcher evaluate the accuracy and quality of the study in the context of a machine learning project? If a researcher were to write a machine learning program that assigns a genre to a subset of books, the researcher would likely gauge the program's accuracy by comparing how well the program's assignments match the way researchers ordinarily classify those books within established literary histories of genre. Machine learning programs are designed to run repeatedly but with slight variations in order to find the algorithm that generates the most accurate results based on preset standards. Thus, as the program repeated its genre assignment process, it would become more accurate at classifying books that follow generic conventions closely, but it wouldn't be able to handle anything that breaks the mold (such as genres that change over time), and it certainly wouldn't be able to detect and suggest new genres.¹⁷ Such a program would still be useful, however, in quantifying how well certain books

¹⁷ This process describes supervised machine learning, the method most often applied to these types of classification problems. Unsupervised machine learning, which researchers usually apply to problems without correct answers such as automatic summarization and topic modeling, remains a possible method to attempt to generate new genres, though it comes with problems as well. Most troublingly, unsupervised machine learning often picks up on superficial differentiations such as time period (because of archaic word usage, for example) rather than more meaningful markers of genre.

fit into established literary genres, and how strictly readers on digital social reading platforms adhere to these established genres. The program might show that readers often mistake science fiction novels for fantasy novels, indicating that those two genres might be less distinct than others to readers in practice, or its results might lead us toward a new hybrid genre. Machine learning can show us patterns like this that we would otherwise miss, and also prompt new questions in moments when it seems to break down: for example, if a machine learning program has trouble classifying a book's genre, one might investigate a subset of the reviews using close reading and learn that different communities and cultures of readers are defining genre differently in their reviews, disrupting the program's ability to assemble a coherent definition. Finally, machine learning allows us to study a larger number of texts and a wider range of reader responses than we could by hand.

Machine learning can do a lot of powerful things, but studies that use machine learning methods often come to a halt when they reach the point of interpretation and understanding. For example, previous machine learning studies that focus on user reviews tend to examine easy-to-quantify aspects like keywords or positive and negative sentiment, but stop short when it comes to understanding the reader's assumptions, priorities, tone, and other factors that call for some degree of interpretation (e.g., Xiong). In part because of its inability to access the meaning of texts, machine learning is rarely applied to literary fiction, and little scholarship exists that addresses and explores alternative methods needed to apply machine learning methods to fiction as opposed to nonfiction. To this end, each project of my dissertation will attend to the unique demands of the literary tradition it interrogates and the fictional texts it examines by building on the long history of the study of social reading, detailed above, to develop appropriate machine learning methods.

The outcomes of my dissertation on social reading within digital environments thus do not take the form of theses about a set of individual texts and writers, but I argue that this limitation is also a strength of the project. In expanding the focus of literary analysis from one or several canonical works to hundreds, my project challenges and complements the methods and arguments of more traditional literary studies of particular genres, authors, and literary-historical periods by confronting them with texts that don't fit their findings and readers that don't agree with their interpretations. This doesn't mean, of course, that I'm abandoning what is productive about traditional literary analysis. I supplement machine reading with case studies and an attention to pre-digital archival materials as well.

Corpora

My results rely on a corpus of user reviews and other literary material from a variety of social reading platforms, including Goodreads, *The New York Times*, *The Los Angeles Review of Books*, and Wikipedia pages for books. Each chapter zeroes in on a subcorpus that pertains to its particular questions. The genre chapters use a corpus that includes metadata about each book review's classification within one or more genre categories, as well as the various genre-related shelves Goodreads users assign to the book. The plot chapter uses a corpus that includes plot summaries from sites including Goodreads, Wikipedia, CliffsNotes, SparkNotes, and various professional book review sites. To collect these reviews, I used the site's application program interface (API) when possible, which is the site's officially maintained protocol for developers to collect publicly available data from a website. When a site did not provide an API, I collected the reviews by scraping individual pages programmatically. I stored reviews in an ordered data structure for easy access via database queries. For each platform, I collected reviews as well as

relevant metadata such as their publication date, star rating, and associated user data. My corpus of reviews was limited to American, English-language versions of the sites, while still recognizing that readers are an international group and they read many books published outside of America that also influence their reviews.¹⁸ Without ignoring this context, my project focuses specifically on the American culture of social reading.

* * *

By grounding my dissertation in digital humanities methods, I enter a controversial debate over the status of these methods within the humanities as a whole. Some humanists reject digital humanities methods as fundamentally at odds with the values of the humanities, and critique the field's seeming immunity from, or complicity with, the crisis in the humanities. In response to digital humanities' singular success, countless articles critiquing digital humanities have sprung up in both popular and academic publications, and the 2013 Modern Language Association conference hosted a roundtable devoted to discussing "The Dark Side of Digital Humanities." Richard Grusin's write-up from the talk expresses a prominent line of critique: "its growth, support, and success can be traced, or is often explained, defended, or promoted, in terms of the very neoliberal values that have been seen to be the cause of the current crisis in, especially, public funding for higher education" (85). Grusin suggests that the methodologies of digital humanities projects enable and encourage the systems of temporary labor and

¹⁸ The reason for restricting the study to English-language reviews is the difficulty of conducting machine learning studies in multiple languages, which requires the daunting task of setting up equivalencies in word choice and usage in order to compare corpora in different languages. The reason for focusing on the United States rests in my own intellectual experience and interest in contributing to the field of American literary studies. While the headquarters of social reading websites are most often in America, and they typically have more American users than users from any other country, I recognize that increasingly, books, authors, and reading communities cross national and geographical boundaries as well as linguistic boundaries.

corporatization that are counter to humanist values and that caused the crisis in the humanities.¹⁹ While critics often frame the controversy as a competition between those with technical skills and those without, my dissertation shifts the terms of this debate by emphasizing the technical skills of the amateur reviewers on social reading websites, fluent in the language and conventions of that digital platform. While many are still excluded from these platforms, or choose not to use them, they remain a place where amateur readers can wield technological tools as a part of their own reading process, just as those tools influence them. In this two-way flow, social reading websites take their place as the newest of many tools readers have used to form communities around books for centuries. As I have shown in my history of social reading since the nineteenth century, I heed Grusin's call to consider "not only how new media technologies reshape or refashion what we mean by a humanities education in the twenty-first century but also how the humanities have always already been engaged with, indeed have coevolved with, technologies of mediation throughout their history" (89). Furthermore, my project contributes to Grusin's critique of the neoliberal values behind social reading websites' design and structure. In suggesting alternative ways to value and categorize books that attend to the content of those books rather than just their covers, my project is firmly a humanistic project, rather than the type of systematic, faceless machine reading that websites like Amazon practice on their users.

Another critique often mounted against digital humanities scholarship, particularly early on, was that it aspired to the ideal of objectivity, putting itself forward as a way of reaching

¹⁹ In a difficult-to-decipher moment, Grusin begins his article with an introductory disclaimer: "I am not now, nor have I ever been, a digital humanist," but then says "I have *programmed* talks and workshops designed to encourage faculty and graduate students to incorporate new digital modes of research and communication into their scholarship. But I have never done digital humanities work myself" (79, my emphasis). Assuming Grusin uses the word "programmed" intentionally, given the topic of his article, we can read this moment in two ways. In one reading, Grusin, despite disclaiming his involvement with the digital humanities, at the same time utilizes the cachet of scientific terms. Alternatively, Grusin is making fun of DH's use (or misuse) of scientific terms for acts that are not, in fact, truly scientific.

verifiable truth about a text.²⁰ Such scholarship was often tautological in its conclusions due to a reliance on the presumption that machines are objective—a presumption that critics have understandably found highly problematic. More effective digital humanities projects find alternative ways to convince their audience that do not rely on sweeping claims of objectivity, and instead find their power in the individual moments of critical insight they enable. My dissertation, despite using machine learning methods to analyze large bodies of texts, attempts to avoid oversimplifying or overpowering these texts. For example, rather than making general claims about readers based on simple statistical queries such as word counts or n-grams, I closely examine and compare the design decisions of a wide variety of social reading websites in order to gauge their concrete impacts on user reviews. My project, far from taking for granted the database structures that underlie these websites, interrogates the assumptions they make, from their binary assignment of genre to their simplified portrayals of the book’s plot in their metadata depictions of a book’s characters and themes.

My dissertation serves to fill two notable methodological gaps. First, digital humanities researchers have so rarely studied books alongside online user reviews of those books that there is no established set of best practices.²¹ Bringing together books and their reviews allows for a more nuanced understanding of reviews than tech companies have achieved by looking at online reviews alone, and a more comprehensive view of social reading and reception than literary studies has conducted previously without access to these reviews, which are a largely

²⁰ Contemporary digital humanities scholars are also guilty of making claims to objectivity. For examples of these types of claims from Moretti’s *Distant Reading*, see Shawna Ross’ review of Moretti’s work, “In Praise of Overstating the Case.”

²¹ The only study I have found that examines in detail online user reviews alongside the books they review is Ed Finn’s doctoral dissertation, “The Social Lives of Books.” However, while my study aims to use the same digital humanities method of machine learning to examine both books and reviews, Finn uses two different methods, complementing network analysis of communities of book reviewers with qualitative analysis of the books they review.

unexamined body of material ripe for study, but too numerous to read individually. The machine learning methods I use in this project allow me to make sense of many more responses to books than has been possible using previous analog methods of studying reader response. Second, my dissertation attempts to address a lack of methodology in computer science designed specifically for literary material. Computer scientists who use machine learning methods rarely study literary fiction because of its lack of business applications and difficult questions of voice, metaphor, and interpretation. My project addresses this gap by developing machine learning methods that I explicitly designed to allow a nuanced account of the literary complexity apparent in fiction reviews.

My project also helps to extend the history of social reading into the twenty-first century by accounting for digital social reading platforms' influence on how we read and understand books. These platforms impact the reading process throughout its course, from choosing which book to read to encouraging and discouraging certain reading and reviewing behaviors in their users, and yet literary scholars rarely discuss this influence. In adding online book reviews to the body of literary material available for analysis, as well as filling several methodological gaps in digital humanists' ability to study these reviews, my dissertation makes a substantial contribution to our understanding of social reading in the digital age.

Chapter 2: The Genres of Amateur and Professional Reviews

As the crisis in humanities funding highlights the field's declining prestige, humanists have rushed to clarify the benefits of literary training over untrained, casual reading. The comparison of these two types of reading, it seems, has only shown how little the methods of professional critics and amateur readers have in common. It is clear that these "two cultures" of reading continue to diverge—in literary theorist John Guillory's estimation, "Professional reading and lay reading have become so disconnected that it has become hard to see how they are both reading" (34). Nowhere is this more clear than in the field of genre fiction, in which the popularity of books like *Twilight* and *The Da Vinci Code* repudiates the general professional disapproval, or at best dismissal, of these bestsellers. Amateur reviewers are fully aware of the gap as well, and in their reviews it's common to see sentiments like this Goodreads user's toward the romance novel *Gone with the Wind*: "It would never win the Pulitzer Prize today, and I don't know that it would deserve it, but that does not mean it's not a great book" (Melissa).

When amateur and professional readers talk about genre, moreover, they're not talking about the same thing. For professionals, genre has a long literary history as a set of terms for organizing discourses that share certain traits, whether formal, stylistic, rhetorical, or aesthetic. For amateurs, genre encompasses both traditional genres and more informal tags and publishing categories, and most often serves to help readers articulate their identities to others and connect with larger communities. These definitions of genre are encoded into the platforms each group uses to publish reviews and interact with other readers. To see the stark differences between amateur and professional book reviewing platforms concretely, we can visually compare the number of books from each genre that these websites feature in order to determine which genres

receive the most (and least) attention on each platform. Below is a breakdown of the landscape of genre on four different book reviewing websites, two focused on amateur book reviewing (Amazon and Goodreads, in the top row) and two focused on professional book reviewing (*The Los Angeles Review of Books* and *Kirkus*, in the bottom row):

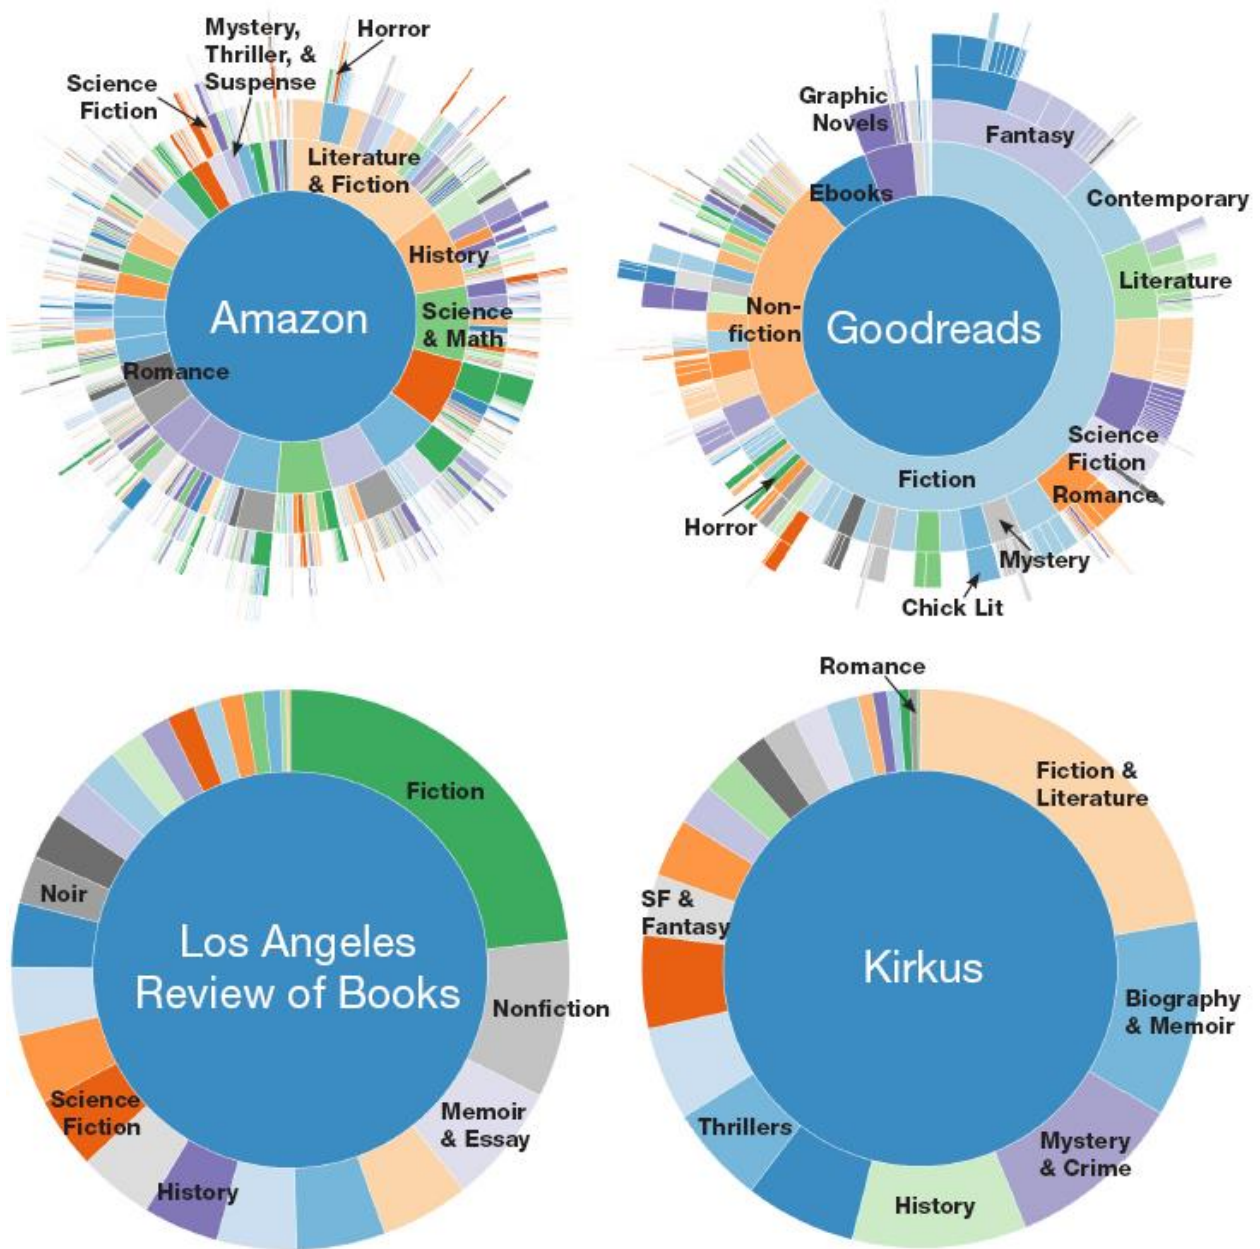


Figure 2.1: Navigation categories for book reviewing websites, both amateur (top row) and professional (bottom row). The size of each colored slice represents the total number of reviews tagged with that category. The platforms use different taxonomies: for example, some include subgenres within the “fiction” label, while some distinguish genre fiction from “literary fiction.” Data collected in January 2017.

At first glance, it's clear that amateur reviewing sites have a much more granular view of genre than professional review platforms. While the professional review sites in the bottom row display only one hierarchical level of genres as labels for their reviews, the amateur sites in the top row, with their series of outer rings, offer several layers of increasingly specific subgenres to categorize their reviews. Concretely, when you visit a professional review website, every review belongs to one of a clear list of genres; on an amateur review website, you can navigate more deeply into each genre to explore increasingly narrow subgenres. In this way, amateur review sites allow readers to discover niche and cult interests, read more books from those subgenres, and meet other users with the same interests.

Turning to the generic makeup of each site, Amazon's chart at the top left shows how many books the site offers from each literary category. The bulk of the books on the site fall under "Literature & Fiction," but it sells a wide variety of books, including selections of romance, mystery, science fiction, and horror novels. In contrast to Amazon's representation of the number of books on sale in each genre, the chart for Goodreads at the top right shows users' engagement with each genre, measured in terms of how often users classify any book as belonging to a given genre. From the Goodreads chart, we can see that some genres receive disproportionate attention compared to their presence in the Amazon marketplace. Goodreads users are far more interested in narrative fiction than nonfiction, and break fiction down into diverse subcategories. They discuss fantasy most of all, but give plenty of attention to other types of genre fiction. Turning to the bottom two charts, we can see that the same isn't true of professional review sites. For professional review sites, genre fiction's low status is apparent in its absence. *The Los Angeles Review of Books* and *Kirkus* most often review literary fiction, which they set apart as mutually exclusive from genre fiction through their category labels.

When they do discuss genre fiction, only mysteries and science fiction make the cut. *Kirkus* has a healthy mystery section, but *The Los Angeles Review of Books* only reviews noir, and both have a meager science fiction section, with *Kirkus* merging science fiction with fantasy. *Kirkus* has a mere 600 romance reviews—amounting to less than 0.5% of its total collection—and *The Los Angeles Review of Books* does not have a romance section. Neither professional review site has a section for horror novels. While amateur reviews discuss all types of fiction agnostic of genre or subgenre, professional reviews are much more selective in the books they will review, accepting some types of genre fiction (mystery and science fiction) but avoiding others (romance and horror). What is it about genre fiction that makes it so interesting to amateur reviewers, yet so polarizing for professional reviewers? And what might their reviews reveal about how these two groups define and evaluate these genres?

To answer these questions, I use a series of computational methods that compare the language of professional and amateur book reviews across four types of genre fiction: science fiction, romance, mystery, and horror. I identify the words that distinguish these reviewers' descriptions of the genres, and the topics they find important in evaluating each genre. For all four genres I examine, I find that professional reviewers follow a strict editorial code that limits their use of sentimental language, instead focusing on the purportedly objective topics of plot and character. Amateur reviewers, in contrast, have more personal and diverse ways of discussing and evaluating genre: they judge science fiction by its subgenres and awards to determine its quality, identify with the character of the detective to solve a mystery, and allow the affect of romance and horror novels to guide their reactions. We can see, then, how readers' location within either a professional or amateur set of genre expectations fundamentally alters their definition of genre, whether they consider it a structural feature of the text, a product of its

literary context, or an affective response. In short, reviewers' position within an amateur or professional reading context determines how they talk about books, and therefore how they define genre.

As I will show, these differences between professional and amateur genre definitions are a symptom of a much more fundamental change in the online reviewing landscape. Professional reviewers, when they discuss genre in terms of its structure, treat genre as a characteristic of books. Amateur reviewers, however, with their more personal view of genre as a way to shape identities and form communities, reveal a shift in genre that's occurred since reviewing has opened up to the public. Technology companies use genre not to categorize books, but to categorize consumers into interest groups that they can easily target and model with predictive algorithms. In their reviews, amateurs show this new way of defining genre, as a personal category of fandom that is then negotiated to determine whether or not it applies to a given book. By comparing professional and amateur book reviews, we can see genre's transition from literary category to consumer category firsthand.

Genre on Social Reading Platforms

Online book reviews appear in carefully designed spaces that shape the ways reviewers can express their opinions. Professional reviewers write for magazines and newspapers, publishing their reviews alongside contextual information about the book and its author. Amateur reviewers post their reviews on social reading websites, the largest of which, in the

American context, is Goodreads.²² Goodreads allows its users to rate books, review them with minimal restrictions on length or style, and respond to others' reviews. Readers on Goodreads can carry out conversations on discussion boards, join public book clubs, and even ask authors questions about their work directly. Despite the differences between professional and amateur book review websites, genre is consistently their most important organizing feature, as sites frequently prompt users to sort through reviews by genre or to read "similar" reviews from the same genre. Goodreads features genre directly on its homepage as a way for readers to find new book recommendations, and offers groups, clubs, and profile badges for fans of specific genres. Professional review websites like *Kirkus* typically offer no way to browse their reviews other than by genre.

Why has genre so overwhelmingly become the organizing structure for these book review websites? McGurl articulates a common explanation for the dominance of genre categories in twenty-first-century fiction, suggesting that genre "implies an audience ready to be pleased again and again within the terms of an implicit contract" ("Everything and Less" 460). Genre gives readers information about a book's themes, tropes, style, and other important attributes, and readers will often visit book review websites (both professional and amateur) in order to decide whether or not to read a book. While agreement over what each genre entails is useful to authors and readers, it is also useful to the companies that run book review websites. Like gender, age, and race, genre fandom—which often correlates with these identities—provides companies with valuable demographic data about their customers that they can use to better tailor their products and sell targeted advertisements. These online book review websites are not neutral platforms:

²² While the largest segment of Goodreads' user base is from the United States (40.77%), Goodreads is an international platform, with the next largest segments of users in India (7.73%) and the United Kingdom (6.43%) (Quantcast).

they use genre to place users into categories rather than as an objective descriptor of a book's content. When sites like Amazon and Goodreads feature increasingly specific subgenres, it is not to better categorize its product offerings, but to better categorize its customers. We will see the impact of these sites' focus clearly in the divide between how professional and amateur reviews discuss genre fiction, with professionals defining genre through the lens of a book's formal features, and amateurs defining genre in terms of their own preferences and identification with the genre.

To understand the differences between professional and amateur reading, then, we can see genre as a place where the differing goals of these two reading contexts become more explicit, in a mutual interplay between the financial gain of publishers and the community gain of users. Professional readers use genre to situate themselves in a clear, established discourse, as reflected in their websites, which organize reviews into non-overlapping genre categories and follow a standard formula for providing at least baseline information about a book's plot and characters. At the same time, websites tailored to non-professional readers attempt to address their priorities by allowing new types of reciprocal engagement with genre: users can increasingly reshape traditional genres and categorize books by their own personalized standards. Goodreads allows readers to create their own genres via its "shelving" system: readers can place books on self-named shelves that can range from traditional genre categories like "fantasy" and "romance" to non-traditional genres like "magic" or "weird" based on readers' own needs and experiences of books.²³ In this way, while professional platforms rely on the formal traits of the book to define its genre, amateur platforms leave much more room to allow their users to shape

²³ It's important to note that the shelving system does not just benefit Goodreads users. Through the shelving system, Goodreads gains crowd-sourced data on the genre of millions of books in its database as well as user data about specific readers' genre preferences. Normally you would have to pay people to classify books into genres manually, but Goodreads gets that information free from its users. The shelving system also provides them with information about their users' preferences for better-targeted advertising.

personal, user-defined genres to express their own reactions and identities. I will discuss these non-traditional genres and how they function on Goodreads in the next chapter; for now, my focus will be on four popular types of genre fiction—science fiction, romance, mystery, and horror novels—and how amateur and professional reviewers alternately interpret them.

Science Fiction and the Negotiation of Status

Veteran editors Martin and Dozois (*Old Mars*) assemble an entertaining array of SF stories recalling the heady days of the pulps while exploring provocative themes of alienation, morality, and discovery. (*Publishers Weekly*, “Old Venus”)

As lauded as Asimov is, his name bandied about whenever people bring up the “true” SF authors, and often when discussing strenuously scientific works of fiction, I found it rather surprising that Asimov was so liberal with his use of fantastically un-scientific technology: atomic blaster guns, personal force fields, nuclear power plants the size of a walnut, unexplained hyperspace FTL, mutants with psionic abilities...etc.). I don’t mind any of that, but most authors today get blasted by purists for the inclusion of any of that. Use any one of those in a story and the Hard SF police will declare that your book is fantasy, not SF. (Goodreads user Sisk)

From “the heady days of the pulps” to the wild success of internationally acclaimed writers like Isaac Asimov, science fiction’s critical fortunes have risen considerably over the past century, but the collection of genres gathered under the umbrella of science fiction still struggles to prove its legitimacy. In the first half of the twentieth century, science fiction appeared in mass-market magazines like *Amazing Stories* published on newly available, inexpensive wood pulp. These “pulp” came to define scifi through the types of stories they accepted, leading many critics to rely on definitions of scifi that tied it to its publication medium. As scifi gained recognition through the success of several prolific authors during the “Golden Age” of science fiction like Asimov and Arthur C. Clarke, scifi moved from pulps to hardcover volumes, and definitions shifted in focus from a work’s publication medium to the tropes it employed, most

notably robots (for Asimov) and aliens (for Clarke). The “New Wave” of science fiction in the 1960s and 1970s was an avant-garde movement against the overuse of tropes in science fiction, rejecting common science fiction clichés and declaring that “science fiction should turn its back on space, on interstellar travel, extra-terrestrial life forms, galactic wars and the overlap of these ideas that spreads across the margins of nine-tenths of magazine s-f” (Ballard 197). Theorists began to develop more sophisticated definitions, notably literary critic Darko Suvin’s characterization of science fiction as a text which contains a “novum.” Suvin defines the novum as a “point of difference” from our reality that is possible within the constraints of science yet nevertheless produces “estrangement” as the reader confronts an unfamiliar world (*Positions and Presuppositions in Science Fiction* 37). Through these definitions, science fiction has carved a place for itself as a cognitively complex and socially provocative collection of genres with institutional backing through awards like the Hugo and Nebula. But as we can see from the Goodreads user quoted above, readers still argue over the legitimacy of various representatives of science fiction, as if fighting to maintain its hard-won legitimacy and separation from still-maligned genres like fantasy.

How do literary historical accounts of science fiction compare to the terms in which contemporary amateur and professional reviewers each apprehend science fiction in the twenty-first century? To contrast the language of amateur and professional reviewers, I first gathered a collection of over one hundred thousand reviews of science fiction—as well as romance, mystery, and horror—from Goodreads as well as professional sites like *Kirkus* and *The Los Angeles Review of Books*.²⁴ By using book reviews of these genres instead of the text of books in

²⁴ I created a corpus that could compare professional and amateur reviews side by side. I distinguished the two types of review by their publication platform: professional reviews came from online newspapers and magazines, while amateur reviews came from Goodreads. In all, I considered 122,579 English-language reviews. The 98,405 professional reviews came from several sources: *Kirkus Review* (87,564), *The Los Angeles Review of Books* (2,609),

each genre, my results reflect readers' reactions to books rather than formal qualities of the book. An initial glimpse into the results using the relatively simple text analysis technique of frequency analysis reveals which words appear most disproportionately in science fiction reviews as opposed to reviews of other genres. Recognizable tropes like "space," "planet," "earth," "human," and "future" characterize science fiction reviews. The fact that these words are common, though, does not mean that they define the genre for readers. Textual frequency does not necessarily indicate textual significance: for example, in "The Purloined Letter," Poe withholds the contents of the letter from the narrative entirely. We might conclude that the letter is unimportant to the story's plot, but in fact, taking into account a more thorough context of veiled references and reactions to the letter, its outsized importance becomes clear. To see more explicitly how readers think about genres of science fiction, consider the phrase's context, or more concretely, the words that appear next to it—for example, "*classic* science fiction" or "*hard* scifi." In these examples, "classic" and "hard" serve as collocations, a term in linguistics for words that appear often beside another word. When you examine the collocations for science fiction book reviews, you see a sharp divide in how professional and amateur reviewers modify and make further distinctions within it:

USA Today Book Reviews (1,988), *The Washington Independent Review of Books* (1,229), *The Austin Chronicle Book Reviews* (944), and *Boston Review* (118). The 28,127 amateur reviews from Goodreads were reviews of the top 100 books (as voted by Goodreads users) from four different genres: horror (including "gothic"), mystery (including "noir" and "thriller"), romance, and science fiction.




Figure 2.2: Collocations appearing directly beside the terms “science fiction,” “scifi,” and “sf” in professional and amateur reviews.

The chart above shows the collocations for science fiction that each group favors most in their reviews. The distance of each word from the center indicates the likelihood of that word appearing beside “science fiction” in either a professional review (on the left side of the chart) or an amateur review (on the right side).²⁵ Some collocations appear in both types of reviews—these have two points, one on each side. The words on the top of the chart best represent the words professional reviewers use to modify science fiction, the words on the bottom represent amateur reviewers, and both groups use the words in the middle. While both professionals and amateurs use certain subgenre terms like “hard scifi,” “military scifi,” and “classic scifi,” amateur reviewers use these terms much more frequently. In addition, the collocations associated with amateur reviewers on the bottom of the chart much more clearly indicate further ways of

²⁵ The value along the y-axis is the collocation’s likelihood ratio calculated separately for each corpus of professional and amateur reviews.

subdividing science fiction than the professional collocations at the top of the chart, which are most often adjectives or publishing-related descriptors that give the book context like “debut,” “sequel,” “trilogy,” and “magazines.” This context most often comes in the form of anecdotes or background about the author and their previous work, the publisher, the genre, or its history of scholarly criticism. For example, a science fiction review in *The Los Angeles Review of Books* begins,

The Lifecycle of Software Objects by Ted Chiang is itself a lovely object, a slim volume with a linen-textured cloth cover and heavy paper dust jacket. The cover and inside illustrations are done in grays with touches of maroon, and the pages are of high-quality paper clearly printed. If one somehow wanted to quibble about getting value for money with such a short novel, surely the careful and aesthetically pleasing production would mute such complaints. But people who are familiar with Ted Chiang’s brilliant novellas would be expecting a novel-length work to be worth the price. And it is. (Gordon)

The professional reviewer, Joan Gordon, an editor of *Science Fiction Studies*, invites her audience into a cohort of informed readers who know Ted Chiang’s work well, and gives them a glimpse into the physical construction of the work as a proxy for its value. Gordon’s cultural and publication context—exaggerated in this case because of the focus on materiality of Chiang’s book itself—is representative of the professional reviewer’s tendency to focus on the context surrounding a science fiction book. Even though Gordon seems to have had a personal reaction to the book as “aesthetically pleasing,” she couches this reaction in the book’s formal qualities.

In contrast, amateur reviews feature the reviewer’s personal reaction front and center. Reviewing the same book, an amateur reviewer begins, “This novella made me want to hunt down a Nintendo console and give it a long hug. The title may sound like some dry technical manual, but the story is warm, human, touching and funny. It’s the best type of SF story: one that makes you think, makes you smile, and leaves you with a glimpse of deeper understanding of human nature” (Roy). Where the professional review described the book’s material body, the

amateur review immediately gives personal insight into the reviewer's own bodily reaction to the book: it made him want to "give it a long hug," it was "touching and funny," it made him "think" and "smile." In Roy's review, we see a characterization of science fiction as a genre that prompts certain types of reactions in readers, while Gordon's review focuses on the formal qualities of the book that might have inspired those reactions.

In addition to the distinctive intimacy of the amateur review, we can see by Roy's delineation of the proper "type of SF story" for Chiang's book that amateurs are highly concerned with subdividing science fiction. Continually, when reading amateur reviews of science fiction, you will see debates over genre and subgenre membership, whether enforced by readers or by some entity external to the review like the "Hard SF police" in the Goodreads quote opening this section. Returning to the collocates chart, amateur reviewers frequently modify "science fiction" with the subdividing terms "pure scifi," "soft scifi," "modern scifi," and "old scifi," as well as a series of words setting apart quality science fiction from the rest ("greatest," "great," "finest," and "masterpiece"). The addition of these evaluative terms suggests that the subdivision of science fiction might help readers judge a book's value or articulate it to their audience. It is surprising that amateur reviewers use these evaluative terms more often than the professionals whose job it is to judge a book's value. Amateur reviewers are clearly not willing to cede that authority to professionals figured as the "Hard SF police," and instead make genre distinctions that rely not on the formal features important to professional critics, but on their own personal judgments of a book's quality within the genre or subgenre to which it belongs.

Amateur reviewers' use of collocates to subdivide science fiction suggest that they are highly concerned about genre and subgenre—but are they more concerned than amateur reviewers of other genres? To find out, we need to expand the scope of our analysis of science

fiction from words to topics. A topic is an area of literary interest that reviewers use a wide range of words to talk about, all of which center around the same type of literary concern, such as plot summary, discussion of characters, or evaluation of the book's overall quality. The typical digital humanities approach when discussing literary topics is "topic modeling," a machine learning method that algorithmically identifies groups of words that tend to appear together in a collection of text documents. For example, the Stanford Literary Lab uses topic modeling to track the "thematic focus" of literary paragraphs, determining that the paragraph is the most natural scale of analysis to surface a text's larger themes (Algee-Hewitt, Heuser, et al.). Topic modeling is an "unsupervised method," meaning that it produces results that are often unexpected and difficult to interpret. For the Lab's study, topic modeling's unpredictability was no hurdle, since the content of the topics was less important than the results of their comparisons across variously sized segments of the text. In this case, I was interested in existing literary topics and genres, so I needed a method that would use the terms of the ongoing literary conversation rather than inventing its own.²⁶

To identify the most important topics in the collection of book reviews, a colleague and I read and annotated dozens of randomly selected reviews and marked, sentence by sentence, the topics these reviews discussed. Manual annotation, or "coding," is a well-documented practice in the social sciences in which researchers assign "a word or short phrase that symbolically assigns a summative, salient, essence-capturing, and/or evocative attribute for a portion of language-based or visual data" (Saldana 3). Our use, of course, was less rigorous than most social

²⁶ While it is often thought that supervised methods cannot teach us anything new or unexpected, they certainly can, like when your experiment proves precisely the opposite of what you expected. In this case, I chose to use a supervised method because I was testing a hypothesis I developed from earlier experiments and through close readings of the reviews, so I had a body of prior work and knowledge in which I wanted to ground my results. The strength of supervised methods is that they allow digital humanities projects to interact with the disciplinary conversation of a field without having to invent an entirely new vocabulary or set of definitions for concepts the field has already mapped out.

scientific coding applications, but we did go through a multi-step process to settle on our final categories, as well as ensuring that we had at least 90% agreement in the topics we marked.²⁷ Book reviews discuss a wide range of topics, but some topics appeared in a majority of the reviews. We settled on a bank of nine categories that best encapsulated the most common topics for each sentence in a review: plot, character, style, sentiment, identification, genre, evaluation, context, and a category called “other” for sentences that did not fit into any of the categories (see Table 2.1 below). While some of the topic judgments were easy to make, such as when the reviewer explicitly mentioned the topic in the sentence, often we had to use judgment to determine whether the reviewer was, for example, describing a character or their identification with a character. When not explicitly mentioned, we used the following rough definitions for each topic:

²⁷ Other digital humanities projects also use manual coding, for example James English’s project tracking the time period a book is set in: “My team and I have had to chase down information about the novels by looking at reviews, blurbs, sample pages, and readers’ plot summaries and, in some cases, by locating a physical book and actually reading it” (English 407). Interestingly, this is also the method that technology companies use to determine which product categories their products belong to—just on a much smaller scale. For example, Netflix uses hundreds of “raters” to watch and assign genre tags to each of its movies and TV shows (Nosowitz), and Rotten Tomatoes uses “curators” to classify movie reviews and resolve disagreements using panels of multiple curators (Barnes).

Topic	Topic Guidelines and Examples
Plot	Sentences that commented on the book’s plot overall or gave plot summary, e.g. “Morrie dies, and Mitch attends the funeral.”
Character	Discussion of individual characters as well as the characterization of the book as a whole, e.g. “He is married with no kids.”
Style	Discussion of writing style, pace, and the aesthetic quality of the prose, e.g. “The writing is wonderful, the story is slow-moving but engaging.”
Sentiment	References to a book’s emotional impact on the reviewer, e.g. “I’m still disturbed by this story and if you’re squeamish don’t read it.”
Identification	When the reader explicitly compared himself to a character or discussed his ability to empathize with a character or the book’s author, e.g. “Even getting some background on their families didn’t engender any sympathy in me for them.”
Genre	Discussions of the book’s genre or subgenre generally, or specifically how well the book fit into that genre, e.g. “It’s a gothic classic and I think most people are familiar with the story.”
Evaluation	Opinions on the book’s overall quality, statements of whether the reviewer would recommend the book, and discussion of star ratings, e.g. “Omg, this was literally one of the best books i’ve ever read.”
Context	Includes tangential topics, such as discussion of the book’s sequels, the movie version, book clubs that recommended it, or awards it won, e.g. “I can’t wait for the movie to come out next year.”

Table 2.1: Topic guidelines and examples.

The “other” topic captured anything that didn’t fit into these categories. Surprisingly, only 24% of the total text of the reviews fell into the “other” topic, indicating that the topics above make up a majority of the surprisingly standardized topic vocabulary of amateur and professional book reviews. In fact, many of these topics are the standard fare of book reviewing guides and lesson plans. The chart below shows, for the top five Google search results for “how to write a book review,” whether or not the guide recommended discussing each topic:

	Writing-World.com (Asenjo)	Purdue OWL: Writing a Book Review (Brizee)	How to Write a Book Review (wikiHow, “How to Write a Book Review”)	Book Reviews Handout (UNC Writing Center)	Scholastic Teacher’s Activity Guide (Philbrick)
Plot					
Character					
Style					
Sentiment					
Identification					
Genre					
Evaluation					
Context					

Table 2.2: Comparison of topics addressed in book review writing guides.

For example, the guide “How to Write a Book Review” from Writing-World.com asks readers to ponder, “What is the author’s style? Formal? Informal? Suitable for the intended audience?” (Asenjo), and the Purdue Online Writing Lab guide to “Writing a Book Review” asks you to consider, “Who are the principal characters? How do they affect the story? Do you empathize with them?” (Brizee). While most of the guides addressed a majority of the topics, the UNC Writing Center guide focused on nonfiction, so it avoided topics suited explicitly to fiction reviewing such as character and style. In general, the guides were less likely to recommend discussing the most subjective topics of sentiment and identification in book reviews. Already, we can see how examining trends in the types of topics reviews use might reveal their assumptions about what belongs in a fiction or nonfiction review, or in a review of a given genre.

We could not categorize over a hundred thousand reviews manually, so the next step was to train a predictive model to recognize these nine topics and classify each sentence in the rest of

the reviews.²⁸ Machine learning is a common method in computer science for replicating human classifications using a series of algorithmic rules and features, in this case the most and least distinctive words characterizing each topic. The chart below shows some of the words the model used as a guide as it categorized review sentences:

Topic	Most Predictive Words
Plot	plot, about, climax, spoiler, after, end, life, goes
Character	characters, mother, heroine, hero, relationships, parents, her, him, his
Style	style, prose, beautifully, writing, narrative, description, poetic, wordy, lyrical
Sentiment	felt, feel, emotional, scared, creepy, pleasure, haunting, me
Identification	connected, invested, fan, related, sympathy, I
Genre	genre, classic, sci-fi, romance, horror, thriller, mystery, fantasy
Evaluation	stars, best, worst, recommend, awesome, loved, favorite, awful, it
Context	author, audiobook, kindle, movie, film, series, sequel

Table 2.3: Most predictive words for each topic.

The most predictive words for each topic reflect the words reviewers most often use to discuss the topic, with words signaling the topic itself like “plot” and “genre” leading the pack, while the other predictive words seem to pertain to those types of discussions, like “favorite” in readers’ evaluation of the book’s quality. Using these words as a guide, the model predicted the topic of each sentence in the full corpus of reviews. By visualizing these predictions, we can see an overview of the amount amateur and professional reviewers talk about each topic, and how widely those amounts differ within a given genre. For science fiction, then, we can see that amateur and professional reviewers use topics in noticeably different proportions:

²⁸ The model used a machine learning algorithm called a support vector machine, which plots each review as a point in multidimensional space and then draws lines to separate them into classes based on the words they use. The accuracy of the model was 54%, a significant improvement over random chance (11%) or a model that simply predicted the most common category, “other” (18%). The model performed better when reviews explicitly mentioned terms related to the topics such as the word “style” indicating the “style” topic, and less well when none of the most characteristic words of that topic appeared. Mistakes in the model were most often ambiguous cases, for example the model choosing “character” when a review mentions both plot and character.

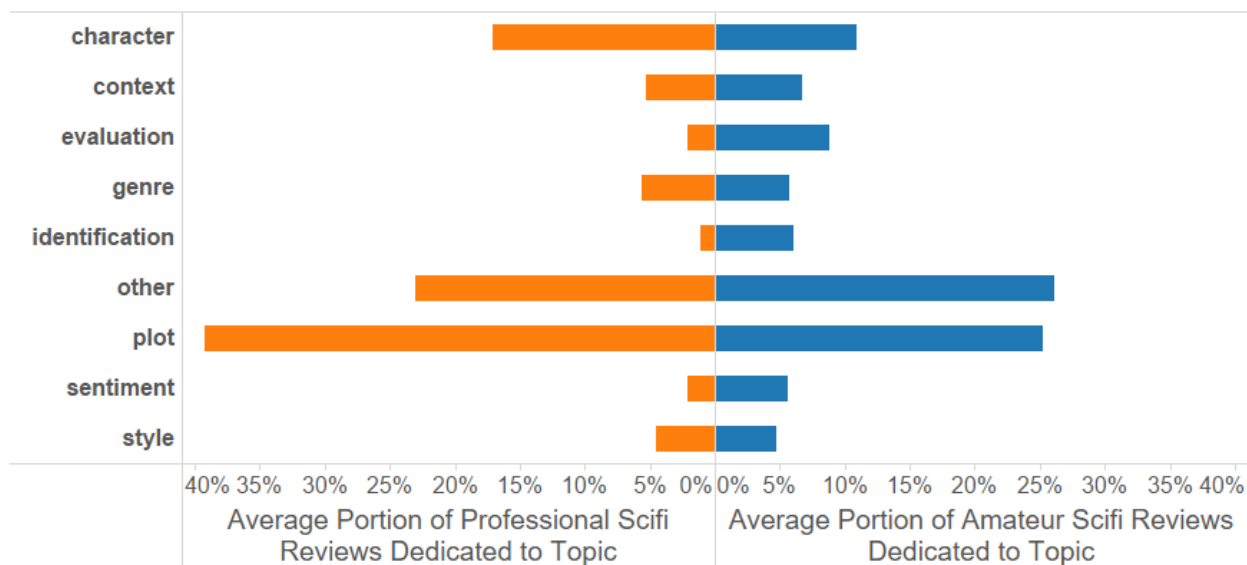


Figure 2.3: Average portion of professional and amateur science fiction reviews dedicated to each topic.

Professional reviews are more likely to discuss a science fiction book’s plot and character, while amateur reviews favor the more personal topics of evaluation, identification, and sentiment when responding to science fiction. These results might suggest that professional reviewers have a formal definition of sci-fi based on plot progression and characterization, contrasted with amateur reviewers’ reader-centric definition based on their personal reaction. But while professional and amateur sci-fi reviews discuss genre in equal amounts, when we compare amateur sci-fi reviews to amateur reviews of other genres, they do seem more concerned about genre:²⁹




Figure 2.4: Percent of amateur reviews discussing genre for four types of genre fiction.

²⁹ Despite the small values in these results, the large sample size of my corpus resulted in an ANOVA test showing that the difference between the genres is statistically significant ($p < 0.05$).

Amateur reviewers indeed discuss genre more in science fiction reviews than in any of the other genres I examined, reinforcing the reading that genre—and, as we saw, subgenre—is an essential part of how amateurs review scifi.

Based on amateurs' greater use of subgenre divisions, one could say that amateur reviewers have a collocation-based understanding of genre: rather than viewing scifi as a single genre unified under a single umbrella, they think of it in terms of adjective pairings that increasingly fracture it into more and more specific subsets. We might be tempted to conclude, seeing that amateurs use more diverse subgenre categories, that amateur reviews offer a more diverse range of perspectives than professional reviewers on what types of themes, settings, and tropes can constitute a given genre. But more a finely-divided field of genres and subgenres does not necessarily mean genre has grown more diverse—only that it has become better organized. In fact, the very proliferation of these genre divisions makes it increasingly difficult to see books outside of the existing categorical structure, allowing established genres to become even more entrenched.³⁰ When Netflix recommends increasingly specific subgenres like “irreverent dark TV comedies,” it reassures us of the accuracy of genre-based recommendation so that we are less and less likely to venture outside of the genres we know, content that Netflix's categories encompass every type of media we might want to see. By interrogating not just genre terms like “scifi,” but also their context of paired subgenre terms, we can see how amateur reviewers increasingly subdivide the genres of science fiction until they have become even more restricted in their definitions than professionals. When we consider the context of social reading websites using these micro-targeted subgenre categories to market to increasingly specific customer

³⁰ Incidentally, these narrower categories further Goodreads' interests by providing more specific categories to better define and differentiate their users for the purposes of market research and targeted advertisements. As we will continue to see, the platform a reviewer uses has a strong influence on that review's depiction of genre.

demographics, we can see that the subdivision of genre does not just restrict readers' definitions of genre, but also the range of literary material that these sites will recommend to them.

Digital humanities scholar Ed Finn has shown that a wider range of co-occurrence than we have used so far can help to illuminate book reviews in particular. His work tracks proper nouns that appear in the same paragraph to form a network of links between cultural objects of discussion in reviews ("Becoming Yourself" 4). Using these expanded collocations, Finn shows that readers make sense of difficult David Foster Wallace books through comparison with other novels and authors. What if we, too, examined a wider range of a word's context? We can expand our view from words that appear directly beside "science fiction" to those within 10 words of the term. Like the earlier collocation chart, the chart below shows the collocations for science fiction in professional and amateur reviews, this time appearing within a range of 10 words:




Figure 2.5: Collocations appearing within 10 words of the terms “science fiction,” “scifi,” and “sf” in professional and amateur reviews.

Professional reviewers, in the words on the left side of the chart, again tend to use the language of publishing conventions to give context to the books they review: for example, “debut,” “ya” (young adult), “anthologies,” “sequel,” and “magazines,” reinforcing a conception of genre rooted in the book and its context within systems of authorship, publishing, and advertising. Amateurs, however, consistently use the language of awards and value to lend science fiction literary legitimacy. While both professionals and amateurs use the term “award” and refer to the “Hugo” and “Nebula” awards for science fiction, amateurs are much more likely to use those terms. They mention “Locus,” a magazine that offers awards to science fiction books, as well as “nominee” and “winner.” They are preoccupied with “status,” “quality,” and what a given book “deserves,” and show an explicit concern over how science fiction books are “categorized” and “regarded.” When amateur reviewers mention awards, it’s typically in the context of asserting a book’s quality against any claims to the contrary:

This is the book you may have heard just got onto the Hugo ballot, following the sturm und drang and reshuffling of this year’s Hugo drama. I think it’s a worthy entry on the ballot; while it may not be the best SF novel I’ve read this year, it’s certainly got everything a Hugo-winner should have: an epic scope, an intriguing set-up, and a lot of speculative science.” (Goodreads user David)

Amateur reviewers like David disproportionately use the language of labels and awards that lend science fiction legitimacy as a label that can they can neatly categorize into genres and subgenres, and that has its own institutional structure of recognition. Interestingly, one might expect the professional reviewers who confer these awards to talk about them more often, but in fact amateurs do, suggesting that amateurs use the language of awards strategically in order to confer legitimacy on a genre they know to be high-quality literature, with the award providing confirmation. Amateur scifi reviewers seem to assert that they can evaluate scifi better than professionals, and they commend the professional award committees when they recognize what

amateurs have seen all along. But when they disagree, amateur reviewers are not shy to reject professional judgments:

I am at a loss to understand what made this book such a “groundbreaking masterpiece” it won both the Hugo & Nebula Awards. Let alone having it compared to LOTR. WTF?! It’s just another proof to me that what “literary experts” deem of quality has nothing to do with reader satisfaction or good storytelling, but with the position of stars in the sky and the quality of those experts’ breakfast the day of the vote. (Goodreads user Sicoe)

In this user’s critique of “literary experts,” she reasserts that scifi is best adjudicated by its fans, not by cloistered professionals. The user’s definition ultimately rests in the judgment of amateur readers—“reader satisfaction”—rather than the more formal qualities that tend to occupy professional reviewers.

Looking closely at the non-traditional categories amateur reviewers create for science fiction books on Goodreads further emphasizes the perceived divide between professional and amateur judgments of science fiction’s quality. In my examination of the personal shelves that Goodreads users created for science fiction novels on Goodreads, a clear division emerged between books that won awards (shelves like “award winning,” “Hugo,” and “Nebula”) and books that prompted an emotional reaction from readers (shelves like “tear-jerker,” “funny,” and “scary”)—a characteristic often associated with pulps, young adult novels, and zombie fiction, typically the lowest-valued constituents of science fiction. The chart below shows the proportion of the total shelves readers placed a book on that had to do with either awards a book won (the blue lines), or sentimental reactions a book elicited (the orange lines), for a selection of science fiction novels that readers placed on shelves falling into at least one of those two shelf types. The lines track the proportions of each of the two shelf types over time for each book:³¹

³¹ Again, despite the small values in these results, the large sample size of my corpus resulted in an ANOVA test showing that the difference between the genres is statistically significant ($p < 0.05$).

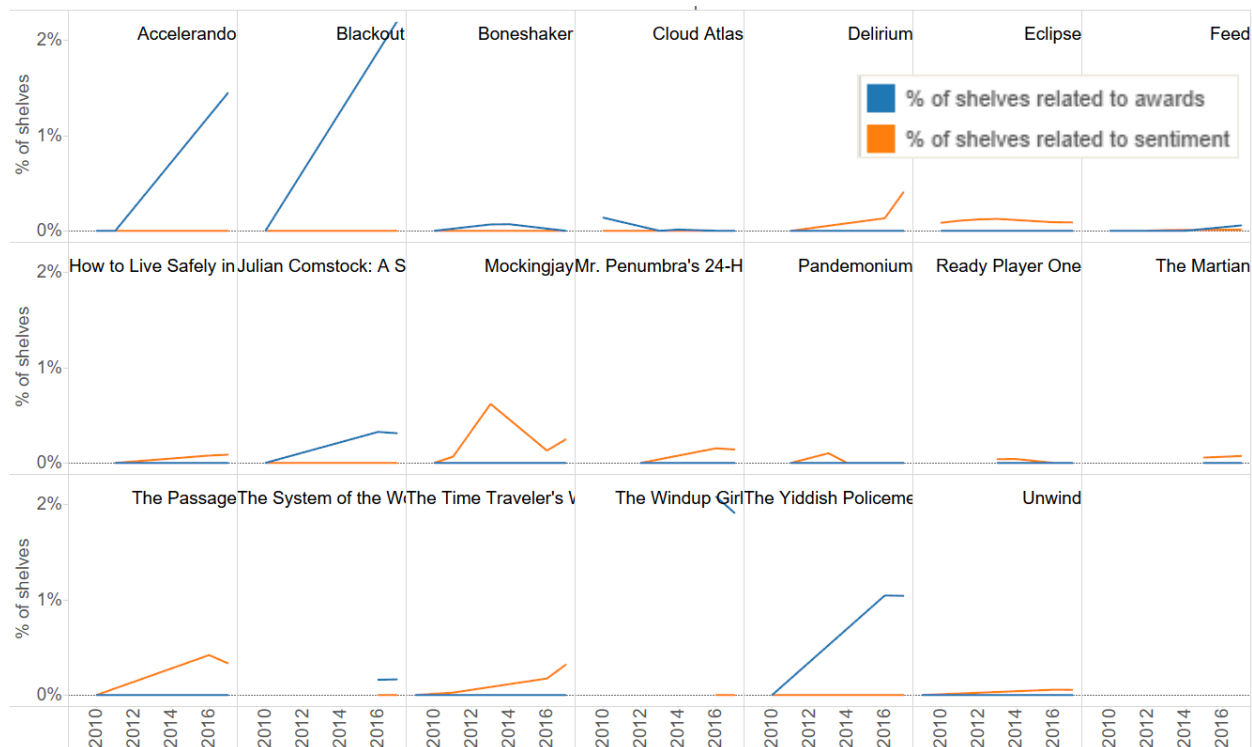


Figure 2.6: For 20 popular science fiction books, the proportion of shelves Goodreads users place them on that are related to awards won (in blue) versus sentiments elicited (in orange), plotted over time.

Not one of the science fiction books in the chart above, nor in any of over 70 popular science fiction books I surveyed, was placed by more than 1% of readers on both types of shelves, underlining the seeming lack of compatibility between books that users remember for the awards they win and books that users remember because they elicit emotion. Particularly striking in the chart above is *Mockingjay*, a popular young adult dystopian science fiction novel that Goodreads users most often categorize based on their emotional response—most commonly placing it on the shelf “made me cry”—and not awards, despite the novel having won several prominent awards. Here, it seems that the popular treatment of the young adult novel as emotional and low-quality supersedes any literary merit it might have in the eyes of Goodreads users. This tension has informed the reception of literary fiction and canonical novels as well. Toni Morrison described the struggle she faced in having her work accepted as both award-winning and emotionally salient: “I would like my work to do two things: be as demanding and sophisticated as I want it

to be, and at the same time be accessible in a sort of emotional way to lots of people, just like jazz” (qtd. in Leypoldt 380). Literary critic Günter Leypoldt explains the unique set of circumstances that made this goal possible for Morrison, including the induction of the literary avant-garde into traditional institutions and the “Oprah effect” (Leypoldt 381). But for science fiction, such a balance doesn’t seem to be possible. Amateur reviewers on Goodreads have a conception of science fiction that pits award-winning scifi against sentimental scifi, two categories they treat as mutually exclusive. On Goodreads, where readers are free to categorize books based on criteria more varied than just “science fiction,” we can see that this increased freedom sometimes carries with it equally strict constraints on the shelves books can occupy. If we consider, then, that amateur review platforms push readers to see genre as a consumer category rather than a literary category, the lack of overlap between sentimental scifi and “quality” scifi can constrain the way scifi fans categorize not just their books but also themselves. The particularly status-aware genre of science fiction highlights how genre becomes a stratifying trait for amateur and professional reviewers eager to establish their credentials as true scifi fans. While this debate over status is certainly not new, online platforms allow amateur reviewers a newfound ability to challenge the assumptions and judgments of literary professionals, reaching a wide audience in the process.

Romance and the Aesthetic Response

Though this novel occasionally resorts to some of the clichéd pitfalls that readers have come to expect from the supernatural romance genre (woman in danger, a vampire who only feeds on “bad guys”), readers who appreciate a romance tale flush with emotion, as well as some climactic action scenes, will find a lot to like. (*Kirkus*, “Ebon City”)

Lust or love at first sight can either flourish or flail in romance books but here it totally flourished. You felt their instant connection and their continuous draw to one another. I loved these two together, but as we all know in the romance genre things don't always go according to plan causing some angst. (Goodreads user Lauren)

Romance enjoys a rich history as a prose genre boasting influential chivalric, heroic, and pastoral tales, but by the eighteenth century, the rising success of the novel began to overshadow the romance: as a contemporary critic notes, “the Modern novel sprung up out of [the romance’s] ruins” (Reeve 8). Literary theorist Richard Chase distinguishes the romance from the novel by defining the former as the work with

an assumed freedom from the ordinary novelistic requirements of verisimilitude, development, and continuity; a tendency towards melodrama and idyl; a more or less formal abstractness and, on the other hand, a tendency to plunge into the underside of consciousness; a willingness to abandon moral questions or to ignore the spectacle of man in society, or to consider these things only indirectly or abstractly. (ix)

These formal and stylistic traits of the romance, for Chase, contribute to the critical distinction between novels, which reflect objective reality, and the romance, which is subjective and free to stray from reality. The romance novels of the twentieth century share little in common with the works Chase describes, which modern readers began to consider overwrought, inaccessible, and long-winded: today’s romance novels are instead page-turning, mass-produced tales with relatable female characters. Following the tradition of serialized mystery and detective fiction, in the 1950s publishers began to produce material that would appeal to women more than the gritty crime fiction dominating the literary market at the time, starting with gothic romance novels like *Mistress of Mellyn*, a fast bestseller. By the 1970s, romance novels were a significant part of the publishing business, and romance publishers like Harlequin were selling hundreds of millions of romances per year, not just in bookstores but also supermarkets and drugstores (Radway, *Reading the Romance* 40–41).

With its newfound popularity came a new wave of criticisms of the genre. Critics like Ann Douglas decried the popular romance as damaging to women, blaming a patriarchal culture for pushing women to “co-sponsor male fantasies about themselves” (28). While such critiques rely on a formal conception of the romance genre, like Chase’s, in asserting that the genre’s effect on readers is measurable as the sum of its formal features apparent on the page, reader response critics of the 1980s offered a rebuttal to these arguments rooted in surveys and interviews with readers, which attempted to gauge the book’s impact on readers. Rather than a formal definition of the genre, reader response critic Janice Radway asserts that romance is “never defined solely by its constitutive set of functions, but...by a set of characters whose personalities and behaviors can be ‘coded’ or summarized through the course of the reading process” (*Reading the Romance* 120). By looking at readers’ reactions, Radway challenges the notion that readers accept the ideology of what they read unquestioningly. Instead, Radway demonstrates that romance readers are far from passive recipients of romance stereotypes. The several dozen Midwestern bookstore patrons she interviews value above all a strong, intelligent heroine over one who is feminine, and a caring, tender hero, counter to the common trope of the gruff, inscrutable hero and the passive heroine (*Reading the Romance* 123–24). Radway’s interviews show that the romance genre allows these women freedom, not only through the act of reading, which lets them escape the demands of everyday life, but through plots flexible enough to allow responses that both confirm and reject female disempowerment. Radway’s key shift is to define romance not by its content, which often seems to reinforce gendered stereotypes, but by how readers receive it. Indeed, Radway’s interviews repeatedly show women questioning stereotypes and inferring progressive lessons from even the most traditional romance novel.

Do romance reviewers view the genre formally, or do they follow Radway's response-based definition? Looking at the collocations that reviewers use to modify the romance genre, we see that professional reviews pair "romance" with subgenre terms that relate to publication context and audience demographics such as "Harlequin" (a prominent romance publishing company), "pulp," "debut," and "adult." These words signal the professional reviewer's role as a literary authority, providing background and context to their readers. In a Los Angeles Times review of *Middle Age: A Romance*, the reviewer begins, "Author of 29 novels, 19 short story collections and some 36 other books, Joyce Carol Oates is not merely a writer; she is an issue. It is impossible to review a new book by her without wondering if she publishes too much, or without wondering if it's right to wonder whether she publishes too much. And yet one can hardly, at this point, approach each new book as if it were an isolated event" (Siegel). Even as this review recounts a personal debate over what to include in the review, the professional voice remains impersonal ("one can hardly") and stays focused on the book, its publishing context, and its author. As part of the professional reviewer's didactic role, the focus on plot and character we have seen so far makes sense, as their reviews provide the information necessary for readers to fully understand the book and its context under the assumption that the most essential information about the romance genre resides in the formal aspects of the book.

Amateur reviewers' collocations surrounding "romance" are instead topical, giving information about the romance's setting, themes, and style such as "paranormal," "historical," "college," "cheesy," "sappy," "cute," "clean," and "erotic." For example, an amateur review of *Middle Age* begins, "I like other novels by her far better but she will always receive at least 3 stars because she is a beautiful writer, even if she wrote about a goldfish swimming round and round it would still be somewhat fascinating. It wasn't my favorite Oates novel. I think you

always learn something about life when you pick of her novels as her characters are always very human (flawed, beautiful, terrible) and you swear you know someone like them” (Dandeneau). While both the professional and amateur reviews of *Middle Age* acknowledged Oates’ prolific literary output, the amateur review focused not on Oates herself but on the reviewer’s own reaction to the novel and evaluation of the book in relation to the other novels. Again, we see that amateur reviewers use more personal language compared to the more factual language of plot and character that professional reviewers use, revealing that for amateur reviewers, the key to the romance genre lies in their personal reaction to the book, not the book itself.

When we examine the topics romance reviews discuss, these trends in the emphases of professional and amateur reviews hold true:




Figure 2.7: Average portion of professional and amateur romance reviews dedicated to each topic.

The largest portions of professional reviews are devoted to plot and character, and the smallest portions to the more personal topics of sentiment, evaluation, or identification with the characters. In contrast, amateur reviews show notably more use of personal topics, and less discussion of plot and character. The divide is clear: professional reviewers are matter-of-fact,

sticking to the who, what, and when of the novel, while amateur reviewers delve into more subjective topics. Notably, the romance genre shows the largest difference of any genre between how amateur and professional reviews discuss plot:




Figure 2.8: Percent of professional and amateur reviews discussing plot for four types of genre fiction.

Professionals talk about plot in romance reviews more than any other genre I examined (the orange bar extends furthest to the left for romance), while amateurs talk the least about plot in romance (the blue bar for romance extends the least far to the right). While professional reviewers base their discussion of romance on its plot, amateurs discuss plot relatively little, instead focusing on their sentimental reaction. We might conclude that professional reviewers rely on a formal conception of the romance genre as a way of couching their affective reactions to romance novels in the more impersonal topic of plot, while amateurs embrace more personal topics in accordance with Radway’s reader response theory of the romance, which defines genre in terms of readers’ reactions rather than its formal features.

We learned from Goodreads user Lauren, quoted at the beginning of this section, that “Lust or love at first sight can either flourish or flail in romance books.” The romance is a polarizing genre, and reviewers have different vocabularies for discussing romances that they loved or hated. By further breaking down reviews by positive (4-5 stars) and negative (1-2 stars)

ratings, we can determine whether professional and amateur reviewers have different standards of evaluation for romance novels. The chart below shows the words reviewers use most often in positive and negative reviews:




Figure 2.9: Most common words in reviews, broken down by professional vs. amateur reviewers and low-rated vs. high-rated reviews.

The location of each word in the chart is determined by how well it characterizes that group of reviews, whether professional or amateur (along the x-axis), or low-rated or high-rated (along the

y-axis).³² Each quadrant of the chart shows words that most distinguish of one of those four groups: the top right shows words distinctive of highly-rated amateur reviews, and the bottom left shows words distinctive of low-rated professional reviews. The chart helps to visualize the vocabularies of each of these types of review. Professional reviewers talk relatively factually about romance novels: the words furthest to the left, and therefore most distinctive of professional reviews, are necessary building blocks of a discussion of plot and character such as “family,” “woman,” “mother,” “sister,” and “father.” In contrast, the words most distinctive of amateur reviews are emotional: “loved,” “want,” “feel,” and “felt.” Along the y-axis, too, we can gain insight into what constitutes a good or bad romance for the two groups. For professional reviewers, we can tell that plot is a key criterion for evaluating the romance genre, since “plot” appears near the bottom of the chart—when professional reviewers mention plot, they are very likely to review the book negatively. If plot were not a strong factor in their evaluations, it would appear near the middle of the y-axis. Plot does not share the same evaluative function in amateur reviews, which instead use words like “perfect,” “beautiful,” “want,” and “right” in high-rated reviews as opposed to the blunter “sex” in low-rated reviews. It seems that, for amateur reviewers, romance is less about plot than an aesthetic judgment based on readers’ desire for and feeling of beauty and perfection.

Literary critic Rita Felski describes the process by which readers, in defiance of the constraints placed on professional critics, allow themselves to become enchanted by a book: “Once we face up to the limits of demystification as a critical method and a theoretical ideal, once we relinquish the modern dogma that our lives should become thoroughly disenchanted, we

³² The scores for each word are based on a method for finding keywords called Term Frequency-Inverse Document Frequency or TF-IDF, which emphasizes words that appear often in the reviews of one genre, but de-emphasizes them if they also appear frequently in reviews of other genres. The chart shows the results of performing TF-IDF on the corpus of reviews broken down both by status and by rating. Highly-rated indicates that the review had 4 or 5 stars; low-rated indicates 1 or 2 stars.

can truly begin to engage the affective and absorptive, the sensuous and somatic qualities of aesthetic experience” (Felski 76). As Felski would be first to remind us, this is not to say that professional critics pay no attention to aesthetic experience, but only that they are more limited in how they can express that judgment while amateurs are free to express the “sensuous and somatic” side of their aesthetic reaction. At first glance, it seems obvious that professionals would judge books on the basis of plot while amateur reviewers focus on their emotional response, but it is actually surprising—we pay professional critics because we assume that their reaction to a book is more valuable than that of an average reader. Why, then, do professionals focus so much on the matter-of-fact topic of plot? To make sense of these results, we might hypothesize that professional critics offer their trained perspective on the book’s aesthetic impact through the lens of plot. Rather than simply recording their own reaction, professionals ground their aesthetic response to books in plot rather than emotion, helping their audience understand why a book makes them happy or bored, or decide if it is worth their time. We might characterize this tendency as professional reviewers using plot as a purportedly objective touchstone that allows them to exert critical control over their emotional reactions to a text. In contrast, Goodreads shelves offer a variety of ways of expressing the full range of a book’s emotional impact through alternative genres like “made me cry,” “funny,” “creepy,” “disturbing,” and “guilty pleasure,” all shelves found on Goodreads, which allow users greater latitude to express their emotional reactions. For now, we can conclude that for amateur readers of romance novels, enchantment is one of the main criteria they use for evaluation, rather than the more traditional standards of plot or character that professional reviewers use, whether or not they consider (or conceal) their own enchantment in some other form. Seen in another way, both amateur and professional reviewers have aesthetic reactions to romance novels, but while

amateur reviewers express these reactions, professional reviewers trace them to their origins in formal features of the novel. Professional reviewers, then, use genre to categorize books based on their contents, in contrast with amateur reviewers, who use genre as a way to express and categorize their own reactions. As we will see in the following two sections, this amateur focus on personal reactions to a book manifests most clearly in the genres of mystery and horror.

Mystery and the Thrill of the Case

The point of a detective novel, after all, is for a mystery to be investigated, picked apart, and solved. Successful mysteries rely on unknowns, tension, and suspense. (Semnani for *The Los Angeles Review of Books*)

What I expect from a mystery novel is a good mystery which is not too obvious, clues which are revealed across the narrative which both lead me (the reader) toward and away from the correct conclusion, and a protagonist who also uses those clues toward the solution of said mystery. (Goodreads user J Austill)

Mystery, like science fiction, has benefitted from the weakening of traditional literary hierarchies. Originating in the detective tales of Edgar Allan Poe and Arthur Conan Doyle, the mystery had a rich history in periodicals and magazines, where many successful mysteries originally appeared. Serial publications like *Black Mask* and *Detective Fiction Weekly* generated suspense for their mystery tales as readers waited for the next installment. Critics disagree over what made some mystery series, like Doyle's Sherlock Holmes, so successful: was it a suspenseful plot that made full use of serialization (Priestman 90–91), the intriguing character of the detective (Kayman 49), or a superior use of the formal element of clues (Moretti, *Graphs, Maps, Trees* 74)? As we can see from the mystery reviews quoted above, what ties together the mystery for both amateur and professional reviewers is the act of solving it, and the suspense, characters, and clues are valuable only insofar as they contribute to (or frustrate) that goal. Does

this reading of the mystery hold up at scale? Can we find evidence for this definition of mysteries in thousands of reviews?

Just like science fiction, professional reviewers use a relatively standardized vocabulary to review mysteries and thrillers, with collocations that range from stereotypical descriptors of plot and characters (“complex” and “gritty”) to common subgenres (“classic” and “literary”). Amateur reviews, however, focus on the cognitive pleasures and challenges of reading a mystery, which entails an experience of “solving” and “unraveling” that is “fast-paced” and “entertaining.” Based on the way they talk in reviews, amateur readers see themselves as detectives attempting to solve the crime in the book and evaluating the quality of that crime-solving experience. Genre thus shifts from something inherent in the book to something outside the book—its genre is defined not by an internal characteristic but by a way of relating to the book that occurs when the reader attempts to search for clues and solve a case. Amateur reviewers have a different way of relating to genre than we have seen before: mysteries are not just words placed in certain formal patterns, but they allow the reader to take on the role of a detective and the fast-paced, thrilling affect that goes along with it. In turn, Goodreads shelves provide an outlet for this desire, emphasizing readers’ identification with characters in shelves such as “realistic,” “girl power,” and “kick-ass heroines,” and emphasizing their engagement with the book in shelves such as “couldn’t put down,” “edge of my seat,” and alternatively “lost interest.” These shelves offer an alternative way of grouping books based on the personal connection readers have to these books rather than a set of textual characteristics said to define the mystery genre.

If amateur readers are invested in the feeling of becoming a detective to solve the crime, we might expect them to use language that shows they are identifying with characters in the mystery novel. When we break down the topics mystery reviews discuss, this is indeed the case:




Figure 2.10: Average portion of professional and amateur mystery reviews dedicated to each topic.

Not only do amateur reviews use the topic of identification more than professional reviews—as we might expect of readers identifying with a detective as they solve the case themselves—but they also favor the other personal topics of evaluation and sentiment. Turning to professional reviews, they focused much more on plot and character than amateur reviews. This professional focus is explicit in the Los Angeles Times review of the mystery novel *Wife of the Gods*, which begins: “Top-drawer detective fiction is composed in equal parts of plot, character and an interesting setting evoked in such a way that it, too, becomes a narrative presence” (Rutten). This conception of mystery as the sum of its structural elements of plot, character, and setting indicates that professional reviewers tend toward a formal definition of the mystery genre, evaluating it based on what appears in the text. We see a much more personal account of the genre in amateur reviews, with a Goodreads review of *Wife of the Gods* instead emphasizing the

reviewer's identification with the plot and characters: "This brought back such vivid memories of Ghana! Well told, well told and so accurately reflecting the Ghana I knew but with cell phones added" (Helene). As we saw from the language of collocates amateurs use to describe the mystery genre, amateurs tend to be more invested in the experience of solving the mystery than the plot being "complex" or "gritty." Amateurs, then, seem to view genre as located in the reader's experience of the plot or identification with a character or setting. While both amateurs and professionals acknowledge the centrality of the mystery itself to the mystery genre, amateurs view the mystery through the lens of their personal reaction to the book, and professionals frame the mystery in terms of the complex plots and narrative suspense that prompt that reaction. But suspenseful plots are not the only way books inspire reviewers to share their personal reactions, as we will see next.

Horror and the Ambiance of Genre

Not surprisingly, in a horror collection in the classic vein (a gray, acrid atmosphere; surreal appearances; and a final scream...or silence), there's bound to be a stomach-turning few featuring wayward human parts: a woman's hand grows finger-by-finger in someone's wallet; within just a few pages, the "nightmare flower" crunches down a baby; and what is in grandma's homemade jelly? Most of the stories, though, are concerned with the deadly fog of intent that slowly creeps into the obsessing human brain. (*Kirkus*, "Nightmare Flower")

Somewhere along the line the horror genre got a bad rap because of all those slasher movies like *Saw* or *Hostel*. The horror I grew up with, which still exists by the way, wasn't about the gore. Not to say there wasn't gore sometimes, but it was about watching people just like you and me doing disturbing things in response to extraordinary situations. The horror aspect comes when you realize that, given those same conditions, you might do terrible things too. It is about watching others find that dark place and realizing you have one too, as well as understanding the "monster" and sympathizing with it. (Goodreads user Sharon)

Modern horror has its roots in the late eighteenth-century Gothic novel, a genre ranging from supernatural tales of suspense to psychological explorations of the uncanny. What tied together Gothic novels was their ability to inspire terror (and conversely, a sense of the sublime) in readers: “In order to replicate the ‘sense of self-annihilation’ associated with Gothic buildings, a novel—it was argued—had to arouse in its readers nothing less than metaphysical dread: some version—however fleeting or artificial—of that ‘universal apprehension of superior agency’ commonly associated with the supernatural” (Castle 690). The Gothic was a genre wholly defined by the reader, specifically the individual effect a text has on a reader. The affective impact of the Gothic continues in today’s horror novels and films. Noël Carroll describes horror as a narrative form that produces a particular affect in its audience. What makes horror so appealing, he argues, is its ability to transform our terror into pleasure: since we know what we are seeing isn’t real, “what would, by hypothesis, ordinarily distress, disturb, and disgust us, can also be the source of pleasure, interest, and attraction” (*The Philosophy of Horror* 189). Thus, for Carroll, horror is located in both plot and affect. For Castle, Carroll, and many other horror theorists, affect characterizes the genre. Do professional and amateur reviewers define horror in the same way?

Indeed, amateur reviews of horror novels expect them to create a certain mood, evidenced in reviews that frequently use terms like “breathless,” “imagery,” “setting,” “style,” and “ambiance.” This last word is particularly interesting given the typical digital humanities approach to genre, which focuses on word-based methods like frequency analysis and topic modeling—does this focus cause us to miss those aspects of genre that words cannot capture? What about imagery and other stylistic devices that reach past the words on the page to affect readers? Goodreads users create shelves to reflect this broad interest in characteristics of books

that “horror” cannot fully express, such as “creepy,” “weird,” “dark,” “spooky,” “Lovecraftian,” “quiet,” “speechless,” and “that feeling.” If horror is an ambiance created by moods, images, and feelings—words in particular patterns and contexts that have certain effects on individual readers—then a method that looks at words alone, or at books but not readers, or even at the horror genre but not at its many adjacent categories, misses these aspects of genre.

We saw with romance reviews that amateur reviewers used much more sentimental language than professional reviewers. In fact, this trend transcends genre—professional reviewers are distinctively reluctant to discuss a book’s emotional impact:




Figure 2.11: The overall portion of all professional and amateur reviews that discussed the book’s sentiment, broken down by the genre of the book under review.

In each of these four genres, amateurs use sentiment language much more often than professionals. But if it is a characteristic of professional style to avoid sentiment language, what kind of language do professionals use instead to discuss horror novels, which are so often defined by the sentiment they arouse? One topic professionals turn to instead is the context surrounding horror novels’ publication, author, and audience:




Figure 2.12: The overall portion of all professional and amateur reviews that discussed the book’s context, broken down by the genre of the book under review.

The orange bar extending furthest left shows professionals’ 6% greater use of the context topic in horror reviews compared to other genres. The most common types of context reviewers discussed included the movie version of the book, the book’s author, and other books in the series. From collocates, we learned that amateur reviewers were looking for an ambiance that is difficult to express in words. To express this ambiance, professionals, unwilling to use affective language to describe horror’s impact on them, might turn to a more impersonal and intellectualized discussion of the book’s publication context and the author’s oeuvre. Especially in the case of the contextual discussions of the book’s movie version, professional reviewers unable to express the indescribable ambiance of a horror novel in words might find it easier to describe the imagery, special effects, or soundtrack of the film version.

Professional reviewers are also more likely to discuss a horror novel’s characters:




Figure 2.2: The overall portion of all professional and amateur reviews that discussed the book’s characters, broken down by the genre of the book under review.

Professional reviews of the horror genre are more likely than other genres to discuss a book’s characters or characterization more generally. Much like a book’s publication context, its characters are a relatively matter-of-fact topic, allowing professional reviewers to list the book’s objective characteristics rather than their more subjective responses. Breaking down reviews by their topics reveals that professional horror reviews, of the four genres, focus the most on character and context, unwilling to use the more personal language amateur reviewers use to describe horror’s uncanny affect. We can see this abstract trend concretely in a pair of representative reviews. Take, for example, Stewart O’Nan’s horror novel *The Night Country*, an interesting case study because it could easily be classified as horror simply by its tropes, including a jack-o’-lantern on the cover, a Halloween setting, and ghosts as characters. The *Kirkus* review of *The Night Country* begins, “O’Nan (*Wish You Were Here*, 2001, etc.), who’s made a career exploring the dark side, welcomes Halloween with a ‘ghost story’ that soars when the supernatural lets good old-fashioned character take center stage” (*Kirkus*, “The Night Country”). While the review does acknowledge the book’s tropes, they aren’t its main focus—it immediately gives readers context about the author’s previous work and career trajectory, as well as insight into the novel’s characterization, asserting that these are two key aspects for

understanding the work and whether readers will enjoy it. In contrast, the typical amateur review of the same novel focuses much more on the novel's affect:

This book is not “scary”, but it **will** haunt you. It is sad, and at times funny, and just feels incredibly “true”. The actions of all the teenagers (living and otherwise), the downward spiral of Brooks (a cop involved in the tragedy), and the quiet but unshakable strength of the mother of one of the kids who suffered severe brain damage in the tragedy all feel spot-on, and weave together beautifully as the countdown towards Halloween and the 1-year anniversary of the tragedy approaches. I really felt like I was in the story, which is a fairly unique take on the ghost story. I originally read a print version of the book, but have since purchased the audio version, which I listen to in late October every year. It just wouldn't be Halloween without this haunting, sad, absolutely brilliant book! (Goodreads user Minge)

This Goodreads review offers a nuanced account of the sentiments the novel aroused in him (“feels,” “felt,” “sad”), his identification with the characters (“I really felt like I was in the story”), and his personal evaluation of the novel (“beautifully,” “brilliant”). Horror's dependence as a genre on a characteristic response from its readers makes such personal language necessary, leaving professional reviews seeming overly detached when they avoid it. In embracing personal language, amateur horror reviews advance an account of genre rooted in readers and their reactions to books, rather than the formal and contextualizing impulse of professional reviews. Like mystery, horror encourages amateur reviewers to commit fully to their personal reactions. Whether readers become immersed in a mystery by assuming the role of the detective, or in a horror novel by identifying with the victim, these two genres show the differences between amateur and professional reviewers most clearly: professional reviewers discuss plot and character, while amateur reviewers throw themselves into plots and characters.

The Topics of Amateur and Professional Reviews, at Scale

For each of these four genres, it is clear that while professional reviewers rely on formal features to define each genre, amateur reviewers use a variety of different criteria to judge them: for science fiction, how well the book fits into various subgenres and award hierarchies; for romance, how enchanting a relationship it portrays; for mystery, how exciting the process of solving the crime is; and for horror, how well the book creates an eerie ambiance for the reader. Amateur reviewers have a diverse and reader-focused set of expectations for each genre, setting their reviews apart from the more formal reviews of professionals. We can see these persistent differences between amateur and professional reviews by comparing how often they discuss each topic in their reviews, averaged across all four genres:




Figure 2.14: Differences in topic proportions between amateur and professional reviews.

The chart shows, for each colored topic, the overall proportion of reviews that amateur (top) and professional (bottom) reviewers devoted to that topic. While the amount of space that amateur and professional reviewers spent on average discussing context, genre, style, and other topics was roughly equal, other topics showed significant differences. Professional reviews devote more sentences on average to plot and character than do amateur reviews. However, professional reviewers do little explicit evaluation of the book, they rarely discuss the idea of identifying with themes or characters, and they largely ignore the sentiments or emotions a book caused them to feel. Amateur reviewers talk much more in these personal terms. In emphasizing more empirical

topics, professional reviewers not only do less evaluative work but also less interpretive work on the books they review under genre fiction headings than amateur reviewers. A further study might compare a third group—professional literary critics writing in academic journals—to see if the stylistic formality of professional writing influences the evaluative and interpretive content of literary scholarship.

In visualizing the focus of professional and amateur reviewers at scale, we see our earlier conclusions about individual genres reinforced, but we also see in more detail the zero-sum game of book reviewing: when professional reviewers focus more on plot, it is the more personal topics, rather than a discussion of style or “other” topics, that they leave out. Even though both groups are writing book reviews, the differences between the two groups are pronounced: in each case, professionals hold their own views close to the chest, while amateurs are not afraid to express their personal reactions.³³ Though this difference could simply be a matter of style and professional review conventions, Alan Liu would call it a symptom of professionals’ “emotional labor management,” in which systems of depersonalization beginning during the industrial revolution and continuing today in modern human resources departments attempt to eradicate undisciplined emotions from the work life of the white-collar professional (*The Laws of Cool* 90). The fact that professional critics talk little about the three most personal topics—sentiment, identification, and evaluation, each presupposing access to an individual’s opinion and difficult to disguise as objective or universal—reinforces the idea that professional book reviews

³³ For example, when professionals evaluate books, they use phrases like “A tedious trek through a footnote to history” compared to amateurs’ “This is literally one of the best books I’ve ever read.” Instead of expressing their own identification with the book, professionals predict how readers will react: “Graham’s awareness...goes a long way toward increasing reader sympathy” compared to amateurs’ “At times, I wanted to smack her. But I felt for her at the same time.” And when professionals mention sentiment at all, again they hide behind the hypothetical reader: “Readers inclined to lament their own circumstances may brighten up when considering the odds Ollison has overcome” as opposed to amateurs’ “I’m still disturbed by this story and if you’re squeamish don’t read it.”

discourage a focus on the reviewer's personal perspective, instead turning to formal aspects of the book to characterize its genre.

Not only did amateur reviewers talk more than professionals about personal topics, but they also had much more freedom in the proportions they used for each topic:




Figure 2.15: Variation in topic proportions between amateur and professional reviews for three topics.

In the chart above, each point represents a single book review, distributed along the horizontal axis depending on how much of the review discussed evaluation, identification, or sentiment. The width of the gray box summarizing each reviewer group shows the range of proportions reviews most often fall into for each topic. A larger width indicates that reviews have a wider range of possible proportions—rather than reviews typically adhering to a standardized formula that allows only certain amounts of discussion for each topic, reviews can fall at any point along a spectrum from very few to quite a lot of sentences focused on that topic. For each of the three personal topics in the chart above, amateur reviews are much less standardized in how much of each topic they discuss. On average, they tend to discuss these three topics more often, and they have more flexibility to vary the amount of each topic that they use. In contrast, professional reviewers tend to stay within a narrow (and low) range of discussion of evaluation, identification, and sentiment. When professional reviewers write about a particular genre, they can only use a narrow range of terms and access limited areas of experience to describe their perspective on the

genre. These restrictions are perhaps due to editorial pressures, but the effect amounts to much more than simply a stylistic difference in formality. The editorial constraints on professional reviews result in a treatment of genre that focuses overwhelmingly on formal aspects of the novel like plot and character, and leaves little room for one's personal experience of the book. These constraints have a measurable impact on how professionals construct genre, limiting the scope of their definitions largely to the narrative structures and tropes that one can discern from a book's plot and characterization and leaving no room for reader-focused accounts of the book's affective impact or its ability to inspire readers to identify with its characters.

When reviewing in a professional context means focusing on the supposedly objective measures of plot and character, reviewers define genre in those terms. And when amateur reviewers overwhelmingly express their emotions and personal preferences, their view of genre becomes a function of the mood and feelings it arouses in them as readers. As I have shown, professional platforms place stricter limitations on reviewers than their amateur counterparts. When professional publishing standards constrain a reviewer's ability to express a book's affective impact, they leave out what best distinguishes genres for non-professional readers. Professional criticism needs to rethink its constraints on style, such as leaving out emotions and speaking in terms of a generalized reader rather than one's personal reaction, lest it miss what makes genre so poignant for modern readers.

But at the same time, amateur readers' focus on the more personal side of genre is not without its own problems. I have shown in each of the four genres I examined that amateur reviews treat genre as something determined internally, in a reader's reaction to a book and their personal identity. Amateur reviewing platforms encourage this transformation of genre from a formal literary category into a consumer category because it helps them produce more accurate

demographic data and recommendation algorithms. For users, though, these increasingly specific consumer categories are not necessarily more flexible, as we will see in the following chapter, which looks closely at this process of consumer segmentation as it functions on Goodreads. I have hinted at the wide variety of personal shelves that Goodreads users create to broaden the scope of genre on the platform, from “kick-ass heroine” to “guilty pleasures,” but in the next chapter I will examine Goodreads’ shelving system in detail to determine how it shapes amateur reviewers’ conceptions of genre. As I have shown, platforms have a significant and measurable effect on how their users discuss and define genre, and a more detailed analysis of a single, influential platform will clarify how this influence operates, and the effect it has on its users’ conceptions of genre.

Chapter 3: The Genres of Goodreads

As we saw in the previous chapter, in amateur social reading communities like Goodreads, users talk more personally than professional reviewers about genre. Goodreads encourages its users to express their individuality through personally named bookshelves, membership in a wide range of book clubs and fan groups, and access to an exhaustive catalog of books. But at the same time, Goodreads offers prominent genre labels for each book and bases its recommendations on clear genre categories, using genre prescriptively to separate both books and people in a way that undermines the flexibility of genres on the site. Goodreads, then, serves as a contradictory space where its users' diverse conceptions of genre clash with a website that encourages users to compromise on shared genre definitions.

Which impulse wins out? This chapter uses computational methods to track the expansion and consolidation of genre on Goodreads at scale. First, I map the diverse field of user-created categories on Goodreads by clustering bookshelves that readers use similarly, showing the wide range of literary and para-literary labels that Goodreads users apply to books. By examining the clusters that take shape from users' interactions on the site, we can observe an emergent landscape of genre based on readers' actual use of genre labels rather than pre-existing or theoretical genre hierarchies. As I will argue, these clusters are more apt than traditional genre terms at characterizing the way genre functions on social reading platforms like Goodreads.

Despite the diverse categories and clusters users have access to on Goodreads, when you examine the way users discuss genre in their reviews, that diversity fades. In fact, by training a computer to guess which genre a review is discussing, it becomes clear that reviews grow noticeably more predictable in how they talk about genre over the decade that Goodreads has

operated. Goodreads encourages this predictability in its design and underlying algorithms, an effect I will examine in detail.

Like many twenty-first-century technology companies, Goodreads' product is its users and their personal preferences. Goodreads uses genre as another consumer category, alongside age, gender, and location, to better target its recommendations and advertisements. As I will show, Goodreads' transformation of genre from literary to consumer category has a clear standardizing effect on the way readers discuss genre on the site. But this transformation is not unprecedented. To fully understand genre's most recent transformations, it is helpful to view genre in its broader historical context as a category with a wide range of functions in the literary market.

Genre Theory on the Literary Market

In the modern literary market, with more books published every year than we will ever have time to read, one of genre's most common uses is to help us find books we are likely to enjoy. But genre's utility in telling us what to expect from a book has the consequence of limiting the people who might read the book and the ways they might interpret it. The act of labeling a book with a genre can cut off a significant portion of its possible audience or restrict its potential impact. This discriminative function of genre is not new, but while genre has served for millennia as a literary term for classifying books, the way we define genre and the models we have built to organize genres have changed significantly since the first theories of genre. From early taxonomies that marked some genres as valuable and others as frivolous or even dangerous,

to modern systems that outwardly reject prescriptive hierarchies but underneath run on rigidly defined database systems, genre has long dictated the limits of our literary experiences.

The Western tradition of genre theory is largely built on Aristotle's *Poetics*. Aristotle's genre system was hierarchical at every turn, elevating certain types of writing, genres, and poetic forms above the rest. Aristotle places poetry above other forms, calling it "more philosophical and more ethically serious than history" (sec.9). Within poetry, Aristotle gives special attention to three genres—tragedy, epic, and comedy—which he argues most effectively accomplish the pleasure of poetic unity by achieving unity of time, place, and action in their construction. Though other poetic genres certainly existed at the time, Aristotle excludes them from his *Poetics*, leaving out major poets from lesser genres including the didactic poet Hesiod, the elegiac writer Solon, and the choral lyricist Pindar (Halliwell 154). Of Aristotle's three privileged genres, tragedy is at the top of the hierarchy, and in fact, Aristotle spends the most time outlining his theory of the genre. Tragedy comes closest to achieving the unity that Aristotle sees as the ultimate goal of poetry. Aristotle describes tragedy as a work where "the structural union of the parts being such that, if any one of them is displaced or removed, the whole will be disjointed and disturbed. For a thing whose presence or absence makes no visible difference, is not an organic part of the whole" (sec.8). For a work to be unified, it can have no superfluous plots, and in the same way, Aristotle's system of genre admits no superfluous genres. For Aristotle, genre is a system in which every part has a place and is essential to the whole, and tragedy marks itself as essential through its demonstration of a unified structure. Epic is tragedy's prototype, aspiring to tragedy's status but falling short in its ability to produce the emotional and cognitive pleasure of unity that tragedy achieves. Comedy he places below the rest, as a less serious genre with base characters and trivial plots (sec.5). In his differentiation

between comedy (a fragmented genre) and tragedy (a unified whole), Aristotle constructs a hierarchy of genres, with unity as the standard for distinction. Within each genre, too, Aristotle charts a progression from primitive initial experimentation to an ultimate perfection of the genre, as he traces through the development of early hymns and encomia into Homer's dramatic epic, followed by the dithyramb, ultimately leading to Attic tragedy (sec.4). In each of these moves, Aristotle places literary works into hierarchical relationships to each other, with clear statements of each work's relative value.

In practice, the classical hierarchy of genre allowed its adopters and curators to determine not just which works were most valuable, but also which works belong in the hierarchy at all. According to the widely-cited Greek geographer Strabo, Aristotle was "the first man, so far as I know, to have collected books and to have taught the kings in Egypt how to arrange a library," and his works were some of the first that the Library of Alexandria acquired (13.1.54). The library organized its books using plaques marked with poetic genres like tragedy and epic, as well as non-poetic genres that Aristotle discussed in other works like philosophy, history, and medicine, and betrayed the influence of Aristotle's system of value judgments in excluding comic poets from the collection (Witty 135). Aristotle's hierarchy of genre provided an effective means of determining which works belonged in the library and which it should exclude. In this way, basing genre on hierarchical principles serves as an effective organizing tool, but leads directly to an exclusive system that relegates books that don't fit into the system to obscurity.

Rather than correcting the favoritism that the classical genre system lent to certain genres, Renaissance thinkers perpetuated this hierarchical system. While Aristotle only considered fictional verse within his account of poetics, Renaissance thinkers developed Aristotle's system to make it more comprehensive, with humanists like Giovanni Giorgio Trissino, Sperone Speroni

and Giraldi Cintio expanding Aristotle's account of the tragic and applying his standards of value to the newer genres of *elocutio* and *romanzi* (Cottino-Jones 575). Many Renaissance scholars bolstered their theories with a scientific authority derived from Linnaeus' taxonomy of biological organisms by using terms from biology like "species," "root," and "stem" to describe novels. These theorists argued over whether certain new genres truly belonged in the same species as more respected genres, with thinkers like Etienne Jodelle attempting to lend legitimacy to novels by calling them "a species of historiographic composition" (Norton 310), and Francesco Bonciani arguing that prose narrative was not truly "a species of poetics" (Norton and Cottino-Jones 326). These debates over species membership used the clout of scientific metaphor to justify their literary judgments, further solidifying a system in which cultural elites determine which literary works have value while adding a sense of taxonomic predictability to the progress of literary innovation. During this period, however, the invention of movable type was driving increasing literacy and literary output as well as the rise of the novel with its many subgenres, which resulted in a wave of new books that didn't easily fit into any existing categories and a proliferation of new categories as critics tried to keep up, throwing Renaissance attempts at constructing a clear genre system into turmoil.

From the classical period through the Renaissance, one of genre's most important functions was discriminative, a way to distinguish good books from bad and valuable genres from drivel. Critics' most important function in this system was to inform the public how to read each genre properly, usually through didactic prefaces and treatises accompanying the great works of literature (Lipking). In the Romantic period, however, we see a dominant push to create a theory of genre that did not aim to exclude or compare, but rather treated every genre individually to determine its own unique aesthetics. Romantic writer Victor Hugo urged not to

judge writers by “rules and genres” (107), critic Friedrich Schlegel declared that “every poem is a genre in itself” (*Literary Notebooks* 116), and philosopher Friedrich Schelling called the *Divine Comedy* “so completely self-enclosed that any theory abstracted from more individual forms is completely inadequate for it; as its own world it also requires its own theory” (240). Rather than assembling ever-more-comprehensive taxonomies, Romantic thinkers aspired to an aesthetic egalitarianism in which all poetry would be evaluated based on its beauty rather than its genre classification. The Romantic idea of genre is most clear in Schlegel’s conception of Romantic poetry: resisting the typical practice of using modifiers like ‘Romantic’ to divide poetry into genres, Schlegel uses the term to refer to an idea that transcends genre, describing Romantic poetry as “a progressive universal poetry. Its destiny is not merely to reunite all of the different genres....For, in a certain sense, all poetry is or should be Romantic” (“Athenaeum Fragments”). Schlegel uses the concept of Romantic poetry not to divide poetry but instead to emphasize what unites all poetry. In rejecting the historical practice of discussing poetry in terms of its generic affiliation and instead embracing the aesthetic and unifying characteristics of poetry, the Romantics showed the shallowness of any account of genre that was strictly hierarchical.

But while Romantic thinkers perfected their aesthetics, the literary world kept expanding, and the sheer volume of literary output by the nineteenth century demanded a system of organization that could help booksellers track their sales and segment their customers.³⁴ Not giving up on genre, critics in the late nineteenth century looked for a way to classify books that could account for the increasing diversity and messiness of the literary market. Darwin’s evolutionary theory proved to be a compelling model for both critics and booksellers, since it

³⁴ For a detailed treatment of the ever-increasing output and diversity of the literary market leading up to the nineteenth century, see: (L. J. Miller 25; Long, “The Cultural Meaning of Concentration in Publishing” 18). Franco Moretti’s edited collection *The Novel* also includes several essays detailing the literary market during this period, notably: (Austin).

accounted for the constant change in forms and styles each genre underwent over time, allowing for a fluid conception of genre that tracked natural variations rather than an *a priori* definition that would become out of date with the next innovative bestseller. The most notable of these evolutionary genre theorists was French critic Ferdinand Brunetière, who traced the history of genres and their development over time as a linear evolution from one novelistic genre to the next, from the historical romance (*chanson de geste*), to the epic novel (*roman épique*), to the novel of manners (*roman de moeurs*) (Dames 517). Evolutionary models of genre were influential, and have persisted through the twentieth century in David Fishelov's conception of a genre's survival³⁵ and Franco Moretti's vision of a world literature shaped like an evolutionary tree.³⁶ While the evolutionary model is more flexible than previous hierarchical conceptions of genre, it still holds that modern genre categories represent genre's most advanced and sophisticated incarnation. In this way, evolutionary models of genre resemble earlier hierarchical systems in privileging one type of text over others, namely the commercially successful text over

³⁵ In *Metaphors of Genre*, David Fishelov aims to develop the evolutionary comparison specifically for the case of literature. Fishelov details four common "metaphors of genre": biological species, families, social institutions, and speech acts. He argues that a genre's survival relies on "the relationships between the production of the texts 'belonging' to the genre and the literary and cultural environment and how this environment coerces (or thwarts) the genre's production" (39). In this way, Fishelov deepened the evolutionary metaphor by adding an account of a text's production process and its environment, two elements impacting its "evolution." His theory also shows a greater understanding of biological evolution by indicating that the evolution of genres isn't smooth like a tree, but rather staggered and uneven (46–48). While Fishelov's evolutionary perspective on genre was a popular way of looking at genre, much of his argument relies on traditional genres and canonical authors, leaving out a large portion of literary material that does not fit these constraints.

³⁶ Franco Moretti attempts to modernize evolutionary theories of genre by applying the model of the evolutionary tree to literary history. His stated goal is to transform the act of studying world literature to look more like "Take a form, follow it from space to space, and study the reasons for its transformations" (*Graphs, Maps, Trees* 90)—a process meant to sound similar to evolutionary science. While Moretti's work is provocative, later digital humanists have developed competing models for the literary market. For example, Moretti notes a generational pattern in the publication of genre fiction over time that Moretti notices in *Graphs, Maps, Trees*, but when Ted Underwood looked at the language of genre fiction, he found that the pattern wasn't a series of generational shifts in genre or even a gradual consolidation of genre but instead an initial composition that doesn't change much over a century (Underwood).

the abandoned form, and limiting the realm of possible genre definitions to only include what can be explained linearly through small formal changes resulting in direct market responses.

Despite continued attempts at modeling the evolutionary development of genres, as the literary field continued to expand and diversify, evolutionary theories based on a natural order were soon overwhelmed by exceptions—even Moretti, despite his extensive use of evolutionary trees for mapping the rise and fall of the novel, called the state of genre taxonomy in the twentieth century “a total disaster” given the diversity of nineteenth-century literary material and narrative genres not yet catalogued even a century later (“Narrative Markets, ca. 1850” 152). Turning away from theories of genre that relied on a natural or scientific order, twentieth-century literary critics began to reject genre taxonomies, asserting that scientific structures could never fully account for literature’s diversity. Philosopher Jacques Derrida expresses the twentieth-century apathy for strict genre taxonomies: “As soon as the word *genre* is sounded, as soon as it is heard, as soon as one attempts to conceive it, a limit is drawn. And when a limit is established, norms and interdictions are not far behind: ‘Do’, ‘Do not’, says ‘genre’, the word *genre*, the figure, the voice, or the law of genre” (56). Instead, post-structuralist critics led the creation of a new theory of genres that saw them as historically contingent, mutable, and characterized by weak and overlapping boundaries. One strain of theorists attempted to dive deeper into the text to discover how each text uses genre conventions to achieve its artistic goals, as post-structuralist critic Roland Barthes does in *S/Z*, demonstrating that “the single text is good for all the texts of literature...[it is not] the access to a Model, but an entry into a network with a thousand entries” (12). Each book, for Barthes, can activate many different genre conventions that only emerge through deep analysis, not sweeping classifications of the external structures of genre. Another wave of theorists pushed for a wider critical perspective, based on the insight that “genres are

never, as frequently perceived, objects which already exist in the world and which are subsequently studied by genre critics, but fluid and tenuous constructions made by the interaction of various claims and practices by writers, producers, distributors, marketers, readers, fans, critics and other discursive agents” (Bould and Vint 48). Looking with renewed vigor at forces outside the text, these scholars considered a range of cultural and historical influences, including that of societal conventions like communications scholar Carolyn Miller, who declared that “genre study is valuable not because it might permit the creation of some kind of taxonomy, but because it emphasizes some social and historical aspects of rhetoric that other perspectives do not (151).³⁷ Fields emerged to examine in detail the many non-formal factors that influence genre, including reception theory, which assessed the role of a reader’s expectations, references, and mindset in shaping their understanding of a genre.³⁸ In this case, the increased focus on readers (rather than literary critics and cultural elites) allowed a newly invigorated academic focus on less well-studied, culturally devalued genres with nevertheless dedicated fanbases,

³⁷ Carolyn Miller places genre outside the text in a social interplay of rhetorical situations that inspire regular actions in their participants. She argues that “a rhetorically sound definition of genre must be centered not on the substance or the form of discourse but on the action it is used to accomplish” (151). Understanding genre in terms of repeated situations and responses has inspired a focus on forms of writing outside of typical literary genres, such as government reports, business documents, and communally written work. More recently, her work has focused on the online genre of the blog, and her analysis of bloggers shows that the genre is driven most commonly by self-expression and community development (Miller and Shepherd). Miller’s work on blogging underlines the need for a separate analysis of new genres that arise online, such as those on social reading websites like Goodreads.

³⁸ Reception theorist Hans Robert Jauss argued that genre causes a range of expectations for a written work in an interplay between the reader and the work’s historical context. He called this area of negotiation the “horizon of expectation,” referring to the reader’s expectations, references, and mindset for a given work, which it was the critic’s duty to “objectify” in order to analyze a text. In Jauss’ theory, a work cannot be without genre, and genre is essential to understanding literature: “Just as there is no act of verbal communication that is not related to a general, socially or situationally conditioned norm or convention, it is also unimaginable that a literary work set itself into an informational vacuum, without indicating a specific situation of understanding. To this extent, every work belongs to a genre” (79). Genre, then, as an essential part of every literary work, makes the horizon of expectations explicit: for example, *Don Quixote* references the genre of tales of knighthood, which tells readers how they’re meant to read it and what to compare it to. This sense of genre as a way to ascertain a reader’s expectations became foundational in modern marketing strategies that see genre as a way to clarify audience expectations for a given product so that companies can ensure that their products both create the right expectations and fulfill them for their customers.

including romance,³⁹ the Gothic,⁴⁰ and science fiction.⁴¹ Concluding that taxonomies can never fully explain books, critics largely relinquished the role of policing genre terms and boundaries, and instead moved to develop more thorough accounts of the many factors that influence literary works, including author, publisher, and reader.

At the same time, booksellers began to incorporate these extra-novelistic factors into their sales practices. Facing increased competition and a more diverse and unpredictable customer base, they needed a way to track their customers' reading preferences, and they began to collect more and more data about books' authors, publishers, and readers. But unlike critics, who used this more complex picture of genre to dismantle literary taxonomies, booksellers used their data-driven perspective on the literary field to double down on taxonomy. We see taxonomic ways of thinking about books throughout twentieth-century publishing practices in book-of-the-month

³⁹ Reader response critic Janice Radway argues for the importance of studying the undervalued romance genre in *Reading the Romance*. In it, Radway pushes back against previous arguments framing the popular romance as damaging to women. This position, Radway argues, assumes that readers accept the ideology of what they read unquestioningly. Instead, Radway demonstrates that romance readers are far from passive recipients of romance stereotypes. The readers she interviews value above all a strong, intelligent heroine over one who is feminine, and a caring, tender hero, counter to the common trope of the gruff, inscrutable hero and the passive heroine (*Reading the Romance* 119–24). Radway's interviews show that the romance genre allows the women freedom, not only through the act of reading, which lets them escape the demands of everyday life, but through a plot that enables responses that can both confirm and reject female disempowerment. Radway's key shift is to separate the content of these popular romance books, which often seems to reinforce gendered stereotypes, from how readers receive them, which can question these stereotypes and place the books in a more progressive context.

⁴⁰ In "The Gothic Novel," Terry Castle traces and explains the resurgence of interest in the Gothic in the twentieth century. Gothic was popular in the late eighteenth century, she argues, not because of a specific mood or identifiable formal trait, but its ability to inspire terror in its readers. Further demonstrating the shift in twentieth-century genre criticism from formal and taxonomic theories to a full account of readers and their reactions to genre, Castle's explanation for the resurgence of interest in the Gothic is focused on readers: "It should not surprise us that in our own age – in which demonic images from the past haunt us sleeping and waking – we should find in the Gothic resurgence of the late eighteenth century such a powerful reflection of our own aspirations and fears" (678).

⁴¹ As one example of science fiction's recent reinvigoration as a site of genre study, ecocritic Ursula K. Heise explores how science fiction novelists use the genre's ability to present alternative futures to illuminate modern debates over environmental concerns. Science fiction posits worlds where characters confront the ethical implications of biological scarcity ("Reduced Ecologies"), imagines posthuman futures where humans are just one perspective among many including aliens, androids, and animals ("The Android and the Animal"), and uses other planets as settings to address the question of how to build new ecologies on our own planet ("Marian Ecologies"). Like Radway's recuperation of the romance and Castle's argument for the relevance of the Gothic, Heise reminds us of the real impact even a traditionally maligned genre can have on readers and their conceptions of the environment.

clubs, bookstore chains, and now online book retailers thriving in recent decades. Reader response critic Janice Radway's account of the Book of the Month Club, founded in 1926, emphasizes the club's primacy of categorization as a necessary first step before evaluating a book's quality: Radway cites the club's first editor-in-chief Henry Canby, who claimed, "'There is no help except to set books upon their planes and assort them into their categories—which is merely to define them before beginning to criticize.' The crucial move in the evaluative practice of the Book-of-the-Month Club judges was not judgment at all but the activity of categorization, of sorting onto different planes" (*A Feeling for Books* 273). In prioritizing genre categorization above determining a book's value, Radway argues that the club put readers' expectations above their traditional duty as an institution to provide the public with exceptional literature and avoid frivolous genre fiction and bestsellers. Continuing this commercial focus on the reader's desires above literary value judgments, the bookstore chains that rose to prominence in the 1960s began to use state-of-the-art computer systems to catalog sales and customer data in order to better predict customer demand, using not just data about a book's genre but also its "author, title, edition, publisher, binding, publication date, language, etc." (Striphas 92). The extensive cataloging of books throughout the twentieth century led online commerce website Amazon to start by selling books, since they were already well-cataloged by systems like the ISBN and barcode. In each of these developments of the literary market, booksellers incorporated more and more data into their orbit in the name of increased personalization and better anticipation of their customers' desires. The taxonomic, one-size-fits-all organization systems developed at the beginning of the twentieth century, like the Dewey Decimal and Library of Congress Classification systems, look simplistic and inflexible when compared to the extensive databases and predictive models of the late twentieth century. Though these older systems were designed to

incorporate all of literary material into one system, they are unable to adjust to the changing literary market and consumer demands like a constantly-updating digital database can. Compare, for example, the Dewey Decimal system's outdated treatment of religion with eight sections devoted to Christianity and one section for "Other religions," to Amazon's extensive categories for religious books including Christianity, Judaism, Hinduism, New Age, Buddhism, and the occult, with the ability to add new categories as they arise. While twentieth-century literary critics have rejected genre taxonomies, the literary market uses genre as just one element within ever-expanding and ever-updating databases that aim to incorporate as much para-literary material as possible into their predictive systems.

What role has genre been relegated to, now that it is just one feature among many that companies use to predict customers' reactions to books? Marxist critic Fredric Jameson has argued that as literature has become more commodified, genres turn "into a brand-name system against which any authentic artistic expression must necessarily struggle" (93).⁴² In the place of well-known and trusted critics, twenty-first-century publishers have used branding to mark their books' value. In branding, publishers use features external to the book itself to mark its value and tell readers what kind of experience it offers. In this way, external features like its cover design or an endorsement by Oprah begin to mark its importance rather than internal features like its style or theme. As genre theorists have shifted from formalist to historically contextualized

⁴² In *The Political Unconscious*, Fredric Jameson details the transformation of genre from cultural institution to mere brand. Defining genres as "social contracts between a writer and a specific public, whose function is to specify the proper use of a particular cultural artifact" (92), he notes that genre has largely fallen out of critical favor, declaring it "thoroughly discredited by modern literary theory and practice" (91). He cites several reasons for genre's decline, including that the increasingly mediated social lives of modern texts have proved too unwieldy to fit standardized generic rules, and that the cultural institutions reinforcing genre have fallen prey to market forces. While Jameson's theory of genre as brand describes the twentieth-century literary market well, with new technologies that support amateur reviewing and social tagging, it is necessary to update Jameson's conception of genre to describe twenty-first-century genre as it functions online.

accounts of genre, then, outside of the academy genre has made a similar shift, becoming an external feature that companies use as a brand to tell consumers what to expect. We see this concretely in genre-focused brands like Harlequin Romance and many publishers' genre-specific imprints, but also more generally in many stores' treatment of books as decorative objects meant to make a statement rather than meant to be read. In the first half of the twentieth century, books were commonly loss leaders in department stores, but stores kept selling them because they gave the store an air of refinement, attracted wealthy customers, and inspired purchasing through their genre's brand appeal: "Not only were department store managers urged to sprinkle books throughout different departments and in store windows, but book departments were told to display merchandise that tied in with particular books, such as luggage with travel titles, cameras with photography books, and so forth. Even fiction could serve this purpose: 'A society novel will make every reader wish to dress well'" (L. J. Miller 36–37). These stores treated a book's genre as external to the book, acting as merely a stamp of approval for a set of products, functionally offering the genre's brand endorsement.

As the literary market has moved online, new models have arisen to complicate the branding of genre. The Internet's promise of an open and democratic platform offers models like social tagging, where a website's users can choose for themselves what genre a book belongs to rather than only allowing authors, publishers, or booksellers to dictate its genre. Goodreads is the clearest example of this trend: the site adds features like personal shelf titles that allow readers to label genres freely and fluidly, seeming to answer the call for a more democratized critical apparatus. But rather than creating more diverse genres, often these systems simply create even more comprehensive taxonomies, as this chapter will show. Instead of technology making genre classification more transparent and flexible, online genre systems often make classifications

seem less ambiguous than they truly are, while obscuring any biases or mistakes. They make genre seem objective by hiding the historical and social context that comes with it. But by reading book reviews on Goodreads, we can return context to these choices—and see how Goodreads users are adopting or creating taxonomies of their own—to better understand how people are thinking about genre today.

Technology companies are already using algorithms to make genre classifications and recommend books based on genre preferences. This chapter conducts a case study of what's happening to genre on Goodreads: what does the twenty-first-century landscape of genre look like, and what happens when these digital genres receive their definitions from users and platforms rather than experts and professional critics? Digital humanities offers a way to push genre theory forward by providing a fuller account of the cultural factors shaping genre online and providing computational means to measure how genre is changing at scale.

The current trend is for technology to fuel even more restrictive taxonomies. But it's possible to use the affordances of the digital to support a more flexible model of genre than we've seen before. Recent critics have called “for a more fluid reading across forms, genres, and periods than is the prevailing norm in academic criticism today” (Ngai 7), and suggested that genre offers an opportunity for a more diverse curriculum through its ability to span time and place (Dimock 1384). But before these visions of a more fluid twenty-first-century genre theory can become concrete, we must first understand how genre functions today on one of the most popular sites of book discussion online.

Reading Genre on Goodreads

As artificial intelligence has grown more advanced, we have begun to cede more and more of our research and decision-making to algorithms. They help us find new music, TV shows, restaurants, and even relationships. The promise of accurate algorithmic recommendation has transformed the publishing industry, with companies like Amazon using big data to guide customers to books they'll want to buy. But what drives these recommendations? For Goodreads, a book discussion and recommendation platform, everything revolves around genre. The site uses genre as its most meaningful organizational category, standing in as shorthand for its users' reading preferences. Goodreads' emphasis on genre is visible from its most basic design features to its recommendation system and community construction.

Even a cursory glance at the website shows that genre dominates its design and organization. Just visit the homepage (see Figure 1 below), which advertises its book recommendations based on "what titles or genres you've enjoyed in the past," encourages users to browse books by genre, offers Goodreads Deals on books grouped by genre, and displays Goodreads Choice Awards for each genre. After logging in, users who select the "Browse" menu see a list of their favorite genres and a popular book from one of those genres, and Goodreads offers them personal recommendations for each genre. Every book's webpage has a prominent section titled "Genres" which lists the most popular labels readers have given that book.

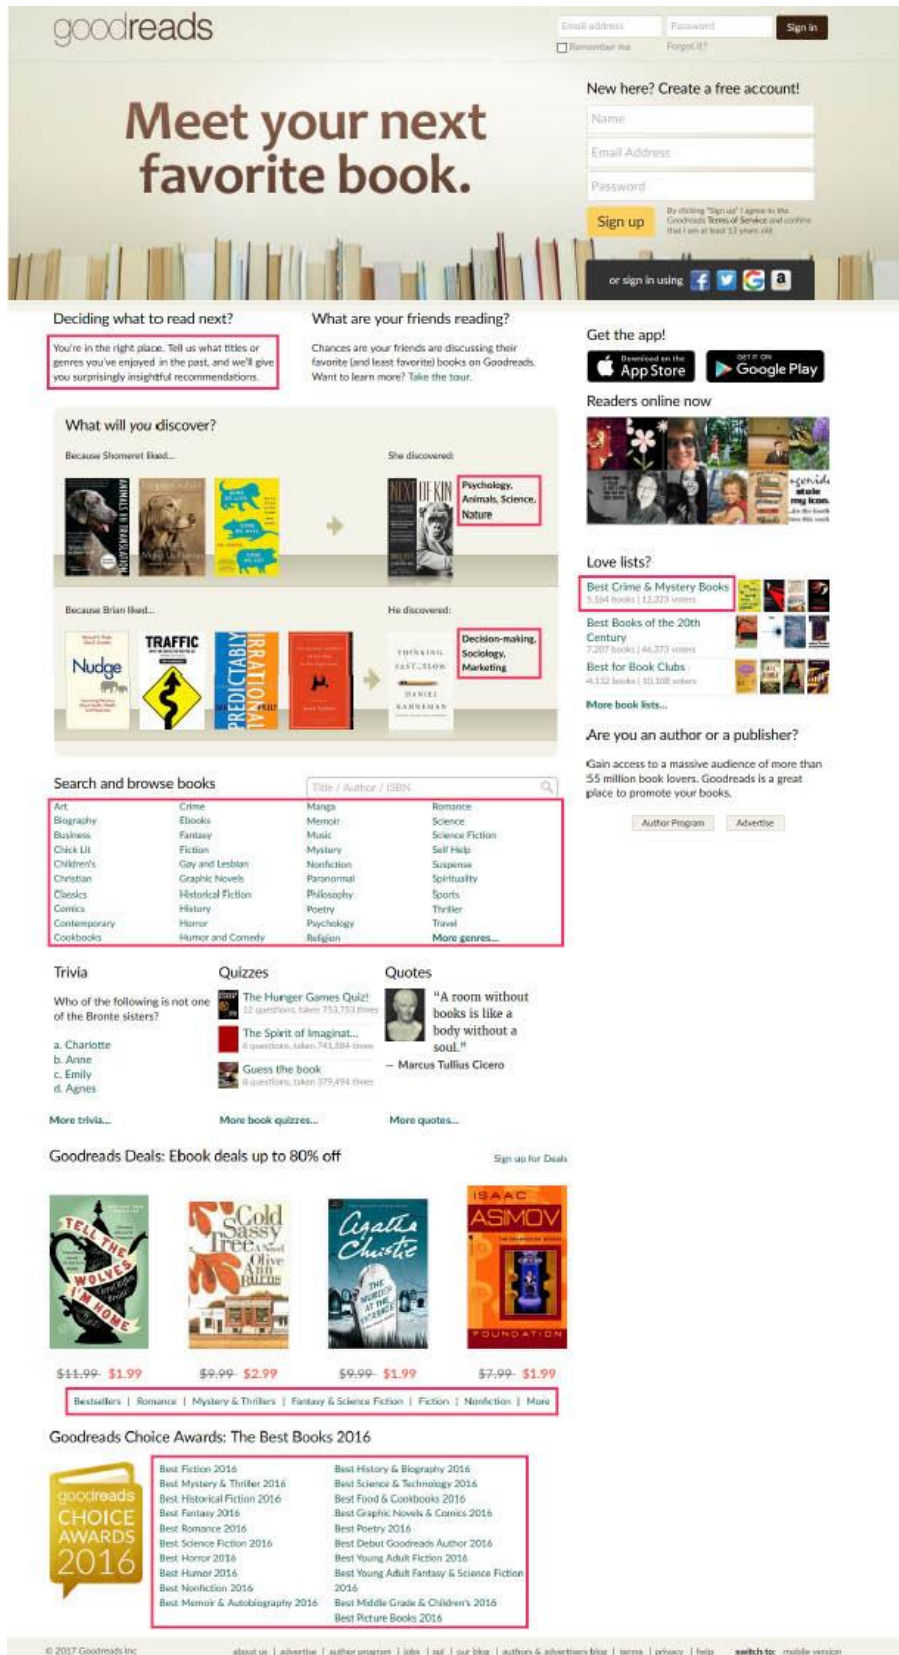


Figure 3.1: Goodreads homepage with elements related to genre boxed in red.

In this intense focus on genre as the basis for the site’s organization, Goodreads assumes that its readers want to find books by browsing genres rather than by other possible characteristics including a book’s time period, publisher, national origin, or author. Most significantly, Goodreads’ use of genre for browsing books mirrors e-commerce sites that allow users to browse goods by product categories—users navigate through genres in the same way they navigate through “Clothing” or “Electronics” on Target.com. The utilitarian function of genre on Goodreads as a navigational tool rather than a privileged term with a long literary history reveals that genre has transformed in its migration online. On Goodreads, genre is just another product category.

One of the site’s most vaunted features is its personalized book recommendation system. This, too, relies on the assumption that genre is an accurate predictor of users’ preferences. Users can ask the site for recommendations based on a favorite genre, subgenre, or other product category they create. For example, if a user has given a high rating to the social psychology books *Nudge* and *Predictably Irrational*, Goodreads will suggest the novel *Thinking Fast and Slow* based on the shared genres—or, more accurately, product categories labeled as and functioning as genres—“Decision-making,” “Sociology,” and “Marketing”:



Figure 3.2: Example of a Goodreads recommendation based on genre from the homepage.

The recommendations take into account the book's genre, the user's ratings for other books in the genre, and the ratings of other users who like the same genres. In a similar way, genre informs Goodreads' targeted advertisements. A 2013 announcement explains, "Publishers and authors are able to buy ads on Goodreads that reach readers of certain books and genres to promote their upcoming books. So, if you're a fan of Lee Child, you might see ads for a new thriller by a debut author" (Goodreads, "FAQs on Amazon's Acquisition of Goodreads"). Like recommendations, Goodreads targets advertisements to users based on their genre fandom and reading history. For both recommendations and advertisements, two features central to Goodreads' value proposition, genre is essential to providing accurately targeted results. While genre may not even be the most predictive factor in a reader's enjoyment, and Goodreads' proprietary algorithms certainly consider other criteria such as a book's author, popularity, and publishing date, genre is by far the most visible part of its recommendation and advertisement system.

Goodreads also promotes itself as a place where readers can come together to discuss books and build online communities of readers. In designing these communities, again Goodreads implies that genre is a coherent and legible category with which to bring together like-minded people. Goodreads creates and features genre-based book clubs, discussion groups, and "Favorite Genre" sections in every user profile, assuming that its users want to primarily identify and interact with other genre fans rather than forming groups and friendships based on different characteristics, such as their nationality or age. In groups like "The Mystery, Crime, and Thriller Group," "The History Book Club," and "Horror Aficionados," readers discuss their favorite books and scenes in a community that requires nothing in common except a shared

genre fandom. For example, the historical fiction group describes the type of reader it aims to attract:

Welcome to Historical Fictionistas! We want to experience all different kinds of HF with all different kinds of people. The more diverse, the better. If you're looking to expand your HF horizons, you've come to the right place. (Goodreads, "Historical Fictionistas")

The group explicitly welcomes any fan of the genre, whether or not they might fit in with other group members in any other way. Similarly, the history group touts its "worldwide members from over 171 countries." For many of the genre groups on Goodreads, genre fandom is not only the leading feature tying members together, but the only feature that matters for group membership. At the same time as these genre-based groups bring fans of the same genre together, they also emphasize how diverse their members are by other metrics, solidifying the idea that genre functions on Goodreads as the glue that ties readers together in a fundamental way that transcends racial, ethnic, or gender, or other barriers. The site's genre-focused communities assume that genre is the most meaningful separation of the reading landscape and that these groups gather readers who are more similar than fans of other genres, despite any other differences they might have. In this way, Goodreads effects the transformation of genre from literary category to consumer category, used to separate readers into groups based on their affiliation with various genre communities in order to better target its recommendations and advertisements.

The more predictable a user's genre preferences, the more predictable their purchases. It is not surprising, then, that the Goodreads platform actively attempts to make its users' preferences more predictable. When I examined reader responses on Goodreads at scale, I expected more diverse responses to genre than I found. Despite the wide range of genres Goodreads allows on its platforms, within each genre, reviews became more similar in

vocabulary over the decade since Goodreads launched. Technology companies rely on machines to do work previously done by people, but it's not a simple replacement of human with machine: when algorithms take over the task of predicting readers' reactions, the drive to optimize predictability changes our reactions to make them more legible to algorithms.

From Genres to Clusters

In saying that Goodreads is a site focused on genre, it's necessary to clarify that genre is broadly defined on Goodreads. The site offers categories far more diverse than well-established genres like science fiction and horror. Goodreads' shelving system, which allows readers to place books on any number of individually-named shelves, creates a series of new literary categories to study and compare to established genres. These new categories are one of the advantages of studying reader response on Goodreads, since the site provides a more diverse landscape of genre than that found in anthologies or on bookstore shelves. As I will show, depicting genre not as a collection of static and well-established forms, but in terms of clusters of shelves that readers use for similar or different purposes, is a more authentic way to understand genre as it functions on Goodreads.

On Goodreads, every user has a personal "bookshelf" where they can create and name shelves that hold portions of their online book collection. Goodreads harvests these personal shelf names as crowdsourced genre tags for each book, taking the most popular shelves and featuring them alongside the book as a list of the book's genres:

Recommend It | Stats | Recent Status Updates

READERS ALSO ENJOYED

ISACASINOV OF FOUNDATION TRILOGY | THE WELLS SEVEN NOVELS | THE SCHIFF

GENRES

Science Fiction	2,639 users
Fiction	2,005 users
Humor	1,231 users
Fantasy	939 users
Classics	563 users

[See top shelves...](#)

The Ultimate Hitchhiker's Guide to the Galaxy (Hitchhiker's Guide to the Galaxy #1-5 & short story)
by Douglas Adams

★★★★★ 4.38 · Rating details · 240,121 Ratings · 4,319 Reviews

At last in paperback in one complete volume, here are the five novels from Douglas Adams's Hitchhiker series.

"The Hitchhiker's Guide to the Galaxy"

Seconds before the Earth is demolished for a galactic freeway, Arthur Dent is saved by Ford Prefect, a researcher for the revised Guide. Together they stick out their thumbs to the stars and begin a wild journey through time an ...more

GET A COPY

Kindle Store \$13.99 | Amazon | Stores ▾ | Libraries

Paperback, 815 pages
Published April 30th 2002 by Del Rey Books (first published January 17th 1996)

Original Title The Ultimate Hitchhiker's Guide: Five Complete Novels and One

Figure 3.3: The Goodreads book page for *The Ultimate Hitchhiker's Guide to the Galaxy*, with its list of genres highlighted.

Since Goodreads uses readers' personalized shelf names to label each book's genre, we might expect to see books labeled with idiosyncratic genre names that reflect the diversity of Goodreads' user base, but in fact, Goodreads' design ensures that each book's listed genre conforms to established genre conventions. The genre list displayed beside each book filters out shelves that Goodreads has determined not to be proper genre terms: for example, the page for *The Ultimate Hitchhiker's Guide to the Galaxy* neatly lists five genres, but a look at the most popular shelves for the book shows that the genre list skipped over shelves like "to-read," "currently-reading," and "favorites," and included the more formal "science-fiction" while excluding "sci-fi":

GENRES			
Science Fiction	2,601 users	to-read	74,990 people
Fiction	1,959 users	currently-reading	8,441 people
Humor	1,215 users	favorites	3,227 people
Fantasy	932 users	science-fiction	2,601 people
Classics	540 users	sci-fi	2,449 people
		fiction	1,959 people
		humor	1,215 people
		own	962 people
		fantasy	932 people
		owned	822 people
		classics	540 people

[See top shelves...](#)

Figure 3.4: *Hitchhiker's Guide* genre list from book page (left), compared to a list of the most popular shelves for the book (right).

By curating the list of genres that appears on the main book page to hide shelves that are less clearly genres, Goodreads encourages the equation of shelves with genres. By using an established system like genre as a critical piece of metadata about each book on its platform, Goodreads can provide more legible categories to its users who know what to expect from “science fiction” but wouldn’t know what a personalized category like “desserts” meant.

Despite the platform’s emphasis on established genre terms, Goodreads users are free to create shelves to help organize each book not just by genre, but by when they read it, where they bought it from, how it made them feel, or any other system of organization they can imagine. While genres remain some of the most popular shelves on Goodreads, the shelving system allows users to construct a much more expansive collection of genres if they so choose. In a Goodreads discussion board thread asking readers how they use shelves, one reader responds:

As for how I use shelves...my main purpose initially was keeping track of genre. Now, I read romance novels, and a lot of people will tell you a romance novel is just a romance novel. But for those who read the genre, they’ll tell you it breaks down to much more - paranormal romance, contemporary romance, historical romance, romantic suspense, series romance, etc. And I wanted to keep track of which genre books fell into. Then I broke some down further...historical-romance--regency, historical-romance--frontier, hq-intrigue, hq-blaze, etc (HQ=Harlequin).

Most of my shelves are these kinds of genre shelves. But I also have author shelves for several authors where I've read a lot of their books, and series shelves for series that have maybe 7 or more books, year read shelves, a library shelf. (jenjn79)

Goodreads seems to encourage an obsessive categorization in some users, prompting them to create shelves that initially include established genres but grow to cover every potential grouping of books they can think of, no matter how specific. In another discussion board response to a question about how people use shelves on Goodreads, Cait explains, "When I look at a book I particularly like, I look at how other people have shelved it to see what they have to say. If someone has an interesting-looking shelf, I'll often check out what else they've shelved that way to see if there's something else there that I'll want to read" (Cait). Readers like Cait use a book's list of assigned shelves as a way to navigate to similar books, a practice that Goodreads facilitates with easy-to-navigate pages for each book listing all of the shelves readers assigned to it, not just its official genres. Goodreads legitimizes even the most esoteric and personal of these shelves by assigning each of them an official landing page where users can view every book users have placed on that shelf, as well as discussion groups, videos, and quotes that other users have tagged as related to the shelf. The list of top shelves on every book's homepage is titled "Genres," not "Shelves," and the URL for every shelf landing page calls them "genres" as well. Despite Goodreads' declaration that these personal shelf names are genres, readers like jenjn79 and Cait use Goodreads to create systems of personalized categories extending well outside established genre categories. If we restricted our conception of genre on Goodreads only to include traditional genres, we would miss the diverse and surprising ways that readers use shelves for their own purposes.

Goodreads' profile viewing functionality also encourages readers to reach for shelf names outside of established genre categories to serve as a status signal to other readers. Digital

culture and critical race studies scholar Lisa Nakamura connects the past practice of surveying a host's bookshelves at a party to what Goodreads encourages with its publicly visible shelves, "inviting users to fill their virtual shelves with images of books for others to see" ("Words with Friends" 240). In keeping with Nakamura's assessment of the performative function of the books on a user's shelf, the names of these shelves are equally important, allowing users to broadcast their literary taste. As Goodreads user Lisa explains in a discussion board post, "I create the shelves for myself but there's a couple ones that I think would be fun/different for other members to see, and some day I'll add more of those. I haven't been nearly as creative or informative as many members but I do enjoy being able to name shelves whatever I want" (Lisa). Readers like Lisa use Goodreads shelves to perform an identity, recommend groups of books to others, and choose friends based on shared interests, all functions which rely on a flexible reading of Goodreads shelves rather than a restrictive vision of shelves that correspond to clearly-defined genres.

Unlike genres, which are popularly viewed as distinct and easily identifiable, shelves assert that there are many potential categories for a book, and that individuals will have different opinions on which categories fit best. Instead of trying to determine the correct genre for every book, Goodreads allows its users the freedom to assign books to one or more shelves according to their own personal system of categories. While publishers often limit a book to a single genre—consider the anthology, which chooses examples of "The Best Science Fiction of 2016" while allowing no room for science-fiction/romance or "sort of science fiction"—overlapping categories are possible, even probable, with shelves. Readers can use the shelving system to express contradictory reactions to a book rather than settling on a single category. For example, one Goodreads user put the romance novella *Sidebarred* onto both of the seemingly

contradictory shelves “cute” and “will forever haunt me” (Carissa). While literary critics recognize that genre is fluid and sometimes contradictory, publishers’ need for unambiguous marketing categories and libraries’ and bookstores’ physical need to place books in a single, findable location have too often forced readers to simplify their conceptions of genre to match these organizing authorities. On Goodreads, readers can disagree with others, and even themselves, about a book’s rightful category.

The fluidity of Goodreads shelves as well as their sheer magnitude—users have created more than 15,000 unique shelves—frustrates any attempt to map shelves onto an established system of genres. Goodreads shelves demand a new way of thinking about genre that can account for the unwieldy social tagging models common on online platforms like Goodreads, YouTube, and Flickr that are expanding the boundaries of genre to include not just formal traits but extra-textual components like format, affect, and a reader’s personal context. Recently, DH scholars have begun to represent genre using a method called clustering, a type of machine learning that attempts to automatically categorize items (in this case, books) based on their shared features, such as length, publication date, or use of certain words. Clustering works particularly well for genre because it forms categories that are not strict, but allow for overlapping and hazy boundaries as well as change over time. The most recent work uses clustering to represent genre based on the textual features of novels, finding that genres like detective fiction and Gothic fiction are distinguishable as coherent clusters of novels (Underwood; Wilkens). Since the clustering method is unsupervised, meaning there is no “right answer” for the clusters it forms, the strength of these studies is their ability to try out different assumptions about how genre works and see how the clusters change as a result. This project uses the flexibility of the clustering method not to represent clusters of literary texts as genres,

but to model how people use Goodreads shelves in ways that replicate existing genre categories, modify them, and extend beyond them to form new types of literary categories.

In discussing clusters of shelves rather than genres, this project asserts that genres aren't something that exists in the text of books, but instead something that is produced on Goodreads' platform and through readers' use of these categories, some recognizable and others provocatively new. A full account of genre would track its formation and revision by authors, publishers, readers, and a variety of other entities, but this study focuses specifically on genre as a negotiation between Goodreads and its users over the definition and literary membership of each genre, with the Goodreads platform offering a framework that users then fill in. This project's method of clustering shelves allows us to see patterns of similarity in shelving at a larger scale than we could through manual examination. The results of clustering don't depend on an a priori hypothesis about which shelves belong together—in other words, the method doesn't just reproduce established ideas of genre by feeding it existing categories and teaching it to replicate them. By forming clusters of shelves that share books in common rather than sharing a relationship to an existing genre, this method can find clusters of shelves that challenge the existing genre landscape. Its results are not prescriptive genre definitions, but fuzzy groupings that a researcher can grow or shrink to encompass one or many shelves in order to compare different hypotheses. Clustering based on Goodreads shelves reveals how readers are using genre terms day-to-day, rather than merely reinforcing a top-down prescription of genre created by experts or organizations.

I ran the clustering algorithm using shelving data from 18,239 random books on Goodreads, filtering out little-used shelves as well as very popular shelves that applied to many

different books like “to read” and “owned.”⁴³ The resulting plot, shown below, represents each shelf as a point, where a shelf’s distance from another shelf is determined by how many books users have placed on both of the shelves—concretely, the shelves located closest together share many books in common, and distant shelves rarely see a user put the same book on both of them. I’ve used colored boxes to highlight some of the most compelling clusters that emerged, giving them subjective titles to attempt to represent the dominant theme of the shelves that compose them:




Figure 3.5: A visualization of the result of clustering Goodreads shelves based on the books they have in common.

⁴³ I used the t-SNE algorithm to reduce the dimensionality of the data to visualize it in a two-dimensional scatterplot. The algorithm is especially well suited to high-dimensional, non-linear data like the shelving data in this application. It works by clustering points when it determines that they are both close together in high-dimensional space (as measured by Euclidean distance) and similar to each other (as measured by their conditional probability of similarity).

Many of the clusters that surfaced visibly were recognizable as genres, like scifi-fantasy, mystery, history, horror, romance, and drama, as well as the subgenre “cozy mystery.” The gray expanse in the center of the chart contains the shelves that shared the most books with other shelves, typically indicating more general shelves like “favorites” and “library books.” In the chart above, I’ve highlighted shelves that fit an apparent theme, but these judgments are subjective: the clusters are merely a tool to begin to examine what kinds of shared traits can cause shelves to function similarly for users (and thus appear close together on the chart) and how we might determine membership in each cluster of shelves. As the chart makes clear, many of the clusters do not have clear boundaries. For example, when we look more closely at the scifi-fantasy cluster, we can see two shelves marked with gray points as opposed to orange:



Figure 3.6: Close-up of the scifi-fantasy cluster.

The two shelves on the edges of the scifi-fantasy cluster, “to read mystery” and “cold war,” emphasize the porous boundaries of clusters as compared to genres. While scifi-fantasy and mystery are typically considered to be distinct genres, the clusters that most resembled these

genres are located close together, meaning they share books in common, and the “to read mystery” shelf appears right at the edge of the sci-fi-fantasy cluster. The proximity of the “cold war” shelf, too, demonstrates how tropes and settings can become incorporated into genre definitions, or at least lurk at their edges. By visualizing sci-fi-fantasy as a cluster of similarly-used shelves rather than a genre, we can see these blurry boundaries and adjacent themes clearly. Within each cluster, too, surprising shelves often appear. Some of the clusters added a telling shelf or two to the mix. For example, the drama cluster highlights the cultural centrality of Shakespeare in drama with the addition of the “Shakespeare” shelf to a grouping of shelves denoting the dramatic mode more generally:




Figure 3.7: Close-up of the drama cluster.

The drama cluster was the most closely-grouped cluster that I observed, indicating that readers had little disagreement over which books belonged on these shelves. In the same way as we examined the sci-fi-fantasy cluster’s edges, we can also examine the dispersal of a cluster to gain insight into how readers use it, comparing compact clusters like drama with widely spread clusters like romance, which seems to inspire much more diverse uses from readers. We can also focus on sections of the chart where clusters are less clear, and diverse shelves mix in evocative ways. The area surrounding the horror cluster is one such place, with a variety of related and seemingly unrelated shelves coming together based on readers’ similar usage of them:



Figure 3.8: Close-up of the horror cluster.

I’ve highlighted the shelves “horror,” “horror-thriller,” “ghosts,” and “ghost stories” as one cluster, but nearby we see the shelves “anthologies” and “shorts,” reflecting the format of many horror pieces. Between those two shelves, we see “not for me,” a revealing shelf that reflects the polarizing nature of the horror genre. We also see “angels” nearby, a surprising addition to a genre typically seen as frightening. In fact, the shelf “terror” is relatively far from the horror cluster, instead lying closer to the shelf “dark fantasy.” From this placement, we might argue that horror as it functions on Goodreads isn’t as much about inspiring fear as it is about tropes like ghosts and angels as well as an abbreviated format, characteristics that happen to be easier to identify and agree upon than more nebulous categories like terror and dark fantasy. This close-up of the horror cluster also highlights the peculiarity of shelves on Goodreads, with interloping shelves like “need to get” appearing nearby the ghost-related shelves and “favourites” appearing to cluster with dystopian shelves. Several other clusters included these less meaningful shelves that host a wide variety of different books, and whose names have less to do with the book itself than its reader’s individual context, like “good books,” “read in 2013,” and “unsorted.” These types of shelves made up a majority of the gray points in the scatterplot but added little insight

into how readers use shelves on Goodreads. One of the strengths of the clustering method is that it doesn't produce clear lists of shelves that map cleanly onto established genres. Rather, as the horror cluster demonstrates, clusters are simply visual groupings of shelves that one can read narrowly or expansively, analyzing the boundary cases and surprising inclusions for what they tell us about how people are using shelves in practice rather than according to theoretical expectations of their behavior. These clusters better reflect how genre works on Goodreads than does the concept of genre, which in popular usage assumes clear boundaries and stable definitions that simply don't exist on Goodreads. The idea of a cluster, which is loosely defined and has blurry edges, is a more fitting way to think about genre in a social reading context. Clusters are not based on predefined notions of genre curated by literary critics, but on how readers actually use them. As much as I might expect the "horror" and "terror" shelves to be related, their relative separation is a real effect on Goodreads. By visualizing and analyzing these clusters, we can create a more accurate picture of how readers use shelves.

In addition to clusters that echo and modify established genres, a variety of non-genre-focused clusters emerged, including:

- clusters based on a book's format, like comics and audiobooks
- clusters based on tropes like the paranormal, including shelves related to vampires and werewolves
- clusters based on the age of the book's audience, including shelves for young adult literature and children's books
- clusters based on subjects like the humanities and food

My argument is not that we should use these clusters instead of established genres, or that they are in some way superior to genres, but that these clusters exist and readers are using them, and we should study them with the same attention we give well-established genres.

By examining the variety of shelves that the clustering method groups together, we can see the massive expansion of the limits of genre that has taken place on Goodreads. Not only do readers use well-established genres like mystery and romance, but they use clusters of shelves based on a wide variety of factors from tropes to personal identity. On Goodreads, genre is limited not by professional conventions but by the imaginations of its users. But as we will see, the great variety of shelf clusters on Goodreads does not mean it's a place where conceptions of genre become more diverse.

Predicting Genre

In examining the clusters of shelves that form through readers' activity on Goodreads, we saw that genre is not the only common way of classifying books. Clusters of shelves sometimes map onto established genres like romance, but they have a much wider range of types than just genres and subgenres, including shelves based on a book's format, tropes, audience, and subject. Do readers use these clusters of shelves differently from how they use clusters based on genre? Is there something uniquely coherent about genre as a way to tie together a cluster of shelves that makes it more cohesive than other types of clusters, or is genre a more fluid category, able to hold more diverse perspectives and reactions to books?

To answer these questions, we need some way to measure the cohesiveness of readers' responses to different clusters so that we can compare how effectively clusters tie together

groups of books. I decided to measure the similarity of readers' responses to books in a given cluster: do most readers respond in the same way, using similar language, or do readers use a varied vocabulary to talk about the cluster? To frame this question in another way, how predictable are readers' responses to a given cluster? While predictability may not be the measure a literary scholar might use to gauge readers' responses to a genre, it's an appropriate measure for responses on Goodreads, since predictability is Goodreads' greatest purpose as a recommendation and advertising platform. If we can measure how predictable reviews of various clusters of shelves are on Goodreads, we can see how successful Goodreads has been in shaping readers' responses into predictable forms that further its goals of accurate organization, recommendation, and online community founded on consumer categories like genre. By determining how well Goodreads has achieved its own metrics of success, we pave the way for more detailed conclusions about the precise ways various reading communities respond to specific genre categories and clusters.

My simple question was, how well can you teach a computer to recognize if a review is discussing a specific cluster of shelves? To measure the predictability of readers' responses to a cluster, I created a computer program that uses a machine learning algorithm to "learn" the words that most and least characterize reviews from each cluster. The program looks for the most useful words it can find to predict the cluster a review is discussing, whether or not those words accord with a typical reader's understanding of the topic or genre of that cluster. Machine learning as a method is not capable of producing a definition of each cluster, but it does give us a measure to gauge whether the language reviewers use resembles the language of other reviewers discussing the same cluster. It treats each cluster as a shared vocabulary between readers, and accordingly reveals how well-defined the vocabulary is for each cluster of reviews. A method

which measures clusters in terms of the predictability of reader responses to that cluster best reflects the new social reading model of sites like Goodreads, in which online booksellers care most about literary categories (whether based on a topic, genre, or subgenre) being something people will react in similar ways to in their reviews, and that will please a predefined community—or at least not surprise them.

Several previous digital humanities projects have used machine learning to predict the genre a book belongs to, but none have attempted to predict the genre a review is discussing. For example, the Stanford Literary Lab trained a model to recognize the Gothic novel and the bildungsroman, finding that it was actually easier to define a genre by what it's not (Allison, Heuser, et al. 18). Recently, Ted Underwood compared the cohesiveness of several different sources of genre corpora to see how well a machine learning model could recognize individual books as members of genres like science fiction and detective fiction (Underwood). These projects rely on “correct” genre assignments for books that are consistent and detectable in a book's language. In contrast, the method I use in this project doesn't try to define or teach a computer to recognize genre itself, but the conversation around genre. This method can detect similar responses to books even if their content is drastically different. Furthermore, this project follows Goodreads' lead in defining genre more expansively than most digital humanities projects, and considering a wider range of literary and para-literary categories based on the clusters of shelves readers use on the site. Projects that look for genre in the content of books run the risk of finding signals that are not related to the genre under study but instead incidental, like in Underwood's acknowledgment that he might have been tracking the language of crime rather than detective fiction as he intended. By instead studying reader reactions that are intentionally

labeled with a specific genre or another type of shelf name, as reviews on Goodreads are, we can be more confident that what we are tracking is readers' responses to that shelf.

My corpus contained every accessible English-language Goodreads review of 28,681 randomly selected books,⁴⁴ totaling 262,900 reviews dating from the site's launch in 2007 to 2016:




Figure 3.9: Number of reviews in dataset broken down by year.

Over 60% of the books had only one review in the dataset.⁴⁵ *Twilight* had the most reviews in the dataset of any single book, with 418 reviews. Rather than taking the same number of reviews from each book, this corpus provides a representative cross-section of reviewing activity on Goodreads, with heavier representation for books with more reviews. A different corpus (of only the most popular or the most canonical books, for example) would give a different set of shelf clusters to look into, but my random sample provides a broad perspective on the overall

⁴⁴ To randomly select books for this study, I used Goodreads' built-in random book finder: <https://www.goodreads.com/book/random>. While some books had thousands of reviews, the Goodreads interface only allows readers to view several hundred results, so I collected every review that is accessible via the interface, reflecting the reviews that Goodreads users would conceivably be able to read.

⁴⁵ The most likely reason for such a high proportion of single-review books is a quirk in Goodreads' database that treats different editions of the same book as separate books, leading to many one-review editions of the same book as a reviewer adds their review under their specific edition.

reviewing activity of Goodreads users on books both popular and unknown. For this experiment, I used only the linguistic content of reviews, rather than images, emojis, the number of likes or comments the review received, profile information about the reviewer like gender or the size of their friend network, or any information about the sentiment or reading level of the review. Limiting the analysis to just the text of the review is appropriate for this project tracking the discourse around clusters of shelves, but another study could examine these extra-textual factors to see if they function differently than the raw text of reviews.

I took several steps to prepare the reviews for machine reading. Many reviews mention a book's shelf or genre explicitly, or use names or titles that tie them closely to a single shelf or genre. By leaving these clues in reviews, the program would decide a book's cluster based on these overwhelmingly dominant signals, but by removing shelf and genre terms, common book titles, author names, and character names from reviews, the program must find less obvious words to distinguish each cluster. While this deletion almost certainly eliminated some interesting information, such as which books or authors are touchstones for a genre like Asimov is for science fiction, removing these features ensured that the program was judging predictability based on a truly common discourse, rather than becoming simply a metric for how often reviewers use a shelf's name or how dominant a particularly popular character was in reviews. I also removed "stopwords," or common words like "and" and "or" that typically don't tell us much about the writer's meaning. Even after removing these distinctive words, the number of total words of reviews was in the millions, so I narrowed the remaining words based on how statistically distinctive they are for their cluster.⁴⁶ After this process, all that remained in

⁴⁶ The formula I used for distinctiveness is called TF-IDF, which takes a ratio of the number of times reviews from a given shelf use that word divided by the total number of times that word appears in all reviews from all shelves. The result is a measure of that word's importance to the given shelf. I could then compare the importance of each word within each shelf and keep only the most important words for each shelf.

the reviews were words that not only appear often when discussing a given cluster, but rarely appear when discussing other clusters.

Using this highly edited set of reviews, I trained a machine learning model to find the words that best distinguish a given cluster from reviews of all other clusters.⁴⁷ Practically, the model receives a set of reviews labeled either “Romance” or “Not Romance,” it looks for keywords that appear most and least in the “Romance” reviews, and it builds a set of words it can use to predict future reviews using the same rules. After building this model of predictive words, you can then give it new reviews and see how well it predicts their cluster.

⁴⁷ I used a linear SVC model with 10-fold cross-validation. This model attained the most accurate results for Goodreads reviews of the 7 algorithms I tried.

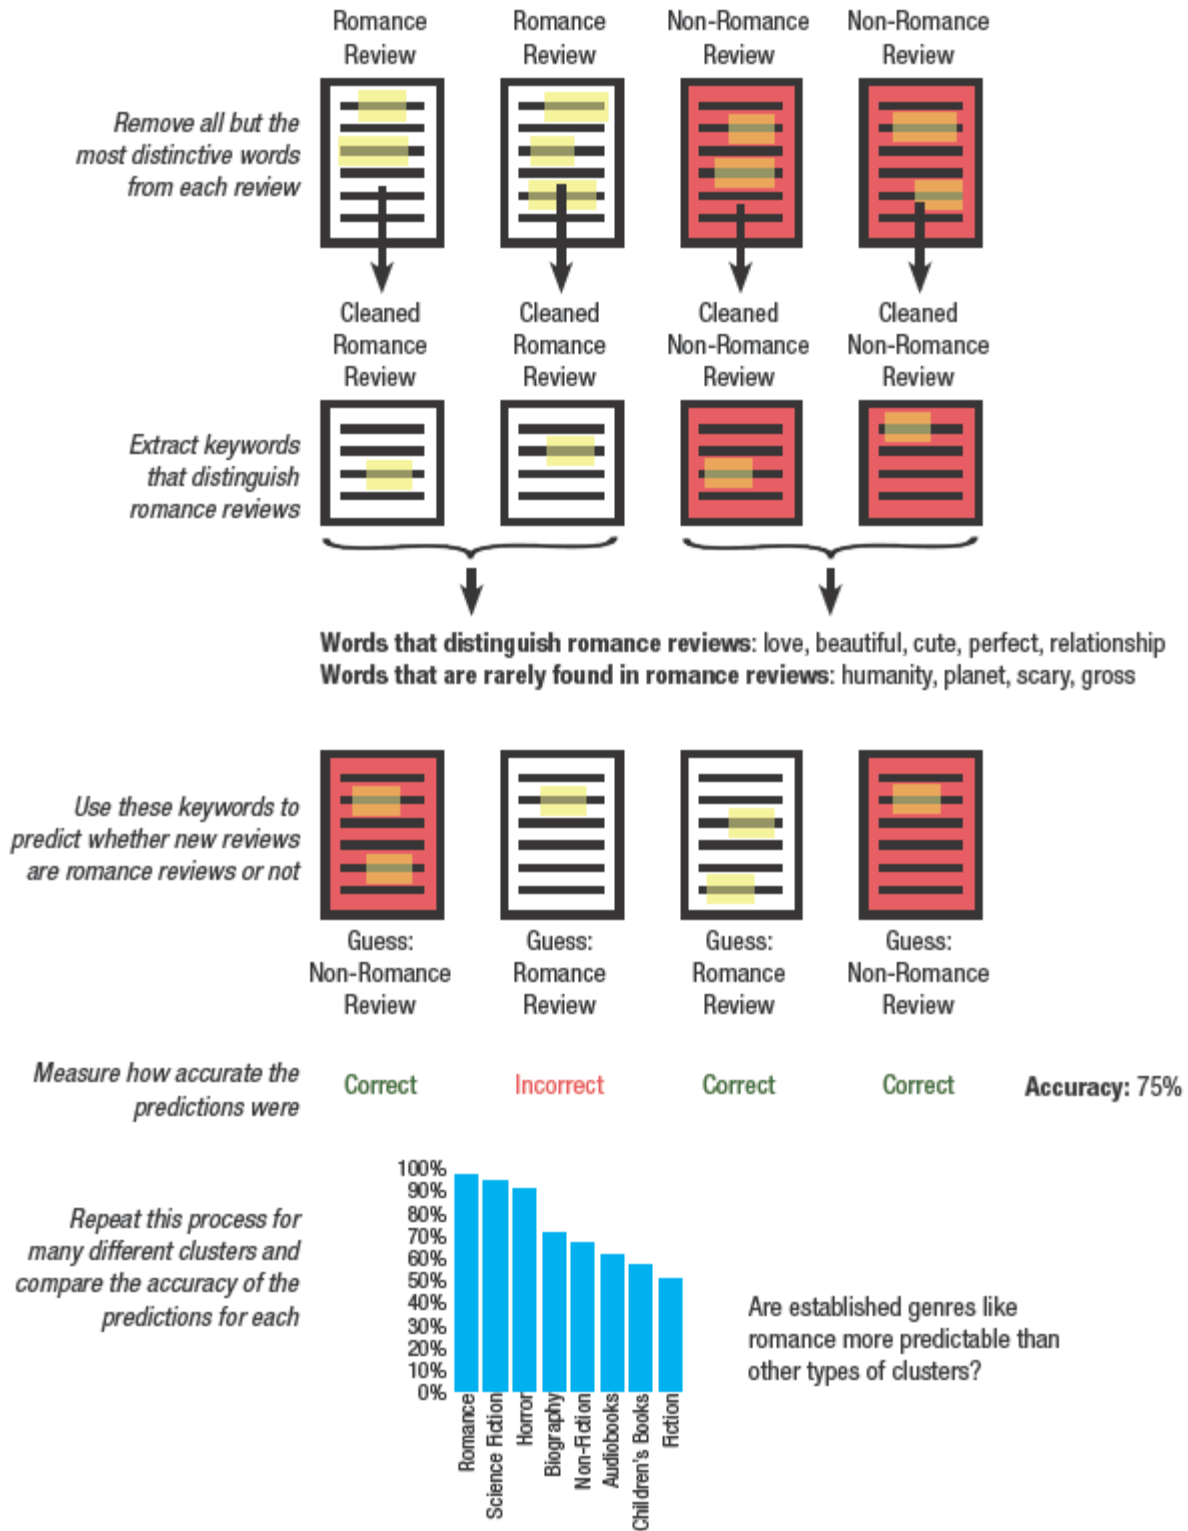


Figure 3.10: Diagram of the prediction method.

The model produces an accuracy score from 0% to 100% that reflects how easy it was for the model to distinguish romance reviews from other reviews.⁴⁸ This score gives a rough metric of how clearly separated the cluster is from others. A high score means that the words in that cluster's reviews are distinctive of that cluster. For the clusters I looked at, accuracy scores ranged from 45% to 95%, or from slightly less accurate than a coin flip to almost always correct, but these raw scores are less meaningful on their own than as a comparison to determine which clusters are more predictable than others. A different classification algorithm might make the scores as a whole 5% more or less accurate, but if reviews of science fiction are 20% more predictable than reviews of romance, no matter which algorithm you use, we can see that readers are using more coherent language to discuss science fiction.

But what clues help the computer classifier recognize a cluster? By looking at the words the classifier uses to predict, for example, whether a review is discussing a science fiction book, you can see that the method is not so different from how humans might classify the genre. If the word “space,” “Mars,” “alien,” “planet,” or “robot” appears, the classifier is likely to predict the review is discussing science fiction. Interestingly, the words “girl,” “family,” “love,” and “wife” prompted the classifier to say a review was likely *not* science fiction. This result doesn't mean that science fiction has a negative view of girls or families, but that they simply aren't a part of reviewers' discourse for the genre: science fiction fails to prompt readers to think about the idea of family or discuss the presence or absence of girls—that conversation simply never arises. Readers describe science fiction heroes as “men,” “humans,” and “characters,” but rarely in terms of their family ties, even to note that they have no family. Similarly reasonable results of the classifier mark horror as “scary” and “creepy” but not “funny” or “cute,” and note that

⁴⁸ This percentage reflects the model's F1 score, the weighted average of its precision and recall for classifications of the given genre.

romance discusses “love,” “relationships,” and “sex” but not “crime” or “monsters.” Clearly, the classifier is good at picking out words that many people consider markers of the discussion around these established genres.

Are certain clusters easier for the classifier to recognize than others? Using the diverse range of categories we find from clustering shelves together, I ran the classifier on each cluster to see how well it could predict a review’s cluster from the reviewer’s language. The chart below shows the predictability scores for each cluster I tested, roughly grouped into types like “established genre clusters,” “subgenre clusters,” and “audience clusters.” The chart includes every cluster in my corpus with at least 100 reviews, so this is a selection of the more popular clusters on Goodreads:




Figure 3.11: Comparing the predictability of reviews from different types of clusters

Format Clusters: fiction, non-fiction, kindle, ebook, audiobook, comics & graphic novels

Personal Use Clusters: to read, currently reading, owned, library, borrowed, book club, reviewed, favorites, abandoned

Audience Clusters: chick lit, adult, young adult, children's books

Status Clusters: literary classics, mainstream, cult classics, fluff

Response Clusters: made me cry, guilty pleasures, boring, thought-provoking, disappointing, cute, sweet, page turner, creepy, disturbing, depressing

Subgenre Clusters: urban fantasy, cozy mystery, ww2 history, erotic romance, crime mysteries, historical fiction, cyberpunk, new weird, space opera

Trope Clusters: post-apocalyptic/dystopian, paranormal, vampires, time travel, zombies, aliens

Established Genre Clusters: biography, science fiction, horror, romance, mystery, history

Each point on the chart shows the classifier's accuracy in identifying reviews from that cluster. For clusters with scores near 50%, like the "fiction" shelf, the classifier did no better than flipping a coin. Higher scores showed that reviews from that cluster share a common vocabulary, allowing the classifier to make better predictions based on the words reviewers use to talk about that cluster. The chart shows that reviews of books in established genre clusters, like romance and science fiction, are the most predictable, followed by shelves centered on a specific textual trope like vampires or zombies, and then subgenre clusters like erotic romance and World War II history. These results are not surprising, since teachers and marketers train reviewers from a young age in the language they should use to talk about established genres and subgenres, and tropes come with a clear topic and set vocabulary related to that topic. The least predictable types of clusters are those based on external features like a book's format—any book can be an ebook or audiobook, so the label says little about how to respond to the book—and a reader's personal relation to the book, such as whether they own it or borrowed it from the library. Clusters that explicitly identify a particular response from their readers vary in their predictability: some, like "disappointing," prompt very predictable responses, while others, like "page-turner," result in more varied responses.

It's important to remember that this chart does not reflect the predictability of a cluster, but rather the predictability of *responses* to that cluster. Digital humanities scholars have found patterns in the conformity of literature in different genres, but we wouldn't necessarily expect these results for shelf clusters to follow the same patterns. Instead, we would expect less similarity in reviews of heterogeneous clusters that have nothing to do with the content of the book, like "kindle," or those that are based on an individual's personal circumstances and draw together many different books and responses, like "currently reading." We would expect the

most coherent clusters to be those that imply a certain response, like “awesome” or “made me cry,” since reviews would have a more similar (and thus more predictable) response. As it turns out, clusters based on established genres prompt even more cohesive responses than clusters that imply a specific response, which is surprising: why would responses to books that made readers cry be less similar than responses to science fiction, which we might expect to inspire a wider range of possible reactions?

The Science Fiction Cluster

The science fiction genre is well-established, particularly on Goodreads. By clicking a link on the homepage, users can view the landing page for the science fiction genre, which defines science fiction, highlights new releases and popular books in the genre, and links to pages for related subgenres, science fiction book clubs, discussion groups, and giveaways.

Genres > Fiction >

Science Fiction

Science fiction (abbreviated SF or sci-fi with varying punctuation and capitalization) is a broad genre of [fiction](#) that often involves speculations based on current or future science or technology. Science fiction is found in books, art, television, films, games, theatre, and other media. In organizational or marketing contexts, science fiction can be synonymous with the broader definition of [speculative fiction](#), encompassing creative works incorporating imaginative elements not found in contemporary reality; this includes [fantasy](#), [horror](#) and related genres.

Although the two genres are often co [...more](#)

NEW RELEASES TAGGED "SCIENCE FICTION"

RELATED GENRES

Aliens	Near Future
Alternate History	Planetary Romance
Alternate Universe	Robots
Apocalyptic	Space Opéra
Cyberpunk	Space Opera
Dying Earth	Steampunk
Dystopia	Sword and Planet
Fiction	Time Travel
Hard Science Fiction	Utopia
Military Science Fiction	Utopia

FEATURED GOODREADS AUTHORS TAKING QUESTION

Robert Jackson Bennett
Author of City of Stairs

Robert Jackson Bennett is a two-time award winner of the Shirley Jackson Award for Best Novel, an Edgar Award winner for Best Paperback Original, and is also the 2010 recipient of the Sydney J Bounds Award for Best Newcomer, and a Philip K Dick Award Citation of Excellence. His fifth novel, City of Stairs, is in stores now.

Ask a Question

Figure 3.12: Goodreads science fiction landing page (goodreads.com/genres/science-fiction).

At a glance, it seems like Goodreads is a place where people can experience a wide range of perspectives on the genre and discuss their diverse opinions with other readers. But in fact, science fiction reviews are becoming more predictable over time:




Figure 3.13: Predictability of science fiction reviews on Goodreads.

In the decade since Goodreads launched, reviews of science fiction books have become 17% more predictable. In practical terms, this means that science fiction reviews developed over time

a set of keywords reviewers use to talk about the genre that appear often, but rarely appear in reviews of other genres. Science fiction reviews, then, are becoming more similar in language. Since people writing reviews on Goodreads are likely to know already what science fiction is, we might have expected the reviews to remain at a flat level of predictability, showing that reviewers maintain a relatively consistent idea of what science fiction is and write reviews using language that varies little from year to year. Or perhaps we might expect reviews of the genre to become less predictable as people use Goodreads to find new examples of a wide range of science fiction, read reviews from people much different from them, and ultimately broaden the vocabulary they associate with the genre. Since instead, predictability increases, this is evidence that an effect of the Goodreads platform is to encourage science fiction reviews to become more rhetorically similar. To narrow down what might be causing this increased predictability, we can look at the words the model found most useful in predicting a review as science fiction.

Throughout the decade, abstract language characterizes science fiction reviews: words like “future,” “space,” “universe,” “world,” “technology,” “Earth,” “society,” “humanity,” and “time.” Science fiction tropes also mark its reviews, like “alien” and “robot.” A review is unlikely to be about science fiction if it includes language relating to families or relationships like “family,” “parents,” “mother,” “wife,” “husband,” “sister,” “love,” or “relationship.” Another constant is the absence of the “vampire” trope: science fiction could be defined, according to its reviews, as the genre without vampires. It appears that science fiction reviewers are especially careful not to use the term in their reviews. This result could arise from an actual lack of vampires in the content of science fiction books in relation to other genres, or perhaps instead from fans of science fiction—a genre that frequently finds itself defending its legitimacy—distancing their genre from the even more culturally maligned genre of vampire

fiction. The lack of familial and vampire terms in science fiction reviews is merely an observation, insufficient to determine whether families and vampires are disappearing at the point of science fiction book production or reception. But by tracking the types of words that distinguish science fiction reviews over time, we can begin to pinpoint what might be the cause of the genre's increasingly predictable language.

Subgenres are easy to find and use on Goodreads, since you can place a book on shelves for both science fiction and any other subgenre without having to choose between the two. This flexibility also allows readers to sidestep debates over competing terms and definitions for a genre, such as the debate between “science fiction” versus “speculative fiction,” since readers can simply choose the term they prefer (or assign a book both genres) without needing to argue their case for the correct term to an authority or another reader. Starting in 2011, names of science fiction subgenres like “classic,” “military,” and “dystopian” begin to mark reviews clearly as science fiction. The main page for the science fiction genre includes links to more specific subgenres, many of which seem to echo review language: “Aliens,” “Dystopia,” “Military Science Fiction,” “Robots,” and “Time Travel.” It is possible that these explicitly listed subgenres are guiding reviewers to use this language in their reviews. By codifying science fiction tropes like aliens and robots as subgenres, Goodreads encourages its users to reference them in science fiction reviews. This is one way that Goodreads, through the design of its platform, might influence the language of reviews to become more predictable.

If the Goodreads platform encourages reviews of science fiction to become more linguistically consistent by explicitly featuring specific subgenres, we might expect that these subgenres might become a place where review language can grow more diverse. That is, if readers feel constrained in how they can talk about science fiction, they might use science fiction

subgenres as a place to share the responses that don't fit the more restrictive mold of science fiction proper. However, for each of the science fiction subgenre shelves I looked at, its predictability—like that of science fiction reviews as a whole—increased over time:




Figure 3.14: Predictability of science fiction-related subgenre reviews on Goodreads.

Like the broader science fiction genre, science fiction subgenres are growing more predictable, by margins ranging from 5% to 32%. These results also contradict the idea that subgenres begin

and remain predictable because they have a specific trope or style tying them together—instead, we again see an early period when reviews use greater language diversity followed by increasingly predictable language.

The subgenre shelf that experienced the largest increase in predictability was dystopian fiction. For the first two years that readers used the shelf on Goodreads, it was barely more predictable than flipping a coin: the words that most distinguished the shelf were unspecific, personal reaction words like “amazing,” “whoa,” and “ok.” These words fade by 2010, when the shelf becomes increasingly distinguished by tropes like “zombie” in 2009 and “government” in 2011. Like with the vampire trope in the science fiction genre, it is unclear if the increase in these trope terms is due to a change in production—authors writing more zombie books—or reception—readers more interested in discussing zombies—or some combination of the two, but in either case, we can learn from the trend. By using a trope term like “zombies” much more frequently than other common terms that were equally popular in past years but less clearly summed up as a trope, readers show an affinity for standardized vocabulary in their reviews. While zombies have become familiar figures in reviews since 2009, they are no more common than other zombie-related terms like death or viruses, but reviewers seize on the word in their reviews as representative of a discourse that distinguishes dystopian fiction.

While tropes become positive indicators that a review is discussing dystopian fiction, the names of other genres begin to mark reviews that are *not* discussing a dystopia: starting in 2009, “mystery” becomes a signal that a review is not discussing a dystopia: in 2010, “history”; in 2013, “cyberpunk”; in 2014, “fantasy”; and in 2016, “horror.” It seems that as dystopia becomes increasingly recognizable in the language of reviews, it is its difference from other genres that drives this consolidation. Rather than the subgenre becoming a proving ground for more diverse

conceptions of genre on Goodreads, dystopian readers are united in their lack of discussion of other genres—the discourse of dystopia is increasingly isolated. For each of the science fiction subgenres I examined, a similar vocabulary (or absent vocabulary) gradually emerged to set reviews of that subgenre apart from the rest.

Given that science fiction reviewers use a variety of diverse subgenre categories, we might have expected that these subgenres would have diverse discourses. But more subgenres do not make for more diverse genres—only better organized and partitioned genres. In fact, the very proliferation of these subgenres makes it increasingly difficult to see books outside of this categorical structure, allowing established genres to become even more entrenched. When Netflix recommends increasingly specific subgenres like “irreverent dark TV comedies,” it reassures us of the accuracy of genre-based recommendation so that we are less and less likely to venture outside of the categories they recommend. By interrogating not just genre terms like science fiction, but also their context of subgenres and trope-based adjacent genres, we can see how the increasing subdivision of the science fiction genre on Goodreads creates even more restricted discourse around these genres.

If readers are using genre as the primary basis for recommending and selecting books, a predictable discourse helps them know what to expect: a science fiction review will reliably discuss which tropes the book uses and which abstract themes it addresses without venturing into the more concrete topics of families and relationships. Goodreads benefits from a predictable discourse because it means their recommendation algorithms work well, and it encourages tight-knit communities of people who all generally agree and are similar. But as a result, readers are less likely to be exposed to surprising ideas or to broaden the range of possible responses they associate with the genre.

The Creepy Cluster

As well-established genres like science fiction and its related subgenres became less diverse on Goodreads, what happens to clusters that are not explicitly generic? One non-traditional cluster that is unique to Goodreads relates to an affective response in readers: the “creepy” genre. We might have expected the clustering method to include it within the horror genre, but it stands distinct as its own cluster and includes books that are not typically considered horror such as *Gone Girl* and *Miss Peregrine’s Home for Peculiar Children*. Is “creepy” simply a cluster identifying slightly less horrifying horror?

Though Goodreads launched in 2007, users didn’t begin using the shelf “creepy” until 2010. While there is a lot of overlap between creepy and established genres like horror, fantasy, and mystery, the creepy genre has a distinct set of associated books, videos, quotes, and linked book lists on its Goodreads landing page. Classic novels like *Dracula*, *Frankenstein*, *The Strange Case of Dr. Jekyll and Mr. Hyde*, and *The Turn of the Screw* are overwhelmingly shelved as “horror” and very rarely “creepy,” suggesting that either their age or classroom context detracts from their creepy affect. Books that became movies are more likely to be creepy—the five books most commonly shelved as “creepy” are *The Shining*, *Carrie*, *Unwind*, *Pet Sematary*, and *Warm Bodies*, all of which have movie versions. Children’s and young adult books are also disproportionately creepy, including *Doll Bones*, *Miss Peregrine’s Home for Peculiar Children*, *Coraline*, and *The Graveyard Book*.

Why do people create and maintain a shelf for creepy books when they could use horror instead? The label “creepy” allows readers to include books that don’t match the traditional definition of horror, as well as exclude books that are traditionally considered horror but don’t creep them out. Response-related shelves like creepy, sad, and gross allow readers an alternative

system of categorization that isn't answerable to experts. No one can argue that a particular book isn't really creepy like they can with horror—if they felt goosebumps as they read it, it's creepy. These response genres give readers an outlet for affective responses that are often elided from top-down definitions of genre while they validate affective responses as meaningful. Since readers that label a book “creepy” all claim the same affective response, we might expect reviews of creepy books to use similar language. But again, we see an initial diversity of response followed by a consolidation of language that makes creepy reviews more predictable:




Figure 3.15: Predictability of creepy reviews on Goodreads.

Throughout the decade, creepy reviews disproportionately use the words “scary,” “horror,” “gothic,” and “ghost.” In the period before 2011 when the cluster’s language was more diverse, the word “movie” characterizes the cluster. As reviewers began to reference movies less, the creepy cluster became more predictable. Reviewers began to avoid the words “cry,” “crying,” “sad,” and “heart” in creepy reviews in 2010, and starting in 2012, reviewers seem to settle on the terms “spooky,” “strange,” and “weird” to describe the cluster. Over time, reviewers of the creepy cluster seem to develop a distinct vocabulary to describe it, commonly using the language of strangeness or weirdness above and beyond what exists in horror. The consistent use of this vocabulary makes reviews of the cluster less diverse and more predictable. New, affective clusters on Goodreads like the creepy cluster might have served as a place for readers to stretch the boundaries of genre to include a broader selection of terms to describe the reception of creepy books. Instead, just as established genres like science fiction begin to consolidate over the decade that Goodreads has operated, readers who use the creepy cluster develop a coherent language of related terms that they predictably use to describe books in the cluster. This effect may seem to undermine Goodreads’ usefulness as a site for sharing diverse viewpoints and encountering a wide range of new books, but in fact, a coherent language for each product category benefits Goodreads by making it easier to predict users’ preferences. If two users have different definitions and expectations for the same genre, Goodreads’ recommendation and advertisement targeting algorithms will be less effective—they work best when categories are well-defined and when people act in predictable ways. When Goodreads’ goal of showcasing diverse viewpoints interferes with its motivation as an advertising company, profit wins. The effect is a narrower discourse of genre on Goodreads.

Predicting Genre Over Time

We have seen that both science fiction and creepy books become more predictable over time, but is this true of all clusters on Goodreads? We can do the same type of analysis for many other clusters to see how the predictability of that cluster's reviews changes over time. The chart below shows the predictability scores of established genres and other clusters on Goodreads (the two charts on the top), contrasted with the scores of reviews from various professional book review websites (bottom left) and from academic articles discussing different genres (bottom right).⁴⁹

⁴⁹ The professional book review platforms included *The Austin Chronicle*, *Boston Review*, *Kirkus*, *The Los Angeles Review of Books*, *The New York Times*, *Publishers Weekly*, *USA Today*, *The Wall Street Journal*, and *The Washington Independent Review of Books*. The academic articles were compiled from searches for genre terms on Project Muse, from a variety of publications. For each genre, whether from Goodreads, a professional publication, or academic articles, I included every genre for which my corpus included at least 100 reviews. Results are cumulative, so results for 2010 include reviews from years before 2010, beginning with the first year with 100 reviews. This is because readers experience reviews not one year at a time, but in the context of all the reviews that have been published previously as well—all of these reviews combine to form a perception of the genre.




Figure 3.16: Comparison of the predictability of reviews on Goodreads, professional, and academic platforms.

The top two charts show that, as a whole, reviews on Goodreads grow more predictable over time, while professional and academic reviews stay equally easy to predict. With traditional

machine learning problems, such as training an algorithm to recognize images that contain a cat, a trend of increasing predictability as you give the system more data would make sense. In the case of recognizing clusters, though, since there is no objective “cat” for the model to find, we should expect the opposite trend, since as the classifier adds more reviews each year, it will face more diverse responses and exceptions to the rule of which words characterize a cluster. But what happens is the opposite: for both established genres and unconventional clusters on Goodreads, as people write more reviews, it becomes easier to predict which cluster those reviews belong to—the new reviews pick up on the language that distinguishes a given cluster and continue to use it. The difference between Goodreads and the professional and academic platforms shows that this increase in predictability isn’t a general property of all book reviewing communities, but instead a characteristic distinctive of Goodreads. Again, this is surprising because Goodreads prides itself as a place where millions of people can express their diverse opinions of books, unlike professional critics constrained by the demands of editors, publishing standards, and the market. Instead, the Goodreads platform consolidates and conforms the discussion for established genres and unconventional clusters alike.

But could something outside of Goodreads have caused this increase, like a major site redesign or an influx of users? If that were the case, we might expect to see a more abrupt shift upward rather than the gradual but irregular trend we see across a wide range of clusters. For example, when Amazon bought Goodreads in 2013, we might have expected to see a change in review predictability resulting from changes Amazon makes to the site’s design, marketing, or recommendation algorithms. Amazon’s purchase of Goodreads raised alarms for many Goodreads users, but a common reassurance was that Amazon promised to change little about the successful social reading platform. Hugh Howey, an author who was hugely successful self-

publishing through Amazon, predicted: “My guess is that we won’t see many changes at all. I’m betting that the real acquisition here is all the data behind the scenes. The algorithms that tell me what to buy (and almost always nail it) are going to get better. The social networks that feed my reading habit are going to get stronger” (Flood). Howey predicts improvements in Goodreads’ two most core selling points, its recommendation and community features. If Howey were correct, we would expect to see reviews become more predictable starting in 2013 as Amazon implemented changes to the reviewing process in order to make their recommendations more accurate. While we do see reviews increasing in predictability, it is a slow increase over Goodreads’ decade-long tenure rather than a drastic improvement that might have resulted from a single event.

Cultural theorists have long known that genre isn’t static. For example, film critic Rick Altman describes the creation of the horror genre in film not as a pre-planned and intentional break from the past, but as an outgrowth of several previous genres that we can only recognize retrospectively (Altman). Since a reader’s response to a genre is a significant factor in how it is defined, in addition to its formal features, a book can be one genre in 2000 and a different genre in 2010 after readers’ perceptions and social circumstances change. Thus projects that study the text of novels to find gradual shifts in the definitions and limits of genre are bound to miss the concurrent—or competing—shifts in how readers receive these genres. For example, literary critic Franco Moretti notices a generational pattern in the publication of genre fiction over time (*Graphs, Maps, Trees*), but when Underwood looked at the language of genre fiction, he found that the pattern wasn’t generational shifts in genre or even a gradual consolidation of genre but instead an initial composition that doesn’t change much over a century (Underwood). The pattern of reception is different: even if the books changed little, people’s reception on Goodreads did

consolidate gradually over the past decade. What we have seen on Goodreads is a solidification of generic and non-generic clusters, where readers' responses are becoming more similar in language over time. Looking just at the text of books published over time wouldn't show this trend—for the full picture, we have to account for reception.

One potential way the Goodreads platform may encourage this growing similarity in language is by bringing reviews that are most exemplary of a genre or cluster to the first page of review results. The sort order of reviews for a given book on Goodreads follows no clear pattern: it's not in order of date written, length, number of likes or comments, or any other obvious metric. The most detail Goodreads has revealed is that it uses its “proprietary algorithm” to sort reviews. A review on the first page has a much higher chance of being read and therefore influencing genre definitions for everyone who reads and later writes reviews of the genre. So by putting reviews that use the standard language of the genre on the first page, Goodreads can encourage reviewers to think about the genre in those terms, and further solidify those terms as exemplary of the genre when readers use them in their own reviews. By comparing the predictability of reviews on the first page of results to those on later pages, we can see if this is happening on Goodreads, whether intentionally or as an unintended side effect of the algorithm Goodreads uses to sort reviews. The chart below shows the results of classifying the genre of reviews that appear on the first page of results compared to those on the second and third pages:




Figure 3.17: Predictability of reviews appearing on the first, second, and third page of Goodreads results for a given book.

For each of six established genres, reviews on the second page of results are slightly less recognizable as members of that genre than reviews on the first page. We see an even steeper drop in predictability on the third page of results, showing that these de-emphasized reviews are even more diverse in the language they use. By bringing reviews that use consistent language to the first page of results, the Goodreads sorting algorithm effectively encourages reviewers to replicate their language, making these genres more predictable and similar in discourse.

* * *

Goodreads is not a neutral platform for readers to discuss genre freely. Goodreads shapes how its users talk about genre, encouraging similarity and agreement rather than a diverse discourse. But where exactly is the push toward similarity coming from? I've found evidence in at least two aspects of the Goodreads platform: its review sorting algorithm, and its community-centric design.

Goodreads' review sorting algorithm highlights reviews that use a common discourse by bringing reviews that are more predictive to the first page of results. In this way, Goodreads influences its users' perception of a genre to better fit the perceptions of other users, thereby making its predictions—which rely on genre—more accurate. Every user becomes an unwitting advertiser, using a standardized vocabulary to recommend books to other users. While this effect is not necessarily intentional, Goodreads' design passively encourages and rewards this consolidation through its reliance on genre for its navigation, recommendation, and community features.

Goodreads' community-centric design is another factor causing users' responses to genre to become more predictable. Goodreads creates landing pages for each genre that serve as central hubs for content related to the genre, including links to subgenres, videos, related quotes, and interviews. I saw a clear effect for science fiction, where these linked subgenre terms like "aliens" and "military science fiction" became words that distinguished reviews of that genre from other genres, showing that the design of these landing pages creates an established vocabulary that later comes to distinguish the genre in reviews. Also linked on these genre landing pages are genre-devoted book clubs and discussion groups that are easy for users to join and advertise their membership on their profile. These groups institutionalize Goodreads genres by providing places where people talk together and consolidate the language they use to discuss the genre. The genre-based design of the site also makes it easy for users to focus only on one favorite genre and become immersed in content related to that genre, or to completely avoid a genre they don't like. Both situations cause a filter bubble effect which may account for some of the increase in predictability when users who don't like a genre—who might have more diverse opinions on it—stop writing reviews, leaving only dedicated fans of the genre who are more likely to think similarly about it.

Ultimately, while Goodreads promises a democratization of voices through its massive user base and the ability to create unlimited genres, its platform supports and encourages conformity within each of those genres. Conformity is good for Goodreads because it makes its algorithms more accurate: if everyone agrees on what a genre is and how to talk about it, making their responses more predictable to algorithms, then Goodreads will provide better-focused communities, better book recommendations, and better-targeted advertisements. The more accurate its predictions become, the stronger Goodreads' case is that its platform—which is

largely based on genre—is an accurate reflection of its users’ desires. But of course readers’ desires are too complex and diverse to be perfectly represented by any platform. Goodreads insists that prediction is powerful enough to anticipate readers’ every desire, and it shapes its platform to encourage users to act more predictably. However, what it predicts is not readers’ unmediated desires, but desires filtered through Goodreads’ algorithms. What readers see on the site isn’t an impartial representation of its users’ views on genre, but instead an altered view of readers’ depictions of genre through the lens of design decisions that emphasize keywords and tropes while avoiding disagreement and complexity. Goodreads is not a neutral site for readers to hash out what a genre is; rather, it shapes the language readers use to describe genre to become more consistent and less predictable. Instead of a thriving and democratic marketplace of genre ideas, Goodreads has created a well-oiled machine for the production and fulfillment of readers’ genre desires.

Chapter 4: Modes of Plot Summary

Plot summary appears wherever novels are discussed or sold: it's on the back of every book, at the beginning of every literary encyclopedia entry, and in nearly every book review and literary article. In Chapter One, we saw that the most common topic book reviews discuss is plot. Together with genre, plot summary is one of the most important ways reviewers recommend books and publishers advertise and sell books, since it helps readers predict whether they'll enjoy a book or not. But publishers, booksellers, literary critics, and individual readers talk about plot in different ways and for different purposes. Take, for example, these excerpts from summaries of Ruth Ozeki's *My Year of Meats*:

My Year of Meats was about a television show aired in Japan sponsored by American meat industry, whose main purpose was to increase US meat consumption in Japan. The show featured wholesome American wives cooking meat (“*pork is possible, but beef is best!*”) and overall promoting the wholesomeness of beef and American culture.

From a Goodreads review (Anny)

The leisure of even temporary unemployment holds little appeal for an ambitious documentary filmmaker named Jane Takagi-Little, the heroine of Ruth L. Ozeki's jaunty first novel, “*My Year of Meats*.” So when the producer of a new Japanese television series calls, Jane hits the road.

From a *New York Times* review (Funderburg)

MYM is centered on the story of Jane Takagi-Little, a documentary filmmaker who is employed by a U.S. beef lobby to produce *My American Wife!* (*MAW!*), a television show promoting American family values and meat to a Japanese audience, and who ultimately produces a documentary exposé of meat industry practices that reveals surprising things about cattle production and the status of race and gender in the United States.

From a scholarly article in *The Journal of Asian American Studies* (Cheng 191)

Each excerpt is a plot summary of the same novel, but each highlights different aspects of the book based on its goals and audience: the Goodreads review explains the meaning of the book's

unusual title to curious readers giving the review a quick skim, the *New York Times* review offers readers who appreciate literary language a sense of the book's playful style, and the academic article gestures to broader themes like race and gender to make a case to other scholars for the book's importance. While these are summaries of the same book, they look surprisingly different: Anny's is matter-of-fact and focuses on the TV show rather than the main character; the *New York Times* review is full of descriptive language and includes the contextual detail that this is the author's first novel; the academic article favors long sentences and abstract language that places the novel in a scholarly conversation.

What these reviews make clear is that an objective plot summary is impossible. There are too many matters of opinion intervening: which characters to mention, which events to focus on, what style of language to use, which abstract themes to mobilize—the list goes on. With so many sources, audiences, purposes, and styles of plot summary, it's clear that we need a more nuanced way of describing the practice of creating plot summaries for diverse readers and reading communities. Such a vocabulary would clarify how these communities express their viewpoints on plot within social reading platforms that have their own purposes and goals, sometimes complementary and sometimes competing with those of their users.

This chapter characterizes four different “modes” of plot summary, each with different priorities, purposes, and stylistic choices. I find that different social reading platforms use these modes to achieve their various goals, from creating a balanced portrayal of the plot to driving user engagement through suspense and emotion. Through this wider perspective on plot summary, we can see that while Goodreads emphasizes one particular type of plot summary, that is just one mode among many other possibilities. While it benefits Goodreads to propagate that mode on their platform, ultimately it narrows the discourse around plot. Unlike prior work on

computational models of narrative, my goal is not to define narrative plot, or even to reveal the plot structure of individual books, but to explore why readers respond to certain plot events and not others, and to determine how online social reading platforms guide these responses. Studying readers' plot summaries can help us outline the characteristics of plot for contemporary American readers, understand the underlying assumptions driving our choices of which plot events to include in reviews, and appreciate the extent to which our interpretation and enjoyment of novels depends on the values of the technology companies whose platforms we use and are used by.

Narrative, Plot, Event

Plot summary is just one aspect of the more comprehensive field of narrative fiction. Narratology has offered a range of definitions of narrative, but nearly all of them rest on a bipartite construction separating the story itself from how it is told. We can see both of these aspects in the *Cambridge Introduction to Narrative Theory*'s definition of narrative as "the representation of an event or a series of events" (Abbott 13). Narratologists have developed different terms for this distinction between events and their representation, but the most dominant are story and discourse, where "story" refers simply to the events that occur, and "discourse" involves all the complications of how it is narrated, from causality and ellipsis to frames and focalization. Plot describes how the story unfolds in order: Aristotle defines plot as an "arrangement of incidents" (sec.1450a), and literary critic Seymour Chatman calls it "the events of a story" (43). One might argue that the true substance of literature is in the artistry of the arrangement rather than its raw materials, but in focusing specifically on these raw materials,

we can observe the plot in its most basic, unadorned form. It is plot that we outline when we write plot summaries.

Despite the proliferation of new platforms and technologies of storytelling in recent decades, characterizing the many forms plot can take is not a new problem. Aristotle formulated the most enduring theory of plot, in which he argues that of all the parts of a tragedy, plot is the most important: “The plot then is the first principle and as it were the soul of tragedy” (sec.1450a). A proper plot, he argues, must be unified and organized around a single action with a clear beginning, middle, and end, and include no extraneous events. While these values remain dominant in how we informally evaluate plot, more recently, narratologists have developed a detailed vocabulary to describe a wide range of plots, many of which would not fit Aristotle’s prescriptions. Though the terms they use vary, each theory relies on some notion of the event as plot’s basic narrative unit. For formalist critic Vladimir Propp, the basic unit of the Russian fairy tale was the function: “Function is understood as an act of a character, defined from the point of view of its significance for the course of the action” (21). He outlines an exhaustive list of the 31 functions that make up the Russian fairy tale, associating each function with the character types who enact them, such as the hero or the villain. While later narratologists offer more complex visions of plot, such as the structuralist critic Gérard Genette’s nuanced account of plot based on grammatical tense, mood, and voice, they tend to share with Propp the same understanding that a story rests on a basic narrative unit: “For me, as soon as there is an action or event, even a single one, there is a story because there is a transformation, a transition from an earlier state to a later and resultant state” (Genette 19). Whether we call the basic unit a “function,” “action,” or “event,” formalist and structuralist narratologists offer us a vocabulary to simplify and break down the units of narrative.

The idea of basic narrative units has fueled cross-disciplinary work in fields outside of literary studies. Cognitive psychologists developed theories to describe the way people reconstruct stories in their minds, such as “script theory,” which treats human behavior as a scripted pattern (Schank and Abelson), and “story grammars,” which attempt to model the construction of stories as mathematical equations that they claim reflect how people remember stories (Black and Wilensky). These formal definitions of plot paved the way for computer scientists to develop automatic ways of characterizing plot in the recent fields of auto-summarization and story generation. Beginning with formalists and continuing in modern computational methods, the distillation of narratives into their basic units has been at the foundation of how we understand plot.

And yet we cannot fully account for plot by remaining within the confines of the page, as later narratologists have made clear. Literary theorist Roland Barthes led the shift from text to reader as the ultimate source and interpreter of narrative devices. For Barthes, texts have multiple meanings:

To interpret a text is not to give it a (more or less justified, more or less free) meaning, but on the contrary to appreciate what plural constitutes it. Let us first posit the image of a triumphant plural, unimpoverished by any constraint of representation (of imitation). In this ideal text, the networks are many and interact, without any one of them being able to surpass the rest; this text is a galaxy of signifiers, not a structure of signifieds; it has no beginning; it is reversible; we gain access to it by several entrances, none of which can be authoritatively declared to be the main one; the codes it mobilizes extend as far as the eye can reach, they are interminable (meaning here is never subject to a principle of determination, unless by throwing dice); the systems of meaning can take over this absolutely plural text, but their number is never closed, based as it is on the infinity of language. (Barthes, *S/Z* 5–6)

For Barthes, there is not just one interpretation of a text—each person contributes their own meanings to create a “plural text,” with many possible perspectives and emphases. Similarly, literary critic Peter Brooks finds formalist theories of narrative insufficient, instead offering a

more dynamic theory that takes human thought processes into account. Rather than an activity wholly enacted in the text, Brooks defines plot in terms of its effect on the reader: "Plot as we have defined it is the organizing line and intention of narrative, thus perhaps best conceived as an activity, a structuring operation elicited in the reader trying to make sense of those meanings that develop only through textual and temporal succession" (37). While Brooks turned to Freud to provide a theory of human motivations, more recently, scholars in cognitive narratology have begun to study how readers make sense of plot neurologically. The critical shift in these theories of plot is that they locate it within the reader rather than the text itself.

Given this turn from book to reader, then, it is a natural move to shift our focus from plot to plot summary. Plot summary allows us to see these reactions and interpretations concretely, as the reader reformulates the plot in their mind in order to reconstruct the plot in their summary. However, plot summary is a narrower category than plot, with its own topical and stylistic constraints. Take, for example, this plot summary from the *Publishers Weekly* review of *My Year of Meats*:

Japanese-American documentary filmmaker Jane Takagi Little seems to have found the perfect job producing *My American Wife*, a program sponsored by American beef exporters that introduces Japanese housewives to "typical, wholesome" American wives, their families and their beef recipes. Jane and her crew travel around the U.S., filming wives and their families as they make beef dinners. Meanwhile, in Tokyo, shy Akiko has been driven to bulimia by her domineering and abusive husband, John, who works with the beef exporters on the show. John insists that Akiko watch the show, cook, gain weight and get pregnant. Over the course of the "year of meats," Jane begins to feel guilty about exploiting the wives, confused about her romantic life and disturbed by the sordid secrets she uncovers about meat production. Inspired by Jane's increasingly subversive episodes (particularly the segment on lesbian vegetarians), Akiko gradually realizes what she wants out of life and finds the courage to reach for it. (*Publishers Weekly*, "My Year of Meats")

Like most plot summaries, this summary is organized grammatically around characters and events: each sentence's subject is a main character (Jane, Akiko, and John), and each recounts an

event or series of events in the book. It leaves out some of the more complicated aspects of the book's plot that narratologists might want to study, such as its style, narration, and framing. While it may seem that plot summaries give us only a pale representation of the book's characters, story, and complexity, they give us valuable insight into individuals' understanding of the plot. For example, we can ask: with so much of the novel seemingly off limits in the conventions of the plot summary format, how did *Publishers Weekly* decide what to include in their summary? Plot summary is a restricted format that leaves a lot out, but within that restricted format we can see the results of a strict process of selection and evaluate people's priorities by what they choose to include.

While this short introduction has shown how closely literary scholars have dissected plot, plot summary is much less popular an object of study. Perhaps the closest parallel is the study of narrative description, which is similar to plot summary in that it is condensed, accelerated, and selective. However, plot summary is a different kind of description, located outside the story and written by diverse authors. Scholars who study description characterize it as a process of selection that brings attention to specific aspects of a given object. Plot summary, in fact, demands precisely this: to choose certain parts of a text to focus on and interpret. This process of selective attention is increasingly relevant in the twenty-first century, with new forms of media and economies predicated on capturing our attention. In literary studies, this concern manifests in terms of the archive, where scholars are reclassifying the types of digital objects we should preserve as "attention data" that reflects not necessarily the canonical material but the material that has succeeded in the battle for our attention (Summers). As technology companies continue to measure their success in clicks and page views, the effects are clear in the way we describe books on their platforms. More and more, technology companies are thumbing the scale of what

types of attention they encourage. As we will see, plot summary is another “means of attention control” which Goodreads leverages to promote certain types of responses (Houser). We need a new theorization of plot summary for the digital age of attention control.

Methods for Studying Plot Summary at Scale

Building on Barthes’ reader-response theory of plot, this project turns to plot summaries in order to capture the many diverse ways readers interpret plot. I treat plot not as an objective account of what happens in a book discoverable by programmatic means, but as an ensemble of possible readings and reactions that we can begin to understand through plot summaries. My goal is not to determine what plot is, but what plot does for readers: why do they draw out certain aspects of the plot and not others? And how does the platform on which their response appears influence that response?

Even in the relatively limited sphere of plot summary, leaving out the vast context available to narrative through chronology, ellipsis, focalization, and countless other devices, we can use plot summary as the basis for a more nuanced look at the differences in conventions on different social reading platforms we now have access to on Goodreads and other social reading platforms. But the vast scale of these platforms—the sheer quantity of material now available on social reading websites—demands digital humanities methods, which will allow us to sample a wider audience and range of sites, and notice large-scale patterns in their language and content. Are there guiding rules that they use, such as including events that are most pivotal to the progression of the plot, or events that are emotionally salient? By noticing which platforms’ summaries highlight key moments and which highlight idiosyncratic or peripheral moments, we

can begin to interrogate the values motivating the writers on each platform. A plot summary exists on a social reading platform as an argument about what is important about the book. Through the lens of plot summary, we can infer the values of different social reading platforms, observing how each one defines and shapes what counts as plot to its audience.

To notice these patterns, my methods isolate the individual events that constitute plot summaries. Using a character list, either from a summary website or automatically identified, my program extracts every subject-verb-object group in a summary that includes any of the book's characters. These subject-verb-object groups, or simple events, tell us which characters the summaries on different social reading platforms favor, whether they use descriptive or emotional verbs, and any other patterns we can find in the types of events various platforms favor.

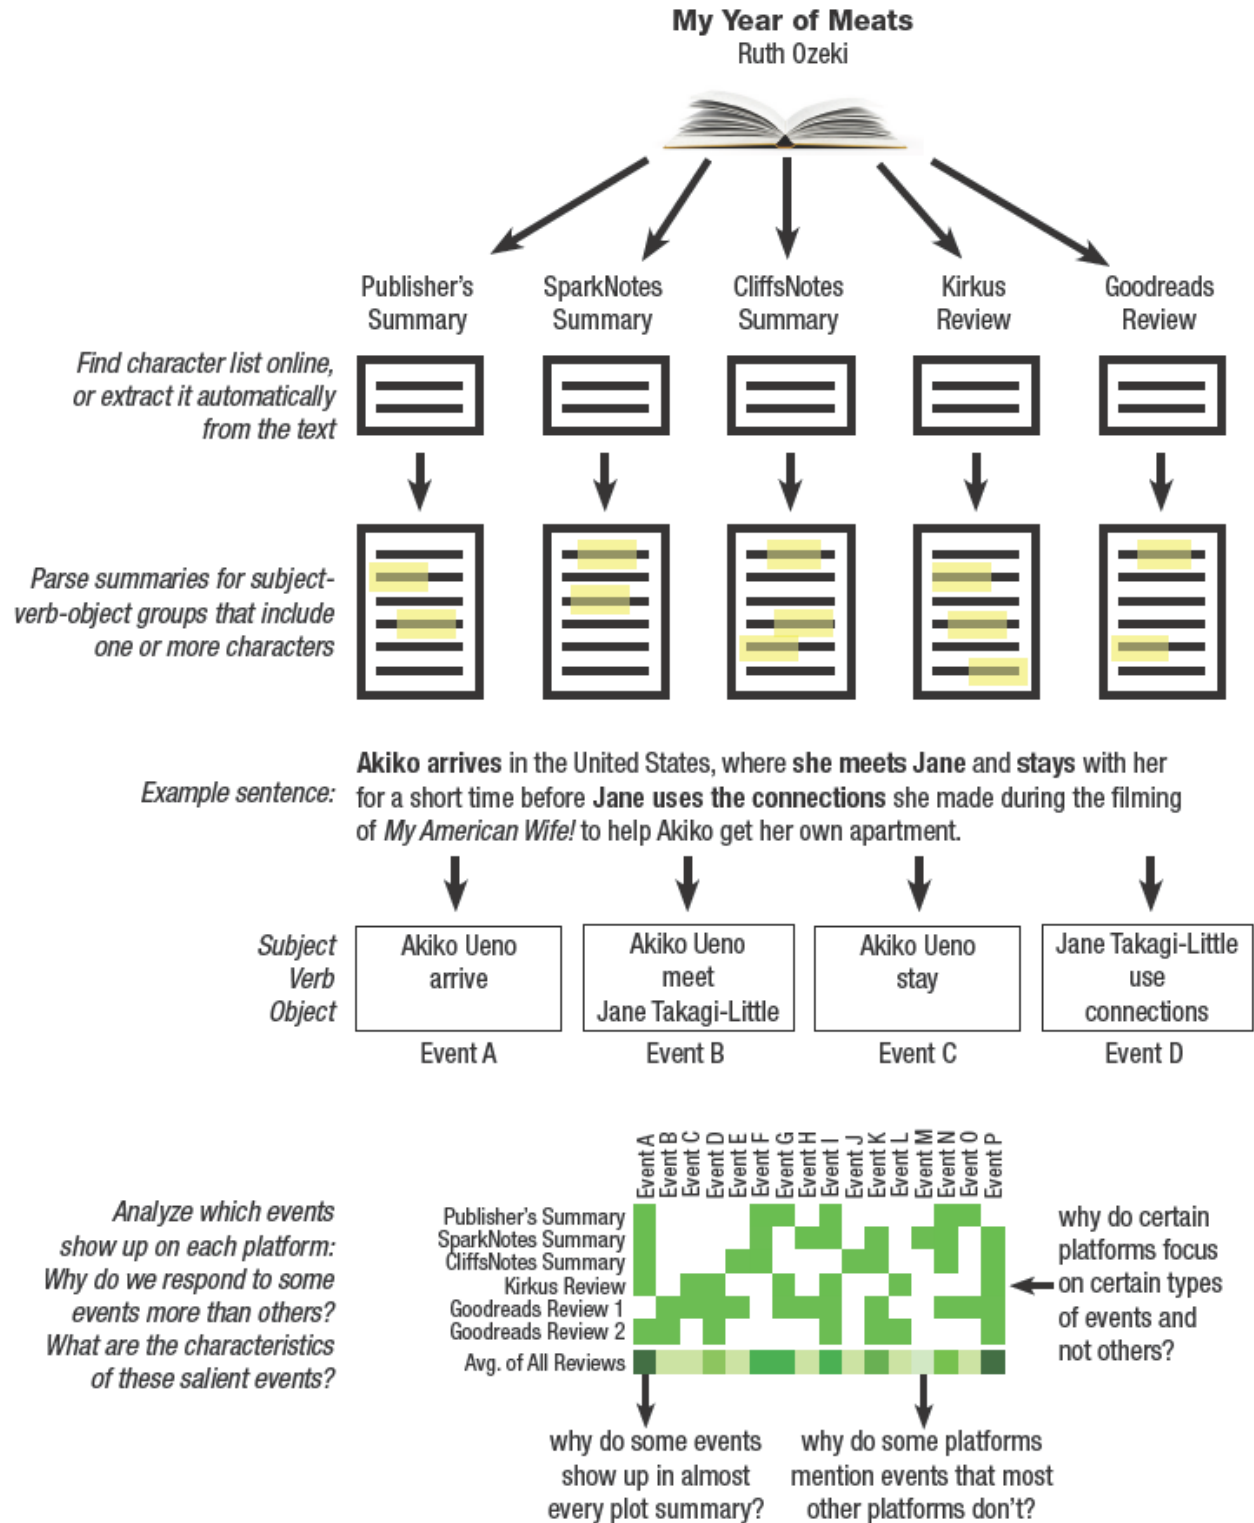


Figure 4.1: Diagram detailing this project's methods.

In requiring events to involve a character, I have raised a potential issue: what exactly counts as a character? Many novels include ghosts, natural disasters, and other entities that might test the limits of any definition of character. The typical procedure for computational projects that detect plot and characters is to use a “gold standard” character list from an open encyclopedia like Wikipedia (Bandari et al.) or a database of metadata like Freebase (Bamman et al. 353). Some projects attempt to detect characters automatically: for example, one group finds characters in novels by looking for nouns that perform character-like actions, such as “rejoice,” “accost,” and “frown” (Vala et al. 771). My project uses a mixed method: if the book has one or more character lists available online, I combine them and use all of the characters from those lists; if not, I automatically detect people and organizations that modify or are modified by a verb in the summary. In either case, the resulting characters do not just include those from an official character list or that recognizably denote a person or organization. When a noun that is not clearly a character acts on or is acted on by a character, it appears in my results, as I will discuss in more detail below. For example, the CliffsNotes summary of Zora Neale Hurston’s *Their Eyes Were Watching God* includes the sentence, “There [Janie] met new people, Tea Cake’s fun-loving friends, and experienced another community” (CliffsNotes). While only Janie and Tea Cake are on CliffsNotes’ character list for the book, one of the events my method finds is “Janie experienced community,” with “community” functioning as a character with a relationship to Janie. Still, a different method could more effectively pick up these inanimate, unconventional potential characters even when they appear on their own. When I tried these methods, they produced too many results that didn’t look to me like events. For this reason, I specifically focus on events that involve characters, rather than all grammatically possible events.

In this project, then, events are any grouping of a subject, verb, and object that include a character. In the words of literary critic Tzvetan Todorov, “To combine a noun and a verb is to take the first step towards narrative” (119). Nouns and verbs alone, of course, paint a highly simplified picture of narrative, but grammatical definitions of “event” have their roots in the theory and practice of several prominent narratologists. When formalist critic Vladimir Propp analyzed the plot of Russian folktales, he laid out a series of “functions,” or actions of a *dramatis personae*, that compose each folktale, from “interdiction” and “violation” to “punishment” and “wedding” (26–64). While I don’t claim to have assembled a complete catalog of the functions that make up novels, my project follows Propp in moving fluently between individual verbs and the narratives they belong to. I depart from Propp, however, in my treatment of characters. While Propp’s *dramatis personae* were interchangeable within their categories, such as “villain” or “helper,” my project allows a much wider range of activity for each character. Literary theorist Algirdas Greimas helpfully distinguishes between “*actants*, having to do with narrative syntax, and *actors*, which are recognizable in the particular discourses in which they are manifested” (Greimas 106). To put my method in Greimas’ terms, I aggregate all of the grammatical *actants* that involve one or more *actors* from the novel. This distinction allows me to narrow my project’s scope to include only actors who are characters within the text—rather than authors, historical figures, or characters from other novels—without missing the unexpected or abstract *actants* with whom these characters interact in plot summaries.

My simplification of events as groups of subjects, verbs, and objects, however, offers only a partial and selective vision of the events discussed in plot summaries. It leaves out much of the complexity of narrative, such as context, focalization, and narrative reliability. However, the goal of this project is narrower than narrative: it aims merely to survey at scale what kinds of

events people include in their plot summaries. Events offer a simplified depiction of a book, but what they show is still useful, such as which events appear consistently and which rarely or not at all, the quality of those events, and how summaries arrange them chronologically.

I collected plot summaries from a variety of social reading websites: Goodreads, including both the publisher’s plot summary at the top of the page and the plot summaries included within book reviews on the site; Wikipedia, which includes a plot summary section for most books; student-focused summary sites GradeSaver, Shmoop, SparkNotes, and CliffsNotes; and professional book review websites *Kirkus* and *Publishers Weekly*. The books I ultimately included in my analysis were American novels, plays, and short stories with at least three summaries available online from any of these social reading platforms. In total, this list included 75 books.⁵⁰ Not all of the websites had summaries for every book:




Figure 4.2: Breakdown of this project’s corpus, which includes books that have plot summaries on at least three different platforms. The chart shows how many total summaries each platform had in the corpus.

⁵⁰ The full list of books appears in the Appendix.

The 75 books I examined are not a random sample: to have summaries on multiple sites, books must be either popular or canonical. As a result, they have an earlier average publishing date than a random sample would, and a broader, less self-selected audience of summary-writers on Goodreads, some of whom read the book as a school assignment.

For each book, I looked for a list of characters from the summary websites in the corpus. If there were multiple lists, I included every character on every list, removing duplicates. If character names were similar but not exact, as in the case of nicknames, I favored the longer, most complete version of the name. For example, I found several different character lists for Margaret Atwood's *Oryx and Crake*, and combined them in the following way:

Goodreads: Snowman, Oryx, Crake

Wikipedia: Snowman, Crake, Oryx, Sharon, Jimmy's father, Ramona

GradeSaver: Amanda Payne, Bernice, CorpSeCorps, Crake, Crakers, Crake's mother, Glenn, God's Gardeners, Jack, Jimmy, Jimmy's father, Jimmy's mother, MaddAddam, Oryx, Ramona, Sharon, Snowman, Susu, Uncle En, Uncle Pete

Combined list: Snowman, Oryx, Crake, Sharon, Jimmy's father, Ramona, Amanda Payne, Bernice, CorpSeCorps, Crakers, Crake's mother, Glenn, God's Gardeners, Jack, Jimmy, Jimmy's mother, MaddAddam, Susu, Uncle En, Uncle Pete

If none of the sites offered a character list, I used Stanford CoreNLP's named entity recognition parser to automatically identify people and organizations in the summary as potential characters (Manning et al.). Ruth Ozeki's *My Year of Meats* had no official character list available online, so the parser found these characters automatically:

Jane Takagi-Little, Akiko Ueno, John Ueno, BEEF-EX, Sloan

While BEEF-EX stands out as an organization distinct from the other characters, which are all human, it does act in some ways like a character to encourage the growth of the American meat market in Japan and motivate many of the other characters' actions throughout the novel. While in this case, the character list looks relatively accurate, some books had characters that were

particularly difficult for the parser to recognize, such as Cormac McCarthy's *The Road*, which simply calls its characters "the boy" and "the man," causing confusion between multiple men who appeared in the novel. It was also common for the parser to identify the book's author as a character, so I had to filter out names from outside the book by hand.

Goodreads reviews, since they often include extraneous material that isn't strictly plot summary, posed a few problems of their own. The method sometimes picked up extraneous material like evaluations and general statements about the book itself, such as treating statements like "I loved Jane" as an event from the novel. To minimize these mistakes, I removed quotations from the reviews, as well as any event that involved "I," "me," "my," or a proper name from outside the book as its subject or object. By only including events with characters, we effectively filter Goodreads reviews for just the parts that are plot summary. Some noise, however, certainly still exists in the results. For example, the sentence "The characters included Jane and Akiko" would create the event "include Jane and Akiko," which is not an event from the novel. However, my conclusions rely on the most common events that appear in many reviews, and these mistakes did not occur often enough to overshadow more legitimate events from the summaries.

To identify the events that compose each plot summary, I needed to first identify each sentence's subjects, verbs, and objects. I used a method fleshed out by computer scientist David Bamman in two projects, one using Wikipedia biographies and another using Wikipedia movie plot summaries, to find the "agent" and "patient" of verbs using linguistic rules (Bamman and Smith; Bamman et al.). Following Bamman, I used Stanford CoreNLP to divide summaries into sentences, mark each word's part of speech, and identify the referent of each pronoun (Manning et al.). For example, take the final sentence of the Wikipedia summary of Ruth Ozeki's *My Year*

of *Meats*: “Akiko arrives in the United States, where she meets Jane and stays with her for a short time before Jane uses the connections she made during the filming of *My American Wife!* to help Akiko get her own apartment” (Wikipedia, “My Year of Meats”). The parser marks each word with its part of speech:

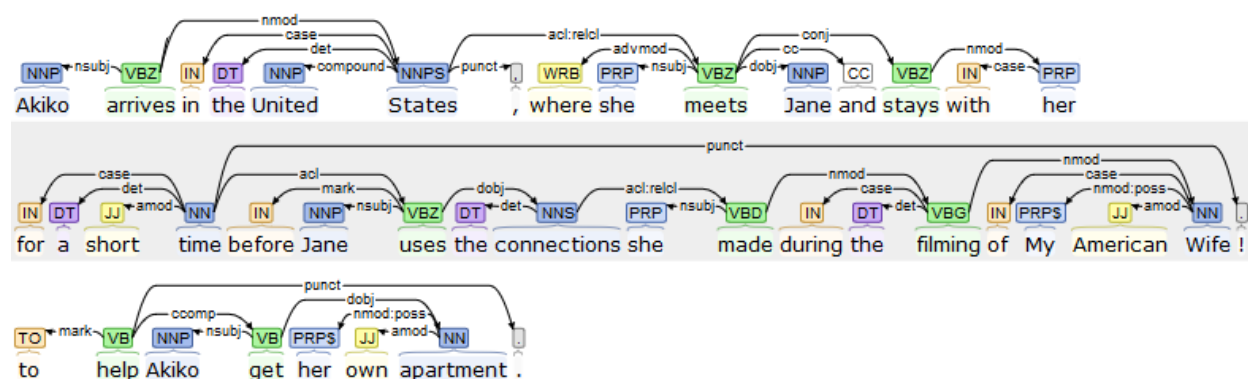


Figure 4.3: Stanford CoreNLP dependency parse results for a sentence from the Wikipedia summary of Ruth Ozeki’s *My Year of Meats*.

In this diagram, the green tags mark verbs, with arrows connecting words to each other according to linguistic rules for dependency parsing. Since natural language is complex, the parse is not always perfectly accurate. In this case, the parser cuts the sentence in two because of the exclamation point in the show title “*My American Wife!*,” separating the ending of the sentence, “to help Akiko get her own apartment.” This mistake made no difference in the final results, however, since my method still captured the relevant subjects, verbs, and objects in each portion of the sentence. I also substituted pronouns for their referent using CoreNLP’s coreference parser, in effect transforming “where she meets Jane and stays with her” to “where Akiko meets Jane and stays with Jane.” This final “her” is ambiguous—it could refer to either Akiko or Jane—and the parser can also make mistakes, but the results are typically accurate. Finally, I lemmatized the verbs, which treats “arrives,” “arrived,” and “arriving” as the same verb—“arrive”—in order to consolidate the results despite small differences in reviewers’

grammatical and stylistic choices. After proceeding through all of these steps, the program ultimately extracts events as subject-verb-object groupings from each sentence:

Akiko arrives in the United States, where **she meets Jane** and **stays** with her for a short time before **Jane uses the connections** she made during the filming of *My American Wife!*....

- Event 1: Akiko Ueno arrive
- Event 2: Akiko Ueno meet Jane Takagi-Little
- Event 3: Akiko Ueno stay
- Event 4: Jane Takagi-Little use connections

In this example, the program excluded one verb—“filming”—since it didn’t have a character as its subject or object. Below is a breakdown of how many events each platform contributed to the total number of events in the dataset:

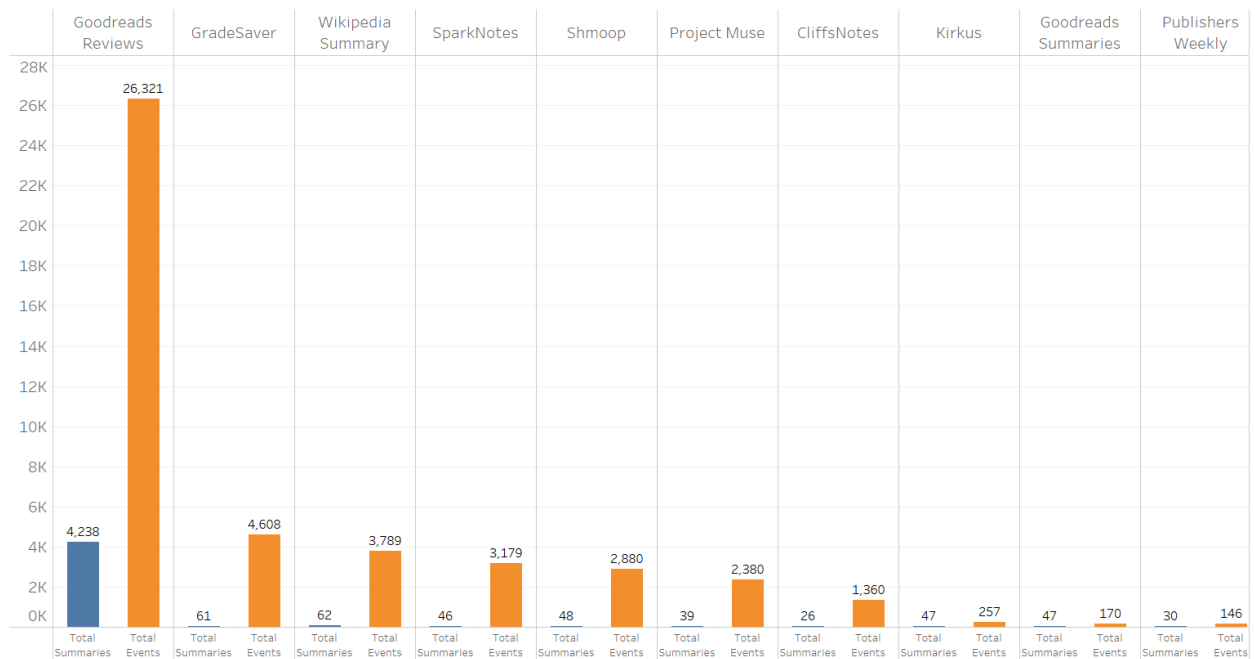


Figure 4.4: Total number of plot summaries and events from each social reading platform.

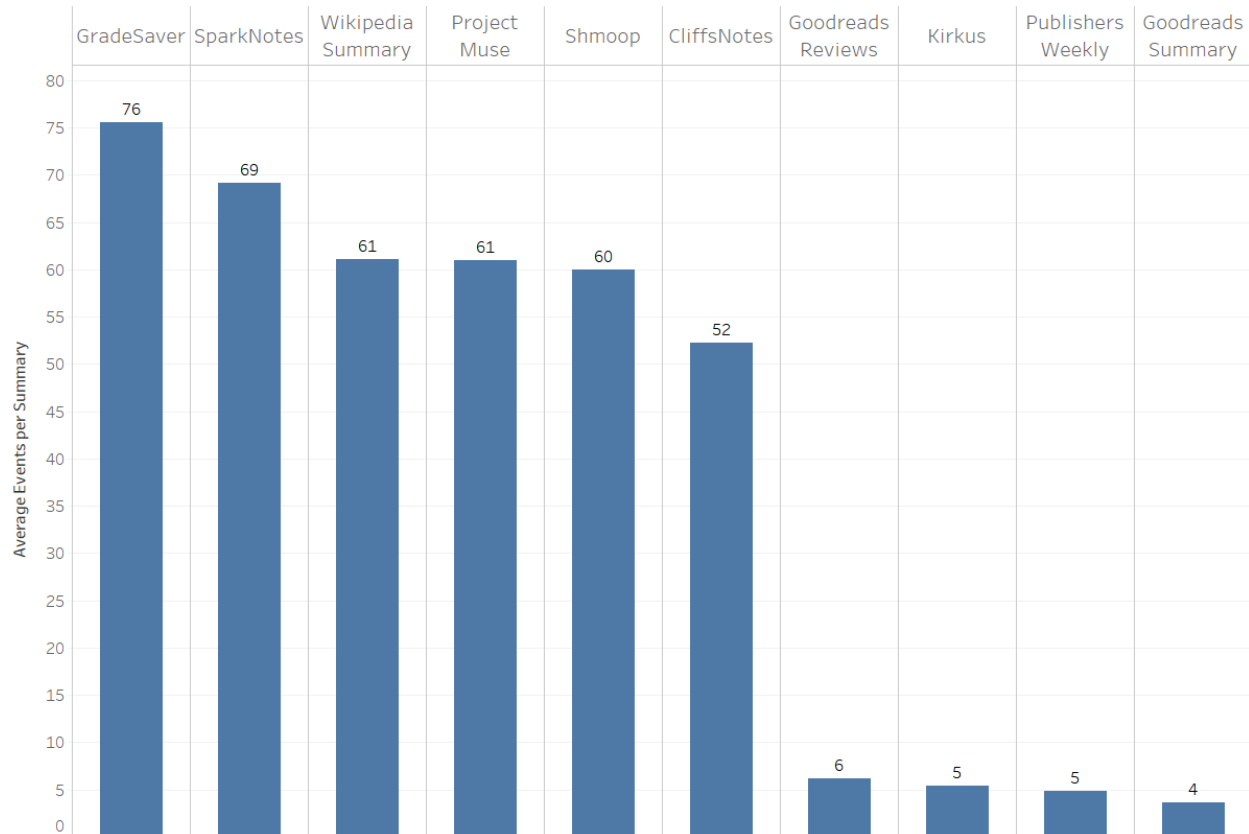


Figure 4.5: Average number of events in each summary, broken down by platform.

Goodreads provided the largest raw number of plot summaries to analyze, since many Goodreads reviews included plot summaries anywhere between one and twenty sentences long. The average Goodreads review, however, only contained four events, while literary-professional platforms included many more events, with the highest average number of events in summaries from GradeSaver. These dramatic differences in length and style from platform to platform makes comparison difficult, but highlights the unique goals of each platform. As we will see, GradeSaver’s high number of events makes sense when we consider its goal of providing students comprehensive coverage of a novel so that they can pass a reading comprehension test. Each platform has a similar pairing of audience and goals that makes its plot summary style more comprehensible. By isolating the event as a unit of analysis, we can begin to understand the logic behind these patterns in the types of events each platform emphasizes.

The Modes of Plot Summary

The provision of this kind of summary is only one of dozens of distinguishable techniques for providing facts, most of which—perhaps fortunately—have never been named. (Booth 171)

At the risk of dismaying literary critic Wayne C. Booth, who in his discussion of narrative summary decries the idea of naming every possible technique, in this section I will attempt to distinguish and characterize the modes of plot summary found on social reading platforms. By delineating the types of characters and events each mode favors, we can see how different platforms use different modes to achieve their goals, whether to sell a book, inspire emotion, or help students pass a test. Each mode points to a different use for plot summary, which then shapes what's included and excluded in the summary.

The Major/Minor Mode

Perhaps the most intuitive metric for which events to include in a plot summary is “importance.” In the major/minor mode, summaries favor main characters over minor ones, and central plots over subplots. To justify these distinctions, we would have to assume that narratives are composed of important and unimportant events that readers can recognize. Fortunately, both Barthes and Chatman offer theories for how to delineate which events are important. They each distinguish events that are integral to advancing the plot from events that serve a less crucial purpose. Barthes calls important events “nuclei” and minor events “catalyses,” while Chatman calls them “kernels” and “satellites.” For both theorists, an important event “advances the plot by raising and satisfying questions” (Chatman 53), while minor events “do no more than ‘fill in’ the narrative space” (Barthes, “Introduction” 247). This distinction rests on the story-discourse divide, or the divide between the events in a text and how they are represented. Important events

would change the story noticeably if removed, while minor events merely supplement the discourse: a minor event “can be deleted without disturbing the logic of the plot, though its omission will, of course, impoverish the narrative aesthetically” (Chatman 54). For both Barthes and Chatman, important events are those that are integral to the plot, and move it forward toward its logical conclusion.

The distinction between important and minor events does not just exist in literary theory. It’s widely acknowledged in popular discussions of narrative, and appears as an underlying assumption in the most prominent guides for writing plot summaries. Take, for example, two of the top results for the Google search “how to write plot summary”:

The goal of a plot summary is to give readers a basic understanding of the story, which includes introducing them to characters and essential plot points. You try to avoid specific details, such as a characters' appearance, unless it plays a pivotal role in advancing the story. (Reference.com)

Necessary detail must be maintained....the *Odyssey* contains various scenes where people recount myths to each other, and other such scenes of little importance to the main plot. If most of these get left out, or mainly consist of a sentence or two, that is not a problem, and helps keep the focus on the main story.” (Wikipedia, “How to Write a Plot Summary”)

These guides use importance explicitly as a way to determine which events to include in a plot summary. But do readers actually follow these guidelines when they set out to write plot summaries? Social science research suggests that readers do consider the importance of events they include in their plot summary, but they have a nuanced definition of importance. Linguists Walter Kintsch and Teun van Dijk found that readers use two standards to evaluate which events to include in a plot summary: “textual relevance in which importance is defined in terms of what the author considers important,” and “contextual relevance, where importance derives from the reader’s personal interests or background knowledge” (qtd. in Winograd 406). In writing summaries, readers must choose between two different arbiters of importance: the author and

themselves. Can we see evidence of these two competing theories of importance in online plot summaries from different platforms?

One possible way to judge an event's importance might be whether or not it involves a main character as opposed to a minor character. To determine which characters qualify as main characters, several digital humanities projects have developed algorithmic measures of character centrality. One Stanford Literary Lab project attempts to find the most important characters in *Hamlet* by measuring the "character space"—or lines of dialogue—devoted to each character in the play (Moretti, "Network Theory, Plot Analysis" 3). The resulting network diagrams, while they generated productive questions, were difficult for researchers to decipher.⁵¹ A more recent project compares the representation of male and female characters in fiction from 1800-2007. This project evaluates characters by how many words the text uses to describe them, finding that the representation of female characters has declined since the nineteenth century (Underwood and Bamman). Both of these methods use relatively simple metrics to approximate a character's importance to the text, but they affirm that quantitative methods are especially apt at comparing the importance of different characters at scale.

Rather than measuring the prominence of characters in literature, this project measures the importance of characters to readers, as measured by their representation in plot summaries. A summary that mentions only one or two characters makes a fundamentally different argument about what is important in the novel than one that details ten characters, or an entirely different set of characters. When I reviewed the platforms targeting students and classrooms, including CliffsNotes, GradeSaver, Shmoop, and SparkNotes, it was much more common to see

⁵¹ As the pamphlet concludes, "Now, there is no question that these figures contained much more information than Figure 1 of the pamphlet: they showed, not just who had talked to whom, but also whether the exchange had been mutual, and how extensive it had been (measured in the number of words). All this was new. Was it also *visible*? Clearly the answer was no" (Moretti, "Network Theory, Plot Analysis" 12).

summaries explicitly name the main characters at the outset, distinguishing them from other, minor characters. Take, for example, Cormac McCarthy's *No Country for Old Men*. The GradeSaver summary begins, "The central plot of the novel follows the experiences and thoughts of three central characters, Sheriff Ed Tom Bell, Hunter Llewelyn Moss, and psychopathic hitman Anton Chigurh, and how their paths intersect over the course of a series of highly disturbing and violent events" (GradeSaver, "No Country for Old Men"). The three characters GradeSaver highlights all happen to be male—but we see a different trend when we look at plot summaries in Goodreads reviews. These amateur summaries explicitly include minor characters in their plot summaries, such as Steven's declaration that "For me, the much more interesting characters were Sheriff Bell of Jewel County and Llewyn Moss' wife, Carla Jean. Carla Jean was a fascinating character in the book (I thought she was played brilliantly in the movie too) with much more depth than you might imagine" (Steven). Steven explicitly acknowledges that Carla Jean, as a minor character, might have been expected to have little depth, and that he was pleasantly surprised to find the opposite. These examples suggest that Goodreads reviews are more willing to include minor characters in their plot summaries. But does this trend hold at scale?

By extracting events from plot summaries in my corpus, I had at hand a rough representation of the makeup of each plot summary and the characters it involved. I could then see, for each book, which characters received the most (and least) attention. When I looked at how each platform discussed the characters in *No Country for Old Men*, I saw that two distinct styles separated the platforms—one, a distributed focus on multiple characters, and the other, an intense focus on one character:




Figure 4.6: For each character in *No Country for Old Men*, the chart depicts the percent of all total character mentions attributed to that character as either subject or object.

The chart shows, for each plot summary platform, how many of the total character references cited each character. Thus the *New York Times* summary of *No Country for Old Men* includes only events that involve Sheriff Ed Tom Bell, resulting in 100% representation, while out of 58 total events involving characters in the GradeSaver summary, only 20 of them include Sheriff Bell, resulting in a 34% score. The four platforms at the top of the chart each discuss at least three different characters from the novel. The four at the bottom, however, only include events that involve a single character. These four platforms represent professional literary critics' and publishers' summaries. "Goodreads Summary," at the bottom, refers to the summary at the top of the *No Country for Old Men* webpage on Goodreads, which most often comes from the back cover of the book or the publisher's website (Goodreads, "Help Topic: Description").

Interestingly, these four literary-professional summaries don't agree on the book's main character: while three focus on Sheriff Bell, the publishers' summary from Goodreads instead

cites Llewelyn Moss. That the first two platforms are the most evenly distributed is not surprising, since they are the two multiply-authored platforms, and they reflect the diverse interests of their wider authorship base. In the middle, the two platforms focused on students—Shmoop and GradeSaver—more resemble the amateur-written sites than the literary-professional sites, despite being professionally authored. GradeSaver’s tagline, “Getting You the Grade,” offers a likely explanation: in order to prepare their audience of students for tests and class discussions, their summaries need broad coverage of as many characters as possible.

Moving beyond *No Country for Old Men*, we can expand our analysis to see if these trends hold across many different books. If we look at the percentage of each review devoted to the most common character—taking, from the chart above, only the character with the highest percentage value for each platform—are there platforms that frequently focus solely on a single character, and other platforms that tend toward a more balanced distribution?




Figure 4.7: For each platform, the box-and-whisker plot shows the distribution of the percent of total character mentions attributed to the characters on that platform.

In the chart above, each point shows the percent representation of the main character in a book summary from the platform.⁵² High values, toward the right of the chart, show the intense focus on a single character that we saw in literary-professional summaries above, while low values, toward the left, show more balanced attention to characters in the plot summary. We can compare platforms in terms of their average, marked by the line at the center of the box plot, and their variance, or the width of each horizontal box. Overall, the results reflect what we saw in *No Country for Old Men*: classroom-focused plot summaries from SparkNotes, GradeSaver, Shmoop, and CliffsNotes focus less intently on a single main character than literary and publisher-created summaries. It's very common for publishers' summaries to mention only one main character—in fact, that's true of one- to two-thirds of publishers' summaries from Goodreads, *Publishers Weekly*, and *Kirkus*.⁵³

This hyper-focus on main characters in publishers' summaries means that they rarely talk about minor characters, who were most often female in my corpus—we saw this clearly in their avoidance of Carla Jean Moss, the lone female character in *No Country for Old Men*. Classroom-focused platforms have a more distributed focus on a larger number of characters, as well as the multiple-authored Wikipedia summaries and Goodreads review summaries. These platforms give us a more balanced representation of the book itself, inflected by the interests and priorities of the summary writers. Particularly on Goodreads, the disproportionate focus we see on Carla Jean Moss—the most of any platform—shows us that readers were interested enough in the character to include her disproportionately in their plot summaries, counter to the editorial judgment of literary-professional summaries, which were more restrictive.

⁵² For Goodreads reviews, the points represent the average representation of each character across every review from that book.

⁵³ 18 out of 30 publishers' summaries from Goodreads (60%) have only one character, 6 of 19 (31%) from *Publishers Weekly*, and 10 of 33 (30%) from *Kirkus*.

In the same way that we looked for main characters by measuring how much attention they receive in plot summaries, we can look for major events by examining the most common verbs that appear in plot summaries. Take, for example, the most common verbs in summaries of Daniel Keyes' *Flowers for Algernon*:

Flowers for Algernon	
Goodreads Reviews	have begin realize want become learn write get see treat
Wikipedia Summary	try keep lose fire not have realize remember start believe
CliffsNotes	have reduce
GradeSaver	realize take try begin die feel fire remember confront
Shmoop	get take hope not have become begin believe call die
SparkNotes	realize get become begin believe call die take write ask
Goodreads Summary	begin change tell

Figure 4.8: The most frequent verbs in plot summaries of *Flowers for Algernon* by platform.

The chart above shows the verbs that appeared most frequently in plot summaries of the novel on each platform. If we judge an event's importance by how many different summaries it appeared in, the most important events in *Flowers for Algernon* are moments of realization, marked by the verb "realize." The book has one key moment of realization—when the main character, Charlie, learns that an operation that has increased his intelligence will eventually lose its effect—as well as a series of smaller realizations that underline the book's message about the benefits and drawbacks of intelligence. As a whole, the summaries frame the book as a series of realizations:

“He realizes that Nemur's hypothesis contains an error and that there is a possibility that his intelligence gain will only be temporary.” (SparkNotes, “Flowers for Algernon”)

“When Charlie travels with his doctors to Chicago to announce the results of his surgery, he realizes that not only has intelligence leap-frogged them all, but they will never see him as more than a patient, a human Algernon.” (Goodreads user Bruns)

“Charlie realizes that whatever he does with his time left must be meaningful for other people.” (GradeSaver, “Flowers for Algernon”)

“Turns out teacher is way cuter than he ever realized, which makes him eager for some extra tutoring.” (Shmoop, “Flowers for Algernon”)

“He realizes his co-workers at the factory, who he thought were his friends, only liked him around so that they could make fun of him.” (Wikipedia, “Flowers for Algernon”)

Despite the wide range of interpretations readers assign to each moment of realization, to “realize” is clearly important to the story, as evidenced by its frequent use in plot summaries on a variety of platforms. But when we look more closely at the types of events readers are referring to in their summaries, we can see a wider range of events that we might have expected. It’s not just one pivotal moment of realization, but many small ones. What is it that made these events stand out to readers as they wrote plot summaries? What is notable about them, enough to merit mention in one’s plot summary? The next three modes explore rationales for including events that might not be important, but deserve inclusion by another standard, whether they are unusual, suspenseful, or emotionally salient.

The Unusual/Unremarkable Mode

When we choose events to focus on in our plot summaries, a common practice is to value the unusual over the unremarkable. Education scholar Karl K. Taylor observed a trend in the types of plot points that grade-school students include when they sit down to summarize an article: “Often they pointed out the most unusual and ignored what they considered common knowledge” (390). In the context of science fiction, literary theorist Darko Suvin gives us a term for such a distinctive moment in the plot: a “novum,” which he uses to mean “novelty, innovation” (*Metamorphoses of Science Fiction* 63). The novum can be a setting, an invention, a phenomenon, or a character, as long as it separates the world of the novel from the real world.

Given this lack of specificity, it is difficult to make rules for recognizing the novum in science fiction texts. One digital humanities project has attempted this feat, asking, “can we quantify the novum?” (Simeone et al.). Inspired by Suvin’s theory, the project uses machine learning to identify distinctively “strange” words in science fiction texts. My project has a different focus: rather than finding words a science fiction author uses to mark the novum in fiction, this project instead examines plot summaries for the terms that people repeatedly use to express what surprised them or stood out to them about the book—a reader-response perspective on the novum.

More concretely, I look for the novum among the most unusual verbs in each plot summary. To quantify the unusual, rather than using raw counts like in the previous section, we can use the TF-IDF algorithm, which assigns a high score to words that are generally uncommon but appear frequently in a specific context—in this case, words that appear disproportionately often in summaries of a given book. Can we detect evidence of a novum by finding the events that characterize plot summaries of, say, Kurt Vonnegut’s science fiction novel *Cat’s Cradle*? In fact, when we look at the most distinctive verbs in plot summaries from different platforms, we can see that the platforms use fairly similar verbs to describe the events of the novel:

	Cat’s Cradle
Goodreads Reviews	interview invent marry travel learn meet write end
Wikipedia Summary	interview cleanse marry travel learn meet inquire
GradeSaver	marry travel learn meet write tell surprise succeed
Shmoop	interview marry learn meet commit decide offer
SparkNotes	interview marry travel learn meet divide write end
Goodreads Summary	leave

Figure 4.9: The most distinctive verbs in plot summaries of *Cat’s Cradle* by platform.

Each row in the chart above shows the most distinctive verbs for the plot summary of *Cat's Cradle* on a given platform. In almost every summary except the very brief publisher's summary at the top of the book's Goodreads page, the verbs "interview" and "marry" are in the top three most distinctive verbs for the book. These words clearly mark important acts, like "realization" in *Flowers for Algernon*—in this case, "interview" refers to the narrator's task of interviewing several of the characters for a book he's writing, setting the plot in motion: "He's there to interview the charitable doctor Julian Castle, but by either coincidence or fate, he also meets a host of other characters including Mr. and Mrs. Crosby, the Mintons, and all three of the Hoenikker children" (Shmoop, "Cat's Cradle"). That so many of the plot summaries mention the act of interviewing underlines its importance to the plot—we are still, here, in the major/minor mode. Similarly, marriage appears characteristically often in the book's summaries, again highlighting a key plot event—the narrator's decision to stay involved with the novel's unconventional characters because of a compelling incentive: "He accepts the job kind of to help the poor people of San Lorenzo but mostly so he can marry the beautiful Mona Monzano" (Shmoop, "Cat's Cradle"). By studying the verbs that most distinguish plot summaries of a given book, we have another avenue for finding important events that move the plot forward. But what about events that are not necessarily important, but unusual? Can we detect a novum?

Returning to Figure 4.9 above, one specific verb stands out. It is the second-most-distinctive verb in plot summaries within Goodreads reviews, yet it doesn't appear on any other platform's summary: "invent." The plot of *Cat's Cradle* hinges on two inventions: one from our world, the atomic bomb, which Vonnegut assigns a fictional creator in the novel; and another that only exists in the world of *Cat's Cradle*, Ice-9, the novum that brings us into Vonnegut's new reality. "Invent" is the word that ties these two worlds together. The only plot summaries in

which “invent” appears are in Goodreads reviews, which refer to both inventions—the atomic bomb and Ice-9. Other platforms refer obliquely to these inventions, but don’t use “invent” explicitly. For example, the summary at the top of the Goodreads page for *Cat’s Cradle*, which comes from the book’s publisher, Penguin, begins cryptically: “Dr Felix Hoenikker, one of the founding ‘fathers’ of the atomic bomb, has left a deadly legacy to the world” (Goodreads, “Cat’s Cradle”). While the publisher’s summary refers once, enigmatically, to the book’s novum, Goodreads reviews bring invention front and center. It seems that science fiction readers on Goodreads are particularly attuned to the unusual.

As we have seen with *Cat’s Cradle*, some plot summaries pay particular attention to the unusual. But does that mean they ignore the unremarkable? What would it look like for a summary to focus on ordinary, everyday events? Put another way, instead of finding a novum, can we find an anti-novum, an event so ordinary as to be unremarkable in plot summaries? We might find the most obvious candidate in mystery and detective novels, where the murder is so commonplace that it barely merits mention. But by the standards of the major/minor mode, the murder is certainly important. When writing a plot summary of a mystery, writers must choose whether it’s more important to favor important events or unusual events. Which mode wins out?

There was only one mystery novel in the corpus, Dan Brown’s *The Da Vinci Code*. Nevertheless, it proved a compelling test case. The novel’s inciting event is the murder of Jacques Saunière, which brings professor Robert Langdon to the scene of the crime to decode the murderer’s coded message. The chart below shows the most distinctive verbs in plot summaries of *The Da Vinci Code*:

The Da Vinci Code	
Goodreads Reviews	decipher find murder tell call arrive leave star lay meet
Wikipedia Summary	summon find arrest tell shoot inform leave erase kill take
GradeSaver	find murder arrest tell call threaten arrive not fulfill
SparkNotes	find murder arrest tell shoot call trick threaten arrive
Kirkus	enlighten reach have drive flee
Goodreads Summary	decipher match receive join learn

Figure 4.10: The most distinctive verbs in plot summaries of *The Da Vinci Code* by platform.

“Murder” appears in the top ten most distinctive verbs on three of the six platforms. The platforms that mention murder are Goodreads reviews and the classroom-focused summary sites GradeSaver and SparkNotes. When they discuss the murder, they frame it as an event that sets the plot in motion:

In the Louvre late at night, the curator Jacques Sauniere is murdered, and his body is disfigured with strange signs. (GradeSaver, “The Da Vinci Code”)

The latter has been murdered by Silas, an albino monk who belongs to the Opus Dei, and is following the commands of a man he knows only as the Teacher and who is in search of a ‘keystone’ that can lead to the Holy Grail. (Goodreads user Vani)

After murdering Sauniere, Silas calls the ‘Teacher’ and tells him that, according to Sauniere, the keystone is in the Church of Saint-Sulpice in Paris. (SparkNotes, “The Da Vinci Code”)

Each of these summaries sets the scene clearly and unambiguously. “Murder” serves, in these summaries, as a starting point for the plot. For GradeSaver and SparkNotes, the explicit discussion of the murder at the outset makes sense for their audience of students looking for a thorough guide to help them pass a test. Some Goodreads reviewers follow this trend, underlining that they typically see the event of murder in the mystery novel as a necessary inclusion in plot summaries on classroom-focused platforms as well as Goodreads. In contrast, the literary-professional platforms have a roundabout way of discussing the act of murder that

doesn't bring it to the attention of my automatic event-spotting algorithm. For example, *Kirkus'* plot summary skips any explicit statement that a murder has occurred, leaving it implied: "When Harvard professor of symbology Robert Langdon—in Paris to deliver a lecture—has his sleep interrupted at two a.m., it's to discover that the police suspect he's a murderer, the victim none other than Jacques Saumière, esteemed curator of the Louvre" (*Kirkus*, "The Da Vinci Code"). As the style of this *Kirkus* summary makes clear, publishers are less concerned than classroom-focused summary sites with explicit statements and clear context, instead leaving obvious events like the act of murder stylistically elided. Literary-professional summaries like *Kirkus* and the back-of-the-book summary at the top of the Goodreads page are most concerned with bookselling, whether they're acting as arbiters of quality like *Kirkus* or simply promoters like the book's publisher. Given limited space and a self-selected audience of readers who enjoy the genre well enough to pick up the book, or at least a review of it, literary-professional summaries skip the obvious—in this case, the murder—and instead focus on what makes the book stand out among its peers.

What is unusual, then, about *The Da Vinci Code*, if it isn't the murder? Returning to the chart above of the novel's most distinctive verbs, we can see that summaries appearing on Goodreads—both in user reviews and in the publisher's summary at the top of the page—disproportionately discuss the act of "deciphering." This term refers to Langdon's task of deciphering the coded message left by the murder victim. Thus the act of deciphering occurs immediately after—and because of—the act of murder, but it is a much more unusual plot point, setting *The Da Vinci Code* apart from other murder mysteries. The publisher's summary at the top of the book's Goodreads page mentions deciphering but not murder, two events that are closely tied but only one of which is unconventional within the genre, indicating that publishers

seem to favor the unusual over the commonplace in their plot summaries. We have seen that, unlike classroom-focused platforms, publishers tend to avoid the unremarkable, as they tread the fine line between fulfilling readers' expectations for a given genre, and highlighting an unusual twist that arouses the reader's interest. But how do publishers choose which twists to include? In the next section, we will examine why deciphering in particular might have been the focus of their summary.

The Suspenseful/Spoiler Mode

As we saw with *The Da Vinci Code*, some summaries explicitly mention murder, while literary-professional summaries don't. This omission is surprising, since conventional wisdom suggests that murder sells. What is it that they emphasize over murder? Take these excerpts from three literary-professional summaries of *The Da Vinci Code*:

As their search [for the Holy Grail] moves from France to England, Neveu and Langdon are confounded by two mysterious groups—the legendary Priory of Sion, a nearly 1,000-year-old secret society whose members have included Botticelli and Isaac Newton, and the conservative Catholic organization Opus Dei. Both have their own reasons for wanting to ensure that the Grail isn't found. (*Publishers Weekly*, "The Da Vinci Code")

By the time Sophie and Langdon reach home base, everything—well, at least more than enough—has been revealed. (*Kirkus*, "The Da Vinci Code")

Unless Langdon and Neveu can decipher the labyrinthine puzzle in time, the Priory's ancient secret -- and an explosive historical truth -- will be lost forever. (Goodreads, "The Da Vinci Code")

Notably, we see the term "decipher" in the last quote from the publisher's summary of the book.

We can see, now, the context in which this summary brings up deciphering: as part of a suspenseful cliffhanger to entice readers to buy the book. In fact, all three of these summaries leave cryptic gaps in the plot to pique the reader's interest. The main mechanism they seem to

use to create suspense is timing: they repeat “by the time” and “in time,” and play with chronology to leave gaps and stop abruptly so that readers are left wondering what happens next.

In fact, most theories of suspense emphasize that it is a feeling—sometimes anxiety, sometimes pleasure—that the book creates in readers through formal means, most often in the chronological structure of the text. For Barthes, suspense is the desire “to know, to learn the origin and the end” (*The Pleasure of the Text* 10). In his theory, texts build suspense in two main ways: through the hermeneutic code, by creating unexplained enigmas in the story; and through the proairetic code, by including actions that imply further actions and provoke readers’ imaginations. Suspense, and the anxiety it produces in readers, can also emerge through the formal quality of repetition, as film scholar Eugenie Brinkema notes of the shark-attack film *Open Water*: “Anxiety has neither to do with the appearance of the sharks as threat nor with the disappearance of their absence. . . . Instead, it is the pulsation of seeing and then not seeing, not seeing and then seeing, that constitutes a form of time that is intermittent” (219). In their analyses of the formal features of chronology and repetition, literary and film scholars characterize suspense as a recognizable structural aspect of the plot with a characteristic effect on the reader.

While narratologists have developed complex theories of suspense, the popular understanding is comparatively simple. It most commonly relies on the plot diagram, a simplification of story structure that orders the plot in terms of the events that occur at the beginning, middle, and end of a story. According to the plot diagram, most stories begin with an exposition, build to a climax, and finally end with a resolution. In this model, suspense occurs in the events leading to the climax:



Figure 4.11: Diagram of the “Rising Action” section of the plot diagram from 7th-grade lesson plan (Jordan)

In the diagram above, from a 7th-grade language arts lesson on plot, the “rising action” is characterized as a series of plot events that create increasing uncertainty and fear in the reader. This structure is the way many students learn to diagram plot in classrooms, and indeed how many writing guides advise students to write summaries. Take, for example, selections from two of the top results for the Google search “How to write a plot summary”:

Identify the sections in a plot diagram. One of the more traditional ways to structure a story is to use the triangle plot diagram, also known as Freytag’s Pyramid. Freytag’s Pyramid is broken into six sections: the set up, the inciting incident, the rising action, the climax, the falling action, and the resolution. (wikiHow, “How to Write a Plot Outline”)

When summarizing the events in a story, focus on the main points of the narrative arc....The inciting incident is a singular event that ‘kicks off’ the story and leads to the major conflict within the novel. This leads to the rising action, in which the story continues to build and eventually comes to a point where the main character might have to take drastic action -- or might miss her opportunity to do this....The climax is an event that changes the course of a story, for better or worse. (Duczeminski)

As these excerpts indicate, teachers and online guides overwhelmingly suggest the structure of the plot diagram to frame plot summaries.

This structural view of plot exists in several prominent digital humanities projects as well. Matthew Jockers' software program, called *Syuzhet* after the Russian Formalist term for narrative, constructs charts for novels based on the average sentiment of words in each section of the text. His initial results tend to mimic the plot diagram, showing increasingly negative sentiment leading to the climax and an abrupt positive shift leading to the resolution (Jockers). In another project, the Stanford Literary Lab trained a neural network to predict whether a passage was suspenseful or not, using human readers' ratings as a baseline. They found that topics related to the military and physical pain predict suspenseful passages, while drawing room conversation and sentimental romanticism predict the opposite (Algee-Hewitt). Their project produces charts of suspense over time in novels, tending to show spikes in suspensefulness at the end of chapters and gradually increasing over the course of the novel. Both of these projects use the model of the plot diagram as a hypothesis to test their results against, noting novels that fit or break from its structure.

However, both of these projects use a literary corpus, looking for patterns in novelistic language that reflect the overall action of the plot. When summary-writing guides suggest that writers use the plot diagram as a template for their summary, they assume not just that novels will follow the plot diagram, but that summaries will, too. Using our plot summary corpus, we can test this hypothesis. Do summarizers follow these guides and attempt to recreate the rising and falling suspense of the plot in their summaries?

This question becomes more complicated, however, when we consider that everyone who writes a plot summary has already read the book. How can a plot summary be suspenseful if the

writer already knows the outcome? The answer to this “paradox of suspense” might come from more closely distinguishing surprise from suspense. The previous mode discussed unusual events, which rely on the element of surprise, but surprise only works once—as soon as you know what happens, it doesn’t surprise you the second time through. Unlike surprising events, suspenseful events can still be suspenseful even on the second or third reading. Theorists offer several possible explanations for how this works. Noël Carroll suggests that when we watch a film the second time, we are merely *entertaining* uncertainty about what will happen rather than truly experiencing it (“The Paradox of Suspense”). Another explanation, called the “desire-frustration theory of suspense,” suggests that “the frustration of a strong desire to affect the outcome of an imminent event is necessary and sufficient for suspense” (Smuts 281). Both of these theories help to explain the frequent presence in Goodreads reviews of “spoilers,” or plot details that reveal surprising or suspenseful moments in the book. It would seem that including a spoiler in a review meant for an audience that hasn’t yet read the book makes little sense, since the spoiler would ruin the experience for future readers. However, if readers can maintain a feeling of uncertainty despite knowing what will happen next, spoilers are not so damaging to the reader’s experience. In fact, some studies have found that spoilers don’t actually harm the reading experience (Leavitt and Christenfeld). Spoilers are common not just in Goodreads reviews but also in Wikipedia summaries, where the plot summary style guide advises, “Information should not be intentionally omitted from summaries in an effort to avoid ‘spoilers’ within the encyclopedia article. (Spoiler warnings were used early in the project but the consensus of editors was that this practice was unencyclopedic so their use has been discontinued.)” (Wikipedia, “How to Write a Plot Summary”). The common practice of including spoilers in plot summaries emphasizes the distinctive place of surprising events in the

plot summary: too integral to leave out, but not integral enough to ruin the experience of reading when revealed.

Inspired by the plot diagram, we can track the kinds of events that appear at the beginning, middle, and end of plot summaries to see how writers are wielding suspenseful events in their summaries. Each column below corresponds to an individual plot summary of Barbara Kingsolver’s *The Poisonwood Bible*, with a spike at the moment that summary uses a particular verb:




Figure 4.12: Chart showing the use of four different verbs over time in plot summaries of *The Poisonwood Bible* from seven platforms.

The four verbs I’ve highlighted above refer to suspenseful moments in the book, such as Nathan Price’s refusal to allow his family to leave Africa amid growing violence, and whom his daughter Rachel will end up marrying. The chart emphasizes that suspenseful verbs follow a clear pattern of appearing either at the beginning or the end of a summary, or not at all—but rarely in the middle. This rule is particularly strict, it seems, in literary-professional summaries from *Kirkus*, *Publishers’ Weekly*, and *The New York Times*.

We can see this general pattern again in summaries of E.B. White’s children’s novel *Charlotte’s Web*:




Figure 4.13: Chart showing the use of four different verbs over time in plot summaries of *Charlotte’s Web* on four platforms.

Again, we see that the suspenseful verbs generally appear toward the end of the summaries, but there are exceptions in Goodreads review summaries and on the classroom-focused sites GradeSaver and Shmoop. Consistently, literary-professional summaries seem to adhere most strictly to a formal rule pushing suspenseful verbs to the beginning or the end of their plot summaries.

It’s notable that these suspenseful events aren’t simply appearing in the same place they occurred in the story. The verbs in the two charts above appear in a wide variety of locations throughout the summaries rather than in a consistent order. Placement in a plot summary isn’t chronological—summary writers bring events to the beginning or end of their summary for stylistic reasons. This effect is most consistent in professional summaries from platforms most focused on attracting readers’ curiosity in order to sell the book or generate interest in their review. These results point to the existence of a formal rule that guides the creation of suspense, particularly in publisher-driven summaries.

The Emotional/Restrained Mode

Suspense is only one of the possible affective states readers experience as they read. The emotional/restrained mode—a play on the paradox of “emotional restraint”—considers emotional content in summaries more generally, from happiness and excitement to sadness and despair. Like suspense, emotional content appears frequently in reviews to attract readers’ interest and investment. It can also serve as a framework for structuring a plot summary. In fact, work in the computational and social sciences often casts emotion as a way of characterizing plot. Computer scientist Wendy Lehnert models plot as a structure whose basic unit consists of positive and negative events, defined by whether they please or displease the characters in the story (294–95). Positive states include achieving a goal and getting good news; negative states include plans failing and the end of a positive interpersonal relationship (326–27). While Lehnert’s work attempts to glean these mental states from narrative descriptions of the characters, plot summaries offer a view of readers’ mental states as they recount individual plot events. Their mental states are visible in the types of events they mention, the verbs they use to characterize those events, and where in the review they locate the events.

Nevertheless, online guides for writing plot summaries typically discourage emotional content. As a Wikipedia guide dictates, “Do not attempt to recreate the emotional impact of the work through the plot summary” (Wikipedia, “How to Write a Plot Summary”). Despite this guideline, emotional content frequently makes its way into plot summaries. Take, for example, the most frequent verbs in plot summaries of Margaret Mitchell’s romance novel *Gone with the Wind*:

	Gone with the Wind
Goodreads Reviews	love marry have rape flex tell hate not care represent
Wikipedia Summary	love marry have rape tell refuse leave not marry
GradeSaver	love marry have propose tell not care plunge refuse
SparkNotes	love marry have propose tell refuse leave vow see
Kirkus	post
Goodreads Summary	find use

Figure 4.14: The most frequent verbs in plot summaries of *Gone with the Wind* by platform.

While a majority of the platforms focus on the major event of marriage, summaries from Goodreads reviews and Wikipedia—the two platforms open to amateur writers—include the more controversial and highly emotionally-charged event of “rape.” The Wikipedia summary recounts, “He then takes her in his arms and carries her up the stairs to her bedroom, where it is strongly implied that he rapes her—or, possibly, that they have consensual sex following the argument” (Wikipedia, “Gone with the Wind”). It seems that the authors of this Wikipedia plot summary aren’t following Wikipedia’s “How To Write a Plot Summary” guide that discourages emotional content. Goodreads reviews, too, ignore advice about restraining emotional content in plot summaries:

As for it’s treatment of women, there is a point in the book where (by modern day standards) Scarlet is raped by Rhett. (Goodreads user Rishonda)

Also: Scarlett almost getting attacked and raped; GWTW’s racy social commentary, all of the men partaking in early KKK activities. (Goodreads user Fabian)

Perhaps because of the emotional intensity of such an act, each of these invocations of the rape scene is brief and ambivalent. In contrast, however, there is no mention of rape in the literary-professional summaries from GradeSaver, SparkNotes, *Kirkus*, or the publisher’s summary at the

top of the book's Goodreads page. This pattern exists for other books and genres as well, such as the science fiction novel *Ender's Game*, where the most common verbs in plot summaries are consistent across platforms except for Goodreads review summaries, which are the only summaries to use emotionally-charged verbs like "love" and "hurt." While professional summaries seem to restrain emotional content in their summaries, amateur summaries offer insight into the events that most impacted them emotionally. In the books I examined, amateur summaries from Goodreads and Wikipedia were much more likely to discuss events in the emotional mode, in contrast to the restrained mode of professional summaries.

* * *

Each platform makes use of the four modes of plot summary in pursuit of different goals. Literary-professionals, including professional reviewing platforms and publishers' back-of-the-book summaries, tend toward the modes that attract readers' interest by highlighting the most unusual and suspenseful events. Classroom-focused platforms like GradeSaver and SparkNotes attempt to depict as many of the important events as possible so that their audience of students will have a balanced understanding of the text. Wikipedia and Goodreads summaries, as the two platforms with amateur writers, are distinct in focusing on emotionally-charged events involving love and even rape.

All summaries have to pick which events they will include, and we can see what's important to the summarizer by what they choose. As we have seen, there is no objective plot summary—we can see different priorities and preferences in every platform's plot summaries, and while professional platforms tend to produce similar summaries, it's clear that professional

summary is just one mode of plot summary among many, and amateurs embrace different but equally valid theories of plot summary. Unlike the detailed attention literary scholars have paid to narrative plot, plot summary is much less studied. In fact, plot summary is a reader-response goldmine of information about how people interpret narrative plot. In the classroom, too often we think of plot summary as something to spend as little time on as possible in order to reach the more worthwhile mode of literary analysis. Studying plot summaries on their own, however, opens up a new angle from which we can measure and observe the range of possible interpretations and retellings of a single plot. The study of plot summaries offers an opportunity for a radical shift in what we consider to be important about a book. Sites like Goodreads serve as a source of alternative summaries with a different focus than professional summaries. We can deepen our current understanding of readers and how they react to books with these quantitative methods. The goal of a Goodreads plot summary isn't a well-balanced representation of the book's plot. By instead focusing on characters they identified with and emotionally-charged events, their plot summary seems more focused on describing why they liked the book to others, or more precisely, how they experienced it cognitively and affectively. Goodreads summaries offer a highly skewed version of the plot—but by studying this skew, we can see what is most important to readers.

The Goodreads Effect

Rather than taking these amateur summaries at face value, however, we should first consider how the Goodreads platform influences readers' perceptions of plot. Plot summaries on Goodreads are far from an objective glimpse into readers' minds. On the website, reviews appear

in a default organization that follows no explicit logic: these are reviews selected by a black-box algorithm that Goodreads has decided are the most useful reviews for readers. Since many readers decide whether or not to read a book based on reading the first few reviews, this organization has a significant impact on whether readers choose to read that book, and frames their perception of the book's plot before they start reading. By comparing the way these default reviews discuss plot to the way non-featured reviews discuss plot, we can clarify the characteristics of plot that Goodreads deems most appropriate for its users to find in book reviews.

Goodreads immediately presents users with thirty algorithmically-sorted reviews, with the option to continue browsing through up to ten pages of results. If the reviews appeared randomly, we would expect each page of reviews to exhibit similar characteristics—for example, the use of verbs from page to page should remain roughly similar. While most verbs indeed seem randomly distributed across each page of reviews, in several books, certain verbs stood out. For example, when you observe the frequency at which the words “love” and “hate” appear in reviews, for most books in the corpus, the words appear no more than four times on each page of review results. But for three books, the two words appeared unusually often on the first page of reviews:




Figure 4.15: Use of the verbs “hate” and “love” in Goodreads review plot summaries of books in this project’s corpus.

For a single book—*Gone with the Wind*—love and hate are over-represented on the first page of reviews. No other page has more than four uses of the words; the first page has thirteen uses of “love” and eight uses of “hate.” These mostly come from a few reviews that use the terms repeatedly: for example, “Yet, that is what makes Melanie different from Scarlett. One woman manipulated and hated another for the love of a man, where another one loved and cared out of the kindness of her heart” (Goodreads user Lina). Again, my methods filtered out statements like “I love this book,” so these uses of “love” and “hate” reference the actions of characters. Scarlett and Rhett do a lot of loving and hating, but this still doesn’t explain why Goodreads pushes the summaries that recount these actions to the first page of results.

We can see a similar emphasis when we look at the verb “die” in the young adult bestsellers *Allegiant* and *The Hunger Games*:




Figure 4.16: Use of the verb “die” in Goodreads review plot summaries of books in this project’s corpus.

Reviews of *The Hunger Games*, the first book in a trilogy by Suzanne Collins, stand out for their high use of the verb “die” on the first page of results. The publisher’s summary at the top of the book’s Goodreads page starts with the bold lines: “Winning will make you famous. Losing means certain death”—and it ends similarly: “When sixteen-year-old Katniss Everdeen steps forward to take her sister's place in the games, she sees it as a death sentence. But Katniss has been close to death before. For her, survival is second nature” (Goodreads, “The Hunger Games”). When we see, then, that reviews on the first page are far more likely to include the verb “die,” could it be because the publisher’s summary primes readers to think about death as they write their own summaries? It’s true that *The Hunger Games* features more deaths than the average young adult novel. But like “love” and “hate” in *Gone with the Wind*, it’s still unclear why these verbs would become so over-represented on the first page of results and not later

pages. In fact, the effect, like that in *Gone with the Wind*, is dramatic on the first page but not any of the following pages, indicating an all-or-nothing push to place these reviews on the first page of results.

The effect is even more pronounced for *Allegiant*, the third book of Veronica Roth's *Divergent* trilogy. The publisher's summary at the top of the Goodreads page doesn't explicitly talk about death, but instead life: "Perhaps beyond the fence, she and Tobias will find a simple new life together, free from complicated lies, tangled loyalties, and painful memories" (Goodreads, "Allegiant"). The reason so many reviews of *Allegiant* use the verb "die" is that—spoiler alert—the main character, Tris, dies. One review complains, "i was with it every step of the way.. until tris died" (Goodreads user Emma). This review—and the other reviews that mention Tris' death—contains a significant spoiler, raising the question of why Goodreads is bringing these controversial reviews to the first page of results. The priority, it seems, is less on preserving any sense of surprise for potential readers, and more on emphasizing the central themes of the book, allowing readers to express their distaste for key plot events, or surfacing controversial aspects of the book so that readers can make their reading decision based on that information.

For both of these books with disproportionate use of the verb "die" on the first page of results, the conversation seems inspired by the publisher's summary at the top of the book's Goodreads page. Perhaps the algorithm is meant to highlight reviews that resemble the publisher's summary in language or themes, or perhaps reviews that happen to mention events like "love," "hate," and "die" naturally receive the most attention on the site, bringing them to the top. In either case, it's clear that for certain books, Goodreads is making its distinction from

other platforms even more pronounced by bringing reviews that emphasize certain plot events to the first page.

There are more types of plot summary than we study in classrooms. The growth of social reading platforms in particular has brought many more types of summary into circulation, written not just by literary professionals or publishers, but by diverse readers. Goodreads is a valuable place to see readers trying out new modes and unofficial, alternative perspectives on plot. But at the same time, just as classroom-focused sites encourage broad-based coverage of the plot and publishers emphasize unusual and suspenseful events, Goodreads highlights dramatic events like “love” and “death” because of its unique set of priorities as a platform that achieves its success through passionate user engagement. By more explicitly characterizing the many possible modes of plot summary and their competing purposes and audiences, we can achieve a more nuanced understanding of the way platforms and algorithms are shaping our understanding of narrative plot.

Conclusion

To truly grapple with the age of the algorithm and our growing entanglement with computational cultural processes, we need to take action as scholars, teachers, and most of all performers of humanistic inquiry. We need an experimental humanities, a set of strategies for direct engagement with algorithmic production and scholarship, drawing on theories of improvisation and experimental investigation to argue that a culture of process, of algorithmic production, requires a processual criticism that is both reflexive and playful. (Finn, *What Algorithms Want* 13)

Perhaps this passage helps to explain why this project looks so different from other literary dissertations. I have attempted to answer Ed Finn’s call for an “experimental humanities” that reflexively engages with social reading platforms by analyzing the algorithms they employ using algorithmic methods. To return to my discussion of models in the introduction, algorithms are a type of model that enacts its view of the world at the same time as it describes it. Algorithms are abstracted from their creators, which makes it difficult to assign responsibility for their results, an especially thorny problem when their results favor or disfavor a certain demographic. The problem is that algorithms are black boxes, composed of hundreds of thousands of lines of proprietary code within far-reaching systems applied distinctly across millions of users. Without access to the underlying source code, the only way to determine how an algorithm works is to reverse-engineer it. ProPublica has conducted some of the most striking investigations of algorithmic bias, including a project revealing how Facebook’s advertising algorithms allowed employers to discriminate against older workers (Angwin, Scheiber, et al.), and criminal sentencing algorithms assigned harsher penalties to black defendants (Angwin, Larson, et al.). While my examination of social reading websites revealed nothing quite so disturbing, I did see a clear “algorithmic effect” pushing the content of book reviews in directions favorable to the tech companies running the online platforms publishing the reviews.

In each of my experiments, I have shown measurable changes in readers' book reviews due to the algorithmic design of social reading websites. In my analysis of genre on Goodreads, I discovered that the review ranking algorithm brings reviews that are more coherent—that use a shared language—to the top of the page. The result is an exaggerated conformity in the reviews users read, leading them to believe that everyone reacts in similar ways to books, and reducing the range of different perspectives users come across. In my analysis of plot summaries, I found that a platform's goals and audience influence even the characters and events that appear in that platform's plot summaries, revealing several different modes of writing plot summary, each with its own distinct characteristics. The methods I used in each examination—unsupervised clustering, comparative classification, and natural language processing—allow for complexity and ambiguity in the results they provide. None of these methods has an objectively correct output; rather, they present results that require interpretation and analysis. Each method reverse-engineers the Goodreads algorithm to find out, as Finn would say, what that algorithm wants: in the case of Goodreads, for readers to speak the same language, and to be passionately invested in the plot.

Traditional literary genre categories need reimagining in the digital age to account for the ways both readers and tech companies have begun to break the rules and create new types of categories to organize their books and their communities. Plot, too, has become a way for companies to commodify readers through plot summaries that use the language, characters, and events companies deem most likely to encourage readers to buy the book. While I did find some evidence showing that readers at times subvert these corporate reformulations and assert their own new genre categories and perspectives on plot, it's clear that these impulses have too little outlet within the corporate platforms that dominate the book reviewing landscape. These

platforms promise a democratic and consumer-focused medium for users, regardless of their expertise, to express their responses to books. At the same time, social reading platforms cultivate conformity in these responses to drive their underlying goals of increased user engagement with advertisements in the service of increased profit.

Using computational methods to reach these conclusions is far from a repudiation of the traditional methods of literary studies. Instead, I hope that my project has called attention to the moments when algorithmic methods fall short. Given, for example, Google's mission to "organize the world's information" via algorithms, humanistic methods allow us to ask whether that's a meaningful, feasible, or even desirable goal. The humanities provides the tools to realize that this goal was never possible to begin with, whether by questioning the objectivity of taxonomies or even tech companies' conception of "information," which is too often a narrow sub-category of the personal data that is relevant to advertisers. In short, humanistic methods can account for ambiguity and complexity in ways that computational methods cannot. However, the undeniable power and influence tech companies wield means that literary scholars cannot simply call their algorithms biased and be done. Instead, we have to use all the tools at our disposal—including, at times, tech companies' own algorithms—to determine the concrete effect they've had, and continue to have, on how we read.

Appendix: List of Texts in Plot Summary Corpus

A Connecticut Yankee in King Arthur's Court by Mark Twain (1889)

A Wrinkle in Time by Madeleine L'Engle (1962)

"A&P" by John Updike (1961)

Allegiant by Veronica Roth (2013)

"Bartleby the Scrivener" by Herman Melville (1853)

Blood Music by Greg Bear (1983)

Cat's Cradle by Kurt Vonnegut (1963)

Catch-22 by Joseph Heller (1961)

Catching Fire by Suzanne Collins (2009)

Cathedral by Raymond Carver (1983)

Ceremony by Leslie Marmon Silko (1977)

Charlotte's Web by E. B. White (1952)

Death of a Salesman by Arthur Miller (1949)

Divergent by Veronica Roth (2011)

Do Androids Dream of Electric Sheep? by Philip K. Dick (1968)

Ella Enchanted by Gail Carson Levine (1997)

Ender's Game by Orson Scott Card (1985)

Fahrenheit 451 by Ray Bradbury (1953)

Flowers for Algernon by Daniel Keyes (1966)

Franny and Zooey by J. D. Salinger (1961)

Gone with the Wind by Margaret Mitchell (1936)

Holes by Louis Sachar (1998)

I, Robot by Isaac Asimov (1950)

Jurassic Park by Michael Crichton (1990)

Mockingjay by Suzanne Collins (2010)

My Year of Meats by Ruth Ozeki (1998)

No Country for Old Men by Cormac McCarthy (2005)

One Flew Over the Cuckoo's Nest by Ken Kesey (1962)

Oryx and Crake by Margaret Atwood (2003)

Slaughterhouse-Five by Kurt Vonnegut (1969)

Something Wicked This Way Comes by Ray Bradbury (1962)

Stories of Your Life and Others by Ted Chiang (2002)

Stranger in a Strange Land by Robert A. Heinlein (1961)

Super Sad True Love Story by Gary Shteyngart (2010)

The Adventures of Huckleberry Finn by Mark Twain (1885)

The Adventures of Tom Sawyer by Mark Twain (1876)

The Brief Wondrous Life of Oscar Wao by Junot Díaz (2007)

The Da Vinci Code by Dan Brown (2003)

The Giver by Lois Lowry (1993)

The Great Gatsby by F. Scott Fitzgerald (1925)

The Hunger Games by Suzanne Collins (2008)

The Joy Luck Club by Amy Tan (1989)

The Jungle by Upton Sinclair (1906)

The Left Hand of Darkness by Ursula K. Le Guin (1969)

The Poisonwood Bible by Barbara Kingsolver (1998)

The Road by Cormac McCarthy (2006)

The Windup Girl by Paolo Bacigalupi (2009)

“The Yellow Wallpaper” by Charlotte Perkins Gilman (1892)

Their Eyes Were Watching God by Zora Neale Hurston (1937)

To Kill a Mockingbird by Harper Lee (1960)

Winesburg, Ohio by Sherwood Anderson (1919)

Wool by Hugh Howey (2011)

Bibliography

- Abbott, H. Porter. *The Cambridge Introduction to Narrative*. Cambridge University Press, 2008.
- Algee-Hewitt, Mark, Sarah Allison, et al. "Canon/Archive: Large-Scale Dynamics in the Literary Field." *Stanford Literary Lab Pamphlet 11*, Jan. 2016, pp. 1–13.
- Algee-Hewitt, Mark, Ryan Heuser, et al. "On Paragraphs. Scale, Themes, and Narrative Form." *Stanford Literary Lab Pamphlet 10*, Oct. 2015, pp. 1–22.
- Algee-Hewitt, Mark. "Virtual Readers." *Stanford Literary Lab*, 25 Aug. 2016, <https://litlab.stanford.edu/virtual-readers/>.
- Allison, Sarah, Ryan Heuser, et al. "Quantitative Formalism: An Experiment." *Stanford Literary Lab Pamphlet 1*, Jan. 2011, pp. 1–26.
- Allison, Sarah, Marissa Gemma, et al. "Style at the Scale of the Sentence." *Stanford Literary Lab Pamphlet 5*, June 2013, pp. 1–29.
- Altman, Rick. *Film/Genre*. British Film Institute, 1999.
- Angwin, Julia, Noam Scheiber, et al. "Dozens of Companies Are Using Facebook to Exclude Older...." *ProPublica*, 20 Dec. 2017, <https://www.propublica.org/article/facebook-ads-age-discrimination-targeting>.
- Angwin, Julia, Jeff Larson, et al. "What Algorithmic Injustice Looks Like in Real Life." *ProPublica*, 25 May 2016, <https://www.propublica.org/article/what-algorithmic-injustice-looks-like-in-real-life>.
- Anny. "Anny (Indonesia)'s Review of My Year of Meats." *Goodreads*, 3 Feb. 2016, <https://www.goodreads.com/review/show/1536755272>.

- Archer, Jodie. *Reading the Bestseller: An Analysis of 20,000 Novels*. Stanford University, 2014, <http://search.proquest.com/pqdtglobal/docview/1550513714/1091F7F7B2244803PQ/1>.
- Aristotle. *The Poetics of Aristotle: Translation and Commentary*. Translated by Stephen Halliwell, University of North Carolina Press, 1987.
- Asenjo, Bill. "How to Write a Book Review." *Writing-World.Com*, 2002, <http://www.writing-world.com/freelance/asenjo.shtml>.
- Augst, Thomas. *The Clerk's Tale: Young Men and Moral Life in Nineteenth-Century America*. University of Chicago Press, 2003.
- Austin, John. "United States, 1780-1850." *The Novel, Volume 1: History, Geography, and Culture*, edited by Franco Moretti, Princeton University Press, 2006, pp. 455–65.
- Ballard, J. G. *A User's Guide to the Millennium: Essays and Reviews*. Picador, 1996.
- Bamman, David, et al. "Learning Latent Personas of Film Characters." *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, 2013, pp. 352–61, <http://www.anthology.aclweb.org/P/P13/P13-1035.pdf>.
- Bamman, David, and Noah A. Smith. "Unsupervised Discovery of Biographical Structure from Text." *Transactions of the Association for Computational Linguistics*, vol. 2, 2014, pp. 363–76.
- Bandari, Roja, et al. *What Do You Do With A Million Readers?* 2015, http://dh2015.org/abstracts/xml/TANGHERLINI_Timothy_Roland_What_Do_You_Do_With_A_/TANGHERLINI_Timothy_Roland_What_Do_You_Do_With_A_Millio.html.
- Barnes, Brooks. "Attacked by Rotten Tomatoes." *The New York Times*, 7 Sept. 2017, <https://www.nytimes.com/2017/09/07/business/media/rotten-tomatoes-box-office.html>.

- Barthes, Roland. "An Introduction to the Structural Analysis of Narrative." *New Literary History*, translated by Lionel Duisit, vol. 6, no. 2, 1975, pp. 237–72.
- . *S/Z: An Essay*. Translated by Richard Miller, Blackwell, 1974.
- . *The Pleasure of the Text*. Translated by Richard Miller, Macmillan, 1975.
- Baym, Nina. *Novels, Readers, and Reviewers: Responses to Fiction in Antebellum America*. Cornell University Press, 1987.
- Black, John B., and Robert Wilensky. "An Evaluation of Story Grammars." *Cognitive Science*, vol. 3, no. 3, July 1979, pp. 213–29.
- Blair, Amy L. *Reading Up: Middle-Class Readers and the Culture of Success in the Early Twentieth-Century United States*. Temple University Press, 2012.
- Booth, Wayne C. *The Rhetoric of Fiction*. University of Chicago Press, 1961.
- Bould, M., and S. Vint. "There Is No Such Thing as Science Fiction." *Reading Science Fiction*, edited by James Gunn et al., Palgrave Macmillan, 2008, pp. 43–51, <http://eprints.uwe.ac.uk/8057/>.
- Brinkema, Eugenie. *The Forms of the Affects*. Duke University Press Books, 2014.
- Brizee, Allen. "Writing a Book Review." *Purdue Online Writing Lab (OWL)*, 13 Jan. 2016, <https://owl.english.purdue.edu/owl/resource/704/01/>.
- Brodhead, Richard H. *Cultures of Letters: Scenes of Reading and Writing in Nineteenth-Century America*. University of Chicago Press, 1993.
- Brooks, Cleanth. "The New Criticism." *The Sewanee Review*, vol. 87, no. 4, 1979, pp. 592–607.
- Brooks, Peter. *Reading for the Plot: Design and Intention in Narrative*. Knopf, 1984.
- Bruns, David. "Review of Flowers for Algernon." *Goodreads*, 22 Dec. 2013, <https://www.goodreads.com/review/show/797119798>.

- Buell, Lawrence. *The Dream of the Great American Novel*. Belknap Press, 2014.
- Cait. "A Question about How People Use Shelves." *Goodreads*, 6 Jan. 2009,
<https://www.goodreads.com/topic/show/91679-a-question-about-how-people-use-shelves>.
- Cantor, Geoffrey, and Sally Shuttleworth, editors. *Science Serialized: Representations of the Sciences in Nineteenth-Century Periodicals*. MIT Press, 2004.
- Carissa. "Carissa's Review of Sidebarred." *Goodreads*, 30 Apr. 2016,
<https://www.goodreads.com/review/show/1625393923>.
- Carroll, Noel. "The Paradox of Suspense." *Beyond Aesthetics: Philosophical Essays*, Cambridge University Press, 2001, pp. 254–70.
- . *The Philosophy of Horror: Or, Paradoxes of the Heart*. Routledge, 1990.
- Castle, Terry. "The Gothic Novel." *The Cambridge History of English Literature, 1660-1780*, edited by John Richetti, Cambridge University Press, 2005, pp. 673–706.
- Cavallo, Guglielmo, and Roger Chartier. *A History of Reading in the West*. Translated by Lydia Cochrane, University of Massachusetts Press, 2003.
- Chartier, Roger. *The Order of Books: Readers, Authors, and Libraries in Europe Between the 14th and 18th Centuries*. Stanford University Press, 1994.
- Chase, Richard. *The American Novel and Its Tradition*. Johns Hopkins UP, 1980.
- Chatman, Seymour. *Story and Discourse: Narrative Structure in Fiction and Film*. Cornell University Press, 1978.
- Cheng, Emily. "Meat and the Millennium: Transnational Politics of Race and Gender in Ruth Ozeki's *My Year of Meats*." *Journal of Asian American Studies*, vol. 12, no. 2, May 2009, pp. 191–220.

- Chun, Wendy Hui Kyong. *Control and Freedom: Power and Paranoia in the Age of Fiber Optics*. The MIT Press, 2008.
- . *Programmed Visions: Software and Memory*. The MIT Press, 2011.
- CliffsNotes. "Their Eyes Were Watching God: Book Summary." *CliffsNotes*,
<https://www.cliffsnotes.com/literature/t/their-eyes-were-watching-god/book-summary>.
Accessed 9 Dec. 2017.
- Cordón García, José Antonio, et al., editors. *Social Reading: Platforms, Applications, Clouds and Tags*. Chandos Publishing, 2013.
- Cottino-Jones, Marga. "Literary-Critical Developments in Sixteenth- and Seventeenth-Century Italy." *The Cambridge History of Literary Criticism, Vol. 3: The Renaissance*, edited by Glyn P. Norton, Cambridge University Press, 1999, pp. 566–77.
- Dames, Nicholas. "Theories of the Novel." *The Cambridge History of Literary Criticism: Volume 6, The Nineteenth Century, c.1830-1914*, edited by M. A. R. Habib, Cambridge University Press, 2013, pp. 506–23.
- Dandeneau, Lolly K. "Review of Middle Age." *Goodreads*, 16 Apr. 2012,
<https://www.goodreads.com/review/show/313597491>.
- David. "Review of The Three-Body Problem." *Goodreads*, 11 Nov. 2014,
<https://www.goodreads.com/review/show/1103706545>.
- Derrida, Jacques. "The Law of Genre." *Critical Inquiry*, translated by Avital Ronell, vol. 7, no. 1, Oct. 1980, pp. 55–81.
- Dimock, Wai Chee. "Introduction: Genres as Fields of Knowledge." *PMLA*, vol. 122, no. 5, 2007, pp. 1377–88.
- Douglas, Ann. "Soft-Porn Culture." *New Republic*, Aug. 1980, pp. 25–29.

- Duczeminski, Matt. "How to Write a Plot Summary." *The Pen and The Pad*,
<https://penandthepad.com/write-plot-summary-6535877.html>. Accessed 19 Dec. 2017.
- Elfenbein, Andrew. "Cognitive Science and the History of Reading." *PMLA*, vol. 121, no. 2,
2006, pp. 484–502.
- Emma. "Review of Allegiant." *Goodreads*,
<https://www.goodreads.com/review/show/744501240>. Accessed 16 Jan. 2018.
- English, James F. "Now, Not Now Counting Time in Contemporary Fiction Studies." *Modern
Language Quarterly*, vol. 77, no. 3, Sept. 2016, pp. 395–418, doi:10.1215/00267929-
3570667.
- Fabian. *Review of Gone with the Wind*. <https://www.goodreads.com/review/show/181273653>.
Accessed 16 Jan. 2018.
- Farr, Cecilia Konchar. *Reading Oprah: How Oprah's Book Club Changed the Way America
Reads*. State University of New York Press, 2005.
- Felski, Rita. *Uses of Literature*. Blackwell Publishing, 2008.
- Finn, Ed. "Becoming Yourself: The Afterlife of Reception." *Stanford Literary Lab Pamphlet 3*,
Sept. 2011, pp. 1–25.
- . *The Social Lives of Books: Literary Networks in Contemporary American Fiction*. Stanford
University, May 2011.
- . *What Algorithms Want: Imagination in the Age of Computing*. MIT Press, 2018.
- Fish, Stanley. *Is There a Text in This Class? The Authority of Interpretive Communities*. Harvard
University Press, 1980.
- Fishelov, David. *Metaphors of Genre: The Role of Analogies in Genre Theory*. Pennsylvania
State University Press, 1993, <http://www.psupress.org/books/titles/0-271-02325-2.html>.

- Flood, Alison. "Amazon Purchase of Goodreads Stuns Book Industry." *The Guardian*, 2 Apr. 2013, <https://www.theguardian.com/books/2013/apr/02/amazon-purchase-goodreads-stuns-book-industry>.
- Foucault, Michel. *The Birth of Biopolitics: Lectures at the Collège de France, 1978-1979*. Edited by Michel Senellart, Translated by Graham Burchell, Palgrave Macmillan, 2010.
- Frye, Northrop. *Anatomy of Criticism: Four Essays*. Princeton University Press, 1957.
- Funderburg, Lise. "My Year of Meats." *The New York Times*, 26 July 1998, <http://www.nytimes.com/books/98/07/26/bib/980726.rv140750.html>.
- Genette, Gérard. *Narrative Discourse Revisited*. Translated by Jane E. Lewin, Cornell University Press, 1988.
- Goodreads. "Allegiant." *Goodreads*, https://www.goodreads.com/work/best_book/15524549-allegiant. Accessed 16 Jan. 2018.
- . "Cat's Cradle." *Goodreads*, https://www.goodreads.com/work/best_book/1621115-cat-s-cradle. Accessed 17 Dec. 2017.
- . "FAQs on Amazon's Acquisition of Goodreads." *Goodreads*, 22 May 2013, <https://www.goodreads.com/topic/show/1337250-faqs-on-amazon-s-acquisition-of-goodreads>.
- . "Help Topic: Description." *Goodreads*, <https://www.goodreads.com/help/show/404>. Accessed 14 Dec. 2017.
- . "Historical Fictionistas." *Goodreads*, <https://www.goodreads.com/group/show/22454-historical-fictionistas>. Accessed 18 July 2017.
- . "The Da Vinci Code." *Goodreads*, https://www.goodreads.com/work/best_book/2982101-the-da-vinci-code. Accessed 17 Dec. 2017.

---. "The Hunger Games." *Goodreads*, https://www.goodreads.com/work/best_book/2792775-the-hunger-games. Accessed 16 Jan. 2018.

Gordon, Joan. "Winsome Ghosts in the Machine: Joan Gordon's 'The Lifecycle of Software Objects.'" *Los Angeles Review of Books*, Apr. 2012, <https://lareviewofbooks.org/article/winsome-ghosts-in-the-machine-joan-gordons-the-lifecycle-of-software-objects/>.

GradeSaver. "Flowers for Algernon Summary." *GradeSaver*, <http://www.gradesaver.com/flowers-for-algernon/study-guide/summary>. Accessed 17 Dec. 2017.

---. "No Country for Old Men Summary." *GradeSaver*, <http://www.gradesaver.com/no-country-for-old-men/study-guide/summary>. Accessed 17 Dec. 2017.

---. "The Da Vinci Code Summary." *GradeSaver*, <http://www.gradesaver.com/the-da-vinci-code/study-guide/summary>. Accessed 17 Dec. 2017.

Greimas, Algirdas Julien. *On Meaning: Selected Writings in Semiotic Theory*. Translated by Frank Collins and Paul Perron, University of Minnesota Press, 1987.

Grusin, Richard. "The Dark Side of Digital Humanities: Dispatches from Two Recent MLA Conventions." *Differences*, vol. 25, no. 1, Jan. 2014, pp. 79–92.

Guillory, John. "The Ethical Practice of Modernity: The Example of Reading." *The Turn to Ethics*, edited by Marjorie Garber et al., Routledge, 2000.

Halliwell, Stephen. "Aristotle's Poetics." *The Cambridge History of Literary Criticism, Vol. 1: Classical Criticism*, edited by George A. Kennedy, Cambridge University Press, 1989, pp. 149–83.

- Hayles, N. Katherine. "How We Read: Close, Hyper, Machine." *ADE Bulletin*, no. 150, 2010, pp. 62–79.
- . "Hyper and Deep Attention: The Generational Divide in Cognitive Modes." *Profession 2007*, 2007, pp. 187–99.
- Heise, Ursula K. "Martian Ecologies and the Future of Nature." *Twentieth-Century Literature*, vol. 57, no. 3–4, Sept. 2011, pp. 447–71.
- . "Reduced Ecologies: Science Fiction and the Meanings of Biological Scarcity." *European Journal of English Studies*, vol. 16, no. 2, Aug. 2012, pp. 99–112.
- . "The Android and the Animal." *PMLA*, vol. 124, no. 2, Mar. 2009, pp. 503–10.
- Helene, Mary. "Review of Wife of the Gods." *Goodreads*, 26 Nov. 2014, <https://www.goodreads.com/review/show/1116182746>.
- Hickey, Walt. "'Ghostbusters' Is A Perfect Example Of How Internet Movie Ratings Are Broken." *FiveThirtyEight*, 14 July 2016, <http://fivethirtyeight.com/features/ghostbusters-is-a-perfect-example-of-how-internet-ratings-are-broken/>.
- . "Men Are Sabotaging The Online Reviews Of TV Shows Aimed At Women." *FiveThirtyEight*, 18 May 2016, <http://fivethirtyeight.com/features/men-are-sabotaging-the-online-reviews-of-tv-shows-aimed-at-women/>.
- Hickey, Walt, and Dhrumil Mehta. "Al Gore's New Movie Exposes The Big Flaw In Online Movie Ratings." *FiveThirtyEight*, 7 Sept. 2017, <https://fivethirtyeight.com/features/al-gores-new-movie-exposes-the-big-flaw-in-online-movie-ratings/>.
- Horton, Russell, et al. "Mining Eighteenth Century Ontologies: Machine Learning and Knowledge Classification in the Encyclopédie." *Digital Humanities Quarterly*, vol. 3, no.

- 2, 2009. *Digital Humanities Quarterly*,
<http://www.digitalhumanities.org/dhq/vol/3/2/000044/000044.html>.
- Houser, Heather. *Description as Data in Literary Studies*. 2014,
<http://arcade.stanford.edu/blogs/description-data-literary-studies>.
- Hugo, Victor. "Preface." *Cromwell*, translated by Tilottama Rajan, Garnier-Flammarion, 1968.
- Hunter, Jane H. *How Young Ladies Became Girls: The Victorian Origins of American Girlhood*.
 Yale University Press, 2002.
- Iser, Wolfgang. *How to Do Theory*. Blackwell, 2005.
- J Austill. "Review of The Hound of the Baskervilles." *Goodreads*, 24 Apr. 2016,
<https://www.goodreads.com/review/show/1620009568>.
- Jackson, H. J. *Romantic Readers: The Evidence of Marginalia*. Yale University Press, 2005.
- Jameson, Fredric. *The Political Unconscious*. Routledge, 1981.
- Jauss, Hans Robert. *Toward an Aesthetic of Reception*. University of Minnesota Press, 1982.
- jenjn79. "A Question about How People Use Shelves." *Goodreads*, 6 Jan. 2009,
<https://www.goodreads.com/topic/show/91679-a-question-about-how-people-use-shelves>.
- Jockers, Matthew. "Introduction to the Syuzhet Package." *Introduction to the Syuzhet Package*,
 13 Dec. 2017, <https://cran.r-project.org/web/packages/syuzhet/vignettes/syuzhet-vignette.html>.
- Jockers, Matthew L., and David Mimno. "Significant Themes in 19th-Century Literature."
Poetics, vol. 41, no. 6, Dec. 2013, pp. 750–69.
- Jordan, Lori. "Elements of Suspense in Literature." *Prezi*, 17 Oct. 2016,
<https://prezi.com/wyt6zmamrm9w/elements-of-suspense-in-literature/>.

- Kayman, Martin A. "The Short Story from Poe to Chesterton." *The Cambridge Companion to Crime Fiction*, edited by Martin Priestman, English Language edition, Cambridge UP, 2003, pp. 41–58.
- Kirkus. "Review of Ebon City by Leo Featherstone." *Kirkus Review*, Feb. 2014, <https://www.kirkusreviews.com/book-reviews/leo-featherstone/ebon-city/>.
- . "Review of Nightmare Flower by Elizabeth Engstrom." *Kirkus Review*, Aug. 1992. www.kirkusreviews.com, <https://www.kirkusreviews.com/book-reviews/elizabeth-engstrom/nightmare-flower/>.
- . "Review of The Da Vinci Code by Dan Brown." *Kirkus Review*, Jan. 2003, <https://www.kirkusreviews.com/book-reviews/dan-brown/the-da-vinci-code/>.
- . "Review of The Night Country by Stewart O’Nan." *Kirkus Review*, Aug. 2003, <https://www.kirkusreviews.com/book-reviews/stewart-onan/the-night-country/>.
- László, János. *The Science of Stories: An Introduction to Narrative Psychology*. Routledge, 2008.
- Lauren. "Review of Worked Up." *Goodreads*, 1 Aug. 2016, <https://www.goodreads.com/review/show/1715879737>.
- Leavitt, Jonathan D., and Nicholas J. S. Christenfeld. "Story Spoilers Don’t Spoil Stories." *Psychological Science*, vol. 22, no. 9, Sept. 2011, pp. 1152–54.
- Lehnert, Wendy G. "Plot Units and Narrative Summarization." *Cognitive Science*, vol. 4, 1981, pp. 293–331.
- Lehuu, Isabelle. *Carnival on the Page: Popular Print Media in Antebellum America*. The University of North Carolina Press, 2000.

- Levine, Lawrence. *Highbrow/Lowbrow: The Emergence of Cultural Hierarchy in America*. Harvard University Press, 1990.
- Leyboldt, Günter. "Degrees of Public Relevance: Walter Scott and Toni Morrison." *Modern Language Quarterly*, vol. 77, no. 3, Sept. 2016, pp. 369–93.
- Lina. "Review of *Gone with the Wind*." *Goodreads*,
<https://www.goodreads.com/review/show/290651709>. Accessed 16 Jan. 2018.
- Lipking, Lawrence. "Literary Criticism and the Rise of National Literary History." *The Cambridge History of English Literature, 1660-1780*, edited by John Richetti, Cambridge University Press, 2005, pp. 471–97.
- Lisa. "A Question about How People Use Shelves." *Goodreads*, 7 Jan. 2009,
<https://www.goodreads.com/topic/show/91679-a-question-about-how-people-use-shelves>.
- Liu, Alan. "From Reading to Social Computing." *Literary Studies in the Digital Age*, 2013,
<https://dlsanthology.commons.mla.org/from-reading-to-social-computing/>.
- . "The Big Bang of Online Reading." *Advancing Digital Humanities: Research, Methods, Theories*, edited by Paul Longley Arthur and Katherine Bode, Palgrave Macmillan, 2014, pp. 113–29.
- . *The Laws of Cool: Knowledge Work and the Culture of Information*. University of Chicago Press, 2004.
- Long, Elizabeth. *Book Clubs: Women and the Uses of Reading in Everyday Life*. University of Chicago Press, 2003.
- . "The Cultural Meaning of Concentration in Publishing." *Book Research Quarterly*, vol. 1, no. 4, Winter -86 1985, pp. 3–27.

- Machor, James L., editor. *Readers in History: Nineteenth-Century American Literature and the Contexts of Response*. Johns Hopkins University Press, 1992.
- . *Reading Fiction in Antebellum America: Informed Response and Reception Histories, 1820-1865*. Johns Hopkins University Press, 2011.
- Manning, Christopher, et al. "The Stanford CoreNLP Natural Language Processing Toolkit." *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, 2014, pp. 55–60.
- McGurl, Mark. "Everything and Less: Fiction in the Age of Amazon." *Modern Language Quarterly*, vol. 77, no. 3, Sept. 2016, pp. 447–71.
- . *The Program Era: Postwar Fiction and the Rise of Creative Writing*. Harvard University Press, 2011.
- McHenry, Elizabeth. *Forgotten Readers: Recovering the Lost History of African American Literary Societies*. Duke University Press, 2002.
- McLachlan, Cameron. "'Too Perfectly Historical for Words': Reading Sociably at the Katherine Mansfield Birthplace/Te Puakitanga." *Journal of New Zealand Literature*, no. 27, 2009, pp. 70–89.
- Melissa. "Melissa's Review of Gone with the Wind." *Goodreads*, 7 Nov. 2007, <https://www.goodreads.com/review/show/8815697>.
- Miller, Carolyn R. "Genre as Social Action." *Quarterly Journal of Speech*, vol. 70, no. 2, May 1984, pp. 151–67.
- Miller, Carolyn R., and Dawn Shepherd. "Blogging as Social Action: A Genre Analysis of the Weblog." *Into the Blogosphere*, 2004, <http://conservancy.umn.edu/handle/11299/172818>.

- Miller, Laura J. *Reluctant Capitalists: Bookselling and the Culture of Consumption*. University Of Chicago Press, 2007.
- Mingee, Robert. "Review of *The Night Country*." *Goodreads*, 25 Sept. 2014, <https://www.goodreads.com/review/show/1064332241>.
- Moretti, Franco. "Conjectures on World Literature." *New Left Review*, no. 1, Feb. 2000, pp. 54–68.
- . *Distant Reading*. Verso, 2013.
- . *Graphs, Maps, Trees: Abstract Models for Literary History*. Verso, 2007.
- . "Narrative Markets, ca. 1850." *Review (Fernand Braudel Center)*, vol. 20, no. 2, 1997, pp. 151–74.
- . "Network Theory, Plot Analysis." *Stanford Literary Lab Pamphlet 2*, May 2011, pp. 1–31.
- . *The Bourgeois: Between History and Literature*. Verso, 2013.
- Nakamura, Lisa. *Digitizing Race: Visual Cultures of the Internet*. University of Minnesota Press, 2007, <https://www.upress.umn.edu/book-division/books/digitizing-race>.
- . "'Words with Friends': Socially Networked Reading on Goodreads." *PMLA*, vol. 128, no. 1, Jan. 2013, pp. 238–43, doi:10.1632/pmla.2013.128.1.238.
- Ngai, Sianne. *Ugly Feelings*. Harvard University Press, 2005.
- North, Michael. *Reading 1922: A Return to the Scene of the Modern*. Oxford University Press, 1999.
- Norton, Glyn P. "Theories of Prose Fiction in Sixteenth-Century France." *The Cambridge History of Literary Criticism, Vol. 3: The Renaissance*, edited by George Alexander Kennedy, Cambridge University Press, 1993, pp. 305–13.

- Norton, Glyn P., and Marga Cottino-Jones. "Theories of Prose Fiction and Poetics in Italy: Novella and Romanzo (1525-1596)." *The Cambridge History of Literary Criticism, Vol. 3: The Renaissance*, edited by George Alexander Kennedy, Cambridge University Press, 1993, pp. 322–36.
- Nosowitz, Dan. "Netflix's 'Elite Raters' Have Rated Over 50,000 Movies." *Fast Company*, 17 Sept. 2010, <https://www.fastcompany.com/1689560/netflixs-elite-raters-have-rated-over-50000-movies>.
- Pawley, Christine. *Reading on the Middle Border: The Culture of Print in Late-Nineteenth-Century Osage, Iowa*. University of Massachusetts Press, 2001.
- Phegley, Jennifer, and Janet Badia, editors. *Reading Women: Literary Figures and Cultural Icons from the Victorian Age to the Present*. University of Toronto Press, 2005.
- Philbrick, Rodman. "Writing a Book Review." *Writing with Writers*, 2017, <http://teacher.scholastic.com/writewit/bookrev/index.htm>.
- Priestman, Martin. *Detective Fiction and Literature: The Figure on the Carpet*. Palgrave Macmillan, 1990.
- Propp, Vladimir. *Morphology of the Folk Tale*. Edited by Louis A. Wagner, Translated by Laurence Scott, University of Texas Press, 1968.
- Publishers Weekly. "Fiction Book Review: Old Venus." *Publishers Weekly*, 10 Nov. 2014, <https://www.publishersweekly.com/9780345537287>.
- . "Fiction Book Review: The Da Vinci Code." *PublishersWeekly.Com*, 3 Feb. 2003, /978-0-385-50420-1.
- . "My Year of Meats." *Publishers Weekly*, 3 Nov. 2003, /978-0-7861-2478-7.

- Quantcast. "Goodreads.Com Audience Insights and Demographic Analytics." *Quantcast*,
<https://www.quantcast.com/goodreads.com?country=US>. Accessed 18 July 2017.
- Radway, Janice A. *A Feeling for Books: The Book-of-the-Month Club, Literary Taste, and Middle-Class Desire*. University of North Carolina Press, 1997.
- . *Reading the Romance: Women, Patriarchy, and Popular Literature*. University of North Carolina Press, 1984.
- Reeve, Clara. *The Progress of Romance and The History of Charoba, Queen of Aegypt*. The Facsimile Text Society, 1930, <https://catalog.hathitrust.org/Record/002041663>.
- Reference.com. "How Do You Write a Plot Summary?" *Reference.Com*,
<https://www.reference.com/education/write-plot-summary-b8d09ed2e77c0c6f>. Accessed 14 Dec. 2017.
- Rishonda. "Review of Gone with the Wind." *Goodreads*,
<https://www.goodreads.com/review/show/1573731082>. Accessed 16 Jan. 2018.
- Roy, Daniel. "Review of The Lifecycle of Software Objects." *Goodreads*, 9 May 2011,
<https://www.goodreads.com/review/show/167038391>.
- Rubin, Joan Shelley. *The Making of Middlebrow Culture*. The University of North Carolina Press, 1991.
- Rutten, Tim. "Inspector Darko Blazes a Bright New Tale in Ghana." *Los Angeles Times*, 13 Aug. 2009, p. D.1.
- Saldana, Johnny. *The Coding Manual for Qualitative Researchers*. Sage Publications, 2009.
- Schank, Roger C., and Robert P. Abelson. *Scripts, Plans, Goals and Understanding: An Inquiry into Human Knowledge Structures*. Lawrence Erlbaum, 1977.

- Schelling, F. W. J. *The Philosophy of Art*. Edited by Douglas W. Stott, Univ of Minnesota Press, 1989.
- Schlegel, Friedrich. "Athenaeum Fragments." *Friedrich Schlegel's Lucinde and the Fragments*, translated by Peter Firchow, University of Minnesota Press, 1971, pp. 161–240.
- . *Literary Notebooks, 1797–1801*. Edited by Hans Eichner, Athlone Press, 1957.
- Scorsese, Martin. "Martin Scorsese on Rotten Tomatoes, Box Office Obsession and Why 'Mother!' Was Misjudged." *The Hollywood Reporter*, Oct. 2017, <http://www.hollywoodreporter.com/news/martin-scorsese-rotten-tomatoes-box-office-obsession-why-mother-was-misjudged-guest-column-1047286>.
- Semnani, Neda. "A Mormon Mystery." *Los Angeles Review of Books*, Feb. 2015, <https://lareviewofbooks.org/article/mormon-mystery/>.
- Sharon. "Review of Scattered Ashes." *Goodreads*, 27 Aug. 2012, <https://www.goodreads.com/review/show/401670256>.
- Shmoop. "Cat's Cradle Summary." *Shmoop*, <https://www.shmoop.com/cats-cradle-vonnegut/summary.html>. Accessed 17 Dec. 2017.
- . "Flowers for Algernon Summary." *Shmoop*, <https://www.shmoop.com/flowers-for-algernon/summary.html>. Accessed 17 Dec. 2017.
- Sicherman, B. "Sense and Sensibility: A Case Study of Women's Reading in Late-Victorian America." *Reading in America: Literature and Social History*, edited by Cathy N. Davidson, Johns Hopkins University Press, 1989, pp. 201–25.
- Sicoe, Veronica. "Review of The Left Hand of Darkness." *Goodreads*, 31 Aug. 2013, <https://www.goodreads.com/review/show/708000528>.
- Siegel, Lee. "Love's Labor's Lost." *Los Angeles Times*, 16 Sept. 2001, p. BR.3.

- Simeone, Michael, et al. "Towards a Poetics of Strangeness: Experiments in Classifying Language of Technological Novelty." *Journal of Cultural Analytics*, Sept. 2017, <http://culturalanalytics.org/2017/09/towards-a-poetics-of-strangeness-experiments-in-classifying-language-of-technological-novelty/>.
- Sisk, Micah. "Review of Foundation and Empire." *Goodreads*, 10 Nov. 2014, <https://www.goodreads.com/review/show/1102768100>.
- Smuts, Aaron. "The Desire-Frustration Theory of Suspense." *The Journal of Aesthetics and Art Criticism*, vol. 66, no. 3, 2008, pp. 281–90.
- So, Richard Jean, and Hoyt Long. "Network Analysis and the Sociology of Modernism." *Boundary 2*, vol. 40, no. 2, 2013, pp. 147–82.
- SparkNotes. "Flowers for Algernon: Plot Overview." *SparkNotes*, <http://www.sparknotes.com/lit/algernon/summary.html>. Accessed 17 Dec. 2017.
- . "The Da Vinci Code: Plot Overview." *SparkNotes*, <http://www.sparknotes.com/lit/davincicode/summary.html>. Accessed 17 Dec. 2017.
- Steven. "Review of No Country for Old Men." *Goodreads*, 25 Mar. 2008, <https://www.goodreads.com/review/show/18581494>.
- Strabo. *The Geography of Strabo*. Edited by H. L. Jones, Harvard UP, 1924.
- Striphas, Ted. *The Late Age of Print: Everyday Book Culture from Consumerism to Control*. Reprint edition, Columbia University Press, 2011.
- Summers, Ed. "Archival Attention: Visualizing Archival Appraisal." *Medium*, 5 Dec. 2017, <https://medium.com/@edsu/archival-attention-df4991aec716>.
- Suvin, Darko. *Metamorphoses of Science Fiction: On the Poetics and History of a Literary Genre*. Yale University Press, 1979.

- . *Positions and Presuppositions in Science Fiction*. Macmillan, 1988.
- Taylor, Karl K. "Teaching Summarization Skills." *Journal of Reading*, vol. 27, no. 5, 1984, pp. 389–93.
- Todd, Emily B. "Antebellum Libraries in Richmond and New Orleans and the Search for the Practices and Preferences of 'Real' Readers." *American Studies*, vol. 42, no. 3, 2001, pp. 195–209.
- Todorov, Tzvetan. *The Poetics of Prose*. Translated by Richard Howard, Cornell University Press, 1977.
- UNC Writing Center. "Book Reviews." *The Writing Center, University of North Carolina at Chapel Hill*, 2017, <http://writingcenter.unc.edu/tips-and-tools/book-reviews/>.
- Underwood, Ted. "The Life Cycles of Genres." *Journal of Cultural Analytics*, May 2016, <http://culturalanalytics.org/2016/05/the-life-cycles-of-genres>.
- Underwood, Ted, and David Bamman. "The Gender Balance of Fiction, 1800-2007." *The Stone and the Shell*, 29 Dec. 2016, <https://tedunderwood.com/2016/12/28/the-gender-balance-of-fiction-1800-2007/>.
- Vala, Hardik, et al. "Mr. Bennet, His Coachman, and the Archbishop Walk into a Bar but Only One of Them Gets Recognized: On The Difficulty of Detecting Characters in Literary Texts." *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2015.
- Vani. "Review of The Da Vinci Code." *Goodreads*, 18 Nov. 2014, <https://www.goodreads.com/review/show/1109975204>.
- Wadsworth, Sarah. *In the Company of Books: Literature and Its "Classes" in Nineteenth-Century America*. University of Massachusetts Press, 2006.

Wald, Priscilla, and Michael A. Elliott, editors. *The American Novel 1870-1940*. Vol. 6, Oxford University Press, 2014.

wikiHow. "How to Write a Book Review." *WikiHow*, 13 Oct. 2017,
<https://www.wikihow.com/Write-a-Book-Review>.

---. "How to Write a Plot Outline." *WikiHow*, <https://www.wikihow.com/Write-a-Plot-Outline>.
Accessed 19 Dec. 2017.

Wikipedia. "Flowers for Algernon." *Wikipedia*, 14 Dec. 2017,
https://en.wikipedia.org/w/index.php?title=Flowers_for_Algernon&oldid=815317273.

---. "Gone with the Wind." *Wikipedia*, 3 Jan. 2018,
[https://en.wikipedia.org/w/index.php?title=Gone_with_the_Wind_\(novel\)&oldid=818494121](https://en.wikipedia.org/w/index.php?title=Gone_with_the_Wind_(novel)&oldid=818494121).

---. "How to Write a Plot Summary." *Wikipedia*, 20 Nov. 2017,
https://en.wikipedia.org/w/index.php?title=Wikipedia:How_to_write_a_plot_summary&oldid=811288328.

---. "My Year of Meats." *Wikipedia*, 7 July 2017,
https://en.wikipedia.org/w/index.php?title=My_Year_of_Meats&oldid=789382334.

Wilkins, Matthew. "Genre, Computation, and the Varieties of Twentieth-Century U.S. Fiction." *Journal of Cultural Analytics*, Nov. 2016, <http://culturalanalytics.org/2016/11/genre-computation-and-the-varieties-of-twentieth-century-u-s-fiction/>.

Wimsatt, W. K., and Monroe C. Beardsley. "The Affective Fallacy." *The Sewanee Review*, vol. 57, no. 1, 1949, pp. 31–55.

---. "The Intentional Fallacy." *The Verbal Icon: Studies in the Meaning of Poetry*, University of Kentucky Press, 1954, pp. 3–18, <http://faculty.smu.edu/nschwartz/seminar/fallacy.htm>.

- Winograd, Peter N. "Strategic Difficulties in Summarizing Texts." *Reading Research Quarterly*, vol. 19, no. 4, 1984, pp. 404–25.
- Witty, Francis J. "The Pinakes of Callimachus." *The Library Quarterly: Information, Community, Policy*, vol. 28, no. 2, 1958, pp. 132–36.
- Wolf, Maryanne. *Proust and the Squid: The Story and Science of the Reading Brain*. HarperCollins, 2007.
- Xiong, Wenting. *Helpfulness-Guided Review Summarization*. University of Pittsburgh, 2014, <http://search.proquest.com/pqdtglobal/docview/1667378636/abstract/9A36B5D13A4F48ABPQ/1>.
- Zboray, Ronald J., and Mary Saracino Zboray. *Literary Dollars and Social Sense: A People's History of the Mass Market Book*. Routledge, 2005.
- Zunshine, Lisa. *Why We Read Fiction: Theory of Mind and the Novel*. Ohio State University Press, 2006.