Patterns of Quality: Comparing Reader Reception Across Fanfiction and Commercially Published Literature

Mia Jacobsen[1,*], Yuri Bizzoni[1], Pascale Feldkamp Moreira[1,2] and Kristoffer L. Nielbo[1]

[1]*Center for Humanities Computing, Aarhus University, Denmark*

[2]*Department of Comparative Literature, Aarhus University, Denmark*

Abstract

Recent work on the textual features linked to literary quality has primarily focused on commercially published literature, such as canonical or best-selling novels, that are systematically filtered by editorial and market mechanisms. However, the biggest repositories of fiction texts currently in existence are free fanfiction websites, where fans post fictional stories about their favorite characters for the pleasure of writing and engaging with others. This makes them a particularly interesting domain to study the patterns of perceived quality "in the wild", where text-reader relations are less filtered. Moreover, since fanfiction is a community-built domain with its own conventions, comparing it to published literature can more generally provide insights into the reception and perceived quality of published literature itself. Taking a novel approach to the study of fanfiction, we observe whether three textual features associated with perceived literary quality in published texts are also relevant in the context of fanfiction. Using different reception proxies, we find that despite the differences of fanfiction from published literature, some "patterns of quality" associated with positive reception appear to hold similar effects in both of these contexts of literary production.

Keywords

fanfiction, literary quality, reader appreciation, canon, fandoms,

1. Introduction

Throughout literary history, the question of literary quality has garnered answers in various forms, from complete (and competing) aesthetic theories to creative writing manuals [25, 31]. A small number of studies have also attempted, in recent years, to tackle the topic from a quantitative perspective, identifying stylistic and narrative patterns that might hold a connection with a positive reader experience and response [20, 10, 23]. Such studies have converged on general elements at the level of style and structure, which appear connected to different kinds of literary reception, i.a., linguistic complexity [44], emotional articulation [64], and "basic" stylometric parameters [4]. Such analyses, however, have been primarily focused on com-

mercially published literature,[1] which comes with its inevitable filters of marketing, editorial processing, as well as more general cultural and social promotion, such as the profile and fame of an author – factors that have also been shown to have a significant impact for sales figures [83]. In this respect, the world of online self-publishing can be an ideal "experimental setting" for observing how and whether observations made on published literature hold when filters and constraints are diminished.

Despite the existence of large repositories of non-fanfiction works,[2] the largest part of online fiction is arguably constituted of fanfiction.[3] A fanfic is defined as a fictional story written by fans that centers around pre-existing characters, plots and/or entire imaginary worlds [3, 78, 69, 68]. While this definition is broad enough to include texts from the literary tradition – Euripides' *The Trojan Women* is, in this sense, a fanfic of *The Iliad* – fanfiction is generally understood in a narrower sense as the production of groups of non-professional writers expanding commercially published novels or shows. In the digital age, fanfiction is predominantly posted to online websites [30], such as Fanfiction.net[4] and Archiveofourown.org[5], also known as AO3. There are currently over 13 million posts on AO3 from more than 65,000 fan communities.

Fanfiction is often considered a 'lesser' form of literature: derivative, unedited, written mostly to elicit strong affective responses, and so forth [72, 67]; and despite the ready availability of large amounts of fanfiction online, relatively little research has examined these descriptions critically and explored its textual profile quantitatively. To bridge the gap between the study of textual features of perceived literary quality of published literature and of fanfiction, this paper explores and compares selected textual features from 9,000 fanfics with those from a corpus of ca. 9,000 published works. Specifically, we utilize readability, nominal style, and the Hurst exponent of stories' sentiment arcs. We chose these measures as they have been related to reader appreciation and popularity in the context of commercially published and established literature.

2. Related Works

2.1. Literary Quality

The question of what makes creative writing 'good' is perhaps one of the most debated in literary history. Over time there have been many rules and recommendations about how to write better, supposedly applicable across genres: from detailed suggestions about which parts of speech one ought to avoid to overall recommendations on style. Generally, one might distinguish between two main lines of reasoning: on one hand, the idea that good literature is simple or "more direct", and on the other hand, the concept that good literature is complex and demanding .

---

[1] With 'commercially published', we refer to fiction as it has traditionally been published on the market in a book form (at least in a Western context). Henceforth, we use the phrase 'published literature' to refer to commercially published literature.

[2] We refer to non-fanfiction works of literature online, such as those found on Wattpad.

[3] Fanfiction is often abbreviated to just *fanfic* or even *fic*. Here, we'll use the term *fanfiction* for the genre more generally, while we refer to individual stories as *fanfics*.

[4] https://www.fanfiction.net

[5] https://archiveofourown.org

Regarding good writing as more 'direct', we observe scholars and authors promoting a direct style, such as Sherman [73] who suggested that simplicity – i.e., shorter sentences – should be a marker of "better" literature. Some simplicity laws for literature have traditionally been set forth by critics and authors alike – for example, Hemingway recommends plain and understated prose [38], Stephen King has famously advocated more readable texts King [48], and Strunk, White, and Angell [76]'s influential book, *The Elements of Style*, advised avoiding 'embellishment'. Measures adopted from linguistics, so-called readability indices – essentially based on estimating difficulty by sentence or word length – are also used to estimate both the accessibility and, implicitly, the "quality" of a text. They are widely implemented in more recent creative writing and publishing aids.[6] Generally, more recent studies that seek to predict perceived literary quality or success also include textual features related to readability (i.a., sentence-length, vocabulary richness and redundancy) [16, 23, 24, 49, 52, 1].

Conversely, others have promoted "purple prose", characterized as complex and challenging writing, "rich, succulent and full of novelty" [84]. The profile of canonic literary works has often been associated with such more challenging style, whether because canonic works exhibit lower readability [7], higher textual entropy [1] or higher perplexity and cognitive demand for the reader [10, 85]. Recent studies that have compared features of texts across various "proxies of quality" – i.e., comparing books that won prizes to bestsellers – show that the preference for easier/ more challenging books varies across proxies. For example, more readable books are more likely to have a high number of ratings on Goodreads but are less likely to win awards [7]. Similarly, more prestigious literature appears to elicit higher perplexity (i.e., perplexity or novelty as measured through Large Language Models) than popular literature [85]. What the textual features of "quality" are, then, depends largely on how we conceptualize quality, especially since defining literary success as "popularity" or spread (i.e., number of Goodreads' ratings) appears to exhibit opposed tendencies to defining it as "prestige" or expert choice (i.e., canonicity as measured by the presence of a work on college syllabi). For the present study, we consider both polarities, operationalizing quality as both popularity and appreciation (see Section 3.3).

Going beyond traditional stylometric assessments, recent studies have focused on gauging the complexity of a text and its relation to literary success by looking at the deeper dynamics of their emotional development – intensity, fluctuations, and trajectory [53, 10]. Most of these works have centered on tracing so-called sentiment arcs, i.e., arcs that represent how the valence of words or sentences fluctuates across a narrative text [46]. Focusing on the *affect* of texts in this way introduced a new lens for understanding the impact of texts on readers [28, 18], with the potential for moving beyond the stylistic level when modeling perceptions of literary quality [66]. While most studies have focused on the visual shapes of sentiment arcs [71, 47], others have sought to gauge their mathematical properties [54, 11] based on the idea that readers tend to appreciate a certain balance in the complexity and predictability of the narrative flow. Hu, Liu, Thomsen, Gao, and Nielbo [40] and Bizzoni, Peura, Nielbo, and Thomsen [11] have modeled the persistence, coherence, and predictability of arcs through measures like the Hurst exponent to measure the global complexity of sentiment arcs [8] – a measure that appears to be applicable for predicting the relative success of a work [13]. This perspective aligns

---

[6]Such as the Hemingway or Marlowe applications

with theories that go beyond simple stylistic complexity emphasizing the capacity of narratives to engage and challenge readers at an affective level [2]. Such approaches foreground narrative or sentiment complexity as a key determinant for perceived literary quality [40], drawing on the role of predictability for aesthetic attraction [22, 57] in other domains, such as in music or the visual arts [56, 15].

2.2. Fanfiction as a medium

Unlike most contemporary fiction, fanfics are often posted serially with authors posting one chapter of a story at a time. Despite the option of planning the story ahead, most fanfiction writers also choose to write on an ongoing basis. Depending on the platform, readers can leave comments and likes for specific stories to indicate their appreciation and encourage the authors to continue, which also allows them to respond and incorporate readers' feedback into future chapters [32, 14]. Many fanfiction writers are also more focused on creating emotional situations and experiences[7] that center on the characters of their favorite media. As such, developing a structured and "complete" narrative is less important for writers [3], which means it is not uncommon for fanfics to have a thin or non-existent plot (known as the *PWP* genre or *Plot, What Plot?*) and to be left unfinished.

Early research into fan communities and fanfiction often compared fanfics to their respective source material, focusing primarily on the similarities and contrasts in the content of the fanfics compared to the source. A form of power struggle between producers and fans was observed [45] in the way fans created their version of the source texts through re-interpreting and re-writing certain narrative elements or events [77]. Fanfiction became, in this sense, a medium through which fans could "fix" their favorite stories – which later became a genre of its own, the "fix-it-fics" [69]. In more recent years, the understanding of the power dynamic between producers and fans has become more nuanced [78], and focused more on the internal structure of the community and how it's reflected in the texts [77, 26, 17].

2.3. Quantitative Studies of Fanfiction

An early quantitative study of fanfiction compared the attention allocated to different characters in fanfics to the corresponding original text [60]. Comparing the source texts of ten different canons to their corresponding fanfiction, they found that fanfiction deprioritizes main characters in favor of secondary ones and that it devotes more attention to female characters than the original texts. There have since been multiple studies that have investigated the features of popular fanfiction. Concerning stylometrics, popular fanfics have been found more likely to have a simpler syntactic structure and plainer writing style, but a wider vocabulary [55, 63]. When comparing fanfiction to source texts, fans prefer longer fanfics with a greater romantic focus and emotional arcs that are dissimilar to the original novel [75, 63]. Fans also appear to generally prefer more character interaction over narrative exposition [55], while different communities might prefer the characters' dynamics to be more similar to or more deviant from the original text [75]. This indicates that some aspects of well-liked fanfiction may hold across fandoms, such as a preference for a novel story arc and features related to the

---

[7]Or for characters to have sex [61].

simplicity/complexity of the text, whereas other aspects, such as specific character interaction dynamics, are more community-dependent.

3. Methods

We compare a corpus of commercially published fiction to a corpus of fanfiction using three different textual measures and three different conceptions of "quality".

The three textual measures are related to complexity at two theoretically distinct levels: the *stylistic level* where we measure complexity through readability and nominal ratio, and the *narrative* or *structural level* where we measure complexity via the Hurst exponent of the sentiment arc (described in depth in the following sections).

We then relate these three measures to our three conceptions of perceived literary quality. The first we call *spread* which relates to the popularity of a given work. The second is *appreciation* which is based on crowd-sourced opinions of the given works. The final conception is based on the idea of *canonicity* which is related more to expert opinions and prestige for the commercially published novels (also described in depth in the following sections).

3.1. Corpora

We compare a corpus of published novels – the so-called *Chicago Corpus* – to a corpus of fanfics, both comprising around 9,000 works (see Table 1).

The Chicago Corpus   is a selection of 9,089 English-language novels from different genres, published in the US between 1880 and 2000, and covering around 3,150 authors (see Bizzoni, Moreira, Lassen, Thomsen, and Nielbo [9] for details).

It has been used in recent studies focusing on the textual properties of books deemed of "high quality" [85].[8] As it is also a unique dataset in terms of size.[9] and diversity, we consider it an especially good match for our research. Texts were selected based on the number of libraries holding a copy of the novels and the corpus spans many genres across high- and low-brow fiction. It also lists both prestigious and popular works ranging from Nobel prize winners to Science Fiction classics [51].[10]

The fanfiction corpus   consists of a sample of 9,000 fanfics from three different fandoms: *Percy Jackson and the Olympians, Harry Potter*, and *The Lord of the Rings.* These three fandoms were chosen because they constitute some of the biggest fan groups based on literary works.[11]

---

[8]The annotated Chicago corpus dataset – though including only a selection of full texts – is available at https://github.com/centre-for-humanities-computing/chicago_corpus.

[9]Often, studies on reader appreciation rely on < 1,000 books [35, 49] Though larger corpora exist (e.g. the Hathi), to the best of our knowledge this is the largest curated corpus of narrative fiction (without common OCR errors, spurious texts, noise segments such as introductions, etc.).

[10]The corpus has no reference publication, though other studies are based on it [79, 19]. See the corpus description at the Textual Optics Lab.

[11]An exhaustive study of fanfiction is nearly impossible at this stage given the sheer number of texts and fan groups. We decided to pick these three fandoms due to their size and popularity, taking them as a good starting point for this initial study.

Although narrowing the corpus in this way limits the generalizability of the findings, it allows for a more controlled comparison between the two corpora as well as across fandoms. Especially when considering the limited previous work on quantitative studies of fanfiction, a more controlled corpus is important for robustness while also opening the door for greater variety in future research.

The fanfics were randomly sampled and scraped from the online fanfiction site *Archive of Our Own* from the three fandom tags – "Harry Potter – J. K. Rowling", "Percy Jackson and the Olympians – Rick Riordan", and "Lord of the Rings – J. R. R. Tolkien" – and 3,000 fanfics from each tag were retrieved.[12] Only fanfics that were written in English and had no crossovers (characters from other fandoms) were included to facilitate a more controlled comparison across fandoms. 14 texts were excluded due to artifacts in the scraping process.

Table 1
Summary of the Chicago and Fanfiction corpora

Corpus	Texts	N. words
Chicago	9,089	1,059,783,918
Fanfiction	8,986	101,075,646

3.2. Textual measures

For the textual features, we have selected three different measures that have been used in the study of literary quality and reception in published literature: the Dale-Chall New Readability score, measuring textual simplicity at the stylistic level; the "nominal ratio", measuring the nominality of the writing (presence of nominal style); and the Hurst exponent, measuring the complexity of the sentiment trajectory of a story, i.e., the sentiment arc.

We chose these measures based on their effectiveness in previous quantitative analyses of literary quality, as well as their coverage of different layers of text: surface stylometry (Dale-Chall), grammatical patterns (nominal ratio) and, finally, sentimental-narrative profile (Hurst exponent) [7, 85, 10]. In other words, these measures have a well-established relationship to different conceptions of literary quality, and allow us to assess the texts on different levels of complexity.

Dale-Chall Readability The Dale-Chall readability score was developed in the 1940s by linguists Edgar Dale and Jean Chall, who attempted to measure the difficulty of a text from easy (low score) to hard (high score). Like other readability metrics, it proposes a sentence and word length combination calibrated with specific constants.

$$\text{Raw Score} = 0.1579 \left(\frac{\text{Difficult Words}}{\text{Total Words}} \times 100 \right) + 0.0496 \left(\frac{\text{Total Words}}{\text{Total Sentences}} \right) \qquad (1)$$

The score also adjusts for the presence of "difficult words" – defined as words which do not appear on a list of words which 80% of fourth-graders would know.[13]

---

[12]The texts were collected at the beginning of 2024 and are published between January 2002 and December 2023.
[13]See the Dale-Chall word-list.

$$\text{Adjusted Score} = \begin{cases} \text{Raw Score} & \text{if } \frac{\text{Difficult Words}}{\text{Total Words}} \leq 0.05 \\ \text{Raw Score} + 3.6365 & \text{if } \frac{\text{Difficult Words}}{\text{Total Words}} > 0.05 \end{cases} \tag{2}$$

Among several tested formulae, the Dale-Chall appeared to be one of the best predictors of spread and popularity in literary fiction [7].

Nominal Ratio   We call "nominal ratio" the ratio between nouns and adjectives over verbs. We estimated the number of nouns, adjectives, and verbs in each work by annotating the texts with parts-of-speech tags using spacy's [39] large English pre-trained model.[14]

$$\text{Nominal Ratio} = \frac{\text{Nouns} + \text{Adjectives}}{\text{Verbs}} \tag{3}$$

This metric has proved, in published literature, to hold strong correlations with LLM-based perplexity as well as reception and quality proxies, and to be in accordance with the effects of "optimized" communication in non-literary domains [85, 27]. It can be understood as a measure of how demanding the writing style is, and be used to nuance the score from the readability measure.

Hurst exponent   The Hurst exponent (H) measures the long-term memory of a time series, indicating whether it is trending, mean-reverting, or exhibiting a random walk behaviour. A value of $H = 0.5$ suggests a random walk (no correlation), $H > 0.5$ indicates persistent or trending behavior, and $H < 0.5$ suggests anti-persistent or mean-reverting behavior.

The formula for estimating the Hurst exponent H using rescaled range analysis is:

$$\frac{R(n)}{S(n)} = a \cdot n^H$$

where $R(n)$ is the range of the first n values, $S(n)$ is the standard deviation of the first n values, and a is a constant [41, 70].[15] More persistent arcs tend to be connected to more predictable narratives [6], while mean-reverting arcs are connected to more complex narratives.

Recent studies in the dynamics of sentiment arcs for novels and short stories alike have found that different conceptions of quality are related differently to the Hurst exponent [40, 8]. Lower Hurst exponents (i.e., more complex narratives) have been connected to more 'highbrow' fiction, and higher Hurst exponents (i.e., more predictable dynamics), to more widely spreading works (like bestsellers) [5].

We compute the Hurst exponent of both novels and fanfics from their sentiment arcs – i.e., the consecutive highs and lows in valence across a narrative – computed through valence annotation with VADER [42] on a sentence base.

---

[14]https://github.com/explosion/spacy-models/releases/tag/en_core_web_lg-3.7.1

[15]For a more detailed version see for example Hu, Liu, Thomsen, Gao, and Nielbo [40]

3.3. Quality measures

Naturally, there are several different approaches to the understanding of literary quality. In this paper, we use a combination of crowd-sourced and expert-based proxies. The crowd-sourced quality measures are based on two proxies of reader reception:[16] reader appreciation and spread. In addition to the reader reception measures, we also use conceptions of *canonical literature* as a separate expert-based quality proxy, as defined below.

Published Fiction Quality Measures

To gauge the appreciation and spread of published fiction, we use ratings from the popular online social platform Goodreads. With its more than 90 million users, Goodreads catalogues books from a wide spectrum of genres and derives book-ratings from a heterogeneous pool of readers in terms of background, gender, age, native language and reading preferences [62].[17]. While Goodreads' ratings and rating count do not present an absolute measure of literary appreciation, they do offer a valuable perspective on a title's overall reception among a diverse population of readers, whose preferences appear to differ from expert critics [81].

Appreciation: Goodreads average rating   We used the average number of stars (from 1 to 5) assigned to a book by Goodreads' users as our measure of reader appreciation for published novels. This measure has the benefit of being independent from the number of ratings the book received.

Spread: Goodreads rating count   Complementary to the average rating, the number of ratings on Goodreads indicates how many users have taken the time to assign a score to a given novel, independently from the score-value. As such, we use it as a metric of spread or popularity – the book might be infamous, but it did manage to reach enough readers to get rated online.[18]

Canonicity: Prizes and Penguin Classics   Defining what constitutes canonical literature is a complicated task, and different literary scholarships have held extreme views on whether the conception of the canon is arbitrary or universal [37, 36]. Still, recent studies have shown that, at the large scale, readers seem to converge on their perception of what is canonic, classic or "literary" [49, 82, 34], categories which also appear to exhibit a distinct textual profile [16, 4, 12, 5].

---

[16]With "quality proxy", we mean an approximation of the concept of literary quality or reader appreciaiton, i.e., a specific operationalization of reader appreciation *among many potential ones*. For example, what is rated high on Goodreads may not be a book held in many libraries or may have very few translations; library holding numbers and translation counts being two other ways of approximating reader appreciation.

[17]Still, we see the continuation of established patterns on the platform: for example, works that are often assigned on college syllabi are also perceived as 'classics' on Goodreads [82]

[18]Feldkamp, Bizzoni, Thomsen, and Nielbo [34] show how Goodreads rating count seems to correlate with other proxies that tend to measure dissemination or popularity over appeciation, such as translation and library holdings, the Wikipedia rank of the author, and so forth.

For gauging the feature values of more "canonical" fiction in the published corpus we created two subsets of canonical fiction based on Feldkamp, Bizzoni, Thomsen, and Nielbo [34]: i) a *Prizes* subset of those titles that have been long-listed for the Pulitzer and National Book Award, or are by a Nobel-winning author, and ii) a *Canon* subset of those titles that are canonical insofar as they are included in the Penguin Classics series,[19] appear among the top 1,000 titles on college syllabi for English Literature,[20] or are by authors who are mentioned in the Norton Anthology of English and American literature.[21,22]

Fanfiction Quality Measures

To gauge the appreciation and spread of fanfiction, we use metrics available on AO3.

Appreciation: The kudos/hits ratio   The kudos/hits ratio is based on the number of likes a story has received – the so-called *kudos* – and the number of people who have opened the fanfic – the number of *hits*. It is computed as the number of kudos divided by the number of hits, which amounts to a measure of how many fans who opened and read a fanfic also enjoyed it. Although it doesn't account for repeating readers – thus penalizing fanfics that update more often – it does account for the fact that older fanfics will have a greater number of kudos and hits merely because they are older [65]. It also has the benefit of normalizing the number of kudos based on hits, and thus makes it independent of its popularity.

Spread: Number of hits   The popularity/spread proxy for fanfiction was measured as the number of hits. Fanfics get a hit each time a user opens the given story.

"Canonical" fanfiction   Since there is no comparable measure of canonicity as it pertains to fanfiction, this divide is primarily used for testing whether fanfics with a high kudos/hits ratio behave like canonical fiction. Thus, to have a comparable split of "canonical" fanfiction, we used the kudos/hits ratio to divide the fanfiction corpus into three. The first split divided the fanfiction into the bottom 50% and upper 50% kudos/hits ratio scores. The 50% cut occurred at a kudos/hits ratio of 5.49. This meant 4,493 fanfics in the "non-canonical" group (i.e., bottom 50%) and 4,493 fanfics in the "canonical group" (i.e., upper 50%). The second split was at the upper quantile of the kudos/hits ratio, meaning a kudos/hits ratio of 8.29. 2,240 fanfiction are "canonical" using this split while 6,737 fanfics are "non-canonical". For the final group, a split at 87.5% of the distribution was used. This meant fanfics were split at a kudos/hits ratio of 10.76, meaning 1,124 "canonical" fanfics and 7,862 "non-canonical" fanfics.

---

[19]https://www.penguin.co.uk/penguin-classics
[20]We used the data of the OpenSyllabus project: https://www.opensyllabus.org
[21]https://www.norton.com/books/9780393543902
[22]For more on the collection of these quality-proxies, see the Chicago Corpus Dataset documentation.

4. Results

4.1. Comparing fanfiction to the Chicago Corpus

The difference in textual features across corpora was tested using a Wilcoxon Rank Sum test. Compared to the Chicago corpus, the fanfiction texts have a significantly higher Dale Chall Readability score ($r = .68$, $p < .001$), meaning they are generally less readable than published literature. A qualitative inspection of the fanfics with the highest Dale Chall score showed that this effect is to some degree explained by very long and run-on sentences with few full stops.[23]
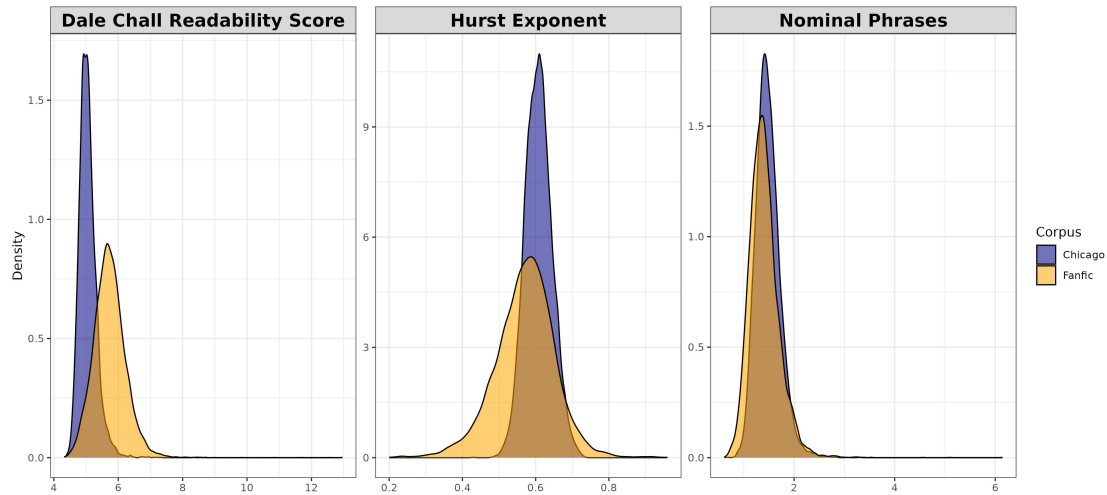

Figure 1: Distribution of scores for each corpus across textual measures.

Table 2

Mean values, * indicates a significant difference across corpora ($p < .001$).

	Fanfiction mean	Chicago mean	Difference
Dale Chall	5.73	5.06	-0.67*
Nominal Phrases	1.41	1.48	0.07*
Hurst Exponent	0.57	0.61	0.04*

Fanfiction is also found to have significantly lower nominal ratio ($r = .16$, $p < .001$) making the writing style generally less demanding compared to published fiction, confirming that the lower readability is not linked to an overall more sophisticated style, but simply to a harder-to-read style. Differences were evident at the level of sentiment arcs' structure as well:[24] fanfiction

---

[23]For example: "Better than watching her though he cannot help hearing her voice, low in song, raised unceasing as it has risen and fallen these long slow hours, past time when mortal throat would silence in hoarseness, untiring, in effort of power on which more than life shall depend, rising and falling, now to cajole, coax into quietude, lulling into stillness, then strong to command, bind into submission, that over which she sings": https://archiveo fourown.org/works/4233628

[24]something discussed in qualitative analyses as well, see for example Kustritz [50]

Table 3
Spearman's ρ between each textual feature and **spread**. $^*p < .001$

	Dale-Chall	Nominal ratio	Hurst exponent
Fanfiction	-0.24*	-0.17*	0.092*
Chicago	-0.13*	-0.05*	0.072*

Table 4
Spearman's ρ between each textual feature and **appreciation**. $^*p < .001$

	Dale-Chall	Nominal ratio	Hurst exponent
Fanfiction	0.18*	0.063*	-0.12*
Chicago	-0.20*	-0.089*	0.16*

has a significantly lower Hurst exponent than Chicago ($r = .29$, $p < .001$), meaning they are less coherent, but also less predictable, in their story arcs.

With respect to published literature, these results suggest a hybrid scenario in fanfiction, where some features normally associated with high-brow or canonical literature (such as a lower Hurst exponent) are mixed with others linked to low-brow or popular literature (such as a lower nominal ratio). These differences are also visualized on Figure 1.

4.2. Reader Appreciation and Spread

Concerning correlations between the spread of texts and their different textual features, there is a similar pattern for the Chicago corpus and the fanfic corpus. As reported in Table 3), fanfics' number of hits exhibits the same pattern as the Chicago corpus' number of ratings: both hits and number of ratings have a negative correlation with Dale Chall scores and nominal style, while there is a slight positive correlation with Hurst (Table 3). This indicates that dense, demanding, and less readable writing styles may slow the spread of narratives regardless of its type.

Concerning reader appreciation, fanfics and published literature contrast in several ways. As reported in Table 4, the kudos/hits ratio is found to be positively correlated with Dale Chall scores. This means that as fanfics become harder to read, reader appreciation increases. The opposite pattern applies to the Chicago corpus, where the average rating on Goodreads has a negative correlation with Dale Chall scores (Table 4). Similarly for nominal ratio, the fanfiction corpus is found to have a positive correlation between nominal ratio and appreciation, while the opposite is found in the Chicago corpus. For the Hurst exponent, the direction of the correlation is again opposite across the two corpora, with fanfiction having a negative correlation between reader appreciation and Hurst, while Chicago has a positive one.

This could indicate two differing interpretations. The first being that what is appreciated in fanfiction differs from what is appreciated in published novels. In other words, while the Goodreads users show a preference for a more predictable narratives that are easier to read both stylistically and grammatically, fans exhibit the opposite preference. The second interpretation concerns what is actually being measured. As mentioned earlier, there are multiple ways to
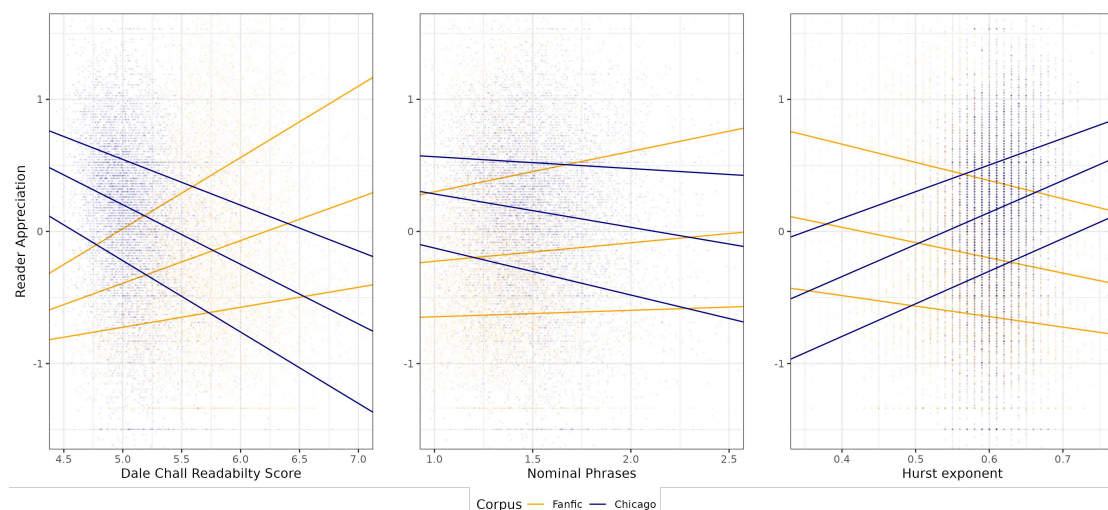
Figure 2: Quantile regression lines between textual feature scores and **reader appreciation** proxies across the two corpora. Reader appreciation proxies have been standardized for the visualization.
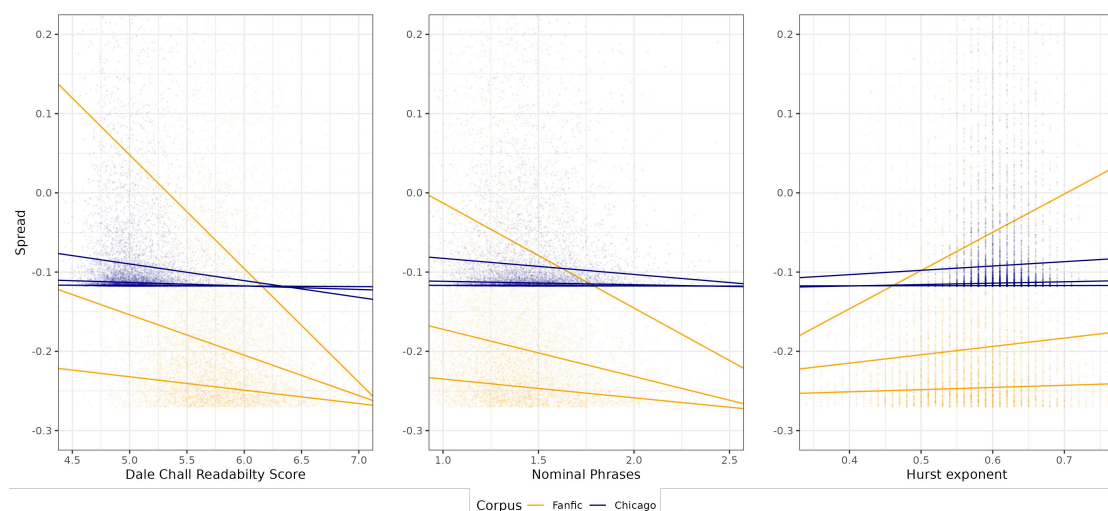

Figure 3: Quantile regression lines between textual feature scores and **spread** proxies across the two corpora. Spread proxies have been standardized for the visualization.

conceptualize "quality" and these different conceptions are related to different preferences - users on Goodreads have different preferences than literary critics. As such, the kudos/hits ratio might be more similar to other kinds of quality proxies than the Goodreads average rating.

4.3. Canonical Fiction and Fanfiction

While appreciation shows this complete opposite pattern (also visualized on Figure 2 and Figure 3), these seemingly divergent trends in reader appreciation are not necessarily a departure

729

Table 5

Mean textual feature score for each canon group in **Fanfiction**. All differences between non-canon and canon are significant at $p < .001$ (using the Mann-Whitney U test).

	50% split		75% split		87.5% split	
	Non-canon	Canon	Non-canon	Canon	Non-canon	Canon
Dale Chall	5.66	5.81	5.69	5.84	5.71	5.87
Nominal Ratio	1.40	1.42	1.40	1.44	1.40	1.46
Hurst	0.58	0.56	0.57	0.55	0.57	0.55

of fanfiction from edited literature patterns. Looking at Wu, Moreira, Nielbo, and Bizzoni [85] and Bizzoni, Moreira, Lassen, Thomsen, and Nielbo [9], from which we took some of these metrics, it seems that fanfictions' reader appreciation might be displaying patterns that are similar to another category of reception, that of *canonical fiction*.

As we detailed in Section 3.3, for the second part of our study we compared canonical fiction, as defined in Bizzoni, Moreira, Lassen, Thomsen, and Nielbo [9], with the most appreciated fanfictions (that we call, for comparison purposes, "canonical" fanfiction).

We show the main results in Table 6 and Table 5. When comparing increasingly exclusive subsets of highly-appreciated fanfiction, the Dale-Chall Readability score and the nominal ratio appear consistently higher for the canonical groups. In fact, they become higher the more exclusive the "quality group" becomes. There is consistently a higher Dale Chall readability score and nominal ratio in the canonical group, and this score also increases as the 'canon' group becomes more exclusive (i.e., the threshold increases). The same is evident for published literature.

Surprisingly, the Dale-Chall score for "non-canonical" fanfiction also increases as the threshold increases, whereas the nominal ratio remains the same. This indicates that the Dale Chall score has a more linear relationship with the kudos/hits ratio, whereas the mean nominal ratio for canonical fanfiction might be driven by few, greatly appreciated fanfics with a high nominal ratio.

For the Hurst exponent, the similarities between published and unpublished literature are less obvious, as there is no difference between canonical and non-canonical fiction in Hurst score, but canonical fanfics consistently have a lower Hurst exponent than the non-canonical. Previous works found a link between Hurst exponent and perceived literary quality when comparing bestsellers to high-brow texts [5], but the split between canon and non-canon might be slightly too crude to pick it up.

4.4. Comparing fandoms

So far, we have treated the fandom corpus as one whole. However, as fandoms are known to develop unique characteristics linked to their original texts [33, 21], the same textual features could be used differently in different fandoms within our corpus.

As reported on Table 7, both the Dale Chall Readability score and the nominal ratio show a difference between the three groups, while the Hurst exponent is equal across groups. The Dale

Table 6

Mean textual feature score for each canon group in **Chicago**.

	Canon		Prizes	
	Non-Canonical	Canon	Non-Prizes	Prizes
Dale Chall	5.08	5.29	5.10	5.15
Nominal Ratio	1.51	1.64	1.51	1.56
Hurst	0.61	0.61	0.61	0.60

Table 7

Mean textual feature score for each fandom.

	Percy Jackson	*Harry Potter*	*Lord of the Rings*
Dale Chall	5.63	5.74	5.82
Nominal Ratio	1.33	1.36	1.52
Hurst	0.57	0.57	0.57

Chall score and nominal ratio both show a similar pattern: *Percy Jackson* fanfiction has the most readable and least demanding style, while *Lord of the Rings* fanfiction appears to be the least readable and most demanding in style. This could indicate a sort of 'tide effect', meaning the writing style of the source texts gets integrated into the fanfiction. *Lord of the Rings* has been described as prose-heavy and using flowery descriptions [29], while *Percy Jackson* is written in very casual language that is also quite slang-heavy [58]. The Hurst exponent does indicate, though, that fanfiction as a medium also has textual features that transcend the writing style of the original author. The equal Hurst exponent across groups might be a product of the medium of fanfiction which lends itself to a less predictable story arc as compared to published fiction.

As such, not only are these features useful in illuminating aspects of literary quality, they can also show in which ways fanfics have both community-specific trends and traits that hold across fan groups.

5. Discussion

Research on literary quality has identified its relevant patterns working almost exclusively on published literature, with its socioeconomic filters of editorial houses, marketing campaigns, newspapers, anthologies, professional criticism, and so forth [80]. We have looked for some of the same patterns in the world of online fanfiction. Less constrained by these filters, online platforms tend to level access to literary production, where the distinction between writers and readers is more blurred, and where the production and consumption of texts is faster [32, 43]. Our main findings are the following:

1. Fanfics and published literature have a different overall textual profile when it comes to readability, nominal ratio, and Hurst exponent of their sentiment arcs
2. Despite these differences, the same features that appear to correlate with spread in published literature can be found in fanfiction: more readable texts with a lower nominal

ratio and a more coherent/predictable arc have a larger spread.

3. Fanfics showing higher levels of reader appreciation behave similarly to novels included in the literary canon and long-listed for high-brow awards, displaying a more challenging prose and higher nominal ratio. They also exhibit a reduced Hurst exponent for their sentiment arcs, a pattern found in other works looking at bestsellers vs high-brow literature.

4. Fanfics seem to mirror the expected level of complexity of the originals. LOTR fanfics lean towards the canonical style, while Percy Jackson fanfics lean towards the popular/robust strategies, and Harry Potter fanfics fall somewhere in the middle.

Overall, fanfics tend to show signs of a different super-style (beyond the style of individual authors) as compared to published novels, mixing traits that are usually distinct (i.e., they exhibit less nominal style but are harder to read). On the other hand, just as in published books, fanfics that use "robust" communication strategies, i.e., more readable and less cognitively challenging writing spread more. These findings first of all support the idea often put forward in qualitative analyses, that fanfiction differs from traditional fiction in its overall traits [59, 74]. Secondly, despite these overall differences, it supports the general interpretation for literary reception set forward in Wu, Moreira, Nielbo, and Bizzoni [85]: Different communicative strategies are used by popular and high-brow texts, both relating to robust communication through "noisy channels" and capitalizing on increased linguistic and narrative complexity at the price of higher cognitive loads. Despite the presence of confounding, spurious effects that inform the fanfiction domain as a whole (such as the run-on sentences and unstructured storylines), these same mechanisms might be in place when it comes to successfully achieving popularity and appreciation.

It is worth noting that we are not interpreting these phenomena in an absolute sense - e.g. showing whether fanfictions are "better" or "worse" than published literature. What we are seeing are parallel tendencies that mirror each other within the fanfiction and published corpora, so that the same stylistic and narrative features seem to point to similar reader behaviours, in terms of reception and perceived quality, despite the vastly different characteristics of the texts.

In the future, we would like to expand our analysis to larger and more diverse fanfiction corpora, as well as corpora of original fiction posted on online platforms (user-published). On the other hand, it would be greatly interesting to expand our set of linguistic and narrative measures beyond the three we have currently selected, to gauge more representative profiles of narrative styles in different domains (e.g. published vs. posted texts). With the present study and in future works, we aim to contribute towards a more nuanced understanding of narrative styles across diverse textual domains, potentially challenging prevalent notions of 'inferiority' attributed to self-published texts compared to established literature, and blurring the distinctions between these categories.

Acknowledgments

References

[1] M. Algee-Hewitt, S. Allison, M. Gemma, R. Heuser, F. Moretti, and H. Walser. *Canon/Archive. Large-scale Dynamics in the Literary Field.* Stanford Literary Lab, 2016. URL: https://litlab.stanford.edu/LiteraryLabPamphlet11.pdf.

[2] E. C. O. Alm. "Affect in*text and speech". PhD thesis. University of Illinois at Urbana-Champaign, 2008.

[3] J. L. Barnes. "Fanfiction as imaginary play: What fan-written stories can tell us about the cognitive science of fiction". In: *Poetics* 48 (2015), pp. 69–82.

[4] J. Barré, J.-B. Camps, and T. Poibeau. "Operationalizing Canonicity: A Quantitative Study of French 19th and 20th Century Literature". In: *Journal of Cultural Analytics* 8.3 (2023). DOI: 10.22148/001c.88113.

[5] Y. Bizzoni, P. Feldkamp, I. M. Lassen, M. Jacobsen, M. R. Thomsen, and K. Nielbo. *Good Books are Complex Matters: Gauging Complexity Profiles Across Diverse Categories of Perceived Literary Quality.* 2024. DOI: 10.48550/arXiv.2404.04022. URL: http://arxiv.org/abs/2404.04022.

[6] Y. Bizzoni, P. Feldkamp, and K. L. Nielbo. "Global Coherence, Local Uncertainty - Towards a Theoretical Framework for Assessing Literary Quality". In: *Computaitonal Humanities Research 2024*. Tallinn: CEUR-WS.org, 2024.

[7] Y. Bizzoni, P. Moreira, N. Dwenger, I. Lassen, M. Thomsen, and K. Nielbo. "Good Reads and Easy Novels: Readability and Literary Quality in a Corpus of US-published Fiction". In: *Proceedings of the 24th Nordic Conference on Computational Linguistics (NoDaLiDa)*. Tórshavn, Faroe Islands: University of Tartu Library, 2023, pp. 42–51. URL: https://aclanthology.org/2023.nodalida-1.5.

[8] Y. Bizzoni, P. Moreira, M. R. Thomsen, and K. Nielbo. "Sentimental Matters - Predicting Literary Quality by Sentiment Analysis and Stylometric Features". In: *Proceedings of the 13th Workshop on Computational Approaches to Subjectivity, Sentiment, & Social Media Analysis*. Toronto, Canada: Association for Computational Linguistics, 2023, pp. 11–18. URL: https://aclanthology.org/2023.wassa-1.2.

[9] Y. Bizzoni, P. F. Moreira, I. M. S. Lassen, M. R. Thomsen, and K. Nielbo. "A Matter of Perspective: Building a Multi-Perspective Annotated Dataset for the Study of Literary Quality". In: *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*. Ed. by N. Calzolari, M.-Y. Kan, V. Hoste, A. Lenci, S. Sakti, and N. Xue. Torino, Italia: ELRA and ICCL, 2024, pp. 789–800. URL: https://aclanthology.org/2024.lrec-main.71.

[10] Y. Bizzoni, P. F. Moreira, M. R. Thomsen, and K. L. Nielbo. "The Fractality of Sentiment Arcs for Literary Quality Assessment: the Case of Nobel Laureates". In: *Journal of Data Mining & Digital Humanities* Nlp4dh (2023). DOI: 10.46298/jdmdh.11406.

[11] Y. Bizzoni, T. Peura, K. Nielbo, and M. Thomsen. "Fractal Sentiments and Fairy Tales-Fractal scaling of narrative arcs as predictor of the perceived quality of Andersen's fairy tales". In: *Journal of Data Mining & Digital Humanities* Nlp4dh (2022). DOI: 10.46298/jdmdh.9154.

[12] Y. Bizzoni, T. Peura, K. Nielbo, and M. Thomsen. "Fractality of sentiment arcs for literary quality assessment: The case of Nobel laureates". In: *Proceedings of the 2nd International Workshop on Natural Language Processing for Digital Humanities*. Taipei, Taiwan: Association for Computational Linguistics, 2022, pp. 31–41. URL: https://aclanthology.org/2022.nlp4dh-1.5.

[13] Y. Bizzoni, T. Peura, M. R. Thomsen, and K. Nielbo. "Sentiment Dynamics of Success: Fractal Scaling of Story Arcs Predicts Reader Preferences". In: *Proceedings of the Workshop on Natural Language Processing for Digital Humanities*. NIT Silchar, India: NLP Association of India (NLPAI), 2021, pp. 1–6. URL: https://aclanthology.org/2021.nlp4dh-1.1.

[14] R. W. Black. "Language, culture, and identity in online fanfiction". In: *E-learning and Digital Media* 3.2 (2006), pp. 170–184.

[15] A. Brachmann and C. Redies. "Computational and Experimental Approaches to Visual Aesthetics". In: *Frontiers in Computational Neuroscience* 11 (2017), p. 102. DOI: 10.3389/fncom.2017.00102.

[16] J. Brottrager, A. Stahl, A. Arslan, U. Brandes, and T. Weitin. "Modeling and Predicting Literary Reception". In: *Journal of Computational Literary Studies* 1.1 (2022), pp. 1–27.

[17] K. Busse. *Framing fan fiction: Literary and social practices in fan fiction communities.* Iowa City: University of Iowa Press, 2017.

[18] E. Cambria, D. Das, S. Bandyopadhyay, and A. Feraco. "Affective computing and sentiment analysis". In: *A practical guide to sentiment analysis.* Cham: Springer, 2017, pp. 1–10.

[19] J. Cheng. "Fleshing out models of gender in English-language novels (1850–2000)". In: *Journal of Cultural Analytics* 5.1 (2020), p. 11652. DOI: 10.22148/001c.11652.

[20] K. Cheng, J. Li, J. Tang, and H. Liu. "Unsupervised sentiment analysis with signed social networks". In: *Thirty-First AAAI Conference on Artificial Intelligence.* San Francisco, 2017.

[21] F. Coppa. *The Fanfiction Reader: Folk tales for the digital age.* Ann Arbor, Michigan: University of Michigan Press, 2017.

[22] J. Cordeiro, P. R. M. Inácio, and D. A. B. Fernandes. "Fractal Beauty in Text". In: *Progress in Artificial Intelligence.* Ed. by F. Pereira, P. Machado, E. Costa, and A. Cardoso. Lecture Notes in Computer Science. Cham: Springer International Publishing, 2015, pp. 796–802. DOI: 10.1007/978-3-319-23485-4_80.

[23] A. van Cranenburgh and R. Bod. "A Data-Oriented Model of Literary Language". In: *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers.* Valencia, Spain: Association for Computational Linguistics, 2017, pp. 1228–1238.

[24]   T. Crosbie, T. French, and M. Conrad. "Towards a Model for Replicating Aesthetic Literary Appreciation". In: *Proceedings of the Fifth Workshop on Semantic Web Information Management*. Swim '13. New York, New York: Association for Computing Machinery, 2013, pp. 1–4. DOI: 10.1145/2484712.2484720. URL: https://doi.org/10.1145/2484712.2484720.

[25]   J. D. Culler. *The literary in theory*. Stanford: Stanford University Press, 2007.

[26]   J. S. Curwood, A. M. Magnifico, and J. C. Lammers. "Writing in the wild: Writers' motivation in fan-based affinity spaces". In: *Journal of Adolescent & Adult Literacy* 56.8 (2013), pp. 677–685.

[27]   S. Degaetano-Ortlieb and E. Teich. "Toward an optimal code for communication: The case of scientific English". In: *Corpus Linguistics and Linguistic Theory* 18.1 (2022), pp. 175–207.

[28]   I.-A. Drobot. "Affective Narratology. The Emotional Structure of Stories". In: *Philologica Jassyensia* 9.2 (2013), p. 338.

[29]   M. D. Drout. "Tolkien's prose style and its literary and rhetorical effects". In: *Tolkien Studies* 1.1 (2004), pp. 137–163.

[30]   J. Duggan. "Who writes Harry Potter fan fiction? Passionate detachment,"zooming out," and fan fiction paratexts on AO3". In: *Transformative Works and Cultures* 34 (2020), pp. 1–25.

[31]   T. Eagleton. *Literary theory: An introduction*. Malden: John Wiley & Sons, 2011.

[32]   S. Evans, K. Davis, A. Evans, J. A. Campbell, D. P. Randall, K. Yin, and C. Aragon. "More than peer production: Fanfiction communities as sites of distributed mentoring". In: *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing*. New York, 2017, pp. 259–272.

[33]   J. Fathallah. *Fanfiction and the author: How fanfic changes popular cultural texts*. Amsterdam: Amsterdam University Press, 2017.

[34]   P. Feldkamp, Y. Bizzoni, M. R. Thomsen, and K. L. Nielbo. *Measuring Literary Quality. Proxies and Perspectives*. Report. Darmstadt, 2024. DOI: 10.26083/tuprints-00027391. URL: https://tuprints.ulb.tu-darmstadt.de/27391/.

[35]   V. Ganjigunte Ashok, S. Feng, and Y. Choi. "Success with Style: Using Writing Style to Predict the Success of Novels". In: *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*. Seattle, Washington, USA: Association for Computational Linguistics, 2013, pp. 1753–1764. URL: https://aclanthology.org/D13-1181.

[36]   J. Guillory. *Cultural Capital: The Problem of Literary Canon Formation*. Chicago, IL: University of Chicago Press, 1995. URL: https://press.uchicago.edu/ucp/books/book/chicago/C/bo3634644.html.

[37]   R. von Hallberg. "Editor's Introduction". In: *Critical Inquiry* 10.1 (1983), pp. iii–vi. URL: https://www.jstor.org/stable/1343403.

[38]   E. Hemingway. *On Writing*. Ed. by L. W. Phillips. New York: Touchstone, 1999.

[39] M. Honnibal and M. Johnson. "An Improved Non-monotonic Transition System for Dependency Parsing". In: *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Ed. by L. Màrquez, C. Callison-Burch, and J. Su. Lisbon, Portugal: Association for Computational Linguistics, 2015, pp. 1373–1378. DOI: 10.18653/v1/D15-1162. URL: https://aclanthology.org/D15-1162.

[40] Q. Hu, B. Liu, M. R. Thomsen, J. Gao, and K. L. Nielbo. "Dynamic evolution of sentiments in Never Let Me Go: Insights from multifractal theory and its implications for literary analysis". In: *Digital Scholarship in the Humanities* 36.2 (2020), pp. 322–332. DOI: 10.1093/llc/fqz092.

[41] H. E. Hurst. "Long-term storage capacity of reservoirs". In: *Transactions of the American society of civil engineers* 116.1 (1951), pp. 770–799.

[42] C. Hutto and E. Gilbert. "VADER: A parsimonious rule-based model for sentiment analysis of social media text". In: *Proceedings of the international AAAI conference on web and social media*. Vol. 8. 1. 2014, pp. 216–225. DOI: 10.1609/icwsm.v8i1.14550.

[43] A. Jamison. *Fic: Why fanfiction is taking over the world*. BenBella Books, Inc., 2013.

[44] K. Jautze, C. Koolen, A. van Cranenburgh, and H. de Jong. "From high heels to weed attics: a syntactic investigation of chick lit and literature". In: *Proceedings of the Workshop on Computational Linguistics for Literature*. Ed. by D. Elson, A. Kazantseva, and S. Szpakowicz. Atlanta, Georgia: Association for Computational Linguistics, 2013, pp. 72–81.

[45] H. Jenkins. *Textual poachers: Television fans and participatory culture*. Routledge, 1992.

[46] M. Jockers. *A Novel Method for Detecting Plot*. 2014. URL: https://www.matthewjockers.net/2014/06/05/a-novel-method-for-detecting-plot/.

[47] M. Jockers. *Revealing Sentiment and Plot Arcs with the Syuzhet Package*. 2015. URL: https://www.matthewjockers.net/2015/02/02/syuzhet/.

[48] S. King. *On Writing: A Memoir of the Craft*. Anniversary. New York: Scribner, 2010.

[49] C. Koolen, K. van Dalen-Oskam, A. v. Cranenburgh, and E. Nagelhout. "Literary Quality in the Eye of the Dutch Reader: The National Reader Survey". In: *Poetics* 79 (2020), pp. 1–13. DOI: 10.1016/j.poetic.2020.101439.

[50] A. Kustritz. "They All Lived Happily Ever After. Obviously.: Realism and Utopia in Game of Thrones-Based Alternate Universe Fairy Tale Fan Fiction". In: *Humanities* 5.2 (2016), p. 43.

[51] H. Long and T. Roland. *US Novel Corpus*. Tech. rep. Textual Optic Labs, University of Chicago, 2016. URL: http://icame.uib.no/brown/bcm.html.

[52] S. Maharjan, J. Arevalo, M. Montes, F. A. González, and T. Solorio. "A Multi-task Approach to Predict Likability of Books". In: *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*. Valencia, Spain: Association for Computational Linguistics, 2017, pp. 1217–1227. URL: https://aclanthology.org/E17-1114.

[53] S. Maharjan, S. Kar, M. Montes, F. A. González, and T. Solorio. "Letting Emotions Flow: Success Prediction by Modeling the Flow of Emotions in Books". In: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*. New Orleans, Louisiana: Association for Computational Linguistics, 2018, pp. 259–265. DOI: 10.18653/v1/N18-2042.

[54] S. Maharjan, S. Kar, M. Montes, F. A. González, and T. Solorio. "Letting Emotions Flow: Success Prediction by Modeling the Flow of Emotions in Books". In: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Volume 2, Short Papers*. New Orleans, Louisiana: Association for Computational Linguistics, 2018, pp. 259–265. URL: https://aclanthology.org/N18-2042.

[55] A. Mattei, D. Brunato, and F. Dell'Orletta. "The Style of a Successful Story: a Computational Study on the Fanfiction Genre". In: *Computational Linguistics CLiC-it 2020* (2020), p. 284.

[56] J. McDonough and A. Herczyński. "Fractal patterns in music". In: *Chaos, Solitons & Fractals* 170 (2023), p. 113315. DOI: 10.1016/j.chaos.2023.113315.

[57] L. H. McGavin. "Creativity as Information: Measuring Aesthetic Attractions". In: *Nonlinear Dynamics, Psychology, and Life Sciences* 1.3 (1997), pp. 203–226. DOI: 10.1023/a:1022342915622. URL: https://doi.org/10.1023/A:1022342915622.

[58] R. Mead. "The Percy Jackson Problem". In: *The New Yorker* (Oct. 22, 2014). URL: https://www.newyorker.com/culture/cultural-comment/percy-jackson-problem.

[59] K. Miłkowska-Samul. ""How come you're not shipping them??? They're canon"". In: *Kwartalnik Neofilologiczny* 2 (2019).

[60] S. Milli and D. Bamman. "Beyond canonical texts: A computational analysis of fanfiction". In: *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Austin, Texas, 2016, pp. 2048–2053.

[61] L. M. Mixer. "And then they boned: An analysis of fanfiction and its influence on sexual development". MA thesis. Humboldt State University, 2018.

[62] L. Nakamura. ""Words with Friends": Socially Networked Reading on Goodreads". In: *Pmla* 128.1 (2013), pp. 238–243. DOI: 10.1632/pmla.2013.128.1.238.

[63] D. Nguyen, S. Zigmond, S. Glassco, B. Tran, and P. J. Giabbanelli. "Big data meets storytelling: using machine learning to predict popular fanfiction". In: *Social Network Analysis and Mining* 14.1 (2024), p. 58.

[64] E. Öhman and R. Rossi. "Affect as Proxy for Mood". In: *Journal of Data Mining and Digital Humanities* Special Issue: Natural Language Processing for Digital Humanities (2023).

[65] F. Pianzola, A. Acerbi, and S. Rebora. "Cultural accumulation and improvement in online fan fiction". In: *CEUR Workshop Proceedings (Vol. 2723)*. Amsterdam, 2020. DOI: 10.31219/osf.io/4wjnm.

[66] F. Pianzola, S. Sharma, and F. Tsiwah. *A Computational Analysis linking the Emotion Arcs of Books and Reader Response*. 2023. URL: https://discourse.igelsociety.org/t/a-computational-analysis-linking-the-emotion-arcs-of-books-and-reader-response/426.

[67] Z. Pilz. *Bad Fiction and the Brain. The Effect of Intentionally Bad Written Fiction on the Brain*. Vol. 6. Darmstadt: Universitäts-und Landesbibliothek Darmstadt, 2023.

[68] D. Pimenova. "Fan Fiction: Between Text, Conversation, And Game Daria Pimenova". In: *Internet Fictions* (2008), p. 44.

[69] S. Pugh. *The democratic genre: Fan fiction in a literary context*. Brigend: Seren, 2005.

[70] B. Qian and K. Rasheed. "Hurst exponent and financial market predictability". In: *IASTED conference on Financial Engineering and Applications*. Proceedings of the IASTED International Conference Cambridge, MA. 2004, pp. 203–209.

[71] A. J. Reagan, L. Mitchell, D. Kiley, C. M. Danforth, and P. S. Dodds. "The Emotional Arcs of Stories Are Dominated by Six Basic Shapes". In: *EPJ Data Science* 5.1 (2016), pp. 1–12. DOI: 10.1140/epjds/s13688-016-0093-1. URL: https://epjdatascience.springeropen.com/articles/10.1140/epjds/s13688-016-0093-1.

[72] A. Sereda. "'Dirty stories saved my life': Fanfiction as a source of emotional support". MA thesis. Charles University, 2019.

[73] L. A. Sherman. *Analytics of Literature: A Manual for the Objective Study of English Prose and Poetry*. Athenaeum Press. Ginn, 1893.

[74] M. G. Sindoni. ""I really have no idea what non-fandom people do with their lives." A multimodal and corpus-based analysis of fanfiction". In: *Lingue e Linguaggi* 13 (2015), pp. 277–300.

[75] Z. Sourati Hassan Zadeh, N. Sabri, H. Chamani, and B. Bahrak. "Quantitative analysis of fanfictions' popularity". In: *Social Network Analysis and Mining* 12.1 (2022), p. 42.

[76] W. Strunk, E. B. White, and R. Angell. *The Elements of Style*. Ed. by T. Editor. 4th edition. New York, Munich: Pearson, 1999.

[77] B. Thomas. "Canons and fanons: Literary fanfiction online". In: *Dichtung Digital. Journal für Kunst und Kultur digitaler Medien* 9.1 (2007), pp. 1–11.

[78] B. Thomas. "What is fanfiction and why are people saying such nice things about it??" In: *Storyworlds: A Journal of Narrative Studies* 3 (2011), pp. 1–24.

[79] T. Underwood, D. Bamman, and S. Lee. "The transformation of gender in English-language fiction". In: *Journal of Cultural Analytics* 3.2 (2018), p. 11035. DOI: 10.22148/16.019.

[80] W. Van Peer. *The quality of literature: Linguistic studies in literary evaluation*. Vol. 4. John Benjamins Publishing, 2008.

[81] M. Verboord. "Female bestsellers: A cross-national study of gender inequality and the popular–highbrow culture divide in fiction book production, 1960–2009". In: *European Journal of Communication* 27.4 (2012), pp. 395–409. DOI: 10.1177/0267323112459433.

[82]   M. Walsh and M. Antoniak. "The Goodreads 'Classics': A Computational Study of Readers, Amazon, and Crowdsourced Amateur Criticism". In: *Journal of Cultural Analytics* 4 (2021), pp. 243–287.

[83]   X. Wang, B. Yucesoy, O. Varol, T. Eliassi-Rad, and A.-L. Barabási. "Success in Books: Predicting Book Sales Before Publication". In: *EPJ Data Science* 8.1 (2019). DOI: 10.1140/epjds/s13688-019-0208-6.

[84]   P. West. "In Defense of Purple Prose". In: *The New York Times* (1985). URL: https://www.nytimes.com/1985/12/15/books/in-defense-of-purple-prose.html.

[85]   Y. Wu, P. F. Moreira, K. L. Nielbo, and Y. Bizzoni. "Perplexing Canon: A study on GPT-based perplexity for canonical and non-canonical literary works". In: *To appear in: Proceedings of the 8th Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature*. St. Julians, Malta: Association for Computational Linguistics, 2024.