



What if fan-fiction, but also coding:
**How does fan-fiction differ in style to its original canon and does it
affect its success?**

Rohan Cyr Lambert¹
Supervisors: Chenxu Hao¹, Ivan Kondyurin¹
Responsible Professor: Hayley Hung¹
¹EEMCS, Delft University of Technology, The Netherlands

A Thesis Submitted to EEMCS Faculty Delft University of Technology,
In Partial Fulfilment of the Requirements
For the Bachelor of Computer Science and Engineering
January 26, 2025

Name of the student: Rohan Cyr Lambert
Final project course: CSE3000 Research Project
Thesis committee: Hayley Hung, Ivan Kondyurin, Elmar Eisenmann

An electronic version of this thesis is available at <http://repository.tudelft.nl/>.

Abstract

Natural language processing, specifically style, is not explored significantly in the context of fan-fiction. By using function word frequency analysis, this paper explores the similarity in style between original works and fan-fictions derived from them as well as the impact of those stylistic similarities on the fan-fictions' successes. Investigating the works of Worm and Narnia and utilising a control set, this hypothesis is examined. Similarity in style seems to be variable, but it is suggested that certain stylistic features lead to success more than other regardless of the original work's style. More research would be needed to confirm those numbers.

1 Introduction

Natural language processing, also known as NLP, is the field of study of text with computers, from speech recognition to generative AIs. While there is a tremendous amount of work that has been performed on classical literature, journals, speeches, and even transcripts of discussions, fan-fiction (works of literature based on existing works of fiction, though they often change components to make unique stories [1, 2, 3]) is one of lesser explored types of text in the field of NLP.

This study focuses on stylometry, a field of NLP focused on measuring the "style" of writing in a methodical and quantifiable way. It has mostly been explored in the context of Authorship Attribution, with features such as word frequencies, errors, character types, sentence structures or semantic dependencies [4].

This study aims to look at the difference between fan-fictions and their original source (which can be referred to as canon) and answer the question: **How does fan-fiction differ in style to its original canon and does it affect its success?** This is interesting to look at because style hasn't often been studied in its relation to the success of its works, as it is usually quite unique to its author. Yet in the context of fan-fiction, this makes sense to look at. Most often people tend to read fan-fiction because they want more of the original material. The hypothesis is then that fan-fiction closer to their original work in style may garner more success as it may fulfil that need of more of the original work that fan-fiction readers are looking for.

To answer our main question: **How does fan-fiction differ in style to its original canon and does it affect its success?** We will try to answer these two sub-questions: **Does fan-fiction inherit style from its canon?** and **Does a similar style to the canon increase a fan-fiction's success?**

In this paper we hope to contribute to the current study of style and fan-fiction in NLP by joining these two fields together. It posits that measuring style may be used for more than just authorship attribution and has potential in identifying the inspiration of authors and a style's links to the success of written works.

The background for this research will be described in section 2 alongside the related works. In section 3 the method for the experiment is described. The results obtained from the experiment are presented in section 4. section 5 will discuss the ethical considerations of this study and what steps were taken to act upon them. Further analysis of the data and discussion on its relevance will happen in section 6. Finally, section 7 will contain the final conclusions of this paper and suggestions for future works on the topic.

2 Background and related works

2.1 Background

2.1.1 Stylometry

The field of stylometry aims to rationalise writing using methodical and quantifiable methods. It establishes a style of a work which has been used primarily in the recent past to identify authors of works as style has been shown consistent in multiple texts by the same author.

Stylometry has been used much less in associating style with success, nor has it been used significantly when looking at the influences that affect an author's writing. It however looks for quantifiable data and this makes it suitable for this study.

Function word frequency is one of the most successful features when measuring the style of works in the context of authorship attribution. Although it can seem unintuitive it has been shown to provide successful results since the 1960s [5, 6].

Function words are the words that support sentences. While content words such as "cat" or "speak" tell us the content of a sentence, function words such as "yes", "before", or "or" serve to bind that sentence together and usually give context to the content of said sentence. These consist of adpositions, articles, conjunctions... etc.

2.1.2 Fan-fiction

"Fan-fiction" refers to works of fiction created using parts of an original work, which can include characters, settings or even concepts. However, while these works do use some features from their original source, they tend to differ in a lot of ways. They are called fan-fiction as they are usually created by fans, and while some may be considered aspiring or even full-fledged writers, a lot of fan-fiction authors are amateurs and hobbyist, often writing their first stories. This, as well as the absence of traditional publishers, means that while a lot of fan-fiction may be of a high quality, many other works may be less well written or developed.

2.2 Related Work

While fan-fiction is less studied than regular literature, there are still some studies that are relevant to this paper:

- **Changes to the original work.** It has been shown that there are often quite some differences between the original work and their fan-fictions, such as emphasis on certain character or a change in setting, relationship or even genre of the work [7, 1, 2, 3]. People do not want to read or write the same exact story.
- **Features of fan-fiction tend to have an impact of its success.** Correlations have been shown between the success of a fan work and some features such as themes or emotions [8].

The main study relevant to this paper in stylometry is:

- **Style has been shown to correlate to success.** Research has been done on the style of successful fan-fiction suggesting that certain stylistic characteristics have significant correlation with success of a work. This was done without consideration for

the original work, however, and rather suggests that certain stylistic characteristics are just more attractive to readers of fan-fiction in general [9].

3 Method

This section describes the methods used to obtain our data as well as the specifications of our dataset.

As mentioned earlier in this paper, stylometry has mostly been used in authorship attribution. The features which are successful in that context often are not particularly intuitive [6]. Stylistic similarity is measured in this paper using solely functional word frequency, one of the most common and effective measure of style.

3.1 Function word frequency

The list of function words used in this experiment consists of 277 words sourced from Semantic Similarity¹ from which the pronouns were excluded. Pronouns, contrary to most other function words, can be related to genre of content of a work and even to other aspects like the gender of the author [10, 11], hence their exclusion.

The function words will be counted as ngrams of size 1, assuming "bag of words", meaning that the order and position of these words in the text do not have a significant impact on its style.

Looking at the frequency of these function-words has been a successful way to measure stylometry and differentiate authors in modern stylistics studies.

3.2 Stylo

The tool that will be used for measuring function word frequency is "Stylo", a package for high level computational text analysis in R. [12] Stylo counts the presence of these function words in a text and produces vectors from this data which is then translated to a "distance" in style between all the works that were processed. This means that for any given two works we can obtain a number that describes how far apart they are in style, the lower the number, the closer stylistically. These distance numbers are measured relatively, they are only valid when looking at the set of works that is measured together. This means that only comparison between the works processed is possible.

The vectors for each work were calculated by Stylo using the 100 most frequent function words found in the works. While it can be done with all of the function words in the list used, this was done to save on processing time and because the rarer words such as "hereupon" or "whyever" are unlikely to appear more than a handful of times, even in a dataset of around 8000 works. This should have lead to a negligible loss in precision of stylistic analysis, especially given a lot of the works are short.

3.3 Dataset

The dataset used for this experiment first consists of two original works of fiction: The book "Worm" by John C. "Wildbow" McCrae and the series of books The Chronicles of Narnia by C. S. Lewis. These books are referenced in this paper as the "canon" of their respective fan-fictions, and The Chronicles of Narnia will also be referenced as simply "Narnia".

¹<https://semanticsimilarity.wordpress.com/function-word-lists/>

The rest of the dataset consists of 3 sets of fan-fictions, a set of Worm fan-fiction, a second set of Narnia fan-fiction and a control set made of fan-fictions from various sources. The control set serves to make a point of comparison with the greater ensemble of existing fan-fiction.

3.3.1 Worm

Worm is a science fiction/superhero webnovel published from 2011 to 2013². This novel was first chosen as the original work for this experiment because it both had a large enough amount of fan-fictions and only exists in the form of a book. Since we are looking at stylistic similarities between the fan-fictions and the canon, we want the author of the fan-fiction to have experienced said original work as a book rather than as a game, series or movie, which would be possible with works like the Witcher, or Lord of the Rings.

Worm was retrieved from its website using a web-scraper made specifically for it³, then simply converted from epub (electronic publication, a file containing a book for book reading programs) to plaintext. This version makes some fixes to the original text, mostly centred on typos, punctuation, capitalisation and formatting. These changes should not affect our results as, while these above mentioned features are used in stylometry, they are not used in this study.

The choice of Worm might have an effect on our data, however. Worm shares a lot of similarities with fan-fiction: It is split in an array of short chapters which were uploaded one by one over time, online and free to access. There was also contact between the author and the work's community while it was being written, which is common in fan-fiction.

3.3.2 Narnia

The Chronicles of Narnia is a series of 7 portal fantasy/coming of age books published from 1950 to 1956. These novels were added to this study partway through. While making the first iteration of the control set of fan-fiction used in this study, fan-fictions of Narnia were selected. It was then decided that, given the shared characteristics between Worm and fan-fiction, a second data point of more established literature would be useful.

The Narnia series of books was obtained online as a single pdf which was then converted to plaintext. Narnia is used in this study as a single work, assuming that the authors style did not change significantly during the 6 years between the first and last novel of the series and that fan-fiction of Narnia would be influenced by the whole series.

Narnia also has a few characteristics that may affect our data. It was written 60 years before Worm and the fan-fictions of our dataset, which, if general style of writing evolved over time, might affect our results. Contrary to Worm, Narnia has movie adaptations of the first 3 books, so some people writing fan-fiction on it potentially might have never read the books to begin with. Finally, Narnia has been translated to numerous languages, and so the writers of fan-fiction may not have read the series in English.

Even with these downsides, this dataset is useful as Worm itself has issues and having a second set will allow for a more balanced study.

²<https://parahumans.wordpress.com/>

³<https://github.com/domenic/worm-scraper>

3.3.3 Fan-fiction works

The sets of fan-fiction used in this study were sourced from the website Archive Of Our Own, also known as AO3⁴. They were obtained using a dedicated web-scraper⁵, which was modified after obtaining the set of Worm fan-fiction as it had stopped working properly when the control and Narnia sets were downloaded.

The first set of fan-fiction consists of 2911 Worm fan-works. 2920 works were taken from AO3 on the 26th of November 2024. The selection parameters were all the fan-fictions of Worm in English and excluding crossovers. We requested only English works as the canon is in English as well, and comparing style would make little sense across language. Crossovers are excluded to avoid variables of additional original works. After obtaining those 2920 fan-fictions, 9 were removed because they were empty of text, containing either only pictures, or a single link. Removing the fictions that were short was considered but decided against. Notably, the shortest fan-fiction of the set⁶ consists of only 4 words but is still a valid piece of fan-work that makes sense in the context of Worm.

The second set of fan-fiction contains 2577 Narnia fan-works, 2614 were obtained from AO3 on the 17th of January. The selection parameters were all the fan-fictions of Narnia in English, excluding crossovers, released between January 2014 and November 2024 included. These parameters were selected to make this dataset as close as possible to the set of Worm fan-fiction. Similarly to the Worm dataset, out of the 2614 works that were retrieved, 37 were removed for being empty, links or pictures.

The control set of fan-fiction is made of 2482 fan-works. It is formed from 5 subsets all obtained on the 17th of January, the selection parameters were a random sample of 500 fan-fictions from their respective canons in English, excluding crossovers, released between January 2014 and November 2024 included. These parameters were chosen to make this control set as comparable as possible to the Worm and Narnia datasets.

The original works for the control set were each taken from a different category on the AO3 website. Those categories being: Anime & Manga, Cartoons & Comics & Graphic Novels, Movies, TV Shows, and Video Games. For each of these categories, the original work with the corresponding amount of fan-fiction closest to Worm’s 2920 was chosen. This was done assuming that a similar amount of fan-fiction would equate to a similar amount of community involvement. Similarly to the previous sets, after obtaining each set of 500 fan-fictions, the empty ones that contained only pictures or links were removed. The information on that control set is detailed in Table 1.

Table 1: Distribution of fan-fictions in the control set

Canon	Genre	Medium	Year of release	Number of works
Batman: Under the Red Hood	Action/Superheroes	Cartoon/Movie	2010	489
Black clover	Shounen/Fantasy/Action	Anime/Manga	2015	500
Castle	Buddy-cop/Crime/Comedy/Drama	TV series	2009	497
Doki Doki Literature Club	Visual Novel/Horror	Video Game	2017	498
Transformers (Michael Bay Movies)	Action/Blockbuster/Science fiction	Movies	2007	498

⁴<https://archiveofourown.org/>

⁵<https://github.com/radiolarian/A03Scraper>

⁶<https://archiveofourown.org/works/29594592>

3.4 Processing results

Once the distance table between the 8000 works of fiction was created, the data needed to be comparable between different sets. This meant normalising to obtain values that, while they mean nothing on their own, are accurate cross set.

For each of the 7971 fan-fiction, the collected data was:

- Kudos. "Kudos" are the equivalents of social media "likes" on A03. They are the main way this paper measures the success of a piece of fan-fiction.
- Hits. "Hits" are the number of views/accesses to the page that a fan-fiction has gotten. While they could also be used as a measure of success, they are more volatile and so were not used for this.
- Distance to Narnia and Worm. These are the distances in style to Narnia and Worm for any given fan-fiction. The higher the value, the less they have in common stylistically.

From this, the data was normalised by using percentages of Kudos and Hits. By dividing the amount of kudos by the total amount of kudos in the sets (or subsets for the controls), the percentage of Kudos and Hits of a given fan-fiction was obtained. This is equivalent to saying "This fan-fiction has taken this percentage of this community's attention.". Then, the number obtained was multiplied by the amount of fan-fictions to account for the different set sizes. This final number, which will be referenced as "Normalised Kudos", can be used to compare the different sets. This value was calculated for the control set, both when looking at it as a whole and for each subset separately.

Fan-fiction, like many other online mediums, has an attention based economy, where few of the works proportionally get a lot of the attention. This is seen with the averages consistently being higher than the medians for interaction values, hence the uses of median as the preferred measure in this paper.

Kudos are used to measure the success of a fan-fiction. It was shown that 99.3% of interactions in online fan-fiction is positive [9], so comments could have been used, but Kudos were more useful as a set with more granularity, and very few fictions have no Kudos, while a large amount of them have no comments.

4 Results

This section presents the results of this study.

4.1 Results

The results of this experiment can be split in two parts: similarity of fan-fictions to their canons, and success compared to that similarity. These correspond to our two subquestions.

4.1.1 Similarity of fan-fictions to their canons

For each of the three sets, we have the difference in style, measured as "distance" to both Narnia and Worm. They are themselves measured at a distance in style of 0.2685, which is a high similarity. This is cemented in Figure 1 which shows that for all pieces of fan-fiction in each of the datasets, there is a high correlation between being similar to Worm and being similar to Narnia.




Figure 1: Comparison of distance in style to Narnia and Worm for each dataset


Figure 1 also shows other interesting points:

- Fan-fictions are on average more similar to Worm than Narnia, even Narnia fan-fictions.
- The fan-fictions most similar to Worm are Worm fan-fictions and the ones most similar to Narnia are Narnia fan-fictions.
- Narnia fan-fictions, as shown by the scale of Figure 1b, seem to have a greater variety in style than Worm fan-fictions or fan-fictions from the control set.


These points of data can be seen as well in Figure 2, which clearly shows shifted bell

curves, the higher variety of Narnia fan-fiction leading to a lower bell curve, leading Worm fan-fiction and the control set to be more similar in style to Narnia than Narnia fan-fiction.

On the other hand, Worm fan-fiction is significantly closer in style to Worm than either of the other sets, which leads it to also be much more similar in style to Narnia than the control and Narnia set of fan-fiction.



(a) Ranges of distance to Narnia



(b) Ranges of distance to Worm

Figure 2: Distribution of fan-fiction for each range of distance to Narnia and Worm

4.1.2 Success of fan-fictions relative to the similarity to their canons

The graphs which compare success for each set and subset of fan-fiction can be found in Appendix A.

Figure 3 Is a set of 3 out of 6 graphs which show the success of fan-fiction by range of distance to their original works. The 3 reciprocal graphs showing Worm fan-fiction success for ranges of distance to Narnia, the Narnia fan-fiction success for ranges of distance to Worm and the control set fan-fiction success for ranges of distance to Worm can be found in Appendix A. Due to the similarity between Worm and Narnia and the correlation in distance shown in Figure 1, the 3 graphs chosen have the same shape as the 3 in Appendix A


respectively.

Diving into Figure 3a shows a clear correlation between similarity to the original in style and success of the fiction with a few small spikes in the later values. This is the same when comparing worm fan-fiction success to their similarity to Narnia.


Figure 3b on the other hand tells a different story. While being more similar to Narnia seems to lead fan-fictions to be more successful, it is not reflected in the 3 fan-fictions that get within 0.2 distance, and the curve is much flatter. The spikes in the higher distance values that are minor when looking at Worm or the control set's fan-fictions are much more pronounced, reaching a success comparable to the success of the fan-fictions more similar to the original work.

The control set in Figure 3c is a similar shape to the Worm fan-fictions in Figure 3a, showing that for fan-fictions that belong to neither set, being similar to Worm and Narnia seems to still lead to a higher success.


It is important to note that for all 3 of these graphs, the number of data-points drop significantly in the higher distances, this is not shown in Figure 2 which stops as the numbers fall deeply, with each range past 1.1 dropping to numbers below 10, except for Narnia fan-fiction where the full drop-off happens at around 1.3.



(a) Worm fan-fiction compared with Worm



(b) Narnia fan-fiction compared with Narnia



(c) Control fan-fiction compared with Narnia

Figure 3: Median success of fan-fictions for each range of distance in style to their canons

5 Responsible Research

There are several points that must be discussed to examine potential ethical problems with this study.

- **Using other’s work.** The datasets used are taken from freely available works online. Reading through the data in this study would allow one to identify any given work, but all the data was already public to begin with. One issue that exists, however, is that to ensure reproducibility, it would be preferable for the works to be stored online in case the website became inaccessible in the future. However, that may interfere with an author’s preferences of if they wanted to delete their work from the internet. Hence, all the fictions used in this study will be kept online as plaintext and all the authors will be contacted and given the means to have their work removed from that archive.
- **Generative AI.** Generative AI was used in the making of the code for this project. Specifically, GitHub copilot was used to speed up the coding of some of the python scripts to manage the data and process it as CSV files. No generative AI was involved in the writing of this paper, nor in the thought process that lead to the decisions taken during this research process.
- **Dataset choices.** The original book (Worm) that this study was gonna be based on was arbitrarily chosen by the author. The criteria when choosing this book were that it was available in text easily, solely existed as a book and not another medium and that it had a significant number of fan-fictions. No systematic research to find suitable works was performed, the author instead thought of book series they had read or heard about and settled on the only one they could think of that covered these criteria. The subsequent datasets were chosen systematically instead by collecting every suitable original work listed on AO3 and taking the closest in set size.
- **Author’s bias.** During the analysis of the results in this study, it is important to note that this is written by a single author who has biases that they may not be aware of as well as some that they are aware of. Notably, the author has read the novel Worm used in this study, and regards it with high esteem. They have similarly read the first 3 Chronicles of Narnia books (in French) and have enjoyed them. These opinions are likely to have influenced their analysis.

6 Discussion

The aim of this paper is to answer the question: "How does fan-fiction differ in style to its original canon and does it affect its success?". While the results seem to give some definite answer to this question, there are some caveats that need to be discussed.

6.1 Does fan-fiction inherit style from its canon?

The answer to this question is inconclusive, while Worm fan-fiction is closer to Worm than an average fan-fiction, this is not the case for Narnia fan-fiction and Narnia. This could be due to a few factors, first Narnia was written in the 1950s, at least 60 years apart from the rest of this dataset and Worm, an hypothesis could be that style in English literature in general evolved. This seems unlikely to be the reason for this however as Worm and Narnia instead are very similar stylistically. Another factor that could be the case is the existence of the 3 Narnia movies, if the author of Narnia fan-fictions have never read the books and only seen the movies then there wouldn’t be a way for them to be stylistically influenced. This also seems unlikely a reason, after all the control set contains fan-fiction

from varied original mediums as well and is still more stylistically similar to Narnia than Narnia fan-fiction. This seems to be caused by the high variance in Narnia fan-fiction style. This variance is supported by Narnia having the proportionally largest amount of fan-fiction that had to be removed from the set for being images, audio fiction, videos or links. This paper posits that this difference is rather due to Narnia having a community from a different demographic as the other original work. As a coming of age book series from the 50s, the people who have experienced it might be quite different to the people that experienced the other original works all around 2010 to 2017.

The similarity to Worm still supports this question’s answer being yes.

Another point that is worth discussing is that while the main part of the corpus did not fit this inheritance for Narnia, the extremes in similarity do. The fan-fictions that were the closest to Worm were Worm fan-fictions, and those closest to Narnia were Narnia fan-fictions. This suggests that while it may not be true for all fan-fictions of a work, some of them may inherit a large amount of stylistic similarity.

6.2 Does a similar style to the canon increase a fan-fiction’s success?

This question is slightly muddled as well. Given the similarity between Worm and Narnia and the fact that the control set followed the same trend as the Worm fan-fiction set, the conclusion that we can take from our data is that it is not being similar to the canon that leads to the success of a fan-fiction, but rather being similar to the stylistic features shared by Worm and Narnia. Especially those of Worm. This comes back to the potential problems of using Worm mentioned earlier in this paper. Worm as a book was released in a way and format quite similar to fan-fiction, and given its success as a web serial, it could be considered as a very successful work in the category fan-fiction falls in. The similarity to Narnia could then either be a coincidence or a symptom of some stylistic features being an important factor in a book’s success, Narnia of course being a very successful book series. Stylistic features being a major factor in a fan-fiction success correspond to some of the results of previous studies [9]. It was considered that Worm and Narnia’s similarity in style be linked to their size and that a lot of fan-fiction are somewhat short, this was dismissed as both the longer and shorter fan-fiction were placed across the spectrum of stylistic similarity and difference.

It seems to be that given the shape of our results graphs, sharing the stylistic characteristics of Worm leads to greater success which decreases with increased distance in style and that while less prevalent, the few fan-fiction that are wildly different in style to it may also find some success even if diminished in comparison to the ones most similar.

7 Conclusions and Future Work

7.1 Conclusion

This study attempted to find correlation in the style of original works and the fan-works made from them. While a correlation seems to exist for one of the two datasets used, it is not the case for the other. The only correlation that seems shared by the two sets is the fact that the fan-fictions most similar in style to a given original work were fan-fictions from said original work.

The second part of this paper’s research was to then ask if works more similar in style would have greater success. This seemed to be the case for both of the main datasets, but this characteristic was also shared by the control set, suggesting that, more important than similarity to the original work, certain stylistic characteristics would be the cause of greater success in fan-fiction. Furthermore, that those characteristics may have simply been shared by the chosen original works.

7.2 Future Work

This study would benefit to be performed using different sets, notably ones where the original works that are looked at are not stylistically similar. Since Narnia as a set also seemed to differ quite a lot to both the Worm dataset and the control set, it might be worth investigating if it is an anomaly within fan-fiction.

Another line of research that may be worth looking into is using a variety of methods to measure style. This study originally wanted to use Event Density [13] as a more uncommon style feature, but there are a lot of options that may be looked at which might give a more comprehensive distinction of style.


Finally it may be worth expanding on using fan-fiction to compare author’s styles as the works that have inspired fan-fiction authors should intuitively be easy to find. This may lead to more discoveries in the formation and evolution of author’s styles.

References


- [1] J. L. Barnes, “Fanfiction as imaginary play: What fan-written stories can tell us about the cognitive science of fiction,” *Poetics*, vol. 48, pp. 69–82, Feb. 2015.
- [2] V. Van Steenhuyse, “The Writing and Reading of Fan Fiction and Transformation Theory,” *CLCWeb: Comparative Literature and Culture*, vol. 13, Dec. 2011.
- [3] B. Thomas, “What Is Fanfiction and Why Are People Saying Such Nice Things about It?,” *Storyworlds: A Journal of Narrative Studies*, vol. 3, pp. 1–24, 2011. Publisher: University of Nebraska Press.
- [4] E. Stamatatos, “A survey of modern authorship attribution methods,” *Journal of the American Society for Information Science and Technology*, vol. 60, no. 3, pp. 538–556, 2009. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/asi.21001>.
- [5] F. J. Tweedie, S. Singh, and D. I. Holmes, “Neural Network Applications in Stylometry: The “Federalist Papers”,” *Computers and the Humanities*, vol. 30, no. 1, pp. 1–10, 1996. Publisher: Springer.
- [6] M. Kestemont, “Function Words in Authorship Attribution. From Black Magic to Theory?,” in *Proceedings of the 3rd Workshop on Computational Linguistics for Literature (CLFL)* (A. Feldman, A. Kazantseva, and S. Szpakowicz, eds.), (Gothenburg, Sweden), pp. 59–66, Association for Computational Linguistics, Apr. 2014.
- [7] 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* (J. Su, K. Duh, and X. Carreras, eds.), (Austin, Texas), pp. 2048–2053, Association for Computational Linguistics, Nov. 2016.

- [8] D. Nguyen, S. Zigmond, S. Glassco, B. Tran, and P. J. Giabbanelli, “Big data meets storytelling: using machine learning to predict popular fanfiction,” *Social Network Analysis and Mining*, vol. 14, p. 58, Mar. 2024.
- [9] A. Mattei, D. Brunato, and F. Dell’Orletta, “The Style of a Successful Story: a Computational Study on the Fanfiction Genre,” in *Proceedings of the Seventh Italian Conference on Computational Linguistics CLiC-it 2020 : Bologna, Italy, March 1-3, 2021* (F. Dell’Orletta, J. Monti, and F. Tamburini, eds.), Collana dell’Associazione Italiana di Linguistica Computazionale, pp. 284–289, Torino: Accademia University Press, 2020. Code: Proceedings of the Seventh Italian Conference on Computational Linguistics CLiC-it 2020 : Bologna, Italy, March 1-3, 2021.
- [10] S. C. Herring and J. C. Paolillo, “Gender and genre variation in weblogs,” *Journal of Sociolinguistics*, vol. 10, no. 4, pp. 439–459, 2006. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-9841.2006.00287.x>.
- [11] M. L. Newman, C. J. Groom, L. D. Handelman, and J. W. Pennebaker, “Gender Differences in Language Use: An Analysis of 14,000 Text Samples,” *Discourse Processes*, vol. 45, pp. 211–236, May 2008.
- [12] M. Eder, J. Rybicki, and M. Kestemont, “Stylometry with R : a package for computational text analysis,” 2016.
- [13] M. Sims, J. H. Park, and D. Bamman, “Literary Event Detection,” in *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, (Florence, Italy), pp. 3623–3634, Association for Computational Linguistics, 2019.


A Additional Visualizations



(a) Worm fan-fiction compared with Narnia



(b) Narnia fan-fiction compared with Worm



(c) Control fan-fiction compared with Worm

Figure 4: Median success of fan-fictions for each range of distance in style to their canons counterparts




Figure 5: Enter Caption




Figure 6: Enter Caption