The impact of emotional journeys on fanfiction popularity
A computational analysis of linear correlations between emotional behavior and popularity

Julian van der Weijden[1]

Supervisor(s): Hayley Hung[1]**, Chenxu Hao**[1]**, Ivan Kondyurin**[1]

[1]EEMCS, 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: Julian van der Weijden
Final project course: CSE3000 Research Project
Thesis committee: Hayley Hung, Chenxu Hao, Ivan Kondyurin, Elmar Eisemann

Abstract

Fanfiction writers always look for ways to make their stories more engaging. Analyzing what influences the popularity of fanfiction provides insights into readers' preferences and allows writers to tailor to these. This paper attempts to find linear correlations between fanfiction stories and the emotional journey of their characters. It does so by computationally extracting these journeys from 319 Good Omens fanfiction stories, defining and extracting several features from them and using simple linear regression to determine their correlation to fanfiction popularity. Five features were found to have a significant influence on fanfiction popularity. It was also determined that readers prefer characters that have low emotional fluctuations in their behavior.

1 Introduction

Predicting popularity has always been a topic of interest among researchers. They have been predicting the popularity of social media accounts[1], news items[2] and online videos[3], just to name a few. This is because creating a good predictor gives crucial insights into people's preferences. And knowing people's preferences allows tailoring to these preferences, in order to increase the popularity of a product. Depending on the context, this has various benefits:

1. Social media popularity can lead to better social standing and higher networking possibilities.
2. News item popularity increases the reach of news, spreading information to more people.
3. Online video popularity has the same benefits as above, and can also increase profits (if the video is monetized).

This work attempts to find features relevant to predicting fanfiction popularity, to give insights into reader's preferences and allow writers to write more engaging stories. In other words, it attempts to find a linear correlation between certain features of a fanfiction story and its popularity. It does this by computationally analyzing 319 fanfiction stories and extracting features related to one of the most important parts of any story - its characters.

This work particularly focuses on characters' emotional behavior and its change over time, as this is a crucial part of a character's identity[4]. The change over time of their emotional behavior will be referred to as a character's "emotional journey". This term is used rather than "emotional arc" as this is generally used to refer to the overarching emotional change of a text as a whole, not of individual characters.

So more specifically, this paper aims to answer the following research question:

> How do the emotional journeys of its characters influence the popularity of a fanfiction story?

The following sub-questions will aid in answering this question:

1. How do we computationally extract a character and their corresponding emotional journey from a story?
2. How do we determine the correlation between a character's emotional journey and a fanfiction's popularity?

Section 2 will talk about similar research and the gaps this paper aims to fill. Section 3 explores the corpus and answers the sub-questions as stated in the previous paragraph. It constructs a pipeline that takes a fanfiction text as input, and outputs a variety of features related to characters' emotional journeys. After this, section 4 runs the pipeline on 319 fanfictions. It attempts to find linear correlations by using simple linear regression and determines the significance of the correlations found. Section 5 attempts to explain the cause behind the results. It also places the results in a larger context, relating them back to previous works. Finally, section 6 gives a brief conclusion and summary, and talks about future works. Section 7 covers questions regarding the ethics of this paper.

2 Related work

This section explores various existing works in the domain of predicting (fan)fiction popularity based on extracted emotion, and discusses the gaps they leave that this paper attempts to fill.

Nguyen et al. [5] extracts various features from Supernatural fanfiction texts and uses these to predict fanfiction popularity. Some of these features are related to the emotions found in fanfiction, like overall emotion of the story or the average emotions of the 2 main characters. There are no emotional features related to the progression of emotion throughout the story.

Sourati Hassan Zadeh et al. [6] explores the correlation between emotional arc similarity and the popularity of fanfiction. They define emotional arc similarity as the similarity between the emotional arc of the original text and of the fanfiction text. The emotional arcs as described only take into account the happiness score of individual words, losing some nuance. It does not explore how individual characters' emotional journeys influence the popularity.

Wang et al. [7] does not explore fanfiction, but rather analyzes fiction texts. They generate emotional traces (emotional journeys) of fiction texts using 8 emotions: anger, disgust, fear, sadness, anticipation, joy, trust and surprise. They use neural networks on this data to predict popularity. It only explores the overall emotion and does not distinguish between characters.

This leaves a gap in using the emotional progression of individual characters to predict popularity. Although it's a rather specific gap to fill, characters and their emotions form an important part of stories[4]. Therefore it's an interesting and worthwhile aspect to the problem of predicting fanfiction popularity.

3 Methodology

This section explains the methods used to gather results that help answer the research question. It talks briefly about the data collection process and corpus. After this, it describes the methods for the 2 core steps involved in determining the correlation between character journeys and fanfiction popularity: Extracting the characters' emotional journeys and determining the correlation between these journeys and fanfiction popularity. It then briefly talks about how these are combined to answer the main research question. Figure 1 shows an overview of the pipeline architecture and also the order in which these components will be discussed in this section.

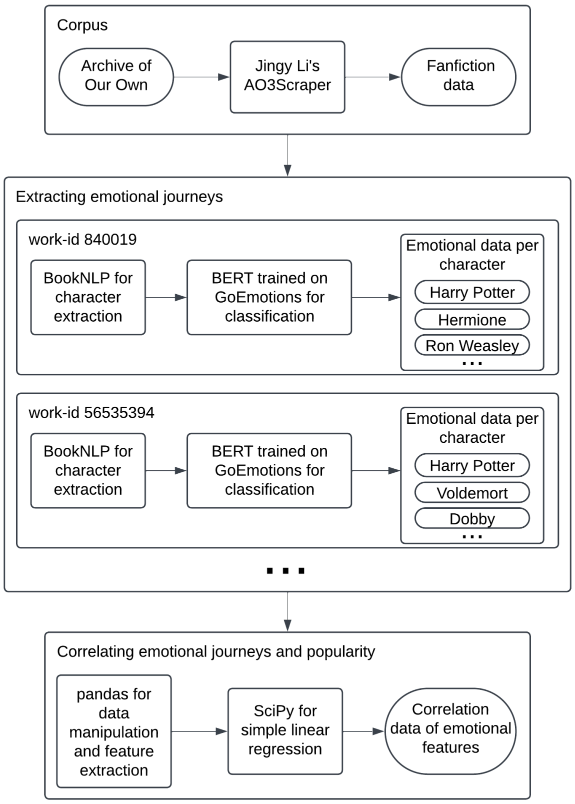

Figure 1: The pipeline architecture. "Emotional data" consists of 27 emotional journeys per character, one for each emotion (see Figure 2)

3.1 Corpus from Ao3

The full dataset consists of thousands of fan fiction works, along with their metadata, of which 319 long-form (> 45000 words) "Good Omens"[1] fanfictions are used for analysis in this paper. This dataset was collected by me and my peer group as part of the Research Project (CSE3000) course at the Delft University of Technology. It consists of ethically scraped fan fiction texts, scraped from Archive Of Our Own (Ao3)[2] - a fanfiction archive - in accord with their terms and conditions. Within this dataset, popularity is represented by the amount of "Kudos" (Ao3's version of likes) a fanfiction has received.

[1]https://en.wikipedia.org/wiki/Good_Omens

[2]Fanfictions were scraped using Jingy Li's AO3 Scraper https://github.com/radiolarian/AO3Scraper from Archive Of Our Own: https://archiveofourown.org/

3.2 Extracting a character's emotional journey

To computationally extract a character and their corresponding emotional journey from a text, the MARCUS pipeline from Bhyravajjula et al. [8] is adapted to extract these. MARCUS is a Natural Language Processing (NLP) pipeline that extracts all character actions (referred to as "events") and extracts their sentiment and related emotions. It then goes on to combine these into a single value per event (referred to as "circumstance"), and plotting these events with relation to where they appear in a text. This creates a graph that represents a character's journey.

The present work only concerns itself with the emotions of characters' actions, and therefore has no need for their sentiment. It also does not combine different emotions into a single "circumstance" value. Instead, it creates a curve for each different emotion. This is to explore the possibility that various emotions have varying performance as predictors of popularity.

Because this work analyzes character actions, all emotional data in this paper represents the amount an emotion is displayed through the actions of a character (referred to as "emotional behavior"). It does *not* represent how much a character feels an emotion. Although one could argue that these are strongly related.

This subsection explains the steps the pipeline goes through to extract emotional journeys, and the tools required to do so.

Character extraction with BookNLP

BookNLP[9] is a Natural Language Processing(NLP) pipeline[3] that scales to longer texts like fan fiction. BookNLP is used to extract all characters (that are mentioned more than 100 times). It groups the various ways they are referred to into a single entity (i.e. "Harry Potter", "Harry" and "Mr. Potter" are joined into one entity). It also associates each character as an agent or patient to every event they participate in, indicating whether they do the event (agent) or the event is being done to them (patient). In the sentence "He hit her": "He" would be the agent and "her" would be the patient. This work only considers "agent" events. The extracted events contain information about where in the text they were taken from, which are sorted to generate the progression.

Besides being a part of the MARCUS pipeline, BookNLP was chosen over the alternatives due to it's scalability, accuracy (when compared to NLTK[10]) and ease of use.

Emotional classification with BERT and GoEmotions

BERT[11] is a language representation model that can be trained to perform a variety of tasks. In this paper a specific pre-trained model[4] is used. This model is trained on the GoEmotions[12] dataset, which contains 27 emotion labels (plus a "neutral" emotion, which is ignored in this work).

This model performs multi-label emotion classification on a per-sentence basis, to extract the emotions from all sentences a character is an agent in. Table 1 shows 3 example sentences, and the emotions assigned to them by the model.

[3]https://github.com/booknlp/booknlp

[4]https://github.com/monologg/GoEmotions-pytorch

Sentence	Emotions
She hit him furiously, for his mistakes had become too many.	Anger
He said sorry, expressed his regrets and started crying.	Remorse, Sadness
She thanked him for acknowledging his misdoings, and asked him curiously why he did all those things.	Curiosity, Gratitude

Table 1: Example sentences with their emotions as assigned by a BERT model trained on the GoEmotions dataset

This model is applied per character, on all relevant events as extracted by BookNLP. Each event is part of a sentence in the text, which this model assigns emotions to. This data is treated as a time series, to which a rolling window is applied. This is done per emotion, leaving each character with an emotional progression for each of the 27 emotions (Figure 2).

Besides being a part of the MARCUS pipeline, BERT is used because at the start of writing this paper (November 2024) BERT was the state-of-the-art language model for a variety of tasks. This includes the task of classification as used in this paper. There have been some developments since then, which will be discussed in section 6. GoEmotions is used because it provides more granular information (27 emotions) than standard sentiment analysis which classifies sentences on a scale from negative to positive.

3.3 Correlating emotional journeys and popularity

To determine the correlation between a character's emotional journey and a fanfiction's popularity, various features are extracted from the emotional journeys of characters. Because fanfiction stories generally have a variable number of characters, the average of these features are used to compare fanfictions. This section describes what features are extracted from the emotional journeys and why. It also describes how these features are then used to answer the research question.

Feature engineering

Simple linear regression needs an independent variable and a dependent variable. The dependent variable will be the kudos of a fanfiction story. The independent variables need to be extracted from the emotional journeys. Right now each fanfiction has a variable number of characters with 27 emotional journeys. This subsection will explain which features are extracted from these to be used as the independent variable.

Note that all features represent emotional *behavior* - the emotions as displayed through their action. It does not represent the emotions the characters feel.

For each of the 27 emotional journeys of a character, 3 features are extracted:

1. The average of an emotion (avrge_<emotion>). This represents the average amount a specific emotion was displayed through their behavior. This feature is not related to emotional progression and only serves as a comparison, to determine if emotional progression has better predictive capabilities than average emotion. High values means this specific emotion is displayed a lot by this character.
2. The change in emotion (delta_<emotion>). This compares the average emotion of the first 35% (the rolling window size) to the last 35%. This value can be negative. A strongly positive value means this emotion has increased a lot over the story. A strongly negative value means this emotion has decreased a lot over the story. A near-zero value means this emotions remained mostly the same.
3. The standard deviation of an emotion (stdev_<emotion>). This represents the amplitude of the fluctuations of a certain emotion. In more intuitive terms, this is how variable a characters emotions are. A high value means this character has periods of very strong emotion and periods of very low emotion. Low values mean this character displays this emotion (roughly) the same amount throughout the entire story.

These are averaged across a fanfiction's characters to obtain avrge_avrge_<emotion>, avrge_delta_<emotion> and avrge_stdev_<emotion> respectively. This is done for every emotion, for a total of 81 features. "Word count" (word_count) and "Character count" (char_count) are added to this list of features to compare predictive performance. This makes a total of 83 features per fanfiction, of which 54 relate to emotional journeys.

Linear regression

Simple linear regression (using SciPy[5]) is then used with these averaged features as independent variables and a fanfiction's popularity - "Kudos" - as the dependent variable. Section 4 shows the results and significance of the correlation found by using linear regression on these features. Since this data represents the correlation between emotional journeys and fanfiction popularity, it is suitable to answer the research question.

4 Results

This section will show the results obtained by running the pipeline as described in section 3. It will show the most significant correlations from performing single linear regression on 83 features. It will also provide insights into what conclusions can be drawn from the results.

4.1 Regression lines

Plotting the regression lines in a scatter plot gives insights into the strength of the correlation, and can provide insights into the significance if there is a strong linear correlation. Having manually looked through all 83, none seem to have a strong linear correlation. Figure 3 shows this for remorse. Due to the slope being non-zero, this figure seems to imply a correlation for avrge_avrge_remorse, avrge_delta_remorse and maybe even avrge_stdev_remorse (although much smaller). However, in these figures there is

[5] https://scipy.org/

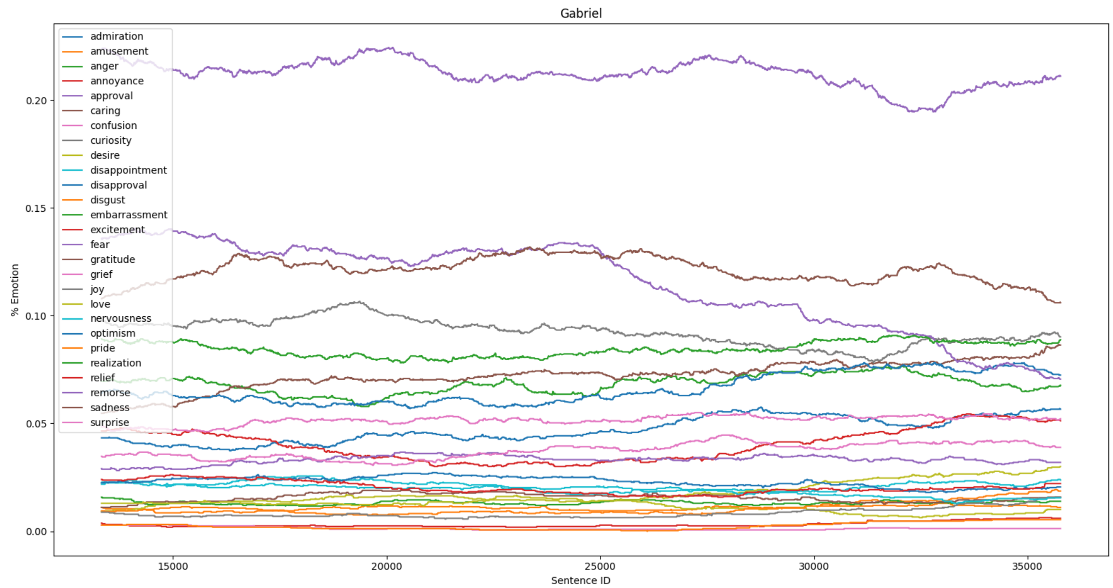

Figure 2: Example of a character's (named Gabriel) change in emotions over the course of a fanfiction text. The y axis represents what percentage of the previous 35% (the rolling window size) of the text was classified with that emotion. The x axis is the position in the text as given by their sentence number.

no way to tell if these correlations are significant or if they were produced by chance.

4.2 Significance testing

Using a null-hypothesis significance test, it is possible to determine how likely it is that the correlations observed are due to chance. SciPy conveniently does this automatically. It gives a p-value for each regression line produced. Using the null-hypothesis that the correlation is actually 0 (Slope = 0.0), this p-value represents the probability that the correlation observed was due to chance. Table 2 shows all features with a p-value lower than 0.05, as this is a common threshold in significance tests.

"char_count" - the amount of characters in a fanfiction - shows the most significant correlation. This means that there is quite certainly a correlation between character count and fanfiction popularity. However, the Pearson correlation coefficient ("r") has quite a low absolute value. This means that predicting the amount of kudos a fanfiction will receive based on character count alone is not going to lead to accurate predictions. This is because the Pearson correlation coefficient is a measure of how close data points lie to the regression line. A value of 1.0 or -1.0 would make for a perfect predictor (the sign only indicates the positive or negative nature of the correlation). Creating a scatter plot for char_count (See figure 4) clearly shows that points generally fall pretty far off the line, but even with the naked eye it is clear that there is some positive correlation between char_count and kudos.

The same goes for features related to emotional journeys in Table 2. None have a high enough r-value to be accurate predictors by themselves. However, they all influence the popularity in a statistically significant way. Interpreting these gives us the following insights (negative correlations have been verbally flipped for consistency):

feature	slope	r	p
char_count	118.328	0.183	0.001
avrge_avrge_nervousness	-40527.254	-0.154	0.006
avrge_delta_disgust	26738.014	0.148	0.008
avrge_stdev_nervousness	-61477.016	-0.147	0.009
word_count	0.005	0.145	0.010
avrge_delta_realization	7189.748	0.132	0.018
avrge_stdev_excitement	-35675.607	-0.123	0.027
avrge_delta_gratitude	14306.669	0.123	0.028
avrge_avrge_excitement	-18475.564	-0.114	0.042

Table 2: All features with a p-value lower than 0.05. "slope" is the slope of the regression line. "r" is the Pearson correlation coefficient. "p" is the p-value. Underlined features are features related to emotional journeys, the others are there for comparison.

1. Fanfictions in which characters' emotional behavior shows more "disgust" towards the end (avrge_delta_disgust) generally have a higher popularity.
2. Fanfictions in which characters show a static amount of nervousness (avrge_stdev_nervousness) generally have a higher popularity.
3. Fanfictions in which characters have more "realizations"

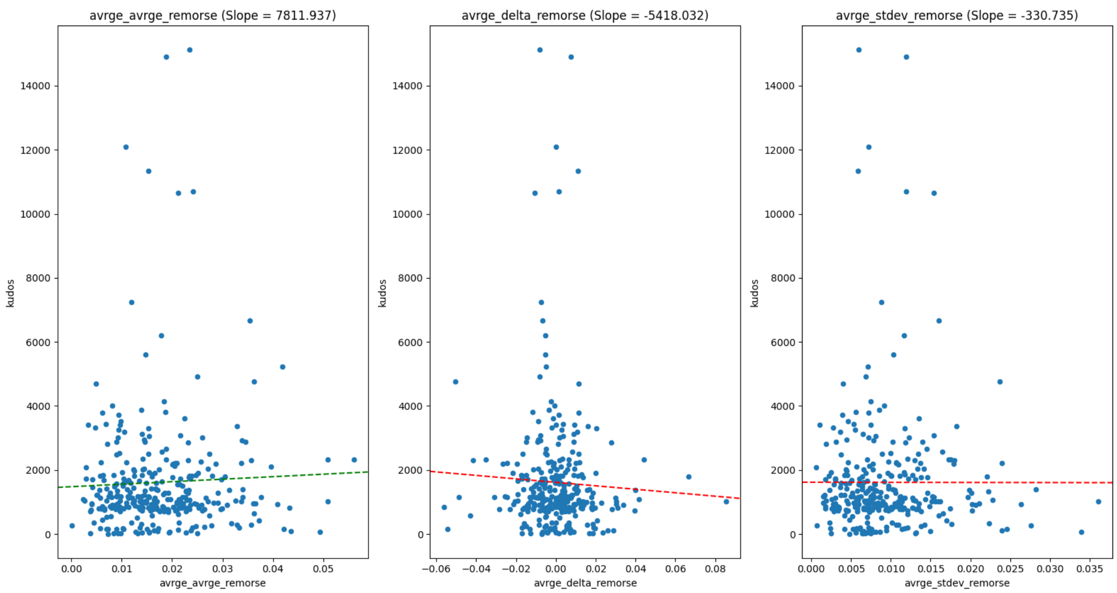

Figure 3: The 3 features related to the remorse emotion, plotted against the amount of kudos of a fanfiction. The line is the least-squares regression line. A green line indicates a positive correlation, red indicates a negative correlation. These figures have no way to tell the significance of these correlations.

towards the end (avrge_delta_realization) generally have a higher popularity.

4. Fanfictions in which characters show a static amount of excitement (avrge_stdev_excitement) generally have a higher popularity.

5. Fanfictions in which characters grow more grateful (avrge_delta_gratitude) generally have a higher popularity.

Plotting these 5 emotional journey-related features (See figure 5) allows for visual confirmation that they do in fact have some correlation.

Taking a closer look at the features related to characters' emotional fluctuation, table 2 shows negative correlations for both avrge_stdev_nervousness and avrge_stdev_excitement. This begs the question if this is by coincidence or if this holds for all emotions and readers prefer characters that behave in an emotionally static way. If we take into consideration all features related to emotional fluctuation regardless of their p-value (see table 3), we see that the 13 most significant features all have a negative correlation with regards to fanfiction popularity. Some of these features by themselves are insignificant, as they have a high probability of being caused by chance. However, these features taken together suggest readers prefer characters that behave in an emotionally static/stable manner.

5 Discussion

This work has shown a pipeline to extract features from a fanfiction that relate to the emotional journeys of its characters. It performed linear regression on these features to determine their correlation and a null-hypothesis significance test (with a null-hypothesis of 0 correlation) to determine these correlations' significance. This resulted in 5 statistically significant features related to emotional journeys. From these it could be concluded that emotional journeys do indeed influence the popularity of a fanfiction story, and showed some insights into how. It was also concluded that fanfiction readers prefer characters that behave in an emotionally stable manner.

It is possible that additional significant features exist within the 54 emotional journey-related features extracted. If this is the case then these have a rather low Pearson correlation coefficient, which would require a significantly bigger dataset to determine their significance.

Some of the results are intuitively correct; Characters that grow more grateful (avrge_delta_gratitude) and characters that realize they were wrong, realize they had feelings for eachother all along etc. (avrge_delta_realization) intuitively would have a positive influence on popularity. For a growth in behavior displaying "disgust" (avrge_delta_disgust), it is harder to intuitively understand how this contributes to popularity. It could be due to quirks within the demographic that reads Good Omens, or even fanfiction in general. It could also be a result of the GoEmotions dataset mislabeling certain - more positive - emotions.

That fanfiction readers prefer emotionally stable characters could be explained by looking at the nature of fanfiction; Fanfiction stories are derivative works that take gaps in an original work, and fill them in [13]. Because characters are such an important part of stories, the emotions of these characters are usually well-explored by the original text. This leaves

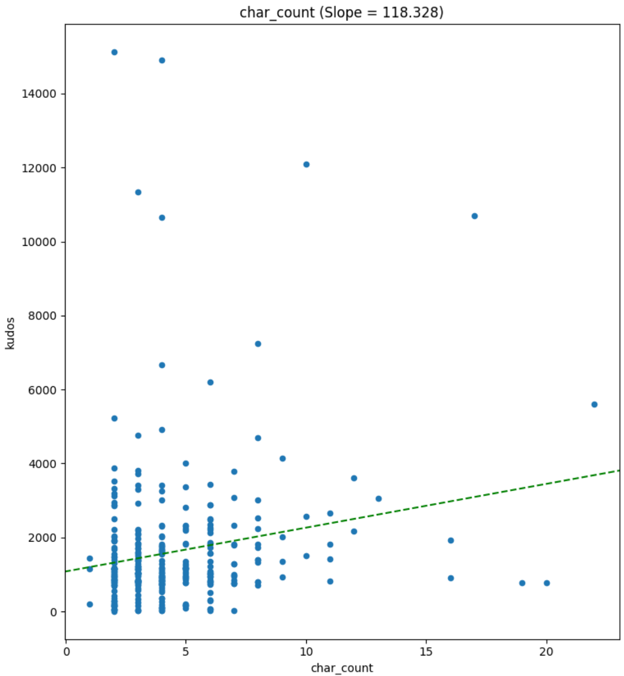

Figure 4: The regression line of character counts, plotted against the kudos of a fanfiction.

few gaps fanfiction writers can try to fill. It is easier to put existing characters in new situations than it is to continue a well-developed character's emotional journey. This is because there is a near-infinite amount of situations characters can find themselves in, and plenty of gaps in the continuity of stories for writers to fill. This makes characters with strongly fluctuating emotions rare within fanfiction, and therefore attracts readers that just want "more of the same".

None of the features found in this paper perfectly predict fanfiction popularity. This was never expected since human preferences are inherently variable. No two people will have the exact same preferences, and even a single person's preferences will vary over time. It is therefore not a surprise that fanfiction popularity doesn't have a perfect linear correlation with regards to the emotional journeys of their characters. Taking human preference as averaged over a group *can* provide some valuable insights. Since this paper analyzes a large group of fanfictions with an even larger group of people rating them, this also applies to the insights as brought forth by this paper.

Previous works have made various predictive models and determined their accuracy. Because this paper does not create any predictive models but rather highlights features that could be used in these models, it is hard to compare this work to previous works. However, considering that 5 out of 9 of the statistically significant features relate to characters' emotional journeys - rather than average emotion - makes it safe to assume that these features could be worthwhile additions to predictive models for fanfiction popularity.

avrge_stdev_<emotion>	coef	r	p
nervousness	-61477.016	-0.147	0.009
excitement	-35675.607	-0.123	0.027
embarrassment	-38837.403	-0.091	0.103
joy	-22414.176	-0.087	0.122
caring	-16493.158	-0.079	0.159
fear	-18374.771	-0.075	0.179
sadness	-14980.688	-0.069	0.219
admiration	-11351.231	-0.060	0.284
realization	-9747.693	-0.051	0.364
pride	-31780.538	-0.049	0.385
anger	-8097.312	-0.043	0.442
amusement	-7447.921	-0.037	0.510
disappointment	-10433.427	-0.033	0.560
grief	54206.033	0.024	0.667
desire	-9065.308	-0.024	0.675
gratitude	7461.003	0.023	0.680
love	-4493.530	-0.022	0.700
disgust	7133.385	0.019	0.736
approval	2511.794	0.017	0.758
relief	-7738.020	-0.015	0.792
annoyance	1452.468	0.008	0.885
curiosity	519.385	0.004	0.949
confusion	715.109	0.003	0.953
surprise	-650.525	-0.003	0.961
optimism	629.766	0.002	0.967
disapproval	-464.576	-0.002	0.971
remorse	-330.735	-0.001	0.986

Table 3: All features representing the fluctuation of character emotion. slope" is the slope of the regression line. "r" is the Pearson correlation coeffi- cient. "p" is the p-value.

6 Conclusions and Future Work

The question this paper set out to answer was:

> How do the emotional journeys of its characters influence the popularity of a fanfiction story?

Five significant features (p-value < 0.05) related to characters' emotional journeys were found to influence popularity. From this it can be concluded that emotional journeys do play a significant role in determining a fanfictions popularity. The five features that seem to influence popularity the most (in order) are:

1. The growth of character "disgust" during a fanfiction story (p = 0.008, r = 0.148)
2. The fluctuation of character "nervousness" during a fanfiction story (p = 0.009, r = -0.147)
3. The growth of character "realization" during a fanfiction story (p = 0.018, r = 0.132)
4. The fluctuation of character "excitement" during a fanfiction story (p = 0.027, r = -0.123)
5. The growth of character "gratitude" during a fanfiction story (p = 0.028, r = 0.123)

It was also shown that fanfiction readers prefer characters that behave in an emotionally stable manner. This was done by

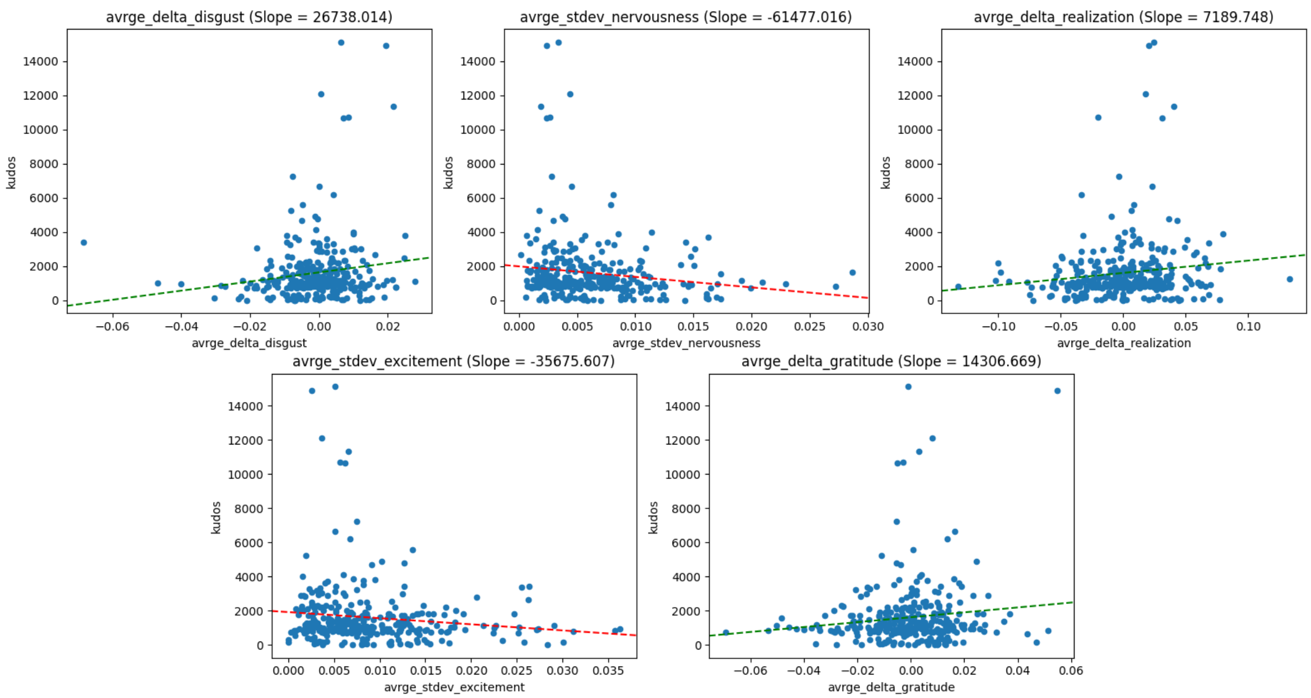

Figure 5: The regression lines of the significant features related to emotional journeys, plotted against the kudos of a fanfiction.

observing that the 13 most significant features related to emotional fluctuation had a negative impact on popularity.

The GoEmotions[12] dataset used in this paper is a hand-labeled dataset constructed from Reddit comments. The task of the raters was to "identify the emotions expressed by the writer of the text" [12, p. 4043]. Because of this it might not be suited to identify the emotions displayed by the characters in the text. In some manual testing, it seemed to perform adequately at this task, but it is likely the GoEmotions dataset sometimes picks up on the emotions of the writer of the text, rather than the characters depicted in it. Using (or creating) a dataset specifically tailored to the emotional behavior of characters in stories could be worthwhile direction for future work.

B. Warner, A. Chaffin, B. Clavié, et al.[14] is a recent paper that introduced ModernBERT. This renewed version of BERT claims to be better than BERT in (almost) every way. Exploring how ModernBERT can improve the results in this paper could therefore prove to be a promising line of research.

Other future work can explore the non-linear correlation of the features extracted by this paper, construct larger predictive models based on these features, explore how differences between characters within a story influences popularity, explore alternative features related to emotional journeys (for example with the help of tools such as tsfresh[6]) and extend the corpus to include fanfictions based on other original works, fanfictions in other languages and shorter fanfictions.

[6]https://tsfresh.com/

7 Responsible Research

This section discusses various ethical aspects of the work performed in this paper.

7.1 Data collection

The data collected is from AO3 - an online fanfiction archive. They forbid scraping for commercial purposes[7], but don't prohibit it for research purposes. In regards to the privacy of the authors, all scraped data is publicly available information. The data scraped is not published alongside this paper, and any derived results in this paper do not mention any of the authors' information.

7.2 Consequences of false conclusions

If the conclusions drawn by this paper turn out to be false - whether by chance or by error - the consequences of this should be relatively tame. A writer might base some decisions on the insights provided by this paper, which could lead to disappointment if the popularity does not turn out as high as they expect. However, due to the erroneous nature of predicting human preference - which has been extensively displayed in this paper - it is impossible to guarantee popularity, no matter if the conclusions drawn in this paper are true or false.

7.3 Biases of NLP techniques and datasets

It is important for anyone reading this paper to be aware of the biases of the tools and datasets used, as these can influ-

[7]https://archiveofourown.org/tos

ence the results. This includes the BookNLP pipeline[8], the GoEmotions dataset[12] and the language model BERT[11] (and its gender biases [15][16]). It is hard to determine the exact effects of these biases on the results of this paper due to the sequential nature of the pipeline used to extract these results.

7.4 Reproducibility

This paper should serve as a sufficient basis for reproducing the same results and conclusions, provided the dataset is the same. However, since the time of scraping the data (27-11-2024), the exact fanfictions available on Archive Of Our Own might have changed. This could result in slight deviations from the exact results as displayed in this paper. The conclusions should not change much as only results with a low probability of being due to chance (low p-value) are used to draw these conclusions.

[8]https://github.com/booknlp/booknlp

References

[1] M. Meghawat, S. Yadav, D. Mahata, Y. Yin, R. R. Shah, and R. Zimmermann, "A multimodal approach to predict social media popularity," in *2018 IEEE Conference on Multimedia Information Processing and Retrieval (MIPR)*, Apr. 2018, pp. 190–195. DOI: 10.1109/MIPR.2018.00042. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/8396998 (visited on 01/20/2025).

[2] K. Lerman and T. Hogg, "Using a model of social dynamics to predict popularity of news," in *Proceedings of the 19th international conference on World wide web*, ser. WWW '10, New York, NY, USA: Association for Computing Machinery, Apr. 26, 2010, pp. 621–630, ISBN: 978-1-60558-799-8. DOI: 10.1145/1772690.1772754. [Online]. Available: https://dl.acm.org/doi/10.1145/1772690.1772754 (visited on 01/20/2025).

[3] T. Trzciński and P. Rokita, "Predicting popularity of online videos using support vector regression," *IEEE Transactions on Multimedia*, vol. 19, no. 11, pp. 2561–2570, Nov. 2017, Conference Name: IEEE Transactions on Multimedia, ISSN: 1941-0077. DOI: 10.1109/TMM.2017.2695439. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/7903630 (visited on 01/20/2025).

[4] W.-p. Su, B. Pham, and A. Wardhani, "Personality and emotion-based high-level control of affective story characters," *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, no. 2, pp. 281–293, Mar. 2007, Conference Name: IEEE Transactions on Visualization and Computer Graphics, ISSN: 1941-0506. DOI: 10.1109/TVCG.2007.44. [Online]. Available: https://ieeexplore.ieee.org/abstract/document/4069237 (visited on 01/18/2025).

[5] 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, no. 1, pp. 1–17, Dec. 1, 2024, Company: Springer Distributor: Springer Institution: Springer Label: Springer Number: 1 Publisher: Springer Vienna, ISSN: 1869-5469. DOI: 10.1007/s13278-024-01224-x. [Online]. Available: https://link.springer.com/article/10.1007/s13278-024-01224-x (visited on 01/22/2025).

[6] Z. Sourati Hassan Zadeh, N. Sabri, H. Chamani, and B. Bahrak, "Quantitative analysis of fanfictions' popularity," *Social Network Analysis and Mining*, vol. 12, no. 1, pp. 1–11, Dec. 1, 2022, Company: Springer Distributor: Springer Institution: Springer Label: Springer Number: 1 Publisher: Springer Vienna, ISSN: 1869-5469. DOI: 10.1007/s13278-021-00854-9. [Online]. Available: https://link.springer.com/article/10.1007/s13278-021-00854-9 (visited on 01/22/2025).

[7] X. Wang, S. Zhang, and I. Smetannikov, "Fiction popularity prediction based on emotion analysis," in *2020 International Conference on Control, Robotics and Intelligent System*, Xiamen China: ACM, Oct. 27, 2020,

pp. 169–175, ISBN: 978-1-4503-8805-4. DOI: 10.1145/3437802.3437831. [Online]. Available: https://dl.acm.org/doi/10.1145/3437802.3437831 (visited on 01/23/2025).

[8] S. Bhyravajjula, U. Narayan, and M. Shrivastava, "MARCUS: An event-centric NLP pipeline that generates character arcs from narratives," presented at the Text2Story@ECIR, 2022. [Online]. Available: https://www.semanticscholar.org/paper/MARCUS%3A-An-Event-Centric-NLP-Pipeline-that-Arcs-Bhyravajjula-Narayan/cea59b3ecc85b78a4bb7d1a5321d5fa2d1bda5a6 (visited on 11/21/2024).

[9] D. Bamman, T. Underwood, and N. A. Smith, "A bayesian mixed effects model of literary character," in *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, K. Toutanova and H. Wu, Eds., Baltimore, Maryland: Association for Computational Linguistics, Jun. 2014, pp. 370–379. DOI: 10.3115/v1/P14-1035. [Online]. Available: https://aclanthology.org/P14-1035 (visited on 12/13/2024).

[10] S. Bird, E. Loper, and E. Klein, "Natural language processing with python," *O'Reilly Media Inc.*, 2009.

[11] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, *BERT: Pre-training of deep bidirectional transformers for language understanding*, version: 2, May 24, 2019. DOI: 10.48550/arXiv.1810.04805. arXiv: 1810.04805[cs]. [Online]. Available: http://arxiv.org/abs/1810.04805 (visited on 12/04/2024).

[12] D. Demszky, D. Movshovitz-Attias, J. Ko, A. Cowen, G. Nemade, and S. Ravi, "GoEmotions: A dataset of fine-grained emotions," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, D. Jurafsky, J. Chai, N. Schluter, and J. Tetreault, Eds., Online: Association for Computational Linguistics, Jul. 2020, pp. 4040–4054. DOI: 10.18653/v1/2020.acl-main.372. [Online]. Available: https://aclanthology.org/2020.acl-main.372 (visited on 12/02/2024).

[13] V. Cuntz-Leng, "Queering Harry, slashing Potter: Between latent meanings and resistant readings," *GENDER – Zeitschrift für Geschlecht, Kultur und Gesellschaft*, vol. 9, no. 2, Dec. 6, 2017, Number: 2, ISSN: 2196-4467. [Online]. Available: https://www.budrich-journals.de/index.php/gender/article/view/28018 (visited on 11/21/2024).

[14] B. Warner, A. Chaffin, B. Clavié, *et al.*, *Smarter, better, faster, longer: A modern bidirectional encoder for fast, memory efficient, and long context finetuning and inference*, Dec. 19, 2024. DOI: 10.48550/arXiv.2412.13663. arXiv: 2412.13663[cs]. [Online]. Available: http://arxiv.org/abs/2412.13663 (visited on 01/04/2025).

[15] M. Bartl, M. Nissim, and A. Gatt, "Unmasking contextual stereotypes: Measuring and mitigating BERT's gender bias," 2020.

[16] S. Jentzsch and C. Turan, "Gender bias in BERT - measuring and analysing biases through sentiment rating in a realistic downstream classification task," in *Proceedings of the 4th Workshop on Gender Bias in Natural Language Processing (GeBNLP)*, C. Hardmeier, C. Basta, M. R. Costa-jussà, G. Stanovsky, and H. Gonen, Eds., Seattle, Washington: Association for Computational Linguistics, Jul. 2022, pp. 184–199. DOI: 10.18653/v1/2022.gebnlp-1.20. [Online]. Available: https://aclanthology.org/2022.gebnlp-1.20/ (visited on 01/15/2025).