Good Omens: A Collaborative Authorship Study

Leonardo Grotti, Mona Allaert and Patrick Quick

Universiteit Antwerpen, Faculty of Arts, Prinsstraat 13, B-2000, Antwerp

Abstract

Good Omens is a collaborative novel written by Terry Pratchett and Neil Gaiman. Rising interest in the book, amplified by the success of the recent screen adaptation, has aroused curiosity regarding its realization. We use Rolling Delta and Rolling Classify to detect stylistic signals from each author as these methods reveal authorial takeovers. The same techniques are applied to compare the screenplay of the show to the novel. The results indicate that *Good Omens* resembles Pratchett's work more closely. The screenplay is correctly attributed to Gaiman, its sole author, and the comparison reveals that Gaiman may have relied less on the source material over the course of the narrative arc.

Keywords

Good Omens, Rolling Stylometry, PCA, collaborative authorship

1. Introduction

In 1983, Terry Pratchett published *The Colour of Magic*, the first book in his forty-one-book *Discworld* series. Although Pratchett is now recognized as one of the most popular fantasy writers of the past two decades [9], during the early 1980s he was far from the level of success he would come to enjoy. As noted by Shanahan [26], Pratchett was still working as a newspaper journalist and would not become a full-time writer until 1987. In 1985, in the early stages of his writing career, Pratchett granted his first interview as an author to *Space Voyager Magazine* to promote his series [14].

It was on this occasion that Pratchett met Neil Gaiman, who at the time was working for *Space Voyager Magazine* and conducted the interview. The two stayed in touch due to 'a shared delight and amazement at the sheer strangeness of the universe, in stories, in obscure details, in strange old books in unregarded bookshops' [14, p. 488]. Five years later, they co-wrote *Good Omens: The Nice and Accurate Prophecies of Agnes Nutter, Witch*, which became an international bestseller.

Unlike other notable literary collaborations (e.g., Conrad and Ford, [25], 17th century French playwrights, [7]) that between Pratchett and Gaiman was rather unproblematic. A cause for their successful partnership can be found in their similar backgrounds: both authors operated

CHR 2022: Computational Humanities Research Conference, December 12 – 14, 2022, Antwerp, Belgium
leonardo.grotti@student.uantwerpen.be (L. Grotti); mona.allaert@uantwerpen.be (M. Allaert); patrick.quick@student.uantwerpen.be (P. Quick)
https://github.com/corvusMidnight (L. Grotti); https://github.com/MonaDT (M. Allaert); https://github.com/patrickquick (P. Quick)
0000-0001-7914-3191 (L. Grotti)

in the field of fantasy and science fiction. They also professed a love for comedy and claimed that the main objective of writing *Good Omens* was 'to make the other one laugh' [14, p. 484].

The collaboration seemed to work well in that they were both equally invested in the writing process:

> We wrote the first draft in about nine weeks. Nine weeks of gloriously long phone calls, in which we would read each other what we'd written, and try to make the other one laugh. We'd plot, delightedly, and then hurry off the phone, determined to get to the next good bit before the other one could. We'd rewrite each other, footnote each other's pages, and sometimes even footnote each other's footnotes. [2]

Even though both Pratchett and Gaiman remained playfully evasive about attributing specific aspects to one another, it is clear that Gaiman initiated the project. He wrote 5000 words in which he created one of the main characters, Crowley, and wrote a passage regarding a baby swap, which would come to be the premise of *Good Omens*. The draft was then sent to Pratchett for feedback, who suggested writing it together as a novel. In the beginning, they wrote separately: Pratchett during the day, Gaiman during the night, with a short overlap in the afternoon to compare notes. However, towards the end of the writing process Gaiman moved into Pratchett's spare room to polish the final parts before publication [2].

Both authors kept looking back with fondness on their project and remained in touch for potential cinematic adaptations of the novel. In 2008, however, Pratchett was diagnosed with early-onset Alzheimer's and passed away in 2015. At Pratchett's request, Gaiman took it upon himself to write the screenplay for a six-part television show, which was released in 2019 and is currently awaiting its second season.

When *Good Omens* was written, both authors were only at the start of their careers: '[i]n those days Neil Gaiman was barely Neil Gaiman and Terry Pratchett was only just Terry Pratchett' [14, p. 475]. This changed for Pratchett when the *Discworld* series achieved great renown. Gaiman gained popularity in the United States, predominantly due to his work as a graphic novelist. As the reputation of the authors flourished, so did that of *Good Omens*, which earned the status of 'cult classic' [14, p. 478]. Consequently, interest in the creative process of *Good Omens* bloomed. Even though being cryptic about the writing process of the novel was a deliberate choice [2], a stylometric approach to *Good Omens* could shine a light on the question of who wrote what.

2. Material

Good Omens takes an unlikely setting for a comedy, namely, the end of the world. At the center of the story are two main characters: the angel Aziraphale and the demon Crowley (who was loosely based on Gaiman [11]), who operate as agents of heaven and hell on earth. The story follows their friendship that spans thousands of years and their attempts to prevent the apocalypse, which will come as a result of the birth of Satan's son, Adam. The characters' inability to distinguish the motives of heaven and hell in bringing about, or, rather, ensuring the end of the world, functions as the framework for the novel's comedy.

The novel is composed of six chapters plus an appendix. The latter[1] contains a fictional interview regarding the collaboration between the two authors. Here, Pratchett and Gaiman make some comments regarding who wrote what. They report that most of the scenes between Adam and Anathema (witch and owner of the fictional book of prophecies *The Nice and Accurate Prophecies of Agnes Nutter*) were written by Pratchett, and the passages with the Four Horsemen of the Apocalypse[2] were of Gaiman's hand. Moreover, they claim that Gaiman was the dominant author at the opening of the novel, whereas Pratchett had more control towards the end [14, p. 477].

Scholars have not concerned themselves with the stylometric analysis of *Good Omens*. One exception is Callaway [8]. In her blog post, Callaway uses the function Rolling Classify (using 50 features and 5000 word-long segments), in combination with a linear Support Vector Machine (SVM), to detect authorial takeovers in *Good Omens*. The resulting graph highlighted that Gaiman is indeed more present at the beginning of the book[3] and in sections that include the Four Horsemen. In Callaway [8]'s analysis, however, Pratchett is still the predominant author and she concluded that 'even in areas where one of the two author's signal dominates, the other author is present. Both Gaiman and Pratchett are detectable all over their shared work' [8].

Regarding the appendix, it is important to consider to what extent these attributions may be constructed. As previously mentioned, the two often reworked each other's material, with Pratchett professing that 'both can write passably in the other one's style' [14, p. 477]. The appendix, although written as an interview to a third unnamed person, was authored by both Pratchett and Gaiman. It is also worth mentioning that the appendix is effectively part of the novel. For that reason, it may be challenging to try and attribute certain passages to a specific author or verify their claims. Despite its limited length and scope, the appendix of *Good Omens* remains the most reliable source. Interviews found online often echo the same information or refer directly to it.

On the other hand, Callaway's [8] study does not reference any source material and does not relate its results to any of the author's claims. Additionally, the results presented there are not replicable or comparable since (i) the reference corpus and the 50 features used to run Rolling Classifier are not provided and (ii) chapters markers have not been added to the graph. Thus, our expectations are based on the content of the appendix and we only partially compare our results to Callaway's [8].

First, we anticipate that Pratchett will be the most dominant author of *Good Omens* since he took upon himself the role of editor (see the interview in [5]). Also, we expect Gaiman's style to emerge in sections of the novel involving characters (e.g. Four Horsemen of the Apocalypse) attributed to him. Finally, we forecast Pratchett's idiom to be predominant in the later sections of the book and Gaiman's in the earlier ones.

Regarding the screenplay, not many hypotheses can be formulated within the scope of this study. Firstly, a screenplay belongs to a different text type and cannot be reliably compared to novels. As such, its validity as a control text is limited. Secondly, neither author publicly

[1] Perhaps fittingly titled *Good Omens, The Facts (or, at least, lies that have been hallowed by time).*
[2] For reference, War, Famine, Pollution, and Death.
[3] Callaway [8] enriches the graph with short notes summarizing the content of each segment in the novel.

Table 1
Reference corpus for Pratchett and Gaiman with year of publication and n° of word tokens. Titles in bold are part of the Rolling Stylometry sub-corpus

Title	Author	Year of Publication	Tokens
Small Gods	Terry Pratchett	1982	93 042
The Colour of Magic	Terry Pratchett	1983	66 073
Light Fantastic	Terry Pratchett	1986	65 389
Equal Rites	Terry Pratchett	1987	67 691
Mort	Terry Pratchett	1987	74 089
Sorcery	Terry Pratchett	1988	79 805
Wyrd Sisters	Terry Pratchett	1988	86 428
Pyramids	Terry Pratchett	1989	88 422
Moving Pictures	Terry Pratchett	1990	98 887
Lords and Ladies	Terry Pratchett	1992	90 240
Neverwhere	Neil Gaiman	1996	100 714
Hogfather	Terry Pratchett	1996	96 104
Stardust	Neil Gaiman	1997	60 552
Jingo	Terry Pratchett	1997	107 876
American Gods	Neil Gaiman	2001	184 086
Thief	Terry Pratchett	2001	103 881
Coraline	Neil Gaiman	2002	31 504
Anansi Boys	Neil Gaiman	2005	111 164
Thud!	Terry Pratchett	2005	113 531
Fragile Things	Neil Gaiman	2007	108 356
M is for Magic	Neil Gaiman	2007	47 773
The Graveyard Book	Neil Gaiman	2008	69 440
Unseen Academicals	Terry Pratchett	2009	137 861
The Ocean at the end of the Lane	Neil Gaiman	2013	56 287
Trigger Warning	Neil Gaiman	2015	101 004
Raising Steam	Terry Pratchett	2015	126 088
Total	/	/	2 366 257

commented on the novel's potential for screen adaptation, and Gaiman never shared his process in writing it. We still expect Gaiman's style to be overwhelmingly predominant since he was the sole author. Thirdly, because the show is faithful to the narrative arc of the novel, we anticipate the screenplay to take liberally from the book.

3. Methodology

As a preliminary step, we first compiled a comprehensive corpus, consisting of ten novels for Gaiman and sixteen for Pratchett. Table 1 above summarizes the structure of our corpus.

The corpus consists of 2,366,257 word tokens (see Table 1). It must be noted that Gaiman is slightly underrepresented in the dataset since he wrote fewer novels, and we did not include any of his non-fiction work.[4] The novels in bold are part of the sub-corpus selected to run Rolling

[4]Pratchett's texts make up 1,495,407 word tokens, whereas Gaiman's texts consist of 870,850 word tokens.

Stylometry functions. The novel *Good Omens* consists of 110,935 tokens while the screenplay of the television adaptation consists of 86,425 tokens. The entire screenplay, rather than just the dialogue, was considered since it includes detailed character and scene descriptions.

Modern stylometry studies often do not limit themselves to the use of a single technique [27]. Rather, scholars (e.g. see [25], [20]) have shown how the implementation of different methodologies yields better and more reliable results. Thus, to better assess the stability of our results, the present paper proposes a combination of three different methods: Principal Components Analysis (PCA), Rolling Delta, and Rolling Classify [12].

Before proceeding further, it is worth noting that here and throughout we calculate stylistic distance using 'Burrows's Delta' [6]. Burrow's Delta is a metric which combines z-transformation (i.e., standardization) of frequency with Manhattan distance [13]. Roughly, to calculate delta given the x Most Frequent Words (MFW)[5] in n texts, we first compute the relative frequency of each word in each document. By doing so, we obtain a x-scores-long representation of each document. The δ (standard deviation) of each term's frequency across the whole corpus is then calculated. The distance between two documents *n1* and *n2* is expressed as the absolute difference between each individual word's relative frequency in *n1* and *n2* divided by the same word's δ across the corpus. Finally, the resulting deltas are collected in a distance table which is used as the basis for the cluster analysis.[6]

PCA is an unsupervised dimensionality reduction technique, i.e. a method that does not require ground-truth labels for the data. As such, it is often considered ideal for exploratory purposes [18]: instead of being driven to a specific solution by the researcher, PCA results are data-driven [27]. The documents are first vectorized into a 67×117 matrix (sixty-seven segments by two authors and 117 MFW). Then, we normalize the resulting matrix (following L1 norm, see [18]) and scale it. PCA operates by dimensional reduction: 'it transforms to new set of variables, the principal components (PCs), which are uncorrelated, and which are ordered so that the first few retain most of the variation present in all of the original variables' ([4, p. 447], originally in [17, p. 1]). In our case, PCA is ideal compared to other techniques[7] since (i) it offers more reliable results for smaller sets of authors and (ii) allows us to visualize the stylistic characteristics from which it was built [27].

Rolling Delta and Rolling Classify are both part of 'Rolling Stylometry' [12]. Rolling Stylometry is a sequential classification technique in that it operates by means of a rolling window across fixed segments of text. In other words, both functions split the text into overlapping, same-length fragments [25] and roll over it. For instance, if we take a 'window size' of 5000 words and a 'step size'[8] of 1000, the first analyzed segment will cover the range 1–-5000, the second 1001–-6001, etc. Both functions allow a maximum of twelve texts in the reference corpus (e.g. in our case, six texts for each of the two authors) and a separate test set (usually one

[5]We here talk about words, but it worth noting that Burrows Delta works not only on words but also on most frequent items (e.g. n-grams).

[6]The above explanations echoes those found in Stover and Kestemont [27] and Karsdorp, Kestemont, and Riddell [18]. For a more technical, in-depth explanation of Burrow's Delta, see Burrows [6].

[7]Such as Agglomerative Clustering Analysis, see Müllner [22].

[8]Note that these are the parameters for Rolling Delta and not Rolling Classify, which specifies 'slice size' and 'slice overlap'. Although the name and configuration of the two parameters are different, they can be considered equivalents to those of Rolling Delta. E.g., for a 'slice size' of 5000 with a 'slice overlap' of 4000, the first analyzed segment will cover the range 1–5000, the second 1001–-6001, and so on [12].

text, *Good Omens* in this case).

The difference between Rolling Delta and Classify lies in the way the segments are analyzed: Rolling Delta calculates the Burrow's Delta distances of each segment in the test sets from the segments of texts from the reference corpus. Rolling Classify, on the other hand, uses the texts in the reference corpus to train a classifier and then classifies the text segments from the test set. Also, while Rolling Classify allows the user to select a custom set of most frequent features, Rolling Delta does not: rather, it automatically selects an X number of most frequent features.

A preliminary analysis revealed that the upper tail (i.e., 250 MFW) of the extracted features contained many author-related lemmas, such as *Rincewind* (one of the main characters of the *Discworld* series) and *black* and *white* (both common collocations of the word *magic*, strongly related to the fantasy genre). Following Binongo [3], using only function words yields undeniable advantages: (i) because of their scarce semantic content, they are less context-dependent compared to content words, (ii) since they are not inflected, they are often found in only one form, and (iii) their usage is often unaffected by a writer's stylistic choices. As such, we removed content words and culled personal pronouns, often considered too indicative of a specific genre or narrative style. After the process, the final MFW list, which has been used to conduct further analyses and experiments, consisted of 117 function words.[9]

We use PCA for two reasons: first, as a tool to narrow down our corpus. As noted above, both Rolling Stylometry functions work using a restricted corpus of twelve texts. As such, we want our novel selection to be as representative and comprehensive of each author's style as possible. PCA allows us to identify stylistic clusters[10] and to make an informed decision while also identifying distinctive features for both Pratchett and Gaiman. Second, PCA is useful in giving an explorative representation of where *Good Omens* lies stylistically.

However, PCA has a major drawback: when it comes to collaborative authorship, static visualizations may be misleading. i.e., a book may be attributed in its entirety to one author through clustering or classification. Thus, PCA cannot help scholars in assessing whether the other author had an influence on the writing process and to what extent. Rolling Stylometry enables the model to identify authorial dominance and takeovers throughout the text rather than attributing an entire text to a specific author [12]. As such, it is especially fitting for collaborative authorship attribution [21] and was selected to give an in-depth insight into *Good Omens*' style. Both Rolling Delta and Rolling Stylometry are implemented and can be accessed in the environment for statistical computing R [23].

To further improve the quality of our results obtained through Rolling Delta and Classify, we also performed a preliminary analysis of the twelve-text sub-corpus using SVM. SVM is specifically fit for text categorization due to its inductive bias and the linearly separable nature of the task [16]. Using SVM in combination with Terms Frequency-Inverse Dictionary Frequency (TFIDF) vectorizer, we were able to test how different parameters (i.e. MFW and text segment size) affected the model's ability to correctly distinguish between the two authors. SVM was also compared with other models (Logistic Regression and KneighborsClassifier), which

[9]Here we apply a broader definition of function words; i.e., all those words that belong to a closed class [10]. For instance, we do not remove auxiliary modal verbs. For an extensive discussion of the role of function words and their uses in authorship attribution, see [19].

[10]Here and throughout the paper we use the term cluster to refer to the visual clusters that can be observed in the PCA visualization.

showed that SVM, in combination with a MFW $\geq$ 250 and segment size $\geq$ 1000, reached an accuracy of 1.0.[11]

4. Results

Fig. 1 shows a PCA of ten novels by Gaiman (blue) and sixteen by Pratchett (black), plus the novel *Good Omens* (red). The plot was obtained by slicing the texts into 30,703 word segments (the length of *Coraline*, by Gaiman, the shortest novel in the corpus, following [18]). The plot represents how the works of the two authors cluster together, i.e., how different (or similar) they are from one another. Interestingly, PCA places the three segments of *Good Omens* in the middle of the plot. Such placement is partially misleading: although in the middle, *Good Omens* clusters closer to Pratchett's novels than Gaiman's. This is also confirmed by predicting the segments' author using Burrow's Delta, which attributes all three of them to Pratchett.

With Pratchett's novels, we can clearly observe a distinct group of works within the bottom-left quadrant. These include some of Pratchett's late-80s to mid-90s works, such as *Mort*, *Pyramids*, *Hogfather*, *Wyrd Sisters*, and *Equal Rites*. In contrast, his novels from the early 2000s cluster together across the bottom and top-left quadrants. It is also worth noting that all novels written after 2007 (*Unseen* and *Raising*), the year in which Pratchett was diagnosed with Alzheimer's, are clustered far away from the rest of the novels in the top-left quadrant. Furthermore, *The Colour of Magic*, Pratchett's first breakthrough novel, gets clustered together with Gaiman's early novels, which is not surprising, considering that Gaiman admittedly read *The Colour of Magic* and has said on many occasions to have been influenced by it [15].

Following these observations, we select one novel closer in style to *Good Omens* (*Pyramids*), and two novels distinctive of that same first cluster, *Wyrd Sisters* and *Moving Pictures*, for our corpus. We actively include *The Colour of Magic*, as it is not only an interesting case because of its relation to Gaiman, but also because it represents Pratchett's early style. From the second cluster, we pick novels written between the late 90s (*Jingo*) and early 2000s (*Thud!*). We exclude novels that were written after 2007 as they represent Pratchett's style after his disease diagnosis and are clustered together far from the rest.

For Gaiman, although the temporal span of publication is not as wide as Pratchett's, we follow the same procedure. We pick one of the first novels, which was said to be influenced by Pratchett's style (*Neverwhere*), two from the early-mid (*American Gods* and *Anansi Boys*) and mid-late 2000s (*M is for Magic* and *Graveyard Book*) periods, and one of his latest novels (*Trigger Warning*). For Gaiman, *Anansi Boys* represents the closest novel to *Good Omens* (stylistically), while the four novels mentioned earlier come from the central, densest cluster (between the top and bottom-right quadrants) of Gaiman's style. As a rule of thumb, our selection tries to be both representative of novel clusters in the PCA, time of publication, and stylistic proximity to *Good Omens*.

Figure 2 represents the diagram outputted using Rolling Delta on the twelve novels (*Jingo*, 1997, *Pyramids*, 1989, *Colour of Magic*, 1983, *Wyrd Sisters*, 1988, *Moving Pictures*, 1990, *Thud*, 2005, for Pratchett; and *M is for Magic*, 2007, *Anansi Boys*, 2005, *American Gods*, 2001, *The Graveyard Book*, 2008, *Neverwhere*, 1996, *Trigger Warning*, 2015, for Gaiman) selected through the

[11] See Table 2 and Table 3 in Appendix A.2 for a summary of the SVM set-up results.

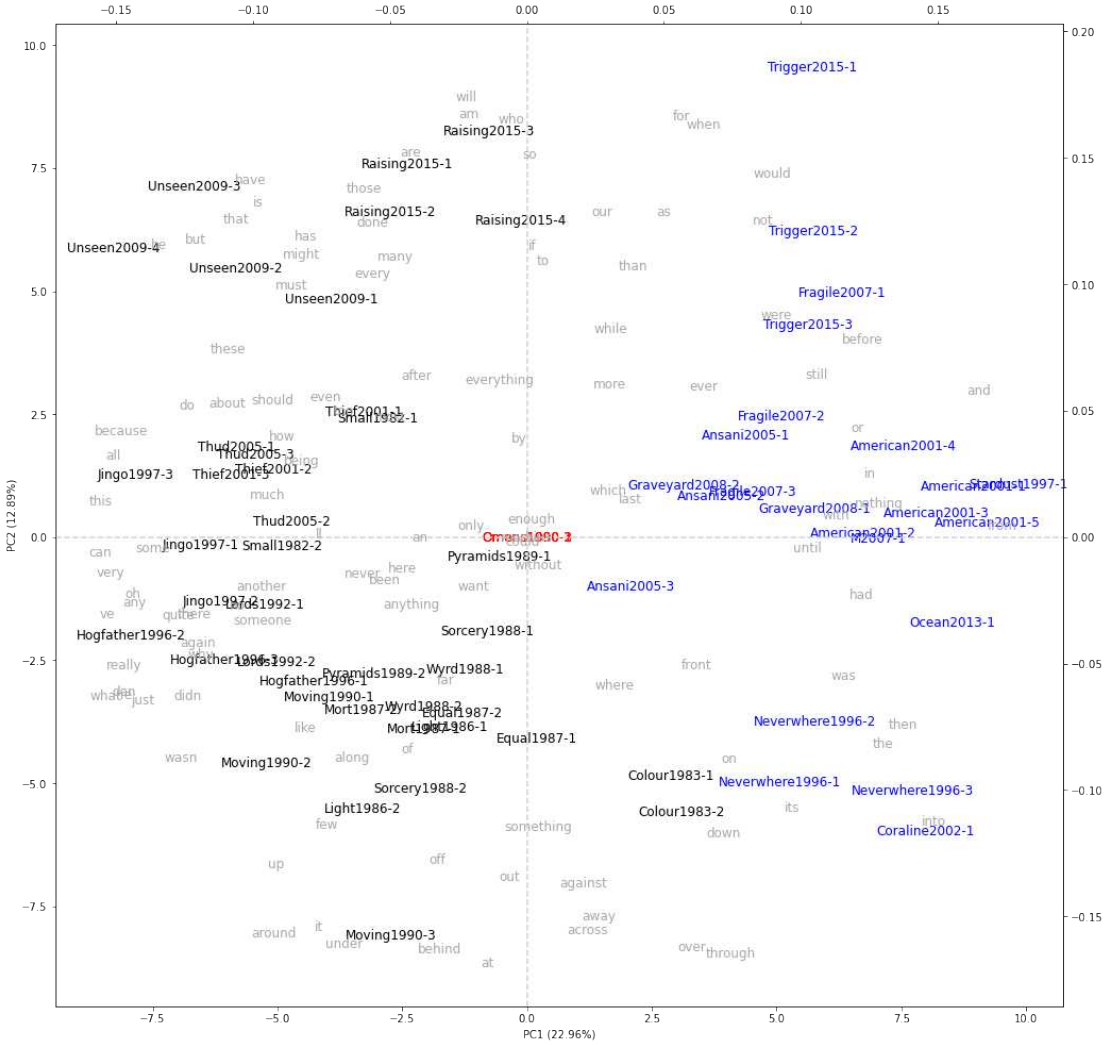

Figure 1: Principal Component Analysis of the reference corpus for Pratchett (black) and Gaiman (blue)

above-described PCA. Pratchett's novels are highlighted through warm colours, and Gaiman's through cold. The horizontal axis represents *Good Omens*' segments, while the vertical axis represents the delta distance for each segment compared to the reference novels. The closer a line comes to the x-axis, the more similar the novel represented by that line will be to the segment. The vertical lines represent the end of each of the six chapters.[12] Finally, the seven vertical lines delimit the six chapters of the book. The text was split into 5000-word-long seg-

[12]It is worth noting that *Good Omens*' very first chapter (titled *In the Beginning*) is shorter than the segment length selected to run Rolling Delta and Rolling Classify. As such, the first line (*a*) in Fig. 2 and 3 appears to be outside of the text. Although this is slightly more visible in Fig. 2, we decided not to move it to retain the original partition of the book and to make Fig. 2 and 3 more comparable.

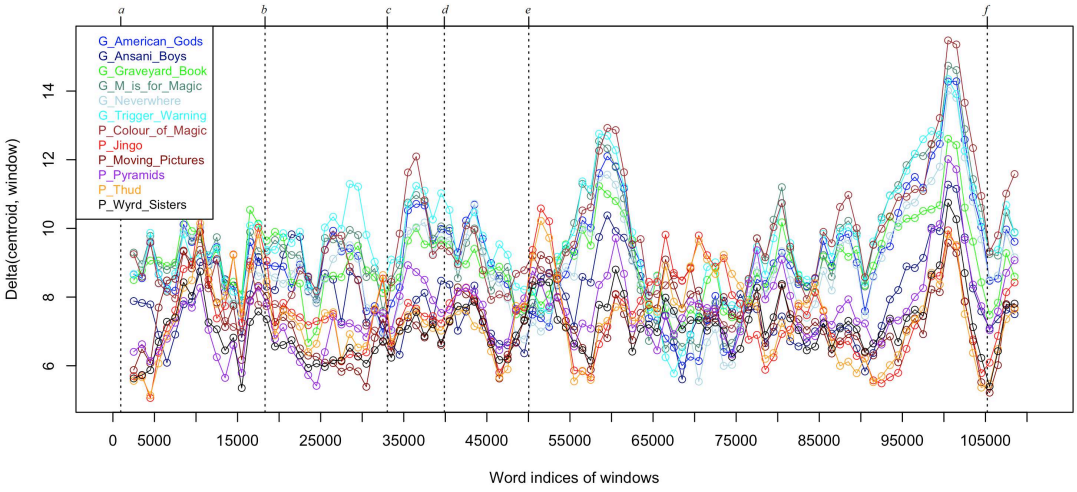

Figure 2: Rolling Delta diagram for Good Omens. Vertical lines indicate the end of each of the six chapters

ments (window size), with a rolling window (step size) of 1000 words. Each segment was then analyzed using the 250 MFW.

The warm-coloured lines come significantly closer to the horizontal line for most sections of *Good Omens*: i.e., Pratchett's style (mostly *Jingo*, *Pyramids*, and *Moving Pictures*) is predominant throughout the whole book. This is perhaps not surprising, considering that Pratchett has declared that he and Gaiman agreed that Pratchett would take on the role of editor and finalize the novel. Gaiman's style appears to be predominant in only two sections:[13] around the beginning of the 6th chapter (line *e*) and between the 65000- and 75000-word marks, with only two smaller contributions at the very beginning of the second chapter and around the 85000-word mark. This trend only partially fits with the authors' claims. As noted before, we expected Gaiman to be far more present at the beginning of the novel. Here we get a short glimpse of *Anansi Boys*' style, which shows up again further down in the middle of Chapter 6 (*e*–*f*). However, Pratchett's style (*Jingo* and *Pyramids*) is prevalent throughout Chapter 2 (*a*–*b*). Regarding the two segments in which Gaiman's idiom (specifically, *Trigger Warning*, *M is for Magic*, and *American Gods*) is present for longer sections of the novel, they both coincide with two significant scenes involving the Four Horsemen of the Apocalypse. The very start of Chapter 6 introduces Death, arguably the most important Horseman, and involves Pollution while the next span (65000–75000) corresponds to the coming together of the Four Horsemen.[14] The fact that these two segments are attributed to Gaiman aligns with statements from the authors:

[13]For a more intuitive visualization of the novels see Fig. 6 in Appendix A.2. The chromatic distinction is here by author rather than by novel. Such visualization allows to better distinguish how Pratchett's novels (as a whole) are closer to *Good Omens*' style compared to Gaiman's.

[14]The scene of the Four coming together happens at exactly word n° 70209. However, the previous scene is still related to the Four Horsemen, who are being followed by the other four characters.

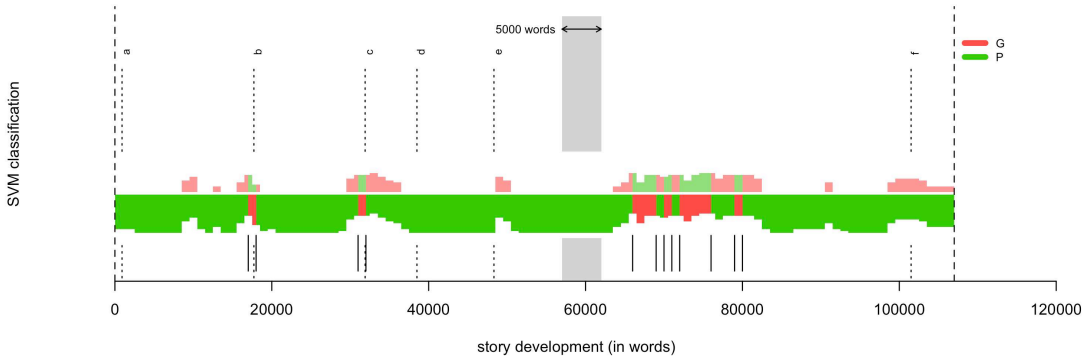

Figure 3: Rolling Classify diagram for Good Omens. Vertical lines indicate the end of each of the six chapters

'... the Four Horsemen and anything with maggots started with Neil' [14, p. 478].

Fig. 3 is the visualization of the results of Rolling Classify[15] on the novel. In distinction to Delta, Rolling Classify does not plot stylometric information per novel. The horizontal *x*-axis represents the word count of *Good Omens*, with Pratchett's presence delineated in green and Gaiman's in red. The vertical lines on the underside of the x-axis represent to whom the segment has been attributed. The upper lines, on the other hand, denote whether the second author's style is present and to what extent. The height of the lines on both sides indicates the degree of certainty with which the classification has been made. The vertical dotted lines represent chapter markers. The Rolling Classify method was configured using our list of 117 MFW to analyze text segments of 5000 words using 1000-word steps.

The Rolling Classify results generally confirm Rolling Delta's output: *Good Omens* is predominantly composed in Pratchett's style. We again observe that Gaiman's style is most discernible between the 65,000- and 75,000-word marks, with a small additional contribution at 80,000. Across the intersection of Chapters 2–3 (*b*) and Chapters 3–4 (*c*), we see short instances of text segments attributed to Gaiman's style, too. These do not correspond to the results of the Delta. Where Rolling Delta found a predominant presence of Gaiman (see Fig. 2) at the beginning of Chapter 6 (*e*–*f*), Rolling Classify attributes the segment to Pratchett, with Gaiman's presence being detected in the background. This may be related to the difference in MFW used. While Rolling Classify allows using a custom MFW list, this is not possible for Rolling Delta.[16] As such, the latter analysis may have been influenced by content words (e.g. characters' names, see Section 3) present in the unculled 250 MFW.

Compared to Callaway [8], our results attribute significantly larger chunks of *Good Omens* to Pratchett. The segments attributed to Gaiman in Callaway [8] are detected as Gaiman's

[15]As a reminder, we here use the same 12 novels used for Rolling Delta: *Jingo*, 1997, *Pyramids*, 1989, *The Colour of Magic*, 1983, *Wyrd Sisters*, 1988, *Moving Pictures*, 1990, *Thud*, 2005, for Pratchett and *M is for Magic*, 2007, *Anansi Boys*, 2005, *American Gods*, 2001, *The Graveyard Book*, 2008, *Neverwhere*, 1996, *Trigger Warning*, 2015, for Gaiman.

[16]In our Rolling Delta experiments, we tick the option which allows the user to cull personal pronouns.

Figure 4: Rolling Classifier diagram for Good Omens screenplay with episode markers

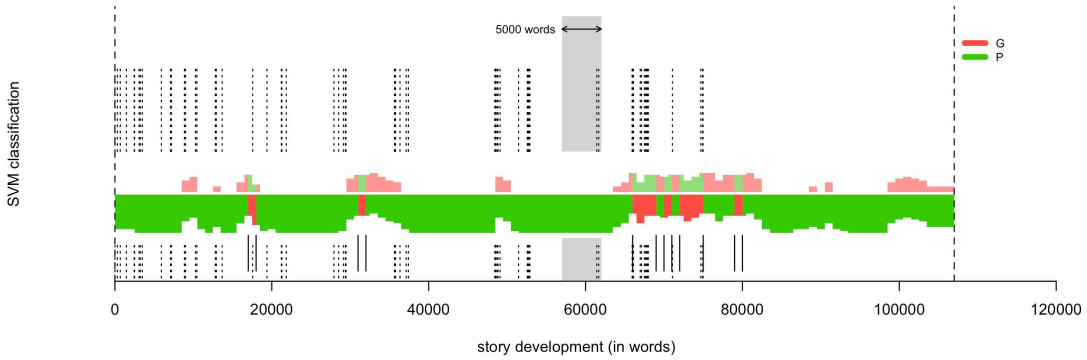

Figure 5: Rolling Classify diagram for Good Omens novel with text-match markers

authorial signals in Fig. 3, but are not attributed to him (except some smaller segments).[17] Additionally, our results show a higher degree of certainty in attributing the segments to each author, i.e., less overlay between the two authors is present throughout the novel in Fig. 3 compared to Callaway's [8] results.

Fig. 4 and 5 test our hypotheses for the screenplay. Both figures were obtained using a slice size of 5000 with an overlap of 4000. However, while Fig. 5 uses our list of 117 MFW, the screenplay was analyzed using 1000 MFW. Fig. 4 shows that our classifier correctly attributes the screenplay to Gaiman, with a segment between Episode 5 and 6 (e)[18] attributed to Pratchett.

[17]For instance, Callaway [8] results show that a large portion of text before the 2000-word mark is attributed to Gaiman while ours indicate that only a smaller section at the end of Chapter 2 is to be attributed to Gaiman. A similar pattern can be observed at the very end of the novel.

[18]Vertical lines denote episodes.

The classification results shown in Fig. 5 are the same as Fig. 3; however, the plot of the novel is here overlaid with sixty-five vertical dotted lines representing identical passages in the screenplay.[19] The screenplay was compared to the novel using the text reuse detection tool Text-Matcher, which yielded the list of matches [24]. The matches comprise both dialogue and character and scene descriptions.

Interestingly, the matching passages occur relatively frequently up until the 20,000 word mark of the novel. Then, it progressively diminishes until the 75,000 word mark. From this point of the novel onward there are no matches between the book and the screenplay. This pattern leads us to assume that over time Gaiman has relied less on the source material. These results are compatible with the observations that can be made by comparing the novel to the series: some of the most important scenes and characters from the book have been excluded from the screen adaptation[20] while others are solely present in the show.[21] It is worth noting that further analysis is needed to explore the style of the screenplay and that our conclusions' reliability is limited by the scope of this study.

5. Conclusions

The present paper aimed to explore authorial takeovers in *Good Omens* by Terry Prachett and Neil Gaiman. Additionally, we also compared the novel to the screenplay of the show written by Gaiman and based on the book.

The application of stylometric techniques to the works of the two authors yields interesting results. From the PCA, we can observe how Pratchett's novels written after 2007, the year of his Alzheimer's diagnosis, cluster differently from most of his works. This pattern denotes a shift in Pratchett's writing style, which may be related to his neurological disease.[22] Interestingly, PCA also locates *The Colour of Magic*, Prattchett's breakthrough work, next to *Neverwhere*—one of Gaiman's first novels. The clustering suggests that *The Colour of Magic* may have influenced Gaiman's early idiom.[23]

Rolling Delta partially confirms our expectations of the novel. Pratchett's idiom is predominant in the book. Instances of Gaiman's style, especially throughout Chapter 6, can largely be attributed to the presence of the Four Horsemen in those sections, characters that he authored [14].

[19] The chapter markers are not present in this visualisation.

[20] For instance, the highway chase at the beginning of Chapter 6, where four bikers decide to follow the Four Horsemen in their ride through the M25 highway is not present in the show

[21] E.g., the finale, during which Aziraphale and Crowley switch bodies to survive the punishments of Heaven and Hell, is absent from the novel

[22] Our conclusion is here derived from an observation of stylometric patterns and does not account for the complex nature of Alzheimer's. There is little academic research regarding the effect of Alzheimer's on Pratchett's writing style. The only article on the issue was published on the Pratt School of Information's website by one of the institute's students (see [28]). Here, the author outlines how vocabulary complexity has not diminished but rather increased throughout Pratchett's last novels, thus concluding (with the necessary reservations) that his neurological condition did likely not affect his writing style.

[23] This was the first of Pratchett's works read by Gaiman (Gaiman, 2018). Critical literature on Pratchett notes that many writers "have found after lengthy exposure to Pratchett's prose that it has worn grooves in their heads" [1, p.148].

Rolling Classify generally confirms the results of Rolling Delta, except for two additional shorter segments being attributed to Gaiman. Compared to the results obtained by previous studies [8], we find Pratchett to be far more predominant throughout the novel. Our results reveal a higher degree of confidence in attributing segments to each author, showing fewer overlays between Gaiman's and Pratchett's styles compared to Callaway's [8].

The screenplay analysis further shows that the classifier can correctly attribute the text almost entirely to Gaiman despite the difference in genre. Based on text matches between the screenplay and the novel, we speculate that Gaiman may have relied less on the source material towards the end of the screenplay. It is worth noting that the use of the screenplay as control for the efficacy of our classifier is limited as the two texts do not belong to the same genre. Further research could explore the issue of the screenplay by retrieving other screenplays written by Pratchett and Gaiman and that of the upcoming second season of *Good Omens*.

6. Code and data availability

Code and datasets are available at https://zenodo.org/record/7257715

7. Acknowledgments

A special thanks to Prof. Mike Kestemont and Dr. Wouter Haverals, who supported and encouraged us during the making of this project. We also want to thank Eveline C. for allowing us to use her living room as our office.

References

[1] A. H. Alton, W. C. Spruiell, and D. Palumbo. *Discworld and the Disciplines: Critical Approaches to the Terry Pratchett Works (Critical Explorations in Science Fiction and Fantasy, 45)*. annotated edition. McFarland & Company, 2014.

[2] BBC News. "Good Omens: How Neil Gaiman and Terry Pratchett wrote a book". In: (2014). URL: https://www.bbc.com/news/magazine-30512620.

[3] J. N. G. Binongo. "Who Wrote the 15th Book of Oz? An Application of Multivariate Analysis to Authorship Attribution". In: *Chance* 16.2 (2003), pp. 9–17. DOI: 10.1080/09332480.2003.10554843. eprint: https://doi.org/10.1080/09332480.2003.10554843. URL: https://doi.org/10.1080/09332480.2003.10554843.

[4] J. N. G. Binongo and M. W. A. Smith. "The application of principal component analysis to stylometry". In: *Literary and Linguistic Computing* 14.4 (1999), pp. 445–466. DOI: 10.1093/llc/14.4.445.

[5] L. Breebaart. *The Annotated Pratchett File v9.0 - Words from the Master*. 2016. URL: https://www.lspace.org/books/apf/words-from-the-master.html.

[6] J. Burrows. "'Delta': a Measure of Stylistic Difference and a Guide to Likely Authorship". In: *Literary and Linguistic Computing* 17.3 (2002), pp. 267–287. DOI: 10.1093/llc/17.3.267.

[7] F. Cafiero and J. Camps. "'Psyché' as a Rosetta Stone? Assessing Collaborative Authorship in the French 17th Century Theatre". In: *Proceedings of the Conference on Computational Humanities Research, CHR2021, Amsterdam, The Netherlands, November 17-19, 2021*. Ed. by M. Ehrmann, F. Karsdorp, M. Wevers, T. L. Andrews, M. Burghardt, M. Kestemont, E. Manjavacas, M. Piotrowski, and J. van Zundert. Vol. 2989. CEUR Workshop Proceedings. CEUR-WS.org, 2021, pp. 377–391. URL: http://ceur-ws.org/Vol-2989/long%5C%5Fpaper51.pdf.

[8] E. Callaway. *Good Omens Stylometry*. Elizabeth Callaway. URL: http://www.elizabethcallaway.net/good-omens-stylometry.

[9] J. B. Croft. "Nice, Good, or Right: Faces of the Wise Woman in Terry Pratchett's "Witches" Novels". In: *Mythlore: A Journal of J.R.R. Tolkien, C.S. Lewis, Charles Williams, and Mythopoeic Literature* 26.3 (2008).

[10] M. Deuchar. "Are function words non-language-specific in early bilingual two-word utterances?" In: *Bilingualism: Language and Cognition* 2.1 (1999), pp. 23–34. DOI: 10.1017/s1366728999000127.

[11] G. Dougary. "Good Omens: Neil Gaiman reveals what he and Terry Pratchett shared". In: (2019). URL: https://www.smh.com.au/culture/tv-and-radio/good-omens-neil-gaiman-reveals-what-he-and-terry-pratchett-shared-20190603-p51u1y.html.

[12] M. Eder. "Rolling stylometry". In: *Digital Scholarship in the Humanities* 31.3 (2016), pp. 457–469. DOI: 10.1093/llc/fqv010.

[13] S. Evert, T. Proisl, T. Vitt, C. Schöch, F. Jannidis, and S. Pielström. "Towards a better understanding of Burrows's Delta in literary authorship attribution". In: *Proceedings of the Fourth Workshop on Computational Linguistics for Literature*. Denver, Colorado, USA: Association for Computational Linguistics, 2015, pp. 79–88. DOI: 10.3115/v1/W15-0709. URL: https://aclanthology.org/W15-0709.

[14] N. Gaiman and T. Pratchett. *Good Omens: The Nice and Accurate Prophecies of Agnes Nutter, Witch (Cover may vary)*. William Morrow, 1990.

[15] N. Gaiman [neilhimself]. *The Colour of Magic*. [Tweet]. 2018. URL: https://twitter.com/neilhimself/status/1023385399694163969.

[16] T. Joachims. "Text categorization with Support Vector Machines: Learning with many relevant features". In: Berlin, Heidelberg: Springer Berlin Heidelberg, 1998, pp. 137–142. DOI: 10.1007/bfb0026683.

[17] I. T. Jolliffe. "Principal Component Analysis and Factor Analysis". In: *Principal Component Analysis* (1986), pp. 115–128. DOI: 10.1007/978-1-4757-1904-8_7.

[18] F. Karsdorp, M. Kestemont, and A. Riddell. *Humanities Data Analysis: Case Studies with Python*. Princeton University Press, 2021.

[19] M. Kestemont, S. Moens, and J. Deploige. "Collaborative authorship in the twelfth century: A stylometric study of Hildegard of Bingen and Guibert of Gembloux". In: *Digital Scholarship in the Humanities* 30.2 (2013), pp. 199–224. DOI: 10.1093/llc/fqt063.

[20] M. Kestemont, J. Stover, M. Koppel, F. Karsdorp, and W. Daelemans. “Authenticating the writings of Julius Caesar”. In: *Expert Systems with Applications* 63 (2016), pp. 86–96. DOI: 10.1016/j.eswa.2016.06.029.

[21] T. Litvinova and O. Litvinova. “Analysis and Detection of a Radical Extremist Discourse Using Stylometric Tools”. In: *Digital Science 2019* (2019), pp. 30–43. DOI: 10.1007/978-3-030-37737-3_3.

[22] D. Müllner. “Modern hierarchical, agglomerative clustering algorithms”. In: *arXiv* (2011).

[23] R Core Team, Vienna, Austria. “R: A language and environment for statistical computing”. In: *R Foundation for Statistical Computing* (2020). URL: https://www.R-project.org/.

[24] J. Reeve. “Text-matcher”. In: *Github journal* (2020). DOI: 10.5281/zenodo.3937738.

[25] J. Rybicki, D. Hoover, and M. Kestemont. “Collaborative authorship: Conrad, Ford and Rolling Delta”. In: *Literary and Linguistic Computing* 29.3 (2014), pp. 422–431. DOI: 10.1093/llc/fqu016.

[26] J. Shanahan. “Terry Pratchett: Mostly Human”. In: *Twenty-First-Century Popular Fiction*. 1st ed. Amsterdam, Netherlands: Amsterdam University Press, 2017, p. 31.

[27] J. Stover and M. Kestemont. “Reassessing The Apuleian Corpus: A computational Approach To Authenticity”. In: *The Classical Quarterly* 66.2 (2016), pp. 645–672. DOI: 10.1017/s0009838816000768.

[28] *Were Terry Pratchett’s Final Works Affected by Alzheimer’s Disease?: An Analysis into Vocabulary Trends within the Discworld Series, Post Diagnosis*. Tech. rep. 2016. URL: https://studentwork.prattsi.org/dh/2016/05/08/were-terry-pratchetts-final-works-affected-by-alzheimers-disease-an-analysis-into-vocabulary-trends-within-the-discworld-series-post-diagnosis/.

A. Additional Figures and Tables

A.1. SVM set-up

Tables 2 and 3 report the results obtained with different SVM set-ups. The experiments were carried out using a 0.70–0.15–0.15 train-validation-test split. The validation set was utilized to verify the consistency of the results. The training data was limited to the 12-novel sub-corpus since the aim of these experiments was to understand whether the classifiers available in Rolling Classify could correctly attribute different text segments to each author. Because Rolling Stylometry functions only allow a maximum of twelve texts (six for each author), using a greater training set for our models would have contradicted the experiments' goal. I.e., even if a classification algorithm trained on all available texts had reached better results, they would have not been a reliable foundation for the Rolling Classify experiments.

Table 2
Performance of different classifiers for different sentence lengths on the 12-novels sub-corpus. MFW set to 1000

Classifier	Segment Length	MFW	Accuracy	Marco-avg
svm.SVC	250	1000	0.99	0.99
svm.SVC	500	1000	1.00	1.00
svm.SVC	1000	1000	1.00	1.00
KNeighborsClassifier	250	1000	0.97	0.97
KNeighborsClassifier	500	1000	0.99	0.99
KNeighborsClassifier	1000	1000	1.00	1.00
LogisticRegression	250	1000	0.99	0.99
LogisticRegression	500	1000	0.99	0.99
LogisticRegression	1000	1000	1.00	0.99

Table 3
svm.SVC performance for different MFW on the 12-novels sub-corpus. Segments length set to 1000 following the results in Table 2

Classifier	Segment Length	MFW	Accuracy	Marco-avg
svm.SVC	1000	50	0.92	0.93
svm.SVC	1000	100	0.98	0.98
svm.SVC	1000	117	0.99	0.99
svm.SVC	1000	250	1.00	1.00
svm.SVC	1000	500	1.00	1.00
svm.SVC	1000	1000	1.00	1.00

A.2. Rolling delta with color coding per author

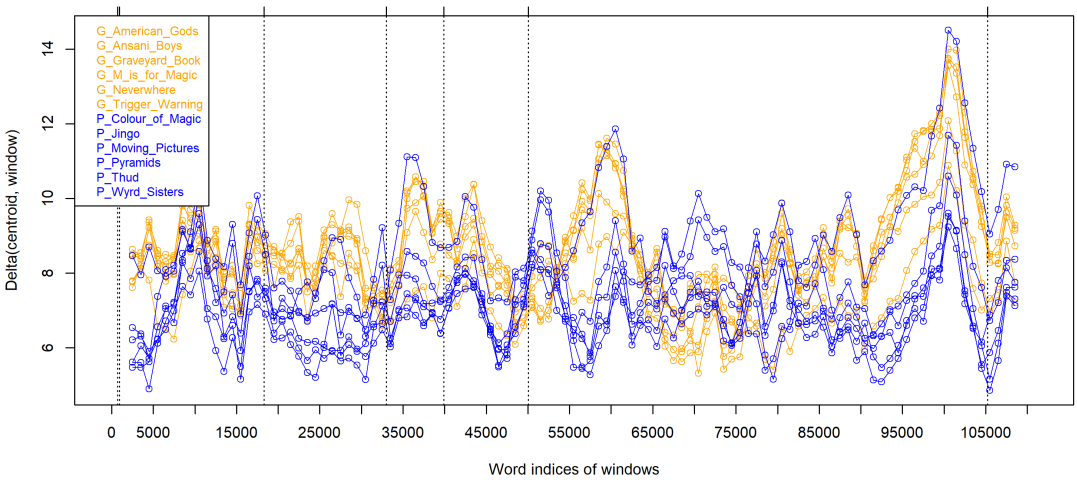

Figure 6: Rolling Delta diagram for Good Omens 250 MFW, window size 5000, step size 1000. Gaiman novels are in amber, Pratchett's are in blue