The emotional arcs of horror: a distant reading of Stephen King's novels

Delaney Woods

Abstract: Sentiment analysis, the computational inference of emotion in text through Natural Language Processing, is increasingly used to analyze social and cultural trends. In this thesis, we create narrative time-series and word-shift graphs for each of Stephen King's novels using the Hedonometer, quantifying the lexical changes responsible for emotional arcs found in each story. Our results suggest King's work has increasingly shifted in genre from horror to science fiction. The work contributes to a growing science of stories being developed by the Computational Story Lab.

Background

Stories are the primary way we learn and share information with others, conveying wisdom over generations and between cultures. Stories are important to politics and news, and they shape the way we feel about events in our daily life. For example, the story that COVID-19 is fake has influenced millions of people and cost our society countless lives. The tools of computational social science are revealing insights about the power of stories. Natural Language Processing of literature, for example, allows us to quantify previously theoretical constructs about culture, and reveal values of the time.

Stephen King is a very influential and popular author in the horror genre. He has published 62 novels and 11 short story collections. He published his first book, *Carrie*, in 1974 and published his most recent book, *Billy Summers*, in 2021. His long career allows for temporal analysis and raises the question of how the emotional language found in his work developed over his career. Many aspects of his novels stay similar. For example, a lot of King's books take place in Maine, where he creates a natural setting that is "hostile and savage" and where "malefic energies reside in secret" (Magistrale 2002). King typically writes novels that fit in the following genres: Gothic horror, dystopian technology, epic fantasy, and the journey quest" (Magistrale 2002). Another common factor across Stephen King's books according to Anthony Magistrale is that the evil King creates exists because of small town pressure to conform and a lack of compassion (Magistrale 2002).

Magistrale notices some changes throughout Stephen King's career. He notes that in the 1990s, King's novels become shorter and show a different, more realistic depiction of women. Magistrale writes, "*Misery* holds a pivotal position in King's canon; the novel signals a transition that begins to emphasize a new significance for women characters" (Magistrale 2002). *Misery* is about a woman, Annie Wilkes, who holds a popular romance author captive. Annie plays a very important role in this novel and Magistrale believes this trend for women characters continues. Magistrale also reports that "in his early fiction, individuals move out into a confrontation with the Gothic world; his later work, in contrast, features a domesticized Gothic, where individuals are under assault within their own homes" (Magistrale 2002).

One goal for this thesis is to explore whether any of these changes are reflected in the emotional arcs of the novels. Magistrale believes that King's novels end in an "overly optimistic way" and it will be interesting to see if the sentiment analysis will illustrate this (Magistrale 2002). The happy endings may be due to editorial pressure to not end on a bad note, especially in his early career. King may also just personally like having his novel's end on a higher note.

Within Stephen King novels, there is the Gothic horror that fans want, but this is typically paired with a discussion of social issues and problems within American society. He uses horror to reveal truths about society and the bad effect they have on individuals (Magistrale 2002). This social commentary is an important part of his work and most likely contributes to his overall success and popularity. It reveals the anxieties of the time and allows us to learn about the adjacent possibilities for our evolving culture.

Analyzing books through sentiment analysis can reveal patterns in the language found in books that one would not be able to learn through other kinds of analysis. Using sentiment analysis gives a satellite view of the books to see patterns, rather than a closer look at characters or dialogue. Stories are often looked at for themes, symbols, and characters, but this analysis will add a layer of distant reading that is novel. This field of statistics with literature is a new way to analyze a story. Sentiment analysis is a very hot topic currently and there is a lot of research being done to analyze it. There are many articles about applying sentiment analysis to different data sets. For example, in my research, I found some investigators using the dataset of children's books and fairytales (Alm 2005), a book by Virginia Woolf (Elkins 2019), individual genres (Kim 2017), one playwright (Schmidt 2018), and even YouTuber vlogs (Kleinberg 2018). There are endless data sets to choose from which is why this field is so interesting and why there is so much new work that can still be done.

For example, the Hedonometer is a project that measures the daily collective mood in ten languages. Researchers at the Computational Story Lab use the Hedonometer to look at average happiness on Twitter. In Figure 1, the higher spikes have higher happiness that day and the lower ones have less happiness. To do this, the Hedonometer looks at a random 10% of all tweets from the 500 million that are posted every day to find the average happiness for each day. To make this estimate, they use a dictionary with 10,000 words that have been rated on a continuous sad/happy scale by people. This graph shows trends that we can analyze. For example, every year Christmas Day is very happy and has a higher plot point. Certain events are much lower like the protests against police brutality. The highest point on this graph, Christmas Day of 2020, has a happiness score of a little more than 6.2. The lowest point on this graph, protests against police brutality over the summer of 2020, has a happiness score of a little more than 5.6. These two points only have a difference of about 0.6. This might not seem like a lot but with the Hedonometer this difference is significant.

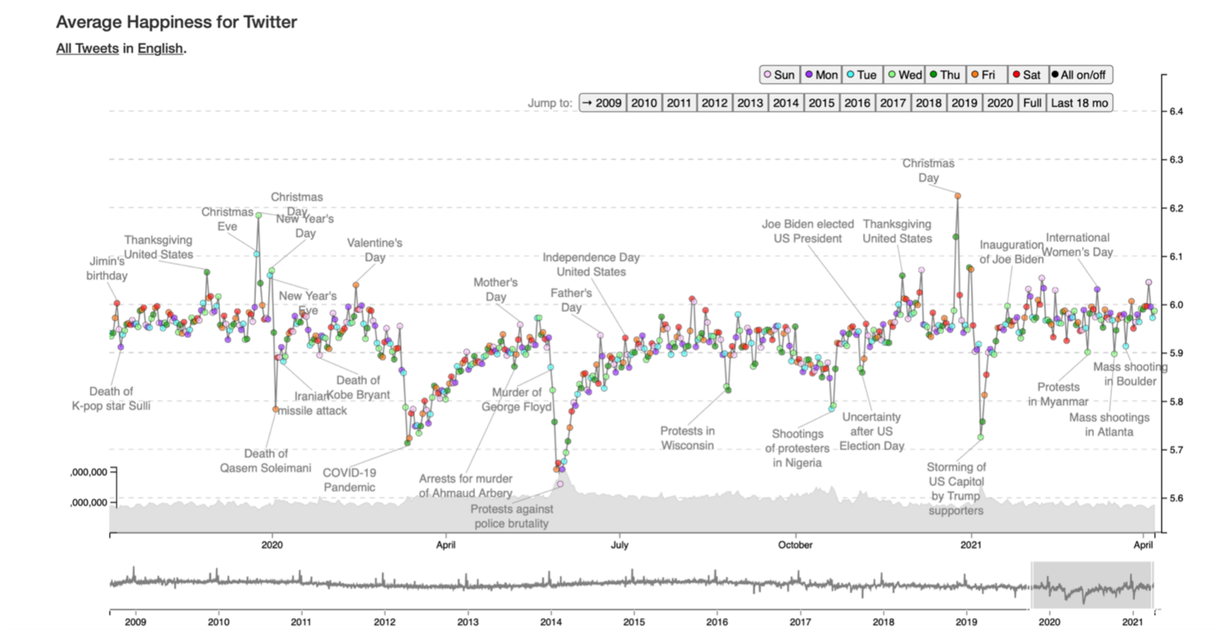

Figure 1: Average Happiness for Twitter. Snapshot of the daily estimate of English tweet happiness, estimated using sentiment analysis and provided by http://hedonometer.org

An important article that my work will be based on is the work done by Andrew Reagan and colleagues on the emotional arcs of stories. This paper reveals the findings from doing sentiment analysis on a large set of fiction books. They used the Hedonometer to find the average happiness for sections of each book and then plotted this data. This research found that from all their story arcs, there are six basic shapes that dominated (Reagan et. al. 2016). The researchers named these arcs “rags to riches,” “man in a hole”, “Cinderella”, “tragedy”, “Icarus”, and “Oedipus”. Over ninety percent of books followed this set of six arcs, supporting Kurt

Vonnegut's theory from the 1970s (Vonnegut 1995). Vonnegut suggested that stories had shapes and that we could use computers to find these shapes. He was correct! The dominant arcs are shown in Figure 2 (Reagan et. al. 2016).

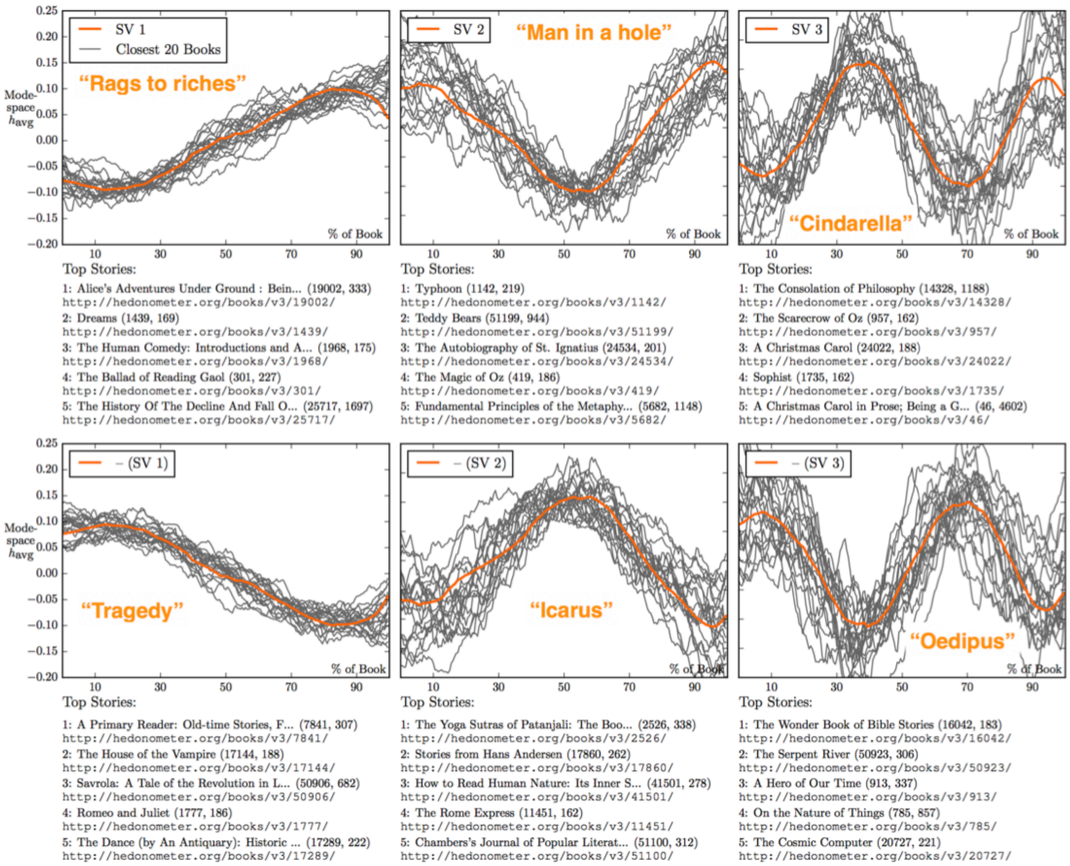

Reagan et al "The emotional arcs of stories are dominated by six basic shapes" EPJ Data Science 2016

Figure 2: The 6 emotional arcs of fiction. Six basic shapes dominate the emotional arcs found in thousands of stories. Reagan et. al. (2016) used sentiment analysis to estimate the emotion experienced by a reader, inspired by Kurt Vonnegut's 'Shapes of Stories' lecture.

As an example comparison between the experience of the reader and the machine estimated sentiment, in Figure 3 we see human ratings of their feelings while reading Romeo & Juliet (black) along with the Hedonometer estimated sentiment as a function of narrative time (blue). The two lines follow a similar emotional arc and have the same highs and lows, showing that the machine estimation of sentiment is a good estimation. Matthew Jockers created emotional arcs using machine estimated sentiment and then annotated his graph with high and low points throughout the story (Jockers 2014). These points lined up well with the graph created by the computer, again showing that these graphs are accurate as they match human readings of the books.

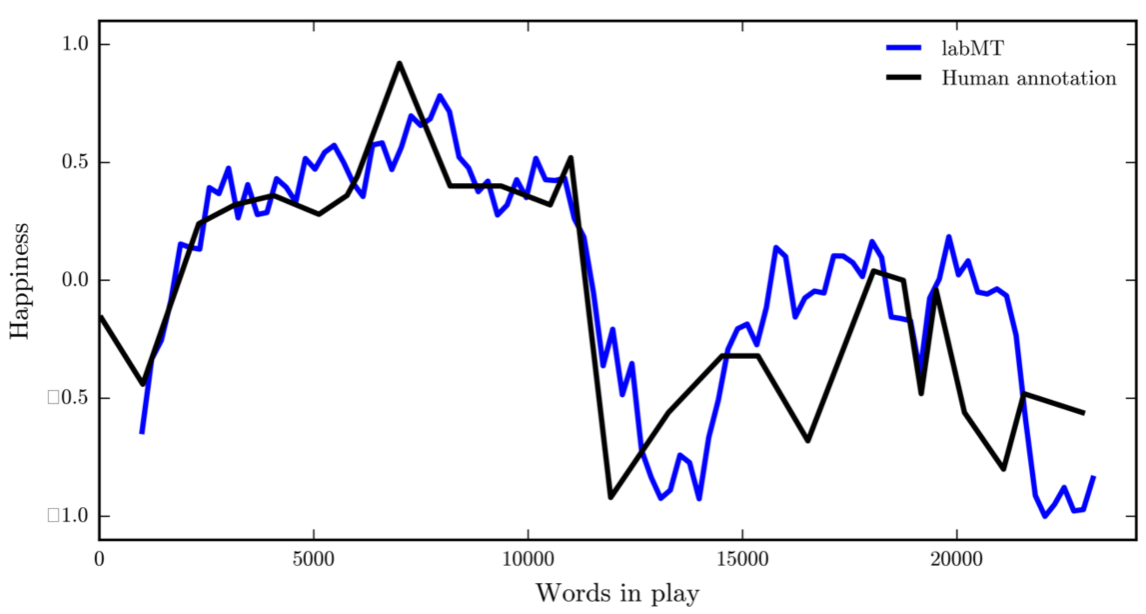

Figure 3: Romeo and Juliet happiness ratings human versus computer. This graph compares the emotional arc made by the Hedonometer to the emotional arc made by humans rating their feelings. Underwood (2015)

However, books are often more complicated than the curves pictured in Figure 2. Some books cannot be simplified into any of these forms and have more variation. For example, Figure 4 was made for the last Harry Potter book (Reagan et. al. 2016). This graph would not fit into any of the arcs in Figure 2. It has many more curves and to try to simplify it to a simpler curve would cause us to lose a lot of valuable information. This graph was also annotated with events throughout the story that corelate with certain points on the graph.

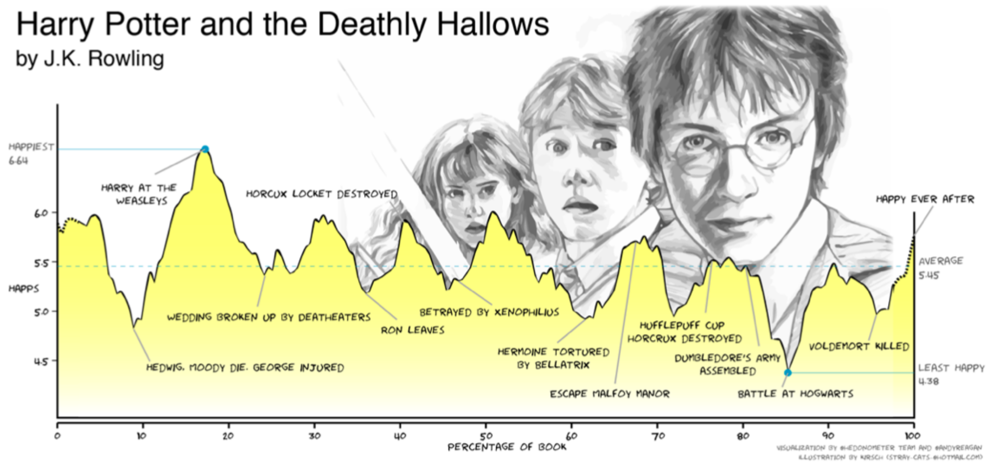

Reagan et al "The emotional arcs of stories are dominated by six basic shapes" EPJ Data Science 2016

Figure 4: Emotional Arc for Harry Potter and the Deathly Hallows. Not all stories can be simplified into one of the dominant arcs. Longer novels often stitch together multiple emotional arcs.

David McClure looked at the distribution of individual words throughout a novel at the Stanford literary lab to see if certain words are associated with beginnings, middles, and ends. To do this, he looked at 27,266 novels. He showed his results in a time series plot with the x axis being narrative time. McClure found that the word “death” is used more often towards the end of novels. He also found that “athletic” is used at the beginning and he suggests that this is because this word is used to describe people, and characters are typically described at the beginning of novels. It is also fair to assume that death typically occurs at the ends of novels as part of the conclusion. Despite these clear trends, some words do not have trends that are as obvious, such as the word “irony”. McClure writes, “it looks like – beginnings are about youth, education, physical size, (good) appearance, color, property, hair, and noses? And endings – forgiveness, criminal justice, suffering, joy, murder, marriage, arms, and hands?” (McClure 2017).

As stated earlier, sentiment analysis is important to learn about culture and is being used on many different data sets with different applications. For example, one article asked why certain narratives or texts were more successful than others and looked at pace, volume, and circuitousness. The data set used by these researchers were TV shows, movies, and research articles (Toubia 2021). They found that TV shows and movies that are faster paced are more successful, but this is not the case for research articles (Toubia 2021). Knowing this is helpful when trying to create media that people will enjoy. Other researchers also looked at arcs in movies to predict success. This article was based off Andrew Reagan's article and attempts to create emotional arcs for movies. They found that emotional arcs in movies can be separated into 6 basic shapes, just like fiction stories (Del Vecchio 2020). They also came to the conclusion that the Man in the Hole story arc leads to the most revenue (Del Vecchio 2020). This is true for movies despite their genre which is very interesting. Hipson and his team looked at emotion dynamics in movie dialogues (Hipson 2021). These researchers wanted to trace the emotional arcs of individual characters in movies. This research found that negative words used by characters increased as the movie went along. Characters used more positive words in general than negative words and the researchers also split positive and negative into more specific categories such as trust and disgust. This research can extend past literary studies and be used to help public health.

There are also historic reasons to use sentiment analysis. One article looked at diary entries from soldiers in Australia during WWI (Dennis-Henderson). Historic data is now often available in digital text formats which allows for more analysis of them, when before this was not possible. The researchers found that the authors wrote more about everyday experiences than war experiences (Dennis-Henderson). Their overall sentiment was also positive throughout the war (Dennis-Henderson). This positive overall sentiment reflects the idea that there is a positivity bias in human language (Dodds et. al. 2015). This study confirms the Pollyanna hypothesis which is the idea that people remember positive things better. Another study tested to see if this positive bias was consistent for children's literature (Jacobs 2020). They looked at German and English children literature and found that these data sets also followed the Pollyanna hypothesis (Jacobs 2020).

Method

First, I needed to get all of Stephen King's books in the form of text files. I used Library Genesis to download the 62 books and 11 short story collections. I then downloaded the Hedonometer word list in English that has the happiness scores for the 10,000 most common words in English (Dodds et. al. 2015). This labMT data set was created by using Twitter, Google Books, music lyrics, and the New York Times to find words used most frequently (Dodds et. al. 2011). The top 5,000 words from each were taken and combined to find 10,222 unique words that now make up the labMT data set (Dodds et. al. 2011). To find the happiness scores for each word, users on Mechanical Turk were given words from the data set and asked to rate how a word made them feel on an integer scale of 1-9 where 1 is very sad, 5 is neutral, and 9 is very happy (Dodds et. al. 2011). They were given 100 tasks where 100 words were rated at a time and each unique word was rated by 50 individuals (Dodds et. al. 2011). Each word would then have an average happiness score at the end of this process.

Next, I used code to split the texts into windows based on word count and used the Hedonometer to create a happiness score for each window. Each word in the window could be given a score from the Hedonometer and then the average for that window could be found. I excluded words that were rated between 4 and 6 because these words are neutral or context dependent and not important to the overall happiness score. These scores were saved to a file for each book. I did this for the window sizes of 1000 words, 2000 words, and 5000 words. I ended up choosing the window size of 5000 words to use for the rest of my research. Once I had these files with the happiness scores for each window, I used them to create a time series graph for each book that plotted the happiness throughout the book. I used Jupyter Lab to create the files and the graphs. I did not include the short story collections in this section of data.

To find the overall happiness score for each book, I took the files that had the happiness score for each window. I used the window size with 5000 words and used code in Jupyter lab. I added all the happiness scores together and then divided by the number of windows to find the average which is the average happiness score for the whole novel. I did this for all 62 novels and 11 short story collections. I created a table that includes all 73 of these, sorted from saddest to happiest. This table then allowed me to create a timeseries graph where the x axis is the year, and the y axis is the happiness. Each book was represented on this scatterplot with a point and a least squares regression line was fitted to the graph.

To create the word-shift graphs, I downloaded Ryan Gallagher's "shifterator" package on GitHub (Gallagher et. al. 2021). This package allowed me to create word-shifts that are vertical bar charts to compare books. The code found the overall happiness score for each book and which words were used most often with whether the words were positive or negative. I also excluded words with happiness scores between 4 and 6 in this data set because of their neutrality and dependence on context.

Results

Figure 5 shows three word-shift graphs that demonstrate some different comparisons. The first word-shift (far left) compares the language found in *Elevation* and *The Shining*. *The Shining* was published in 1977, closer to the beginning of King's career, and *Elevation* appeared in 2018, much later in King's career. This word shift reveals that *Elevation* has an average happiness score of 5.97 compared to *The Shining* which has a score of 5.69*.* The positive words "good", "smile", "first", "new" and "Christmas" are used more often in *Elevation*. More negative words like "against", "dead", "hurt", "fell", "broken", "blood", and "kill" appear less often in *Elevation* and more often in *The Shining*, contributing to *Elevation*'s relatively higher happiness score. There are many more negative words used more often in *The Shining.* We do find that positive words like "daddy", "like", "mommy" and "great" are used more in *The Shining* and that the less negative words "not", "bill", "lose", and "zero" appear less often. However, the word "bill" may be referring to the character Bill, the protagonist's cat in *Elevation.* All words were converted to lowercase before making the word-shifts, which is why it is being counted here. The words were scored out of context, so some words will contribute imperfectly. The most common use of the word "bill" in everyday language is for money that is owed, which has a negative sentiment. Although this is being scored incorrectly, we will leave it since *Elevation* already has a high score and bill is not extremely negative.

The middle word-shift compares the first 10 years of King's career (1974-1984) and the most recent 10 years of his career (2011-2021). The first 10 years have an average score of 5.72 and the last 10 years have an average score of 5.84. The top four bars indicate four separate contributions to the difference between decades; namely (1) relatively positive words appearing more frequently in the most recent 10 years (longest yellow bar) and (2) relatively negative words appearing less often in the last decade (light blue bar). These both contribute to the most recent decade being happier. Going against the trend are (3) relatively positive words appearing less often in the most recent decade, and (4) relatively negative words appearing more often in the last 10 years. The positive words “she”, “mom”, “kids”, and “couple” appear more often in the most recent decade. The substantive negative words “dead” and “fear” appear less often, as do the more mundane negations “don’t”, “can’t”, and “couldn’t”. Going against the finding that the initial decade is sadder, we see the relatively positive words “father”, “sunlight”, and “smiled” appear more often in the first 10 years, and the relatively negative words “never”, and “not” appear less often during this period. The more negative words “suicide” and “prison” also appear less often in the first 10 years. Another important finding is that words associated with women such as “she”, “mom”, “woman”, and “women” appear more often in the last 10 years, showing that King writes more significant woman characters in more recent years than the beginning of his career.

The far right word-shift compares two sections from the book *Firestarter*. It compares the first 20% to the final 20% and finds that the first 20% has an average score of 5.75 and the last 20% has an average score of 5.69. In the first 20%, positive words like “all”, “smiled”, “laughed”, and “green” appear more often, and substantive negative words like “gun”, “kill”, and “shot” are used less often as well as less negative words like “not” and “no”. These contribute to the first 20% having a higher happiness score. In contradiction, words like “she”, “father, “computer”, and “power” are positive words that appear more often in the last 20%. Negative words like “bad”, “pain”, “afraid”, and “hurt” appear less often in this last 20%.

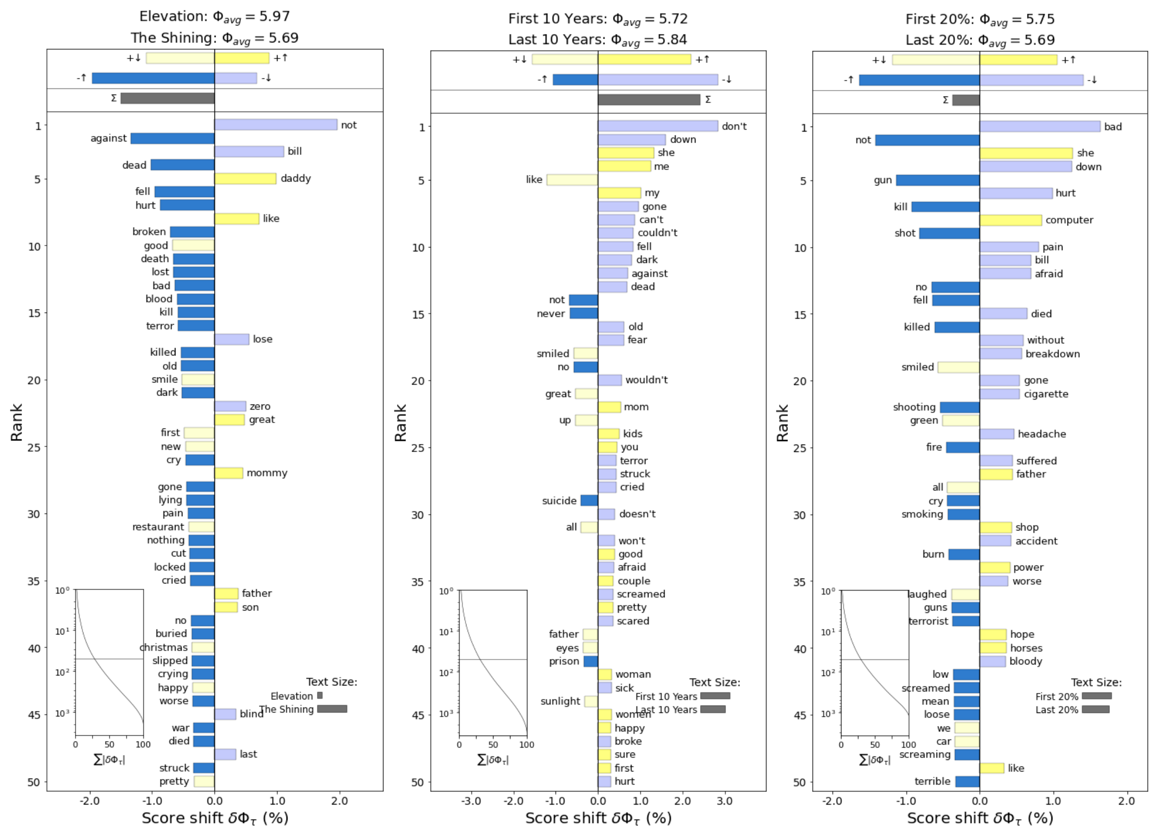

Figure 5: Word Shifts. The left word-shift compares two novels: *Elevation* and *The Shining*. The middle word-shift compares the first 10 years of King's career to his last (most recent) 10 years. The right word-shift compares one novel: the first 20% of *Firestarter* to the last 20% of *Firestarter*.

We also looked at timeseries graphs of Stephen King books to visualize their emotion as seen in Figure 6 which has timeseries for three of King's books: *Elevation* (2018), *Firestarter* (1980), and *The Shining* (1977). The x axis shows the percentage of the book, and the y axis shows the happiness score for that point of the book. *Elevation* has a much flatter curve compared to the other two and hovers between 5.8 and 6.0. This contributes to the high happiness score of *Elevation* revealed in the first word-shift of Figure 5. *Firestarter* hovers between abut 5.6 and 5.8 at the beginning and has a very low point, its lowest point, right before the end of the book followed by a large upward curve. This low point contributes to why the last

20% has a lower happiness score than the first 20% that we found in Figure 5. This low point is at 5.211 and the book ends at 5.894. *The Shining* also has more curves and variation. Just like *Firestarter*, the book starts out around 5.6 and the lowest point is right before the end at about 5.2. The highest point for *The Shining* happens early in the book and is around 6.1.

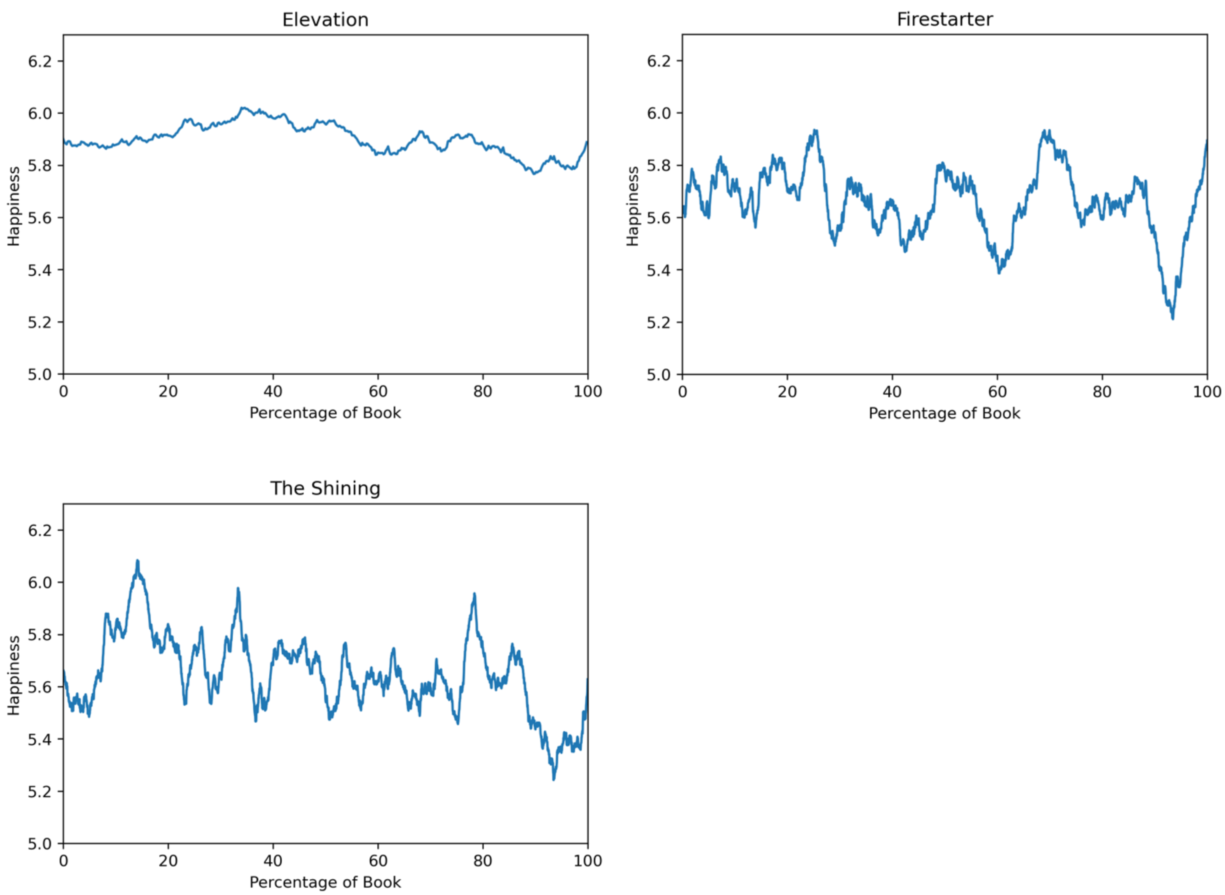

Figure 6: Timeseries graphs for 3 of Stephen King's novels. The timeseries are for *Elevation*, *Firestarter*, and *The Shining*. The x axis shows the percentage of the book, sometimes referred to as "narrative time", and the y axis shows the happiness as calculated by sentiment analysis.

We were able to find an average happiness score for each text (Table 1). This table contains the name of the book, the year it was published, and the average happiness score for that book. This table includes the short story collections that Stephen King wrote for a total of 73 rows. They are sorted from saddest score to happiest score. The saddest book is *The Dark Tower*

II: The Drawing of the Three (1987) with a happiness score of 5.550. That same year of 1987 has the second saddest book which is *Misery* with a score of 5.563. The happiest book is *Elevation* (2018) with a happiness score of 5.970. The second happiest book is *Gwendy's Button Box* (2017) with a happiness score of 5.924. The average happiness of all these books combined is 5.715.

Name	Year	Score
The Dark Tower II: Drawing of the Three	1987	5.550
Misery	1987	5.563
The Green Mile	1996	5.587
Desperation	1996	5.606
IT	1986	5.610
Dreamcatcher	2001	5.611
The Long Walk	1979	5.612
The Dark Tower I: The Gunslinger	1982	5.621
The Running Man	1982	5.625
The Eyes of the Dragon	1987	5.635
The Regulators	1996	5.651
The Dark Half	1989	5.654
The Tommyknockers	1987	5.654
Thinner	1984	5.657
Firestarter	1980	5.659
The Dark Tower III: The Waste Lands	1991	5.659
Cujo	1981	5.660
Under the Dome	2009	5.664
From A Buick 8	2002	5.666
The Stand	1978	5.666
The Dark Tower VII: The Dark Tower	2004	5.673
Gerald's Game	1992	5.682
Salem's Lot	1975	5.685
Night Shift	1978	5.687
The Shining	1977	5.690
Pet Sematary	1983	5.691
Four Past Midnight	1990	5.692
Sleeping Beauties	2017	5.694
The Talisman	1984	5.694
The Girl Who Loved Tom Gordon	1999	5.696
Needful Things	1991	5.698
Insomnia	1994	5.700
The Dark Tower: Wind Through the Keyhole	2012	5.700
Carrie	1974	5.702
Everything's Eventual: 14 Dark Tales	2002	5.702
Skeleton Crew	1985	5.703
Just After Sunset	2008	5.705
End of Watch	2016	5.712
The Outsider	2018	5.712
Nightmares and Dreamscapes	1993	5.713
The Dark Tower IV: Wizard and Glass	1997	5.718
Dolores Claiborne	1992	5.719
Christine	1983	5.723
Different Seasons	1982	5.726
The Dark Tower VI: Song of Susannah	2004	5.727
Black House	2001	5.728
Full Dark, No Stars	2010	5.728
Finders Keepers	2015	5.730
Cell	2006	5.732
The Dark Tower V: Wolves of the Calla	2003	5.738
The Institute	2019	5.738
Blaze	2007	5.742
Rose Madder	1995	5.749
Roadwork	1981	5.755
Billy Summers	2021	5.756
The Bazaar of Bad Dreams	2015	5.759
The Dead Zone	1979	5.767
Hearts in Atlantis	1999	5.768
Rage	1977	5.768
Lisey's Story	2006	5.771
Cycle of the Werewolf	1983	5.781
Mr. Mercedes	2014	5.787
Doctor Sleep	2013	5.788
Later	2021	5.797
The Colorado Kid	2005	5.798
Bag of Bones	1998	5.801
If It Bleeds	2020	5.803
11/22/1963	2011	5.810
Duma Key	2008	5.823
Revival	2014	5.840
Joyland	2013	5.891
Gwendy's Button Box	2017	5.924
Elevation	2018	5.970

Table 1: Happiness Scores. Stephen King’s novels and short story collections are included in this table, along with their year of publication, sorted by average happiness score.

The information from Table 1 was used to create a scatterplot shown in Figure 7. This is a scatterplot of happiness by year. The x axis is year while the y axis is happiness, and each point is a book. The scores stay between about 5.55 and 6. As described in Figure 1, this difference of about 0.5 is large when using the Hedonometer. The graph in Figure 7 also uses a trendline to show the relationship between year and average happiness. There looks like there is a slight positive relationship between year and average happiness.

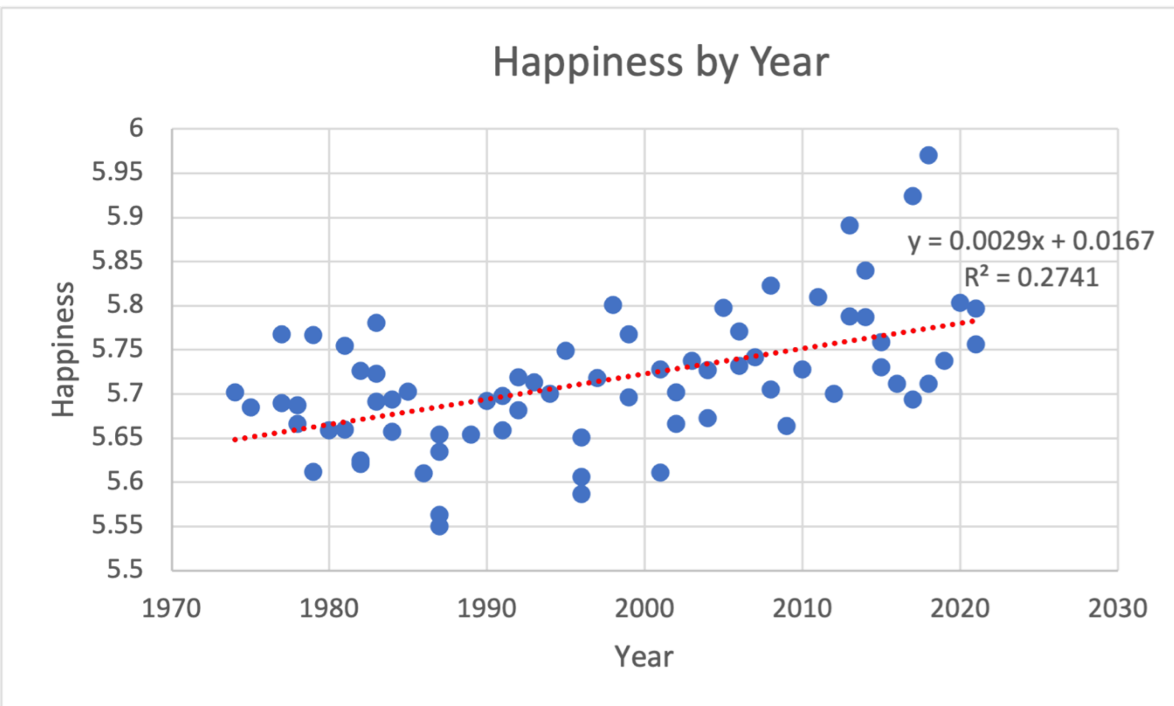

Figure 7: Scatterplot of happiness by year. Each dot represents a book and the linear trendline shows the relationship.

Figures 8 and 9 split the above scatterplot into two with Figure 8 including books published before 2000 and Figure 9 including books published after 2000. In 1999, King was out for a walk when he was hit by a van. He was extremely injured and spent three weeks in the hospital. This experience may have influenced his writing. We cannot be sure how many books were already written pre accident but published post 2000. Despite that, the two graphs seem to differ. In Figure 8, the red trendline is relatively flat and does not seem to indicate any relationship between year and happiness score. The r value or the correlation coefficient for this graph is 0.028 which means there is no relationship. Therefore, pre-2000, the year and the average happiness score did not have any relationship. However, Figure 9 suggests there is a relationship between year and average happiness score. The red trendline on this graph has an upward, positive slope and the r value for this line is 0.45. This is a low positive association but still proposes some relationship.

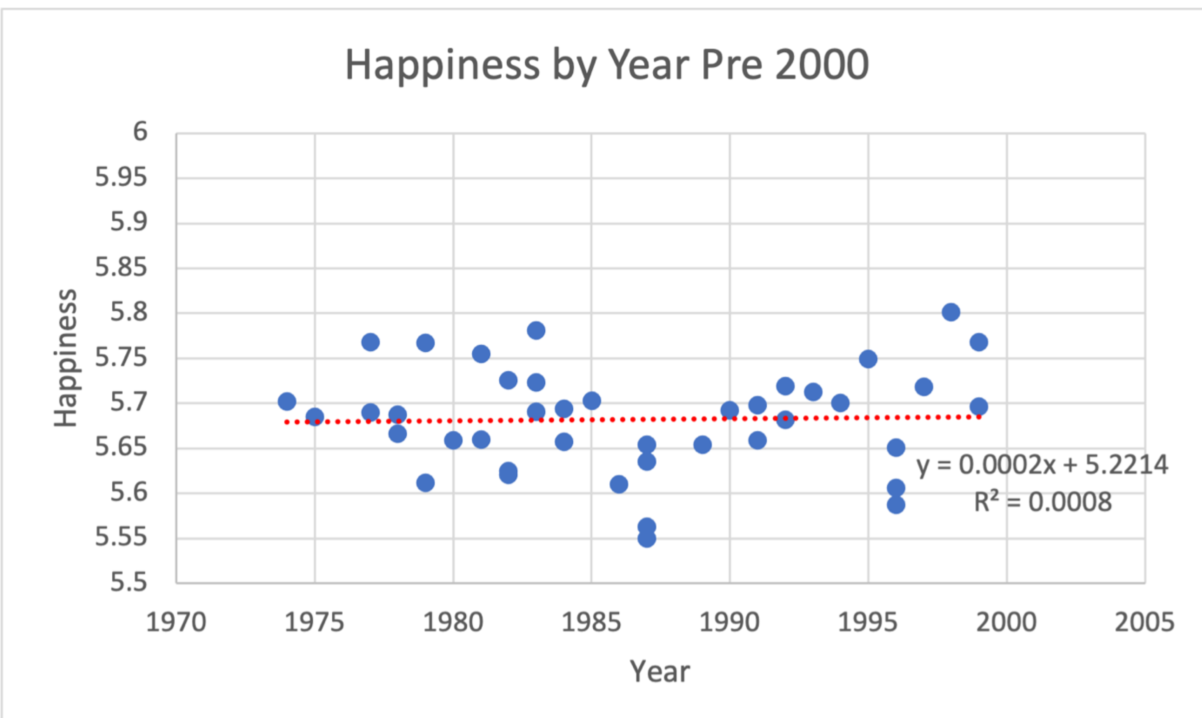

Figure 8: Scatterplot of Books Pre 2000. The data set was split, and books published before 2000 are included in this scatterplot.

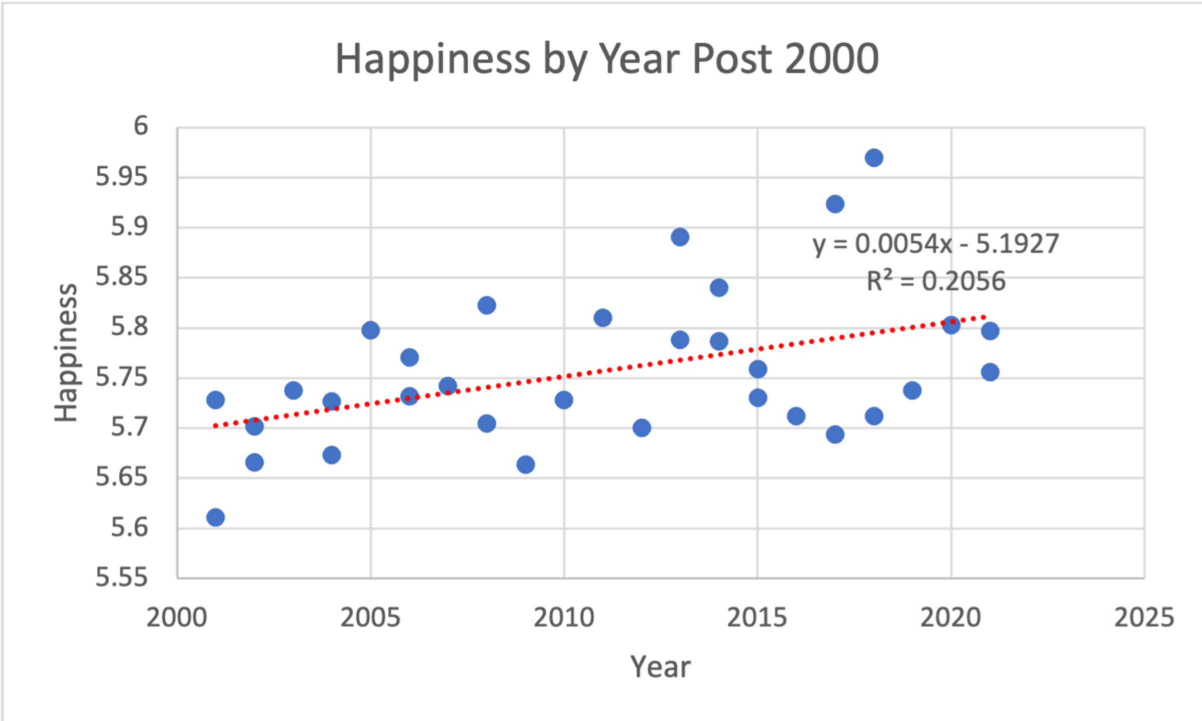

Figure 9: Scatterplot of Books Post 2000. The data set was split, and books published after 2000 are included in this scatterplot.

Figure 10 shows the time series graphs for all 62 Stephen King novels and the 4 short stories from his collection *Different Seasons* which include "Rita Hayworth and the Shawshank Redemption", "Apt Pupil", "The Body", and "The Breathing Method". The x axis is the percentage of the book, and the y axis is the happiness. The y axis for each graph is the same, ranging from 5.0 to 6.3. As noted from Figure 6, some graphs have flatter curves and less variation similar to *Elevation* like *Cycle of the Werewolf* and *Gwendy's Button Box.* Most of them have more dramatic highs and lows like *Dreamcatcher* with a low of 5.117 and *IT* with a low of 5.160. Some of the books with dramatic highs are *11/22/1963* with a high of 6.192 and *The Dark Tower 4: Wizard and Glass* with a high of 6.106.

Figures 11 and 12 follow the style of Figure 4 where the emotional arc is annotated. Included in these graphs are the highest point, the lowest point, and other important moments throughout the story. Figure 11 shows the emotional arc for the short story "The Body". The happiest moment in “The Body” is Gordie looking back on a happy memory he has of his brother. Gordie looked up to his brother, so this moment has lots of happy language. The saddest moment is when Teddy and the worker at the junkyard get into an argument about Teddy’s father and there is lots of negative language used during this scene. Figure 12 shows the emotional arc for the short story “Rita Hayworth and the Shawshank Redemption”. The happiest scene in this narrative is when Andy begins working at the library and remakes it into an impressive collection. The saddest scene is less defined but is when Red describes the way that Andy escaped. This scene may not appear sad to the reader since we are happy Andy escaped but the language used is very negative. This is an example of an instance where the human sentiment may not line up perfectly with the machine estimated sentiment.

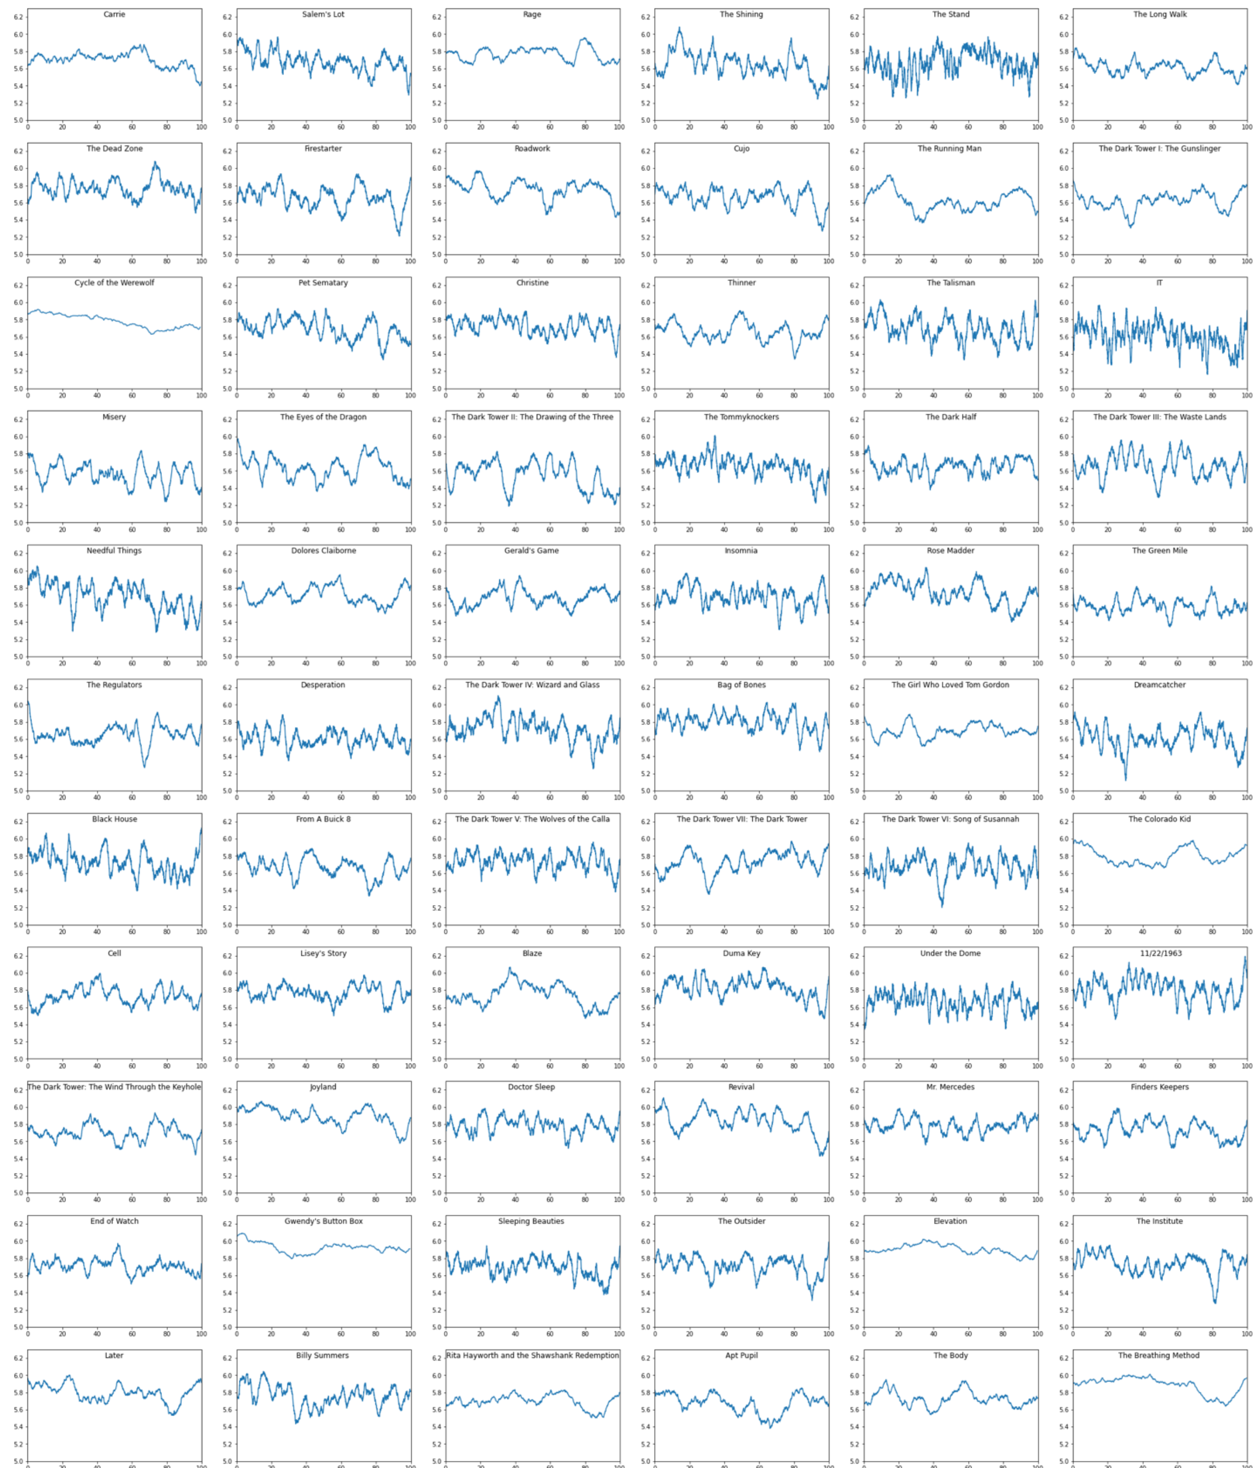

Figure 10: Timeseries Graphs for all 66 stories. The first 62 graphs show the emotional arc for all of King's novels. The last 4 are the 4 short stories in the short story collection *Different Seasons.*

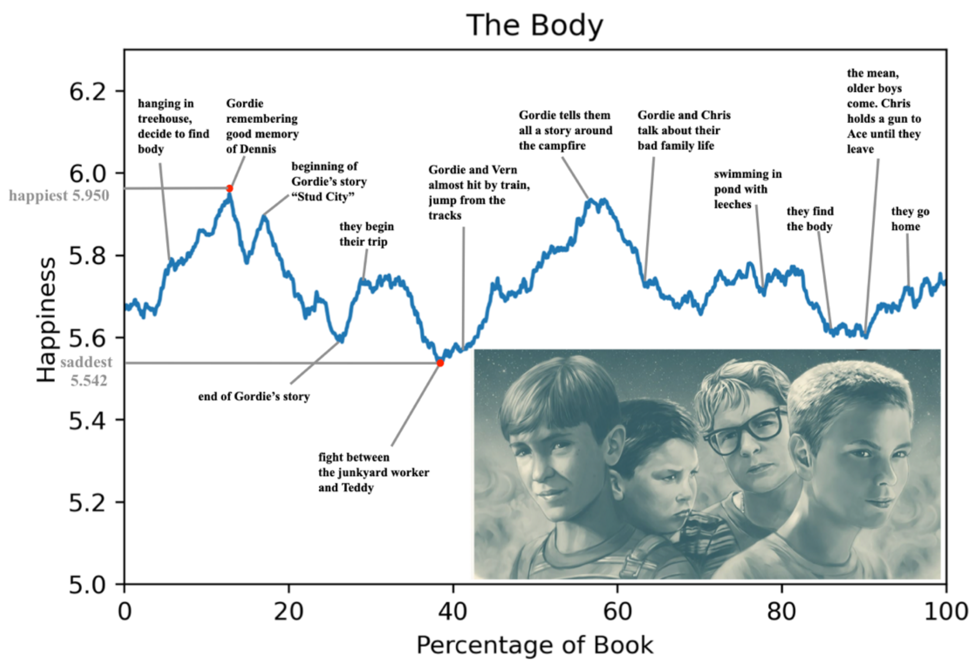

Figure 11: Emotional Arc for "The Body". The arc was labeled with the happiest point, the saddest point, and other important plot points.

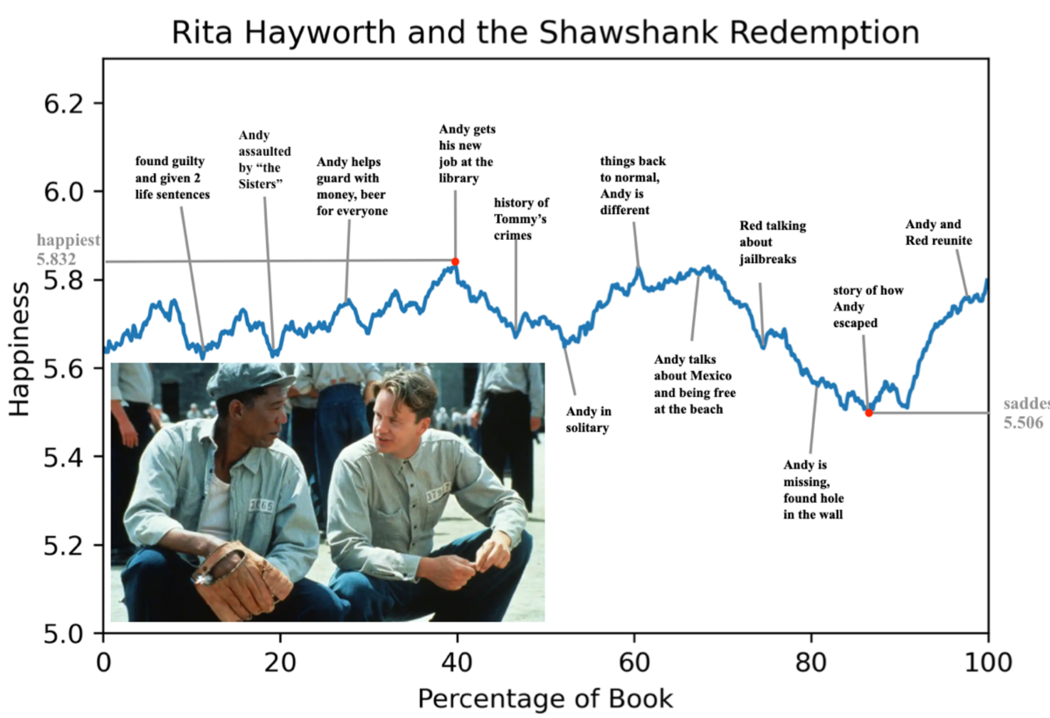

Figure 12: Emotional Arc for "Rita Hayworth and the Shawshank Redemption". The arc was labeled with the happiest point, the saddest point, and other important plot points.

Discussion

The graph in Figure 7 suggests that there may be a relationship between year and average happiness score. It appears that the later the year, which is the further into King's career, the happier the average score. The trendline appears to have a positive, upward slope and has an equation of $y = 0.0029x + 0.0167$. The r value or the correlation coefficient for this graph is 0.5235. This r value reveals there is a moderate positive association between year and average happiness score. Therefore, it is moderately associated that as year increases the happiness will also increase. Figure 9 also suggests a relationship. Interestingly, when only including books from 2000 on, the r value is 0.45 which is a less strong association. This correlation coefficient implies a low positive association.

Figure 5 also suggests that earlier books from his career are sadder than later books in his career or use sadder language. The middle word-shift from this figure compares the first 10 years of King's career (1974-1984) to the last 10 years (2011-2021). The average happiness score for the first 10 years is 5.72 while the average happiness score for the last 10 years is 5.84. The more recent years have a higher average score than the first 10 years with a difference of 0.12. The word-shift on the left of Figure 5 also relates to this. This word-shift compared *The Shining* (1977) to *Elevation* (2018). The average score for *The Shining* is 5.69 and the average score for *Elevation* is 5.97. *Elevation*'s score is slightly less than .3 larger than *The Shining*'s score. *Elevation* is the happiest book with this score.

From Table 1, we found that the average score of all Stephen King novels and short story collections is 5.715. If we take the happiest book, *Elevation* with a score of 5.97 and subtract the saddest book, *The Dark Tower II: The Drawing of the Three* with a score of 5.55, we get a difference of 0.42. There is a large difference between these two scores, so they are significantly different in happiness scores. Another important thing to notice from Table 1 is the four books with the lowest happiness score and the four books with the highest happiness score. The four books with the lowest happiness scores are *The Dark Tower II: The Drawing of the Three* (1987), *Misery* (1987), *The Green Mile* (1996), and *Desperation* (1996). These books are all published in the first half of King's career. The four books with the highest happiness scores are *Elevation* (2018), *Gwendy's Button Box* (2017), *Joyland* (2013), and *Revival* (2014). These books are all published in the second half of King's career. Therefore, it seems to suggest that King began to use happier language later in his career.

Overall, this work seems to imply that books from his earlier career are sadder than books in his later career or that books from his earlier career use sadder language. These results suggest King's work has increasingly shifted in genre from horror to science fiction. For example, *Elevation* (the book with the happiest language) follows a man named Scott who is losing weight every day for no apparent reason, no matter what he does. This plot is more thriller, suspense, or science fiction rather than horror. The saddest book, *The Dark Tower II: The Drawing of the Three,* is a dark fantasy. *Misery*, the second saddest book, is a psychological horror story about an author who is being held captive by a fan of his. A closer look at the genres of all the books may help in verifying the shift in genre.

Our findings would benefit from interpretation by a scholar more familiar with King's career trajectory, putting the computational analysis of his language into context. The work in this research looks at King's novels from a distance so a closer reading along with the findings from this research could lead to some more conclusions.

Conclusion

Sentiment analysis and Natural Language Processing present a new way to perform distant reading of stories. With careful attention, we may notice how we feel when we read novels, but that is not typically what people are looking for when they are analyzing a book. Critics look at themes, relationships between characters, and plot devices. Sentiment analysis allows the emotion to be quantified and investigated. Specifically, looking at Stephen King's books, we can see if there is an emotion trend that his books seem to follow. We can also compare findings of emotion with popularity of books. For example, *The Shining* is a very popular Stephen King book and there is a lot of negative language used in this novel. I did not include any research about sales of books in my thesis, but that could be a follow up project to see if popularity and average happiness score have any association.

Sentiment analysis allows us to analyze all of King's novels in seconds, complementing the qualitative, careful, close reading traditionally performed by critics. This overview allows us to understand patterns in King's career as a writer and see how the emotional language in his writing has changed over time. In a way, we are observing his work the way a telescope might observe distant stars and compare aspects of their nature (e.g., mass, energy, momentum). This gives us an understanding of King that we may not have been able to conclude in another way. Sentiment analysis allows us to quantify emotion more systematically.

As mentioned earlier in my thesis, Magistrale believes that Stephen King has changed in writing more toward a focus on the "domestic gothic". This may be reflected in the emotional language becoming happier overtime. It appears that books from his later career use happier language, so the domestic gothic may include less extreme negative words than the gory horror found in his earlier books.

As I mentioned in the discussion of the annotated emotional arc in Figure 12 for the short story "Rita Hayworth and the Shawshank Redemption", sometimes the machine estimated sentiment comes into contrast with what we feel as human beings experiencing the story. The machine estimated sentiment is solely based on word usage, and therefore without context the instrument can lead to different interpretations. As humans, we typically do not focus on individual words to ascertain how we feel at a point throughout a story; our emotional reactions depend more on the whole sentence/paragraph/scene. Therefore, there can be some misclassification with sentiment analysis, but this is not as common and typically the machine sentiment lines up with the human sentiment. One advantage of word-shift graphs when compared with more sophisticated word-embedding and machine learning techniques is that we can see the words responsible for observed changes, and therefore identify when context has become important.

Overall, the happiness score does seem to be a relatively good barometer of positivity. There are other dimensions that could be added into this research such as safety/danger that would continue to improve our models. However, the work here with just the happiness scale is a good first step on its own. The happiness index does not imply much about the cultural importance about a novel. If we think about *The Shining* again, the text is not very positive, yet it is very culturally important. In this case, since King is a horror author and readers expect scary books from him, a more negative text may actually be more successful and culturally important. My work did not go into this type of detail, but it would be interesting to consider. The research I've done with looking at emotional arcs does not inherently suggest any value about Stephen King's fiction.

Acknowledgments

I would like to thank my thesis advisor, Chris Danforth, for his support and help throughout this thesis. I would also like to thank my committee members Peter Dodds and Tony Magistrale for taking the time to be on my committee. Another big thank you to Mikaela Fudolig for all her coding help; this project would not exist without her help.

References

Alm, C. O., Roth, D., & Sproat, R. (2005). Emotions from text: Machine Learning for text-based emotion prediction. *Proceedings of Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing,* 579-586. ACL Anthology. https://aclanthology.org/H05-1073

Del Vecchio, M., Kharlamov, A., Parry, G., & Pogrebna, G. (2020). Improving productivity in Hollywood with data science: Using emotional arcs of movies to drive product and service innovation in entertainment industries. *Journal of the Operational Research Society*, 1–28. https://doi.org/10.1080/01605682.2019.1705194

Dennis-Henderson, A., Roughan, M., Mitchell, L., & Tuke, J. (n.d.). Life still goes on: Analysing Australian WW1 diaries through Distant Reading. https://aclanthology.org/2020.latechclfl-1.11.pdf

Dodds PS, Clark EM, Desu S, Frank MR, Reagan AJ, Williams JR, Mitchell L, Harris KD, Kloumann IM, Bagrow JP, Megerdoomian K, McMahon MT, Tivnan BF, Danforth CM (2015) Human language reveals a universal positivity bias https://doi.org/10.1073/pnas.1411678112

Dodds PS, Harris KD, Kloumann IM, Bliss CA, Danforth CM (2011) Temporal Patterns of Happiness and Information in a Global Social Network: Hedonometrics and Twitter. PLoS ONE 6(12): e26752. https://doi.org/10.1371/journal.pone.0026752

Elkins, K., & Chun, J. (2019). Can Sentiment Analysis Reveal Structure in a Plotless Novel? *ArXiv:1910.01441 [Cs]*. http://arxiv.org/abs/1910.01441

Gallagher, R. J., Frank, M. R., Mitchell, Lewis, Schwartz, A. J., Reagan, A. J., Danforth, C. M., Dodds, P. S. (2021). Generalized Word Shift Graphs: A Method for Visualizing and Explaining Pairwise Comparisons Between Texts. *EPJ Data Science*, 10(4).

Hipson, W. E., & Mohammad, S. M. (2021). Emotion Dynamics in Movie Dialogues. *ArXiv:2103.01345 [Cs]*. http://arxiv.org/abs/2103.01345

Jacobs, A. M., Herrmann, B., Lauer, G., Lüdtke, J., & Schroeder, S. (2020). *Sentiment analysis of children and youth literature: Is there a Pollyanna effect?* Frontiers. https://doi.org/10.3389/fpsyg.2020.574746

Jockers, M. L. (2014). *A Novel Method for Detecting Plot.* Matthew L Jockers. https://www.matthewjockers.net/2014/06/05/a-novel-method-for-detecting-plot/

Kim, E., Padó, S., & Klinger, R. (2017). Investigating the Relationship between Literary Genres and Emotional Plot Development. *Proceedings of the Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature*, 17–26. https://doi.org/10.18653/v1/W17-2203

Kleinberg, B., Mozes, M., & van der Vegt, I. (2018). Identifying the sentiment styles of YouTube's vloggers. *ArXiv:1808.09722 [Cs]*. http://arxiv.org/abs/1808.09722

Magistrale (2002) Chp. “Stephen King” in Thomson, D. H., Voller, J. G., & Frank, F. S. (2002). *Gothic writers: A critical and bibliographical guide*. Westport, CT: Greenwood Press.

McClure, David. (2017). Distributions of words across narrative time in 27,266 novels. https://litlab.stanford.edu/distributions-of-words-27k-novels/

Reagan, A. J., Mitchell, L., Kiley, D., Danforth, C. M., & Dodds, P. S. (2016). The emotional arcs of stories are dominated by six basic shapes. *EPJ Data Science*, *5*(1), 1–12. https://doi.org/10.1140/epjds/s13688-016-0093-1

Schmidt, T., & Burghardt, M. (2018). *An Evaluation of Lexicon-based Sentiment Analysis Techniques for the Plays of Gotthold Ephraim Lessing*. https://aclanthology.org/W18-4516/

Toubia, Olivier, Jonah Berger, and Jehoshua Eliashberg. “How Quantifying the Shape of Stories Predicts Their Success.” *Proceedings of the National Academy of Sciences* 118, no. 26 (June 29, 2021): e2011695118. https://doi.org/10.1073/pnas.2011695118.

Underwood, T. (2015). *Free research question about plot*. The Stone and the Shell. Retrieved March 28, 2022, from https://tedunderwood.com/2015/04/01/free-research-question-about-plot/

Vermont Complex Systems Center. (n.d.). *Average Happiness for Twitter*. Hedonometer. https://hedonometer.org/timeseries/en_all/?from=2020-09-23&to=2022-03-22

Vonnegut, K. (1995). Shapes of Stories. https://www.youtube.com/watch?v=oP3c1h8v2ZQ

Appendix

Name	Year	Score
Carrie	1974	5.702
Salem's Lot	1975	5.685
Rage	1977	5.768
The Shining	1977	5.690
Night Shift	1978	5.687
The Stand	1978	5.666
The Long Walk	1979	5.612
The Dead Zone	1979	5.767
Firestarter	1980	5.659
Roadwork	1981	5.755
Cujo	1981	5.660
The Running Man	1982	5.625
The Dark Tower I: The Gunslinger	1982	5.621
Different Seasons	1982	5.726
Cycle of the Werewolf	1983	5.781
Pet Sematary	1983	5.691
Christine	1983	5.723
Thinner	1984	5.657
The Talisman	1984	5.694
Skeleton Crew	1985	5.703
IT	1986	5.610
Misery	1987	5.563
The Eyes of the Dragon	1987	5.635
The Dark Tower II: The Drawing of the Three	1987	5.550
The Tommyknockers	1987	5.654
The Dark Half	1989	5.654
Four Past Midnight	1990	5.692
The Dark Tower III: The Waste Lands	1991	5.659
Needful Things	1991	5.698
Dolores Claiborne	1992	5.719
Gerald's Game	1992	5.682
Nightmares and Dreamscapes	1993	5.713
Insomnia	1994	5.700
Rose Madder	1995	5.749
The Green Mile	1996	5.587
The Regulators	1996	5.651
Desperation	1996	5.606
The Dark Tower IV: Wizard and Glass	1997	5.718
Bag of Bones	1998	5.801
The Girl Who Loved Tom Gordon	1999	5.696

Hearts in Atlantis	1999	5.768
Dreamcatcher	2001	5.611
Black House	2001	5.728
From A Buick 8	2002	5.666
Everything's Eventual: 14 Dark Tales	2002	5.702
The Dark Tower V: The Wolves of the Calla	2003	5.738
The Dark Tower VI: Song of Susannah	2004	5.727
The Dark Tower VII: The Dark Tower	2004	5.673
The Colorado Kid	2005	5.798
Cell	2006	5.732
Lisey's Story	2006	5.771
Blaze	2007	5.742
Just After Sunset	2008	5.728
Duma Key	2008	5.823
Under the Dome	2009	5.664
Full Dark, No Stars	2010	5.705
11/22/1963	2011	5.810
The Dark Tower: The Wind Through the Keyhole	2012	5.700
Joyland	2013	5.891
Doctor Sleep	2013	5.788
Revival	2014	5.840
Mr. Mercedes	2014	5.787
Finders Keepers	2015	5.730
The Bazaar of Bad Dreams	2015	5.759
End of Watch	2016	5.712
Gwendy's Button Box	2017	5.924
Sleeping Beauties	2017	5.694
The Outsider	2018	5.712
Elevation	2018	5.970
The Institute	2019	5.738
If It Bleeds	2020	5.803
Later	2021	5.797
Billy Summers	2021	5.756

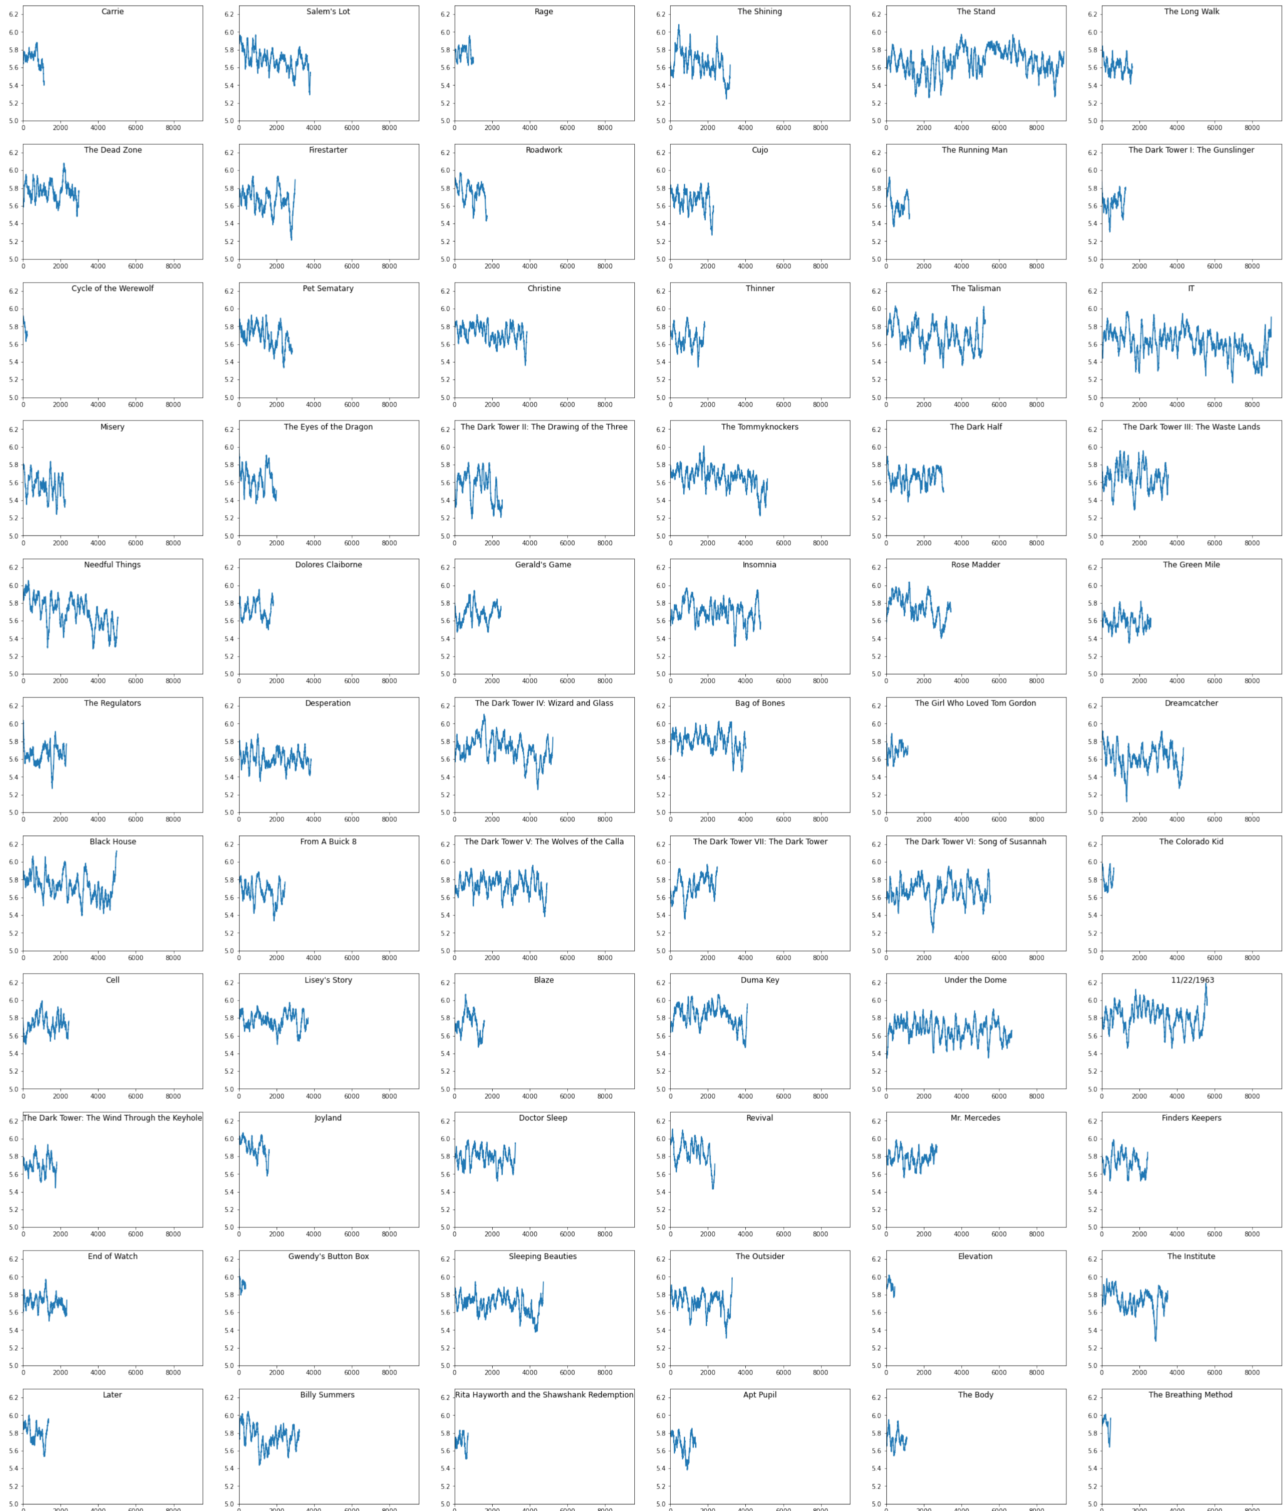
Carrie
Salem's Lot
Rage
The Shining
The Stand
The Long Walk
The Dead Zone
Firestarter
Roadwork
Cujo
The Running Man
The Dark Tower I: The Gunslinger
Cycle of the Werewolf
Pet Sematary
Christine
Thinner
The Talisman
IT
Misery
The Eyes of the Dragon
The Dark Tower II: The Drawing of the Three
The Tommyknockers
The Dark Half
The Dark Tower III: The Waste Lands
Needful Things
Dolores Claiborne
Gerald's Game
Insomnia
Rose Madder
The Green Mile
The Regulators
Desperation
The Dark Tower IV: Wizard and Glass
Bag of Bones
The Girl Who Loved Tom Gordon
Dreamcatcher
Black House
From A Buick 8
The Dark Tower V: The Wolves of the Calla
The Dark Tower VII: The Dark Tower
The Dark Tower VI: Song of Susannah
The Colorado Kid
Cell
Lisey's Story
Blaze
Duma Key
Under the Dome
11/22/1963
The Dark Tower: The Wind Through the Keyhole
Joyland
Doctor Sleep
Revival
Mr. Mercedes
Finders Keepers
End of Watch
Gwendy's Button Box
Sleeping Beauties
The Outsider
Elevation
The Institute
Later
Billy Summers
Rita Hayworth and the Shawshank Redemption
Apt Pupil
The Body
The Breathing Method