Authorship Attribution of English Poetry using Sentiment Analysis

Lubomir Ivanov
Iona University, New Rochelle, NY 10801
livanov@iona.edu

Felix Perez
Iona University, New Rochelle, NY 10801
felixalejandroperez533@gmail.com

Abstract

We present a basic methodology and share some interesting results from experiments on using sentiment analysis for authorship attribution of poetry. We demonstrate that sentiment analysis can be effectively used to determine the authorship of poetic works given a sufficiently large training corpus. We also share some promising preliminary results from sentiment-analysis-based attribution of non-poetry works. Most results compare well with traditional authorship attribution approaches. Moreover, adding sentiment analysis to a traditional-feature-based ensemble classifier improved the accuracy of attribution. The strengths and limitations of our methodology and directions for further research are outlined at the end of the paper.

Introduction

Authorship attribution is the task of determining, with a high degree of confidence, the authorship of an unattributed text or document. It is an important field not in literary and historical studies, but in many other areas such as forensics, political science, economics, etc. Automated authorship attribution is based on using machine learning models or ensembles trained to recognize a set of stylistic features used in unique ways by candidate authors. The selection of machine learning models and stylistic features determines, to a large extent, the accuracy of the generated attribution hypothesis. Multilayer perceptrons (MLP), support vector machines (with sequential minimal optimization) (SMO), random forests (RF), and logistic model trees (LMT) usually provide the highest attribution accuracy when paired with stylistic features such as function words, character- and word-n-grams, part-of-speech (PoS) tags, prepositions, suffixes, etc. For an overview of authorship attribution, the reader is referred to (Juola 2009, Stamatatos 2009 & 2016).

Several recent papers (Gaston et all, 2018, Ivanov 2023, Ivanov 2019, Ivanov, Aebig, and Meerman 2018, Ivanov 2016) have demonstrated that the attribution accuracy can be increased by augmenting the attribution model ensembles with non-traditional features such as prosody, linguistic inquiry and word count (LIWC), topic models, and abstractness/concreteness. These features complement the strength of traditional stylistic indicators by focusing on specific traits of authors' writing styles, such as their use of lexical stress, beats, pauses, the choice of abstract or concrete vocabulary, and reliance on literary prosodic techniques such as alliteration, assonance, and consonance.

The goal of the work presented in this paper is to explore the usefulness of sentiment analysis (SA) as another non-traditional stylistic feature for authorship attribution. All authors express emotions in their writings - to varying degrees. The question we attempt to address is whether these emotions – if correctly detected – can be used to differentiate authors writing styles. Much depends, of course, on the accuracy of the sentiment analysis methodology as well as on the text genre (poetry, fiction, etc.) and the personality traits of the individual author. In this paper, we focus specifically on the use of sentiment analysis in authorship attribution of poetry, where emotions are usually much more strongly expressed compared to, for example, in a business document.

This paper begins with a brief review of the essence of sentiment analysis and notes some earlier attempts to use SA for authorship attribution. Next, we describe our SA-based attribution methodology and discuss the results of a multitude of attribution experiments using the Gutenberg Poetry Corpus (Jacob 2018). The results are compared with those obtained using traditional stylistic features. We also present results from enhancing an authorship attribution ensemble classifier based on traditional features with SA-based attribution. The results from several more experiments with two non-poetry corpora are also presented and compared to the results from the poetry experiments. We share our conclusions drawn from the experiments and discuss some strengths and limitations of our methodology. Finally, we outline future research directions for a broader and more in-depth approach to SA-based authorship attribution.

Copyright © 2024 by the authors.
This open access article is published under the Creative Commons Attribution-NonCommercial 4.0 International License.

Sentiment Analysis

Overview of Sentiment Analysis

Sentiment Analysis is the exploration and development of methodologies for extracting, quantifying, and analyzing emotive information from various contexts such as text, video and/or audio. SA straddles the boundaries of natural language processing (NLP), computational linguistics, psychology and other disciplines, which study the expression of human emotions in creative works and human activities in general. Our focus is on text-based SA, whose primary objective is evaluating emotive content in written texts.

The most widely studied aspect of text-based SA is polarity detection – the classification of emotions as positive, negative or neutral using either a discrete or a continuous scale. The seminal work on the topic was done by Peter Turney (Turney, 2002) and Bo Pang et al (Pang et al, 2002). Since then, a wide variety of polarity detection methodologies have emerged: Rule-based methods use predefined rules, which assign a polarity score to every word or phrase in the text. The overall score is determined by adding the polarity scores of the individual text elements. The main drawback of rule-based methods is that word/phrase semantics and, therefore, the associated emotion is highly dependent on the context. Even with the use of cutting-edge word embedding methods, it is difficult to assign an appropriate sentiment value for all possible contexts. Additionally, complexities such as negations, diminishers, intensifiers, and the author's use of sarcasm are difficult to define in terms of all-encompassing rules. As an alternative to rule-based SA methods, a number of machine-learning (ML) based methodologies have evolved. Some of these are trained to assign sentiment scores through supervised learning from pre-labeled data. Unsupervised methods, on the other hand, cluster and rank text elements (or whole texts) according to their polarity based on specific inherent features of text elements, such as word frequencies. Hybrid methodologies, which combine rule-based and ML methods have emerged as well.

Besides polarity detection, other types of SA have been developed to evaluate subjectivity or emotion type (e.g., joy, sadness, anger, fear, etc.) of words, phrases, or texts. It is beyond the scope of this paper to describe the broad research in the field of sentiment analysis. An excellent review of current status of SA is presented in (Wankhade et al, 2022).

Sentiment Analysis Based Attribution Research

There have been several attempts to use SA for authorship attribution. In (Schneider, 2015), the accuracy of SA-based attribution is studied using K-Nearest-Neighbor (KNN) and Naive Bayes (NB) classifiers trained on sci-fi text corpora. Two-author and multi-author experiments were performed. The reported results are in the 40%-65% range for the two-author experiments and significantly lower (20%-35% range) in the multi-author experiments. Similar SA-based attribution results are reported by (Gaston et al, 2018) using multimodal machine learning applied to the CASIS-25 dataset. Yet another study (Biringa et al, 2020) used NB, Support Vector Machines (SVM) and RF classifiers trained on SA data extracted from a Victorian Era Authors dataset.

Attribution Methodology

As human beings, all author - consciously or unconsciously - express a variety of emotions in the texts they create. The primary question we are attempting to address is: "Do authors exhibit *sufficiently unique sentiment patterns* in their writing to determine authorship with a high degree of confidence?". The key phrase in the above statement is "sentiment patterns". There are many ways to define a sentiment pattern. Our initial approach defines it simply as the frequency of a particular sentiment polarity range in a given text. At the end of the paper, we briefly discuss alternative, more complex ways to define sentiment patterns, which may produce even stronger results.

Corpora Used

Our goal was to assess the effectiveness of SA in attribution of poetry, which is usually rich in emotive content.

The primary corpus used in most of our experiments is the Gutenberg Poetry Corpus (GPC) (Jacobs, 2018). It consists of approximately three million lines extracted from the much larger (Project Gutenberg) collection. GPC is distributed in JSON format where each line has an attached id that can be used to look up the author and the work from which the line was taken. We wrote software, which reconstructed the original texts from the individual poetry lines in GPC and then distributed the works into 500-line files.

To test our methodology on non-poetry texts, we used two additional corpora:

- The Reuters-RCV1 corpus (Lewis et al 2004, NIST) includes English-language news stories from the 1990s.
- The Corpus of English Novels (CEN) (De Smet, 2008) consists of English-language novels by twenty-five British and North American writers between 1881 and 1922.

Our Methodology

Our approach is based on computing the sentiment polarity value of each text element in every document/text. We define a text element as either a sentence or a paragraph. The polarity values of the text elements are extracted using two Python libraries, TextBlob (TextBlob 2013) and VADER (Valence Aware Dictionary for Sentiment Reasoning) (Hutto & Gilbert, 2014), which are highly regarded for sentiment analysis. We also used the FLAIR library (Akbik et al, 2019), but it did not significantly impact the results while increasing the computation time of the experiments.

TextBlob and VADER are pre-trained SA tools, which compute sentiment polarity values as floating-point numbers between -1 and 1, where 1 is a "very positive" and -1 is a "very negative" sentiment. Both libraries have strengths and weaknesses: VADER is geared towards social media texts and does well with sarcasm and irony. TextBlob is more general and, in addition to SA, offers a variety of other functions such as part-of-speech tagging, tokenization, parsing, etc. To leverage the strengths of TextBlob and VADER, we opted to define, for each text in the corpus, a single vector, which combines the polarity values computed by the two libraries. To form the texts' vectors, we define a set of polarity ranges in increments of 0.1 for each of the two libraries, i.e. {"TB: -1 to -0.9", "TB: -0.9 to -0.8", … "TB: 0.9 to 1", "V: -1 to -0.9", "V: -0.9 to -0.8", … "V: 0.9 to 1"}. We then perform two sets of experiments, - one with paragraph-level polarities and one with sentence-level polarities. In each experiment, we process the texts one at a time using TextBlob and VADER to compute the polarity value of every text element (paragraph or sentence) and then incrementing the element of the polarity vector in whose range the computed value falls. After processing all text elements of a document, all vector element counts are divided by the total number of text elements in the specific text. To break a text into sentences, we use the PySBD sentence boundary disambiguation package (Sadvilkar & Neumann, 2020). Paragraphs are determined simply by scanning for newline characters in the text. The computed vectors are stored in WEKA (Hall et al, 2009) format for training machine learning classifiers.

Results

Gutenberg Poetry Corpus Experiments

Most experiments were performed using the texts in the Gutenberg Poetry Corpus. We picked out all authors from GPC with at least 12 texts and prepared several 15-authors sets randomly selected from the created sub-corpus. Several authors appeared in multiple (but not all) sets. Similarly, we prepared several 10-authors and 7-authors sets. For each set of authors, we conducted two sets of experiments – one using paragraph polarity detection and one using sentence polarity detection. The WEKA files generated in these experiments were used to train a variety of machine learning classifiers using leave-one-out cross-validation. The strongest results were obtained using MLP, SMO, LMT, and RF classifiers, though RF performed worse than the other three classifiers in almost all experiments. Table 1 below lists the average and maximum accuracies obtained from all 15-, 10-, and 7-authors experiments. While the results were close, the sentence polarity experiments proved to provide a slightly higher accuracy in almost all 10- and 7-author experiments. Interestingly, all 15-authors paragraph-based experiments yielded slightly higher accuracies than their sentence-based counterparts. To test this further, we conducted an additional set of experiments with a single set of randomly selected 20 authors. The results, presented in Table 2, appear to confirm that the paragraph-based approach produces a slightly higher accuracy than the sentence-based approach when the number of candidate authors is large, but further testing is needed to understand the reasons for this anomaly.

Number of Authors / **Classifier/ Text-Element**	**15 Authors**		**10 Authors**		**7 Authors**	
	Avg %	**Max %**	**Avg %**	**Max %**	**Avg %**	**Max %**
SMO/Paragraphs	78.07	80.51	75.85	85.96	93.81	93.99
MLP/Paragraphs	77.48	79.72	75.59	85.46	94.93	96.24
LMT/Paragraphs	79.06	82.48	80.56	84.46	94.24	94.36
RF/Paragraphs	71.04	74.41	71.37	83.21	89.91	90.60
SMO/Sentences	77.98	80.12	84.11	85.96	95.73	95.86
MLP/Sentences	77.48	80.51	86.43	87.72	96.05	96.99
LMT/Sentences	77.49	79.92	85.93	86.72	96.40	96.57
RF/Sentences	74.01	77.17	79.72	84.46	92.02	92.86
Best Accuracy	*82.48*		*87.72*		*96.99*	

Table 1: Average and maximum accuracies from the GPC sentiment polarity authorship attribution experiments.

Number of Authors / **Classifier/ Text-Element**	**20 Authors**
SMO/Paragraphs	72.82 %
MLP/Paragraphs	73.46%
LMT/Paragraphs	72.98%
RF/Paragraphs	67.96%
SMO/Sentences	72.49%
MLP/Sentences	71.84%
LMT/Sentences	72.33%
RF/Sentences	69.90%
Best Accuracy	*73.46%*

Table 2: Average and maximum accuracies from the sentiment polarity 20-authors experiments.

Next, we wanted to compare SA-based authorship attribution with traditional-feature-based attribution on the same sets of authors. We used the most common and highly performing traditional stylistic features: character-2-grams (C2G), function words (FW), part-of-speech tags (PoS), prepositions (PREP), suffixes (SUF), first-word-in-a-sentence (FWiS), and vowel-initiated words (VIW). The results from the traditional features experiments are presented in Table 3 below. We have appended the average and maximum SA-based attribution results for comparison.

Num. of Authors / Features/Classifiers	15 Authors	10 Authors	7 Authors
C2G/SMO	87.88%	90.41%	88.35%
C2G/MLP	81.41%	84.96%	86.93%
FW/SMO	83.05%	88.49%	90.63%
FW/MLP	76.20%	82.80%	89.16%
PoS/SMO	87.84%	89.50%	93.75%
PoS/MLP	82.00%	85.44%	91.19%
PREP/SMO	74.04%	76.82%	80.23%
PREP/MLP	64.29%	69.93%	81.82%
SUF/SMO	83.04%	85.69%	88.19%
SUF/MLP	76.12%	79.93%	85.10%
FWiS/SMO	64.87%	73.75%	72.79%
FWiS/MLP	58.20%	62.60%	74.74%
VIW/SMO	83.19%	83.80%	84.21%
VIW/MLP	76.83%	82.20%	87.85%
SA Polarity Avg	*77.91%*	*81.41%*	*95.19%*
SA Polarity Max	*82.48%*	*87.72%*	*96.99%*

Table 3: A comparison of the accuracies from the traditional stylistic features experiments and the SA polarity experiments.

It is clear that sentiment polarity performs well on the Gutenberg Poetry Corpus in comparison with many traditional stylistic features, outperforming all traditional features in the 7-author experiments. Of course, more testing with other poetry corpora is needed before any conclusions can be reached, but the results of the exhaustive testing with the GPC corpus appear promising for using sentiment polarity as a stylistic feature in poetry attribution.

Sentiment-Based Weighted-Voting Attribution

Analogously to a human expert who uses a broad set of features to attribute a text, it has been demonstrated that a weighted voting ensemble methodology, which combines different learning methods with a variety of stylistic features through weighted voting, usually outperforms individual machine-model/stylistic-feature classifier pairs (Petrovic, 2014). The performance improvement is often most significant when testing a large number of candidate authors. We added sentiment polarity to a weighted average ensemble implemented based on the (Petrovic, 2014) methodology, and performed two experiments – one with the 20-author set and the other with one of the 15-author sets used in the stand-alone SA experiments. For the 15-author experiment, the ensemble, without sentiment polarity, produced a 97.7% accuracy. Adding sentiment polarity to the ensemble increased the accuracy to 98.1%. In the 20-author experiment, the accuracy without sentiment polarity was 94.80%. Adding sentiment polarity produced a slightly improved accuracy of 94.82%. These results were not unexpected: It has been demonstrated (Ivanov et al, 2018) that non-traditional features often serve to highlight a specific trait in the author's writing style, which the more traditional features fail to capture. In our case, a close examination of the individual author accuracies revealed that adding sentiment polarity improved the attribution accuracy for one specific author, which led to the improved overall accuracy. Clearly, the author in question exhibits a more emotional writing style, which is captured in the high sentiment polarity value. Interestingly, two other authors had very low sentiment polarity average accuracy values in the 15-author experiment (44% and 33% respectively). Yet, the traditional features had no problem recognizing the writing styles of these authors, yielding a 98% and a 92% average accuracy in the 15-author experiment. Since the SA values were under the 66% threshold of the ensemble, they were disregarded in computing the overall ensemble attribution accuracy. However, this is an indication that there are authors – even in poetry – who are less emotive in their writing and, therefore, present a challenge for our SA-based attribution methodology.

Non-Poetry Results

While our focus was on SA-based poetry attribution, we were curious how our methodology would perform on non-poetic texts. We set up a few small experiments using the CEN and Reuters corpora. Both corpora have not been pre-processed for attribution and contain text, which may slightly skew the results – titles, tables of contents, ellipses and chapter separators, etc. However, our goal was to get a general sense of whether the methodology works at all on non-poetry texts. The results are presented in Table 4 below.

The experiments with the CEN corpus indicated that SA may, indeed, be useful for author attribution of non-poetry texts. While not as strong as the results from our poetry attribution experiments, the CEN corpus attribution experiments indicated that a sufficiently high accuracy (as high as 80.4% and probably higher after pre-processing the corpus) can be achieved on works of fiction.

Classifier/ Text-Element \ Corpus	CEN	Reuters	
	7 Authors	7 Authors	10 Authors
SMO/Paragraphs	75.3%	37.7%	31.6%
MLP/Paragraphs	75.3%	45.1%	34.0%
LMT/Paragraphs	72.2%	40.6%	31.4%
RF/Paragraphs	71.1%	53.7%	44.2%
SMO/Sentences	78.4%	43.1%	38.4%
MLP/Sentences	80.4%	47.8%	38.2%
LMT/Sentences	71.1%	46.6%	38.2%
RF/Sentences	76.3%	58.3%	49%
Best Accuracy	*80.4%*	*58.3%*	*49%*

Table 4: Accuracy from non-poetry corpora experiments

The attribution accuracy in the Reuters corpus experiments was significantly lower. This is not surprising because the Reuters texts are short news stories, often reporting business-related facts and numerical data. The reduced emotive charge of the narrative in these stories is the most obvious explanation for the disparity between the CEN and Reuters corpus results. However, another significant reason is the structure of the training/testing texts: The CEN corpus contains, on average, about 12 texts per author. For the most part, these texts are full-length novels with hundreds of paragraphs and thousands of sentences. This allows the attribution methodology to form very precise sentiment polarity vectors, which appear to truly capture the individual writing styles of the authors. The Reuters corpus, on the other hand, includes fifty texts per author but each text is only a few short paragraphs long. Thus, there isn't a sufficient amount of SA data to accurately place the generated polarity vector within the vector space characterizing the specific author. As a result, the classifiers fail to learn the proper separation of the authors' vector space, leading to the lower observed accuracy in the experiments. This suggests that SA-based attribution may be useful primarily for attributing longer texts and less useful for short content.

The results we observed in the non-poetry experiments are mostly in agreement with the SA-based attribution results reported in the literature: (Schneider, 2015) reports accuracies in the 40%-65% range depending on the number of sentences (500 to 1500) and the number of authors used in the experiments. The reported accuracy is understandably higher in the two-author experiments and lower (as low as the 20th percentile) in the multi-author experiments. (Gaston et al, 2018) reports a 20-percentile SA attribution accuracy using MLP and SVM classifiers on the CASIS-25 dataset. The authors point out that the accuracy can be increased through feature selection but indicate that, by itself, SA frequency-based attribution has a limited application. Our experiments, however, indicate that SA can be a valuable addition to the set of stylistic features – by itself or in an ensemble.

The quality of the SA-based attribution depends on a number of factors including the text genre, number of texts in the training corpus, the individual text sizes, and the accuracy of the sentiment analysis performed by the chosen libraries. For our two selections, TextBlob and VADER, the current (as of 2023) average accuracy of sentiment analysis has been demonstrated to be in the range 75% to 85%. As the field of sentiment analysis matures and more accurate SA methods are implemented, the accuracy of the SA-based authorship attribution will likely increase as well.

Conclusion and Future Work

In this paper, we have presented an SA-based methodology for authorship attribution. We have demonstrated that sentiment polarity works well on poetry both as a stand-alone stylistic feature and in combination with traditional stylistic features in an ensemble classifier. We have also demonstrated that sentiment polarity may work well for genres other than poetry as long as the texts representing the candidate authors are sufficiently long.

Our next focus will be to study the usefulness of the sentiment polarity attribution methodology in other genres. We will "clean up" the CEN corpus and run a full set of experiments on it using different sized author subsets with both paragraph and sentence polarities. We will explore the use of additional libraries such as FLAIR and BERT (Turc et al, 2019), which also provide sentiment analysis tools. The methodology will be applied to other literary corpora as well, including a corpus of 18^{th} century American and British texts that we have previously used for attribution studies. This corpus should be particularly well-suited for SA-based attribution given the raw emotions expressed in the writings from the time of the American Revolution.

A different direction for SA-based attribution research is to explore the *sequences* of sentiment change in the writings of a given author. We would like to determine if the works of a particular author exhibit similar patterns of sentiment change and whether such patterns can be captured by an LSTM or another deep learning model.

Yet another research direction is to augment sentiment polarity with other aspects of sentiment such as emotion types (happiness, fear, anger, etc.) and sentiment subjectivity/objectivity. These features may provide additional information to help us capture even better the elusive notion of an author's writing style.

References

Akbik, A.; Bergmann, T.; Blythe, D.; Rasul, K.; Schweter, S.; Vollgraf, R. 2019. “FLAIR: An easy-to-use framework for state-of-the-art NLP”, 2019 Annual Conference of the North American Chapter of the Association for Computational Linguistics (NAACL), pages 54-59.

Biringa, C.; Anirudh, R. S.; Richard, D.; Vikas, J. 2020. Text Mining and Authorship Attribution. Technical Report, University of Massachusetts, Dartmouth.

De Smet, H. 2008. Corpus of English novels (last accessed January 22, 2024): https://perswww.kuleuven.be/~u0044428/

Gaston, J., Narayanan, M., Dozier, G. & Cothran, D. L., Arms-Chavez, C., Rossi, M., King, M., Xu, J. 2018. Authorship Attribution via Evolutionary Hybridization of Sentiment Analysis, LIWC, and Topic Modeling Features. 933-940. 10.1109/SSCI.2018.8628647.

Hall M., Frank E., Holmes G., Pfahringer B., Reutemann P., Witten I., 2009. *The WEKA Data Mining Software: An Update; SIGKDD Explorations*, Volume 11, Issue 1

Hutto, C.J.; Gilbert, E.E. 2014. VADER: A Parsimonious Rule-based Model for Sentiment Analysis of Social Media Text. Eighth International Conference on Weblogs and Social Media (ICWSM-14). Ann Arbor, MI, June 2014.

Ivanov, L., “A Multi-Dictionary Approach to Abstractness/Concreteness-Based Authorship Attribution “, Proc. of FLAIRS-36, Clearwater Beach, FL, 5/2023

Ivanov, L. 2019. “Learning Patterns of Assonance for Authorship Attribution of Historical Texts”, FLAIRS-32, Sarasota, FL, USA, pp. 191-196.

Ivanov, L.; Aebig A.; Meerman S. 2018. “Lexical Stress-Based Authorship Attribution with Accurate Pronunciation Patterns Selection”, TSD 2018, Brno, Czech Rep., pp. 67-75.

Ivanov, L. 2016. “Using Alliteration in Authorship Attribution of Historical Texts, Text, Speech, and Dialogue”. TSD 2016. Lecture Notes in Computer Science, vol. 9924. Springer

Ivanov, L.; Petrovic, S. 2015. “Using Lexical Stress in Authorship Attribution of Historical Texts”, Chapter, LNCS: TSD, v.9302, pp.105-113

Jacobs, A. 2018. “The Gutenberg English Poetry Corpus: Exemplary Quantitative Narrative Analyses”. Frontiers in Digital Humanities. 5. 5. 10.3389/fdigh.2018.00005.

Juola P. 2009. “JGAAP: A System for Comparative Evaluation of Authorship Attribution”. Journal of Chicago Colloquium on Digital Humanities and Comp. Science, 1(1): 1-5.

Lewis, D. D.; Yang, Y.; Rose, T.; and Li, F. 2004. RCV1: A New Benchmark Collection for Text Categorization Research. *Journal of Machine Learning Research*, 5:361-397

Martins, R.; Almeida, J.; Henriques, P.; Novais, P. 2018. “Increasing authorship identification through emotional analysis”, Trends and advances in information systems and technologies, pp. 763–772.

Martins, R.; Almeida, J.; Henriques, P.; Novais, P. 2019. ”A sentiment analysis approach to improve authorship identification”, Expert Systems, https://doi.org/10.1111/exsy.12469

Narayanan M. *et al*., "Adversarial Authorship, Sentiment Analysis, and the AuthorWeb Zoo," *2018 IEEE Symposium Series on Computational Intelligence (SSCI)*, Bangalore, India, 2018, pp. 928-932, doi: 10.1109/SSCI.2018.8628806.

Pang, B.; Lee, L. 2005. "Seeing stars: Exploiting class relationships for sentiment categorization with respect to rating scales". Proceedings of the Association for Computational Linguistics (ACL). pp. 115–124.

Pang, B.; Lee, L.; Vaithyanathan, S. 2002. "Thumbs up? Sentiment Classification using Machine Learning Techniques". Proceedings of the Conf. on Empirical Methods in Natural Language Processing (EMNLP). pp. 79–86.

Petrovic, S.; Berton, G.; Schiaffino, R.; Ivanov, L. 2014. “Authorship Attribution of Thomas Paine Works”. International Conference on Data Mining DMIN’14. pp. 182-188. Springer

Project Gutenberg. (n.d.). (last accessed January 20, 2024), from www.gutenberg.org.

Schneider, M. 2015. ”A Study on the Efficacy of Sentiment Analysis in Author Attribution”. Electronic Theses and Dissertations. Paper 2538.

Sadvilkar, N.; Neumann, M. 2020. “PySBD: Pragmatic Sentence Boundary Disambiguation”, 2nd Workshop for Natural Language Processing Open-Source Software (NLP-OSS) at EMNLP, https://doi.org/10.48550/arXiv.2010.09657

Stamatatos E. 2016. “Authorship Verification: A Review of Recent Advances”, Research in Computer Science, 123, pp.9-25, IPN.

Stamatatos E. 2009. “A Survey of Modern Authorship Attribution Methods”, Journal of the American Society for Information Science and Technology, 60(3), pp. 538-556, Wiley

TextBlob. 2013. (last accessed January 19, 2024): https://textblob.readthedocs.io/en/dev/

Turc, I.; Chang, M.W.; Lee, K.; Toutanova, K. 2019. “Well-Read Students Learn Better: On the Importance of Pre-training Compact Models”, arXiv preprint arXiv:1908.08962v2

Turney, P. 2002. "Thumbs Up or Thumbs Down? Semantic Orientation Applied to Unsupervised Classification of Reviews". Proceedings of the Association for Computational Linguistics. pp. 417–424. *arXiv:cs.LG/0212032*.

Wankhade, M.; Rao, A.C.S.; Kulkarni, C. A. 2022. “Survey on sentiment analysis methods, applications, and challenges”. *AI Rev* 55, 5731–5780, https://doi.org/10.1007/s10462-022-10144-1