Style as Signature: Profile-Based Authorship Verification of Mihai Eminescu's Journalistic Corpus

Ioana-Roxana Boriceanu[1], **Liviu P. Dinu**[1,2]
[1]Faculty of Mathematics and Computer Science
[2]Human Language Technologies Research Center
University of Bucharest, Romania
ioana-roxana.boriceanu@s.unibuc.ro, ldinu@fmi.unibuc.ro

Abstract

Authorship verification aims to assess whether a questioned text is stylistically compatible with an author's known writings, a task that is particularly challenging in historical corpora with partial ground truth. We address this problem in the context of Mihai Eminescu's journalistic corpus, a historically grounded collection comprising published articles, manuscripts, and texts of uncertain authorship. Using a profile-based framework with character n-grams and function words, we examine how stylistic compatibility behaves across different profile construction settings and temporal splits. The results show that character trigram profiles consistently accept verified texts while producing a small and stable set of rejections among disputed items, whereas function word profiles show near complete acceptance across the corpus. A qualitative analysis shows that rejected texts exhibit meaningful differences in discourse structure and communicative purpose. These findings illustrate how authorship verification can support literary scholarship through stable signals for close reading.

1 Introduction

Mihai Eminescu (1850-1889) is widely regarded as the most influential poet of Romania and one of the central figures of its cultural history. Often described as a literary genius, Eminescu played a decisive role in shaping modern Romanian poetry and literary language. He is best known for the originality of his poetic work, which is commonly associated with European Romanticism and shows strong influences from German literature and philosophy. Eminescu's literary work addresses themes such as love, nature, time, history, and the human condition, often combining lyrical expression with philosophical reflection. Drawing on folklore, mythology, and classical sources, his poetry has had a lasting influence on Romanian literature and remains central to literary scholarship (Iorga, 1969).

Eminescu's career spanned several distinct stages, ranging from early work in theatre as a prompter, to periods of study in Vienna and Berlin, and later to administrative and editorial positions such as library director and journalist. These experiences shaped both his intellectual outlook and his writing. Beyond his literary activity, Eminescu also played a significant role in the development of Romanian journalism. His journalistic work, produced between 1870 and 1889, forms a large and influential corpus that reflects both the intellectual climate of the period and his distinctive authorial voice. His journalistic texts address a wide range of topics, including politics, education, culture, economics, and social reform, offering a sustained commentary on the tensions produced by the rapid modernisation of society (Mocanu, 2020).

An important aspect of this journalistic corpus is that a substantial portion of the articles were published without signature, while others appeared under pseudonyms such as *Fantasio* or *Varro*. Eminescu rarely signed his work explicitly, and for most of his journalistic career he avoided attaching his name to published texts (Mocanu, 2020). Although editorial scholarship has clarified the authorship of many articles, a small group of texts remains of disputed authorship, making this body of journalistic writing a particularly suitable case for authorship verification (AV).

Authorship verification (AV) (Stamatatos, 2016) is a task in natural language processing that aims to determine whether a given text was written by a particular author. In contrast to authorship attribution (AA) (Stamatatos, 2009), which assigns a text to one author from a predefined set of candidates, AV examines whether a questioned document is stylistically compatible with an author's known body of work. This problem arises in a variety of contexts, including the resolution of literary authorship disputes (Juola, 2013), the authentication of historical manuscripts (Tuccinardi, 2017), and forensic text

Proceedings of SIGHUM 2026, pages 102–110

analysis (Juola, 2021).

In the case of Mihai Eminescu, his writings have been extensively studied from a literary perspective, with particular attention to their cultural, political, and ideological context. His writing style has also been the subject of early quantitative and statistical analyses, primarily focused on his poetic corpus (Marcus, 1970; Seche, 1974; Irimia, 2009). By contrast, systematic studies of authorship and stylistic variation in his journalistic corpus have received limited attention. To address this gap, we apply authorship verification techniques to Mihai Eminescu's journalistic corpus in order to support literary scholarship in the analysis of texts of uncertain provenance. Our contributions can be summarised as follows:

- We present a detailed case study of profile-based authorship verification on a historically grounded Romanian journalistic corpus with verified and uncertain texts.
- We show that rank-based character trigram profiles achieve a stable balance between acceptance of verified texts while retaining structured variation among disputed items, whereas function word profiles tend to over-accept in this setting.
- We demonstrate that verification rejections are interpretable and motivate closer reading of differences in discourse structure and communicative purpose, supporting the use of profile-based authorship verification as an exploratory analytical tool and not only as a purely classificatory mechanism.

2 Related Work

A variety of approaches have been explored for authorship verification (AV). Early work includes the unmasking method (Koppel et al., 2007), which assesses authorship by progressively removing discriminative features and observing changes in classification performance. Another line of research relies on profile-based methods (Potha and Stamatatos, 2014), where an author's known texts are aggregated into a single stylometric representation and compared with questioned documents using distance based measures. Compression based methods (Halvani et al., 2017) avoid explicit feature engineering and rely on similarity scores derived from text compression. Another widely used approach is the impostors method (Koppel and Winter, 2014), which evaluates authorship by contrasting a questioned text against large sets of distractor documents using randomly selected feature subsets.

More recent studies have employed language model based techniques for authorship verification (Kim et al., 2025; Hung et al., 2023). While these approaches often achieve strong empirical performance, their reliance on opaque representations makes them less suitable for settings where interpretability and transparency are required, such as literary scholarship and forensic analysis. In contrast, we focus on a transparent profile-based setting where decisions are based on explicit ranked feature profiles and where outcomes can be interpreted in relation to textual structure.

3 Data and Methods

3.1 Mihai Eminescu's feuilletons

Our dataset consists of Mihai Eminescu's journalistic writings as published in the *Opere* edition of the Romanian Academy (1989). This corpus is organized into five volumes, each corresponding to a distinct period of his journalistic activity and containing three categories of texts: (1) published articles (PUB) attributed to Eminescu in contemporary newspapers, (2) manuscript texts (MAN) preserved in his own handwriting, and (3) texts of disputed authorship, referred to as dubia (DUB) in this work. The distribution of texts across volumes is summarized in Table 1.

Volume	PUB	MAN	DUB
Vol. 9 (1870-1877)	510	52	12
Vol. 10 (1877-1878)	264	14	44
Vol. 11 (1880)	323	11	7
Vol. 12 (1881)	356	31	14
Vol. 13 (1882-1883, 1888-1889)	242	96	26
Total	**1695**	**204**	**103**

Table 1: Distribution of texts across the five journalistic volumes

The corpus exhibits substantial variation in document length, with texts ranging from 10 to 12,767 words and a median length of 536 words. Manuscript texts tend to be shorter, with a median of 119 words, whereas published articles and texts of disputed authorship are generally longer. During preprocessing, we removed metadata such as headers, editorial notes, and footnotes where present.

We applied standard text normalization, including lowercasing and the removal of artefacts introduced by optical character recognition (OCR)

in the transcription of manuscript materials. In particular, lowercasing was applied to reduce sparsity and to avoid mixing up stylistic signals with editorial or contextual variation in capitalization, particularly for proper names, titles, and institutional references, which are unevenly distributed across journalistic texts. Original punctuation and spacing were preserved, as these features are relevant for the character n-gram representations used in our experiments.

3.2 Methods

We adopt a profile-based authorship verification framework inspired by Potha and Stamatatos, 2014 and integrate the rank distance proposed by Dinu, 2003, which has been shown to be effective in computational stylistics (Popescu and Dinu, 2008; Dinu et al., 2012). Texts with verified authorship, including both published articles and manuscript materials, are used to construct the author profile, while texts of disputed authorship are treated as the evaluation set. Each questioned text is represented in the same feature space as the profile, and authorship is assessed by computing its distance from the author profile.

In this study, we treat style operationally as recurring patterns in low-level textual features, following standard practice in stylometry. We consider two types of features: function words (FW) and character n-grams (CNG), both extracted directly from the normalized text. Character n-grams are widely used in stylometric analysis because they capture local orthographic and morphological regularities while remaining robust for short and medium-length documents. Function words are generally assumed to be processed non-consciously by authors and are therefore well suited for capturing stable stylistic patterns that are less influenced by topic or content (Chung and Pennebaker, 2011).

To measure stylistic similarity, we primarily employ the rank-based distance proposed in Dinu, 2003:

$$D_{\text{rank}}(P,Q) = \sum_{i=1}^{n} \left| r_P(f_i) - r_Q(f_i) \right|, \quad (1)$$

where f_i denotes the features under comparison, and $r_P(f_i)$ and $r_Q(f_i)$ represent their ranks in the author profile P and the questioned text Q, respectively. In cases of equal feature frequencies, tied features are assigned the average of the corresponding rank positions, following the standard formulation of rank distance.

For comparison, we also evaluate the asymmetric frequency-based CNG dissimilarity function introduced in Potha and Stamatatos, 2014, which serves as a strong profile-based baseline in our experiments:

$$D_{\text{freq}}(Q,P) = \sum_{g \in Q} \left(\frac{2\big(f_Q(g) - f_P(g)\big)}{f_Q(g) + f_P(g)} \right)^2, \quad (2)$$

where Q denotes the questioned text, P the author profile, and $f_Q(g)$ and $f_P(g)$ are the normalised frequencies of the character n-gram g in Q and P, respectively.

For the CNG representation, we retain the top L_{AUTHOR} trigrams as the author vocabulary. For each document, we also extract its top L_{DOC} trigrams. In our experiments, we set $L_{\text{AUTHOR}} = 1500$ and $L_{\text{DOC}} = 1000$, corresponding to the number of most frequent character trigrams retained for the author profile and for each questioned document, respectively.

The function word list used in our experiments, shown in Table 2, includes frequent words that mainly serve grammatical and discursive functions. These include conjunctions, prepositions, auxiliary verbs, pronouns, determiners, and discourse particles. The selection is grounded in prior work on Romanian computational stylistics and authorship analysis (Dinu et al., 2012) and was slightly adjusted using word-frequency analyses of Eminescu's literary corpus (Seche, 1974), as well as the most frequent items observed in the constructed author profiles.

și în să se cu o la nu a ce mai din pe un că ca mă fi care era lui fără ne pentru el ar dar îl tot am mi însă într cum când toate al după până decât ei nici numai dacă eu avea fost le sau spre unde unei atunci mea prin ai atât au chiar cine iar noi sunt acum ale are asta cel fie fiind peste această cele face fiecare nimeni încă între aceasta aceea acest aceasta acestei avut ceea cât da făcut noastră poate acestui alte celor cineva către lor unui altă ați dintre doar foarte unor vă aceste astfel avem aveți cei ci deci este suntem va vom vor de cari

Table 2: The function words used in our experiments

Distances are computed by comparing the rank distributions of the author profile and the questioned document over a shared feature space defined as the union of their respective vocabularies. For each text of uncertain authorship, we compute its distance to the author profile using both feature types and both distance formulations. The acceptance threshold is derived from the distribution of

distances obtained for texts with verified authorship. We define the threshold as

$$T = \mu + k\sigma, \tag{3}$$

where μ and σ denote the mean and standard deviation of the distances for genuine texts. We set $k = 2$ as a conservative choice that retains the vast majority of verified texts under approximately normal distance variation, while still allowing informative rejections among uncertain items. A questioned text is considered stylistically compatible with the author profile if its distance satisfies $d \leq T$. We deliberately avoid optimizing the threshold for maximal separation, as our objective is not classification performance but the identification of stable and interpretable deviations within an author's stylistic space.

Our analysis also considers continuous distance distributions across splits and volumes, and we evaluate threshold stability under resampling. In this sense, the threshold does not define a sharp stylistic boundary, but provides an interpretable reference point for identifying consistent outliers.

Within this framework, the evaluation focuses on internal stylistic structure in a historically grounded corpus under partial ground truth. The primary analyses are therefore restricted to verified and disputed Eminescu texts, using verified material to estimate the expected range of authorial variation. This design targets the stability and interpretability of deviations inside a single authorial space.

4 Experiments and Results

4.1 Experiments

The experiments are designed to answer two questions: whether verification behaviour is stable across reasonable parameter choices, and whether observed rejections correspond to interpretable textual differences. We evaluate two profile construction settings. In the first setting, the author profile is constructed from all published texts (PUB), while thresholds are estimated from the manuscript texts (MAN), which provide an independent set of verified writings. In the second setting, the profile is constructed from a random subset of the published texts (1356 documents, 80%) together with all manuscript texts, and thresholds are estimated using the held-out published texts (the remaining 20%, 339 documents). This second setting serves as a robustness check in which profile composition and threshold estimation are varied using disjoint subsets of the published texts.

Our primary distance measure is the rank-based distance (Dinu, 2003), which has been used in computational stylistics as an interpretable method for comparing ranked feature profiles. For comparison, we also report results obtained with the asymmetric frequency-based CNG dissimilarity function proposed by Potha and Stamatatos, 2014. This allows us to assess whether verification behaviour depends on the specific distance formulation or remains consistent across different measures.

Parameter choices for CNG based profiles

We investigate character n-gram representations with $n \in \{3, 4, 5\}$ under varying profile size configurations. In this context, profile size refers to the number of distinct n-gram types retained after sorting by frequency. For each setting, we construct an author profile from the published texts (PUB) and evaluate acceptance rates for manuscript texts (MAN) and dubia (DUB) using the rank-based distance. Parameter selection is guided by two criteria: high acceptance of verified manuscript texts and avoidance of configurations that yield uniform acceptance across text categories, including those of uncertain authorship. In practice, we treat MAN acceptance rates above 95% as acceptable.

Detailed results for all parameter combinations are reported in Table 3. Each row corresponds to a specific combination of author profile size (L_{AUTHOR}) and questioned document profile size (L_{DOC}), while columns report results for character 3-grams, 4-grams, and 5-grams, respectively. Values are reported as pairs of percentages in the form (MAN%, DUB%), indicating the proportion of texts classified as stylistically compatible with the Eminescu profile under the corresponding setting.

Across the parameter grid, character trigram profiles change gradually with profile size and yield a stable set of rejected texts across a wide range of profile sizes. While larger n-gram (4-grams and 5-grams) representations can produce stronger separation between MAN and DUB texts, they also alter the composition of the rejected set, leading to rejections that are less consistent across parameter settings. In the present study, we prioritise stability and interpretability of verification outcomes over maximal discrimination. Trigram representations maintain high acceptance of verified manuscript texts while producing a limited and interpretable

L_{AUTHOR}	L_{DOC}	3-grams	4-grams	5-grams
1000	800	(95.6, 91.3)	(99.0, 97.1)	(100.0, 100.0)
1000	1000	(96.6, 93.2)	(98.5, 95.1)	(100.0, 100.0)
1500	800	(98.0, 98.1)	(97.5, 95.1)	(99.5, 99.0)
1500	1000	(96.6, 94.2)	(97.1, 94.2)	(99.5, 99.0)
1500	1500	(92.6, 85.4)	(97.5, 91.3)	(98.5, 98.1)
2000	800	(99.0, 100.0)	(97.1, 93.2)	(99.5, 99.0)
2000	1000	(98.5, 98.1)	(96.6, 94.2)	(99.5, 98.1)
2000	1500	(97.5, 93.2)	(95.6, 89.3)	(98.5, 95.1)
2000	2000	(95.6, 92.2)	(95.6, 84.5)	(95.6, 92.2)
3000	800	(100.0, 100.0)	(97.1, 93.2)	(99.0, 97.1)
3000	1000	(100.0, 100.0)	(95.6, 92.2)	(99.0, 98.1)
3000	1500	(100.0, 100.0)	(95.6, 89.3)	(98.0, 95.1)
3000	2000	(100.0, 100.0)	(94.1, 83.5)	(96.6, 91.3)
3000	3000	(100.0, 100.0)	(91.7, 77.7)	(91.7, 79.6)

Table 3: Acceptance rates (MAN%, DUB%) across n-gram lengths and profile sizes.

set of rejections among uncertain items. For this reason, we adopt character trigrams for the remaining experiments.

For each value of n, we vary the number of retained character n-grams for the author profile ($L_{\text{AUTHOR}} \in \{1000, 1500, 2000, 3000\}$) and for each questioned document ($L_{\text{DOC}} \in \{800, 1000, 1500, 2000, 3000\}$), subject to $L_{\text{DOC}} \leq L_{\text{AUTHOR}}$. As a compromise between robustness and discriminative capacity, we select $L_{\text{AUTHOR}} = 1500$ and $L_{\text{DOC}} = 1000$ for the trigram representation. This configuration satisfies the MAN acceptance criterion while preserving variation in DUB acceptance and yields a stable and interpretable set of rejected texts, making it suitable for the remaining experiments.

Stability analyses

We evaluate the robustness of the proposed verification framework by testing whether its behaviour remains stable under variation in data splits, temporal coverage, and threshold estimation. These analyses are intended to assess whether verification outcomes reflect consistent stylistic properties of the corpus and not artefacts of a particular sampling or configuration.

We first perform five-fold cross-validation on the manuscript set (MAN). Manuscript texts provide an independently verified set of writings that are stylistically close to the published articles, while remaining separate from the texts used to construct the author profile in our primary setting. In each fold, the acceptance threshold is estimated from four fifths of the MAN texts and evaluated on the held-out subset. Acceptance remains consistently high for both feature types (CNG: 97.6–100%; FW: 95.0–100%), indicating that the manuscript texts form a stylistically coherent set under the proposed distance based criteria.

Next, we examine the distribution of distances across journalistic volumes and text categories (PUB, MAN, DUB) to assess potential period specific effects. For each split and volume specific subset, we compute descriptive statistics including mean, standard deviation, and range. For both CNG and FW representations, distance distributions are comparable across volumes, with no abrupt shifts that would suggest strong temporal drift within the verified material.

To further probe temporal sensitivity, we construct separate author profiles for earlier and later phases of Eminescu's journalistic activity and evaluate manuscript texts from the opposite period. We define the *early* period as Vols. 9–10 (1870–1878) and the *late* period as Vols. 12–13 (1881–1883, 1888–1889), excluding Vol. 11 (1880) to increase temporal separation. Distance distributions remain similar across early and late configurations for both feature sets, suggesting that verification outcomes are not driven by a narrow temporal slice of the verified corpus.

Finally, we assess the stability of the decision threshold under sampling variation by bootstrapping the manuscript set used for threshold estimation and recomputing the threshold over 200 trials. For both feature representations, the resulting bootstrap distributions exhibit limited variance, indicating that the acceptance criterion is stable under resampling.

Method	Profile	Threshold	PUB held-out (339)	MAN (204)	DUB (103)
CNG rank	PUB	MAN	–	198	97
FW rank	PUB	MAN	–	204	103
CNG rank	80% PUB + MAN	20% PUB	332	–	100
FW rank	80% PUB + MAN	20% PUB	338	–	102
CNG freq	PUB	MAN	–	204	103

Table 4: Number of accepted texts across feature types and profile configurations

Title	Volume	Words	Summary	Rejected
Fata mamei Ango și Giroflé-Girofla	Vol. 9	567	theatre commentary with moral reflection	CNG
Nefericitul X	Vol. 9	376	humorous anecdotal sketch	CNG
Ziarele din Viena	Vol. 10	475	biographical news report on Francisc Schuselka	CNG
Asaltul Angelescu	Vol. 10	649	political–military critique	CNG
Din istoria calului	Vol. 10	555	historical essay on horses in antiquity	CNG
O serată literară. Despot Vodă	Vol. 10	183	cultural news report on a literary salon event	FW
Printr-o indiscrețiune	Vol. 12	350	diplomatic news brief containing a French telegram	CNG

Table 5: DUB texts rejected by at least one rank-based profile

4.2 Results

Table 4 reports, for each feature type and profile construction setting, the number of texts accepted as stylistically compatible with the author profile. The numbers in parentheses in the column headers indicate the number of texts evaluated in each category for the corresponding setting: manuscript texts (MAN) and held-out published texts (PUB) used for threshold estimation, depending on the profile construction strategy, and the full set of disputed texts (DUB). Across all feature types and profile construction settings, only 7 of the 103 disputed texts are rejected at least once, while the remaining 96 are accepted under all feature types and profile construction settings.

We observe that a small number of published or manuscript texts also fall marginally outside their respective acceptance thresholds. Such deviations are expected, as thresholds are derived statistically from empirical distance distributions and are not intended to include every individual verified text. Importantly, these cases are rare and do not indicate systematic instability.

For rank-based CNG representations, acceptance remains high for verified texts while avoiding complete acceptance of DUB texts. Although most disputed texts are accepted, the observed rejections indicate that the method maintains discriminative sensitivity instead of converging toward uniform acceptance. In the PUB profile with MAN threshold setting, CNG rank accepts 97 of 103 DUB texts (94%), while in the second profile setting it accepts 100 of 103 (97%). The three DUB texts rejected in the second setting are a subset of the six rejected in the first. This balance aligns with the goals of authorship verification, which aims to accommodate natural stylistic variation while remaining sensitive to meaningful deviations.

In contrast, the function word based rank distance yields uniformly high acceptance across all text categories, including DUB texts. Under this feature representation, the second profile setting accepts 102 of the 103 disputed texts, while the first setting accepts all of them. This behaviour suggests that, in the present setting, function word distributions might be insufficiently discriminative. Two interacting factors contribute to this effect. First, the newspaper register limits how much function word use can vary across texts, leading to very similar rank profiles throughout the corpus. Second, short texts introduce some local rank noise, but because the threshold is estimated from comparable verified texts, this variation is treated as normal rather than leading to rejections. As a result, distance values are tightly clustered and acceptance is high.

Finally, the frequency-based CNG dissimilarity function, used here as a baseline, produces a highly permissive profile that accepts all DUB texts. This confirms that the rank-based formulation provides a stricter and more informative decision criterion in this authorship verification setting.

Results analysis

Table 5 lists the DUB texts rejected by the stricter rank-based criterion, together with their volume of publication, length, brief content summaries, and the feature representation under which rejection occurs. To understand what drives these rejections,

we examine this small subset qualitatively and consider factors known to affect profile-based stylometric distances.

A close reading of Eminescu's verified journalistic texts reveals a recurrent structural tendency. Even when prompted by concrete events or documents, his articles rarely remain at the level of description. Reporting is typically followed by interpretive expansion, causal explanation, or normative critique, often extending the immediate topic into broader reflections on institutions, social dynamics, or political principles. Several rejected DUB texts, by contrast, remain confined to the informational layer and do not develop this progression from occasion to diagnosis.

Some rejected texts illustrate this contrast particularly clearly. The humorous sketch *Nefericitul X* consists of a sequence of loosely connected anecdotes built around a generic figure, without developing a sustained argument or interpretive trajectory. *Din istoria calului*, a historical essay on the horse, similarly presents a factual overview, compiling information without developing a broader interpretive or argumentative claim. Both are rejected by both CNG profiles and differ from the dominant mode of the verified corpus, where historical or biographical material usually serves as the basis for polemical or diagnostic argumentation.

The remaining rejected DUB texts also converge in discourse organisation. They rely on episodic or self-contained forms of exposition, without developing a sustained argumentative trajectory. Meaning is produced through a sequence of scenes, illustrative examples, or isolated incidents, without the progressive accumulation of causal or institutional analysis that characterises much of Eminescu's verified journalism. This pattern appears across genres, from *Ziarele din Viena*, a biographical news report, to the pedagogical theatre commentary *Fata mamei Ango și Giroflé-Girofla*, and the political military commentary *Asaltul Angelescu*, which addresses a concrete incident without extending it into a broader institutional critique.

Other rejected cases highlight more specific sensitivities of the feature representations. The only text rejected under the function word profile, *O serată literară. Despot Vodă*, is an extremely short cultural notice reporting on a literary salon event. Approximately one third of the text consists of a list of personal names. In such a short document, this concentration substantially alters the relative distribution of function words, producing an atypical profile. Notably, this text is not rejected by the character n-gram based profile, suggesting that CNG features are less sensitive to this specific distributional effect.

Another illustrative case is *Printr-o indiscrețiune*, which consists of a diplomatic news report reproducing and commenting on telegraphic exchanges between European political figures. A substantial portion of the article is made up of verbatim quotations in French, reflecting diplomatic conventions of the period. This text is rejected under both CNG based profile configurations. However, a close qualitative reading suggests stylistic compatibility with Eminescu's journalistic writing in the Romanian framing passages. To probe this further, we repeated the analysis after removing the French-language segment. Under this condition, the remaining Romanian text was no longer rejected by the CNG profile. This indicates that the rejection is primarily driven by the presence of extended foreign-language quotations, which substantially alter character n-gram distributions, with the surrounding text remaining stylistically compatible.

Taken together, the rejected DUB texts display recurring characteristics that help explain their distance from the author profile. These properties are relevant for authorship analysis because they intersect with recurring aspects of Eminescu's journalistic practice, which typically integrates reporting with interpretive or argumentative development. The verification outcomes therefore bear on authorial compatibility, while remaining sensitive to the possibility that register or compositional factors may also contribute to observed deviations.

From a methodological perspective, these findings suggest that profile-based authorship verification should not be evaluated solely in terms of aggregate acceptance rates. Instead, the explanatory value of rejections and their alignment with known discourse and genre properties are important. In this respect, the present study illustrates how verification outcomes function as analytic signals that prompt closer textual examination, not as categorical judgments about authorship.

5 Conclusion

This paper investigated profile-based authorship verification in a historically grounded journalistic corpus, focusing on the writings of Mihai Eminescu. Using rank-based distance measures over

character n-grams and function words, we showed that character trigram profiles provide stable verification behaviour across different profile constructions, temporal splits, and threshold estimation procedures, while remaining tolerant of natural stylistic variation within verified texts.

The proposed framework does not aim to maximize discrimination, but to preserve stylistic continuity while identifying meaningful deviations. The rank-based profile representation yields decisions that are stable under resampling and can be related to concrete textual properties, making them suitable for interpretive analysis. Our qualitative examination shows that texts rejected by the model systematically differ from the verified corpus in discourse structure, communicative purpose, and textual composition, with differences that are not attributable to random variation. These findings suggest that profile-based authorship verification is well suited for exploratory analysis of stylistic boundaries within an author's body of work.

More broadly, this study argues that in historical and literary contexts with partial ground truth, authorship verification methods should be evaluated not only by acceptance rates, but by the stability, interpretability, and coherence of their outcomes. Our approach offers these properties and is therefore well suited to corpus-driven literary scholarship in settings with partial and uncertain ground truth. Future work may extend this approach to comparative author profiles, mixed author corpora, or other languages, and explore how verification signals interact with genre, register, and editorial practices in diachronic corpora.

Limitations

We focus on a single author and a single historical corpus, so generalization to other authors, languages, and genres remains to be established. Evaluation on comparable journalistic texts by other authors would provide an explicit estimate of false acceptance under domain-matched negative controls. Document length variation may affect rank stability, especially for shorter texts. While rank-based distances are less sensitive to absolute frequency differences than frequency-based measures, a more systematic assessment using controlled text segmentation would further clarify the influence of document length on verification behaviour. Finally, the method assesses stylistic compatibility, and acceptance and rejection should be interpreted as analytic signals, not as definitive authorship judgments.

Acknowledgments

This work was supported by Ministry of Research, Innovation and Digitization, CNCS-UEFISCDI, project SIROLA, number PN-IV-P1-PCE-20231701, within PNCDI IV and by the project "Romanian Hub for Artificial Intelligence - HRIA", Smart Growth, Digitization, and Financial Instruments Program, 2021-2027, MySMIS no. 334906.

References

Cindy Chung and James Pennebaker. 2011. The psychological functions of function words. In *Social communication*, pages 343–359. Psychology Press.

Liviu P Dinu. 2003. On the classification and aggregation of hierarchies with ifferent constitutive elements. *Fundamenta Informaticae*, 55(1):39–50.

Liviu P Dinu, Vlad Niculae, and Octavia-Maria Şulea. 2012. Pastiche detection based on stopword rankings. exposing impersonators of a romanian writer. In *Proceedings of the Workshop on Computational Approaches to Deception Detection*, pages 72–77.

Oren Halvani, Christian Winter, and Lukas Graner. 2017. Authorship verification based on compression-models. *arXiv preprint arXiv:1706.00516*.

Chia-Yu Hung, Zhiqiang Hu, Yujia Hu, and Roy Lee. 2023. Who wrote it and why? prompting large-language models for authorship verification. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 14078–14084.

Nicolae Iorga. 1969. *Studii literare. Scriitori români*, volume 1–2. Lyceum, Bucureşti.

Dumitru Irimia, editor. 2009. *Dicţionarul limbajului poetic eminescian. Concordanţele prozei antume*. Editura Universităţii din Bucureşti, Bucureşti.

Patrick Juola. 2013. How a computer program helped show jk rowling write a cuckoo's calling. *Scientific American*, 20.

Patrick Juola. 2021. Verifying authorship for forensic purposes: A computational protocol and its validation. *Forensic Science International*, 325:110824.

Junghwan Kim, Haotian Zhang, and David Jurgens. 2025. Leveraging multilingual training for authorship representation: Enhancing generalization across languages and domains. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 34855–34880.

Moshe Koppel, Jonathan Schler, and Elisheva Bonchek-Dokow. 2007. Measuring differentiability: Unmasking pseudonymous authors. *Journal of Machine Learning Research*, 8(6).

Moshe Koppel and Yaron Winter. 2014. Determining if two documents are written by the same author. *Journal of the Association for Information Science and Technology*, 65(1):178–187.

Solomon Marcus. 1970. *Poetica matematică*. Editura Academiei, Bucureşti.

Mihaela Mocanu. 2020. Eminescu's journalistic works–editing approaches and reading patterns. *Philobiblon*, 25(1):43–61.

Marius Popescu and Liviu P Dinu. 2008. Rank distance as a stylistic similarity. In *Coling 2008: Companion volume: Posters*, pages 91–94.

Nektaria Potha and Efstathios Stamatatos. 2014. A profile-based method for authorship verification. In *Hellenic Conference on Artificial Intelligence*, pages 313–326. Springer.

Luiza Seche. 1974. *Lexicul artistic eminescian în lumină statistică*. Editura Academiei R.S.R., Bucureşti.

Efstathios Stamatatos. 2009. A survey of modern authorship attribution methods. *Journal of the American Society for information Science and Technology*, 60(3):538–556.

Efstathios Stamatatos. 2016. Authorship verification: A review of recent advances. *Res. Comput. Sci.*, 123(1):9–25.

Enrico Tuccinardi. 2017. An application of a profile-based method for authorship verification: Investigating the authenticity of pliny the younger's letter to trajan concerning the christians. *Digital Scholarship in the Humanities*, 32(2):435–447.