# Analysis of persona assigned LLMs

**Ioana-Călina Pascu**
National University of Science and Technology Politehnica Bucharest
313 Splaiul Independenței Bucharest, Romania
ioanacalinap@gmail.com

**Ștefan Trăușan-Matu**
National University of Science and Technology Politehnica Bucharest
313 Splaiul Independenței Bucharest, Romania
and
Research Institute for Artificial Intelligence ”Mihai Draganescu”
and
Academy of Romanian Scientists
stefan.trausan@upb.ro

**ABSTRACT**

As generative language models become more capable of imitating authorial style, traditional ideas about originality and literary expressiveness must be reconsidered. This paper aims to evaluate, through a computational approach, the extent to which a large language model (LLM) can reproduce the stylistic patterns of two well-known authors: James Joyce and E.M. Forster. The comparative analysis relies on the extraction of features grouped into three categories: stylometric (statistics on structure and vocabulary), idiolectal (unusual words, neologisms, sentence types), and rhythmic-prosodic (sound-based rhetorical devices). These features are examined in both original and machine-generated texts using quantitative measures such as Euclidean distance, cosine similarity, Jensen–Shannon divergence, and the Jaccard index. The results reveal subtle but significant differences between what can be algorithmically reproduced and what remains unique to human literary expression.

**Author Keywords**

stylometry, idiolect, rhetorical devices, generative AI

**ACM Classification Keywords**

I.2.7 Natural Language Processing

**DOI:** 10.37789/icusi.2025.15

## INTRODUCTION

The emergence of large language models (LLMs) has fundamentally brought a major shift in how machines generate language, bringing generative capabilities closer to human-like expression. Built upon deep neural architectures and trained on massive corpora, LLMs now can be used to produce not only grammatically sound text, but also surprisingly convincing stylistic imitations.

This advancement raises questions about the boundaries between human creativity and algorithmic mimicry. Recent studies [1] have demonstrated that machine-generated poetry may be perceived as more “human” than works authored by established writers, including canonical figures like T.S. Eliot. Such findings challenge traditional notions of authenticity, originality, and literary value. If a generative model can convincingly replicate an author's voice, what then remains uniquely human in literary creation?

This research addresses the above question by conducting a comparative stylistic analysis between original literary excerpts and their LLM-generated counterparts. Specifically, the study investigates the ability of a state-of-the-art generative model to imitate the stylistic fingerprint of two distinct authors: James Joyce and E.M. Forster. The comparison is based on three levels of feature analysis: stylometric, idiolectal, and rhythmic. Each feature set is quantified using established distance and similarity measures — Euclidean distance, cosine similarity, Jensen–Shannon divergence, and Jaccard index — to compare original and generated texts. The experiment includes one chapter from *Ulysses* by James Joyce (i.e. Sirens) and *A Room with a View* written by E.M. Forster, alongside machine-generated texts of similar length, prompted to imitate the style of the respective author. The objective is not only to measure textual similarity, but to interrogate the creative limits of generative models and their ability to reproduce the expressive complexity of human-authored literature.

## STATE OF THE ART

### Stylometric Analysis and Distant Reading

The concept of *distant reading* refers to the use of computational and statistical methods to analyze literary texts at scale, bypassing close interpretive reading. Rather than aiming for hermeneutic depth, distant reading identifies patterns, themes, frequencies, or affective tones based on explicit processing instructions [2]. This form of “non-reading” is highly relevant to the present study, which examines whether a large language model (LLM) can

replicate the structural and stylistic patterns of human-authored texts.
Stylometry, a core approach within distant reading, focuses on quantifiable linguistic traits such as sentence length, word frequency, and function word usage. These features tend to manifest unconsciously and are more resistant to intentional manipulation, making them reliable stylistic indicators [3]. Stylometric studies often combine data preprocessing, ad hoc feature extraction, statistical analysis, and visualization techniques to identify authorial signatures across large text corpora.

Early work in this field includes Mendenhall's graphical analysis of word-length distributions in Shakespeare [4] and Bacon's works, suggesting that in this way they could distinguish authors. A foundational contribution came from Mosteller and Wallace [5], who applied Bayesian inference to analyze the frequency of function words in *The Federalist Papers* — words like *the*, *of*, or *and* that carry little semantic load but reveal stable stylistic habits. Other influential studies include Brinegar's statistical attribution of anonymous letters to Mark Twain [6] and Thisted and Efron's probabilistic evaluation of a disputed Shakespeare poem based on vocabulary diversity [7].

Despite concerns about potential superficiality ("*literary statistics can be charged with crudeness and shallowness*" [8]) stylometry remains an indispensable tool for comparative literary analysis. It provides access to the *signifiants* (linguistic signs) but not directly to the *signifiés* (their conceptual meaning), thus complementing rather than replacing traditional literary interpretation.

### Natural Language Processing (NLP)

Natural Language Processing (NLP) is a branch of artificial intelligence concerned with modeling human–machine linguistic interaction. Its goal is to interpret, manipulate, and generate natural language with both semantic accuracy and discursive fluency [9]. NLP pipelines typically involve preprocessing (e.g., tokenization, case normalization, stopword removal, and stemming), feature extraction (e.g., lexical and syntactic markers), and textual analysis using machine learning, sentiment analysis, or semantic relation extraction.

A breakthrough was the introduction of the Transformer architecture, initially designed only for machine translation. Inspired by encoder–decoder structures, it incorporates self-attention mechanisms to capture complex linguistic dependencies, such as pronoun resolution or long-range syntactic links [10]. Based on this foundation, models like BERT (Bidirectional Encoder Representations from Transformers) [11] and GPT (Generative Pre-trained Transformer) [12] have become standard across NLP tasks.

Still, human language remains fundamentally ambiguous —rich in polysemy, figurative structures, and contextual nuance. This complexity is captured in the observation that "*the invention of natural languages cannot be credited to purely rational minds*" [8]. As a result, even the most advanced NLP systems struggle to fully resolve the intricacies of meaning, particularly in literary texts.

### LLMs and Ethical Considerations

LLMs are deep learning models capable of producing new textual content by analyzing and modeling large-scale language data. They emulate stylistic, grammatical, and narrative logic through probabilistic text generation, selecting each next word based on its likelihood given previous context. GPT-based models, for instance, predict tokens autoregressively, while BERT uses a masked language modeling objective, similar to cloze tasks [11][12][13].

The rise of LLMs in literary contexts raises significant ethical concerns. Because they are not guided by creative intent or boundaries, these models often produce content that, while readable, lacks originality, forcing us to rethink what makes a text truly "authored". Moreover, AI generated literature may be even accused of a kind of plagiarism, since unlike human storytellers, guided by intention, culture, and aesthetic vision, AI relies on pattern recognition and reuse across vast datasets. These concerns lead us to ask: If machines can convincingly mimic literary style, what remains uniquely human in creative writing? Does reliance on AI-generated content risk diluting artistic diversity, reducing expression to trends and data-driven conformity? [14]. Such concerns suggest the need for responsible integration of AI in literature. Rather than replacing human creativity, LLMs should serve as tools for augmentation, experimentation, and critical reflection — enhancing, but not displacing, the rich multiplicity of human voices.

## IMPLEMENTATION

### General Architecture

The project follows a linear yet modular architecture, consisting of four main components:

- Input and selection of texts,
- Text preprocessing,
- Feature extraction,
- Comparison and analysis.

Each stage is implemented as an independent Python module to facilitate reuse and clarity in future research (e.g. across different authors or models). The tool compares two primary inputs: a literary text and its stylistic counterpart generated by a GPT-based model. The full processing pipeline begins with cleaning and annotating the texts, continues with feature extraction, and ends with comparative evaluation and visualization.

### Tools and Libraries

The system is developed in Python, chosen for its robust NLP and statistical libraries and its compatibility with Jupyter Notebooks for rapid prototyping. Key libraries include:

- **spaCy** – tokenization, part-of-speech tagging, syntactic parsing, sentence segmentation;

- **NLTK** – n-gram generation, stop word filtering, syllable estimation, and frequent phrase detection;
- **CMU Pronouncing Dictionary** via the pronouncing package – phonetic transcription and stress pattern recognition;
- **Scikit-learn & SciPy** – for computing cosine similarity and performing statistical comparisons;
- **WordNet** (via NLTK) – to identify neologisms and out-of-vocabulary terms.

### Text Preprocessing
Preprocessing is crucial for standardizing and preparing text data. The raw input is first cleaned to remove non-linguistic artifacts such as HTML tags, markdown syntax, or excessive whitespace. The cleaned text is segmented into sentences and tokenized into words using NLTK's standard tools.

The spaCy library is then applied to annotate the text syntactically and morphologically. Each token receives information such as part-of-speech (POS), dependency relation, and lemma. This linguistic structure enables extraction of advanced features such as tree depth, branching factor, and POS distributions.

### Feature Extraction
This study uses a hybrid feature extraction approach combining computational stylometry, idiolectal analysis and the analysis of rhythmic and sound-based patterns. Features are categorized into three distinct layers: stylometric, idiolectic, and rhythmic-prosodic, each targeting different levels of linguistic structure.

### Stylometric Features
Stylometric indicators are used to assess lexical, syntactic, and structural properties of texts — features often left unconsciously by authors and considered to form a "stylistic fingerprint." Extracted attributes include:

- **Character and word frequency;**
- **N-grams**: contiguous sequences of *n* units (unigrams, bigrams, trigrams) that capture recurring patterns at different granularities;
- **Stopword ratio**: frequency of function words (e.g., *the*, *and*, *of*), which carry minimal semantic load but are stable stylistic markers;
- **Type-Token Ratio (TTR)**: a lexical diversity metric calculated as the ratio of unique words (types) to total word tokens. Higher values indicate richer vocabulary use;
- **Syntactic tree depth and branching factor**: derived from dependency parses, these features reflect structural complexity. Tree depth measures the longest syntactic path in a sentence, while the branching factor quantifies local syntactic density.

These features, implemented using spaCy and NLTK, offer interpretable metrics for characterizing surface style and syntactic architecture.

### Idiolectal Features
The term *idiolect* describes the personal and distinctive way an individual uses language — from the sounds they favor to the way they construct ideas in writing. Theoretically, no two people, even if they speak the same language, share the same linguistic patterns. In this study, idiolectal features are of particular interest, as they reflect whether a generative model can reproduce not just structural or lexical norms, but also the distinctive stylistic habits of a given author — such as recurrent expressions, culturally specific constructions, or preferred lexical items. *„In terms of philology, an individual (writing) style is a complex concept reflecting ones sociohistorical nature, ethnic, psychological, moral, and ethical peculiarities."* [15].

To model authorial distinctiveness, two feature types were extracted:

- **Sentence type distribution**: proportion of declarative, interrogative, and exclamatory sentences, determined via punctuation. This serves as a proxy for tonal and rhetorical variation.
- **Neologism detection**: identification of words not found in lexical resources such as WordNet. These include invented words, lexical blends, and creative morphological constructions.

These features highlight whether the generative model can emulate the expressive idiosyncrasies of individual authors.

### Rhythmic and Prosodic Features
Rhythmic-prosodic features focus on phonological repetition and auditory aesthetics, important stylistic components, especially in experimental prose. This study extracts:

- **Anaphora/Epiphora/Anadiplosis**: repetition patterns across sentence boundaries, detected by tracking first and last words;
- **Epizeuxis**: immediate repetition of the same word (e.g., *"Yes, yes, yes"*);
- **Alliteration**: repetition of initial phonemes, identified using phonetic transcriptions;
- **Stress patterns**: binary stress profiles per sentence, built from CMU Pronouncing Dictionary data, summarized via average stress length;
- **Syllable density**: average syllables per word/sentence, reflecting phonological weight:
- **Syntactic parallelism**: POS sequence overlap across adjacent sentences, capturing pattern-based sentence repetition;
- **Onomatopoeia**: detection of predefined sound-imitating words (e.g., *buzz*, *snap*, *crash*).

Together, these features aim to capture text rhythm, sound patterns, and rhetorical repetition, dimensions difficult to mimic through statistical learning alone.

### Distance, Similarity and Divergence Measures

To quantify stylistic (dis)similarity between original and generated texts, different measures were applied based on the data type of each feature:

### Numerical Features - Euclidean Distance

Scalar features (e.g., average word length, TTR, syllable density, syntactic depth) were compared using Euclidean distance, defined as:

$$d(x, y) = \sqrt{\sum_{i=1}^{n} (x_i - y_i)^2}$$

As a metric, it captures absolute deviations in a standardized real-valued space; small values indicate proximity in the underlying stylistic scalars. This choice assumes approximately well-behaved (often near-Gaussian) variability and offers direct interpretability as magnitude of difference [16].

### Vector Features - Cosine Similarity

Frequency-based features (e.g., character counts, POS tag distributions, n-grams) were treated as high-dimensional vectors. Their angular similarity was measured via:

$$\mathrm{sim}_{cos}(x, y) = \frac{x \cdot y}{||x|| \cdot ||y||}$$

The measure emphasizes directional alignment of sparse profiles and is invariant to overall length, making it suitable when composition rather than magnitude carries stylistic signal. Values near 1 indicate strong directional alignment (i.e. similar stylistic profiles) while values near 0 suggest orthogonal, unrelated feature distributions [17].

### Distributional Features - Jensen–Shannon Divergence (JSD)

Proportional features (e.g., punctuation types, syllable stress distributions, sentence types) were modeled as discrete probability distributions and compared using:

$$\mathrm{JSD}(P \parallel Q) = H\left(\frac{P + Q}{2}\right) - \frac{1}{2}H(P) - \frac{1}{2}H(Q)$$

where

$$H(P) = -\sum_{i=1}^{n} p_i \log p_i$$

is the Shannon Entropy of distribution P. Intuitively, JSD quantifies how much the entropy of the average distribution exceeds the average entropy of the individual distributions. In a stylometric context, this reflects much more uncertainty a reader would experience when encountering a hybrid text (an average between the original and the generated one) compared to reading each separately. JSD is symmetric and bounded in [0,1], with 0 indicating identical distributions and higher values reflecting increased stylistic divergence [18].

### Set Features - Jaccard Similarity

Categorical features (e.g., neologisms, rhetorical devices, onomatopoeia) were compared using **Jaccard similarity**, defined as:

$$J(A, B) = \frac{|A \cap B|}{|A \cup B|}$$

This similarity measure captures an overlap between sets without considering frequency. It is particularly effective for rare, expressive markers that may appear only once per text but carry high stylistic weight [19].

## RESULTS

This section presents the results of a comparative stylometric analysis between original literary texts and their machine-generated counterparts, with the goal of assessing the ability of a large language model (LLM) to replicate complex and distinctive literary styles. The results highlight both the model's strengths in reproducing surface-level patterns and its limitations in mimicking deeper stylistic and creative structures. This paper presents only a subset of the results obtained by the first author in her graduation thesis [20].

The first text selected for analysis is the *Sirens* chapter from *Ulysses* by James Joyce, chosen for its experimental prose, use of stream-of-consciousness, fragmented syntax, sound-based rhythmic patterns, and polyphony. The original text contains approximately 12,000 words. Its structure is marked by syntactic complexity, short repetitive phrases, and idiosyncratic linguistic inventions. The counterpart text was generated using GPT-4.0 (public version), prompted with: *"Write a story in the style of James Joyce's Ulysses, chapter Sirens. Write at least 1000 words."* Due to token and memory limitations in the API (Application Programming Interface) and the non-premium account, generation was performed iteratively using repeated *"Continue the story"* prompts. The process was repeated 15 times until the generated text reached a comparable length to the original.

To test robustness, the same comparative analysis was applied to a second novel with contrasting stylistic features: *A Room with a View* by E.M. Forster. The original text was approximately 379,000 characters in size. Generation was performed using ChatGPT-4.0 with the prompt: *"Write a story in the style of E.M. Forster's A Room with a View. Write at least 1000 words."*. Because of size constraints, generation was split into 56 iterations: one initial prompt followed by 55 consecutive *"Continue the story"* prompts.

This dual-text setup, combining one experimental modernist text and one classical narrative, offers a valuable perspective on the stylistic range of generative models. It also provides a benchmark for evaluating whether the stylometric, idiolectic, and rhythmic-prosodic metrics introduced earlier remain meaningful across divergent literary styles.

**Case Study: E.M. Forster**

The comparison between *A Room with a View* and the LLM-generated version reveals that the model handles certain surface-level stylistic features with notable precision. The generated text appears fluent and well-structured, capturing many of the formal characteristics that define Forster's prose style. This is not unexpected — such features are highly regular and well-represented in training data, making them easier to replicate through statistical learning.

In fact, Forster's prose follows a relatively stable narrative structure and relies on stylistic norms that are well represented in the model's training data. These elements can be convincingly imitated without requiring deeper understanding of narrative intent or literary context. In particular, the model closely aligns with the original text in several measurable aspects, for example, the average word length and type-token ratio (see Figure 1).

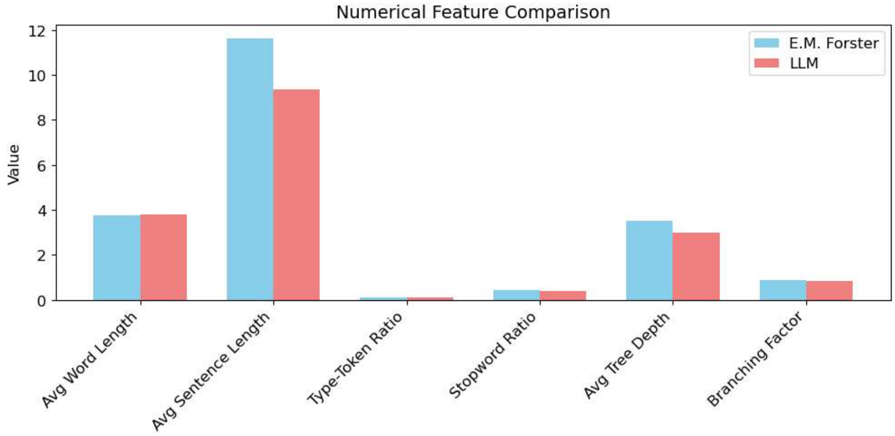


**Figure 1: Numerical features comparison (Forster vs LLM)**

However, divergences emerge in stylistic areas requiring intentional authorial control. Structurally, for example, the generated text displays:

- Shorter average sentence lengths.
- Shallower syntactic parse trees.
- More uniform rhythm and narrative pacing.

In addition to the narrative structure, there are differences in terms of rhetorical devices (see Figure 2). These findings confirm that while LLMs can approximate stylistic tendencies at a statistical level, they fall short of replicating the literary identity encoded in deeper creative decisions. These differences reflect the probabilistic nature of generative modeling: LLMs reproduce frequently observed linguistic patterns but lack access to the cultural or aesthetic intentions underlying an author's work.

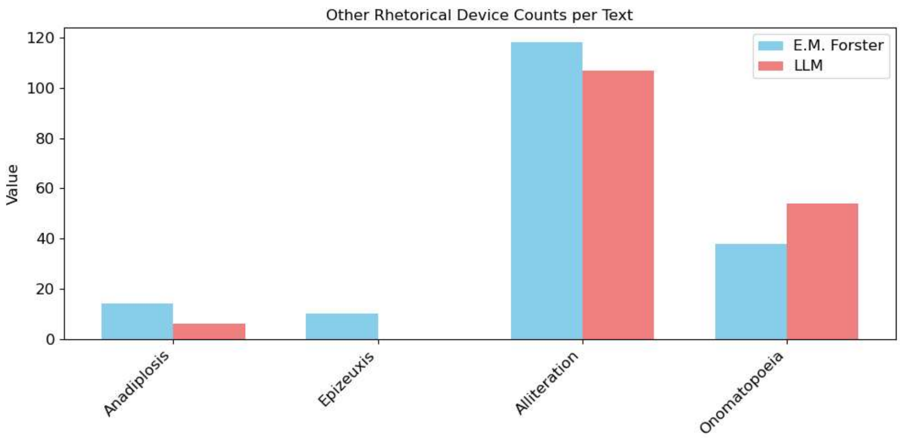


**Figure 2: Comparison of rhetorical devices (Forster vs LLM)**

**Case Study: James Joyce**

At a superficial level, GPT-4 demonstrated strong performance in reproducing the stylistic signature of Joyce's *Sirens* chapter:

- High cosine similarity in character and trigram frequency.
- Low Jensen–Shannon divergence on phonological features.

These results indicate high alignment with orthographic and phonetic regularities. However, more advanced stylistic dimensions show consistent failure (see Figure 3):

- The type–token ratio was lower than the original.
- Syntactic complexity was reduced, with shorter sentences and simpler structures.
- Lexical diversity was limited.

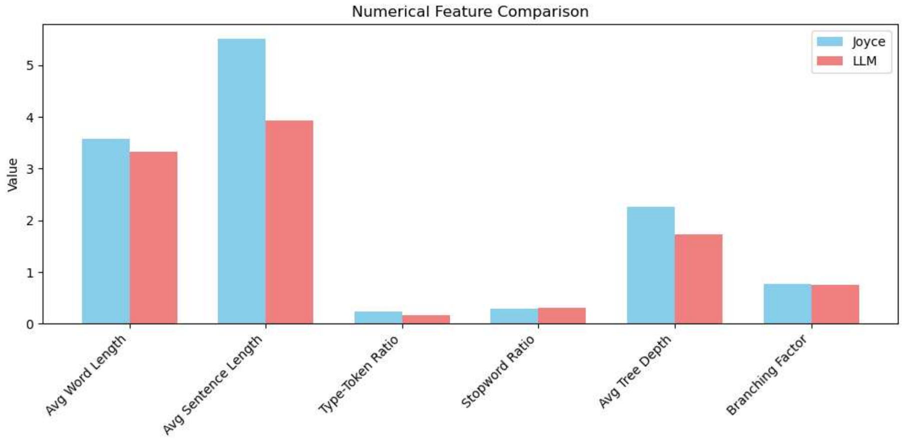


**Figure 3: Numerical features comparison (Joyce vs LLM)**

The idiolectic analysis most clearly revealed the model's limitations: the generated text contained no neologisms, failing to replicate Joyce's signature style of linguistic innovation. Phonetic distortions, invented words, and unconventional constructions, all essential to Joyce's text in *Sirens* chapter, were absent. On the rhythmic-prosodic level (see Figure 4), the LLM could detect and reproduce some devices (e.g., anaphora, epiphora, and onomatopoeia), but their distribution and contextual embedding were less convincing. Stress patterns were shorter and more uniform; syntactic parallelism was frequent but overly regularized, suggesting reliance on internal templates rather than adaptive creativity.

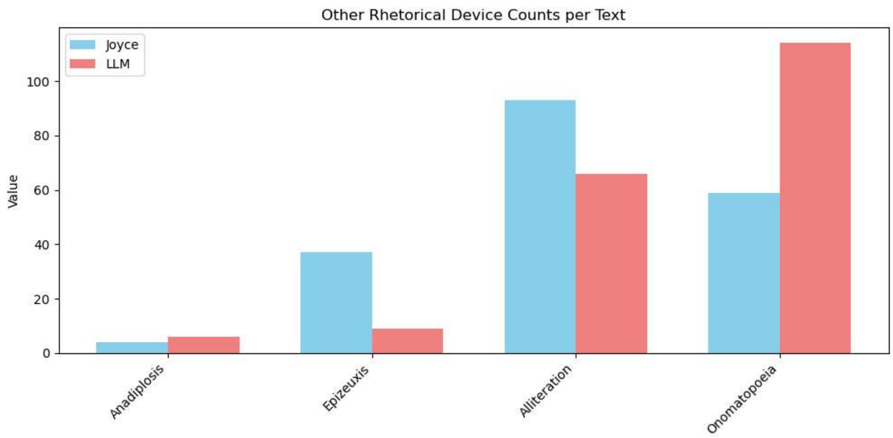


**Figure 4: Comparison of rhetorical devices (Joyce vs LLM)**

In fact, rhythm is a feature that relates text and music [21] and is induced by rhetorical devices [22]. *Sirens* chapter is recognized as having a major musical dimension. Even Joyce himself declared that this chapter is like a "fuga per canonem" [23]. A fugue is a polyphonic music piece, which has equivalents in novels [24], in which ideas/concepts/words ("voices", in music) occur in parallel, independent, threads entering in divergences/dissonances and convergences/consonances, inducing a sense of creativity and life [24, 25]. As Bakhtin mentioned, life has a polyphonic character [24] and, consequently, literary texts able to express in a higher degree the "flavor" of life should have, in our opinion, a similar character [22].

## DISCUSSION

The pattern across measures is consistent: lower-level cues (e.g., character n-grams, function-word ratios, basic POS distributions) align closely between original and generated texts, while higher-level structure (e.g., sentence-type proportions, dependency depth, rhetorical devices) diverges. Generated passages tend toward shorter sentences, shallower dependency trees, and a flatter rhythmic profile. These statements are descriptive; statistical significance is not claimed here.

Practical use follows from this split. For non-native writers, they highlight which features diverge from a target style (e.g., overly short sentences, few subordinate clauses, low use of discourse markers), enabling targeted revision. For educators, side-by-side distributions (POS, sentence types, n-grams) make feedback specific and explainable, pointing to missing connectors, flat sentence rhythm, or overuse of patterns in student work.

A methodological limitation concerns the generation configuration. Texts were produced with the free ChatGPT-4.0 web interface using a simple "*Continue the story*" prompt; this configuration offers limited control over decoding parameters and guidance, so outputs may be lower in quality than those from a pro/API setup.

## CONCLUSIONS AND FUTURE WORK

This study investigated the extent to which large language models (LLMs) can replicate the literary styles of human authors by comparing original texts with machine-generated imitations. Through a layered computational approach that includes stylometric, idiolectic, and prosodic features, the analysis showed that LLMs often capture surface style but fall short of deeper creative expression.

Quantitative metrics showed high similarity in lexical features such as word length, character frequency, and n-grams. However, in areas involving expressive intent — such as syntactic complexity, lexical innovation, and rhetorical devices — the generated texts were notably flatter, more uniform, and less contextually adapted. The clearest limitations emerged in idiolectic features: LLMs failed to generate neologisms, idiomatic expressions, or stylistic deformations characteristic of authors like Joyce. Even when stylistic figures (e.g., anaphora, epiphora) appeared, they were inserted in generic, pattern-driven ways, lacking the contextual nuance of the originals.

As expected, stylistic regularity in Forster's prose enabled better model reproduction, while the experimental nature of Joyce's writing highlighted the model's creative limitations. This contrast underscores the current asymmetry between pattern replication and genuine stylistic embodiment.

Future work should address these gaps by expanding the corpus to cover a broader range of genres and styles, integrating semantic and discourse-level features (e.g., narrative structure, topic coherence), and incorporating human evaluations to complement automated metrics. These directions could provide a more comprehensive understanding of what LLMs can and cannot achieve in reproducing literary expression.

Ultimately, while LLMs demonstrate impressive fluency, their imitation of style remains bound by statistical surface learning. True literary voice — rooted in intention, context, and cultural imagination — remains out of reach for artificial generation. Understanding this boundary is essential, not only for NLP evaluation, but for the broader cultural implications of algorithmic authorship. The capability of AI to generate complex polyphonic literary texts is also a research direction to be investigated.

## ACKNOWLEDGMENTS

This research is supported by the project "Romanian Hub for Artificial Intelligence - HRIA", Smart Growth, Digitization and Financial Instruments Program, 2021-2027, MySMIS no. 334906.

## REFERENCES

1. Sun, D. (2025, February 4). *AI is taking down T.S. Eliot. I blame us. The Washington Post*. https://www.washingtonpost.com/opinions/2025/02/04/chatgpt-ai-poetry-preference-study/. Accessed July 11, 2025.
2. J. Drucker, Why Distant Reading Isn't, *PMLA*, vol. 132, no. 3, pp. 628–635, 2017, doi: 10.1632/PMLA.2017.132.3.628.
3. R. D. Peng and N. W. Hengartner, Quantitative analysis of literary styles, *American Statistician*, vol. 56, no. 3, pp. 175–185, 2002, doi: 10.1198/000313002100.
4. T. Mendenhall, A mechanical solution of a literary problem, *Popular Science Monthly*, vol. 60, pp. 97–105, 1901.
5. F. Mosteller and D. Wallace, *Inference and disputed authorship: The federalist*, Stanford University Center for the Study, 2008.
6. C. Brinegar, Mark Twain and the Quintus Curtius Snodgrass letters: A statistical test of authorship,

*Journal of the American Statistical Association*, vol. 58, pp. 85–96, 1963.

7. B. Efron and R. Thisted, Estimating the number of unseen species: How many words did Shakespeare know? *Biometrika*, vol. 63, no. 3, pp. 435–447, 1976.
8. C. Delcourt, Stylometry, *Revue belge de Philologie et d'Histoire*, vol. 80, no. 3, pp. 979–1002, 2002, doi: 10.3406/RBPH.2002.4651.
9. N. Patwardhan, S. Marrone, and C. Sansone, Transformers in the Real World: A Survey on NLP Applications, *Information*, vol. 14, no. 4, p. 242, 2023, doi: 10.3390/INFO14040242.
10. A. M. P. Brasoveanu and R. Andonie, Visualizing Transformers for NLP: A Brief Survey, *Proc. Int. Conf. on Information Visualisation*, vol. 2020-September, pp. 270–279, 2020, doi: 10.1109/IV51561.2020.00051.
11. J. Devlin, M. W. Chang, K. Lee, and K. Toutanova, BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, *NAACL HLT*, vol. 1, pp. 4171–4186, 2018. [Online]. Available: https://arxiv.org/pdf/1810.04805
12. A. Radford, K. Narasimhan, T. Salimans, and I. Sutskever, Improving Language Understanding by Generative Pre-Training, [Online]. Available: https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf. Accessed July 11, 2025.
13. Y. Hu, Literature in the Age of Artificial Intelligence, pp. 1781–1787, 2023, doi: 10.2991/978-2-38476-092-3_228.
14. M. Swathi, R. Dhayalakrishnan, Bots and Books: How Artificial Intelligence is Shaping Contemporary Literature, *Contemporaneity of English Language and Literature in the Robotized Millennium*, vol. 3, no. 2, pp. 1–4, 2024, doi: 10.46632/CELLRM/3/2/1.
15. Lagutina, K., Lagutina, N., Boychuk, E., Vorontsova, I., Shliakhtina, E., Belyaeva, O., Paramonov, I., & Demidov, P. G. (2019). *A survey on stylometric text features*. In *2019 25th Conference of Open Innovations Association (FRUCT)* (pp. 184–195). IEEE. https://doi.org/10.23919/FRUCT48121.2019.8981504.
16. Y. Jie, J. Amores, N. Sebe, and T. Qi, A new study on distance metrics as similarity measurement, *IEEE Int. Conf. on Multimedia and Expo, ICME*, vol. 2006, pp. 533–536, 2006, doi: 10.1109/ICME.2006.262443.
17. A. R. Lahitani, A. E. Permanasari, and N. A. Setiawan, Cosine similarity to determine similarity measure: Study case in online essay assessment, *Proc. Int. Conf. on Cyber and IT Service Management, CITSM*, 2016, doi: 10.1109/CITSM.2016.7577578.
18. J. Briët and P. Harremoës, Properties of classical and quantum Jensen-Shannon divergence, *Phys Rev A*, vol. 79, no. 5, p. 052311, 2009, doi: 10.1103/PHYSREVA.79.052311.
19. S. Niwattanakul, J. Singthongchai, E. Naenudorn, and S. Wanapu, Using of Jaccard Coefficient for Keywords Similarity, in *Int. MultiConf. of Engineers and Computer Scientists, IMECS*, 2013, p. 578.
20. I.-C. Pascu (2025) Analiza stilului LLM-urilor cu o persona asignată, *Engineer Bachelor Thesis*, Computer Science Department, N.S.T. University Politehnica Bucharest
21. S. Trausan-Matu, Connections between human music and natural language, in Milena Slavcheva (ed.) *Knowledge, Language & Models*, INCOMA Ltd. Shoumen, ISBN 978-954-452-062-5, 2020, pp. 302-317.
22. S. Trausan-Matu and A.-D. Urse (2024) Computer-support for the analysis of rhetorical devices in texts, in Puchianu, C.M., Datcu, D. (eds.) *International Conference on User-System Interaction*, MatrixRom, Bucuresti, pp. 95-98.
23. J. McCreedy, A Musiquely Concrète Analysis of 'Sirens': Joyce's Use of the Harmonic Series and His Anticipation of Serialism, https://www.researchgate.net/publication/368190296_A_MUSIQUELY_CONCRETE_ANALYSIS_OF_SIRENS_JOYCE'S_USE_OF_THE_HARMONIC_SERIES_AND_HIS_ANTICIPATION_OF_SERIALISM, Accessed July 11, 2025.
24. M. M. Bakhtin, *Problems of Dostoevsky's Poetics* (C. Emerson, Ed. and Trans.). Minneapolis, MN: University of Minnesota Press, 1983.
25. S. Trausan-Matu, The Polyphonic Model of Collaborative Learning, in Mercer, N., Wegerif, R., & Major, L. (eds.), *The Routledge international handbook of research on dialogic education*, Routledge, New York, NY,2020, pp. 454-468.