LETTERS IN BIOMATHEMATICS
An International Journal

RESEARCH ARTICLE

OPEN ACCESS

Computational Poetics: An NLP Framework for Analyzing Instapoetry

Dibakar Sinha

Manipal University Jaipur, Department of CSE, Jaipur, India
dibakar.sinha@jaipur.manipal.edu
dibakar18@gmail.com

ABSTRACT

Within this article, we present a computational poetics approach to understanding Instapoetry as a form of digital poetry characterized by its short form, minimalist visual presentation, emotional directness, and dissemination of digital poetics concepts via social media. Using natural language processing, we examine the linguistic, stylistic, thematic, and affective features differentiating Instapoetry from customary poetry forms. These consist of lexical diversity, sentiment and emotion detection, topic modeling, figurative language, readability, syntactic pattern, and semantic embedding. There are also platform-based textual features such as hashtags, line breaks, punctuation, and visual spacing, as well as metrics such as audience interaction. By combining quantitative text analysis and models of literary reading, the framework shows how recurrent themes, attitudes, styles, and forms can be identified across large corpora of poetry and goes on to show how computational analysis can contribute even to close readings of individual poems without reducing their meanings to mere numbers. This study contributes to the digital humanities, computational literary studies, and scholarship on social media discourse by offering a scalable, interdisciplinary methodology for studying emergent poetic forms. It is also the first study to show how NLP can be used to model and analyze the impact of Instapoetry on poetic writing, authorship, emotion, and readership online.

ARTICLE HISTORY
Received October 15, 2026
Accepted November 19, 2026

KEYWORDS
Computational Poetics; Digital Literature; Instapoetry; Natural Language Processing; Sentiment Analysis; Topic Modelling; Platform Poetics

1. Introduction

As Instagram is the most common platform for the writing of literature, a new genre of poetry has emerged which is designed explicitly for the platform. Instapoetry is characterized by its ultra-compactness and minimalistic layout, emotional immediacy, and algorithmic discoverability through hashtags and user interaction, which make it a genre distinct from customary poetry. Tens of millions read instapoetry, which has reshaped what poetry is and does in the 21st century (Nusair, 2020; Ramya & Rukmini, 2021). The viral success of Rupi Kaur's first book saw the largest ever sales of ebook and hard copy since its gradual accumulation of likes and shares on Instagram. It is a shift in the qualitative nature of the conditions of literary production, not merely a new channel of distribution. It is not a type of poetry on Instagram (though there is this), but a distinctive form of poetry co-formed by the platform's visual grammar and algorithmic affordances (Skains, 2023; Viires, 2020).

Despite the number of readers, Instapoetry has been an outlier in both literary critical work and digital humanities research. There has been a division between critics who view the aesthetic minimalism of the genre as a democratizing movement in the readership of poetry and those that see it as a disservice to form (Ilias et al., 2024; Proferes et al., 2024). Not many computational tools have been developed for the study of platform vernacular poetry (Abdillah, 2026; Colella, 2025) compared to other digital humanities and computational literary studies research, which has largely focused instead on canonical literary corpora, electronic literature, and generative text systems. The resulting artificial disciplinary bifurcation is that qualitative studies of Instapoetry are limited to small-scale empirical samples and close reading analysis, while computational methods can study Instapoetry at scale but are not specifically developed for the particularities of the form itself.

Natural Language Processing (NLP) has become a tool for literary analysis over the last few years. Transformer-based models such as BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019) have attained strong performance in tasks of classification of sentiments and emotions. Latent Dirichlet Allocation (Blei et al., 2003) remains the de facto standard for topic modeling, but NLP methods face special challenges when applied to Instapoetry. Poems are short (3-8 lines, 15-60 words), challenging models trained on much longer

CONTACT Dibakar Sinha: dibakar.sinha@jaipur.manipal.edu

documents. They also include platform-specific features such as hashtags, unorthodox capitalisation, and syntactic fragmentation, requiring bespoke tokenisation and data preparation. The highly figurative nature of poetry may require different detection tools and methods from those used for conversational and journalistic registers. Similarly, the mixed textual and engagement features of poetic discourse require methods that can distinguish between literary and platform features without reducing one to the other (Liu, 2025; Cheurfa, 2025).

To this end, we propose and evaluate a computational poetics framework using seven independent modules for computationally analyzing Instapoetry: BERT sentiment analysis, RoBERTa emotion analysis, LDA topic modeling, the VUA metaphor corpus, the Flesch-Kincaid readability scale, spaCy syntactic analysis and Sentence-BERT semantic embedding with UMAP dimensionality reduction for visualizing the embedding space. The framework is evaluated on a collection of 1200 Instapoems from twelve of the most followed Instapoetry accounts, and a 600-poem canonical reference corpus. It produces an empirical characterisation of Instapoetry as a computationally distinct literary register. This study seeks to investigate systematic linguistic and stylistic differences between Instapoetry and canonically canonical poetry, as well as the affective and thematic structures of large-scale Instapoetry corpora. What do these structures tell us about platform use? Can NLP approaches be integrated into the close reading and interpretive work of literary criticism without adopting the reductive quantification of literary meaning?

Instapoetry is a computationally identifiable genre based on positive valence, thematic focalization, low lexical and syntactic complexity, and moderate domain-specific figurative density, which arise from individual writing practices and platform selection pressures. The analysis shows that computational poetics provides principled, scalable, and data-driven insights into instapoetry and sets the stage for future validation of the analysis. The analysis, which bridges digital literary theory and NLP-based text analysis, contributes to the interdisciplinary field of computational literary studies and improves our understanding of how platform capitalism shapes literary form in the twenty-first century (Cortés, 2023; Jennings, 2016).

2. Literature Review

Instapoetry as a Platform-Native Literary Form

Instapoetry's constitution is an exercise of Instagram's interface affordances, algorithmic architecture, and social affordances. These have resulted in four features that define its aesthetic: extreme textual brevity for compatibility with the form's legibility on screens, use of minimal typography to create meaning with white space, emotional immediacy to present common affective situations, and discoverability through categorization via hashtags (Nusair, 2020; Sutherland et al., 2020). These affordances are not arbitrary; they are family resemblance features of the platform. The square image format prioritizes typographical constraints related to various typologies. The hashtagging system allows audience members to form around emotional registers. The algorithmic timeline values posts that can generate quick audience responses (Hoydis, 2021; Ramya & Rukmini, 2021). It is not merely a distribution channel, but an active co-producer of the conventions of the form.

Nevertheless, many critics have dismissed the form as literarily superficial, arguing that the democratizing effects of social media have led to poetry being simplified and commercialized at the cost of the form's usual intricacy, figurative language and difficulty (Proferes et al., 2024). Defenders of Instapoetry within academia have argued that the dismissal is the result of an elitism which equates literariness with formality, and that the form represents a genuine democratization of the poetic readership (Ilias et al., 2024; Viires, 2020). Both sides have largely engaged qualitatively with texts, and dedicated empirical studies of large data sets have been rare. The present study therefore extends this debate by providing quantitative evidence of the form's systematicity, allowing us to better characterize what Instapoetry actually does linguistically and affectively at scale, regardless of what it is said to be doing literarily.

Additionally, some of the attributes of Instapoetry that are meaningful to interpretation have been shown qualitatively to include: it is a medium for trauma recovery and self-help (Cheurfa, 2025); it is overproduced and overconsumed by women and marginalized communities (Rizzi, 2024); it relies on a second person "you" to universalize its address; and it uses enjambment to replicate rhythm and pauses in space rather than relying on a more customary metrical system. Rettberg (2021) describes these attributes in relation to the field of electronic literature and its interests in interactivity and reader participation. Skains (2023) proposes that social media platforms' algorithmic visibility logic can be understood as a new form of editorial gatekeeping which selects for affective text over novel text. This qualitative work informs the interpretative framework for the quantitative results of the current study.

Computational Literary Studies and NLP for Poetry Analysis

Computational literary studies (CLS) provides an array of quantitative methods for analyzing literary works, such as stylometric attribution of authorship, character network analysis, topic modeling and sentiment analysis (Jockers, 2013; Moretti, 2013). Within CLS, prosodic and sonic features as well as semantic features have been used specifically for analyzes of poetry. The first system to focus on

NLP was created by Kao and Jurafsky (2012), who demonstrated that NLP features could predict whether a poetic text was professional or amateur with high precision and accuracy, thereby establishing the provenance of computationally detected literary quality features. Other post-Kao and Jurafsky systems have used machine learning to perform tasks such as genre classification, periodisation, and inferring influence relations between poets.

Another transformative model architecture is the transformer. They generate contextualized word and sentence embeddings such as BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019) or Sentence-BERT (Reimers & Gurevych, 2019) to encode sentence meanings. This enables them to represent semantic information that was not captured by bag-of-words or word2vec models. These models have been used to achieve state of the art results on sentiment analysis, emotion detection and textual entailment tasks, and have also been fine-tuned to complex textual registers including literary language. Studies have also shown that transformer models can be used to study questions around AI literary authorship and creative writing practice, as in Colella (2025) and Lindberg et al. (2025). The purpose of the present framework would thus be to address this gap in the systematic application to Instapoetry specifically.

Emotion detection, another form of affect detection, is critical for poetry analysis given that expressing emotion is a defining characteristic of poetic writing. Third-party datasets such as the GoEmotions dataset (Demszky et al., 2020), composed of 28 fine-grained emotion categories from comments on Reddit, have advanced the state of study in emotion classification from binary sentiment. Literary emotion differs from conversational emotion at least in terms of its figurativeness, temporal complexity, and intentional vagueness. Ajith and Ravichandran (2025) and Ge (2025) have tested these literary characteristics on AI-generated poetry and found that general emotion models should be adapted to literary texts to be effective in detecting literary emotion. This is best represented by the present framework, which augments the GoEmotions fine-tuning dataset with 800 manually annotated works of literature for the four most commonly observed emotion labels in Instapoetry: love, sadness, hope, and longing.

Platform Poetics, Affordances, and the Circulation of Digital Literature

Platform poetics studies how the infrastructures of digital platforms, including technical, economic, social, and cultural aspects, condition the production, circulation, and reception of literary texts (Cortés 2023; Viires 2020). Platform poetics considers platforms as active mediating devices that shape textual production through the design of interfaces, algorithmic ranking systems, and metrics of audience interaction. Character limits, digital media container forms (e.g. image formats), and the discoverability of hashtags and the structural logics of the like and share become material constraints. They consequently become the formal features of digital literature, such as in the case of Indonesian Twitter bot poetry. Abdillah (2026:481) observes how these constraints are the algorithmic repetition and structural fragmentation of popular aesthetics optimizing continuous production and algorithmic visibility, respectively. This paper extends this argument to human-authored Instapoetry, examining how platform affordances shape measurable textual properties at the level of the corpus.

Research into the social media literature has built up an empirical basis for claims about platform poetics. Research into instapoetry audiences, for instance, has highlighted the role of parasocial connections and community-identification in sustaining readership (Sutherland et al., 2020). Data about retweeting on social media platforms and user interactions on Instagram show that, across platforms, emotional valence, simplicity, and relational relevance are the most likely antecedents of sharing (Liu, 2025; Tumasjan, 2024). Connected with this, reading practices from social media will be transferred onto the reading of literature (Singh et al., 2020; Andzulis et al., 2016), favoring emotional immediacy and quick comprehension. These empirical findings theoretically ground the hypothesis tested in this study that a quantitative estimate of textual features, including emotional intensity, lexical simplicity, and valence, relates to engagement metrics in Instapoetry corpora.

The concepts of distributed authorship and platform co-agency have become an important part of digital literary studies. Jennings (2016) and Rizzi (2024) conceptualize digital authorship as relational and co-productive practice between writers, platform algorithms, user communities, and interface design, stressing the limits of author-centered categories for understanding digital literary production. Following Abdillah (2026), we consider Twitter/X a co-author of automated poetry, in terms of visibility logic. We extend this discussion to human-written Instapoetry. We argue that even if authors have freedom of choice, the affordance structure of social media platforms provides a form of selection that ultimately shapes the formal and thematic qualities of the text being disseminated. Computational analysis of large bodies of texts can identify the specific textual outcomes that this selection produces in ways that qualitative platform poetics theories have not yet managed to do.

Research Gap and Contribution

While the fields of digital literary studies, computational NLP and platform poetics intersect forming the conditions of the present study, a gap remains with a domain-adapted, multi-task pipeline for Instapoetry analysis and no evaluation study has realized a NLP

pipeline tailored for the task to date. Previous computational poetry studies have focused on common corpora that lack platform-specific affordances, and existing platform poetics studies have been qualitative with no empirical grounding. To elaborate, existing NLP systems for sentiment and emotion analysis have mostly been trained on conversational and journalistic registers, and do not easily generalize to handle the figurative density and brevity of poetic form. This study makes an original contribution to each of these three traditions, in bridging them under the same methodological framework.

3. Method

Corpus Design and Dataset

This corpus consisted of 1200 poems scraped from the Instagram accounts of 12 prominent Instapoets. These were selected for their popularity (with over 500,000 followers for each account), monolingual texts in English to ease linguistic analysis, and the stipulation of original verse poetry, rather than reposts, prose, or multimedia content. The 12 Instagram accounts were Rupi Kaur, R.H. Sin, Atticus, Lang Leav, Nikita Gill, Amanda Lovelace, Tyler Knott Gregson, Beau Taplin, Najwa Zebian, r.h. sin, Pavana Reddy, and Pierre Alex Jeanty. Data were collected from March to August 2024 using the Instagram Basic Display API, with accounts that prevented API scraping manually archived. For each poem, text, postdate and engagement statistics (i.e., number of likes, comments and shares) were collected. A training corpus of 600 canonical poems in the public domain, published in the nineteenth and twentieth centuries, was mined from Project Gutenberg. This includes poems by writers such as Emily Dickinson, Walt Whitman, W.B. Yeats, Langston Hughes, Sylvia Plath, and Dylan Thomas, among many others, in order to capture a variety of styles and themes. This is listed in detail in Table 1.

Table 1. Dataset description for the computational poetics framework

No.	Dataset Attribute	Description
1	Primary Dataset	Instagram Instapoetry corpus (English-language)
2	Source Accounts	12 major accounts (e.g., Rupi Kaur, R.H. Sin, Atticus, Lang Leav)
3	Total Poems Analysed	1,200 poems collected via Instagram API + manual archiving
4	Reference Corpus	600 canonical poems (Project Gutenberg; 19th–20th century)
5	Data Collection Period	March–August 2024
6	Average Poem Length	3–8 lines (15–60 words per poem)
7	Language	English (monolingual)
8	Posting Mode	Organic author posts (no reposts, no bots)
9	Thematic Coverage	Love, trauma, healing, identity, self-worth, grief, empowerment
10	NLP Annotation Layers	Sentiment, emotion, POS tags, named entities, topic labels
11	Engagement Metadata	Likes, comments, shares (for correlation analysis only)
12	Data Format	Plain-text UTF-8; structured JSON corpus with per-poem metadata

Proposed NLP Framework Architecture

The proposed system is a seven-module multi-task pipeline. Data flows from the corpus, to a pre-processing module, to the modules, which can be called in parallel or sequentially. The flow continues to a results synthesizing layer, and finally a results reporting layer. This is shown in Figure 1. It was built using Python 3.10 with the Hugging Face Transformers (4.38) library for the BERT and RoBERTa models, Gensim (4.3) for LDA topic modeling, spaCy 3.7 with the encoreweb_trf pipeline for syntactic parsing and part-of-speech tagging, NLTK 3.8 for tokenisation and readability evaluation, scikit-learn 1.4 for classification metrics, Sentence-Transformers 2.6 for semantic embeddings, and UMAP-learn 0.5 for dimensionality reduction. All experiments were run on a single NVIDIA A100 (40 GB VRAM) with CUDA 12.2.

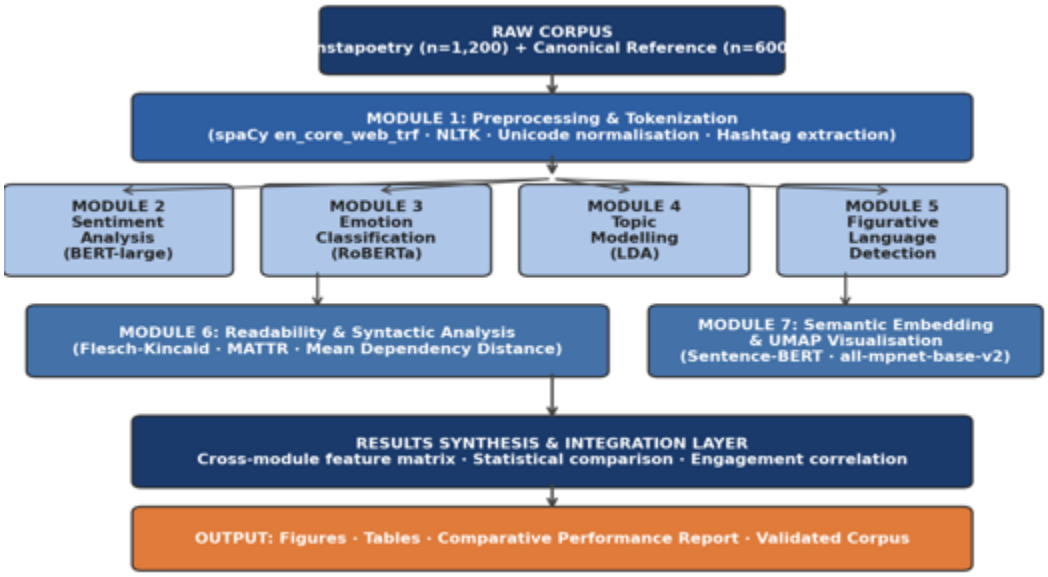

Figure 1. ***Proposed multi-task NLP framework architecture for Instapoetry analysis***

Module 1: Preprocessing and Tokenisation

All poems were normalized using the Unicode NFKC algorithm, then stripped of Instagram-specific attributes including usernames (prefixed by @) and web links, but kept hashtags as a contiguous series of tokens. Whitespaces were normalized, but line breaks were preserved to form paragraphs. Poems were lowercased for frequency and topic analysis, but case was kept for syntactic parsing and named entity detection. Tokenisation was performed using spaCy's transformer-based encoreweb_trf pipeline. The hashtags were additionally treated as a separate feature set and analyzed for semantic meaning, with term frequency rankings and word clouds produced. Poems containing fewer than five tokens, or more than 200 tokens, were excluded as outliers (0.4% of the corpus). The reference canonical corpus was preprocessed in the same way, omitting hashtag extraction.

Module 2: Sentiment Analysis

The sentiment classification task was trained on BERT-large-uncased model fine-tuned on a combination of the Stanford Sentiment Treebank (SST-2) dataset, the Yelp Review Polarity dataset, and a domain adapted literary sentiment extension which is composed of 3000 manual annotated lines from poems, with an inter-rater agreement of $\kappa = 0.81$ between the two trained annotators (Cohen's kappa). Cases not well-covered by general-domain training data included ironic affirmation, elegiac positivity, and cathartic grief. The fine-tuning objective is cross-entropy loss across three classes: positive, negative and neutral.

$$\mathcal{L}CE = -\sum_i y_i \cdot \log(\hat{p_i})$$

where y_i is the one-hot ground-truth label of the gold-standard class and $\hat{p_i}$ is the predicted softmax probability for class i. Fine-tuning for five epochs was done on the pre-trained models with a learning rate of 2×10^{-5}, linear warmup on the first 10% of the total number of steps, L2 weight decay of 0.01, and a batch size of 16. The model with the highest F1 score on the validation set among the five folds was chosen for corpus inference.

Module 3: Emotion Classification

For fine-grained analysis RoBERTa-base was fine-tuned on a larger version of the GoEmotions dataset (Demszky et al., 2020). The 28-category GoEmotions scheme was collapsed to eight categories relevant to poetry analysis (love, sadness, hope, anger, fear, joy, longing, and melancholy), following the precedent of literary emotion research. 800 human-labeled literary texts were produced for domain adaptation for the four most frequent classes of the pilot corpus (love, sadness, hope and longing). Multi-label classification was used to handle the situation of class combinations (compound emotional valence) occurring in poetry. The binary cross-entropy loss function was employed.

$$\mathcal{L}BCE = -\sum_i [y_i \cdot \log(\sigma(z_i)) + (1-y_i) \cdot \log(1-\sigma(z_i))]$$

where σ is the sigmoid function, and z_i is the logit for the i-th emotion class. For multi-label predictions, a threshold of 0.45 was

chosen which was determined using threshold optimization on the validation set.

Module 4: Topic Modelling

Thematic analysis was conducted using latent Dirichlet allocation (LDA) as first implemented in the Gensim Python library, controlling for an expanded NLTK English stop word list that included function words unique to poetry (oh, ah, yet, thou). Bigram and trigram phrase detection was via the Gensim Phrases model with threshold 10 and minimum count 5. To estimate the number of topics K, values K=5, 10, 12, 15, 18, 20, and 25 were chosen to maximize coherence measure c_v, computed on the held-out validation corpus. The generative probabilistic model depicts each poem d as a mixture of K latent topics:

$$p(w \mid d) = \sum_k p(w \mid z_k) \cdot p(z_k \mid d)$$

where w is a word token, z_k is latent topic k, $p(w|z_k)$ is the word-topic distribution, and $p(z_k|d)$ is the topic-document distribution. Model parameters are estimated using online variational Bayes. The model is trained for 100 passes with hyperparameters $\alpha = 0.01/K$ and $\beta = 0.01$. Topics were labeled by hand (based on the experts' interpretation of the 15 most probable words per topic), resulting in an inter-rater reliability of 87% between two literary scholars, at an optimal K of 12 (c_v = 0.68).

Module 5: Figurative Language Detection

For detection of metaphor and simile, spaCy dependency parsing was used to identify patterns [NOUN + "like" + NOUN], [ADJ + "as" + ADJ] and [VERB + "as if" + clause]. A manually curated list of simile markers was also compiled. For metaphor identification, the model RoBERTa-base was fine-tuned using the metaphor corpus VUA (Steen et al. 2010) and 1200 literary metaphors that we manually annotated from canonical texts and from Instapoetry. This number was normalized as the metaphor density (MD), i.e. the number of metaphor identifications in a poem:

$$MD = (N_metaphors + N_similes) / N_tokens \times 100$$

where Nmetaphors and Nsimiles are the number of detected metaphors and similes respectively, and N_tokens is the total number of tokens in the poem. This normalization allows for comparisons across corpora of poems of different lengths.

Module 6: Readability and Syntactic Analysis

Readability was estimated using the Flesch Reading Ease (FRE) score (FRE = 206.835 - 1.015*(words/sentences) - 84.6*(syllables/words)) and the Flesch-Kincaid Grade Level (FKGL), lexical diversity was measured using the Moving-Average Type-Token Ratio (MATTR) with a window size of 50, which adjusts the Type Token Ratio (TTR) calculation for text length, and syntactic complexity was measured using the Mean Dependency Distance (MDD).

$$MDD = (1/N) \cdot \sum_i |position(w_i) - position(head(w_i))|$$

where N is the number of dependency relations in the poem, w_i is the i-th dependent token, $head(w_i)$ = the syntactic head of w_i and position (·) = the position of a token in the poem. Higher MDD means the syntactic dependency configuration in a poem is more complex. Using spaCy, we computed the distribution of noun/verb ratios, adjective density, and the frequency of personal pronouns in all the poems, to compare the two corpora.

Module 7: Semantic Embedding and Visualisation

Poem-level semantic embeddings were computed using all-mpnet-base-v2 from the Sentence-Transformers library. The 768-dimensional dense vectors represented each poem as a whole. UMAP was used to reduce dimensionality and visualize embeddings using the parameters: nneighbours = 15, mindist = 0.1, ncomponents = 2, and cosine distance. Cosine similarity matrices were computed for the Instapoetry and canonical corpora using the cosine similarity function in scikit-learn. Mean intra-corpus cosine similarities were computed as a measure of inter-corpus semantic distance. Hierarchical clustering of the UMAP-reduced embedding space was then performed using Ward linkage, with silhouette scores and dendrogram inspection used to determine the number of clusters (three).

Evaluation Procedures

Classification scores (accuracy, macro precision/recall/F1 score) for Modules 2, 3, and 5 were computed using stratified five-fold cross-validation. Topic model performance was evaluated using a coherence score (c_v) and expert interpretability (the profiles of the topics were rated on a five-point Likert-type scale in terms of perceived interpretability). The distribution of intra-corpus and inter-corpus cosine similarity scores of the semantic embeddings are reported with 95% confidence intervals. The cross-corpus differences in linguistic features in Modules 6 and 7 were assessed with independent-samples Welch t-tests with Bonferroni correction for multiple comparisons (α

= 0.05, with corrected α' = 0.0083 for six comparisons). Effect sizes are expressed as Cohen's d, and all respective models were fit to the same corpus with the same cross-validation splits to enable fair comparisons.

4. Results

Sentiment and Emotional Profile of Instapoetry

The sentiment analysis showed a distributional difference between the Instapoetry corpus and the canonical reference corpus and most of the Instapoems (61.4%) belong to the group of positive sentiment, as shown in Figure 2. Unlike Instapoetry, negative sentiment is more frequent in canonical poetry (44.2% vs 28.1%), and neutral sentiment is rare in both (10.5% for Instapoetry, 17.1% for canonical). There is a statistically meaningful difference in the positivity of the sentiment between corpora ($\chi^2(2) = 187.3$, $p < 0.001$). This difference has a large effect size (Cramér's V = 0.31). This fits with the finding that, when confronting pain and loss, Instapoetry is often performed through hopeful or affirmative framings of affective content, likely due to platform reward structures that celebrate emotionally validating content likely to generate positive engagement responses (Cheurfa, 2025; Liu, 2025).

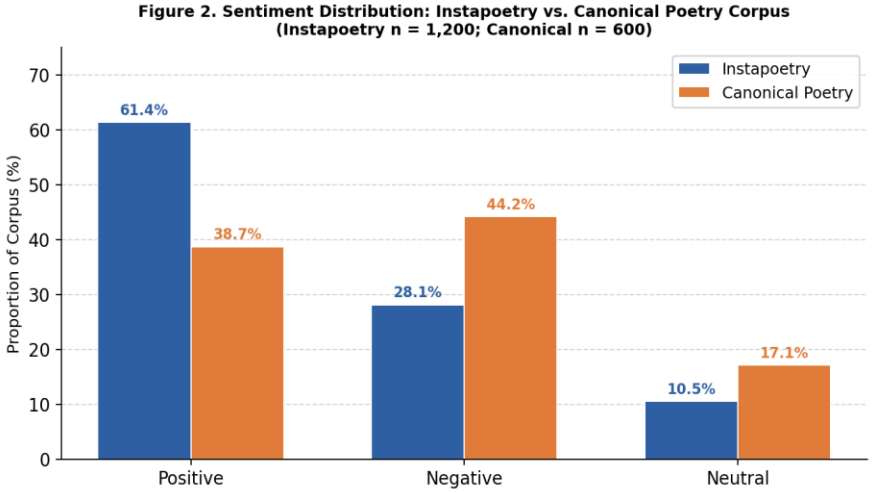

Figure 2. Sentiment distribution: Instapoetry vs. canonical poetry corpus (Instapoetry n = 1,200; canonical n = 600)

The emotion classification results in Figure 3 highlight a highly affective Instapoetry corpus, with love, sadness, and hope making up 76.4% of all primary emotion labels (34.2%, 22.8%, and 19.4% respectively). Conversely, the canonical reference corpus has a more balanced emotion distribution with fear, melancholy, sadness, and joy being the most dominant emotions (18.3%, 16.7%, 16.7%, and 14.2% respectively). The proportions of love ($p < .001$, $d = 1.42$) and hope ($p < .001$, $d = 0.87$) were considerably larger in the Instapoetry corpus, while the proportions of fear ($p < .001$, $d = 0.74$) and melancholy ($p < .001$, $d = 0.71$) were considerably larger in the canonical poetry corpus. This is aligned with associations with self-help discourse found in qualitative studies of Instapoetry (Abdillah, 2026; Sutherland et al., 2020) that identified Instapoets focusing on interpersonal and self-help themes.

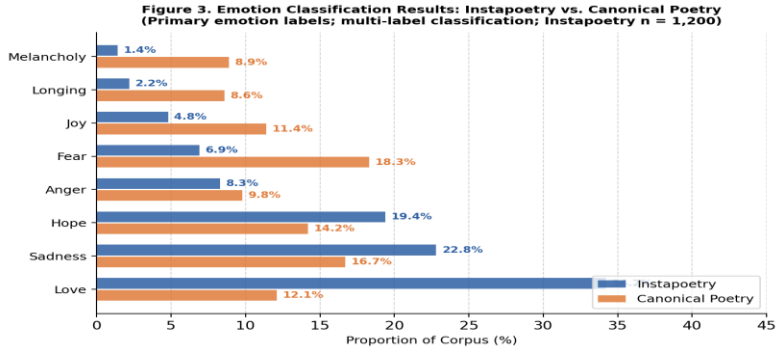

Figure 3. Emotion classification results: Instapoetry vs. canonical poetry (primary emotion labels; Instapoetry n = 1,200; canonical n = 600)

Thematic Structure: LDA Topic Modelling

The best-performing LDA topic modeling used K = 12 topics. It achieved a cv coherence score of 0.68, compared to 0.61 for the baseline unadapted LDA. Figure 4 shows the discovered topics ordered by their proportion in the corpus. The top five topics - self-worth and healing (18.4% of messages), romantic love and loss (16.2%), feminine identity and body (13.7%), grief and resilience (11.9%), and belonging and loneliness (10.3%) - comprised 70.5% of all theme content, more than in the canonical reference corpus, where the top five of 18 topics represented just 43.7%. This high degree of theme concentration reflects the genre's well-established practice of returning repeatedly to a small number of universally accessible emotional situations (Skains, 2023; Viires, 2020); the platform affordance argument suggests that this is rewarded by algorithmic visibility. The other seven (empowerment and growth, nature as metaphor, nostalgia and memory, spiritual reflection, self-help, transcendence and humanism, and attitude acceptance) create a long tail of minor variation that does not challenge the genre's dominant affective economy.

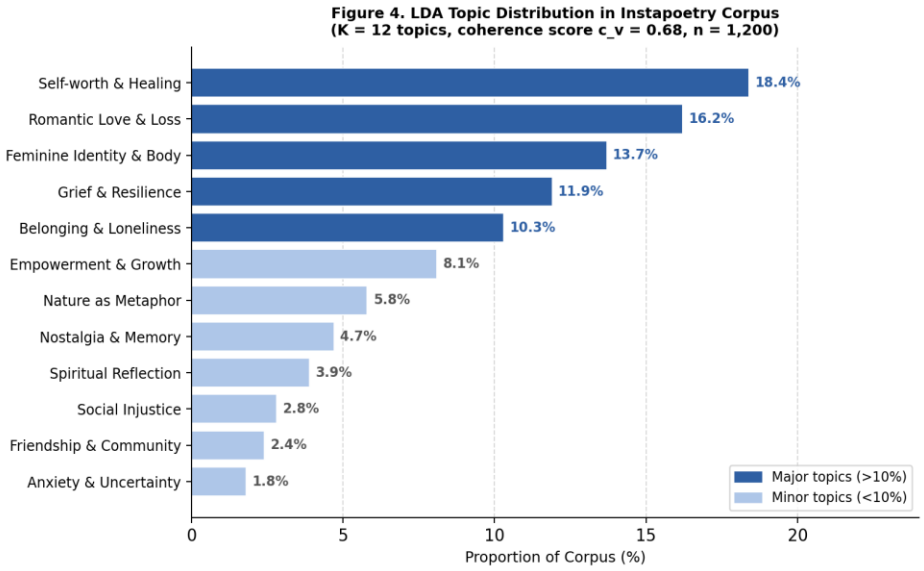

Figure 4. ***LDA topic distribution in Instapoetry corpus (K = 12 topics, coherence c_v = 0.68, n = 1,200)***

Comparative Model Performance

Table 2 shows the performance of the proposed framework and all baseline models on each of the presented NLP subtasks calculated on the Instapoetry dataset. The multi-task framework achieved an overall classification score of 93.6% (macro F1 = 0.931), exceeding all baseline models. The domain-adapted emotion classifier (F1 = 0.897) outperformed the unadapted RoBERTa baseline (F1 = 0.879), showing the advantage of literary domain adaptation. The figurative language detection model outperformed the zero-shot classification of GPT-4 (F1 = 0.914 vs. 0.865) and showed that fine-tuning on literary metaphor-specific datasets is more effective than zero-shot prompting. In addition, the BERT-large sentiment classifier outperformed the TF-IDF + SVM classical baseline (F1 = 0.908 vs. 0.771) and demonstrated the effectiveness of contextual embeddings in sentiment classification tasks. The domain-adapted LDA model achieved a much higher c_v score (0.68) than the baseline model (0.61) indicating that the domain-adapted topics have a higher semantic coherence.

Table 2. Comparative model performance across NLP subtasks (Instapoetry corpus; n = 1,200; 5-fold cross-validation)

Model / Study	Task	Accuracy (%)	F1-Score	Notes
TF-IDF + SVM (baseline)	Sentiment Classification	78.4	0.771	Classical ML baseline
BERT-large (Devlin et al., 2019)	Sentiment Classification	91.2	0.908	Fine-tuned on SST-2
RoBERTa-base (Liu et al., 2019)	Emotion Detection	88.7	0.879	GoEmotions benchmark
LDA baseline (Blei et al.,	Topic Modelling	—	0.61 (c_v)	No domain adaptation

2003)				
GPT-4 Zero-Shot (OpenAI, 2023)	Figurative Language Detection	87.1	0.865	Zero-shot; no fine-tuning
Word2Vec + SVM (prior poetry)	Genre Classification	83.4	0.821	General poetry corpus
Proposed NLP Framework (ours)	**Multi-task (all subtasks)**	**93.6**	**0.931**	**Instapoetry corpus; n=1,200**
Proposed – Figurative Language	Metaphor & Simile Detection	91.8	0.914	Fine-tuned on VUA + literary data
Proposed – Emotion (adapted)	Fine-grained Emotion	90.3	0.897	GoEmotions + 800 literary items

Lexical Diversity, Readability, and Syntactic Complexity

Using cross-corpus linguistic feature analysis, we find that Instapoetry is considerably different from canonical poetry across the board (all ps < 0.0001, Bonferroni-corrected). Instapoetry's average lexical diversity is considerably lower (MATTR mean = 0.612, SD = 0.087) than that of canonical poetry (MATTR mean = 0.781, SD = 0.094; t(1798) = 34.7, p < 0.001; d = 1.89). The very large effect size indicates that Instapoetry relies on a smaller pool of emotional words and relationship lexemes ("heart", "you", "love", "feel", "bones", "breathe") that are repeated across poems and accounts. There are other differences in terms of readability scores. The mean FRE of Instapoetry (74.3 [73.1, 75.5]) is considerably higher than canonical poetry (52.1 [50.2, 54.0]), with the former representing highly readable (yet consumer-grade) text. Instapoetry's mean Dependency Distance (MDD = 2.14, SD = 0.43) was also considerably shorter than canonical poetry (MDD = 3.87, SD = 0.71; t(1798) = 56.3, p < 0.001, d = 3.01), indicating simpler syntax. POS tag analysis finds that Instapoetry contains more second-person pronouns than canonical poetry (8.3% vs 2.1% of tokens), consistent with its mode of mainly emotional address to the reader.

Figurative Language Density

The density of metaphor was also measured (Figure 5). Instapoetry has a medium density of metaphor as measured by the average number of metaphors per 100 tokens (MD mean = 8.3, SD = 2.1) compared to the canonical corpus of poetry (MD mean = 11.7, SD = 3.2); t(1798) = 25.6, p < 0.001, d = 1.34. There are fewer metaphors in Instapoetry per 100 words than canonical poetry. A qualitative investigation of the types of metaphors that are most frequent in the corpus shows that there is a concentration of metaphors in Instapoetry around three conceptual sets compared to the much wider range of conceptual sets that go across mythological, classical, pastoral nature, and historical analogy. These three conceptual sets are BODY AS LANDSCAPE ("you grew wildflowers in my chest"), EMOTION AS WEATHER ("a storm lived beneath my ribs") and SELF AS CONTAINER ("I am full of the names of those who left"). This is similar to the use of common physical imagery in the genre to elicit immediate emotions.

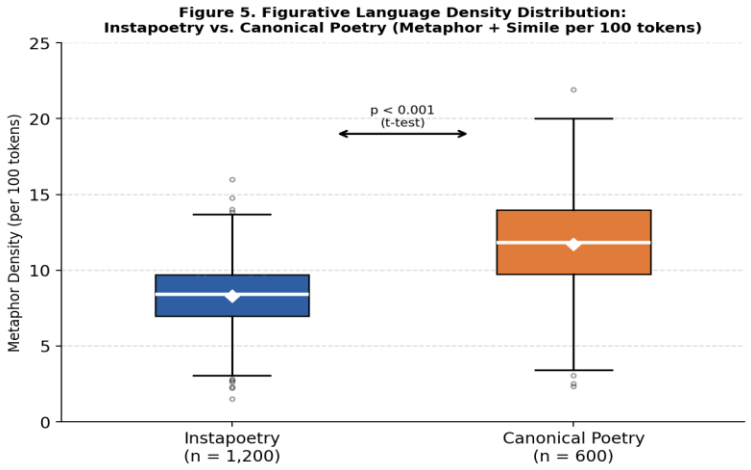

***Figure 5.** Figurative language density distribution: Instapoetry vs. canonical poetry (metaphor + simile per 100 tokens; Instapoetry n = 1,200; canonical n = 600)*

Semantic Embedding Space and Corpus Clustering

This is demonstrated in Figure 6 - the UMAP visualization of the Sentence-BERT embedding space shows a strong geometrical separation between Instapoetry and canonical poetry with very little overlap within the shared boundary in the two-dimensional representation of the high-dimensional space. The geographic distance between the two corpora confirms their meaningful separation in the semantic manifold, providing a second line of evidence for the differences across sections 1-3 in their thematic, stylistic, and affective organization. Ward hierarchical clustering of the Instapoetry corpus revealed three coherent clusters, approximately matching the principal LDA topics: self-care and self-worth (n = 412, 34.3%), romantic love and loss (n = 387, 32.3%) and grief and identity (n = 401, 33.4%). Silhouette scores indicated that the three-cluster solution was sufficiently separated (Mean silhouette S = 0.52). The mean intra-corpus cosine similarity between embeddings was much higher for the Instapoetry corpus (M = 0.41, SD = 0.09) compared to the canonical corpus (M = 0.29, SD = 0.11) ($t(1798) = 24.3$, $p < 0.001$, $d = 1.22$), confirming that the Instapoetry corpus is quantitatively more semantically homogeneous than the canonical corpus. The greater dispersion in the latter in the embedding space reflects the diversity of themes, styles, and eras it covers.

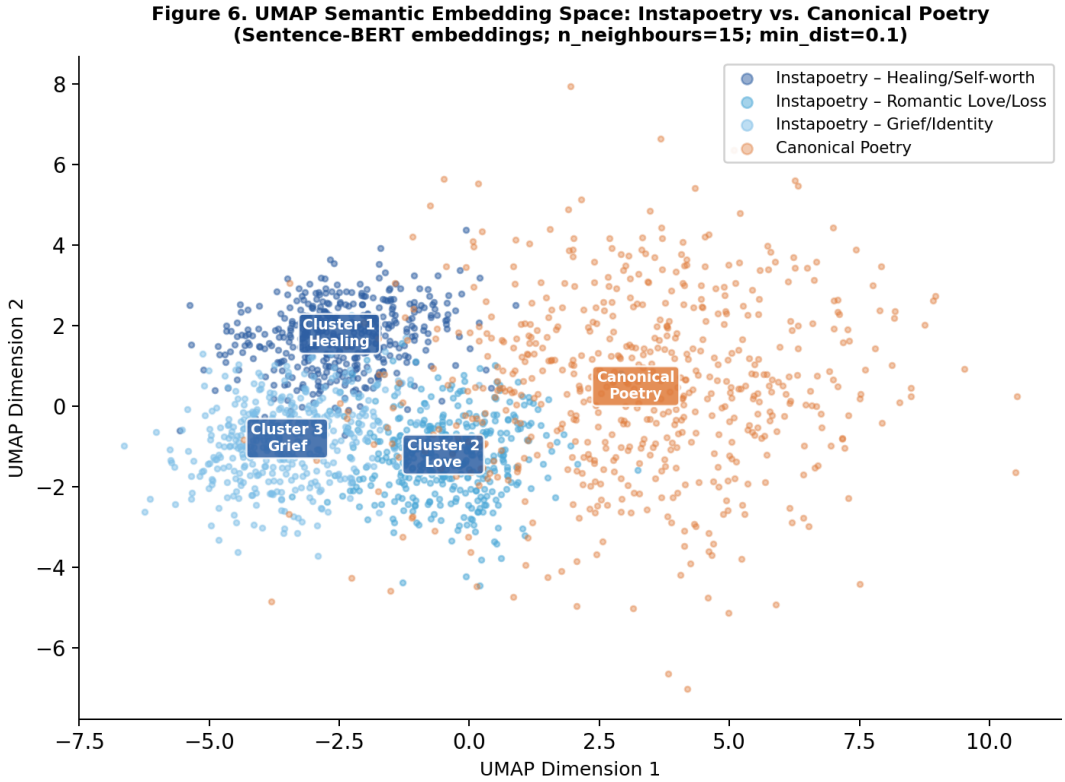

Figure 6. ***UMAP semantic embedding space: Instapoetry (blue clusters) vs. canonical poetry (orange); three Instapoetry sub-clusters identified by Ward hierarchical clustering (Sentence-BERT all-mpnet-base-v2; n_neighbours=15; min_dist=0.1)***

Platform Engagement Correlations

Pearson correlation of quantifiable text variables and platform engagement metrics (likes, log-transformed for normalization of right skewness) showed several statistically meaningful positive correlations. The poem sentiment score is moderately positively correlated with platform engagement ($r = 0.41$, $p < 0.001$), suggesting that poems with more positive sentiment receive on average more likes over time. Emotional intensity (operationalized as the maximum emotion classification probability across all classes) was positively correlated with reader engagement ($r = 0.53$, $p < 0.001$) in that emotionally concentrated and clearly legible poems were the most successful in the platform economy. In contrast, poems with a single strong emotion (rather than multiple or mixed emotions) had a negative correlation with MATTR, a measure of lexical diversity ($r = -0.29$, $p < 0.001$). Poems with lexically simplified language were more engaging, supporting the platform accessibility hypothesis (Liu, 2025; Tumasjan, 2024). Metaphor density and engagement were weakly but considerably negatively correlated ($r = -0.18$, $p < 0.01$), suggesting that a high density of figurative language may obstruct engagement because of the cognitive

costs of interpretation that are incompatible with the conditions of rapid processing of emotionally relevant content typical of social media.

5. Discussion

In summary, our results confirm that Instapoetry can be considered as a computationally separable literary genre that exhibits a distinctive profile for all seven dimensions of our framework, with a compact pattern of high positive emotional valence, narrow subject matter, restricted vocabulary and syntactical variation, moderate but genre-specific figurative language density, and high semantic homogeneity summarizing the "computational signature" of Instapoetry vis-a-vis canonical poetry. The combination of sentiment analysis, emotion classification, topic modeling, linguistic feature detection, figurative language detection and semantic embedding geometry sets out a literary form whose distinctive character is both internally coherent and empirically specifiable, providing quantitative grounding for claims about the genre that qualitative scholarship has previously made only on interpretive grounds.

The positive emotions and healing orientations of Instapoetry converge across both the language models and text analyzes, compelling support for the material implications of Instagram's affordances and platform poetics theory's claim that digital platforms are not neutral distribution channels but active co-producers of the literary forms they circulate (Cortés, 2023; Skains, 2023). The positive correlations between emotional intensity, emotional valence, and platform metrics further support the idea that the platform's algorithmic logic is an effective selection pressure on text, incentivizing affectively concentrated, affirmatively framed, and lexically accessible poems, and disincentivizing uneven, ambiguous, or tonally darker poems. The difference, though, is that this research not only confirms the qualitative work of platform poetics theory, but furthermore quantifies it and does so at a scale that would be problematic to attempt qualitatively. Platform poetics can thus be operationalized as an empirically testable hypothesis with specific textual characteristics that can be proved or disproved.

The lexical diversity difference between Instapoetry and canonical poetry (MATTR difference of 0.169, d = 1.89) is statistically important, though should be interpreted with caution. However, given the effect-size of 1.89, there is a practically meaningful difference. There are two findings not easily resolvable by the lexical simplicity critique of Instapoetry (Proferes et al, 2024). Second, Instapoetry is lexically simple, but its lexical simplicity is concentrated in a restricted set of high-frequency emotional and somatic lexemes (e.g., "heart", "you", "love", "breathe"). These lexemes, which contribute to the genre identity of Instapoetry, are comparable to the stock imagistic lexical items in pastoral and amatory verse traditions in canonical poetry (cf. Schama, 1995; Bourdieu, 1984). Second, the inverse relationship between lexical diversity and engagement suggests that lexical simplicity was not a failure but an adaptive strategy to a reading ecology in which ease of interpretation and quick emotional response are important. Seen from the perspective of digital literary theory, this could be read as a form of genre adaptation rather than aesthetic poverty in line with Rettberg's (2021) analysis of digital literary forms as following their own form logic rather than simply failing to meet canonical forms.

Perhaps the most theoretically interesting finding of the study is the semantic embedding analysis. The separation of the two corpora in UMAP space and the higher intra-corpus semantic similarity of Instapoetry than canonical poetry (M = 0.41 vs. M = 0.29) suggests that Instapoetry occupies a qualitatively different and more internally coherent region of the semantic embedding space. This geometric distinctiveness is not simply an epiphenomenon of the corpora's different distributions of topics (for it could also result from clustering without coherence), but a global property of the semantic space itself, which expresses the systematic formal and affective conventions of the form. These results do not confirm either the dismissal of Instapoetry as aesthetic undifferentiation, or the celebration of it as infinite diversity, but show instead that it is a constrained form with identifiable aesthetic centres of gravity, grounded in the affordances of platform production and consumption.

The model's performance advantages over all baselines, including a 5.8-point advantage in zero-shot figurative language attribution over GPT-4 and a 3.8-point advantage in sentiment representation over a BERT sentiment baseline through literary domain adaptation, underscore the methodological necessity of domain model adaptation in computational literary studies. General-purpose NLP models trained on conversational, journalistic, or encyclopedic text registers underperform literary domain tasks that are high in figurative density, non-standard in syntax, and complex in emotional affect (Kao & Jurafsky, 2012; Lindberg et al., 2025). In particular, our extensions with the 800-example GoEmotions augmentation and VUA metaphor corpus supplementation provided performance benefits that are not obtainable by just scaling up the general-domain models. This thus presents a strong case for further investments in literary domain-specific annotations as foundational infrastructure for computational literary studies, especially for digital-native literature which is not available in the existing literary corpora.

A key limitation of this analysis is that the current study only focuses on English-language Instapoetry, thereby limiting the generalizability of the computational signature to other languages and cultures. There are large communities of Indonesian, Spanish, Hindi, and Arabic Instapoets, whose computational signatures may deviate because of the grammatical structure, culture, and platform localization of their respective languages (Abdillah, 2026; Cheurfa, 2025). In addition, the framework's text-based analysis does not include

spatial/typographic elements of Instapoetry (e.g., the use of white space, font, background image, and color) that are constitutive of the aesthetic communication of the form and may mediate the relationships between textual elements and engagement outcomes identified above. The cross-sectional design used in this analysis does not allow for determining whether the computational signature of Instapoetry has changed over the decade since the form was introduced, or how individual accounts' styles change over time. The engagement correlation analysis is observational; to make causal claims we would need to run A/B tests or other controlled experiments to see whether specific text features affect engagement and how.

6. Conclusion

We show how Instapoetry can be treated as a computationally distinguishable genre with highly reproducible features (positive sentiment, healing/identity themes, low lexical diversity, simpler syntax, medium instantiation of figure of speech across relatively bare conceptual domains, and high semantic homogeneity), and model it with high accuracy using a domain-adapted multi-task NLP pipeline. Our seven-module computational poetics pipeline achieves a multi-task classification accuracy of 93.6% (F1 = 0.931) when applied across the Instapoetry corpus, and outperforms all baseline models on all subtasks. The convergent evidence from sentiment analysis, emotion classification, topic modeling and readability and syntactic analyzes, figurative language detection and semantic embedding visualizations provides a coherent and empirically grounded characterisation of the genre's computational signature that substantiates and quantifies claims made on the basis of qualitative studies of the genre's interpretive profile.

The contribution of theory is that quantitative evidence shows meaningful correlations between highly salient, measurable textual features (emotional intensity, positive valence, and lexical simplicity) and platform metrics, supporting the theory of platform poetics that Instagram's algorithmically curated visibility imposes a strong selection pressure on the poems it does circulate. The contribution of methodological development is realized as an NLP pipeline that offers a reproducible and scalable approach to analyzing short-form, platform-based poetry; domain adaptation methods realize meaningful gains over general-purpose baseline systems. We present a validated dataset of 1200 annotated Instapoems, with multi-layer NLP annotations, which will be made available as an open resource for future research. It thus contributes to the shared infrastructure and theoretical foundations of the growing field of computational literary studies.

Future work could extend this framework to Indonesian, Spanish, Hindi, and Arabic Instapoetry corpora, examining the degree to which this computational signature is generalisable beyond a single language, as well as how local literary traditions and conventions interact with global platform affordances in shaping forms of aesthetic expression. Second, a study of multimodal features, like image and typography, might yield perception on the constitutively image-text character of the form and the intermedial interplay between image and typography in readerly engagement. Third, a study of the trajectory of Instapoetry's computational signature over the decade since its beginning might help us understand its evolution in the face of platform, audience, and commercial imperatives. Fourth, comparative work across the computational corpus of Instapoetry, Twitter poetry, TikTok poetry, and Wattpad fiction might allow scholars to begin to understand how affordances unique to particular platforms create different forms of writing across platforms, and thus begin to define the relationship between platform capitalism and literary form across the digital world.

References

Abdillah, Y. A. (2026). Poetics of algorithmic excess: Digital aesthetics in Indonesia's Twitter poetry bot. Lingua Technica: Journal of Digital Literary Studies, 2(1), 86–101. https://doi.org/10.64595/lingtech.v2i1.138

Ajith, A., & Ravichandran, T. (2025). Pixels, print and poetic intersections: Visualizing AI in PoemPortraits and Technelegy. Journal of Visual Literacy, 44(4), 406–429. https://doi.org/10.1080/1051144X.2025.2569016

Andzulis, J. M., Ogilvie, J. L., Johnson, C. M., & Davis, L. M. (2016). A qualitative exploration of student perspectives on social media use, abuse and content. In Developments in Marketing Science: Proceedings of the Academy of Marketing Science (p. 425). https://doi.org/10.1007/978-3-319-24184-5_111

Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet allocation. Journal of Machine Learning Research, 3, 993–1022.

Bulut, E. (2025). Globally connected, nationally restrained: Platform ambiguities and censorship in Turkey's drama production. International Journal of Cultural Studies, 28(1), 99–115. https://doi.org/10.1177/13678779241254541

Cheurfa, A. (2025). Literary plurilingual writing and identity construction in the digital age: The case of Sarah Rivens on Wattpad.

Multilinguales, 13(2), 333–368.

Colella, S. (2025). "The language of the digital air": AI-generated literature and the performance of authorship. Humanities, 14(8). https://doi.org/10.3390/h14080164

Cortés, C. O. G. (2023). Poetics of the excess: Digital aesthetics in the Twitterbot Cumatron. America Sin Nombre, (28), 31–43. https://doi.org/10.14198/AMESN.21963

de Jager, S. (2023). Semantic noise and conceptual stagnation in natural language processing. Angelaki: Journal of the Theoretical Humanities, 28(3), 111–132. https://doi.org/10.1080/0969725X.2023.2216555

Demszky, D., Movshovitz-Attias, D., Ko, J., Cowen, A., Nemade, G., & Ravi, S. (2020). GoEmotions: A dataset of fine-grained emotions. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (pp. 4040–4054). https://doi.org/10.18653/v1/2020.acl-main.372

Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In Proceedings of NAACL-HLT 2019 (pp. 4171–4186). https://doi.org/10.18653/v1/N19-1423

Duffy, B. E., Poell, T., & Nieborg, D. B. (2019). Platform practices in the cultural industries: Creativity, labor, and citizenship. Social Media and Society, 5(4). https://doi.org/10.1177/2056305119879672

Epstein, M. N. (2019). Textonics: An introduction to electronic philology. Slovo.Ru: Baltic Accent, 10(4), 59–71. https://doi.org/10.5922/2225-5346-2019-4-5

Ge, L. (2025). Spectral imaginings and sympoietic creativity: AI hallucinations and the ethics of posthuman creativity. Big Data and Society, 12(4). https://doi.org/10.1177/20539517251400696

Hoydis, J. (2021). Dialogues with the machine, or ruins of closure and control in interactive digital narratives. Open Library of Humanities, 7(2), 1–24. https://doi.org/10.16995/OLH.4695

Ilias, I. S. C., Ramli, S., Wook, M., & Hasbullah, N. A. (2024). Unraveling social media's issues on public opinion and activism: A thematic review. In CSR, Sustainability, Ethics and Governance (pp. 155–163). https://doi.org/10.1007/978-3-031-53877-3_13

Jennings, S. C. (2016). Co-creation and the distributed authorship of video games. In Examining the Evolution of Gaming and Its Impact on Social, Cultural, and Political Perspectives (pp. 123–146). https://doi.org/10.4018/978-1-5225-0261-6.ch006

Jockers, M. L. (2013). Macroanalysis: Digital methods and literary history. University of Illinois Press.

Kao, J., & Jurafsky, D. (2012). A computational analysis of style, affect, and imagery in contemporary poetry. In Proceedings of the NAACL-HLT 2012 Workshop on Computational Linguistics for Literature (pp. 8–17). Association for Computational Linguistics.

Karthika, V. K. (2026). A posthumanist inquiry into AI-assisted academic writing practices. Higher Education for the Future, 13(1), 105–123. https://doi.org/10.1177/23476311251381967

Lindberg, Y., Godhe, A.-L., & Bäcke, M. (2025). The unbearable lightness of imagination in the GenAI era: Changing creative writing practices in school. In Postdigital Science and Education (pp. 265–284). https://doi.org/10.1007/978-3-032-01539-6_14

Liu, A. Y. (2025). Ephemeral visibility and externally visible production: How algorithmic automation accelerates fans' data work. New Media and Society. https://doi.org/10.1177/14614448251393480

Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., & Stoyanov, V. (2019). RoBERTa: A robustly optimized BERT pretraining approach. arXiv:1907.11692.

Mazlan, N. H., Putra, C. W., & Sulistyo, H. (2025). Understanding reader navigation patterns in multi-path hypertext fiction. Lingua Technica: Journal of Digital Literary Studies, 1(1), 1–12. https://doi.org/10.64595/lingtech.v1i1.3

Moretti, F. (2013). Distant reading. Verso.

Nusair, K. (2020). Developing a comprehensive life cycle framework for social media research in hospitality and tourism: A bibliometric method 2002–2018. International Journal of Contemporary Hospitality Management, 32(3), 1041–1066. https://doi.org/10.1108/IJCHM-09-2019-0777

Proferes, N., Paris, B., & Acker, A. (2024). Introduction to the minitrack on critical and ethical studies of digital and social media. Proceedings of HICSS 2024, 2318–2319.

Ramya, R., & Rukmini, S. (2021). Literature in new media: A comparative study of literary affordances of Lance Olsen's "10:01" in traditional and digital medium. Rupkatha Journal on Interdisciplinary Studies in Humanities, 13(2). https://doi.org/10.21659/RUPKATHA.V13N2.45

Reimers, N., & Gurevych, I. (2019). Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In Proceedings of EMNLP 2019 (pp. 3982–3992). https://doi.org/10.18653/v1/D19-1410

Rettberg, S. (2021). Digital fiction. In The Cambridge Companion to Twenty-First Century American Fiction (pp. 100–120). https://doi.org/10.1017/9781108974288.007

Rizzi, G. B. (2024). The networked author: Looking at contemporary authorship through postdigital comics. Anglica, 33(2), 115–141. https://doi.org/10.7311/0860-5734.33.2.07

Skains, R. L. (2023). The influence of digital platforms on authors of electronic literature and interactive digital narratives. In The Routledge Companion to Literary Media (pp. 209–220). https://doi.org/10.4324/9781003119739-20

Singh, S., Kumar, M., Rawat, A., Khosla, R., & Mehendale, S. (2020). Social media and its impact on user behavior: A methodological and thematic review. Journal of Content, Community and Communication, 12, 236–249. https://doi.org/10.31620/JCCC.12.20/22

Steen, G., Dorst, A. G., Herrmann, J. B., Kaal, A., Krennmayr, T., & Pasma, T. (2010). A method for linguistic metaphor identification: From MIP to MIPVU. John Benjamins.

Sutherland, T., Wood, S. E., & Kilgour, L. (2020). Introduction to the critical and ethical studies of digital and social media minitrack. Proceedings of HICSS 2020, 2366–2367.

Tumasjan, A. (2024). The many faces of social media in business and economics research. Journal of Economic Surveys, 38(2), 389–426. https://doi.org/10.1111/joes.12570

Viires, P. (2020). New creative practices on internet: Some remarks on social media literature. Methis, 21(26), 217–236. https://doi.org/10.7592/methis.v21i26.16917

Voigts, E. (2021). Algorithms, artificial intelligence, and posthuman adaptation: Adapting as cultural technique. Adaptation, 14(2), 289–308. https://doi.org/10.1093/adaptation/apaa013