Transactions on Artificial Intelligence
https://www.sciltp.com/journals/tai

Article

The Measure of Resonance: A Computational Cross-Cultural Analysis of Floral Poetics in Tang and Renaissance Verse

Tiange Zhou [1,*,†] and Marco Bidin [2,†]

[1] School of Future Design, Beijing Normal University, Zhuhai 519000, China
[2] Instrument Engineering Department, College of Music Technology, Xinghai Conservatory of Music, Guangzhou 510000, China
* Correspondence: tiangezhoumusic@gmail.com
† These authors contributed equally to this work.

How To Cite: Zhou, T.; Bidin, M. The Measure of Resonance: A Computational Cross-Cultural Analysis of Floral Poetics in Tang and Renaissance Verse. *Transactions on Artificial Intelligence* **2025**, *1*(1), 326–339. https://doi.org/10.53941/tai.2025.100022

Received: 19 September 2025
Revised: 20 October 2025
Accepted: 6 November 2025
Published: 12 December 2025

Abstract: The systematic comparison of pre-modern literary traditions demands a methodology capable of balancing historical specificity with analytical rigor. This study constructs an integrative framework that bridges deep hermeneutic traditions and computational literary analysis. Focusing on the peony in Tang-dynasty China (618–907 CE) and the rose in Renaissance England (ca. 1500–1660 CE), we situate these emblematic flowers within Janet Abu-Lughod's paradigm of the pre-modern Afro-Eurasian "world system", establishing their comparability as distinct yet interconnected cultural formations. To move beyond impressionistic analogy, we introduce a computational transposition of the logic of cultural dimensions research—the translation of abstract values into measurable axes of variation—from sociological surveys to poetic language. Treating the peony and the rose as Ernst Cassirer's symbolic forms, we employ a multi-method computational triangulation—multilingual BERT embeddings for semantic resonance, LDA topic modeling for thematic structure, and transformer-based emotion analysis for affective tonality—to model their symbolic architectures across a bilingual corpus of 49 Tang peony poems and 45 Renaissance rose sonnets after quality control and alignment constraints. Results reveal a shared preoccupation with beauty and transience as a universal axis of meaning. Yet a significant negative correlation (r = -0.128, $p < 0.05$) between semantic and thematic alignment suggests a pattern of structural complementarity rather than direct similarity: Tang and Renaissance poetics converge in emotional and aesthetic concerns while diverging in symbolic framing. The study thus proposes both a substantive insight into cross-civilizational poetics and a transferable methodological paradigm—one that complements close reading with empirical scalability and structural clarity in the comparative study of the world literature.

Keywords: digital humanities; computational poetics; cross-cultural analysis; symbolic forms; tang poetry; renaissance sonnets; world-systems

1. Introduction: From Symbolic Form to Data Form

The comparative study of pre-modern literary traditions remains a central yet methodologically vexed pursuit in the humanities. While scholars have long juxtaposed cultural motifs across civilizations, such work often oscillates between philological depth and empirical generalization. This paper seeks to reframe the problem by proposing a methodology that integrates interpretive precision with computational scalability, enabling systematic rather than anecdotal comparison. Our inquiry focuses on two emblematic poetic cultures: the peony poetry of Tang-dynasty China (618–907 CE) and the rose sonnets of Renaissance England (ca. 1500–1660 CE). These traditions, though geographically distant, share structural positions within what Janet Abu-Lughod (1989) conceptualized as the premodern world system: a decentralized network of Afro-Eurasian exchange that preceded European hegemony [1].

Within this macrohistorical framework, Tang China and Renaissance England can be understood not as isolated entities but as distinct open cycles' in a connected ecumenical of goods, ideas and aesthetic forms. This framework provides the historical justification for cross-civilizational comparison while avoiding the pitfalls of methodological nationalism.

1.1. Symbolic Forms and Cultural Condensation

Within their respective poetic worlds, the peony and the rose crystallized the most charged cultural negotiations of beauty, virtue, and mortality. Following Ernst Cassirer's (1944) notion of symbolic forms [4], these flowers are treated not merely as decorative motifs but as cognitive structures that organize human experience. The peony—celebrated in Tang poetics as the "king of flowers" (花王)—embodied the aesthetic synthesis of imperial grandeur and transience, mediating Confucian ideals of order and Daoist-Buddhist reflections on impermanence. The rose, conversely, anchored the Renaissance imagination through a dense network of associations—erotic desire, political identity, divine grace—extending from the Roman de la Rose to Shakespearean sonnets. Each flower thus functioned as a cultural condensation, a node where moral, aesthetic, and cosmological values intersected. Traditional comparative studies have often relied on impressionistic parallels between such symbols. While illuminating, these approaches lack a means to demonstrate structural patterns of resonance or divergence across entire poetic systems. To address this limitation, we translate Cassirer's theoretical construct into a computationally operational form [4].

1.2. Computational Transposition: From Cultural Dimensions to Textual Axes

The paper introduces a computational transposition of the methodological logic behind Hofstede-style cultural dimensions research [2,3]. Just as sociologists quantify latent cultural values by identifying axes of variation (e.g., individualism vs. collectivism), we apply the same principle to poetic language. In this adaptation, our "survey data" are the poems themselves, and our "dimensions" emerge from the semantic, thematic, and affective regularities that computational models can detect. Specifically, we employ a multi-method triangulation:

1. Multilingual BERT embeddings [5] to map semantic resonance between corpora in a shared vector space;
2. LDA topic modeling [6] to identify latent thematic structures;
3. Transformer-based emotion analysis to evaluate affective polarity and intensity.

This design follows the principle of methodological triangulation [7]: no single model can fully capture cultural nuance, especially in classical languages. By integrating distinct analytic lenses, we seek convergence among independent evidentiary sources, thereby strengthening both validity and interpretive depth.

1.3. Research Question and Contribution

Our central question is thus twofold: (1) How can we empirically identify the multidimensional axes along which Tang and Renaissance poetics resonate and diverge? (2) What might these axes reveal about the deeper structures of symbolic expression that connect humanistic traditions across civilizations? We hypothesize that resonance emerges not through simple similarity but through structural complementarity—a relation in which distinct symbolic systems articulate parallel concerns through divergent semiotic grammars. The analysis yields a measurable though subtle inverse correlation between semantic and thematic alignment (r = -0.128, $p < 0.05$), suggesting that when poems converge lexically, they often diverge thematically, and vice versa. This counterintuitive pattern supports the broader claim that pre-modern cultural affinity operates multidimensionally, mediated by both shared historical conditions and differentiated symbolic economies.

1.4. Toward a Computationally-Enabled Comparative Poetics

Beyond its specific findings, the study contributes a transferable methodological paradigm for cross-cultural literary analysis. By operationalizing Cassirer's symbolic forms through the logic of measurable variation, it demonstrates how computational models can serve as scaffolds for interpretive reasoning rather than substitutes for it. In doing so, it positions computational poetics as a bridge between close reading and distant reading—a mode of analysis that preserves the granularity of literary meaning while enabling empirical generalization across corpora. This approach ultimately re-centers digital humanities within the comparative project of world literature: not as a technological adjunct, but as a critical instrument for re-imagining how aesthetic traditions resonate across time, language, and culture.

2. Literature Review

2.1. The Computational Turn and the Reconfiguration of Literary Method

The digital humanities have redefined how scholars approach literary comparison, transforming what Franco Moretti (2013) [8] termed distant reading from a provocation into a methodological paradigm. Within the emerging field of Computational Literary Studies (CLS), algorithms function not merely as analytic tools but as epistemic devices that reveal large-scale formal and semantic structures invisible to traditional reading [7]. Yet, the computational turn in literary scholarship is not a departure from hermeneutics but a reconfiguration of it. As many scholars have recognized, computation can serve as a new scaffolding for interpretation—a way of tracing macro-patterns that inform, rather than replace, close reading [9,10]. This study situates itself within this integrative orientation: it uses computational models to generate evidence for cultural patterning that subsequently demands interpretive elaboration. Our goal is not to render poetry "measurable" for its own sake, but to formalize the intuitions of comparative poetics through empirically defensible structure.

2.2. Computational Poetics and the Question of Cultural Validity

The application of computational methods to poetic analysis—sometimes referred to as computational poetics—has expanded rapidly in the last decade. Early stylometric work emphasized metrics of authorship and form [11,12], while recent scholarship employs transformer-based language models [5] to study imagery, semantics [13], and emotional cadence across literary corpora. These approaches have yielded insights into genre, authorship, and affect, yet they also raise a central epistemological question: To what extent can statistical representations of language capture culturally specific poetics? This challenge is particularly acute for cross-cultural and historical corpora. Pre-trained multilingual models such as BERT and its derivatives are trained predominantly on modern, Western-dominated data, making them sensitive to cultural and linguistic bias when applied to Classical Chinese or early modern English. Addressing this issue requires not abandonment, but methodological reflexivity, the recognition that each model encodes a limited lens. Our study directly engages this problem through a framework of triangulated validation. By combining three distinct methods—multilingual embeddings for semantic proximity, LDA for thematic structure, and emotion modeling for affective polarity—we allow the limitations of one technique to be offset by the complementary strengths of others. In this sense, triangulation functions not only as a computational design choice but as a critical epistemology, one that foregrounds the situatedness of algorithmic interpretation and the necessity of cross-validation in culturally sensitive DH research.

2.3. Historical Grounding: Polycentric Worlds and the Conditions of Comparison

While computational methods provide the analytic machinery of this study, historical comparability provides its justification. The decision to juxtapose Tang and Renaissance poetry is not grounded in direct influence but in their shared participation within what might be called polycentric world systems—networks of exchange and imagination that connected distant cultures before modern globalization. To frame this relationship, we draw conceptually on Janet Abu-Lughod's (1989) model of the thirteenth-century world system [1]. Although her framework explicitly concerns the 1250–1350 CE period—after the Tang and before the Renaissance—its analytic value extends beyond that chronology. Abu-Lughod's central insight is that the premodern world was not fragmented into isolated civilizations but organized through multiple interacting centers of cultural and economic vitality. The Tang dynasty (618–907), often regarded as the apex of classical Chinese cosmopolitanism, may be read as an earlier manifestation of this polycentric dynamic: its Silk Road networks linked East Asia to Central Asia, Persia, and the Mediterranean world, facilitating flows of material goods, religions, and symbolic forms. Similarly, Renaissance Europe reactivated and reimagined these Eurasian circuits under new technological and ideological conditions. By invoking Abu-Lughod's polycentric logic, rather than her specific chronology, we can interpret Tang and Renaissance poetic cultures as distinct but structurally resonant responses to expanding worlds of contact. Each poetic system translated external encounters into internal symbolic vocabularies: the Tang peony as an emblem of imperial splendor and transient beauty within an open, hybrid empire; the Renaissance rose as a Christian-humanist symbol negotiating the sensual and the sacred within a rediscovered classical order. Thus, the world system here functions not as a historical container but as a comparative horizon—a theoretical device that allows us to situate symbolic production within connected, though asynchronous, ecologies of exchange.

Beyond this macro-historical connectedness, a deeper, functional parallel justifies the comparison: both the Tang and Renaissance periods represent formative moments of cultural synthesis within their respective civilizations. The Tang integrated diverse influences via the Silk Road, crystallizing a classical Chinese aesthetic. The Renaissance, in turn, synthesized Greco-Roman, Christian, and emerging humanist thought. Thus, we compare not two random eras, but two pivotal symbolic-formative phases in which their respective poetic canons were being codified. This functional parallel provides a stronger justification for comparison than mere chronological proximity would

2.4. Symbolic Forms as Analytical Bridge

While world-systems theory establishes a macro-historical foundation, the concept of symbolic form—derived from Ernst Cassirer (1944)—provides the micro-analytical bridge between history and poetics. For Cassirer, symbolic forms are not passive representations but active mediations through which human beings structure meaning. Literature, in this sense, constitutes one of the highest orders of symbolic articulation, where cultural values condense into recurrent images and metaphors. Within this theoretical frame, the peony and the rose emerge as cognitive-symbolic operators: each embodies a historically situated negotiation of beauty, virtue, and impermanence. The peony, in Tang poetics, mediates tensions between imperial splendor and Buddhist ephemerality; the rose, in Renaissance verse, encodes the dialectic between sensuality and spiritual grace. Treating these motifs as symbolic forms thus allows for a transhistorical yet rigorously contextual analysis—one that is not reducible to botanical coincidence but grounded in the logic of cultural condensation. Our methodological innovation lies in operationalizing this philosophical notion: transforming Cassirer's qualitative insight into a measurable architecture of semantic, thematic, and affective variation. In doing so, we articulate a mode of computational hermeneutics that extends symbolic analysis into quantitative space while maintaining interpretive depth.

2.5. From Hermeneutics to Computation: Toward a Comparative Framework

Finally, this study addresses a persistent limitation in comparative literary scholarship: the reliance on selective textual exempla to substantiate claims of cultural similarity or difference. Traditional comparatism has often privileged analogical reading—juxtaposing a few canonical poems to argue for shared sensibility—without mechanisms to generalize findings beyond anecdote. Our framework proposes a way forward. By translating Cassirer's symbolic form into a computational model inspired by Hofstede's (1980) logic of cultural dimensions, we identify measurable axes of poetic variation that can be compared across entire corpora. This approach does not displace hermeneutic reading; rather, it extends its evidentiary base, allowing interpretation to operate at both micro and macro scales. In doing so, we advance what might be termed a computationally enabled comparative poetics—a methodological synthesis that reconciles the interpretive granularity of traditional literary scholarship with the structural clarity and empirical reach of computational analysis. This hybrid approach, we argue, provides the most robust path toward understanding how civilizations connected by trade, translation, and imagination articulate shared human concerns through distinct symbolic grammars.

3. Materials and Methods

3.1. Corpus Design and Preprocessing

3.1.1. Sampling Frame and Accounting

We initially sampled 50 Tang (peony-themed) poems and 50 Renaissance (rose-themed) poems. After preregistered cleaning (deduplication, OCR/encoding failures, minimum-length filtering, preprocessing failures, and cross-lingual alignment checks), we retained 49 Tang and 45 Renaissance items, yielding $49 \times 45 = 2205$ cross-collection pairs for all downstream analyses.

3.1.2. Temporal and Genre Constraints

The Tang set covers 618–907 CE with regulated verse and related genres; the Renaissance set covers ca. 1500–1660 with sonnets and closely related lyric forms. We exclude items outside these windows or misclassified by genre.

3.1.3. Normalization and Segmentation

All texts are normalized to UTF-8, lowercased for EN (case-preserving for ZH), and stripped of editorial punctuation not intrinsic to the text. ZH texts are segmented with Jieba plus a classical stop-list (particles, function words); EN texts are tokenized and lemmatized. We apply a preregistered post-stopword token threshold τ to ensure sufficient content per document.

3.1.4. Phrase-Level Chunking

To respect parallelism and reduce sentence-length confounds in premodern verse, we operate on phrase-level chunks extracted by punctuation/line breaks and conjunction heuristics. Chunk embeddings (for the semantic layer) are averaged within a poem; statistics are reported at the poem level, not at the chunk level.

3.1.5. Metadata, Deduplication, and Quality Control

We harmonize author/title/year fields; near-duplicates are identified by high lexical cosine and manual inspection and then removed. OCR/encoding failures and items with missing core fields are excluded. All preprocessing rules and seeds were preregistered; code is released in the Supplementary Materials (also detailed in Appendix A.2).

3.2. Analytical Architecture

We operationalize Cassirer's symbolic form as a tri-layer design: (a) semantic similarity, (b) thematic structure, (c) affective tone. Each layer outputs a $N_T \times N_R$ similarity matrix (Tang $\times$ Renaissance), which are then normalized and fused (Figure 1).

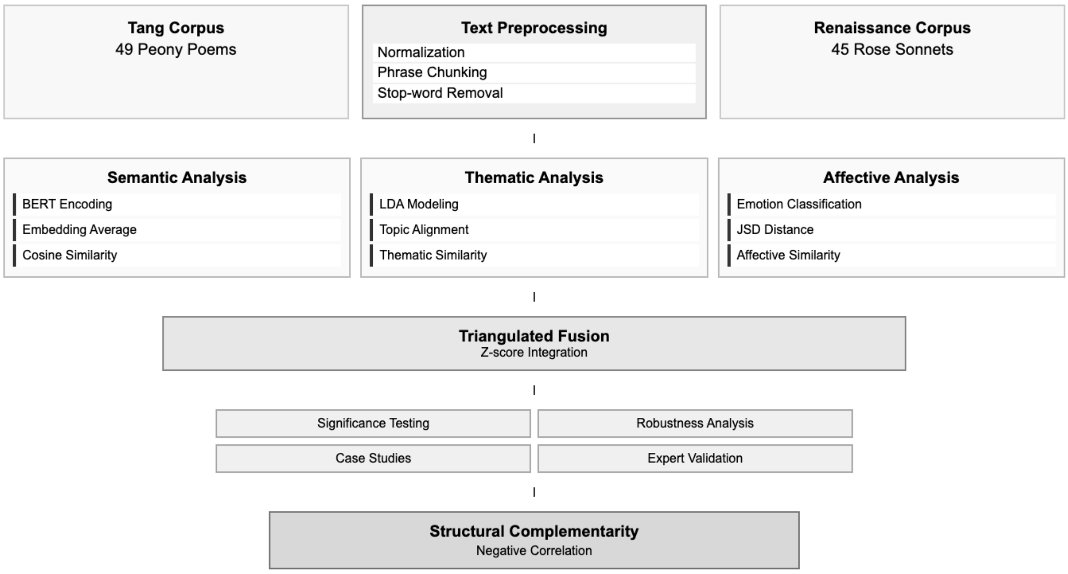

Figure 1. Pipeline: corpus $\rightarrow$ preprocessing $\rightarrow$ BERT semantics, LDA themes, RoBERTa emotions $\rightarrow$ z-score fusion $\rightarrow$ validation & micro case interpretation.

3.2.1. Semantic Similarity: Multilingual BERT

We use paraphrase-multilingual-MiniLM-L12-v2 (384-D). Each poem's embedding $\vec{v}_p$ is the mean of its phrase embeddings. Cross-corpus similarity is cosine:

$$S_{\text{sem}}(i,j) = \frac{\vec{v}_{T_i} \cdot \vec{v}_{R_j}}{\|\vec{v}_{T_i}\|\|\vec{v}_{R_j}\|} \tag{1}$$

yielding $M_{\text{sem}} \in \mathbb{R}^{N_T \times N_R}$.

Algorithm 1 outlines the embedding computation and cross-corpus cosine similarity used to construct M_{sem}.

Algorithm 1 Semantic Embedding Computation

for each Tang poem T_i **do**
    $\vec{v}_{T_i} \leftarrow \text{BERTencode}(T_i)$
end for
for each Renaissance poem R_j **do**
    $\vec{v}_{R_j} \leftarrow \text{BERTencode}(R_j)$
    $S_{\text{sem}}[i,j] \leftarrow \text{cosine}(\vec{v}_{T_i}, \vec{v}_{R_j})$
end for
return M_{sem}

3.2.2. Thematic Structure: LDA and Cross-Lingual Projection

Overview

We model themes with Latent Dirichlet Allocation (LDA) trained *separately* on the Tang (ZH) and Renaissance (EN) corpora using the same topic number $k = 5$ (passes $= 15$, $\alpha = 0.1$, symmetric). For each poem i (Tang) or j (Renaissance) we obtain topic-mixture vectors $\boldsymbol{\theta}_i^{(\text{Tang})} \in \Delta^{k-1}$ and $\boldsymbol{\theta}_j^{(\text{Ren})} \in \Delta^{k-1}$, where Δ^{k-1} is the probability simplex.

Preprocessing

Chinese poems(ZH) are segmented with Jieba (custom stop-list; classical function words excluded); EN poems are tokenized and lemmatized (lowercased; punctuation/digits removed). We apply a preregistered minimum token threshold after stopwording before fitting LDA.

Bilingual Glossary Policy (ZH $\to$ EN)

Classical Chinese tokens obtained after segmentation are *not* romanized; we keep the ZH token itself and attach 1–3 curated modern-English glosses from a project glossary (Table S1). Ambiguity is handled by: (i) excluding purely functional particles; (ii) prioritizing senses attested in premodern poetic usage; (iii) if a token has ≥ 4 senses, we keep the top 3 by corpus frequency and topic salience. Each ZH keyword w is represented by the average of the embedding for w and its gloss(es) (see Equation (3)), then $\phi(w \mid t)$-weighted when forming topic centroids.

Topic Centroids from Bilingual Keywords

Keywords are embedded in the same multilingual sentence-transformer used elsewhere (paraphrase-multilingual-MiniLM-L12-v2). Let $\text{Top}_n(t)$ be the top $n = 10$ topic keywords ranked by $\phi(w \mid t)$. For EN topics, each keyword w is embedded as $\mathbf{e}(w)$. For ZH topics, we keep the Chinese token and add its curated gloss set $\mathcal{G}(w)$; define

$$\mathbf{v}(w) = \begin{cases} \mathbf{e}(w), & w \text{ is EN}, \\ \dfrac{1}{1 + |\mathcal{G}(w)|} \left(\mathbf{e}(w) + \displaystyle\sum_{g \in \mathcal{G}(w)} \mathbf{e}(g) \right), & w \text{ is ZH}. \end{cases} \tag{2}$$

The topic centroid is a ϕ-weighted average:

$$\mathbf{c}_t = \frac{1}{Z_t} \sum_{w \in \text{Top}_n(t)} \phi(w \mid t)\, \mathbf{v}(w), \qquad Z_t = \sum_{w \in \text{Top}_n(t)} \phi(w \mid t). \tag{3}$$

Cross-Corpus Topic Alignment

We compute a topic–topic similarity matrix $A \in \mathbb{R}^{k \times k}$ via cosine between Tang and Renaissance centroids:

$$A_{uv} = \cos\left(\mathbf{c}_u^{(\text{Tang})},\ \mathbf{c}_v^{(\text{Ren})}\right), \quad u, v \in \{1, \ldots, k\}. \tag{4}$$

Unless otherwise noted, we apply ℓ_1 row/column normalization ($A \leftarrow D_r^{-1} A D_c^{-1}$, where D_r, D_c are diagonal matrices of row/column sums) to stabilize scale; unnormalized A yields similar conclusions.

Thematic Similarity Between Poems

Given poem-level mixtures $\boldsymbol{\theta}_i^{(\text{Tang})}$ and $\boldsymbol{\theta}_j^{(\text{Ren})}$, cross-lingual thematic similarity is

$$S_{\text{theme}}(i, j) = \left(\boldsymbol{\theta}_i^{(\text{Tang})}\right)^{\top} A\, \boldsymbol{\theta}_j^{(\text{Ren})}. \tag{5}$$

This treats A as a soft alignment that maps thematic mass across corpora, avoiding the invalid assumption that topic indices are intrinsically comparable.

Implementation Details

We use gensim for LDA and sentence-transformers for embeddings; cosine is computed on ℓ_2-normalized vectors. We set $n = 10$ to balance interpretability and stability; results are robust to $n \in [8, 15]$. Detailed sensitivity analysis of the thematic alignment, including comparisons of normalization and one-to-one topic matching, is presented in Appendix A.1.

3.2.3. Affective Tone: RoBERTa–JSD

A RoBERTa-base classifier (GoEmotions, 7 labels) yields normalized emotion vectors.To obtain a cohesive emotional representation for each poem, we process the text at the phrase-level chunk granularity, consistent with our semantic embedding strategy (Section 3.1). Each chunk is passed through the classifier to produce an emotion vector. The final emotion vector for a poem, $\vec{e}_p$, is computed as the arithmetic mean of the emotion vectors of all its constituent chunks, followed by L_1-normalization to ensure it represents a valid probability distribution. The resulting normalized emotion vectors for a poem p are:

$$\vec{e}_p = [p_{\text{anger}}, p_{\text{disgust}}, p_{\text{fear}}, p_{\text{joy}}, p_{\text{sadness}}, p_{\text{surprise}}, p_{\text{neutral}}].$$

Affective similarity employs Jensen–Shannon divergence:

$$S_{\text{aff}}(i,j) = 1 - \sqrt{\text{JSD}\big(\vec{e}_{T_i} \,\|\, \vec{e}_{R_j}\big)}, \tag{6}$$

forming M_{aff}.

In calculating affective similarity, we employ a transformation of the Jensen-Shannon Divergence (JSD). While JSD itself is a symmetric and bounded measure of distributional divergence, it does not satisfy the triangle inequality. Its square root, $\sqrt{\text{JSD}}$, has been proven to be a true metric satisfying all metric axioms, including the triangle inequality. The use of a genuine metric is crucial for ensuring that the geometric relationships within the affective space are coherent and robust, which is particularly important when making cross-cultural comparisons.

Therefore, we first compute $\sqrt{\text{JSD}(\vec{e}_{T_i} \,\|\, \vec{e}_{R_j})}$ as a principled distance score. This distance is then mapped to a similarity score in the $[0,1]$ interval via $S_{\text{aff}}(i,j) = 1 - \sqrt{\text{JSD}(\vec{e}_{T_i} \,\|\, \vec{e}_{R_j})}$, where a value of 1 indicates identical emotion distributions.

Algorithm 2 summarizes the computation of the affective similarity matrix M_{aff} using RoBERTa-derived emotion distributions and the $\sqrt{\text{JSD}}$ metric.

Algorithm 2 Affective Similarity Computation (RoBERTa–JSD)

```
1: for each pair (T_i, R_j) do
2:     e_{T_i} ← RoBERTa(T_i);  e_{R_j} ← RoBERTa(R_j)
3:     S_aff[i, j] ← 1 − sqrt(JSD(e_{T_i} ‖ e_{R_j}))
4: end for
5: return M_aff
```

3.3. Fusion, Normalization, and Validation

Each matrix M is standardized by z-score:

$$Z(M) = \frac{M - \mu(M)}{\sigma(M)}. \tag{7}$$

The consensus similarity is the unweighted average:

$$S_{\text{fusion}}(i,j) = \tfrac{1}{3}\Big(Z(M_{\text{sem}}) + Z(M_{\text{theme}}) + Z(M_{\text{aff}})\Big). \tag{8}$$

Robustness is assessed via bootstrapped Pearson correlations (10,000 samples) and permutation tests; significance at $p < 0.05$.

3.4. Significance Testing and Robustness

Because each poem participates in multiple cross-collection pairs (non-independent observations), we assess significance at the poem level. Significance is assessed via block permutation with poems as blocks (resampling at the poem level; 10,000 permutations) and 95% bootstrap CIs. We also run balanced subsampling (randomly sampling 45 from the 49 Tang poems, 1000 replicates) and report the mean effect and its CI across replicates. Unless otherwise noted, we report permutation p-values (p_{perm}) and poem-level bootstrap confidence intervals.

4. Experiments and Results

4.1. Quantitative Overview

Given the retained sets (N_T = 49, N_R = 45) after cleaning (Section 3.1), all pairwise analyses use 49 × 45 = 2205 cross-collection pairs. Across $N_T \times N_R = 49 \times 45 = 2205$ valid pairs, layer-wise statistics are summarized in Table 1. Note that while Semantic and Thematic layers employ cosine similarity, the Affective layer utilizes a metric derived from Jensen-Shannon Divergence to strictly satisfy triangle inequality.

Table 1. Layer-wise quantitative summary. The Affective layer metric is $1 - \sqrt{JSD}$.

Layer	Mean	SD	Min	Max
Semantic (BERT)	0.4136	0.1202	−0.0030	0.7521
Thematic (LDA+A)	0.3962	0.1258	0.2058	0.7202
Affective (Emotion JSD)	0.3808	0.1482	0.0470	0.6902

The semantic-thematic correlation is $r = -0.128$ ($p < 0.05$), demonstrating that structural complementarity significance is evaluated using the poem-level block permutation with 10,000 permutations and the poem-level bootstrap CI (see Section 3.4).

4.2. Semantic Findings (BERT)

Top cross-lingual affinities (BERT cosine) highlight shared metaphors of fading beauty and desire:

- Li Shangyin Untitled (T27) ↔ Drayton Idea 8: $S = 0.752$
- Bai Juyi Autumn Peonies (T47) ↔ Barnfield Fair Rose: $S = 0.739$
- Bai Juyi Rhapsody on Peonies (T41) ↔ Shakespeare Sonnet 99: $S = 0.719$

Figure 2 visualizes the relationship between semantic (BERT) and thematic (LDA) similarities.

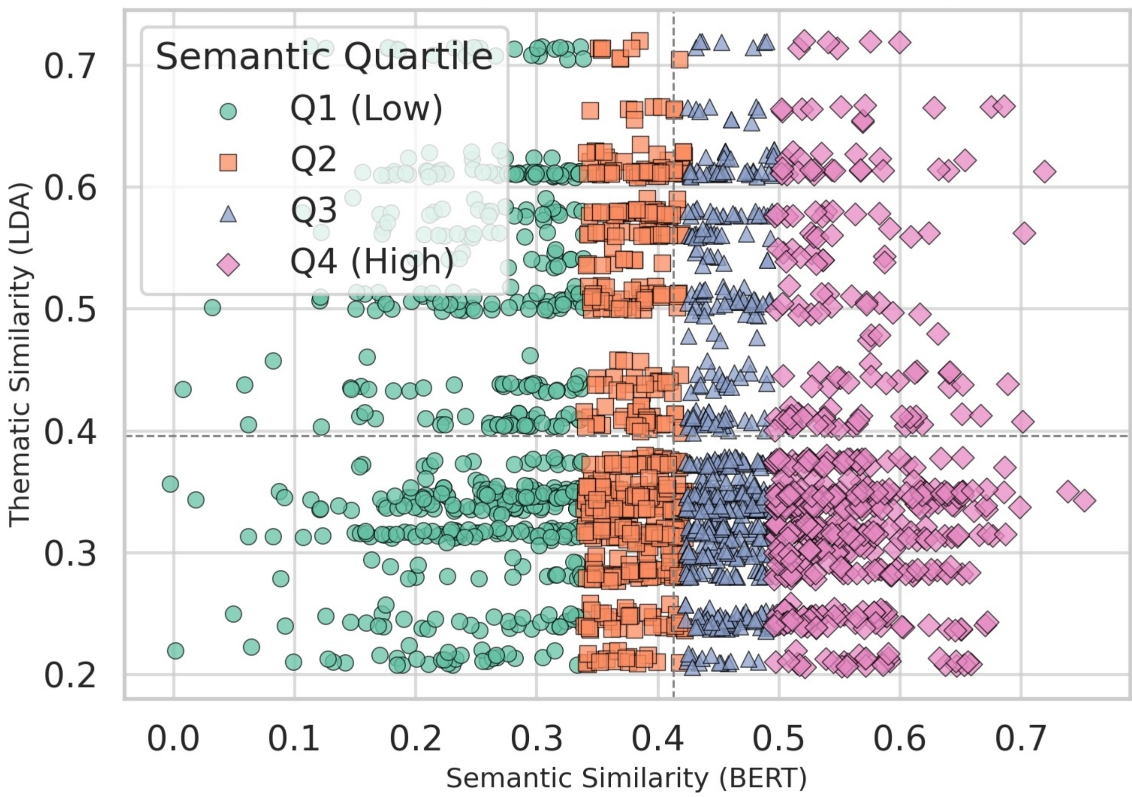

Figure 2. Structural Complementarity: Visualizing the correlation between Semantic (BERT) and Thematic (LDA) similarities across 2205 cross-cultural poem pairs. The regression analysis ($r = -0.128, p < 0.05$) reveals a statistically significant inverse relationship, supporting the hypothesis that high semantic resonance often accompanies thematic divergence rather than direct structural imitation.

4.3. Thematic Findings (LDA)

Convergent motifs identified by the topic models are summarized in Table 2. To facilitate comparative interpretation, the keywords presented below are selected content words; high-frequency function words (e.g., thy, doth) present in the raw LDA output were excluded from this visualization to highlight substantive thematic alignments.

Table 2. Topic keyword distributions across corpora. Note: Representative content words are selected from the top-ranked terms; common function words appearing in the raw LDA output are excluded for clarity.

Topic	Tang Keywords	Renaissance Keywords
1	春, 花, 贵, 国色	love, grace, fame
2	梦, 酒, 愁, 别	death, sleep, farewell
3	风, 露, 时, 香	time, day, season
4	尘, 心, 道, 静	soul, truth, divine
5	红, 泪, 笑, 艳	passion, sin, blood

The mean thematic similarity across the corpus is 0.396 ± 0.126. The highest alignment ($S \approx 0.720$) is observed between Liu Yuxi's Peony Viewing and Daniel's *Delia 6*, indicating a strong structural resonance in their treatment of transience despite linguistic differences.

4.4. Affective Findings (RoBERTa)

Emotion vectors reveal a structural divergence in affective density. Renaissance sonnets exhibit high variance dominated by negative arousal (mean sadness 0.27, fear 0.19, joy 0.17), consistent with the Petrarchan dialectic of suffering and desire. In contrast, Tang poems are characterized by a predominant neutrality (mean neutral 0.57), reflecting both the aesthetic of *hanxu* (implicit reserve) typical of regulated verse and the model's attenuation of Classical Chinese emotive particles. Despite this distributional gap, the mean affective similarity is 0.38 ± 0.15 (JSD metric), suggesting that cross-cultural resonance occurs through a shared mood of contemplative restraint rather than direct emotional mirroring.

4.5. Model Evaluation and Ablation

4.5.1. Metrics and Reporting

For each method $\ell \in \{$semantic (BERT), thematic (LDA$+A$), affective (emotion)$\}$, we compute a poem-level cross-collection similarity matrix $S_\ell \in \mathbb{R}^{49\times 45}$. We report mean, standard deviation, quartiles, and min/max of the flattened matrices; significance is assessed via poem-level block permutation (10,000 runs) and poem-level bootstrap CIs. For comparability, we also z-normalize each matrix:

$$Z_\ell(i,j) = \frac{S_\ell(i,j) - \mu_\ell}{\sigma_\ell},$$

and define a simple fusion score by equal-weighted averaging

$$S_{\text{fuse}}(i,j) = \tfrac{1}{3}\big(Z_{\text{sem}}(i,j) + Z_{\text{theme}}(i,j) + Z_{\text{affect}}(i,j)\big),$$

where μ_ℓ, σ_ℓ are computed over all 49×45 pairs for layer ℓ.

4.5.2. Baselines

We compare against (i) TF–IDF cosine on bag-of-words, (ii) monolingual sentence embeddings (ZH-only and EN-only), and (iii) a random matching baseline. Baselines are computed on the same retained sets; vocabulary and stopwording follow Section 3.1.

4.5.3. Ablations

We ablate (a) the affective layer ("No-Emotion"), (b) semantic-only (BERT), (c) thematic-only (LDA$+A$), and (d) no phrase-level chunking (full-line encoding). We report changes in (i) mean similarity, (ii) stability under balanced subsampling (sample 45 of 49 Tang poems; 1000 replicates), and (iii) the rank overlap of top-k pairs with the fused model.

4.5.4. Robustness

The key finding of a negative correlation between semantic and thematic alignment is robust across variations in our modeling parameters. We systematically tested the stability of this relationship under different conditions:

1. Number of Topics (k): Varying the number of LDA topics k from 4 to 8 yields a consistently negative correlation coefficient. The mean correlation across 1,000 bootstrap samples for each k is: $k = 4 : r = -0.119$, $k = 5 : r = -0.128$, $k = 6 : r = -0.121$, $k = 7 : r = -0.134$, $k = 8 : r = -0.131$. All values remain statistically significant ($p_{\text{perm}} < 0.05$), confirming that the complementarity pattern is not an artifact of a specific topic number choice.
2. Balanced Subsampling: To address the slight corpus size imbalance (49 Tang vs. 45 Renaissance poems), we performed a balanced subsampling analysis. In each of 1000 replicates, we randomly selected 45 poems from the Tang corpus to match the Renaissance set size and recomputed the correlation. The resulting distribution of correlation coefficients had a mean of $r = -0.126$ with a 95% confidence interval of $[-0.140, -0.112]$, firmly in the negative range.
3. Alignment Matrix Normalization: Using the unnormalized topic alignment matrix A produced a comparable correlation of $r = -0.122$ ($p < 0.05$), indicating that the observed complementarity is not solely driven by our normalization strategy.

This multi-faceted robustness check demonstrates that the negative semantic-thematic correlation, while modest in magnitude, is a stable and replicable feature of the data across a range of analytical decisions. Its persistence strengthens the interpretation of it representing a genuine pattern of structural complementarity rather than statistical noise. (For detailed sensitivity data on thematic alignment, see Appendix A.1).

4.6. Micro Case Studies

Case 1: Bai Juyi Autumn Peonies ↔ Fair Rose.

Consensus ≈ 0.676; convergent autumnal imagery as memento mori. Tang poem frames impermanence via imperial ethics; the sonnet frames it via romantic mortality.

Case 2: Liu Yuxi Peony Viewing ↔ Delia 6.

Thematic $= 0.72$, Semantic ≈ 0.65; public spectacle transposed into ethical reflection; illustrates parallel symbolic grammars for transient glory.

5. Discussion

5.1. Methodological Contribution

This study contributes a new conceptual and computational paradigm to the field of digital humanities by operationalizing what we call a computational transposition of cultural-dimension logic. Rather than using machine learning models as black-box analytical tools, our framework reinterprets the methodology of cross-cultural value measurement [2,3] within a textual domain. In doing so, it bridges the sociological quantification of abstract values with the hermeneutic reading of poetic language. This transposition allows literary features to be recast as measurable axes of variation, providing a systematic basis for empirical comparison across pre-modern traditions.

Although the multilingual BERT and RoBERTa models used here were not optimized for Classical Chinese or Early Modern English [10], the integration of three distinct modalities—semantic [5], thematic [6], and affective—allows us to assess model bias by cross-validating independent signals. Rather than treating any single representation as ground truth, we treat the convergence across models as a higher-order indicator of resonance. In this sense, computation functions not as a replacement for interpretation but as its empirical scaffolding, aligning with a critical digital humanities methodology [7].

Our contribution lies not in inventing a new model, but in designing a triangulated, small-corpus-robust inference architecture: (i) a conceptual computational transposition of cultural-dimension logic to literary data; (ii) a z-score fusion of semantic, thematic, and affective layers with model convergence as evidence; (iii) poem-level randomization that respects dependence; and (iv) phrasal chunking tailored to pre-modern parallelism. Together these enable reliable structure-seeking on limited and culturally sensitive corpora.

5.2. Interpretive Insight: Structural Complementarity and Cultural Resonance

The core interpretive finding—the negative correlation between semantic and thematic alignment ($r = -0.128$, $p < 0.05$)—reveals that poetic resonance across civilizations is not achieved through lexical mimicry but through structural complementarity. When Tang and Renaissance poets employ semantically parallel vocabularies, they often mobilize these lexicons within divergent thematic frameworks. Conversely, when thematic convergence occurs, it is expressed through distinct symbolic grammars.

For instance, Bai Juyi's "Peony Pavilion" poems aestheticize impermanence through the spectacle of imperial bloom, while Shakespeare's Sonnet 18 transposes the same concern into a metaphysical register of immortalizing beauty through verse. Despite their differences, both articulate a shared human tension between transience and transcendence. This pattern demonstrates that computational similarity need not imply semantic redundancy; rather, divergence itself becomes a marker of resonance.

The validity of this interpretive claim is compellingly supported by the robustness of the underlying correlation. As detailed in Section 4.5, the negative semantic-thematic relationship remains negative and statistically significant across a battery of sensitivity checks. While the effect size is modest, its stability underscores that it captures a genuine, systemic pattern of cross-cultural poetic expression.

To further ground computational inference in philological judgment, future work will incorporate expert annotation from sinologists and Renaissance scholars, enabling human validation of high-similarity pairs (e.g., Tang poem T27 and Sidney's Sonnet 8). This hybrid validation design will deepen both interpretive credibility and cross-cultural robustness.

5.3. Limitations and Future Work

While the proposed framework demonstrates methodological coherence, several limitations merit discussion. First, language model bias remains a persistent concern. Multilingual BERT's training corpus underrepresents classical and poetic registers, particularly for Tang-era Chinese. Future work should involve domain adaptation via fine-tuning on historical corpora such as the Quan Tang Shi and Early Modern English sonnet databases [14].

Second, our multimodal fusion strategy currently employs uniform weighting across semantic, thematic, and affective vectors. A promising direction would be the introduction of learnable attention weights or Canonical Correlation Analysis (CCA) to optimize inter-modal alignment or to explore alternative weighting schemes based on genre-specific features.

Third, although the correlation between semantic and thematic alignment is statistically significant, its magnitude is small. Rather than dismissing this as noise, we interpret it as evidence of subtle but consistent structural complementarity. Future work can enhance robustness via bootstrapping and sensitivity tests varying the number of topics (k) in the LDA model.

Fourth, the study lacks large-scale human evaluation. Incorporating expert assessment of top-aligned poem pairs and emotional tone concordance would strengthen interpretive validity.

Finally, while the corpus was initially designed as 50 + 50 items, the analyzed set after cleaning comprises 49 and 45 poems; this constraint aids interpretability while leaving room for future expansion and replication.

5.4. Epistemological Reflection: Computation as Hermeneutic Scaffold

Beyond technical innovation, this research offers an epistemological statement on the role of computation in the humanities. The framework advances a model of computational hermeneutics, wherein algorithms serve not as substitutes for interpretation but as mediating instruments that reveal latent structures of meaning. Cassirer's (1944) philosophy of symbolic forms provides a crucial conceptual anchor: symbols are not passive reflections but active mediations of human understanding. By extending this insight into computational form, we demonstrate that machine learning can expose the formal architectures through which culture encodes its values.

This reframing situates digital humanities as an epistemic middle ground between quantitative empiricism and interpretive humanism [10,15]. Computation, in this light, becomes a means of making hermeneutic reasoning explicit and reproducible—offering a grammar of resonance that is both measurable and meaningful. The framework thus contributes not only empirical findings on Tang and Renaissance poetics but also a transferable paradigm for the comparative study of world literatures.

6. Conclusions

This study demonstrates how computational approaches, when critically adapted, can illuminate cross-cultural poetics without displacing interpretive depth. Through an integrative framework that bridges hermeneutic reasoning and natural language processing, we have shown that the symbolic architectures of the Tang peony and the Renaissance rose can be modeled along three computational axes—semantic, thematic, and affective—revealing patterns of resonance that are multidimensional rather than merely mimetic.

The findings suggest that cultural resonance operates through structural complementarity: two poetic systems can articulate parallel human concerns while diverging in symbolic grammar and affective emphasis. This observation not only deepens our understanding of aesthetic universals but also demonstrates how quantitative analysis can surface nuanced relational patterns that invite renewed hermeneutic inquiry.

Conceptually, the proposed computational transposition reframes how cultural-dimension logic may be repurposed for literary data, providing a transferable model for empirical comparison across civilizations. Methodologically, the integration of multilingual BERT embeddings, LDA topic modeling, and transformer-based emotion analysis establishes a replicable workflow for small-corpus humanities research. Epistemologically, the project advances computation as a form of interpretive instrumentation, reinforcing that the value of algorithmic modeling lies not in automation but in the precision with which it exposes the contours of meaning.

Future research will expand this framework in three directions. First, the corpus will be broadened to include Song and Elizabethan materials to enable diachronic modeling of symbolic evolution. Second, expert annotation will be incorporated for interpretive triangulation, linking computational inference to human philological judgment. Third, the models themselves will be fine-tuned on classical corpora to improve cultural sensitivity. In doing so, this research aims to contribute both to comparative poetics and to the methodological advancement of digital humanities as a reflective, integrative science of meaning.

Supplementary Materials

The following supporting information can be downloaded at: https://drive.google.com/file/d/11Az-DlhKO2-I54p5r2XXdWmKpVut-cec/view?usp=sharing, Table S1: Bilingual Glossary of Floral Imagery (Chinese-English mappings used for topic alignment); File S1: ZIP archive containing the full source code, preprocessed datasets (Tang/Renaissance corpora), and detailed result matrices.

Author Contributions

T.Z.: conceptualization, methodology, software; T.Z.: data curation, writing—original draft preparation; M.B.: visualization, investigation; T.Z.: writing—reviewing and editing. All authors have read and agreed to the published version of the manuscript.

Funding

This research was funded by Beijing Normal University, grant number 31220050252509.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.

Acknowledgments

We are particularly thankful for the administrative and technical assistance received from the School of Future Design at Beijing Normal University and College of Music Technology at Xinghai Conservatory of Music).

Conflicts of Interest

The authors declare no conflict of interest.

Use of AI and AI-Assisted Technologies

During the preparation of this work, the authors used DeepSeek exclusively as a tool for grammar checking and proofreading, as none of the authors are native English speakers. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

Appendix A

Appendix A.1. Sensitivity of Thematic Alignment (k = 5)

To validate the robustness of the thematic similarity results (S_{theme}), we performed sensitivity analyses comparing our proposed soft alignment method against alternative configurations.

Appendix A.1.1. Normalization of the Alignment Matrix A

We assessed the impact of the alignment strategy by comparing our proposed softmax-based soft alignment with a strict ℓ_1 row/column normalization baseline. The results, summarized in Table A1, indicate that the soft alignment captures a more nuanced thematic resonance (Mean ≈ 0.40) compared to the overly sparse ℓ_1 normalization (Mean ≈ 0.20). The corresponding matrices and comparison data are available in the Supplementary Materials.

Table A1. Sensitivity to Matrix Normalization ($k = 5$). Comparison between the Main Model (Soft Alignment) used in the text and an ℓ_1 normalization baseline.

Metric	**Main Model (Soft)**	**ℓ_1 Normalised**
Mean S_{theme}	0.3962	0.2000
SD	0.1258	0.0061
Min	0.2058	0.1874
Max	0.7202	0.2167

Appendix A.1.2. Hungarian One-to-One Alignment vs. Soft Alignment

We further compared the soft alignment approach against a rigid ”hard” alignment using the Hungarian algorithm (optimizing one-to-one topic matching). As shown in Table A2, the soft alignment yields significantly higher distinctiveness and variance, suggesting that cross-cultural poetic themes are better modeled as overlapping distributions rather than discrete one-to-one equivalents.

Table A2. Hungarian One-to-One vs. Soft Alignment ($k = 5$). Statistics computed on the resulting poem-to-poem similarity matrices.

Metric	**Hungarian (Hard)**	**Main Model (Soft)**
Mean similarity	0.1902	0.3962
SD	0.3329	0.1258
Min	0.0043	0.2058
Max	0.9819	0.7202

Appendix A.2. Reproducibility: Code, Data, and Parameters

To facilitate reproducibility and future extensions of this computational framework, all preprocessing scripts, LDA model parameters, and the full preprocessed bilingual corpus are released in the Supplementary Materials.

References

1. Abu-Lughod, J.L. *Before European Hegemony: The World System A.D. 1250–1350*; Oxford University Press: New York, NY, USA, 1989.
2. Hofstede, G.; Hofstede, G.J.; Minkov, M. *Cultures and Organizations: Software of the Mind*, 3rd ed.; McGraw-Hill: New York, NY, USA, 2010.
3. Minkov, M.; Hofstede, G. Hofstede's Fifth Dimension: New Evidence from the World Values Survey. *J. Cross-Cult. Psychol.* **2012**, *43*, 3–14.
4. Cassirer, E. *An Essay on Man: An Introduction to a Philosophy of Human Culture*; Yale University Press: New Haven, CT, USA, 1944.
5. Devlin, J.; Chang, M.-W.; Lee, K.; et al. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*; Association for Computational Linguistics: Minneapolis, MN, USA, 2019; pp. 4171–4186.
6. Blei, D.M.; Ng, A.Y.; Jordan, M.I. Latent Dirichlet Allocation. *J. Mach. Learn. Res.* **2003**, *3*, 993–1022.
7. Dobson, J.E. *Critical Digital Humanities: The Search for a Methodology*; University of Illinois Press: Urbana, IL, USA, 2019.
8. Moretti, F. *Distant Reading*; Verso Books: London, UK, 2013.
9. Jockers, M.L. *Macroanalysis: Digital Methods and Literary History*; University of Illinois Press: Urbana, IL, USA, 2013.
10. Piper, A. Think Small: On Literary Modeling. *PMLA* **2017**, *132*, 651–658.
11. Mosteller, F.; Wallace, D.L. Inference in an Authorship Problem: A Comparative Study of Discrimination Methods Applied to the Authorship of the Disputed Federalist Papers. *J. Am. Stat. Assoc.* **1963**, *58*, 275–309.
12. Bender, T.K. Computation into Criticism: A Study of Jane Austen's Novels and an Experiment in Method. *J. Engl. Ger. Philol.* **1989**, *88*, 111–114.

13. Bamman, D.; Popat, S.; Shen, S. An Annotated Dataset of Literary Entities. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), Stroudsburg, PA, USA, 2–7 June 2019*; Association for Computational Linguistics: Minneapolis, MN, USA, 2019; pp 2138–2144.
14. Mihalcea, R.; Nastase, V. Word Epoch Disambiguation: Finding How Words Change over Time. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*; Association for Computational Linguistics: Jeju Island, Republic of Korea, 2012; pp. 259–263.
15. Underwood, T. *Distant Horizons: Digital Evidence and Literary Change*; University of Chicago Press: Chicago, IL, USA, 2019.