

Reconfiguring the Ideal Order: Ablation and Field Formation in the Twentieth-Century Nigerian Novel in English

Grant Hamilton <hamilton_at_cuhk_dot_edu_dot_hk>, Chinese University of Hong Kong  <https://orcid.org/0000-0002-5567-995X>

Eli Park Sorensen <eliparks_at_cuhk_dot_edu_dot_hk>, Chinese University of Hong Kong  <https://orcid.org/0000-0002-9912-0840>

DOI: pending

Abstract

Computational literary studies has developed sophisticated methods for detecting textual relationships – similarity networks, text reuse detection, semantic comparison – yet these approaches remain largely constrained to pairwise analysis. They reveal which texts resemble one another but cannot capture how individual works contribute to the collective semantic architecture of a corpus. This paper introduces a methodology that addresses this gap by shifting focus from text-to-text relationships to field formation. Treating a corpus as an integrated semantic space constructed through GloVe word embeddings, the method employs ablation testing – the systematic removal of individual texts followed by retraining – to measure each text’s structural contribution to that space. Procrustes analysis quantifies the geometric difference between full and ablated embeddings, and an iterative protocol with consistency-based exponential decay filtering distinguishes genuine semantic disruption from algorithmic noise. The resulting “disruption scores” identify texts that are structurally necessary to a corpus’s semantic field, capturing a dimension of textual significance irreducible to similarity, centrality, or explicit citation. Drawing on T.S. Eliot’s concept of the “ideal order” as a heuristic for the relational and achronological dynamics the model quantifies, the paper demonstrates the methodology on a corpus of 142 twentieth-century Nigerian novels in English. Results validate the approach against established literary-historical scholarship while also surfacing critically overlooked texts whose structural significance warrants further attention. The methodology is portable across corpora and offers scholars a data-driven means of engaging intertextuality’s foundational insight: that literary significance is relational, systemic, and collectively constituted.

I. Introduction

Since the late 1960s, literary criticism has increasingly abandoned the model of literary influence in favour of intertextuality as the dominant framework for understanding relationships between texts. This shift responded to growing unease with the way literary influence implied hierarchical, causal links — typically framed as a movement from source text A to derivative text B — and privileged authorial intention and canonized achievement [Juvan 2008, pp. 55–56]. Intertextuality emerged as a more distributed and dynamic model. Introduced to Western literary theory by Julia Kristeva and Roland Barthes, it reconceptualized textual production as a process of absorption and transformation. As Kristeva explains, “any text is constructed as a mosaic of quotations; any text is the absorption and transformation of another” [Kristeva 1986, p. 37]. Barthes similarly declared that “a text is made of multiple writings, drawn from many cultures and entering into mutual relations of dialogue, parody, contestation” [Barthes 1977, p. 148]. These formulations replaced the linear causality of literary influence with the notion of infinite textual networks. In doing so, they liberated literary analysis from biographical intentionality and unidirectional causation, foregrounding instead the idea that literary meaning emerges through a text’s embeddedness in broader discursive systems.

This reconceptualization of textual meaning has found new expression in computational literary studies (CLS), where large-scale methods now trace the discursive entanglements that Kristeva and Barthes theorized. Over the past decade or so, a substantial body of scholarship has demonstrated the increasing interpretive potential of quantitative approaches to literary corpora. Matthew Jockers's *Macroanalysis* (2013) introduced large-scale text mining to literary history, using topic modelling and stylometric analysis to identify patterns of influence and thematic clustering across thousands of nineteenth-century novels. Ted Underwood's *Distant Horizons* (2019) employed machine learning to trace the differentiation of literary genres and the transformation of gender roles across two centuries of English-language fiction. Andrew Piper's *Enumerations* (2018) brought distributional semantics and vector space models to bear on elementary literary categories including plot, character, and fictionality, while Katherine Bode's *A World of Fiction* (2018) foregrounded the critical importance of corpus construction and bibliographic rigour for data-rich literary history. Daniel Shore's *Cyberformalism* (2018) offered a complementary perspective, using regex-based pattern matching in digital archives to trace the circulation of linguistic forms — an approach that privileges syntactic structure over the lexical frequencies that dominate most quantitative methods. However, perhaps the most pertinent monograph to the present study is Michael Gavin's *Literary Mathematics* (2022), which deploys vector-space models on the Early English Books Online corpus to map semantic fields and conceptual structures, engaging directly with Kristevan intertextuality as a theoretical framework for distributional semantics. Gavin's insight that corpora reveal measurable relationships between language and the conditions of textual production provides a crucial precedent for the approach developed here, though Gavin's analysis of the topology of semantic fields gives way in our method to a measure of the structural contribution of individual texts to that topology through their removal.

Within this broader tradition, specific computational methods have addressed distinct dimensions of intertextual connection. Network analysis has offered critics new ways to visualize structural relationships across literary corpora — notably in Franco Moretti's work on plot and character networks [Moretti 2011] and in Frank Fischer and Daniil Skorinkin's application of social network analysis to Russian literature [Fischer and Skorinkin 2020]. Text reuse detection algorithms, such as those developed by the Tesserae project and Passim, have traced explicit quotations and borrowings across time and genre, from Classical allusion [Coffee et al. 2013] to nineteenth-century newspaper reprints [Smith, Cordell, and Mullen 2015]. More recently, word embedding models have expanded the scope of intertextual analysis by identifying semantic affinities beyond surface-level lexical overlap. Patrick Burns and his team [Burns et al. 2021] have demonstrated this approach in their study of Latin literature, while Charlotte Schubert (2020) has employed Word Mover's Distance to detect paraphrases in ancient Greek texts, successfully identifying relationships between Thucydides and later historians like Cassius Dio. Such approaches have operationalized key aspects of intertextuality, making visible the latent textual patterns that remain stubbornly inaccessible to traditional close reading.

However, as Paul Molitor and Jörg Ritter note, "Most aspects of intertextuality are difficult to grasp in mathematical terms" [Molitor and Ritter 2020, p. 49]. While current methods excel at detecting explicit textual connections — quotations, citations, lexical overlaps — they are largely constrained to pairwise comparisons. Such techniques illuminate which texts resemble one another or occupy central positions in similarity networks, but they fall short of revealing how individual works shape the broader semantic terrain of a corpus. To address this limitation, this paper introduces a computational methodology that measures a text's contribution to the semantic field in which it participates. The approach treats a corpus as an integrated semantic space constructed through word embeddings, then employs ablation testing to assess how that space changes when individual texts are removed. Using Procrustes analysis to quantify geometric differences between the full corpus embedding and each ablated version, we generate "disruption scores" that identify the texts that are the most consequential to the corpus's semantic architecture. Importantly, this is not a measure of similarity, centrality, or explicit citation. It is a measure of field formation that captures something closer to the concerns of intertextuality, which is to say participation in collective meaning-making systems.

We demonstrate this methodology using a corpus of 142 twentieth-century Nigerian novels in English. The corpus is well-suited to the task since it is large enough to yield meaningful semantic patterns yet compact enough to remain computationally tractable. Critical work on the Nigerian novel is also sufficiently robust to support validation of computational findings against established literary scholarship. Because disruption scores identify texts that shape the semantic field of a corpus irrespective of their canonical status, this approach offers scholars a data-driven entry point

for re-examining literary histories and the opportunity to surface works that are structurally significant to the discursive system in which they participate. While our case study focuses on the twentieth-century Nigerian novel in English, the methodology is portable across literary corpora. By shifting attention from resemblance to relational impact in this way, this approach expands the interpretive repertoire of computational literary studies and opens new pathways for engaging the distributed dynamics of textual significance.

II. Methodology

The reconceptualization of textual significance beyond resemblance or citation requires a conceptual framework for understanding what it means to measure a text's contribution to a semantic field. While the methodology itself is grounded in distributional semantics and linear algebra, we find a productive heuristic in T.S. Eliot's early essay "Tradition and the Individual Talent" (1919). Though Eliot could not have anticipated computational approaches to literary analysis, his notion of a literary field that reconfigures itself as new texts are introduced to it offers a remarkably apt way of thinking about the dynamics our method seeks to quantify. Writing decades before Kristeva and Barthes articulated intertextuality, Eliot proposed that works of literature exist within a continually reconfiguring system of relations, which he called the "ideal order." His vision sharply diverged from the romantic-humanist paradigm that privileged the solitary genius of the author and construed literary relations as linear inheritance [Arnold 1865] [Saintsbury 1900–1904] [Stephen 1874–1879]. Instead, Eliot articulated a dynamic understanding of literary tradition that anticipates the architecture of intertextuality theory. Commenting on how texts acquire literary meaning, he writes:

The existing monuments [literary works] form an ideal order among themselves, which is modified by the introduction of the new (the really new) work of art among them. The existing order is complete before the new work arrives; for order to persist after the supervention of novelty, the *whole* existing order must be, if ever so slightly, altered; and so the relations, proportions, values of each work of art toward the whole are readjusted. [Eliot 1919, p. 55]

] Two aspects of this formulation prove essential for the methodology developed in this paper. First, Eliot insists that literary significance is determined relationally rather than intrinsically. "No poet, no artist of any art, has his complete meaning alone," Eliot states in the same essay. He continues, "His significance, his appreciation is the appreciation of his relation to the dead poets and artists" [Eliot 1919, p. 55]. Understood like this, although a text may possess qualities specific to the literary work itself, the meaning of that text only becomes fully legible when it is put into conversation with other texts. It is this emphasis on relationality that anticipates perhaps the most radical claim of intertextuality — that the literary meaning of a text is something that is constituted through its dialogic interplay with other texts. Here, it is not necessarily what a text *is* but what it *does* within a field of relations that is important.

The second aspect of Eliot's model worth dwelling on is its insistence on the achronological nature of literary influence. As he explains, "the past should be altered by the present as much as the present is directed by the past" [Eliot 1919, p. 55]. That the order of composition or publication of literary works does not determine the magnitude of reorientation, that a later work can retroactively reshape our understanding of earlier texts just as earlier works inform our reading of the later, is extremely important. With this insight alone, one can decisively distinguish Eliot's model of literary inheritance from those dominant at the turn of the century which conceived influence as a strictly chronological process in which earlier texts exerted causal force upon later ones [Clayton and Rothstein 1991, p. 5]. Eliot's ideal order operates through mutual constitution where texts derive meaning from their simultaneous participation within the literary field rather than from the date of their introduction to the field.

A. Word Embeddings

To build a model that captures something like the systemic reconfigurations described by Eliot requires reconceiving the corpus not as a collection of discrete texts but as a unified semantic field. Word embeddings offer precisely this kind of representation. Rooted in distributional semantics, which posits that words occurring in similar contexts tend to share meanings [Harris 1954] [Firth 1957], embeddings map lexical items as vectors in a multidimensional space, where proximity reflects shared patterns of usage across a corpus [Pennington, Socher, and Manning 2014] [Mikolov et al.

2013]. Unlike network-based approaches that model texts as nodes linked by pairwise similarity, embeddings construct a single, integrated semantic field in which lexical meaning emerges relationally from the corpus's internal distributional structure [Turney and Pantel 2010] [Erk 2012]. This distinction is important because while networks aggregate bilateral relationships, embeddings generate a continuous semantic topology — one that enables more nuanced forms of interpretation, comparison, and conceptual mapping across the corpus.

The embeddings used in this study are generated *de novo*, trained exclusively on a corpus of twentieth-century Nigerian novels in English. We implement the GloVe (Global Vectors for Word Representation) algorithm [Pennington, Socher, and Manning 2014], ensuring that the resulting vector space reflects only the linguistic textures and conceptual rhythms of these texts. Because the embedding algorithm constructs its representation from co-occurrence patterns across the entire dataset without privileging texts based on their publication chronology, the resulting semantic field is intentionally ignorant of temporal sequence. Thus, the embedding represents an integrated semantic field where publication sequence does not determine structural importance. In this sense, the embedding space functions as a computational analogue to Eliot's "ideal order" — a mathematical structure that exhibits the key properties of mutual constitution, achronological operation, and the relational determination of significance.

9

B. Ablation Testing

Once the corpus is modelled as an integrated semantic field in this way, ablation testing offers a mechanism for assessing the contribution of individual texts to its construction. The procedure is straightforward. Each text is systematically removed from the corpus, the word embedding is retrained without it, and the resulting change in the semantic space to the full corpus embedding is measured.

10

Borrowed from machine learning practices that evaluate feature importance through removal [Meyes et al. 2019], this method captures a dynamic fundamentally distinct from those modelled by similarity metrics or network analysis. While similarity calculations assess how closely one text resembles another [Forstall and Scheirer 2019], network centrality measures the number and strength of a text's connections within a graph of texts [Moretti 2011] [Trilcke et al. 2016]. Both approaches quantify presence — what a text shares or how it links with others. Ablation, by contrast, measures absence or, rather, what is lost when a text is removed. It challenges the intuition that a text's value lies in its typicality or centrality. In a semantic field constructed from distributional patterns across all texts, a highly idiosyncratic work may introduce distinctive collocations, vocabulary, or syntactic structures that reshape the geometry of the space in ways that more conventional texts do not. Conversely, a text that closely resembles others may occupy a position easily approximated by its peers, making its removal less disruptive to the field's structure. For this reason, ablation testing cannot be regarded as a proxy for measures of similarity or connectivity. It measures the extent to which a text's unique linguistic patterns have shaped the architecture of the semantic field that the corpus collectively produces. This is precisely the phenomenon Eliot describes when he notes that the ideal order must be "altered" by the addition of literary works. Ablation captures the magnitude of that forced reorientation — not as a historical sequence, but as a structural shift within an integrated semantic space constructed from the corpus's distributional patterns.

11

C. Procrustes Analysis

Measuring the difference between two multidimensional semantic spaces presents a significant methodological challenge. Indeed, the choice of metric by which to understand the difference between original and recast embeddings is of utmost importance because it is a choice that fundamentally shapes what the comparison can reveal. Although one might begin by trying to compare word frequencies between the full and reduced corpora, such an approach would fail to account for the relational semantics that word embeddings model. A more sophisticated approach might look to calculate average cosine similarity between corresponding word vectors in the two embeddings. However, this encounters a deeper structural issue with embeddings that makes such direct vector-wise comparisons untenable. Word embeddings are not anchored to a universal coordinate system. The optimization process involves random initialization and converges on non-unique solutions. As a result, independently trained embeddings may differ by arbitrary linear transformations such as rotation, reflection, scaling, or translation, even when trained on nearly identical data [Antoniak and Mimno 2018] [Wendlandt, Kummerfeld, and Mihalcea 2018]. To be clear, these transformations are

12

not semantically meaningful — they are byproducts of the embedding algorithm’s stochastic nature and the geometry of high-dimensional space. As such, if we directly compare vector positions across two independently trained embeddings, we risk measuring these arbitrary differences rather than genuine semantic change.

A further complication arises from vocabulary differences between embeddings. The full corpus embedding contains words from the removed text that do not appear in the ablated version, making direct comparison impossible even after alignment. To address this, we restrict the Procrustes comparison to a common vocabulary of high-frequency words that appear in every text in the corpus. For our test corpus, 141 such words exist (such as “people,” “man,” “woman,” “house,” “day,” “time”). Such restriction serves both a technical and a conceptual purpose. Technically, it ensures that both embeddings contain exactly the same words, making geometric comparison possible. Conceptually, it refocuses the analysis on how texts reshape the relational architecture of shared linguistic building blocks rather than measuring the impact of unique or rare vocabulary. Thus, the insight to be drawn from this Procrustes analysis is that disruptive texts reshape the semantic field not through introducing new vocabulary but through realising distinctive patterns that reposition the relationships between such fundamental words.

13

Procrustes analysis works by aligning two matrices through optimal rotation, scaling, and translation before measuring their difference [Gower and Dijksterhuis 2004]. Named after the mythological figure who stretched or disfigured his guests to fit the dimensions of his bed, the method addresses a fundamental challenge in geometric comparison — given two matrices representing geometric configurations, how can we determine whether they differ meaningfully (at the level of structure) or merely arbitrarily (just in terms of their relative positions)? In answering this, Procrustes finds the transformation that minimizes the distance between two configurations, isolating genuine structural differences from superficial ones [Schönemann 1966]. The formulation has proven valuable across various disciplines. Biologists use it to compare organism shapes across specimens, archaeologists to match artifact measurements, and psychologists to align factor structures across populations [Rohlf and Slice 1990] [Kendall 1989]. Regardless of the demands of each discipline, each application of Procrustes shares the common need to compare representations of similar underlying structures that may appear different due to arbitrary coordinate system placement. We adapt it here to semantic space comparison. Given two 141×50 matrices representing the anchor words in the full and ablated embeddings, Procrustes removes arbitrary transformations that result from the stochastic optimization process, allowing the critic to measure genuine geometric incongruence. After optimal alignment, any remaining difference reflects genuine structural change in how such fundamental words relate to each other within the semantic field rather than meaningless repositioning caused by algorithmic randomness.

14

Of course, alternative methods for comparing embedding spaces exist but they measure different phenomena to the one described above. For example, canonical correlation analysis (CCA) identifies the maximal correlation between linear combinations of variables across two spaces [Hotelling 1936]. However, it assesses statistical association rather than geometric congruence. This means that two spaces may exhibit high canonical correlation, indicating that variables co-vary in similar ways, yet still diverge markedly in their spatial configuration. Put simply, the actual positions and distances among points may differ substantially even when canonical correlations are strong. For our purposes, it is the geometric divergence itself that must be measured. It is important to know whether removing a text has reshaped the semantic architecture of the field, not merely whether variables still correlate. Procrustes provides this by quantifying the irreducible structural difference between two embeddings after all arbitrary transformations have been optimally removed.

15

However, that is not to suggest that Procrustes alignment eliminates all sources of variability. Even after removing arbitrary transformations, some residual noise persists due to the stochastic nature of the optimization process. The embedding algorithm may converge to slightly different local optima across runs, producing embeddings that differ in ways unrelated to the semantic impact of text removal. Thus, to distinguish genuine semantic disruption from this algorithmic noise, we repeat the ablation testing five times per text using different random seeds. Each iteration produces an independent Procrustes distance. Texts that generate consistently high distances across all five iterations provide reliable evidence of field disruption while texts with highly variable distances across runs suggest that measurements are capturing stochastic noise rather than consistent semantic impact. For this reason, our iterative approach rewards consistency. Because simply averaging the five Procrustes distances would treat all measurements

16

as equally reliable, ignoring the signal that variability provides about measurement quality, we instead calculate both the mean distance and the range of distances across iterations. The range, expressed as a percentage of the mean, quantifies measurement consistency. An exponential decay function then adjusts each text's score based on this consistency metric. Texts with low variability retain their full disruption scores, while texts with high variability receive progressively stronger penalties. This approach formalizes the intuition that disruption scores should reflect stable, replicable geometric changes rather than volatile measurements that fluctuate unpredictably across algorithmic runs.

In sum, the architecture of the model thus addresses the comparison challenge through a two-stage process. Procrustes analysis removes the arbitrary transformations that make direct comparison impossible, isolating genuine structural differences from meaningless repositioning. Then multiple iterations with consistency-based adjustment separate reliable signals of semantic disruption from residual algorithmic noise. Together, these techniques produce disruption scores that quantify each text's contribution to the formation of the semantic field. Put simply, the scores measure how much the corpus's relational architecture depends on the distinctive linguistic patterns each text introduces, capturing a dimension of textual impact that neither similarity metrics nor network centrality can assess.

17

D. Disruption Scores and the Principle of Corpus Dependence

Given that comparative analysis yields a quantitative measure of the change in the geometric structure of the corpus's semantic field, the central methodological claim presented here is somewhat simple: The resulting score — which we term the “disruption score” — represents the structural impact exerted by the ablated text on the corpus as a whole. A high disruption score indicates that the removed text contributed substantially to the architecture of the semantic space, not by virtue of lexical similarity to other texts or by occupying a central position within a network, but through its capacity to shape the corpus's semantic configuration in ways that are not recoverable from the remaining texts.

18

Thus, the disruption score of a text is not a measure of its literary influence in the traditional sense. Neither is it a measure of literary quality, formal innovation, or canonical importance. Rather, it constitutes a computationally tractable measure of a text's structural necessity to the semantic field — that is to say, the degree to which the field's architecture depends on the distinctive linguistic patterns that a text contributes to it. A text can be highly disruptive because it uses fundamental vocabulary in distinctive ways. These usage patterns reshape the geometric relationships among the shared linguistic building blocks through which all texts in the corpus construct meaning. Removing such a text forces the semantic field to reconfigure around its absence, producing measurable geometric change.

19

The results of such comparative analysis are entirely corpus-dependent. Changing the dimensions of the corpus — by adding or removing texts — will inevitably affect the disruption scores assigned to individual works since these scores are contingent upon the specific condition of the (semantic) context in which a given text is situated. While one might claim that this constitutes a significant methodological limitation, in fact quite the reverse is true. Indeed, this constitutive feature of the analysis resonates with foundational insights of intertextuality. As scholars have argued, the significance of a literary work is not intrinsic or universal but relational and contextual [Allen 2000, p. 5]. A text's meaning emerges from its position within networks of other texts, not from qualities it possesses in isolation. Thus, the same literary work may register as highly disruptive in one corpus while found to be somewhat marginal in another. By making corpus boundaries an explicit parameter of analysis, our model foregrounds what traditional criticism often leaves implicit — that all claims about literary importance are necessarily bounded by the archives from which they are derived [Moretti 2000].

20

In light of this corpus-dependence, the caveat which necessarily precedes all results in this essay — “*Given this corpus, our results show...*” — is therefore not to be understood as a compromising disclaimer but as a recognition that all literary analysis operates within such bounds, whether acknowledged or not. Canonical judgments about which authors matter, which texts are central, which works are influential all depend on implicit decisions about what to include in the comparison set. Traditional literary criticism naturalizes these boundaries, treating them as given rather than constructed. But computational methods, by requiring explicit corpus definition, make these boundaries visible as methodological choices. The disruption model outlined here embraces such visibility in that it measures contribution to a specific semantic field, defined by specific texts, and produces results that change if those boundaries in turn change.

21

III. Implementation

A. Corpus Construction

Clearly, the way in which one builds a literary corpus involves choices that inevitably impact the shape of results [Bode 2020] [Algee-Hewitt et al. 2016] [Schöch et al. 2021]. The literary corpus analysed in this study concerns the twentieth-century Nigerian novel in English, drawn from a larger archive of twentieth-century African prose fiction in English (1900–1999) held at the Chinese University of Hong Kong. As a literary field born in the crucible of decolonisation, the Nigerian novel exhibits a dense, complex, and contested field of intertextuality, making it an ideal and challenging test case. Our working corpus, which stands at 142 novels (c.10 million tokens), is large enough for robust statistical analysis yet computationally manageable. Indeed, the seminal bibliographic work of Wendy Griswold and Misty Bastian [Griswold and Bastian 1990] [Griswold 2000] provides the crucial context for this scale, revealing that our corpus of 142 novels represents c.28% of all novels in English published in Nigeria during this period. This high percentage of coverage provides a dataset of rare density, lending our analysis a degree of empirical weight seldom seen in computational literary studies. Furthermore, the rich body of existing criticism, from scholars such as Abiola Irele, Biodun Jeyifo, and Simon Gikandi, provides a crucial set of established interpretive benchmarks against which our computational discoveries can be validated.

22

B. Preprocessing and Tokenization

Preprocessing decisions directly affect the semantic space that word embeddings construct. The pipeline employed in this study implements several operations designed to isolate meaningful linguistic patterns while reducing noise. Hyphens are replaced with spaces to separate compound words, allowing the embedding algorithm to capture relationships between component words. Punctuation, symbols, and numbers are removed since they contribute little to the distributional semantics that word embeddings model [Pennington, Socher, and Manning 2014]. All text is converted to lowercase to prevent spurious distinctions based on capitalization [Bird, Klein, and Loper 2009].

23

Stopword removal requires particular attention. Standard English stopwords — high-frequency function words such as articles, prepositions, and pronouns — are typically removed on the grounds that they carry minimal semantic content. However, the Procrustes alignment procedure requires a set of ultra-high-frequency words that appear consistently across all texts to serve as stable reference points for geometric comparison. To ensure these words survive preprocessing, twelve specific words are absented from the standard stopwords lists: “the,” “and,” “of,” “in,” “to,” “is,” “was,” “said,” “had,” “he,” “she,” and “they.” These words were selected because they represent the most frequent tokens across the corpus while also spanning different grammatical categories (articles, conjunctions, prepositions, verbs, pronouns). Beyond these preserved words, standard English stopwords are removed using the Snowball stopwords list via `quanteda’s stopwords()` function [Benoit et al. 2018], as are single-letter tokens that typically represent OCR artifacts or formatting remnants.

24

The resulting tokenized texts form the input by which a feature co-occurrence matrix (FCM) is constructed. The FCM records how frequently words co-occur within specified context windows, defining the distributional patterns from which semantic relationships are derived. Following standard practice in word embedding construction, we employ a symmetric context window of five words in each direction [Pennington, Socher, and Manning 2014]. Co-occurrences are weighted by distance. Thus, words appearing adjacent to the target word receive full weight (1.0), while words appearing at greater distances receive inversely proportional weights (1/2, 1/3, 1/4, 1/5 for distances of 2, 3, 4, and 5 words respectively). It is a weighting scheme that reflects the linguistic intuition that words appearing closer together are more strongly associated [Firth 1957].

25

C. GloVe Implementation

The Global Vectors for Word Representation (GloVe) algorithm constructs word embeddings by factorizing the word co-occurrence matrix to produce vectors that capture semantic relationships through geometric properties [Pennington, Socher, and Manning 2014]. Unlike word2vec [Mikolov et al. 2013], which uses neural network architectures, GloVe

26

directly models global corpus statistics, making it particularly suitable for capturing corpus-level semantic structures.

Several hyperparameters govern the GloVe training process. Embedding dimensionality is set to 50. This choice balances expressiveness and computational efficiency; dimensions that are too low cannot capture the complexity of semantic relationships in a literary corpus, while dimensions that are too high introduce computational costs and, more importantly, risk overfitting to noise. The `x_max` parameter, which caps the maximum value that co-occurrence counts can contribute to the loss function, is set to 10, preventing extremely frequent word pairs from dominating the embedding space [Pennington, Socher, and Manning 2014, p. 1535]. Learning rate and other optimization parameters follow the defaults established by Pennington and his colleagues, which have proven robust across diverse corpora [Pennington, Socher, and Manning 2014, p. 1538].

For the full corpus embedding, optimization runs 15 iterations. Inspection of the loss function during training reveals that convergence typically occurs within 10–15 iterations, with minimal improvement thereafter. For computational efficiency, ablated embeddings run only 3 iterations. Because each ablated corpus retains 141 of the 142 texts, the co-occurrence statistics driving the optimization remain largely stable, allowing the algorithm to reach convergence rapidly. Importantly, rapid convergence does not indicate minor geometric change between embeddings. The iteration count simply reflects how quickly the algorithm finds an optimal configuration for the reduced corpus, not how different that configuration is from the original. All embeddings use identical random seeds to ensure reproducibility and minimize stochastic variation between runs. This process produces a word embedding matrix of dimensions $V \times 50$, where V is the vocabulary size. For the full corpus embedding, $V = 63,701$ unique words after preprocessing, reflecting the collective vocabulary across all 142 texts. Individual text embeddings are considerably smaller, ranging from $\approx 1,000$ to $\approx 11,000$ words depending on each novel's length and lexical diversity. Each row in these matrices represents a word, while each column represents a dimension in the semantic space. Words that appear in similar contexts occupy nearby positions in this 50-dimensional space, and semantic relationships manifest as geometric relationships.

D. Ablation Testing Protocol

The ablation testing procedure requires constructing 142 modified embeddings, one for each text in the corpus. For each text, the following steps are executed:

1. **Text removal:** All tokens belonging to the target text are identified within the full corpus token sequence and removed, producing a reduced corpus containing all tokens from the remaining 141 texts.
2. **Embedding reconstruction:** The GloVe algorithm is then applied to the reduced corpus using identical parameters to those used for the full corpus embedding (`rank=50`, `x_max=10`, context window 1:5 with distance weighting) but, as discussed, optimization runs 3 rather than 15 iterations.
3. **Iterative repetition:** Steps 1-2 are repeated five times for each text, each time using a different random seed for the GloVe optimization. This addresses the stochastic nature of the embedding algorithm, which can produce slightly different results across runs due to random initialization [Antoniak and Mimno 2018] [Wendlandt, Kummerfeld, and Mihalcea 2018].
4. **Procrustes comparison:** Each ablated embedding is compared to the full corpus embedding using Procrustes analysis, yielding five distance measurements per text. The choice of five iterations represents a considered balance between computational cost and statistical robustness. With 142 texts in the corpus, each additional iteration multiplies the computational cost substantially ($142 \text{ texts} \times 5 \text{ iterations} = 710$ independent embedding training runs). However, the iteration count is made more palatable through its coupling to a consistency-based decay function that penalizes measurements exhibiting high stochastic variance (see below). Together, this iterative measurement and variance-weighted aggregation distinguishes reliable signals of semantic disruption from algorithmic noise without requiring computationally prohibitive iteration counts.

This iterative approach distinguishes signal (the consistent change caused by text removal) from noise (stochastic variation across runs) given that a single comparison between the full embedding and the recast ablated embedding might reflect algorithmic variability rather than the true semantic impact of the removed text. By repeating the process

five times and aggregating results, we isolate the consistent geometric change attributable to the ablated text.

E. Procrustes Implementation

The Procrustes analysis follows the standard orthogonal Procrustes procedure implemented through the `vegan` package in R [Oksanen et al. 2024] [Schönemann 1966] [Gower and Dijksterhuis 2004]. For each ablated embedding, the procedure extracts the 141 anchor words from both the full corpus embedding (X) and the ablated embedding (Y), producing matrices `X_anchor` and `Y_anchor`, each with dimensions 141×50 . The `procrustes()` function then computes the optimal rotation, scaling, and translation that minimizes the sum of squared differences between the two configurations. The function returns a “sum of squares” statistic quantifying the residual difference after optimal alignment. This statistic, averaged across five iterations, constitutes the raw Procrustes distance for each text. As mentioned earlier, an exponential decay function based on measurement consistency then produces the final disruption score.

31

F. The Decay Function

Because ablation testing is repeated five times per text, each text receives five raw Procrustes distances. The variability across these five measurements reflects the stochastic nature of the embedding algorithm. Texts that produce highly variable Procrustes distances across iterations suggest that the measurements are capturing algorithmic noise rather than consistent semantic disruption. Conversely, texts that produce consistent Procrustes distances across iterations provide more reliable evidence of genuine field disruption. To account for this variability, we apply a decay function that adjusts each text’s disruption score based on the consistency of its measurements:

32

1. For each text, calculate the mean Procrustes distance across the five iterations: $\mu = \text{mean}(P_1, P_2, P_3, P_4, P_5)$.
2. Calculate the range of Procrustes distances: $\text{range} = \max(P_1, P_2, P_3, P_4, P_5) - \min(P_1, P_2, P_3, P_4, P_5)$.
3. Express this range as a percentage of the mean: $\text{variability} = (\text{range} / \mu) \times 100$.
4. Apply an exponential decay function: $\text{adjusted_score} = \mu \times \exp(-0.01 \times \text{variability})$.

A decay factor of 0.01 was selected to provide moderate penalization of high variability without eliminating the contribution of genuinely disruptive texts that happen to produce some measurement variance. Thus, a text with 20% variability (range equal to 20% of its mean) would have its score multiplied by $\exp(-0.2) \approx 0.82$, retaining 82% of its original value, while a text with 50% variability would retain $\exp(-0.5) \approx 0.61$, or 61% of its value. This exponential relationship ensures that increasing variability produces progressively stronger penalties but does not reduce any text’s score to zero. The resulting adjusted Procrustes scores — renamed *disruption scores* — provide the primary measure of each text’s contribution to corpus-level semantic field formation. Higher scores indicate texts whose removal produces large and consistent changes to the semantic space, even after accounting for algorithmic variability.

33

G. Validation and Interpretability

The validity of disruption scores as measures of semantic impact rests on demonstrating that they capture field-level contribution rather than superficial document properties. The most obvious potential confound is text length: Longer texts contribute more tokens to the corpus, which might mechanically inflate their disruption scores regardless of semantic impact. To test this, we calculated the correlation between text length (measured in token count) and disruption scores across all 142 texts in the corpus. The analysis reveals no meaningful relationship between these variables. The Pearson correlation coefficient is -0.01 ($p = 0.90$), indicating that text length explains virtually none of the variance in disruption scores. Figure 1 visualizes this independence. Here, the regression line is essentially flat, and texts ranging from approximately 30,000 to over 200,000 tokens show comparable disruption potential.

34

Text Length vs. Disruption Score

Pearson $r = -0.01$, $p = 0.9031$




Figure 1. Independence of Disruption Scores from Text Length

Figure 2 reinforces this finding by comparing disruption scores across length quartiles. The median disruption scores remain remarkably stable across all four quartiles, with overlapping interquartile ranges demonstrating that shorter texts (Q1) disrupt the semantic field as readily as longer texts (Q4).

Disruption Scores by Text Length Quartile




Figure 2. Disruption Score by Text Length Quartile

This result validates the core theoretical claim that ablation testing measures structural necessity rather than lexical volume or typicality. The near-zero correlation demonstrates empirically what the methodology predicts conceptually — that disruption scores capture how a text's distinctive usage patterns reposition fundamental vocabulary within the semantic space, not simply how many tokens a text contributes or how many unique words it introduces to the corpus.

36

External validation comes from comparison with established literary-historical knowledge. Texts identified as highly disruptive by this model should align, at least partially, with texts that scholars have recognized as significant through traditional methods. Obviously, perfect alignment would be neither expected nor desirable, as the computational approach captures different dimensions of significance to canonical literary criticism. However, that said, complete divergence from the literary-historical record would cast doubt on the interpretability of disruption scores. That the results identify some expected figures like Chinua Achebe and Buchi Emecheta while also surfacing less-discussed writers like Eddie Iroh and Obi Egbuna suggests that the methodology captures meaningful patterns while also revealing new critical insights.

37

Here, a note on replicability is warranted. The full ablation pipeline — training 710 independent GloVe embeddings (142 texts × 5 iterations), performing Procrustes alignment for each, and applying the decay function — is computationally intensive, and we wish to be candid about the demands that this places on replication. For the corpus analysed here (142 texts of approximately 10 million tokens), the procedure is tractable on standard computing infrastructure, but scaling to substantially larger corpora (thousands of texts or hundreds of millions of tokens) would require significant computational resources and, potentially, methodological adaptation. However, several components of the framework are independently portable. The Procrustes comparison of full and single-text-ablated embeddings can be performed

38

with any embedding algorithm. The conceptual architecture — measuring structural contribution through absence rather than presence, and quantifying the geometric change that removal produces — is adaptable to a range of implementations and corpus sizes. Equally important, the disruption scores themselves are best understood as entry points for further enquiry rather than as definitive measures of literary significance. What the method offers is identification: a principled, data-driven means of flagging texts whose structural role in a corpus warrants critical attention. The close reading and interpretive work that follows — the kind of analysis that ultimately determines what a high or low disruption score means for literary history — remains the domain of the traditionally trained literary scholar, and requires no computational apparatus to undertake. In this sense, the methodology's most replicable output is not the score itself but the question it generates: How does this text reshape the literary field in which it sits?

IV. The Twentieth-Century Nigerian Novel in English

The disruption scores generated by the ablation methodology range from 16.21 to 18.33 across the 142 texts in the corpus, with a mean of 17.08 and standard deviation of 0.35. This relatively narrow distribution is a direct consequence of the methodology's design for statistical robustness. Because each text undergoes five independent ablation tests, with disruption scores calculated as the mean of these iterations and then adjusted by an exponential decay function that penalizes high stochastic variance, the methodology explicitly filters out noise while retaining only stable, replicable signals. Texts exhibiting inconsistent geometric changes across iterations — those likely reflecting algorithmic noise rather than genuine semantic impact — receive progressively stronger penalties, compressing their scores toward the mean. This conservative approach means that texts achieving high disruption scores have demonstrated consistent, statistically reliable geometric impact across multiple independent measurements. Despite this deliberate constraint on variance, statistically significant differentiation emerges. The highest-scoring text exceeds the mean by 3.6 standard deviations, placing it in the 99.98th percentile ($p < 0.0001$). While the absolute numerical range appears narrow, these Procrustes distances represent geometric reconfigurations of 50-dimensional semantic space, where small numerical differences correspond to substantial structural changes in how fundamental vocabulary relates across the corpus. As such, the distribution captures the important insight that field formation is largely collective, with most texts contributing comparably to the semantic architecture, while a minority introduce distinctive and statistically reliable linguistic patterns that measurably reshape the field's geometry. Table 1 gives the twenty-five texts with the highest disruption scores, representing the works most consequential to the formation of the corpus's integrated semantic field.

Rank	Author	Title	Year	Disruption Score
1	Amos Tutuola	Ajaiyi and His Inherited Poverty	1967	18.33
2	Buchi Emecheta	The Moonlight Bride	1980	17.89
3	Amos Tutuola	Feather Woman of the Jungle	1962	17.86
4	Chinua Achebe	No Longer at Ease	1960	17.84
5	Theophilus Chijioke Ugwu	A Changed Man	1985	17.69
6	Flora Nwapa	Efuru	1966	17.68
7	Obi Egbuna	The Rape of Lysistrata	1980	17.68
8	Buchi Emecheta	The Rape of Shavi	1983	17.63
9	Augustus Adebayo	I Am Directed	1991	17.62
10	Ben Okri	Astonishing the Gods	1995	17.62
11	Elechi Amadi	Estrangement	1986	17.61
12	Flora Nwapa	Idu	1970	17.59
13	Ken Saro-Wiwa	Segi Finds the Radio	1991	17.54
14	Amos Tutuola	The Palm Wine Drinkard	1952	17.52
15	Chukwuemeka Ike	Toads for Supper	1965	17.50
16	Ogali A. Ogali	The Juju Priest	1977	17.49
17	Buchi Emecheta	Kehinde	1994	17.46
18	Ken Saro-Wiwa	Mr B Again	1989	17.45
19	Festus Iyayi	The Contract	1982	17.45
20	Cyprian Ekwensi	An African Nights Entertainment	1962	17.44
21	John Chukwu	Destined to Live	1992	17.43
22	Eddie Iroh	The Siren in the Night	1982	17.43
23	Dan Fulani	Flight 800	1983	17.40
24	T.M. Aluko	Kinsman and Foreman	1966	17.39
25	Ken Saro-Wiwa	Mr B	1987	17.39

Table 1. Top 25 Texts by Disruption Score

Several patterns emerge immediately from these results. First, literary works by certain authors appear multiple times: Tutuola (three texts), Emecheta (three texts), Nwapa (two texts), and Saro-Wiwa (three texts). This repeated presence suggests that these writers consistently produce work that shapes the semantic field in distinctive ways, contributing unique linguistic patterns that the corpus integrates. Second, the temporal distribution spans the entire century — from Tutuola’s *The Palm Wine Drinkard* (1952) to Okri’s *Astonishing the Gods* (1995). This demonstrates well the achronological nature of the methodology. Third, the top tier includes both canonically central figures (Tutuola,

Emecheta, Achebe, etc.) and less-discussed writers (Ugwu, Egbuna, Adebayo, etc.). This suggests that although the methodology gives results that align at times with the established literary history of the Nigerian novel, it also captures significant patterns that diverge from such narratives.

The dominance of Amos Tutuola in the disruption hierarchy provides strong validation that the methodology captures the way in which linguistically idiosyncratic texts reposition common vocabulary within the semantic field. Tutuola occupies the first, third, and fourteenth positions, with *Ajaiyi and His Inherited Poverty* (1967) achieving the highest disruption score in the corpus. This result aligns precisely with what the methodology should identify: linguistic innovation so distinctive that its removal fundamentally alters the semantic space. As Simon Gikandi observes, “Tutuola’s language was his own creation. It was made up of broken-up rules of grammar, a limited vocabulary, unusual sentence structures, and neologisms” [Gikandi 2003, p. 753]. But what the analysis shows is that these stylistic features manifest as distinctive distributional patterns. Tutuola’s “broken” grammar and unusual structures position even fundamental vocabulary — that is, words shared across all the novels of the corpus — in ways that contribute unique dimensions to the embedding space. When Tutuola’s texts are removed from the corpus, the ablated embeddings must reconstruct the geometric relationships among fundamental vocabulary without the distinctive usage patterns through which Tutuola positioned common words, forcing ordinary vocabulary to occupy different semantic positions in his absence. In short, the high disruption scores quantify what critics have long recognized about Tutuola’s writing — that his linguistic innovation is not merely stylistic flourish [Ogundipe 1969] [Hogan 2000] [Zabus 2007]. Rather, it is fundamental to the character of the twentieth-century Nigerian novel in English.

41

The prominence of Buchi Emecheta and Flora Nwapa similarly validates the methodology. Three of Emecheta’s novels rank in the top twenty-five, with *The Moonlight Bride* (1980) achieving the second-highest disruption score, while Nwapa’s *Efuru* (1966) and *Idu* (1970) both rank as high disruptors. Both writers are widely recognized as having fundamentally reshaped the representation of women in Nigerian literature. As Maureen Eke notes:

42

Among the writers whose works helped to transform the presence of African women writers in literature is Buchi Emecheta. Her depiction of the experiences of African, specifically Nigerian, women has challenged the stereotyped and idealized images of African women found in male texts. No discussion of African or black women’s writing can be complete without her. [Eke 2003, p. 240]

Similarly, Marie Umeh writes that Nwapa was the first woman writer in the male-dominated space of Nigerian literature to dismantle representations of African women as “largely marginal figures playing stereotypical roles with little or no authenticity,” and thereby replace the “Law of the Father” with “the Law of the Mother” [Umeh 2003, p. 554]. The disruption scores for both writers suggest that this reshaping operated not only at the level of theme and representation but at the level of language itself. Emecheta’s distinctive patterns for describing women’s experiences, domestic spaces, familial relationships, and social constraints, and the way in which Nwapa wrote about female agency and maternal authority contribute unique geometric configurations to the embedding space of the twentieth-century Nigerian novel in English. Removing these texts eliminates not just “women’s perspectives” but the specific linguistic patterns through which those perspectives were articulated, forcing the semantic field to reconfigure in non-trivial ways around their absence.

Chinua Achebe’s *No Longer at Ease* (1960) ranks fourth, a result that warrants careful interpretation. One might initially expect Achebe to dominate the disruption hierarchy given his towering canonical status. As Ngugi wa Thiong’o wrote in tribute to Achebe after his passing, “Every African novel became *Things Fall Apart*, and every [African] writer some sort of Chinua Achebe” (thiongo 2013, p. 52). A comment on the significance of Achebe to the evolution of the Nigerian novel, Ngugi’s observation also helps to explain why *No Longer at Ease* achieves a high but not dominant disruption score. If Achebe’s linguistic patterns — his careful balance of English and Igbo linguistic structures, his deployment of proverbs, his particular cadences — have been so widely absorbed by subsequent Nigerian writers, then it is possible that the semantic field of the Nigerian novel has become saturated with Achebe-like patterns. Thus, although removing *No Longer at Ease* would still produce a measurable disruption to the semantic field of the corpus, it would perhaps do so not to the degree one might expect given Achebe’s critical reputation.

43

Indeed, the analysis shows that other texts in the corpus broadly compensate for the absence of some of the most notable literary works from the embedding space. Table 2 shows that Wole Soyinka's *Season of Anomy* (1973), Cyprian Ekwensi's *Beautiful Feathers* (1963), and Onuora Nzekwu's *Blade Among the Boys* (1962) place amongst the least disruptive novels in the corpus. At first puzzling, such divergence between canonical standing and disruption scores in fact confirms that the methodology measures field-constitutive integration rather than literary quality or historical importance. Such a result suggests that widespread early influence — specifically *early* influence — over a literary corpus can, paradoxically, reduce the disruption score of a text. That is to say, when the innovations of a novel become so integrated that they appear in the work of many other writers, then those innovations of *the individual novel* become less structurally significant to the embedding space. If the importance of the publication date of a literary work makes itself felt in this model, then it is in this aspect of the analysis.

Rank	Author	Title	Year	Disruption Score
133	Jerry N. Njoku	Vengeance is Sweet	1985	16.57
134	Richard Anieke	Shameful Sacrifice	1985	16.56
135	Wole Soyinka	Season of Anomy	1973	16.54
136	Cyprian Ekwensi	Beautiful Feathers	1963	16.49
137	Martina Nwakoby	A House Divided	1985	16.46
138	Buchi Emecheta	Gwendolen	1989	16.45
139	Amos Tutuola	The Witch Herbalist of the Remote Town	1981	16.45
140	Adaora Lily Ulasi	Many Thing Begin for Change	1971	16.42
141	Eddie Iroh	Forty Eight Guns for the General	1976	16.24
142	Onuora Nzekwu	Blade Among the Boys	1962	16.21

Table 2. Bottom 10 Texts by Disruption Score

The appearance of Theophilus Chijioke Ugwu's *A Changed Man* (1985) at rank five in table 1 exemplifies the methodology's capacity to surface overlooked texts. Ugwu is not a widely-discussed figure in Nigerian literary criticism, yet the disruption score indicates the novel's substantial contribution to the formation of the semantic field of the corpus. This result does not claim that *A Changed Man* is "better" or "more important" than canonically recognized works. Rather, it suggests that the novel introduces linguistic patterns consequential to the corpus's overall architecture. However, what those patterns are, or how they manifest in this novel, is only something that can be determined through close reading. What the disruption score accomplishes here, then, is identification: It flags Ugwu's *A Changed Man* as a text that warrants further critical attention based on the structural role it plays in the semantic field of the twentieth-century Nigerian novel in English. Indeed, this capacity for identification may prove to be the methodology's most portable and productive contribution to literary studies more broadly. For scholars working with large or under-explored corpora — where the sheer volume of texts makes comprehensive close reading impractical — disruption analysis offers a principled means of directing critical attention toward works whose structural significance might otherwise

remain invisible, opening pathways for the kind of sustained interpretive engagement that computational methods complement but cannot replace.

V. Conclusion

This paper began by observing that computational literary studies has developed sophisticated methods for detecting textual relationships — mapping similarity networks, tracing explicit quotations, identifying semantic resemblances — yet these methods operate fundamentally through pairwise comparison. Such methods reveal which texts connect to other texts, but they cannot capture how individual texts contribute to the collective semantic architecture within which all texts in a corpus operate. While intertextuality has long insisted that texts gain significance through participation in larger discursive systems rather than through bilateral influence, computational methods have not produced tools for operationalizing this insight at the level of the corpus itself. The disruption methodology introduced here addresses this gap by shifting focus from text-to-text relationships to field formation. By treating a corpus as an integrated semantic space constructed through word embeddings, then employing ablation testing to measure how that space changes when individual texts are removed, the approach outlined in this essay is able to quantify each text's contribution to the collective linguistic architecture of a corpus. The resulting disruption scores identify texts that are structurally necessary to the field's formation not because those literary works resemble other literary works or even because they occupy central network positions, but because their distinctive linguistic patterns have been consequential in shaping the semantic space that makes meaning possible within the corpus.

46

For this reason, the disruption methodology occupies a specific niche within computational literary studies. Where network analysis reveals clusters, text reuse detection traces explicit borrowing, and diachronic embeddings track semantic change, this approach addresses a different question: *Which texts have been most consequential to the formation of this corpus's semantic architecture?* This question matters when the research goal is understanding field construction rather than bilateral relationships, when critics seek to identify structurally necessary texts rather than merely central or similar ones, and when the theoretical framework conceptualizes literary significance as collective participation rather than sequential influence. While the method does not identify biographical influence, trace allusions, or reveal thematic content — since to do so requires very different approaches — it does measure how distinctive linguistic patterns contribute to the semantic space that makes meaning possible within a corpus. In so doing, it offers scholars working on corpora a data-driven entry point for identifying overlooked texts of critical significance, and for those interrogating established canonical works a means of challenging received narratives about which texts shaped linguistic possibility spaces.

47

Indeed, the methodology opens several productive research directions. Comparative analysis across national literatures, genres, or historical periods could reveal whether linguistic field formation follows universal patterns or culturally specific dynamics — comparing twentieth-century novels from Nigeria, Kenya, and South Africa, for instance, would illuminate how postcolonial African literatures constitute themselves relationally. Questions of canonicity become empirically tractable since systematic comparison could determine whether celebrated texts tend to be field-constitutive or whether canonical status operates through mechanisms orthogonal to linguistic architecture. Translation studies could also benefit. Asking whether translated texts disrupt semantic fields in distinctive ways or whether they assimilate to receiving-language norms seems a significant question to address. Similarly, one might use this methodology to build a response to the question of how French Symbolism reshaped Anglophone verse, or Latin American magical realism made its presence felt in English-language fiction. Periodization, too, becomes measurable: Clusters of high disruption scores might indicate linguistic watershed moments, while low mean scores signal consolidation, offering a degree of empirical grounding for claims about movements like Modernism or the Harlem Renaissance. Indeed, a temporally iterative variant of the method that performs the ablation procedure on successive chronological subsets of a corpus rather than on the corpus as a synchronic whole could capture the recursive dimension of literary field formation, revealing how disruption scores shift as the field accumulates new texts over time, thereby modelling the historically layered process through which tradition constitutes itself. The methodology might also be able to distinguish individual from collective innovation in a corpus — was British Romanticism's semantic field constructed by canonical figures alone, or did minor writers contribute beyond their critical status? Such questions, intractable through existing computational approaches, are opened up through disruption analysis.

48

Thus, by operationalizing intertextuality's core insight that significance is relational and systemic rather than intrinsic and individual, the methodology presented in this essay demonstrates how computational approaches can engage productively with literary theory. It makes visible the collective dimension of literary production. Eliot concluded "Tradition and the Individual Talent" by noting that the poet's significance lies not in isolated achievement but in participation within the ideal order — that "very few know when there is expression of *significant* emotion, emotion which has its life in the poem and not in the history of the poet" [Eliot 1919, p. 73] and that criticism should therefore attend to the literary work rather than the biography of the writer. A century later, computational methods offer new ways to realize this critical vision, to study texts as participants in collective systems rather than as expressions of individual genius.

49

Data Availability

The corpus of twentieth-century Nigerian novels in English addressed in this paper is drawn from a larger archive of twentieth-century African prose fiction in English held at the Chinese University of Hong Kong. At this time, the archive's holdings are not publicly accessible due to copyright law. However, the code for this paper is available for inspection at <https://github.com/Prof-Hamilton/Code-for-Essays>, a public GitHub repository.

50

Funding

This research, part of "The African Literary Archive – A Computational Analysis of Literary Influence" (2023–25), was funded by Hong Kong's University Grants Committee (UGC) through the General Research Fund (project code: 14614522).

51

Works Cited

- Algee-Hewitt et al. 2016** Algee-Hewitt, M. et al. (2016) "Canon/archive: Large-scale dynamics in the literary field", *Stanford Literary Lab Pamphlets*, 11. Available at: <https://litlab.stanford.edu/LiteraryLabPamphlet11.pdf> (Accessed: 20 January 2026).
- Allen 2000** Allen, G. (2000) *Intertextuality*. London: Routledge.
- Antoniak and Mimno 2018** Antoniak, M. and Mimno, D. (2018) "Evaluating the stability of embedding-based word similarities", *Transactions of the Association for Computational Linguistics*, 6, pp. 107–119.
- Arnold 1865** Arnold, M. (1865) *The function of criticism at the present time*. London: Smith, Elder & Co.
- Barthes 1977** Barthes, R. (1977) *Image-music-text*, trans. by S. Heath. New York: Hill and Wang.
- Benoit et al. 2018** Benoit, K. et al. (2018) "quanteda: An R package for the quantitative analysis of textual data", *Journal of Open Source Software*, 3(30), p. 774. Available at: <https://doi.org/10.21105/joss.00774> (Accessed: 20 January 2026).
- Bird, Klein, and Loper 2009** Bird, S., Klein, E. and Loper, E. (2009) *Natural language processing with Python: Analyzing text with the Natural Language Toolkit*. Sebastopol: O'Reilly Media. Available at: <https://www.nltk.org/book/> (Accessed: 20 January 2026).
- Bode 2020** Bode, K. (2020) "Why you can't model away bias", *Modern Language Quarterly*, 81(1), pp. 95–124.
- Bode 2018** Bode, K. (2018) *A world of fiction: Digital collections and the future of literary history*. Ann Arbor: University of Michigan Press.
- Burns et al. 2021** Burns, P.J. et al. (2021) "Profiling of intertextuality in Latin literature using word embeddings", in *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, pp. 4900–4907. Available at: <https://aclanthology.org/2021.naacl-main.389.pdf> (Accessed: 20 January 2026).
- Clayton and Rothstein 1991** Clayton, J. and Rothstein, E. (1991) "Figures in the corpus: Theories of influence and intertextuality", in J. Clayton and E. Rothstein (eds) *Influence and intertextuality in literary history*. Madison: University of Wisconsin Press, pp. 3–36.
- Coffee et al. 2013** Coffee, N. et al. (2013) "The Tesserae Project: Intertextual analysis of Latin poetry", *Literary and Linguistic Computing*, 28(2), pp. 221–228.

- Eke 2003** Eke, M.N. (2003) "Emecheta, Buchi", in S. Gikandi (ed.) *Encyclopedia of African literature*. London: Routledge, pp. 239–244.
- Eliot 1919** Eliot, T.S. (1919) "Tradition and the individual talent", *The Egoist*, 6(4–5), pp. 54–55 & 72–73.
- Erk 2012** Erk, K. (2012) "Vector space models of word meaning and phrase meaning: A survey", *Language and Linguistics Compass*, 6(10), pp. 635–653.
- Firth 1957** Firth, J.R. (1957) "A synopsis of linguistic theory 1930–1955", in J.R. Firth (ed.) *Studies in linguistic analysis*. Oxford: Blackwell, pp. 1–32.
- Fischer and Skorinkin 2020** Fischer, F. and Skorinkin, D. (2020) "Social network analysis in Russian literary studies", in D. Gritsenko, M. Wijermars, and M. Kopotev (eds) *The Palgrave handbook of digital Russia studies*. London: Palgrave Macmillan, pp. 517–536.
- Forstall and Scheirer 2019** Forstall, C.W. and Scheirer, W.J. (2019) *Quantitative intertextuality: Analyzing the markers of information reuse*. Cham: Springer.
- Gavin 2022** Gavin, M. (2022) *Literary mathematics: Quantitative theory for textual studies*. Stanford: Stanford University Press.
- Gikandi 2003** Gikandi, S. (2003) "Tutuola, Amos", in S. Gikandi (ed.) *Encyclopedia of African literature*. London: Routledge, pp. 752–754.
- Gower and Dijksterhuis 2004** Gower, J.C. and Dijksterhuis, G.B. (2004) *Procrustes problems*. Oxford: Oxford University Press.
- Griswold 2000** Griswold, W. (2000) *Bearing witness: Readers, writers, and the novel in Nigeria*. Princeton: Princeton University Press.
- Griswold and Bastian 1990** Griswold, W. and Bastian, M. (1990) "A bibliographic listing of Nigerian novels: 1952–1990", *Journal of Commonwealth Literature*, 25(2), pp. 214–227.
- Harris 1954** Harris, Z.S. (1954) "Distributional structure", *WORD*, 10(2–3), pp. 146–162.
- Hogan 2000** Hogan, P.C. (2000) "Understanding The Palm-Wine Drinkard", *ARIEL*, 31(4), pp. 33–58.
- Hotelling 1936** Hotelling, H. (1936) "Relations between two sets of variates", *Biometrika*, 28(3–4), pp. 321–377.
- Jockers 2013** Jockers, M.L. (2013) *Macroanalysis: Digital methods and literary history*. Champaign: University of Illinois Press.
- Juvan 2008** Juvan, M. (2008) *History and poetics of intertextuality*. West Lafayette: Purdue University Press.
- Kendall 1989** Kendall, D.G. (1989) "A survey of the statistical theory of shape", *Statistical Science*, 4(2), pp. 87–99.
- Kristeva 1986** Kristeva, J. (1986) "Word, dialogue and novel", in T. Moi (ed.) *The Kristeva reader*. New York: Columbia University Press, pp. 34–61.
- Meyes et al. 2019** Meys, R. et al. (2019) "Ablation studies in artificial neural networks". arXiv. Available at: <https://arxiv.org/pdf/1901.08644> (Accessed: 20 January 2026).
- Mikolov et al. 2013** Mikolov, T. et al. (2013) "Distributed representations of words and phrases and their compositionality", *Advances in Neural Information Processing Systems*, 26, pp. 3111–3119.
- Molitor and Ritter 2020** Molitor, P. and Ritter, J. (2020) "Digital methods for intertextuality studies", *it — Information Technology*, 62(2), pp. 49–51.
- Moretti 2011** Moretti, F. (2011) "Network theory, plot analysis", *New Left Review*, 68, pp. 80–102.
- Moretti 2000** Moretti, F. (2000) "The slaughterhouse of literature", *Modern Language Quarterly*, 61(1), pp. 207–227.
- Ngugi wa Thiong'o 2013** Ngugi wa Thiong'o (2013) "Tribute to Chinua Achebe", *Leeds African Studies Bulletin*, 75, pp. 51–53.
- Ogundipe 1969** Ogundipe, M. (1969) "The Palm-Wine Drinkard: A reassessment of Amos Tutuola", *Présence Africaine*, 71, pp. 99–108.
- Oksanen et al. 2024** Oksanen, J. et al. (2024) "vegan: Community ecology package", R package version 2.6–8. Available

at: <https://CRAN.R-project.org/package=vegan> (Accessed: 20 January 2026).

- Pennington, Socher, and Manning 2014** Pennington, J., Socher, R. and Manning, C. (2014) “GloVe: Global vectors for word representation”, in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Doha: Association for Computational Linguistics, pp. 1532–1543.
- Piper 2018** Piper, A. (2018) *Enumerations: Data and literary study*. Chicago: University of Chicago Press.
- Rohlf and Slice 1990** Rohlf, F.J. and Slice, D.E. (1990) “Extensions of the Procrustes method for the optimal superimposition of landmarks”, *Systematic Zoology*, 39(1), pp. 40–59.
- Saintsbury 1900–1904** Saintsbury, G. (1900–1904) *A history of criticism and literary taste in Europe from the earliest texts to the present day*. 3 vols. London: William Blackwood and Sons.
- Schöch et al. 2021** Schöch, C. et al. (2021) “Creating the European Literary Text Collection (ELTeC): Challenges and perspectives”, *Modern Languages Open*, 1, pp. 1–19.
- Schönemann 1966** Schönemann, P.H. (1966) “A generalized solution of the orthogonal Procrustes problem”, *Psychometrika*, 31(1), pp. 1–10.
- Schubert 2020** Schubert, C. (2020) “Intertextuality and digital humanities”, *it — Information Technology*, 62(2), pp. 53–59.
- Shore 2018** Shore, D. (2018) *Cyberformalism: Histories of linguistic forms in the digital archive*. Baltimore: Johns Hopkins University Press.
- Smith, Cordell, and Mullen 2015** Smith, D.A., Cordell, R. and Mullen, A. (2015) “Computational methods for uncovering reprinted texts in antebellum newspapers”, *American Literary History*, 27(3), pp. E1–E15.
- Stephen 1874–1879** Stephen, L. (1874–1879) *Hours in a library*. 3 vols. London: Smith, Elder & Co.
- Trilcke et al. 2016** Trilcke, P. et al. (2016) “Theatre plays as ‘small worlds’? Network data on the history and typology of German drama, 1730–1930”, in *Digital Humanities 2016*. Kraków: Jagiellonian University & Pedagogical University, pp. 385–387. Available at: <https://dh2016.adho.org/abstracts/360> (Accessed: 20 January 2026).
- Turney and Pantel 2010** Turney, P.D. and Pantel, P. (2010) “From frequency to meaning: Vector space models of semantics”, *Journal of Artificial Intelligence Research*, 37(1), pp. 141–188.
- Umeh 2003** Umeh, M.L. (2003) “Nwapa, Flora”, in S. Gikandi (ed.) *Encyclopedia of African literature*. London: Routledge, pp. 553–556.
- Underwood 2019** Underwood, T. (2019) *Distant horizons: Digital evidence and literary change*. Chicago: University of Chicago Press.
- Wendlandt, Kummerfeld, and Mihalcea 2018** Wendlandt, L., Kummerfeld, J.K. and Mihalcea, R. (2018) “Factors influencing the surprising instability of word embeddings”, in *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. New Orleans: Association for Computational Linguistics, pp. 2092–2102.
- Zabus 2007** Zabuz, C. (2007) *The African palimpsest: Indigenization of language in the West African Europhone novel*. Amsterdam: Rodopi.



This work is licensed under a Creative Commons Attribution-NonCommercial 4.0 International License.