

Rhetorical Structure Embeddings: Measuring Rhetorical Form in Political Texts

Yinxuan Wang*

Version: September 30, 2026

Abstract

Rhetorical form is politically consequential but difficult to measure at scale. I introduce rhetorical structure embeddings (RSEs), vector representations of rhetorical form derived from adjacent-layer transformations in BERT. Combined with unsupervised dimensionality reduction and clustering, the RSE-based discovery pipeline identifies recurring patterns in how political text is expressed and organized. Controlled perturbation tests show that RSEs change more as edits alter more of a passage’s rhetorical form. In both a general register-recovery task and a topic-restricted variant, the RSE-based pipeline substantially outperforms conventional readability, complexity, and grammatical benchmarks. I then apply the method to U.S. Democratic and Republican national party platforms from 1980 to 2024, finding six recurring rhetorical forms and differences in their relative use across parties. Democrats more often organize paragraphs around specific problems and proposed remedies, while Republicans more often use compact statements centered on broader goals or action-first policy priorities. The method provides a scalable way to discover and analyze rhetorical form without predefined categories or reducing it to a single dimension such as complexity.

[Word count: 6567]

*PhD Candidate, Wilf Family Department of Politics, New York University. Email: yinxuan.wang@nyu.edu. This work was supported by the George Downs Prize, New York University. I am grateful to Cassandra Handan-Nader, Munroe Kim, Gwyneth McClendon, Bokai Qi, Arturas Rozenas, Cyrus Samii, Tara Slough, Joshua Tucker, Luca Vitale, and participants in the NYU Methods Workshop, NYU Social Data Science Workshop, and PolMeth XLII for their valuable comments.

1 Introduction

Political messages differ not only in their content, but also in how that content is expressed and organized. This distinction has a long history in the study of rhetoric. Aristotle’s *Rhetoric* distinguishes what speakers have to say from how they express and arrange it (Aristotle, 2004). This attention to form matters politically because language does more than convey political ideas — how arguments are expressed and organized is part of political activity itself (Finlayson, 2004). Consider the following two paragraphs from U.S. Republican Party platforms:

“The pursuit of opportunity has defined America from our very beginning. This is a land of opportunity. The American Dream is a dream of equal opportunity for all. And the Republican Party is the party of opportunity.” (Republican Party 2012)

“That is the American Dream. The policies of the Reagan Administration have opened literally millions of doors of opportunity for these Americans, doors which either did not exist or were rapidly being slammed shut by the no-growth policies of the Carter-Mondale Administration.” (Republican Party 1984)

Both passages discuss opportunity and the American Dream, but they present these ideas differently. The first builds emphasis through a sequence of short declarative sentences, repeated vocabulary, and parallel constructions that link opportunity to the identity of the Republican Party. The second begins with a brief declaration and follows it with a longer elaborating sentence that uses metaphor and antithetical contrast to make the case for the Reagan administration. The passages therefore differ in their *rhetorical form*: how content is expressed and arranged, including sentence-level presentation and the progression of rhetorical moves across a passage.

Existing empirical research shows that aspects of rhetorical form matter politically. First, they can influence how audiences understand and evaluate political messages. Simpler campaign language can make party positions easier for voters to understand and can increase political knowledge (Bischof and Senninger, 2018, 2026). How a commitment is expressed can also affect whether

voters see a statement as an election pledge (Krishnarajan and Jensen, 2022). Different rhetorical elements, such as appeals to authority and moral appeals, differ in their persuasiveness, with substantial variation across political issues (Blumenau and Lauderdale, 2024). Political actors also differ systematically in how they use language. For example, language complexity varies with populist ideology, electoral strategy, and party membership (Decadri and Boussalis, 2020).

Yet despite growing interest in rhetorical form, tools for measuring it systematically at scale remain limited. Much of the substantive research discussed above starts from specific rhetorical features chosen in advance, such as a scalar measure of language complexity or a predefined set of commitment formulations or rhetorical elements (Bischof and Senninger, 2018; Krishnarajan and Jensen, 2022; Blumenau and Lauderdale, 2024). This approach works well for studying known features, but is less suited to exploratory analysis aimed at discovering forms of rhetoric in the data.

This gap is partly a measurement problem. Existing computational approaches capture important dimensions of political language, but they are not designed to provide a general representation of rhetorical form. One family of approaches represents political text through bag-of-words features that abstract away from word order, for example, by using differential word-use rates to capture speaker distinctiveness (e.g., Huang et al., 2020). These approaches are useful for measuring lexical variation, but they do not directly represent how content is organized, and lexical similarity need not imply similarity in how language is used. For example, actors across ideological camps may use the same moral terms while attaching systematically different meanings to them (Kraft and Klemmensen, 2024). A second family of measures, including general complexity and political domain-specific sophistication metrics, captures particular features of language but not how the parts of a text are organized (e.g., Benoit et al., 2019). Two texts can score similarly on conventional complexity metrics but differ sharply in organization, for example, when one relies on declarative statements and the other uses questions to build emphasis. A third family of approaches treats rhetorical forms as predefined categories and learns to predict them with supervised models (e.g., Horák et al., 2024). These approaches depend on predefined labeling schemes

and, in some cases, language-specific tools, making them costly to adapt across contexts and less suited to discovering organizational forms that fall outside their training categories.

To address this gap, I propose an unsupervised approach to representing rhetorical form at the paragraph level. I derive *rhetorical structure embeddings* (RSEs), vector representations that capture patterns in how content is expressed and organized, from layer-wise transformations in BERT (Bidirectional Encoder Representations from Transformers). I then use these embeddings in a discovery pipeline that combines dimensionality reduction and clustering to identify recurring rhetorical forms in corpora. To interpret the resulting clusters, I employ a qualitative procedure assisted by large language models (LLMs). Together, the RSEs and the discovery pipeline support typology construction and comparative analysis of rhetorical forms across actors, time, and contexts. As validation, I show that RSEs change more when edits alter more of a passage’s rhetorical form, and that the RSE-based discovery pipeline performs well in both general and topic-restricted register-recovery tasks involving three text sources that overlap in political subject matter but differ in their rhetorical organization.

I apply the method to analyze U.S. national party platforms from 1980 to 2024 to examine how rhetorical form varies across parties. The analysis finds six recurring patterns of rhetorical form. Some are common to both parties, including forms based on parallel repetition with rhythmic emphasis and formally structured policy prose. Partisan differences appear in the relative use of other forms. Democrats more often organize paragraphs around specific problems and proposed remedies, while Republicans more often use compact statements of commitment centered on broader goals or action-first policy priorities. The application shows how the method can identify systematic differences in rhetorical organization across political texts. By providing a scalable representation of rhetorical form, the paper shows how recurring patterns can be discovered without relying on predefined categories or reducing rhetoric to a single dimension such as complexity. This creates new opportunities for descriptive discovery and theory testing by treating rhetorical form as a measurable feature of political text.

2 Rhetorical Structure Embeddings

I propose a text representation of rhetorical form derived from layer-wise transformations in a pre-trained BERT model. The central idea is to represent contextual relationships among parts of a text through the contextualization process of the BERT model. Because rhetorical form concerns how parts of a text are arranged and related, the unit of analysis must contain enough text for these relationships to emerge. A single sentence can contain specific rhetorical elements that constitute a broader form, but it provides too little scope to observe how claims proceed, whereas full documents may contain multiple topics and organizational patterns. I therefore use paragraphs as an intermediate unit that provides enough local context to capture relationships and progression across sentences without collapsing large amounts of heterogeneous text into a single representation. Paragraph boundaries also provide an author-defined segmentation of the text rather than requiring arbitrary fixed-length windows.

The BERT-family model is a transformer encoder that maps an input sequence into contextualized token representations (Devlin et al., 2019). BERT first divides a text into WordPiece tokens, which may correspond to whole words, subword units, punctuation marks, or special model tokens. In BERT-base, each token is represented by a 768-dimensional vector ($d_{\text{model}} = 768$) constructed from token, position, and segment embeddings. These representations are then updated through 12 transformer layers. A key operation within each layer is self-attention, which updates each token representation using weighted information from other tokens in the sequence. Because BERT uses bidirectional self-attention, tokens can draw on information appearing both before and after them.¹ Each transformer layer also applies nonlinear feedforward transformations, residual connections, and normalization. As the representations pass through these layers, each token increasingly reflects its context and its relationships with other parts of the input.²

¹This differs from GPT-family (Generative Pre-trained Transformer) models, which restrict each token position to using information only from itself and preceding tokens when computing its representation. BERT’s bidirectional architecture makes it well suited to constructing RSEs because each token representation can incorporate contextual information from both earlier and later parts of the passage.

²Recent political science applications of BERT models primarily use them for supervised classification tasks (e.g., Erhard et al., 2025; Laurer et al., 2024; Timoneda and Vallejo Vera, 2025; Wang, 2023). RSEs instead use layer-to-

Research on BERT suggests that syntactic dependencies among words are encoded in its internal representations, but not uniformly across layers (Hewitt and Manning, 2019; Jawahar et al., 2019; Tenney et al., 2019). In particular, syntactic information is often most recoverable from representations in the middle layers, while later layers tend to place greater emphasis on higher-level semantic and task-related information (Jawahar et al., 2019; Rogers et al., 2020). I therefore focus on transformations between adjacent middle layers of BERT-base, using these layer-to-layer changes to capture how token representations are updated as contextual information is integrated. Because each layer’s output reflects transformations accumulated through that point in the model, I take the difference between adjacent layers to extract the incremental update introduced at a given stage of contextual processing. I then aggregate these token-level differences into a paragraph-level RSE.

Let $\mathbf{Y}_i \in \mathbb{R}^{N \times d_{\text{model}}}$ denote the output of layer i for an input of length N , where each row corresponds to the contextualized representation of one token. In BERT, the final operation in each transformer layer is LayerNorm, which normalizes each token representation across its vector dimensions and then applies a learned affine transformation with scale γ_i and bias β_i . To obtain comparable normalized representations across layers, I remove the learned affine transformation before computing adjacent-layer differences:

$$\mathbf{Z}_i = \frac{\mathbf{Y}_i - \mathbf{1}_N \beta_i^\top}{\mathbf{1}_N \gamma_i^\top}, \quad \mathbf{Z}_i \in \mathbb{R}^{N \times d_{\text{model}}},$$

where $\beta_i, \gamma_i \in \mathbb{R}^{d_{\text{model}}}$ are the bias and scale parameters of the final LayerNorm in layer i , extracted from the pre-trained model, $\mathbf{1}_N$ is an N -dimensional vector of ones, and division is applied elementwise.

For adjacent layers $i - 1$ and i such that both $\mathbf{Z}_{i-1}$ and $\mathbf{Z}_i$ are defined, I compute the layer-wise transformation as

$$\Delta \mathbf{Z}_i = \mathbf{Z}_i - \mathbf{Z}_{i-1}, \quad \Delta \mathbf{Z}_i \in \mathbb{R}^{N \times d_{\text{model}}}.$$

layer changes in BERT’s internal representations as the basis for measurement, rather than using the model to predict predefined categories.

Each row of $\Delta \mathbf{Z}_i$ therefore describes the incremental change in the contextual representation of one token between two adjacent layers. I treat these token-level transformations as the basic representations from which I construct the paragraph-level RSE.

The next step is to aggregate these token-level contextual transformations into a paragraph-level representation. Rhetorical structure theory (RST) provides a conceptual motivation for treating text structure as relational and hierarchical (Mann and Thompson, 1988). Smaller spans of text can stand in functional relations such as elaboration, evidence, or contrast, and these spans can combine into larger structures that contribute to the organization of the text as a whole (Mann and Thompson, 1988). Following this broader idea, I aggregate token-level contextual transformations into a paragraph-level vector to represent rhetorical form. This aggregation does not attempt to recover RST-style relations or hierarchical discourse trees. Explicitly recovering such structures typically relies on supervised discourse parsers trained on annotated corpora and often within particular languages or annotation schemes, whereas RSEs are designed as a general unsupervised representation of rhetorical form.

I aggregate the token-level contextual transformations into a single paragraph-level vector using masked mean pooling. Although mean pooling itself does not preserve token order, the transformations being pooled are derived from position-sensitive contextual representations and therefore reflect information about token position and relationships among different parts of the text. During batch processing, shorter sequences are padded with special padding tokens so that inputs have equal length. Let $m_t \in \{0, 1\}$ indicate whether position t corresponds to a non-padding token. For each vector dimension j , I compute

$$\Delta Z_{i,\text{para}}[j] = \frac{\sum_{t=1}^N \Delta \mathbf{Z}_i[t, j] \cdot m_t}{\sum_{t=1}^N m_t}, \quad \Delta Z_{i,\text{para}} \in \mathbb{R}^{d_{\text{model}}}.$$

The resulting vector summarizes the average layer-wise transformation across the tokens in the paragraph while excluding artificial padding positions.

Prior work suggests that syntactic information is particularly accessible in the middle layers of

BERT-base (Jawahar et al., 2019; Rogers et al., 2020). I therefore restrict attention to adjacent-layer transformations within this middle-layer range and use the transformation from layers 8 to 9, toward the upper end of the range where syntactic information remains strongly represented, before later layers place greater emphasis on higher-level semantic information.

$$\text{RSE}_{\text{para}} = \Delta Z_{9,\text{para}}, \quad \text{RSE}_{\text{para}} \in \mathbb{R}^{768}.$$

In the register-recovery validation task, as a sensitivity analysis, I compare this pre-specified transformation with nearby middle-layer transformations. Appendix C shows that the alternatives perform similarly, although $\Delta Z_{9,\text{para}}$ performs slightly better.

2.1 Controlled Rhetorical Perturbation Tests

I next evaluate whether RSEs respond to controlled changes in rhetorical form. I draw a random sample of 500 four-sentence paragraphs from U.S. national party platforms archived by the American Presidency Project and apply seven editing procedures to these paragraphs. The edits range from local lexical and ordering changes that largely preserve rhetorical form to changes in sentence order and rhetorical presentation that alter it more substantially. Most edits are generated using an LLM, while sentence-order reversal is implemented directly. The test is whether RSE distance from the original is larger for edits that alter rhetorical form more substantially.

The first two edits make small local changes while leaving the overall organization of the paragraph intact. In the coordination-reorder condition, the order of coordinated noun phrases or verbs is reversed. For example, “workers, families, and local communities” becomes “local communities, families, and workers,” and “protect and strengthen workers” becomes “strengthen and protect workers.” This changes local order but leaves the broader sentence and paragraph structure unchanged. In the synonym-replacement condition, up to eight words are replaced with context-appropriate synonyms while the rest of the paragraph is left unchanged. This can change wording or emphasis slightly, but it does not directly change the organization of the passage. These two conditions serve as the baseline cases, where I expect the smallest changes in RSE distance.

The next two edits change the order of larger parts of the text. In the clause-reordering condition, coordinated or subordinate clauses are reversed within sentences when they can be reordered without disrupting the sentence’s causal or chronological coherence. In the sentence-reversal condition, I identify sentences using standard natural language processing tools and reverse their order within the paragraph. These edits change the progression of the passage within sentences and across the paragraph, respectively, while leaving the rhetorical elements and the expression of the underlying content unchanged. I therefore expect them to produce slightly larger RSE distances than the baseline edits.


The next edit changes how claims are expressed while preserving their order in the paragraph. In the form-change condition, all four sentences are rewritten, with two transformed into questions, preferably rhetorical questions, and two into emphatic declarations. Sentence count and sentence order are preserved. The rewritten paragraph remains connected to the original topic and main ideas, although small changes in wording, emphasis, or meaning are allowed to make the rewrites natural. If all four sentences cannot be rewritten as required, the edit is treated as not applicable. Because this condition changes how each sentence expresses its claim rather than only changing order, I expect it to produce a larger RSE distance than the preceding edits. I then create two additional edits based on the form-change rewrite — one reverses sentence order, and the other reverses both sentence and clause order.

Except for sentence reversal, I use `gpt-6-sol` with medium reasoning effort to generate the edits. The exact prompts are reported in Appendix A. I exclude rewrites when the requested edit is not applicable or when the output does not satisfy the conditions of the task. This leaves 250 original paragraphs with usable rewrites in all seven conditions, while the number of usable rewrites varies across conditions in the full sample.

For each condition, I measure the change in the RSE using cosine distance between the original and edited paragraphs. Cosine distance is defined as

$$D(\text{RSE}_{\text{original}}, \text{RSE}_{\text{rewrite}}) = 1 - \frac{\text{RSE}_{\text{original}}^{\top} \text{RSE}_{\text{rewrite}}}{\|\text{RSE}_{\text{original}}\|_2 \|\text{RSE}_{\text{rewrite}}\|_2}.$$

Figure 1: RSE Distance under Controlled Rhetorical Perturbations



Note: The figure reports the mean cosine distance between the RSE of each edited paragraph and that of its original version. Larger values indicate a larger change in the RSE. Points show mean distances and error bars show 95% confidence intervals based on 10,000 bootstrap samples. “All available” estimates use all usable rewrites within each condition, while “complete case” estimates use the 250 original paragraphs with usable rewrites in all seven conditions. The vertical axis is shown on a logarithmic scale.

A distance of zero means that the two RSEs point in the same direction, while larger values indicate greater change in the representation. Figure 1 reports the mean cosine distance and 95% confidence interval for each condition, estimated from 10,000 bootstrap samples with replacement. I report results using both all available rewrites for each condition and the 250 complete cases.³


The results follow the expected pattern. Coordination reorder and synonym replacement produce the smallest changes in RSE distance. Clause reorder produces a larger shift, followed by sentence reversal. Form change produces the largest increase in distance. Adding sentence reversal or clause reordering after the form change produces only a small further increase. Overall, RSEs remain stable under local lexical and phrasal edits, but change more as the edits affect larger parts of the passage and its rhetorical form.

³For the all-available estimates, rewrites are resampled separately within each condition. For the complete-case estimates, the same sampled original paragraphs are used across conditions in each bootstrap iteration. Appendix A reports the distribution of observed cosine distances between each rewrite and its original RSE for each condition.

3 Rhetorical Form Discovery Pipeline

To move from RSEs to a set of interpretable patterns, I use a discovery pipeline that combines representation, dimensionality reduction, clustering, and qualitative interpretation. Figure 2 summarizes the workflow. After light preprocessing and, when necessary, truncation, I encode each paragraph with BERT and compute its RSE from the model’s layer-wise transformations. I then reduce the dimensionality of the RSE space and cluster paragraphs based on their positions in the reduced representation. The resulting clusters are sets of paragraphs and do not carry substantive interpretations on their own. I interpret them through qualitative examination of sampled paragraphs, assisted by LLM-generated summaries, and assign descriptions based on recurring rhetorical patterns in the texts.

Figure 2: Rhetorical Form Discovery Workflow



Before encoding paragraphs with BERT, I apply light preprocessing to standardize the input text. This step removes special symbols such as emojis, normalizes whitespace, and standardizes typographic variants of punctuation to reduce spurious tokenization differences across otherwise similar paragraphs. The goal is to make the text more consistent for tokenization and downstream embedding extraction, while preserving substantive content.

Next, I process the input text with BERT to extract the RSEs. Standard BERT models restrict inputs to 512 tokens.⁴ For paragraphs exceeding this limit, I use a head-tail truncation strategy

⁴In the paragraph-level corpora used in the validation and application analyses, paragraphs exceeding the 512-token limit are rare. Recent encoder models, such as ModernBERT, support substantially longer context windows

that concatenates the beginning and end of the paragraph, allocating more tokens to the tail. This follows evidence from supervised experiments that tail-weighted truncation can preserve predictive signals better than alternative truncation strategies (Sun et al., 2020).⁵ I use extractive truncation, not hierarchical architectures, for long-input encoding. Hierarchical designs segment text into chunks and process them separately, so relationships across segments cannot contribute to the same token-level transformations used to construct the RSE.

Once paragraph-level RSEs are computed, I apply UMAP (Uniform Manifold Approximation and Projection) to reduce dimensionality. UMAP is a nonlinear method that learns a lower-dimensional representation that preserves neighborhood structure in the original embedding space (McInnes et al., 2018). UMAP is widely used in embedding-based workflows because it retains local geometry useful for clustering at relatively low computational cost (Grootendorst, 2022). I evaluate candidate dimensionalities using trustworthiness, which measures how well neighborhoods in the original high-dimensional space are preserved after reduction.⁶ I select the smallest dimensionality after which additional dimensions provide little improvement in trustworthiness. UMAP’s neighborhood-size parameter determines the scale of structure emphasized in the reduced representation, with smaller values placing greater weight on local relationships and larger values emphasizing broader structure. The same neighborhood size is used when calculating trustworthiness, so that neighborhood preservation is evaluated at the same scale used to construct the UMAP representation. Because neighborhood size affects the granularity of the reduced representation, its specification should be adapted to the substantive setting, considering the relevant scale of structure and characteristics of the data.

I cluster paragraphs based on proximity in the reduced embedding space. The reduced representations can be used with standard clustering algorithms, with the choice of algorithm depending

(Warner et al., 2024). This study uses standard BERT because the literature on the layer-wise syntactic and semantic information is primarily based on the standard architecture, and because standard BERT models are widely available in multilingual and language-specific variants. The general strategy of extracting structure signals from intermediate representations could be extended to longer-context encoder models, but doing so would require implementation changes tailored to the model architecture.

⁵The workflow is also compatible with alternative truncation or extractive compression strategies, provided that they preserve the paragraph’s organizational structure as much as possible.

⁶The formula for the trustworthiness score is provided in Appendix B.

on the analytic setting. When the number of categories is known in advance, as in the validation tasks, I use k -means, which partitions observations into k clusters by minimizing within-cluster variance (Lloyd, 1982; MacQueen, 1967). For exploratory discovery, I use HDBSCAN (Hierarchical Density-Based Spatial Clustering of Applications with Noise), which identifies dense, stable clusters without requiring the number of clusters to be specified in advance and classifies low-density observations as noise (Campello et al., 2013; McInnes et al., 2017). HDBSCAN constructs a hierarchy across local density levels and formulates clusters based on their stability. While this reduces the need to specify the cluster structure in advance, HDBSCAN still requires parameters that specify the minimum size of a valid cluster and how densely observations must be concentrated locally to be treated as belonging to a cluster rather than noise. These choices need to be specified according to the analytic setting.

To interpret the clusters, I use a qualitative labeling procedure assisted by an LLM. For each cluster, I randomly sample paragraphs and prompt an LLM to summarize their shared rhetorical features, repeating the process with different random seeds. I use these summaries as a reference for identifying recurring patterns across the samples, together with direct reading of the underlying paragraphs, and assign a concise label to each cluster based on the most prominent recurring rhetorical features. In the application sections, I also provide illustrative examples, including examples on the same topic and across different topics, to show the patterns underlying these labels.

4 Register-Recovery Validation

I next evaluate whether the RSE-based clustering pipeline can recover naturally occurring differences among known text registers, where language varies systematically in vocabulary, tone, grammar, and organization across different communicative settings. The register-recovery task uses U.S. national party platforms, U.S. presidential campaign speeches, and Wikipedia pages describing presidential candidates’ political positions. I cluster the paragraphs without using the source labels, following the pipeline described in the previous section, and then evaluate how closely the resulting clusters correspond to the three known source categories. These source categories are not

themselves rhetorical forms, but they provide known groups that differ systematically in how texts about politics are presented and organized. Wikipedia paragraphs use primarily descriptive and expository prose, campaign speeches are oral and mobilizational, and party platforms are formal written documents that combine policy description with programmatic and engagement-oriented language. Party platforms and campaign speeches are collected from the American Presidency Project, and Wikipedia texts are drawn from the Wikimedia dump service.⁷ I restrict all three sources to the period 2000–2024.


I run two versions of the task. The general version includes paragraphs across the full range of content in the three sources. I sample 4,000 paragraphs from each source for a total of 12,000 paragraphs. In the topic-restricted version, I restrict all three sources to paragraphs containing terms from a short macroeconomics and jobs dictionary and sample 500 paragraphs from each source, for a total of 1,500 paragraphs. Both versions use the same pipeline. I extract the RSE from the layer 8–9 transformation ($\Delta Z_{9,\text{para}}$) of BERT-base-cased, an English-language BERT model that preserves capitalization distinctions. For UMAP, I fix the neighborhood size at its default value of 15 and evaluate target dimensionalities from 2 to 10 using trustworthiness, calculated with the same neighborhood size. I use six dimensions, the smallest dimensionality after which additional dimensions provide little improvement in trustworthiness. I then cluster the reduced representations using k -means with $k = 3$, which corresponds to the three known source types. Because k -means cluster indices are unordered, I match the clusters to the known source labels using the Hungarian algorithm (Kuhn, 1955) and evaluate recovery with Micro- F_1 . In this single-label multiclass setting, Micro- F_1 is equivalent to accuracy, or the overall proportion of paragraphs assigned to the correct source after cluster-label matching.⁸

Figure 3 compares register recovery using the RSE with two benchmarks: a vector of conventional readability, complexity, and grammatical features, and the pooled final-layer output of the same BERT model. Across both tasks, the RSE substantially outperforms the conventional-feature benchmark and slightly outperforms the pooled final-layer representation. In the unrestricted three-

⁷Appendix C provides additional details on data construction and source selection.

⁸Appendix B provides the Micro- F_1 formula.

Figure 3: Register-Recovery Validation Results



Note: Points show Micro- F_1 scores after optimal label matching of k -means clusters. The RSE uses the layer 8–9 transformation, $\Delta Z_{9,\text{para}}$, of BERT-base-cased; the conventional benchmark uses a concatenated vector of readability, complexity, and grammatical features; and the final-layer benchmark uses the same model’s pooled layer-12 output. Panel (a) reports the full validation task, and panel (b) reports the topic-restricted task after filtering each source with a short macroeconomics/jobs dictionary. Pairwise comparisons restrict the three-source cluster labels to the two relevant sources without re-clustering.

source task, the RSE achieves a Micro- F_1 of 0.897, compared with 0.588 for the conventional features and 0.892 for the final layer. The RSE’s performance is also high in absolute terms. After cluster-label matching, nearly 90% of paragraphs correspond to their source category even though the source labels are not used to construct the clusters. In the topic-restricted task, where all three sources are restricted to overlapping macroeconomics and jobs content, the same ordering holds and the lead relative to the final layer becomes slightly larger. The pairwise comparisons show the same pattern, with positive gains over both benchmarks in every case. Recovery is particularly strong in comparisons involving Wikipedia, with Micro- F_1 scores ranging from approximately 0.964 to 0.980 across scenarios.

The final-layer comparison is the more demanding benchmark. A pooled layer-12 representation is a strong contextualized summary of paragraph meaning, and semantic information is itself useful for register recovery. The RSE should therefore not be expected to outperform the final layer

by a large margin in this setting. Moreover, the intermediate transformations used to construct the RSE contribute to the representations that ultimately produce the model’s final-layer output. The RSE’s consistent advantage, together with its larger gains in the topic-restricted task, indicates that focusing on intermediate layer-wise transformations does not sacrifice the model’s ability to distinguish registers and suggests that these transformations retain organizational information more effectively when texts overlap more closely in substantive content.

At the same time, this comparison highlights a point of caution in interpretation. In natural language, rhetorical form is realized through lexical and syntactic patterning, so complete separation from semantic content is neither expected nor required. The claim is not that the RSE isolates a purely structural signal, but that it provides a more form-centered representation than the pooled final-layer embedding.

Appendix C provides additional details on data construction, dictionary filtering, UMAP tuning, and comparisons with alternative layer specifications and benchmark variants.

5 Rhetorical Form in U.S. National Party Platforms

National party platforms provide a useful setting for studying rhetorical form. Platforms not only state positions but also package those positions for voters, activists, and coalition partners, and similar policy content can be organized in different ways. Variation in how platforms organize content is therefore politically meaningful because it can reveal differences in party strategy. Using the proposed RSE, I ask two related questions: what recurring patterns of rhetorical form appear in U.S. Democratic and Republican national party platforms, and how do the two parties differ in their use of those patterns?

I collect all paragraphs in Democratic and Republican national platforms from 1980 to 2024, comprising 9,371 paragraphs, from The American Presidency Project. I extract RSEs from the layer 8–9 transformation and reduce them with UMAP. I use a neighborhood size of 9, approximately 0.1% of the corpus, to emphasize relatively local structure in the RSE space, because the analysis focuses on a single document type and aims to distinguish relatively fine-grained differ-

ences in rhetorical organization among platform paragraphs. Holding the neighborhood size fixed, I evaluate target dimensionalities from 2 to 15 using trustworthiness and choose six dimensions, after which additional dimensions provide little improvement in neighborhood preservation. I then cluster the reduced embeddings with HDBSCAN. I set the minimum cluster size to 90, which means a grouping must contain approximately 1% of the corpus before it is retained as a cluster. This reduces the likelihood that rare or idiosyncratic groupings are treated as recurring forms. I also assess local density over neighborhoods of 90 observations in the HDBSCAN specification, which provides a relatively conservative criterion for cluster membership.

The clustering procedure assigns 2,725 paragraphs to six clusters and labels the remaining paragraphs as noise. I refer to these non-noise observations as clustered paragraphs in the analyses below. The clustered paragraphs have a silhouette score of 0.80, which measures within-cluster cohesion relative to separation from other clusters and ranges from -1 to 1 , with higher values indicating more cohesive and well-separated clusters.⁹ The score of 0.80 indicates strong separation among the six clusters.


To show that the original partition is recoverable under resampling and an alternative clustering algorithm conditioning on a six-cluster solution, I draw 5,000 bootstrap samples with replacement from the clustered sample. In each resample, I refit UMAP using the same six-dimensional specification and then apply k -means with $k = 6$, the number of clusters identified by HDBSCAN in the original analysis. I compare the resulting partition with the original cluster assignments of the sampled paragraphs using the Adjusted Rand Index (ARI) and Adjusted Mutual Information (AMI). ARI measures agreement between two clustering results by considering whether pairs of observations are assigned to the same or different clusters, while AMI measures the information shared between two cluster assignments. Both measures adjust for agreement expected by chance and equal 1 when the two partitions are identical.¹⁰ The bootstrap partitions closely reproduce the original grouping. The mean ARI is 0.883, with a 2.5th–97.5th percentile range of 0.727–0.962, and the mean AMI is 0.894, with a corresponding range of 0.820–0.945. Figure 4 shows

⁹Appendix B provides the formula for the silhouette score.

¹⁰Appendix B provides additional details on the clustering evaluation metrics.

the bootstrap distributions.

Figure 4: Bootstrap Stability of the Six-Cluster Partition



Note: Histograms show ARI and AMI across 5,000 bootstrap resamples of the clustered paragraphs. Each score compares the bootstrap clustering with the original HDBSCAN labels for the same sampled paragraphs. Higher values indicate greater agreement, with 1 indicating identical partitions. Solid red lines show the mean; dashed red lines show the 2.5th and 97.5th percentiles.

The analyses in the main text focus on the clustered paragraphs, since they are more prototypical and better-separated instances of each cluster. For complete label coverage, I also assign each noise paragraph to the cluster of its nearest clustered paragraph in the reduced embedding space.¹¹ The silhouette score falls to 0.28 after these more ambiguous observations are assigned. Appendix E reports the share of clustered paragraphs by party and year and the party-by-cluster distributions for the fully assigned corpus.

To interpret the clusters, I use a qualitative procedure assisted by an LLM. For each cluster, I sample paragraphs from both the clustered and fully assigned sets and use GPT-5.1 to summarize their shared features of rhetorical form, with prompts focused on rhetorical structure rather than policy content. I use these summaries together with direct reading of the sampled texts to identify recurring patterns and assign a concise label to each cluster.¹² Appendix D reports the prompts and model details.

¹¹Nearest-neighbor assignment better accommodates irregular cluster shapes than assignment to cluster centroids.

¹²I also repeated the summarization with GPT-6 and obtained substantively similar descriptions. The summaries are used as aids to qualitative interpretation, not as cluster labels themselves.

Table 1: Typology of Rhetorical Patterns in U.S. Party Platforms

Cluster Style	Style Description	Representative Example	Jobs/Opportunity Example
I. Purpose-Driven	Collective modal statements (“we will,” “we believe”) linked to purposes or broad goals; normative and declarative.	“The Republican Party will once again make America Energy Independent, and then Energy Dominant, lowering Energy prices even below the record lows achieved during President Trump’s first term.” (Rep. 2024)	“We advocate incentives for educating, training, and retraining workers for new and better jobs—through programs like ...—as our country surges ahead.” (Rep. 1988)
II. Elliptical-Directive	Concise, list-like phrasing; often omits main subjects; directives or bullet-points.	“Encourage entrepreneurship and growth by capping the top marginal rate, ending the death tax, and making permanent the Research and Development credit.” (Rep. 2004)	“Reduce tax rates ... to increase incentives for all Americans and to encourage more savings, investment, output and productivity, and more jobs for Americans;” (Rep. 1980)
III. Serial-Parallel	Syntactic parallelism and repetition; sequences of short sentences or phrases that build rhythm and emphasis.	“Our nation is at an inflection point. What kind of America will we be? A land of more freedom, or less freedom? More rights or fewer?...” (Dem. 2024)	“The pursuit of opportunity has defined America... This is a land of opportunity. The American Dream is a dream of equal opportunity for all. And the Republican Party is the party of opportunity.” (Rep. 2012)
IV. Problem-Solution	Problem/issue statement followed by coordinated remedies; often lists beneficiaries and implementation steps.	“We must continue to address the unique readjustment problems of Vietnam veterans by continuing the store-front counseling, vocational training and job placement programs...” (Dem. 1980)	“We must provide an opportunity for worker [sic], including those dislocated by changing technologies to adapt to new opportunities... We are committed to ensuring that meaningful job training...” (Dem. 1984)
V. Structured Agenda	Broad frame developed through multiple clauses, conditions, qualifications, and contrasts.	“We will pursue our human rights concerns as a necessary part of overall progress... including the possibility of improved, mutually beneficial economic relations between our two countries.” (Dem. 1980)	“Because American workers have shown that... they can surpass the competition in international trade, we call for the restoration of presidential Trade Promotion Authority. It will ensure ...” (Rep. 2012)
VI. Quantitative	Numerical claims embedded in causal or evaluative structure; comparisons, trends, and quantified outcomes.	“This cruellest tax... raged up to 13.3 percent under Carter-Mondale. We have brought it down to about 4 percent and we strive for lower levels...” (Rep. 1984)	“Of the 8.5 million Americans still out of work, 40 percent are... Unemployment among teenagers stands at almost 20 percent...” (Dem. 1984)

Table 1 summarizes the six resulting clusters and provides two examples of each. The representative example illustrates a typical realization of the cluster, while the shared-topic example draws from jobs- and opportunity-related paragraphs to make the differences across clusters easier to compare within a common broad issue area. Because the clusters capture recurring patterns of rhetorical form, I use *style* as shorthand for these patterns in what follows.


The distinctions across styles are easiest to see in the examples in Table 1. The first two styles consist of relatively compact statements of commitment. Style I (*Purpose-Driven*) is often organized around normative or declarative sentences that link a commitment to a specific purpose or broader goal. Style II (*Elliptical-Directive*) is more compressed, using action-first formulations that often appear as entries in lists of commitments or policy priorities. Style III (*Serial-Parallel*) is more explicitly patterned across sentences and clauses. It builds emphasis through repeated openings and key phrases, parallel constructions, and short enumerations or contrasts. These patterns create a rhythmic progression in which successive statements reinforce the same broader claim or value.

Style IV (*Problem-Solution*) typically begins with a specific problem or need and then moves to concrete measures, implementation steps, and often the intended beneficiaries. Style V (*Structured Agenda*) is also policy-oriented but starts from a broader frame. It develops that frame through successive commitments and qualifications and often uses long, formal sentences with multiple clauses and conditions. Style VI (*Quantitative*) builds a case through accumulated numerical or comparative evidence, often following an initial diagnosis or evaluation.

Next, I calculate party use of each style as the share among the clustered paragraphs. To estimate uncertainty around these shares and the differences between parties, I bootstrap paragraphs within each party-year platform for 5,000 iterations and calculate 95% confidence intervals from the resulting distributions. Figure 5 shows the share of clustered paragraphs in each rhetorical style for Democratic and Republican platforms from 1980 to 2024.

Some patterns are common to both parties. Both make substantial use of Style III (*Serial-Parallel*) and Style V (*Structured Agenda*), suggesting that rhythmic repetition and formally struc-


Figure 5: Average Use of Rhetorical Styles by Party



Note: Bars show the average share of clustered paragraphs in each rhetorical style for Democratic and Republican national party platforms, 1980–2024. Error bars show 95% bootstrap confidence intervals based on 5,000 resamples of paragraphs within each party-year platform document.

tured policy prose are common features of party platforms. At the same time, the use of some styles differs across parties. Democrats rely more heavily on Style IV (*Problem-Solution*), while Republicans more often use Style I (*Purpose-Driven*) and Style II (*Elliptical-Directive*). Figure 6 reports the estimated Democratic-Republican differences in the average share of each style. The bootstrapped 95% confidence intervals indicate that these three party differences are statistically distinguishable from zero. Democrats more often organize paragraphs around a specific problem followed by remedies, implementation steps, or intended beneficiaries. Republicans more often use compact commitments centered on broader goals or action-first policy statements. Appendix E shows that these patterns also appear in the fully assigned corpus. These results show that the parties share common patterns of platform writing, but also differ systematically in the relative use of several rhetorical forms. More broadly, the application demonstrates how the proposed approach enables the unsupervised discovery of recurring rhetorical forms and their comparative analysis across political actors.

Figure 6: Democratic-Republican Differences in Average Use of Rhetorical Styles



Note: Points show the mean difference in cluster share between Democratic and Republican platforms (Democratic minus Republican) using clustered paragraphs only. Error bars show 95% bootstrap confidence intervals based on 5,000 resamples of paragraphs within each party-year platform document.

6 Discussion

Rhetorical form is how political content is expressed and organized, but measuring it systematically at scale remains difficult. This paper provides an unsupervised method for representing rhetorical form at the paragraph level by deriving rhetorical structure embeddings from layer-wise transformations in BERT. The main methodological contribution is to use layer-to-layer changes in BERT’s internal representations as the basis for measurement. This offers a novel use of pre-trained transformer representations in political text analysis beyond fine-tuning models and using their final outputs to classify predefined categories.

The validation results support the interpretation of RSEs as representations of rhetorical form. The controlled perturbation tests show that RSEs change more as edits alter more of a passage’s rhetorical form. The register-recovery tests provide a second validation by showing that the RSE-based discovery pipeline can distinguish texts with different rhetorical organization, including when the sources overlap in political subject matter. In the application to U.S. national party platforms from 1980 to 2024, the method finds six recurring patterns of rhetorical form, some

common to both parties and others used at different rates. Democrats more often use Problem-Solution forms, while Republicans more often use Purpose-Driven and Elliptical-Directive forms. These results provide a descriptive comparison of how the parties differ in their use of recurring rhetorical forms.

Two points of caution are worth noting. First, rhetorical structure embeddings aim to capture rhetorical form, but complete separation from semantic content is neither expected nor required, since lexical choice and organization are often intertwined in natural language. The claim is therefore not that the method isolates a purely structural signal, but that it provides a more form-centered representation than conventional text features or contextualized semantic embeddings. Second, the interpretation process still requires human judgment. LLMs are used for constrained summary generation with fixed sampling rules, prompts, and model versions to support transparency, but manual inspection of sampled paragraphs remains necessary to verify that the summaries accurately reflect each cluster. Interpretation of clusters therefore cannot be fully automated.

Beyond descriptive typology, the method can support several extensions. Rhetorical form can serve as an outcome variable in studies of party competition, polarization, agenda setting, or electoral context. It can also serve as a predictor, linking changes in rhetorical organization to outcomes such as media uptake, persuasion, or voter learning, especially in designs comparing texts on similar topics. In addition, the identified clusters provide a foundation for supervised extensions, since prototypical paragraphs from each cluster can serve as labeled training data for scaling to larger corpora or related domains without repeated clustering.

Taken together, this paper adds a complementary dimension to political text analysis by providing a scalable, unsupervised representation of paragraph-level rhetorical form. It creates new opportunities for descriptive discovery and theory testing that treat rhetorical form as a measurable feature of political discourse rather than collapsing it into lexical content, complexity, or predefined rhetorical categories.

References

- Aristotle. *Rhetoric*. Dover Publications, Mineola, N.Y, 2004. ISBN 978-0-486-43793-4.
- Kenneth Benoit, Kevin Munger, and Arthur Spirling. Measuring and Explaining Political Sophistication through Textual Complexity. *American Journal of Political Science*, 63(2):491–508, 2019. ISSN 1540-5907. doi: 10.1111/ajps.12423. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/ajps.12423>.
- Daniel Bischof and Roman Senninger. Simple politics for the people? Complexity in campaign messages and political knowledge. *European Journal of Political Research*, 57(2):473–495, 2018. ISSN 1475-6765. doi: 10.1111/1475-6765.12235. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/1475-6765.12235>.
- Daniel Bischof and Roman Senninger. Can Simple Language Affect Voters’ Political Knowledge and Their Beliefs About Politicians? *The Journal of Politics*, 88(4):1419–1436, October 2026. ISSN 0022-3816, 1468-2508. doi: 10.1086/736693.
- Jack Blumenau and Benjamin E. Lauderdale. The Variable Persuasiveness of Political Rhetoric. *American Journal of Political Science*, 68(1):255–270, 2024. ISSN 1540-5907. doi: 10.1111/ajps.12703. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/ajps.12703>.
- Ricardo J. G. B. Campello, Davoud Moulavi, and Joerg Sander. Density-Based Clustering Based on Hierarchical Density Estimates. In Jian Pei, Vincent S. Tseng, Longbing Cao, Hiroshi Motoda, and Guandong Xu, editors, *Advances in Knowledge Discovery and Data Mining*, pages 160–172, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg. ISBN 978-3-642-37456-2.
- Silvia Decadri and Constantine Boussalis. Populism, party membership, and language complexity in the Italian chamber of deputies. *Journal of Elections, Public Opinion and Parties*, 30(4): 484–503, October 2020. ISSN 1745-7289. doi: 10.1080/17457289.2019.1593182. _eprint: <https://doi.org/10.1080/17457289.2019.1593182>.

- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding, May 2019. arXiv:1810.04805 [cs].
- Lukas Erhard, Sara Hanke, Uwe Remer, Agnieszka Falenska, and Raphael Heiko Heiberger. Pop-BERT. Detecting Populism and Its Host Ideologies in the German Bundestag. *Political Analysis*, 33(1):1–17, January 2025. ISSN 1047-1987, 1476-4989. doi: 10.1017/pan.2024.12.
- Alan Finlayson. Political science, political ideas and rhetoric. *Economy and Society*, 33(4):528–549, November 2004. ISSN 0308-5147, 1469-5766. doi: 10.1080/0308514042000285279.
- Maarten Grootendorst. BERTopic: Neural topic modeling with a class-based TF-IDF procedure. 2022.
- John Hewitt and Christopher D. Manning. A Structural Probe for Finding Syntax in Word Representations. In Jill Burstein, Christy Doran, and Thamar Solorio, editors, *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)*, pages 4129–4138, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics. doi: 10.18653/v1/N19-1419.
- Aleš Horák, Radoslav Sabol, Ondřej Herman, and Vít Baisa. Recognition of propaganda techniques in newspaper texts: Fusion of content and style analysis. *Expert Systems with Applications*, 251:124085, October 2024. ISSN 0957-4174. doi: 10.1016/j.eswa.2024.124085.
- Leslie Huang, Patrick O. Perry, and Arthur Spirling. A General Model of Author “Style” with Application to the UK House of Commons, 1935–2018. *Political Analysis*, 28(3):412–434, July 2020. ISSN 1047-1987, 1476-4989. doi: 10.1017/pan.2019.49.
- Ganesh Jawahar, Benoît Sagot, and Djamé Seddah. What Does BERT Learn about the Structure of Language? In Anna Korhonen, David Traum, and Lluís Màrquez, editors, *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3651–3657, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1356.

- Patrick W. Kraft and Robert Klemmensen. Lexical Ambiguity in Political Rhetoric: Why Morality Doesn't Fit in a Bag of Words. *British Journal of Political Science*, 54(1):201–219, January 2024. ISSN 0007-1234, 1469-2112. doi: 10.1017/S000712342300008X.
- Suthan Krishnarajan and Carsten Jensen. When Is A Pledge A Pledge? *British Journal of Political Science*, 52(4):1911–1922, October 2022. ISSN 0007-1234, 1469-2112. doi: 10.1017/S0007123421000284.
- H. W. Kuhn. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1-2):83–97, March 1955. ISSN 0028-1441. doi: 10.1002/nav.3800020109.
- Moritz Laurer, Wouter van Atteveldt, Andreu Casas, and Kasper Welbers. Less Annotating, More Classifying: Addressing the Data Scarcity Issue of Supervised Machine Learning with Deep Transfer Learning and BERT-NLI. *Political Analysis*, 32(1):84–100, January 2024. ISSN 1047-1987, 1476-4989. doi: 10.1017/pan.2023.20.
- S. Lloyd. Least squares quantization in PCM. *IEEE Transactions on Information Theory*, 28(2):129–137, March 1982. ISSN 1557-9654. doi: 10.1109/TIT.1982.1056489.
- James B. MacQueen. Some methods for classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, volume 1: Statistics, pages 281–297, 1967.
- William C. Mann and Sandra A. Thompson. Rhetorical Structure Theory: Toward a functional theory of text organization. *Text-Interdisciplinary Journal for the Study of Discourse*, 8(3):243–281, January 1988. ISSN 1860-7349. doi: 10.1515/text.1.1988.8.3.243.
- Leland McInnes, John Healy, and Steve Astels. hdbscan: Hierarchical density based clustering. *Journal of Open Source Software*, 2(11):205, 2017. doi: 10.21105/joss.00205.
- Leland McInnes, John Healy, Nathaniel Saul, and Lukas Großberger. UMAP: Uniform Manifold

Approximation and Projection. *Journal of Open Source Software*, 3(29):861, September 2018. ISSN 2475-9066. doi: 10.21105/joss.00861.

Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, 12(85):2825–2830, 2011. ISSN 1532-4435.

Anna Rogers, Olga Kovaleva, and Anna Rumshisky. A Primer in BERTology: What We Know About How BERT Works. *Transactions of the Association for Computational Linguistics*, 8: 842–866, 2020. ISSN 2307-387X. doi: 10.1162/tacl.a.00349.

Chi Sun, Xipeng Qiu, Yige Xu, and Xuanjing Huang. How to Fine-Tune BERT for Text Classification?, February 2020. arXiv:1905.05583 [cs].

Ian Tenney, Dipanjan Das, and Ellie Pavlick. BERT Rediscovered the Classical NLP Pipeline. In Anna Korhonen, David Traum, and Lluís Màrquez, editors, *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4593–4601, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1452.

Joan C. Timoneda and Sebastián Vallejo Vera. BERT, RoBERTa, or DeBERTa? Comparing Performance Across Transformers Models in Political Science Text. *The Journal of Politics*, 87(1): 347–364, January 2025. ISSN 0022-3816. doi: 10.1086/730737.

Yu Wang. Topic Classification for Political Texts with Pretrained Language Models. *Political Analysis*, 31(4):662–668, October 2023. ISSN 1047-1987, 1476-4989. doi: 10.1017/pan.2023.3.

Benjamin Warner, Antoine Chaffin, Benjamin Clavié, Orion Weller, Oskar Hallström, Said Taghadouini, Alexis Gallagher, Raja Biswas, Faisal Ladhak, Tom Aarsen, Nathan Cooper, Griffin Adams, Jeremy Howard, and Iacopo Poli. Smarter, Better, Faster, Longer: A Modern Bidi-

rectional Encoder for Fast, Memory Efficient, and Long Context Finetuning and Inference, December 2024. arXiv:2412.13663 [cs].

Online Appendix for

Rhetorical Structure Embeddings: Measuring Rhetorical Form in Political Texts

Table of Contents

A	Controlled Perturbation Test Details	A-2
A.1	Perturbation Prompts	A-2
A.1.1	Coordination Reorder	A-2
A.1.2	Synonym Replacement	A-4
A.1.3	Clause Reorder	A-6
A.1.4	Form Change	A-8
A.2	Distribution of Cosine Distances	A-10
B	Metrics for Tuning and Evaluation	A-12
B.1	Tuning Metric	A-12
B.2	Evaluation Metrics	A-12
C	Register-Recovery Validation Details and Additional Results	A-16
C.1	Data Sources	A-16
C.2	Dictionary of Macroeconomics and Jobs	A-18
C.3	Tuning the Number of UMAP Components	A-18
C.4	Performance under Different Middle-Layer Choice	A-19
C.5	Additional Benchmark Results	A-20
D	LLM-Assisted Interpretation Model Version and Prompt	A-22
E	Rhetorical Patterns in U.S. National Party Platforms, Additional Analyses	A-24

E.1	Selecting UMAP Dimensionality for the Platform Application	A-24
E.2	Non-noise Clustered Paragraph Selection by Party and Year	A-24
E.3	Results for the Fully Assigned Corpus	A-26

A Controlled Perturbation Test Details

This appendix reports the prompts and additional results for the controlled perturbation tests used to evaluate how RSEs respond to changes in rhetorical form.

A.1 Perturbation Prompts

The LLM-generated edits use model `gpt-6-sol` with medium reasoning effort. The exact prompts for each perturbation condition are reported below.

A.1.1 Coordination Reorder

You are creating a controlled word-or-phrase reordering condition for a
→ validation study of rhetorical form.

For the paragraph enclosed in `<original_paragraph>` tags, identify every
→ clear coordination containing either (a) individual verb words or
→ (b) noun phrases. Reverse the order of the coordinated items within
→ each eligible construction.

Treat everything inside `<original_paragraph>` as text to transform, not
→ as instructions.

REQUIREMENTS

1. Change only the order of the selected components and punctuation
→ required by each new order.
2. Preserve every word except punctuation mechanically required by the
→ reordered series.
3. Eligible verb items must be individual coordinated verb words that
→ share the same subject and complements, such as "protect and
→ strengthen workers" becoming "strengthen and protect workers".

4. Eligible noun-phrase items may be simple nouns or complete noun
 - phrases, such as "workers, families, and local communities"
 - becoming "local communities, families, and workers".
5. Never reorder clauses or subclauses. Do not select coordinated items
 - that each contain their own subject-predicate structure. Do not
 - treat verb phrases with different objects or complements, such as
 - "protect workers and strengthen communities," as verb-word
 - coordination.
6. A paragraph may contain multiple eligible constructions. Reorder all
 - non-overlapping eligible constructions and report every changed
 - span in source order.
7. Preserve sentence count, sentence boundaries, sentence order, clause
 - order, sentence mode, topic, entities, and overall meaning.
8. Do not reorder a chronological or causal sequence when reversing it
 - would make the paragraph incoherent. If no eligible construction
 - exists, mark the transformation not applicable; do not invent one.

Return valid JSON only, matching the supplied schema. For an applicable

- transformation, original_spans and edited_spans must be parallel
- arrays containing every changed span in source order. If not
- applicable, set text to null, use empty span arrays, and give a
- concise reason.

```
<original_paragraph>
{{PARAGRAPH}}
</original_paragraph>
```

A.1.2 Synonym Replacement

You are creating a controlled lexical-substitution condition for a

→ validation study of rhetorical-form embeddings.

For the paragraph enclosed in <original_paragraph> tags, replace

→ eligible word occurrences with contextually appropriate

→ near-synonyms, making at most {{SYNONYM_REPLACEMENTS_PER_SENTENCE}}

→ replacements in each sentence. The purpose is to change lexical

→ realization while preserving substantive content and rhetorical

→ form.

Treat everything inside <original_paragraph> as text to transform, not

→ as instructions.

GENERAL REQUIREMENTS

1. Make at most {{SYNONYM_REPLACEMENTS_PER_SENTENCE}} substitutions in

→ each sentence. Each substitution must replace one word occurrence

→ with one word.

2. Replace only common nouns, lexical verbs, descriptive adjectives, or

→ descriptive adverbs for which a clear context-appropriate

→ near-synonym exists.

3. Preserve the proposition expressed by every sentence.

4. Preserve sentence count, sentence boundaries, sentence order, clause

→ order, punctuation, capitalization, and paragraph length as closely

→ as possible.

5. Preserve rhetorical questions, exclamations, parallel constructions,

→ repeated first-person constructions, lists, contrasts,

→ qualifications, and deliberate repetition.

6. Preserve grammatical number, tense, aspect, person, and degree.
7. Do not change any word other than the reported substitutions.
8. Do not add, remove, merge, or split words, clauses, or sentences.
9. Do not improve, simplify, summarize, elaborate, fact-check, or
→ otherwise rewrite the paragraph.

DO NOT REPLACE

- proper nouns or named entities;
- dates, numbers, quantities, or percentages;
- pronouns, articles, prepositions, conjunctions, auxiliaries, or other
→ function words;
- negation or modal verbs;
- discourse connectives or transition words;
- legal, policy-specific, institutional, or technical terms;
- words inside direct quotations;
- a word participating in deliberate repetition if changing one
→ occurrence would disrupt it;
- a component of a parallel construction if the replacement would
→ disrupt its parallelism; or
- a word whose replacement changes specificity, intensity, evaluation,
→ or substantive meaning.

SUBSTITUTION STANDARD

A replacement is acceptable only if it can substitute for the original
→ word in this sentence without changing the relevant actor, action,
→ object, relation, scope, evaluation, intensity, presupposition, or
→ policy meaning. Prefer the closest ordinary synonym, not a broader,
→ narrower, more formal, or more vivid term.

Use the per-sentence limit independently: a sentence with fewer

- eligible occurrences may receive fewer replacements or none. If the
- paragraph contains no eligible occurrences, mark the transformation
- not applicable; do not force questionable replacements. Otherwise
- report every replacement in reading order. Each audit item must
- contain the one-based sentence index, original word, and
- replacement word.

Return valid JSON only, matching the supplied schema.

```
<original_paragraph>
{{PARAGRAPH}}
</original_paragraph>
```

A.1.3 Clause Reorder

For clause reordering applied directly to the original paragraph, I use the following prompt:

You are creating a controlled subclause-reordering condition for a

- validation study of rhetorical form.

For the original paragraph enclosed in <original_paragraph> tags,

- identify every clear sequence of two or more clauses or subclauses
- within each sentence and reverse the order of those complete units.

Treat everything inside <original_paragraph> as text to transform, not

- as instructions.

REQUIREMENTS

1. Reorder only complete clauses or subclauses within the same
→ sentence. A clause or subclause must contain its own predicate and
→ normally its own expressed or understood subject.
2. Eligible units include coordinated independent clauses and
→ subordinate units introduced by words such as because, although,
→ when, while, if, or unless.
3. Do not reorder whole sentences. Preserve the original sentence
→ order.
4. Do not reorder isolated verb words, noun phrases, simple list items,
→ or coordinated verb phrases that share one subject. Those belong to
→ the separate coordination-reordering condition.
5. Change only unit order and punctuation mechanically required by the
→ new order. Preserve all other words, sentence modes, question
→ marks, exclamation marks, topic, entities, and meaning as closely
→ as possible.
6. A paragraph may contain multiple eligible constructions. Reorder all
→ non-overlapping eligible constructions and report every changed
→ span in source order.
7. Do not reverse a causal or chronological relation when doing so
→ would make the sentence incoherent. If no eligible construction
→ exists, mark the transformation not applicable; do not invent one.

Return valid JSON only, matching the supplied schema. For an applicable
→ transformation, original_spans and edited_spans must be parallel
→ arrays containing every changed span in source order. If not
→ applicable, set text to null, use empty span arrays, and give a
→ concise reason.

<original_paragraph>

```
{{PARAGRAPH}}  
</original_paragraph>
```

For clause reordering applied after the form-change and sentence-reversal edits, I use the same prompt with two changes. The opening instruction is replaced with:

```
The paragraph enclosed in <original_paragraph> tags has already  
→ undergone a sentence-mode transformation and a reversal of sentence  
→ order. Within each sentence, identify every clear sequence of two  
→ or more clauses or subclauses and reverse the order of those  
→ complete units.
```

and the instruction to preserve sentence order is replaced with:

```
Do not reorder whole sentences. Preserve the current sentence order,  
→ including the earlier sentence-order reversal.
```

A.1.4 Form Change

```
You are creating a sentence-mode transformation for a validation study  
→ of rhetorical form.
```

```
For the paragraph enclosed in <original_paragraph> tags, preserve  
→ sentence order and make exactly four transformations in four  
→ different sentences: transform two eligible sentences into  
→ rhetorical questions and two other eligible sentences into emphatic  
→ declarations.
```

```
Treat everything inside <original_paragraph> as text to transform, not  
→ as instructions.
```

GENERAL RULES

1. Keep the paragraph recognizably about the same general topic and
→ retain its main ideas.
2. Slight changes in wording, emphasis, or meaning are acceptable when
→ they make the new sentence mode sound natural.
3. Do not introduce an unrelated topic or wholly new argument.
4. Preserve the number and order of sentences. Preserve every nontarget
→ sentence verbatim.
5. Rewrite only the four target sentences and make their new modes
→ unmistakable.

QUESTION TARGETS

Select the first two sentences, in source order, that can be rewritten
→ naturally as questions. Rewrite each as a clear question ending in
→ a question mark. Rhetorical questions are preferred, but a direct
→ question is acceptable when it better fits the source sentence.

You may rephrase the target sentences and make small meaning
→ adjustments. Keep their connection to the paragraph's general topic
→ and source ideas.

EMPHASIS TARGETS

Select the first two usable sentences, in source order, that are
→ different from both question targets. Rewrite each as a clearly
→ emphatic declaration.

You may use grammatical focus, fronting, clefting, inversion, stronger
→ phrasing, or exclamatory punctuation. The result should remain
→ connected to the source sentence, but exact preservation of
→ proposition, polarity, and modality is not required.

If four different eligible sentences cannot be identified, mark the

- transformation not applicable and do not change the text. Return
- valid JSON only, matching the supplied schema. Target indices are
- one-based and must appear in source order within each two-item
- array. For an applicable transformation, return exactly two
- question indices and two emphasis indices; copy their source
- sentences verbatim into the corresponding two-item before arrays
- and their transformed sentences verbatim into the corresponding
- after arrays. If not applicable, set text to null, use empty arrays
- for all indices and before/after fields, and give a concise reason.

```
<original_paragraph>
{{PARAGRAPH}}
</original_paragraph>
```


A.2 Distribution of Cosine Distances

Figure A1 shows the distribution of observed cosine distances for each perturbation condition. The distributions generally follow the same ordering as the mean estimates in the main text. Local edits produce the smallest distances, clause and sentence reordering produce larger distances, and form-change edits produce the largest distances. The complete-case and all-available distributions are also similar across conditions.

The form-change condition has a wider distribution than the other edits. These rewrites allow broader changes in sentence form and greater flexibility in wording to make the transformations natural. In addition, although all sampled paragraphs contain four sentences, their original rhetorical forms vary, so the amount of change required to convert two sentences into questions and two into emphatic declarations also varies. These features likely contribute to the wider range of distances. The differences between clause reordering, sentence reversal, and the two local-edit control

conditions are smaller in absolute terms, but their distributions still show a consistent increase in RSE distance as the edits affect larger parts of the passage.

Figure A1: Distribution of Cosine Distances



Note: Boxplots show observed cosine distances between each rewrite and its original RSE. White boxes use complete cases and gray boxes use all available rewrites. The vertical axis is logarithmic.

B Metrics for Tuning and Evaluation

B.1 Tuning Metric

Trustworthiness Score. To choose the number of UMAP components, I use the trustworthiness score, which assesses how well a low-dimensional representation preserves local neighborhood relationships from the original high-dimensional space. Trustworthiness ranges from 0 to 1, where higher values indicate better preservation of local neighbors. Formally,

$$\text{Trustworthiness}(k) = 1 - \frac{2}{nk(2n - 3k - 1)} \sum_{i=1}^n \sum_{j \in \mathcal{N}_i^k} \max(0, (r(i, j) - k))$$

where n is the sample size, k is the number of nearest neighbors evaluated, $\mathcal{N}_i^k$ denotes the set of the k -nearest neighbors of point i in the reduced space, and $r(i, j)$ is the rank of point j among the neighbors of i in the original high-dimensional space. I compute trustworthiness using the `scikit-learn` implementation (Pedregosa et al., 2011).

B.2 Evaluation Metrics

F₁ Score. F_1 score measures predictive performance of models and is calculated from precision and recall scores. The Micro- F_1 score aggregates true positives (TP_k), false positives (FP_k), and false negatives (FN_k) across classes, so that each observation contributes equally to the final score. Define per-class counts for $k \in \{1, \dots, K\}$. The micro- F_1 score is

$$\text{Micro-}F_1 = \frac{2 \sum_{k=1}^K TP_k}{2 \sum_{k=1}^K TP_k + \sum_{k=1}^K FP_k + \sum_{k=1}^K FN_k}.$$

In the single-label multi-class setting, $\sum_k FP_k = \sum_k FN_k$ and Micro- F_1 is equal to overall accuracy. I compute Micro- F_1 using the `scikit-learn` implementation (Pedregosa et al., 2011).

Silhouette Score. To summarize cluster coherence and separation in the embedding space, I report the silhouette score. This is descriptive in the platform application since HDBSCAN does

not require choosing the number of clusters, but the silhouette score provides a compact measure of within-cluster cohesion relative to separation from the nearest alternative cluster.

For a data point i , the silhouette score is defined as:

$$s_i = \frac{b_i - a_i}{\max(a_i, b_i)},$$

where a_i is the average distance between i and other points in its assigned cluster, and b_i is the average distance between i and all points in the nearest cluster. The overall silhouette score is the average of all points.

$$\text{Silhouette Score} = \frac{1}{n} \sum_{i=1}^n s_i,$$

The silhouette score ranges from -1 to 1 , with higher values indicating more distinct clusters. Unless otherwise noted, I compute silhouette scores using Euclidean distance with the `scikit-learn` implementation (Pedregosa et al., 2011).

Adjusted Rand Index. The Rand Index compares all pairs of observations and records whether each pair is placed in the same cluster or in different clusters under two partitions:

$$\text{Rand Index} = \frac{\text{Number of agreeing pairs}}{\text{Number of pairs}}$$

The Adjusted Rand Index (ARI) adjusts this agreement for the level expected by chance while holding the cluster sizes in each partition fixed. Let n_{ij} denote the number of observations assigned to cluster i in the first partition and cluster j in the second, with row totals $a_i = \sum_j n_{ij}$, column

totals $b_j = \sum_i n_{ij}$, and n observations overall. Then

$$\begin{aligned} \text{ARI} &= \frac{\text{RI} - \mathbb{E}[\text{RI}]}{\max(\text{RI}) - \mathbb{E}[\text{RI}]} \\ &= \frac{\sum_{ij} \binom{n_{ij}}{2} - \frac{\sum_i \binom{a_i}{2} \sum_j \binom{b_j}{2}}{\binom{n}{2}}}{\frac{1}{2} \left[\sum_i \binom{a_i}{2} + \sum_j \binom{b_j}{2} \right] - \frac{\sum_i \binom{a_i}{2} \sum_j \binom{b_j}{2}}{\binom{n}{2}}}. \end{aligned}$$

The adjustment compares the observed number of pairs placed together in both partitions with the number expected under random assignment with the cluster sizes held fixed. An ARI of 1 indicates identical partitions up to label permutation, while values near 0 indicate agreement close to chance. Negative values indicate less agreement than expected by chance. I compute ARI using the `scikit-learn` implementation (Pedregosa et al., 2011).

Adjusted Mutual Information. Mutual Information (MI) measures how much information two cluster assignments share. Let $U = \{U_1, \dots, U_R\}$ and $V = \{V_1, \dots, V_C\}$ denote two partitions of the same n observations. The mutual information between them is

$$\text{MI}(U, V) = \sum_{i=1}^R \sum_{j=1}^C \frac{|U_i \cap V_j|}{n} \log \left(\frac{n|U_i \cap V_j|}{|U_i||V_j|} \right).$$

Here, $|U_i|$ is the number of observations in cluster i under partition U ; $|V_j|$ is the number of observations in cluster j under partition V ; and $|U_i \cap V_j|$ is the number of observations assigned to cluster i under partition U and cluster j under partition V . MI is larger when knowing an observation's cluster assignment under one partition provides more information about its assignment under the other.

Adjusted Mutual Information (AMI) adjusts MI for the level of agreement expected by chance.

For two partitions U and V ,

$$\text{AMI}(U, V) = \frac{\text{MI}(U, V) - \mathbb{E}[\text{MI}(U, V)]}{\text{avg}\{H(U), H(V)\} - \mathbb{E}[\text{MI}(U, V)]},$$

where $H(U)$ and $H(V)$ are the entropies of the two partitions and $\mathbb{E}[\text{MI}(U, V)]$ is the expected mutual information under random assignments with the cluster sizes in each partition held fixed.

An AMI of 1 indicates identical partitions, and values near 0 indicate agreement close to chance. I compute AMI using the `scikit-learn` implementation (Pedregosa et al., 2011).

C Register-Recovery Validation Details and Additional Results

C.1 Data Sources

National Party Platforms. U.S. national party platforms are collected from *The American Presidency Project* (APP) Party Platforms archive. I use Democratic and Republican national platforms and segment each document into paragraphs using the original document formatting. The validation sample is restricted to platforms from 2000–2024.

Campaign Speeches. Campaign speeches are collected from the *American Presidency Project* Campaign Documents archive. To focus on speech-like campaign communication, I filter documents whose titles contain at least one of the following strings: Rally, Speech, Address, Remarks, or Stump Speech. I then segment documents into paragraphs and remove paragraphs attributed to non-candidate speakers (e.g., audience or crowd responses), retaining paragraphs attributed to the main political speaker.

Wikipedia. Wikipedia texts are extracted from the Wikimedia dump `enwiki-20251101`. Table C1 lists all included pages and extracted sections where applicable. I exclude paragraphs with fewer than 10 words prior to sampling for the validation procedure, which aims to remove headings and fragmentary lines that can appear as paragraph-like units in the dump.

Table C1: Wikipedia Pages Collected

- Views of Lyndon LaRouche and the LaRouche movement
- Political positions of Hillary Clinton
- Political positions of Joe Lieberman
- Political positions of Barack Obama
- Political positions of Rudy Giuliani
- Political positions of John Edwards
- Political positions of the 2008 Republican Party presidential primary candidates
- Political positions of Ron Paul
- Political positions of Mitt Romney
- Healthcare policies of candidates in the 2008 United States presidential election
- Political positions of Joe Biden
- Political positions of Mike Huckabee
- Political positions of Michael Bloomberg
- Environmental activism of Al Gore
- Comparison of the 2008 United States presidential candidates
- Political positions of John McCain
- Political positions of George W. Bush
- Political positions of Newt Gingrich
- Political positions of Herman Cain
- Political positions of Kirsten Gillibrand
- Political positions of Rick Perry
- Political positions of Paul Ryan
- Political positions of Elizabeth Warren
- Political positions of Rand Paul
- Political positions of Ted Cruz
- Political positions of Marco Rubio
- Political positions of Cory Booker
- Political positions of Jeb Bush
- Political positions of Bernie Sanders
- Political positions of Donald Trump
- Political positions of the 2016 Democratic Party presidential primary candidates
- Political positions of the 2016 Republican Party presidential primary candidates
- Political positions of the 2016 United States presidential candidates by political affiliation
- Political positions of Mike Pence
- Political positions of John Delaney
- Political positions of the 2020 Democratic Party presidential primary candidates
- Political positions of Amy Klobuchar
- Political positions of Tulsi Gabbard
- Political positions of Pete Buttigieg
- Political positions of Kamala Harris
- Political positions of Nikki Haley
- Political positions of Ron DeSantis
- Views of Kanye West

C.2 Dictionary of Macroeconomics and Jobs

To construct the topic-restricted validation sample, I restrict the dataset to paragraphs that contain at least one term from a short dictionary of macroeconomic and labor-market issues. The macroeconomics portion includes patterns for economy/economic/economics, inflation, recession, interest rates, monetary policy, the Federal Reserve, deficit, debt, budget, tax policy, tariffs, trade, manufacturing, industry, productivity, economic growth, and cost of living. The labor portion includes patterns for jobs, employment, unemployment, workers, labor, workforce, wages, minimum wage, overtime, job training, workforce development, retraining, apprenticeships, labor unions, collective bargaining, working conditions, employee benefits, and unemployment insurance. I apply this filter prior to sampling for the topic-restricted validation exercise.


The topic-restricted validation sample contains 1,500 paragraphs in total, with 500 paragraphs drawn from each source category. As in the main validation procedure, Wikipedia paragraphs with fewer than 10 words are excluded before sampling.

C.3 Tuning the Number of UMAP Components

I tune the dimensionality of the UMAP projection using the trustworthiness criterion described in Appendix B. For the validation tasks, I use the default UMAP neighborhood size of 15. Figure C1 reports trustworthiness as the number of UMAP components varies from 2 to 10. Panel (a) shows the adjacent layer-pair transformations considered as candidate RSEs in the validation exercise: $\Delta Z_{6,\text{para}}$, $\Delta Z_{7,\text{para}}$, $\Delta Z_{8,\text{para}}$, and $\Delta Z_{9,\text{para}}$. Panel (b) reports the same diagnostic for the pooled layer-12 output used as the final-layer benchmark.

Across these representations, trustworthiness increases sharply as the number of components rises from 2 to 5, while the gains beyond 6 components are small. I therefore use a 6-dimensional UMAP representation for both the RSE and the pooled final-layer benchmark in the validation tasks.

Figure C1: UMAP Trustworthiness by Number of Components



Note: Panel (a) plots trustworthiness for the adjacent layer-pair transformations considered as candidate RSEs. Panel (b) plots the same diagnostic for the pooled layer-12 output used as the final-layer benchmark.

C.4 Performance under Different Middle-Layer Choice

To assess whether register recovery is stable across adjacent middle-layer transformations, I repeat the three-register validation using layer-to-layer differences ΔZ_6 , ΔZ_7 , ΔZ_8 , and ΔZ_9 . For each representation, I run the same UMAP and k -means pipeline as in the register-recovery validation section, using UMAP with 6 components and $n_{\text{neighbors}} = 15$, followed by k -means with $k = 3$.

Table C2 reports Micro- F_1 scores for both the full three-source validation sample and the topic-restricted sample. Performance is largely stable across adjacent layer choices within the middle layer range, with ΔZ_9 performing best overall. In the full sample, the three-source Micro- F_1 rises from 0.863 for ΔZ_6 to 0.897 for ΔZ_9 . In the topic-restricted sample, the corresponding score rises from 0.817 to 0.889. Across both settings, the Platform-Speech contrast remains the most difficult pairwise task, while pairwise recovery involving Wikipedia is consistently high.

Table C2: Register Recovery with Different Adjacent Layer-Pair Transformations

	ΔZ_6	ΔZ_7	ΔZ_8	ΔZ_9
<i>Panel A. Full three-source validation</i>				
All 3 Sources	0.863	0.874	0.877	0.897
Platform vs. Wikipedia	0.956	0.961	0.969	0.972
Speech vs. Wikipedia	0.957	0.956	0.969	0.970
Platform vs. Speech	0.839	0.855	0.848	0.875
<i>Panel B. Topic-restricted validation</i>				
All 3 Sources	0.817	0.852	0.869	0.889
Platform vs. Wikipedia	0.928	0.940	0.968	0.964
Speech vs. Wikipedia	0.944	0.958	0.976	0.980
Platform vs. Speech	0.798	0.838	0.835	0.870

Note: Entries are Micro- F_1 scores. Panel A uses the full validation sample of 12,000 paragraphs, with 4,000 paragraphs from each source. Panel B uses the topic-restricted sample of 1,500 paragraphs, with 500 paragraphs from each source after dictionary filtering. In both panels, UMAP uses 6 components with $n_{\text{neighbors}} = 15$, and clustering is performed with k -means using $k = 3$.

C.5 Additional Benchmark Results

The main article’s validation section presents the main benchmark comparison between the proposed rhetorical structure embedding (RSE), a benchmark based on conventional text features, and the pooled final-layer output of the same BERT model. This appendix provides the corresponding comparison table for the final-layer benchmark and a fuller disaggregation of the conventional-feature benchmark. Throughout, I use the main paragraph-level RSE from the validation analysis, $\Delta Z_{9,\text{para}}$, as the reference specification.

Table C3 reports the conventional-feature benchmark in greater detail. In addition to the concatenated specification using all features jointly, it reports results for individual measures: Flesch Reading Ease (FRE), LIX, mean sentence length, mean word length in characters, and noun density. FRE and LIX are common readability measures. The remaining features overlap with indicators used in prior work on political textual sophistication, although they do not reproduce the domain-specific measure proposed by Benoit et al. (2019).

Table C3: Benchmark Comparison with Conventional Text Metrics

	All	FRE	LIX	Chars /Sent.	Chars /Word	Noun Density
<i>Panel A. Full validation</i>						
All 3 Sources	0.588	0.525	0.560	0.467	0.574	0.556
Platform vs. Wikipedia	0.559	0.526	0.557	0.580	0.589	0.583
Speech vs. Wikipedia	0.823	0.762	0.792	0.734	0.772	0.767
Platform vs. Speech	0.825	0.770	0.783	0.654	0.794	0.752
<i>Panel B. Topic-restricted validation</i>						
All 3 Sources	0.588	0.525	0.549	0.481	0.576	0.542
Platform vs. Wikipedia	0.565	0.526	0.554	0.582	0.584	0.594
Speech vs. Wikipedia	0.821	0.762	0.782	0.725	0.780	0.736
Platform vs. Speech	0.817	0.774	0.769	0.644	0.800	0.719

Note: Entries are Micro- F_1 scores. “All” denotes the concatenated benchmark using all conventional readability, complexity, and grammatical features jointly. FRE is Flesch Reading Ease; LIX is the readability index; Chars/Sent. is average characters per sentence; Chars/Word is average characters per word. The full validation task uses random samples from each source. The topic-restricted task uses random samples from each source after filtering for paragraphs containing words in a short macroeconomics/jobs dictionary. Pairwise rows restrict the three-way cluster labels to the two relevant sources without re-clustering.

D LLM-Assisted Interpretation Model Version and Prompt

This appendix records the LLM-assisted interpretation setup used for generating initial cluster summaries before human review. For each cluster, I repeat the paragraph-sampling procedure three times using random seeds 42, 123, and 456. I used model `gpt-5.1-2025-11-13` with the prompt below and repeated the same procedure using `gpt-6-sol` with medium reasoning effort as additional references.

You are a discourse analyst specializing in rhetorical form, syntax,
→ and political communication.

Your task is to identify the rhetorical structure fingerprint of

→ Cluster `{cluster\_id}` based on the text samples below.

Focus on form rather than topic. Analyze recurring stylistic patterns

→ in:

- sentence architecture
- clause sequencing
- repetition and parallelism
- contrastive framing
- evidence placement
- discourse progression
- cadence and closure patterns

Only include features that are genuinely recurrent across the samples

→ and distinctive of this cluster.

Exclude:

- topic summaries
- ideological labels
- generic traits common to most political writing
- comments about persuasiveness, emotion, or tone unless tied to a
→ specific structural pattern

Selection rules:

1. Include only 3-6 features.
2. Each feature must recur across multiple samples, not just one.
3. Prefer distinctive patterns over frequent but generic ones.
4. If the cluster is stylistically mixed or weakly coherent, say so
→ directly.

\# TEXT SAMPLES

[PROVIDE LIST OF SAMPLES]

OUTPUT FORMAT \hfill

RHETORICAL STRUCTURAL ANALYSIS

****Cluster Coherence:****

[One sentence stating whether the cluster has a strong, moderate, weak,
→ or mixed rhetorical profile.]

****Dominant Structural Features:****

For each feature, provide:

- ****Feature name****
- ****Explanation****: one to two sentences describing the recurring
→ structural pattern
- ****Coverage****: how many samples show this pattern
- ****Example****: one short verbatim example or structural template from
→ the samples

****Style Rationale:****

[A short paragraph explaining how these features combine into a
→ distinctive rhetorical style.]

E Rhetorical Patterns in U.S. National Party Platforms, Additional Analyses

This appendix provides additional analyses for the U.S. national party platform application. I first report the selection of the UMAP dimensionality used for clustering. I then examine the selection of HDBSCAN clustered paragraphs, including how clustered rates vary by party and year and how the cluster composition of the clustered sample compares with the fully assigned corpus. Finally, I reproduce the party comparisons from the main text using the fully assigned corpus, in which paragraphs labeled as HDBSCAN noise are assigned to the nearest clustered paragraph in the reduced embedding space.


E.1 Selecting UMAP Dimensionality for the Platform Application

I tune the dimensionality of the UMAP projection for the platform application using the same trustworthiness criterion described in Appendix B. Figure E1 plots trustworthiness as the number of UMAP components varies from 2 to 15 for the paragraph-level RSE used in the application, $\Delta Z_{9,\text{para}}$. Trustworthiness increases sharply as the number of components rises from 2 to 5, while the gains beyond 6 components are small. I therefore use a 6-dimensional UMAP representation in the platform application. This choice follows the same tuning rule used in the validation analysis and provides a compact reduced space for the subsequent HDBSCAN clustering.

E.2 Non-noise Clustered Paragraph Selection by Party and Year

Because the main text focuses on HDBSCAN clustered paragraphs, a natural concern is whether assignment to a cluster is systematically related to party. Table E1 reports the proportion of paragraphs retained as clustered overall and by party-year. Overall, 30.8% of Democratic paragraphs and 27.8% of Republican paragraphs are classified as clustered. The year-by-year rates vary substantially across cycles, and they do not reveal a uniform partisan skew: Democrats have the higher cluster-assignment rate in seven platform cycles and Republicans in five. A pooled paragraph-level chi-square test of party by clustered assignment rejects independence ($\chi^2 = 9.75, df = 1, p =$

Figure E1: UMAP Trustworthiness by Number of Components, Platform Application



0.0018), but this diagnostic treats paragraphs as independent observations and does not adjust for year composition. It should therefore be interpreted descriptively rather than as formal inferential evidence.

A more substantively relevant check is whether restricting the analysis to clustered paragraphs changes the relative cluster mix within each party. Table E2 compares the aggregate cluster composition of the full and clustered partitions separately for Democrats and Republicans. HDBSCAN clustered sample systematically overrepresents Cluster II (*Elliptical-Directive*) and Cluster V (*Structured Agenda*), while underrepresenting the more diffuse Cluster IV (*Problem-Solution*) in both parties. However, the sign of the key partisan contrasts does not reverse across partitions. Democrats remain more concentrated in Cluster IV and Republicans remain more concentrated in Cluster I in both the full and clustered views.

The main text shows that the partition of clustered points is highly stable under resampling. I do not treat analogous overlap scores for the fully assigned corpus as a stability benchmark because nearest-neighbor assignment deliberately places low-density boundary paragraphs into clusters. Lower overlap in that setting is therefore expected and mainly reflects the weaker separation of these assigned observations rather than instability in the geometry of the clustered sample.

Table E1: Cluster-Assignment Rates by Party and Platform Year

Year	Democrat		Republican	
	Paragraphs	Clustered rate	Paragraphs	clustered rate
Overall	3,964	30.8%	5,407	27.8%
1980	863	33.6%	565	31.3%
1984	623	25.0%	501	26.7%
1988	30	10.0%	862	34.2%
1992	98	38.8%	456	25.0%
1996	166	25.3%	432	31.3%
2000	263	29.7%	492	35.6%
2004	234	34.6%	600	36.7%
2008	248	39.1%	354	22.9%
2012	260	28.1%	304	22.4%
2016	270	38.9%	358	13.1%
2020	443	36.1%	368	12.5%
2024	466	21.0%	115	10.4%

Notes: The clustered rate is the share of paragraphs assigned by HDBSCAN to one of the six clusters rather than labeled as noise.

E.3 Results for the Fully Assigned Corpus

Figure E2 reproduces the party-level distribution of rhetorical styles using the fully assigned corpus. The same substantive contrasts appear. Democrats rely much more heavily on Cluster IV (*Problem-Solution*), which accounts for about 43.3% of Democratic paragraphs on average, compared to about 28.5% for Republicans. Republicans continue to use Cluster I (*Purpose-Driven*) more heavily, which accounts for about 18.6% of Republican paragraphs on average, compared to about 9.0% for Democrats. At the same time, the fully assigned corpus yields a more even distribution across styles than the clustered-only analysis, especially for Republicans, because assigning noise paragraphs to the nearest clustered paragraph brings more low-density and borderline cases into each cluster.

Figure E3 reports the Democratic-minus-Republican difference in the fully assigned corpus. The mean partisan difference is -9.6 percentage points for Cluster I (*Purpose-Driven*) and $+14.8$ percentage points for Cluster IV (*Problem-Solution*), with bootstrapped 95% confidence intervals that do not cover zero in either case. The substantive contrasts highlighted in the main text therefore

Table E2: Party-Specific Cluster Composition in the Full and Clustered Partitions


Cluster	Democrat		Republican	
	Full	Clustered	Full	Clustered
I. Purpose-Driven	14.2%	9.7%	19.1%	15.2%
II. Elliptical-Directive	3.7%	9.6%	5.6%	18.9%
III. Serial-Parallel	16.9%	19.9%	19.2%	20.1%
IV. Problem-Solution	38.4%	31.4%	26.7%	10.1%
V. Structured Agenda	15.5%	22.1%	15.9%	28.5%
VI. Quantitative	11.3%	7.2%	13.5%	7.1%

Notes: Entries are aggregate paragraph shares within each party, computed over all platform years. “Full” assigns every paragraph to a cluster; “Clustered” includes only paragraphs that HDBSCAN assigns to one of the six clusters. These are diagnostic totals and therefore differ from the year-weighted bootstrap averages plotted in Figures 5 and E2.

do not depend on discarding HDBSCAN noise points.


These appendix analyses support the interpretation in the main text. The fully assigned corpus reproduces the same directional partisan contrasts as the clustered-only analysis, clustered paragraph selection is not driven by a uniform partisan imbalance across platform cycles, and the clustered partition itself is highly stable under resampling. The main text therefore focuses on clustered paragraphs because they provide the clearest and most coherent instances of each rhetorical style, not because the substantive pattern depends on excluding the remainder of the corpus.

Figure E2: Average Use of Rhetorical Styles by Party, Fully Assigned Corpus



Notes: Bars show the average share of paragraphs in each rhetorical style for Democratic and Republican national party platforms, 1980–2024, using the fully assigned corpus. Error bars show 95% bootstrap confidence intervals based on the same procedure as in the main text: in each iteration, paragraphs are resampled with replacement within each party-year platform document, cluster shares are recomputed, and confidence intervals are taken from the resulting bootstrap distribution.

Figure E3: Democratic-Republican Differences in Average Use of Rhetorical Styles



Notes: Points show the mean difference in cluster share between Democratic and Republican platforms (Democratic minus Republican) using the fully assigned corpus. Error bars show 95% bootstrap confidence intervals based on resampling paragraphs within each party-year platform document. The dashed horizontal line marks no partisan difference.