

A Manifold-Aware Topic Modeling Approach via Rank-Based Prototypes

Thiago César Castilho Almeida and Daniel Carlos Guimarães Pedronette

Department of Statistics, Applied Mathematics and Computing

State University of São Paulo (UNESP), Rio Claro, Brazil

{tc.almeida, daniel.pedronette}@unesp.br

Abstract

Recent topic models leverage pretrained embeddings, but neural architectures produce latent representations without grounding in specific texts, and clustering-based pipelines assign representative documents only post hoc, relying on absolute distances distorted by hubness and anisotropy in high-dimensional spaces. We introduce MARETopic, a training-free framework that casts topic discovery as rank-based prototype selection. After projecting embeddings onto a low-dimensional manifold, MARETopic builds ranked lists encoding ordinal neighborhood structure. A greedy algorithm selects exactly K exemplar documents, real corpus texts, whose neighborhoods cover the corpus. Two variants share this criterion. MARETopic_{Corr} scores candidates with a query performance predictor and a rank correlation measure, leading Purity and NMI on the two benchmarks with the most categories, ahead of both neural and clustering-based topic models. MARETopic_{Diff} scores them with a rank-based diffusion matrix, needs neither measure, and runs 1.7 to 1.9 times faster. Without a single gradient update, MARETopic leads topic coherence on two of three datasets. A novel inter-topic Maximal Marginal Relevance step raises vocabulary diversity at little cost in coherence. Our code is available at <https://github.com/thcastilho/maretopic>.

1 Introduction

Topic modeling aims to discover the latent thematic structure of a document collection: given N documents, the goal is to identify K topics—each represented by a distribution over words—and to estimate how strongly each document relates to each topic. These representations support tasks such as classification, novelty detection, corpus summarization, and similarity and relevance judgments (Blei et al., 2003).

Two paradigms dominate the field. *Generative* models represent each topic as an explicit

distribution over the vocabulary, learned through probabilistic inference or optimal transport (Blei et al., 2003; Dieng et al., 2020; Wu et al., 2024a). *Clustering-based* models group documents in an embedding space and extract topic words from each cluster (Grootendorst, 2022). Despite their strengths, both paradigms leave two gaps.

The first gap is the absence of a document-level anchor for each topic. Human inspection remains central to topic model evaluation because automated coherence scores can disagree with human judgments (Hoyle et al., 2021); yet neural models represent topics as latent vectors or distributions over the vocabulary, leaving the question “*which document most represents this topic?*” without an architectural answer. Recent approaches incorporate prototypes or exemplars to improve topic learning and interpretability: ProtoXTM averages learned document representations within topic clusters as contrastive prototypes (Seo and Kwon, 2025), whereas E-LDA constructs candidate topic distributions from exemplar keywords (Breuer, 2025). Neither makes an actual corpus document the semantic definition of a topic.

The second gap concerns the clustering-based paradigm. BERTopic (Grootendorst, 2022) partially addresses interpretability by exposing *representative documents* per topic,¹ but these are selected *post-hoc* rather than being the architectural definition of the topic. Moreover, HDBSCAN (Campello et al., 2013) does not take the number of topics as an input: the cluster count instead emerges from the density structure of the data, making it difficult to produce exactly K topics without dataset-specific tuning or post-hoc merging.

Contextual analysis of neighborhoods offers a natural alternative to relying on individual pairwise scores. In high-dimensional embedding spaces, hubness causes some points to occur disproportion-

¹<https://maartengr.github.io/BERTopic/index.html>




Figure 1: Spatial illustration of a ranked list for a query document d_i in an embedding space. Rank-based methods operate on this relative ordering rather than on absolute distances, providing robustness to the geometric distortions of high-dimensional spaces.

ately often in nearest-neighbor lists (Radovanović et al., 2010), while contextualized language-model representations are strongly anisotropic (Ethayarajh, 2019). Because they retain only ordinal comparisons, ranked lists are unchanged by strictly monotone transformations of the underlying distance function, a property central to ordinal embedding (Terada and Luxburg, 2014). Related rank-based similarity measures have also improved over cosine similarity in specific NLP evaluations (Zhelezniak et al., 2019; Santus et al., 2018). As illustrated in Figure 1, ranked lists provide a strong similarity representation structure, capable of summarizing the region surrounding a query in the embedding space. A ranked list describes an element through its context in the collection, which is what makes the representation manifold-aware: ranking is one of the families of unsupervised manifold learning (Pereira-Ferrero et al., 2024), and spectral formulations make its link to the underlying manifold explicit (Iscen et al., 2018). Indeed, neighborhood information has been exploited across a range of challenging representation tasks, including the evaluation of cross-modal alignment in text–visual models (Huh et al., 2024). In Representation Learning, rank-based methods such as RaDE (de Fernando et al., 2022) and GRaCE (Almeida et al., 2025) exploit this principle to greedily select representative data points whose neighborhoods collectively cover the data manifold.

In this paper, we propose **MARETopic** (Manifold-Aware and Rank-based Exemplars Topic Modeling), a training-free topic model that

operates entirely on ranked lists. Documents are encoded with Sentence-BERT (Reimers and Gurevych, 2019), projected onto a low-dimensional manifold via UMAP (McInnes et al., 2020), and indexed into ranked lists. A greedy algorithm then selects K exemplar documents by maximizing a scoring function while penalizing redundancy with previously selected leaders. Two complementary variants are unified under a single framework: **MARETopic_{Corr}**, which scores candidates via a query performance predictor (QPP) and penalizes redundancy with a rank correlation measure, and **MARETopic_{Diff}**, which scores via a rank-based diffusion matrix. Document–topic affinities are then converted to topic keywords. By construction, each topic is anchored on a real document in the corpus.

The main contributions of this work are:

- *Topic Modeling via Ranked Lists.* We introduce **MARETopic**, the first topic model built entirely on rank-based similarity information. By exploiting manifold-aware similarity over pre-trained embeddings, it discards both absolute distances and the geometric pathologies of high-dimensional spaces. Its greedy prototype selection requires no gradient updates or neural training, yielding exactly K topics.
- *Exemplar-Anchored Topic Representation.* **MARETopic**’s topics are anchored in real corpus documents, directly inspectable without reference to any latent-space representation.
- *Inter-Topic Vocabulary Diversification.* We introduce an inter-topic Maximal Marginal Relevance (MMR) step that penalizes vocabulary overlap among clusters, effectively increasing Topic Diversity at a minor cost in coherence.

We evaluate **MARETopic** on three benchmarks (20 Newsgroups, NYT, and Web of Science) against seven state-of-the-art baselines under five metrics. The two variants rank first in seven of the fifteen dataset–metric cells, and **MARETopic_{Corr}** beats every neural topic model in Purity and NMI on all three corpora without any training.

The remainder of this paper is organized as follows. Section 2 surveys related work on topic modeling and rank-based methods. Section 3 presents the **MARETopic** framework. Section 4 describes the experimental evaluation and presents the results, and Section 5 concludes the work.

2 Related Work

This section reviews two areas of literature foundational to our approach: advances in topic modeling architectures (Section 2.1), and rank-based similarity and manifold learning (Section 2.2).

2.1 Topic Modeling

Latent Dirichlet Allocation (LDA; [Blei et al., 2003](#)) and Non-negative Matrix Factorization (NMF; [Lee and Seung, 1999](#)) laid the foundations of the field, both learning an explicit topic–word distribution—respectively by Bayesian inference and by matrix factorization of the document–term matrix.

The **generative paradigm**—models that learn an explicit topic–word distribution—has evolved through several stages of neural integration. ETM ([Dieng et al., 2020](#)) embeds words and topics in the same continuous space; CombinedTM ([Bianchi et al., 2021](#)) shows that pre-trained contextual document embeddings substantially improve topic coherence; ECRTM ([Wu et al., 2023](#)) regularizes word–topic embeddings via clustering; and FASTopic ([Wu et al., 2024a](#)) replaces variational inference with optimal transport, achieving state-of-the-art coherence and diversity. More recently, [Nguyen et al. \(2024\)](#) introduced setwise contrastive learning to improve latent topic representations. Across these developments, however, topics remain latent vectors or distributions over the vocabulary.

The **clustering-based paradigm**, exemplified by BERTopic ([Grootendorst, 2022](#)) and Top2Vec ([Angelov, 2020](#)), treats topic discovery as document clustering in an embedding space. BERTopic stores *representative documents* per topic, but these are selected *post-hoc* by ranking a random sample of cluster members against the topic’s c-TF-IDF vector; they are interpretability artifacts, not the definition of the topic itself. Moreover, HDBSCAN provides no direct control over the number of topics, and empirical studies report overproduction of topics and difficulty obtaining target counts ([Egger and Yu, 2022](#)). LLM-guided clustering has been proposed as an alternative ([Liu et al., 2025](#)), and CAST ([Ma et al., 2025](#)) improves topic word selection via corpus-aware self-similarity.

An emerging direction explores **prototype-inspired mechanisms** for topic modeling. ProtoXTM ([Seo and Kwon, 2025](#)) uses document-level prototypical contrastive learning for cross-lingual alignment, while E-LDA ([Breuer, 2025](#)) formulates LDA topic assignment as submodular optimization

over topic–document links, drawing on exemplar clustering and facility location. However, neither method uses an actual corpus document as the architectural representation of a topic.

2.2 Ranking and Manifold Learning

Standard nearest-neighbor operations rely on absolute distances in the embedding space, yet high-dimensional spaces exhibit geometric pathologies that degrade distance-based methods. *Hubness*—whereby a few points appear disproportionately often in others’ k -NN lists—is an inherent property of high-dimensional geometry ([Radovanović et al., 2010](#)), confirmed in SentenceBERT spaces ([Nielsen and Hansen, 2024](#)). Concurrently, *anisotropy*—the concentration of Transformer embeddings in a narrow cone—causes even unrelated representations to exhibit high cosine similarity ([Ethayarajh, 2019](#)), with subsequent work identifying isolated clusters and low-dimensional manifolds within these globally anisotropic spaces ([Cai et al., 2021](#)). Ordinal embedding theory provides a principled response: [Terada and Luxburg \(2014\)](#) prove that, in the large-sample limit, local ordinal comparisons suffice to recover the underlying geometry up to similarity transformations, and rank-based similarity measures have been shown to outperform cosine similarity on semantic textual similarity ([Zhelezniak et al., 2019](#)) and outlier detection ([Santus et al., 2018](#)). In semi-supervised NLP, mutual k -NN graphs—which require reciprocal neighborhood agreement—consistently outperform standard k -NN graphs in the evaluated document-classification task ([Ozaki et al., 2011](#)), supporting the robustness of reciprocal neighborhood structures.

Ranking here is more than a robustness device. [Pereira-Ferrero et al. \(2024\)](#) surveys the field into diffusion-based, ranking-based and deep methods: a ranked list relates a query to the whole collection, so it carries neighborhood geometry that a single distance does not, and spectral formulations make the link to the underlying manifold explicit ([Iscen et al., 2018](#)). Working on ranked lists is therefore already working on the manifold. In Representation Learning, these principles have been operationalized for representative samples selection. RaDE ([de Fernando et al., 2022](#)) constructs a log-weighted affinity matrix from ranked lists and performs diffusion to select leaders whose neighborhoods collectively cover the data manifold. GRaCE ([Almeida et al., 2025](#)) refines this

by replacing diffusion with a query performance prediction (QPP) estimator and rank correlation measures that penalize redundant selections. Both methods produce interpretable, exemplar-anchored embeddings in which each dimension corresponds to a real data point.

Manifold learning provides a complementary perspective. Graph Laplacian regularization, incorporated into non-negative matrix factorization (Cai et al., 2011), was later adopted in topic models to enforce smoothness of topic posteriors over document neighborhoods (Li et al., 2019). Graph-based models such as GTM (Zhou et al., 2020) and GNTM (Shen et al., 2021) apply neural message passing over document or word graphs. In both cases, however, graph structure is used to inform topic distributions, not to select exemplar documents. Leticio et al. (2024) show that neighbor-embedding projections improve image-retrieval tasks, and BERTopic (Grootendorst, 2022) uses UMAP as its default dimensionality-reduction step.

To the best of our knowledge, rank-based prototype selection has not been applied to topic modeling. MARETopic addresses this gap by adopting a greedy leader-selection strategy that operates on ranked lists, which represent the input space through ordinal neighborhood relations rather than raw Euclidean distances; projection sharpens those neighborhoods, following prior work on retrieval.

3 The MARETopic Framework

Figure 2 presents an overview of the proposed MARETopic (Manifold-Aware and Rank-based Exemplars Topic Modeling) framework. Two variants perform greedy exemplars selection: MARETopic_{Corr} relies on QPP and rank correlation measures, while MARETopic_{Diff} uses a diffusion matrix. Both produce a document–topic matrix; topic words are extracted via c-TF-IDF with inter-topic MMR to yield K leader documents and their representative word lists.

3.1 Problem Formulation

Let $\mathcal{D} = \{d_1, d_2, \dots, d_N\}$ be a corpus of N documents. Topic modeling seeks to identify K latent topics, each represented by (i) a set of characteristic words, and (ii) a document–topic affinity matrix $\Theta \in \mathbb{R}^{N \times K}$ whose entry θ_{ij} quantifies how strongly document d_i pertains to topic j .

MARETopic frames this as a *rank-based prototype selection* problem: it identifies K leader

documents $\mathcal{L} = \{\ell_1, \dots, \ell_K\} \subset \mathcal{D}$ whose neighborhoods collectively cover the corpus minimizing redundancy. Each leader anchors one topic as a real, inspectable exemplar whose neighborhood defines the topic boundary, without probabilistic modelling or neural training. The greedy selection criterion generalizes the framework of GRaCE (Almeida et al., 2025) and RaDE (de Fernando et al., 2022) to the topic modeling setting; our contributions lie in the integration with manifold-aware text representations, the out-of-sample inference mechanism, and the inter-topic MMR diversification.

3.2 Encoding and Indexing

Documents in $\mathcal{D}$ are encoded with Sentence-BERT (Reimers and Gurevych, 2019), producing a d_{SBERT} -dimensional feature vector $\mathbf{x}_i \in \mathbb{R}^{d_{\text{SBERT}}}$ per document. Before constructing ranked lists, UMAP (McInnes et al., 2020) projects each $\mathbf{x}_i$ from $\mathbb{R}^{d_{\text{SBERT}}}$ to $\mathbb{R}^{d_{\text{UMAP}}}$. The vectors are L2-normalized, and, for each document d_i , a traditional nearest-neighbor index (Omohundro, 1989) is queried for the k nearest neighbors (self excluded), yielding a ranked list τ_i ordering the corpus by increasing distance:

$$\tau_i(d_j) < \tau_i(d_m) \iff \rho(d_i, d_j) \leq \rho(d_i, d_m), \quad (1)$$

where $\rho(\cdot, \cdot)$ denotes Euclidean distance and $\tau_i(d_j)$ is the rank of d_j in the list of d_i . The top- k neighborhood of d_i is then:

$$\mathcal{N}(d_i, k) = \{d_j \in \mathcal{D} \mid \tau_i(d_j) \leq k\}. \quad (2)$$

3.3 Rank-Based Topic Discovery

Both variants select leaders via the same greedy criterion. Let $\mathcal{L}_{t-1}$ be the leaders chosen up to step $t-1$. The next leader is:

$$\ell_t = \arg \max_{d \notin \mathcal{L}_{t-1}} \frac{\eta(d)}{1 + \sum_{\ell \in \mathcal{L}_{t-1}} s(\tau_d, \tau_\ell)}, \quad (3)$$

where $\eta(d)$ measures the topological centrality of candidate d , and $s(\tau_d, \tau_\ell)$ penalizes candidates whose neighborhoods overlap with already-selected leaders to encourage topical diversity. The two variants differ only in how η and s are defined.

MARETopic_{Corr}

Influence estimation. Each candidate is scored with a Query Performance Prediction (QPP) measure. In this work we use **Reciprocal Density** (Pedronette and Torres, 2015), which measures how




Figure 2: Overview of the MARETopic framework.

many neighbors of d are also reciprocal neighbors of each other within $\mathcal{N}(d, k)$:

$$\eta(d) = \frac{1}{k^4} \sum_{\substack{d_j \in \mathcal{N}(d, k) \\ d_i \in \mathcal{N}(d, k)}} f_r(d_j, d_i) w_r(d, d_j) w_r(d, d_i), \quad (4)$$

where $f_r(d_j, d_i) = |\mathcal{N}(d_j, k) \cap \{d_i\}|$ checks whether (d_j, d_i) are reciprocal neighbors, and $w_r(d_j, d_i) = k + 1 - \tau_j(d_i)$ is a rank-decay weight. High values indicate documents whose neighbors agree on neighborhood membership — a sign of locally coherent manifold structure.

Diversity penalization. Neighborhood overlap between candidate and selected leaders is measured by a rank correlation measure. Specifically, this work uses **JaccardMax** (Valem et al., 2022), which measures the overlap between two rankings by computing the maximum Jaccard coefficient across all prefix depths $l < k$:

$$s(\tau_d, \tau_\ell) = \max_{l=1}^{k-1} \frac{|\mathcal{N}(d, l) \cap \mathcal{N}(\ell, l)|}{|\mathcal{N}(d, l) \cup \mathcal{N}(\ell, l)|}. \quad (5)$$

When JaccardMax between d and a selected ℓ is high, their neighborhoods are largely redundant, and the denominator of Eq. 3 suppresses d ’s score.

Any QPP estimator or rank correlation measure can replace Reciprocal Density or JaccardMax, respectively, without any changes to the other steps.

MARETopic_{Diff}

Affinity matrix construction. From the ranked lists $\{\tau_i\}$, a sparse affinity matrix $\mathbf{W} \in \mathbb{R}^{N \times N}$ is built, where each entry encodes the rank position of d_j in the list of d_i :

$$w_{ij} = \begin{cases} 1 - \log_k(\tau_i(d_j)), & \text{if } \tau_i(d_j) \leq k, \\ 0, & \text{otherwise.} \end{cases} \quad (6)$$

Diffusion scoring. A p -step diffusion $\mathbf{A} = \mathbf{W}^p$ (with $p=2$, following de Fernando et al. (2022)) propagates local affinities across the neighborhood graph and smooths out noise in individual rank positions. The diagonal entry a_{dd} captures the *reciprocal affinity* of d after global propagation, playing an analogous role to the QPP measure in MARETopic_{Corr}. In Eq. 3, $\eta(d) = a_{dd}$ and $s(\tau_d, \tau_\ell) = a_{d\ell}$, penalizing candidates that share strong diffusion paths with already-selected leaders.

3.4 Document-Topic Representation and Inference

After selecting $\mathcal{L} = \{\ell_1, \dots, \ell_K\}$, the document-topic matrix Θ is computed as:

$$\theta_{ij} = s(\tau_i, \tau_{\ell_j}), \quad z_i = \arg \max_j \theta_{ij}, \quad (7)$$

where s is the scoring function used for selection (JaccardMax for MARETopic_{Corr}; column ℓ_j of $\mathbf{A}$ for MARETopic_{Diff}). Hard assignment z_i is used for topic-word extraction and clustering evaluation.

Out-of-sample inference reuses the fitted pipeline with no retraining: the UMAP reducer projects a new embedding into the training manifold, and the index returns its ranked list τ^* against the training corpus. For MARETopic_{Corr}, s (JaccardMax) is computed directly between τ^* and each leader’s list. For MARETopic_{Diff}, the diffusion matrix $\mathbf{A} = \mathbf{W}^2$ is defined only over the training graph, so a new document is instead scored by its single-step affinity $w = 1 - \log_k(\tau^*(\ell_j))$ to each leader ℓ_j — a one-hop approximation of the diffusion column that requires no re-diffusion.

3.5 Topic Word Extraction

Given hard assignments $\{z_i\}$, topic words are extracted in two steps. First, c-TF-IDF (Grootendorst,

2022) scores each word w for topic c :

$$\text{c-TF-IDF}(c, w) = \text{tf}(c, w) \cdot \log\left(1 + \frac{\bar{n}}{\text{tf}_{\mathcal{D}}(w)}\right), \quad (8)$$

where $\text{tf}(c, w)$ is the frequency of w in the documents assigned to topic c , $\text{tf}_{\mathcal{D}}(w)$ is the total frequency of w across all topics, and $\bar{n}$ is the mean number of tokens per topic.

Second, an inter-topic MMR (Maximal Marginal Relevance) step diversifies the final word lists. Topics are processed in decreasing order of cluster size. For each topic c , each word w is re-scored as:

$$\text{score}(c, w) = (1 - \lambda) \cdot \text{c-TF-IDF}(c, w) - \lambda \cdot \max_{c' \prec c} \text{c-TF-IDF}(c', w), \quad (9)$$

where λ controls the diversity weight, $|c|$ denotes the cluster size, and $c' \prec c$ ranges over the topics already processed — those with $|c'| \geq |c|$ — so the penalty is the running maximum c-TF-IDF of w over the larger topics, suppressing high-frequency words that appear in them and producing per-topic vocabularies with less lexical overlap.

Placing this step after Θ has two consequences for the evaluation. It re-ranks word lists once the assignment $z_i = \arg \max_j \theta_{ij}$ is fixed, so it cannot change that assignment: Purity and NMI do not depend on λ at all, and only the word-level metrics do. It also never looks at the model that produced Θ , so any clustering-based topic model can use it, as we verify in Appendix A.

4 Experimental Evaluation

This section evaluates MARETopic experimentally. Section 4.1 describes the experimental setup, including datasets, baselines, evaluation metrics, and implementation details. Section 4.2 evaluates MARETopic against seven baselines across three standard benchmarks. Section 4.3 discuss the computational cost of the pipeline. Section 4.4 investigates the sensitivity of the method to key design choices: UMAP projection, k size, and MMR’s λ .

4.1 Experimental Setup

Datasets. **20 Newsgroups (20NG)** (Lang, 1995) and **New York Times (NYT)**² are obtained directly from the TopMost toolkit (Wu et al., 2024b): 20NG contains 18,846 documents across 20 topically distinct categories, and NYT comprises 9,172 articles from 12 news sections. **Web of Science**

(WoS) (Kowsari et al., 2017) consists of 11,967 scientific abstracts organized into 7 disciplines, with a stratified 80/20 train/test split. All corpora are pre-processed by TopMost (lowercased, stopwords removed, minimum token length of three characters) and represented as bag-of-words with vocabulary sizes of 5,000 (20NG) and 10,000 (NYT and WoS), following the protocol of Wu et al. (2024a).

Baselines. We compare against seven topic models spanning both paradigms: **LDA** (Blei et al., 2003), **NMF** (Lee and Seung, 1999), **ETM** (Dieng et al., 2020), **ECRTM** (Wu et al., 2023), and **FASTopic** (Wu et al., 2024a) from the generative paradigm, and **BERTopic** (Grootendorst, 2022) and **Top2Vec** (Angelov, 2020) from the clustering-based paradigm. All methods (except Top2Vec) are evaluated through the TopMost framework (Wu et al., 2024b) default hyperparameters (200 training epochs for neural models).

Evaluation Metrics. Five standard metrics are reported. Topic coherence is measured under the two formulations. C_v combines normalized pointwise mutual information and cosine similarity over word co-occurrence windows (Röder et al., 2015); its agreement with human judgment has been questioned (Hoyle et al., 2021). **NPMI** normalizes pointwise mutual information to $[-1, 1]$ (Bouma, 2009), with a value of zero corresponding to statistical independence for a word pair. Both are computed with `gensim`³ using the training corpus as the reference collection. **TD** (Topic Diversity) measures the fraction of unique words across all topic top-word lists (Dieng et al., 2020). **Purity** and **NMI** assess clustering quality by comparing the hard assignment z_i against ground-truth document labels.

Implementation Details. All experiments use $K = 50$ topics and three independent runs, following (Grootendorst, 2022), reporting means and additionally standard deviations ($\pm$). All methods are evaluated under the same number of topics to ensure a controlled comparison; Appendix B repeats the whole evaluation at $K = 100$. The default configuration for both variants sets the neighborhood size to $k = 100$, the UMAP target dimensionality to $d_{\text{UMAP}} = 5$ (with $n_{\text{neighbors}} = 15$, $\text{min_dist} = 0.0$, cosine metric, mirroring BERTopic’s defaults), and the MMR diversity parameter to $\lambda = 0.3$. Each topic is represented by $m = 15$ top words

²<https://github.com/bobxwu/TopMost/blob/main/data/NYT.zip>

³<https://pypi.org/project/gensim/>

Table 1: Topic modeling evaluation across three datasets (mean $\pm$ std of 3 runs, $K = 50$). Columns are grouped by paradigm. Best value per metric per dataset in **bold**; runner-up underlined. Time is end-to-end wall clock in seconds on an Apple M4 with 16 GB of memory, averaged over three runs, and lower is better.

Metric	Generative					Clustering-based			
	LDA	NMF	ETM	ECRTM	FASTopic	BERTopic	Top2Vec	Ours	
	(Blei et al., 2003)	(Lee and Seung, 1999)	(Dieng et al., 2020)	(Wu et al., 2023)	(Wu et al., 2024a)	(Grootendorst, 2022)	(Angelov, 2020)	MARETopic Diffusion	MARETopic Correlation
20 Newsgroups									
C_v	0.5352 $\pm$.007	0.5573 $\pm$.016	0.4522 $\pm$.004	0.6116 $\pm$.013	0.5993 $\pm$.019	0.5942 $\pm$.007	0.6649 $\pm$.011	0.6708 $\pm$.015	0.6818 $\pm$.003
NPMI	0.0472 $\pm$.003	0.0613 $\pm$.007	-0.0171 $\pm$.003	-0.0372 $\pm$.021	-0.0689 $\pm$.002	0.0842 $\pm$.003	0.1032 $\pm$.003	0.0958 $\pm$.005	<u>0.1027</u> $\pm$.002
TD	0.5196 $\pm$.008	0.5093 $\pm$.006	0.8471 $\pm$.005	0.8689 $\pm$.017	0.9849 $\pm$.008	0.7733 $\pm$.002	0.8662 $\pm$.003	0.8613 $\pm$.013	<u>0.8787</u> $\pm$.001
Purity	0.4204 $\pm$.022	0.2153 $\pm$.014	0.3390 $\pm$.019	<u>0.6114</u> $\pm$.003	0.5738 $\pm$.004	0.4493 $\pm$.009	0.5812 $\pm$.003	0.4144 $\pm$.022	0.6227 $\pm$.005
NMI	0.3763 $\pm$.009	0.1800 $\pm$.015	0.3032 $\pm$.018	<u>0.5348</u> $\pm$.002	0.5202 $\pm$.005	0.3797 $\pm$.005	0.5265 $\pm$.004	0.4011 $\pm$.016	0.5419 $\pm$.003
Time (s)	<u>8.6</u>	5.5	114.3	184.5	128.7	37.1	52.6	49.6	92.6
New York Times									
C_v	0.4534 $\pm$.011	0.4919 $\pm$.010	0.4696 $\pm$.003	0.4457 $\pm$.023	0.5864 $\pm$.017	0.6189 $\pm$.009	0.6944 $\pm$.013	0.6570 $\pm$.012	0.6496 $\pm$.018
NPMI	0.0002 $\pm$.005	0.0468 $\pm$.005	0.0109 $\pm$.001	-0.0997 $\pm$.015	-0.0587 $\pm$.005	0.1166 $\pm$.002	0.1295 $\pm$.006	0.1079 $\pm$.006	0.1060 $\pm$.004
TD	0.5769 $\pm$.022	0.4560 $\pm$.011	0.8427 $\pm$.009	<u>0.9595</u> $\pm$.008	0.9996 $\pm$.001	0.7951 $\pm$.004	0.9071 $\pm$.008	0.8831 $\pm$.006	0.8702 $\pm$.002
Purity	0.5454 $\pm$.024	0.4612 $\pm$.003	0.5556 $\pm$.011	0.6529 $\pm$.002	0.6583 $\pm$.018	0.6184 $\pm$.006	<u>0.6874</u> $\pm$.006	0.6046 $\pm$.016	0.7084 $\pm$.003
NMI	0.2661 $\pm$.020	0.1952 $\pm$.007	0.2815 $\pm$.005	0.3549 $\pm$.004	0.3670 $\pm$.012	0.3221 $\pm$.006	<u>0.3890</u> $\pm$.002	0.3305 $\pm$.011	0.3975 $\pm$.002
Time (s)	<u>7.2</u>	4.5	148.2	223.7	123.9	36.9	44.2	31.6	53.6
Web of Science									
C_v	0.4797 $\pm$.006	0.4869 $\pm$.014	0.4198 $\pm$.006	0.4655 $\pm$.004	0.6000 $\pm$.005	0.6920 $\pm$.005	0.7018 $\pm$.011	0.7198 $\pm$.004	<u>0.7150</u> $\pm$.005
NPMI	-0.0287 $\pm$.011	0.0396 $\pm$.006	-0.0045 $\pm$.005	-0.1076 $\pm$.001	-0.0308 $\pm$.009	0.1244 $\pm$.007	0.1344 $\pm$.003	<u>0.1359</u> $\pm$.003	0.1447 $\pm$.002
TD	0.6636 $\pm$.017	0.4876 $\pm$.007	0.8938 $\pm$.010	<u>0.9996</u> $\pm$.001	1.0000 $\pm$.000	0.7782 $\pm$.010	0.8796 $\pm$.012	0.8569 $\pm$.013	0.8547 $\pm$.001
Purity	0.6529 $\pm$.009	0.6086 $\pm$.012	0.6456 $\pm$.008	0.7733 $\pm$.010	0.7765 $\pm$.012	0.7934 $\pm$.004	0.8293 $\pm$.009	0.6082 $\pm$.021	<u>0.8106</u> $\pm$.003
NMI	0.3875 $\pm$.005	0.2962 $\pm$.010	0.3853 $\pm$.011	0.4518 $\pm$.008	0.4565 $\pm$.005	0.4473 $\pm$.003	0.4901 $\pm$.003	0.3660 $\pm$.008	<u>0.4664</u> $\pm$.002
Time (s)	<u>6.5</u>	5.2	170.7	253.7	137.2	35.0	49.6	37.7	65.2

extracted via c-TF-IDF followed by inter-topic MMR diversification. All methods that operate on document embeddings — FASTopic, BERTopic, MARETopic_{Corr}, and MARETopic_{Diff} — use the all-MiniLM-L6-v2 SBERT encoder, ensuring that differences in topic quality reflect the modeling mechanism rather than the encoder. Robustness to this choice is verified in Appendix C, where three SBERT backbones yield consistent results.

Code Availability. The implementation of MARETopic with full integration with the Top-Most framework (Wu et al., 2024b), is available at <https://github.com/thcastilho/maretopic>.

4.2 Main Results

Table 1 presents the results across all three datasets and five metrics. Qualitative results are provided in Appendices D and E.

Clustering quality. MARETopic_{Corr}, which requires no gradient updates, achieves the highest Purity and NMI on 20NG and NYT — the two corpora with the most categories — surpassing neural models trained for 200 epochs. The largest margin is on NYT, where it gains +5.0 pp in Purity over FASTopic and +2.1 pp over Top2Vec, the strongest clustering-based baseline. On WoS the ordering reverses and Top2Vec leads both metrics, a boundary we return to in Section 5. MARETopic_{Diff} trails

the correlation variant in clustering throughout; on 20NG its Purity falls below several baselines, reflecting that the diffusion scoring separates fine-grained categories less sharply. Its strengths lie instead in coherence.

Topic coherence. Under C_v the two variants are most complementary: MARETopic_{Corr} achieves the highest value on 20NG (0.6818, ahead of the nearest baseline Top2Vec by 1.7 pp) and MARETopic_{Diff} the highest on WoS (0.7198, ahead of Top2Vec by 1.8 pp), with the other variant immediately behind in both cases; on NYT, Top2Vec leads at 0.6944, 3.7 pp above MARETopic_{Diff}.

NPMI orders the methods differently, dividing them by paradigm. Every generative neural model scores non-positive NPMI on at least two corpora, FASTopic and ECRTM on all three: the words they group into a topic co-occur no more often than chance would predict, which is exactly what a globally fitted distribution risks. Every clustering-based method stays well above zero. Inside that group the margins are small: MARETopic_{Corr} leads on WoS (0.1447), sits within 0.0005 of Top2Vec on 20NG, inside the run-to-run deviation, and falls behind clustering-based baselines on NYT.

Topic diversity. Diversity follows paradigm lines: FASTopic’s Sinkhorn regularization enforces near-perfect TD by construction, a structural ad-

vantage that no clustering-based method currently matches. Within the clustering-based paradigm, however, the proposed inter-topic MMR step effectively mitigates the vocabulary overlap inherent to these approaches: with the default penalty, both variants cleanly exceed BERTopic in TD (0.85–0.88 vs. 0.77–0.80). Top2Vec occupies the same band and edges the two variants on NYT and WoS (0.9071 and 0.8796).

4.3 Computational Cost

MARETopic trains nothing, so it pays no epoch cost, and the Time row of Table 1 reflects that. The figures cover encoding, projection, indexing, topic discovery, out-of-sample inference and word extraction. MARETopic_{Diff} sits in the same band as the other embedding-based methods, while MARETopic_{Corr} costs $1.7\text{--}1.9\times$ the diffusion variant, since accumulating a JaccardMax affinity over the neighborhood is more expensive than reading a diffusion diagonal. Both variants undercut every neural topic model on every corpus, where the cost goes to 200 training epochs.

Indexing, the step a rank-based method adds to the pipeline, is not where the cost lies. Building the exact nearest-neighbor index takes 0.18–0.41 s across the three corpora, under 1% of end-to-end runtime in every configuration; the time is dominated by document encoding and by leader selection instead. Selection uses only rank positions, never the distances behind them, so the exact index can be swapped for an approximate one at sublinear query cost without changing the method.

4.4 Ablation Studies

Sections 4.4.1–4.4.3 analyze sensitivity to three internal design choices of MARETopic: the UMAP pre-processing step, the neighborhood size k , and the MMR diversity parameter λ . All ablations follow the same protocol from Section 4.1.⁴ Each study reports C_v , TD, Purity and NMI, except where a design choice cannot affect one of them.

4.4.1 Effect of UMAP Pre-processing

Table 2 isolates the contribution of UMAP by reporting each metric alongside its delta from raw SBERT embeddings, and the two variants respond differently. For MARETopic_{Corr}, the projection

Table 2: Effect of UMAP pre-processing. Each cell shows the UMAP-enabled value with the delta from raw SBERT (green: improvement; red: degradation).

	20NG	NYT	WoS
MARETopic_{Corr}			
C_v	0.6817 −0.026	0.6449 −0.053	0.7189 −0.028
TD	0.8854 +0.025	0.8764 +0.016	0.8542 −0.014
Purity	0.6161 +0.017	0.7102 +0.061	0.8094 −0.002
NMI	0.5398 +0.031	0.3982 +0.041	0.4697 +0.002
MARETopic_{Diff}			
C_v	0.6750 +0.039	0.6458 −0.044	0.7177 −0.009
TD	0.8787 −0.019	0.8831 +0.011	0.8756 −0.002
Purity	0.4103 −0.004	0.5908 −0.027	0.6089 −0.097
NMI	0.3984 +0.020	0.3241 −0.002	0.3623 −0.044

consistently improves clustering on 20NG and NYT at the cost of lower C_v , because the compression partly disrupts vocabulary co-occurrence patterns. For MARETopic_{Diff} it improves C_v and NMI on 20NG but degrades Purity on NYT and substantially on WoS, where the seven broad categories are already well-separated in raw space, and the projection introduces unnecessary distortion.

This asymmetry reflects a fundamental difference between the two scoring mechanisms. The QPP and rank correlation measures in MARETopic_{Corr} rely on fine-grained neighborhood structure and benefit from the tighter local neighborhoods that UMAP provides. The diffusion operator in MARETopic_{Diff}, by contrast, already smooths local noise through matrix multiplication and gains less from the projection. We retain $d_{\text{UMAP}} = 5$ as the default because it yields consistent clustering gains for MARETopic_{Corr} on 20NG and NYT; for MARETopic_{Diff} or coarser-grained corpora, disabling UMAP is a reasonable choice.

4.4.2 Neighborhood Size (k)

Figure 3 shows the effect of k on all four metrics. Purity and NMI rise steeply from $k = 20$ to $k = 100$ for both variants. MARETopic_{Corr} levels off near $k = 100$, while MARETopic_{Diff} keeps gaining beyond it, because the diffusion operator benefits from denser neighborhoods — on 20NG and WoS Purity its $100 \rightarrow 200$ gain still exceeds the previous step. The word-level metrics are far less sensitive to k : for $k \geq 50$, both C_v (peaking around $k = 50\text{--}100$) and TD (declining gently) vary by at most 0.06, the only pronounced movement being the rise in C_v from $k = 20$ to $k = 50$. We set $k = 100$ as the default: it is the elbow of the clustering curves for MARETopic_{Corr}, and the cost of that variant

⁴Ablation experiments are conducted as independent evaluation runs from the main comparison in Table 1; minor numerical deviations are within the observed standard deviation and reflect UMAP stochasticity.




Figure 3: Sensitivity to neighborhood size k across all four metrics. Solid lines: $\text{MARETopic}_{\text{Corr}}$; dashed lines: $\text{MARETopic}_{\text{Diff}}$. Vertical dotted line marks the recommended $k = 100$.

Table 3: Sensitivity to the MMR diversity parameter λ . Only C_v and TD are reported, as MMR acts on word extraction and does not affect cluster assignments. Colored deltas show change relative to the previous λ (green: improvement; red: degradation). Recommended $\lambda = 0.3$ highlighted (gray).

Method	λ	20 Newsgroups		New York Times		Web of Science	
		C_v	TD	C_v	TD	C_v	TD
MARETopic Correlation	0.0	0.6777	0.6627	0.6398	0.6845	0.7299	0.6938
	0.3	0.6674 -0.010	0.8813 $+0.219$	0.6384 -0.001	0.8729 $+0.188$	0.7206 -0.009	0.8600 $+0.166$
	0.5	0.6539 -0.014	0.9538 $+0.073$	0.6349 -0.004	0.9582 $+0.085$	0.6850 -0.036	0.9533 $+0.093$
	0.7	0.6257 -0.028	0.9925 $+0.039$	0.6043 -0.031	0.9893 $+0.031$	0.6689 -0.016	0.9902 $+0.037$
MARETopic Diffusion	0.0	0.6854	0.6502	0.6516	0.7080	0.7287	0.6906
	0.3	0.6676 -0.018	0.8800 $+0.230$	0.6677 $+0.016$	0.8747 $+0.167$	0.7080 -0.021	0.8569 $+0.166$
	0.5	0.6487 -0.019	0.9564 $+0.076$	0.6506 -0.017	0.9569 $+0.082$	0.6816 -0.026	0.9582 $+0.101$
	0.7	0.6322 -0.017	0.9920 $+0.036$	0.6322 -0.018	0.9893 $+0.032$	0.6577 -0.024	0.9938 $+0.036$

grows with k , since JaccardMax scans every prefix $l < k$.

4.4.3 MMR Diversity Parameter (λ)

MMR λ is varied to map the coherence–diversity trade-off. Table 3 reports C_v and TD for $\lambda \in \{0.0, 0.3, 0.5, 0.7\}$. TD is sensitive to λ : without MMR, TD falls below BERTopic because exemplar-based clusters share much of their vocabulary. At $\lambda = 0.3$, TD rises by roughly $+0.2$ on every dataset for both variants, while C_v drops by at most 0.021 ; we adopt this as the default. At $\lambda = 0.7$, TD nearly matches FASTopic (≈ 0.99) but C_v decreases more noticeably. The step acts on word lists after Θ is fixed, so Purity and NMI do not move with λ and are omitted from the table. Nor is the step specific to MARETopic: Appendix A attaches it to BERTopic’s own c-TF-IDF, where it raises Topic Diversity by a comparable margin.

5 Conclusion

This work introduced MARETopic, a training-free framework that reframes topic modeling as rank-based prototype selection. By operating on ranked-list neighborhoods over a low-dimensional manifold, MARETopic sidesteps the hubness and

anisotropy pathologies that distort absolute distances in high-dimensional embedding spaces, and inherently anchors each topic to a real, inspectable corpus document. Coupled with a novel inter-topic Maximal Marginal Relevance (MMR) extraction step, the framework ranks first in seven of the fifteen dataset–metric cells, and $\text{MARETopic}_{\text{Corr}}$ beats every neural topic model in Purity and NMI on all three corpora without any training. The diversity gap to optimal-transport models such as FASTopic belongs to the clustering-based paradigm rather than to this method, and the MMR step that narrows it works on any competitor. On Web of Science, the most separable of the three corpora, Top2Vec reaches higher Purity and NMI; where the manifold is already simple, a centroid and an exemplar may describe the same region, which puts the advantage of rank-based selection on more irregular collections. Future work will extend this rank-based formulation to hierarchical topic structures and dynamic topic modeling over time, complement the evaluation with downstream and intrusion-based protocols, and assess scalability on larger and multilingual corpora.

Limitations

Our evaluation covers three English-language corpora of moderate size (9K–19K documents) with 7–20 categories. Although these are standard benchmarks in topic modeling (Wu et al., 2024a,b), the effectiveness on multilingual or domain-specific corpora (e.g., legal, biomedical) remains to be verified. Section 4.3 measures the nearest-neighbor index that MARETopic adds and finds its cost negligible, but the exact index still scales with corpus size, and very large collections (100K+ documents) would need an approximate one. We have not evaluated that regime. The runtime comparison is also limited by our hardware: the neural baselines have no GPU path here and run CPU-bound. What we claim is architectural, that MARETopic needs neither training nor an accelerator. We do not claim to be faster in wall-clock terms. Topic Diversity remains below generative models such as FASTopic, a gap inherent to the clustering-based paradigm rather than specific to our method; the inter-topic MMR step that narrows it is method-agnostic (Appendix A). Finally, a downstream evaluation would strengthen the practical impact claim, measuring the impact of the proposal in tasks such as document classification and word/topic intrusion.

Acknowledgments

The authors are grateful to the National Council for Scientific and Technological Development—CNPq (grant #313193/2023-1), the São Paulo Research Foundation—FAPESP (grant #2024/04890-5), and Petrobras (grant #2023/00095-3) for their financial support.

References

- Thiago César Castilho Almeida, Gustavo Rosseto Letício, Lucas Pascotti Valem, André Freitas, and Daniel Carlos Guimarães Pedronette. 2025. [Effective graph and rank-based contextual embeddings for textual and multimedia data](#). In *2025 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8.
- Dimo Angelov. 2020. [Top2vec: Distributed representations of topics](#). *ArXiv*, abs/2008.09470.
- Federico Bianchi, Silvia Terragni, and Dirk Hovy. 2021. [Pre-training is a hot topic: Contextualized document embeddings improve topic coherence](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 759–766. Association for Computational Linguistics.
- David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. [Latent dirichlet allocation](#). *The Journal of Machine Learning Research*, 3:993–1022.
- Gerlof Bouma. 2009. Normalized (pointwise) mutual information in collocation extraction. *Proceedings of the Biennial GSCL Conference 2009*.
- Adam Breuer. 2025. [E-LDA: Toward interpretable LDA topic models with strong guarantees in logarithmic parallel time](#). In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 5531–5557. PMLR.
- Deng Cai, Xiaofei He, Jiawei Han, and Thomas S. Huang. 2011. [Graph regularized nonnegative matrix factorization for data representation](#). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(8):1548–1560.
- Xingyu Cai, Jiaji Huang, Yuchen Bian, and Kenneth Church. 2021. [Isotropy in the contextual embedding space: Clusters and manifolds](#). In *International Conference on Learning Representations*.
- Ricardo J. G. B. Campello, Davoud Moulavi, and Jörg Sander. 2013. [Density-based clustering based on hierarchical density estimates](#). In *Advances in Knowledge Discovery and Data Mining*, pages 160–172. Springer Berlin Heidelberg.
- Filipe Alves de Fernando, Daniel Carlos Guimarães Pedronette, Gustavo José de Sousa, Lucas Pascotti Valem, and Ivan Rizzo Guilherme. 2022. [RaDE+: A semantic rank-based graph embedding algorithm](#). *International Journal of Information Management Data Insights*, 2(1):100078.
- Adji B. Dieng, Francisco J. R. Ruiz, and David M. Blei. 2020. [Topic modeling in embedding spaces](#). *Transactions of the Association for Computational Linguistics*, 8:439–453.
- Roman Egger and Joanne Yu. 2022. [A topic modeling comparison between lda, nmf, top2vec, and bertopic to demystify twitter posts](#). *Frontiers in Sociology*, 7.
- Kawin Ethayarajh. 2019. [How contextual are contextualized word representations? Comparing the geometry of BERT, ELMo, and GPT-2 embeddings](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 55–65. Association for Computational Linguistics.
- Maarten Grootendorst. 2022. [Bertopic: Neural topic modeling with a class-based tf-idf procedure](#). *ArXiv*, abs/2203.05794.
- Alexander Hoyle, Pranav Goel, Andrew Hian-Cheong, Denis Peskov, Jordan Boyd-Graber, and Philip Resnik. 2021. [Is automated topic model evaluation broken? the incoherence of coherence](#). In *Advances in Neural Information Processing Systems*,

- volume 34, pages 2018–2033. Curran Associates, Inc.
- Minyoung Huh, Brian Cheung, Tongzhou Wang, and Phillip Isola. 2024. [Position: The platonic representation hypothesis](#). In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 20617–20642. PMLR.
- Ahmet Iscen, Yannis Avrithis, Giorgos Tolias, Teddy Furon, and Ondrej Chum. 2018. [Fast spectral ranking for similarity search](#). In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7632–7641.
- Kamran Kowsari, Donald E. Brown, Mojtaba Heidarysafa, Kiana Jafari Meimandi, Matthew S. Gerber, and Laura E. Barnes. 2017. [Hdltext: Hierarchical deep learning for text classification](#). In *2017 16th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 364–371.
- Ken Lang. 1995. [Newsweeder: Learning to filter net-news](#). In *Machine Learning Proceedings 1995*, pages 331–339. Morgan Kaufmann.
- Daniel D Lee and H Sebastian Seung. 1999. [Learning the parts of objects by non-negative matrix factorization](#). *Nature*, 401(6755):788–791.
- Gustavo Rosseto Leticio, Vinicius Sato Kawai, Lucas Pascotti Valem, Daniel Carlos Guimarães Pedronette, and Ricardo da S. Torres. 2024. [Manifold information through neighbor embedding projection for image retrieval](#). *Pattern Recognition Letters*, 183:17–25.
- Ximing Li, Jiaojiao Zhang, and Jihong Ouyang. 2019. [Dirichlet multinomial mixture with variational manifold regularization: topic modeling over short texts](#). In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence and Thirty-First Innovative Applications of Artificial Intelligence Conference and Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, AAAI’19/IAAI’19/EAAI’19*. AAAI Press.
- Jiangnan Liu, Ziyu Shang, Wenjun Ke, Peng Wang, Zhizhao Luo, Jiajun Liu, Guozheng Li, and Yining Li. 2025. [LLM-guided semantic-aware clustering for topic modeling](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 18420–18435. Association for Computational Linguistics.
- Yanan Ma, Chenghao Xiao, Chenhan Yuan, Sabine N van der Veer, Lamiece Hassan, Chenghua Lin, and Goran Nenadic. 2025. [CAST: Corpus-aware self-similarity enhanced topic modelling](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 7548–7561. Association for Computational Linguistics.
- Leland McInnes, John Healy, and James Melville. 2020. [Umap: Uniform manifold approximation and projection for dimension reduction](#). *ArXiv*, abs/1802.03426.
- Thong Thanh Nguyen, Xiaobao Wu, Xinshuai Dong, Cong-Duy Nguyen, See-Kiong Ng, and Anh Tuan Luu. 2024. [Topic modeling as multi-objective contrastive optimization](#). In *International Conference on Learning Representations*, volume 2024, pages 55760–55779.
- Beatrix Miranda Ginn Nielsen and Lars Kai Hansen. 2024. [Hubness reduction improves sentence-BERT semantic spaces](#). In *Proceedings of the 5th Northern Lights Deep Learning Conference (NLDL)*, volume 233 of *Proceedings of Machine Learning Research*, pages 181–204. PMLR.
- Stephen M. Omohundro. 1989. [Five balltree construction algorithms](#). Technical Report TR-89-063, International Computer Science Institute.
- Kohei Ozaki, Masashi Shimbo, Mamoru Komachi, and Yuji Matsumoto. 2011. [Using the mutual k-nearest neighbor graphs for semi-supervised classification on natural language data](#). In *Proceedings of the Fifteenth Conference on Computational Natural Language Learning*, pages 154–162. Association for Computational Linguistics.
- Daniel Carlos Guimarães Pedronette and Ricardo da S. Torres. 2015. [Unsupervised effectiveness estimation for image retrieval using reciprocal rank information](#). In *Proceedings of the 2015 28th SIBGRAPI Conference on Graphics, Patterns and Images, SIBGRAPI ’15*, page 321–328. IEEE Computer Society.
- V. H. Pereira-Ferrero, T. G. Lewis, L. P. Valem, L. G. P. Ferrero, D. C. G. Pedronette, and L. J. Latecki. 2024. [Unsupervised affinity learning based on manifold analysis for image retrieval: A survey](#). *Computer Science Review*, 53:100657.
- Miloš Radovanović, Alexandros Nanopoulos, and Mirjana Ivanović. 2010. [Hubs in space: Popular nearest neighbors in high-dimensional data](#). *Journal of Machine Learning Research*, 11(86):2487–2531.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992. Association for Computational Linguistics.
- Michael Röder, Andreas Both, and Alexander Hinneburg. 2015. [Exploring the space of topic coherence measures](#). In *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining, WSDM ’15*, page 399–408. Association for Computing Machinery.

- Enrico Santus, Hongmin Wang, Emmanuele Chersoni, and Yue Zhang. 2018. [A rank-based similarity metric for word embeddings](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 552–557, Melbourne, Australia. Association for Computational Linguistics.
- Seung-Won Seo and Soon-Sun Kwon. 2025. [ProtoXTM: Cross-lingual topic modeling with document-level prototype-based contrastive learning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 20340–20354, Suzhou, China. Association for Computational Linguistics.
- Dazhong Shen, Chuan Qin, Chao Wang, Zheng Dong, Hengshu Zhu, and Hui Xiong. 2021. [Topic modeling revisited: A document graph-based neural network perspective](#). In *Advances in Neural Information Processing Systems*, volume 34, pages 14681–14693. Curran Associates, Inc.
- Yoshikazu Terada and Ulrike Luxburg. 2014. [Local ordinal embedding](#). In *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 847–855. PMLR.
- Lucas Pascotti Valem, Vinicius Atsushi Sato Kawai, Vanessa Helena Pereira-Ferrero, and Daniel Carlos Guimarães Pedronette. 2022. [A novel rank correlation measure for manifold learning on image retrieval and person re-id](#). In *2022 IEEE International Conference on Image Processing (ICIP)*, pages 1371–1375.
- Xiaobao Wu, Xinshuai Dong, Thong Thanh Nguyen, and Anh Tuan Luu. 2023. [Effective neural topic modeling with embedding clustering regularization](#). In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 37335–37357. PMLR.
- Xiaobao Wu, Thong Nguyen, Delvin Ce Zhang, William Yang Wang, and Anh Tuan Luu. 2024a. [Fastopic: Pretrained transformer is a fast, adaptive, stable, and transferable topic model](#). In *Advances in Neural Information Processing Systems*, volume 37, pages 84447–84481. Curran Associates, Inc.
- Xiaobao Wu, Fengjun Pan, and Anh Tuan Luu. 2024b. [Towards the TopMost: A topic modeling system toolkit](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, pages 31–41. Association for Computational Linguistics.
- Vitalii Zhelezniak, Aleksandar Savkov, April Shen, and Nils Hammerla. 2019. [Correlation coefficients and semantic textual similarity](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 951–962. Association for Computational Linguistics.
- Deyu Zhou, Xuemeng Hu, and Rui Wang. 2020. [Neural topic modeling by incorporating document relationship graph](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3790–3796. Association for Computational Linguistics.

A Transferability of the Inter-Topic MMR

This section evaluates whether the inter-topic MMR step is specific to MARETopic, by attaching it to a competing model. The step re-ranks word lists without looking at the model that produced them, so any clustering-based topic model can use it. Table 4 applies it to the c-TF-IDF scores of BERTopic, leaving that method’s clustering, vocabulary and scoring formula untouched, so word selection is the only thing that changes. Topic Diversity rises by 0.132–0.172, carrying BERTopic from the 0.77–0.81 of its native word lists to 0.92–0.94, at a cost of at most 0.037 in C_v . On NYT, C_v improves by 0.010 instead. That is above the 0.85–0.88 MARETopic itself reaches on the same corpora. Diversification is therefore a swappable component rather than an advantage of MARETopic.

The same experiment answers a second question: how much of the reported advantage comes from the rank-based selection of leaders, and how much from the machinery MARETopic and BERTopic share, since both extract topic words with c-TF-IDF. Giving BERTopic the same word pipeline MARETopic uses, c-TF-IDF followed by inter-topic MMR at $\lambda = 0.3$, MARETopic_{CORR} still reaches a higher C_v on all three corpora: 0.6818 against BERTopic’s 0.5991 on 20NG, and 0.7150 against 0.6531 on WoS. On NYT the margin is narrow, 0.6496 against 0.6445, well inside the standard deviation of either measurement, so the two are better read as tied.

B Number of Topics

This section repeats the evaluation at a larger topic budget to assess whether the quality of greedy selection degrades as K grows. All results in Section 4.2 fix $K = 50$, the standard setting under which the baselines are published. Because MARETopic selects its topics greedily, raising K extends a sequence it has already begun rather than refitting from scratch: later leaders are chosen from whatever earlier ones left uncovered, and could in principle be weaker. Table 5 reports the outcome at $K = 100$. Among the methods that reach

Table 4: The proposed diversification applied to BERTopic’s own c-TF-IDF (mean of 3 runs, $K = 50$). Deltas are against BERTopic’s native word lists (**green**: increase; **red**: decrease). At $\lambda = 0$ the step reduces to selecting the highest-scoring words, which is what BERTopic does natively, and reproduces its word lists exactly on every run. Purity and NMI are unchanged by construction, since the diversification acts on word lists after Θ is fixed. This is an independent run, so BERTopic’s native values differ slightly from those in Table 1.

Configuration	20 Newsgroups				New York Times				Web of Science			
	C_v	TD	Purity	NMI	C_v	TD	Purity	NMI	C_v	TD	Purity	NMI
BERTopic, native words	0.6120	0.7666	0.4261	0.3770	0.6341	0.8067	0.6187	0.3155	0.6903	0.7707	0.7909	0.4488
+ inter-topic MMR, $\lambda = 0.3$	0.5991 -0.013	0.9382 +0.172	0.4261	0.3770	0.6445 +0.010	0.9382 +0.132	0.6187	0.3155	0.6531 -0.037	0.9195 +0.149	0.7909	0.4488

Table 5: The evaluation repeated at $K = 100$ topics (mean of 3 runs). Each cell carries the value at $K = 100$ and, in small type, its change from $K = 50$ (**green**: increase; **red**: decrease). Best value per metric per dataset in **bold**; runner-up underlined.

Metric	Generative					Clustering-based				
	LDA	NMF	ETM	ECRTM	FASTopic	BERTopic	Top2Vec	Ours		
								MARETopic Diffusion	MARETopic Correlation	
20 Newsgroups										
C_v	0.4817 -0.053	0.5444 -0.013	0.4508 -0.001	0.5751 -0.037	0.5584 -0.041	0.5874 -0.007	0.5995 -0.065	0.6020 -0.069	0.6123 -0.070	
TD	0.5678 +0.048	0.4129 -0.096	0.8111 -0.036	0.8098 -0.059	0.9160 -0.069	0.7218 -0.051	0.8548 -0.011	0.8413 -0.020	0.8138 -0.065	
Purity	0.4220 +0.002	0.2222 +0.007	0.3724 +0.033	0.6111 0.000	0.6202 +0.046	0.4916 +0.042	0.6123 +0.031	0.5875 +0.173	0.6537 +0.031	
NMI	0.3527 -0.024	0.1785 -0.002	0.3133 +0.010	0.5149 -0.020	0.5178 -0.002	0.4042 +0.025	0.5362 +0.010	0.4878 +0.087	0.5269 -0.015	
New York Times										
C_v	0.4327 -0.021	0.4874 -0.005	0.4418 -0.028	0.4853 +0.040	0.5492 -0.037	0.6651 +0.046	0.7001 +0.006	0.6432 -0.014	0.6244 -0.025	
TD	0.6320 +0.055	0.3920 -0.064	0.8025 -0.040	0.9087 -0.051	0.9911 -0.009	0.7389 -0.056	0.8822 -0.025	0.8422 -0.041	0.8293 -0.041	
Purity	0.5744 +0.029	0.5254 +0.064	0.6086 +0.053	0.6924 +0.039	0.6950 +0.037	0.6594 +0.041	0.6964 +0.009	0.7273 +0.123	0.7280 +0.020	
NMI	0.2897 +0.024	0.2409 +0.046	0.3223 +0.041	0.3728 +0.018	0.3796 +0.013	0.3432 +0.021	0.3858 -0.003	0.3926 +0.062	0.3929 -0.005	
Web of Science										
C_v	0.4572 -0.023	0.4436 -0.043	0.3756 -0.044	0.4992 +0.034	0.5412 -0.059	0.6447 -0.047	0.6568 -0.045	0.6218 -0.098	0.6359 -0.079	
TD	0.7020 +0.038	0.4142 -0.073	0.8902 -0.004	0.9829 -0.017	0.9940 -0.006	0.7044 -0.074	0.8703 -0.009	0.8340 -0.023	0.8156 -0.039	
Purity	0.6929 +0.040	0.6041 -0.005	0.6043 -0.041	0.7831 +0.010	0.7971 +0.021	0.7927 -0.001	0.8352 +0.006	0.8140 +0.206	0.8336 +0.023	
NMI	0.3772 -0.010	0.2854 -0.011	0.3036 -0.082	0.4293 -0.022	0.4336 -0.023	0.4176 -0.030	0.4837 -0.006	0.4364 +0.070	0.4474 -0.019	

that budget, MARETopic_{Corr} still leads Purity and NMI on all three corpora, with MARETopic_{Diff} as runner-up on NYT and WoS. Doubling the budget, therefore, does not erode the advantage.

Top2Vec is the one method that cannot be held to the budget: its topic count follows from the density structure of the embedding space, and its reduction step merges topics but never splits them, so a request for 100 returns 84, 80 and 71 topics. Its numbers are therefore not measured at the same budget as the others. On NYT the two MARETopic variants take first and second place in both Purity and NMI, ahead of Top2Vec on each.

The deltas show what raising K changes, and little of it is specific to MARETopic. C_v falls for 23 of the 27 method–corpus pairs and TD for 24, as expected when a fixed vocabulary is divided among twice as many topics, while Purity rises for all but four. The three largest Purity gains belong to MARETopic_{Diff} (+0.21 on WoS, +0.17 on 20NG, +0.12 on NYT), which closes most of its gap to the correlation variant: its weakness at $K = 50$ is

a matter of resolution rather than of the diffusion scoring itself.

C Encoder Robustness

This section evaluates the sensitivity of both MARETopic variants to the choice of SBERT backbone. Table 6 reports all four metrics across three encoders with all other parameters fixed at default ($k = 100$, $d_{\text{UMAP}} = 5$, $\lambda = 0.3$, $K = 50$, mean of 3 runs). Both methods are robust to encoder choice: for MARETopic_{Corr}, Purity varies by at most 1.4 pp and NMI by at most 1.1 pp across encoders; MARETopic_{Diff} is also stable, though its spreads are wider: up to 4.3 pp in Purity and 2.6 pp in NMI. Topic diversity is slightly more sensitive (up to 2.3 pp), which is expected given c-TF-IDF’s dependence on the vocabulary structure of each embedding space; still, all encoders keep TD in the 0.85–0.88 range. This stability is a consequence of operating on relative neighborhood order rather than absolute distances: as long as the embedding

Table 6: Robustness to SBERT backbone (all-MiniLM-L6-v2, all-mpnet-base-v2, all-distilroberta-v1), with all other parameters fixed at default ($k = 100$, $d_{\text{UMAP}} = 5$, $\lambda = 0.3$, $K = 50$, mean of 3 runs). Best value per method per dataset per metric in **bold**.

Method	Encoder	20 Newsgroups				New York Times				Web of Science			
		C_v	TD	Purity	NMI	C_v	TD	Purity	NMI	C_v	TD	Purity	NMI
MARETopic Correlation	MiniLM	0.6818	0.8787	0.6227	0.5419	0.6496	0.8702	0.7084	0.3975	0.7150	0.8547	0.8106	0.4664
	MPNet	0.6875	0.8653	0.6141	0.5405	0.6633	0.8840	0.7062	0.3950	0.7033	0.8715	0.8019	0.4647
	DistilRoBERTa	0.6686	0.8609	0.6170	0.5326	0.6310	0.8658	0.6997	0.3874	0.7097	0.8591	0.8158	0.4753
MARETopic Diffusion	MiniLM	0.6708	0.8613	0.4144	0.4011	0.6570	0.8831	0.6046	0.3305	0.7198	0.8569	0.6082	0.3660
	MPNet	0.6754	0.8707	0.4040	0.3932	0.6684	0.8649	0.6078	0.3341	0.6995	0.8800	0.5894	0.3687
	DistilRoBERTa	0.6567	0.8582	0.3837	0.3752	0.6752	0.8711	0.5930	0.3266	0.7127	0.8605	0.6328	0.3838




Figure 4: Ten representative topics from MARETopic_{Corr} on 20NG. Each panel is a leader document; bars show its top-5 c-TF-IDF terms. Topic indices follow the ordering from the greedy selection.

spaces are of comparable quality, the ranked lists — and therefore the selected leaders — remain largely the same. No single encoder dominates across all settings; MiniLM is retained as default for consistency with the BERTopic baseline.

D Qualitative Topic Examples

This section presents qualitative examples of the topics discovered by MARETopic_{Corr}. Figure 4 shows ten representative topics on 20NG. The topics cover diverse domains such as computing, sports, law enforcement, vehicles, and medicine, with highly concise keywords and little vocabulary overlap, demonstrating the quality of the greedy leader-selection.

E Document–Topic Matrix Visualization

This section visualizes the document–topic affinity matrix to assess whether the rank-based assignment generalizes to unseen data. Figure 5 shows Θ for all 7,532 20NG test documents, computed via a single `transform()` call with no retraining. The block-diagonal structure emerging on held-out data confirms that the rank-based affinity generalises beyond the training corpus: each leader attracts documents from its corresponding class while remaining near-zero elsewhere.




Figure 5: Document–topic affinity matrix Θ for MARETopic_{Corr} on the 20NG test set ($N=7,532$, $K=50$). Rows: topic leaders sorted by dominant class, labelled by top-3 words. Columns: test documents sorted by ground-truth class; red lines mark class boundaries. Color: θ_{ij} (JaccardMax affinity).