

# Bridging the Evaluation Gap: Leveraging Large Language Models for Topic Model Evaluation\*

Zhiyin Tan<sup>1,\*</sup>, Jennifer D'Souza<sup>2</sup>

<sup>1</sup>L3S Research Center, Leibniz University Hannover, Hannover, Germany

<sup>2</sup>TIB Leibniz Information Centre for Science and Technology, Hannover, Germany

## Abstract

This study presents a framework for automated evaluation of dynamically evolving topic taxonomies in scientific literature using Large Language Models (LLMs). In digital library systems, topic modeling plays a crucial role in efficiently organizing and retrieving scholarly content, guiding researchers through complex knowledge landscapes. As research domains proliferate and shift, traditional human-centric and static evaluation methods struggle to maintain relevance. The proposed approach harnesses LLMs to measure key quality dimensions—such as coherence, repetitiveness, diversity, and topic-document alignment—without heavy reliance on expert annotators or narrow statistical metrics. Tailored prompts guide LLM assessments, ensuring consistent and interpretable evaluations across various datasets and modeling techniques. Experiments on benchmark corpora demonstrate the method's robustness, scalability, and adaptability, underscoring its value as a more holistic and dynamic alternative to conventional evaluation strategies.

## Keywords

Topic modeling, Large language models, Evaluation, Natural language processing

## 1. Introduction

Topic taxonomies of Science have been traditionally used to simplify literature search, to study the structure and dynamics of scientific disciplines, or to facilitate bibliometric research evaluations [1]. These taxonomies, often hierarchical and multidisciplinary, provide a framework for categorizing knowledge and can significantly influence the dissemination and evolution of scientific information [2]. They play a crucial role in the organization of academic databases, directly impact the efficiency of information retrieval systems, and serve as essential tools for structuring and navigating vast repositories within digital library systems. As scientific output continues to grow exponentially, the need for effective and dynamically adaptable taxonomic systems [3, 4] becomes increasingly important, not just for academic researchers but also for policymakers and funding agencies aiming to identify and support pivotal research areas [5].

As we shift towards dynamically updatable taxonomies to manage the growing volumes of scientific literature, developing robust evaluation methods is essential to ensure their effectiveness [3]. These methods must verify that taxonomies adapt to the rapid evolution of scientific domains while consistently producing meaningful and coherent topics. Evaluating these dynamic systems involves assessing their accuracy in reflecting current research trends and their capacity to interlink related disciplines seamlessly [6]. This process is crucial for maintaining the integrity and utility of taxonomies in facilitating efficient research discovery and supporting informed decision-making in scientific policy and funding strategies. Human evaluation is laborious and time-consuming; thus, we also need automated evaluation systems to efficiently manage and validate these complex, dynamic structures [7].

---

*IRCDL 2025: 21st Conference on Information and Research science Connecting to Digital and Library science, February 20-21 2025, Udine, Italy*

\*This work is jointly supported by the “HybrInt - Hybrid Intelligence through Interpretable AI in Machine Perception and Interaction” project (Zukunft Nds, Niedersächsisches Ministerium für Wissenschaft, Grant ID: ZN4219) and the SCINEXT project (BMBF, German Federal Ministry of Education and Research, Grant ID: 01IS22070).

\*Corresponding author.

✉ zhiyin.tan@l3s.de (Z. Tan); jennifer.dsouza@tib.eu (J. D'Souza)

🆔 0009-0002-4166-5810 (Z. Tan); 0000-0002-6616-9509 (J. D'Souza)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

Building on the premise of automated evaluation systems, Large Language Models (LLMs) [8, 9] are well-suited to function as evaluators of dynamic taxonomies due to their advanced natural language understanding and generation capabilities [10]. Trained on extensive corpora, these models excel in discerning linguistic patterns and semantic relationships within complex datasets, making them ideal for assessing the coherence and relevance of topics generated by taxonomies. Moreover, LLMs offer a scalable, consistent, and context-sensitive approach to evaluation, overcoming key challenges of traditional methods such as reliance on human annotators or narrowly focused statistical metrics. With the ability to simulate nuanced human-like reasoning, LLMs can evaluate multiple dimensions of topic quality—such as coherence, repetitiveness, diversity, and topic-document alignment—while providing detailed, interpretable feedback.

In this work, we propose a novel framework that leverages LLMs as evaluators for topic model outputs, addressing key limitations in existing evaluation methodologies. The contributions of this work are threefold. First, we introduce a comprehensive set of metrics—coherence, repetitiveness, diversity, and topic-document alignment—that collectively capture multiple facets of topic model quality. Second, we design and implement tailored LLM prompts for each metric, ensuring consistency, interpretability, and adaptability across different datasets and topic modeling techniques. Third, we validate the framework through extensive experiments on benchmark datasets 20 Newsgroups (20NG)<sup>1</sup> and a subset of scholarly documents from the International System for Agricultural Science and Technology (AGRIS)<sup>2</sup>, demonstrating its robustness and effectiveness. This study not only provides a scalable and holistic solution for topic model evaluation but also paves the way for broader applications of LLMs in addressing dynamic and context-sensitive challenges in natural language processing. The code is available here<sup>3</sup>.

## 2. Related Work

### 2.1. Topic Modeling Approaches and Their Evolution

The field of topic modeling has evolved significantly, advancing from early matrix factorization techniques to modern probabilistic and neural architectures. One foundational work, Latent Semantic Indexing (LSI) [11], utilized singular value decomposition to uncover conceptual associations between words and documents through co-occurrence patterns, though it lacked a fully generative probabilistic model. Building on these concepts, Probabilistic Latent Semantic Indexing (pLSI) [12] introduced a probabilistic framework where words are generated from a mixture of topics. This probabilistic perspective paved the way for LDA [13], a seminal contribution that introduced a fully generative Bayesian model capable of inferring a corpus-wide set of latent topics and their associated per-document distributions. LDA’s flexible yet tractable variational inference techniques (as well as alternative inference algorithms like Gibbs sampling, as demonstrated in [14]) solidified it as a cornerstone in topic modeling research.

In more recent years, researchers have embraced neural variational inference frameworks—inspired by VAE [15]—to develop neural topic models (NTMs) such as NVDM [16], GSM/GSB/RSB [17], ProdLDA/AVITM [18], and ETM [19]. These models employ continuous latent representations and deep neural networks to capture richer semantic structures and often rely on embeddings to represent words and documents. Another emerging direction leverages contextualized embeddings and large language models (LLMs) as building blocks or alternatives to traditional topic models. Methods such as clustering-based approaches [20], CombinedTM [21], and BERTopic [22] have demonstrated that directly clustering embeddings of documents and words can yield coherent, diverse topics. More recently, research efforts like [23, 24, 25] propose employing LLMs directly to generate and refine topics, showing promise for overcoming certain limitations of classical topic modeling frameworks.

---

<sup>1</sup><http://qwone.com/~jason/20Newsgroups/>

<sup>2</sup><https://agris.fao.org/>

<sup>3</sup><https://github.com/zhiyintan/topic-model-LLMjudgment>

## 2.2. Evaluation of Topic Models: Methods and Limitations

Topic model evaluation has shifted from basic statistical measures to methods that capture human interpretability. Early work used held-out likelihood, introduced in [13] and consistently employed by [26], measure how well a trained model can predict unseen data, although it can be computationally demanding. Similarly, log probability [27] aggregates observed words and documents likelihood, providing an intuitive gauge of model fit but potentially favoring complexity over interpretability. Perplexity normalized held-out likelihood [13, 26, 28, 27, 29, 18, 30, 31, 32, 19]. However, low perplexity does not necessarily translate to coherent or meaningful topics [33], highlighting a fundamental disconnect between statistical quality and human interpretability. To bridge this gap, human-centered evaluations like word intrusion and topic intrusion tasks [33] were introduced, where human judges identify out-of-place words or topics. Human-rated coherence, first explored in [34] and later adopted by [35, 36, 37] asking annotators score topics on an ordinal scale. Although these methods closely reflect human understanding, their reliance on manual annotation constrains scalability and efficiency. In response, automated metrics were then introduced to approximate human judgments. Topic words-based coherence metrics measure how strongly top-ranked topic words co-occur in the underlying data. Early approaches, such as coherence  $C_{UCI}$  [38] and coherence  $C_{UMass}$  [35], rely on word co-occurrence frequencies and statistics derived from the training corpus. More refined metrics, like NPMI [39], normalize mutual information coherence scores  $C_{NPMI}$  to better align with human judgments [36, 40, 18, 30, 31, 41, 42, 22, 43], while coherence  $C_v$  [44] uses a variation of  $C_{NPMI}$  to compute topic coherence over a sliding window of size and adds a weight to assign more strength to more related words. Moreover, embedding-based coherence used by [45, 46, 30, 21] further improves the match to human judgement [46]. Other metrics assess different aspects of topic quality beyond coherence. Diversity metrics ensure that the discovered topics are distinctive, not redundant or overlapping. For instance, topic diversity [19] counts the proportion of unique top words across topics, while topic redundancy [47] and topic uniqueness [48] measure how frequently top words appear across multiple topics. Similarly, inverted ranked-biased overlap [21] and embedding-based diversity metrics [49, 50] compare ranked word lists or semantic distances to ensure substantial topic variety. Document-level evaluations measure how well topics capture document’s content. [51] ask annotators to rate each topic’s relevance to a given document. [52] vectorize documents associated with the selected topic and calculate a coherence score based on the document vectors. [53] tests whether an outlier topic can be identified given a document and a few topics. Supervised coverage-based methods [54] match model-generated topics to a fixed, human-defined topics, though these methods are resource intensive. More recently, LLM-based evaluations have emerged as a promising new paradigm. Studies [37, 55] demonstrate that large language models can simulate human reasoning, providing nuanced judgments of topic coherence, word intrusions. [56] proposes a set of metrics to quantify the agreement between keywords generated from documents using LLM and topic words generated from documents by a topic model. By leveraging LLMs’ extensive world knowledge and contextual reasoning, this approach overcomes the limitations of statistical co-occurrence, embedding similarities that often fail to capture semantic quality, and avoids the resource intensity of human-centric evaluations. However, current LLM-based methods often focus on a single aspect, such as coherence. In contrast, our approach integrates multiple LLM-based metrics—including coherence, repetitiveness, diversity, and topic–document alignment—into a unified framework. Rather than simply measuring co-occurrence, our framework provides interpretable evidence (e.g., flagged outlier words and identified duplicate concepts) that explains topic flaws. Following [33, 57], topic model evaluations should assess the model’s practical capabilities rather than rely on legacy metrics detached from its intended usage. Our novel topic–document alignment metrics explicitly reveal discrepancies between topic words and document content, which is crucial for applications like recommendation, summarization, and classification. By integrating document-level and topic words-based assessments, our comprehensive, adversarially validated framework bridges the gap between statistical measures and human-centered evaluations, offering actionable insights to improve topic model performance.

**Prompt 1:** Prompt for evaluating coherence.

**Prompt for coherence rating metric  $C_{\text{rate}}$**

Given a list of words [TOPIC WORDS] representing a topic, assess the degree of semantic consistency among the words in the context of the topic. Consider the remaining words in the list as the contextual basis for each word’s semantics. Rate the coherence of the topic on a scale of 1 to 3, where 1 indicates that the words are mostly unrelated, and 3 indicates that the words are highly related and form a clear, unified theme.

The rate is: [RATE]

---

**Prompt for outlier detection metric  $C_{\text{outlier}}$**

Given a list of words [TOPIC WORDS] representing a topic. Consider the remaining words in the list as the contextual basis for each word’s semantics. Which words are not semantically consistent with the remaining words and put them into a comma-separated list.

The semantically inconsistent words are: [WORD LIST]

---

**Prompt for adversarial test of outlier detection metric  $AdvT_{\text{outlier}}$**

Which words from this list [TOPIC WORDS] are not semantically consistent with the remaining words?

The semantically inconsistent words are: [WORD LIST]

### 3. Our Solution: Evaluation Metrics and LLM as Evaluator

In this section, we present our comprehensive evaluation framework that leverages LLMs as evaluators for topic models. We begin with topic words-based evaluation (Section 3.1), where we introduce metrics to evaluate topic coherence  $C_{\text{rate}}$  and  $C_{\text{outlier}}$ , and repetitiveness  $R_{\text{rate}}$  and  $R_{\text{duplicate}}$ , as well include adversarial tests to validate metric robustness. Next, cross-topic evaluation (Section 3.2) focus on topic diversity  $D_{\text{rate}}$  that captures thematic distinctiveness across topics. Finally, topic–document alignment (Section 3.3) introduces novel metrics irrelevant topic words  $A_{\text{ir-topic}}$  and missing themes  $A_{\text{missing-theme}}$  to evaluate the correspondence between topic words and document content. Together, these subsections form an integrated, interpretable, and scalable framework for evaluating topic model performance.

#### 3.1. Topic words-based Evaluation

We evaluate individual topics by examining two key dimensions: (1) coherence, which measures the semantic consistency of the top-ranked topic words using the metrics  $C_{\text{rate}}$  and  $C_{\text{outlier}}$ , and (2) repetitiveness, which assesses potential redundancy within the topic words using the metrics  $R_{\text{rate}}$  and  $R_{\text{duplicate}}$ . This two-pronged evaluation enables us to quantify both the semantic integrity and the diversity of the generated topic words.

**Coherence** Topic coherence measures how well the top-ranked topic words form a semantically unified theme. Inspired by [37], we first employ an **coherence rating metric**,  $C_{\text{rate}}$ , which asks the LLM to assess the overall semantic consistency of the topic words on a 3-point scale (with 1 indicating minimal alignment and 3 indicating strong coherence). While  $C_{\text{rate}}$  yields an overall numerical score, it does not reveal which and how many specific words are responsible for any lack of coherence. To enhance interpretability, we introduce an auxiliary **outlier detection metric**,  $C_{\text{outlier}}$ , that explicitly identifies semantic outlier words. In this procedure, the LLM extracts candidate outliers over 5 iterations, and a word flagged in at least 3 out of 5 iterations is deemed a semantic outlier. We then count the number of outliers as the final evaluation result, and the outlier words themselves are saved for later case studies. In addition, we perform an **adversarial test**  $AdvT_{\text{outlier}}$  to validate the reliability of the outlier detection inspired by established word intrusion methodologies [37, 33]. A semantically unrelated term (e.g., “Shakespeare”) is inserted into the topic list, and the LLM is expected to correctly identify the inserted term. Even if the LLM flags additional words along with the inserted term, the detection is considered successful. Each successful detection is assigned a score of 1 and each unsuccessful attempt a score of 0. The final result is calculated as the percentage of successful detections over the total number of tests. Prompt 1 presents the prompts for the two evaluation metrics as well as for the adversarial test.

**Prompt 2:** Prompt for evaluating repetitiveness.

**Prompt for repetitiveness rating metric  $R_{\text{rate}}$**

Given a list of words [TOPIC WORDS] representing a topic, evaluate if there are words that are semantically equivalent. Rate the repetitive on a scale of 1 to 3, where 1 indicates highly repetitive with significant semantic overlap, and 3 indicates minimal repetition with diverse and distinctive words.

The rate is: [RATE]

---

**Prompt for duplicate concept detection metric  $R_{\text{duplicate}}$**

Given a list of words [TOPIC WORDS] representing a topic, identify pairs of words that refer to the exact same concept or idea (not just related or similar). Provide each pair as a tuple in a comma-separated list.

The word pairs are: [WORD LIST]

---

**Prompt for adversarial test of duplicate concept detection metric  $AdvT_{\text{duplicate}}$**

Given a list of words [TOPIC WORDS] representing a topic, which words from the list have the exact same concept or idea (not just related or similar)?

The word pair with [ANCHOR] is ('None' or a word): [None/WORD]

**Prompt 3:** Prompt for evaluating diversity.

**Prompt for diversity rating metric  $D_{\text{rate}}$**

Given two groups of topic words: [TOPIC WORDS 1], [TOPIC WORDS 2], analyze the themes represented by the two groups. Rate on a scale of 1-3 based on the degree of thematic distinctiveness between the two groups: Rate 1: Partial overlapping themes. Rate 3: Highly distinctive themes.

The rate is: [RATE]

**Repetitiveness** While coherence focuses on thematic alignment, we introduce a **repetitiveness rating metric  $R_{\text{rate}}$**  to assess whether the perceived coherence is due to redundant topic words.  $R_{\text{rate}}$  is rating on a 3-point scale, where a rating of 1 indicates high repetitiveness with significant semantic overlap, and a rating of 3 indicates minimal repetition with diverse and distinctive words. To further elucidate these ratings, we introduce an auxiliary **duplicate concept detection metric**,  $R_{\text{duplicate}}$ , which explicitly identifies exact semantic repetitions in the topic word list.  $R_{\text{duplicate}}$  is critical as it helps distinguish genuine topic coherence from inflated scores due to redundancy. For each topic word, we compute a binary indicator: if a word has at least one other conceptual repetition in the list, it is assigned a value of 1, otherwise, 0. The final  $R_{\text{duplicate}}$  score for a topic is the sum of these indicators, reflecting the number of topic words that have at least one duplicate in the list. In addition, we perform an **adversarial test  $AdvT_{\text{duplicate}}$**  to validate the reliability of using the LLM for duplicate concept detection. In this test, we randomly select an anchor word from the topic word list and manually choose a conceptually identical word to serve as its duplicate. We then insert this duplicate into the topic word list. The LLM is expected to identify the inserted duplicate given the anchor word. Each successful detection is scored as 1, while an unsuccessful one is scored as 0. Prompt 2 presents the prompts designed for both  $R_{\text{rate}}$  and  $R_{\text{duplicate}}$ , as well as for the associated adversarial test.

### 3.2. Cross-topic Evaluation

**Diversity** Diversity quantifies the uniqueness among generated topics by assessing the thematic distinctiveness of their associated top words. Inspired by [50]’s word embedding-based pairwise distance, we exhaustively extract all possible pairs of topics, with each topic represented by its corresponding topic word list. For each pair, the LLM rates the thematic distinctiveness on a 3-point scale, where a rating of 1 denotes partial overlap (low diversity) and a rating of 3 denotes minimal overlap (high distinctiveness). Finally, the average of all pairwise ratings is computed to yield the overall **diversity rating**,  $D_{\text{rate}}$ . The prompt for diversity evaluation is provided in Prompt 3.

### 3.3. Topic-document Alignment

Document-level evaluation focuses on assessing how effectively topics capture the underlying themes of documents. Early methods relied on human annotations [51, 53] or supervised matching against curated references [54] to measure topic relevance. However, these approaches are resource-intensive and

**Prompt 4:** Prompt for evaluating topic-document alignment.

**Prompt for irrelevant topic words detection metric  $A_{\text{ir-topic}}$**

Given a document: [DOCUMENT] and a topic word list [TOPIC WORDS], identify which topics in the word list are not relevant to the document.

Return these extraneous topics, or [ ] if all topics in the word list are relevant to the document.

Return the extraneous topics list or [ ]: [TOPIC WORDS/[ ]]

---

**Prompt for missing themes detection metric  $A_{\text{missing-theme}}$**

Given a document: [DOCUMENT] and a topic word list [TOPIC WORDS], identify which themes present in the document are not included in the topic word list.

Return these missing themes, or [ ] if all themes from the document are included in the word list.

Return the missed themes list or [ ]: [MISSING THEMES/[ ]]

lack scalability. Recent work by [56] introduces a set of metrics that quantify the agreement between keywords generated by LLMs from documents and the topic words produced by topic models. Although these metrics capture similarity, they do not quantify the degree of mismatch between the two sets. This unaccounted discrepancy is critical for evaluating how well a topic model covers less frequent or nuanced themes, which are often key to understanding long-tail phenomena. By incorporating measures of mismatch, we can gain a more complete picture of the model’s limitations and identify specific areas where the topic representation may require further improvement. Motivated by these limitations, we propose two novel LLM-based metrics for topic-document alignment: irrelevant topic words detection metric  $A_{\text{ir-topic}}$  and missing themes detection metric  $A_{\text{missing-theme}}$ . These metrics leverage the contextual understanding of LLMs to assess both overrepresentation (irrelevant topic words) and underrepresentation (missing themes) in topic-document relationships, providing a more comprehensive evaluation of how well topics capture document content. Prompts for these evaluations are provided in Prompt 4.

**Irrelevant Topic Words Detection** Metric  $A_{\text{ir-topic}}$  assesses the extent to which a topic contains words that are not relevant to the content of its associated documents. For each topic-document pair, we instruct the LLM to identify topic words that are not explicitly or implicitly related to the document. The number of extraneous words is tallied for each document and then averaged across all pairs, providing a precise measure of overrepresentation.

**Missing Themes Detection** Conversely, metric  $A_{\text{missing-theme}}$  quantifies the extent to which a topic fails to capture key themes present in the documents. For each topic-document pair, the LLM extracts significant themes from the document that are absent in the topic word list and counts these missing themes. The resulting counts are then averaged across all pairs, yielding a measure of underrepresentation.

## 4. Experimental Setup

### 4.1. Data

**Datasets** The 20NG dataset is a widely used benchmark comprising approximately 20,000 newsgroup posts organized into 20 categories. We adopt the pre-processed version from OCTIS<sup>4</sup>, which contains 11,415 training and 4,894 test documents. Known for its diverse topics, 20NG has been extensively employed in topic model evaluations [29, 17, 32, 43]. Besides, from a repository of approximately 14 million records of food and agricultural scholarly documents—we extract a subset by excluding non-English documents, those with titles shorter than five tokens or abstracts shorter than 40 tokens, and duplicate records (by DOI, and named it **AGRIS**. The final dataset comprises 50,067 documents (45,060 for training and 5,007 for testing). For each document, we retain the title and abstract. To support

---

<sup>4</sup><https://github.com/MIND-Lab/OCTIS/>

sentence-level analysis, abstracts are segmented using SaT [58], yielding 454,850 training and 50,703 test entries. This granularity allows multiple topics to be assigned to a single document, reflecting the multi-faceted nature of scholarly texts, where a single work often spans diverse thematic areas.

**Domain-Specific Stopword Removal** Stopword removal is a critical preprocessing step, particularly for domain-specific data. Stopwords are frequent but low-information terms, and their removal, guided by Zipf’s law [59], reduces token counts while preserving vocabulary diversity, optimizing computational efficiency without compromising semantic integrity. Generic stopwords lists often overlook contextually irrelevant terms in specialized domains. For AGRIS, we employed an information-theoretic framework [60] to identify and remove domain-specific stopwords. The 20NG dataset, pre-processed by OCTIS, required no further stopwords refinement.

## 4.2. Topic Models

We evaluated four topic models chosen for their methodological diversity and proven performance, spanning traditional probabilistic approaches to neural and embedding-based methods, enabling a comprehensive comparison. Their key characteristics and implementations are detailed below.

- **Latent Dirichlet Allocation (LDA)** [13]: a foundational probabilistic topic model that represents documents as mixtures of topics, with each topic modeled as a distribution over words. We use the Gensim implementation<sup>5</sup>.
- **Product of Experts Latent Dirichlet Allocation (ProLDA)** [18]: a neural adaptation of LDA that leverages variational autoencoders to enhance scalability and improve topic coherence. We adopt the code provided by the TopMost toolkit<sup>6</sup>.
- **CombinedTM** [21]: it integrates contextual embeddings from pre-trained transformers into the LDA framework, effectively capturing semantic nuances through deep neural embeddings. We used the official implementation<sup>7</sup>.
- **BERTopic** [22]: combines document embeddings with a class-based TF-IDF procedure to generate coherent and interpretable topics. For this study, we configured BERTopic with UMAP for dimensionality reduction and HDBSCAN for clustering, following its standard pipeline<sup>8</sup>.

For each model, we conducted an extensive parameter tuning process to identify the optimal settings for two key evaluation metrics:  $C_v$  [44] for topic coherence and  $D_{\text{unique}}$  [19] for topic diversity. Once the optimal configurations were determined, we obtain the top 10 topic words and the topic-document pairs from each model on both the 20NG and AGRIS datasets with number of topic  $K = 50$  and  $K = 100$ . Each configuration was run ten times to account for variability in probabilistic and neural-based outputs. This resulted in ten aggregated sets of results for each model and dataset, ensuring a robust and statistically sound evaluation. This rigorous evaluation protocol not only ensures a fair comparison across the diverse modeling paradigms but also provides comprehensive insights into each model’s strengths and limitations in capturing thematic content across varied datasets.

## 4.3. Evaluation

**Metrics** We employ two widely recognized automated metrics as baselines: the coherence metric  $C_v$  [44], which measures the semantic consistency of top-ranked topic words, and the diversity metric  $D_{\text{unique}}$  [19], which quantifies the proportion of unique words across all topics. To address the limitations of automated metrics, we also use the suite of LLM-based metrics described in Section 3 for a more nuanced evaluation of topic quality. These include: (1) **Coherence Metrics**: coherence rating metric  $C_{\text{rate}}$  and outlier detection metric  $C_{\text{outlier}}$ . (2) **Repetitiveness Metrics**: repetitiveness rating metric

<sup>5</sup><https://radimrehurek.com/gensim/models/ldamodel.html>

<sup>6</sup><https://github.com/BobXWu/TopMost>

<sup>7</sup><https://github.com/MilaNLPProc/contextualized-topic-models>

<sup>8</sup><https://maartengr.github.io/BERTopic/>

$R_{\text{rate}}$  and duplicate concept detection metric  $R_{\text{duplicate}}$ . (3) **Diversity Metric**: diversity rating metric  $D_{\text{rate}}$ . (4) **Topic-document Alignment Metrics**: irrelevant topic words detection metric  $A_{\text{ir-topic}}$  and missing themes detection metric  $A_{\text{missing-theme}}$ . For the topic-document alignment metrics, a sample set was constructed by randomly selecting one iteration of the model’s output and sampling up to 100 associated documents per topic—yielding 59,499 samples from AGRIS and 38,321 from 20NG—thus ensuring a comprehensive evaluation. These metrics harness LLMs’ deep semantic understanding to provide a comprehensive, multi-dimensional evaluation of topic quality.

**LLM as Evaluators** We selected three open-source LLMs as evaluators for the proposed metrics: *Mistral-7B-Instruct-v0.3*<sup>9</sup> (referred to as Mistral), *Meta-Llama-3.1-8B-Instruct*<sup>10</sup> (referred to as Llama), and *Qwen2.5-14B-Instruct*<sup>11</sup> (referred to as Qwen). These LLMs, chosen for their diverse pretraining corpora and instruction-tuning objectives, exhibit robust semantic understanding. Their complementary strengths ensure reliable, scalable, and reproducible evaluations, while promoting transparency and facilitating replication in the research community.

**Efficiency and Scalability** We evaluated the computational efficiency and scalability of our LLM-based evaluation framework. All experiments were conducted on an NVIDIA A100 GPU, with approximately 40 GB allocated for coherence, repetitiveness, and diversity evaluations, and about 70 GB for topic-document alignment. The combined evaluation (three LLMs) time for coherence, repetitiveness, and diversity ranged from 15 minutes for  $K = 50$  to 35 minutes for  $K = 100$ , while topic-document alignment evaluation required between 2 and 3 hours. On GPUs with lower performance and memory, reducing the batch size enables smooth operation at the expense of increased processing time, while parallel computing can reduce runtime proportionally to the number of GPUs available. These results demonstrate that our framework is both computationally efficient and scalable, making it suitable for extensive evaluations of topic models.

## 5. Results and Discussion

### 5.1. Quantitative Results

**Adversarial Test** We sampled 100 topics (each with 10 words) from four topic models applied to the 20NG and AGRIS datasets. For adversarial test of outlier detection  $AdvT_{\text{outlier}}$ , the success rates on 20NG are 77% (Mistral), 81% (Llama), and 90% (Qwen), and on AGRIS, 82% (Mistral), 85% (Llama), and 93% (Qwen). For duplicate concept detection  $AdvT_{\text{duplicate}}$ , success rates on 20NG are, 37% (Mistral), 81% (Llama), and 84% (Qwen), and on AGRIS, 29% (Mistral), 74% (Llama), and 81% (Qwen). These results highlight significant variability among LLM evaluators and underscore the importance of using multiple evaluators to reliably assess topic quality.

**Coherence** Tables 1 and 2 indicate that, based on the coherence rating metric  $C_{\text{rate}}$ , *BERTopic* consistently outperforms the other topic models on both datasets, which aligns with the automated metric  $C_v$ . With respect to the outlier detection metric  $C_{\text{outlier}}$ , for  $K = 50$  two of the three LLMs report the fewest outliers for *BERTopic*, supporting its superior coherence. At  $K = 100$ , Qwen maintain its preference for *BERTopic* while Mistral finds that *LDA* and *CombinedTM* are comparable. In contrast, Llama’s  $C_{\text{outlier}}$  suggest that *BERTopic* has more outliers. These discrepancies are further analysed in Section 5.3.

**Repetitiveness** To discern whether high coherence reflects genuine semantic unity or is driven by redundant topic words, we examine  $R_{\text{rate}}$  and  $R_{\text{duplicate}}$ . A robust coherence result should be accompanied by a high  $R_{\text{rate}}$  (indicating minimal repetition) and a low  $R_{\text{duplicate}}$  (indicating few duplicate

<sup>9</sup><https://huggingface.co/mistralai/Mistral-7B-Instruct-v0.3>

<sup>10</sup><https://huggingface.co/meta-llama/Llama-3.1-8B-Instruct>

<sup>11</sup><https://huggingface.co/Qwen/Qwen2.5-14B-Instruct>



**Table 1**Evaluation results of coherence for 20NG: automated metric  $C_v$  vs. LLM-based metrics  $C_{rate}$  and  $C_{outlier}$ .

Coherence	$K = 50$						$K = 100$							
	$C_v$	$C_{rate} \uparrow$			$C_{outlier} \downarrow$			$C_v$	$C_{rate} \uparrow$			$C_{outlier} \downarrow$		
		Mistral	Llama	Qwen	Mistral	Llama	Qwen		Mistral	Llama	Qwen	Mistral	Llama	Qwen
<i>LDA</i>	0.558	2.140	2.390	2.047	0.678	<b>1.274</b>	2.104	0.508	2.111	2.282	1.881	<b>0.700</b>	1.344	2.334
<i>ProdLDA</i>	0.514	2.078	2.204	1.730	1.152	1.308	2.318	0.470	2.053	2.132	1.550	1.205	<b>1.239</b>	2.376
<i>CombinedTM</i>	0.560	2.088	2.366	2.086	0.778	1.376	2.108	<b>0.538</b>	2.060	2.329	1.955	0.754	1.399	2.228
<i>BERTopic</i>	<b>0.585</b>	<b>2.277</b>	<b>2.487</b>	<b>2.421</b>	<b>0.602</b>	1.471	<b>1.915</b>	0.534	<b>2.159</b>	<b>2.416</b>	<b>2.160</b>	0.842	1.423	<b>2.137</b>

**Table 2**Evaluation results of coherence for AGRIS: automated metric  $C_v$  vs. LLM-based metrics  $C_{rate}$  and  $C_{outlier}$ .

Coherence	$K = 50$						$K = 100$							
	$C_v$	$C_{rate} \uparrow$			$C_{outlier} \downarrow$			$C_v$	$C_{rate} \uparrow$			$C_{outlier} \downarrow$		
		Mistral	Llama	Qwen	Mistral	Llama	Qwen		Mistral	Llama	Qwen	Mistral	Llama	Qwen
<i>LDA</i>	0.411	2.106	2.390	2.113	0.778	1.32	2.436	0.353	2.029	2.309	1.894	0.961	<b>1.336</b>	2.462
<i>ProdLDA</i>	0.557	2.214	2.356	2.442	1.162	1.428	2.088	0.552	2.148	2.356	2.38	1.207	1.407	2.265
<i>CombinedTM</i>	0.507	2.172	2.442	2.424	0.924	<b>1.318</b>	2.234	0.571	2.186	2.486	2.457	<b>0.715</b>	1.366	2.173
<i>BERTopic</i>	<b>0.655</b>	<b>2.544</b>	<b>2.672</b>	<b>2.878</b>	<b>0.679</b>	1.558	<b>1.665</b>	<b>0.627</b>	<b>2.452</b>	<b>2.634</b>	<b>2.777</b>	0.805	1.477	<b>1.684</b>

**Table 3**Evaluation results of repetitiveness for 20NG: LLM-based metrics  $R_{rate}$  and  $R_{duplicate}$ .

Repetitiveness	$K = 50$						$K = 100$					
	$R_{rate} \uparrow$			$R_{duplicate} \downarrow$			$R_{rate} \uparrow$			$R_{duplicate} \downarrow$		
	Mistral	Llama	Qwen	Mistral	Llama	Qwen	Mistral	Llama	Qwen	Mistral	Llama	Qwen
<i>LDA</i>	2.000	1.652	2.091	2.333	3.586	<b>1.147</b>	2.000	1.641	2.120	2.208	3.673	<b>1.139</b>
<i>ProdLDA</i>	2.000	<b>1.698</b>	<b>2.204</b>	<b>2.112</b>	<b>3.566</b>	1.682	2.000	<b>1.723</b>	<b>2.248</b>	<b>2.022</b>	3.597	1.585
<i>CombinedTM</i>	2.000	1.650	2.132	2.348	3.578	1.442	2.000	1.670	2.131	2.259	<b>3.518</b>	1.291
<i>BERTopic</i>	2.000	1.595	2.034	2.598	3.686	1.379	2.000	1.639	2.064	2.459	3.575	1.377

**Table 4**Evaluation results of repetitiveness for AGRIS: LLM-based metrics  $R_{rate}$  and  $R_{duplicate}$ .

Repetitiveness	$K = 50$						$K = 100$					
	$R_{rate} \uparrow$			$R_{duplicate} \downarrow$			$R_{rate} \uparrow$			$R_{duplicate} \downarrow$		
	Mistral	Llama	Qwen	Mistral	Llama	Qwen	Mistral	Llama	Qwen	Mistral	Llama	Qwen
<i>LDA</i>	2.000	<b>1.710</b>	<b>2.122</b>	<b>2.102</b>	3.626	<b>1.436</b>	2.000	<b>1.693</b>	<b>2.168</b>	<b>1.966</b>	<b>3.486</b>	<b>1.457</b>
<i>ProdLDA</i>	2.000	1.660	2.102	2.240	<b>3.494</b>	1.628	2.000	1.643	2.099	2.195	3.537	1.663
<i>CombinedTM</i>	2.000	1.622	2.084	2.324	3.644	1.792	2.000	1.660	2.054	2.278	3.595	1.694
<i>BERTopic</i>	<b>2.005</b>	1.482	1.9984	2.577	3.639	1.913	<b>2.004</b>	1.498	2.011	2.595	3.540	1.836

concepts). For instance, although *BERTopic* shows high coherence, Tables 3 and 4 reveal that, most LLM evaluators report a low  $R_{rate}$  and a high  $R_{duplicate}$  for it, suggesting its apparent coherence may be inflated by redundant word selection. In contrast, *ProdLDA* on 20NG and *LDA* on AGRIS tend to exhibit relatively high  $R_{rate}$  and low  $R_{duplicate}$  scores, implying that their coherence is less likely to be enhanced by redundant word selection. These findings underscore the importance of assessing repetitiveness alongside coherence to ensure that high coherence truly reflect meaningful topic quality.

**Diversity** Table 5 compares the automated diversity metric  $D_{unique}$  with the LLM-based diversity rating  $D_{rate}$  for both the 20NG and AGRIS datasets. For 20NG, *ProdLDA* consistently shows the lowest diversity across both metrics, while *LDA* exhibits relatively high diversity. *BERTopic* and *CombinedTM* yield intermediate scores. In AGRIS, however, the results are less consistent: for  $K = 50$ , Mistral and Qwen rate *LDA* as highly diverse, whereas Llama assigns it a lower diversity, and for  $K = 100$ , Mistral again favors *LDA* while Llama and Qwen indicate lower diversity for *LDA*. Furthermore, although LLM-based evaluations generally rate *BERTopic* as highly diverse, its  $D_{unique}$  scores remain moderate—suggesting that even when topics share common words, contextual nuances preserve thematic distinctiveness. Overall, the divergence between  $D_{unique}$  and  $D_{rate}$  on AGRIS underscores the importance of considering both lexical uniqueness and semantic context when assessing topic diversity.

**Table 5**Evaluation results of diversity for 20NG & AGRIS: automated metric  $D_{\text{unique}}$  vs. LLM-based metric  $D_{\text{rate}}$ .

Diversity	20NG									AGRIS								
	$K = 50$						$K = 100$			$K = 50$						$K = 100$		
	$D_{\text{unique}} \uparrow$	$D_{\text{rate}} \uparrow$			$D_{\text{unique}} \uparrow$	$D_{\text{rate}} \uparrow$			$D_{\text{unique}} \uparrow$	$D_{\text{rate}} \uparrow$			$D_{\text{unique}} \uparrow$	$D_{\text{rate}} \uparrow$				
		Mistral	Llama	Qwen		Mistral	Llama	Qwen		Mistral	Llama	Qwen		Mistral	Llama	Qwen		
LDA	<b>0.636</b>	1.701	2.907	2.884	<b>0.614</b>	1.676	2.915	<b>2.917</b>	0.822	<b>1.882</b>	2.764	<b>2.917</b>	<b>0.795</b>	<b>1.807</b>	2.725	2.405		
ProdLDA	0.307	1.450	2.793	2.741	0.218	1.417	2.792	2.731	0.504	1.566	2.887	2.731	0.293	1.509	2.872	2.826		
CombinedTM	0.454	1.568	2.912	2.882	0.314	1.555	2.909	2.871	<b>0.871</b>	1.619	<b>2.904</b>	2.871	0.395	1.755	2.867	<b>2.837</b>		
BERTopic	0.597	1.581	<b>2.940</b>	<b>2.904</b>	0.366	1.580	<b>2.936</b>	2.903	0.532	1.804	2.825	2.903	0.503	1.739	<b>2.878</b>	2.808		

**Table 6**Evaluation of topic-document alignment for 20NG & AGRIS (irrelevant topic words): LLM-based metric  $A_{\text{ir-topic}}$ .

Irrelevant Topic Words ↓	20NG						AGRIS					
	$K = 50$			$K = 100$			$K = 50$			$K = 100$		
	Mistral	Llama	Qwen	Mistral	Llama	Qwen	Mistral	Llama	Qwen	Mistral	Llama	Qwen
LDA	<b>4.125</b>	5.844	7.676	4.460	6.011	7.526	5.287	<b>5.535</b>	<b>8.164</b>	5.441	5.821	8.022
ProdLDA	4.746	<b>5.282</b>	<b>7.627</b>	<b>3.996</b>	<b>4.605</b>	<b>6.440</b>	5.121	5.602	8.202	5.118	5.583	8.242
CombinedTM	4.164	5.668	7.895	4.152	5.922	8.144	5.132	5.693	8.183	<b>4.597</b>	<b>5.286</b>	<b>7.791</b>
BERTopic	4.895	7.362	9.840	5.187	7.433	9.910	<b>5.108</b>	6.560	8.578	5.244	6.748	8.680


**Table 7**Evaluation of topic-document alignment for 20NG & AGRIS (missing themes): LLM-based metric  $A_{\text{missing-theme}}$ .

Missing Themes ↓	20NG						AGRIS					
	$K = 50$			$K = 100$			$K = 50$			$K = 100$		
	Mistral	Llama	Qwen	Mistral	Llama	Qwen	Mistral	Llama	Qwen	Mistral	Llama	Qwen
LDA	<b>6.962</b>	7.475	7.481	<b>6.967</b>	<b>5.128</b>	7.694	<b>4.309</b>	<b>5.128</b>	<b>2.126</b>	4.458	<b>5.167</b>	<b>2.297</b>
ProdLDA	7.866	7.739	7.845	7.682	5.156	7.664	4.419	5.156	2.676	4.639	5.438	2.669
CombinedTM	7.398	7.634	6.982	7.347	5.221	6.974	4.492	5.221	2.509	<b>4.285</b>	5.230	2.666
BERTopic	7.582	<b>7.117</b>	<b>5.114</b>	7.641	7.213	<b>4.578</b>	4.662	5.369	2.814	4.747	5.466	2.729

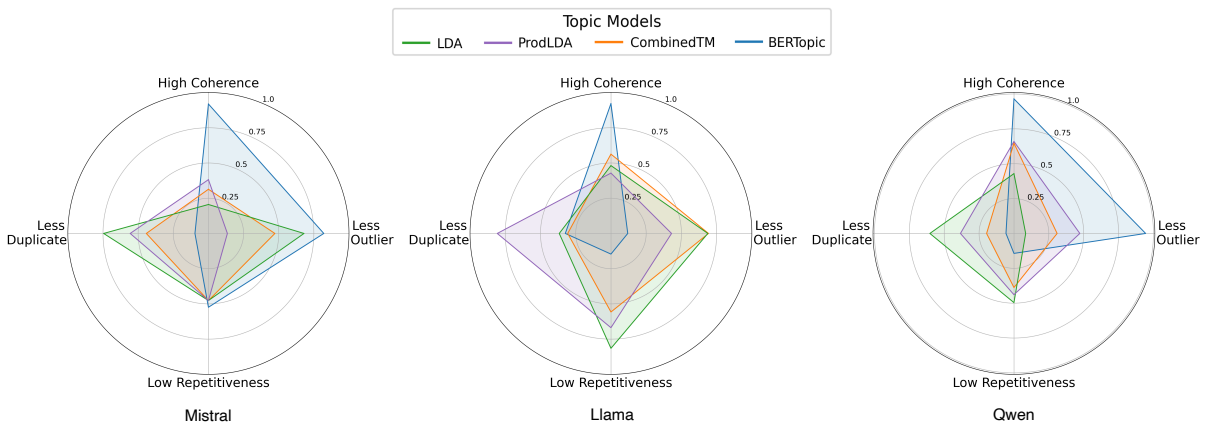
**Irrelevant Topic Words** Table 6 presents the evaluation of the irrelevant topic words detection metric  $A_{\text{ir-topic}}$ , where lower counts indicate better topic-document alignment. On 20NG, *LDA*, *ProdLDA*, and *CombinedTM* consistently yield lower counts compared to *BERTopic*, indicating that their topic words are more closely aligned with document content. In AGRIS, at  $K = 50$  two of the three LLM evaluators favor *LDA* for having the fewest irrelevant words, whereas at  $K = 100$ , *CombinedTM* achieves the lowest count, suggesting its superior ability to capture nuanced document themes—likely due to its effective integration of contextual embeddings.

**Missing Themes** Table 7 reports the missing themes detection metric  $A_{\text{missing-theme}}$ , which quantifies the number of key document themes are omitted from the topic word list, with lower counts indicating better thematic coverage. For 20NG at  $K = 50$ , two out of three LLM evaluators rate *BERTopic* as having the lowest missing theme counts, suggesting that its topic words more comprehensively represent the document themes. At  $K = 100$ , Qwen continues to favor *BERTopic*, while both Mistral and Llama indicate that *LDA* provides the best coverage. In AGRIS, however, the differences across topic models are minimal. Overall,  $A_{\text{missing-theme}}$  provides valuable insight into the extent to which topic models may fail to capture less frequent or nuanced themes from documents, which is vital for understanding long-tail phenomena and enhancing downstream applications.

**Divergent LLM Evaluation Patterns** The evaluation results reveal distinct evaluation tendencies among the three LLMs, offering valuable insights for researchers using LLMs as evaluators. In terms of coherence, Qwen consistently flags a higher number of outliers  $C_{\text{outlier}}$ , suggesting a stricter criterion for semantic consistency, while Mistral reports lower outlier counts, indicative of a more lenient evaluation; Llama’s results generally fall between these extremes. For repetitiveness, Qwen detects fewer duplicate concepts  $R_{\text{duplicate}}$  compared to Llama, with Mistral’s assessments again falling in between—demonstrating variable sensitivity to lexical redundancy across evaluators. In topic-document alignment, Qwen registers higher counts of irrelevant topic words  $A_{\text{ir-topic}}$  yet lower counts of missing



**Figure 1:** Radar plots comparing the evaluation trends of three LLMs in the results ( $K = 50$ ) of the 20NG



**Figure 2:** Radar plot comparing the evaluation trends of three LLMs in the results ( $K = 50$ ) of the AGRIS

themes  $A_{\text{missing-theme}}$  than Mistral and Llama. Moreover, all evaluators report higher missing theme counts for 20NG than for AGRIS, implying that 20NG documents exhibit greater thematic diversity and complexity. These insights underscore the influence of evaluator-specific biases on metric outcomes and highlight the importance of carefully selecting an LLM evaluator based on the intended application.

## 5.2. Visualization

**Standardization of Metrics** To enable fair comparisons across LLM evaluators (Mistral, Llama, Qwen), all metrics are normalized to the  $[0, 1]$  range using the following piecewise function:

$$X_{\text{norm}} = \begin{cases} 0.5 + \frac{X - X_{\text{mean}}}{X_{\text{max}} - X_{\text{min}}}, & \text{if higher values indicate better performance,} \\ 1 - \left(0.5 + \frac{X - X_{\text{mean}}}{X_{\text{max}} - X_{\text{min}}}\right), & \text{if higher values indicate poorer performance.} \end{cases}$$

Here,  $X_{\text{norm}}$  is the normalized score, while  $X_{\text{mean}}$ ,  $X_{\text{max}}$ , and  $X_{\text{min}}$  are the mean, maximum, and minimum values of  $X$  within each evaluator group.

**Visualization and Analysis** The radar plots (Figures 1 and 2) show clear discrepancies in Llama's coherence scores. A high coherence rate  $C_{\text{rate}}$  should imply fewer outliers  $C_{\text{outlier}}$ , yet Llama rates *BERTopic* high while flagging more outliers than other models. Similarly, Mistral gives low  $C_{\text{rate}}$  scores to *CombinedTM* and *LDA* but paradoxically finds fewer outliers in their topics.

**Table 8**Examples of outlier detection metric  $C_{\text{outlier}}$  from 20NG

Topic words	Mistral	Llama	Qwen
interested book advance fax printer print email address mail mailing	fax	fax, printer, print	advance, fax
keyboard window output problem work time run input response drug	drug	drug	drug
science evidence theory scientific observation scientist fact explain bug claim	bug	bug	bug, claim

**Table 9**Examples of duplicate concept detection metric  $R_{\text{duplicate}}$  from 20NG

Topic words	Mistral	Llama	Qwen
faith scripture religion moral good point christian church belief doctrine	(faith, belief), (christian, church)	(faith, religion), (christian, church), (doctrine, scripture)	(faith, belief), (scripture, doctrine)
disease patient health medical child year drug treatment adult number	(patient, adult), (child, adult)	(disease, health), (patient, adult), (child, patient), (treatment, drug)	(disease, treatment), (medical, drug)
client search directory package software file mail fax database project	(client, customer), (mail, email)	(client, directory), (search, package), (software, project)	None

### 5.3. Qualitative Analysis

In this section, we provide a qualitative analysis of representative examples to explore discrepancies and patterns in outlier detection, duplicate concept detection.

**Outlier Detection Discrepancies** Outlier detection is a crucial aspect of evaluating topic coherence, as it identifies semantically inconsistent words in a topic’s word list. Across the examples, outliers identified by the models often intersect but also reflect unique insights. Table 8 shows the examples of outlier detection discrepancies across different LLMs. Compared to the other two models, Mistral is more cautious in detecting outliers in topic words. On the contrary, Qwen is relatively more aggressive in detecting words with unclear semantic pointing from topic words and considering them as outliers.

**Duplicate Concept Detection Contradictions** The extracted duplicate pairs often differ significantly among the LLMs, showcasing varying thresholds for identifying conceptual overlap. Table 9 shows that Mistral treats semantically related nouns (e.g., “*christian*” and “*church*”, collective nouns where there is intersection (e.g., “*patient*” and “*adult*”), and nouns that belong to the same category (e.g. “*child*” and “*adult*”) as conceptually identical. It has also had hallucinations (e.g., detecting a non-existent repetition of the word “*customer*” for “*client*” and a non-existent repetition of the word “*email*” for “*mail*”). Llama treats grammatically related words (e.g., the verb “search” and its potential object “package”), semantically opposite words (e.g., “disease” and “health”) as conceptually identical.

## 6. Conclusion

In this work, we have introduced a comprehensive framework for evaluating topic models using LLM-based metrics that complements traditional automated metrics by incorporating nuanced measures of coherence, repetitiveness, diversity, and topic–document alignment. We designed specific evaluation protocols—including adversarial tests—to reveal not only the strengths and weaknesses of various topic models but also the intrinsic biases and judgment tendencies of different LLM evaluators. Our experiments on both the 20NG and AGRIS datasets demonstrate that LLMs can provide rich, context-sensitive insights into topic quality, while also highlighting evaluator-specific variations that are crucial for informed application in downstream tasks.

These findings illustrate the potential of our framework to expand the boundaries of topic model assessment by emphasizing both interpretability and practical application needs. Future work will focus on further refining these metrics, exploring additional LLM evaluators, and assessing how evaluator biases impact downstream tasks, thereby fostering more robust and actionable topic model assessments.

## References

- [1] L. Waltman, N. J. Van Eck, A new methodology for constructing a publication-level classification system of science, *Journal of the American Society for Information Science and Technology* 63 (2012) 2378–2392.
- [2] M. F. Moura, R. M. Marcacini, B. M. Nogueira, M. d. S. CONRADO, S. O. Rezende, A proposal for building domain topic taxonomies., In: *WORKSHOP ON WEB AND TEXT INTELLIGENCE, 1.; SIMPÓSIO BRASILEIRO EM ...*, 2008.
- [3] Y. Dauxais, U. Zaratiana, M. Laneuville, S. D. Hernandez, P. Holat, C. Grosman, Towards automation of topic taxonomy construction, in: *International Symposium on Intelligent Data Analysis*, Springer, 2022, pp. 26–38.
- [4] S. Kotitsas, D. Pappas, N. Manola, H. Papageorgiou, Scinobo: a novel system classifying scholarly communication in a dynamically constructed hierarchical field-of-science taxonomy, *Frontiers in Research Metrics and Analytics* 8 (2023) 1149834.
- [5] M. H. Bhat, S. M. Shafi, et al., Taxonomies in knowledge organisation—need, description and benefits, *Annals of Library and Information Studies (ALIS)* 61 (2014) 102–111.
- [6] J. Shang, X. Zhang, L. Liu, S. Li, J. Han, Nettare: Automated topic taxonomy construction from text-rich network, in: *Proceedings of the web conference 2020*, 2020, pp. 1908–1919.
- [7] P. Langlais, T. L. Gao, Rate: a reproducible automatic taxonomy evaluation by filling the gap, in: *Proceedings of the 15th International Conference on Computational Semantics*, 2023, pp. 173–182.
- [8] X. Amatriain, A. Sankar, J. Bing, P. K. Bodigutla, T. J. Hazen, M. Kazi, Transformer models: an introduction and catalog, *arXiv preprint arXiv:2302.07730* (2023).
- [9] J. D’Souza, A catalog of transformer models, 2023. URL: <https://orkg.org/comparison/R609337/>. doi:10.48366/R609337.
- [10] D. Stambach, V. Zouhar, A. Hoyle, M. Sachan, E. Ash, Revisiting automated topic model evaluation with large language models, in: *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, 2023, pp. 9348–9357.
- [11] C. H. Papadimitriou, H. Tamaki, P. Raghavan, S. Vempala, Latent semantic indexing: A probabilistic analysis, in: *Proceedings of the seventeenth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems*, 1998, pp. 159–168.
- [12] T. Hofmann, Probabilistic latent semantic indexing, in: *Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval*, 1999. URL: <https://sigir.org/wp-content/uploads/2017/06/p211.pdf>.
- [13] D. M. Blei, A. Y. Ng, M. I. Jordan, Latent dirichlet allocation, *Journal of machine Learning research* 3 (2003) 993–1022. URL: <https://www.jmlr.org/papers/volume3/blei03a/blei03a.pdf?ref=http://githubhelp.com>.
- [14] T. L. Griffiths, M. Steyvers, Finding scientific topics, *Proceedings of the National Academy of Sciences* 101 (2004) 5228–5235. URL: <https://dyurovsky.github.io/learning-humans-machines/class/24-class/papers/griffiths2004.pdf>. doi:10.1073/pnas.0307752101.
- [15] D. P. Kingma, M. Welling, Auto-Encoding Variational Bayes, in: *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014. arXiv:<http://arxiv.org/abs/1312.6114v10>.
- [16] Y. Miao, L. Yu, P. Blunsom, Neural variational inference for text processing, in: M. F. Balcan, K. Q. Weinberger (Eds.), *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, PMLR, New York, New York, USA, 2016, pp. 1727–1736. URL: <https://proceedings.mlr.press/v48/miao16.html>.
- [17] Y. Miao, E. Grefenstette, P. Blunsom, Discovering discrete latent topics with neural variational inference, in: D. Precup, Y. W. Teh (Eds.), *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, PMLR, 2017, pp. 2410–2419. URL: <https://proceedings.mlr.press/v70/miao17a.html>.
- [18] A. Srivastava, C. Sutton, Autoencoding variational inference for topic models, in: *International Conference on Learning Representations*, 2017. URL: <https://openreview.net/forum?id=BybtVK9lg>.

- [19] A. B. Dieng, F. J. R. Ruiz, D. M. Blei, Topic modeling in embedding spaces, *Transactions of the Association for Computational Linguistics* 8 (2020) 439–453. URL: <https://aclanthology.org/2020.tacl-1.29>. doi:10.1162/tacl\_a\_00325.
- [20] S. Sia, A. Dalmia, S. J. Mielke, Tired of topic models? clusters of pretrained word embeddings make for fast and good topics too!, in: B. Webber, T. Cohn, Y. He, Y. Liu (Eds.), *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Association for Computational Linguistics, Online, 2020, pp. 1728–1736. URL: <https://aclanthology.org/2020.emnlp-main.135>. doi:10.18653/v1/2020.emnlp-main.135.
- [21] F. Bianchi, S. Terragni, D. Hovy, Pre-training is a hot topic: Contextualized document embeddings improve topic coherence, in: C. Zong, F. Xia, W. Li, R. Navigli (Eds.), *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)*, Association for Computational Linguistics, Online, 2021, pp. 759–766. URL: <https://aclanthology.org/2021.acl-short.96>. doi:10.18653/v1/2021.acl-short.96.
- [22] M. Grootendorst, Bertopic: Neural topic modeling with a class-based tf-idf procedure, *arXiv preprint arXiv:2203.05794* (2022).
- [23] C. Pham, A. Hoyle, S. Sun, P. Resnik, M. Iyyer, TopicGPT: A prompt-based topic modeling framework, in: K. Duh, H. Gomez, S. Bethard (Eds.), *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, Association for Computational Linguistics, Mexico City, Mexico, 2024, pp. 2956–2984. URL: <https://aclanthology.org/2024.naacl-long.164>. doi:10.18653/v1/2024.naacl-long.164.
- [24] Y. Mu, C. Dong, K. Bontcheva, X. Song, Large language models offer an alternative to the traditional approach of topic modelling, in: N. Calzolari, M.-Y. Kan, V. Hoste, A. Lenci, S. Sakti, N. Xue (Eds.), *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, ELRA and ICCL, Torino, Italia, 2024, pp. 10160–10171. URL: <https://aclanthology.org/2024.lrec-main.887>.
- [25] Y. Mu, P. Bai, K. Bontcheva, X. Song, Addressing topic granularity and hallucination in large language models for topic modelling, 2024. URL: <https://arxiv.org/abs/2405.00611>. arXiv:2405.00611.
- [26] D. M. Blei, J. D. Lafferty, A correlated topic model of science, *The Annals of Applied Statistics* 1 (2007) 17–35. URL: <https://projecteuclid.org/journals/annals-of-applied-statistics/volume-1/issue-1/A-correlated-topic-model-of-Science/10.1214/07-AOAS114.full>.
- [27] D. Newman, A. Asuncion, P. Smyth, M. Welling, Distributed algorithms for topic models., *Journal of Machine Learning Research* 10 (2009).
- [28] C. Wang, D. Blei, D. Heckerman, Continuous time dynamic topic models, in: *Proceedings of the Twenty-Fourth Conference on Uncertainty in Artificial Intelligence, UAI’08*, AUAI Press, Arlington, Virginia, USA, 2008, p. 579–586.
- [29] G. E. Hinton, R. R. Salakhutdinov, Replicated softmax: an undirected topic model, in: Y. Bengio, D. Schuurmans, J. Lafferty, C. Williams, A. Culotta (Eds.), *Advances in Neural Information Processing Systems*, volume 22, Curran Associates, Inc., 2009. URL: [https://proceedings.neurips.cc/paper\\_files/paper/2009/file/31839b036f63806c3a3f47b93af8ccb5-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2009/file/31839b036f63806c3a3f47b93af8ccb5-Paper.pdf).
- [30] R. Ding, R. Nallapati, B. Xiang, Coherence-aware neural topic modeling, in: E. Riloff, D. Chiang, J. Hockenmaier, J. Tsujii (Eds.), *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Brussels, Belgium, 2018, pp. 830–836. URL: <https://aclanthology.org/D18-1096>. doi:10.18653/v1/D18-1096.
- [31] D. Card, C. Tan, N. A. Smith, Neural models for documents with metadata, in: I. Gurevych, Y. Miyao (Eds.), *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Association for Computational Linguistics, Melbourne, Australia, 2018, pp. 2031–2040. URL: <https://aclanthology.org/P18-1189>. doi:10.18653/v1/P18-1189.
- [32] H. Zhang, B. Chen, D. Guo, M. Zhou, Whai: Weibull hybrid autoencoding inference for deep topic modeling, *arXiv preprint arXiv:1803.01328* (2018).

- [33] J. Chang, S. Gerrish, C. Wang, J. Boyd-graber, D. Blei, Reading tea leaves: How humans interpret topic models, in: Y. Bengio, D. Schuurmans, J. Lafferty, C. Williams, A. Culotta (Eds.), *Advances in Neural Information Processing Systems*, volume 22, Curran Associates, Inc., 2009. URL: [https://proceedings.neurips.cc/paper\\_files/paper/2009/file/f92586a25bb3145facd64ab20fd554ff-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2009/file/f92586a25bb3145facd64ab20fd554ff-Paper.pdf).
- [34] D. Newman, J. H. Lau, K. Grieser, T. Baldwin, Automatic evaluation of topic coherence, in: R. Kaplan, J. Burstein, M. Harper, G. Penn (Eds.), *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, Association for Computational Linguistics, Los Angeles, California, 2010, pp. 100–108. URL: <https://aclanthology.org/N10-1012>.
- [35] D. Mimno, H. Wallach, E. Talley, M. Leenders, A. McCallum, Optimizing semantic coherence in topic models, in: R. Barzilay, M. Johnson (Eds.), *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Edinburgh, Scotland, UK., 2011, pp. 262–272. URL: <https://aclanthology.org/D11-1024>.
- [36] N. Aletras, M. Stevenson, Evaluating topic coherence using distributional semantics, in: A. Koller, K. Erk (Eds.), *Proceedings of the 10th International Conference on Computational Semantics (IWCS 2013) – Long Papers*, Association for Computational Linguistics, Potsdam, Germany, 2013, pp. 13–22. URL: <https://aclanthology.org/W13-0102>.
- [37] D. Stambach, V. Zouhar, A. Hoyle, M. Sachan, E. Ash, Revisiting automated topic model evaluation with large language models, in: H. Bouamor, J. Pino, K. Bali (Eds.), *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Singapore, 2023, pp. 9348–9357. URL: <https://aclanthology.org/2023.emnlp-main.581>. doi:10.18653/v1/2023.emnlp-main.581.
- [38] D. Newman, Y. Noh, E. Talley, S. Karimi, T. Baldwin, Evaluating topic models for digital libraries, in: *Proceedings of the 10th Annual Joint Conference on Digital Libraries, JCDL '10*, Association for Computing Machinery, New York, NY, USA, 2010, p. 215–224. URL: <https://doi.org/10.1145/1816123.1816156>. doi:10.1145/1816123.1816156.
- [39] G. Bouma, Normalized (pointwise) mutual information in collocation extraction, *Proceedings of GSCL 30 (2009)* 31–40.
- [40] J. H. Lau, D. Newman, T. Baldwin, Machine reading tea leaves: Automatically evaluating topic coherence and topic model quality, in: S. Wintner, S. Goldwater, S. Riezler (Eds.), *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, Association for Computational Linguistics, Gothenburg, Sweden, 2014, pp. 530–539. URL: <https://aclanthology.org/E14-1056>. doi:10.3115/v1/E14-1056.
- [41] R. Wang, D. Zhou, Y. He, Atm: Adversarial-neural topic model, *Information Processing & Management* 56 (2019) 102098. URL: <https://www.sciencedirect.com/science/article/pii/S0306457319300500>. doi:<https://doi.org/10.1016/j.ipm.2019.102098>.
- [42] R. Wang, X. Hu, D. Zhou, Y. He, Y. Xiong, C. Ye, H. Xu, Neural topic modeling with bidirectional adversarial training, in: D. Jurafsky, J. Chai, N. Schluter, J. Tetreault (Eds.), *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, Online, 2020, pp. 340–350. URL: <https://aclanthology.org/2020.acl-main.32>. doi:10.18653/v1/2020.acl-main.32.
- [43] X. Wu, X. Dong, T. T. Nguyen, A. T. Luu, Effective neural topic modeling with embedding clustering regularization, in: *International Conference on Machine Learning*, PMLR, 2023, pp. 37335–37357.
- [44] M. Röder, A. Both, A. Hinneburg, Exploring the space of topic coherence measures, in: *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining, WSDM '15*, Association for Computing Machinery, New York, NY, USA, 2015, p. 399–408. URL: <https://doi.org/10.1145/2684822.2685324>. doi:10.1145/2684822.2685324.
- [45] T. Schnabel, I. Labutov, D. Mimno, T. Joachims, Evaluation methods for unsupervised word embeddings, in: L. Màrquez, C. Callison-Burch, J. Su (Eds.), *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Lisbon, Portugal, 2015, pp. 298–307. URL: <https://aclanthology.org/D15-1036>. doi:10.18653/v1/

- [46] S. I. Nikolenko, Topic quality metrics based on distributed word representations, in: Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '16, Association for Computing Machinery, New York, NY, USA, 2016, p. 1029–1032. URL: <https://doi.org/10.1145/2911451.2914720>. doi:10.1145/2911451.2914720.
- [47] S. Burkhardt, S. Kramer, Decoupling sparsity and smoothness in the dirichlet variational autoencoder topic model, *Journal of Machine Learning Research* 20 (2019) 1–27. URL: <http://jmlr.org/papers/v20/18-569.html>.
- [48] F. Nan, R. Ding, R. Nallapati, B. Xiang, Topic modeling with Wasserstein autoencoders, in: A. Korhonen, D. Traum, L. Màrquez (Eds.), Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, Association for Computational Linguistics, Florence, Italy, 2019, pp. 6345–6381. URL: <https://aclanthology.org/P19-1640>. doi:10.18653/v1/P19-1640.
- [49] F. Bianchi, S. Terragni, D. Hovy, D. Nozza, E. Fersini, Cross-lingual contextualized topic models with zero-shot learning, in: P. Merlo, J. Tiedemann, R. Tsarfaty (Eds.), Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, Association for Computational Linguistics, Online, 2021, pp. 1676–1683. URL: <https://aclanthology.org/2021.eacl-main.143>. doi:10.18653/v1/2021.eacl-main.143.
- [50] S. Terragni, E. Fersini, E. Messina, Word embedding-based topic similarity measures, in: E. Métais, F. Meziane, H. Horacek, E. Kapetanios (Eds.), Natural Language Processing and Information Systems, Springer International Publishing, Cham, 2021, pp. 33–45.
- [51] S. Bhatia, J. H. Lau, T. Baldwin, An automatic approach for document-level topic model evaluation, in: R. Levy, L. Specia (Eds.), Proceedings of the 21st Conference on Computational Natural Language Learning (CoNLL 2017), Association for Computational Linguistics, Vancouver, Canada, 2017, pp. 206–215. URL: <https://aclanthology.org/K17-1022>. doi:10.18653/v1/K17-1022.
- [52] D. Korenčić, S. Ristov, J. Šnajder, Document-based topic coherence measures for news media text, *Expert Systems with Applications* 114 (2018) 357–373. URL: <https://www.sciencedirect.com/science/article/pii/S0957417418304883>. doi:<https://doi.org/10.1016/j.eswa.2018.07.063>.
- [53] S. Bhatia, J. H. Lau, T. Baldwin, Topic intrusion for automatic topic model evaluation, in: E. Riloff, D. Chiang, J. Hockenmaier, J. Tsujii (Eds.), Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Brussels, Belgium, 2018, pp. 844–849. URL: <https://aclanthology.org/D18-1098>. doi:10.18653/v1/D18-1098.
- [54] D. Korenčić, S. Ristov, J. Repar, J. Šnajder, A topic coverage approach to evaluation of topic models, *IEEE access* 9 (2021) 123280–123312.
- [55] H. Rahimi, D. Mimno, J. Hoover, H. Naacke, C. Constantin, B. Amann, Contextualized topic coherence metrics, in: Y. Graham, M. Purver (Eds.), Findings of the Association for Computational Linguistics: EACL 2024, Association for Computational Linguistics, St. Julian's, Malta, 2024, pp. 1760–1773. URL: <https://aclanthology.org/2024.findings-eacl.123>.
- [56] X. Yang, H. Zhao, D. Phung, W. Buntine, L. Du, Llm reading tea leaves: Automatically evaluating topic models with large language models, arXiv preprint arXiv:2406.09008 (2024).
- [57] D. M. Blei, Probabilistic topic models, *Commun. ACM* 55 (2012) 77–84. URL: <https://doi.org/10.1145/2133806.2133826>. doi:10.1145/2133806.2133826.
- [58] M. Frohmann, I. Sterner, I. Vulić, B. Minixhofer, M. Schedl, Segment any text: A universal approach for robust, efficient and adaptable sentence segmentation, in: Y. Al-Onaizan, M. Bansal, Y.-N. Chen (Eds.), Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Miami, Florida, USA, 2024, pp. 11908–11941. URL: <https://aclanthology.org/2024.emnlp-main.665>.
- [59] K. George, Zipf human behavior and the principle of least effort, 1949.
- [60] A. L. Gerlach M, Shi H, A universal information theoretic approach to the identification of stopwords, *Nature Machine Intelligence* 1 (2019) 606–612. doi:10.1038/s42256-019-0112-6.