

# Efficient Topic Extraction via Graph-Based Labeling: A Lightweight Alternative to Deep Models

Salma MEKAOUI<sup>1,\*</sup>, Hiba SOFYAN<sup>2</sup>, Imane AMAAZ<sup>2</sup>, Imane BENCHRIF<sup>2</sup>,  
Arsalane ZARGHILI<sup>3</sup>, Ilham CHAKER<sup>3</sup>, Nikola S. Nikolov<sup>1</sup>

<sup>1</sup>University of Limerick, Department of Computer Science and Information Systems, Ireland

<sup>2</sup>Euromed University Of Fez, School of Digital Engineering and Artificial Intelligence, Morocco

<sup>3</sup>Faculty of Sciences and Technology, University Sidi Mohamed Ben Abdellah, Morocco

<sup>1,\*</sup>Corresponding author: mekaoui.salma@ul.ie

## Abstract

Extracting topics from text has become an essential task, especially with the rapid growth of unstructured textual data. Most existing works rely on highly computational methods to address this challenge. In this paper, we argue that probabilistic and statistical approaches, such as topic modeling (TM), can offer effective alternatives that require fewer computational resources. TM is a statistical method that automatically discovers topics in large collections of unlabeled text; however, it produces topics as distributions of representative words, which often lack clear interpretability. Our objective is to perform topic labeling by assigning meaningful labels to these sets of words. To achieve this without relying on computationally expensive models, we propose a graph-based approach that not only enriches topic words with semantically related terms but also explores the relationships among them. By analyzing these connections within the graph, we derive suitable labels that accurately capture each topic's meaning. We present a comparative study between our proposed method and several benchmarks, including ChatGPT-3.5 (OpenAI, 2023), across two different datasets. Our method achieved consistently better results than traditional benchmarks in terms of BERTScore and cosine similarity and produced results comparable to ChatGPT-3.5, while remaining computationally efficient. Finally, we discuss future directions for topic labeling and highlight potential research avenues for enhancing interpretability and automation.

**Keywords:** Topic Modeling, Topic Labeling, Graph-based Methods, ConceptNet, Natural Language Processing

## 1. Introduction

Extracting topics from text is a rapidly growing area of research, driven by the enormous and continuously increasing volume of textual data across various domains. This surge has created a demand for methods that can accurately and efficiently identify the underlying themes within large corpora. Many recent studies focus on the application and advancement of Natural Language Processing (NLP) techniques, particularly leveraging sophisticated methods such as topic extraction, semantic understanding, and text organization, often using Large Language Models (LLMs) (Wang et al., 2024; Nakanishi, 2025; Bitaraf et al., 2024; Reynolds, 2025; Li and Chen). However, these approaches are typically computationally expensive and require substantial resources to achieve high-quality results.

In contrast, topic modeling (TM) provides a well-established and efficient alternative. TM is an unsupervised machine learning (ML) technique that automatically discovers latent topics in large collections of unlabelled text, without requiring prior training data (Churchill and Singh, 2022). By identifying the underlying thematic structure of a corpus, TM enables researchers to generate representa-

tive topics in a resource-efficient manner. While TM can reveal the underlying structure of a corpus, a major limitation lies in its inability to directly produce human-readable topics (Churchill and Singh, 2022). Typically, TM outputs are a distribution of words associated with each topic, which may indicate the most relevant terms but often lacks interpretability for human readers. To address this challenge, topic labeling (TL) has been introduced. TL (Mekaoui et al., 2025) is an approach for assigning comprehensive and semantically meaningful labels to TM results, ensuring that they are interpretable and understandable from a human perspective. TL has gained increasing importance and has been applied across diverse domains. A notable application is presented in (Nolasco and Oliveira, 2016) who applied TL on a scholarly dataset derived from KDD conference proceedings (Knowledge Discovery and Data Mining), demonstrating that their method can effectively generate topic labels with efficiency comparable to human-generated labels. In (Tiwari et al., 2023a), TL was applied to extract topics from short texts. Similarly, Hagerer et al. (Hagerer et al., 2021) utilized TL on a dataset of online user comments about organic food, collected from multiple platforms such as Reddit, Quora, Disqus, and major news websites including The Wash-

ington Post and The New York Times.

Although topic labeling has been widely adopted for labeling topics in various domains, numerous studies emphasize persistent challenges and limitations in this field. A central concern relates to computational complexity. Neural-based approaches (Ramón-Ferrer et al., 2023; Mao et al., 2016; Tiwari et al., 2023a; Mu et al., 2024) show that resource-intensive models, such as ChatGPT, require considerable computational power and often suffer from limited accessibility. Several works point to methodological limitations in TL, including the inability to create novel labels, difficulties in selecting an appropriate model, and the tendency to produce labels that are either too generic or insufficiently discriminative for meaningful content interpretation (Alokaili et al., 2020; Flondor and Frincu, 2023; Ramón-Ferrer et al., 2023; Tang et al., 2016; Nolasco and Oliveira, 2016; Sanjaya et al., 2018; Kozbagarov et al., 2021; Li et al., 2016, 2019). Ontology and graph-related challenges are also widely recognized, as these approaches typically involve elaborate, multi-stage pipelines that integrate TM outputs into ontological or graph structures and align them with domain-specific concepts to derive suitable labels for TM results (Allahyari and Kochut, 2015; Hagerer et al., 2021; Adhitama et al., 2017; Scelsi et al., 2021; Chaudhary et al., 2024). Furthermore, studies involving human labeling highlight the substantial manual effort required in the labeling process (Flondor and Frincu, 2023; Sanjaya et al., 2018; Zheng and Li, 2022). Finally, issues of coherence, granularity, precision, and objectivity in label interpretation are consistently reported in the literature (Mao et al., 2016; Sanjaya et al., 2018; Davoudi and An, 2015; He et al., 2019; Gourru et al., 2018; Mao et al., 2012).

In this paper, we introduce a novel TL method that leverages graph structures to generate representative labels for sets of words which represent a topic. Our method addresses several limitations identified in previous studies. Specifically, we focus on integrating TM outputs into a graph structure and identifying the most representative nodes within that graph through a simpler and more transparent procedure. We further hypothesize that adopting a graph-based approach, rather than relying on computationally intensive neural models, can alleviate issues related to high resource demands, model selection complexity, and the tendency of existing methods to produce labels that are overly generic or insufficiently discriminative.

This paper is structured as follows. Section 2 reviews the existing works on TL. Section 3 presents the methodology adopted in this study. Section 4 discusses and analyzes the experimental results, including comparisons with benchmark methods. Finally, Section 5 concludes the paper and outlines

future research directions.

## 2. Related Works

This section reviews the main approaches proposed in the literature for TL. Existing methods can be broadly categorized into four groups: traditional methods, neural-based methods, ontology and graph-based methods, and hybrid approaches.

### 2.1. Traditional Methods

Earlier attempts to solve the TL task mainly relied on classical machine learning (ML) algorithms and information retrieval (IR) strategies. These include well-known classifiers such as support vector machines (SVM) (Effendi et al., 2023), k-nearest neighbors (KNN) (Mao et al., 2016), in addition to extensions of TM methods like non-negative matrix factorization (NMF) and latent dirichlet allocation (LDA) (Li et al., 2022; Zha and Li, 2019; Mao et al., 2015; Liu and Tang, 2019; Wood et al., 2017). From the IR perspective, methods have been designed to extract candidate terms or phrases using a predefined dictionary of meaningful words (Tang et al., 2016), candidate selection algorithms (Nolasco and Oliveira, 2016), or truth discovery frameworks (Sanjaya et al., 2018). Finally, some studies rely on human expertise for topic annotation. In these cases, domain experts or annotators validate and assign labels manually to the TM results (Flondor and Frincu, 2023; Korenčić et al., 2015). Traditional TL approaches remain limited by manual intervention, slow processing, and reliance on external knowledge sources, leading to inconsistent and less adaptive labels across datasets.

### 2.2. Neural-based Methods

With the rise of deep learning, neural models have gained increasing attention in TL research. Some studies directly employ large language models (LLMs) to generate suitable labels for TM results (Ramón-Ferrer et al., 2023; Mu et al., 2024), while others adapt or fine-tune them for this specific task. For instance, the Bidirectional and Auto-Regressive Transformer (BART) (Lewis et al., 2019) has been extended into a TL variant (BART-TL) through a weakly supervised fine-tuning process (Popa and Rebedea, 2021). Similarly, the Seq2Seq architecture has been applied using the Wikipedia dataset to produce meaningful topic labels (Alokaili et al., 2020). Contextual embeddings have also been investigated, such as in the contextualized topic model (CTM) (Zheng and Li, 2022). Other approaches rely on word embedding strategies to enrich the semantics of the label (Kozbagarov et al., 2021; Tiwari et al., 2023b).

Along the same line, neural embedding topic labeling (NETL) exploits neural embeddings, specifically word2vec (Church, 2017) for word representations and doc2vec (Mikolov et al., 2013) for document/title embeddings, to identify the most relevant labels for each topic (Bhatia et al., 2016). Neural-based approaches leverage deep architectures to capture semantic relationships beyond TM resulted words. However, they often require large annotated datasets, high computational resources, and struggle with multilingual adaptation and long-text processing.

### 2.3. Ontology and Graph Based Methods

Another area of research leverages ontologies to map TM results (typically word lists) to domain concepts and categories. This knowledge-driven approach has been applied in several studies, demonstrating its effectiveness in producing semantically meaningful topic labels (Kinariwala and Deshmukh, 2021; Allahyari and Kochut, 2015; Adhitama et al., 2017; Kim and Rhee, 2019; Davoudi and An, 2015; Magatti et al., 2009; Slabbekoorn et al., 2016; Zosa et al., 2022).

Building on this semantic perspective, graph-based representations have also been explored to capture structural relationships among topics or terms. Such approaches include social graph modeling (Hagerer et al., 2021), graph-based ranking (He et al., 2019), semantic enrichment using the ConceptNet knowledge graph (Chaudhary et al., 2024), and hierarchical topic construction with Bayesian Rose Trees (BRTs), which discover topic labels by organizing words and concepts into semantic hierarchies (Ogawa and Saga, 2021). Ontology and graph-based methods leverage structured knowledge sources to enhance the interpretability of TL. However, they face challenges related to ontology construction, concept selection, and the accurate mapping of topic words to semantic hierarchies, which can be complex.

### 2.4. Hybrid approach

These methods begin by extracting relevant knowledge from external resources or specialized domain dictionaries, such as Wikipedia (Vathi et al., 2017; Hossain et al., 2019; Scelsi et al., 2021; Lau et al., 2011). This foundational knowledge is then combined with machine learning algorithms, neural networks, or other computational techniques to improve topic identification. For example, Gourru et al. (Gourru et al., 2018) demonstrate a hybrid approach by merging multiple unsupervised n-gram-based labelers. Similarly, Li et al. (Li et al., 2016, 2019) showcase the effectiveness of integrating various methods within their work. However, these hybrid approaches are often limited by evaluations

on small or narrow datasets, raising concerns about their generalizability. Additionally, while combining multiple methods can improve coverage, issues such as label precision and the selection of the most appropriate labeler remain largely unresolved.

## 3. Methodology

Our approach differs from prior TL methods in three fundamental ways. First, instead of treating TM words as independent unordered terms, we model their distribution as a single sequential textual unit obtained by concatenation. We argue that this preserves contextual interactions between topic words and yields a more coherent semantic representation of the topic. Second, whereas previous work relies on multiple graph-based metrics, including counting, distance measures, degree statistics, similarity variants, and centrality scores, to rank candidate nodes, we argue that a carefully designed similarity algorithm alone is sufficient to identify the most representative node. This substantially simplifies the scoring mechanism while maintaining effectiveness. Third, prior approaches often return to the original source documents to compute additional statistics such as frequency or similarity. In contrast, our objective is to design a lightweight method suitable for large-scale documents; therefore, revisiting the original corpus would contradict this goal. For this reason, our method operates exclusively on the extracted topic-word distribution without reprocessing the full documents.

Our proposed methodology to address the problem of TL is divided into 5 parts, as shown in Figure 1. First, the input consists of a set of words representing the topic to be discovered, which corresponds to the results of TM. Using this input, we construct a single text sequence and transform it into its embedding (Section 3.1). Then, an algorithm specifically designed to extract the most suitable label for each topic set, is applied to generate the final representative label (Section 3.2). Each component of the process is detailed in the following subsections.

### 3.1. Single Text Sequence Construction and Embedding

For each topic, we compute a single text sequence embedding by encoding all topic words together. Formally, let a topic be represented by a set of words

$$T = \{w_1, w_2, \dots, w_n\}.$$

We construct a single text sequence by concatenating the words with commas and spaces:

$$S = w_1, w_2, \dots, w_n$$

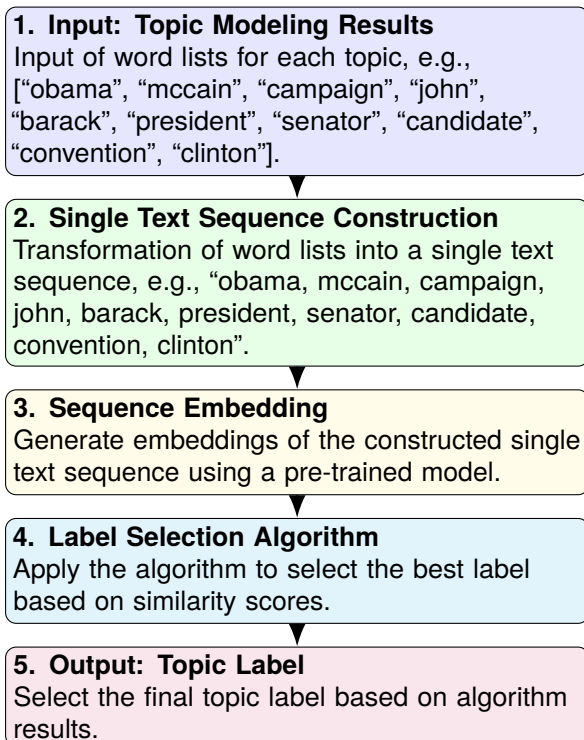


Figure 1: Visual schema of the proposed methodology for TL, illustrating the main steps from input topic words to the selection of the final representative label.

and compute its embedding:

$$E_{\text{topic}} = \text{Embed}(S),$$

where  $\text{Embed}$  denotes a pre-trained embedding model. This embedding captures the semantic representation of the topic as a whole.

We adopt this method because the words representing a topic are semantically interdependent and must be considered together to deliver the intended meaning. For instance, as shown in Table 2, in the second row, the word "president" alone does not clearly indicate that the topic concerns the President of the United States. However, when combined with related words such as "Barack", "Obama", "John", "McCain", "senator", and "Clinton", the topic becomes unambiguous. Similarly, the inclusion of the word "candidate" further reinforces the interpretation that the topic relates to U.S. presidential nominees. This observation motivates us to treat the set of topic words as a single text sequence and compute its embedding accordingly, rather than embedding words separately.

To ensure robustness and identify the most suitable embedding model for our method, we evaluate a diverse set of pre-trained models, including those listed in Table 1. These models are selected based on three key criteria: they are open-source, have manageable parameter sizes (all under 400M), and demonstrate strong performance in

benchmark evaluations, as reported by the *Massive Text Embedding Benchmark* (Hugging Face, 2025). This selection enables a comprehensive evaluation of embeddings that balance efficiency and accuracy, reinforcing the robustness of our method.

Models	Parameters
gtr-t5-large (Ni et al., 2021)	335M
GIST-large-Embedding-v0 (Solatorio, 2024)	335M
all-mpnet-base-v2 <sup>1</sup>	109M
GIST-small-Embedding-v0 (Solatorio, 2024)	33M
all-MiniLM-L12-v2 <sup>3</sup>	33M
all-MiniLM-L6-v2 <sup>2</sup>	22M
GIST-all-MiniLM-L6-v2 (Solatorio, 2024)	22M

Table 1: List of pre-trained embedding models considered in our evaluation, including their respective parameter sizes.

Each pre-trained model produces its own topic embedding  $E_{\text{topic}}$ , which serves as a reference in the subsequent experiments. In particular, individual candidate word embeddings  $E_{c_{w_i}} = \text{Embed}(c_{w_i})$  from each experiment are compared to the corresponding  $E_{\text{topic}}$  to determine the most representative word for the topic. By evaluating multiple embedding models, we aim to compare their effectiveness and select the one that yields the most accurate and semantically meaningful labels in our setting.

### 3.2. Label Selection Algorithm

In this subsection, we present two algorithms that were adopted in our method. These algorithms are closely related; however, the first algorithm relies solely on the set of TM words, whereas the second incorporates graph-based information. We hypothesize that integrating the graph will add greater semantic richness to the set of words and help produce more representative and accurate topic labels. Nevertheless, there is also a possibility that the graph may introduce additional noise, making it unnecessary if the TM words alone are sufficient. Therefore, we designed two separate algorithms to evaluate both scenarios and later compare their performance to determine which achieves the best results.

<sup>1</sup><https://huggingface.co/sentence-transformers/all-mpnet-base-v2>

<sup>2</sup><https://huggingface.co/sentence-transformers/all-MiniLM-L6-v2>

<sup>3</sup><https://huggingface.co/sentence-transformers/all-MiniLM-L12-v2>

### 3.2.1. Algorithm 1: Direct Similarity Labeling (DSL)

We compare each single candidate word embedding  $E_{c_{w_i}}$  from the original word set that represents the TM results to  $E_{\text{topic}}$  and select the word with the highest similarity as the topic label. The objective of this method is twofold: (i) to evaluate how well individual embeddings and the aggregated topic embedding perform in identifying the most representative label, and (ii) to validate our general assumption that the most representative label is the word whose embedding is closest to the topic embedding  $E_{\text{topic}}$ .

---

#### Algorithm 1: Direct Similarity Labeling (DSL)

---

**Input:** Embedding of single text sequence  
 $E_{\text{topic}}$ , Topic Keywords  
 $\{W_1, W_2, \dots, W_n\}$   
**Output:** Final label  $L^*$   
**foreach** Topic keyword  $W_i$  **do**  
    Compute embedding  $E_{w_i}$  for  $W_i$ ;  
    Compute cosine similarity  
     $S_i = \cos(E_{\text{topic}}, E_{w_i})$ ;  
Select topic label  
 $L^* = W_j$  where  $j = \arg \max_i S_i$ ;  
**return**  $L^*$

---

### 3.2.2. Algorithm 2: Graph-Enhanced Labeling (GEL)

**Graph Generation Process** For graph generation, we adopted the method proposed by Akhil *et. al* (Chaudhary *et al.*, 2024). Starting from the words output by TM, we use the ConceptNet API (Liu and Singh, 2004) to construct the most highly connected graph achievable. The intuition is that words belonging to the same topic are likely to be closely related within the ConceptNet knowledge graph. We chose ConceptNet (Liu and Singh, 2004) as our external knowledge source because it is freely available, contains rich semantic relations, and has shown strong results in previous studies (Chaudhary *et al.*, 2024). Our procedure involves two steps: (1) querying each word individually to retrieve its associated concepts and (2) iteratively expanding the graph by querying the neighbors of these concepts. We limited the expansion to three hops, following the findings of Akhil *et. al* (Chaudhary *et al.*, 2024) that three iterations are sufficient to obtain a well-connected graph without introducing excessive noise.

The motivation behind this method is our hypothesis that the words generated by the TM represent a coherent theme. Therefore, when mapped into a concept graph, these words should eventually

become interconnected after a few hops. This process allows us to construct a connected graph that captures semantic links among topic words. As illustrated in Figure 2, words from the first topic (e.g., “server”, “infrastructure”, “virtualisation”, and “virtual”) become connected through ConceptNet, enriching the vocabulary and providing a clearer understanding of the relationships between them.

**Algorithm description** Formally, We compare each single candidate word embedding  $E_{c_{w_i}}$  from the nodes of the constructed graph to  $E_{\text{topic}}$ , and select the word with the highest similarity as the topic label. The objective of this algorithm is to evaluate whether incorporating the graph structure improves label quality. The graph provides an enriched vocabulary that is closely related to the topic, allowing candidate labels to be drawn not only from the original set of words but also from graph-connected nodes.

---

#### Algorithm 2: Graph-Enhanced Labeling (GEL)

---

**Input:** Embedding of single text sequence  
 $E_{\text{topic}}$ , Topic Keywords  
 $\{W_1, W_2, \dots, W_n\}$   
**Output:** Final topic label  $L^*$   
Generate a ConceptNet-associated graph  
 $G$  from the topic  
keywords  $\{W_1, W_2, \dots, W_n\}$ ;  
Compute Embedding of all nodes of  $G$  to  
obtain node embeddings  
 $\{En_1, En_2, \dots, En_m\}$ ;  
**foreach** node  $n_j$  with embedding  $En_j$  in  $G$   
**do**  
    Compute cosine similarity  
     $S_j = \cos(E_{\text{topic}}, En_j)$ ;  
Select Topic label  
 $L^* = n_k$  where  $k = \arg \max_j S_j$  **return**  
 $L^*$

---

## 4. Experiments, Results and Discussions

In this section, we present and analyze the experimental results obtained from the proposed algorithms, DSL and GEL. The section is divided into two main parts. The first part evaluates DSL and GEL on the Topic\_Bhatia dataset, allowing comparison with prior studies that have shown strong performance in topic labeling for this corpus (Chaudhary *et al.*, 2024; Alokaili *et al.*, 2020). The second part reports results on the 20 Newsgroups dataset, demonstrating the generalizability of our method across different domains and enabling comparison

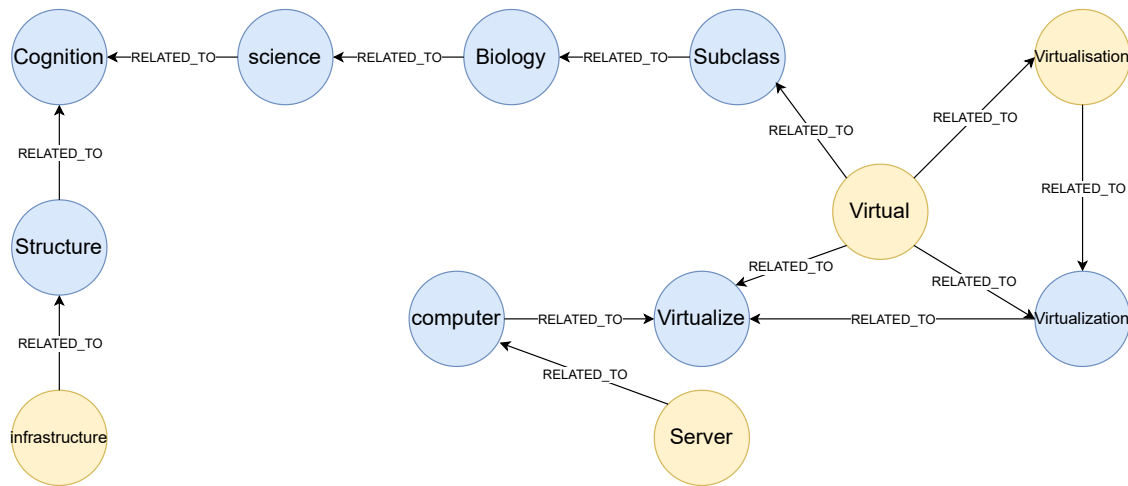


Figure 2: Visualization of how topic words (e.g., “server”, “infrastructure”, “virtualisation”, and “virtual”) become interconnected within ConceptNet after iterative expansion, resulting in a well-connected graph.

with previous works on this benchmark (Ramón-Ferrer et al., 2023).

In addition to evaluating our algorithms, we apply a large pre-trained model specifically fine-tuned for TL, BART-TL (Popa and Rebedea, 2021), to the Topic\_Bhatia dataset. While benchmark results from prior studies are already reported, BART-TL has not previously been applied to this dataset. Therefore, we conduct this experiment to assess the performance of its two variations (BART-TL-ng and BART-TL-all) directly and to compare it with our algorithms as well as with the reported benchmarks. This setup allows us to evaluate both the effectiveness of our method and the potential advantages of applying a pre-trained model in this context (Popa and Rebedea, 2021).

All datasets, experimental setups, and results are available in a repository<sup>1</sup>.

#### 4.1. Experiments on Topic\_Bhatia Dataset

The evaluation of the quality of the generated labels is conducted using BERTScore (Zhang et al., 2019). BERTScore computes the similarity between predictions and references based on contextual embeddings, and has been shown to correlate strongly with human judgments, as demonstrated in its original paper (Zhang et al., 2019). One of the main limitations of traditional TL evaluation methods is their reliance on exact word matches, such as the accuracy, precision, recall, F1 score, and macro F1 score, which overlook semantically related labels (Mekaoui et al., 2025). To address this, we adopt BERTScore since it accounts for semantic

similarity rather than requiring exact lexical overlap. Furthermore, as the baselines also employed BERTScore for evaluation, using the same metric ensures a fair and consistent comparison across all methods.

##### 4.1.1. Dataset Description

The Topic\_Bhatia dataset contains 228 topics, each associated with 19 candidate labels. The topics are drawn from four domains: blogs articles, English language books, New York Times news articles, and PubMed biomedical abstracts (Lau et al., 2011). Human ratings for these labels were collected through a crowdsourcing task on Amazon Mechanical Turk (MTurk), where annotators evaluated labels on a scale from 0 (lowest) to 3 (highest). Only labels with an average rating of 2 or above were retained, resulting in a refined dataset of 219 topics and 1,156 topic label pairs, compared to the original 4,332 combinations. Examples of this dataset are mentioned in Table 2.

##### 4.1.2. Results and Discussion

As shown in Table 3, the proposed algorithms achieve high performance across all BERTScore metrics. Prior to these experiments, the highest recorded F1 score on this dataset was 0.936, achieved by Top2Label (Chaudhary et al., 2024) (see Table 4). In comparison, DSL reached an F1 score of 0.955 using the all-MiniLM-L6-v2 embedding, demonstrating that, in some cases, a single word from the topic modeling results is sufficient to identify the most accurate topic label. GEL also achieved a top F1 score of 0.955, this time using gtr-t5-large. Notably, GEL performed consistently well across almost all embeddings, indicating

<sup>1</sup>[https://github.com/SalmaMekaoui/Enhancing\\_LLms\\_through\\_knowledge\\_graphs](https://github.com/SalmaMekaoui/Enhancing_LLms_through_knowledge_graphs)

Dataset	Example Terms	Possible Labels
Topic_Bhatia	vmware, server, virtual, oracle, update, virtualisation, application, infrastructure, management, microsoft	cloud computing, microsoft exchange server, vmware, web application, virtualization, operating system
	obama, mccain, campaign, john, barack, president, senator, candidate, convention, clinton	conservative, democrat, democratic party (US), presidential nominee, bill clinton
20 Newsgroups	game, team, play, hockey, player, win, goal, season, fan, playoff	sport hockey
	god, religion, atheist, moral, claim, point, objective, good, belief, argument	religion atheism

Table 2: Example topic terms and corresponding possible labels for the two datasets used in our study: *Topic\_Bhatia* and *20 Newsgroups*.

that incorporating graph-based relationships between words enriches the contextual information and produces labels that are semantically closer to the gold standard.

These results suggest that both DSL and GEL provide significant improvements over existing methods, while the main difference between DSL and GEL lies in the balance between precision and recall. In DSL, precision and recall are more closely aligned, suggesting that the candidate label captures most of the reference information without introducing irrelevant content. In GEL, precision exceeds recall, indicating that while the predicted label is semantically accurate, it covers only a limited portion of the reference. Although, the numerical difference is relatively small (0.02).

Turning to pretrained models, as presented in Table 4, BART-TL-ng achieved an F1 score of 0.892. Across all metrics, the results of this experiment did not surpass those of DSL and GEL, nor did they exceed the benchmark. This highlights limitations of pre-trained models in generating semantically representative labels for TM results.

Overall, these results indicate that using relatively simple methods based on the TM words alone is sufficient to achieve high-quality labels, and there is no need for more complex or computationally costly approaches. Although DSL and GEL achieved very similar scores, determining which method performs better would require a manual comparison between the predicted and gold labels, which remains resource-intensive. Furthermore, the *Topic\_Bhatia* dataset includes multiple gold reference labels for each topic word set as shown in Table 2. Accordingly, the evaluation algorithm computes BERTScore between each generated label and all reference labels separately, selecting the highest similarity score as the final evaluation value. This setup helps explain the relatively high scores obtained by both DSL and GEL.

To further validate our method, we applied the same set of experiments to the 20 Newsgroups

dataset. The results from this evaluation are reported in Section 4.2.

## 4.2. 20 Newsgroups Dataset

For the 20 Newsgroups dataset, we evaluated the quality of the generated labels using the cosine similarity metric. This evaluation method differs from the previous experiments because our primary objective here is to directly compare our method with an established benchmark (Ramón-Ferrer et al., 2023). Virginia et al. (Ramón-Ferrer et al., 2023) employed the same dataset and assessed multiple conversational models to determine their effectiveness in producing semantically coherent labels based on cosine similarity. The authors find out that the best method was chatGPT 3.5 with a cosine similarity score of 0.655 (Ramón-Ferrer et al., 2023).

To ensure a fair and consistent comparison, we adopted the same evaluation strategy and applied the same embedding method, namely `all-MiniLM-L6-v2`, for computing cosine similarity scores. Another objective of this experiment is to evaluate the impact of DSL and GEL across multiple datasets. The experiments on the *Topic\_Bhatia* dataset were unable to provide an unbiased and conclusive comparison of DSL and GEL. This limitation motivated us to extend the evaluation to additional datasets to better assess which algorithm produces the most effective results.

### 4.2.1. Dataset Description

This dataset consists of more than 20,000 posts organized into 20 distinct newsgroups and is widely used in document clustering research. In this work, we use a version where each of the 20 topics is represented by its label and a set of the 10 most representative words. Examples of this dataset are presented in Table 2.

Experiment	Embedding Method	BERT Precision	BERT Recall	BERT F1
DSL	all-MiniLM-L6-v2	<b>0.958</b>	<b>0.953</b>	<b>0.955</b>
	all-MiniLM-L12-v2	0.948	0.944	0.945
	all-mpnet-base-v2	0.943	0.944	0.943
	gtr-t5-large	0.944	0.943	0.942
	GIST-small-Embedding-v0	0.951	0.951	0.951
	GIST-large-Embedding-v0	0.940	0.943	0.941
	GIST-all-MiniLM-L-v2	0.952	0.949	0.949
	GEL	all-MiniLM-L6-v2	<b>0.964</b>	0.945
all-MiniLM-L12-v2		0.961	0.944	0.952
all-mpnet-base-v2		0.961	0.945	0.952
gtr-t5-large		0.963	<b>0.948</b>	<b>0.955</b>
GIST-small-Embedding-v0		0.955	0.939	0.946
GIST-large-Embedding-v0		0.961	0.942	0.951
GIST-all-MiniLM-L6-v2		0.956	0.937	0.946

Table 3: Performance comparison of the application of the adopted method. Experiments with DSL and GEL, corresponding to algorithms 1 and 2, report results across multiple embedding models.

Category	Method	BERT Precision	BERT Recall	BERT F1
Baselines	Alokaili topics_bhatia	0.919	0.926	0.919
	Alokaili topics_bhatia_tfidf	0.930	0.933	0.929
	Top2Label_all_relations	<b>0.935</b>	<b>0.937</b>	<b>0.936</b>
BART-TL Benchmark	BART-TL-all	0.873	0.900	0.886
	BART-TL-ng	<b>0.880</b>	<b>0.905</b>	<b>0.892</b>
Proposed Method	DSL(all-MiniLM-L6-v2)	0.958	<b>0.953</b>	<b>0.955</b>
	GEL(gtr-t5-large)	<b>0.963</b>	0.948	<b>0.955</b>

Table 4: Performance comparison of baseline, proposed, and pretrained methods. For the proposed algorithms, DSL and GEL, only the results corresponding to the best-performing embeddings are reported. The table also includes a pretrained benchmark model (BART-TL) for comparison.

#### 4.2.2. Results and Discussion

From the results illustrated in Figure 3, it is evident that GEL outperforms both DSL and Bart-TL. This suggests that incorporating graphs provides significantly better results compared to using only the original set of words or a pre-trained model. Furthermore, although DSL and GEL do not surpass the benchmark performance achieved by ChatGPT (cosine similarity = 0.655), they still achieve strong results. Specifically, GEL achieves 0.627 with the GIST-all-MiniLM-L6-v2 embedding, while DSL achieves a cosine similarity of 0.578 using the all-MiniLM-L12-v2 embedding. Notably, these two experiments, which are highly interpretable and explainable, outperform BART-TL, which achieved a very low cosine similarity of 0.31 (as shown in Table 5), highlighting the effectiveness of simple and explainable methods. Moreover, these results are obtained with a relatively simple architecture and minimal computational resources.

For the *Topic\_Bhatia* dataset, each topic could

be associated with multiple possible labels, which made it challenging to determine a consistent gold standard for evaluation. As for the 20 Newsgroups dataset, despite its well-established reputation in classification tasks, it remains quite limited for topic labeling purposes. The number of topic modeling keywords and labels is restricted to only 20, which constrains the diversity and scalability of experiments.

## 5. Conclusion

In this paper, we proposed two methods for topic labeling that assign representative labels to topic modeling outputs using only the topic modeling results themselves and their corresponding representative graphs. Experiments conducted on two datasets demonstrated strong performance in terms of BERTScore and cosine similarity when compared to both a fine-tuned pre-trained model for topic labeling and existing benchmarks.




Figure 3: Cosine similarity scores across the application adopted method on the 20 Newsgroups dataset. The red dashed line represents the benchmark (ChatGPT, 10 words).

Table 5: Cosine similarity scores across the application adopted method on the 20 Newsgroups dataset.

Embedding	DSL	GEL	BART-TL-all	BART-TL-ng
No Embedding	-	-	0.310	0.295
all-MiniLM-L6-v2	0.550	0.619	-	-
all-MiniLM-L12-v2	0.578	0.624	-	-
all-mpnet-base-v2	0.512	0.587	-	-
gtr-t5-large	0.401	0.392	-	-
GIST-small-Embedding-v0	0.564	0.620	-	-
GIST-large-Embedding-v0	0.550	0.532	-	-
GIST-all-MiniLM-L6-v2	0.567	0.627	-	-

In terms of F1 score, DSL and GEL achieved 0.955 using the *all-MiniLM-L6-v2* and *gtr-T5-large* embeddings, respectively, surpassing the highest previously reported benchmark of 0.936. GEL, which incorporates graph representations to enrich the topic word sets with semantically informative context, produced excellent results on the 20 Newsgroups dataset. Using the *GIST-all-MiniLM-L6-v2* embedding model, it achieved a cosine similarity score that was only 0.028 lower than that of GPT-3.5, despite requiring significantly fewer computational resources.

These findings suggest that simple and interpretable methods, such as combining topic model-

ing results with graph-based representations, can outperform more complex pre-trained models in generating semantically meaningful topic labels, while remaining computationally efficient. Furthermore, results on the 20 Newsgroups dataset confirmed the robustness and generalizability of our approach across different domains.

Although the experiments yielded strong results, several limitations should be acknowledged regarding the size and nature of the datasets used. Future work could explore the creation of larger, high-quality labeled datasets specifically designed for the topic labeling task, especially considering the growing importance of topic labeling in extracting meaningful topics from large-scale text corpora.

Moreover, the promising results obtained through graph-based expansion of topic keywords open new research directions. Future studies could explore alternative knowledge graphs beyond ConceptNet, leverage graph structures to identify the most informative nodes, or apply advanced graph representation techniques, such as graph2vec, to derive more representative and semantically rich topic labels.

## 6. Acknowledgements

This publication has emanated from research conducted with the financial support of Taighde Éireann - Research Ireland under Grant number 18/CRT/6223.

## 7. References

- Rifki Adhitama, Retno Kusumaningrum, and Rahmat Gernowo. 2017. Topic labeling towards news document collection based on latent dirichlet allocation and ontology. In *2017 1st International Conference on Informatics and Computational Sciences (ICICoS)*, pages 247–252. IEEE.
- Mehdi Allahyari and Krys Kochut. 2015. Automatic topic labeling using ontology-based topic models. In *2015 IEEE 14th International Conference on Machine Learning and Applications (ICMLA)*, pages 259–264. IEEE.
- Areej Alokaili, Nikolaos Aletras, and Mark Stevenson. 2020. Automatic generation of topic labels. In *Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval*, pages 1965–1968.
- Shraey Bhatia, Jey Han Lau, and Timothy Baldwin. 2016. Automatic labelling of topics with neural embeddings. *arXiv preprint arXiv:1612.05340*.
- Ehsan Bitaraf, Maryam Jafarpour, Sina Shool, and Reza Saboori Amleshi. 2024. Unveiling medical insights: Advanced topic extraction from scientific articles. In *Digital Health and Informatics Innovations for Sustainable Health Care Systems*, pages 944–948. IOS Press.
- Akhil Chaudhary, Evangelos Milios, and Enayat Rajabi. 2024. Top2label: Explainable zero shot topic labelling using knowledge graphs. *Expert Systems with Applications*, 242:122676.
- Kenneth Ward Church. 2017. Word2vec. *Natural Language Engineering*, 23(1):155–162.
- Rob Churchill and Lisa Singh. 2022. The evolution of topic modeling. *ACM Computing Surveys*, 54(10s):1–35.
- Heidar Davoudi and Aijun An. 2015. Ontology-based topic labeling and quality prediction. In *International Symposium on Methodologies for Intelligent Systems*, pages 171–179. Springer.
- Fellyca Effendi, Muhammad Rizky Pribadi, Wijang Widhiarso, and Sherly Devyanes. 2023. Hybrid model for topic classification of english learning accounts on instagram using lda and svm. In *2023 10th International Conference on Electrical Engineering, Computer Science and Informatics (EECSI)*, pages 649–653. IEEE.
- Elena Flondor and Marc Frincu. 2023. Analysing android apps classification and categories validation by using latent dirichlet allocation. In *International Conference on Computational Collective Intelligence*, pages 282–297. Springer.
- Antoine Gourru, Julien Velcin, Mathieu Roche, Christophe Gravier, and Pascal Poncelet. 2018. United we stand: Using multiple strategies for topic labeling. In *International Conference on Applications of Natural Language to Information Systems*, pages 352–363. Springer.
- Gerhard Johann Hagerer, Martin Kirchhoff, Hannah Danner, Robert Pesch, Mainak Ghosh, Archishman Roy, Jiayi Zhao, and Georg Groh. 2021. Socialvistum: An interactive visualization toolkit for correlated neural topic models on social media opinion mining. *arXiv preprint arXiv:2110.10575*.
- Dongbin He, Minjuan Wang, Abdul Mateen Khatkhat, Li Zhang, and Wanlin Gao. 2019. Automatic labeling of topic models using graph-based ranking. *IEEE Access*, 7:131593–131608.
- Syeda Sumbul Hossain, Md Rezwan Ul-Hassan, and Shadikur Rahman. 2019. Polynomial topic distribution with topic modeling for generic labeling. In *International Conference on Advances in Computing and Data Sciences*, pages 409–419. Springer.
- Hugging Face. 2025. Massive text embedding benchmark (mteb) leaderboard. <https://huggingface.co/mteb>. Accessed: 2025-09-16.
- Hyon Hee Kim and Hey Young Rhee. 2019. An ontology-based labeling of influential topics using topic network analysis. *Journal of Information Processing Systems*, 15(5):1096–1107.
- Supriya Ashish Kinariwala and Sachin Deshmukh. 2021. Onto\_tml: Auto-labeling of topic models. *Journal of Integrated Science and Technology*, 9(2):85–91.
- Damir Korenčić, Strahil Ristov, and Jan Šnajder. 2015. Getting the agenda right: measuring media agenda using topic models. In *Proceedings of the 2015 Workshop on Topic Models: Post-Processing and Applications*, pages 61–66.
- Olzhas Kozbagarov, Rustam Mussabayev, and Nenad Mladenovic. 2021. A new sentence-based interpretative topic modeling and automatic topic labeling. *Symmetry*, 13(5):837.
- Jey Han Lau, Karl Grieser, David Newman, and Timothy Baldwin. 2011. Automatic labelling of topic models. In *Proceedings of the 49th annual meeting of the association for computational linguistics: human language technologies*, pages 1536–1545.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer.

2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.
- Chenliang Li, Shiqian Chen, and Yan Qi. 2019. Filtering and classifying relevant short text with a few seed words. *Data and Information Management*, 3(3):165–186.
- Chenliang Li, Jian Xing, Aixin Sun, and Zongyang Ma. 2016. Effective document labeling with very few seed words: A topic model approach. In *Proceedings of the 25th ACM international on conference on information and knowledge management*, pages 85–94.
- Pengyu Li, Christine Tseng, Yaxuan Zheng, Joyce A Chew, Longxiu Huang, Benjamin Jarman, and Deanna Needell. 2022. Guided semi-supervised non-negative matrix factorization on legal documents. *arXiv preprint arXiv:2201.13324*.
- Xiaofei Li and Guo Chen. Automatic literature review generation by integrating large and small models.
- Hugo Liu and Push Singh. 2004. Conceptnet—a practical commonsense reasoning tool-kit. *BT technology journal*, 22(4):211–226.
- Lin Liu and Lin Tang. 2019. A label distribution topic model for multi-label classification. In *Proceedings of the 4th International Conference on Intelligent Information Processing*, pages 52–57.
- Davide Magatti, Silvia Calegari, Davide Ciucci, and Fabio Stella. 2009. Automatic labeling of topics. In *2009 Ninth International Conference on Intelligent Systems Design and Applications*, pages 1227–1232. IEEE.
- Xian-Ling Mao, Zhao-Yan Ming, Zheng-Jun Zha, Tat-Seng Chua, Hongfei Yan, and Xiaoming Li. 2012. Automatic labeling hierarchical topics. In *Proceedings of the 21st ACM international conference on information and knowledge management*, pages 2383–2386.
- Xian-Ling Mao, Yixuan Xiao, Qiang Zhou, Jun Wang, and Heyan Huang. 2015. Ehlda: a supervised hierarchical topic model. In *China National Conference on Chinese Computational Linguistics*, pages 215–226. Springer.
- Xian-Ling Mao, Yi-jing Zhao, Qiang Zhou, Wen-Qing Yuan, Liner Yang, and He-Yan Huang. 2016. A novel fast framework for topic labeling based on similarity-preserved hashing. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 3339–3348.
- Salma Mekaoui, Ilham Chaker, Arsalane Zarghili, and Nikola S Nikolov. 2025. Systematic literature review of topic labeling. *IEEE Access*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. *Advances in neural information processing systems*, 26.
- Yida Mu, Chun Dong, Kalina Bontcheva, and Xingyi Song. 2024. Large language models offer an alternative to the traditional approach of topic modelling. *arXiv preprint arXiv:2403.16248*.
- Takafumi Nakanishi. 2025. Fusion of xai and llm: Automatic topic name extraction from text data using approximate inverse model explanations (aime). In *2025 19th International Conference on Semantic Computing (ICSC)*, pages 203–208. IEEE Computer Society.
- Jianmo Ni, Chen Qu, Jing Lu, Zhuyun Dai, Gustavo Hernandez Abrego, Ji Ma, Vincent Y Zhao, Yi Luan, Keith B Hall, Ming-Wei Chang, et al. 2021. Large dual encoders are generalizable retrievers. *arXiv preprint arXiv:2112.07899*.
- Diogo Nolasco and Jonice Oliveira. 2016. Detecting knowledge innovation through automatic topic labeling on scholar data. In *2016 49th Hawaii International Conference on System Sciences (HICSS)*, pages 358–367. IEEE.
- Takuro Ogawa and Ryosuke Saga. 2021. Text-based causality modeling with a conceptual label in a hierarchical topic structure using bayesian rose trees.
- OpenAI. 2023. Introducing chatgpt. <https://openai.com/blog/chatgpt>. Accessed: 2025-10-07.
- Cristian Popa and Traian Rebedea. 2021. Bart-tl: Weakly-supervised topic label generation. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1418–1425.
- Virginia Ramón-Ferrer, Carlos Badenes-Olmedo, and Oscar Corcho. 2023. Automatic topic label generation using conversational models. In *Proceedings of the 12th Knowledge Capture Conference 2023*, pages 17–24.
- Michael Reynolds. 2025. Research on the application of large language models in classification and indexing.
- Ngurah Agus Sanjaya, Mouhamadou Lamine Ba, Talel Abdessalem, and Stéphane Bressan. 2018. Harnessing truth discovery algorithms on the

- topic labelling problem. In *Proceedings of the 20th International Conference on Information Integration and Web-based Applications & Services*, pages 8–14.
- Thomas Scelsi, Alfonso Martinez Arranz, and Lea Frermann. 2021. Principled analysis of energy discourse across domains with thesaurus-based automatic topic labeling. In *Proceedings of the 19th Annual Workshop of the Australasian Language Technology Association*, pages 107–118.
- Kristian Slabbekoorn, Tomoya Noro, and Takehiro Tokuda. 2016. Ontology-assisted discovery of hierarchical topic clusters on the social web. *Journal of Web Engineering*, pages 361–396.
- Aivin V. Solatorio. 2024. [Gistembed: Guided in-sample selection of training negatives for text embedding fine-tuning](#). *arXiv preprint arXiv:2402.16829*.
- Wenling Tang, Xu Wu, Yuxiao Li, and Jin Xu. 2016. A topic label extraction method for the university bbs. In *2016 IEEE First International Conference on Data Science in Cyberspace (DSC)*, pages 678–682. IEEE.
- Paras Tiwari, Ashutosh Tripathi, Avaneesh Singh, and Sawan Rai. 2023a. Advanced hierarchical topic labeling for short text. *IEEE access*, 11:35158–35174.
- Paras Tiwari, Ashutosh Tripathi, Avaneesh Singh, and Sawan Rai. 2023b. Advanced hierarchical topic labeling for short text. *IEEE access*, 11:35158–35174.
- Eleni Vathi, Georgios Siolas, and Andreas Stafylopatis. 2017. Mining and categorizing interesting topics in twitter communities. *Journal of Intelligent & Fuzzy Systems*, 32(2):1265–1275.
- Yiming Wang, Jindong Zhang, Zhiyao Yang, Bing Wang, Jingyi Jin, and Yitong Liu. 2024. Improving extractive summarization with semantic enhancement through topic-injection based bert model. *Information Processing & Management*, 61(3):103677.
- Justin Wood, Patrick Tan, Wei Wang, and Corey Arnold. 2017. Source-lda: Enhancing probabilistic topic models using prior knowledge sources. In *2017 IEEE 33rd international conference on data engineering (ICDE)*, pages 411–422. IEEE.
- Daochen Zha and Chenliang Li. 2019. Multi-label dataless text classification with topic modeling. *Knowledge and Information Systems*, 61(1):137–160.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. *arXiv preprint arXiv:1904.09675*.
- Bing Zheng and Hua Li. 2022. Automatic labeling of sdn controller defect text based on neural topic model. In *2022 IEEE 22nd International Conference on Software Quality, Reliability, and Security Companion (QRS-C)*, pages 402–410. IEEE.
- Elaine Zosa, Lidia Pivovarova, Michele Boggia, and Sardana Ivanova. 2022. Multilingual topic labelling of news topics using ontological mapping. In *European Conference on Information Retrieval*, pages 248–256. Springer.