

Neural Topic Modeling with Large Language Models in the Loop

Anonymous ACL submission

Abstract

Topic modeling is a fundamental task in natural language processing, allowing the discovery of latent thematic structures in text corpora. While Large Language Models (LLMs) have demonstrated promising capabilities in topic discovery, their direct application to topic modeling suffers from issues such as incomplete topic coverage, misalignment of topics, and inefficiency. To address these limitations, we propose LLM-ITL, a novel LLM-in-the-loop framework that integrates LLMs with Neural Topic Models (NTMs). In LLM-ITL, global topics and document representations are learned through the NTM. Meanwhile, an LLM refines these topics using an Optimal Transport (OT)-based alignment objective, where the refinement is dynamically adjusted based on the LLM’s confidence in suggesting topical words for each set of input words. With the flexibility of being integrated into many existing NTMs, the proposed approach enhances the interpretability of topics while preserving the efficiency of NTMs in learning topics and document representations. Extensive experiments demonstrate that LLM-ITL helps NTMs significantly improve their topic interpretability while maintaining the quality of document representation.

1 Introduction

Topic modeling is an essential task in natural language processing that uncovers hidden thematic structures within text collections in an unsupervised way. The ability to automatically extract topics has proven to be invaluable across a range of disciplines, such as bioinformatics (Liu et al., 2016), marketing research (Reisenbichler and Reutterer, 2019), and information retrieval (Yi and Allan, 2009). Topic models are conventionally based on probabilistic frameworks that generate a set of interpretable global topics, each represented as a distribution over vocabulary terms as Latent Dirich-

let Allocation (LDA) (Blei et al., 2003). These topics are then used to represent individual documents as mixtures of topics, providing a structured and interpretable view of the corpus. Recently, research on Neural Topic Models (NTMs) (Zhao et al., 2021; Churchill and Singh, 2022; Abdelrazek et al., 2023; Wu et al., 2024a) has been popular, which uses deep neural networks to model document-topic distributions, enabling more expressive and flexible representations compared to their probabilistic counterparts.

While Large Language Models (LLMs) (OpenAI, 2022; Touvron et al., 2023a,b) have redefined the landscape of natural language processing, topic models can still be valuable tools for text analysis. Specifically, LLMs can provide a fine-grained understanding of a document; however, given a large collection of domain-specific documents, topic models are more suitable to obtain a clear global view of the topics in a more interpretable way with much less computational cost. Unsurprisingly, it has been a trending research direction to use LLMs to improve topic modeling (Rijcken et al., 2023; Wang et al., 2023; Pham et al., 2023; Mu et al., 2024; Doi et al., 2024; Chang et al., 2024). Despite the promising performance of these initial studies, most existing methods involve prompting LLMs to generate topics for each document in the corpus, which may lead to several limitations. As LLMs are asked to focus on a document individually, they may be unable to cover global topics across all the documents in the corpus (Doi et al., 2024), which is critical in topic modeling. Moreover, although LLMs excel at capturing local context, they usually struggle with long documents with multiple interrelated topics, which may evolve or shift throughout the text. With their limited window of focus, LLMs may miss key topics of a document that are necessary to fully understand its content. Finally, it is computationally expensive as LLMs have to do inference for every document in the corpus or sub-

sets; thus, existing methods usually scale poorly with large datasets.

To overcome the aforementioned limitations, we propose **LLM-ITL**, a framework that integrates LLMs into NTMs and enhances the overall quality and interpretability of the learned topics, while maintaining computational efficiency. Specifically, to enhance the interpretability of the topics learned by an NTM, we introduce an LLM-based refinement step. The representative words for each topic, as generated by the NTM, are provided to the LLM, which suggests improved words that better describe the semantic meaning of the topic. The refinement process is guided by a novel plug-in objective based on Optimal Transport (OT), which ensures that the topics learned by the NTM align closely with the LLM’s refinements. Additionally, to mitigate potential hallucinations from the LLM (i.e., the generation of inaccurate or irrelevant suggestions) for hard or incoherent topics, we introduce a confidence-weighted mechanism that adjusts the influence of the LLM’s suggestions based on their confidence scores. Finally, to balance learning from the input corpus by the NTM and refinement by the LLM, a warm-up phase is adopted to ensure that the NTM component learns corpus-relevant topics and document representations before the LLM’s refinement is applied. Our proposed LLM-ITL framework offers the following key contributions:

- **Boosting topic quality while maintaining document representation quality:** With the LLM’s refined topics and OT-based alignment, the topics generated are more interpretable and semantically coherent. At the same time, LLM-ITL ensures that the document-topic distributions, as learned by the NTM, remain high-quality and reflective of the document’s content.
- **Efficiency in using LLMs:** Unlike most existing LLM-based approaches that rely on document-level LLM analysis, LLM-ITL uses LLMs at the word level, significantly reducing computational overhead.
- **Flexibility:** LLM-ITL is a modular framework that can integrate with a variety of NTMs and LLMs, offering flexibility in model selection depending on the application and computational constraints.

2 Background

2.1 Problem Setup for Topic Modeling

Given a document collection $\mathcal{D} := \{d_1, \dots, d_N\}$, a topic model learns to discover a set of global topics $\mathcal{T} := \{t_1, \dots, t_K\}$, each of which is a distribution over the V vocabulary words $t \in \Delta^V$ (Δ denotes the probability simplex). Ideally, each topic represents a semantic concept that can be interpreted with its top-weighted words. At the document level, the topic model represents each document as a distribution over the K topics, i.e., $z \in \Delta^K$, which indicates the topic proportion of each topic within the document. The interpretability of topic models derives from both the corpus-level topics \mathcal{T} , and the document-level topical representation z for each document.

2.2 Neural Topic Models

A Neural Topic Model (NTM) is typically trained by modeling $p(z|x)$ and $p(x|z)$, where $x \in \mathbb{N}^V$ represents the Bag-of-Words (BOWs) of a document. NTMs, which employ deep neural networks for topic modeling, are commonly based on Variational Auto-Encoders (VAEs) (Kingma and Welling, 2013) and Amortized Variational Inference (AVI) (Rezende et al., 2014). For VAE-NTMs, $p(x|z)$ is modeled by a decoder network ϕ , i.e., $x := \phi(z)$. The posterior $p(z|x)$ is approximated by $q(z|x)$, which is modeled by an encoder network θ , i.e., $z := \theta(x)$. The training objective of VAE-NTMs is to maximize the Evidence Lower Bound (ELBO):

$$\max_{\theta, \phi} (\mathbb{E}_{q_{\theta}(z|x)} [\log p_{\phi}(x|z)] - \mathbb{KL}[q_{\theta}(z|x) \parallel p(z)]), \quad (1)$$

where the first term encourages the reconstruction of the document, and the second is the Kullback–Leibler divergence between the approximate posterior and the prior distribution. By implementing a single linear layer for the decoder $\phi \in \mathbb{R}^{V \times K}$, the k -th topic distribution t_k can be obtained by normalizing the k -th column of the decoder’s weight matrix:

$$t_k := \text{softmax}(\phi_{:,k})^T. \quad (2)$$

Its topic words w_k are obtained by taking the top-weight words of t_k , written as:

$$w_k := \mathcal{V}[f_{\text{topn}}(t_k, N)], \quad (3)$$

where $f_{\text{topn}}(a, N)$ defines a function that returns the indices of the top- N values of vector a ; \mathcal{V} denotes the vocabulary set of the corpus.

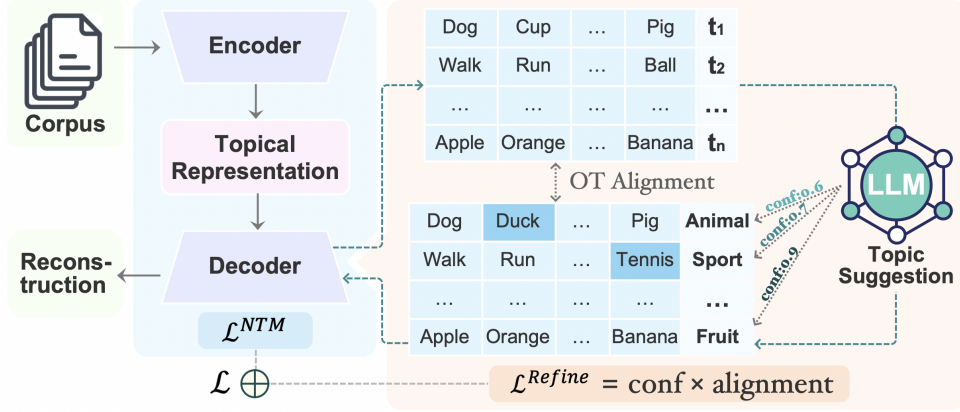


Figure 1: LLM-ITL overview. The topics and document representations are learned by the Neural Topic Model (NTM) component. After a warm-up stage, a Large Language Model (LLM) suggests better topic words for the learned topics (e.g., decoder) from the NTM. An Optimal Transport (OT)-based topic alignment objective is proposed to align the word distribution between the topics from the NTM and those suggested by the LLM. The alignment is further weighted by the confidence of the LLM in providing the suggestions. This confidence-weighted topic refinement objective is plugged into the standard training of an NTM as the overall objective of LLM-ITL.

2.3 Optimal Transport

Optimal Transport (OT) has been widely used for comparing probability distributions (Cuturi, 2013; Frogner et al., 2015; Seguy et al., 2017; Peyré et al., 2019). Let $\mu(\mathbf{x}, \mathbf{a}) := \sum_{i=1}^N a_i \delta_{x_i}$ and $\mu(\mathbf{y}, \mathbf{b}) := \sum_{j=1}^M b_j \delta_{y_j}$ be two discrete distributions, where $\mathbf{a} := [a_1, \dots, a_N]$ and $\mathbf{b} := [b_1, \dots, b_M]$ are the probability vectors; $\mathbf{x} := \{x_1, \dots, x_N\}$ and $\mathbf{y} := \{y_1, \dots, y_M\}$ are the supports of these two distributions. The OT distance between $\mu(\mathbf{x}, \mathbf{a})$ and $\mu(\mathbf{y}, \mathbf{b})$ is obtained by finding the optimal transport plan \mathbf{P}^* that transports the probability mass from $\mathbf{a} \in \Delta^N$ to $\mathbf{b} \in \Delta^M$, written as following:

$$d_{\text{OT}}(\mu(\mathbf{x}, \mathbf{a}), \mu(\mathbf{y}, \mathbf{b})) := \min_{\mathbf{P}} \sum_{i=1}^N \sum_{j=1}^M C_{i,j} P_{i,j}, \quad (4)$$

subject to $\sum_{j=1}^M P_{i,j} = a_i, \forall i = 1, \dots, N$ and $\sum_{i=1}^N P_{i,j} = b_j, \forall j = 1, \dots, M$. Here, $\mathbf{P} \in \mathbb{R}_{\geq 0}^{N \times M}$ is the transport plan, with entry $P_{i,j}$ indicating the amount of probability mass moving from a_i to b_j ; $\mathbf{C} \in \mathbb{R}_{\geq 0}^{N \times M}$ denotes the cost matrix, with entry $C_{i,j}$ specifying the distance between supports x_i and y_j . Various OT solvers (Flamary et al., 2021) have been proposed to compute the OT distance.

3 Method

In this work, we propose LLM-ITL, an LLM-in-the-loop framework that efficiently integrates the LLM with the training of NTMs, offering a more interpretable and comprehensive topic modeling pipeline. An overview of LLM-ITL is illustrated in Figure 1. LLM-ITL involves the following key

components: LLM-based topic suggestion, OT distance for topic alignment, and confidence-weighted topic refinement.

3.1 LLM-based Topic Suggestion

During the training of an NTM, it typically generates a set of topics, where each topic is represented by a distribution over words, with the highest-probability words forming the core “meaning” of the topic. While these words offer a rough semantic grouping, they often lack clarity or precision, leading to difficulties in interpretation. For instance, topics may contain words that are too general, too specific, or semantically ambiguous, making it hard for users to derive clear labels or understand the thematic focus of the topic.

To address this, LLM-ITL proposes to use LLMs to suggest better words or labels that more clearly express the same underlying concept. The LLM is prompted with the top words from each topic, and it generates two outputs: a topic label, which is a concise and interpretable summary of the topic; and a set of refined topic words, which better represent the underlying semantic concept of the topic. This process capitalizes on the LLM’s ability to grasp language nuances and provide more semantically rich suggestions for the topic. The LLM’s extensive pre-training on diverse and large datasets allows it to capture subtle relationships between words that may not be apparent in the purely statistical or neural-based methods employed by NTMs.

To obtain the topic label and refined words in a structured manner, *chain-of-thought (CoT) prompt-*

ing (Wei et al., 2022) is employed. CoT prompting encourages the LLM to reason step-by-step through the task, ensuring that it carefully considers the topic words before generating a label and refinement. The LLM’s output sequence s includes both the *topic label* and *refined words*, extracted as follows for each set of topic words:

$$s := \theta^{\text{llm}}(\text{Prompt}(w)),$$

$$\text{Topic label } w^l : (s_{\text{start of label}}, \dots, s_{\text{end of label}}),$$

$$\text{Refined words } w' : (s_{\text{start of words}}, \dots, s_{\text{end of words}}), \quad (5)$$

where w represents the original topic words; θ^{llm} denotes the LLM model; the topic label w^l and the refined words w' are extracted as subsequences from the LLM’s output s . Notably, the words in w' are further filtered based on the corpus vocabulary to ensure that no out-of-vocabulary (OOV) words are included. Our used prompt is illustrated in Figure A1. A study of prompt variants is provided in Appendix H.

3.2 OT-based Topic Alignment

A key innovation in LLM-ITL is the use of Optimal Transport (OT) distance to align the topic word distributions generated by the NTM with the refined topic word distributions provided by the LLM. OT is a mathematical framework that computes the “cost” of transforming one probability distribution into another, and has shown its effectiveness in measuring the alignment between two sets of words (Kusner et al., 2015; Yang et al., 2024).

Formally, given a set of original topic words $w := \{w_1, w_2, \dots, w_N\}$ with probability vector $t := [t_1, t_2, \dots, t_N]$ as obtained by Eq. 3 and Eq. 2, respectively; as well as refined topic words $w' := \{w'_1, w'_2, \dots, w'_M\}$ with probability vector $u := [u_1, u_2, \dots, u_M]$ ¹ from the LLM, the OT distance between these two word distributions can be formulated as:

$$d_{\text{OT}}(\mu(w, t), \mu(w', u)) = \min_P \sum_{i=1}^N \sum_{j=1}^M C_{i,j} P_{i,j}, \quad (6)$$

where $P \in \mathbb{R}_{\geq 0}^{N \times M}$ is the transport plan, with entry $P_{i,j}$ denoting the amount of probability mass transported from t_i to u_j ; $C \in \mathbb{R}_{\geq 0}^{N \times M}$ is the cost matrix, where $C_{i,j}$ represents the cost of transporting mass between word w_i and w'_j .

The cost matrix C is constructed using the cosine distance between pre-trained word embeddings $\mathcal{E}^w := \{e^{w_1}, e^{w_2}, \dots, e^{w_N}\}$ (for the original topic words) and $\mathcal{E}^{w'} := \{e^{w'_1}, e^{w'_2}, \dots, e^{w'_M}\}$

¹We assume each of the refined topic words from the LLM is equally important, thus u is a uniform probability vector.

(for the refined topic words). The cosine distance for each entry $C_{i,j}$ is computed as:

$$C_{i,j} := d_{\text{cos}}(e^{w_i}, e^{w'_j}), \quad (7)$$

where $d_{\text{cos}}(a, b)$ denotes the cosine distance between the embedding vectors a and b .

By minimizing this OT distance, the learned topic words from the NTM become aligned with the refined words suggested by the LLM, leading to more semantically coherent topics. This OT-based refinement loss is incorporated into the overall training objective, guiding the NTM to adjust its learned topics to match the LLM’s refined representations.

3.3 Confidence-Weighted Topic Refinement

LLMs, despite their powerful language capabilities, can sometimes produce hallucinated outputs—irrelevant or incorrect suggestions that do not align with the input data (Ji et al., 2023). To mitigate the impact of such hallucinations, LLM-ITL introduces a confidence-weighted refinement mechanism. The confidence mechanism assesses the reliability of the LLM’s refinements and adjusts their influence on the NTM’s training accordingly. This ensures that high-confidence refinements have a greater impact on the final topic representation, and vice versa.

We propose two methods for calculating topic labeling confidence of the LLM, considering whether the LLM is open-source or not: (1) Label token probability, applicable for open-source LLMs where the token probability of their generation is accessible; (2) Word intrusion confidence, available for both open and closed-source LLMs.

Label Token Probability This method computes the product of the token probabilities for the topic label generated by the LLM. It reflects the LLM’s certainty in generating the specific topic label:

$$\text{Conf}(w^l)^{\text{prob}} := \prod_{i=\text{sol}}^{\text{eol}} p(s_i | s_{<i}, c), \quad (8)$$

where “sol” and “eol” denote the indices of “start of label” and “end of label” token, respectively; $p(s_i | s_{<i}, c)$ denotes the token probability of the i -th token; c denotes the input context to the LLM.

Word Intrusion Confidence This method evaluates the proportion of irrelevant or “intruder” words removed by the LLM during suggestion. A topic label generated based on a higher rate of intruder removal indicates that it is harder for the LLM to identify the topic from the original topic words, leading to lower confidence:

$$\text{Conf}(\mathbf{w}^l)^{\text{intrusion}} := 1 - \frac{N^{\text{intruder}}}{N^{\mathbf{w}}}, \quad (9)$$

where $N^{\mathbf{w}}$ denotes the number of words in the given topic; N^{intruder} denotes the number of intruders identified by the LLM.

By incorporating the topic labeling confidence (Eq. 8 or Eq. 9) as a weight for the topic alignment loss, we adaptively adjust the impact of the LLM’s suggestion based on the confidence score. We write our confidence-weighted topic refinement objective as follows:

$$\min_{\phi} \sum_{k=1}^K \text{Conf}(\mathbf{w}_k^l) \text{dOT}(\mu(\mathbf{w}_k, \mathbf{t}_k), \mu(\mathbf{w}_k', \mathbf{u}_k)). \quad (10)$$

3.4 Integration with Warm-Up

While the LLM’s refinement can lead to more coherent topics, over-reliance on it may bias the topics toward the LLM’s knowledge rather than reflecting the global information of the input corpus, thereby harming the topical representation of documents. To address this issue, we propose integrating the NTM with the LLM after a warm-up stage. This ensures that the refinement process begins only after the NTM has learned a stable topic representation, allowing the model to capture the core structure of the corpus before fine-tuning the topics with LLM guidance. By integrating the topic refinement objective with the training of an NTM with warm-up, we obtain the overall objective of LLM-ITL:

$$\min_{\Theta} (\mathcal{L}^{\text{ntm}} + \gamma \cdot \mathbf{I}(t > T^{\text{refine}}) \cdot \mathcal{L}^{\text{refine}}), \quad (11)$$

where $\Theta := \{\theta, \phi\}$ denotes model parameters; \mathcal{L}^{ntm} and $\mathcal{L}^{\text{refine}}$ denote the NTM loss and refinement loss in Eq. 1 and Eq. 10, respectively; γ controls the strength of focusing on the LLM’s refinements; t and T^{refine} denote the current training step and the start of topic refinement, respectively; $\mathbf{I}(\cdot)$ denotes the indicator function. We provide the algorithm of LLM-ITL in Appendix B.

4 Related Work

Topic Models Classical topic models, such as Latent Dirichlet Allocation (LDA) (Blei et al., 2003) and its variants (Blei and Lafferty, 2006; Rosen-Zvi et al., 2012; Yan et al., 2013), are Bayesian probabilistic models with various generative assumptions about the documents. Neural Topic Models (NTMs) (Miao et al., 2017; Srivastava and Sutton, 2017; Card et al., 2017; Dieng et al., 2020; Bianchi et al., 2020; Zhao et al., 2020; Nguyen and

Luu, 2021; Shen et al., 2021; Duan et al., 2021; Li et al., 2022; Xu et al., 2023a,b; Yang et al., 2023; Miyamoto et al., 2023; Wu et al., 2024b; Chen et al., 2025a,b) use deep neural networks to learn topics and document representations. Clustering-based topic models (Sia et al., 2020; Grootendorst, 2022; Angelov and Inkpen, 2024) discover topics using clustering algorithms based on embeddings from pre-trained language models. Ours is not a specific topic model, but rather a general framework designed to integrate LLMs with a wide range of NTMs.

LLMs in Topic Modeling LLMs have been involved in topic modeling in various ways. Rijken et al. (2023) investigate the use of ChatGPT (OpenAI, 2022) to generate descriptions for topic words and found the effectiveness of these topic descriptions. Recent works leverage LLMs for topic model evaluation in different ways, such as applying LLMs for word intrusion or topic rating for topics (Rahimi et al., 2023; Stambach et al., 2023), or keyword generation for documents (Yang et al., 2024). LLM-based topic models have emerged (Wang et al., 2023; Pham et al., 2023; Mu et al., 2024; Doi et al., 2024), which prompt LLMs to generate topics and assign topics to documents. Unlike existing methods that focus on document-level analysis relying on LLMs, our approach prompts LLMs to analyze a group of topic words, which are then used to guide the training of NTMs. More recently, Chang et al. (2024) show that LLMs are effective at refining topic words, leading to improved topic coherence. However, their method refines the topic words of trained topic models in a post-hoc manner, whereas ours incorporates refinement as a regularization term during training. Our method is also loosely related to uncertainty estimation of LLMs, and we omit the discussion on this in Appendix C.

5 Experiments

5.1 Experimental Setup

Datasets We conduct experiments on four widely used and publicly available datasets for topic modeling, including 20NewsGroup (Lang, 1995) (20News), Reuters-21578 (Aletras and Stevenson, 2013) (R8), DBpedia (Auer et al., 2007) and AG-News (Zhang et al., 2015). Further details of these datasets are described in the Appendix D.1. The number of mined topics (i.e., K) is commonly regarded as a hyper-parameter for the dataset (Zhao

et al., 2020; Wu et al., 2024a). For datasets containing long documents, such as 20News and R8, we set the number of topics to 50. For datasets with short documents, such as DBpedia and AGNews, we set the number to 25. We also run experiments at different K values, which are reported in Appendix E.

Baselines The LLM-ITL framework is highly modular and can be seamlessly integrated with a wide range of NTMs. We integrated LLM-ITL with 8 commonly benchmarked NTMs, as listed in Table 1. We also evaluate other types of topic models, including probabilistic models such as LDA (Blei et al., 2003), cluster-based models like BERTopic (Grootendorst, 2022), and the latest LLM-based model, TopicGPT (Pham et al., 2023). Further details about these baselines and their settings are provided in Appendix D.2.

Settings of LLM-ITL We use LLAMA3-8B-Instruct (Dubey et al., 2024) in LLM-ITL for main experiments. For LLM generation, we use greedy decoding to enable deterministic output and set the maximum new generation tokens to 300. For OT computation, we use GloVe (Pennington et al., 2014) word embeddings pre-trained on Wikipedia to construct the OT cost matrix, and compute the OT distance using the POT (Flamary et al., 2021) package. For topic labeling confidence, we use label token probability in Eq. 8 for our main experiments. For hyper-parameters of LLM-ITL, we maintain consistent settings across all datasets and base models it integrates with: the topic refinement strength γ is set to 200; the number of words for the topic label is set to 2 when prompting the LLM; the warm-up steps are set as $T^{\text{refine}} = T^{\text{total}} - 50$, where T^{total} denotes the total training steps of the base model without integration. Each component and hyper-parameter of LLM-ITL is studied in the following sections. Each trial² of LLM-ITL in our experiment takes a few hours on a single 80GB A100 GPU.

Evaluation Metrics We evaluate both the topic quality and the document representation quality for topic models. For topic quality, we apply the widely used **topic coherence** metric, C_V (Röder et al., 2015). We report the average C_V values of all learned topics. For document representation

quality, we evaluate the alignment between a document’s true label and the top-weighted topic of its topical representation, known as **topic alignment** (Chuang et al., 2013; Pham et al., 2023). This is commonly evaluated using external clustering metrics, Purity and Normalized Mutual Information (NMI). Since Purity and NMI are considered equally important, fall within the same range (from 0 to 1), and are often reported together, we report their average as **PN**, serving as an overall indicator of topic alignment performance. Detailed results for Purity and NMI are provided in Appendix E.5. Intuitively, topic coherence (C_V) reflects how coherent the learned topic words are, while topic alignment (PN) indicates how well the model represents the documents through the learned topics. Further details on the calculation of these metrics are provided in Appendix D.3. In addition, we compute the performance improvement of LLM-ITL over its base model as: $\frac{S(\text{base+LLM-ITL}) - S(\text{base})}{S(\text{base})}$, where $S(\cdot)$ represents the evaluation metric. As for other topic model evaluation metrics, including topic diversity (TD) (Dieng et al., 2020) and overall topic quality (TQ) (Wang et al., 2022) are reported in Appendix E.3 and E.4, respectively.

5.2 Results

Topic Coherence and Alignment We show the performance of topic coherence and topic alignment for different models in Table 1. We summarize the following remarks based on the results: (1) For topic coherence, using LLM-ITL significantly improves performance across all cases, with minimum gains of +7.4% (NSTM as the base model on DBpedia) and maximum gains of +70.6% (ECRTM as the base model on 20News) over the base model. (2) In terms of topic alignment, integrating LLM-ITL demonstrates comparable performance to its base model, with changes ranging from -3.5% to +4.2%. (3) Moreover, for long-document corpora such as 20News and R8, applying LLM-ITL outperforms existing LLM-based topic models like TopicGPT in terms of topic alignment in most cases. This suggests that LLMs alone struggle to fully capture the topics of long documents, highlighting the advantages of integrating NTMs to enhance topical representation for document collections. Additional results with different settings and evaluation metrics are presented in Appendix E.

Balancing Topic Coherence and Alignment Over-reliance on the LLM’s refinement may intro-

²All experiments are conducted five times with different model random seeds throughout this work. The mean and standard deviation values of performance are reported.

Model	20News		R8		DBpedia		AGNews	
	C_V	PN	C_V	PN	C_V	PN	C_V	PN
LDA (Blei et al., 2003)	0.529 ± 0.006	0.489 ± 0.009	0.384 ± 0.006	0.700 ± 0.005	0.537 ± 0.010	0.762 ± 0.013	0.544 ± 0.013	0.596 ± 0.008
BERTopic (Grootendorst, 2022)	0.404 ± 0.005	0.342 ± 0.008	0.323 ± 0.010	0.697 ± 0.001	0.545 ± 0.015	0.720 ± 0.009	0.506 ± 0.016	0.450 ± 0.010
TopicGPT (Pham et al., 2023)	NA	0.363 ± 0.000	NA	0.410 ± 0.000	NA	0.706 ± 0.000	NA	0.634 ± 0.000
NVDM (Miao et al., 2017)	0.261 ± 0.008	0.141 ± 0.006	0.298 ± 0.010	0.359 ± 0.010	0.329 ± 0.017	0.175 ± 0.009	0.355 ± 0.010	0.252 ± 0.009
+ LLM-ITL	0.336 ± 0.014	0.138 ± 0.007	0.374 ± 0.016	0.360 ± 0.011	0.479 ± 0.025	0.170 ± 0.006	0.455 ± 0.016	0.248 ± 0.013
	↑ 28.7%	↓ 2.1%	↑ 25.5%	↑ 0.3%	↑ 45.6%	↓ 2.9%	↑ 28.2%	↓ 1.6%
PLDA (Srivastava and Sutton, 2017)	0.368 ± 0.011	0.272 ± 0.008	0.375 ± 0.004	0.525 ± 0.006	0.502 ± 0.014	0.664 ± 0.009	0.551 ± 0.012	0.481 ± 0.010
+ LLM-ITL	0.525 ± 0.010	0.272 ± 0.011	0.466 ± 0.011	0.522 ± 0.008	0.627 ± 0.020	0.660 ± 0.005	0.625 ± 0.017	0.486 ± 0.008
	↑ 42.7%	↑ 0.0%	↑ 24.3%	↓ 0.6%	↑ 24.9%	↓ 0.6%	↑ 13.4%	↑ 1.0%
SCHOLAR (Card et al., 2017)	0.479 ± 0.019	0.582 ± 0.010	0.386 ± 0.004	0.680 ± 0.013	0.608 ± 0.019	0.825 ± 0.015	0.576 ± 0.011	0.638 ± 0.003
+ LLM-ITL	0.591 ± 0.012	0.568 ± 0.010	0.439 ± 0.008	0.680 ± 0.012	0.678 ± 0.017	0.828 ± 0.013	0.655 ± 0.012	0.639 ± 0.002
	↑ 23.4%	↓ 2.4%	↑ 13.7%	↑ 0.0%	↑ 11.5%	↑ 0.4%	↑ 13.7%	↑ 0.2%
ETM (Dieng et al., 2020)	0.491 ± 0.006	0.404 ± 0.010	0.433 ± 0.006	0.679 ± 0.018	0.513 ± 0.004	0.762 ± 0.011	0.534 ± 0.013	0.568 ± 0.008
+ LLM-ITL	0.578 ± 0.008	0.398 ± 0.010	0.571 ± 0.010	0.686 ± 0.012	0.704 ± 0.015	0.742 ± 0.016	0.644 ± 0.017	0.569 ± 0.005
	↑ 17.7%	↓ 1.5%	↑ 31.9%	↑ 1.0%	↑ 37.2%	↓ 2.6%	↑ 20.6%	↑ 0.2%
NSTM (Zhao et al., 2020)	0.444 ± 0.015	0.373 ± 0.005	0.412 ± 0.004	0.665 ± 0.014	0.652 ± 0.005	0.767 ± 0.011	0.588 ± 0.022	0.593 ± 0.010
+ LLM-ITL	0.521 ± 0.015	0.360 ± 0.006	0.549 ± 0.011	0.673 ± 0.006	0.700 ± 0.013	0.752 ± 0.014	0.675 ± 0.012	0.585 ± 0.010
	↑ 17.3%	↓ 3.5%	↑ 33.3%	↑ 1.2%	↑ 7.4%	↓ 2.0%	↑ 14.8%	↓ 1.3%
CLNTM (Nguyen and Luu, 2021)	0.490 ± 0.014	0.575 ± 0.011	0.361 ± 0.008	0.691 ± 0.005	0.500 ± 0.015	0.683 ± 0.040	0.558 ± 0.028	0.607 ± 0.014
+ LLM-ITL	0.612 ± 0.010	0.576 ± 0.005	0.436 ± 0.008	0.691 ± 0.005	0.612 ± 0.019	0.684 ± 0.039	0.655 ± 0.020	0.594 ± 0.011
	↑ 24.9%	↑ 0.2%	↑ 20.8%	↑ 0.0%	↑ 22.4%	↑ 0.1%	↑ 17.4%	↓ 2.1%
WeTe (Wang et al., 2022)	0.495 ± 0.017	0.318 ± 0.009	0.529 ± 0.011	0.649 ± 0.002	0.546 ± 0.003	0.721 ± 0.012	0.560 ± 0.006	0.516 ± 0.004
+ LLM-ITL	0.583 ± 0.028	0.316 ± 0.008	0.599 ± 0.015	0.656 ± 0.005	0.613 ± 0.015	0.703 ± 0.012	0.623 ± 0.007	0.517 ± 0.004
	↑ 17.8%	↓ 0.6%	↑ 13.2%	↑ 1.1%	↑ 12.3%	↓ 2.5%	↑ 11.2%	↑ 0.2%
ECRTM (Wu et al., 2023)	0.323 ± 0.014	0.516 ± 0.005	0.308 ± 0.006	0.673 ± 0.008	0.581 ± 0.026	0.668 ± 0.202	0.438 ± 0.035	0.554 ± 0.019
+ LLM-ITL	0.551 ± 0.025	0.519 ± 0.002	0.364 ± 0.010	0.667 ± 0.007	0.684 ± 0.013	0.696 ± 0.142	0.505 ± 0.025	0.553 ± 0.020
	↑ 70.6%	↑ 0.6%	↑ 18.2%	↓ 0.9%	↑ 17.7%	↑ 4.2%	↑ 15.3%	↓ 0.2%

Table 1: Topic coherence (C_V) and topic alignment (PN). “NA” indicates the evaluation is not applicable. The performance improvement of LLM-ITL over its base model is computed.

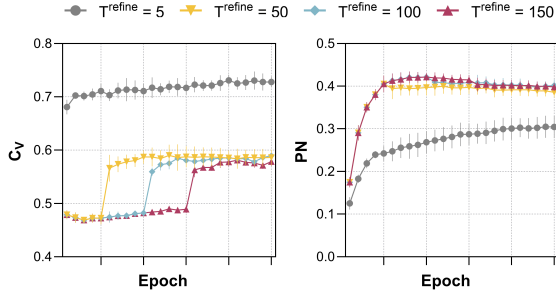


Figure 2: Learning curves of LLM-ITL (ETM as the base model) with different T^{refine} in terms of topic coherence (C_V) and topic alignment (PN) on 20News.

duce out-of-corpus information, thereby harming topical representation of documents for the corpus (i.e., topic alignment). Here, we demonstrate the effectiveness of using warm-up integration to balance topic coherence and alignment. From the learning curves illustrated in Figure 2, we observe that starting topic refinement earlier (e.g., $T^{\text{refine}} = 5$) can result in greater improvements in topic coherence but leads to poorer topic alignment performance. For larger values of T^{refine} , the performance in terms of both metrics becomes better balanced and comparable, indicating the effectiveness of warm-up integration and low sensitivity within a certain

range. For more hyper-parameter studies of LLM-ITL, see Appendix G.

Qualitative Analysis During the inference phase, LLM-ITL infers the topic proportion for a given document from the NTM component, and obtains the topic label from the LLM component, as shown in Table 2. We can observe that (1) Compared to topic models with top-words topics such as LDA, LLM-ITL provides more coherent topic words and offers topic labels, making the semantic meaning of the topics easier to identify. (2) Compared to the LLM-based topic model TopicGPT, LLM-ITL can obtain topic proportions as an indicator of the importance or relevance of topics to the document, offering more practical usage. For example, for TopicGPT, “Internet Culture” should be less relevant for the example document than “Software Development” if a good topic proportion is available.

Flexibility with different LLMs LLM-ITL is a framework compatible with most LLMs. Here, we examine the flexibility of LLM-ITL by integrating it with various LLMs. Apart from LLAMA3-8B-Instruct (Dubey et al., 2024), we implement LLM-ITL with the latest open-sourced LLMs, includ-

Document	Model	Topic	Proportion
Here’s a listing that I came across a while ago. This question seems to come up often enough that I figured this would be of interest. Note that the server “X Appeal” for DOS is available in demo form on the internet via anonymous ftp. This is one way of quickly checking out the feasibility of using your system as an X server. Enjoy! - Pete *** Many words omitted here *** 1280x960x1 (TT, SM194) color 320x200x4 color 640x200x2 color 640x480x4 color 320x480x8 Ethernet Card: Atari Card (Mega or VME bus) Riebl/Wacker (Mega or VME bus) — End Enclosure —	LDA	[1] window file program problem run work running machine time version	0.23
		[2] card driver monitor color video mode vga window screen problem	0.19
		[3] image data graphic package software program format tool file processing	0.16
		[4] mb mac mhz bit chip card scsi ram cpu memory	0.10
	
	TopicGPT	[1] Software Development – The document provides a list of X servers that can be used on non-UNIX networked machines.	NA
		[2] Internet Culture – The document mentions the availability of X servers for various operating systems and provides information on how to access the file via anonymous ftp.	NA
	LLM-ITL	[1] Software Development – application model version program software designed tool development implementation window	0.23
		[2] Computer Hardware – computer bios disk pc chip controller ram memory card apple	0.13
		[3] Service Support – support provide offer access available providing includes provides use feature	0.11
	

Table 2: Examples of different topic models’ output for a given document from 20News. Only the document’s top assigned/weighted (≥ 0.1) topics of its topical proportion/representation are listed in LDA and LLM-ITL (with ETM as the base model).




Figure 3: Learning curves of LLM-ITL (ETM as the base model) with different LLMs in terms of topic coherence (C_V) and topic alignment (PN) on 20News.

ing Mistral-7B-Instruct-v0.3 (Jiang et al., 2023), Phi-3-Mini-128K-Instruct (Abdin et al., 2024), Yi-1.5-9B-Chat (Young et al., 2024), Qwen1.5-32B-Chat (Bai et al., 2023) and LLAMA3-70B-Instruct (Dubey et al., 2024). As shown in Figure 3, LLM-ITL consistently improves topic coherence of its base model across different LLMs, and the improvement can be further enhanced when using larger LLMs such as LLAMA3-70B-Instruct.

Ablation Study for Confidence Here, we investigate the effectiveness of including the confidence scores during topic refinement. As illustrated in Figure 4, label token probability and word intrusion confidence consistently yield better performance in terms of PN. This suggests that by including proposed confidence during topic refinement, we reduce the impact of potential noisy suggestions from the LLM and achieve better topical representation

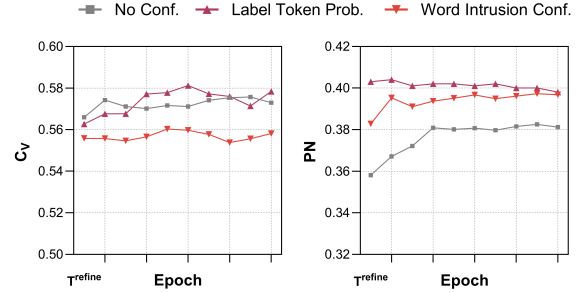


Figure 4: Ablation for confidence. Error bars are omitted for clarity in the figure.

for documents. For further studies on alternative LLM confidence measures, see Appendix F.

6 Conclusion

In this paper, we introduced LLM-ITL, a novel framework that integrates Large Language Models (LLMs) with Neural Topic Models (NTMs) to address the limitations of both traditional topic models and the direct use of LLMs for topic discovery. By incorporating a confidence-weighted Optimal Transport (OT)-based topic alignment, LLM-ITL improves the interpretability and coherence of topics while maintaining the quality of document representations. Our framework effectively leverages the strengths of both LLMs and NTMs, offering a flexible, scalable, and efficient solution for topic modeling. Extensive experiments on benchmark datasets demonstrate that applying LLM-ITL with NTMs significantly boosts topic interpretability while maintaining document representation quality.

7 Limitation

The proposed framework relies on the refinements generated by the LLM. Over-reliance on the LLM during the refinement process may introduce out-of-corpus information or bias the topics toward the LLM’s pre-training knowledge. This can negatively impact the alignment of learned topics with the input corpus, particularly in cases where the corpus diverges significantly from the LLM’s training data.

References

Aly Abdelrazek, Yomna Eid, Eman Gawish, Walaa Medhat, and Ahmed Hassan. 2023. Topic modeling algorithms and applications: A survey. *Information Systems*, 112:102131.

Marah Abidin, Sam Ade Jacobs, Ammar Ahmad Awan, Jyoti Aneja, Ahmed Awadallah, Hany Awadalla, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Harkirat Behl, et al. 2024. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv preprint arXiv:2404.14219*.

Nikolaos Aletras and Mark Stevenson. 2013. Evaluating topic coherence using distributional semantics. In *Proceedings of the 10th international conference on computational semantics (IWCS 2013)–Long Papers*, pages 13–22.

Dimo Angelov and Diana Inkpen. 2024. **Topic modeling: Contextual token embeddings are all you need**. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 13528–13539, Miami, Florida, USA. Association for Computational Linguistics.

Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. 2007. Dbpedia: A nucleus for a web of open data. In *international semantic web conference*, pages 722–735. Springer.

Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Sheng-guang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.

Federico Bianchi, Silvia Terragni, and Dirk Hovy. 2020. Pre-training is a hot topic: Contextualized document

embeddings improve topic coherence. *arXiv preprint arXiv:2004.03974*.

David M Blei and John D Lafferty. 2006. Dynamic topic models. In *Proceedings of the 23rd international conference on Machine learning*, pages 113–120.

David M Blei, Andrew Y Ng, and Michael I Jordan. 2003. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022.

Dallas Card, Chenhao Tan, and Noah A Smith. 2017. Neural models for documents with metadata. *arXiv preprint arXiv:1705.09296*.

Shuyu Chang, Rui Wang, Peng Ren, and Haiping Huang. 2024. Enhanced short text modeling: Leveraging large language models for topic refinement. *arXiv preprint arXiv:2403.17706*.

Chao Chen, Kai Liu, Ze Chen, Yi Gu, Yue Wu, Mingyuan Tao, Zhihang Fu, and Jieping Ye. 2024. Inside: LLMs’ internal states retain the power of hallucination detection. *arXiv preprint arXiv:2402.03744*.

Jiuhai Chen and Jonas Mueller. 2023. Quantifying uncertainty in answers from any language model and enhancing their trustworthiness.

Shangyu Chen, Xiaohao Yang, Pengfei Fang, Mehrtash Tafazzoli Harandi, Dinh Phung, and Jianfei Cai. 2025a. Stereographic projection for embedding hierarchical structures in hyperbolic space. In *International Conference on Pattern Recognition*, pages 307–321. Springer.

Shangyu Chen, He Zhao, Viet Huynh, Dinh Phung, and Jianfei Cai. 2025b. Neural topic model with distance awareness. In *International Conference on Pattern Recognition*, pages 337–352. Springer.

Jason Chuang, Sonal Gupta, Christopher Manning, and Jeffrey Heer. 2013. Topic model diagnostics: Assessing domain relevance via topical alignment. In *International conference on machine learning*, pages 612–620. PMLR.

Rob Churchill and Lisa Singh. 2022. The evolution of topic modeling. *ACM Computing Surveys*, 54(10s):1–35.

Marco Cuturi. 2013. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26.

Adji B Dieng, Francisco JR Ruiz, and David M Blei. 2020. Topic modeling in embedding spaces. *Transactions of the Association for Computational Linguistics*, 8:439–453.

Tomoki Doi, Masaru Isonuma, and Hitomi Yanaka. 2024. Topic modeling for short texts with large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 4: Student Research Workshop)*, pages 21–33.

713	Zhibin Duan, Dongsheng Wang, Bo Chen, Chaojie Wang, Wenchao Chen, Yewen Li, Jie Ren, and Mingyuan Zhou. 2021. Sawtooth factorial topic embeddings guided gamma belief network. In <i>International Conference on Machine Learning</i> , pages 2903–2913. PMLR.	769
714		770
715		771
716		772
717		773
718		774
719	Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. <i>arXiv preprint arXiv:2407.21783</i> .	775
720		776
721		777
722		
723		
724	Chrisantha Fernando, Dylan Banarse, Henryk Michalewski, Simon Osindero, and Tim Rocktäschel. 2023. Promptbreeder: Self-referential self-improvement via prompt evolution . <i>Preprint</i> , arXiv:2309.16797.	
725		
726		
727		
728		
729	Rémi Flamary, Nicolas Courty, Alexandre Gramfort, Mokhtar Z Alaya, Aurélie Boisbunon, Stanislas Chambon, Laetitia Chapel, Adrien Corenflos, Kilian Fatras, Nemo Fournier, et al. 2021. Pot: Python optimal transport. <i>Journal of Machine Learning Research</i> , 22(78):1–8.	778
730		779
731		780
732		781
733		
734		
735	Charlie Frogner, Chiyuan Zhang, Hossein Mobahi, Mauricio Araya, and Tomaso A Poggio. 2015. Learning with a wasserstein loss. <i>Advances in neural information processing systems</i> , 28.	782
736		783
737		784
738		785
739	Xiang Gao, Jiaxin Zhang, Lalla Mouatadid, and Kamalika Das. 2024. Spuq: Perturbation-based uncertainty quantification for large language models. <i>arXiv preprint arXiv:2403.02509</i> .	786
740		787
741		788
742		
743	Jiahui Geng, Fengyu Cai, Yuxia Wang, Heinz Koepl, Preslav Nakov, and Iryna Gurevych. 2024. A survey of confidence estimation and calibration in large language models. In <i>Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)</i> , pages 6577–6595.	789
744		790
745		791
746		792
747		793
748		
749		
750		
751	Maarten Grootendorst. 2022. Bertopic: Neural topic modeling with a class-based tf-idf procedure. <i>arXiv preprint arXiv:2203.05794</i> .	794
752		795
753		796
754	Bairu Hou, Yujian Liu, Kaizhi Qian, Jacob Andreas, Shiyu Chang, and Yang Zhang. 2023. Decomposing uncertainty for large language models through input clarification ensembling. <i>arXiv preprint arXiv:2311.08718</i> .	797
755		
756		
757		
758		
759	Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea Madotto, and Pascale Fung. 2023. Survey of hallucination in natural language generation. <i>ACM Computing Surveys</i> , 55(12):1–38.	798
760		799
761		800
762		801
763		
764	Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. <i>arXiv preprint arXiv:2310.06825</i> .	802
765		803
766		804
767		805
768		
	Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, et al. 2022. Language models (mostly) know what they know. <i>arXiv preprint arXiv:2207.05221</i> .	806
		807
		808
		809
	Diederik P Kingma and Max Welling. 2013. Auto-encoding variational bayes. <i>arXiv preprint arXiv:1312.6114</i> .	810
		811
		812
		813
		814
	Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. 2023. Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation. <i>arXiv preprint arXiv:2302.09664</i> .	815
		816
		817
		818
	Matt Kusner, Yu Sun, Nicholas Kolkin, and Kilian Weinberger. 2015. From word embeddings to document distances. In <i>International conference on machine learning</i> , pages 957–966. PMLR.	819
		820
		821
	Ken Lang. 1995. Newsweeder: Learning to filter netnews. In <i>Machine Learning Proceedings 1995</i> , pages 331–339. Elsevier.	
	Yewen Li, Chaojie Wang, Zhibin Duan, Dongsheng Wang, Bo Chen, Bo An, and Mingyuan Zhou. 2022. Alleviating "posterior collapse" in deep topic models via policy gradient. <i>Advances in Neural Information Processing Systems</i> , 35:22562–22575.	
	Zhen Lin, Shubhendu Trivedi, and Jimeng Sun. 2023. Generating with confidence: Uncertainty quantification for black-box large language models. <i>arXiv preprint arXiv:2305.19187</i> .	
	Lin Liu, Lin Tang, Wen Dong, Shaowen Yao, and Wei Zhou. 2016. An overview of topic modeling and its current applications in bioinformatics. <i>SpringerPlus</i> , 5(1):1–22.	
	Potsawee Manakul, Adian Liusie, and Mark JF Gales. 2023. Selfcheckgpt: Zero-resource black-box hallucination detection for generative large language models. <i>arXiv preprint arXiv:2303.08896</i> .	
	Yishu Miao, Edward Grefenstette, and Phil Blunsom. 2017. Discovering discrete latent topics with neural variational inference. In <i>International Conference on Machine Learning</i> , pages 2410–2419. PMLR.	
	Nozomu Miyamoto, Masaru Isonuma, Sho Takase, Junichiro Mori, and Ichiro Sakata. 2023. Dynamic structured neural topic model with self-attention mechanism. In <i>Findings of the Association for Computational Linguistics: ACL 2023</i> , pages 5916–5930.	
	Yida Mu, Chun Dong, Kalina Bontcheva, and Xingyi Song. 2024. Large language models offer an alternative to the traditional approach of topic modelling. <i>arXiv preprint arXiv:2403.16248</i> .	
	Thong Nguyen and Anh Tuan Luu. 2021. Contrastive learning for neural topic model. <i>Advances in neural information processing systems</i> , 34:11974–11986.	

- OpenAI. 2022. Chatgpt: Optimizing language models for dialogue. <http://web.archive.org/web/20230109000707/https://openai.com/blog/chatgpt/>.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Gabriel Peyré, Marco Cuturi, et al. 2019. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607.
- Chau Minh Pham, Alexander Hoyle, Simeng Sun, and Mohit Iyyer. 2023. Topicgpt: A prompt-based topic modeling framework. *arXiv preprint arXiv:2311.01449*.
- Hamed Rahimi, Jacob Louis Hoover, David Mimno, Hubert Naacke, Camelia Constantin, and Bernd Amann. 2023. Contextualized topic coherence metrics. *arXiv preprint arXiv:2305.14587*.
- Martin Reisenbichler and Thomas Reutterer. 2019. Topic modeling in marketing: recent advances and research opportunities. *Journal of Business Economics*, 89(3):327–356.
- Jie Ren, Jiaming Luo, Yao Zhao, Kundan Krishna, Mohammad Saleh, Balaji Lakshminarayanan, and Peter J Liu. 2022. Out-of-distribution detection and selective generation for conditional language models. In *The Eleventh International Conference on Learning Representations*.
- Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. 2014. Stochastic backpropagation and approximate inference in deep generative models. In *International conference on machine learning*, pages 1278–1286. PMLR.
- Emil Rijcken, Floortje Scheepers, Kalliopi Zervanou, Marco Spruit, Pablo Mosteiro, and Uzay Kaymak. 2023. Towards interpreting topic models with chatgpt. In *The 20th World Congress of the International Fuzzy Systems Association*.
- Michael Röder, Andreas Both, and Alexander Hinneburg. 2015. Exploring the space of topic coherence measures. In *Proceedings of the eighth ACM international conference on Web search and data mining*, pages 399–408.
- Michal Rosen-Zvi, Thomas Griffiths, Mark Steyvers, and Padhraic Smyth. 2012. The author-topic model for authors and documents. *arXiv preprint arXiv:1207.4169*.
- Vivien Seguy, Bharath Bhushan Damodaran, Rémi Flamary, Nicolas Courty, Antoine Rolet, and Mathieu Blondel. 2017. Large-scale optimal transport and mapping estimation. *arXiv preprint arXiv:1711.02283*.
- Dazhong Shen, Chuan Qin, Chao Wang, Zheng Dong, Hengshu Zhu, and Hui Xiong. 2021. Topic modeling revisited: A document graph-based neural network perspective. *Advances in neural information processing systems*, 34:14681–14693.
- Suzanna Sia, Ayush Dalmia, and Sabrina J Mielke. 2020. Tired of topic models? clusters of pretrained word embeddings make for fast and good topics too! *arXiv preprint arXiv:2004.14914*.
- Akash Srivastava and Charles Sutton. 2017. Autoencoding variational inference for topic models. *arXiv preprint arXiv:1703.01488*.
- Dominik Stambach, Vilém Zouhar, Alexander Hoyle, Mrinmaya Sachan, and Elliott Ash. 2023. Re-visiting automated topic model evaluation with large language models. *arXiv preprint arXiv:2305.12152*.
- Katherine Tian, Eric Mitchell, Allan Zhou, Archit Sharma, Rafael Rafailov, Huaxiu Yao, Chelsea Finn, and Christopher D Manning. 2023. Just ask for calibration: Strategies for eliciting calibrated confidence scores from language models fine-tuned with human feedback. *arXiv preprint arXiv:2305.14975*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023a. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023b. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Dennis Ulmer, Martin Gubri, Hwaran Lee, Sangdoo Yun, and Seong Joon Oh. 2024. Calibrating large language models using their generations only. *arXiv preprint arXiv:2403.05973*.
- Dongsheng Wang, Dandan Guo, He Zhao, Huangjie Zheng, Korawat Tanwisuth, Bo Chen, and Mingyuan Zhou. 2022. Representing mixtures of word embeddings with mixtures of topic embeddings. *arXiv preprint arXiv:2203.01570*.
- Han Wang, Nirmalendu Prakash, Nguyen Khoi Hoang, Ming Shan Hee, Usman Naseem, and Roy Ka-Wei Lee. 2023. Prompting large language models for topic modeling. In *2023 IEEE International Conference on Big Data (BigData)*, pages 1236–1241. IEEE.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.

- Xiaobao Wu, Xinshuai Dong, Thong Thanh Nguyen, and Anh Tuan Luu. 2023. Effective neural topic modeling with embedding clustering regularization. In *International Conference on Machine Learning*, pages 37335–37357. PMLR.
- Xiaobao Wu, Thong Nguyen, and Anh Tuan Luu. 2024a. A survey on neural topic models: methods, applications, and challenges. *Artificial Intelligence Review*, 57(2):18.
- Xiaobao Wu, Fengjun Pan, Thong Nguyen, Yichao Feng, Chaoqun Liu, Cong-Duy Nguyen, and Anh Tuan Luu. 2024b. On the affinity, rationality, and diversity of hierarchical topic modeling. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 19261–19269.
- Miao Xiong, Zhiyuan Hu, Xinyang Lu, Yifei Li, Jie Fu, Junxian He, and Bryan Hooi. 2023. Can llms express their uncertainty? an empirical evaluation of confidence elicitation in llms. *arXiv preprint arXiv:2306.13063*.
- Weijie Xu, Wenxiang Hu, Fanyou Wu, and Srinivasan Sengamedu. 2023a. [DeTiME: Diffusion-enhanced topic modeling using encoder-decoder based LLM](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 9040–9057, Singapore. Association for Computational Linguistics.
- Weijie Xu, Xiaoyu Jiang, Srinivasan Sengamedu Hanumantha Rao, Francis Iannacci, and Jinjin Zhao. 2023b. [vONTSS: vMF based semi-supervised neural topic modeling with optimal transport](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 4433–4457, Toronto, Canada. Association for Computational Linguistics.
- Xiaohui Yan, Jiafeng Guo, Yanyan Lan, and Xueqi Cheng. 2013. A biterm topic model for short texts. In *Proceedings of the 22nd international conference on World Wide Web*, pages 1445–1456.
- Xiaohao Yang, He Zhao, Dinh Phung, Wray Buntine, and Lan Du. 2024. Llm reading tea leaves: Automatically evaluating topic models with large language models. *arXiv preprint arXiv:2406.09008*.
- Xiaohao Yang, He Zhao, Dinh Phung, and Lan Du. 2023. Towards generalising neural topical representations. *arXiv preprint arXiv:2307.12564*.
- Xing Yi and James Allan. 2009. A comparative study of utilizing topic models for information retrieval. In *Advances in Information Retrieval: 31th European Conference on IR Research, ECIR 2009, Toulouse, France, April 6-9, 2009. Proceedings 31*, pages 29–41. Springer.
- Alex Young, Bei Chen, Chao Li, Chengen Huang, Ge Zhang, Guanwei Zhang, Heng Li, Jiangcheng Zhu, Jianqun Chen, Jing Chang, et al. 2024. Yi: Open foundation models by 01. ai. *arXiv preprint arXiv:2403.04652*.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. *Advances in neural information processing systems*, 28.
- He Zhao, Dinh Phung, Viet Huynh, Yuan Jin, Lan Du, and Wray Buntine. 2021. Topic modelling meets deep neural networks: A survey. *arXiv preprint arXiv:2103.00498*.
- He Zhao, Dinh Phung, Viet Huynh, Trung Le, and Wray Buntine. 2020. Neural topic model via optimal transport. *arXiv preprint arXiv:2008.13537*.

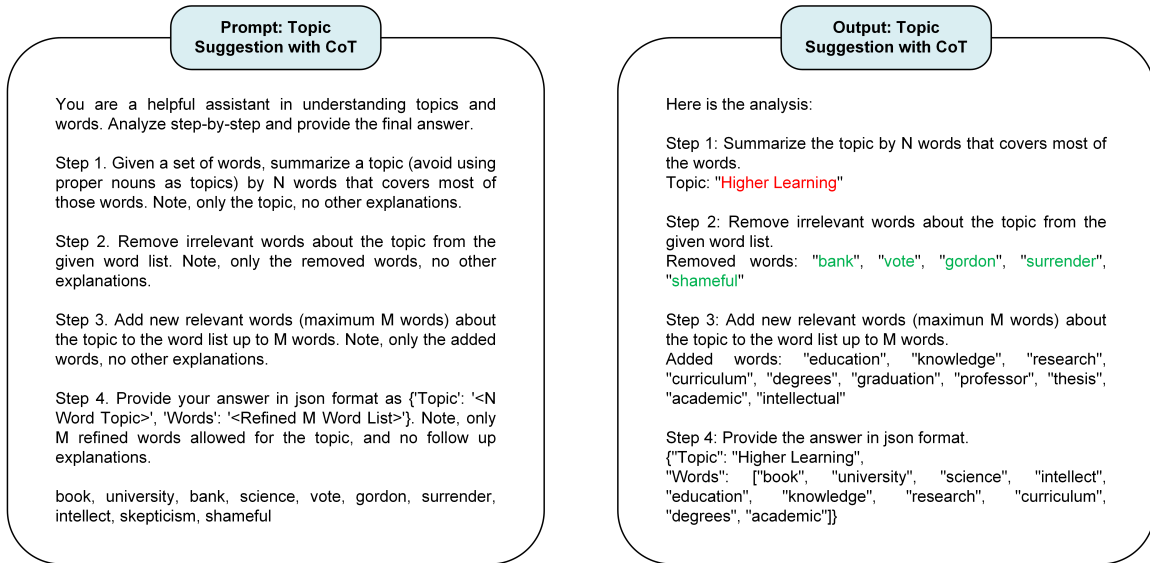


Figure A1: Prompt and output of topic suggestion with CoT. We take the product of token probabilities of topic label (e.g., words in **red** color) as the label token probability. We take the proportion of intruders (e.g., words in **green** color) as the word intrusion confidence. $N = 2$ and $M = 10$ in this example.

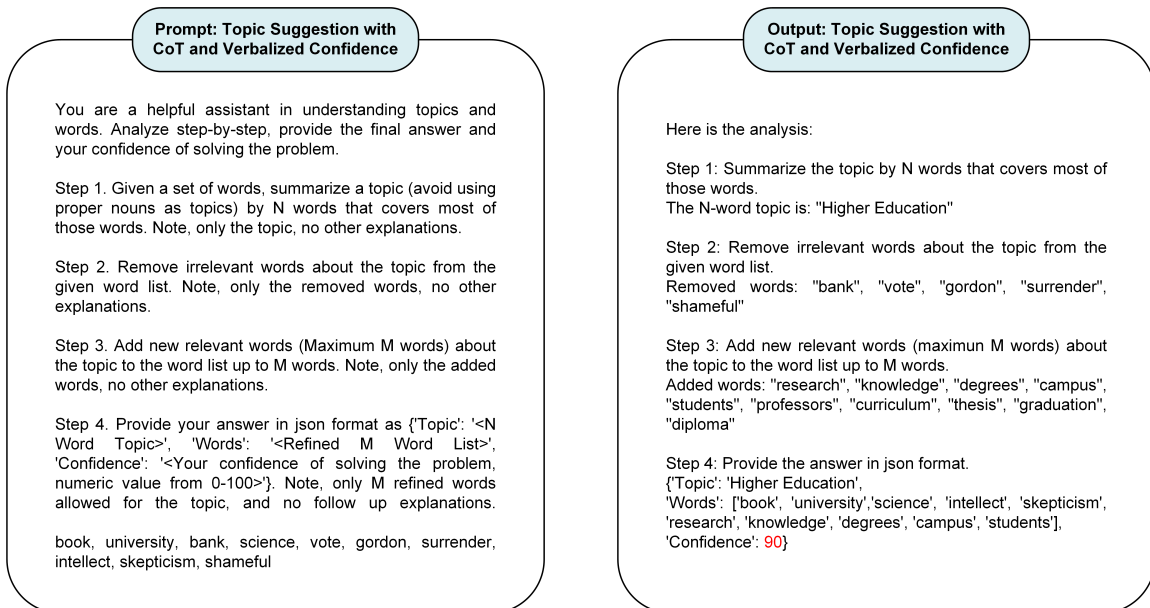


Figure A2: Prompt and output of topic suggestion with CoT and verbalized confidence. The number in **red** color in the LLM's output represents the verbalized confidence. $N = 2$ and $M = 10$ in this example.

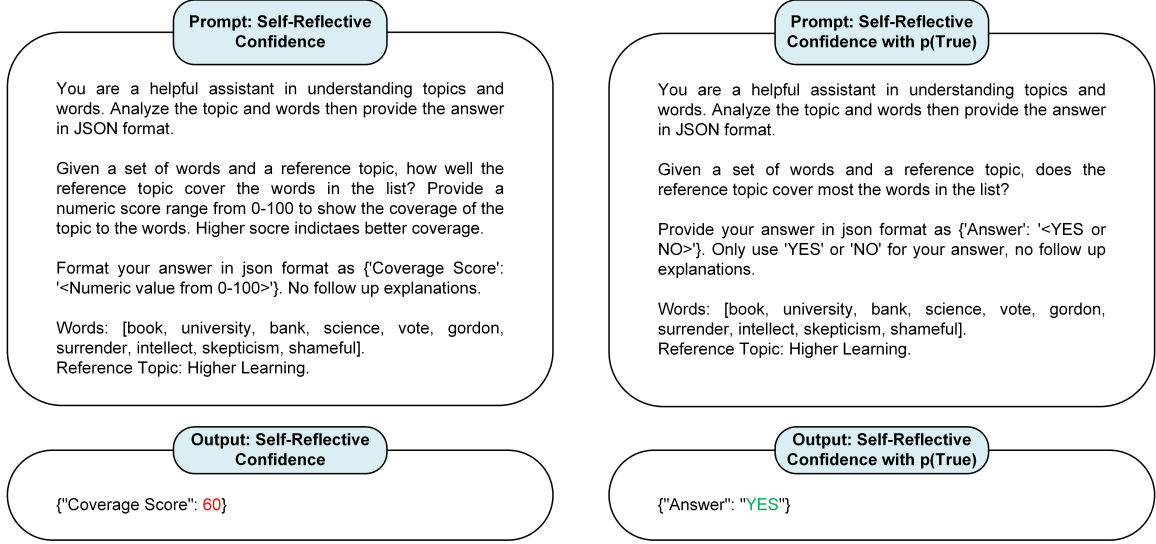


Figure A3: Prompt and output of self-reflective confidence and $p(\text{True})$. **Left:** self-reflective confidence where the number in **red** color represents the confidence. **Right:** $p(\text{True})$ confidence where the token probability of “YES” (p) in **green** color or “NO” ($1-p$) is used as confidence.

B Algorithm

Algorithm 1: Algorithm for LLM-ITL

Input: Train documents; An LLM; Pre-trained word embeddings; Hyper-parameters T^{refine} , γ ; Training iteration I ; Number of topics K .

Initialize: Initialize the parameters θ , ϕ of the NTM.

*/*Warm-up*/*

for $i = 1 : T^{\text{refine}}$ **do**

 Compute NTM loss by Eq. 1;

 Compute gradients w.r.t θ and ϕ ;

 Update θ and ϕ based on the gradients;

end

*/*Topic Refinement*/*

for $i = T^{\text{refine}} : I$ **do**

for $k = 1 : K$ **do**

 Obtain topic distribution t_k by Eq. 2;

 Obtain topic words w_k by Eq. 3;

 Obtain refined words w_k from the LLM by Eq. 5;

 Construct OT cost matrix by Eq. 7;

 Compute OT distance by Eq. 6;

if *Open-Source LLM* **then**

 Compute topic labeling confidence by Eq. 8

end

else

 Compute topic labeling confidence by Eq. 9

end

end

 Compute $\mathcal{L}^{\text{refine}}$ by Eq. 10;

 Compute \mathcal{L}^{ntm} by Eq. 1;

 Compute overall loss by $\mathcal{L}^{\text{ntm}} + \gamma \cdot \mathcal{L}^{\text{refine}}$;

 Compute gradients w.r.t θ and ϕ ;

 Update θ and ϕ based on the gradients;

end

Output: Trained NTM with θ , ϕ .

C Related work: LLM Uncertainty Estimation

Uncertainty estimation for LLMs (Geng et al., 2024) is emerging with the rapid usage of LLMs and their risk of hallucination (Ji et al., 2023). Sequence probability (Ren et al., 2022) leverages token probabilities to measure answer confidence. Verbalized confidence (Tian et al., 2023; Xiong et al., 2023) utilizes the LLM’s own capability to evaluate its answer uncertainty. Consistency-based (Lin et al., 2023; Manakul et al., 2023) approaches sample multiple outputs from the LLM and measure answer consistency as uncertainty. Entropy-based (Kuhn et al., 2023; Hou et al., 2023) approaches estimate the output space from multiple LLM outputs and compute entropy as uncertainty for the answer. Hybrid frameworks (Chen and Mueller, 2023; Gao et al., 2024) combine different approaches for a comprehensive estimation. Internal states (Chen et al., 2024) are another useful source for LLM uncertainty quantification. Unlike those works that estimate uncertainty for LLMs in general natural language generation tasks, ours focuses on task-specific uncertainty of the LLM in suggesting topic words.

D Detailed Experimental Settings

D.1 Details of Dataset

Dataset	# Docs Train	# Docs Test	Voc Size	Avg. Doc Length	# Labels
20News	11778	2944	13925	150	20
R8	5485	2189	5338	102	8
DBpedia	15598	3899	8550	51	14
AGNews	16000	4000	8389	38	4

Table D1: Statistics of the Datasets

We conduct experiments on 20News³, R8⁴, DBpedia⁵, and AGNews⁶. For DBpedia and AGNews, we randomly sample a subset of 20,000 documents. We retain the original text documents for models that accept text as input, and preprocess the documents into Bag-of-Words (BOW) format for models that are trained on BOWs. We convert the documents into BOW vectors through the following steps: First, we clean the documents by removing special characters and stop words, followed by

³https://huggingface.co/datasets/SetFit/20_newsgroups

⁴<https://huggingface.co/datasets/yangwang825/reuters-21578>

⁵https://huggingface.co/datasets/fancyzhx/dbpedia_14

⁶https://huggingface.co/datasets/fancyzhx/ag_news

tokenization. Next, we build the vocabulary by including words with a document frequency greater than five and less than 80% of the total documents. Since we use the pre-trained word embeddings of GloVe (Pennington et al., 2014), we further filter the vocabulary by retaining only the words that are in the GloVe vocabulary. Finally, we transform the documents into BOWs based on the filtered vocabulary set. The statistics of the preprocessed datasets are summarized in Table D1.

D.2 Details of Baselines

We run the following topic models in our experiments, including Latent Dirichlet Allocation (LDA) (Blei et al., 2003), the most popular probabilistic topic model that generates documents by mixtures of topics; Neural Variational Document Model (NVDM) (Miao et al., 2017), a pioneering NTM based on the VAE framework; LDA with Products of Experts (PLDA) (Srivastava and Sutton, 2017), an NTM that uses a product of experts instead of the mixture model in LDA; Neural Topic Model with Covariates, Supervision, and Sparsity (SCHOLAR) (Card et al., 2017), an NTM that leverages extra information from metadata; Embedded Topic Model (ETM) (Dieng et al., 2020), which involves word and topic embeddings in the generative process of documents; Neural Sinkhorn Topic Model (NSTM) (Zhao et al., 2020), an NTM based on an optimal transport framework; Contrastive Learning Neural Topic Model (CLNTM) (Nguyen and Luu, 2021), an NTM that is based on the contrastive learning framework; BERTopic (Grootendorst, 2022), a recent clustering-based topic model that utilizes embeddings from pre-trained language models; WeTe (Wang et al., 2022), which represents a mixture of word embeddings as a mixture of topic embeddings; Embedding Clustering Regularization Topic Model (ECRTM) (Wu et al., 2023), an NTM with embedding clustering regularization; TopicGPT (Pham et al., 2023), a latest LLM-based topic model that leverages an LLM for topic generation and assignment.

As for the implementations of baseline models, we use Mallet⁷ for LDA with Gibbs sampling, and the original implementations for the other models. For NTMs including NVDM, PLDA, SCHOLAR, ETM, NSTM, CLNTM, WeTe and ECRTM, we tune their hyper-parameters for the datasets; For BERTopic, we fine-tune the topic representations

⁷https://radimrehurek.com/gensim_3.8.3/models/wrappers/ldamallet.html

after the topics are learned, as suggested by their implementation⁸. For TopicGPT, we use GPT-4 for topic generation and GPT-3.5 for topic assignment, randomly sampling 600 documents from the training set for each dataset, as suggested by their paper. We run all models except TopicGPT five times in each experiment and report the mean and standard deviation of their performance. For TopicGPT, we run it once for each experiment using a temperature value of zero to enable deterministic output, following the setting of their paper.

D.3 Details of Evaluation Metrics

For topic evaluation, we apply the commonly-used topic coherence metric, C_V (Röder et al., 2015), which evaluates topic coherence based on the co-occurrence of the topic’s top words in a reference corpus. Following standard evaluation protocol, we use Wikipedia as the reference corpus and consider the top 10 words of each topic, with implementation done using the Palmetto package⁹ (Röder et al., 2015). We report the average C_V score of all learned topics as the overall topic coherence performance. For documents’ topical representation (i.e., topic proportion) evaluation, a common practice is to compare the document clusters formed by topic proportions with those formed by the documents’ true labels, known as topic alignment. Following previous works (Chuang et al., 2013; Pham et al., 2023), we assign each test document to a cluster based on the top-weighted topic of its topical representation, and compute Purity and Normalized Mutual Information (NMI) based on the documents’ cluster assignments and their true labels. As Purity and NMI are often reported together and within the same range, we report the average score of both metrics as PN. For all evaluations, we use the model state at the end of the training iteration to compute all evaluation metrics.

E More Results

E.1 Topic Coherence at Different K

We illustrate the topic coherence performance (C_V) of topic models across different numbers of topics (i.e., K) in Table E1. Overall, using LLM-ITL significantly boosts topic coherence across all cases in our experiments, with minimum gains of +10.4% (NSTM as the base model on DBpedia at $K = 50$) and maximum gains of +86.3% (ECRTM as the

base model on 20News at $K = 25$) over the base model.

E.2 Topic Alignment at Different K

We illustrate the topic alignment performance (PN) of topic models across different numbers of topics (i.e., K) in Table E2. Overall, using LLM-ITL inherits the topic alignment performance of its base model, with slight changes ranging from -5.2% (NSTM as the base model on 20News at $K = 100$) to +5.2% (PLDA as the base model on 20News at $K = 75$).

E.3 Topic Diversity (TD)

We evaluate the performance of topic diversity (TD) defined by Dieng et al. (2020), which evaluates the distinctiveness of the top-words within the learned topics. The results across different numbers of topics (i.e., K) is illustrated in Table E3. Overall, applying LLM-ITL results in changes to TD ranging from -34.6% (ECRTM as the base model on 20News at $K = 100$) to +55.5% (NSTM as the base model on 20News at $K = 100$) compared to its base model. A further analysis of overall topic quality, which considers both topic coherence and diversity, is described in the following section.

E.4 Overall Topic Quality (TQ)

Wang et al. (2022) define topic quality (TQ) as $TQ := TC \times TD$, serving as an overall indicator of topic quality that considers both topic coherence and diversity. We illustrate the topic quality performance of topic models across different numbers of topics (i.e., K) in Table E4. Overall, LLM-ITL boosts overall topic quality in all cases in our experiments, with a minimum improvement of 3.7% (SCHOLAR as the base model on R8 at $K = 100$) and a maximum improvement of 90.9% (CLNTM as the base model on DBpedia at $K = 100$) compared to its base mod

E.5 Purity & NMI

As we report PN (the mean of Purity and NMI) as an overall topic alignment metric in previous sections, here we present the detailed performance of Purity and NMI in Table E5 and Table E6, respectively.

⁸<https://maartengr.github.io/BERTopic/index.html>

⁹<https://github.com/dice-group/Palmetto>

Model	20News			R8			DBpedia			AGNews		
	K=25	K=75	K=100	K=25	K=75	K=100	K=50	K=75	K=100	K=50	K=75	K=100
LDA	0.558 ± 0.008	0.525 ± 0.005	0.510 ± 0.004	0.387 ± 0.011	0.373 ± 0.003	0.374 ± 0.006	0.529 ± 0.005	0.520 ± 0.010	0.506 ± 0.004	0.526 ± 0.010	0.502 ± 0.003	0.480 ± 0.003
BERTopic	0.435 ± 0.029	0.403 ± 0.012	0.390 ± 0.005	0.349 ± 0.013	0.329 ± 0.008	0.333 ± 0.007	0.559 ± 0.011	0.537 ± 0.007	0.516 ± 0.006	0.500 ± 0.016	0.485 ± 0.004	0.479 ± 0.004
TopicGPT	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA	NA
NVDM	0.299 ± 0.012	0.241 ± 0.004	0.232 ± 0.004	0.324 ± 0.011	0.283 ± 0.002	0.270 ± 0.002	0.274 ± 0.003	0.260 ± 0.004	0.259 ± 0.005	0.289 ± 0.004	0.266 ± 0.005	0.256 ± 0.007
+ LLM-ITL	0.439 ± 0.019	0.298 ± 0.009	0.292 ± 0.004	0.438 ± 0.017	0.358 ± 0.011	0.340 ± 0.006	0.365 ± 0.008	0.353 ± 0.008	0.346 ± 0.006	0.357 ± 0.018	0.326 ± 0.005	0.325 ± 0.005
	↑ 46.8%	↑ 23.7%	↑ 25.9%	↑ 35.2%	↑ 26.5%	↑ 25.9%	↑ 33.2%	↑ 35.8%	↑ 33.6%	↑ 23.5%	↑ 22.6%	↑ 27.0%
PLDA	0.418 ± 0.020	0.352 ± 0.010	0.336 ± 0.002	0.391 ± 0.008	0.365 ± 0.006	0.348 ± 0.007	0.484 ± 0.008	0.457 ± 0.009	0.438 ± 0.004	0.515 ± 0.007	0.490 ± 0.012	0.470 ± 0.003
+ LLM-ITL	0.557 ± 0.026	0.534 ± 0.010	0.506 ± 0.013	0.477 ± 0.036	0.456 ± 0.011	0.452 ± 0.004	0.598 ± 0.009	0.570 ± 0.021	0.568 ± 0.015	0.595 ± 0.012	0.582 ± 0.012	0.581 ± 0.008
	↑ 33.3%	↑ 51.7%	↑ 50.6%	↑ 22.0%	↑ 24.9%	↑ 29.9%	↑ 23.6%	↑ 24.7%	↑ 29.7%	↑ 15.5%	↑ 18.8%	↑ 23.6%
SCHOLAR	0.520 ± 0.018	0.478 ± 0.009	0.472 ± 0.014	0.399 ± 0.010	0.385 ± 0.005	0.376 ± 0.004	0.434 ± 0.007	0.422 ± 0.015	0.403 ± 0.021	0.577 ± 0.025	0.581 ± 0.010	0.541 ± 0.014
+ LLM-ITL	0.622 ± 0.028	0.584 ± 0.008	0.562 ± 0.011	0.462 ± 0.010	0.437 ± 0.006	0.424 ± 0.009	0.564 ± 0.037	0.565 ± 0.024	0.556 ± 0.020	0.643 ± 0.015	0.647 ± 0.005	0.627 ± 0.014
	↑ 19.6%	↑ 22.2%	↑ 19.1%	↑ 15.8%	↑ 13.5%	↑ 12.8%	↑ 30.0%	↑ 33.9%	↑ 38.0%	↑ 11.4%	↑ 11.4%	↑ 15.9%
ETM	0.511 ± 0.007	0.466 ± 0.005	0.461 ± 0.008	0.447 ± 0.017	0.432 ± 0.011	0.431 ± 0.009	0.507 ± 0.007	0.501 ± 0.006	0.496 ± 0.007	0.497 ± 0.010	0.497 ± 0.004	0.480 ± 0.004
+ LLM-ITL	0.627 ± 0.019	0.566 ± 0.007	0.553 ± 0.014	0.570 ± 0.014	0.565 ± 0.010	0.580 ± 0.004	0.657 ± 0.031	0.631 ± 0.014	0.626 ± 0.009	0.626 ± 0.018	0.620 ± 0.012	0.618 ± 0.006
	↑ 22.7%	↑ 21.5%	↑ 20.0%	↑ 27.5%	↑ 30.8%	↑ 34.6%	↑ 29.6%	↑ 25.9%	↑ 26.2%	↑ 26.0%	↑ 24.7%	↑ 28.8%
NSTM	0.450 ± 0.008	0.425 ± 0.003	0.424 ± 0.008	0.415 ± 0.014	0.407 ± 0.006	0.402 ± 0.013	0.636 ± 0.008	0.621 ± 0.014	0.616 ± 0.006	0.596 ± 0.015	0.583 ± 0.015	0.575 ± 0.018
+ LLM-ITL	0.510 ± 0.027	0.534 ± 0.008	0.518 ± 0.011	0.549 ± 0.011	0.544 ± 0.021	0.534 ± 0.016	0.702 ± 0.009	0.688 ± 0.009	0.690 ± 0.014	0.661 ± 0.021	0.670 ± 0.008	0.649 ± 0.012
	↑ 13.3%	↑ 23.3%	↑ 22.2%	↑ 32.3%	↑ 33.7%	↑ 32.8%	↑ 10.4%	↑ 10.8%	↑ 12.0%	↑ 10.9%	↑ 14.9%	↑ 12.9%
CLNTM	0.501 ± 0.010	0.478 ± 0.007	0.469 ± 0.010	0.387 ± 0.005	0.346 ± 0.004	0.354 ± 0.002	0.414 ± 0.030	0.394 ± 0.014	0.393 ± 0.011	0.549 ± 0.010	0.567 ± 0.012	0.569 ± 0.010
+ LLM-ITL	0.623 ± 0.011	0.587 ± 0.011	0.580 ± 0.006	0.446 ± 0.006	0.422 ± 0.008	0.411 ± 0.004	0.556 ± 0.004	0.540 ± 0.020	0.547 ± 0.008	0.613 ± 0.003	0.628 ± 0.005	0.658 ± 0.013
	↑ 24.4%	↑ 22.8%	↑ 23.7%	↑ 15.2%	↑ 22.0%	↑ 16.1%	↑ 34.3%	↑ 37.1%	↑ 39.2%	↑ 11.7%	↑ 10.8%	↑ 15.6%
WeTe	0.488 ± 0.010	0.505 ± 0.017	0.498 ± 0.013	0.534 ± 0.008	0.534 ± 0.011	0.533 ± 0.011	0.557 ± 0.012	0.526 ± 0.005	0.516 ± 0.015	0.587 ± 0.007	0.562 ± 0.003	0.556 ± 0.005
+ LLM-ITL	0.539 ± 0.030	0.578 ± 0.022	0.556 ± 0.021	0.604 ± 0.015	0.605 ± 0.012	0.601 ± 0.008	0.650 ± 0.009	0.660 ± 0.011	0.665 ± 0.018	0.655 ± 0.007	0.689 ± 0.010	0.693 ± 0.015
	↑ 10.5%	↑ 14.5%	↑ 11.6%	↑ 13.1%	↑ 13.3%	↑ 12.8%	↑ 16.7%	↑ 25.5%	↑ 28.9%	↑ 11.6%	↑ 22.6%	↑ 24.6%
ECRTM	0.350 ± 0.013	0.324 ± 0.007	0.313 ± 0.003	0.279 ± 0.015	0.287 ± 0.013	0.288 ± 0.009	0.520 ± 0.016	0.436 ± 0.014	0.407 ± 0.008	0.408 ± 0.026	0.382 ± 0.011	0.352 ± 0.014
+ LLM-ITL	0.652 ± 0.026	0.518 ± 0.017	0.520 ± 0.023	0.352 ± 0.012	0.354 ± 0.008	0.362 ± 0.008	0.630 ± 0.006	0.563 ± 0.010	0.521 ± 0.010	0.475 ± 0.019	0.451 ± 0.021	0.426 ± 0.012
	↑ 86.3%	↑ 59.9%	↑ 66.1%	↑ 26.2%	↑ 23.3%	↑ 25.7%	↑ 21.2%	↑ 29.1%	↑ 28.0%	↑ 16.4%	↑ 18.1%	↑ 21.0%

Table E1: Topic coherence (C_V) at other settings of K .

1175

Model	20News			R8			DBpedia			AGNews		
	K=25	K=75	K=100	K=25	K=75	K=100	K=50	K=75	K=100	K=50	K=75	K=100
LDA	0.460 ± 0.011	0.488 ± 0.006	0.489 ± 0.012	0.714 ± 0.008	0.699 ± 0.002	0.689 ± 0.006	0.748 ± 0.007	0.738 ± 0.005	0.730 ± 0.007	0.576 ± 0.005	0.570 ± 0.004	0.562 ± 0.005
BERTopic	0.289 ± 0.021	0.372 ± 0.004	0.390 ± 0.004	0.732 ± 0.011	0.689 ± 0.006	0.685 ± 0.006	0.739 ± 0.002	0.716 ± 0.010	0.695 ± 0.003	0.472 ± 0.004	0.484 ± 0.010	0.495 ± 0.004
TopicGPT	0.363 ± 0.000	0.363 ± 0.000	0.363 ± 0.000	0.410 ± 0.000	0.410 ± 0.000	0.410 ± 0.000	0.706 ± 0.000	0.706 ± 0.000	0.706 ± 0.000	0.634 ± 0.000	0.634 ± 0.000	0.634 ± 0.000
NVDM	0.146 ± 0.005	0.145 ± 0.003	0.153 ± 0.003	0.380 ± 0.007	0.353 ± 0.013	0.359 ± 0.003	0.150 ± 0.004	0.151 ± 0.003	0.156 ± 0.004	0.223 ± 0.004	0.212 ± 0.004	0.211 ± 0.004
+ LLM-ITL	0.143 ± 0.005	0.142 ± 0.005	0.153 ± 0.004	0.377 ± 0.008	0.355 ± 0.006	0.358 ± 0.003	0.150 ± 0.004	0.150 ± 0.003	0.152 ± 0.003	0.222 ± 0.001	0.212 ± 0.003	0.211 ± 0.004
	↓ 2.1%	↓ 2.1%	↑ 0.0%	↓ 0.8%	↑ 0.6%	↓ 0.3%	↑ 0.0%	↓ 0.7%	↓ 2.6%	↓ 0.4%	↑ 0.0%	↑ 0.0%
PLDA	0.231 ± 0.003	0.290 ± 0.005	0.303 ± 0.004	0.523 ± 0.014	0.520 ± 0.007	0.520 ± 0.006	0.643 ± 0.004	0.620 ± 0.008	0.610 ± 0.005	0.474 ± 0.005	0.475 ± 0.004	0.471 ± 0.005
+ LLM-ITL	0.241 ± 0.003	0.305 ± 0.010	0.313 ± 0.011	0.533 ± 0.009	0.518 ± 0.008	0.514 ± 0.006	0.641 ± 0.010	0.624 ± 0.006	0.611 ± 0.005	0.472 ± 0.004	0.470 ± 0.003	0.466 ± 0.001
	↑ 4.3%	↑ 5.2%	↑ 3.3%	↑ 1.9%	↓ 0.4%	↓ 1.2%	↓ 0.3%	↑ 0.6%	↑ 0.2%	↓ 0.4%	↓ 1.1%	↓ 1.1%
SCHOLAR	0.572 ± 0.010	0.580 ± 0.008	0.588 ± 0.004	0.717 ± 0.017	0.686 ± 0.006	0.681 ± 0.005	0.827 ± 0.016	0.794 ± 0.010	0.680 ± 0.114	0.587 ± 0.049	0.426 ± 0.048	0.389 ± 0.067
+ LLM-ITL	0.562 ± 0.011	0.578 ± 0.006	0.587 ± 0.007	0.716 ± 0.018	0.682 ± 0.006	0.675 ± 0.004	0.828 ± 0.014	0.773 ± 0.015	0.701 ± 0.055	0.567 ± 0.053	0.410 ± 0.073	0.382 ± 0.076
	↓ 1.7%	↓ 0.3%	↓ 0.2%	↓ 0.1%	↓ 0.6%	↓ 0.9%	↑ 0.1%	↓ 2.6%	↑ 3.1%	↓ 3.4%	↓ 3.8%	↓ 1.8%
ETM	0.363 ± 0.007	0.434 ± 0.008	0.456 ± 0.013	0.679 ± 0.018	0.666 ± 0.007	0.664 ± 0.005	0.751 ± 0.010	0.747 ± 0.005	0.739 ± 0.006	0.553 ± 0.003	0.555 ± 0.006	0.549 ± 0.003
+ LLM-ITL	0.352 ± 0.009	0.426 ± 0.012	0.441 ± 0.012	0.690 ± 0.011	0.673 ± 0.011	0.666 ± 0.019	0.740 ± 0.005	0.738 ± 0.014	0.744 ± 0.008	0.558 ± 0.005	0.560 ± 0.008	0.553 ± 0.008
	↓ 3.0%	↓ 1.8%	↓ 3.3%	↑ 1.6%	↑ 1.1%	↑ 0.3%	↓ 1.5%	↓ 1.2%	↑ 0.7%	↑ 0.9%	↑ 0.9%	↑ 0.7%
NSTM	0.365 ± 0.013	0.382 ± 0.006	0.405 ± 0.007	0.674 ± 0.004	0.662 ± 0.004	0.662 ± 0.007	0.739 ± 0.007	0.739 ± 0.003	0.738 ± 0.004	0.576 ± 0.007	0.570 ± 0.006	0.569 ± 0.005
+ LLM-ITL	0.351 ± 0.012	0.366 ± 0.008	0.384 ± 0.009	0.681 ± 0.012	0.669 ± 0.003	0.669 ± 0.009	0.741 ± 0.009	0.738 ± 0.006	0.736 ± 0.002	0.576 ± 0.006	0.566 ± 0.006	0.566 ± 0.004
	↓ 3.8%	↓ 4.2%	↓ 5.2%	↑ 1.0%	↑ 1.1%	↑ 1.1%	↑ 0.3%	↓ 0.1%	↓ 0.3%	↑ 0.0%	↓ 0.7%	↓ 0.5%
CLNTM	0.536 ± 0.007	0.581 ± 0.008	0.586 ± 0.006	0.691 ± 0.010	0.680 ± 0.009	0.680 ± 0.006	0.542 ± 0.039	0.438 ± 0.019	0.427 ± 0.062	0.578 ± 0.047	0.487 ± 0.055	0.438 ± 0.035
+ LLM-ITL	0.529 ± 0.013	0.585 ± 0.005	0.585 ± 0.007	0.690 ± 0.009	0.680 ± 0.008	0.676 ± 0.006	0.532 ± 0.032	0.438 ± 0.053	0.428 ± 0.034	0.579 ± 0.026	0.468 ± 0.073	0.453 ± 0.104
	↓ 1.3%	↑ 0.7%	↓ 0.2%	↓ 0.1%	↑ 0.0%	↓ 0.6%	↓ 1.8%	↑ 0.0%	↑ 0.2%	↑ 0.2%	↓ 3.9%	↑ 3.4%
WeTe	0.284 ± 0.005	0.329 ± 0.007	0.339 ± 0.007	0.678 ± 0.002	0.670 ± 0.007	0.649 ± 0.002	0.724 ± 0.009	0.726 ± 0.010	0.725 ± 0.010	0.502 ± 0.006	0.502 ± 0.003	0.487 ± 0.005
+ LLM-ITL	0.278 ± 0.003	0.335 ± 0.011	0.336 ± 0.006	0.685 ± 0.003	0.678 ± 0.009	0.653 ± 0.005	0.722 ± 0.014	0.727 ± 0.008	0.727 ± 0.003	0.501 ± 0.005	0.489 ± 0.005	0.500 ± 0.004
	↓ 2.1%	↑ 1.8%	↓ 0.9%	↑ 1.0%	↑ 1.2%	↑ 0.6%	↓ 0.3%	↑ 0.1%	↑ 0.3%	↓ 0.2%	↓ 2.6%	↑ 2.7%
ECRTM	0.520 ± 0.008	0.533 ± 0.004	0.535 ± 0.003	0.652 ± 0.015	0.670 ± 0.009	0.665 ± 0.008	0.750 ± 0.007	0.745 ± 0.013	0.772 ± 0.006	0.543 ± 0.010	0.536 ± 0.009	0.539 ± 0.007
+ LLM-ITL	0.519 ± 0.009	0.527 ± 0.005	0.532 ± 0.003	0.644 ± 0.015	0.666 ± 0.013	0.666 ± 0.009	0.758 ± 0.008	0.776 ± 0.013	0.776 ± 0.019	0.541 ± 0.008	0.539 ± 0.012	0.524 ± 0.061
	↓ 0.2%	↓ 1.1%	↓ 0.6%	↓ 1.2%	↓ 0.6%	↑ 0.2%	↑ 1.1%	↑ 4.2%	↑ 0.5%	↓ 0.4%	↑ 0.6%	↓ 2.8%

Table E2: Topic alignment (PN) at other settings of K .

Model	20News				R8				DBpedia				AGNews			
	25	50	75	100	25	50	75	100	25	50	75	100	25	50	75	100
LDA	0.854	0.752	0.701	0.667	0.413	0.294	0.244	0.223	0.726	0.612	0.555	0.527	0.606	0.446	0.366	0.331
BERTopic	0.971	0.951	0.934	0.905	0.735	0.620	0.537	0.507	0.958	0.936	0.899	0.859	0.949	0.934	0.925	0.905
TopicGPT	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
NVDM	0.993	0.985	0.956	0.931	0.802	0.801	0.789	0.790	0.982	0.964	0.919	0.875	0.962	0.940	0.896	0.847
+ LLM-ITL	0.970	0.966	0.930	0.896	0.834	0.811	0.765	0.714	0.981	0.940	0.878	0.829	0.962	0.929	0.871	0.813
	↓ 2.3%	↓ 1.9%	↓ 2.7%	↓ 3.8%	↑ 4.0%	↑ 1.2%	↓ 3.0%	↓ 9.6%	↓ 0.1%	↓ 2.5%	↓ 4.5%	↓ 5.3%	↑ 0.0%	↓ 1.2%	↓ 2.8%	↓ 4.0%
PLDA	0.854	0.847	0.857	0.852	0.654	0.506	0.454	0.424	0.886	0.886	0.852	0.821	0.837	0.837	0.771	0.717
+ LLM-ITL	0.870	0.816	0.763	0.767	0.626	0.508	0.463	0.422	0.812	0.812	0.752	0.713	0.796	0.796	0.730	0.689
	↑ 1.9%	↓ 3.7%	↓ 11.0%	↓ 10.0%	↓ 4.3%	↑ 0.4%	↑ 2.0%	↓ 0.5%	↓ 8.4%	↓ 8.4%	↓ 11.7%	↓ 13.2%	↓ 4.9%	↓ 4.9%	↓ 5.3%	↓ 3.9%
SCHOLAR	0.986	0.920	0.860	0.814	0.695	0.509	0.411	0.363	0.965	0.491	0.319	0.246	0.951	0.608	0.398	0.263
+ LLM-ITL	0.958	0.863	0.793	0.736	0.662	0.475	0.385	0.332	0.926	0.520	0.392	0.330	0.903	0.605	0.458	0.292
	↓ 2.8%	↓ 6.2%	↓ 7.8%	↓ 9.6%	↓ 4.7%	↓ 6.7%	↓ 6.3%	↓ 8.5%	↓ 4.0%	↑ 5.9%	↑ 22.9%	↑ 34.1%	↓ 5.0%	↓ 0.5%	↑ 15.1%	↑ 11.0%
ETM	0.890	0.786	0.687	0.604	0.688	0.521	0.444	0.373	0.907	0.864	0.762	0.644	0.932	0.828	0.738	0.668
+ LLM-ITL	0.940	0.832	0.754	0.695	0.676	0.518	0.441	0.377	0.974	0.909	0.837	0.770	0.909	0.733	0.636	0.578
	↑ 5.6%	↑ 5.9%	↑ 9.8%	↑ 15.1%	↓ 1.7%	↓ 0.6%	↓ 0.7%	↑ 1.1%	↑ 7.4%	↑ 5.2%	↑ 9.8%	↑ 19.6%	↓ 2.5%	↓ 11.5%	↓ 13.8%	↓ 13.5%
NSTM	0.648	0.507	0.385	0.330	0.460	0.320	0.244	0.227	0.854	0.768	0.679	0.621	0.831	0.743	0.701	0.703
+ LLM-ITL	0.718	0.593	0.564	0.513	0.446	0.311	0.296	0.224	0.866	0.734	0.683	0.602	0.843	0.749	0.731	0.649
	↑ 10.8%	↑ 17.0%	↑ 46.5%	↑ 55.5%	↓ 3.0%	↓ 2.8%	↑ 21.3%	↓ 1.3%	↑ 1.4%	↓ 4.4%	↑ 0.6%	↓ 3.1%	↑ 1.4%	↑ 0.8%	↑ 4.3%	↓ 7.7%
CLNTM	0.978	0.908	0.850	0.790	0.654	0.448	0.355	0.315	0.851	0.452	0.293	0.224	0.889	0.404	0.272	0.220
+ LLM-ITL	0.936	0.846	0.779	0.718	0.699	0.489	0.386	0.345	0.842	0.491	0.358	0.308	0.820	0.427	0.300	0.274
	↓ 4.3%	↓ 6.8%	↓ 8.4%	↓ 9.1%	↑ 6.9%	↑ 9.2%	↑ 8.7%	↑ 9.5%	↓ 1.1%	↑ 8.6%	↑ 22.2%	↑ 37.5%	↓ 7.8%	↑ 5.7%	↑ 10.3%	↑ 24.5%
WeTe	0.769	0.437	0.382	0.309	0.814	0.617	0.467	0.382	0.950	0.750	0.620	0.569	0.962	0.884	0.710	0.691
+ LLM-ITL	0.879	0.617	0.538	0.448	0.834	0.610	0.488	0.424	0.950	0.816	0.739	0.619	0.942	0.872	0.738	0.738
	↑ 14.3%	↑ 41.2%	↑ 40.8%	↑ 45.0%	↑ 2.5%	↓ 1.1%	↑ 4.5%	↑ 11.0%	↑ 0.0%	↑ 8.8%	↑ 19.2%	↑ 8.8%	↓ 2.1%	↓ 1.4%	↑ 3.9%	↑ 6.8%
ECRTM	0.981	0.986	0.962	0.965	0.979	0.970	0.969	0.976	0.966	0.934	0.973	0.986	1.000	1.000	0.999	0.999
+ LLM-ITL	0.769	0.747	0.701	0.631	0.954	0.958	0.950	0.932	0.926	0.898	0.926	0.933	1.000	0.976	0.954	0.928
	↓ 21.6%	↓ 24.2%	↓ 27.1%	↓ 34.6%	↓ 2.6%	↓ 1.2%	↓ 2.0%	↓ 4.5%	↓ 4.1%	↓ 3.9%	↓ 4.8%	↓ 5.4%	↑ 0.0%	↓ 2.4%	↓ 4.5%	↓ 7.1%

Table E3: Topic diversity (TD) at different K . Std. values have been omitted to save space in this table.

Model	20News				R8				DBpedia				AGNews			
	25	50	75	100	25	50	75	100	25	50	75	100	25	50	75	100
LDA	0.477	0.397	0.368	0.340	0.160	0.113	0.091	0.083	0.390	0.324	0.289	0.267	0.329	0.234	0.184	0.159
BERTopic	0.423	0.385	0.376	0.353	0.257	0.200	0.177	0.169	0.523	0.523	0.483	0.443	0.480	0.467	0.449	0.434
TopicGPT	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
NVDM	0.297	0.257	0.230	0.216	0.260	0.239	0.224	0.213	0.323	0.264	0.239	0.226	0.342	0.272	0.238	0.217
+ LLM-ITL	0.426	0.324	0.277	0.262	0.366	0.303	0.274	0.243	0.470	0.343	0.310	0.287	0.438	0.332	0.284	0.264
	↑ 43.4%	↑ 26.1%	↑ 20.4%	↑ 21.3%	↑ 40.8%	↑ 26.8%	↑ 22.3%	↑ 14.1%	↑ 45.5%	↑ 29.9%	↑ 29.7%	↑ 27.0%	↑ 28.1%	↑ 22.1%	↑ 19.3%	↑ 21.7%
PLDA	0.357	0.312	0.302	0.287	0.256	0.190	0.166	0.148	0.479	0.429	0.390	0.360	0.522	0.431	0.378	0.337
+ LLM-ITL	0.484	0.428	0.407	0.388	0.300	0.237	0.211	0.191	0.575	0.485	0.429	0.405	0.558	0.474	0.425	0.400
	↑ 35.6%	↑ 37.2%	↑ 34.8%	↑ 35.2%	↑ 17.2%	↑ 24.7%	↑ 27.1%	↑ 29.1%	↑ 20.0%	↑ 13.1%	↑ 10.0%	↑ 12.5%	↑ 6.9%	↑ 10.0%	↑ 12.4%	↑ 18.7%
SCHOLAR	0.513	0.441	0.411	0.384	0.278	0.197	0.158	0.136	0.587	0.213	0.134	0.099	0.548	0.351	0.231	0.142
+ LLM-ITL	0.597	0.510	0.463	0.413	0.306	0.208	0.168	0.141	0.628	0.293	0.222	0.184	0.591	0.389	0.296	0.183
	↑ 16.4%	↑ 15.6%	↑ 12.7%	↑ 7.6%	↑ 10.1%	↑ 5.6%	↑ 6.3%	↑ 3.7%	↑ 7.0%	↑ 37.6%	↑ 65.7%	↑ 85.9%	↑ 7.8%	↑ 10.8%	↑ 28.1%	↑ 28.9%
ETM	0.455	0.387	0.320	0.278	0.307	0.225	0.192	0.161	0.466	0.438	0.381	0.320	0.498	0.412	0.367	0.320
+ LLM-ITL	0.589	0.481	0.427	0.385	0.385	0.296	0.249	0.219	0.686	0.597	0.528	0.482	0.585	0.458	0.395	0.357
	↑ 29.5%	↑ 24.3%	↑ 33.4%	↑ 38.5%	↑ 25.4%	↑ 31.6%	↑ 29.7%	↑ 36.0%	↑ 47.2%	↑ 36.3%	↑ 38.6%	↑ 50.6%	↑ 17.5%	↑ 11.2%	↑ 7.6%	↑ 11.6%
NSTM	0.291	0.226	0.164	0.140	0.191	0.132	0.099	0.091	0.557	0.488	0.421	0.383	0.489	0.443	0.408	0.404
+ LLM-ITL	0.367	0.309	0.295	0.266	0.245	0.171	0.161	0.119	0.607	0.515	0.470	0.416	0.569	0.495	0.490	0.421
	↑ 26.1%	↑ 36.7%	↑ 79.9%	↑ 90.0%	↑ 28.3%	↑ 29.5%	↑ 62.6%	↑ 30.8%	↑ 9.0%	↑ 5.5%	↑ 11.6%	↑ 8.6%	↑ 16.4%	↑ 11.7%	↑ 20.1%	↑ 4.2%
CLNTM	0.490	0.445	0.406	0.371	0.253	0.162	0.123	0.111	0.425	0.188	0.115	0.088	0.496	0.222	0.154	0.125
+ LLM-ITL	0.584	0.518	0.457	0.417	0.312	0.213	0.163	0.142	0.515	0.273	0.193	0.168	0.537	0.262	0.188	0.180
	↑ 19.2%	↑ 16.4%	↑ 12.6%	↑ 12.4%	↑ 23.3%	↑ 31.5%	↑ 32.5%	↑ 27.9%	↑ 21.2%	↑ 45.2%	↑ 67.8%	↑ 90.9%	↑ 8.3%	↑ 18.0%	↑ 22.1%	↑ 44.0%
WeTe	0.375	0.216	0.193	0.154	0.435	0.326	0.249	0.203	0.519	0.418	0.326	0.294	0.538	0.514	0.399	0.384
+ LLM-ITL	0.474	0.359	0.310	0.249	0.504	0.365	0.295	0.255	0.583	0.530	0.487	0.412	0.587	0.571	0.508	0.512
	↑ 26.4%	↑ 66.2%	↑ 60.6%	↑ 61.7%	↑ 15.9%	↑ 12.0%	↑ 18.5%	↑ 25.6%	↑ 12.3%	↑ 26.8%	↑ 49.4%	↑ 40.1%	↑ 9.1%	↑ 11.1%	↑ 27.3%	↑ 33.3%
ECRTM	0.343	0.319	0.312	0.302	0.273	0.299	0.278	0.281	0.561	0.486	0.424	0.401	0.438	0.408	0.381	0.352
+ LLM-ITL	0.501	0.412	0.363	0.328	0.336	0.349	0.336	0.337	0.633	0.565	0.521	0.486	0.505	0.464	0.430	0.396
	↑ 46.1%	↑ 29.2%	↑ 16.3%	↑ 8.6%	↑ 23.1%	↑ 16.7%	↑ 20.9%	↑ 19.9%	↑ 12.8%	↑ 16.3%	↑ 22.9%	↑ 21.2%	↑ 15.3%	↑ 13.7%	↑ 12.9%	↑ 12.5%

Table E4: Topic quality (TQ := $C_V \times$ TD) at different K . Std. values have been omitted to save space in this table.

Model	20News				R8				DBpedia				AGNews			
	25	50	75	100	25	50	75	100	25	50	75	100	25	50	75	100
LDA	0.457	0.521	0.527	0.531	0.907	0.920	0.927	0.926	0.813	0.833	0.837	0.839	0.818	0.821	0.822	0.822
BERTopic	0.313	0.371	0.404	0.425	0.882	0.875	0.892	0.897	0.748	0.817	0.810	0.795	0.648	0.688	0.716	0.738
TopicGPT	0.336	0.336	0.336	0.336	0.577	0.577	0.577	0.577	0.718	0.718	0.718	0.718	0.819	0.819	0.819	0.819
NVDM	0.166	0.165	0.166	0.170	0.610	0.600	0.596	0.603	0.210	0.193	0.199	0.204	0.432	0.402	0.389	0.389
+ LLM-ITL	0.162	0.163	0.166	0.172	0.610	0.603	0.598	0.602	0.206	0.195	0.197	0.200	0.428	0.401	0.388	0.389
	↓ 2.4%	↓ 1.2%	↑ 0.0%	↑ 1.2%	↑ 0.0%	↑ 0.5%	↑ 0.3%	↓ 0.2%	↓ 1.9%	↑ 1.0%	↓ 1.0%	↓ 2.0%	↓ 0.9%	↓ 0.2%	↓ 0.3%	↑ 0.0%
PLDA	0.272	0.317	0.337	0.346	0.756	0.770	0.775	0.777	0.746	0.749	0.733	0.730	0.716	0.727	0.731	0.730
+ LLM-ITL	0.284	0.318	0.355	0.361	0.763	0.767	0.771	0.771	0.742	0.746	0.739	0.732	0.722	0.725	0.727	0.726
	↑ 4.4%	↑ 0.3%	↑ 5.3%	↑ 4.3%	↑ 0.9%	↓ 0.4%	↓ 0.5%	↓ 0.8%	↓ 0.5%	↓ 0.4%	↑ 0.8%	↑ 0.3%	↑ 0.8%	↓ 0.3%	↓ 0.5%	↓ 0.5%
SCHOLAR	0.597	0.627	0.634	0.650	0.921	0.911	0.930	0.928	0.872	0.864	0.822	0.690	0.855	0.790	0.570	0.533
+ LLM-ITL	0.580	0.607	0.631	0.649	0.920	0.911	0.927	0.923	0.875	0.865	0.798	0.709	0.855	0.765	0.546	0.523
	↓ 2.8%	↓ 3.2%	↓ 0.5%	↓ 0.2%	↓ 0.1%	↑ 0.0%	↓ 0.3%	↓ 0.5%	↑ 0.3%	↑ 0.1%	↓ 2.9%	↑ 2.8%	↑ 0.0%	↓ 3.2%	↓ 4.2%	↓ 1.9%
ETM	0.355	0.410	0.444	0.471	0.872	0.875	0.887	0.877	0.794	0.813	0.817	0.812	0.783	0.781	0.789	0.787
+ LLM-ITL	0.348	0.403	0.436	0.457	0.860	0.875	0.870	0.861	0.760	0.781	0.798	0.811	0.784	0.784	0.794	0.790
	↓ 2.0%	↓ 1.7%	↓ 1.8%	↓ 3.0%	↓ 1.4%	↑ 0.0%	↓ 1.9%	↓ 1.8%	↓ 4.3%	↓ 3.9%	↓ 2.3%	↓ 0.1%	↑ 0.1%	↑ 0.4%	↑ 0.6%	↑ 0.4%
NSTM	0.356	0.375	0.390	0.423	0.883	0.895	0.903	0.908	0.818	0.825	0.840	0.848	0.818	0.819	0.821	0.825
+ LLM-ITL	0.341	0.361	0.371	0.395	0.884	0.896	0.903	0.905	0.797	0.825	0.836	0.843	0.807	0.819	0.817	0.822
	↓ 4.2%	↓ 3.7%	↓ 4.9%	↓ 6.6%	↑ 0.1%	↑ 0.1%	↑ 0.0%	↓ 0.3%	↓ 2.6%	↑ 0.0%	↓ 0.5%	↓ 0.6%	↓ 1.3%	↑ 0.0%	↓ 0.5%	↓ 0.4%
CLNTM	0.551	0.623	0.637	0.650	0.894	0.923	0.926	0.929	0.725	0.574	0.462	0.454	0.821	0.766	0.641	0.564
+ LLM-ITL	0.543	0.623	0.645	0.649	0.893	0.923	0.925	0.928	0.727	0.565	0.461	0.452	0.806	0.765	0.616	0.591
	↓ 1.5%	↑ 0.0%	↑ 1.3%	↓ 0.2%	↓ 0.1%	↑ 0.0%	↓ 0.1%	↓ 0.1%	↑ 0.3%	↓ 1.6%	↓ 0.2%	↓ 0.4%	↓ 1.8%	↓ 0.1%	↓ 3.9%	↑ 4.8%
WeTe	0.284	0.327	0.347	0.360	0.868	0.860	0.870	0.867	0.739	0.765	0.784	0.796	0.722	0.729	0.735	0.724
+ LLM-ITL	0.281	0.327	0.352	0.356	0.872	0.862	0.878	0.872	0.718	0.761	0.790	0.804	0.725	0.732	0.718	0.741
	↓ 1.1%	↑ 0.0%	↑ 1.4%	↓ 1.1%	↑ 0.5%	↑ 0.2%	↑ 0.9%	↑ 0.6%	↓ 2.8%	↓ 0.5%	↑ 0.8%	↑ 1.0%	↑ 0.4%	↑ 0.4%	↓ 2.3%	↑ 2.3%
ECRTM	0.549	0.552	0.574	0.579	0.837	0.886	0.893	0.891	0.715	0.824	0.808	0.833	0.765	0.771	0.763	0.764
+ LLM-ITL	0.549	0.555	0.565	0.574	0.834	0.881	0.890	0.894	0.747	0.827	0.840	0.830	0.765	0.769	0.761	0.725
	↑ 0.0%	↑ 0.5%	↓ 1.6%	↓ 0.9%	↓ 0.4%	↓ 0.6%	↓ 0.3%	↑ 0.3%	↑ 4.5%	↑ 0.4%	↑ 4.0%	↓ 0.4%	↑ 0.0%	↓ 0.3%	↓ 0.3%	↓ 5.1%

Table E5: Purity at different K . Std. values have been omitted to save space in this table.

1177

Model	20News				R8				DBpedia				AGNews			
	25	50	75	100	25	50	75	100	25	50	75	100	25	50	75	100
LDA	0.463	0.456	0.449	0.447	0.521	0.480	0.472	0.452	0.712	0.662	0.639	0.620	0.373	0.332	0.318	0.303
BERTopic	0.266	0.312	0.341	0.355	0.582	0.519	0.486	0.474	0.691	0.662	0.622	0.595	0.252	0.255	0.252	0.251
TopicGPT	0.390	0.390	0.390	0.390	0.244	0.244	0.244	0.244	0.694	0.694	0.694	0.694	0.449	0.449	0.449	0.449
NVDM	0.126	0.117	0.124	0.136	0.149	0.118	0.111	0.115	0.140	0.107	0.104	0.107	0.071	0.045	0.035	0.033
+ LLM-ITL	0.123	0.113	0.119	0.135	0.145	0.117	0.111	0.114	0.134	0.104	0.102	0.104	0.069	0.044	0.035	0.034
	↓ 2.4%	↓ 3.4%	↓ 4.0%	↓ 0.7%	↓ 2.7%	↓ 0.8%	↑ 0.0%	↓ 0.9%	↓ 4.3%	↓ 2.8%	↓ 1.9%	↓ 2.8%	↓ 2.8%	↓ 2.2%	↑ 0.0%	↑ 3.0%
PLDA	0.191	0.228	0.243	0.260	0.291	0.279	0.265	0.262	0.582	0.537	0.506	0.491	0.246	0.222	0.219	0.212
+ LLM-ITL	0.199	0.225	0.255	0.266	0.303	0.277	0.264	0.258	0.578	0.536	0.509	0.489	0.249	0.220	0.214	0.207
	↑ 4.2%	↓ 1.3%	↑ 4.9%	↑ 2.3%	↑ 4.1%	↓ 0.7%	↓ 0.4%	↓ 1.5%	↓ 0.7%	↓ 0.2%	↑ 0.6%	↓ 0.4%	↑ 1.2%	↓ 0.9%	↓ 2.3%	↓ 2.4%
SCHOLAR	0.547	0.538	0.525	0.525	0.512	0.449	0.443	0.433	0.778	0.790	0.766	0.671	0.420	0.383	0.282	0.244
+ LLM-ITL	0.545	0.529	0.525	0.525	0.511	0.449	0.438	0.426	0.781	0.791	0.746	0.692	0.423	0.368	0.274	0.241
	↓ 0.4%	↓ 1.7%	↑ 0.0%	↑ 0.0%	↓ 0.2%	↑ 0.0%	↓ 1.1%	↓ 1.6%	↑ 0.4%	↑ 0.1%	↓ 2.6%	↑ 3.1%	↑ 0.7%	↓ 3.9%	↓ 2.8%	↓ 1.2%
ETM	0.371	0.405	0.424	0.442	0.485	0.463	0.460	0.438	0.728	0.690	0.678	0.667	0.352	0.325	0.321	0.311
+ LLM-ITL	0.357	0.394	0.416	0.425	0.519	0.498	0.476	0.471	0.723	0.698	0.678	0.676	0.354	0.332	0.327	0.315
	↓ 3.8%	↓ 2.7%	↓ 1.9%	↓ 3.8%	↑ 7.0%	↑ 7.6%	↑ 3.5%	↑ 7.5%	↓ 0.7%	↑ 1.2%	↑ 0.0%	↑ 1.3%	↑ 0.6%	↑ 2.2%	↑ 1.9%	↑ 1.3%
NSTM	0.375	0.371	0.375	0.388	0.465	0.435	0.422	0.415	0.717	0.653	0.638	0.628	0.368	0.334	0.319	0.313
+ LLM-ITL	0.362	0.359	0.362	0.373	0.477	0.450	0.435	0.433	0.706	0.656	0.639	0.630	0.362	0.333	0.316	0.311
	↓ 3.5%	↓ 3.2%	↓ 3.5%	↓ 3.9%	↑ 2.6%	↑ 3.4%	↑ 3.1%	↑ 4.3%	↓ 1.5%	↑ 0.5%	↑ 0.2%	↑ 0.3%	↓ 1.6%	↓ 0.3%	↓ 0.9%	↓ 0.6%
CLNTM	0.522	0.526	0.524	0.521	0.489	0.459	0.434	0.427	0.641	0.510	0.414	0.400	0.392	0.390	0.333	0.311
+ LLM-ITL	0.514	0.529	0.525	0.521	0.486	0.459	0.434	0.425	0.640	0.500	0.415	0.403	0.382	0.394	0.321	0.315
	↓ 1.5%	↑ 0.6%	↑ 0.2%	↑ 0.0%	↓ 0.6%	↑ 0.0%	↑ 0.0%	↓ 0.5%	↓ 0.2%	↓ 2.0%	↑ 0.2%	↑ 0.8%	↓ 2.6%	↑ 1.0%	↓ 3.6%	↑ 1.3%
WeTe	0.285	0.309	0.311	0.319	0.488	0.440	0.471	0.431	0.702	0.683	0.667	0.653	0.310	0.275	0.270	0.250
+ LLM-ITL	0.276	0.305	0.318	0.315	0.498	0.449	0.479	0.434	0.688	0.683	0.663	0.650	0.310	0.271	0.260	0.259
	↓ 3.2%	↓ 1.3%	↑ 2.3%	↓ 1.3%	↑ 2.0%	↑ 2.0%	↑ 1.7%	↑ 0.7%	↓ 2.0%	↑ 0.0%	↓ 0.6%	↓ 0.5%	↑ 0.0%	↓ 1.5%	↓ 3.7%	↑ 3.6%
ECRTM	0.492	0.480	0.492	0.492	0.467	0.459	0.448	0.439	0.621	0.677	0.682	0.712	0.342	0.316	0.309	0.314
+ LLM-ITL	0.490	0.482	0.489	0.491	0.455	0.452	0.442	0.439	0.645	0.689	0.712	0.721	0.341	0.313	0.317	0.322
	↓ 0.4%	↑ 0.4%	↓ 0.6%	↓ 0.2%	↓ 2.6%	↓ 1.5%	↓ 1.3%	↑ 0.0%	↑ 3.9%	↑ 1.8%	↑ 4.4%	↑ 1.3%	↓ 0.3%	↓ 0.9%	↑ 2.6%	↑ 2.5%

Table E6: NMI at different K . Std. values have been omitted to save space in this table.

F Study of Confidence Alternatives

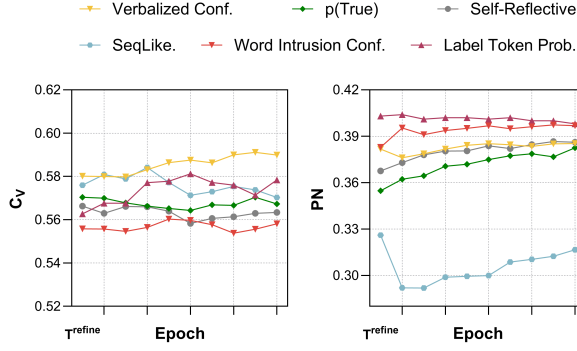


Figure F1: Learning curves of LLM-ITL (ETM as the base model) with different LLM confidences in terms of topic coherence (C_V) and topic alignment (PN) on 20News. Error bars are omitted for clarity in the figure.

Here, we compare our proposed topic labeling confidence with other LLM confidence alternatives within the LLM-ITL framework. We limit our focus to single-sample approaches, where a single round of LLM inference is performed for a given topic. This ensures they can be efficiently applied within the training of LLM-ITL. We consider the following confidence alternatives to our topic labeling confidence in this study: (1) **No Conf.**, where no LLM confidence estimation is included, and $\text{Conf}(w_k^l) = 1$ in Eq. 10 during the training. (2) **Verbalized Confidence** (Xiong et al., 2023), which directly asks the LLM for its confidence in solving a problem. The prompt we used for eliciting verbalized confidence is shown in Figure A2. (3) **Self-Reflective Confidence** (Chen and Mueller, 2023), which prompts the LLM to evaluate its own answer in a two-stage manner. The topic label is obtained in the first stage, and the LLM evaluates this answer in a follow-up question (Figure A3). (4) **p(True)** (Kadavath et al., 2022), which is similar to self-reflective confidence, but asks a true/false question instead. It takes the token probability of the response as the confidence measure (Figure A3). (5) **SeqLike.** (Ren et al., 2022), which computes the length-normalized sequence likelihood of the output from the LLM.

From the results in Figure F1, we can observe that while verbalized confidence results in greater improvements in topic coherence, it can bias the topics toward the LLM’s knowledge rather than the input corpus, leading to reduced topic alignment. In contrast, label token probability and word intrusion confidence consistently achieve the best topic alignment performance, indicating a stronger relevance of the topics to the corpus.

G Study of Hyper-parameters




Figure G1: Learning curves of LLM-ITL (ETM as the base model) with different γ in terms of topic coherence (C_V) and topic alignment (PN) on 20News.

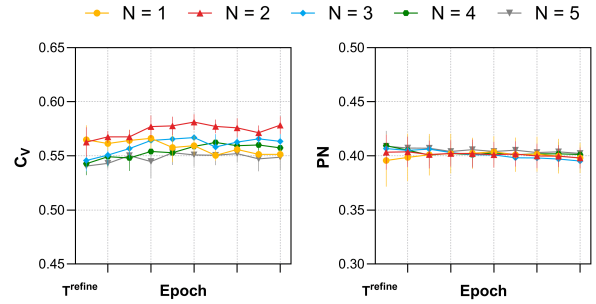


Figure G2: Learning curves of LLM-ITL (ETM as the base model) with different settings of N (i.e., the number of words in the topic label) in terms of topic coherence (C_V) and topic alignment (PN) on 20News.

Here, we study the hyper-parameter of LLM-ITL, focusing on topic refinement strength γ , and the number of words N used for the topic label.

As for γ , we vary its value from 1 to 400 and plot the learning curves in terms of C_V and PN, as shown in Figure G1. We can observe that: (1) In terms of topic coherence, γ values between 100 and 300 yield similar performance, suggesting low sensitivity to γ within this range. (2) In terms of topic alignment, a higher γ leads to slightly reduced performance in the initial phase. This occurs because relying heavily on topic refinement from the LLM causes the topics to bias towards the LLM’s knowledge rather than the information from the input corpus. However, as training progresses, they converge to similar values. These observations suggest that LLM-ITL exhibits low sensitivity to γ within a certain range and offers flexibility in controlling the balance between learning from the corpus and the LLM.

As for N , we vary the number of words for the topic label from 1 to 5 and plot the learning curves in terms of C_V and PN. As illustrated in Figure G2, we can observe that: (1) Using more words as the

Prompt	Success Rate (\uparrow)	N_Input (\downarrow)	N_Output (\downarrow)	Refined TC (\uparrow)
Origin	0.978	224.48	197.86	<u>0.554</u>
Variant_1	0.967	228.48	80.81	0.512
Variant_2	0.935	269.48	167.05	0.498
Variant_3	<u>0.980</u>	223.48	172.51	0.537
Variant_4	0.972	189.48	191.04	0.562
Variant_5	0.956	<u>208.48</u>	<u>159.66</u>	0.543
Iterative Refinement (Chang et al., 2024)	0.993	1872.13	936.70	0.478

Table H1: Study of prompt variants. The best and second-best performance of each column are highlighted in boldface and underlined, respectively.

topic label for LLM-ITL, such as a 5-word topic label (e.g., $N = 5$), results in the least improvement in topic coherence. While using a 2-word topic label (e.g., $N = 2$) achieves the best performance in terms of topic coherence. (2) For topic alignment, the number of words in the topic label demonstrates comparable performance, suggesting low sensitivity to N .

H Study of Prompts

Prompt Variants Here, we study the effectiveness of different prompts for topic suggestion. We obtain variants of topic suggestion prompts in Figure A1 through modifying the technique used in PromptBreeder (Fernando et al., 2023). To be specific, we first create a set of 100 “mutation-prompts” (e.g., “Make a variant of the prompt”) and 100 “thinking-styles” (e.g., “Let’s think step by step”). We generate a set of 50 task prompts by concatenating a randomly drawn “mutation-prompt” and a randomly drawn “thinking-style” to the original prompt, and provide that to the Claude 3.5¹⁰ to produce a continuation, resulting in a different task prompt. Secondly, we randomly select 50 topics over 4 experiment datasets. We run those topics through 50 generated task prompts and filter out the generated prompts that cannot give JSON format in the selected topics or generate above 300 tokens. We are left with 14 topics. We then leverage Claude 3.5 to judge the quality of generated topics and refined topic words. We rank 14 methods by overall topics and refined topic words to get 5 variants of prompts. In addition to the prompt variants generated by those steps, we also investigate the topic refinement prompt used in (Chang et al., 2024) (see Figure 2 of their paper). All the prompt variants for topic refinement in this study are illustrated in Table H2.

Setup We randomly sample 1000 topics learned by topic models, then use different prompts to re-




Figure H1: Learning curves of LLM-ITL (ETM as the base model) with different prompts in terms of C_V and PN on 20News.

fine the topics with LLAMA3-8B-Instruct. We analyze the effectiveness of prompts in different aspects, including **Success Rate** (Ulmer et al., 2024): the proportion of cases where the target answer can be successfully extracted from the LLM’s output; **N_Input** and **N_Output** (Chang et al., 2024): the average number of tokens of input and output of the LLM; and **Refined TC**: the average topic coherence scores of the refined topics.

Results From the results in Table H1, we observe the following: (1) Through prompt optimization, the effectiveness of the prompt can be further enhanced (e.g., Variant_4), where the number of tokens (i.e., the cost) is reduced and the refined topics are more coherent. (2) The iterative refinement (Chang et al., 2024) shows less effectiveness in terms of both cost and refined topic coherence compared with our prompt variants when applied to LLAMA3-8B-Instruct.

Based on the above observations, we further investigate the effectiveness of the improved prompt within the LLM-ITL framework. We plot the learning curves of LLM-ITL using the original prompt and its variant (Variant_4, which shows better performance from Table H1). We observe that the overall performance in terms of both metrics is comparable by using both prompts.

¹⁰<https://www.anthropic.com/news/claude-3-family>

	Prompt
Origin	<p>Analyze step-by-step and provide the final answer.</p> <p>Step 1. Given a set of words, summarize a topic (avoid using proper nouns as topics) by 2 words that covers most of those words. Note, only the topic, no other explanations.</p> <p>Step 2. Remove irrelevant words about the topic from the given word list. Note, only the removed words, no other explanations.</p> <p>Step 3. Add new relevant words (maximum 10 words) about the topic to the word list up to 10 words. Note, only the added words, no other explanations.</p> <p>Step 4. Provide your answer in json format as {"Topic": "<2 Word Topic>", "Words": "<Refined 10 Word List>"}. Note, only 10 refined words allowed for the topic, and no follow up explanations.</p>
Variant_1	<p>Perform the following actions sequentially and provide the final result:</p> <p>Step 1. After examining a set of words, condense a subject (avoid proper nouns) into 2 words that encompass most of those words. (Note: Only the subject, no further elaboration.)</p> <p>Step 2. Eliminate irrelevant words from the given word list based on the subject. (Note: Only the removed words, no further elaboration.)</p> <p>Step 3. Add new pertinent words (maximum 10 words) related to the subject to the word list until it reaches 10 words. (Note: Only the added words, no further elaboration.)</p> <p>Step 4. Present your response in JSON format as {"Topic": "<2 Word Subject>", "Words": "<Refined 10 Word List>"}. Note: Only 10 refined words are permitted for the subject, and no follow-up explanations.</p>
Variant_2	<p>Perform a meticulous examination and furnish the conclusive resolution.</p> <p>Stride 1. Bestowed a catalogue of vocabularies, condense a subject matter (circumvent the employment of proper appellations as subjects) by dual words that envelop the preponderance of those vocabularies. (Heed, solely the subject, devoid of supplemental explication.)</p> <p>Stride 2. Dislodge irrelevant vocabularies concerning the subject from the granted vocabulary catalogue. (Heed, solely the dislodged vocabularies, devoid of supplemental explication.)</p> <p>Stride 3. Amalgamate novel applicable vocabularies (maximal 10 vocabularies) concerning the subject to the vocabulary catalogue up to 10 vocabularies. (Heed, solely the amalgamated vocabularies, devoid of supplemental explication.)</p> <p>Stride 4. Tender your resolution in json format as {"Topic": "<2 Word Subject>", "Words": "<Refined 10 Word Catalogue>"}. Heed, solely 10 refined vocabularies permitted for the subject, and devoid of successive explication.</p>
Variant_3	<p>Step-by-step analysis and final answer:</p> <p>Step 1. Given a set of words, summarize a topic (avoid using proper nouns as topics) by 2 words that covers most of those words. (Note, only the topic, no other explanations.)</p> <p>Step 2. Remove irrelevant words about the topic from the given word list. (Note, only the removed words, no other explanations.)</p> <p>Step 3. Add new relevant words (maximum 10 words) about the topic to the word list, keeping the total word count at 10 words. (Note, only the added words, no other explanations.)</p> <p>Step 4. Provide your answer in JSON format as {"Topic": "<2 Word Topic>", "Words": "<Refined 10 Word List>"}. Note, only 10 refined words allowed for the topic, and no follow-up explanations.</p>
Variant_4	<p>Break down the analysis into steps and give the final response.</p> <ol style="list-style-type: none"> Look at a set of words and identify a 2-word topic that sums up most of those words (don't use proper nouns as topics, just state the topic). Remove words from the list that don't relate to the topic (just list the removed words). Add new relevant words about the topic to the list, up to 10 words total (just list the new added words). Provide your response in JSON format: {"Topic": "<2 Word Topic>", "Words": "<Refined 10 Word List>"}. Only include 10 words for the refined list, no explanations.
Variant_5	<p>Step-by-step analysis and provide the final answer in JSON format:</p> <p>Step 1: Based on the given set of words, summarize a topic using 2 words that encompass most of those words (avoid proper nouns).</p> <p>Step 2: Remove any irrelevant words from the given word list that do not relate to the summarized topic.</p> <p>Step 3: Add new relevant words (up to 10 words) that are related to the summarized topic.</p> <p>Step 4: Present your answer in the following JSON format: {"Topic": "<2 Word Topic>", "Words": "<Refined 10 Word List>"}, where "Topic" contains the 2-word summarized topic, and "Words" contains the refined list of 10 words related to that topic. Do not provide any additional explanations.</p>
Iterative Refinement (Chang et al., 2024)	<p>Please analyze the following tasks and provide your answer in the specified format.</p> <ol style="list-style-type: none"> Determine the common topic shared by these words: [<TOPIC_WORDS >]. Assess whether the word "<WORD>" aligns with the same common topic as the words listed above. <p>Respond with:</p> <ul style="list-style-type: none"> - "Yes", if the given word shares the common topic. - If "No", suggest 10 single-word alternatives that are commonly used and closely related to this topic. These words should be easily recognizable and distinct from the ones in the provided list. <p>Format your response in JSON, including the fields "Topic", "Answer", and "Alternative words" (only if the answer is "No").</p>

Table H2: Prompt variants for topic refinement