

Topics as Entity Clusters: Entity-based Topics from Large Language Models and Graph Neural Networks

Manuel V. Loureiro, Steven Derby, Tri Kurniawan Wijaya

Huawei Ireland Research Centre

Dublin, Ireland

{manuel.loureiro, tri.kurniawan.wijaya}@huawei.com, steven.derby@huawei-partners.com

Abstract

Topic models aim to reveal latent structures within a corpus of text, typically through the use of term-frequency statistics over bag-of-words representations from documents. In recent years, conceptual entities — interpretable, language-independent features linked to external knowledge resources — have been used in place of word-level tokens, as words typically require extensive language processing with a minimal assurance of interpretability. However, current literature is limited when it comes to exploring purely entity-driven neural topic modeling. For instance, despite the advantages of using entities for eliciting thematic structure, it is unclear whether current techniques are compatible with these sparsely organised, information-dense conceptual units. In this work, we explore entity-based neural topic modeling and propose a novel topic clustering approach using bimodal vector representations of entities. Concretely, we extract these latent representations from large language models and graph neural networks trained on a knowledge base of symbolic relations, in order to derive the most salient aspects of these conceptual units. Analysis of coherency metrics confirms that our approach is better suited to working with entities in comparison to state-of-the-art models, particularly when using graph-based embeddings trained on a knowledge base.

Keywords: neural topic modeling, entity clustering, knowledge base, large language model, graph neural network

1. Introduction

Following the seminal work of Blei et al. (2003), topic models have since become the *de facto* method for extracting and elucidating prominent themes from corpora. Traditionally, the semantic content of a document is composed of document-term frequencies or latently through a mixture of distributions of topics, common with probabilistic generative models such as *Latent Dirichlet Allocation* (LDA). Here, individual topics are represented by salient lexical constituents such as words that depict some subjects of the corpora (Blei et al., 2003; Blei and Lafferty, 2006; Li and McCallum, 2006; Teh et al., 2006; Crain et al., 2012). In recent years, the field of *Natural Language Processing* (NLP) has seen a trend toward continuous vector representations of words, which look to capture the paradigmatic relationship between concepts by learning distributional co-occurrence patterns in text. For example, large-scale language models such as *BERT* (Devlin et al., 2018) produce robust contextualized representations that capture an array of linguistic phenomena and implicit real-world knowledge (Peters et al., 2018; Tenney et al., 2019a,b; Petroni et al., 2019; Rogers et al., 2020), paving the way for so-called *Neural Topic Modeling* (Sia et al., 2020; Bianchi et al., 2021; Grootendorst, 2022).

Despite their successes, it becomes evident that certain limitations emerge from conventional topic modeling due to the noisy and superfluous na-

ture of word-level tokens. These methods rely heavily on text processing and data-driven techniques to uncover statistical patterns and infer relevant thematic structure within the corpus, resulting in topics with limited guarantees of expressiveness or interpretability. An alternative approach, which has proven successful, involves extracting conceptual entities from the text and representing them with informative language-agnostic entity identifiers. These identifiers can be unambiguously linked to external lexical resources, encapsulating grounded real-world knowledge (Newman et al., 2006; Chemudugunta et al., 2008; Andrzejewski et al., 2009, 2011; Kim et al., 2012; Hu et al., 2013; Allahyari and Kochut, 2016). Entities, thus, provide rich, interpretable features that facilitate the elicitation of complex topical structure, particularly in multilingual settings, — a limitation of current word-level topic models (Ni et al., 2009; Boyd-Graber and Blei, 2009). However, despite these advantages, there is limited literature on the subject of purely entity-driven neural topic modeling; it is arguable whether conventional topic modeling approaches are suitable for adequately handling these sparse, information-rich data formats (Wang and McCallum, 2006; Zhao et al., 2010; Hong et al., 2010; Quan et al., 2015; Li et al., 2016).

In this paper, we explore entities exclusively as topic features; conceptual entities are distinct, free-form human-derived concepts we represent using encyclopedic-based definitions and several key re-

lational attributes, which offer a better alternative for topic modeling. Here, we propose *Topics as Entity Clusters* (TEC)¹, a novel topic modeling algorithm that can discover meaningful and highly informative topics using vector representation clustering methods. We achieve this by using two sources to represent entities: (1) contextualized text representations constructed from entity definitions and (2) structured graph data extracted from a knowledge base that we use to train a graph neural network to learn node embeddings.

Through the use of experimental procedure, we demonstrate that this combination of implicit contextualised information and explicit semantic knowledge, in conjunction with our novel technique, are important components in overcoming similar data sparsity issues that are observed in other state-of-the-art topic models. Most impressively, we find that simpler graph-based embeddings trained on a knowledge graph of concepts demonstrate considerable improvements in comparison to the large language model embeddings, which we demonstrate across a range of coherency metrics.

2. Literature Review

We are by no means the first to consider entity-based topic models. For instance, Newman et al. (2006) proposed representing documents with salient entities obtained using *Named Entity Recognition* (NER) instead of using the words directly, while entity-centric approaches have also been proposed which focus on word-entity distribution to generate topics (Kim et al., 2012; Hu et al., 2013; Allahyari and Kochut, 2016). Others have attempted to capture the patterns among words, entities, and topics, either by expanding LDA (Blei et al., 2003) or more complex Bayesian topic models — see Alghamdi and Alfalqi (2015), Chauhan and Shah (2021) and Vayansky and Kumar (2020) for a general overview.

2.1. Word embeddings

Researchers have also found success by capitalising on contemporary work in distributional semantics by integrating embedding lookup tables into their frameworks to represent words and documents. For instance, *lda2vec* (Moody, 2016) combines embeddings with topic models by embedding word, document, and topic vectors into a common representation space. Concurrently, Van Gysel et al. (2016) introduces an unsupervised model that learns unidirectional mappings between latent

vector representations of words and entities. Using a shared embedding space for words and topics, Dieng et al. (2020) instead presents the *Embedded Topic Model* (ETM), which merges traditional topic models with the neural-based word embeddings of Mikolov et al. (2013).

2.2. Neural topic models

In recent years, researchers have also looked to incorporate modern deep learning techniques that utilize contextualized representations in contrast to more traditional static embeddings (Grootendorst, 2022; Zhao et al., 2021). Srivastava and Sutton (2016) propose *ProdLDA*, a neural variational inference method for LDA that explicitly approximates the Dirichlet prior. Other models, however, such as *Neural Variational Document Model* (NVDM) (Miao et al., 2017), employ a multinomial factor model of documents that uses amortized variational inference to learn document representations. Bianchi et al. (2021) expand on ProdLDA presenting *CombinedTM*, which improves the model with contextualized embeddings.

2.3. Knowledge extraction

More related to our work, Piccardi and West (2021) — leveraging the self-referencing nature of Wikipedia — define a cross-lingual topic model in which documents are represented by extracted and densified bags-of-links. The adoption of large-scale knowledge base systems has recently gained popularity in NLP as a way to directly inject knowledge into the model (Gillick et al., 2019; Sun et al., 2020; Liu et al., 2020), particularly when concerning conceptual entities (Andrzejewski et al., 2009, 2011; Yang et al., 2015; Terragni et al., 2020a,b).

2.4. Clustering

Clustering techniques have also proved effective for topic modeling. For instance, Sia et al. (2020) introduces clustering to generate coherent topic models from word embeddings, lowering complexity and producing better runtimes compared to traditional topic modeling approaches. Thompson and Mimno (2020) experiment with different pre-trained contextualised embeddings and demonstrate that clustering contextualised representations at the token level is indistinguishable from a Gibbs sampling state for LDA. These findings were also recently corroborated by Zhang et al. (2022) who cluster sentence embeddings and extract top topic words using TF-IDF to produce more coherent and diverse topics than neural topic models.

¹Code available at <https://github.com/manuelvloureiro/topics-as-entity-clusters>

3. Topics as Entity Clusters

In this section, we describe the advantages and challenges associated with performing neural topic modeling with entities, as well as our proposed strategy. First, we begin by motivating our research before discussing our novel approach which can be viewed in Figure 1.

3.1. Background

In the context of topic modeling, entities have proven highly useful to supplement or enrich the data with expert symbolic data and provide factual grounding in the real world. In the literature, these approaches tend to either directly leverage entity data (Newman et al., 2006) or propose models that can effectively utilize prior knowledge in conjunction with word tokens (Kim et al., 2012; Hu et al., 2013; Yang et al., 2015; Terragni et al., 2020a,b). In many ways, entities are well suited for this task, since topic models require the extrapolation of subjects through a small number of distinctly related concepts with variable overlapping themes. Furthermore, the use of words or phrase-level tokens, even in combination with external knowledge, are incompatible in multilingual settings, requiring techniques to link the languages together (Ni et al., 2009; Boyd-Graber and Blei, 2009), something that could be trivial with entity extraction and their language-agnostic representation through entity identifiers.

This approach is not without its caveats. A number of methods have been proposed to better accommodate entities in non-neural models, by either learning word-entity relationships or directly incorporating entity knowledge (Newman et al., 2006; Chemudugunta et al., 2008; Andrzejewski et al., 2009, 2011; Allahyari and Kochut, 2016). For one, documents may only contain a small number of entities, inhibiting the ability of the model to learn important document-term distributions with unsupervised techniques when taking an exclusively entity-based approach (Wang and McCallum, 2006; Zhao et al., 2010; Hong et al., 2010; Quan et al., 2015; Li et al., 2016). Several methods have been proposed to supplement the data in short text topic modeling by adding external information to the text (Hong et al., 2010; Weng et al., 2010), adding heuristic restrictions to the model (Nigam et al., 2000) or altering the models to better handle sparse word co-occurrence patterns (Quan et al., 2015; Li et al., 2016; wang et al., 2021). However, there is limited work pertaining to purely entity-driven neural topic models, which tend to either incorporate LDA into their approach — which struggles in sparse data settings — or rely on term-frequencies for topic inference (Srivastava and Sutton, 2016; Miao et al., 2017; Bianchi et al.,

2021; Zhao et al., 2021; Grootendorst, 2022). Furthermore, measures used to analyse topic quality heavily rely on automated techniques that directly compare *Pointwise Mutual Information* (PMI) with the resulting topic structure, which also depends on co-occurrence statistics. As such, it is important to consider the results of these metrics when performing our analysis².

3.2. Entity representation

Despite the limited research on the subject, pre-trained neural language models should be effective at supplementing topic models in minimal data setting due to their language capabilities and their rich understanding of real-world knowledge (Petroni et al., 2019). However, an approach for generating vector representations for producing the best topics is less clear, particularly when working with sparsely distributed and expressive entity-document features. Broadly, we can consider two methods for constructing entity representations effective for this data structure (Wang et al., 2017): *implicit knowledge* from a pre-trained large language model and *explicit knowledge* extracted directly from a knowledge graph.

Implicit knowledge. Throughout the field of natural language process, language models have been used to construct document representations with great success to express salient knowledge obtained implicitly through a considerable amount of unsupervised learning (Petroni et al., 2019). However, it would be suboptimal to simply compute documents embeddings based on the bag-of-entities approach, as the response elicited from unsupervised models are highly contextualized in nature (Ethayarajh, 2019). We can overcome this challenge by instead extracting descriptions of the entities from an encyclopedic corpus — another benefit of using human-curated conceptual entities.

Explicit knowledge. An advantage of using entities from lexical resources such as a knowledge base is that they provide a systematic framework for organizing and describing human-annotated relationships between concepts. Indeed, the entities confined within these semantic networks exhibit complex relational structures that provide meaningful information about their content, provided in the form of a directed graph. For instance, the triplet <PETRA, CULTURE, NABATAEAN KINGDOM> contains intricate encyclopedic knowledge about the city of Petra that can be difficult to learn with less specialised corpora. Language models may fail to adequately capture this relationship due to

²We also note that there are questions as to whether these PMI measures correlate with human intuition regardless (Hoyle et al., 2021; Doogan et al., 2021)

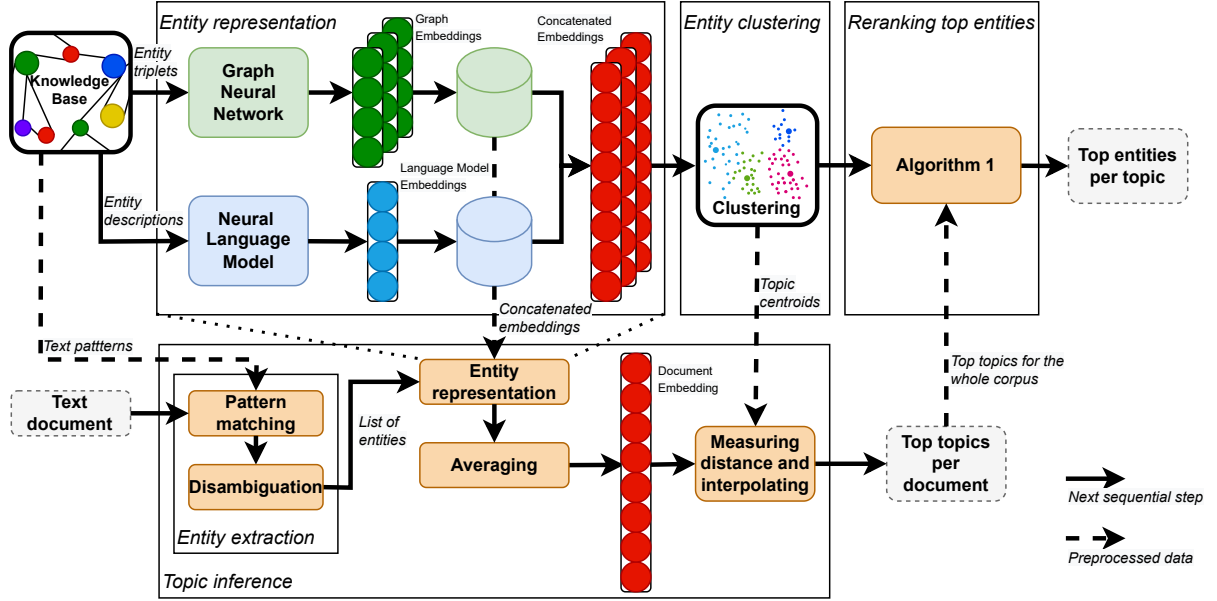


Figure 1: Overview of Topics as Entity Clusters (TEC). The top half illustrates the processing of entity embeddings, topic centroids and top entities per topic, while the bottom half inferring the top topics per document.

the abstract notion of the concept. In addition, explicit knowledge relations afford us supplementary information which is important for sparse bag-of-entities representations.

Combining approaches. Since both methods provide their own distinct advantages to modeling entities, we experiment with combining both sources of knowledge in order to accurately and effectively represent conceptual entities. For some normalized language and graph embeddings $\hat{E}_{LM} \in \mathbb{R}^{d_{LM}}$ and $\hat{E}_G \in \mathbb{R}^{d_G}$, respectively, we weight their contributions using the following concatenation function,

$$\hat{E} = \left[\sqrt{\frac{1}{1+\alpha}} \cdot \hat{E}_{LM}^T, \sqrt{\frac{\alpha}{1+\alpha}} \cdot \hat{E}_G^T \right]^T \quad (1)$$

where $\alpha \in \mathbb{R}$ is the scalar ratio of embedding weights and $\hat{E} \in \mathbb{R}^{d_{LM}+d_G}$ is our final embedding used in entity clustering. We take the square root to guarantee that the final embedding is normalized similarly to the input embeddings.

3.3. Entity clustering

Independent of the specific method, in this work we represent entities in an embedding space and use clustering to construct centroids which we interpret as topic centroids. As such, we model topics based on the innate structure that emerges from the shared embedding space with these entities, similar to other methods in the literature. To this effect, we apply K-Means to the set of enti-

ties contained in a corpus, using the implementation available in FAISS (Johnson et al., 2021). We adopt a two-stage approach to extract entities, which allows us to represent text as a language-agnostic collection of entity identifiers arranged in order of appearance. We first extract candidate entities by finding language-specific text patterns in the original text, followed by a disambiguation process. We include further in-depth detail in Appendix A.

3.4. Topic inference

Topic inference requires that the representation of documents be in the same embedding space as entities and topic centroids. To accomplish this, we extract entities as described in Section 3.3. We then obtain the document representation by calculating the weighted average of those entity embeddings, weighted by term frequency. That is, for each term t in a document d , we generate the document embedding as,

$$E_d = \frac{1}{|d|} \sum_{t \in d} tf_{t,d} \cdot E_t \quad (2)$$

where tf is our term-frequency matrix and E_t is the embedding of the term generated as described in section 3.2. With K representing the number of topics, we can now measure the Euclidean distances $\mathbf{d} = [d_1, d_2, \dots, d_K]^T$ of the document to the topic centroids. Documents are assumed to contain a share of all topics. We infer the topic weight contribution $\mathbf{w} = [w_1, w_2, \dots, w_K]^T$ to the document

using the inverse distance squared weighted interpolation. If we consider the embedding of a document as an interpolation of topic centroids, squaring the distances yields more weight to the closest topic centroids, which helps when working with irregularly-spaced data (Shepard, 1968):

$$w_i = \frac{d_i^{-2}}{\sum_{j=1}^K d_j^{-2}}, \forall i \in \{1, \dots, K\}. \quad (3)$$

3.5. Reranking top entities

A list of highly descriptive entities, weighted by their importance, can be used to express the theme of a topic. However, the closest entities to topic centroids are not necessarily the most descriptive, due to the fact that this naive strategy disregards import entity co-occurrence information within the corpus. To help alleviate issues associated with data sparsity we propose a novel inference-based method to rerank top entities, which assigns the entity frequency of a document to the top topic centroid, as measured by \mathbf{w} (see Algorithm 1).

We start by assigning entities to topics based on their distances weighted by a small value, ϵ (Lines 1-3). We follow by inferring the top topic for each document and updating the top entities in that topic using the document entity frequency. The update is proportional to the inference score, $\max(\mathbf{w})$, as it represents the degree of confidence in the inference (Lines 4-9). To further improve topic diversity in this sparse setting, we only update the top topic. Lastly, we calculate the relative frequencies to obtain the top entities per topic (Lines 10-12).

4. Experiments

We study the performance of TEC and qualitatively compare it to other state-of-the-art topic models using a set of corpora preprocessed into lists of entity identifiers. By contrasting the top entities and measuring results across several metrics, we can infer the quality of each topic model. As it does not pertain to our work, we do not directly compare models trained on entities and word-level tokens, although we do provide an in-depth discussion in Appendix B on why the comparison between word and entity-topic models is challenging, as well as some experimental analysis. Furthermore, since representations are constructed from language-agnostic entity-identifiers, we cannot partition the results by the source languages as our approach is invariant in this setting.

Algorithm 1: Reranking top entities

Input: Number of topics K , number of top entities per topic N , small initialization weight ϵ , documents $Docs$, all entity identifiers in the corpus $entities$, entity embeddings \hat{E}

Output: Lists of top entities per topic $topEntities$, each element is a list of pairs $(entityId, frequency)$

```

1 for topicId ∈ {1, ..., K} do
2   topEntities[topicId] ←
   ClosestEntities(topicId,  $\hat{E}$ , N,  $\epsilon$ )
3 end
4 for doc ∈ Docs do
5    $\mathbf{w}$ , entityFrequency ← TopicInference(doc)
   topTopic ← argmax( $\mathbf{w}$ )
6   for entityId ∈ entities do
7     topEntities[topTopic][entityId] +=
     max( $\mathbf{w}$ ) · entityFrequency[entityId]
8   end
9 end
10 for topicId ∈ [1, ..., K] do
11   topEntities[topicId] ←
   RelativeFrequency(topEntities[topicId])
12 end

```

Table 1: Statistics of the corpora.

Corpus	Vocabulary	Documents	Avg. Entities per Document
WIKIPEDIA	359,507	359,507	44.62
CC-NEWS	94,936	412,731	13.97
MLSUM	89,383	661,422	11.71

4.1. Entity selection and embedding

We build the entity extractor using Wikidata³ as the source of our knowledge base and the English version of Wikipedia⁴ as the encyclopedic corpus. Wikidata entities contain a reference to the corresponding Wikipedia article. Wikidata currently has more than 97 million entities, most of which would be a long tail of entities in a topic model therefore we restrict the entity extractor to only include the top one million entities, as ranked by QRANK⁵ – a public domain project that ranks page views across Wikimedia projects. Out of these entities, we select those matching at least one predicate-object pair from lists of preselected objects for predicates INSTANCE OF, SUBCLASS OF, and FACET OF.

Specifically for the graph neural network, and considering the previously selected entities, we then incorporate all predicates that reference objects as entities existing in the knowledge base.

³Wikidata JSON dump

⁴Collected with Beautiful Soup

⁵QRANK

This strategy not only enhances the graph’s density but also contributes to more robust distinctions between similar entities, resulting in richer embedding representations. For instance, an entity like $\langle \text{JACK DORSEY, OCCUPATION, ENTREPRENEUR} \rangle$ is included since we have $\langle \text{ENTREPRENEUR, INSTANCE OF, PROFESSION} \rangle$ in our knowledge base. Conversely, entries like $\langle \text{JACK DORSEY, EYE COLOR, LIGHT BROWN} \rangle$ are excluded, as they do not match the initial filtering criteria.

Entities are then encoded both using a large language model and a graph neural network. Notwithstanding that other models could be used, we use SBERT⁶ and *node2vec* (Grover and Leskovec, 2016) respectively.

4.2. Corpora

We evaluate all models on various corpora: *Wikipedia*, *CC-News*⁷, and *MLSUM* (Scialom et al., 2020); Table 1 contains a statistics summary. The *Wikipedia* corpus consists of a sample of preprocessed documents, each matching an entity in the vocabulary. *CC-News* consists of monolingual news articles written in English. *MLSUM* on the other hand is a collection of news articles written in multiple languages: German, Spanish, French, Russian, and Turkish. Topic modeling across multiple languages faces challenges due to the language bifurcation problem, which refers to the emergence of thematic structures that are segmented by their source language (Boyd-Graber and Blei, 2009). Despite this, the language-agnostic nature of entities offers a unique advantage, unaffected by such issues, compared to word tokens. For instance, the entity EARTH is given by a tag (Q2) rather than the singular lexical unit “Earth” in English, “Terre” in French or “Erde” in German.

We preprocess the documents according to Section 3.3. The language-specific components for documents in English, German, Spanish and French are *spaCy* lemmatizers (Honnibal and Montani, 2017), for documents in Russian we use *pymorphy2* (Korobov, 2015), and for documents in Turkish we use *zeyrek*⁸.

4.3. Models

Although not strictly a neural topic model, we begin by comparing our approach with LDA (Blei et al., 2003) due to its pervasiveness in the literature. Specifically, we use the *Mallet* implementation of LDA (McCallum, 2002). On top of that, we com-

pare using other state-of-the-art neural topic models from the literature.

NVDM-GSM. *Neural Variational Document Model* (NVDM) is a neural network-based topic model that discovers topics through variational inference training, proposing a number of ways to construct topic distributions, such as a *Gaussian Softmax* (GSM) function (Miao et al., 2017).

ProdLDA. Similar to NVDM-GSM, this model is an autoencoder trained to reconstruct the input embeddings with variational inference-based training (Srivastava and Sutton, 2016).

CombinedTM. This model is a direct extension to ProdLDA that includes pre-trained contextualized embeddings from a pretrained large language model (Bianchi et al., 2021). It extracts contextual vectors for documents using *SBERT*.

WikiPDA. We also consider the Wikipedia-based Polyglot Dirichlet Allocation model, an LDA model trained on entities extracted from Wikipedia (Piccardi and West, 2021). Although it is not a neural approach it is state-of-the-art for working with topic modeling on Wikipedia data, which has its own preprocessing method for further comparison.

4.4. Metrics

We want to evaluate how well models capture the meaning of topics using topic coherence (TC), to estimate the relationship between top entities of a topic (Röder et al., 2015). We follow the implementation of *gensim* (Rehurek and Sojka, 2011),

$$TC = \frac{1}{T} \sum_{t \in \{1..T\}} \left[\frac{2}{N(N-1)} \sum_{\substack{i \in \{1..N\} \\ j \in \{1..i-1\}}} NPMI_{tij} \right]. \quad (4)$$

Topic coherence is defined over the most relevant entities for all topics $t \in \{1..T\}$, where T is the number of topics, N is the number of top entities per topic (we use $N = 10$), and $NPMI_{tij}$ is the normalized pointwise mutual information between entities $e_{t(i)}$ and $e_{t(j)}$, with $e_{t(i)}$ representing top entity of topic t with index i . Normalized pointwise mutual information is given by:

$$NPMI_{tij} = \frac{\log \frac{P(e_{t(i)}, e_{t(j)})}{P(e_{t(i)}) P(e_{t(j)})}}{-\log (P(e_{t(i)}, e_{t(j)}))}. \quad (5)$$

We use normalized pointwise mutual information as it has been shown to correlate well with human judgments of topic coherence (Bouma, 2009).

Topic diversity (TD) is the ratio between the number of unique entities and the total number of entities, considering the top 25 entities per topic (Ding et al., 2020). Topic quality (TQ) is the product of topic coherence and topic diversity.

⁶We use [paraphrase-multilingual-mpnet-base-v2](#).

⁷*CC-News* available at [Hugging Face](#).

⁸*zeyrek* available on <https://github.com/obulat/zeyrek>.

	WIKIPEDIA			CC – NEWS			MLSUM		
Languages	en			en			fr, de, es, ru, tr		
Metric	TC	TD	TQ	TC	TD	TQ	TC	TD	TQ
Number of Topics × 100									
LDA (Entities)	-0.05 (0.01)	0.98 (0.00)	-0.05 (0.01)	-0.13 (0.01)	0.97 (0.00)	-0.13 (0.01)	-0.02 (0.01)	0.96 (0.00)	-0.02 (0.01)
NVDM-GSM	0.06 (0.02)	0.87 (0.02)	0.05 (0.02)	-0.03 (0.02)	0.61 (0.14)	-0.02 (0.01)	0.08 (0.01)	0.59 (0.09)	0.04 (0.01)
ProdLDA	-0.16 (0.03)	0.62 (0.16)	-0.10 (0.03)	-0.30 (0.01)	0.23 (0.01)	-0.07 (0.00)	-0.21 (0.02)	0.36 (0.04)	-0.08 (0.01)
CombinedTM	-0.10 (0.02)	0.22 (0.03)	-0.02 (0.00)	-0.32 (0.01)	0.37 (0.21)	-0.12 (0.07)	-0.22 (0.01)	0.37 (0.11)	-0.08 (0.02)
WikiPDA	0.08 (0.01)	0.73 (0.01)	0.06 (0.00)	—	—	—	—	—	—
TEC E_{LM} ($\alpha = 0$)	0.18 (0.01)	0.95 (0.00)	0.17 (0.01)	0.11 (0.01)	0.79 (0.01)	0.08 (0.01)	0.16 (0.01)	0.79 (0.01)	0.13 (0.01)
TEC $\alpha = 1/2$	0.21 (0.01)	0.95 (0.01)	0.20 (0.01)	0.17 (0.01)	0.81 (0.01)	0.14 (0.01)	0.24 (0.01)	0.82 (0.01)	0.19 (0.01)
TEC $\alpha = 1$	0.21 (0.01)	0.96 (0.01)	0.21 (0.01)	0.19 (0.02)	0.82 (0.01)	0.15 (0.02)	0.24 (0.01)	0.82 (0.01)	0.19 (0.01)
TEC $\alpha = 2$	0.22 (0.01)	0.97 (0.00)	0.22 (0.01)	0.19 (0.01)	0.83 (0.01)	0.16 (0.01)	0.24 (0.01)	0.82 (0.01)	0.20 (0.01)
TEC E_G ($\alpha = \infty$)	0.24 (0.01)	0.97 (0.00)	0.23 (0.01)	0.20 (0.02)	0.83 (0.01)	0.16 (0.02)	0.24 (0.01)	0.83 (0.01)	0.20 (0.01)
Number of Topics × 300									
LDA (Entities)	0.09 (0.01)	0.97 (0.00)	0.09 (0.01)	-0.07 (0.01)	0.91 (0.00)	-0.07 (0.01)	0.03 (0.01)	0.88 (0.00)	0.02 (0.01)
NVDM-GSM	0.06 (0.02)	0.69 (0.03)	0.04 (0.02)	0.04 (0.01)	0.52 (0.05)	0.02 (0.01)	0.13 (0.01)	0.42 (0.04)	0.06 (0.01)
ProdLDA	-0.14 (0.02)	0.44 (0.17)	-0.06 (0.03)	-0.21 (0.01)	0.16 (0.01)	-0.03 (0.00)	-0.16 (0.01)	0.17 (0.01)	-0.03 (0.00)
CombinedTM	-0.13 (0.03)	0.15 (0.03)	-0.02 (0.00)	-0.32 (0.01)	0.41 (0.17)	-0.13 (0.05)	-0.23 (0.01)	0.17 (0.07)	-0.04 (0.02)
WikiPDA	0.06 (0.01)	0.84 (0.01)	0.05 (0.00)	—	—	—	—	—	—
TEC E_{LM} ($\alpha = 0$)	0.25 (0.01)	0.95 (0.00)	0.24 (0.01)	0.11 (0.01)	0.72 (0.00)	0.08 (0.01)	0.14 (0.01)	0.74 (0.01)	0.10 (0.01)
TEC $\alpha = 1/2$	0.29 (0.01)	0.95 (0.00)	0.28 (0.01)	0.18 (0.01)	0.75 (0.01)	0.13 (0.01)	0.22 (0.01)	0.76 (0.00)	0.16 (0.01)
TEC $\alpha = 1$	0.30 (0.01)	0.96 (0.00)	0.29 (0.01)	0.18 (0.01)	0.75 (0.00)	0.14 (0.01)	0.22 (0.01)	0.76 (0.01)	0.17 (0.01)
TEC $\alpha = 2$	0.31 (0.01)	0.96 (0.00)	0.30 (0.01)	0.18 (0.01)	0.76 (0.01)	0.14 (0.00)	0.22 (0.01)	0.76 (0.01)	0.17 (0.01)
TEC E_G ($\alpha = \infty$)	0.31 (0.01)	0.96 (0.00)	0.30 (0.01)	0.18 (0.01)	0.76 (0.01)	0.14 (0.01)	0.23 (0.01)	0.76 (0.01)	0.17 (0.00)

Table 2: Results on all corpora and topic models. We record the results on **TC** : (Topic Coherence), based on the normalized pointwise mutual information of top $N = 10$ entities in a topic, **TD** : (Topic Diversity) the ratio of unique entities to total entities and **TQ** : (Topic Quality) Topic Diversity \times **TC**. The results are reported as the averages of 10 experimental runs and include the margin of error related to 95% confidence intervals.

4.5. Implicit vs Explicit knowledge

To further supplement our findings, we also perform experiments on TEC to determine how the influence of different types of knowledge affects its performance, in this case, the implicit or explicit knowledge from the language model and graph neural network trained on a knowledge base respectively. To achieve this, we train several models with varying values of α , a parameter which effectively controls the impact of a particular source of information as described in Equation 1. As $\alpha \rightarrow \infty$, the influence of the language model representations on TEC will approach zero, leaving TEC to rely exclusively on the graph neural network representations. Conversely, setting $\alpha = 0$ will ensure that TEC relies solely on the language model representations, excluding information from the graph neural network. In this work, we compare the default TEC model, which shares equal influence from both sources of information ($\alpha = 1$), with models trained using varying degrees of information from each of these channels.

4.6. Experiments specifications

For each combination of model, corpus, and the number of topics, i.e., 100 and 300, we compute metrics over 10 runs and present both the averages and 95% confidence interval margin of error in Table 2. We use sequential seeds for the sake of reproducibility. We use implementation defaults

for all models, with the exceptions of NVDM-GSM, where we run 100 epochs, and ProdLDA and CombinedTM, that we run each for 250 epochs. We report metrics for the epoch with higher **TC**.

We run the experiments in a shared Linux machine with 72 CPU cores, 256GB RAM and a Tesla V100-SXM2-16GB GPU.

4.7. Quantitative results

Table 2 shows that both LDA and the neural topic models underperform compared to TEC across all datasets and metrics, except for topic diversity of LDA. Entity extraction leads to sparser representations of corpora, consequently resulting in poorer coherence than typically expected from topic models trained on word tokens — a challenge that underscores the difficulty of constructing practical entity-based topic models. However, our proposed method significantly outperforms similar state-of-the-art approaches. While LDA is quite consistent in terms of topic diversity, the overall topic quality is generally poor, as evidenced by its coherence. Our results suggest that the most valuable source of information for building entity-based topic models is explicit knowledge from a graph neural network as opposed to implicit knowledge from the language model. These findings are supported by the fact that TEC performs consistently worse when purely built using language model embeddings ($\alpha = 0$) in comparison to models constructed exclusively with graph

Model	Sample Topic
LDA	United Nations (Q1065) Teenage Mutant Ninja Turtles (Q12296099) Miles Davis (Q93341) Star Trek (Q1092) United Nations Security Council (Q37470) United Nations Relief and Works Agency for Palestine Refugees in the Near East (Q846656) public health (Q189603) Dizzy Gillespie (Q49575) Greenpeace (Q81307) John Coltrane (Q7346)
NVDM-GSM	bitcoin (Q131723) Apple Inc. (Q312) Halloween [film franchise] (Q1364022) Fisker Inc. [automaker] (Q1420893) IBM (Q37156) Michael Myers (Q1426891) Yakuza [video game series] (Q2594935) Facebook (Q355) cryptocurrency (Q13479982) Vancouver (Q234053)
ProdLDA	Paul McCartney (Q2599) Maxim Gorky (Q12706) Lucy-Jo Hudson (Q1394969) Bob Dylan (Q392) sport utility vehicle (Q192152) FIFA World Cup (Q19317) sedan (Q190578) American football (Q41323) concept car (Q850270) racing automobile (Q673687)
CombinedTM	vocalist (Q2643890) United States of America (Q30) music interpreter (Q3153559) England (Q21) Ryuichi Sakamoto (Q345494) human rights (Q8458) David Tennant (Q214601) Harry Potter (Q76164749) Comedian (Q2591461) Aoni Production (Q1359479)
WikiPDA	a cappella (Q185298) X-Men (Q128452) Marvel Comics (Q173496) To Be [music album] (Q17025795) The Allman Brothers Band (Q507327) proton–proton chain reaction (Q223073) features of the Marvel Universe (Q5439694) Features of the Marvel Cinematic Universe (Q107088537) Uncanny X-Men (Q1399747) member of parliament (Q486839)
TEC $E_{LM} (\alpha = 0)$	Google (Q95) Amazon (Q3884) Microsoft (Q2283) open source (Q39162) Apple Inc. (Q312) Facebook (Q355) Meta Platforms (Q380) Cisco Systems (Q173395) Salesforce.com (Q941127) Citrix Systems (Q916196)
TEC $E_G (\alpha = \infty)$	Mike Tyson (Q79031) World Boxing Organization (Q830940) International Boxing Federation (Q742944) Floyd Mayweather (Q318204) World Boxing Association (Q725676) Tyson Fury (Q1000592) Manny Pacquiao (Q486359) World Boxing Council (Q724450) Evander Holyfield (Q313451) Joe Frazier (Q102301)

Table 3: Example topics using WIKIPEDIA corpus for models trained with 300 topics. Each topic is represented by its top 10 entities. Here, the Q number represents the unique wikidata identifier.

embeddings ($\alpha = \infty$) or some combination of the two ($\alpha \in \{1/2, 1, 2\}$). Furthermore, as the scale of the contribution moves toward graph embeddings (represented by increasing values of α), the overall performance rises steadily, supporting the conclusion that large language model embeddings are generally less important in entity-driven settings. We hypothesise that the increase in performance seen when using explicit knowledge representations can be attributed to their ability to capture semantically related concepts, providing supplementary information to these sparse entity-document terms. Concretely, the graph network is trained on concepts and those within its proximity in the knowledge base, which affords further context to these entity representations when performing clustering which results in better topics. However, we note that in some cases, implicit knowledge from a language model is not always detrimental as in the cases of *CC-News* with 300 topics or *MLSUM* with 100 topics. As such, its not completely inconceivable to image cases where such knowledge would be useful given certain metrics or corpora.

4.8. Qualitative results

We present exemplar topics for the different models in Table 3. Here, we map these entities to their English form although we note that these concepts are language agnostic. Using visual inspection, we find cases where some top entities do not match the general topic theme. These must be attributed to limitations in the model as they all share the same preprocessed corpora, with the exception of WikiPDA. Overall these issues seem less prevalent with TEC. Particularly for ProdLDA and

CombinedTM, we also find unrelated entities that linger across many topics, with the lingering entities varying between runs. We also find topics covering multiple themes, such as the ones resulting from LDA and WikiPDA.

5. Limitations

TEC assumes that documents contain entities, yet this is not necessarily the case. The proposed model is specifically valuable for entity-rich applications such as news articles. A potential solution we are interested in exploring in the future is to train a self-supervised model to generate word embeddings using the bag-of-words as input and the document embedding as the target, to have words represented in the shared embedding space.

We demonstrate that graph neural networks trained on a knowledge base of concepts are highly effective for constructing entity-based topic models. Whilst *node2vec* has proven both practical and sufficient for our work, it is still a shallow neural network that may be unable to learn deeper, more complex relationships between entities. We believe that our results can improve if we obtain embeddings using a multilayer graph neural network with unsupervised training, although we note that the goal of this work is to explore entities as features for neural topic models, for which we demonstrate that graph embeddings trained on an explicit knowledge base performs best.

Another issue, pervading topic modeling literature, is that current automatic coherency metrics may be poorly aligned with human judgment on topic quality (Hoyle et al., 2021; Doogan et al., 2021). Furthermore, these automated approaches

rely heavily on co-occurrence statistics, suboptimal when working with sparse term-document frequencies. It would be interesting to evaluate these models directly using human annotators to judge the performance, for example using intrusion tests.

Lastly, updating the knowledge base will force the retraining of the model, which does not currently guarantee a direct relationship between former and new topics. It requires additional research as this can be a hindrance for some applications.

6. Conclusions

We explore entity-based neural topic models based on the clustering of vector representations. Despite the sparse data structures which result from extracting entities from text data, we find that our TEC model, which performs clustering of implicit and explicit knowledge representations can produce more coherent topics than current state-of-the-art models. TEC represents documents using language-agnostic entity identifiers, resulting in a single set of topics shared across languages. It allows the extension to new languages without sacrificing the performance of the existing ones.

Our results suggest that the implicit knowledge provided by language models is superior to the state-of-the-art in terms of coherence and quality. Nevertheless, these results are surpassed by the explicit knowledge encoded in graph-based embeddings, using human contributed Wikidata knowledge base as a source.

Acknowledgements

We thank Tiziano Piccardi, Federico Bianchi, and our colleagues at Huawei Ireland Research Centre, for helpful comments and discussion.

Bibliographical references

Lund, J., Armstrong, P., Fearn, W., Cowley, S., Byun, C., Boyd-Graber, J. and Seppi, K., [Automatic evaluation of local topic quality](#) *In Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 788-796. 2019.

Doogan, C. and Buntine, W., 2021, June. [Topic model or topic twaddle? re-evaluating semantic interpretability measures](#) *In Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (pp. 3824-3848).

Hoyle, A., Goel, P., Hian-Cheong, A., Peskov, D., Boyd-Graber, J. and Resnik, P., 2021. [Is automated topic model evaluation broken? the incoherence of coherence](#). *Advances in Neural Information Processing Systems*, 34, pp.2018-2033.

Nigam, Kamal, Andrew Kachites McCallum, Sebastian Thrun, and Tom Mitchell. [Text classification from labeled and unlabeled documents using EM](#). *Machine learning* 39 (2000): 103-134.

Wang, Y., Li, X., Zhou, X. and Ouyang, J., 2021, November. [Extracting topics with simultaneous word co-occurrence and semantic correlation graphs: neural topic modeling for short texts](#) *In Findings of the Association for Computational Linguistics: EMNLP 2021* (pp. 18-27)

Weng, Jianshu, Ee-Peng Lim, Jing Jiang, and Qi He. [Twitterrank: finding topic-sensitive influential twitterers](#). *In Proceedings of the third ACM international conference on Web search and data mining*, pp. 261-270. 2010.

Li, Chenliang, Haoran Wang, Zhiqian Zhang, Aixin Sun, and Zongyang Ma. [Topic modeling for short texts with auxiliary word embeddings](#). *In Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*, pp. 165-174. 2016.

Terragni, Silvia, Elisabetta Fersini, and Enza Messina. [Constrained relational topic models](#). *Information Sciences* 512 (2020): 581-594.

Yang, Yi, Doug Downey, and Jordan Boyd-Graber. [Efficient methods for incorporating knowledge into topic models](#). *In Proceedings of the 2015 conference on empirical methods in natural language processing* pp. 308-317. 2015.

Terragni, Silvia, Debora Nozza, Elisabetta Fersini, and Messina Enza. [Which matters most? comparing the impact of concept and document relationships in topic models](#). *In Proceedings of the First Workshop on Insights from Negative Results in NLP*, pp. 32-40. 2020.

Wang, Xuerui, and Andrew McCallum. 2006. [Topics over time: a non-markov continuous-time model of topical trends](#). *In Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining* pp. 424-433. 2006.

Gunel, B., Zhu, C., Zeng, M. and Huang, X., 2020. [Mind the facts: Knowledge-boosted coherent abstractive text summarization](#). *KR2ML Workshop at NeurIPS'19*

Hao, Yanchao, Yuanzhe Zhang, Kang Liu, Shizhu He, Zhanyi Liu, Hua Wu, and Jun Zhao. [An end-to-end model for question answering over knowledge base with cross-attention combining global knowledge](#). *In Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics* (Volume 1: Long Papers), pp. 221-231. 2017.

Quan, Xiaojun, Chunyu Kit, Yong Ge, and Sinno Jialin Pan. 2015

- Short and sparse text topic modeling via self-aggregation. *AAAI Press/International Joint Conferences on Artificial Intelligence* (pp. 2270-2276)
- Zhao, Wayne Xin, Jing Jiang, Jianshu Weng, Jing He, Ee-Peng Lim, Hongfei Yan, and Xiaoming Li. [Comparing twitter and traditional media using topic models](#). In *Advances in Information Retrieval: 33rd European Conference on IR Research, ECIR 2011, Dublin, Ireland, April 18-21, 2011. Proceedings 33, Springer Berlin Heidelberg, 2011.* (pp. 338-349.)
- Hong, L., and Davison, B. D. [Empirical study of topic modeling in twitter](#). In *Proceedings of the first workshop on social media analytics* (pp. 80-88).
- Wang, J., Wang, Z., Zhang, D., and Yan, J. [Combining Knowledge with Deep Convolutional Neural Networks for Short Text Classification](#). In *IJCAI* (Vol. 350, pp. 3172077-3172295).
- Alfred V. Aho and Margaret J. Corasick. 1975. [Efficient string matching: An aid to bibliographic search](#). *Communications of the ACM*, 18(6):333–340.
- Rubayyi Alghamdi and Khalid Alfalqi. 2015. [A survey of topic modeling in text mining](#). *International Journal of Advanced Computer Science and Applications*, 6(1).
- Mehdi Allahyari and Krysztof Kochut. 2016. [Discovering coherent topics with entity topic models](#). In *2016 IEEE/WIC/ACM International Conference on Web Intelligence (WI)*, pages 26–33.
- David Andrzejewski, Xiaojin Zhu, and Mark Craven. 2009. [Incorporating domain knowledge into topic modeling via Dirichlet forest priors](#). In *Proceedings of the 26th Annual International Conference on Machine Learning, ICML '09*, page 25–32, New York, NY, USA. Association for Computing Machinery.
- David Andrzejewski, Xiaojin Zhu, Mark Craven, and Benjamin Recht. 2011. [A framework for incorporating general domain knowledge into latent Dirichlet allocation using first-order logic](#). In *Twenty-Second International Joint Conference on Artificial Intelligence*.
- Federico Bianchi, Silvia Terragni, and Dirk Hovy. 2021. [Pre-training is a hot topic: Contextualized document embeddings improve topic coherence](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 759–766, Online. Association for Computational Linguistics.
- David Blei and John Lafferty. 2006. [Correlated topic models](#). *Advances in Neural Information Processing Systems*, 18:147.
- David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. [Latent Dirichlet allocation](#). *Journal of machine Learning research*, 3(Jan):993–1022.
- Gerlof Bouma. 2009. Normalized (pointwise) mutual information in collocation extraction. *Proceedings of the Biennial GSCL Conference*, 30:31–40.
- Jordan Boyd-Graber and David M Blei. 2009. [Multilingual topic models for unaligned text](#). In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, pages 75–82.
- Uttam Chauhan and Apurva Shah. 2021. [Topic modeling using latent Dirichlet allocation: A survey](#). *ACM Computing Surveys*, 54(7).
- Chaitanya Chemudugunta, America Holloway, Pádraic Smyth, and Mark Steyvers. 2008. Modeling documents by combining semantic concepts with unsupervised statistical learning. In *The Semantic Web - ISWC 2008*, pages 229–244, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Steven P Crain, Ke Zhou, Shuang-Hong Yang, and Hongyuan Zha. 2012. [Dimensionality reduction and topic modeling: From latent semantic indexing to latent Dirichlet allocation and beyond](#). In *Mining Text Data*, pages 129–161. Springer.
- Joachim Daiber, Max Jakob, Chris Hokamp, and Pablo N Mendes. 2013. [Improving efficiency and accuracy in multilingual entity extraction](#). In *Proceedings of the 9th International Conference on Semantic Systems*, pages 121–124.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). *arXiv preprint arXiv:1810.04805*.
- Adji B. Dieng, Francisco J. R. Ruiz, and David M. Blei. 2020. [Topic modeling in embedding spaces](#). *Transactions of the Association for Computational Linguistics*, 8:439–453.
- Kawin Ethayarajh. 2019. [How contextual are contextualized word representations? Comparing the geometry of BERT, ELMo, and GPT-2 embeddings](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 55–65, Hong Kong, China. Association for Computational Linguistics.
- Daniel Gillick, Sayali Kulkarni, Larry Lansing, Alessandro Presta, Jason Baldridge, Eugene Ie, and Diego Garcia-Olano. 2019. [Learning dense representations for entity retrieval](#). In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 528–537, Hong Kong, China. Association for Computational Linguistics.
- Aditya Grover and Jure Leskovec. 2016. [Node2vec: Scalable feature learning for networks](#). In *Proceedings of the 22nd ACM SIGKDD*

- International Conference on Knowledge Discovery and Data Mining*, KDD '16, page 855–864, New York, NY, USA. Association for Computing Machinery.
- Matthew Honnibal and Ines Montani. 2017. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. To appear.
- Linmei Hu, Juanzi Li, Zhihui Li, Chao Shao, and Zhixing Li. 2013. Incorporating entities in news topic modeling. In *Natural Language Processing and Chinese Computing*, pages 139–150, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2021. Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data*, 7(3):535–547.
- Hyungsul Kim, Yizhou Sun, Julia Hockenmaier, and Jiawei Han. 2012. ETM: Entity topic models for mining documents associated with entities. In *2012 IEEE 12th International Conference on Data Mining*, pages 349–358.
- Mikhail Korobov. 2015. Morphological analyzer and generator for russian and ukrainian languages. In Mikhail Yu. Khachay, Natalia Konstantinova, Alexander Panchenko, Dmitry I. Ignatov, and Valeri G. Labunets, editors, *Analysis of Images, Social Networks and Texts*, volume 542 of *Communications in Computer and Information Science*, pages 320–332. Springer International Publishing.
- Wei Li and Andrew McCallum. 2006. Pachinko allocation: DAG-Structured mixture models of topic correlations. In *Proceedings of the 23rd International Conference on Machine Learning*, ICML '06, page 577–584, New York, NY, USA. Association for Computing Machinery.
- Weijie Liu, Peng Zhou, Zhe Zhao, Zhiruo Wang, Qi Ju, Haotang Deng, and Ping Wang. 2020. K-BERT: Enabling language representation with knowledge graph. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(03):2901–2908.
- Tomonari Masada. 2022. AmLDA: A non-VAE neural topic model. In *Information Management and Big Data*, pages 281–295, Cham. Springer International Publishing.
- Andrew Kachites McCallum. 2002. Mallet: A machine learning for language toolkit. <http://mallet.cs.umass.edu>.
- Grootendorst. 2022. BERTopic: Neural topic modeling with a class-based TF-IDF procedure arXiv preprint arXiv:2203.05794.
- Nguyen, Son and Ngo, Bao and Vo, Chau and Cao, Tru. 2019. Hot topic detection on twitter data streams with incremental clustering using named entities and central centroids. In 2019 IEEE-RIVF International Conference on Computing and Communication Technologies (RIVF), pages 1-6.
- Pablo N. Mendes, Max Jakob, Andrés García-Silva, and Christian Bizer. 2011. DBpedia Spotlight: Shedding light on the web of documents. In *Proceedings of the 7th International Conference on Semantic Systems*, I-Semantics '11, page 1–8, New York, NY, USA. Association for Computing Machinery.
- Yishu Miao, Edward Grefenstette, and Phil Blunsom. 2017. Discovering discrete latent topics with neural variational inference. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 2410–2419. PMLR.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S. Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc.
- David Mimno, Hanna Wallach, Edmund Talley, Miriam Leenders, and Andrew McCallum. 2011. Optimizing semantic coherence in topic models. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 262–272, Edinburgh, Scotland, UK. Association for Computational Linguistics.
- Christopher E Moody. 2016. Mixing Dirichlet topic models and word embeddings to make lda2vec. arXiv preprint arXiv:1605.02019.
- David Newman, Chaitanya Chemudugunta, Padhraic Smyth, and Mark Steyvers. 2006. Analyzing entities and topics in news articles using statistical topic models. In *Intelligence and Security Informatics*, pages 93–104, Berlin, Heidelberg. Springer Berlin Heidelberg.
- David Newman, Youn Noh, Edmund Talley, Sarvaz Karimi, and Timothy Baldwin. 2010. Evaluating topic models for digital libraries. In *Proceedings of the 10th Annual Joint Conference on Digital Libraries*, JCDL '10, page 215–224, New York, NY, USA. Association for Computing Machinery.
- Xiaochuan Ni, Jian-Tao Sun, Jian Hu, and Zheng Chen. 2009. Mining multilingual topics from wikipedia. In *Proceedings of the 18th International Conference on World Wide Web*, WWW '09, page 1155–1156, New York, NY, USA. Association for Computing Machinery.
- Matthew E. Peters, Mark Neumann, Luke Zettlemoyer, and Wen-tau Yih. 2018. Dissecting contextual word embeddings: Architecture and representation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1499–1509, Brussels, Belgium. Association for Computational Linguistics.

- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. [Language models as knowledge bases?](#) In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2463–2473, Hong Kong, China. Association for Computational Linguistics.
- Tiziano Piccardi and Robert West. 2021. [Crosslingual topic modeling with WikiPDA](#). In *Proceedings of the Web Conference 2021, WWW '21*, page 3032–3041, New York, NY, USA. Association for Computing Machinery.
- Radim Rehurek and Petr Sojka. 2011. Gensim—python framework for vector space modelling. *NLP Centre, Faculty of Informatics, Masaryk University, Brno, Czech Republic*, 3(2).
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Michael Röder, Andreas Both, and Alexander Hinneburg. 2015. [Exploring the space of topic coherence measures](#). In *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining, WSDM '15*, page 399–408, New York, NY, USA. Association for Computing Machinery.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2020. [A primer in BERTology: What we know about how BERT works](#). *Transactions of the Association for Computational Linguistics*, 8:842–866.
- Thomas Scialom, Paul-Alexis Dray, Sylvain Lamprier, Benjamin Piwowarski, and Jacopo Staiano. 2020. [MLSUM: The multilingual summarization corpus](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8051–8067, Online. Association for Computational Linguistics.
- Donald Shepard. 1968. [A two-dimensional interpolation function for irregularly-spaced data](#). In *Proceedings of the 1968 23rd ACM National Conference, ACM '68*, page 517–524, New York, NY, USA. Association for Computing Machinery.
- Suzanna Sia, Ayush Dalmaia, and Sabrina J. Mielke. 2020. [Tired of topic models? Clusters of pretrained word embeddings make for fast and good topics too!](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1728–1736, Online. Association for Computational Linguistics.
- Akash Srivastava and Charles Sutton. 2016. [Neural variational inference for topic models](#). In *NeurIPS Bayesian deep learning workshop*.
- Yu Sun, Shuohuan Wang, Yukun Li, Shikun Feng, Hao Tian, Hua Wu, and Haifeng Wang. 2020. [ERNIE 2.0: A continual pre-training framework for language understanding](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):8968–8975.
- Yee Whye Teh, Michael I. Jordan, Matthew J. Beal, and David M. Blei. 2006. [Hierarchical Dirichlet processes](#). *Journal of the American Statistical Association*, 101(476):1566–1581.
- Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019a. [BERT rediscovered the classical NLP pipeline](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4593–4601, Florence, Italy. Association for Computational Linguistics.
- Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R Thomas McCoy, Najoung Kim, Benjamin Van Durme, Sam Bowman, Dipanjan Das, and Ellie Pavlick. 2019b. [What do you learn from context? Probing for sentence structure in contextualized word representations](#). In *International Conference on Learning Representations*.
- Laure Thompson and David Mimno. 2020. [Topic modeling with contextualized word representation clusters](#). *arXiv preprint arXiv:2010.12626*.
- Christophe Van Gysel, Maarten de Rijke, and Evangelos Kanoulas. 2016. [Learning latent vector spaces for product search](#). In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management, CIKM '16*, page 165–174, New York, NY, USA. Association for Computing Machinery.
- Ike Vayansky and Sathish A.P. Kumar. 2020. [A review of topic modeling methods](#). *Information Systems*, 94:101582.
- Zihan Zhang, Meng Fang, Ling Chen, and Mohammad-Reza Namazi-Rad. 2022. [Is neural topic modelling better than clustering? An empirical study on clustering with contextual embeddings for topics](#). *arXiv preprint arXiv:2204.09874*.
- He Zhao, Dinh Phung, Viet Huynh, Yuan Jin, Lan Du, and Wray Buntine. 2021. [Topic modelling meets deep neural networks: A survey](#). In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 4713–4720. International Joint Conferences on Artificial Intelligence Organization. Survey Track.

A. Entity extraction

Pattern matching. We first extract candidate entities by finding language-specific text patterns in the original text. For a given language, these text patterns are collections of preprocessed surface forms representing the multiple entities contained in a knowledge base. Inspired by [Mendes et al. \(2011\)](#) and [Daiber et al. \(2013\)](#), we use the deterministic Aho-Corasick algorithm ([Aho and Corasick, 1975](#)) due to its speed and effectiveness in extracting text patterns. The only language-specific components are the preprocessing components, such as lemmatizers, that increase the number of relevant entity matches. These preprocessing components are independent of each other. Consequently, we can expand the model to additional languages without compromising the performance of the others. Assuming the KB contains surface forms in multiple languages, we build finite-state automata, one per language, using lists of preprocessed text patterns. The only language-specific components are the preprocessing components, such as lemmatizers, that increases the number of useful entity matches. This approach provides independence between languages which allows their expansion or update without risking the performance of the incumbent remaining ones.

An alternative method to extract entities is *named entity recognition* (NER) and *entity linking* (EL), but that requires more complex language-specific pipelines with unknown guarantees of success. Similarly to BOW-based methods that have a fixed vocabulary, our approach also assumes a fixed number of entities contained in the KB. For that reason, Aho-Corasick automata are deterministic solutions guaranteed to contain the same entities that form the KB.

Extracting entities does also not guarantee that entities are linked to the ones existing in the KB. It would also require us to either use pretrained NER models or *entity linking* (EL) is an alternative method to obtain entities but, that typically first finds entities and only then attempts to link them to a KB, we have as a requirement that the entities must be represented in the knowledge base. This requirement results from the definition of topic models where the vocabulary must be defined *a priori*.

Disambiguation. As each text pattern can be associated with multiple entities it is necessary to identify which one is the best match and whether that entity is not a spurious match, but is significant to represent the document. The text embeddings result from applying a language model and the entity embeddings are assumed to be previously computed as described in 3.2. We use cosine similarity to rank the prospective entities and

keep the most similar one as long as its score is above some predefined threshold. Notice that if the Language Model supports multiple languages then the entities only require a single embedding to represent it. In practice we can have as many entity embeddings as the languages we want to cover, and so will the entitizer as long as there are language-specific Aho-Corasick automata.

We then select the ones that are above a predefined threshold value. Regarding the entities obtained through pattern matching, the same text pattern may relate to multiple entities. For example, it is not uncommon to find organizations sharing the same acronyms or people with the same name. When a text pattern relates to multiple candidate entities, this method also performs disambiguation as we only keep, at most, the most similar entity⁹.

B. Words vs Entities

We have previously stated that in Section 12 it is not valuable to directly compare word-level topic modeling to entity-based topic modeling due to the sparsity issues associated with PPMI and due to the fact that these metrics have previously received criticism for their inability to accurately determine satisfactory interpretable topics. In the literature, classical topic modeling methods are generally dominated by word or phrase-level approaches in which certain preprocessing steps are used to filter and tokenize sentences into lists of document terms. Common techniques for filtering include approaches that attempt to remove noisy or unhelpful terms from the documents via techniques such as the removal of punctuation, purely syntactically functional terms (for example, stop words) or low-frequency terms. However, in comparison to entity-based extraction methods, these word-level document representations are much denser in terms of raw token count and shallower in terms of actual semantic information. For instance, entity extraction and tokenization methods may isolate an entity such as NEW YORK CITY, from which word-level approaches may similarly filter the terms "New", "York" and "City" (or just "New York" since the disambiguation process can determine this phrase to be the entity NEW YORK CITY) which individually are more fine-grained in terms of syntagmatic/paradigmatic information between documents whilst also being less rich in terms of pure semantic information. Indeed, the term "city" most likely appears many times across documents in a corpus whereas something more specific such as the phrase "New York" or "New York City" which

⁹While this approach may not be the most optimal way of providing entities with both high precision and recall, it does generate a sufficient number for performing our research.

CCNews			
Model	TC	TD	TQ
Number of Topics $\times 100$			
LDA (Entities)	-0.13 (0.01)	0.97 (0.00)	-0.13 (0.01)
LDA (Words)	0.14 (0.00)	0.59 (0.00)	0.08 (0.00)
TEC $E_G(\alpha = \infty)$	0.20 (0.02)	0.83 (0.01)	0.16 (0.02)
Number of Topics $\times 300$			
LDA (Entities)	-0.07 (0.01)	0.91 (0.00)	-0.07 (0.01)
LDA (Words)	0.14 (0.00)	0.43 (0.00)	0.06 (0.00)
TEC $E_G(\alpha = \infty)$	0.18 (0.01)	0.76 (0.01)	0.14 (0.01)

Table 4: Results on CCNews corpus for all topic models. We record the results on **TC** : (Topic Coherence), based on the normalized pointwise mutual information of top $N = 10$ entities in a topic, **TD** : (Topic Diversity) the ratio of unique entities to total entities and **TQ** : (Topic Quality) Topic Diversity \times **TC**. The results are reported as the averages of 10 experimental runs and include the margin of error related to 95% confidence intervals. Our model outperforms all baselines across all metrics except for **TD**.

LDA(Word – Tokens)					
Topic 1	Topic 2	Topic 3	Topic 4	Topic 5	Topic 6
polic	fire	peopl	counti	appl	price
offic	weather	thing	st	app	averag
man	storm	make	hill	devic	target
arrest	hurrican	realli	west	phone	analyst
kill	area	talk	ohio	googl	stock

Table 5: Example of the top 5 terms from six randomly sampled topics from LDA trained on word-level tokens. Given that this method is a common approach to topic modeling, it is unsurprising that the topics are quite coherent.

gets classified as the entity NEW YORK CITY may only appear sparingly, further illustrating the problem with sparsity which any entity-based model (including the model presented in this work) must overcome, and for which models such as LDA have been shown to be less suited. Further issues emerge when directly measuring and comparing the intrinsic performance between word and entity-level-based approaches. Because *Positive Pointwise Mutual Information* (PPMI) is generally employed as a proxy for measuring the coherence of a topic, the sparse structure of entity-level tokens will produce highly distinct distributions to that of word-token PPMI matrices. Consider the fact that n-grams in word tokens will have high joint probabilities, which boosts coherence metrics as described in Equations 4 and 5. However, these higher coherence scores do not necessarily reflect an increase in interpretability. Thus, while coherency metrics may offer a useful simulacrum for gauging interpretability and meaningful thematic structure, they are ultimately only useful when comparing different topic modeling methods within the same tokenization strategy. To further demonstrate this and for completeness, we examine how a model such

as LDA compares when trained on words and entities to further illustrate our discussion.

B.1. Preprocessing

For our experiments, we compare only on *CC-News* text data. Note that we can only perform these experiments on *CC-News* or *Wikipedia* since word-level topic modeling is not directly compatible with the multilingual setting of *MLSUM*, a distinct advantage of using entities. We perform word tokenization before removing stop words, punctuation and tokens that represent digits. For stemming, we use the *NLTK* implementation of the *Porter Stemmer*¹⁰. Finally, we prune the tokens by i) removing tokens that occur less than 15 times in the corpus ii) removing tokens that occur in less than 20 documents, and iii) removing tokens that occur in greater than 50% of documents. Similarly, we use the *Mallet* implementation of LDA (McCallum, 2002) to train the LDA model.

¹⁰https://www.nltk.org/_modules/nltk/stem/porter.html

LDA(Entity – Tokens)		
Topic 1	Topic 2	Topic 3
Puerto Rico FC Audit broadcasting Victoria farming	Las Vegas Nevada Hawaii University of the Pacific Pixel 1 XL	WWE Sunderland A.F.C cycling Orix Buffaloes Brock Lesnar
Topic 4	Topic 5	Topic 6
Aston Villa F.C Birmingham City F.C Burton Albion F.C music interpreter Midland	Russia Ultimate Fighting Championship Vladimir Putin Moscow Kremlin mixed martial arts	White House presidency of Donald Trump Minnesota Melania Trump Executive Office of the President of the United States

Table 6: Example of the top 5 terms from six randomly sampled topics from LDA trained on entity tokens. Here we see that topics are not as clear when compared to a model trained on word-level tokens.

TEC $\alpha = 2$		
Topic 1	Topic 2	Topic 3
Brazil Rio de Janeiro São Paulo Portugal Tony Ramos	Bill Clinton Barack Obama Joe Biden presidential election George W. Bush	France Paris Claude Lelouch Laurent Ruquier Patrice Leconte
Topic 4	Topic 5	Topic 6
Bill Gates Warren Buffett Mark Zuckerberg Steve Jobs Google	Basketball National Basketball Association LeBron James Jason Kidd Kevin Durant	Iran Tehran Ali Khamenei Mahmoud Ahmadinejad Ruhollah Khomeini

Table 7: Example of the top 5 terms from six randomly sampled topics from TEC $\alpha = 2$ trained on entity tokens. We see that our novel approach produces highly interpretable and coherent topics from entities.

B.2. Results

Looking at Table 4, we see that, in comparison to LDA trained on entities, the word-level tokens provide stronger performance on a range of coherency metrics except on topic diversity, indicating its inability to thoroughly extrapolate thematic structure. We attribute this to multiple factors, including the fact that these metrics rely on point-wise mutual information from extremely sparse co-occurrence information, and that LDA itself relies heavily on these term-document distributions in text data. However, when comparing word-level word-based LDA to TEC (although as we state, we must be careful when making direct comparisons) we see that TEC outperforms LDA on topic coher-

ence and topic quality, whilst retaining a high degree of topic diversity further illustrating its performance. With respect to entity-based models, the inability of LDA to build coherent topics is particularly striking, as opposed to TEC. We also see that the LDA model trained on words tends to have less variation across these metrics, highlighting the difficulty of training with sparse, information-rich document terms.

B.3. Topic examples

Next, we provide further examination through visual inspection. We note that the following analysis is purely visual in nature, although they are supported by the empirical findings in the paper.

Observing the results in Table 5, we see that LDA trained on word-level tokens produce fairly reasonably coherent topics from the data, as to be expected from the model. Comparing with results of LDA trained on entities as seen in Table 6, we observe that the topics are quite unclear with a number of noisy terms that do not seem to relate well to the other terms since LDA is not entirely equipped to handle this type of data. However, looking at examples from Table 7 we see that all topics are highly interpretable and coherent, demonstrating the strength of our novel approach. In the future, crowdsourcing tasks such as word intrusion or topic-word matching from human annotators may provide better insight when evaluating these models (Lund et al., 2019; Hoyle et al., 2021; Doogan et al., 2021).