

Experiments on Generalizability of BERTopic on Multi-Domain Short Text

Muriël de Groot*

Albert Heijn,
Ahold Delhaize,
Amsterdam, The Netherlands
muriel.de.groot@ah.nl

Mohammad Aliannejadi

IRLab,
University of Amsterdam,
Amsterdam, The Netherlands
m.aliannejadi@uva.nl

Marcel R. Haas

Business Intelligence
University of Amsterdam
Amsterdam, The Netherlands
m.r.haas@uva.nl

Abstract

Topic modeling is widely used for analytically evaluating large collections of textual data. One of the most popular topic techniques is Latent Dirichlet Allocation (LDA), which is flexible and adaptive, but not optimal for e.g. short texts from various domains. We explore how the state-of-the-art BERTopic algorithm performs on short multi-domain text and find that it generalizes better than LDA in terms of topic coherence and diversity. We further analyze the performance of the HDBSCAN clustering algorithm utilized by BERTopic and find that it classifies a majority of the documents as outliers. This crucial, yet overseen problem excludes too many documents from further analysis. When we replace HDBSCAN with k-Means, we achieve similar performance, but without outliers.

1 Introduction

Currently, one of the most widely used topic models in NLP is Latent Dirichlet Allocation (LDA) (Blei et al., 2003). Although LDA is effective at modeling long conventional text collections such as news articles, it performs poorly on less conventional text such as short documents (Zuo et al., 2016; Hong and Davison, 2010; Zhao et al., 2011; Qiang et al., 2017). Other algorithms like the biterm topic model (BTM) (Yan et al., 2013) are successful in modeling topics in short texts. However, a common weakness between LDA and BTM is that the models cannot capture the context of words. As a result, the models lack generalizability across different domains, where different words are used to describe similar concepts (Zhao and Mao, 2014; Zhuang et al., 2010). BERTopic (Grootendorst, 2022) is a state-of-the-art topic modeling technique that could potentially overcome the mentioned limitations of LDA. BERTopic leverages pre-trained language embeddings and a class-based TF-IDF to create dense clusters, allowing for easily interpretable topics. We study and compare the performance of

BERTopic and LDA on multi-domain short text corpora. Our goal is to compare how these models perform and generalize across multiple domains.

To this aim, we conduct our experiments on a corpus of open-question responses on university course evaluations from various faculties, which ensures that the language used across the corpus is diverse. Our experiments reveal that the HDBSCAN clustering utilized in BERTopic discards 74% of the documents by labeling them as outliers, leading to results with little value in practice. Therefore, we propose using the k-Means clustering algorithm to maintain effectiveness while controlling the model’s sensitivity to outliers.

Our results show that BERTopic outperforms LDA in terms of different performance measures in various domains, confirming that it can generalize better across different domains on short text. Our contributions can be summarized as follows. (i) We apply BERTopic to unconventional data sets, characterized as short documents originating from different domains; (ii) we investigate the impact of replacing the original HDBSCAN clustering by k-Means clustering on the performance of BERTopic.

2 Experimental Methodology

Metrics. We compare the performance of LDA, BERTopic HDBSCAN, and BERTopic k-Means regarding topic coherence and diversity. Topic coherence is evaluated using normalized pointwise mutual information (NPMI) (Bouma, 2009). NPMI ranges from -1 to 1, where a higher NPMI indicates more strongly related words within a topic. Topic diversity is the percentage of unique words in the top-n words of all topics (Dieng et al., 2020) and is in the ranges [0,1]. Diversity near 0 indicates redundant topics, while values close to 1 indicate varied topics.

We calculate these performance measures on the corpora of each faculty separately, as well as from all faculties combined. As such, we evaluate how scalable the LDA and BERTopic models are across domains with potentially varying vocabulary.

*Work done when student at the University of Amsterdam.

Table 1: Results. Ranging from 5 to 30 topics with steps of 5, topic coherence (TC) and topic diversity (TD) are shown for each topic model. All results are averaged across three runs. Highest and lowest scores are marked in green and red, respectively.

		Faculty I	Faculty II	Faculty III	Faculty IV	Faculty V	All faculties	Average	20NG
LDA	TC	-0.022	-0.020	-0.013	0.025	-0.015	0.031	-0.002	0.057
	TD	0.358	0.528	0.414	0.621	0.500	0.718	0.523	0.752
BERTopic HDBSCAN	TC	0.046	0.057	0.030	0.118	0.145	0.091	0.081	0.166
	TD	0.735	0.837	0.855	0.829	0.903	0.880	0.840	0.902
BERTopic k-Means	TC	0.028	0.033	0.028	0.050	0.029	0.033	0.032	0.113
	TD	0.531	0.626	0.598	0.653	0.618	0.571	0.600	0.853

Data. Our dataset consists of open-text comments of university students across five faculties, ranging from computer science to law students. Our dataset contains 62,522 raw responses. The length of the responses varies considerably among the faculties, ranging from 14 to 20 words median. We also experiment on the 20NG dataset,¹ which serves as a benchmark for validation. It comprises 11,096 news articles across 20 categories. We further examine whether vocabulary and document length impact the performance differences between models in the main experiment. To assess the impact of document length, two samples are taken from the 20NG dataset. The sample that consists of short documents has a median document length of 16, which is equal to the median document length of the CER data set. In other words, this sample of short documents from the 20NG data set represents the same document length as the CER data set. The long document sample consists of documents containing between 60 and 100 words.

3 Results

Table 1 lists the results. BERTopic HDBSCAN can successfully create a university-wide topic model, with a topic coherence of 0.091 and a topic diversity of 0.880. In these metrics it outperforms LDA, which has coherence and diversity scores of 0.031 and 0.718, respectively. Moreover, differences in vocabulary among faculties were detected. To illustrate, students from Faculty I tend to use 'tutorial', where students from faculty III use the word 'class' or 'seminar' to describe the same concept. Although vocabulary differences have been detected, there was no impact on the topic generation of BERTopic. Furthermore, we observe that the topic coherence of BERTopic HDBSCAN declines when applied to shorter documents, making the model as susceptible to document length as LDA, see Table 2. BERTopic k-Means is the least susceptible to short documents.

A limitation of BERTopic HDBSCAN is the

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

Table 2: Results of the document length experiment on the 20NG corpus. The performance difference of the models between long and short documents is denoted with δ .

		Long documents	Short documents	δ
LDA	TC	-0.020	-0.180	-0.160
	TD	0.614	0.656	0.042
BERTopic HDBSCAN	TC	0.080	-0.065	-0.145
	TD	0.910	0.972	0.062
BERTopic k-Means	TC	0.058	-0.023	-0.081
	TD	0.851	0.850	-0.001

outlier generation. Approximately 74% of the student responses are classified as outliers, rendering HDBSCAN inappropriate for the analysis of university course evaluations, in which every response matters. In contrast, k-Means does not generate outliers, creates interpretable topics, and is more robust to short documents than HDBSCAN and LDA. However, in terms of generalizability to a university-wide corpus, BERTopic k-Means performs less effectively, compared to BERTopic HDBSCAN, with a coherence of 0.032 and a diversity of 0.571.

4 Conclusions

This paper explores the state-of-the-art BERTopic algorithm on open-question responses from course evaluations, with the goal of creating a topic model architecture that is scalable to a university-wide level. Besides the algorithm's default clustering algorithm, HDBSCAN, we investigated k-Means. We compared the models to a baseline LDA model. Although BERTopic HDBSCAN shows superior performance with regards to topic coherence and diversity in a university-wide model, it labels the majority of the documents as outliers. We propose to replace HDBSCAN with k-Means, because that algorithm does not generate outliers and is more robust to short documents. The absence of an all-encompassing topic model performance measure that takes outlier generation into account is a limitation in this research and should be considered in future work.

References

- 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 GSCL*, 30:31–40.
- 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.
- Maarten Grootendorst. 2022. Bertopic: Neural topic modeling with a class-based tf-idf procedure. *arXiv preprint arXiv:2203.05794*.
- Liangjie Hong and Brian D Davison. 2010. Empirical study of topic modeling in twitter. In *Proceedings of the first workshop on social media analytics*, pages 80–88.
- Jipeng Qiang, Ping Chen, Tong Wang, and Xindong Wu. 2017. Topic modeling over short texts by incorporating word embeddings. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 363–374. Springer.
- 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.
- Rui Zhao and Kezhi Mao. 2014. Supervised adaptive-transfer pls for cross-domain text classification. In *2014 IEEE International Conference on Data Mining Workshop*, pages 259–266. IEEE.
- Wayne Xin Zhao, Jing Jiang, Jianshu Weng, Jing He, Ee-Peng Lim, Hongfei Yan, and Xiaoming Li. 2011. Comparing twitter and traditional media using topic models. In *European conference on information retrieval*, pages 338–349. Springer.
- Fuzhen Zhuang, Ping Luo, Zhiyong Shen, Qing He, Yuhong Xiong, Zhongzhi Shi, and Hui Xiong. 2010. Collaborative dual-plsa: mining distinction and commonality across multiple domains for text classification. In *Proceedings of the 19th ACM international conference on Information and knowledge management*, pages 359–368.
- Yuan Zuo, Jichang Zhao, and Ke Xu. 2016. Word network topic model: a simple but general solution for short and imbalanced texts. *Knowledge and Information Systems*, 48(2):379–398.