

# A Process for Topic Modeling via Word Embeddings

Diego Saldaña-Ulloa

Benemérita Universidad Autónoma de Puebla,  
Facultad Ciencias de la Computación, Puebla,  
Mexico

`diegos.ulloa13@gmail.com`

**Abstract.** This work combines algorithms based on word embeddings, dimensionality reduction, and clustering. The objective is to obtain topics from a set of unclassified texts. The algorithm to obtain the word embeddings is the BERT model, a neural network architecture widely used in NLP tasks. Due to the high dimensionality, a dimensionality reduction technique called UMAP is used. This method manages to reduce the dimensions while preserving part of the local and global information of the original data. K-Means is used as the clustering algorithm to obtain the topics. Then, the topics are evaluated using the TF-IDF statistics, Topic Diversity, and Topic Coherence to get the meaning of the words on the clusters. The results of the process show good values, so the topic modeling of this process is a viable option for classifying or clustering texts without labels.

**Keywords:** Topic modeling, word embedding, dimensionality reduction, TF-IDF.

## 1 Introduction

Text processing in the digital age is widely linked to daily processes occurring in all ambits. These types of tools allow obtaining several features related to the semantics of a text that can later be used for a wide variety of purposes, mainly related to the correct categorization of a document. The processes involving textual categorization are one of the tools with the most applications in different areas [2], considering the growing increase in information generation. This previous statement drives the need to automatically categorize a text (or document) as a fundamental task.

The process of categorizing a document or extracting its related topics is called topic modeling. Topic modeling involves collecting a set of latent variables that help define concepts from documents [3]. This type of process can be helpful in sites with a large amount of information, for example, in databases of journals or articles, and even in content generated in social networks by users [4, 5].

The methods to extract topics from a document look for a way to access the semantic information of the text; in this way, different types of processes

that combine techniques and statistical algorithms can be designed. As a first step, using processes based on the count of words is common. However, these techniques could not consider the dependency between the document's different terms [6]. This is the reason for the necessity of techniques that consider both the words' context and the frequency of terms.

In this work, we propose a set of techniques and algorithms that consider the latent variables of a document, the context of the words (through word embeddings), and the frequency of the different terms present in a text to perform topic modeling. The organization that will be followed will be the one described below: Section 1 describes the background related to the topic modeling area as well as the related works, section 2 introduces the theory of the different methods and algorithms used, section 3 presents the description of the process for the extraction of topics from a document, in section 4 the experimental results are detailed, and finally, the conclusions are handled.

## 2 Related Work

Topic modeling is a set of techniques used to extract information from a document and define a set of ideas that characterize it. There are different methods to carry out this process, but many arise from the Vector Space Model (VSM) [7]. VSM is a mathematical model to represent texts that work considering the relevance (numerical weights) of different terms on a set of documents [8]. Relevance is generally assigned by a function related to the frequency of each term in the entire document. In this way, an  $n$ -dimensional vector type can be defined for each document, formed by the numerical weights of each word.

TF-IDF is one of the first methods used to consider the relevance of a term over the whole document. This method considers the frequencies of the terms (TF) over the total corpus size (IDF). The result of this procedure is precisely a matrix of weights for each of the terms in a document. The TF-IDF method is used by another technique called Latent Semantic Analysis (LSA) which is used in natural language processing tasks such as text classification. LSA factors the TF-IDF matrix using Singular Value Decomposition (SVD) and thus manages to reduce its dimensionality. In this process, part of the semantic information of the terms is captured, so LSA can be used to assign topics to a document based on the numerical weights of each term over the entire text [9].

Another method that is used in topic extraction is PLSA (Probabilistic Latent Semantic Analysis). It is based on LSA, and they differ in that PLSA considers that each term comes from a probability distribution given for each one of the topics. In this way, a set of probability distributions of a fixed set of topics characterizes each document. The model estimates the topic and word distributions that best explain a co-occurrence of terms in the corpus [3].

Latent Dirichlet Allocation (LDA) is another method used for topic modeling. It works similarly to PLSA, each document is considered a mixture of topics, and each topic comes from a probability distribution of the possible words in the

topic. The algorithm works by assigning words to topics and iteratively refining the assignment by maximizing the probability of the data. This is done through a Bayesian estimate of the probability distributions of words and topics given the observed data [9].

The methods described above are commonly used for topic modeling tasks. Several works use them as part of a text categorization process or through comparing methods [9–12]. However, LSA, PLSA, and LDA have several disadvantages, such as fixing the number of topics, lack of capturing non-linear relationships between words, and the assumption that a document can contain different topics, which may not be valid. For this reason, other proposals have been developed that consider the possible dependence on the context of words.

Deep learning is today's most widely used method for natural language processing tasks. Different models and architectures can capture the dependency between the words in a sentence. With this consideration, the task focuses on obtaining an n-dimensional representation of a word or sentence, i.e., an embedding. The use of word embeddings has become increasingly widespread nowadays. The advantage of an n-dimensional representation is that it can be used from different approaches [2].

Some works have focused on using word embeddings for topic modeling tasks [13–15]. They generally work by obtaining the embeddings for the texts of each document and then applying term frequency techniques, clustering, and combinations with LDA. In the present work, this approach is used, combining the extraction of word embeddings through a pre-trained BERT (Bidirectional Encoder Representations from Transformers) model, a UMAP dimensionality reduction technique [16], and the clustering of the elements as the final step to get a set of topics using K-means. Some methods and results of [13] inspired this work.

### **3 Word Embeddings, Dimensionality Reduction, and Topic Modeling**

Word embedding represents a word or text as an n-dimensional vector considering the semantic and syntactic characteristics within the corpus [17]. This type of representation is commonly used in multiple NLP tasks, generally as features that feed some machine learning algorithm. The most usual way to obtain word embeddings is through different types of neural network architectures. It has been verified that better results are obtained using neural networks than techniques based on the frequency of terms [18].

#### **3.1 Transformers and BERT Model for Word Embedding**

One of the most used architectures currently in NLP tasks is the Transformers, particularly a derived model called BERT. Transformers are a type of deep neural network architecture that uses attention mechanisms (weights between input data elements) to select the most important parts of an input sequence,

catching the long-range dependencies of a sentence [19]. They work through an encoder-decoder structure that takes  $X = \{x_1, x_2, \dots, x_N\}$  sequences (tokenized values from the input text) and produces  $Z = \{z_1, z_2, \dots, z_N\}$  representations. The encoder-decoder blocks use multi-head attention mechanisms to capture the dependencies between different tokens in the input sequence [20]. In general, this type of architecture uses a stack of encoder-decoder layer stacks; that is, the information is autoregressive and depends on the previous results of the computation.

BERT comprises a set of Encoder Transformer layers and an attention mechanism to capture relationships between all the words in a sentence. The difference from BERT is that this unit deals with the forwards and backward of a sentence [21]. This model is pre-trained on a large text corpus using two unsupervised learning tasks: masked language modeling and next-sentence prediction [19]. The first part consists of predicting the original masked tokens (tokens generated on the input data) based on the context of the surrounding words. This allows the model to understand relationships between words in a sentence. In the next sentence prediction, the model aims to predict if a pair of sentences follow each other. This helps BERT to understand the relationship between sentences in a document. Since pre-training aims to learn dependencies and relationships between words and sentences, the model can be fine-tuned to perform multiple tasks as word embedding [20].


### 3.2 UMAP for Dimensionality Reduction

The dimensionality reduction plays an essential role in data analysis and visualization tasks. Some of these techniques work through multiple transformations on the input data that are linear combinations of the initial information. Dimensionality reduction techniques generally do not consider the global characteristics of the input data, so the distance of a set of points in the original space will not necessarily be preserved in the reduced space [24].

Uniform Manifold Approximation and Projection (UMAP) is a dimension reduction technique that uses graph layout algorithms to arrange data in a low-dimensional space [16]. The main idea focuses on approximating the high-dimensional data manifold (curved surface embedded in a high-dimensional space). The embedding of the data is computed by searching for a low-dimensional projection of the data with the closest possible equivalent global shape and structure as the original dataset [24], i.e., UMAP preserves the local and global structure of the data. This method considers that there is a manifold on which the data would be uniformly distributed, and the main objective is to preserve the topological structure of this manifold [16].

### 3.3 K-means Clustering

The  $K$ -means algorithm is one of the best-known methods for solving clustering problems.  $K$ -means operates on a set of observations to try to group them into a certain number of sets, considering that the square of the distance between each



**Fig. 1.** Clustering of the reduced embeddings.

element and its cluster center is the minimum possible [25]. We can formally define the  $K$ -means problem as the next optimization problem.


$K$ -means optimization problem: Given a set of elements  $\mathbf{x}_n \in X$ ,  $X \subset \mathbb{R}^D$  and an integer  $k \leq n$  such that there are  $c_k \in C$  subsets. The objective of the  $K$ -means clustering is to partition  $n$  elements into  $k$  sets  $S$  to minimize the within-cluster sum of squares (WCSS), i.e.:

$$\min_{c \in C} \sum_{x \in X} \|x - c\|^2. \quad (1)$$

During each iteration of the algorithm, the following steps are computed: first, the  $k$  cluster centroids are taken randomly from the elements of  $X$ . Next, each data element is assigned to its nearest cluster. New centroids are created by taking the average value of the elements assigned to the cluster. Finally, the difference between the new centroid and the previous one is calculated. The algorithm stops until there are no significant differences between these two values.

## 4 Topic Modeling with Word Embeddings

The process used for topic modeling consists of multiple of steps involving the algorithms described above. A pre-trained BERT model [26, 27] is used as a first



**Fig. 2.** Elbow method to select the number of clusters.

step to obtain the embeddings of a set of texts. Using a pre-trained model has an advantage, such as time and resource efficiency or a broad understanding of the language, because these models are trained on a large amount of text data. Since the pre-trained BERT models to generate embeddings result in high-dimensionality vectors [21], it is necessary to use dimensionality reduction techniques.

That is why as a second step, the UMAP algorithm is applied. The reduced dimensions preserve the original data's local and global structure through nonlinear transformations [16]. BERT generates embeddings considering the context of the words in the corpus; therefore, similar text strings must have similar embeddings in the final space. However, difficulties have been observed in vectors with high dimensionality when using distance metrics [21].

For this reason, using dimensionality reduction techniques such as the one described above is expected. By applying UMAP, this information manages to be preserved due to how the algorithm operates. To this point, generating embeddings with a reduced dimension for the text set to be described is possible. An additional step is needed to use this information to generate topics from a set of documents.

The third step of the process described in this work involves using clustering techniques to assign clusters to the embeddings obtained. K-means is used in this work due to its ease of use and interpretability. The clusters obtained by K-Means on each document's set of embeddings correspond to each text's topics.


## 5 Experimental Results

A dataset of 1212 news items in Spanish extracted from a website [28] was used for the experimental process. As part of the preprocessing, punctuation marks and special characters (\$, #, %, &, among others) were removed. Likewise, common words without significant relevance (stopwords) from Spanish were eliminated in concordance with [29]. For this process, a pre-trained BERT model was used in more than 50 languages, including Spanish. The BERT model, pre-trained in more than 50 languages, is designed to preserve the distance between words with the same meaning in different languages [26, 27]. This model is used because of the limited availability of BERT models trained exclusively in Spanish. Because of the characteristics of the pre-trained model (with English and Spanish languages included), a random sample composed of half of the news was taken, then its translation into English was obtained. This translated sample was incorporated into the original Spanish dataset to assess whether the resulting topics would contain similar words from both languages. The final size of the dataset was 2183 text news.

Once this process was completed, the embeddings were obtained (with the pre-trained BERT model in more than 50 languages) for each news text. The resulting embeddings had a dimensionality of 768, according to what was reported in the literature [21]. Subsequently, the UMAP dimensionality reduction algorithm was used, preserving two dimensions. The reasons for this choice focused on having a direct relationship between the two-dimensional visualization and the subsequent results, figure 1. The K-Means algorithm was used as a third step to obtain clusters on the embeddings of the texts. The heuristic called the elbow method was used to select the number of clusters [30]. The initial centroids required by the method were chosen according to the initialization of the K-Means++ algorithm [22]; this guaranteed that the distance between the initial centroids was distant, which ensured a higher probability of better results. In addition, the method was repeated 100 times, and the best result was selected according to [23]. Figure 2 shows the result of the elbow method applied; it is observed that the best number of clusters is four.

With the clusters obtained, the words of each text were grouped, and the TF-IDF statistic was applied. TF-IDF is a statistic used to evaluate the importance of a term in a document considering a corpus of words [8]. To calculate it, the term frequency TF is used, which is an estimate of the probability of occurrence of a term in the document [31], and the inverse document frequency (IDF), which is the change in the amount of information of a term above all the corpus [31] or the rarity of the term considering the corpus. TF-IDF is the product of TF and IDF components.

With this process, the most important words for each cluster can be obtained. This behavior is shown in 3. The clustering carried out by K-Means does an excellent job since the words seem related in their meaning. The results should also be attributed to the embeddings obtained by BERT and to the dimension reduction process with UMAP that manages to preserve the global and local information of the original BERT vector. Additionally, it is observed




**Fig. 3.** Topics obtained with K-Means. The words and their TF-IDF score are shown.

that some topics contain the same words but in a different language due to the mentioned characteristics of the BERT, model pre-trained in more than 50 languages. Based on these results, topic 0 corresponds to subjects related to digital finance and innovation, topic 1 to economic issues, topic 2 to climate change and alternative energy, and topic 3 to types of government regulations or government agreements.

To statistically evaluate the results obtained by the combination of methods, two metrics extracted from [13] Topic Diversity and Topic Coherence were used. Topic Diversity is used to evaluate the redundancy of words in each document, that is, the number of unique terms with respect to a document or text. The range of values is between 0 and 1, where 1 indicates a great diversity of terms and 0 is more redundant terms or words [13]. Topic Coherence is a metric used to measure the association of terms in a corpus. It works by calculating the probabilities of occurrence of two terms in the same document [32]; this is done for all the terms of the same topic considering the corpus. The range of values goes from  $[-1,1]$ , where 1 indicates a perfect association between the terms; that is, there is a co-occurrence (coherence) of the terms in the document.

Table 1 show the results of Topic Diversity and Topic Coherence for the process described in this work. A comparison with the method described in [13] is also shown. A significant difference with [13] is that the author of that work used a modified version of the DBSCAN clustering technique called HDBSCAN as the algorithm to get the clusters of the embeddings of the texts. Both methods show similar results regarding the Topic Diversity and Topic Coherence metrics. It is also observed that the general context of each topic is similar for both clustering techniques. The parameters selected for HDBSCAN were such that four topics were obtained (just like K-Means). In addition, figure 4 shows the topics created by the HDBSCAN algorithm.





**Fig. 4.** Topics obtained with HDBSCAN. The words and their TF-IDF score are shown.

**Table 1.** Comparison between clustering methods.

Cluster method	Topics	Topic Diversity	Topic Coherence
K-Means	Digital finance, Innovation	0.25	0.53
	Economic issues	0.3	0.50
	Climate change, green energies	0.35	0.31
	Government regulation	0.425	0.11
HDBSCAN	Digital finance, technology	0.24	0.46
	Economic issues	0.32	0.4
	Climate change, green energies	0.38	0.31
	Government regulation	0.4	0.16

## 6 Conclusions

This work combined several algorithms to obtain the topics for a Spanish news dataset. The text was preprocessed, eliminating Spanish stopwords, punctuation marks, and special characters. Likewise, a random sample of half of the dataset was taken and translated into English, and this sample was incorporated into the original dataset. The later was done due to the type of model used to get the text embeddings. With the pre-processed text, a pre-trained BERT model trained on more than 50 languages was used to obtain each text's embeddings. Due to the high dimensionality of the resulting vectors, the UMAP dimensionality reduction technique was applied.

The K-Means algorithm was used to obtain a defined number of clusters. In this case, the number was selected by applying the cubit method. To visualize which sets of words were found within each cluster, the TF-IDF statistic was applied to obtain the most relevant words for each topic. Finally, two metrics related to the diversity of terms and coherence were applied to the topics found.

These results were compared with those obtained in a previous work [13]. A good relationship is observed between the topics and the coherence and diversity of the terms.

Through the process described in this work, good results are obtained in the topic modeling of a set of texts. The described technique could be used in another type of process in real applications where it is necessary to classify a set of unthematic texts.

## References

1. Smith, T.F., Waterman, M.S.: Identification of Common Molecular Subsequences. *J. Mol. Biol.* **147**, 195-197 (1981)
2. Dieng, A., Ruiz, F., Ble, D.: Topic modeling in embedding spaces. *Transactions of the Association for Computational Linguistics* **8**, 439-453 (2020)
3. Vayansky, I., Kumar, S.: A review of topic modeling methods. *Information Systems* **94**, 1-15 (2020)
4. Blei, D.: Probabilistic topic models. *Commun ACM* **55**(4), 77-84 (2012)
5. Hong, L., Davison, B.: Empirical study of topic modeling in twitter. In: *Proceedings of the First Workshop on Social Media Analytics -SOMA*, pp. 80-88 (2010)
6. Ramage, D., Rosen, E., Chuang, J., Manning, C., McFarland, D.: Topic modeling for the social sciences. In: *Workshop on Applications for Topic Models NIPS*, Whistler, pp. 1-4 (2009)
7. Kherwal, P., Bansal, P.: Topic modeling: A comprehensive review. *EAI Endorsed Transactions on Scalable Information Systems* **7**(24), 1-16 (2020)
8. Ferilli, S.: *Automatic digital document processing and management: Problems, algorithms and techniques*. Springer (2011)
9. Alghamdi, R., Alfalqi, K.: A survey of topic modeling in text mining. *International Journal of Advanced Computer Science and Applications* **6**, 147-153 (2015)
10. Kalepalli, Y., Tasneem, S., Teja, P., D, P., Manne, S.: Effective comparison of LDA with LSA for topic modelling. In: *2020 4th International Conference on Intelligent Computing and Control Systems (ICICCS)*, pp. 1245-1250 (2020)
11. George, L., Birla, L.: A study of topic modeling methods, In: *2018 second international conference on intelligent computing and control systems (iciccs)* IEEE, pp. 109-113 (2018)
12. Nallapati, R., Cohen, W.: Link-plsa-lda: A new unsupervised model for topics and influence of blogs. In: *2 (ed.) Proceedings of the International AAAI Conference on Web and Social Media*, pp. 84-92 (2021)
13. Grootendorst, M.: Bertopic: Neural topic modeling with a class-based tf-idf procedure. *Arxiv:2203.05794* (2022)
14. Suhyeon, K., Haecheong, P., L, J.: Word2vec-based latent semantic analysis (w2v-lsa) for topic modeling: A study on blockchain technology trend analysis. *Expert Systems with Applications* **152**, 1-12 (2020)
15. Moody, C.: Mixing dirichlet topic models and word embeddings to make lda2vec. *ArXiv:1605.02019* (2016)
16. McInnes, L., Healy, J.: Umap: Uniform manifold approximation and projection for dimension reduction. *ArXiv:1802.03426* (2018)
17. Yang, L., Zhiyuan, L., Tat-Seng, C., S, M.: Topical word embeddings. In: *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pp. 2418-2424 (2015)

18. Almeida, F., G, X.: Word embeddings: A survey. ArXiv:1901.09069 (2019)
19. Gillioz, A., Casas, J., Mugellini, E., Abou, O.: Overview of the transformer-based models for nlp tasks. In: Proceedings of the Federated Conference on Computer Science and Information Systems 21, pp. 179–183 (2020)
20. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A., Kaiser, U., Polosukhin, I.: Attention is all you need. *Advances in Neural Information Processing Systems* **30**, 1–11 (2017)
21. Devlin, J., Chang, M., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. In: Proceedings of NAACL-HLT, pp. 4171–4186 (2019)
22. Arthur, D., Vassilvitskii, S.: k-means++: the advantages of careful seeding. In: Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1027–1035 (2007)
23. Pasi, F., Sami S.: How much can k-means be improved by using better initialization and repeats?. *Pattern Recognition* **93**, 95–112 (2019)
24. Vermeulen, M., Smith, K., Eremin, K., Rayner, G., Walton, M.: Application of uniform manifold approximation and projection (umap) in spectral imaging of artworks. *Spectrochimica Acta Part A: Molecular and Biomolecular Spectroscopy* **252**, 1–15 (2021)
25. Arthur, D., Vassilvitskii, S.: k-means++: the advantages of careful seeding. In: Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, pp. 1027–1035 (2007)
26. Reimers, N., Gurevych, I.: Making monolingual sentence embeddings multilingual using knowledge distillation. In: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP), Association for Computational Linguistics, pp. 4512–4525 (2020)
27. Reimers, N., Gurevych, I.: Sentence-bert: Sentence embeddings using siamese bert-networks. In: Conference on Empirical Methods in Natural Language Processing, pp. 3982–3992 (2019)
28. Kaggle: Spanish News Classification. <https://www.kaggle.com/datasets/kevinmorgado/spanish-news-classification> Accessed 15 May 2023
29. Steven, B., Loper, E., Klein, E.: *Natural Language Processing with Python*. O'Reilly Media Inc. (2009)
30. Thorndike, R.: Who belongs in the family?. *Psychometrika* **18**, 267–276 (1953)
31. Aizawa, A.: An information-theoretic perspective of tf-idf measures. *Information Processing and Management* **39**, 45–65 (2003)
32. Bouma, G.J.: Normalized (pointwise) mutual information in collocation extraction. In: Proceedings of the Biennial GSCL Conference, pp. 31–40 (2009)