



HAL
open science

ZeroBERTo: Leveraging Zero-Shot Text Classification by Topic Modeling

Alexandre Alcoforado, Thomas Palmeira Ferraz, Rodrigo Gerber, Enzo Bustos, André Seidel Oliveira, Bruno Miguel Veloso, Fabio Levy Siqueira, Anna Helena Reali Costa

► To cite this version:

Alexandre Alcoforado, Thomas Palmeira Ferraz, Rodrigo Gerber, Enzo Bustos, André Seidel Oliveira, et al.. ZeroBERTo: Leveraging Zero-Shot Text Classification by Topic Modeling. Vlória Pinheiro; Pablo Gamallo; Raquel Amaro; Carolina Scarton; Fernando Batista; Diego Silva; Catarina Magro; Hugo Pinto. Computational Processing of the Portuguese Language. 15th International Conference, PROPOR 2022, Fortaleza, Brazil, March 21–23, 2022, Proceedings, 13208, Springer International Publishing, pp.125-136, 2022, Lecture Notes in Computer Science, 978-3-030-98304-8. 10.1007/978-3-030-98305-5_12 . hal-03628242

HAL Id: hal-03628242

<https://telecom-paris.hal.science/hal-03628242>

Submitted on 4 Jun 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

ZeroBERTo: Leveraging Zero-Shot Text Classification by Topic Modeling

Alexandre Alcoforado¹, Thomas Palmeira Ferraz^{1,2}, Rodrigo Gerber¹, Enzo Bustos¹, André Seidel Oliveira¹, Bruno Miguel Veloso³, Fabio Levy Siqueira¹, and Anna Helena Reali Costa¹

¹Escola Politécnica, Universidade de São Paulo (USP), São Paulo, Brazil
{alexandre.alcoforado,rodrigo.gerber,enzobustos,
andre.seidel,levy.siqueira,anna.reali}@usp.br

²Télécom Paris, Institut Polytechnique de Paris, Palaiseau, France
thomas.palmeira@telecom-paris.fr

³Universidade Portucalense & INESC TEC, Porto, Portugal
bruno.m.veloso@inesctec.pt

Abstract

Traditional text classification approaches often require a good amount of labeled data, which is difficult to obtain, especially in restricted domains or less widespread languages. This lack of labeled data has led to the rise of low-resource methods, that assume low data availability in natural language processing. Among them, zero-shot learning stands out, which consists of learning a classifier without any previously labeled data. The best results reported with this approach use language models such as Transformers, but fall into two problems: high execution time and inability to handle long texts as input. This paper proposes a new model, ZeroBERTo, which leverages an unsupervised clustering step to obtain a compressed data representation before the classification task. We show that ZeroBERTo has better performance for long inputs and shorter execution time, outperforming XLM-R by about 12% in the F1 score in the FolhaUOL dataset.

Keywords: Low-Resource NLP · Unlabeled data · Zero-Shot Learning · Topic Modeling · Transformers

1 Introduction

The current success of supervised learning techniques in real-world Natural Language Processing (NLP) applications is undeniable. While these techniques require a good set of labeled data, large corpora of annotated texts are difficult to obtain, as people (sometimes experts) are needed to create manual annotations or revise and correct predefined labels. This problem is even more critical in languages other than English: statistics¹ show that English is used by 63.1 % of the population on the internet, while Portuguese, for instance, is only used by 0.7 %. This scenario has contributed to the rise of Low-Resource NLP, which aims to develop techniques to deal with low data availability in a specific language or application domain [Hedderich et al., 2021].

Recently, the concept of *zero-shot learning* emerged in NLP: a semi-supervised approach in which models can present results equivalent to those of supervised tasks, such as classification in the absence of labeled data. Current approaches to the zero-shot text classification task (0SHOT-TC) make use of the good performance that *Transformers* have demonstrated in text entailment tasks [Yin et al., 2019]. In order to be able to process text in a way that is not uniquely suited to any specific task or data-set, these Transformers are first pre-trained in large general databases (usually taken from Wikipedia)

¹Statistics available at: https://w3techs.com/technologies/overview/content_language.

and then fine-tuned into a small mainstream data-set for the natural language inference task (such as GLUE [Wang et al., 2019] and XNLI [Conneau et al., 2018]). However, the use of models based entirely on Transformers falls into two critical problems: (i) limitation of the maximum size of the input text, and (ii) long run-time for large volumes of data. While there are transformer-based solutions to these problems individually [Beltagy et al., 2020, Sanh et al., 2019, Zaheer et al., 2020], to the best of our knowledge, there is no solution that addresses both, nor even in the context of 0SHOT-TC.

In this paper, we propose a new hybrid model that merges Transformers with unsupervised learning, called ZeroBERTo – Zero-shot BERT based on Topic Modeling –, which is able to classify texts by learning only from unlabeled data. Our contribution not only handles long inputs – not limiting the input size and considering every input token to encode the data – but also offers a faster execution time. We propose an experimental setup with unlabeled data, simulating low-resource scenarios where real-life NLP researchers may find themselves. Then, we compare ZeroBERTo to a fully-Transformer-based zero-shot on a categorization dataset in Portuguese, FolhaUOL². Our results show that our model outperforms the previous one, in the best scenario, with about 12% better *label aware weighted F1-score* and around 13 times faster total time.

The paper is structured as follows: Sect. 2 presents a background of how it is possible to move from data scarcity to zero-shot learning, as well as the related work on getting the best model for the 0SHOT-TC task. Sect. 3 formalizes the ZeroBERTo task and describe its training and inference procedures. Then, Sect. 4 describes the experimental setup that makes it possible to simulate low-resource scenarios to evaluate the proposed model. Finally, the discussion of the results of the experiments along with our final remarks is in Sect. 5.

2 Background and Related Work

The first approach to overcome the shortage of labeled data for classification suggests adopting data augmentation strategies [Jacobs, 1992], relying on methods to generalize from small sets of already annotated data. Still, the problem persists when there is an extreme lack of data. An alternative approach is to treat the task as a topic modeling problem. Topic modeling is an unsupervised learning technique capable of scanning a set of documents, detecting patterns of words and phrases within them, and automatically clustering groups of similar words and expressions that best characterize a set of documents [Chen et al., 2016]. There is usually a later labeling step for these clusters, which can be a problem as their interpretation is sometimes challenging, and a labeling error can affect the entire process. An automatic method for this is in our interest.

The context presented helps explain the growing interest in the field of Low-Resource NLP [Chang et al., 2008, Hedderich et al., 2021], which addresses traditional NLP tasks with the assumption of scarcity of available data. Some approaches to this family of tasks propose semi-supervised methods, such as adding large quantities of unlabeled data to a small labelled dataset [Nigam et al., 2000], or applying cross-lingual annotation transfer learning [Bentivogli et al., 2004] to leverage annotated data available in languages other than the desired one. Other approaches try to eliminate the need for annotated data for training, relying, for example, on pre-trained *task-agnostic* neural language models [Meng et al., 2020], which may be used as language information sources, as well as representation learning models [Ji and Eisenstein, 2014] for word, sentence or document classification tasks.

Recent breakthroughs in pre-trained neural models have expanded the limits of what can be done with data shortage. The Transformer model [Vaswani et al., 2017], which relies solely on attention mechanisms for learning, followed by BERT [Devlin et al., 2019] – a pre-trained Transformer encoder capable of deep bidirectional encoding – offered the possibility of using general-purpose models with previous language understanding. With little to no fine-tuning, BERT-like models have been successfully applied to most natural language comprehension tasks [Logeswaran et al., 2020], and also show a significant reduction in the need for training data [Brown et al., 2020]. Such models tend to work better for the 0SHOT-TC task, as they carry information about the context and semantic attributes within their pre-trained parameters. On the downside, pre-trained Transformers are complex, with millions of trainable parameters and slow processing of large quantities of data, and due to memory issues, most pre-trained Transformers cannot process inputs larger than 512 tokens at a time.

²Available at: <https://www.kaggle.com/marlesson/news-of-the-site-folhauol>.

Also, attention models have another problem related to input size: attention cannot keep track of all information present in a large text, worsening the performance of the models.

In this context, zero-shot learning approaches stand out [Socher et al., 2013]. A simple way to explain zero-shot is to compare its paradigm with humans’ ability to recognize classes of objects by having a high-level description of them without seeing an example of the object previously. Yin et al. [2019] defines that 0SHOT-TC aims to learn a classifier $f : X \rightarrow Y$, whereas classifier $f(\cdot)$, however, does not have access to data X specifically labeled with classes from Y . We can use the knowledge that the model already has to learn an intermediate layer of semantic attributes, which is then applied at inference time to recognize unseen classes during the training stages [Zhang et al., 2019].

Several works that seek to improve the performance of zero-shot learning inspired ours. Li et al. [2015] worked in the image domain, seeking to develop a two-stage model that first learns to extract relevant attributes from images automatically and then learns to assign these attributes to labels. Our proposal performs the same two steps for the text classification problem but does not use any specific knowledge of external data or require any labelled data.

In the text domain, Mekala and Shang [2020] defines weak-supervised learning similar to our definition of zero-shot learning. With unlabeled data and a list of classes as inputs, it applies seed word lists to guide an interactive clustering preview. Meng et al. [2020] uses topic mining to find out which words have the same semantic meaning as the proposed labels, and with that makes a fine-tuning of the language model assuming the implicit category as the presence of these words in the text. Unlike these approaches, our model does not require the user to have any seed word for the labels, and instead of automatically learning them from the labels themselves, ZeroBERTo discovers them from the input data through topic modeling and then assigns them to the labels based on the language model used.

3 Proposed Method

In this section, we introduce ZeroBERTo which leverages Topic Modeling and pre-trained Language Models (LMs) for the task of zero-shot multi-class text classification (0SHOT-TC).

3.1 0SHOT-TC Task Formalization

Given a set of unlabeled documents $\mathcal{D} = \{d_1, d_2, \dots, d_n\}$ and a set of m semantically disjoint and relevant label names $\mathcal{L} = \{l_1, l_2, \dots, l_m\}$, 0SHOT-TC aims to learn $f : \mathcal{D} \times \mathcal{L} \rightarrow \Theta$, $|\Theta| = |\mathcal{L}|$ and Θ defines a probability $\theta_j^i \in [0, 1]$ for each label l_j being the label for d_i [Yin et al., 2019]. A single-label classification of a document d_i may then be carried out as $l_j \in \mathcal{L} \mid j = \operatorname{argmax}_{(j)} (\theta_1^i, \theta_2^i, \dots, \theta_m^i)$ – as a notation simplification, for now on, we mention this as $\operatorname{argmax}_{(l \in \mathcal{L})} (\Theta_i)$.

Standard approaches to the 0SHOT-TC task treat it as a *Recognizing Textual Entailment (RTE)* problem: given two documents d_1, d_2 , we say “ d_1 entails d_2 ” ($d_1 \Rightarrow d_2$) if a human reading d_1 would be justified in inferring the proposition expressed by d_2 (named *hypothesis*) from the proposition expressed by d_1 [Korman et al., 2018]. In the case of 0SHOT-TC, d_2 is the hypothesis $\mathcal{H}(l_j)$, which is simply a sentence that expresses an association to l_j . For example, in a news categorization problem, a label could be “**sports**” and a hypothesis for it could be “This news is about **sports**”. Creating the hypothesis is essential to make it understandable by a Language Model, and allows us to discover the probability $P(l_j | d_i) = P(d_i \Rightarrow \mathcal{H}(l_j))$, as $P(d_i \Rightarrow \mathcal{H}(l_j))$ can easily be inferred by a LM, using d_i and $\mathcal{H}(l_j)$ as inputs. For the zero-shot text classification task, it calculates the textual entailment probability of each possible label. This inference, however, is quite demanding computationally.

3.2 ZeroBERTo

ZeroBERTo works differently: instead of processing the entire document in the LM, it learns a compressed data representation in an unsupervised way and only processes this representation in the LM. Thus, it is possible to obtain better performance with large inputs and shorter total time than the standard model, even considering the training time added by the unsupervised step.

To learn this representation, ZeroBERTo uses a statistical model, named **Topic Modeling (TM)**, which examines documents and discovers, from the statistics of the words that occur in them, which abstract “topics” are covered, discovering hidden semantic structures in a text body. Given a set

of unlabeled documents \mathcal{D} , TM aims at learning a set of topics \mathcal{T} . A topic $t \in \mathcal{T}$ is a list of q words or n -grams that are characteristic of a cluster but not of the entire documents set. Then, TM also learns how to represent any document $d_i \in \mathcal{D}$ as a composition of topics expressed by $\Omega_{TM}(d_i) = (\omega_1^i, \omega_2^i, \dots, \omega_k^i)$, such that ω_k^i denotes the probability of a document d_i belonging to a topic t_k .

With this in place, instead of analyzing the relation between document d_i and label l_j , we determine the entailment between the learned topic representation $\Omega_{TM}(d_i)$ of each document and each label l_j . Topics found are given as input to the LM, as a text list of words/expressions that represent the topic, in order to infer entailment probabilities. If the topic representation was learnt properly, then we can assume independence between l_j and d_i given a topic t_k , therefore $P(l_j|t_k, d_i) = P(l_j|t_k) = P(t_k \Rightarrow \mathcal{H}(l_j))$. We then solve the 0SHOT-TC task by calculating the compound conditional probability

$$\theta_i^j = P(l_j|d_i) = \sum_{t_k \in \mathcal{T}} P(l_j|t_k) * P(t_k|d_i) = \sum_{t_k \in \mathcal{T}} P(t_k \Rightarrow \mathcal{H}(l_j)) * \Omega_{TM}^k(d_i) \quad (1)$$

for each label l_j to determine $\Theta_i = (\theta_i^1, \theta_i^2, \dots, \theta_i^m)$. Classification is then carried out by selecting $\text{argmax}_{l \in \mathcal{L}}(\Theta_i)$.

Algorithm 1: Given a set of documents \mathcal{D} , a set of labels \mathcal{L} , a hypothesis template \mathcal{H} , a topic model TM and a Language model LM as input, ZeroBERTo-training (see Alg. 1) returns a trained model \mathcal{Z} . For that, it trains TM on \mathcal{D} using $TM.FIT$ (line 2), that learns the topic representation of those documents. Then, it applies $LM.PREDICT$ for all topics learned in TM (lines 4 to 7). This function, given a topic t_k , returns the set of probabilities $P(t_k \Rightarrow \mathcal{H}(l_j))$ for all $l_j \in \mathcal{L}$. In the end, the model \mathcal{Z} gathers all information learned from \mathcal{D} .

Algorithm 2: ZeroBERTo-prediction leverages a trained model \mathcal{Z} and a specific document d_i to return the predicted label $l \in \mathcal{L}$ (see Alg. 2). For this, it uses $\mathcal{Z}.TM.TOPICENCODER$ (line 1), that returns the topic representation $\Omega_{TM}(d_i)$ of d_i . This was learned by $\mathcal{Z}.TM$ in Alg. 1. Then, it calculates the equation (1) for all candidate labels (lines 2 to 8), returning the one with maximum probability.

Algorithm 1 ZeroBERTo-training

Require: $\mathcal{D}, \mathcal{L}, \mathcal{H}, TM, LM$

Ensure: \mathcal{Z}

```

1: create  $\mathcal{Z}$            ▷ Instantiate ZeroBERTo
2:  $TM.FIT(\mathcal{D})$        ▷ Topic Model Training
3:  $\mathcal{P} \leftarrow \{\}$ 
4: for each  $t_k \in TM.topics$  do
5:    $p_k \leftarrow LM.predict(t_k, \mathcal{H}(\mathcal{L}))$ 
6:    $\mathcal{P} \leftarrow \mathcal{P} \cup \{p_k\}$ 
7: end for
8:  $\mathcal{Z}.TM \leftarrow TM$ 
9:  $\mathcal{Z}.\mathcal{P}, \mathcal{Z}.\mathcal{L} \leftarrow \mathcal{P}, \mathcal{L}$ 
10: return  $\mathcal{Z}$ 

```

Algorithm 2 ZeroBERTo-prediction

Require: \mathcal{Z}, d_i

Ensure: l

```

1:  $\Omega_{TM}^i \leftarrow \mathcal{Z}.TM.TOPICENCODER(d_i)$ 
2:  $\Theta_i \leftarrow \{\}$ 
3: for each  $l_j \in \mathcal{Z}.\mathcal{L}$  do
4:    $\theta_j^i \leftarrow 0$ 
5:   for each  $t_k \in \mathcal{Z}.TM.topics$  do
6:      $\theta_j^i \leftarrow \theta_j^i + (\mathcal{P}(t_k) * \Omega_{TM}^k(t_k))$ 
7:   end for
8:    $\Theta_i \leftarrow \Theta_i \cup \{\theta_j^i\}$ 
9: end for
10: return  $\text{argmax}_{l \in \mathcal{L}}(\Theta_i)$ 

```

4 Experiments

In this section, we present the experiments performed to validate the effectiveness of ZeroBERTo. Considering that it would be difficult to evaluate our model in a real low-resource scenario, we propose an experimental setup to simulate low-resource situations in labeled datasets. We compare ZeroBERTo with the XLM-R Transformer, fine-tuned only on the textual entailment task. To perform the unsupervised training and evaluation, we use FolhaUOL dataset³.

4.1 Dataset

The FolhaUOL dataset is from the Brazilian newspaper ‘‘Folha de Sao Paulo’’ and consists of 167,053 news items labeled into journal sections (categories) from January 2015 to September 2017.

³Available at: <https://www.kaggle.com/marlesson/news-of-the-site-folhauol>.

Table 1: Number of articles by news category within FolhaUOL dataset after cleaning and organizing the data.

Category	# of articles	Category	# of articles
Poder e Política	22022	Educação	2118
Mercado	20970	Turismo	1903
Esporte	19730	Ciência	1335
Notícias dos Países	17130	Equilíbrio e Saúde	1312
Tecnologia	2260	Comida	828
TV, Televisão e Entretenimento	2123	Meio Ambiente	491

Categories too broad, that do not have a semantic meaning associated with a specific context (as the case of “editorial” and “opinion”), were removed from the dataset keeping only the categories presented in Table 1. For each news article, we take the concatenation of its title and content as input. Table 1 presents the data distribution by category.

4.2 Models

We compare our proposal to the XLM-R model.

XLM-R is the short name for XLM-RoBERTa-large-XNLI, available on Hugging Face⁴, which is state of the art in Multilingual 0SHOT-TC. It is built from XLM-RoBERTa [Conneau et al., 2020] pre-trained in 100 different languages (Portuguese among them), and then fine-tuned in the XNLI [Conneau et al., 2018] and MNLI [Williams et al., 2018] datasets (which do not include the Portuguese language). It is already in the zero-shot learning configuration described by Yin et al. [2019] with template hypothesis as input. The template hypothesis used was “*O tema principal desta notícia é {}*” and texts larger than the maximum size of XLM-R (512 tokens) are truncated.

ZeroBERTo The implementation of our model here makes use of BERTopic [Grootendorst, 2020] with M-BERT-large (Multilingual BERT) [Devlin et al., 2019] as topic modeling step, and the same XLM-R described above as the Transformer for associating the topic representation of each document to labels. Repeating the use of XLM-R seeks to make the comparison fair. BERTopic’s hyperparameters are: interval n for n -grams to be considered in the topic representation ($n_grams_range \in \{1, 2, 3\}$); number of representative words/ n -grams per topic ($top_n_words = 20$); and minimum amount of data in each valid topic ($min_topic_size = 10$). The XLM-R template hypothesis used is “*O tema principal desta lista de palavras é {}*”.

4.3 Evaluation

To simulate real-world scenarios, we propose a variation of stratified k -fold cross-validation [Refaeilzadeh et al., 2009]. First, we split the data into k disjoint stratified folds, *i.e.* the data were evenly distributed in such a way as to make the distribution of the classes in the k folds follow the distribution in the entire dataset. Next, we use these k -folds to perform the following 4 experiment setups:

Exp. 1 - Labeling a dataset: Simulates a situation where one needs to obtain the first labeling of a dataset. ZeroBERTo is trained in $(k - 1)$ folds and has the performance compared to XLM-R in the same $(k - 1)$ folds, in order to assess its ability to label data already seen. Since this is unsupervised learning, evaluating the model’s labeling ability in the training data makes sense as it was not exposed to the ground truth labels.

Exp. 2 - Building a model for additional inferences: Simulates a situation where the researcher wants to create a current model in a real-life application without having data labeled for it. ZeroBERTo is trained in $(k - 1)$ folds and can infer new data compared to XLM-R on the remaining fold.

⁴Available at: <https://huggingface.co/joeddav/xlm-roberta-large-xnli>

Table 2: Table shows the results of the experiments for the FolhaUOL dataset. P is weighted-average Precision, R is weighted-average Recall, and F1 is weighted-average F1-score.

	Exp. 1		Exp. 2		Exp. 3		Exp. 4	
	XLM-R	ZeroBERTo	XLM-R	ZeroBERTo	XLM-R	ZeroBERTo	XLM-R	ZeroBERTo
P	0.47 ± 0.00	0.66 ± 0.01	0.46 ± 0.01	0.16 ± 0.08	0.46 ± 0.01	0.64 ± 0.01	0.47 ± 0.00	0.29 ± 0.17
R	0.43 ± 0.00	0.54 ± 0.01	0.43 ± 0.00	0.21 ± 0.05	0.43 ± 0.00	0.56 ± 0.02	0.43 ± 0.00	0.31 ± 0.12
F1	0.43 ± 0.00	0.54 ± 0.01	0.42 ± 0.01	0.15 ± 0.07	0.42 ± 0.01	0.52 ± 0.02	0.43 ± 0.00	0.19 ± 0.17
Time	61h30min	9h21min	15h22min	6h20min	15h22min	1h10min	61h30min	2h25min




Figure 1: Figure shows text entailment results between topics (X-axis) and labels (Y-axis) for the first 50 Topics (sorted by size) in fold 0 from Experiment 3. In total, 213 topics were generated in this experiment.

Exp. 3 - Labeling a smaller dataset: Simulates situation of scarcity of data in which, besides not having labeled data, little data is present. ZeroBERTo is trained in one fold and compared to XLM-R in the same fold. Considering the topic-representation learning stage, the presence of little data could be a bottleneck for ZeroBERTo since the topic representation may not be properly learned.

Exp. 4 - Building model for additional inferences but with a scarcity of training data: simulates again how the model would behave in a real-life application with few training data. ZeroBERTo is trained in 1 fold and compared to XLM-R in the remaining $k - 1$ folds.

We evaluated the performance of both models for each experiment with the following label-aware metrics: weighted-average Precision (P), weighted-average Recall (R), and weighted-average F1-score (F1). For the k -fold CV, we use $k = 5$. Experiments were run on an Intel Xeon E5-2686 v4 2.3GHz 61 GiB CPU and an NVIDIA Tesla K80 12 GiB GPU using the PyTorch framework. To run XLM-R, we use batches sized 20 to prevent GPU memory overflow.

4.4 Results

Table 2 shows the results of the proposed experiments. Time for ZeroBERTo considers unsupervised training time and inference time. Further, as no training is required, only a single run of XLM-R was done on all data. Thus, the times for XLM-R are estimated. Nevertheless, in all experiments, the total time (training + execution) of ZeroBERTo was much lower than the execution time of XLM-R. Our model surpassed XLM-R in all metrics in the experiments in which the evaluation was performed on the data used in the unsupervised training (Exp. 1 and 3). Figure 1 presents a visualization for the entailment mechanism between topics and labels represented by term $P(l_j|t_k) = P(t_k \Rightarrow \mathcal{H}(l_j))$ in equation (1). The darker the green, the greater the conditional odds.

5 Discussion and Future Work

The experiments simulated low-resource scenarios where a zero-shot text classifier can be useful. The results showed that it is possible to obtain a better performance in the 0SHOT-TC task with the addition of an unsupervised learning step that allows a simplified representation of the data, as proposed by ZeroBERTo. Moreover, the proposed model presents itself as an excellent tool to help researchers deal with low-resource scenarios, such as the need to label an entire dataset without any previously labeled training data. Another interesting feature is that the model showed evidence of robustness for smaller amounts of data. In experiment 3, it was trained with 25% of the data from

experiment 1 and got similar performance metrics in lower time, refuting our concern that little data could be a bottleneck in the model.

However, for configurations where ZeroBERTo was tested simulating real-life applications (Exp. 2 and 4), being exposed to new data, the performance was worse than XLM-R. The results suggest it occurs due to the inability of the embedded topic model to adequately represent new data as a composite of previously learned topics, overfitting training data. This is clear from observing the high variance of the metrics among the k-folds. It allows us to conclude that, for now, the scenarios presented in experiments 1 and 3 are more suitable for using the model.

We have 0.54 of F1-score in the best scenario regarding the metrics obtained. Despite being a positive result considering that it is a multi-class classification, there is much room for improvement. The main reason to be pointed out for low performances is the use of multilingual models that were not fine-tuned in the Portuguese language, which is quite impressive.

A critical remark to be made is concerning the memory and time trade-off. For example, ZeroBERTo was more than 10x faster than XLM-R in Exp. 3. However, the topic model used by ZeroBERTo bases its clustering on the HDBSCAN method, which reduces time taken for data processing but increases the need for memory [McInnes and Healy, 2017], which XLM-R does not do. As the size of input data grows, processing may become unfeasible. XLM-R, on the other hand, does not use any interaction between data and can be processed in parallel and distributed without any negative effect on the final result. It should be noted, however, that ZeroBERTo does not depend on BERTopic and can use other Topic Modeling techniques that address this issue more adequately in other scenarios.

A significant difficulty of this work was that, as far as the authors are aware of, there are no large benchmark datasets for multi-class text classification in Portuguese, nor general use datasets with semantically meaningful labels. In this sense, some future work directions involve the production of benchmark datasets for Portuguese text classification (and OSHOT-TC). It would also be interesting to produce Natural Language Inference datasets in Portuguese, which could, in addition to the existing ones [Fonseca et al., 2016, Real et al., 2020], enable fine-tuning of Transformers 100 % in Portuguese. Then, it would be possible to compare the performance of the models using BERTimbau (BERT-Portuguese) [Souza et al., 2020] both in clustering and classifying. It would also be worthwhile to test the proposed model in other domains: to name one, legislative data present similar challenges [Ferraz et al., 2021]. Another interesting future work would be to enable ZeroBERTo to deal with multi-label classification, where each document can have none, one or several labels.

Acknowledgments

This research was supported in part by *Itaú Unibanco S.A.*, with the scholarship program of *Programa de Bolsas Itaú* (PBI), and by the *Coordenação de Aperfeiçoamento de Pessoal de Nível Superior* (CAPES), Finance Code 001, CNPQ (grant 310085/2020-9), and *USP-IBM-FAPESP Center for Artificial Intelligence* (FAPESP grant 2019/07665-4), Brazil. Any opinions, findings, and conclusions expressed in this manuscript are those of the authors and do not necessarily reflect the views, official policy, or position of the financiers.

References

- Iz Beltagy, Matthew E Peters, and Arman Cohan. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*, 2020.
- Luisa Bentivogli, Pamela Forner, and Emanuele Pianta. Evaluating Cross-Language Annotation Transfer in the MultiSemCor Corpus. In *Proceedings of the 20th International Conference on Computational Linguistics*, COLING ‘04, page 364–es. ACL, 2004.
- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.
- Ming-Wei Chang, Lev-Arie Ratinov, Dan Roth, and Vivek Srikumar. Importance of Semantic Representation: Dataless Classification. In *Proceedings of the 23rd AAAI Conference on Artificial Intelligence*, volume 2, pages 830–835, 2008.

- Qiuxing Chen, Lixiu Yao, and Jie Yang. Short text classification based on LDA topic model. In *2016 International Conference on Audio, Language and Image Processing (ICALIP)*, pages 749–753. IEEE, 2016.
- Alexis Conneau, Ruty Rinott, Guillaume Lample, Adina Williams, Samuel Bowman, Holger Schwenk, and Veselin Stoyanov. XNLI: Evaluating Cross-lingual Sentence Representations. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2475–2485. Association for Computational Linguistics, 2018.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Édouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. Unsupervised Cross-lingual Representation Learning at Scale. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, 2020.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, 2019.
- Thomas Palmeira Ferraz, Alexandre Alcoforado, Enzo Bustos, André Seidel Oliveira, Rodrigo Gerber, Naíde Müller, André Corrêa d’Almeida, Bruno Miguel Veloso, and Anna Helena Reali Costa. DEBACER: a method for slicing moderated debates. In *Anais do XVIII Encontro Nacional de Inteligência Artificial e Computacional*, pages 667–678. SBC, 2021.
- E Fonseca, L Santos, Marcelo Criscuolo, and S Aluisio. ASSIN: Avaliacao de similaridade semantica e inferencia textual. In *12th International Conference on Computational Processing of the Portuguese Language, PROPOR*, pages 13–15, 2016.
- Maarten Grootendorst. BERTopic: Leveraging BERT and c-TF-IDF to create easily interpretable topics., 2020.
- Michael A Hedderich, Lukas Lange, Heike Adel, Jannik Strötgen, and Dietrich Klakow. A Survey on Recent Approaches for Natural Language Processing in Low-Resource Scenarios. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2545–2568, 2021.
- Paul S Jacobs. Joining statistics with NLP for text categorization. In *Third Conference on Applied Natural Language Processing*, pages 178–185, 1992.
- Yangfeng Ji and Jacob Eisenstein. Representation Learning for Text-level Discourse Parsing. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13–24, 2014.
- Daniel Z Korman, Eric Mack, Jacob Jett, and Allen H Renear. Defining textual entailment. *Journal of the Association for Information Science and Technology*, 69(6):763–772, 2018.
- Xin Li, Yuhong Guo, and Dale Schuurmans. Semi-Supervised Zero-Shot Classification with Label Representation Learning. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 4211–4219. IEEE, 2015.
- Lajanugen Logeswaran, Ann Lee, Myle Ott, Honglak Lee, Marc’Aurelio Ranzato, and Arthur Szlam. Few-shot sequence learning with transformers. *arXiv preprint arXiv:2012.09543*, 2020.
- Leland McInnes and John Healy. Accelerated hierarchical density based clustering. In *2017 IEEE International Conference on Data Mining Workshops (ICDMW)*, pages 33–42. IEEE, 2017.
- Dheeraj Mekala and Jingbo Shang. Contextualized weak supervision for text classification. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 323–333, 2020.
- Yu Meng, Yunyi Zhang, Jiaxin Huang, Chenyan Xiong, Heng Ji, Chao Zhang, and Jiawei Han. Text Classification Using Label Names Only: A Language Model Self-Training Approach. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9006–9017, 2020.

- Kamal Nigam, Andrew Kachites McCallum, Sebastian Thrun, and Tom Mitchell. Text classification from labeled and unlabeled documents using EM. *Machine learning*, 39(2):103–134, 2000.
- Livy Real, Erick Fonseca, and Hugo Gonalo Oliveira. The ASSIN 2 Shared Task: a quick overview. In *International Conference on Computational Processing of the Portuguese Language*, pages 406–412. Springer, 2020.
- Payam Refaeilzadeh, Lei Tang, and Huan Liu. Cross-validation. *Encyclopedia of database systems*, 5:532–538, 2009.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*, 2019.
- Richard Socher, Milind Ganjoo, Christopher D Manning, and Andrew Ng. Zero-Shot Learning Through Cross-Modal Transfer. In *Advances in Neural Information Processing Systems*, pages 935–943, 2013.
- Fabio Souza, Rodrigo Nogueira, and Roberto Lotufo. BERTimbau: pretrained BERT models for Brazilian Portuguese. In *Brazilian Conference on Intelligent Systems*, pages 403–417. Springer, 2020.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *7th International Conference on Learning Representations, ICLR 2019*, 2019.
- Adina Williams, Nikita Nangia, and Samuel Bowman. A Broad-Coverage Challenge Corpus for Sentence Understanding through Inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, 2018.
- Wenpeng Yin, Jamaal Hay, and Dan Roth. Benchmarking Zero-shot Text Classification: Datasets, Evaluation and Entailment Approach. 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 3914–3923, 2019.
- Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, et al. Big bird: Transformers for longer sequences. *Advances in Neural Information Processing Systems*, 33, 2020.
- Jingqing Zhang, Piyawat Lertvittayakumjorn, and Yike Guo. Integrating Semantic Knowledge to Tackle Zero-shot Text Classification. In *Proceedings of NAACL-HLT*, pages 1031–1040, 2019.