



Leveraging Large Language Models for Topic Classification in the Domain of Public Affairs

Alejandro Peña¹(✉) , Aythami Morales¹ , Julian Fierrez¹ ,
Ignacio Serna¹ , Javier Ortega-Garcia¹ , Íñigo Puente², Jorge Córdova²,
and Gonzalo Córdova²

¹ BiDA - Lab, Universidad Autónoma de Madrid (UAM), 28049 Madrid, Spain
alejandropenna@uam.es

² VINCES Consulting, 28010 Madrid, Spain

Abstract. The analysis of public affairs documents is crucial for citizens as it promotes transparency, accountability, and informed decision-making. It allows citizens to understand government policies, participate in public discourse, and hold representatives accountable. This is crucial, and sometimes a matter of life or death, for companies whose operation depend on certain regulations. Large Language Models (LLMs) have the potential to greatly enhance the analysis of public affairs documents by effectively processing and understanding the complex language used in such documents. In this work, we analyze the performance of LLMs in classifying public affairs documents. As a natural multi-label task, the classification of these documents presents important challenges. In this work, we use a regex-powered tool to collect a database of public affairs documents with more than 33K samples and 22.5M tokens. Our experiments assess the performance of 4 different Spanish LLMs to classify up to 30 different topics in the data in different configurations. The results shows that LLMs can be of great use to process domain-specific documents, such as those in the domain of public affairs.

Keywords: Domain Adaptation · Public Affairs · Topic Classification · Natural Language Processing · Document Understanding · LLM

1 Introduction

The introduction of the Transformer model [22] in early 2017 supposed a revolution in the Natural Language Domain. In that work, Vaswani *et al.* demonstrated that an Encoder-Decoder architecture combined with an Attention Mechanism can increase the performance of Language Models in several tasks, compared to recurrent models such as LSTM [8]. Over the past few years, there has been a significant development of transformer-based language model architectures, which are commonly known as Large Language Models (LLM). Its deployment sparked

a tremendous interest and exploration in numerous domains, including chatbots (e.g., ChatGPT,¹ Bard,² or Claude³), content generation [2, 16], virtual AI assistants (e.g., JARVIS [20], or GitHub’s Copilot⁴), and other language-based tasks [9–11]. These models address scalability challenges while providing significant language understanding and generation abilities. That deployment of large language models has propelled advancements in conversational AI, automated content creation, and improved language understanding across various applications, shaping a new landscape of NLP research and development. There are even voices raising the possibility that most recent foundational models [1, 12, 13, 21] may be a first step of an artificial general intelligence [3].

Large language models have the potential to greatly enhance the analysis of public affairs documents. These models can effectively process and understand the complex language used in such documents. By leveraging their vast knowledge and contextual understanding, large language models can help to extract key information, identify relevant topics, and perform sentiment analysis within these documents. They can assist in summarizing lengthy texts, categorizing them into specific themes or subject areas, and identifying relationships and patterns between different documents. Additionally, these models can aid in identifying influential stakeholders, tracking changes in public sentiment over time, and detecting emerging trends or issues within the domain of public affairs. By leveraging the power of large language models, organizations and policymakers can gain valuable insights from public affairs documents, enabling informed decision-making, policy formulation, and effective communication strategies. The analysis of public affairs documents is also important for citizens as it promotes transparency, accountability, and informed decision-making.

Public affairs documents often cover a wide range of topics, including policy issues, legislative updates, government initiatives, social programs, and public opinion. These documents can address various aspects of public administration, governance, and societal concerns. The automatic analysis of public affairs text can be considered a multi-label classification problem. Multi-label classification enables the categorization of these documents into multiple relevant topics, allowing for a more nuanced understanding of their content. By employing multi-label classification techniques, such as text categorization algorithms, public affairs documents can be accurately labeled with multiple attributes, facilitating efficient information retrieval, analysis, and decision-making processes in the field of public affairs.

This work focuses on NLP-related developments in an ongoing research project. The project aims to improve the automatic analysis of public affairs documents using recent advancements in Document Layout Analysis (DLA) and Language Technologies. The objective of the project is to develop new tools that allow citizens and businesses to quickly access regulatory changes that affect their

¹ <https://openai.com/blog/chatgpt>.

² <https://blog.google/technology/ai/bard-google-ai-search-updates/>

³ <https://www.anthropic.com/index/introducing-claude>.

⁴ <https://github.com/features/preview/copilot-x>.

present and future operations. With this objective in mind, a system is being developed to monitor the publication of new regulations by public organizations. The block diagram of the system is depicted in Fig. 1. The system is composed of three main modules: *i*) Harvester module based on web scrappers; *ii*) a Document Layout Analysis (DLA) module; and *iii*) a Text Processing module. The Harvester monitors a set of pre-defined information sources, and automatically downloads new documents in them. Then, the DLA module conducts a layout extraction process, where text blocks are characterized and automatically classified, using Random Forest models, into different semantic categories. Finally, a Text Processing module process the text blocks using LLMs technology to perform multi-label topic classification, finally aggregating individual text predictions to infer the main topics of the document.

The full system proposed in Fig. 1 serves us to adapt LLMs to analyze documents in the domain of public affairs. This adaptation is based on the dataset used in our experiments, generated in collaboration with experts in public affairs regulation. They annotated over 92K texts using a semi-supervised process that included a regex-based tool. The database comprises texts related to more than 385 different public affairs topics defined by experts.

From all the analysis tool that can be envisioned in the general framework depicted in Fig. 1, in the present paper we focus in topic classification, with the necessary details of the Harvester needed to explain our datasets and interpret our topic classification results. Other modules such as the Layout Extractor are left for description elsewhere.

Specifically, the main contributions of this work are:

- Within the general document analysis system for analyzing public affairs documents depicted in Fig. 1, we propose, develop, and evaluate a novel functionality for multi-label topic classification.
- We present a new dataset of public affairs documents annotated by topic with more than 33K text samples and 22.5M tokens representing the main Spanish legislative activity between 2019 and 2022.
- We provide experimental evidence of the proposed multi-label topic classification functionality over that new dataset using four different LLMs (including RoBERTa [11] and GPT2 [16]) followed by multiple classifiers.

Our results shows that using a LLM backbone in combination with SVM classifiers suppose an useful strategy to conduct the multi-label topic classification task in the domain of public affairs with accuracies over 85%. The SVM classification improves accuracies consistently, even with classes that have a lower number of samples (e.g., less than 500 samples).

The rest of the paper is structured as follows: In Sect. 2 we describe the data collected for this work, including data preprocessing details. Section 3 describes the development of the proposed topic classification functionality. Section 4 presents the experiments and results of this work. Finally, Sect. 5 summarizes the main conclusions.

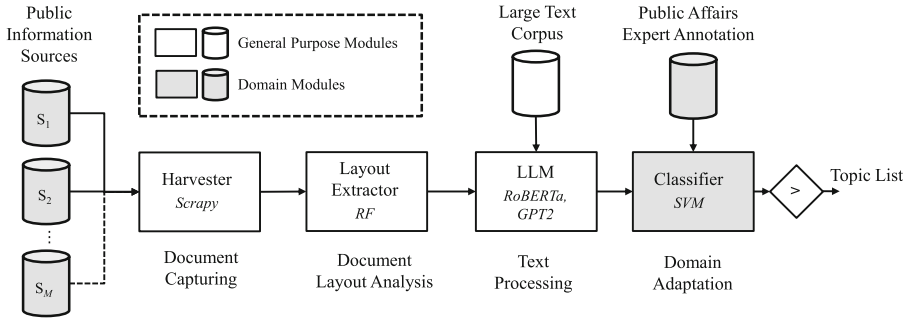


Fig. 1. Block diagram of an automatic public affairs document analysis system. The white blocks represent general-purpose modules, while the grey blocks represent domain-specific modules.

2 Data Collection and Analysis

The major decisions and events resulting from the legislative, judicial and administrative activity of public administrations are public data. Is a common practice, and even a legal requisite, for these administrations to publish this information in different formats, such as governmental websites or official gazettes⁵. Here, we use a regex-powered tool to follow up parliamentary initiatives from the Spanish Parliament, resulting in a legislative-activities text corpora in Spanish. Parliamentary initiatives involve a diverse variety of parliament interactions, such as questions to the government members, legislative proposals, etc.

Raw data were collected and processed with this tool, and comprise initiatives ranging from November 2019 to October 2022. The data is composed of short texts, which may be annotated with multiple labels. Each label includes, among others, topic annotations based on the content of the text. These annotations were generated using regex logic based on class-specific predefined keywords. Both topic classes and their corresponding keywords were defined by a group of experts in public affairs regulations. It is important to note that the same topic (e.g., “*Health Policy*”) can be categorized differently depending on the user’s perspective (e.g., citizens, companies, governmental agencies). We have simplified the annotation, adding a ID number depending on the perspective used (e.g., “*Health Policy_1*” or “*Health Policy_2*”). Our raw data is composed of 450K initiatives grouped in 155 weekly-duration sessions, with a total number of topic classes up to 385. Of these 450K samples, only 92.5K were labeled, which suppose roughly 20.5% of the samples. However, almost half of these are annotated with more than one label (i.e. 45.5K, 10.06% of samples), with a total number of labels of 240K. Figure 2 presents the distribution of the 30 most frequent topics in the data, where we can clearly observe the significant imbalance between classes. The most frequent topic in the raw data is “*Healthcare Situation*”, appearing in more then 25K data samples. Other topics, such as

⁵ <https://op.europa.eu/en/web/forum>.




Fig. 2. Distribution of the top 30 most frequent topics in the raw data.

“*Health Policy*”, have an important presence in the data as well. However, only 8 out of these 30 topics reach 5K samples, and only 5 of them are present in at least 10K. This imbalance, along with the bias towards health-related subjects in the most frequent topics, is inherent to the temporal framework of the database, as the Covid-19 pandemic situation has dominated significant public affairs over the past 3 years. Note that Fig. 2 depicts the thirty most frequent topics, whereas 385 topics are present in the data. To prevent the effects of major class imbalances, we will now focus on the 30 topics of Fig. 2.

2.1 Data Curation

We applied a data cleaning process to the raw corpora to generate a clean version of the labeled data. We started by removing duplicated texts, along with data samples with less than 100 characters. Some works addressing Spanish models applied a similar filtering strategy with a threshold of 200 characters [17, 19, 23] with the aim of obtaining a clean corpus to pre-train transformer models. Here we set the threshold to 100, as our problem here does not require us to be that strict (i.e., we do not want to train a transformer from scratch). Instead, we desired to remove extremely short text, which we qualitative assessed that were mainly half sentences, while retaining as much data as possible. In this sense, we filter text samples of any length starting with lowercase, to prevent half sentences to leak in. We also identified bad quality/noisy text samples to start with “CSV” or “núm”, so we remove samples based on this rule. Finally, given the existence

Table 1. Summary of the parliamentary initiative database after the data cleaning process, which includes 33,147 data samples with multi-label annotations across 30 topics. We include a topic ID, the topic, and the number of samples annotated for each of them.

ID	Topic	#Samples	ID	Topic	#Samples
1	Healthcare Situation	13561	16	Primary Healthcare_1	1425
2	Health Policy_1	12029	17	Sustainability	1370
3	Health Policy_2	8229	18	Wastes	1294
4	Health Policy_3	8111	19	Aids	1216
5	Industrial Emissions	5101	20	Primary Healthcare_2	1189
6	Covid-19 Pandemic	3298	21	Tourism Offer	1181
7	Tourism Policy	2209	22	Labor Equity	1074
8	Tourism Companies	2033	23	Industry	1051
9	Climate Change_1	1930	24	Infrastructures	1029
10	Vaccination	1924	25	Covid-19 Vaccination	997
11	Vaccine	1751	26	National Healthcare System	964
12	Covid-19 Vaccine	1617	27	Climate Change_2	886
13	Tourism Strategy	1533	28	Housing Policy	744
14	Labor Reform	1529	29	Department of Health_1	541
15	Health Innovation	1469	30	Department of Health_2	518

of co-official languages different from Spanish in Spain (e.g., Basque, Galician or Catalan), which are used by a significant percentage of Spanish citizens, we filter data samples from these languages. Due to the lack of reliable language detectors in these co-official languages, and the use of some linguistic, domain-specific patterns in the parliamentary initiatives, we identified a set of words in these languages and use it to detect and filter out potential samples not written in Spanish. We applied this process several times to refine the set of words.

At data sample level, we clean texts by removing excessive white spaces and initiative identifiers in the samples. We then filter URLs and non-alphanumeric characters, retaining commonly used punctuation characters in Spanish written text (i.e., ()-.:?;!_ ;). After applying all the data curation process, we obtain a multi-label corpus of 33,147 data samples, with annotations on the 30 topics commented above. Table 1 presents the number of samples per topic category. Note that the number of samples of each topic has significantly decreased compared to the proportions observed in the raw data (see Fig. 2). The impact of the data curation process is different between topics, leading to some changes in the frequency-based order of the topics. The topic with most data samples in the curated corpus is still “*Healthcare Situation*”, but the number of samples annotated with this topic has been reduced by half. On the other hand, we have several topics with less than 1K samples, setting a lower limit of 518.

3 Methodology and Models

As we previously mentioned in Sect. 2, the samples in our dataset may present more than one topic label. Hence, the topic classification task on this dataset is a multi-label classification problem, where we have a significant number of classes that are highly imbalanced. This scenario (i.e., high number of classes, some of them with few data samples, with overlapped subjects between classes) leads us to discard a single classifier for this task. Instead of addressing the problem as a multi-label task, we break it into small, binary detection tasks, where an individual topic detector is trained for each of the 30 classes in a one vs all setup. This methodology, illustrated in Fig. 3, represents a big advantage, as it provides us a high degree of versatility to select the best model configuration for each topic to deploy a real system. During inference, new data samples can be classified by aggregating the predictions of the individual classifiers [5].

The architecture of the binary topic models is depicted in Fig. 3. We use a transformer-based model as backbone, followed by a Neural Network, Random Forest, or SVM classifier. In this work, we explore different transformer models, pretrained from scratch in Spanish by the Barcelona Supercomputing Center in the context of the MarIA project [7]. We included both encoder and decoder architectures. These model architectures are the following:

- **RoBERTa-base.** An encoder-based model architecture with 12 layers, 768 hidden size, 12 attention heads, and 125M parameters.
- **RoBERTa-large.** An encoder-based model architecture with 24 layers, 71,024 hidden size, 16 attention heads, and 334M parameters.
- **RoBERTalex.** A version [6] of RoBERTa-base, fine-tuned for the Spanish legal domain.
- **GPT2-base.** A decoder-based model architecture with 12 layers, 768 hidden size, 12 attention heads, and 117M parameters.

We listed above the configurations reported in [7] for the open-source models available in the HuggingFace repository of the models.⁶ The RoBERTa models [11] are versions of BERT models [9], in which an optimized pre-training strategy and hyperparameter selection was applied, compared to the original BERT pre-training. The Spanish versions of these models were pre-trained following the original RoBERTa configuration, with a corpus of 570 GB of clean Spanish written text. The RoBERTalex model is a fine-tuned version of Spanish RoBERTa-base, trained with a corpus of 8.9 GB of legal text data. On the other hand, GPT2 [16] is a decoder-based model of the GPT family [2, 12, 13, 15]. As such, the model is aimed to generative tasks (note that modern versions of GPT models, such as InstructGPT [13] or GPT4 [12] are fine-tuned to follow human instructions, so they cannot be considered generative models in the same way as earlier GPT models), different from the RoBERTa family, which is specialized in text understanding. The version used of GPT2 was trained using the same corpus as the RoBERTa models. All the models use byte-level BPE tokenizer [16]

⁶ <https://huggingface.co/PlanTL-GOB-ES>.

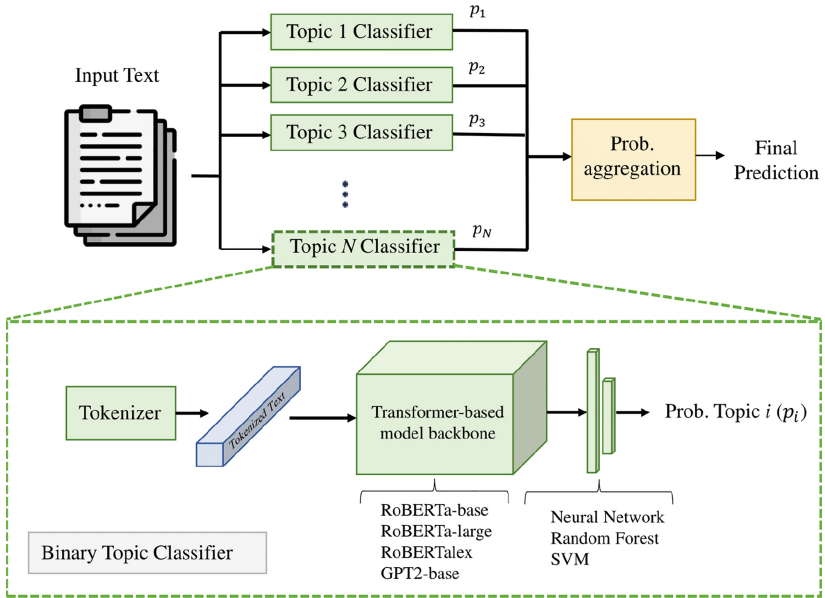


Fig. 3. Proposed multi-label topic classification system, in which an individual topic detector is applied to an input text before aggregating all the predictions, and the architecture of each binary topic classifier.

with vocab size of 50,265 tokens, and have the same length for the context windows, i.e. 512. While left padding is used in the RoBERTa models, right padding is advisable for the GPT2 model.

4 Experiments

As exposed in Sect. 3, due to the nature of the dataset collected for this work, we address multi-label topic classification by training a binary topic classifier for each class (one vs all), and then aggregating the individual predictions on a versatile way (e.g., providing rank statistics, topics over a fixed threshold, etc.). Hence, our experiments will focus on assessing the performance of different topic classifiers configurations, and the potential of the newly available Spanish language models in unconstrained scenarios (i.e., multi-label political data, with subjective annotations based on private-market interest). Section 4.1 will evaluate first the performance of different transformer-based models on our dataset, and then explore the combination of the best-performance model with SVM and Random Forest classifiers.

We conduct all the experiments using a K-fold cross validation setup with 5 folds, and report mean and average results between folds. We select True Positive Rate (TPR), and True Negative Rate (TNR) as our performance measures, due to the class imbalances in the parliamentary dataset. We use in our experiments

the models available in the HuggingFace transformers library⁷, along with several sklearn tools. Regarding the hardware, we conducted the experiments in a PC with 2 NVIDIA RTX 4090 (with 24 GB each), Intel Core i9, 32 GB RAM.

4.1 Topic Classification in the Domain of Public Affairs

Recalling from Fig. 3, our topic detector architecture is mainly composed of *i*) a transformer backbone, and *ii*) a classifier. We train the transformer models with a binary neural network classification output layer. For each topic, we train the detector using Weighted Cross Entropy Loss to address the class imbalance in a “One vs All” setup. Topic classifiers are trained for 5 epochs using a batch size of 32 samples, and freezing the transformer layers. Table 2 presents the results of the topics classifiers using the four transformer models explored in this work (i.e., RoBERTa-base [7], RoBERTa-large [7], RoBERTalex [6], and GPT2-base [7]). We can observe a general behavior across the RoBERTa models. The classifiers trained for the topics with more samples obtain higher TPR means, close to the TNR mean values. In these cases, the classifiers are able to distinguish reasonably well text samples in which the trained topic is present. These results are, in general, consistent across folds, exhibiting moderate deviation values. This behavior degrades from Topic 9 onwards, where the low number of samples (i.e., less than 2K) leads to an increase of the TNR to values over 90% with a decay of TPR. However, we can observe some exceptions in the classifiers using RoBERTa-base as backbone (topics 11, 12, 24), where TNR scales to values close to 100% while preserving TPR performances over 80%. Furthermore, RoBERTa-base classifiers exhibit better results than the RoBERTa-large classifiers (probably due to the constrained number of samples), and even than the RoBERTalex models. Remember that both RoBERTa-base and RoBERTalex are the same models, the latter being the RoBERTa-base model with a fine-tuning to the legal domain that, a priori, should make it more appropriate for the problem at hand. Regarding GPT2-based classifiers, we observe similar trends to those of the RoBERTa models, but exhibiting lower performances. This is not surprising, as the GPT model was trained for generative purposes, rather than text understanding like RoBERTa.

It’s worth noting here the case of Topic 1, which obtains the lowest TNR mean value in all models, with deviation values over 0.15, despite being the topic with more data samples (i.e. a third of the data). We hypothesize that the low performances when detecting negative samples is mostly due to the overlap with the rest of the topics, as this topic focuses on general healthcare-related aspects (remember from Table 1 that half of the topics are related with healthcare).

From the results presented in Table 2, we can conclude that RoBERTa-base is the best model backbone for our task. Now, we want to assess if a specialized classifier, such as Support Vector Machines (SVM) or Random Forests (RF), can be used to fine tune the performance to the specific domain. For these classifiers, we used RoBERTa-base as feature extractor to compute 768-dimensional text

⁷ <https://huggingface.co/docs/transformers/index>.

Table 2. Results of the binary classification for each topic (one vs all), using different transformer models with a Neural Network classifier. We report True Positive Rate (TPR) and True Negative Rate (TNR) as mean_{std} (in parts per unit), computed after a K-fold cross validation (5 folds).

ID	RoBERTa-b [7]		RoBERTa-l [7]		GPT2-b [7]		RoBERTalex [6]	
	TPR	TNR	TPR	TNR	TPR	TNR	TPR	TNR
1	.80 _{.07}	.75 _{.19}	.78 _{.08}	.76 _{0.19}	.58 _{.14}	.60 _{.15}	.79 _{.10}	.70 _{.18}
2	.87 _{.09}	.88 _{.04}	.84 _{.11}	.86 _{.05}	.61 _{.25}	.82 _{.05}	.83 _{.10}	.82 _{.06}
3	.83 _{.08}	.87 _{.04}	.81 _{.09}	.87 _{.04}	.65 _{.18}	.79 _{.07}	.79 _{.09}	.84 _{.05}
4	.86 _{.07}	.89 _{.03}	.83 _{.10}	.88 _{.03}	.69 _{.17}	.79 _{.07}	.80 _{.10}	.86 _{.04}
5	.76 _{.05}	.81 _{.06}	.72 _{.07}	.81 _{.07}	.63 _{.06}	.74 _{.09}	.67 _{.08}	.80 _{.06}
6	.82 _{.05}	.87 _{.02}	.83 _{.05}	.87 _{.03}	.67 _{.04}	.63 _{.08}	.68 _{.06}	.83 _{.04}
7	.85 _{.04}	.93 _{.03}	.83 _{.06}	.91 _{.05}	.64 _{.08}	.78 _{.08}	.75 _{.07}	.94 _{.03}
8	.82 _{.02}	.89 _{.05}	.81 _{.03}	.88 _{.06}	.63 _{.04}	.78 _{.07}	.69 _{.02}	.91 _{.04}
9	.79 _{.10}	.90 _{.04}	.77 _{.11}	.89 _{.06}	.58 _{.08}	.76 _{.08}	.68 _{.07}	.91 _{.03}
10	.76 _{.26}	.96 _{.03}	.67 _{.31}	.95 _{.03}	.49 _{.42}	.91 _{.10}	.62 _{.34}	.95 _{.04}
11	.89 _{.11}	.98 ₀₂	.72 _{.31}	.98 _{.01}	.55 _{.44}	.93 _{.09}	.70 _{.32}	.97 _{.03}
12	.88 _{.12}	.98 _{.02}	.73 _{.30}	.98 _{.01}	.57 _{.41}	.94 _{.09}	.72 _{.30}	.97 _{.02}
13	.76 _{.09}	.89 _{.06}	.75 _{.08}	.86 _{.07}	.33 _{.09}	.79 _{.09}	.58 _{.14}	.91 _{.05}
14	.76 _{.12}	.93 _{.03}	.72 _{.12}	.93 _{.03}	.39 _{.13}	.81 _{.06}	.65 _{.13}	.94 _{.02}
15	.61 _{.09}	.85 _{.04}	.58 _{.10}	.86 _{.03}	.53 _{.06}	.82 _{.05}	.54 _{.08}	.90 _{.02}
16	.75 _{.03}	.90 _{.03}	.71 _{.05}	.88 _{.04}	.43 _{.04}	.77 _{.03}	.64 _{.05}	.91 _{.03}
17	.71 _{.25}	.94 _{.06}	.64 _{.32}	.96 _{.05}	.59 _{.31}	.93 _{.05}	.65 _{.31}	.92 _{.05}
18	.62 _{.08}	.90 _{.03}	.54 _{.10}	.85 _{.05}	.36 _{.07}	.79 _{.07}	.51 ₀₅	.91 ₀₂
19	.69 _{.10}	.92 _{.02}	.69 _{.09}	.91 _{.03}	.45 _{.11}	.86 _{.05}	.49 _{.12}	.95 _{.01}
20	.73 ₀₅	.93 _{.02}	.73 _{.06}	.90 _{.03}	.32 _{.04}	.86 _{.03}	.58 _{.05}	.94 _{.02}
21	.67 _{.04}	.89 _{.05}	.67 _{.06}	.86 _{.06}	.48 _{.05}	.84 _{.05}	.45 _{.03}	.93 _{.03}
22	.71 _{.05}	.95 _{.02}	.66 _{.03}	.94 _{.02}	.40 _{.04}	.89 _{.04}	.51 _{.03}	.97 _{.01}
23	.70 _{.08}	.96 _{.02}	.57 _{.17}	.96 _{.02}	.24 _{.17}	.96 _{.01}	.43 _{.17}	.98 _{.01}
24	.83 _{.08}	.97 _{.04}	.69 _{.11}	.98 _{.01}	.20 _{.24}	.98 _{.01}	.55 _{.18}	.98 _{.01}
25	.80 _{.16}	.97 _{.04}	.54 _{.36}	.97 _{.04}	.44 _{.40}	.98 _{.03}	.57 _{.34}	.97 _{.04}
26	.52 _{.10}	.95 _{.01}	.48 _{.13}	.96 _{.01}	.17 _{.03}	.96 _{.02}	.40 _{.10}	.98 _{.01}
27	.72 _{.08}	.97 _{.02}	.62 _{.08}	.97 _{.02}	.25 _{.07}	.97 _{.01}	.56 _{.05}	.98 ₀₁
28	.44 _{.05}	.97 _{.02}	.32 _{.15}	.96 _{.03}	0 ₀	1 ₀	.20 _{.08}	.99 ₀₁
29	.46 _{.06}	.98 _{.01}	.17 _{.04}	.99 ₀	0 ₀	1 ₀	.18 ₀₄	.99 ₀
30	.43 _{.06}	.98 _{.01}	.15 _{.03}	.99 ₀	0 ₀	1 ₀	.15 _{.03}	.99 ₀

embeddings from each of the text samples. We explored two approaches for these embeddings: *i*) using the embedding computed for the [CLS] token, and *ii*) averaging all the token embeddings (i.e., mean pooling). In the original BERT model [9], and hence the RoBERTa model, the [CLS] is a special token appended at the start of the input, which the model uses during training for the Next

Table 3. Results of the binary classification for each topic (one vs all), using RoBERTa-base [7] in combination with SVM and Random Forest classifiers. We report True Positive Rate (TPR) and True Negative Rate (TNR) as mean_{std} (in parts per unit), computed after a K-fold cross validation (5 folds).

ID	RoBERTa-b [7] + SVM		RoBERTa-b [7] + RF	
	TPR	TNR	TPR	TNR
1	.80 _{.07}	.76 _{.20}	.70 _{.11}	.81 _{.21}
2	.87 _{.09}	.88 _{.04}	.74 _{.18}	.94 _{.03}
3	.83 _{.07}	.88 _{.04}	.64 _{.18}	.97 _{.02}
4	.86 _{.07}	.90 _{.02}	.67 _{.18}	.98 _{.02}
5	.80 _{.05}	.80 _{.06}	.12 _{.06}	.99 _{.01}
6	.85 _{.05}	.85 _{.03}	.23 _{.04}	.99 ₀
7	.90 _{.02}	.90 _{.05}	.49 _{.06}	1 ₀
8	.89 _{.01}	.86 _{.07}	.33 _{.02}	1 ₀
9	.88 _{.07}	.87 _{.04}	.24 _{.05}	1 ₀
10	.84 _{.18}	.94 _{.03}	.51 _{.41}	1 ₀
11	.92 _{.08}	.97 _{.02}	.57 _{.38}	1 ₀
12	.93 _{.07}	.98 _{.02}	.56 _{.36}	1 ₀
13	.87 _{.04}	.85 _{.08}	.08 _{.02}	1 ₀
14	.87 _{.07}	.88 _{.04}	.14 _{.04}	1 ₀
15	.70 _{.08}	.80 _{.06}	.06 _{.02}	1 ₀
16	.89 _{.03}	.88 _{.04}	.13 _{.05}	1 ₀
17	.79 _{.18}	.92 _{.06}	.59 ₃₁	1 ₀
18	.78 _{.06}	.81 _{.05}	.09 _{.03}	1 ₀
19	.87 _{.03}	.85 _{.03}	.14 _{.10}	1 ₀
20	.89 _{.03}	.90 _{.03}	.14 _{.06}	1 ₀
21	.88 _{.02}	.79 _{.08}	.09 _{.02}	1 ₀
22	.90 _{.03}	.88 _{.03}	.16 _{.03}	1 ₀
23	.89 _{.04}	.89 _{.05}	.27 _{.15}	1 ₀
24	.90 _{.05}	.95 _{.02}	.37 _{.23}	1 ₀
25	.90 _{.07}	.95 _{.04}	.41 _{.31}	.99 _{.01}
26	.83 _{.06}	.89 _{.04}	.17 _{.11}	1 ₀
27	.91 _{.04}	.90 _{.04}	.33 _{.04}	1 ₀
28	.87 _{.04}	.86 _{.06}	.06 _{.01}	1 ₀
29	.84 _{.06}	.89 _{.03}	.10 _{.03}	1 ₀
30	.85 _{.05}	.89 _{.03}	.08 _{.03}	1 ₀

Sentence Prediction objective. Thus, the output for this embedding is used for classification purposes, serving the [CLS] embedding as a text representation. We repeated the experiment using both types of representations, and end up selecting the first approach after exhibiting better results. Table 3 presents the results of the topic models using RoBERTa-base text embeddings together with

a SVM and Random Forest classifier. In all cases, we use a complexity parameter of 1 and RBF kernel for the SVM, and a max depth of 1,000 for the Random Forest. We note that these parameters can be tuned for each topic to improve the results. The first thing we notice in Table 3 is the poor performance of the RF-based classifiers, which are the worst among all the configurations. Almost for all the topics under 2K samples, the TNR saturates to 1, and the TPR tends to extremely low values. From this, we can interpret that the classifier is not learning, and just predicting the negative, overrepresented class. However, the performance on the topics over 2K samples is far from the one observed for the RoBERTa models of Table 2. This could be expected, as the RF classifier is not the best approach to work with input data representing a structured vector subspace with semantic meaning, such as text/word embedding subspaces, specially when the number of data samples is low. On the other hand, the SVM performance clearly surpass all previous configurations in terms of TPR. While the results are comparable with those of RoBERTa-base with NN for the first 5 topics, this behavior is maintained for all topics, regardless of the number of data samples. Almost all classifiers achieve a TPR over 80%, except for topics 15, 17 and 18. Nevertheless, the results in these topics increase with the SVM (e.g., for topic 15, where RoBERTa-base with the NN classifier achieved a TPR mean of 61%, here we obtain a 70%). TNR values are, in general, slightly lower, but this could be caused because in previous configurations, topic classifiers tend to exhibit bias towards the negative class as the number of samples falls (i.e., similar to the behavior of the RF classifier). Interestingly, the high deviation observed in the Topic 1 TNR appears too in both SVM and RF classifiers, which could support our previous hypothesis. As we commented before, we suspect that an hyperparameter tuning could improve even more the SVM results on our data.

5 Conclusions

This work applies and evaluates Large Language Models (LLMs) for topic classification in public affairs documents. These documents are of special relevance for both citizens and companies, as they contain the basis of all legislative updates, social programs, public announcements, etc. Thus, enhancing the analysis of public documents using the recent advances of the NLP community is desirable.

To this aim, we collected a Spanish text corpora of public affairs documents, using a regex-powered tool to process and annotate legislative initiatives from the Spanish Parliament during a capture period over 2 years. The raw text corpora is composed of more than 450K initiatives, with 92K of them being annotated in a multi-label scenario with up to 385 different topics. Topic classes were defined by experts in public affairs regulations. We preprocess this corpus and generate a clean version of more than 33K multi-label texts, including annotations for the 30 most frequent topics in the data.

We use this dataset to assess the performance of recent Spanish LLMs [6,7] to perform multi-label topic classification in the domain of public affairs. Our experiments include text understanding models (three different RoBERTa-based

models [11]) and generative models [16], in combination with three different classifiers (i.e., Neural Networks, Random Forests, and SVMs). The results show how text understanding models with SVM classifiers supposes an effective strategy for the topic classification task in this domain, even in situations where the number of data samples is limited.

As future work, we plan to study in more depth biases and imbalances [4] like the ones mentioned before presenting Fig. 2, and compensating them with imbalance-aware machine learning procedures [18]. More recent LLMs can be also tested for this task, including multilingual and instruction-based models, which have shown great capacities in multiple NLP tasks, even in zero-shot scenarios. We will also continue our research by exploring the incorporation of other NLP tasks (e.g. text summarization, named entity recognition) and multimodal methods [14] to our framework, with the objective of enhancing automatic analysis of public affairs documents.

Acknowledgments. This work was supported by VINCES Consulting under the project VINCESAI-ARGOS and BBforTAI (PID2021-127641OB-I00 MICINN/FEDER). The work of A. Peña is supported by a FPU Fellowship (FPU21/00535) by the Spanish MIU. Also, I. Serna is supported by a FPI Fellowship from the UAM.

References

1. Anil, R., Dai, A.M., Firat, O., Johnson, M., et al.: PaLM 2 technical report. [arXiv:2305.10403](https://arxiv.org/abs/2305.10403) (2023)
2. Brown, T., Mann, B., Ryder, N., Subbiah, M., et al.: Language models are few-shot learners. In: NIPS, vol. 33, pp. 1877–1901 (2020)
3. Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., et al.: Sparks of artificial general intelligence: early experiments with GPT-4. [arXiv:2303.12712](https://arxiv.org/abs/2303.12712) (2023)
4. DeAlcala, D., Serna, I., Morales, A., Fierrez, J., et al.: Measuring bias in AI models: an statistical approach introducing N-Sigma. In: COMPSAC (2023)
5. Fierrez, J., Morales, A., Vera-Rodriguez, R., Camacho, D.: Multiple classifiers in biometrics. Part 1: Fundamentals and review. *Inf. Fusion* **44**, 57–64 (2018)
6. Gutiérrez-Fandiño, A., Armengol-Estapé, J., Gonzalez-Agirre, A., Villegas, M.: Spanish legalese language model and corpora. [arXiv:2110.12201](https://arxiv.org/abs/2110.12201) (2021)
7. Gutiérrez-Fandiño, A., Armengol-Estapé, J., Pàmies, M., Llop, J., et al.: MarIA: Spanish language models. *Procesamiento del Lenguaje Nat.* **68** (2022)
8. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* **9**(8), 1735–1780 (1997)
9. Kenton, J., Chang, M., Lee, K., Toutanova, K.: BERT: pre-training of deep bidirectional transformers for language understanding. In: NAACL, pp. 4171–4186 (2019)
10. Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., et al.: BART: denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In: ACL, pp. 7871–7880 (2020)
11. Liu, Y., Ott, M., Goyal, N., Du, J., et al.: RoBERTa: a robustly optimized BERT pretraining approach. [arXiv:1907.11692](https://arxiv.org/abs/1907.11692) (2019)
12. OpenAI: GPT-4 technical report. Technical report (2023)
13. Ouyang, L., Wu, J., Jiang, X., Almeida, D., et al.: Training language models to follow instructions with human feedback. In: NIPS, vol. 35, pp. 27730–27744 (2022)

14. Peña, A., Serna, I., Morales, A., Fierrez, J., et al.: Human-centric multimodal machine learning: recent advances and testbed on AI-based recruitment. *SN Comput. Sci.* **4**, 434 (2023). <https://doi.org/10.1007/s42979-023-01733-0>
15. Radford, A., Narasimhan, K., Salimans, T., Sutskever, I., et al.: Improving language understanding by generative pre-training. Technical report (2018)
16. Radford, A., Wu, J., Child, R., Luan, D., et al.: Language models are unsupervised multitask learners. *OpenAI Blog* **1**(8), 9 (2019)
17. Raffel, C., Shazeer, N., Roberts, A., Lee, K., et al.: Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.* **21**(1), 5485–5551 (2020)
18. Serna, I., Morales, A., Fierrez, J., Obradovich, N.: Sensitive loss: improving accuracy and fairness of face representations with discrimination-aware deep learning. *Artif. Intell.* **305**, 103682 (2022)
19. Serrano, A., Subies, G., Zamorano, H., Garcia, N., et al.: RigoBERTa: a state-of-the-art language model for Spanish. [arXiv:2205.10233](https://arxiv.org/abs/2205.10233) (2022)
20. Shen, Y., Song, K., Tan, X., Li, D., et al.: HuggingGPT: solving AI tasks with ChatGPT and its friends in HuggingFace. [arXiv:2303.17580](https://arxiv.org/abs/2303.17580) (2023)
21. Touvron, H., Lavril, T., Izacard, G., Martinet, X., et al.: LLaMA: open and efficient foundation language models. [arXiv:2302.13971](https://arxiv.org/abs/2302.13971) (2023)
22. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., et al.: Attention is all you need. In: *Advances in Neural Information Processing Systems*, vol. 30 (2017)
23. Xue, L., Constant, N., Roberts, A., Kale, M., et al.: mT5: a massively multilingual pre-trained text-to-text transformer. In: *NAACL*, pp. 483–498 (2021)