

Fine-tuning ClimateBert transformer with ClimaText for the disclosure analysis of climate-related issues in corporates' financial and non-financial reports

Eduardo C. Garrido-Merchán¹, Cristina González-Barthe¹,
María Coronado-Vaca^{1*}

¹Faculty of Economics and Business Sciences, Universidad Pontificia Comillas,
Alberto Aguilera 23, Madrid, 28015, Spain.

*Corresponding author: mcoronado@comillas.edu;

Abstract

There is a growing demand from financial agents (especially particular and institutional investors), for companies to report on climate-related financial risks. A vast amount of information, in text format, is expected to be disclosed in the short term by firms in order to identify these risks in their financial and non-financial reports, particularly in response to the growing regulation that is being passed on the matter. This paper applies state-of-the-art NLP (Natural Language Processing) techniques for the detection of climate change in text corpora. We use transfer learning to fine-tune two transformer models, BERT and ClimateBert -a recently published DistillRoBERTa-based model specifically tailored for climate text classification-. We fine-tune both models on the novel "ClimaText" database, consisting of data collected from Wikipedia, 10K Files Reports, and web-based claims. Our text classification model obtained from the ClimateBert fine-tuning process on ClimaText, outperforms the models created with BERT and the current state-of-the-art transformer in this particular problem. Our study is the first one to implement the recently published ClimateBert algorithm on the ClimaText database. Based on our results, it can be said that ClimateBert fine-tuned on ClimaText is an outstanding tool within the NLP pre-trained transformer models that may and should be used by investors, institutional agents, and companies themselves to monitor the disclosure of climate risk in financial reports,

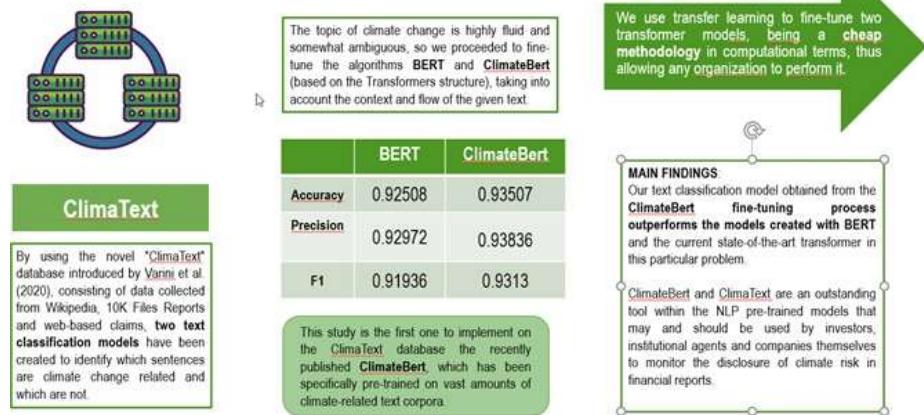
as well as in other textual sources (policies or new legislation passed on this matter. Furthermore, our transfer learning methodology is cheap in computational terms, thus allowing any organization to perform it.

Fig.1. Graphical abstract

Keywords: corporate sustainability in financial reporting, climate-related financial risks, green finance, ClimaText, ClimateBert, Natural Language Processing (NLP)

Fig. 1 Graphical abstract

Climate change topic detection: Fine-tuning ClimateBert on ClimaText



This is the **submitted manuscript** at *Neural Computing and Applications* (Springer Nature).

You can **access a full-text view-only version of the published paper**
by using this link: <https://rdcu.be/eZZtF>

If your institution has a electronic subscription to *Neural Computing and Applications*, you can **download the published paper** from the journal website: Access to the Journal website: <https://doi.org/10.1007/s00521-025-11820-3>

1 Introduction

In recent years, there has been a growing demand from financial agents, especially from particular and institutional investors, for companies to report on climate-related financial risks. There is also a recent trend towards mandatory rather than voluntary reporting worldwide. Thus, a vast amount of information, in text format, can be expected to be disclosed in the short term by firms in order to identify these types of risks in their financial and non-financial reports, particularly in response to the growing regulation that is being passed on the issue.

The Financial Stability Board (FSB) created the Task Force on Climate-Related Financial Disclosures (TCFD) in 2015 at the request of the G-20 leaders, to improve and increase the disclosure of climate-related financial information. In 2017, the TCFD released climate-related financial disclosure recommendations designed to help companies provide better information to support informed capital allocation [1]. The latest TCFD status report describes the steady increase in climate-related financial disclosures since 2017. As of November 2022, the number of TCFD supporters surpasses 4,000 organizations from more than 100 countries with a combined market capitalization of \$27 trillion [2]. Many international organizations have developed various frameworks to help companies in their disclosure of environmental, social, and governance (ESG) issues, including the GRI (Global Reporting Initiative), CDP (Carbon Disclosure Project), SASB (Sustainability Accounting Standards Board), and some others.

The European Union (EU) became the first jurisdiction to require its largest companies to provide information on environmental, personnel and human rights, and sustainability issues [3]. The European Non-Financial Reporting Directive (European NFRD, 2014) [4] obliges companies subject to it to prepare a Report of Non-Financial Information (NFIR), through which to provide this information. Many EU countries have already transposed this NFR Directive into their national law (Austria, Germany, Hungary, Ireland, Italy, The Netherlands, Poland, Slovenia, Spain, etc.). The EU's Sustainable Finance Disclosure Regulation (SFDR) came into force on March 10th, 2021 [5] and EU Taxonomy disclosures have been required starting on Dec 31st, 2021 [6]. On November 16, 2022, the European Financial Reporting Advisory Group (EFRAG) approved the updated versions of the European Sustainability Reporting Standards (ESRS) [7]. The new standards are part of Europe's Corporate Sustainability Reporting Directive (CSRD) which was approved on November 28, 2022. The CSRD is expected to enter into force for the reporting year 2024, with the first submissions due in 2025 [8]. It aims to strengthen sustainability reporting requirements under the existing NFRD to improve corporate accountability, as well as the quality, consistency, and comparability of information disclosed. This new Directive modernizes and strengthens the rules about the social and environmental information that companies have to report. A broader set of large companies, as well as listed SMEs, will now be required to report on sustainability – approximately 50 000 companies in total. In France, institutional investors already face their own disclosure requirements on climate risk management, and how they are contributing to the transition energy. Article 173 of the French Energy Transition Law, came into force in January 2016. It has been

a pioneer in introducing the obligation for institutional investors to inform on the carbon footprint of their investments. Investors with a balance of more than 500 million euros were obliged to submit their first reports for 2016, in June 2017 [9]. The United Kingdom (UK) has become the first G20 country to make it mandatory for Britain's largest businesses to disclose their climate-related risks and opportunities, in line with TCFD recommendations, with new legislation that came into force from April 2022 [10]. New Zealand also makes climate-related disclosures mandatory, from January 2023, for large publicly listed companies, insurers, banks, non-bank deposit takers, and investment managers [11]. The United States (US) Securities and Exchange Commission (SEC) first addressed disclosure of material environmental issues in the early 1970s [12]. The SEC released a proposal in March 2022 that would require companies to disclose climate-related disclosures to investors [13].

On 3 November 2021, at COP26, the International Financial Reporting Standards (IFRS) Foundation announced the creation of the International Sustainability Standards Board (ISSB) to deliver a global baseline of sustainability-related disclosures to meet capital market needs of high-quality, transparent, reliable, and comparable reporting by companies on climate and other environmental, social and governance (ESG) issues. In March 2022, the ISSB published Exposure Draft IFRS S2 Climate-related Disclosures and, by July 2022, received more than 600 responses to it [14]. The ISSB is finalizing requirements for an entity to disclose information about its climate-related risks and opportunities and expects to issue an IFRS Sustainability Disclosure Standard around the end of Q2 2023. The ISSB's work is backed by the G7, the G20, IOSCO (International Association of Securities Commissions), the FSB, African Finance Ministers and by Finance Ministers and Central Bank Governors from over 40 jurisdictions.

According to the World's Economic Forum Global Risk Report 2023, climate and environmental risks are the core focus of global risks perceptions over the next decade – and are the risks for which we are seen to be the least prepared. All the six environmental risks feature in the top 10 risks over the next 10 years, with “failure to mitigate climate change”, “failure of climate-change adaptation”, “natural disasters and extreme weather events”, and “biodiversity loss and ecosystem collapse” having ranked as top four most critical risks [15].

In this context, Carney (2015) already stated in his speech “Breaking the tragedy of the horizon - climate change and financial stability” that climate-related financial risks were being belittled by financial agents, bearing the risk of triggering the next financial crisis [16].

Climate change as a financial risk is a growing concern among several financial agents seeking ways of assessing and measuring this threat. To this effect, Ilhan et al. [17] carried out a survey among institutional investors of which one third of the respondents hold an executive position at their institution and 11% work for firms that amount more than \$100bn in assets. This survey showed came that 51% of these investors believe that reporting on climate risk is as important as financial reporting and almost one third considered the first even more relevant. There is a growing demand from investors for climate risk reporting as well as a higher disposition to engage with organizations that do carry out such a disclosure. Our study is framed within this context of both, growing awareness and demand for financial reporting of climate-related risks, and of the recent trend towards mandatory rather than voluntary

reporting worldwide.

The growing importance of these disclosures, with their intrinsic characteristic of heterogeneity and dispersed features, make the task of studying and analyzing these type of financial and non-financial reports worthy of automation. As a result, in recent years, a growing literature has emerged that relies on AI for the identification of climate-related information [18–29], among others.

Traditional NLP (Natural Language Processing) techniques such as the “Bag-of-Words” (BoW) approaches have been predominant [30–33], although this technique has the significant drawback of assuming that the words in the text are independent of one another what is evidently false. Further literature has relied on word embeddings technology [34] and, even though they are an improvement on the use of BoW, the terminology regarding climate change is likely to change drastically in relation to the sources which are used in the research, so this NLP technique does not offer a completely satisfactory answer either. Furthermore, it should not be forgotten that the vector representation does not take into account the context of the sentence and therefore negations are not considered by this approach. As a result, both methods are very far from understanding sentences in their full extent [20].

However, the field is experiencing an enormous revolution since the implementation of transformer models has been possible due to the rise of computing Meddeb et al. [19]— using a CamemBERT Transformer-based model, a transformer French language model, they classify French news on climate change, identifying those which are fake ones- [20, 22, 24, 25, 27]. Transformers are large language models that capture the dependencies between words, that are encoded in word embeddings whose space represents the meaning of the words. Specifically, transformers’ empirical results dramatically outperform the classical pipeline of machine learning models with a bag-of-words representation of the most common and relevant words of the texts according to algorithms such as TF-IDF (Term Frequency-Inverse Document Frequency) [35]. But transformers can only be accurately estimated by supercomputing centers or organizations that have lots of computing power, being impossible to be trained from zero by small organizations or businesses for small supervised learning tasks. To avoid this limitation for its application, transfer learning can be used to fine-tune a transformer trained in a similar task to the one that needs to be solved [36]. The fine-tuning process adapts the behavior of the transformer to the particular task to be solved and it is cheap in computational terms, thus allowing any organization to perform it. Moreover, the results that a fine-tuned transformer can deliver outperform classical methodologies, or small models trained from zero [20, 25, 35]. For this reason, we propose fine-tuning the recently published pre-trained transformer ClimateBert which has been specifically tailored for climate text classification [24] on the database Cli-maText [37]. The purpose of this paper is to allow better analysis and identification of the increasing amount of climate-related risk disclosure by companies and thus, to further contribute to the literature on the detection of the topic of climate change in text corpora, particularly in order to encourage the identification by economic agents of climate-related financial risks disclosure of firms.

This article contributes to the literature on the applications of neural computing and specifically on the use of pretrained large language models (LLMs) in the realm of financial services and environmental disclosure in several ways:

Our research entails, to the best of our knowledge, the first-ever use of the recently

published ClimateBert algorithm on the ClimaText database, achieving, as developed throughout the paper, better results than those obtained so far from the implementation of other algorithms on ClimaText by Varini et al. [37]. Also, our study marks the first application of the ClimaText database, having been previously used only by its own authors [37] (one more non-published paper includes this reference in its bibliography but does not use it in its research, instead applying another database).

Since its publication in October 2021, the pre-trained transformer model ClimateBert has only been applied so far in 6 other papers which have not been published yet (two of which are from ClimateBert's original authors): [21, 23, 38–41] as well as in two published studies: [42] and the one published by ClimateBert's authors [22]. Hershovich et al. [38] analyze ClimateBert only regarding its energy consumption, in a context of awareness about the environmental impact that NLP pre-trained models present. Focusing on the policies being adopted by governments around the world, Sietsma et al. [23] identify ClimateBert as a tool that allows real-time tracking of adaptation progress. Yu et al. [41] present a database, climateBUG, that provides a framework for detecting implicit information about how banks disclose their climate change-related activities. They use ClimateBERT to check the classification accuracy of this dataset, the study results suggest that ClimateBert outperforms FinBert and Bert uncased due to its higher domain accuracy, although the best performance was obtained by climateBUG-LM (Language Model). Bri'e et al. [39] apply ML techniques to quantify the content of climate disclosure in relation to the European Non-Financial Reporting Directive (NFRD), studying the change in the attitude of companies in their climate-risk reporting. To do this, they use ClimateBERT to accurately identify whether companies are disclosing information related to their risk as a result of climate change. Fard et al. [40] seek to create a language model in such a way that it allows solving different tasks, such as detecting similarities between concepts related to climate and health, checking facts, extracting relationships of health effects in the generation of political texts, etc. Thus, they create CliMedBERT, which is based on ClimateBERT and diseaseBERT. Toetzke et al. [42] develop a machine learning model that classifies development assistance projects regarding their focus on climate finance and they call this model ClimateFinanBert. To train their model, they built on ClimateBert. The surprising conclusion reached by Bingler et al. [22] in their application of ClimateBert to TFCD reports is particularly striking. In the face of the expected growth of disclosure of climate-related risks by TFCD-supporting companies, the study concludes that said support is mostly cheap talk as firms cherry-pick in order to report primarily non-material climate risk information. They continue this study with a new paper, in which Bingler et al. [21] conclude that "institutional ownership, targeted institutional investor engagement, materiality and downside risk disclosures are associated with a less cheap talk".

Numerous stakeholders rely on climate-related corporate disclosures to regularly evaluate risks, impacts, consequences, and opportunities linked to their investments, lending, or insurance portfolios. In this regard, the present study will be relevant and useful to investors and financial agents as it will help establish whether firms are reporting and disclosing climate-related information considered relevant within their industry sector. Likewise, governmental authorities, regulators or policymakers might also be interested in the present project to assess a specific industry's or firm's current state in relation to climate risk disclosure.

Various studies have analyzed various aspects of environmental, social and governance (ESG) disclosure. We refer to Ellili et al. [43] for a bibliometric and systematic review of this literature until 2022. They identify the four most relevant areas in ESG disclosure research up to date: corporate social responsibility (CSR), corporate strategy, financial performance and environmental economics. Our study and obtained results, to the extent that we make possible a better identification and analysis of the climate-related information disclosure, it will also allow advancing in the research on whether the quantity and quality of a company's ESG disclosure affect a great variety of this firm's variables included in these four big areas identified by Ellili et al. [43]. Thus, we extend our paper's target audience to all these strands of the literature related to ESG disclosure.

The remainder of this paper is organized as follows. In section 2, we first show the fundamentals of transfer learning to understand the motivation behind choosing this methodology. Then we describe the particular details of our transfer learning methodology: we review the fundamentals of the models that we are going to fine-tune (BERT and ClimateBert) and we describe the data used to fine-tune the models. In the following section –section 3– we describe our proposed fine-tuning methodology of both the models - the BERT fine-tuning methodology that we use as a baseline and our proposed fine-tuning methodology of the ClimateBERT model. Then, in section, 4 we will analyze the obtained results in the illustrative experiments. Finally, we close our paper with the conclusions and further work section.

2 Fundamentals of large language models for supervised classification

Natural language processing (NLP) is a multidisciplinary area coming from the intersection of artificial intelligence and linguistics that studies methods to process and analyze large amounts of natural language data [44]. In this context, we can fine-tune large language models, like transformers such as BERT (Bidirectional Encoder Representations from Transformers) or the GPT (Generative pretrained transformers) family, to efficiently solve a task that is similar to the one that they were trained on, which is precisely the methodology that we follow in this paper with the ClimateBert model. In this section, we first show the fundamentals of transfer learning to understand the motivation behind choosing this methodology and then we describe the particular details of our transfer learning methodology.

In particular, one of the tasks addressed by NLP is to classify the texts according to a particular label based on a labeled corpus of texts. For example, to determine whether the sentiment of a particular social network publication is positive or negative based on previously labeled social network publications. Concretely, this action is known as supervised learning classification [45] and has been typically solved by machine learning methods that used the different words of the texts as features. However, the field is experiencing an enormous revolution since the implementation of transformer models has been possible due to the rise of computing. Specifically, empirical results shown by transformers dramatically outperform the classical pipeline of machine learning models with a bag-of-words representation of the most common and relevant words of the texts according to algorithms such as TF-IDF [35]. Driven by the amazing results of transformer models, several communities like finance [46], energy

[47] or tourism [48] are starting to use them for a wide variety of tasks.

Transformer models are high-capacity deep learning models, meaning that they have a plethora of parameters to be optimized, that need enormous text corpora to estimate the value of their parameters in order to minimize the estimation of the generalization error of its predictions [49]. Consequently, they can only be accurately estimated by supercomputing centers or organizations that have lots of computing power, being impossible to be trained from zero by small organizations or businesses for small supervised learning tasks. However, although the capacity issue seems to be a major limitation for its application in research problems such as the one targeted by this manuscript, it is eventually not. Most interestingly, we can use transfer learning to fine-tune a transformer trained in a similar task to the one that needs to be solved. In order to understand why the fine-tuning of the last layers of a deep neural network works, it is important to remark that the corpora where the parameters of the transformers are estimated usually contain huge sources of information such as Twitter, GitHub, or online newspapers. Hence, depending on the problem, it is possible that a transformer creates a model of language from all the relevant sources of information, that we can consider as a relevant population for a study. Consequently, it is possible to download open-source transformers that solve tasks such as sentiment analysis and fine-tune them to a particular problem. Some model sources to download transformers include HuggingFace or Github, which are public repositories where we can download these transformers to use them in practice.

The fine-tuning process adapts the behavior of the transformer to the particular task to be solved and, moreover, it is cheap in computational terms, thus, allowing any organization to perform it. Besides, the results that a fine-tuned transformer can deliver outperform classical methodologies, or small models trained from zero [35] when the task to be solved is very similar to the task that the original transformer solves. When fine-tuning is performed to a big model to solve a similar task we call this process transfer learning. As we need to solve a supervised learning problem using a fine-tuned big model we can perform transfer learning to efficiently solve our task.

Nowadays, there is a plethora of large language models, mainly transformers, that differ in their architecture and task solved. Concretely, transformers were first implemented in 2017, as a result of the prestigious paper Attention is all you need [50], where the architecture of BERT was illustrated. In the next section, 2.1., we describe the particular details of the transformer that we use, which is related to the BERT transformer.

2.1 The ClimateBert model

As we already pointed out, in recent years, a growing literature has emerged that relies on AI for the identification of climate-related information. More concretely, BoW (Bag of Words)-based approaches have been predominant, although this technique has the significant drawback of assuming that the words in the text are independent of one another, which is evidently false. On the other hand, transformers are large language models that capture the dependencies between words, that are encoded in word embeddings whose space represents the meaning of the words.

Consequently, we use a transformer model, a BERT-related model called ClimateBert, that has been specifically tailored for climate text classification. This model reads a whole sequence of words at once, allowing it to learn the context of a word

from its surroundings. More specifically, ClimateBert is the state-of-the-art NLP model based on the Transformer architecture which has been specifically pre-trained on climate-related text corpora of over 1.6 billion paragraphs in relation to climate change consisting of news articles, corporate climate reports, and research abstracts [24].

Regarding our research, one of the most significant breakthroughs of ClimateBert is the domain adaptive pre-training, being the first climate domain adaptive pre-trained model which has been made available to the public. In this sense, the model was pre-trained in the following NLP downstream tasks: text classification, sentiment analysis of risk and opportunity of corporations' statements and fact-checking climate-related claims. In this sense, in the present work we have undertaken the ClimateBert fine-tuning process on the ClimaText corpus (described in next section 2.2), related to climate change disclosure, in order to compare the results obtained in the NLP text classification task using fine-tuning of the ClimateBert model with those provided by the BERT model, serving as a baseline because it achieves better performance than the standard BoW-machine learning architecture [35].

2.2 Description of the data used to fine-tune the model

For the purpose of conducting the present study and fine-tuning the ClimateBert model, we use the database named ClimaText [37]. ClimaText is a dataset that intends to advance in the identification of climate change-related topics in text corpora through NLP techniques. The data that compose this dataset represent sentences that are labeled depending on whether they talk about climate change (in which case, they will be labeled with a 1) or not (labeled with a 0). In particular, several sources were employed to create the database, such as Wikipedia, 10-K filings from the U.S. Securities and Exchange Commission (hereinafter, SEC), and a selection of climate change-related claims gathered from the web.

In this regard, in order to train both models (the BERT fine-tuned and the ClimateBert fine-tuned), we have used the "AL-Wiki" dataset formed by a group of sentences extracted from Wikipedia to which we have applied DUALIST, an interactive machine learning (hereinafter, AL) system that enables the building of classifiers for data processing tasks. Concretely, the AL training sentences dataset thus consists of 3,000 sentences, of which 261 are related to the topic of climate change as opposed to 2,739 that are not, making it a tremendously unbalanced dataset. Furthermore, the test set used in the model is the "10 Ks (2018, test)", consisting of sentences that are extracted from Item 1A of the 2018 10-K Filings. It includes 300 sentences of which 67 address climate change compared to 233 that do not.

It is important to emphasize the relevance of assessing the accuracy of both text classification models using data from the 10-K filings. These reports are required by the SEC for publicly listed US companies, which are obliged to provide a detailed annual presentation of their financial results, more specifically the company's history and activity, its financial statements, the risks it faces, its earnings per share, and any other relevant data. For all these reasons, 10-K Filings provide a very useful decision-making tool for investors. Since 2006, the SEC has required 10-K Filings to include Item 1A. For this purpose, Item 1A, on the basis of which the evaluation set of both models have been created, is known as "Risk Factors", thus including the most relevant risks faced by the company, which may be valid for the entire economy,

for the industrial sector in which the firm operates, for a specific geographic area or exclusively for that company [51].

In this manner, the SEC has adopted a so-called principles-based approach, whereby companies must self-identify the climate change-related risks to their business and disclose these risks in Item 1A of the 10-K filings. It is a significant advantage within the SEC's approach that disclosure of such information is mandatory and that failure to disclose climate risks can result in litigation with U.S. authorities [20].

Additional advantages to using this data include the fact that the regulatory reports reflect both physical and transitional risks that are associated with climate change. It should be noted that the 10-K filings provide a description of future risks to the company and may therefore represent a better indicator of the future exposure that the company will face. This forward-looking nature is deliberate as investors often seek to obtain and incorporate forward-looking financial information [20].

Lastly, prior to the explanation of our methodology, it is important to clarify what is meant by climate change in the database, that is to say when a sentence is being classified with a 1. In this sense, the labeling rules to classify a sentence as positive or 1 followed are as follows: the sentence must talk about climate change. To be more precise, simply discussing the topic of the environment is not enough. Moreover, if the sentence is describing a general scientific or climate fact it will only be relevant in the event that it is a cause or effect of climate change (for example, while the sentence "Methane is CH₄" will not be labeled as positive, "Methane increases temperature" will be). Those sentences that speak about clean energy, fossil fuels or emissions, among others, need to be connected to a societal or environmental aspect of climate change. Additionally, environmental aspects such as acid rain or pollution are not considered in this dataset as climate change related terms and the use of acronyms or names of entities (such as EPA -Environmental Protection Agency-) need to be mentioned with a cause or effect of climate change. Also, the sentence can discuss climate change during any time period. If there is any doubt and in all other cases, the sentence is labeled as negative or 0.

3 Proposed Methodology: Fine tuning of BERT and ClimateBERT

Having reviewed the fundamentals of the models that we are going to fine-tune (BERT and ClimateBert) and having illustrated the data, we now describe the fine-tuning process of the models. Recall that we target a binary classification problem where we want to predict whether a text has climate change as a topic or not, according to the previous section. We first give the details of the BERT fine-tuning methodology that we use as a baseline and then those of our proposed fine-tuning methodology of the ClimateBERT model. Then, in a further section, we will analyze the obtained results.

The pre-trained BERT model has been loaded from the TensorFlow Hub repository in order to fine-tune it to our database. The text entries have been set to lowercase prior to carrying out their tokenization, as well as the accents have been removed. We are working with a training batch size of 32 and setting the maximum length of the input token sequences to 128, typical values set by the community. Concretely,

the model has been trained by means of fine-tuning using 3 epochs on the Wikipedia training dataset, previously described.

The pre-trained ClimateBert model has been loaded from the HuggingFace repository [52], in order to fine tune it to our database. In this sense, ClimateBert's tokenizer does not need for the text entries to be previously set into lowercase or for the accents to be removed. We use the default hyper-parameter values of the Trainer library, being aware that an intelligent hyper-parameter tuning of these values, for instance using Bayesian optimization [53], would enhance the predictive out-of-sample performance of the classifier. Lastly, for a fair comparison with the baseline, we also use 3 epochs for fine-tuning.

4 Illustrative experiments and results

In order to estimate the error of generalization of the models regarding their climate change prediction, we use the accuracy metric estimator as we do not give more weight to false positives or false negatives. Concretely, the accuracy of the model indicates the total number of predictions that are correctly predicted with respect to the overall number of predictions performed by the model. As explained in the prior section, our training data is highly unbalanced, so it is necessary that the accuracy achieved by the validation surpasses, at least, the majority rule. Thus, as long as the sentences labeled 0 represent the majority of the training data set - namely 91.3% - the accuracy achieved must be ensured to be above this value.

For a single experiment, the BERT baseline accomplished an accuracy of 95.7% of the validation set which has been reached with the third epoch of fine-tuning. As in the training set the accuracy is slightly higher, we stop the fine-tuning process to avoid overfitting. On the other hand, we obtain an accuracy point estimation of 97.04% using our refined version of ClimateBert, outperforming BERT's performance with the same epochs of fine-tuning and the same setting.

Subsequently, we proceed to evaluate the BERT model using the 2018 10-K Filings test dataset. The model trained by fine-tuning BERT has been able to achieve a significant performance, obtaining a point estimation accuracy of 90%. Analogously, we proceed to evaluate the ClimateBert fine-tuned model using the same Filings test dataset. In this case, the model trained by fine-tuning ClimateBert has been able to outperform BERT, achieving a significant performance, reaching a point estimation accuracy of 93%.

Having achieved these satisfactory statistical results, we proceeded once again to perform a hypothesis test to verify whether the results obtained by the model through the sample data can be considered representative parameters, which is to say, we look to determine whether the results of the fine-tuned ClimateBert model for the classification of text on climate change are truly significant.

For this purpose, we execute the ClimateBert and BERT models 25 times (to ensure the validity of the statistical hypothesis test as a consequence of the central limit theorem) with different random seeds and different splits of the train, validation and test sets. Using this methodology, we provide an interval estimation via a bootstrapping process of the unknown accuracy, precision, F1, recall and specificity

Model	Accuracy	Precision	F1	Recall	Specificity	Deviation
Varini BERT fine-tuned	0.83	0.58	0.71	0.94	Unknown	Unknown
Our BERT fine-tuned	0.925	0.93	0.91	0.93	0.747	0.006
ClimateBERT fine-tuned	0.935	0.94	0.93	0.94	0.786	0.007

Table 1 Performance comparison between the methodologies.

parameters and also include information from the fine-tuned BERT carried out in the paper of Varini et al. [37], obtaining the following results (see Table 1).

We can conclude that our fine-tuned model (ClimateBert) has better performance with respect to the baseline (fine-tuned BERT) due to the fact that if we compute a two-sample t-test with pooled variance and a confidence level of 99% with the previous data, we obtain a p-value of .000003154, representing that the null hypothesis of both models having the same performance is remote.

Finally, it can be claimed that, with the sole exception of the metric of recall, the results we have obtained from the BERT fine tuning process to the aforementioned ClimaText sources have outperformed those previously obtained by Varini et al. [37]. Moreover, these results have been further improved by using the state-of-the-art algorithm, Climatebert, for the NLP task of climate-related text classification.

All in all, it can be said that the fine-tuned ClimateBert model has got a greater capacity to predict not only true positives (sentences that are climate-related) but also to classify true negatives (sentences that are not speaking about climate change). In this sense, ClimateBert has outperformed BERT in the task of climate change text classification.

5 Conclusions and further research

The present study has got the purpose of further contributing to the literature on the detection of the topic of climate change in text corpora, particularly in order to encourage the identification by economic agents of climate-related financial risks disclosure of firms. For this purpose, we have applied NLP techniques that enable us to carry out the text classification task with optimal results with high performance levels. In particular, our text classification model, obtained from the ClimateBert fine-tuning process, outperforms the models created with BERT, both our own and the one introduced by Varini et al. [37]. We believe that this result truly represents a contribution to the literature regarding the detection of climate change in text corpora.

Based on our results, it can be said that the ultimate team of ClimateBert and ClimaText is an outstanding tool within the NLP pre-trained models that may and should be used by investors, institutional agents, and companies themselves to monitor the disclosure of climate risk in financial reports, as well as in other textual sources, such as policies, new legislation or directives passed on this matter. Although in the present paper we have carried out a text classification task in order to compare the fine-tuned BERT and fine-tuned ClimateBert models, the door remains open for the joint use of ClimaText and ClimateBert in other NLP tasks, such as sentiment analysis. Likewise, another area of further research is to jointly apply ClimateBert and ClimaText in order to analyze whether European companies have significantly improved their

climate-related disclosures after the mandatory request of the European Non-Financial Reporting Directive (NFRD).

Statements and Declarations

- Funding The authors declare that no funds, grants, or other support were received during the preparation of this manuscript.
- Competing interests The authors have no relevant financial or non-financial interests to disclose.
- Data availability All data are publicly available from <http://climatefever.ai>
- Authors' contributions All authors contributed equally to this study.

References

- [1] TCFD, Task Force on Climate-Related Financial Disclosures: Recommendations of the Task Force on Climate-related Financial Disclosures: Final report (June 2017). (2017). Available at: <https://assets.bbhub.io/company/sites/60/2021/10/FINAL-2017-TCFD-Report.pdf>
- [2] TCFD, Task Force on Climate-Related Financial Disclosures: Task Force on Climate-related Financial Disclosures: 2022 status report. (2022). Available at: <https://assets.bbhub.io/company/sites/60/2022/10/2022-TCFD-Status-Report.pdf>
- [3] Steffen, B.: A comparative analysis of green financial policy output in OECD countries. *Environmental Research Letters* **16**(7), 074031 (2021) <https://doi.org/10.1088/1748-9326/ac0c43>
- [4] European NFRD: The Non Financial Reporting Directive (NFRD). Regulation DIRECTIVE 2014/95/EU of the European Parliament and of the Council of 22 October 2014 amending Directive 2013/34/EU as regards disclosure of non-financial and diversity information by certain large undertakings and groups. Directive 2014/95/EU. Official Journal of the European Union. European Union. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32014L0095> (2014)
- [5] European SFDR: Regulation (EU) 2020/852 of the European Parliament and of the Council of 18 June 2020 on establishing a framework to facilitate sustainable investments and amending Regulation (EU) 2019/2088 (OJEU-L-198 of 22 June 2020). European Union. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A32019R2088> (2019)
- [6] European Taxonomy Regulation: Regulation (EU) 2020/852 of the European Parliament and of the Council of 18 June 2020 on establishing a framework to facilitate sustainable investments and amending Regulation (EU) 2019/2088

(OJEU-L-198 of 22 June 2020). European Union. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A32020R0852> (2020)

[7] EFRAG, European Financial Reporting Advisory Group: Cover Note for Public Consultation: Draft European Sustainability Reporting Standards (ESRS). April 29. <https://www.efrag.org/Assets/Download?assetUrl=%2Fsites%2Fwebpublishing%2FSiteAssets%2FESRS.CN.pdf>

[8] European CSRD: Corporate Sustainability Reporting Directive. European Union. https://finance.ec.europa.eu/capital-markets-union-and-financial-markets/company-reporting-and-auditing/company-reporting/corporate-sustainability-reporting_en (2022)

[9] French Law on Energy Transition and Green Growth: Article 173 de la Loi relative à la transition énergétique et pour la croissance verte. <https://www.legifrance.gouv.fr/jorf/article/jo/JORFARTI000031045547> (2022)

[10] GovUK.: Mandatory climate-related financial disclosures by publicly quoted companies, large private companies and LLPs. GovUK. https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1056085/mandatory-climate-related-financial-disclosures-publicly-quoted-private-cos-lbps.pdf (2022)

[11] New Zealand Ministry for the Environment: Mandatory Climate-related Disclosures. <https://environment.govt.nz/what-government-is-doing/areas-of-work/climate-change/mandatory-climate-related-financial-disclosures/> (2021)

[12] U.S. Securities and Exchange Commission.: Commission Guidance Regarding Disclosure Related to Climate Change. <https://www.sec.gov/rules/interp/2010/33-9106.pdf> (2010)

[13] U.S. Securities and Exchange Commission.: SEC Proposes Rules to Enhance and Standardize Climate-related Disclosures for Investors, Press Release. <https://www.sec.gov/news/press-release/2022-46> (2022)

[14] IFRS, International Financial Reporting Standard.: Exposure Draft IFRS S2 Climate-related Disclosures. World Economic Forum. <https://www.ifrs.org/content/dam/ifrs/project/climate-related-disclosures/issb-exposure-draft-2022-2-climate-related-disclosures.pdf> (2022)

[15] World Economic Forum.: The Global Risks Report 2023. https://www3.weforum.org/docs/WEF_Global_Risks_Report_2023.pdf (2023)

[16] Carney, M.: Breaking the tragedy of the horizon–climate change and financial stability. Speech given at Lloyd's of London 29, 220–230 (2015). <https://www.bankofengland.co.uk/speech/2015/>

breaking-the-tragedy-of-the-horizon-climate-change-and-financial-stability

- [17] Ilhan, E., Krueger, P., Sautner, Z., Starks, L.T.: Climate risk disclosure and institutional investors. *The Review of Financial Studies* (hhadoo2) (2023) <https://doi.org/10.1093/rfs/hhadoo2>
- [18] Sautner, Z., Lent, L., Vilkov, G., Zhang, R.: Firm-level climate change exposure. *The Journal of Finance* (accepted paper) (2023) <https://doi.org/10.1111/jofi.13219>
- [19] Meddeb, P., Ruseti, S., Dascalu, M., Terian, S.-M., Travadel, S.: Counteracting french fake news on climate change using language models. *Sustainability* **14**(18), 11724 (2022) <https://doi.org/10.3390/su141811724>
- [20] Kölbel, J.F., Leippold, M., Rillaerts, J., Wang, Q.: Ask BERT: How regulatory disclosure of transition and physical climate risks affects the CDS term structure. *Journal of Financial Econometrics* (nbaco27) (2022) <https://doi.org/10.1093/jjfinec/nbaco27>
- [21] Bingler, J.A., Kraus, M., Leippold, M., Webersinke, N.: Cheap talk in corporate climate commitments: The role of active institutional ownership, signaling, materiality, and sentiment. *Swiss Finance Institute Research Paper* (22-01) (2022). Available at SSRN: <http://dx.doi.org/10.2139/ssrn.4000708>
- [22] Bingler, J.A., Kraus, M., Leippold, M., Webersinke, N.: Cheap talk and cherry-picking: What ClimateBert has to say on corporate climate risk disclosures. *Finance Research Letters* **47-B**, 102776 (2022) <https://doi.org/10.1016/j.frl.2022.102776>.
- [23] Sietsma, A.J., Callaghan, M., Biesbroek, R., Theokritoff, E., Thomas, A., Canosa, I.V., Schleussner, C.-F., Ford, J.D., Minx, J.C.: Global tracking of climate change adaptation policy using machine learning: a systematic map protocol. *2022* (2022) <https://doi.org/10.21203/rs.3.pex-1836/v1> . preprint at Research Square
- [24] Webersinke, N., Kraus, M., Bingler, J.A., Leippold, M.: ClimateBert: A pretrained language model for climate-related text (2021). Preprint at arXiv:2110.12010 <https://arxiv.org/abs/2110.12010>
- [25] Friederich, D., Kaack, L.H., Luccioni, A., Steffen, B.: Automated identification of climate risk disclosures in annual corporate reports (2021). Preprint at arXiv:2108.01415 <https://arxiv.org/abs/2108.01415>
- [26] Hsu, A., Rauber, R.: Diverse climate actors show limited coordination in a large-scale text analysis of strategy documents. *Communications Earth & Environment* **2**(1), 30 (2021) <https://doi.org/10.1038/s43247-021-00098-7>

- [27] Luccioni, A., Baylor, E., Duchene, N.: Analyzing sustainability reports using natural language processing (2020). Preprint at arXiv:2011.08073 <https://arxiv.org/abs/2011.08073>
- [28] Biesbroek, R., Badloe, S., Athanasiadis, I.N.: Machine learning for research on climate change adaptation policy integration: an exploratory UK case study. *Regional Environmental Change* **20**, 85 (2020) <https://doi.org/10.1007/S10113-020-01677-8>
- [29] Alashri, S., Alzahrani, S., Tsai, J.-Y., Corman, S.R., Davulcu, H.: “climate change” frames detection and categorization based on generalized concepts. *International Journal of Semantic Computing* **10**(2), 147–166 (2016) <https://doi.org/10.1142/S1793351X16400055>
- [30] Donner, S.D., Kandlikar, M., Webber, S.: Measuring and tracking the flow of climate change adaptation aid to the developing world. *Environmental Research Letters* **11**(5), 054006 (2016) <https://doi.org/10.1088/1748-9326/11/5/054006>
- [31] Loughran, T., McDonald, B.: Textual analysis in accounting and finance: A survey. *Journal of Accounting Research* **54**(4), 1187–1230 (2016) <https://doi.org/10.1111/1475-679X.12123>
- [32] Cody, E.M., Reagan, A.J., Mitchell, L., Dodds, P.S., Danforth, C.M.: Climate change sentiment on twitter: An unsolicited public opinion poll. *PLoS ONE* **10**(8), e0136092 (2015) <https://doi.org/10.1371/journal.pone.0136092>
- [33] Sham, N.M., Mohamed, A.: Climate change sentiment analysis using lexicon, machine learning and hybrid approaches. *Sustainability* **14**(8), 4723 (2022) <https://doi.org/10.3390/su14084723>
- [34] Luccioni, A., Palacios, H.: Using natural language processing to analyze financial climate disclosures. In: *Proceedings of the 36th International Conference on Machine Learning*, Long Beach, California (2019)
- [35] Garrido-Merchán, E.C., Gozalo-Brizuela, R., González-Carvajal, S.: Comparing BERT against traditional machine learning text classification. *Journal of Computational and Cognitive Engineering* **2**(4), 352–356 (2023) <https://doi.org/10.47852/bonviewJCCE3202838>
- [36] Yang, Q., Zhang, Y., Dai, W., Pan, S.J.: *Transfer Learning*. Cambridge University Press, Cambridge (2020). <https://doi.org/10.1017/9781139061773>
- [37] Varini, F.S., Boyd-Graber, J., Ciaramita, M., Leippold, M.: ClimaText: A dataset for climate change topic detection (2020). Preprint at arXiv:2012.00483 <https://arxiv.org/abs/2012.00483>
- [38] Hershovich, D., Webersinke, N., Kraus, M., Bingler, J.A., Leippold, M.: Towards

climate awareness in NLP research (2022). Preprint at arXiv:2205.05071 <https://arxiv.org/abs/2205.05071>

- [39] Bri'e, B., Stouthuysen, K., Verdonck, T.: Mandatory CSR reporting in europe: A textual analysis of firms' climate disclosure practices. Available at SSRN 4231567 (2022)
- [40] Fard, B.J., Hasan, S., Bell, J.: CliMedBERT: A pre-trained language model for climate and health-related text (2022). Preprint at arXiv:2212.00689 <https://arxiv.org/abs/2212.00689>
- [41] Yu, Y., Scheidegger, S., Elliott, J., Löfgren, Å.: Climatebug: A data-driven framework for analyzing bank reporting through a climate lens. *Expert Systems with Applications* **239**, 122162 (2024) <https://doi.org/10.1016/j.eswa.2023.122162>
- [42] Toetzke, M., Stünzi, A., Egli, F.: Consistent and replicable estimation of bilateral climate finance. *Nature Climate Change* **12**(10), 897–900 (2022) <https://doi.org/10.1038/s41558-022-01482-7>
- [43] Ellili, N.O.D.: Bibliometric analysis and systematic review of environmental, social, and governance disclosure papers: current topics and recommendations for future research. *Environmental Research Communications* **4**(9), 092001 (2022) <https://doi.org/10.1088/2515-7620/ac8b67>
- [44] Nadkarni, P.M., Ohno-Machado, L., Chapman, W.W.: Natural language processing: an introduction. *Journal of the American Medical Informatics Association* **18**(5), 544–551 (2011) <https://doi.org/10.1136/amiajnl-2011-000464>
- [45] Murphy, K.P.: *Machine Learning: a Probabilistic Perspective*. MIT press, Boston (2012)
- [46] Zhao, L., Li, L., Zheng, X., Zhang, J.: A BERT based sentiment analysis and key entity detection approach for online financial texts. In: 2021 IEEE 24th International Conference on Computer Supported Cooperative Work in Design (CSCWD), pp. 1233–1238 (2021). IEEE
- [47] Cai, R., Qin, B., Chen, Y., Zhang, L., Yang, R., Chen, S., Wang, W.: Sentiment analysis about investors and consumers in energy market based on BERT-BiLSTM. *IEEE Access* **8**, 171408–171415 (2020) <https://doi.org/10.1109/ACCESS.2020.3024750>
- [48] Chantrapornchai, C., Tunsakul, A.: Information extraction on tourism domain using SpaCy and BERT. *ECTI Transactions on Computer and Information Technology* **15**(1), 108–122 (2021) <https://doi.org/10.37936/ecti-cit.2021151.228621>
- [49] LeCun, Y., Bengio, Y., Hinton, G.: Deep learning. *Nature* **521**(7553), 436–444 (2015) <https://doi.org/10.1038/nature14539>

- [50] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
- [51] U.S. Securities and Exchange Commission.: SEC.gov. How to Read a 10-K/10-Q. U.S. Securities and Exchange Commission. <https://www.sec.gov/fast-answers/answersreada10khtm.html> (2021)
- [52] Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., et al.: Huggingface's transformers: State-of-the-art natural language processing (2019). Preprint at arXiv:1910.03771 <https://arxiv.org/abs/1910.03771>
- [53] Garrido Merchán, E.C.: Advanced methods for bayesian optimization in complex scenarios. Universidad Autónoma de Madrid (UAM), Higher Polytechnic School, Computer Science Department (2021)