GreenAI – An NLP approach to ESG financing

Nicolaas Ruberg[1], Rafael B. Pereira[1], Mauro L. Stein[1]

[1]Area de Tecnologia da Informação – Banco Nacional de Desenvolvimento Econômico e Social (BNDES)
Av. República do Chile, 100 – 20031-917 – Rio de Janeiro – RJ – Brazil

{nicolaas, rbper, mstein}@bndes.gov.br

Abstract. *Environmental, Social, and Governance (ESG) factors are critical for investors and financing institutions like the Brazilian Development Bank (BNDES). Such institutions are currently working on setting up a framework to assess companies' ESG factors in their financing evaluation. In this study, we identify an opportunity to use Natural Language Processing (NLP) to improve the framework. This opportunity stems from the fact that the key documents for ESG analysis, such as the company's activity report (RAA), Environmental Impact Study (EIA), and Environmental Impact Report (RIMA), undergo manual screening and decomposition whilst being analyzed by specialists. By incorporating NLP, we aim to automate the classification of text passages from these reports and enhance the efficiency of the analysis process.*

Disclaimer. *This is a working paper, and hence it represents research in progress. This paper represents the opinions of the authors, and is the product of professional research. It is not meant to represent the position or opinions of the BNDES or its Members, nor the official position of any staff members.*

1. Motivation and Introduction

In recent years, there has been a shift in society toward sustainability. This shift is impelling political pressure, a regulatory push, and technological advancements to build a more sustainable world economy. Consequently, a change in the investors is taking place, thus guiding a gradual capital reallocation (Hildebrand, et al., 2020). In this new investment scenario, it is becoming crucial for investors to assess companies' sustainability (Napoletano, et al., 2021).

Evaluating sustainability involves analyzing a company's decisions and actions related to Environmental, Social, and Governance (ESG) factors that aim to improve long-term outcomes for the planet and future generations (Inderst, et al., 2018). This analysis encompasses a wide range of aspects that include quantitative and qualitative information and is typically carried out by experts. Supporting decision-making in sustainable investing is a challenging task that currently lacks practical tools for automatically assessing ESG factors, particularly regarding analyzing qualitative information in companies, such as the Annual Activity Report (RAA), Environmental Impact Study (EIA), and Environmental Impact Report (RIMA).

To grant a loan or finance a company or a project under this sustainable financing umbrella, banks and government investment agencies must conduct a sustainability assessment of the companies. Hence, this assessment requires an expert understanding of the borrower companies' actions and strategies towards sustainability. Key inputs to this assessment process are the Environmental Impact Assessment (EIA) and Risk and Impact Management Assessment (RIMA) reports, which detail the impact of the project's activities, as well as the Annual Activity Reports (RAA), which outline the company's sustainability actions and initiatives in the preceding year. However, analyzing these reports for ESG information is often time-consuming and susceptible to bias based on the analyst's background and expertise.

Applying natural language processing (NLP) techniques can improve the ESG analysis of available qualitative information in the sustainability assessment process. The work in (Ruberg, 2021) investigated the main English BERT-like architectures, examining them with different model sizes, and a classical NLP technique (Naïve-Bayes classifier) to establish a baseline for a performance comparison.

This work starts from this approach to enable the ESG understanding of Brazilian companies in Portuguese. In such a way, we devised a Portuguese Language Model fine-tuned to classify ESG paragraphs. This ESG understanding is acquired from the Global Reporting Initiative (GRI), an ESG standard for reporting sustainability (Wikipedia, 2023).

In our work, we use BERTimbau, base an large language models, (Souza, 2020) as a Portuguese pre-trained model and DistilBert multilingual (V. Sanh, 2019): all models were fine-tunned with annual activity reports from 20 companies. We achieved an accuracy of 89% and an F1-score of 87% using BERTimbau; an accuracy of 87% and an F1-score of 85,33% using BERTimbau Large; and an accuracy of 93% and an F1-score of 90% using DistilBERT multilingual on our validation data.

2. Sustainable Financing

According to FORBES magazine (Napoletano, et al., 2021), environmental, social, and governance (ESG), are becoming key factors as part of the investment strategy, where socially conscious investors want to put their money in companies that strive to make the world a better place. In June 2022, over 5 thousand institutional investors and asset managers, representing more than US$ 121 trillion in assets, were signatories of the Principles for Responsible Investments (PRI)[1].

ESG investing relies on independent ratings that help investors assess a company's actions and policies regarding environmental performance, social impact, and governance issues. In other words, ESG investing aims to influence positive changes in society with the invested money. Thus, ESG investing fosters companies and corporations with good scores on scales of factors on environmental, societal, and governance responsibility. Independent third-party entities determine those scores.

Those entities approach the three criteria used to evaluate companies for ESG investing (Napoletano, et al., 2021):

- Environment – what kind of impact does a company have on the environment? This can include a company's footprint, toxic chemicals, water usage involved in

[1] https://www.unpri.org/annual-report-2022/signatories

its manufacturing processes, and sustainability efforts that make up its supply chain;

- Social – How does the company improve its social impact, both within the company and in the broader community? Social factors include LGBTQ+ equality, racial diversity within both the executive staff and the general workforce, inclusion programs and hiring practices. It even looks at how a company advocates for social good in the broader world beyond its limited sphere of business; and
- Governance – How does the company's management drives positive change? Governance includes everything from issues surrounding executive pay to diversity in leadership and how well that leadership responds to and interacts with shareholders.

ESG criteria can also help to avoid investing in companies with high financial risks due to environmental or other practices related to social responsibility.

Nowadays, several independent entities provide objective sustainable report standards. A survey of the most relevant ones is presented by (Meager, 2021). Among them, the Global Reporting Initiative (known as GRI) is the oldest, with the most significant number of participants.

The GRI is an international independent standards organization that helps businesses, governments, and other organizations understand and communicate their impacts on climate change, human rights, and corruption (Wikipedia, Global Reporting Initiative, 2021). In other words, GRI assesses companies and organizations on a scale of factors on ESG criteria and produces a score.

Therefore, in September 2020, the World Economic Forum (WEF) developed a core set of standard metrics and disclosures on non-financial factors for their investors and other stakeholders; those metrics became the basis of several standards.

To publicize a company's initiatives and efforts, many enterprises regularly publish a sustainability report, also known as a corporate social responsibility (CSR) or environmental, social, and governance (ESG) report. Companies usually publish this CSR report on their annual activity report (RAA). To guide the production of the report, the GRI's standards helps companies identify, gather, and disclose this information transparently and comparably.

The series of topic-specific standard, GRI is divided into three series. The GRI's 200 series, that establishes how the company should report its economic dimension; the GRI 300 series, which defines how to report the companies environmental situation and strategy; and the GRI 400 series that addresses the company social aspects, in the following perspectives: employees; human rights for native peoples and local communities; suppliers, customer and government relations; and socioeconomic compliance.

Moreover, the Brazilian National Environmental Policy - PNMA (Brazilian law nº 6.938/81[2]) has established environmental licensing as an administrative instrument.

---

[2]https://www2.camara.leg.br/legin/fed/lei/1980-1987/lei-6938-31-agosto-1981-366135-normaatualizada-pl.pdf

The law establishes the environmental governance, in such a way that the environmental administration agency sets control conditions, restrictions, and guidelines that the entrepreneur must follow. This also includes procedures to locate, install, expand, and operate activities that may impact the environment.

Therefore, concerning the environmental aspects, an enterprise must disclose its environmental viability in the Environmental Impact Study (EIA), a technical document covering the company's physical, biological, and socioeconomic aspects. It identifies the direct and indirect; positive and negative; immediate, medium and long-term environmental impacts. Since the Environmental Impact Study (EIA) is a technical document, it may hold sensitive information with privacy restrictions. Hence, to overcome the disclosure of sensitive information, a company produces the Environmental Impact Report (RIMA), which gives transparency to the enterprise's actions and initiatives in a concise document targeting any external entity or stakeholder.

3. ESG analysis process

The effort to develop an ESG framework is crucial for a better and unified assessment of companies and investment projects, emphasizing the sustainability perspective on investment decisions. Moreover, in the design of the framework, the GRI standard is the primary guideline where specialists should assess loaners, beneficiaries, and clients' reports.

Consequently, a vital part of the ESG analysis involves the right specialist for the company evaluation. Next, we present a workflow process for assessing ESG aspects supported by an NLP tool.

The conceived process introduces the screener role, which works like a dispatcher of extracts from the company's reports to specialists for a detailed evaluation. The automatization of this activity should reduce the overall analysis time. In Figure 1, we present the workflow analysis with the screener role performed by an NLP agent for the GRI assessment.

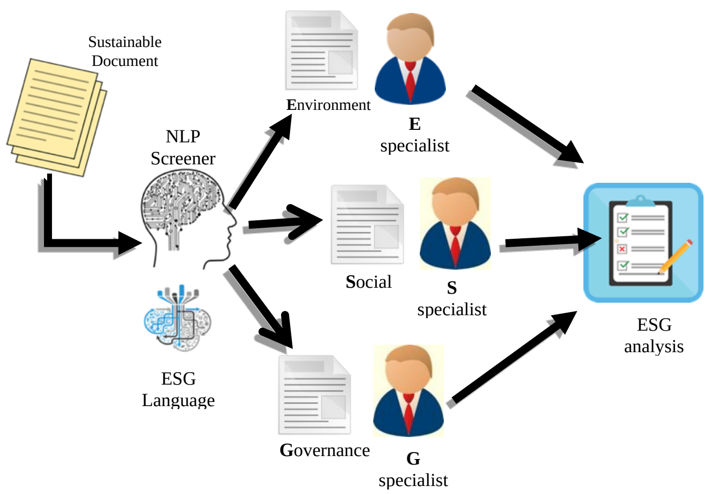

Figure 1 - NLP screener role in ESG analysis.

In this workflow process, the NLP screener is supported by a GRI Language Model, trained on a set of annual activity reports (RAA) in compliance with the GRI standard database. The decision to use the RAA relies on the fact that some companies have their RAA certified by GRI. Hence, providing a *gold standard* (Wissler, 2014) data that can be used as the benchmark for constructing our ESG language model.

4. Natural language processing and text classification

Natural language processing (NLP) exploits knowledge of language and computation by building useful technologies (Bird, Klein, & Loper, 2009). To devise such language knowledge, NLP relies on Artificial Intelligence (AI) to understand and interpret human language, enabling new tools and technologies.

In its use, the field of NLP covers a broad range of tasks and methods like Machine Translation, Text Categorization, Spam Filtering, Information Extraction, Summarization, Dialogue Systems, and others (Diksha Khurana, 2022).

An NLP particular field of our interest is text classification or text categorization, which is the processing that assigns the text documents into a set of predefined text categories; hence, it is an automated process for classifying text into predefined categories. In other words, for a given text, the text classification goal is to assign a discrete label from a set of possible tags. To bring up a more concrete example, recently, on social media and e-commerce websites is vital to understand the customers' sentiment about a product, a post, or a comment. Therefore, classifying customer feedback into good, bad, or neutral is essential for the seller or service provider to make future decisions. So this text classification is called sentiment analysis, an NLP task that, given a text, classifies it into the categories: good, bad, or neutral (the customers' impression or feeling).

4.1. Transformers and text classification

(Dhar, Mukherjee, Dash, & Kaushik, 2021) presents a survey of various algorithms used for categorizing text documents. On the conventional methods, among the different classification algorithms, it is worth mentioning: Naïve Bayes (NB), Support Vector Machine (SVM), and K-Nearest Neighbor (KNN). On the novel techniques, we can list Recurrent Neural Networks (RNN), long short-term memory (LSTM), gated recurrent networks, and Transformers (Vaswani, et al., 2017). Transformers are the building block of a series of powerful transfer learning techniques like Bidirectional Encoder Representations from Transformers (BERT), Generative Pre-trained Transformers 2 (GPT-2), and more recently, Generative Pre-trained Transformers 3 (GPT-3). The first one, BERT, was designed by Google, while OpenAI created the other two, and the last version, GPT-3 was released in 2020 (Brown, et al., 2020).

An essential aspect of natural language processing (NLP) is that communication and language are temporal phenomena; hence, a sequence of input streams needs to be sequentially processed to understand the written text or verbal communication. Traditional machine learning techniques for NLP like Naïve Bayes, Support Vector Machines, and K-Nearest Neighbors don't have this temporal nature, since "the bag of word" model (wikipedia, 2023) and/or similar models, the input is accessed simultaneously.

A Recurrent Neural Network (RNN) is a neural network architecture that captures a preceding sequence. This property is achieved due to a hidden layer with a recurrent link; this link provides some memory (temporality) from previous states. This memory turned RNN into a trend in NLP, particularly in tasks like named-entity recognition

(NER) and part-of-speech (POS) tagging. In Figure 2, we present a standard RNN structure; we can observe that the hidden layers are updated according to the information received from the input layer and the activation from the previous forward propagation (Fang, Xu, & Xu, 2019).

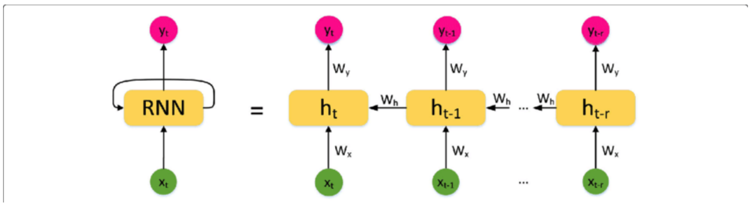

Figure 2 - Standard RNN structure[3].

Another technique to retain the temporal nature in language understanding is the mechanism of self-attention that relates different positions of a single sequence to compute a representation of the sequence. Transformer models (Vaswani, et al., 2017) mainly explore self-attention to extract features of each sequence to determine the word importance concerning the previous word or sentence. The computation can be efficiently parallelized since the architecture does not use recurrent units, just weighted sums, and activations.

The Transformer architecture scales with training data and model size due to its efficient parallel training and ability to capture long-range sequence features (Wolf, et al., 2020). Therefore, this architecture is convenient to pre-train on huge corpora, increasing the accuracy on most NLP tasks like text classification, language understanding, translation, and summarization.

From the presented techniques, we identified the Bidirectional Encoder Representations from Transformers (BERT) architecture best suited for our problem of creating an ESG language model to classify text paragraphs. If we compare the two prominent techniques in NLP, we come across BERT and GPT-2 or 3. The former is a discriminative technique suited for classifying tasks (Jebara, 2004). The latter (GPT) is a generative technique suited for generating similar results as the trained ones (e.g., producing a face with visual features like the ones it was trained at, generating an essay with specific characteristics, etc.) Next, we detailed the BERT architecture and its pretraining strategy. Also, we introduce BERTimbau, the Brazilian Portuguese pretrained Language Model.

4.2. Language models

Next, we introduce BERT model architecture, the reference architecture. And the two architectures used in our experiments: BERTimbau and Distilbert.

4.2.1. BERT reference architecture

The BERT base model is pre-trained on the English language using a masked language modeling (MLM) strategy (Devlin, Chang, Lee, & Toutanova, 2018).

[3] Image from https://www.researchgate.net/figure/A-standard-unfolded-RNN-structure-at-time-t_fig1_333294428.

Model description

The BERT is a transformer model pre-trained on a large corpus of English data, Wikipedia and BooksCorpus (Zhu, et al., 2015), with more than 980 million words from 11.038 books. The pretraining is done on raw text, with no manual labeling, where an automatic process generates inputs and labels from those texts. More precisely, it was pre-trained with two objectives:

- Masked Language Modeling (MLM) to train the model to understand a bidirectional representation of the sentence. 15% of the words are randomly masked in the input, and then run the entire masked sentence through the model, which predicts the masked words.
- Next Sentence Prediction (NSP), where the model is trained to predict if the two input sentences are next to each other. Two masked input sentences are concatenated and provided as inputs, randomly assigned in scenarios of whether they were together.

These two objectives allow the model to learn the English language representation, capturing features that are useful for other NLP tasks.

Specifically for text classification, after a fine-tuning procedure, the final hidden state of the **[CLS]** token will have a fixed-dimensional representation of the sequence (Figure 3), which is fed to the classification layer. The classification layer is the only new parameter added and has a dimension of K x H, where K is the number of classifier labels, and H is the size of the hidden state (Seth, 2019).

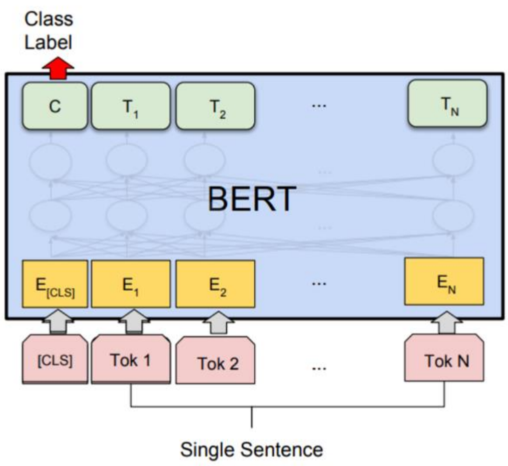

Figure 3 - BERT model for text classification.

4.2.2. BERTimbau Base and Large

The BERTimbau (Souza F. N., 2020) is pre-trained in the Portuguese language using the BERT architecture, providing the same pretraining procedures and objectives as BERT, i.e., Masked Language Modeling (MLM) and Next Sentence Prediction (NSP). This model is case-sensitivity, i.e., it does make a difference between Caps and caps.

BERTimbau "base" has its architecture with 768 feed-forward layers, and BERTimbau large has 1024 feed-forward layers, with 110 and 355 million parameters, respectively. For pretraining data, it used the brWaC - Brazilian Web as Corpus (Wagner Filho, 2018) corpus, a crawl of Brazilian webpages which contains 2.68 billion tokens from 3.53 million documents and is the most extensive open Portuguese corpus to date. On top of its size, brWaC is composed of whole documents, and its methodology ensures high domain diversity and content quality, which are desirable features for pretraining.

4.2.3. DistilBERT base multilingual model (cased)

This model is a distilled version of the BERT base multilingual model. This model is also case-sensitivity. DistilBERT multilingual model is available at HuggingFaces (distilbert-base-multilingual-cased[4]).

Model description

DistilBERT multilingual is trained on the concatenation of Wikipedia in 104 different languages listed here. The model has six layers, 768 dimensions, and 12 heads, totalizing 134M parameters (compared to 177M parameters for BERT-base-multilingual). On average, DistilBERT is two times faster than the English BERT-base-multilingual.

BERT base multilingual model was used as a teacher to generate DistilBERT multilingual. As in the English DistilBERT, it was pre-trained with three objectives:

- Distillation loss: the goal is to achieve a model with the same probabilities as the BERT base model.
- It uses the same strategy as BERT of a Masked language modeling (MLM)
- With cosine embedding loss, the hidden states are trained to be similar to the BERT base model.

Hence, DistilBERT multilingual achieves the same inner language representation as its teacher model, with the advantage of being faster on traditional NLP tasks, fine-tuning, and producing a smaller output model.

5. Experiment

Our experiment uses Brazilian companies' Annual Activity Reports (RAA) to produce a fine-tunned language model to classify ESG paragraphs.

5.1. Training and Validation Data

In our experiments, we manually downloaded twenty verified annual activity reports consistent with the GRI standards from 2016 to 2021. We extracted the training and validation set from those reports to fine-tune the language model that can recognize GRI concepts fitted for Brazilian companies. We experiment with three different pre-trained strategies: BERTimbau, BERTimbau Large and DestilBERT.

For the sake of our training and validation set, the obtained classified paragraphs can be considered a *Gold Standard* (Wissler, 2014) since we extracted the classified text from validated GRI documents, so those had their classification well scrutinized according to the GRI standards.

---

[4] https://huggingface.co/distilbert-base-multilingual-cased

After the data preparation of the twenty Annual Activity Reports, we classified 2,168 paragraphs into the GRI categories (200, 300, and 400). In Figure 4, as we can observe, the distribution text occurrences of those categories are unevenly distributed, which indicates that accuracy is not the best metric to evaluate our model (Qayyum, 2022). Hence in our experiments, we compared our results both with the accuracy, but also with the macro F1-score

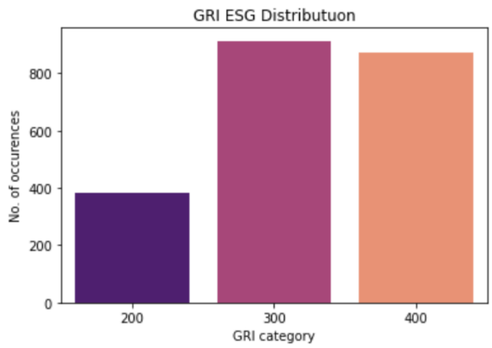

Figure 4 – Dataset GRI categories distributions.

5.1. Running environment

To generate our ESG language model, we run our experiment in Google Colab PRO using TPU with the following hardware:

- Forty processors Intel(R) Xeon(R) CPU @ 2.30GHz, family 6 (model63); and
- A memory of 36 GB with a 2 GB cache.

For the model execution and use, we tested the "on-premises" Jupyter Hub installation. Running with the following hardware:

- Twelve processors Intel(R) Xeon(R) CPU @ 2.50GHz, family 6 (model3f); and
- A Memory of 46 GB.

5.2 Experiment setup

Each model (BERTimbau, BERTimbau Large and DistilBERT) was trained for ten epochs with the same optimizer hyperparameters to minimize the biases in the evaluation phase. Also, we split the documents into paragraphs, and this set of paragraphs is separated into a 90 − 10% train-validation set, which is randomly chosen. We assessed our model performance by analyzing the macro F1-score and accuracy. In the text, we refer to the macro F1-score just as F1-score.

It is also important to remark that the learning rate proved to be extremely important for convergence (all models diverged using Adam's default hyperparameters, a value of $lr = 5 \cdot 10{-}5$, and the default ones for $\beta 1$ and $\beta 2$, proved to be good to make all the transformers converge).

We conducted our experiments with minor text pre-processing (space removal and the GRI token removal); we took inspiration from some NLP forum discussions that have obtained good results in scraping PDF documents without any pre-processing, like the orientation from the PRODIGY support team (Support, 2018).

We fine-tuned all the transformer experiments using the maximum token lengths of 512 on Google Colab's and Colab's PRO TPU.

We evaluated (using a 90–10 train-validation randomly split) each model by accuracy and F1-score. Next, we present our results for each language model.

In, we present the results of our experiments with the BERTimbau (base and large) model and DistilBERT multilingual model. The **DistilBERT** achieved the best accuracy (92.97%) with the best F1-score (89.83%) among the tested models.

Table 1 – BERTimbau (base and large) and DistilBERT model results.

Bert Model	Accuracy	Precision	Recall	F1-score	Size
BERTimbau	89.9497%	86.5942%	88.6509%	87.5297%	415.8MB
BERTimbau Large	87.0000%	83.5315%	89.8709%	85.3321%	1.25Gb
DistilBERT	92.9730%	93.8422%	87.6064%	**89.8381%**	518.6MB

6. User testing

At BNDES we are currently testing a minimum viable project for the ESG specialist. This MVP is simple, since it receives just a paragraph and outputs the GRI macro classification (i.e., 200, 300 and 400). Nevertheless, it gives the specialist an appetizer of the power of the language understanding.

Meanwhile, we are developing a second version of the MVP that will receive a report and split the document into the GRI classes, therefore providing the exact extracts of text to the respective specialist.

7. Conclusion

Sustainability is gaining momentum in society and causing changes in people's and companies' attitude towards Environmental, Social, and Governance (ESG). Therefore, investors are guiding a gradual capital reallocation to more sustainable and social conscious companies. Nevertheless, the question of how to assess companies' sustainability efforts remains an issue since it requires investors to analyze a broad range of aspects regarding initiatives, decisions, and actions on ESG-related factors that could improve long-term outcomes, using both quantitative and qualitative information. Hence, the success of ESG investing – also called socially responsible investing, strongly relies on the adoption of more structured, effective and preferably automatic ways to evaluate these factors. This is specially relevant for an institution like BNDES, which is a public financial entity in charge of long-term financing to enterprises in Brazil and, as such, have to adequately assess the social and environmental risks of these companies.

In this context, our work addressed the problem to establish an automatized ESG framework for assessing companies' sustainability strategies. We proposed an automation approach based on a natural language processing (NLP) method to improve the analysis of ESG factors by exploring the Global Reporting Initiative (GRI), a well-established standard that has been used to structure companies' reports such as the annual activity reports, the Environmental Impact Study (EIA), and the Environmental Impact Report (RIMA) regarding ESG aspects. The proposed solution targets at the analysis carried out by domain experts on these reports and introduces an automatic screener to overcome an important bottleneck in dispatching parts of those documents, for a more refined ESG analysis of companies or projects.

The proposed automation solution is based on Bidirectional Encoder Representations from Transformers (BERT), which relies on the attention mechanism to achieve optimal results on sentence-level analysis tasks. We devised a text classification task to analyze excerpts from the annual activity report of companies considering three categories, according to the GRI standard: Environment; Social; and Governance/Economic.

We evaluated two BERT-like architectures (BERTimbau, BERTimbau Large and multilingual DistilBERT), generating the classifier to the GRI categories. We achieved an accuracy of 89% and an F1-score of 87% using BERTimbau; an accuracy of 87% and an F1-score of 85,33% using BERTimbau Large; and an accuracy of 93% and an F1-score of 90% using DistilBERT multilingual on our validation data.

From our best model, we are currently testing an MVP in a real case analysis conducted at the Brazilian Development Bank (BNDES).

Acknowledgement

First and foremost, we would like to thank Mr. Márcio Onodera, Mr. Fernando Pinhati Jr., and Ms. Vanessa Almeida for providing us with the support and motivation needed to conduct this project. We also like to thank Professor Paolo Torroni from the University of Bologna for Nicolaas' Thesis thorough supervision that inspired this work. We are also grateful to the IT Board at BNDES for their support, contributions and suggestions.

References

Bird, S., Klein, E., & Loper, E. (2009). *Natural Language Processing with Python.* O'ReillyO'Reilly (01 de 04 de 2021). *BNDES Sustainability Bond Framework.* Acesso em 16 de 11 de 2021, disponível em https://www.bndes.gov.br/wps/wcm/connect/site/82a951ca-d413-4abf-8a62-f7337624d4f3/BNDES_011_GEDIT_framework_INGLES_final.29.03.pdf?MOD=AJPERES&CVID=nyed8sy

BNDES. (16 de 11 de 2021). *Sustainable Development*. Fonte: BNDES Homepage: https://www.bndes.gov.br/wps/portal/site/home/desenvolvimento-sustentavel

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Krueger, G. (2020). Language Models are Few-Shot Learners. *Advances in Neural Information Processing Systems*, 1877-1901.

Devlin, J., Chang, M., Lee, K., & Toutanova, K. (2018). BERT: Pretraining of Deep Bidirectional Transformers for Language Understanding. *CoRR*.

Dhar, A., Mukherjee, H., Dash, N., & Kaushik, R. (2021). Text categorization: past and present. *Artificial Intelligence Review volume 54* (pp. 3007–3054). Springer.

Diksha Khurana, A. K. (2022 de 07 de 2022). Natural language processing: state of the art, current trends and challenges. *Multimedia Tools and Applications*, pp. 3713–3744.

Fang, X., Xu, M., & Xu, S. (2019). A deep learning framework for predicting cyber attacks rates. *EURASIP Journal on Information Security*.

Hildebrand, P., Polk, C., Deese, B., & Boivin, J. (02 de 2020). *Sustainability: The tectonic shift transforming investing*. Fonte: BlackRock: https://www.blackrock.com/institutions/en-zz/insights/blackrock-investment-institute/sustainability-in-portfolio-construction

Inderst, G., & Stewart, F. (2018). *Incorporating Environmental, Social and Governance Factors into Fixed Income Investment*. Fonte: World Bank: https://documents1.worldbank.org/curated/en/913961524150628959/pdf/Incorporating-environmental-social-and-governance-factors-into-fixed-income-investment.pdf

Martin, J. H., & Jurafsky, D. (2008). *Speech and Language Processing.* Pearson International Edition.

Meager, E. (31 de 08 de 2021). *Capital Monitor*. Fonte: Capital Monitor AI: https://capitalmonitor.ai/institution/a-guide-to-sustainable-reporting-standards/

Napoletano, E., & Curry, B. (01 de 03 de 2021). *FORBES Advisor*. Fonte: FORBES: https://www.forbes.com/advisor/investing/esg-investing/

Ruberg, N. (3 de 12 de 2021). *Bert goes sustainable: an NLP approach to ESG financing.* Acesso em 30 de 01 de 2023, disponível em AMSLaurea Institutional Theses Repository: http://amslaurea.unibo.it/24787/

Seth, Y. (19 de 07 de 2019). *BERT Explained – A list of Frequently Asked Questions.* Fonte: A blog on data science, machine learning and artificial intelligence: https://yashuseth.blog/2019/06/12/bert-explained-faqs-understand-bert-working/

Souza, F. a. (2020). BERTimbau: Pretrained BERT Models for Brazilian Portuguese. *Intelligent Systems*, pp. 403-417.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Polosukhin, I. (2017). Attention Is All You Need. *Advances in Neural Information Processing Systems 30*, pp. 5998–6008.

wikipedia. (03 de 02 de 2023). *Bag-of-words model.* Fonte: Wikipedia The Free Encyclopedia : https://en.wikipedia.org/wiki/Bag-of-words_model

Wikipedia. (03 de 02 de 2023). *Global Reporting Initiative.* Fonte: Wikipedia- The Free Encyclopedia: https://en.wikipedia.org/wiki/Global_Reporting_Initiative

Wissler, L. &. (2014). The Gold Standard in Corpus Annotation. Passau, Germany: 5th IEEE Germany Student Conference.

Zhu, Y., Kiros, R., Zemel, R., Salakhutdinov, R., Urtasun, R., Torralba, A., & Fidler, S. (22 de 06 de 2015). *Aligning Books and Movies: Towards Story-like Visual Explanations by Watching Movies and Reading Books.* Fonte: arXiv.org: https://arxiv.org/abs/1506.06724