

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Assessing Generative Language Models in Classification Tasks: Performance and Self-evaluation Capabilities in the Environmental and Climate Change Domain

This is a pre print version of the following article:

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/2037963> since 2024-12-13T15:49:36Z

Publisher:

Springer Science and Business Media Deutschland GmbH

Published version:

DOI:10.1007/978-3-031-70242-6_29

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

Assessing Generative Language Models in Classification Tasks: Performance and Self-Evaluation Capabilities in the Environmental and Climate Change Domain

Francesca Grasso^{*[0000–0001–8473–9491]} and Stefano Locci^[0009–0006–9725–2045]

Department of Computer Science, University of Turin
Corso Svizzera 185 - 10149 Turin, Italy
{fr.grasso, stefano.locci}@unito.it

Abstract. This paper examines the performance of two Large Language Models (LLMs) - GPT-3.5-Turbo and Llama-2-13b - and one Small Language Model (SLM) - Gemma-2b, across three different classification tasks within the climate change (CC) and environmental domain. Employing BERT-based models as a baseline, we compare their efficacy against these transformer-based models. Additionally, we assess the models' self-evaluation capabilities by analyzing the calibration of verbalized confidence scores in these text classification tasks. Our findings reveal that while BERT-based models generally outperform both the LLMs and SLM, the performance of the large generative models is still noteworthy. Furthermore, our calibration analysis reveals that although Gemma is well-calibrated in initial tasks, it thereafter produces inconsistent results; Llama is reasonably calibrated, and GPT consistently exhibits strong calibration. Through this research, we aim to contribute to the ongoing discussion on the utility and effectiveness of generative LMs in addressing some of the planet's most urgent issues, highlighting their strengths and limitations in the context of ecology and CC.

Keywords: Large Language Models · Text Classification · Climate Change.

1 Introduction

The advent of Large Language Models (LLMs) in both practical applications and academic research has marked a significant paradigm shift in the Natural Language Processing (NLP) community's focus towards these models and their potential. Recent years have witnessed a surge in studies employing, assessing, evaluating, and leveraging LLMs for various objectives [5]. These include traditional NLP tasks like sentiment analysis [35] and text classification [24], among others. Despite their remarkable ability to generate contextually consistent outputs using natural language and enable a broad range of tasks—from straightforward proofreading to more complex challenges like generating code

^{*} Corresponding Author.

for algorithms—LLMs are not without their flaws, being prone to errors, misunderstandings, and hallucinations [19]. In particular, when it comes to task- or domain-specific performance, especially in classification tasks, LLMs demonstrate conflicting results as compared to other well-grounded machine learning techniques such as transformer-based models and Naive Bayes [34]. If some works highlight LLMs’ proficiency in achieving comparable results to the current state-of-the-art [24, 7], other studies show that their extensive capabilities may lead to suboptimal performance against more specialized models [6, 26]. These observations have driven the NLP community towards exploring also small language models (SLMs), noted for their compact architecture and fewer parameters compared to larger models [34].

When it comes to domain-specific applications, despite the extensive discourse surrounding LLMs, there remains a scarcity of exploration into their application for environmental and climate change (CC)-related texts, a domain of growing significance within the NLP field [25, 12]. The urgency of addressing CC and ecological crisis through NLP techniques has begun to emerge as a crucial point of research. This includes tasks such as sentiment analysis [23], and stance detection [31]. Moreover, the scope of interest should further broaden to encompass a wider array of environmental issues, from deforestation to plastic pollution, underscoring the importance of this research beyond CC to a more holistic environmental perspective [13, 11].

To address these issues, we aim to explore the efficacy of open and close Large and Small Language Models (L/SLMs) in three ecology-related text classification tasks (Eco-Relevance, Environmental Impact Analysis, and Stance Detection)[11]. In particular, we utilize three representative generative models for these categories: GPT-3.5-Turbo-0125 (closed LLM); Llama-2-13b-chat-hf (open mid-LLM), Gemma-2b-it (open SLM), and compare their performance with the baseline consisting of the results for the same tasks of three small, non-generative language models. Moreover, this assessment extends to examining the models’ self-evaluation capabilities through the analysis of their prompt-elicited confidence scores to measure their calibration levels¹. Self-evaluation, the ability of language models to assess the validity of their own outputs, has proven to be crucial for enhancing accuracy and content quality when properly calibrated [21]. In summary, we pose three research questions:

RQ1: How accurately do state-of-the-art generative L/SLMs, both open-source and closed-source, identify texts related to ecology, analyze environmental impacts, and determine the stance on environmental issues?

RQ2: Compared to a baseline set by BERT-based classifiers, do these models demonstrate comparable accuracy in classifying texts within the domain of environment and CC?

RQ3: How well-calibrated are the verbalized confidence scores of LLMs and SLM in these environmental classification tasks?

Our results indicate that while LLMs generally outperform the SLM, BERT-based models still generally surpass both LLMs and SLM. Notably, the perfor-

¹ The code is available here: <https://github.com/stefanolocci/LLMClassification/>

mance of large generative models, especially GPT, is significant, particularly in terms of recall. The calibration assessment revealed mixed results: GPT produced well-calibrated outputs, LLama exhibited moderate calibration, and Gemma struggled to adhere to the prompt template within multilabel settings. These findings aim to contribute to our understanding of the capabilities of L/SLMs in environmental text classification, enriching the ongoing discussion on leveraging NLP for ecological and CC research.

2 Related Works

Given the considerable amount of work and vertiginous growth of material on the topic of LLMs, we report a limited number of studies relevant to our research. **LLMs in Classification Task** Text Classification, a fundamental NLP task, has progressed from traditional methods such as Naïve Bayes [18], to advanced models like BERT [9, 17]. Recent studies on LLMs in text classification present mixed outcomes when compared to traditional supervised methods. For instance, [34] explores the performance of different classes of language models in text classification, often favoring smaller models. [2] investigates LLMs across languages and tasks without a clear best performer, while [6] discusses ChatGPT’s limitations in biomedical fields compared to domain-specific models. Nonetheless, advances in prompting techniques, such as Clues and Reasoning prompts [6], and strategic prompts [7], show potential in improving LLMs’ classification abilities. [24] highlights GPT-4’s comparable performance in propaganda detection to the current state-of-the-art.

LLMs Self Evaluation The significant hype around LLMs has encouraged numerous studies on their capabilities [5] and self-evaluation skills, including calibration—the alignment between a model’s confidence and its prediction accuracy. Calibration is crucial for a model’s reliability, especially in determining when to defer to an expert [29]. Studies [29, 16] indicate LLMs’ verbalized confidences are often better calibrated than conditional probabilities, yet [33, 21] highlight their tendency towards overconfidence. Our research evaluates the accuracy and calibration quality of LLMs’ verbalized confidence scores.

LLMs in Environmental Domain Regarding the intersection of LLMs and the environmental domain, similar to other NLP subfields, efforts in this direction have been focused mostly on climate change (CC) topics rather than the broader environmental and ecological fields (which still include CC). Among the notable works, [4] and [36] propose evaluation frameworks for analyzing LLM responses to CC topics; [14] built a prototype tool for localized climate-related data leveraging LLMs; [15] developed an AI-based tool for fact-checking of CC claims utilizing an array of LLMs; [28] introduced a family of domain-specific LLMs designed to synthesize interdisciplinary research on climate change. Unfortunately, these models require significant hardware resources to operate, making them difficult to access.

3 Methodology

3.1 Baselines and Generative LMs employed

To establish baselines for assessing the added value of Large and Small Language Models (L/SLMs) in accurately classifying texts within the ecological domain, we reference the performance of six pre-trained BERT-based models. These models have been utilized in recent research across three classification tasks: binary Eco-Relevance classification, multilabel Environmental Impact Analysis, and Stance Detection[11]. Key performance metrics from the studies include:

- *BERT and RoBERTa*[10, 17]: Masked language models trained on large English corpora. In the Stance Detection task, RoBERTa achieved a high accuracy of 81.29%, matching DistilRoBERTa, with BERT demonstrating high precision and F1 scores (95.09% and 95.56%, respectively).
- *DistilRoBERTa* [22]: An efficient adaptation of RoBERTa, showing the highest accuracy in the Eco-Relevance task at 89.43%. It maintained competitive performance across all tasks, highlighting its efficiency.
- *ClimateBERT*[32]: Includes variants ClimateBertF, ClimateBertS, and ClimateBertS+D, pre-trained on climate-specific texts. In the Environmental Impact Analysis, ClimateBertS led in accuracy (78.62%) and achieved the best F1 measure (54.67%) among the ClimateBert variants.

These results set a comprehensive baseline for our study, allowing for a direct comparison of the effectiveness of L/SLMs in similar tasks². To account for a heterogeneous range of generative language models, we selected three different types of language models for our evaluation, drawing inspiration from [34], to cover closed, open and small-scale models:

- **GPT-3.5-Turbo-0125**. An advanced iteration within the GPT series developed by OpenAI [1], representing high-capacity closed generative LLMs. Given its widespread use in both public and academic contexts, it serves as our benchmark for closed generative LLMs.
- **Llama-2-13b-chat-hf**. A variant from Meta AI’s Llama-2 model series [30]. Designed for conversational applications, it exemplifies mid-to-large open-domain language models with its capability for human-like text generation and comprehension.
- **Gemma-2b-it**. A newly introduced small language model (SLM) by Google, derived from the Gemma family [27]. Gemma is a family of lightweight, state-of-the-art open models built from the same research and technology used to create the larger Gemini models.

3.2 Classification Tasks and Prompt Design

Tasks Description Numerous studies have showcased the effectiveness of few-shot learning in Language Models [3, 7]. In our study, we conducted a few-shot,

² Here we report just the most informative performance score. The other scores are present in the referenced paper.

Table 1: Prompt for the EIA multilabel classification task, detailing instructions, labels, and a structured response format. Adapted versions were used for Eco-Relevance and Stance Detection tasks with task-specific modifications.

Prompt
<p>Perform a multilabel classification on the provided tweet from an ecological perspective. An instance could discuss topics that are either (potentially or explicitly) harmful or beneficial to the sustainability and well-being of the natural world. Pay extreme attention to the mere content of the text and ignore completely the stance of who is writing the post. Analyzing the content make inferences also on your background knowledge of green and environment topics. Consider paramount that this level is NOT a sentiment analysis.</p> <p>The labels you must use for classification are: "POS": The tweet content is potentially or explicitly beneficial to the sustainability and well-being of the natural world. "NEU": The tweet content is not beneficial nor harmful to the sustainability and well-being of the natural world. "NEG": The tweet content is potentially or explicitly harmful to the sustainability and well-being of the natural world.</p> <p>Your response must strictly follow this format, pay attention to the order of the fields in the answer format independently from the assigned label, the MUST always be as follow: LABEL: <[POS/NEU/NEG]> PROB POS_PROB=<probability of Positive>, NEU_PROB=<probability of Neutral>, NEG_PROB=<probability of Negative>; EXP: <A natural language explanation detailing why the probabilities were assigned as such>.</p> <p>Your answer must includes explicit probabilities for all the "POS", "NEU" and "NEG" labels, reflecting the likelihood of the tweet content being Beneficial (POS), Neutral (NEU) or Negative (NEG) towards the environment. You must select ONE and ONLY ONE LABEL between POS, NEU and NEG for each single tweet.</p> <p>For clarity, here are examples illustrating the expected format (Note: Examples provide the LABEL only for brevity):</p> <p>- <tweet_example.1> [POS] - <tweet_example.2> [POS] - <tweet_example.3> [NEU] - <tweet_example.4> [NEU] - <tweet_example.5> [NEG] - <tweet_example.6> [NEG]</p> <p>Analyze and classify the following tweet according to these guidelines: <tweet></p>

three-layer text classification of tweets using the EcoVerse Dataset [11], comprised of 3k tweets, guiding the L/SLMs to perform classifications on three tasks.

Eco-Relevance: The first task involves binary classification to identify texts related to Ecology. The labels for this level are **eco-related** or **not eco-related**.

Environmental Impact Analysis (EIA): The second multilabel task determines, for eco-related tweets, whether the post conveys behaviors or events with **positive**, **negative**, or **neutral** impacts on the environment.

Stance Detection: The third level identify the stance expressed by the tweet’s author as **supportive**, **neutral**, or **skeptical/opposing** towards environmental causes.

For the few-shot examples, we randomly selected from the training set two examples for each label. Although more complex example sampling strategies have proven effective in scenarios where the demonstrative examples and the input text are significantly semantically divergent [26], we chose random sampling because the tweet topics in this case are relatively homogeneous.

Prompt Design Prompt Engineering is key to enhancing LLMs’ precision, with prompt wording significantly influencing model’s reasoning [7]. By incorporating domain-specific knowledge, this technique helps LLMs match or exceed traditional models in efficiency, often with less data [8]. In text classification, the clarity of prompts is essential for generating categorical outputs. Our objective was to design prompts that yield discrete labels, allowing us to: (i) evaluate LLMs’ classification performance on the EcoVerse dataset metrics like Accuracy and Precision; (ii) benchmark these results against BERT-based models for the same tasks; and (iii) assess the models’ self-evaluation by examining the calibration of their verbalized confidence scores. Accordingly, the prompts were designed to produce direct outputs for quantitative analysis; be comparable with BERT-based models through a few-shot approach; and include verbalized confidence for calibration checks. To facilitate future qualitative analyses, especially regarding classification errors, prompts also requested a rationale for the model’s choices. Table 1 illustrates the comprehensive prompt for the EIA multilabel task. Adaptations of this prompt for the Eco-Relevance and Stance Detection tasks were similarly structured but customized with task-specific wording and examples.

3.3 Experimental Setup

We conducted our experiments on the Paperspace platform³, utilizing a configuration that includes an NVIDIA A100 GPU with 80GB of VRAM, 90GB of RAM, and a 12-core CPU. We employed the vLLM python interface⁴ to load the llama-2-13b-chat-hf and gemma-2b-it models. For the GPT-3.5-turbo model, we utilized OpenAI’s APIs⁵. We set a relatively low temperature of 0.3 for each model, capping the output at 512 tokens. This low temperature setting was chosen to reduce the models’ ”creativity,” aiming for more consistent results on classification side. It’s important to note that even setting the temperature to 0 can still result in some randomness in the models’ outputs. To mitigate the impact of this unpredictability on performance and to obtain statistically significant results, similarly to what performed in other similar studies [20] we executed each experiment 100 times for each of the three tasks and calculated the average metrics, which are presented in Table 2.

4 Results and Discussion

4.1 Models’ Classification Performance

To evaluate the performance of the models, we utilized the following set of metrics: Accuracy, Precision, Recall, and the F1-score. For tasks involving multiple labels, we provided a comprehensive analysis by employing the macro-average versions of these metrics to offer a global perspective on model performance. The results for the three distinct tasks are presented and discussed below.

³ <https://www.paperspace.com/>

⁴ <https://blog.vllm.ai/2023/06/20/vllm.html>

⁵ <https://openai.com/blog/openai-api>

Eco-Relevance Task As shown in Table 2, the GPT-3.5-Turbo-0125 model demonstrates superior performance compared to the other LLMs across all metrics, notably outperforming also BERT-based models in terms of recall. This suggests that GPT-3.5 is not only precise in identifying relevant instances but also excels at recognizing positive instances while minimizing false positives. The Llama-2-13b-chat-hf model follows with commendable results, especially considering the recall measure where, along with GPT-3.5, surpass RoBERTa. This indicates its proficiency in identifying eco-relevant instances. However, its precision is notably lower, which is reflected in a marginally lower F1-score. This indicates that while Llama-2 is effective at detecting relevant instances, it is prone to incorrectly classifying non-eco-relevant instances as relevant. Gemma-2b-it significantly lags behind the other two models across all metrics, possibly given its smaller scale. Despite the notable results of GPT-3.5 in this first classification task, it does not surpass the most performant language model of our selected baseline models, RoBERTa.

Table 2: Comparative results of models across the three different tasks.

Task	Model	Precision	Recall	F-1 score	Accuracy
Eco-Relevance	<i>RoBERTa</i>	88.90%	88.96%	88.93%	88.87%
	Llama-2	62,85%	95,37%	75,76%	73,42%
	Gemma	40,90%	38,49%	39,65%	48,98%
	GPT-3.5	72,12%	96,97%	82,72%	82,34%
EIA	<i>ClimBerts</i>	51.81%	57.78%	54.67%	78.62%
	Llama-2	34,15%	30,73%	28,48%	31,98%
	Gemma	40,42%	39,64%	39,69%	44,87%
	GPT-3.5	68,56%	66,69%	66,75%	84,07%
Stance Detect.	<i>BERT</i>	95.09%	96.04%	95.56%	74.27%
	Llama-2	26,33%	27,14%	26,47%	31,65%
	Gemma	22,80%	37,30%	24,90%	27,60%
	GPT-3.5	56,85%	70,66%	56,28%	65,87%

Environmental Impact Analysis In this task, GPT-3.5 outperforms both the other generative models and the top performing model in EIA ClimateBerts, achieving the highest scores in all evaluated metrics. This notable performance could be attributed to its advanced reasoning capabilities which likely played a crucial role in investigating the complexities of EIA. This task can be challenging as it necessitates a comprehensive understanding that extends beyond linguistic patterns to include extra-linguistic contextual comprehension, meaning that a model must not only parse and understand the text but also relate it to broader environmental contexts.

Stance Detection For the Stance Detection task the performances of all generative models appear to be less optimal, with GPT-3.5 leading among them but drastically underperforming compared to the BERT-like baseline models. Llama-2 and Gemma performed significantly lower across all metrics, with Llama-2 achieving a slightly better balance between precision and recall. The underperformance of GPT-3.5 in this task might suggest that the specific characteristics and the fine-tuning of BERT models lend it an advantage in stance detection. It

also indicates that the advanced reasoning capabilities and broader knowledge base of GPT-3.5, while beneficial for many tasks, may not always translate to superior performance in tasks where domain-specific training and optimization may play a crucial role.

4.2 Calibration Evaluation

For the calibration assessment, we employed the prompt-elicitation method, as detailed in Table 1, to extract verbalized confidence scores from the models, expressed as output tokens. These scores are compared with actual precision measures to assess model calibration effectively [29]. The evaluation process delves into how well the models’ stated confidences align with their empirical precision across three tasks. This method provides insights into the consistency of models’ self-assessed confidences over 100 iterations, highlighting any mismatches between their provided explanations and the probability scores. We categorized the scores into five probability bins: [‘0-0.2’, ‘0.2-0.4’, ‘0.4-0.6’, ‘0.6-0.8’, ‘0.8-1.0’] then calibration was evaluated by comparing the mean probabilities with the observed precision for each bin. Below we detail the calibration performance of each model, focusing specifically on the task that demonstrated the most accurate calibration outcomes.

GPT demonstrates well-calibrated outputs across all tasks, with specific results for the stance detection task presented in Table 3. The stance classification exhibits high confidence levels for the supportive and skeptical/opposing labels, albeit encountering some difficulty with the neutral stance. This results indicates that the probabilities are well-calibrated in each bin, indicating that GPT provides confident responses for correct answers, and similarly, manifests considerable uncertainty when the answers are incorrect.

Table 3: GPT-3.5 calibration results on Stance Detection Task. The table shows for each bin, their mean probability (Pr) for each label [support./neutral/skept-oppos.]. Alongside are presented the Precision (P) scores for each label.

Bins	Bin Mean Pr	P _{sup}	P _{neu}	P _{ske}
0.0 - 0.2	[0,15/0,06/0,03]	0	0	0
0.2 - 0.4	[0,33/0,25/0,2]	0	0	0
0.4 - 0.6	[0,4/0,54/0,60]	0	0	0
0.6 - 0.8	[0,01/0,60/0,60]	0	0,62	0,71
0.8 - 1.0	[0,91/0,01/0,82]	0,83	0,34	0,88

LLaMA’s calibration is notably moderate, with its most reliable performance observed in the stance detection task. Table 4 reveals that model’s probabilities are well-calibrated except for the **neutral** label where we observe poor calibration, most notably in the highest probability bin of 0.8 - 1.0 indicating that the model confidently made incorrect decisions.

Gemma exhibits good calibration at the Eco-Relevance task, as shown in Table 5. However, we were unable to compute it for subsequent levels as it failed to follow the prompt template, displaying too much inconsistency within

Table 4: LLama-2 calibration results on Stance Detection Task

Bins	Bin Mean Pr	P_{sup}	P_{neu}	P_{ske}
0.0 - 0.2	[0,09/0,09/0,06]	0	0	0
0.2 - 0.4	[0,21/0,25/0,2]	0	0	0
0.4 - 0.6	[0,47/0,4/0,0]	0	0	0
0.6 - 0.8	[0,6/0,61/0,80]	0,55	0,28	0,85
0.8 - 1.0	[0,8/0,8/0,8]	0,88	0,57	0,89

the same iteration and across different iterations. This discrepancy compared to its larger counterparts can primarily be attributed to the architectural and training differences inherent to SLMs. To ensure a fair comparative analysis, we deliberately chose to maintain uniform testing conditions across all models. This approach was aimed at assessing each model’s ability to adapt to standardized tasks without necessitating model-specific prompt optimizations. The variations observed in Gemma’s responses, particularly its inconsistent results in multi-label score distribution, underscore a critical insight: prompt engineering may need to be tailored to accommodate the limitations of smaller generative models.

Table 5: Gemma calibration results on Eco-Relevance Task

Bins	Bin Mean Pr	P_{eco_rel}	P_{not_eco_rel}
0.0 - 0.2	[0/0,03]	0	0
0.2 - 0.4	[0,26/0,25]	0	0
0.4 - 0.6	[0,49/0,50]	0,59	0,54
0.6 - 0.8	[0,74/0,74]	0,9	0,88
0.8 - 1.0	[0,92/0,94]	0,96	0,95

5 Conclusion

In this study, we addressed three Research Questions (RQs) regarding the performance of advanced generative language models—GPT-3.5.Turbo-0125, Llama-2-13b-chat-hf, and Gemma-2b-it—in ecological domain tasks: Eco-Relevance, Environmental Impact Analysis (EIA), and Stance Detection. Our findings reveal: **RQ1** showed variable effectiveness across models, with GPT demonstrating particular strength in Eco-Relevance, outperforming others. **RQ2** assessed if these models could exceed a BERT-based benchmark. GPT excelled in the EIA task, yet no model consistently surpassed all baseline metrics. Gemma significantly underperformed compared to the other models. **RQ3** evaluated the calibration of verbalized confidence scores, noting GPT’s consistent reliability. Conversely, LLaMA and Gemma’s calibration varied, indicating a need for refinement. This study highlights the strengths and areas for improvement of generative models in environmental classification, contributing to the dialogue on NLP’s role in addressing ecological and climate issues, and underscores the potential of LLMs in domain-specific tasks.

References

1. Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F.L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al.: Gpt-4 technical report. arXiv preprint arXiv:2303.08774 (2023)
2. del Arco, F.M.P., Nozza, D., Hovy, D.: Leveraging label variation in large language models for zero-shot text classification. ArXiv **abs/2307.12973** (2023)
3. Brown, T.B., et al.: Language models are few-shot learners. ArXiv **abs/2005.14165** (2020)
4. Bulian, J., Schäfer, M.S., Amini, A., Lam, H., Ciaramita, M., Gaiarin, B., Huebscher, M.C., Buck, C., Mede, N., Leippold, M., et al.: Assessing large language models on climate information. arXiv preprint arXiv:2310.02932 (2023)
5. Chang, Y.C., Wang, X., Wang, J., Wu, Y., Zhu, K., Chen, H., Yang, L., Yi, X., Wang, C., Wang, Y., Ye, W., Zhang, Y., Chang, Y., Yu, P.S., Yang, Q., Xie, X.: A survey on evaluation of large language models. ArXiv **abs/2307.03109** (2023)
6. Chen, S., Li, Y., Lu, S., Van, H., Aerts, H.J., Savova, G.K., Bitterman, D.S.: Evaluation of chatgpt family of models for biomedical reasoning and classification. Journal of the American Medical Informatics Association : JAMIA (2023)
7. Clavié, B., Ciceu, A., Naylor, F., Souli’e, G., Brightwell, T.: Large language models in the workplace: A case study on prompt engineering for job type classification. ArXiv **abs/2303.07142** (2023)
8. Deldjoo, Y.: Fairness of chatgpt and the role of explainable-guided prompts. ArXiv **abs/2307.11761** (2023)
9. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. In: North American Chapter of the Association for Computational Linguistics (2019)
10. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: Pre-training of deep bidirectional transformers for language understanding. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1. pp. 4171–4186. Association for Computational Linguistics, Minneapolis, Minnesota (Jun 2019)
11. Grasso, F., Locci, S., Siragusa, G., Caro, L.D.: Ecoverse: An annotated twitter dataset for eco-relevance classification, environmental impact analysis, and stance detection. In: Proceedings of the Fourteenth Language Resources and Evaluation Conference (forth.). European Language Resources Association, Turin, Italy (May 2024)
12. Hershovich, D., Webersinke, N., Kraus, M., Bingler, J.A., Leippold, M.: Towards climate awareness in nlp research. In: Conference on Empirical Methods in Natural Language Processing (2022)
13. Ibrohim, M.O., Bosco, C., Basile, V.: Sentiment analysis for the natural environment: a systematic review. ACM Computing Surveys (2023)
14. Koldunov, N., Jung, T.: Local climate services for all, courtesy of large language models. Communications Earth & Environment **5**(1), 13 (2024)
15. Leippold, M., Vaghefi, S.A., Stambach, D., Muccione, V., Bingler, J., Ni, J., Colesanti-Senni, C., Wekhof, T., Schimanski, T., Gostlow, G., et al.: Automated fact-checking of climate change claims with large language models. arXiv preprint arXiv:2401.12566 (2024)
16. Lin, S.C., Hilton, J., Evans, O.: Teaching models to express their uncertainty in words. Trans. Mach. Learn. Res. **2022** (2022)

17. Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V.: Ro{bert}a: A robustly optimized {bert} pretraining approach (2020)
18. McCallum, A., Nigam, K.: A comparison of event models for naive bayes text classification. In: AAAI Conference on Artificial Intelligence (1998)
19. Mittal, A., Murthy, R., Kumar, V., Bhat, R.: Towards understanding and mitigating the hallucinations in nlp and speech. In: Proceedings of the 7th Joint International Conference on Data Science & Management of Data (11th ACM IKDD CODS and 29th COMAD). pp. 489–492 (2024)
20. Motoki, F., Neto, V.P., Rodrigues, V.: More human than human: measuring chatgpt political bias. *Public Choice* **198**, 3–23 (2023)
21. Ren, J., Zhao, Y., Vu, T., Liu, P.J., Lakshminarayanan, B.: Self-evaluation improves selective generation in large language models. ArXiv **abs/2312.09300** (2023)
22. Sanh, V., Debut, L., Chaumond, J., Wolf, T.: Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. arXiv preprint arXiv:1910.01108 (2019)
23. Sham, N.M., Mohamed, A.H.: Climate change sentiment analysis using lexicon, machine learning and hybrid approaches. *Sustainability* (2022)
24. Sprenkamp, K., Jones, D.G., Zavolokina, L.: Large language models for propaganda detection. ArXiv **abs/2310.06422** (2023)
25. Stede, M., Patz, R.: The climate change debate and natural language processing. Proceedings of the 1st Workshop on NLP for Positive Impact (2021)
26. Sun, X., Li, X., Li, J., Wu, F., Guo, S., Zhang, T., Wang, G.: Text classification via large language models. In: Conference on Empirical Methods in Natural Language Processing (2023)
27. Team, G., Mesnard, T., Hardin, C., Dadashi, R., Bhupatiraju, S., Pathak, S., Sifre, L., Rivière, M., Kale, M.S., Love, J., et al.: Gemma: Open models based on gemini research and technology. arXiv preprint arXiv:2403.08295 (2024)
28. Thulke, D., et al.: Climategpt: Towards ai synthesizing interdisciplinary research on climate change. ArXiv **abs/2401.09646** (2024)
29. Tian, K., Mitchell, E., Zhou, A., Sharma, A., Rafailov, R., Yao, H., Finn, C., Manning, C.D.: Just ask for calibration: Strategies for eliciting calibrated confidence scores from language models fine-tuned with human feedback. ArXiv **abs/2305.14975** (2023)
30. Touvron, H., et al.: Llama 2: Open foundation and fine-tuned chat models. ArXiv **abs/2307.09288** (2023)
31. Upadhyaya, A., Fisichella, M., Nejd, W.: A multi-task model for sentiment aided stance detection of climate change tweets. ArXiv **abs/2211.03533** (2022)
32. Webersinke, N., Kraus, M., Bingler, J.A., Leippold, M.: Climatebert: A pretrained language model for climate-related text. ArXiv **abs/2110.12010** (2021)
33. Xiong, M., Hu, Z., Lu, X., Li, Y., Fu, J., He, J., Hooi, B.: Can llms express their uncertainty? an empirical evaluation of confidence elicitation in llms. ArXiv **abs/2306.13063** (2023)
34. Yu, H., Yang, Z., Pelrine, K., Godbout, J.F., Rabbany, R.: Open, closed, or small language models for text classification? ArXiv **abs/2308.10092** (2023)
35. Zhang, W., Deng, Y., Liu, B.Q., Pan, S.J., Bing, L.: Sentiment analysis in the era of large language models: A reality check. ArXiv **abs/2305.15005** (2023)
36. Zhu, H., Tiwari, P.: Climate change from large language models. ArXiv **abs/2312.11985** (2023)