Understanding Climate Legislation Decisions with Machine Learning

Jeffrey N. Clark
University of Bristol, UK
jeff.clark@bristol.ac.uk

Michelle W.L. Wan
University of Bristol, UK
michelle.wan@bristol.ac.uk

Raúl Santos-Rodríguez
University of Bristol, UK
enrsr@bristol.ac.uk

Abstract

Effective action is crucial in order to avert climate disaster. Key in enacting change is the swift adoption of climate positive legislation which advocates for climate change mitigation and adaptation. This is because government legislation can result in far-reaching impact, due to the relationships between climate policy, technology, and market forces. To advocate for legislation, current strategies aim to identify potential levers and obstacles, presenting an opportunity for the application of recent advances in machine learning language models. Here we propose a machine learning pipeline to analyse climate legislation, aiming to investigate the feasibility of natural language processing for the classification of climate legislation texts, to predict policy voting outcomes. By providing a model of the decision making process, the proposed pipeline can enhance transparency and aid policy advocates and decision makers in understanding legislative decisions, thereby providing a tool to monitor and understand legislative decisions towards climate positive impact.

1 Introduction

Legislation is a key lever in fighting climate change. While climate policies lay out a plan of action, climate legislation is enforceable law. Legislation therefore has the potential to both help and hinder climate-positive progress, impacting societal change, public infrastructure, and industry.

Despite recent commitments, current climate change mitigation efforts remain insufficient [11]. Existing solutions to fight climate change include low-carbon energy generation, tightening emissions regulations, limiting the destruction of carbon sequestering environments, and reducing agricultural emissions; these could halve global emissions, yet remain largely underutilised [12]. In large part, this is an issue of political will and legislative implementation [5]. However, history has shown that national efforts can be rapidly mobilised, such as during World War II, or the COVID-19 pandemic. Political decisions could therefore also rapidly utilise these existing solutions, and accelerate the development and deployment of new technologies, to fight climate change.

Improved transparency around legislative decisions can aid decision makers in progress towards climate change mitigation and adaptation. To this end, policy advocates already attempt to keep track of policy approvals and their driving factors. However, understanding and predicting legislative decisions is highly complex. Machine learning offers techniques towards achieving this, through analysis, modelling, and explaining these decisions.

Recent advances in machine learning have enabled the development of sophisticated tools to help tackle climate change [14], in domains from urban planning [10] to precision agriculture [9]. Of
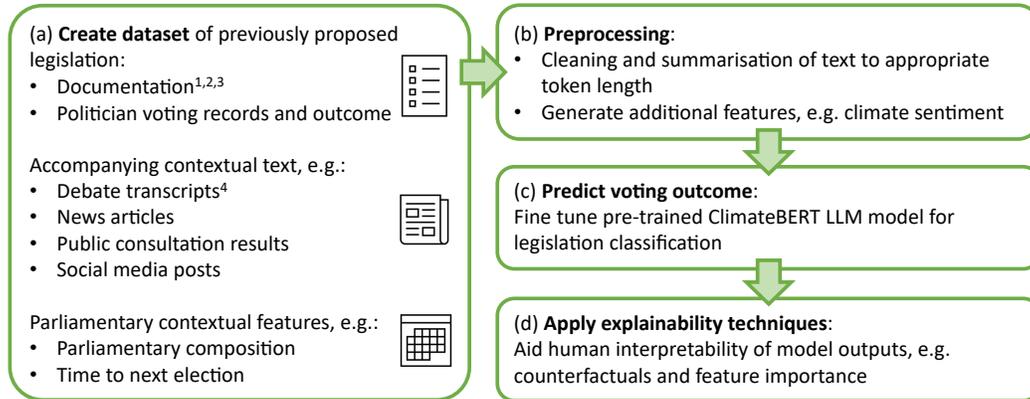
Figure 1: Proposed pipeline. LLM = Large language model.

particular relevance here, natural language processing (NLP) aims to enable machines to understand, process and generate human language. Increased computing power, coupled with transformer-based architectures, have enabled recent large language models (LLMs) such as GPT-4, PaLM 2, and LLaMA, to demonstrate unprecedented performance. Climate change-related applications of NLP have included understanding the effects of temperature on hate speech [17], and efforts are now emerging to utilise NLP specifically around climate legislation and policy, including: automatic summarisation [22], tracking national changes [3], automated policy evaluation [1], monitoring public discussion [7], classification of Paris Agreement Climate Action Plans [4], informing policy relating to climate justice [6], and integrating climate change adaptation policy [2].

In this contribution we propose a pipeline utilising NLP to better understand climate legislation decisions. We aim to investigate the feasibility of NLP for the classification of voting outcomes on proposed legislation. Applying established explainability techniques to the trained classifier can then aid understanding of the model and therefore provide insight into the driving factors of voting decisions themselves. The findings from implementing the pipeline could aid policy experts in legislation advocacy, while improving transparency, accountability, and community engagement; thereby empowering efforts towards climate change mitigation and adaptation.

2 Proposed methods

The proposed approach is comprised of four components:

Create a dataset (Figure 1a). Historically proposed climate legislation and associated parliamentary votes, such as are available online for the UK[12] and European Parliaments[3]. Accompanying texts are collected for climate sentiment and topic analysis. To enhance trustworthiness and reduce bias impacts, these should encompass a range of sources, including: parliamentary debate transcripts (e.g. Hansard[4]), news outlet articles, public consultation results, and social media posts. Utilising elements of computational argumentation [18], could enhance understanding of opinions and climate stances, with greater speed and scalability than current manual forms of gathering opinion data. Features relevant to parliamentary context will be included, such as legislation proposers, upcoming election time frames, and the size of the potential majority in government, which is particularly relevant when whips anticipate voting along party lines.

Summarise the collected text (Figure 1b). Climate legislation documents are extensive, of greater length than many LLMs allow. Summarisation reduces token length [13] and enables input to an LLM, using extractive and hybrid architectures for text summarisation [20]. Additionally, summarisation

---

[1]https://bills.parliament.uk

[2]https://votes.parliament.uk/Votes/Commons

[3]https://www.europarl.europa.eu/plenary/en/home.html

[4]https://hansard.parliament.uk

2

distills the key ideas of legislation texts, by filtering out unwanted noise associated with specific wording and retaining useful signals associated with document content.

Train a classifier   (Figure 1c). To predict the voting outcomes for legislation, an LLM pre-trained for climate tasks (e.g. ClimateBERT [19], or CliMedBERT [8]), is applied after fine tuning. Outcomes are assigned appropriate labels led by experts, such as pass/fail/on hold. LLMs have been shown to perform well for both binary and multiclass classification when trained with suitable legal domain-specific features [16]. Human expertise will be incorporated through feedback to compliment the automated approach.

Apply interpretability and explainability techniques   (Figure 1d). Established techniques for NLP can highlight important elements of the summarised text, determine feature importance for accompanying variables such as public opinion or the size of government majority, and identify counterfactuals which may aid humans in the drafting stage of proposals [21]. Results of explainability analyses may also provide insight into trends in historical climate legislation, by identifying patterns within voting outcomes, thereby aiding transparency and providing insights for the general public.

3   Responsible AI and risks

Responsible implementation of the proposed pipeline must address:

Potential misuse by bad actors   We propose methods to improve understanding of the legislation approval process. Although intended for climate positive purposes, implementation could aid actors with opposing intentions. While this is an inherent risk for existing tools used by policy advocates, AI tools evolve rapidly, often overtaking regulation. Deployment here must therefore prioritise robust ML approaches which are explainable and human-centric. An explainable open-source pipeline can thus aid transparency, building social trust in leglislative decision making processes.

Incomplete representations of legislative process   Documentation and public debates cannot capture the full spectrum of true opinions and intentions of voting politicians. The classifier therefore receives an incomplete picture of the legislative process, in turn affecting performance. Inclusion of the amendments process and gradual adaptation of policies may useful provide insights [15].

Dynamic political environments   Political environments and public opinion are dynamic, affecting the accuracy and validity of classifications. To mitigate inaccuracies during a change in power, models can include voting records and parliamentary composition. Furthermore, the pipeline uses readily available data sources which can be processed automatically, narrowing the inaccuracy window.

4   Expected outputs and beneficiaries

Following development and implementation of the proposed pipeline we intend to publish our findings, and make the code openly available. Legislation is a central factor in environmental progress, ultimately driving the global approach and trajectory for humanity's response to climate change. Improving our understanding of climate legislation decisions could enable improved engagement with the legislative process. By improving transparency and understanding of these decisions, we can build social trust. Furthermore, by analysing policies within the context of social media posts, public consultation transcripts, and news articles, we can additionally gain a better understanding of public sentiment and opinion, which in turn can shape advocacy campaigns by more effectively utilising public support as a policy lever. The pipeline can assist policy experts with advocacy by building an understanding of the characteristics which drive the acceptance of a particular proposal, enabling informed iteration on their proposed legislation to improve chances of adoption. With this in mind, the proposed pipeline seeks to assist lawmakers, and wider society, in work towards cross-sectoral climate mitigation and adaptation, to drive climate positive impact.

Acknowledgments and Disclosure of Funding

References

[1] J. Aitken. Ai-based text analysis for evaluating food waste policies. In *AAAI 2022 Fall Symposium: The Role of AI in Responding to Climate Challenges*, 2022. URL `https://www.climatechange.ai/papers/aaaifss2022/1`.

[2] R. Biesbroek, S. Badloe, and I. N. Athanasiadis. Machine learning for research on climate change adaptation policy integration: an exploratory uk case study. *Regional Environmental Change*, 20(3):85, 2020.

[3] F. N. Cação, A. H. Reali Costa, N. Unterstell, L. Yonaha, T. Stec, and F. Ishisaki. Deeppolicy-tracker: Tracking changes in environmental policy in the brazilian federal official gazette with deep learning. In *ICML 2021 Workshop on Tackling Climate Change with Machine Learning*, 2021. URL `https://www.climatechange.ai/papers/icml2021/35`.

[4] T. Corringham, D. Spokoyny, E. Xiao, C. Cha, C. Lemarchand, M. Syal, E. Olson, and A. Gershunov. Bert classification of paris agreement climate action plans. In *ICML 2021 Workshop on Tackling Climate Change with Machine Learning*, 2021. URL `https://www.climatechange.ai/papers/icml2021/45`.

[5] A. Gore. Al gore at davos 2023, "we have the tech to slow climate change – so where's political will", Jan 2023. URL `https://tinyurl.com/vbjn63eh`.

[6] L. Hampton. Machine learning informed policy for environmental justice in atlanta with climate justice implications. In *NeurIPS 2020 Workshop on Tackling Climate Change with Machine Learning*, 2020. URL `https://www.climatechange.ai/papers/neurips2020/88`.

[7] V. Hätönen and F. Melzer. From talk to action with accountability: Monitoring the public discussion of policy makers with deep neural networks and topic modelling. In *ICML 2021 Workshop on Tackling Climate Change with Machine Learning*, 2021. URL `https://www.climatechange.ai/papers/icml2021/75`.

[8] B. Jalalzadeh Fard, S. A. Hasan, and J. E. Bell. Climedbert: A pre-trained language model for climate and health-related text. In *NeurIPS 2022 Workshop on Tackling Climate Change with Machine Learning*, 2022. URL `https://www.climatechange.ai/papers/neurips2022/110`.

[9] K. G. Liakos, P. Busato, D. Moshou, S. Pearson, and D. Bochtis. Machine learning in agriculture: A review. *Sensors*, 18(8):2674, 2018.

[10] N. Milojevic-Dupont and F. Creutzig. Machine learning for geographically differentiated climate change mitigation in urban areas. *Sustainable Cities and Society*, 64:102526, 2021.

[11] H.-O. Pörtner, D. C. Roberts, H. Adams, C. Adler, P. Aldunce, E. Ali, R. A. Begum, R. Betts, R. B. Kerr, R. Biesbroek, et al. *Climate change 2022: Impacts, adaptation and vulnerability*. IPCC Geneva, Switzerland:, 2022.

[12] H.-O. Pörtner, D. C. Roberts, E. Poloczanska, K. Mintenbeck, M. Tignor, A. Alegría, M. Craig, S. Langsdorf, S. Löschke, V. Möller, et al. *IPCC, 2022: Summary for policymakers*. Cambridge University Press, 2022.

[13] K. Rathi, Y. V. Singh, and S. Raj. A review of state-of-the-art automatic text summarisation. *International Journal of Creative Research Thoughts (2022)*, 10(4), 2022.

[14] D. Rolnick, P. L. Donti, L. H. Kaack, K. Kochanski, A. Lacoste, K. Sankaran, A. S. Ross, N. Milojevic-Dupont, N. Jaques, A. Waldman-Brown, et al. Tackling climate change with machine learning. *ACM Computing Surveys (CSUR)*, 55(2):1–96, 2022.

[15] M. Russell and D. Gover. *Legislation at Westminster: Parliamentary actors and influence in the making of British law*. Oxford University Press, 2017.

[16] D. Song, S. Gao, B. He, and F. Schilder. On the effectiveness of pre-trained language models for legal natural language processing: An empirical study. *IEEE Access*, 10:75835–75858, 2022.

[17] A. Stechemesser, A. Levermann, and L. Wenz. Temperature impacts on hate speech online: evidence from four billion tweets. In *NeurIPS 2022 Workshop on Tackling Climate Change with Machine Learning*, 2022. URL https://www.climatechange.ai/papers/neurips2022/41.

[18] F. Toni. A tutorial on assumption-based argumentation. *Argument & Computation*, 5(1):89–117, 2014.

[19] N. Webersinke, M. Kraus, J. A. Bingler, and M. Leippold. Climatebert: A pretrained language model for climate-related text. *arXiv preprint arXiv:2110.12010*, 2021.

[20] F. Zaman, M. Shardlow, S.-U. Hassan, N. R. Aljohani, and R. Nawaz. Htss: A novel hybrid text summarisation and simplification architecture. *Information Processing & Management*, 57(6): 102351, 2020. doi: 10.1016/j.ipm.2020.102351.

[21] J. E. Zini and M. Awad. On the explainability of natural language processing deep models. *ACM Computing Surveys*, 55(5):1–31, 2022.

[22] A. Żółkowski, M. Krzyziński, P. Wilczyński, S. Giziński, E. Wiśnios, B. Pieliński, J. Sienkiewicz, and P. Biecek. Climate policy tracker: Pipeline for automated analysis of public climate policies. In *NeurIPS 2022 Workshop on Tackling Climate Change with Machine Learning*, 2022. URL https://www.climatechange.ai/papers/neurips2022/59.