IDENTIFYING CLIMATE TARGETS IN NATIONAL LAWS AND POLICIES USING MACHINE LEARNING

Matyas Juhasz*, Tina Marchand†, Roshan Melwani, Kalyan Dutia, Sarah Goodenough, Harrison Pim, and Henry Franks

Climate Policy Radar

ABSTRACT

Quantified policy targets are a fundamental element of climate policy, typically characterised by domain-specific and technical language. Current methods for curating comprehensive views of global climate policy targets entail significant manual effort. At present there are few scalable methods for extracting climate targets from national laws or policies, which limits policymakers' and researchers' ability to (1) assess private and public sector alignment with global goals and (2) inform policy decisions. In this paper we present an approach for extracting mentions of climate targets from national laws and policies. We create an expert-annotated dataset identifying three categories of target ('Net Zero', 'Reduction' and 'Other' (e.g. renewable energy targets)) and train a classifier to reliably identify them in text. We investigate bias and equity impacts related to our model and identify specific years and country names as problematic features. We explore the dataset generated from applying our classifier to the Climate Policy Radar (CPR) dataset, showcasing the potential for automated data collection and research support in climate policy. Our work represents a significant upgrade in the accessibility of these key climate policy elements for policymakers and researchers.

1 INTRODUCTION

Climate law and policy are a primary lever for national and international climate action. Targets – quantified, measurable expressions of prospective policy outcomes – are a cornerstone of effective climate policy. Targets bolster the credibility of countries' commitments by setting quantifiable objectives, and inform policy design, implementation and monitoring (Nachmany & Mangan (2018), Andersen et al. (2021), Haarstad (2020)). Data on existing targets addressing climate change is invaluable for policy analysis, and actors engaged in the design or evaluation of national laws and policies often look to targets as a first 'port of call' to establish national and international progress and ambition in addressing climate change.

Analysis of existing laws and policies is constrained by the (often limited) resources and expertise at the disposal of relevant actors. Common barriers include: (1) key information buried within lengthy, difficult-to-parse documents, and (2) policies and targets being written in many different languages and terminologies (e.g. "reduce emissions by 50% by 2030, against a 2005 baseline" versus "double energy production from renewable sources in the next 5 years"). These constraints on time and resources affect all actors working to understand and improve the efficacy of climate laws and policies, including policymakers, academics, NGOs and UN bodies, and are especially acute for those operating under resource constraints, including in low-income countries and communities.

In this paper we present an approach for extracting mentions of targets from national climate laws and policies, using paragraph-level classification. We publish our model `https://huggingface.co/ClimatePolicyRadar/national-climate-targets` and related dataset `https://huggingface.co/datasets/ClimatePolicyRadar/national-climate-targets`.

*dsci@climatepolicyradar.org

†Was affiliated with Climate Policy Radar when this work was carried out.

2 Related work

Existing databases, including Net Zero Tracker (Lang et al. (2023)), Climate Change Laws of the World (CCLW, Grantham Research Institute at the London School of Economics; Climate Policy Radar (2023)), ClimateWatch (World Resources Institute (2022)), and Climate Action Tracker (Climate Action Tracker (2021)) rely on manual extraction of targets (in some cases relying on volunteers), limiting their scalability and consuming significant organisational resources. As a result, many only collect targets from NDCs and focus only on economy-wide targets (and not on other types of targets like sectoral targets). This results in a fundamental gap in the analysis of global public and private sector commitments.

NLP has been applied to extract climate-related targets before (Schimanski et al. (2023)), but without addressing individual GHGs or economic sectors. Our work extends existing contributions by (1) extending the definitions of emissions reduction and net-zero targets to include those also addressing individual GHGs and sectors of the economy, as these are an important component of a country's ambition and ability to achieve its economy-wide emissions reduction targets (IPCC (2023)); (2) introducing a new 'Other' category of targets to capture quantified targets made by governments with mitigation or adaptation objectives that do not explicitly mention emissions reduction, such as reducing deforestation or scaling up renewable energy capacity; and (3) curating a multi-label dataset, enabling each instance to be associated with zero, one, or multiple designated categories.

3 Data

The data was sourced from the Climate Policy Radar (CPR) database Climate Policy Radar (2024) of national laws and policies and UNFCCC submissions containing over 4,000 documents published by every single national government. We assign the target types "Net Zero", "Reduction" or "Other" to paragraphs in a multi-label classification setting. A target satisfies three criteria: it (1) contains an aim to achieve a specific outcome, (2) is quantifiable, and (3) has been given a deadline. We consider targets set by governments focusing on their specific national objectives and actions, rather than regional or global goals. Reduction targets refer to a reduction of GHG emissions, can be economy-wide or sector-specific, and refer to different types of GHGs. Net Zero targets constitute a commitment to balance GHG emissions with removal, effectively reducing the net emissions to zero. "Other" targets are those that do not fit into the Reduction or Net Zero category, yet satisfy the three criteria (e.g. renewable energy targets). See Appendix A.2 for detailed definitions.

Our approach used CCLW's target summaries (Grantham Research Institute at the London School of Economics; Climate Policy Radar (2023)) as seeds to locate paragraphs for annotation, employing stratified and negative sampling to address sampling challenges and label imbalances. Three expert annotators labeled 2,610 paragraphs with 1,193 targets (Table 1), with a review process and inter-annotator agreement checks to ensure consistency. Active learning was also used for further sampling and annotation.

Net Zero	Reduction	Other	No annotation	Total paragraphs
203	359	631	1584	2610

Table 1: Counts of labels in the final dataset. 'No annotation' counts paragraphs where no target was found.

4 Classifier Training

To accommodate a significant overlap across labels, we used a multi-label text-classification approach. We ran a number of experiments to select the most appropriate base model for our text classification task (Appendix A.3). Selecting climateBERT (Webersinke et al. (2022)), we ran a grid-search to identify the optimal hyperparameters for fine-tuning, set out in Appendix A.4.

Our model effectively predicts all 3 labels with an overall f1 score of 0.849 (Table 2). The lower performance of the NZT label is due to the low prevalence of such targets in climate text, entailing

	NZT	Other	Reduction	all
f1	0.8373 (0.0235)	0.8424 (0.0083)	0.867 (0.036)	0.8488 (0.0124)
precision	0.7767 (0.0428)	0.801 (0.0222)	0.8274 (0.0628)	0.803 (0.02)
recall	0.911 (0.0422)	0.8891 (0.016)	0.9139 (0.0453)	0.9003 (0.003)

Table 2: Classifier performance on annotation classes: net zero, reduction and other target.

a low volume of training data. This category is particularly prone to the model learning erroneous relationships, as discussed in Section 5.

5 Impacts & Equity Considerations

Biases toward countries and round years. Despite the stratified sampling, targets are less prevalent for documents authored by countries in the global south. Models trained on the dataset subsequently attribute higher probabilities to a paragraph containing targets referencing specific country names.

Targets frequently reference years that are multiples of 5 (e.g. 2035 or 2050), and models trained on the dataset can learn these features as predictors. We hypothesise a bias towards round dates in the pretrained RoBERTa model (Appendix A.5).

The effect of machine translation. Our dataset contains English paragraphs, sourced from English documents (65.29%) and Google Cloud Translation API (Google) translated documents (34.71%) from 37 source languages. While there is a drop in Overall F1 score associated with classifying machine translated text (Appendix A.6), this is small (0.023). It would be valuable to investigate whether balancing translated text and original language in the training data could address the observed drops in performance.

6 Model Application

To investigate potential applications of this model we analyzed the text from CPR laws, policies, and UNFCCC submissions, creating a dataset with 24,583 mentions of targets by 201 nations, classified into net zero (5,223), reduction (7,019), and other (13,617) types, using a model threshold mapping to 80% precision and recall. Topic modeling on 'Other' targets (Appendix A.7) with BERTopic (Grootendorst (2022), revealed dominant topics in Renewables (31.4%), Agriculture, Forests & Fisheries (13.4%), and other sectors, This indicates that analysis of 'Other' targets could aid in analysing differences in national climate action in different systems and sectors.

7 Conclusion

In this paper we present an approach for extracting targets from climate law and policy documents, by identifying Net Zero, Reduction, and/or Other types of target, and examine bias and equity in our dataset. This facilitates scalable analysis of climate documents, aiding governments in policy development and enhancing global climate action accountability. Our approach enables rapid target identification within extensive documents, crucial for understanding and addressing climate change. Our model significantly enhances analysis of climate targets by including specific GHGs, sectors, and numerous non-GHG targets.

An important avenue of climate policy analysis that this work enables is identifying discrepancies between the ambition of targets set by national governments in their submissions to the UN Climate Change Secretariat (most commonly in their Nationally Determined Contributions (NDCs)), and of the targets in their national laws and policies. This "implementation gap" is an important (Fransen et al. (2023)) but previously time consuming and manual research challenge. Other future work includes (1) predicting targets made by other actors such as companies and cities, states and regions, (2) further NLP analysis of large, machine-produced datasets such as the one presented in this work, and (3) extracting structured representations of targets for additional analysis, such as extracting target deadlines or segmenting by specific GHG references.

REFERENCES

Inger Andersen, Naoko Ishii, Thomas Brooks, Cynthia Cummis, Gustavo Fonseca, Astrid Hillers, Nicholas Macfarlane, Nebojsa Nakicenovic, Kevin Moss, Johan Rockström, Andrew Steer, Dominic Waughray, and Caroline Zimm. Defining 'science-based targets'. 8(7), July 2021. doi: 10.1093/nsr/nwaa186. URL `https://doi.org/10.1093/nsr/nwaa186`.

Climate Action Tracker. Climate action tracker. `https://climateactiontracker.org`, 2021.

Climate Policy Radar. Climate policy radar database. Online, 2024. URL `https://app.climatepolicyradar.org`.

Taryn Fransen, Jonas Meckling, Anna Stünzi, Tobias S. Schmidt, Florian Egli, Nicolas Schmid, and Christopher Beaton. Taking stock of the implementation gap in climate policy. *Nature Climate Change*, 13(8):752–755, August 2023. ISSN 1758-6798. doi: 10.1038/s41558-023-01755-9. URL `https://doi.org/10.1038/s41558-023-01755-9`.

Google. Google cloud translate api. `https://cloud.google.com/translate`. Accessed: January 24, 2024.

Grantham Research Institute at the London School of Economics; Climate Policy Radar. Climate change laws of the world. Online, 2023. URL `https://climate-laws.org`. Sourced primarily from the Grantham Research Institute at the London School of Economics. Made available under the Creative Commons CC-BY licence.

Maarten Grootendorst. Keybert: Minimal keyword extraction with bert., 2020. URL `https://doi.org/10.5281/zenodo.4461265`.

Maarten Grootendorst. Bertopic: Neural topic modeling with a class-based tf-idf procedure. *arXiv preprint arXiv:2203.05794*, 2022.

H. Haarstad. *Do Climate Targets Matter? The Accountability of Target-setting in Urban Climate and Energy Policy*. Palgrave Pivot, Cham, 2020. doi: 10.1007/978-3-030-26891-6_6. URL `https://doi.org/10.1007/978-3-030-26891-6_6`.

IPCC. Summary for policymakers. pp. 1–34, 2023. doi: 10.59327/IPCC/AR6-9789291691647.001.

John Lang, Camilla Hyslop, Natasha Lutz, Richard Black, Peter Chalkley, Thomas Hale, Frederic Hans, Nick Hay, Niklas Höhne, Angel Hsu, Takeshi Kuramochi, Silke Mooldijk, and Steve Smith. Net zero tracker, 2023.

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized BERT pretraining approach. *CoRR*, abs/1907.11692, 2019. URL `http://arxiv.org/abs/1907.11692`.

Ines Montani, Matthew Honnibal, Adriane Boyd, Sofie Van Landeghem, and Henning Peters. spaCy, October 2023. URL `https://doi.org/10.5281/zenodo.10009823`.

Michal Nachmany and Emily Mangan. Aligning national and international climate targets, 2018. Available at `www.lse.ac.uk/GranthamInstitute/publications`.

Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *ArXiv*, abs/1910.01108, 2019.

Tobias Schimanski, Julia Bingler, Mathias Kraus, Camilla Hyslop, and Markus Leippold. Climatebert-netzero: Detecting and assessing net zero and reduction targets. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pp. 15745–15756, 2023.

Nicolas Webersinke, Mathias Kraus, Julia Bingler, and Markus Leippold. ClimateBERT: A Pretrained Language Model for Climate-Related Text. 2022. doi: https://doi.org/10.48550/arXiv.2212.13631.

World Resources Institute. Climate watch. `https://www.climatewatchdata.org`, 2022.

A Appendix

A.1 Training Data Collection

A.2 Methodology Used for Labelling

The definitions used for labelling are based on existing work by Net Zero Tracker (Lang et al. (2023)) and ClimateBERT-NetZero (Schimanski et al. (2023)) to identify net zero and emissions reduction targets. We extend these definitions to include a new class – 'Other' – to capture all other quantified targets made by national governments in climate policy documents.

We also expand the net zero and emissions reduction targets definitions to include targets for different greenhouse gases (as well as general greenhouse gas targets) and to include sector-specific targets (such as emissions reduction targets for the transport sector).

A.2.1 Definition of a target

A target is

An aim to achieve something, rather than stating something concrete about the future. Often this means that the phrase indicates a level of uncertainty.
Examples

- ✓ Food waste reduced by 10% by 2022 and another 20% by 2030
- ✓ Not less than 25,000 new jobs created in 5 years
- ✓ The Government will endeavour to reach a minimum level of 10% of electrical energy supplied to the grid to be from NRE by a process of facilitation including access to green funding such as CDM.
- ✗ It is anticipated that industrial production will increase by a minimum 4.6% annually.
- ✗ Life expectancy of our people by 2040 will be 80 years due to quality care for older generation, a decent level of pension benefits and a high degree of family care.
- ✗ The Startup & CSI Development Flagship Programme is expected to create about 4,700 additional jobs in existing CSIs, Startups and new CSIs within the 12th FYP period.
- ✗ It is anticipated that industrial production will increase by a minimum 4.6% annually.

Quantifiable: it contains a reference to a measurable quantitative value. This may be numeric or non-numeric. For example, words like all, every, double, halve, eradicate, no, none and independent of refer to measurable quantities.
Examples

- ✓ reduce emissions by 68% by 2030
- ✓ provide piped water supply to all rural households by 2024
- ✗ Credit Guarantee Enhancement Corporation to be set up in 2019-2020.
- ✗ significantly decrease food waste to reduce emissions by the next decade

Given a deadline: it aims to achieve something quantifiable by a certain point in time. It can be expressed through a specific end date or some other representation of an end date in reference to planning cycles or number of years.

- ✓ reduce emissions by 68% by 2030
- ✓ in the next ten years, we will add 100km of new bicycle lanes
- ✓ increase renewable energy capacity by 20% by the end of the current national 5 year planning cycle
- ✗ increase amount of protected areas by a minimum of 4.6%
- ✗ reach energy efficiency savings of at least 2% on an annual basis

A target is not

- A policy action or a commitment to perform one (e.g. *"Publish the government's low carbon transition plan for the period 2020-2025."*).
- An abstract reference with no information about what the target is (e.g. *"Montenegro's compliance schedule will run parallel to that of EU members in the 2020-2030 decade so as to, jointly, reach the international targets established for 2030."*).
- An analysis of a target (e.g. *"It can be seen from Figure 10-3 that while the average RE cost of the MEPU 40% target is higher than the average RE cost of the MEPU 35% target, the average system cost for the 40% target is only marginally higher than the 35% target."*).
- A commitment to set up a vague target in the future (e.g. *"This assumes that the tighter EU ETS cap agreed as part of an EU deal on moving to a 30% target would continue at the same rate of reduction beyond 2020."*).
- A commitment to achieve a target based on the fulfilment of certain conditions (e.g. *"if we receive international finance, then we would be ready to achieve further GHG emissions reduction of 35% by 2040, compared to 2005 levels."*).

A.2.2 TARGET TYPES

Reduction

An emissions reduction target is a claim made by a public institution that refers to a reduction in GHG emissions by a certain point in time. It can be expressed as an absolute or relative reduction of GHG emissions, sometimes benchmarked against a baseline year or a business as usual (BAU) scenario to which the reduction target is compared. It can also be expressed as an emissions intensity reduction target where emissions act as the numerator and something else (typically population, GDP, or revenue) as the denominator. The emissions reduction target can be economy-wide or sectoral, and it can also refer to different types of GHGs (e.g. carbon dioxide, CO_2eq, methane). Must be a national target, not global.

Net zero

A net-zero target is a special type of emissions reduction target where a public institution states to bring its emissions balance down to no additional net emissions by a certain year. The net-zero target can be economy-wide or sectoral. We take particular care with mentions of net-zero technologies, they are not net-zero sectoral targets. Must be a national target, not global. To be considered a net-zero target, the emissions reduction target must contain reference to this scoped language:

- Net zero
- Carbon neutral(ity)
- GHG neutral(ity)
- Greenhouse gas neutral(ity)
- Carbon negative
- Net negative
- Carbon free
- Zero (or 0) emissions
- Zero (or 0) carbon
- Fully decarbonise
- Climate neutral
- Climate positive
- 100% emissions reduction

Other

Refers to cases where a public institution aims to achieve something concrete that is both quantifiable and has a deadline. This could include, but not limited to, non-climate mitigation (emissions reduction or net zero) targets, such as an adaptation, nature-based or renewable energy target. It could also include a policy measure, such as a quantifiable increase in carbon pricing by a certain time as a way to support the achievement of an overall emissions reduction target. Must be a national target, not global.

		ClimateBERT (82.4M)	DistilRoBERTa-base (82.8M)	RoBERTa-base (355M)
NZT	precision	0.777 (0.043)	0.758 (0.047)	0.819 (0.02)
	recall	0.911 (0.042)	0.754 (0.058)	0.799 (0.048)
	f1	0.837 (0.023)	0.754 (0.017)	0.808 (0.021)
Reduction	precision	0.827 (0.063)	0.81 (0.013)	0.843 (0.031)
	recall	0.914 (0.045)	0.9 (0.038)	0.911 (0.022)
	f1	0.867 (0.036)	0.852 (0.014)	0.876 (0.027)
Other	precision	0.801 (0.022)	0.807 (0.019)	0.824 (0.019)
	recall	0.889 (0.016)	0.864 (0.047)	0.895 (0.02)
	f1	0.842 (0.008)	0.834 (0.024)	0.858 (0.004)
all	precision	0.803 (0.02)	0.799 (0.012)	0.829 (0.013)
	recall	0.9 (0.003)	0.856 (0.033)	0.884 (0.021)
	f1	0.849 (0.012)	0.826 (0.014)	0.855 (0.012)

Table 3: Model performances per label and overall

A.3 Model performance comparisons

The models investigated were RoBERTa-base (Liu et al. (2019)), DistilRoBERTa-base (Sanh et al. (2019)) and climatebert/distilroberta-base-climate-f (Webersinke et al. (2022)), with the model performances summarised in Table 3. RoBERTa-base had the best performance, only outperforming ClimateBERT by 0.006 on the overall F1 score. Climatebert's size (more than 4x smaller model) and balanced performance (outperforming RoBERTa-base on the "Net Zero" label by 0.029) were the main factors behind our selection of it as the base model.

A.4 Hyperparameters for Model Training

	value
seed	42
optim	adamw_torch
adam_beta1	0.9
adam_beta2	0.999
model_type	roberta
adam_epsilon	0.0
warmup_steps	100
weight_decay	0.01
learning_rate	0.00002
num_train_epochs	5
lr_scheduler_type	linear
hidden_dropout_prob	0.1
per_device_eval_batch_size	24
gradient_accumulation_steps	1
per_device_train_batch_size	16
attention_probs_dropout_prob	0.1

A.5 Date bias

When predicting masked years, both distilRoBERTa and climateBERT consistently predicted round years more confidently than non-round years. Our analysis shows that "2020", "2021", "2030" and "2050" (and no others between 2020 and 2100) are single tokens in distilRoBERTa's vocabulary, potentially biasing model behaviour.

	Net Zero		Reduction		Other		Overall	
	count	**F1**	**count**	**F1**	**count**	**F1**	**count**	**F1**
Original language	153	0.856	257	0.880	401	0.843	811	0.857
English	50	0.778	102	0.839	230	0.842	382	0.834

Table 4: Classifier performance on original language vs machine-translated text. Column ‘count’ is the number of test samples for each class and in total.

A.6 Machine translation bias

A.7 Topic Modelling on ’Other’ Targets

A.7.1 Pre-Processing Step

To ensure processes heuristics were applied to extract sentence likely to contain a quantified target from each paragraph predicted as containing a target. These heuristics were whether the sentence contained the word ’target’ (not case-sensitive), or any entity expressing a date, amount or measurement (`DATE`, `CARDINAL`, `QUANTITY`, `PERCENT`) as predicted by spaCy’s (Montani et al. (2023)) `en_core_web_lg` model. Paragraphs that metadata extraction were run on were the sentences that the heuristics predicted as containing quantified targets concatenated, or the entire paragraph if no sentences in the original paragraph were predicted by the heuristics.

A.7.2 Topic modelling

BERTopic (Grootendorst (2022), parameters in Table 5) was run on pre-processed paragraphs predicted as ’Other’ to generate 60 topics. Seed topics were iteratively refined, and the final list of topics was grouped into 9 higher-level topics, with irrelevant-seeming topics discarded.

These 9 topics, in descending order of frequency, are:

- Renewables (general)
- Agriculture, forests & fisheries
- Miscellaneous
- Transport
- Electricity, infrastructure & energy efficiency
- Waste, water & plastic
- Social wellbeing (health, education and social housing)
- Wind & solar
- Built environment & construction

Table 5: BERTopic configuration

Parameter	**Value**
nr_topics	60
top_n_words	8
seed topics	”energy efficiency” ”renewable energy” ”energy sources” ”land use” ”forests”, ”forest cover” ”deforestation and reforestation”
embedding model	sentence-transformers/all-MiniLM-L6-v2
representation model	KeyBERT (Grootendorst (2020))
vectorizer model	CountVectorizer
vectorizer ngram_range	(1,3)
vectorizer min_df	5
vectorizer stop_words	”english”