

Beyond Dictionaries: Machine-Learned Normalization of Emission Units in Information Extraction

Lisa Reichenbach^{1,*}, Malte Schierholz^{2,3}, Anna Steinberg Schulten^{2,3}, Simon Oehler¹

¹Deutsche Bundesbank,

²Ludwig-Maximilians-Universität München (LMU Munich),

³Munich Center for Machine Learning (MCML)

Abstract

Sustainability reports denote greenhouse gas emissions with heterogeneous unit strings, such as abbreviations, inconsistent formatting, and scale markers. In emission extraction pipelines, these units must be normalized to a common target, e.g. t CO₂e, to enable reliable comparison and aggregation of emissions across companies, reporting years, scopes, and regions. We propose a hybrid normalization approach: If a raw unit string is not found in a curated dictionary, a machine learning classifier maps noisy unit strings to normalized emission-unit categories that can then be converted to the target unit. The classifier is trained using character-level TF-IDF, lightweight string features, and LLM-assisted data augmentation. On random-splits of the augmented dataset, an ensemble model achieves approximately 95% accuracy. Additional evaluation on rare real-world unit strings indicates that supervised fallback classifiers can handle out-of-dictionary variants in climate-related information extraction.

1 Introduction

Corporate sustainability disclosures are key sources for assessing greenhouse-gas emissions, climate-related risks, and decarbonization progress. As companies disclose growing volumes of Scope 1, Scope 2, and Scope 3 emissions under frameworks such as the GHG Protocol, TCFD, and CDP (GHG Protocol, 2011; TCFD, 2017; CDP, 2024), traditional NLP- and LLM-based systems are increasingly used to extract structured climate-related information from these reports (Bingler et al., 2022; Auzepy et al., 2023; Dave et al., 2024).

A persistent challenge in downstream analyses such as benchmarking, time-series comparisons, and total-emission calculations is that extracted emission-unit strings vary in language, capitalization, punctuation, abbreviation, and scale notation.

*Corresponding author: lisa.reichenbach@bundesbank.de

Raw unit string	Norm. class	Factor
“Tonnes CO ₂ e”	t CO ₂ e	1
“mt CO ₂ e”	t CO ₂ e	1
“kg CO ₂ equivalent”	kg CO ₂ e	0.001
“kilotonnes CO ₂ e”	kt CO ₂ e	1,000
“1,000 t CO ₂ ”	kt CO ₂ e	1,000
“million t CO ₂ e”	Mt CO ₂ e	1,000,000
“US tons CO ₂ e”	USt CO ₂ e	0.907185
“kg CO ₂ /kWh”	Other	–

Table 1: Examples of raw unit strings, normalized classes, and deterministic conversion factors to t CO₂e.

As Table 1 illustrates, superficially similar strings may imply conversion factors that differ by several orders of magnitude. Extracted quantities must therefore be normalized to a common target unit, typically tonnes of CO₂-equivalent (t CO₂e).

Emission-unit normalization can be viewed as a post-processing step after value-unit pairs have been extracted by an upstream RAG-based pipeline (Dimmelmeier et al., 2024; Zou et al., 2025). Dictionary-based normalization handles known variants precisely, but is brittle to unseen spellings, multilingual expressions, abbreviations, and extraction artifacts, reflecting known limitations of rule-based and static-dictionary normalization approaches (Sproat et al., 2001).

In our extraction setting, this brittleness is substantial: in a prior run over 1,760 corporate reports, exact dictionary lookup left 23.8% of extracted value-unit pairs unresolved. Moreover, 18.5% of all pairs were dictionary misses that could be manually assigned to specific emission-unit classes other than Other. This indicates that many unresolved strings are potentially convertible emission units and motivates a supervised fallback component for out-of-dictionary variants.

Yet unit harmonization is rarely treated as a standalone task in ClimateNLP, where prior work has focused mainly on document classification, climate-risk identification, metric extraction, and named entity recognition (Webersinke et al., 2021; Bingler et al., 2022; Dave et al., 2024; Pan et al., 2025).

We address this gap by formulating emission-unit normalization as a supervised text classification task within a hybrid dictionary-classifier pipeline. Dictionary lookup identifies known units, while a lightweight fallback classifier normalizes unknown unit strings. Using a manually curated dictionary augmented with LLM-generated variants (Ding et al., 2024; Dai et al., 2025), we evaluate the approach across complementary evaluation settings covering both augmented and rare real-world unit strings.

2 Related Work

Normalization in information extraction is commonly treated as a post-processing step that maps heterogeneous surface forms to canonical representations. Unit and quantity normalization often relies on curated inventories, rule-based matching, and deterministic conversion tables, which work well for known surface forms but are less robust to noisy, abbreviated, or domain-specific variants (Sproat et al., 2001). Early work on text and temporal normalization illustrates this pattern: CTEMP (Mingli et al., 2005), for example, extracts and normalizes Chinese temporal expressions using a grammar- and constraint-rule-based chart parser. Related value-unit extraction has been studied in biomedical text, including the extraction of laboratory test names, numerical results, units, and contextual attributes from free text (Kang and Kayaalp, 2013). More recent normalization and quantity-extraction systems combine lexical, statistical, or neural components with unit lookup and conversion rules (Lange et al., 2023; Ashmawy et al., 2023; Foppiano et al., 2019; Almasian et al., 2023). However, unit normalization is usually treated as auxiliary, with limited evaluation of fallback mechanisms for out-of-dictionary unit strings.

Climate-related information extraction has primarily addressed climate-document classification, risk identification, metric extraction, and named entity recognition in corporate disclosures (Webersinke et al., 2021; Bingler et al., 2022; Dave et al., 2024; Pan et al., 2025). Emission-unit harmonization is often handled implicitly through static dictionaries or downstream conversion rules, despite the fact that rare or noisy unit strings can lead to large aggregation errors. We address this gap by evaluating a hybrid dictionary-classifier approach for normalizing greenhouse-gas emission units in sustainability-report extraction.




Figure 1: Overview of an emission extraction pipeline and the proposed hybrid extension for unit harmonization. In the evaluated extension, unresolved raw unit strings are passed to a machine learning fallback classifier before the extracted values are converted to t CO₂e and stored in a structured emissions database.

3 Methodology

3.1 Hybrid Harmonization Pipeline

Emission extraction pipelines commonly require a unit-harmonization step after raw value-unit pairs have been extracted. In a deterministic dictionary-based approach, extracted raw unit strings are matched exactly against a curated dictionary. If no match is found, the unit remains unresolved and cannot be reliably converted to a common target unit.

We propose and evaluate a hybrid extension in which unresolved raw unit strings are passed to a supervised machine-learning fallback classifier, as shown in Figure 1. The fallback is designed to recognize unseen lexical variants of known unit classes, while unit categories outside the predefined label space still require dictionary extension or model retraining. In this work, we evaluate the fallback mechanism offline as a standalone unit-normalization component.

3.2 Task Formulation

We formulate the fallback unit-harmonization component as a supervised classification task. Given an extracted raw unit string u that is not resolved by exact dictionary lookup, such as “million metric tonnes of CO₂e”, the model predicts a normalized unit class $c \in C$ such as Mt CO₂e (see Table A.2 for complete class list). Each normalized emission-unit class, except non-convertible catch-all labels such as Other, is associated with a deterministic conversion factor f_c such that 1 unit of the normalized class = $f_c \times \text{t CO}_2\text{e}$. For example, Mt CO₂e corresponds to $f_c = 1,000,000$.

3.3 Data and Feature Representation

The dataset is based on a manually curated unit dictionary in which each instance corresponds to one distinct raw unit-string entry paired with its

normalized unit class and conversion factor. The dictionary also stores occurrence frequencies from previous applications of the emission-extraction pipeline to sustainability-report collections. These frequencies indicate how often a raw string was extracted across reports, reporting years, and scopes, but they do not create duplicate dataset instances and are used only for analysis, not as predictors. All frequency metadata were collected before the machine-learning experiments.

Given the highly imbalanced classes, as evident from the column “Manual” in Table A2, we augment underrepresented classes with GPT-4.1-generated variants (OpenAI, 2025), following prior LLM-based augmentation work under data scarcity and class imbalance (Ye et al., 2024; Ding et al., 2024; Dai et al., 2025). The augmentation is designed to expand existing semantic classes with plausible alternative surface forms, rather than to introduce new unit concepts. Because small changes in prefixes, denominators, or scale expressions can alter unit semantics, all generated candidates are deduplicated against the manual dictionary and manually validated. Lexical overlap among variants should be considered when interpreting performance estimates based on random data splits. The augmentation prompt is shown in Figure A1, and randomly selected examples are provided in Table A1 in Appendix A.1. The final dataset comprises 1,849 unit strings across eight target classes, including 471 manually validated LLM-generated variants. Table A2 reports the resulting class distribution.

The classifiers use only features derived from the raw unit string. Unit strings are lowercased, normalized, and represented using character-level TF-IDF features. Lowercasing was retained because it improved validation performance in preliminary experiments, despite discarding potentially informative case distinctions such as *mt* versus *Mt*; we revisit this trade-off in the error analysis and limitations. Additional regex-based indicators capture scale and quantity expressions such as *kilo*, *million*, *mega*, or *short tons*; simple numerical features such as string length are also included. Textual features are vectorized using character n-grams and combined with scaled numerical features in a ColumnTransformer. Target labels are label-encoded during training and decoded after prediction.

Full implementation details, including preprocessing steps, character n-gram settings, hyperpa-

rameter grids, selected model configurations, software versions, and random seeds, are provided in Appendix A.3. The code and unit dictionary are available in a public GitHub repository¹ to support reproducibility and future work.

3.4 Experimental Setup

As a trivial lower bound, we include a majority-class classifier that always predicts the most frequent normalized unit class in the training data, namely *Other*. Since exact dictionary lookup on the curated dataset is perfect by construction, we do not report it. We compare the majority-class classifier with several supervised models for the fallback task: Logistic Regression, Random Forest, Support Vector Machine (SVM), XGBoost, LightGBM, Multi-layer Perceptron (MLP), and a majority-voting ensemble classifier combining SVM, MLP, and Logistic Regression. Hyperparameters are optimized on the full dataset using grid search with stratified cross-validation.

Evaluation is conducted in two complementary settings. The **random-split dataset** evaluates model performance under a standard stratified split of the full dataset, including both manually curated and LLM-augmented unit strings. The **rare-unit dataset** is a robustness test split comprising manually curated unit strings whose exact raw surface forms occurred at most five times in the prior extraction runs, based on the frequency metadata defined in Section 3.3. The rare-unit test set was constructed from pre-existing dictionary entries and occurrence metadata, independently of model predictions and error analyses. LLM-generated variants are excluded because they were not observed in the extraction output and therefore do not have occurrence counts.

In the rare-unit experiment, the rare-unit test strings are excluded from model training, while LLM-augmented entries remain available in the training data to improve coverage of underrepresented classes. No additional hyperparameter tuning is performed. We instead use the hyperparameters selected in the preceding grid-search cross-validation experiments. Since that cross-validation used random splits of the full augmented dataset, some rare-unit strings may have contributed to hyperparameter selection. Thus, the rare-unit evaluation uses a fixed configuration and excludes the rare test strings from model fitting, but it is not

¹https://github.com/gist-sustainability/unit_harmonization/releases/tag/v0.1.0

Model	10-fold CV			Random-split run		
	Acc.	F1 _w	Macro F1	Acc.	F1 _w	Macro F1
Majority-class	0.3067 ± 0.0026	0.1439 ± 0.0022	0.0587 ± 0.0004	0.3054	0.1429	0.0585
Logistic Regression	0.9524 ± 0.0026	0.9525 ± 0.0143	0.9621 ± 0.0103	0.9595	0.9596	0.9669
Random Forest	0.9421 ± 0.0189	0.9421 ± 0.0189	0.9551 ± 0.0134	0.9351	0.9355	0.9495
SVM	0.9516 ± 0.0116	0.9414 ± 0.0116	0.9542 ± 0.0068	0.9486	0.9489	0.9601
XGBoost	0.9448 ± 0.0149	0.9447 ± 0.0150	0.9570 ± 0.0107	0.9514	0.9515	0.9618
LightGBM	0.9486 ± 0.0165	0.9486 ± 0.0164	0.9598 ± 0.0144	0.9514	0.9515	0.9625
MLP	0.9492 ± 0.0103	0.9493 ± 0.0103	0.9603 ± 0.0081	0.9649	0.9648	0.9751
Voting Ensemble	0.9551 ± 0.0114	0.9552 ± 0.0113	0.9671 ± 0.0082	0.9595	0.9595	0.9708

Table 2: Offline evaluation of fallback classifiers for unit harmonization. Cross-validation results report mean ± standard deviation across 10 stratified folds on the full augmented dictionary-derived dataset. Random-split results are computed from a single stratified 80/20 train-test split on the random-split dataset.

fully independent of the preceding tuning procedure. Its purpose is to assess robustness to naturally observed low-frequency unit variants in the fallback use case.

We report accuracy, weighted F1-score (F1_w), and macro F1-score. In this single-label multi-class setting, accuracy is equivalent to micro-F1, so micro-F1 is not reported separately. Weighted F1 is used as the main imbalance-aware F1 metric because it accounts for class support and reflects the influence of frequent classes such as Other, t C02e, and kt C02e. Because aggregate metrics do not capture the practical severity of different error types, we complement them with per-class results and row-normalized confusion matrices, focusing on whether errors involve distant unit scales, such as t C02e versus Mt C02e, or conservative predictions of Other.

4 Results

4.1 Evaluation of Fallback Classifiers

Table 2 summarizes the offline performance of the fallback classifiers. We report 10-fold stratified cross-validation on the full augmented dictionary-derived dataset using StratifiedKFold with shuffling, fitting the complete preprocessing and classification pipeline separately within each fold. In addition, we evaluate each model in a single stratified 80/20 random-split run. Hyperparameters are fixed to the selected configurations reported in Table A4.

All supervised models clearly outperform the majority-class lower bound, whose weighted F1-score remains below 0.15. The Voting Ensemble achieves the strongest cross-validation performance across all metrics, with a weighted F1-score of 0.9552 ± 0.0113 . Its comparatively low stan-

dard deviation also indicates stable performance across folds. On the random-split dataset, the ensemble remains among the best-performing models, with only MLP achieving higher performance in the single split run. Based on its consistently strongest cross-validation performance and competitive random-split results, we choose the Voting Ensemble as the selected fallback classifier.

The selected ensemble combines SVM, MLP, and Logistic Regression. It was chosen as a compact performance-runtime trade-off, since larger ensembles yielded only marginally higher validation scores while increasing runtime and complexity.

4.2 Per-Class Performance and Error Analysis

To interpret the aggregate results, we examine the per-class performance of the selected fallback classifier, the Voting Ensemble, on the test portion of the random-split dataset. Figure 2 shows the row-normalized confusion matrix for the emission unit categories with a strong diagonal, indicating high recall across classes.

Most remaining errors involve related unit categories or predictions of Other. This distinction is important because the practical severity of misclassifications is not uniform: predicting Other for a convertible unit mainly leads to a missed automatic conversion, whereas confusing distant scale classes, such as t C02e and Mt C02e, would directly induce errors of several orders of magnitude. The confusion matrix shows that such severe scale-related confusions are rare in the random-split dataset. The per-class scores in Appendix A.4 confirm this pattern, with all F1-scores above 0.92. This suggests that the fallback classifier is generally reliable for distinguishing emission unit categories.

The remaining errors in the random-split dataset




Figure 2: Confusion matrix of emission unit categories for the random-split dataset. Each row is normalized to 100%, showing the distribution of predicted unit categories for each true unit category. The diagonal entries correspond to the recall of the respective categories.

Raw string	True	Pred.	Error type
mt CO ₂ e	t CO ₂ e	Mt CO ₂ e	case-sensitive scale ambiguity
metric tons CO ₂ equivalent (based on 2016 data)	t CO ₂ e	Other	unit embedded in context
emissions	Other	t CO ₂ e	no mass unit
kgCO ₂ e/sqm*year	Other	kg CO ₂ e	intensity instead of absolute unit
th. t CO ₂ Eq.	kt CO ₂ e	t CO ₂ e	abbreviated scale marker

Table 3: Representative misclassifications of the selected fallback classifier on the random-split dataset.

are concentrated in linguistically ambiguous or under-specified cases rather than in systematic confusions between distant unit scales. For example, capitalization-sensitive abbreviations can lead to errors: lowercase mt denotes metric tonnes and corresponds to t CO₂e, whereas uppercase Mt denotes megatonnes and corresponds to Mt CO₂e. As shown in Table 3, other errors involve abbreviated scale markers such as th., units embedded in context, missing mass units, or boundary cases such as emission intensities.

4.3 Robustness on the Rare-Unit Dataset

As described in Section 3.4, the rare-unit test set consists of distinct manually curated real-world unit-string entries whose raw surface forms occurred at most five times in prior extraction runs. The resulting split is more imbalanced than the random-split dataset; the full train-test composition

is reported in Appendix A.5, Table A6.

In this setting, the Voting Ensemble achieves an accuracy of 94.85%, a weighted F1-score of 94.84%, and a macro F1-score of 96.93%. The complete results are shown in Appendix Table A7. The weighted F1-score differs from the corresponding random-split result by only 1.11 percentage points, indicating strong performance on rare, naturally occurring unit variants rather than performance driven solely by frequent or LLM-augmented strings. As noted in Section 3.4, however, this evaluation is not fully independent of the preceding hyperparameter selection.

Detailed model comparisons, per-class scores, and the corresponding confusion matrix are provided in Appendix A.5. The rare-unit confusion matrix shows an even more conservative error pattern than the random-split evaluation: almost all remaining misclassifications are predictions of Other, while direct confusions between distant conversion-relevant unit scales are nearly absent.

5 Conclusion

This paper studies emission-unit normalization as a post-processing step for greenhouse gas emission extraction. We propose a hybrid approach that keeps deterministic dictionary lookup as a high-confidence first step and adds a supervised fallback classifier for out-of-dictionary unit strings. Using a curated and LLM-augmented dataset, we train different lightweight classifiers on character-level TF-IDF representations and simple domain-specific string features.

The results show that supervised classifiers can effectively extend dictionary coverage to unseen variants of known emission-unit classes, including rare real-world strings. The best-performing models achieve around 95% accuracy and weighted F1-score on both the random-split dataset and the rare-unit dataset. This suggests that the approach is not limited to frequent or augmented examples and can also handle rare real-world unit strings. Error analysis further suggests that most remaining mistakes stem from genuinely ambiguous abbreviations, missing unit information, or boundary cases such as emission intensities, rather than systematic large-scale conversion errors. The hybrid setup therefore provides a transparent and practical path toward more robust unit harmonization in climate-related information extraction.

6 Limitations

This study evaluates unit normalization as a standalone offline component rather than as part of an end-to-end emission extraction pipeline. Consequently, the evaluation does not capture interactions with upstream extraction errors, document context, table structure, or downstream aggregation. In practice, such errors may propagate, and scale confusions between units such as t CO₂e, kt CO₂e, and Mt CO₂e can lead to conversion errors of several orders of magnitude.

A further limitation is the use of augmented dictionary-derived data. Although all LLM-generated variants were manually validated, random splits may overestimate generalization because highly similar surface forms can appear in both training and test data. The rare-unit dataset mitigates this by excluding LLM-augmented entries from the test set and evaluating only on rare manually curated real-world strings. However, it is not fully independent of model selection, as some rare-unit strings may have appeared in training or validation folds during the preceding random-split grid-search cross-validation. It may therefore be less conservative than a fully untouched test set. Moreover, the rare strings come from previous extraction runs and are not a prospective evaluation on newly extracted units from unseen reports. We do not include a mixed rare-plus-LLM holdout in order to keep the rare-unit test real-world-only, although such a synthetic holdout could be useful for sensitivity analysis on underrepresented classes.

Finally, the classifier is restricted to predefined unit classes and cannot handle genuinely new unit types without dictionary extension or retraining. Ambiguous abbreviations, multilingual expressions, and intensity-based units remain challenging. In particular, lowercasing improves robustness to inconsistent capitalization but removes informative distinctions such as mt versus Mt. Future work should integrate the classifier into the full extraction pipeline, incorporate document or table context, add case-sensitive or confidence-based rules for high-impact abbreviations, and use severity-aware metrics that weight errors by their implied conversion-factor difference.

Use of AI Assistants

AI-based writing assistants, including GPT-based tools, were used to support language refinement, wording suggestions, and minor editorial improve-

ments during manuscript preparation. All AI-generated suggestions were manually reviewed, revised where necessary, and verified by the authors. The authors remain fully responsible for the content, analysis, results, and conclusions presented in this paper.

References

- Satya Almasian, Vivian Kazakova, Philipp Göldner, and Michael Gertz. 2023. *CQE: A Comprehensive Quantity Extractor*. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12845–12859, Singapore. Association for Computational Linguistics.
- Mohamed Ashmawy, Mohamed Fakhr, and Fahima Maghraby. 2023. *Lexical Normalization Using Generative Transformer Model (LN-GTM)*. *International Journal of Computational Intelligence Systems*, 16.
- Alix Auzepy, Elena Tönjes, David Lenz, and Christoph Funk. 2023. *Evaluating TCFD reporting—A new application of zero-shot analysis to climate-related financial disclosures*. *PLOS ONE*, 18(11):1–23.
- Julia Anna Bingler, Mathias Kraus, Markus Leippold, and Nicolas Webersinke. 2022. *Cheap talk and cherry-picking: What ClimateBert has to say on corporate climate risk disclosures*. *Finance Research Letters*, 47:102776.
- CDP. 2024. *CDP Corporate Questionnaire 2024*. <https://www.cdp.net/en/disclose/question-bank>. Accessed: 2026-07-02.
- Haixing Dai, Zhengliang Liu, Wenxiong Liao, Xiaoke Huang, Yihan Cao, Zihao Wu, Lin Zhao, Shaochen Xu, Fang Zeng, Wei Liu, Ninghao Liu, Sheng Li, Dajiang Zhu, Hongmin Cai, Lichao Sun, Quanzheng Li, Dinggang Shen, Tianming Liu, and Xiang Li. 2025. *AugGPT: Leveraging ChatGPT for Text Data Augmentation*. *IEEE Transactions on Big Data*, 11(3):907–918.
- Aditya Dave, Mengchen Zhu, Dapeng Hu, and Sachin Tiwari. 2024. *Climate AI for Corporate Decarbonization Metrics Extraction*. Papers 2411.03402, arXiv.org.
- Andreas Dimmelmeier, Hendrik Doll, Malte Schierholz, Emily Kormanyos, Maurice Fehr, Bolei Ma, Jacob Beck, Alexander Fraser, and Frauke Kreuter. 2024. *Informing climate risk analysis using textual information - A research agenda*. In *Proceedings of the 1st Workshop on Natural Language Processing Meets Climate Change (ClimateNLP 2024)*, pages 12–26, Bangkok, Thailand. Association for Computational Linguistics.
- Bosheng Ding, Chengwei Qin, Ruochen Zhao, Tianze Luo, Xinze Li, Guizhen Chen, Wenhan Xia, Junjie

- Hu, Anh Tuan Luu, and Shafiq Joty. 2024. [Data Augmentation using LLMs: Data Perspectives, Learning Paradigms and Challenges](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 1679–1705, Bangkok, Thailand. Association for Computational Linguistics.
- Luca Foppiano, Laurent Romary, Masashi Ishii, and Mikiko Tanifuji. 2019. [Automatic Identification and Normalisation of Physical Measurements in Scientific Literature](#). In *Proceedings of the ACM Symposium on Document Engineering 2019, DocEng '19*, New York, NY, USA. Association for Computing Machinery.
- GHG Protocol. 2011. [The Greenhouse Gas Protocol: A Corporate Accounting and Reporting Standard: Revised Edition](#). Technical report, World Resources Institute and World Business Council for Sustainable Development. Accessed: 2026-07-01.
- Yanna Shen Kang and Mehmet Kayaalp. 2013. [Extracting laboratory test information from biomedical text](#). *Journal of Pathology Informatics*, 4(1):23.
- Lukas Lange, Jannik Strötgen, Heike Adel, and Dietrich Klakow. 2023. [Multilingual Normalization of Temporal Expressions with Masked Language Models](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 1174–1186, Dubrovnik, Croatia. Association for Computational Linguistics.
- Wu Mingli, Li Wenjie, Lu Qin, and Li Baoli. 2005. [Ctemp: A chinese temporal parser for extracting and normalizing temporal information](#). In *Natural Language Processing – IJCNLP 2005*, pages 694–706, Berlin, Heidelberg. Springer Berlin Heidelberg.
- OpenAI. 2025. [Introducing GPT-4.1 Model Family](#). Accessed: 2026-07-01.
- Huitong Pan, Mustapha Adamu, Qi Zhang, Eduard Dragut, and Longin Jan Latecki. 2025. [ClimateIE: A Dataset for Climate Science Information Extraction](#). In *Proceedings of the 2nd Workshop on Natural Language Processing Meets Climate Change (ClimateNLP 2025)*, pages 76–98, Vienna, Austria. Association for Computational Linguistics.
- Richard Sproat, Alan Black, Stanley Chen, Shankar Kumar, Mari Ostendorf, and Christopher Richards. 2001. [Normalization of Non-Standard Words](#). *Computer Speech & Language*, 15:287–333.
- TCFD. 2017. [Recommendations of the Task Force on Climate-related Financial Disclosures](#). Technical report, Task Force on Climate-related Financial Disclosures. Accessed: 2026-07-02.
- Nicolas Webersinke, Mathias Kraus, Julia Anna Binger, and Markus Leippold. 2021. [Climatebert: A pretrained language model for climate-related text](#). *CoRR*, abs/2110.12010.
- Junjie Ye, Nuo Xu, Yikun Wang, Jie Zhou, Qi Zhang, Tao Gui, and Xuanjing Huang. 2024. [LLM-DA: Data Augmentation via Large Language Models for Few-Shot Named Entity Recognition](#). *ArXiv*, abs/2402.14568.
- Yi Zou, Mengying Shi, Zhongjie Chen, Zhu Deng, Zongxiong Lei, Zihan Zeng, Shiming Yang, Hongxiang Tong, Lei Xiao, and Wenwen Zhou. 2025. [ES-GReveal: An LLM-based approach for extracting structured data from ESG reports](#). *Journal of Cleaner Production*, 489:144572.

A Appendix

A.1 Prompt Template and Examples for LLM-assisted Data Augmentation

To mitigate class imbalance in the unit normalization data, we use an LLM to generate additional variants for underrepresented unit categories. For each target category c , the model is prompted with examples from the well-represented t_CO2e class and instructed to produce realistic unit strings with analogous surface-form variation patterns, including differences in spelling, formatting, abbreviations, whitespace, special characters, and capitalization. The prompt also restricts the output to variants without additional parenthetical descriptions and requests a Python list for straightforward downstream processing.

The prompt template used for the generation is shown in Figure A1. Here, c denotes the target unit category, while $\{tco2e_variants_list\}$ denotes the example list from the well-represented t_CO2e class.

Augmentation prompt

Please generate a list of 100 unique variants for the unit category c , following the style and diversity of the example list below for “ t_CO2e ”. The variants should differ in spelling, formatting, whitespaces, abbreviations, use of special characters and/or capitalization. Do not include any variants with additional descriptions in parentheses, e.g., “(per shipment)”, “(biogenic)”, etc. Please avoid duplicates and ensure that all variants are realistic and could plausibly appear in real-world data. The output should be a Python list. Here is my example list:
 $\{tco2e_variants_list\}$

Figure A1: Prompt template used for LLM-assisted data augmentation. The placeholder c is replaced by the target unit category, while $\{tco2e_variants_list\}$ is replaced by the example variants from the well-represented t_CO2e class.

To illustrate the realized diversity of the augmentation, we manually inspected the LLM-generated variants from the final augmented dictionary. Table A1 presents randomly selected examples from different target classes and summarizes the main types of observed surface-form variation, including capitalization, hyphenation, abbreviations, lexical expansions, and Unicode characters. Although some variants remain close to existing dictionary entries, the examples demonstrate that the augmentation produces a range of plausible surface-form realizations within predefined semantic classes.

Target class	Generated variant	Observed variation
MM Mt CO2e	mm mt CO2-Equivalent	case, hyphen, expansion
kg CO2e	kg CO2-Equiv	hyphen, abbreviation
MM Mt CO2e	mm mt CO ₂ e	case, Unicode subscript
MM Mt CO2e	MMMt CO ₂ -equivalent	abbreviation, Unicode, hyphen
10 kt CO2e	10 KT CO ₂ equivalent	case, expansion
USt CO2e	us ton CO ₂ eq	case, singular, abbreviation
kg CO2e	Kg CO2-Eq	case, hyphen, abbreviation
10 kt CO2e	10 kt CO2-Äq	hyphen, German abbreviation

Table A1: Randomly selected examples of manually validated LLM-generated unit-string variants. The final column summarizes the main types of observed surface-form variation.

A.2 Class Distribution by Data Source

Table A2 reports the class distribution of the final unit-string dataset across all eight target classes, distinguishing between manually curated entries and validated LLM-generated variants. The augmentation was applied selectively to classes with limited manual coverage, while already well-represented classes were kept unchanged. The category Other contains strings that cannot be unambiguously assigned to a target emission-unit class or refer to incompatible quantities, such as “emissions”, “kWh”, or “t CO₂e/million euros”.

A.3 Implementation and Hyperparameter Details

All models were implemented in Python version 3.11.9 using scikit-learn version 1.9.1, with XGBoost version 3.2.0 and LightGBM version 4.7.0. Unless stated otherwise, we used the default parameter settings of the respective libraries. Random

Unit class	Manual (Proportion)	Augmented	Total
t CO ₂ e	343 (24.9%)	0	343
kt CO ₂ e	249 (18.1%)	0	249
Mt CO ₂ e	184 (13.4%)	0	184
kg CO ₂ e	25 (1.8%)	101	126
USt CO ₂ e	4 (0.3%)	112	116
10 kt CO ₂ e	4 (0.3%)	116	120
MM Mt CO ₂ e	2 (0.1%)	142	144
Other	567 (41.1%)	0	567
Total	1378	471	1849

Table A2: Class distribution of unit strings by data source. Manual entries originate from the curated unit dictionary, whereas augmented entries are validated LLM-generated variants.

seeds were fixed to 42 for data splitting and model training where applicable.

Hyperparameters were selected on the training portion using grid search with 5-fold stratified cross-validation and weighted F1-score as the refit criterion. The selected pipeline was then refitted on the full training portion and evaluated on the held-out test set. Raw unit strings were represented by a concatenated text feature and one numerical length feature. The text feature combined: (i) a lowercased raw string in which numeric substrings were replaced by NUMBER; (ii) a regex-based quantity keyword covering markers such as *kilo*, *thousand*, *kt*, *million*, *mio.*, *mega*, *kg* and *US/short tons*; and (iii) a cleaned unit-base string obtained by removing numeric expressions and non-word characters, removing the token *of*, stripping whitespace, and lowercasing. We also tested case-preserving variants of components (i) and (iii), as well as variants preserving case in only one of the two components. These alternatives performed worse in preliminary validation experiments and were therefore not used in the final configuration.

The concatenated text field was vectorized using character-level TF-IDF with analyzer=char_wb, ngram_range=(1,3), min_df=3, max_df=0.95, max_features=1000, and sublinear_tf=True; all other vectorizer parameters were left at their scikit-learn defaults, including L2 normalization. The numerical length feature was standardized using StandardScaler.

The complete hyperparameter search spaces are summarized in Table A3. For the Voting Ensemble, we first evaluated all base-model subsets of size three to six and then selected the final subset based on the trade-off between predictive performance, runtime, and model complexity. Table A4 reports the final model configurations used for the classifier comparison in the main experiments.

Model	Hyperparameter grid
Logistic Regression	Two grids were evaluated. First, $C \in \{0.01, 0.1, 1.0, 10.0, 100.0, 1000.0\}$, $\text{penalty} \in \{l2\}$, $\text{class_weight} \in \{\text{None}, \text{balanced}\}$, $\text{solver} \in \{\text{lbfgs}, \text{newton-cg}\}$, $\text{max_iter} \in \{1000, 2000, 3000\}$. Second, $C \in \{0.01, 0.1, 1.0, 10.0, 100.0\}$, $\text{penalty} \in \{l1, l2\}$, $\text{class_weight} \in \{\text{None}, \text{balanced}\}$, $\text{solver} \in \{\text{saga}\}$, $\text{max_iter} \in \{2000\}$
SVM	$C \in \{0.1, 1, 10, 100\}$, $\text{kernel} \in \{\text{linear}, \text{rbf}\}$, $\text{gamma} \in \{\text{scale}, \text{auto}\}$, $\text{class_weight} \in \{\text{None}, \text{balanced}\}$
Random Forest	$n_estimators \in \{100, 200, 500\}$, $\text{max_depth} \in \{\text{None}, 10, 20\}$, $\text{min_samples_split} \in \{2, 10\}$, $\text{min_samples_leaf} \in \{1, 4\}$, $\text{max_features} \in \{\text{sqrt}, \text{log2}\}$, $\text{bootstrap} \in \{\text{True}, \text{False}\}$, $\text{class_weight} \in \{\text{None}, \text{balanced}\}$
XGBoost	$\text{max_depth} \in \{4, 6\}$, $\text{learning_rate} \in \{0.2\}$, $n_estimators \in \{100, 300\}$, $\text{subsample} \in \{1\}$, $\text{colsample_bytree} \in \{0.8\}$, $\text{gamma} \in \{\text{None}, 0.3\}$, $\text{reg_alpha} \in \{\text{None}, 0.1\}$, $\text{reg_lambda} \in \{\text{None}, 1.5\}$
LightGBM	$\text{num_leaves} \in \{31, 63\}$, $\text{max_depth} \in \{-1, 6, 8\}$, $n_estimators \in \{100, 200, 600\}$, $\text{learning_rate} \in \{0.05, 0.1\}$, $\text{feature_fraction} \in \{0.8, 1.0\}$, $\text{bagging_freq} \in \{0, 1\}$, $\text{reg_lambda} \in \{0.0, 2.0\}$
MLP	$\text{hidden_layer_sizes} \in \{(100,), (200,), (100, 100), (200, 100)\}$, $\text{activation} \in \{\text{relu}, \text{tanh}\}$, $\alpha \in \{10^{-5}, 10^{-4}, 10^{-3}\}$, $\text{learning_rate_init} \in \{0.001, 0.01\}$, $\text{solver} \in \{\text{adam}, \text{lbfgs}\}$, $\text{max_iter} \in \{500, 1000\}$, $\text{early_stopping} \in \{\text{True}, \text{False}\}$
Voting Ensemble	All base-model subsets of size 3 to 6 were evaluated from {XGBoost, LightGBM, Logistic Regression, Random Forest, SVM, MLP}, yielding 42 model combinations. The final subset, consisting of SVM, MLP, and Logistic Regression, was selected based on the trade-off between performance, runtime, and model complexity, as larger ensembles provided only marginal performance gains. For this subset, voting $\in \{\text{hard}, \text{soft}\}$ was tested; $\text{weights}=\text{None}$.

Table A3: Hyperparameter grids used for grid search cross validation.

Model	Selected hyperparameters
Logistic Regression	$C=100$, $\text{penalty}=l2$, $\text{class_weight}=\text{balanced}$, $\text{solver}=\text{lbfgs}$, $\text{max_iter}=1000$
SVM	$C=100$, $\text{kernel}=\text{linear}$, $\text{gamma}=\text{scale}$, $\text{class_weight}=\text{None}$, $\text{probability}=\text{True}$, $\text{random_state}=42$
Random Forest	$n_estimators=200$, $\text{max_depth}=\text{None}$, $\text{bootstrap}=\text{False}$, $\text{class_weight}=\text{None}$, $\text{max_features}=\text{sqrt}$, $\text{min_samples_leaf}=1$, $\text{min_samples_split}=2$, $\text{random_state}=42$, $n_jobs=-1$
XGBoost	$n_estimators=100$, $\text{max_depth}=6$, $\text{learning_rate}=0.2$, $\text{subsample}=1.0$, $\text{colsample_bytree}=0.8$, $\text{gamma}=\text{None}$, $\text{reg_alpha}=\text{None}$, $\text{reg_lambda}=1.5$, $\text{random_state}=42$, $\text{eval_metric}=\text{mlogloss}$
LightGBM	$n_estimators=600$, $\text{max_depth}=6$, $\text{learning_rate}=0.1$, $\text{bagging_freq}=0$, $\text{feature_fraction}=1.0$, $\text{num_leaves}=31$, $\text{reg_lambda}=0.0$, $\text{random_state}=42$, $\text{objective}=\text{multiclass}$, $\text{metric}=\text{multi_logloss}$, $\text{verbosity}=-1$
MLP	$\text{hidden_layer_sizes}=(200, 100)$, $\text{activation}=\text{tanh}$, $\text{solver}=\text{lbfgs}$, $\alpha=0.001$, $\text{learning_rate_init}=0.001$, $\text{max_iter}=500$, $\text{random_state}=42$, $\text{early_stopping}=\text{True}$
Voting Ensemble	$\text{voting}=\text{soft}$, $\text{weights}=\text{None}$; base models: SVM, MLP, Logistic Regression

Table A4: Selected hyperparameters used for the classifier comparison.

A.4 Per-Class Classification Results

Table A5 reports the per-class precision, recall, and F1-scores of the selected fallback classifier on the test portion of the random-split dataset. The results complement the aggregate metrics and the confusion matrix in Section 4.2 by showing that performance is high across the eight unit categories. Most classes achieve F1-scores close to or above 0.95, with perfect scores for USt C02e and MM Mt C02e. Class t C02e exhibits the lowest F1-score, driven by moderately lower precision and recall compared to the other classes. kt C02e shows the lowest recall. The Other class remains robust, suggesting that incompatible or non-convertible units

are usually not forced into a specific emission-unit category.

Class	Precision	Recall	F1	Support
t C02e	0.9275	0.9275	0.9275	69
kt C02e	0.9792	0.9400	0.9592	50
Mt C02e	0.9024	1.0000	0.9487	37
kg C02e	0.9615	1.0000	0.9804	25
USt C02e	1.0000	1.0000	1.0000	23
10 kt C02e	1.0000	1.0000	1.0000	24
MM Mt C02e	1.0000	1.0000	1.0000	29
Other	0.9636	0.9381	0.9507	113

Table A5: Per-class performance of the selected classifier on the test portion of the random-split dataset. Support denotes the number of test instances per class.

A.5 Detailed Results on the Rare-Unit Dataset

The rare-unit evaluation uses a fixed train-test split designed to assess robustness on naturally occurring low-frequency unit strings. The test set consists only of distinct manually curated real-world unit-string entries whose corresponding raw surface forms occurred at most five times in prior extraction runs. These rare-test strings were excluded from training. The training set contains the remaining manually curated entries and all validated LLM-generated variants. LLM-generated variants are not included in the rare-unit test set because this evaluation is intended to focus on naturally observed rare strings. Table A6 summarizes the resulting train-test composition by class.

Class	Manual total	Rare test	Manual train	Aug. train	Total train
t CO2e	343	68	275	0	275
kt CO2e	249	49	200	0	200
Mt CO2e	184	36	148	0	148
kg CO2e	25	19	6	101	107
USt CO2e	4	4	0	112	112
10 kt CO2e	4	1	3	116	119
MM Mt CO2e	2	1	1	142	143
Other	567	113	454	0	454
Total	1378	291	1087	471	1558

Table A6: Train-test composition of the rare-unit evaluation setting. The rare test set contains only manually curated real-world unit strings whose exact raw surface forms occurred at most five times in prior extraction runs. LLM-generated variants are excluded from the test set and used only for training.

Table A7 reports the performance of all candidate fallback classifiers on this rare-unit dataset. The split composition in Table A6 shows that the rare-unit test set is substantially more imbalanced than the random-split dataset, with very small support for USt CO2e, 10 kt CO2e, and MM Mt CO2e. The aggregate metrics should therefore be interpreted together with the per-class results below. The Voting Ensemble achieves the strongest results on this dataset and shows the smallest absolute difference in weighted F1-score relative to the random-split dataset. This is particularly relevant because the MLP achieved the highest score on the random-split dataset, but exhibits a larger drop on the rare-unit dataset. The results therefore support selecting the Voting Ensemble as the more robust fallback classifier.

Table A8 provides the corresponding per-class results for the Voting Ensemble on the rare-unit dataset. Performance remains high across the main classes, with F1-scores above 0.91 for t CO2e, kt CO2e, Mt CO2e, kg CO2e, and Other. The class Other also remains reliable, with a recall of 0.9735, suggesting that rare non-target or incom-

Model	Acc.	F1 _w	Macro F1	$\Delta F1_w$
Log. Regression	0.9416	0.9413	0.9620	0.0183
Random Forest	0.9141	0.9125	0.8228	0.0230
SVM	0.9175	0.9174	0.8260	0.0315
XGBoost	0.8900	0.8892	0.7701	0.0623
LightGBM	0.9038	0.9035	0.8600	0.0480
MLP	0.9347	0.9345	0.9593	0.0303
Voting Ensemble	0.9485	0.9484	0.9693	0.0111

Table A7: Robustness evaluation of the fallback classifiers on the rare-unit dataset. $\Delta F1_w$ denotes the absolute difference in weighted F1-score between the rare-unit dataset and the test portion of the random-split dataset.

patible unit strings are usually not forced into one of the emission-unit classes. For classes with very small rare-test support, especially USt CO2e, 10 kt CO2e, and MM Mt CO2e, the perfect scores should not be overinterpreted. Overall, the per-class results indicate that the classifier generalizes well to rare real-world unit variants and that the strong aggregate performance is not solely explained by frequent classes or by evaluation on LLM-augmented strings.

Class	Precision	Recall	F1	Support
t CO2e	0.9254	0.9118	0.9185	68
kt CO2e	0.9400	0.9592	0.9495	49
Mt CO2e	1.0000	0.9167	0.9565	36
kg CO2e	1.0000	0.9474	0.9730	19
USt CO2e	1.0000	1.0000	1.0000	4
10 kt CO2e	1.0000	1.0000	1.0000	1
MM Mt CO2e	1.0000	1.0000	1.0000	1
Other	0.9402	0.9735	0.9565	113

Table A8: Per-class performance of the selected fallback classifier on the rare-unit dataset. Support denotes the number of test instances per class.

Figure A2 shows the row-normalized confusion matrix for the rare-unit dataset, analogously to Figure 2. Compared to the random-split dataset, the error pattern is more conservative: a lot of the remaining misclassifications are predictions of Other, rather than confusions between alternative specific unit scales. This is relevant because errors are not equally severe in downstream emission aggregation. A false Other prediction prevents automatic conversion and may require manual review, whereas confusing units with very different conversion factors, such as kg CO2e and Mt CO2e, would directly distort normalized emissions by several orders of magnitude. For example, the errors for kg CO2e are entirely assigned to Other. Direct confusions between specific emission units are rare and occur only in small proportions. This suggests that the classifier tends to avoid assigning a potentially

wrong conversion-relevant unit category when uncertain, reducing the risk of severe scale-related conversion errors.




Figure A2: Row-normalized confusion matrix of predicted emission-unit categories on the rare-unit dataset. Values are percentages per true category; diagonal entries correspond to category-wise recall.