

Using Text Classification with a Bayesian Correction for Estimating Overreporting in the Creditor Reporting System on Climate Adaptation Finance

Janos Borst¹, Thomas Wencker² and Andreas Niekler¹

¹Leipzig University

²German Institute for Development Evaluation

Abstract

Development funds are essential to finance climate change adaptation and are thus an important part of international climate policy. However, the absence of a common reporting practice makes it difficult to assess the amount and distribution of such funds. Research has questioned the credibility of reported figures, indicating that adaptation financing is in fact lower than published figures suggest. Projects claiming a greater relevance to climate change adaptation than they target are referred to as "overreported". To estimate realistic rates of overreporting in large data sets over times, we propose an approach based on state-of-the-art text classification. To date, assessments of credibility have relied on small, manually evaluated samples. We use such a sample data set to train a classifier with an accuracy of $89.81\% \pm 0.83\%$ (tenfold cross-validation) and extrapolate to larger data sets to identify overreporting. Additionally, we propose a method that incorporates evidence of smaller, higher-quality data to correct predicted rates using Bayes' theorem. This enables a comparison of different annotation schemes to estimate the degree of overreporting in climate change adaptation. Our results support findings that indicate extensive overreporting ($30.35 \pm 12.37\%$).

1. Introduction

There is international consensus on the need to respond to the global threat posed by climate change (Paris Accord, Article 2). Development funds are essential to finance climate change adaptation and are thus an important part of international climate policy. The 2009 Copenhagen Accord [1] aimed to mobilize USD 100 billion by 2020. Implementation of climate change adaptation measures is one of five targets set to reach the 13th Sustainable Development Goal (SDG): "Take urgent action to combat climate change and its impacts".

The Creditor Reporting System (CRS), maintained by the OECD Development Assistance Committee (DAC), monitors adaptation finance flows from OECD DAC member countries to developing countries. One of the challenges in ensuring valid reporting - or at least comparable figures - across reporting agencies is that the agreements mentioned above lack indicators. To this end, the OECD DAC established in 2009 the Rio markers on climate change adaptation

✉ janos.borst@uni-leipzig.de (J. Borst); Thomas.Wencker@deval.org (T. Wencker); andreas.niekler@uni-leipzig.de (A. Niekler)

© 2022

(CCA). For each aid activity, donors report whether it contributes to (CCA), i.e. reducing “the vulnerability of human or natural systems to the current and expected impacts of climate change, including climate variability, by maintaining or increasing resilience, through increased ability to adapt to, or absorb, climate change stresses, shocks and variability and/or by helping reduce exposure to them” [2, 4]. Activities are eligible for a marker if “a) the climate change adaptation objective is explicitly indicated in the activity documentation; and b) the activity contains specific measures targeting the definition above.” [2, 4]. The Rio marker r can take three values: 2, if CCA is the principal objective; 1, if CCA is a significant objective; and 0, if CCA is neither a principal nor a significant objective.

However, there is increasing evidence that the level of adaptation financing is in fact lower than published figures suggest [3, 4]. One possible reason is that there is no common practice for reporting climate finance [5, 6] and reporting agencies thus follow different reporting rules. This makes it difficult to assess the total amount of CCA finance, to compare commitments between donors, and to assess the geographical and sectoral distribution of funding [5]. Moreover, CCA finance estimates vary among reporting agencies [7]. Hence, aggregate figures of adaptation finance are increasingly considered unreliable given that they comprise thousands of individual aid activity descriptions from the CRS data. Consequently, assessments of credibility have to date relied on analysis of small samples covering a limited period of time [3].

This study applies state-of-the-art machine learning methods to estimate overreporting of CCA finance for all aid activities as reported in the OECD DAC CRS since the introduction of Rio markers. To the best of our knowledge, we are the first to use neural language models for this task.

Our main challenge in applying machine learning methodology is the quality and quantity of available annotated data. We have access to two data sets re-evaluated by experts and published in previous work: The first is small, but following a thorough re-evaluation process, we regard it as high quality. One concern is that the current de facto standard of fine-tuning language models tends to be unstable with very small data sets and is hard to evaluate properly. The second set is much larger, but because it was re-evaluated with access to less information, we regard it as lower quality; nevertheless, its size makes it adequate for training. We propose to combine these two data sets, using the larger for training and extrapolation, and the smaller, higher-quality set to correct first estimates.

2. Related Work

In recent years, several studies have estimated the level of overreporting in CCA finance [8, 3, 4, 9]. These studies are distinguished, first, by the rigor of their methodology to assess overreporting and, second, by the number of aid activities they analyze. Some studies classify multiple aid activities but employ, rather simplistically, keyword searches only on short descriptions of aid activity [8, 10, 4]. Other studies examine only a few aid activities by scrutinizing extant project documentation against in-country expert assessments [9]. Weikmans et al. (2017) strike a balance by manually assessing a large amount of short project descriptions.

This study breaks new ground in terms of method and sample size. We believe we are the first to apply state-of-the-art machine learning - which allows us to code all aid activities reported

in the OECD DAC CRS database - to fully automate the process of detecting overreporting of CCA finance. Moreover, our method can be easily applied to future data releases of the OECD DAC CRS data as well as comparable text data. However, machine-learning has already been used to classify the closely related SDGs automatically [11]. Here, we rely on textual resources to automatically assign Rio markers to CRS Reports. Project reports typically contain both short and long descriptions of the project goals, ranging from one to a few sentences. Currently neural networks produce virtually every state-of-the-art result in text classification, either by training task-specific architectures, e.g. [12] or adapting pretrained language models to a given task [13, 14, 15, 16, 17]. Also, recent works have achieved both higher overall performance [13, 15] and greater sample efficiency, achieving better results with less data [18]. The typical problems of previous classical machine learning approaches in text classification, like out-of-vocabulary or ambiguity, are directly handled by the language model, which is especially important in this case, because not all of the texts we deal with is free of orthographic and syntactical anomalies. We experiment with these models in various combinations to find the best fit for the task at hand.

3. Automatic Classification of Climate Change Adaptation Markers

3.1. Data

The CRS tracks OECD DAC member countries' aid activities. This study works with the original CRS data and two re-evaluated data sets:

Creditor Reporting System (CRS): The publicly available CRS contains harmonized data on aid activities. We use CRS data from 2006 to 2019 containing 1,529,984 aid activities. It includes up to 91 fields of data for each aid activity. The most important information in our context is flagged by: donor, recipient, Rio marker, project title, and short and long description.

WK: Weikmans et al. (2017) sampled 4,757 aid activities from 2012 CRS data and manually re-evaluated the Rio markers based on the aid activity descriptions. The re-evaluation includes a new marker (99) to indicate insufficient information for determining the Rio marker. [3] argue that label 99 can be treated as 0 (not climate adaptation related) because the Rio marker methodology explicitly requires a CCA objective to be indicated in the aid activity documentation.

CARE: Schramek and Harmeling (2021) of the Cooperative for Assistance and Relief Everywhere (CARE) sampled 117 aid activities from the CRS. Each case was re-evaluated and assigned a new Rio marker by experts with access to detailed project-level information beyond the data contained in the CRS.

3.2. Approach

We consider CARE as a high quality re-evaluation with very few samples. The WK data set has substantially more observations, but the re-evaluation had access only to CRS information. Since information from the CRS can be very limited, especially in cases where CCA is not the primary goal, this likely leads to a higher proportion of projects being considered overreported. However, the WK data set is substantially larger than the CARE data set and can be used to train a classifier, which is why we use WK for the training and CARE to estimate a correction factor to extrapolate

the CARE annotations implicitly.

Our approach is as follows: First, we train a high-quality classification model on the WK data set using information only from the CRS meta fields and the re-evaluated Rio markers. We mark a project as overreported if the classifier predicts a lower Rio marker than reported. Second, we calculate the classifier’s overreporting rate on the CARE data set. By comparing overreporting with the high-quality re-evaluation, we can estimate an error factor between the two annotation schemes in a Bayesian framework. Finally, we extrapolate to the complete CRS database and estimate overreporting rates for both annotation schemes.

3.3. Model Training and Model Selection

The WK data in the CRS provides us with text descriptions and the corresponding Rio markers (0, 1, 2 and 99), which we consider input and target of the classifier respectively. To find the best model, we test various language models with standard finetuning and in combination with a CNN architecture to find the best model. The CNN architecture follows [12] and comprises four 1D-convolutions with kernel sizes 3,4,5 and 6, with 100 filters each. The resulting vectors are max-pooled and projected by a linear layer onto the number of classes.

We conduct experiments with a RoBERTa [14] base model, a BERT [13] base model and a distilled version of RoBERTa [19] as published in the Hugging Face ’transformers’ library [20]. We use the AdamW [21] optimizer with a learning rate of $5e-6$, a batch size of 32, 25 epochs, and check-pointing to restore the best model with regards to average macro F1 score. To ensure stability and quality, we test these hyper-parameters for all models with tenfold cross-validation to identify the best model. The results of the cross-validation experiments for all model combinations are shown in Table 1. We used one Tesla V100 32GB for these experiments, one tenfold cross-validation for one model combination (one row in Table 1) took around five hours to complete for BERT and RoBERTa, and around 2.5 h for the distilled RoBERTa. After ensuring an average performance, we randomly split the data 80/20 and train a model with the same parameters.¹

		accuracy	macro		
			P	R	F1
CNN	roBERTa-base	89.81 ± 0.83	84.83 ± 2.0	80.54 ± 2.03	82.31 ± 1.8
	BERT	88.89 ± 1.3	83.35 ± 2.13	79.64 ± 2.81	80.97 ± 2.28
	distilroberta	89.16 ± 1.3	83.47 ± 2.24	81.28 ± 2.96	82.12 ± 2.21
Transformer	RoBERTa-base	89.64 ± 1.61	84.1 ± 2.33	80.58 ± 3.74	82.0 ± 2.87
	BERT	89.27 ± 1.5	84.57 ± 3.39	79.59 ± 2.41	81.52 ± 2.36
	distilroberta	89.14 ± 1.19	83.93 ± 2.76	79.47 ± 1.98	81.35 ± 1.93

Table 1
Aggregated results of the tenfold cross-validation for all tested models. Best results for each metric are highlighted in bold.

As shown in Table 1, the combination of CNN and RoBERTa not only reaches the highest average scores in accuracy and macro F1, but also the lowest standard deviation in the tenfold cross-validation. This leads us to believe that this model will not only generalize well but also

¹Code and final model will be made available upon publication to not jeopardize anonymity of the review.

will less likely deviate from the performance, which is why we choose this model as our classifier. Table 2 shows detailed results for the final model per label. It shows that predicting the label 1

	F1	P	R
0	94.17 ± 0.64	92.83 ± 1.13	95.57 ± 0.95
1	59.62 ± 7.28	64.68 ± 9.55	56.15 ± 8.38
2	86.74 ± 2.09	86.11 ± 4.2	87.52 ± 2.02
99	88.72 ± 3.09	95.69 ± 4.19	82.94 ± 5.04

Table 2

Detailed per-class results of the crossvalidation for the chosen model (CNN + RoBERTa).

is most difficult. Label "99" can be predicted with an F1 value of around 83%. We follow the argumentation in Weikmans et al. (2017) and regard these predictions as as Rio marker 0. The influence of these examples is negligible as, ultimately, 99 is predicted in only < 0.04% of the CRS in the end.

The Rio marker classification model is a proxy to identify overreported cases. We are interested in those cases where our classification algorithm differs from the reported Rio marker, specifically classifying lower than the reported value We define overreporting o of activity x as

$$\mathbf{o}(\mathbf{x}) = \begin{cases} 1 & \text{if reported}(\mathbf{x}) > \text{classifier}(\mathbf{x}) \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

This leads to three cases of overreporting: The classifier predicts 0 and the Rio marker reports 1 or 2, and the much harder case where Rio marker reports 2 and the classifier predicts 1.

3.4. CARE Data

We apply our classifier to the CARE data set and compare the findings to the manual CARE annotations. The set of examples in CARE data set are distinct from those in the WK data. We create an overreported flag for the CARE data by comparing the reported Rio marker to their re-evaluation markers using Equation (1). This marks 21.80% of the data as overreported. Our classifier predicts an overreporting rate of 54.14% on the same data, indicating a significant difference in the annotation schemes. WK annotations appear stricter, leading to higher rates of overreporting than the CARE re-evaluation.

3.5. Extrapolation of CARE Annotation using the Bayesian Formula

Using the Bayesian formula, we estimate the difference in annotation scheme and extrapolate it. As discussed above, our training relies on the WK data set. The approach of training and classifying new data ultimately transfers their annotation scheme to other data sets. Given the high number of hand-coded aid activities, the WK data set is well suited for training purposes. However, comparing the resulting classification with CARE data, we find that this might overestimate overreporting. We estimate the probability that if the classifier would mark any sample as overreported according to the WK data annotations, CARE annotations would agree, and vice versa. Using the Bayesian formula, we update the estimation of our classification. In mathematical

formulation we define two events: W (classifier marks sample as overreported) and C (CARE annotates sample as overreported). The Bayesian formula is:

$$P(C) = \frac{P(C|W)}{P(W|C)} \cdot P(W) \quad (2)$$

We calculate $P(W|C)$, i.e. the probability that our classifier would agree with CARE, and $P(C|W)$, i.e. the probability that CARE would agree with our classifier, from the data. Since the calculation is based on a small sample, we consider the uncertainty of the estimate using a Gauss distribution to approximate this factor and report the 95% confidence interval. We then use this factor to estimate the rate of overreporting according to CARE on the basis of the classifier by Equation (2). Since the sample is rather small, both estimates have quite large uncertainties. To combine these into $\frac{P(C|W)}{P(W|C)}$, we must propagate the individual uncertainties. For a general function $f = \frac{A}{B}$, where A and B have an error margin, it holds for the standard deviation σ_f of f that

$$\sigma_f^2 \approx f^2 \left[\left(\frac{\sigma_A}{A} \right)^2 + \left(\frac{\sigma_B}{B} \right)^2 - 2 \frac{\sigma_{AB}}{AB} \right] \quad (3)$$

For simplicity, to calculate an upper bound, we assume $\sigma_{AB} = 0$. CARE reports an overreporting rate of 21.8%. Our classifier predicts 54.14%. We find $P(C|W) = (27.78 \pm 14.44)\%$ and $P(W|C) = (68.97 \pm 9.47)\%$, which results in

$$\frac{P(C|W)}{P(W|C)} = (40.28 \pm 16.42)\% \quad (4)$$

To estimate the overreporting rate by using a CARE annotation scheme on the large CRS data set, we multiply the rate of overreporting found by our classifier with this correction factor.

3.6. Extrapolation and Exploration

We can now apply the classification algorithm to the CRS data. We restrict the CRS to projects that have a Rio marker higher than 0, otherwise, by definition, they cannot be overreported in Equation (1) and we consider only at the top five DAC donors: France, Germany, Japan, the United Kingdom and the United States. We use the fastText [22] language detection to classify the language of the descriptions. While Germany, Japan, the United Kingdom and the United States report almost all their projects in English, France tends to report in French. The classifier was also trained on French descriptions from the WK set, however, we predict Rio markers for these projects using both the original French descriptions and also automated translations into English using Google Translate (the influence of which is discussed below). After that, the data set contains 46,280 projects from 2010 to 2019 with short and long textual descriptions and a reported Rio marker of 1 or 2. This also complies with how data was sampled in the WK and the CARE data sets.

Figure 1 shows the results of the extrapolation. Table 5 in Appendix A shows the underlying values and the number of observations. Extrapolation is done, again, by concatenating the project title and long description into a text string and feeding it into the network. The network assigns a Rio marker prediction to every project. After that we use Equation (1) to mark all activities

with a flag for overreporting. The classifier detects an overall overreporting rate of **75.35 %** in the CRS in terms of WK annotations and an estimated **$30.35 \pm 12.37\%$** in terms of CARE annotation.

3.7. The Influence of Input Length

Systematic variation of text length by donor or year might bias our results. Longer descriptions usually contain more information and thus improve the validity of the classification. Elaborating on the difficulty of classifying short texts is beyond the scope of this paper and is its own established field of research, e.g. [23, 24]. This should specifically pertain to cases where CCA is a ‘significant’ but not a ‘principal’ objective (i.e., $r = 1$). Here, descriptions might not mention CCA because it is not the main motivation of the aid activity. Moreover, it seems likely that very short descriptions are mostly classified as $r = 0$ for lack of information.

We find evidence that classification quality correlates with description length. As we do not know the true Rio marker for aid activities, we use the rate of agreement between our classification and the assigned Rio marker as an indicator of classification quality. We assume that very low rates of agreement indicate poor performance of classification. As Figure 2 shows, the share of cases where classification results and Rio marker are identical increases with description length.

Description lengths systematically differ by donor and year (for both: $p < 0.00$, Kruskal-Wallis rank sum test), although the distribution of description lengths across donors and years shows that, overall, absolute differences are not large. Regarding donors, Japan is an exception, with considerably shorter descriptions than the median of the other donors (78 vs. 315 characters, respectively). The time series shows an increasing trend where the median of description length increased from 231 to 345 characters between 2010 and 2019 (see also Appendix B, Figure 4). If classification quality depends on description lengths, and description lengths vary by donor and year, this could introduce confounding bias distorting the comparison of overreporting rates across years and donors. More specifically, an increase in description lengths could be interpreted erroneously as a decrease in overreporting. As an example, Japan’s high rates of overreporting (see Figure 3) could be partly explained by the brevity of aid descriptions in the CRS.

To account for possible distortion of our results, we rerun our analysis excluding short descriptions from our estimation of overreporting. We used the interquartile range (IQR) method to identify outliers [25, p. 12], i.e. we excluded descriptions with lengths below the threshold of $Q1 - 1.5 \times IQR$ (in our case, 62 characters; logarithmic: 4.1). This seems appropriate as indicated by an increase in agreement between Rio marker and classifier (see Figure 2).

We rerun the analysis that created Table 5, but excluded all projects with fewer than 62 characters. We also excluded from the CARE data set data points shorter than 62 characters, when calculating the correction factor. A comparison of the results is presented in Figure 1. The overall overreporting rate per year drops slightly, while the estimation based on CARE increases. This stems from the fact that the classifier agrees with the high-quality CARE re-evaluation more often for longer texts. The correction factor and uncertainty in this case slightly increases from 40.28 ± 16.42 to 42.03 ± 16.03 . Overall this leads to slightly lower overreporting rate according to WK but a slightly higher estimate for CARE-corrected overreporting.

In summary, we do find evidence that systematic variation of text length by donor or year can bias results. However, robustness tests indicate that the bias does not change overall conclusions.




Figure 1: Estimated overreporting rate extrapolation of WK data by year (dots) and estimated correction based on CARE data (lines). Black dots and solid lines show estimates based on the full sample. Hollow circles and dashed lines show estimates based on samples excluding very short descriptions.

3.8. Extrapolation Per Donor Per Year

Year	France	Germany	Japan	U.K.	U.S.
2010	74.8%	74.66%	78.99%	73.11%	90.09%
2011	88.1%	84.16%	79.94%	71.95%	87.63%
2012	76.32%	77.52%	90.68%	74.47%	85.76%
2013	90.87%	74.57%	91.75%	69.9%	84.17%
2014	82.61%	75.3%	89.36%	56.46%	83.07%
2015	90.66%	74.78%	89.62%	58.52%	78.43%
2016	96.46%	74.41%	88.8%	55.26%	81.68%
2017	97.91%	72.92%	92.15%	49.76%	79.27%
2018	95.87%	72.04%	92.62%	41.12%	74.44%
2019	94.31%	69.23%	94.28%	45.61%	69.01%

Table 3

Overreporting rates - extrapolation of WK data split by donor country. Note that colored cells have fewer than 500 data points. Darker color corresponds to fewer data points.

The values in Table 3 show overreporting as identified by the classifier based on the WK data set from 2010 to 2019. France had the lowest number of projects in CRS reported with at least Rio marker 1 in 2011 and 2012 with 84 and 114 projects respectively. The rates in Table 3 for 2012 coincide rather well with the findings in [3] for Germany (81%), Japan (92%) and the United Kingdom(82%). The United States was not part of the sample in [3].




Figure 2: Agreement of classifier with Rio marker by lengths of activity description. Random noise is added to the point locations in the y-direction for better visibility. Circles show averages of binned data. The vertical dotted line indicates the cut-off for outliers based on the IQR method for outliers.

The classifier detects significantly lower overreporting rates for France (76.32%) than in the original paper (92%) in 2012. We found that, of the 114 projects reported in 2012 by France, only 58 have unique descriptions, which produces very high uncertainty. Manual evaluation shows that there are cases of the same description occurring up to 11 times and is marked as not overreported. This alone can account for around a 10% difference in overreporting measure when classified differently. We therefore marked the fields where there are a fewer data points than 500 in orange in Table 3 to show where these effects could have larger impact.

Figure 3 also shows the classification results when translating the French descriptions to English (dotted line). There are significant differences in overreporting rates in the period 2013-2015, while the two lines are reasonably close between 2016 and 2019. The year 2014 in particular shows a difference of around 50%, from 33.7% in the translated case to 82.6% when using the original French descriptions. We argue that this happens for two reasons: First, the years with the biggest differences when translating descriptions coincide with the years when very few projects reported, thus the influence of a single misclassification is higher. Second, there is considerable noise in the data, which further increases variance in prediction when switching language.

In 2014, France reported 184 projects, but there are only 53 unique descriptions. The most frequent description was used 36 times, each of the occurrences having a unique ID in the CRS. Half of these 36 projects are reported with a Rio marker 1 and half with Rio marker 2. Every classifier predicting a Rio marker on the basis of these descriptions will therefore differ from the reported value at least half the time. This alone accounts for a difference in overreporting of 10% in that year. This argument also holds for the second and third most frequent project descriptions, which were used 19 and 18 times respectively. When considering only the unique descriptions, the predictions' detection of overreporting in English and French agree in 77%

year	Over-reported	Over-reported (translated)	Count	Unique	Unique (relative)
2010	74.80%	81.71%	246	133	54.07%
2011	88.10%	82.14%	84	80	95.24%
2012	76.32%	64.04%	114	57	50.00%
2013	90.87%	66.54%	263	186	70.72%
2014	82.61%	33.70%	184	53	28.80%
2015	90.66%	66.54%	257	89	34.63%
2016	96.46%	95.29%	594	171	28.79%
2017	97.91%	94.91%	1100	477	43.36%
2018	95.87%	93.74%	847	372	43.92%
2019	94.31%	86.91%	1054	669	63.47%

Table 4

Overreporting for France following the WK data set by years. The second and third columns denote the estimated overreporting rate for French descriptions and English translations, respectively. The following columns show: number of projects, number of unique descriptions, and the ratio of these two.

of cases, while only in 51% of cases overall. In general, the larger the number of projects, the smaller the influence of this phenomenon. However, of the projects that France reported, only around 40% of the descriptions are unique (see Table 4 for more details).

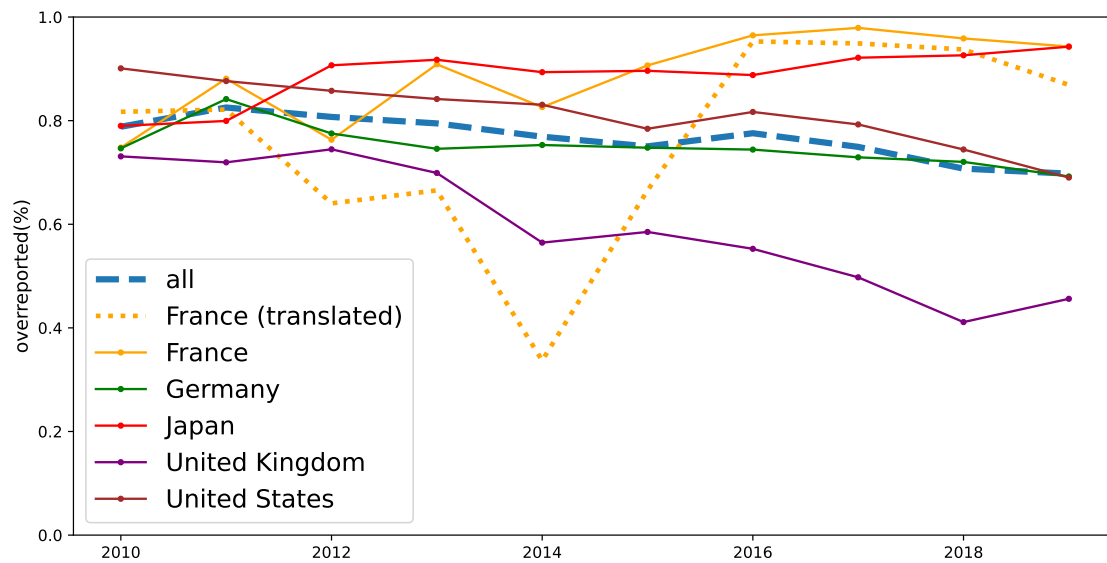


Figure 3: Rates of overreporting by donor country, 2010 to 2019

4. Discussion

Our re-evaluation of aid activities reported as contributing to CCA indicates a lack of quality in the self-reporting of donors. A substantial share of reported adaptation aid activities does not explicitly mention CCA in project descriptions. This is problematic because valid indicators are required to assess whether the international community is meeting its climate policy obligations as described in, for example the Paris Accord. Our finding indicates an overestimation of adaptation aid. Even after downward adjustment of our estimates to account for insufficient information from short project descriptions, our best estimate suggests that about every third activity categorized by donors as adaptation aid is not adaptation related. However, we cannot say whether this is due to a lack of clear reporting standards [6], a lack of compliance with reporting standards, or even incentives to report more than is actually delivered [8]. Moreover, our estimates are subject to considerable uncertainty because an unambiguous classification of aid activities based on the Rio Marker methodology requires extensive knowledge of individual aid activities.

Although the estimates are somewhat uncertain, our results confirm earlier findings of a substantial discrepancy between the figures reported by donors and re-evaluations by independent researchers [8, 3, 4, 9]. This can be partly explained by the fact that our approach draws on earlier classifications as training data and is thus not completely independent. However, our study also goes well beyond existing research in temporal and geographical scope: We assess every adaptation aid activity reported by OECD DAC donors since the Rio marker on adaptation aid was introduced in 2010. If donors had changed their reporting practice, we would likely see this in our data, yet, overall, we find no indication that reporting practices have changed significantly since 2010. The share of overreporting remains at a high level between 2010 and 2019. Although the classifier indicates a slight decrease of overreporting, the fluctuations are within the range of uncertainty of our adjusted estimates. Nevertheless, the results are mainly driven by a significant decrease in overreporting by the United Kingdom. However, we are careful to infer from our data substantial differences in reporting standards between countries; Country-specific results should be examined more closely in future research. We have no reason to suspect that the classifier has particular problems with data reported by the UK based on the results of the cross-validation and given the fact that most project descriptions are in English.

5. Conclusion

In this paper we propose an automated way of detecting overreporting of climate adaptation finance based on CRS project descriptions. Our approach is based on state-of-the-art text classification using finetuning neural language models. We consider the quality of the annotations of our training data when estimating overall overreporting rates, and propose a Bayesian approach to estimate an extrapolation of high-quality annotations. Our approach indicates significant overreporting throughout the study period.

There are two key challenges with this approach: The quality and quantity of annotation scheme data, and the quality of the textual input. While, ultimately, the first challenge can be overcome by extending data with higher-quality annotations, the second proves trickier. Unfortunately, the

CRS data does not have uniform quality built into its textual descriptions. We have discussed the influence of description length and language on the quality of our classifier, but while future work may incorporate techniques from short-text classification research, in many cases the information the CRS contains will likely not suffice to use techniques like augmentation or conceptualization. To improve on this, additional external data would be necessary, to which, at this point, we had no access. Especially deciding if the CCA aspects of a project comprises a "significant" or "principal" object should benefit from this. Another way to improve would be to pay more attention to multi lingual classification research and either incorporate techniques for multi lingual text classification or utilize a high-quality pipeline for translation.

References

- [1] UNFCCC, Copenhagen accord, 2009. URL: <https://unfccc.int/resource/docs/2009/cop15/eng/11a01.pdf>, (Accessed 2022-01-10).
- [2] OECD DAC, OECD DAC Rio Markers for Climate: Handbook, ???? URL: https://www.oecd.org/dac/environment-development/Revised%20climate%20marker%20handbook_FINAL.pdf, (Accessed 2022-01-07).
- [3] R. Weikmans, J. T. Roberts, J. Baum, M. C. Bustos, A. Durand, Assessing the credibility of how climate adaptation aid projects are categorised, *Development in Practice* 27 (2017) 458 – 471.
- [4] L. Junghans, S. Harmeling, Different tales from different countries: A first assessment of the oecd 'adaptation marker', Germanwatch Briefing Paper (2012). URL: <https://germanwatch.org/sites/default/files/publication/7083.pdf>, (Accessed 2022-01-07).
- [5] R. Weikmans, J. T. Roberts, The International Climate Finance Accounting Muddle: Is There Hope on the Horizon?, *Climate and Development* 11 (2019) 97–111. URL: <https://www.tandfonline.com/doi/full/10.1080/17565529.2017.1410087>. doi:10.1080/17565529.2017.1410087.
- [6] R. Weikmans, J. T. Roberts, S.-A. Robinson, What counts as climate finance? Define urgently, *Nature* 588 (2020). URL: <https://www.nature.com/articles/d41586-020-03481-0>.
- [7] S. Yeo, Climate Finance: The Money Trail, *Nature* 573 (2019) 328–331. URL: <https://media.nature.com/original/magazine-assets/d41586-019-02712-3/d41586-019-02712-3.pdf>.
- [8] A. Michaelowa, K. Michaelowa, Coding error or statistical embellishment? the political economy of reporting climate aid, *World Development* 39 (2011). URL: <https://www.sciencedirect.com/science/article/pii/S0305750X11001951>.
- [9] C. Schramek, S. Harmeling, Adaptation finance : Fact or fiction ?, 2021. URL: https://www.care-international.org/files/files/CARE_Climate_Adaptation_Finance_Fact_or_Fiction.pdf, (Accessed 2022-01-07).
- [10] J. T. Roberts, K. Starr, T. Jones, D. Abdel-Fattah, The Reality of Official Climate Aid (2008). URL: <https://www.oxfordenergy.org/wpcms/wp-content/uploads/2011/01/November2008-ClimateAid-BenitoMuller-.pdf>, (Accessed 2022-07-01).
- [11] A. Pincet, S. Okabe, M. Pawelczyk, Linking Aid to the Sustainable Development Goals – a machine learning approach, *OECD Development Co-operation Working Papers* 52, OECD Publishing, 2019. URL: <https://ideas.repec.org/p/oec/dcdaaa/52-en.html>. doi:10.1787/4bdaeb8c-en.
- [12] Y. Kim, Convolutional Neural Networks for Sentence Classification, in: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Association for Computational Linguistics, Doha, Qatar, 2014, pp. 1746–1751. URL: <http://aclweb.org/anthology/D14-1181>. doi:10.3115/v1/D14-1181.
- [13] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, 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 (Long and Short Papers)*, Association for Computational Linguistics, 2019, pp. 4171–4186. URL: <https://aclanthology.org/N19-1423>.
- [14] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer,

- V. Stoyanov, Roberta: A robustly optimized bert pretraining approach, Computing Research Repository arxiv:1907.11692 (2019).
- [15] Z. Yang, Z. Dai, Y. Yang, J. G. Carbonell, R. Salakhutdinov, Q. V. Le, XLNet: Generalized Autoregressive Pretraining for Language Understanding, in: H. M. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché Buc, E. B. Fox, R. Garnett (Eds.), Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems (NeurIPS), 2019, pp. 5754–5764. URL: <http://papers.nips.cc/paper/8812-xlnet-generalized-autoregressive-pretraining-for-language-understanding>.
- [16] R. Aly, S. Remus, C. Biemann, Hierarchical Multi-label Classification of Text with Capsule Networks, in: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop, Association of Computational Linguistics, Florence, Italy, 2019, pp. 323–330. URL: <https://www.aclweb.org/anthology/P19-2045>.
- [17] A. Pal, M. Selvakumar, M. Sankarasubbu, Magnet: Multi-label text classification using attention-based graph neural network, Proceedings of the 12th International Conference on Agents and Artificial Intelligence (2020). URL: <http://dx.doi.org/10.5220/0008940304940505>. doi:10.5220/0008940304940505.
- [18] K. Halder, A. Akbik, J. Krapac, R. Vollgraf, Task-aware representation of sentences for generic text classification, in: Proceedings of the 28th International Conference on Computational Linguistics, International Committee on Computational Linguistics, Barcelona, Spain (Online), 2020, pp. 3202–3213. URL: <https://aclanthology.org/2020.coling-main.285>. doi:10.18653/v1/2020.coling-main.285.
- [19] V. Sanh, L. Debut, J. Chaumond, T. Wolf, Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter, in: The 5th Workshop on Energy Efficient Machine Learning and Cognitive Computing (NeurIPS), 2019.
- [20] T. Wolf, L. Debut, V. Sanh, J. Chaumond, C. Delangue, A. Moi, P. Cistac, T. Rault, R. Louf, M. Funtowicz, J. Davison, S. Shleifer, P. von Platen, C. Ma, Y. Jernite, J. Plu, C. Xu, T. L. Scao, S. Gugger, M. Drame, Q. Lhoest, A. M. Rush, Transformers: State-of-the-art natural language processing, in: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, Association for Computational Linguistics, 2020, pp. 38–45. URL: <https://www.aclweb.org/anthology/2020.emnlp-demos.6>.
- [21] I. Loshchilov, F. Hutter, Decoupled weight decay regularization, in: International Conference on Learning Representations (ICLR), 2019.
- [22] A. Joulin, E. Grave, P. Bojanowski, T. Mikolov, Bag of tricks for efficient text classification, in: Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Short Papers, volume 2, Association for Computational Linguistics, Valencia, Spain, 2017, pp. 427–431. URL: <https://aclanthology.org/E17-2068>.
- [23] J. Chen, Y. Hu, J. Liu, Y. Xiao, H. Jiang, Deep short text classification with knowledge powered attention, Computing Research Repository arxiv:1902.08050 (2019). URL: <https://arxiv.org/pdf/1902.08050.pdf>.
- [24] J. Wang, Z. Wang, D. Zhang, J. Yan, Combining knowledge with deep convolutional neural networks for short text classification, in: International Joint Conference on Artificial Intelligence (IJCAI), 2017.
- [25] I. F. Ilyas, X. Chu, Data cleaning, number 28 in ACM Books, 1st ed., Association for

A. Overreporting per Year

Tables 5 and 6 show detailed numerical results on estimating overreporting split by year. These two tables were summarized in Figure 3, comparing the resulting rates.

Year	Classifier	CARE Estimated	Count
2010	78.82%	26.54 – 36.96%	1818
2011	82.54%	27.79 – 38.7%	3207
2012	80.72%	27.17 – 37.85%	2894
2013	79.47%	26.75 – 37.26%	3731
2014	76.9%	25.89 – 36.06%	3745
2015	75.03%	25.26 – 35.18%	4314
2016	77.57%	26.11 – 36.37%	5306
2017	74.94%	25.23 – 35.14%	7306
2018	70.74%	23.81 – 33.17%	6654
2019	69.77%	23.49 – 32.71%	7710

Table 5

Overreporting rates - extrapolation of WK data and the estimated correction of CARE data.

year	Classifier	CARE Estimated	count
2010	77.12%	27.22 – 37.61%	1683
2011	81.83%	28.88 – 39.9%	3076
2012	81.11%	28.62 – 39.55%	2858
2013	79.83%	28.17 – 38.93%	3698
2014	76.89%	27.14 – 37.5%	3739
2015	75.0%	26.47 – 36.57%	4308
2016	77.55%	27.37 – 37.82%	5300
2017	74.86%	26.42 – 36.51%	7284
2018	70.65%	24.93 – 34.46%	6631
2019	67.22%	23.72 – 32.78%	7111

Table 6

Overreporting rates - extrapolation of WK data and the estimated correction of CARE data after eliminating all inputs with fewer than 62 characters.

B. Analysis of Text Length

Figure 4 illustrates increasing average descriptions lengths over the study period in a standard box plot, for the argument in section 3.7. Also, Japan tends to use fewer characters in descriptions compared to the other donors. This supplements Figure 2 and informs the decision to use the *IQR* range to quantify the lower limit for the number of characters. This also checks that cutting off at 62 characters (logarithmic: 4.1) does not introduce a bias for a particular year or donor.




Figure 4: Activity description lengths (logarithmic) by donor (top figure) and by years (bottom figure).