

Toxic Bias: Perspective API Misreads German as More Toxic

Gianluca Nogara^{*1}, Francesco Pierri^{*2}, Stefano Cresci³, Luca Luceri⁴,
Petter Törnberg⁵, Silvia Giordano¹

¹University of Applied Sciences and Arts of Southern Switzerland, Switzerland

²Politecnico di Milano, Italy

³IIT-CNR, Italy

⁴USC Information Sciences Institute, USA

⁵University of Amsterdam, Netherlands

gianluca.nogara@supsi.ch, francesco.pierri@polimi.it, stefano.cresci@iit.cnr.it,

lluceri@isi.edu, p.tornberg@uva.nl, silvia.giordano@supsi.ch

Abstract

Proprietary public APIs play a crucial and growing role as research tools among social scientists. Among such APIs, Google’s machine learning-based Perspective API is extensively utilized for assessing the toxicity of social media messages, providing both an important resource for researchers and automatic content moderation. However, this paper exposes an important bias in Perspective API concerning German language text. Through an in-depth examination of several datasets, we uncover intrinsic language biases within the multilingual model of Perspective API. We find that the toxicity assessment of German content produces significantly higher toxicity levels than other languages. This finding is robust across various translations, topics, and data sources, and has significant consequences for both research and moderation strategies that rely on Perspective API. For instance, we show that, on average, four times more tweets and users would be moderated when using the German language compared to their English translation. Our findings point to broader risks associated with the widespread use of proprietary APIs within the computational social sciences.

Introduction

In recent years, abusive interactions have become a fixture of the online experience, with approximately 40% of U.S. adults self-reporting as victims of online abuse (Vogels 2021; Pierri 2020). The prevalence of such language on social media has become a major concern for both researchers and the public, often under the label of online “toxicity”. Currently, the leading approach to detecting toxicity is the Perspective API¹, offered by the Jigsaw unit of Google (Lees et al. 2022) and designed to “help mitigate toxicity and ensure healthy dialogue online”. The API is widely used both within academic research efforts and for content moderation on news and social media platforms. As of January 2024, there are over 1,400 documents on Google Scholar mentioning the term “Perspective API”, and the API is used by major

platforms such as Reddit, The New York Times, El Pais, and Faceit.

Perspective API’s machine learning model is developed through supervised learning, utilizing a vast dataset of millions of comments gathered from diverse online platforms, including forums like Wikipedia and The New York Times, and encompassing more than 20 languages. Perspective API defines *toxic* messages as texts that use “rude, disrespectful, or unreasonable language that is likely to make someone leave a discussion” (Gehman et al. 2020; Welbl et al. 2021). The API provides several different toxicity scores, namely *Severe toxicity*, *Insult*, *Profanity*, *Identity attack*, *Threat*, and *Sexually explicit*. In a 2019 SemEval Task, Perspective API has been shown to perform better than other transformer-based models on canonical datasets for hate speech detection (Pavlopoulos et al. 2019). More recent work found that its scores differ from human labels (Welbl et al. 2021). The toxicity score, which ranges between 0 and 1, does not have a meaning in absolute terms, and the traditional approach to identify toxic or hateful content in online social platforms involves defining a threshold on the toxicity score, which is generally in the region of 0.5-0.7 (Hua, Ristenpart, and Naaman 2020; Hua, Naaman, and Ristenpart 2020), above which content is considered to be toxic.

While technologies such as Perspective API promise a safer, more respectful digital environment, their effectiveness is contingent on their accuracy and impartiality across diverse languages and cultural contexts. This paper aims to address the following research question:

- **RQ: Is Perspective API Biased Towards the German Language?**

By leveraging two large-scale datasets of multilingual Twitter conversations and random Wikipedia summaries, we identify and report on a systemic bias within Perspective API when dealing with German texts, showing a tendency of the tool to assign higher toxicity scores to German content compared to other languages. We show that toxicity scores for texts in German are consistently significantly higher than those for any other language, *regardless of the content*. Our findings suggest a strong negative bias against the German language, potentially due to artifacts or biases in the training

^{*}These authors contributed equally.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹<https://www.perspectiveapi.com/>

dataset or the model itself, which is challenging to investigate further given the black-box nature of the tool. Our findings not only question the reliability of a tool employed by several large online platforms – potentially leading to unjust censorship or misrepresentation in online spaces – but also aggravate the growing reliance on proprietary, black-boxed APIs like Perspective within the academic realm. These systems present a fundamental problem as researchers and users cannot scrutinize or understand the inner workings of the algorithms, leading to a substantial risk of systemic biases being embedded and perpetuated through their employment across news outlets, social media, and various online platforms.

The outline of this paper is the following: we first examine recent work employing Perspective API for toxicity research; next, we describe the datasets employed in our analysis; then, we show results for a number of different analyses. Finally, we outline our findings and discuss implications, limitations and future directions.

Perspective API and Toxicity Research

Perspective API has become a foundation for academic research on online abuse and incivility (Rieder and Skop 2021), as well as a relevant case study for legal and social scholars interested in the regulatory nuances of socially relevant AI systems (Friedl 2023). A diverse and rich research field relies on the capacity of the API to reliably assess these aspects of online text. Rather than reviewing the diverse contributions that employ Perspective API to study online toxicity, we will here examine some examples of research that draws on the API concerning the analysis of German language messages, to assess the potential implications of the findings of this paper.

There are many examples of research using Perspective API for cross-national comparison of toxic communication on social media. Rye et al. (Rye, Blackburn, and Beverly 2020), for instance, analyze 14 months of data activity on the far-right platform Dissenter, using Perspective API to measure the toxicity of messages and derive the propensity of specific content to elicit toxic comments. They find that “deutschland.de” is the second most toxic domain, with most angry comments being related to the Muslim diaspora in Europe. Similarly, Hoseini et al. (Hoseini et al. 2023), analyze Telegram messages related to the QAnon conspiracy shared between September 2019 and March 2021, finding that German and Portuguese messages are more toxic than those in English.

Even when studies are not explicitly comparative, they often include cross-language analysis of toxicity, which may potentially be vulnerable to the type of model biases discussed in this paper. Smirnov et al. (Smirnov, Oprea, and Strohmaier 2023), for instance, study the effect of toxic comments directed to Wikipedia’s volunteer editors. Analyzing over 57 million comments on the six most active language editions of Wikipedia (English, German, French, Spanish, Italian, and Russian), they find that toxic speech could lead to a substantial loss of productivity among editors. Differently from the majority of other contributions, Ghosh Roy et al. (Ghosh Roy et al. 2020) do not directly use

Perspective API to identify toxic tweets, but rather use toxicity scores from the API as features in a transformers-based hate speech classifier for English, German, and Hindi. This usage, however, could leak Perspective API’s biases into the trained classifier.

Many studies focus specifically on German-language debates, seeking to draw out the dynamics of toxicity within these social spaces. Lasser et al. (Lasser et al. 2023) study more than 1M tweets posted on Twitter from 2015 to 2018 in German, a period characterized by a “large influx of refugees into Germany which sparked heated discussions and political polarization”. They measure the interaction between incivility (hatefulness and toxicity) and extremity (of speech and speakers) with other dimensions of discourse. For this purpose, they use Perspective API to measure the overall toxicity score as the average of five attributes (toxicity, severe toxicity, profanity, insult, and identity attack). They report very high values of average toxicity and highlight which behaviors can reduce hate speech. Another example is Czymara et al.’s analysis of the effect of Islamist terror attacks on ethnic insulting comments on YouTube (Czymara et al. 2023), showing how terror attacks boost interest in immigration-related topics in general, and lead to a disproportional increase in hate speech in particular. Focusing on the “Identity attack” attribute of Perspective API, they classify insulting comments as those with a score surpassing 0.9. The study reveals varying proportions of ethnically insulting comments, ranging from 2.7% in the French data to 5.4% in the German data and 6.3% in the UK data.

While Perspective API has seen widespread use, few studies have systematically evaluated the accuracy and reliability of the model either across languages or specifically in German. Among them, (Gargee et al. 2022) highlights a discrepancy in Perspective API’s toxicity assessments of sentences that are semantically equivalent, but with different arrangements of words. In detail, they show that switching from active to passive voice, as well as reordering words in a sentence, could lead to significantly different toxicity scores. Then, Wich et al. study toxic German-language speech on Telegram (Wich et al. 2022). They develop a custom classifier and compare its performance with Perspective API, finding that their model exhibits a higher F1 score (54.75% vs. 53.50%) and macro F1 score (73.20% vs. 70.51%). Mihaljevic et al. (Mihaljević and Steffen 2023) use antisemitism in the German language on Telegram to examine how Perspective API scores correlate with different sub-forms of antisemitism. They also demonstrate that manipulating messages with widely known antisemitic codes can substantially reduce toxicity scores, making it easy to bypass content moderation based on Perspective API. While such studies suggest the need for more research on potential biases of the model, the black-box nature of the model limits the capacity of researchers to understand the limitations and dynamics of the classification. Research therefore has to draw on observational data to seek to identify biases in model classification.

Austria	Belgium	France
Germany	Ireland	Italy
Netherlands	Poland	Spain
Sweden	Switzerland	

Table 1: Countries included in the Dataset 1 (representative sample of international tweets across languages).

Data

Our analyses are based on a combination of real-world data from Twitter and Wikipedia, as described in the following subsections.

Dataset 1: Representative Sample of International Tweets Across Languages

Our first Twitter dataset contains a random sample of tweets shared between 2018 and 2021 in 14 countries (see Table1), collected using the Twitter Academic API on Sept. 12, 2022. The collection was done by repeatedly selecting a random time point between 1/1/2018 and 1/1/2022, and fetching the 500 most recent messages sent in a particular country with a particular language, until approximately 250 000 messages were collected for each region. The country and language were filtered using the Twitter API *place_country* parameter, combined with the *lang* parameter to specify the language. The resulting dataset represents a roughly representative sample of Twitter messages uniformly distributed over time. The procedure thus produces a dataset with a pseudo-random sample of Twitter messages from 11 countries over the specified time frame. The total number of tweets equals 2 666 696, while those shared in German-speaking countries, such as Austria (AT), Germany (DE), and Switzerland (CH), are 733 451 (27.5%). Having multiple countries and languages enables showing the unique biases of German classification, and including several German-speaking countries allows verifying that the results do not follow from the dynamics of online conversation in a specific country.

Dataset 2: COVID-19 Tweets in Italian and German

The second Twitter dataset contains tweets related to COVID-19 vaccines in the Italian and German languages, shared between November 1, 2020, and June 30, 2021 – a time when vaccination campaigns kicked off in many European countries. The resulting number of tweets is 1 453 450 for Italian and 1 219 445 for German. The data were collected through Twitter’s Historical Search API in July 2021 and are described in detail in (Giovanni et al. 2022). These larger-scale data enable a deeper comparison of Perspective API’s labeling dynamics between German and another language.

Dataset 3: Random Sample of Wikipedia Texts

The third dataset consists of a random sample of summaries from Wikipedia pages in English and German language. Specifically, we collected 12 279 summaries of a random sample of pages available in both languages leveraging

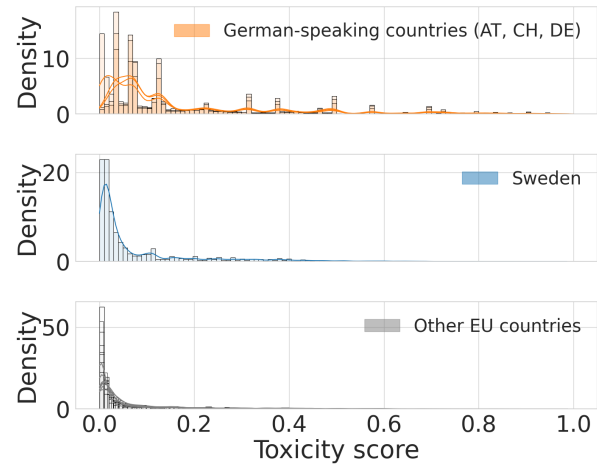


Figure 1: Distribution of toxicity scores for tweets shared in German-speaking countries (Austria, Switzerland, and Germany) versus those in other EU countries, from Dataset 1. Histograms are built with 100 equal-width bins. Distributions are statistically different at $\alpha = 0.05$ according to a Kruskal-Wallis test. The median toxicity of tweets shared in German-speaking countries is 0.075, and ~ 0.023 when considering the overall distribution of tweets in other EU countries.

Wikipedia Python library². We further collected a similar amount of random texts in Arabic, Chinese and Japanese³ to test whether languages with non-ASCII characters might exhibit biases in toxicity scores. Including random texts allows verifying that the biases identified do not apply only to Twitter messages but to any text.

Analyses and Results

This section presents the results of the analyses conducted on the three aforementioned datasets. In each case, we ran Perspective API on all available texts focusing on the attribute “Toxicity” – the estimated probability for a comment to be perceived as toxic – as it is the most general and widely employed in the literature. Specifically, we employed Perspective API in September 2022 to obtain toxicity scores for Dataset 1, whereas we ran it between September and December 2023 for labeling Datasets 2 and 3.

Our analyses first investigate the distribution of toxicity scores in a random sample of multilingual tweets (Dataset 1). Next, we perform a similar analysis for tweets in Italian and German language related to COVID-19 vaccines (Dataset 2). For this dataset, we also assess the potential impact of Perspective API’s language bias when it is used for automatically flagging toxic messages for removal – a common content moderation strategy (Rieder and Skop

²<https://pypi.org/project/wikipedia/>

³They do not correspond to the same texts across the five languages, as we were unable to retrieve summaries that satisfied this condition with the Wikipedia API.




Figure 2: Distribution of toxicity scores for tweets shared in German-speaking countries (Austria, Switzerland, and Germany) from Dataset 1, separating texts that contain at least one non-ASCII character from those that only contain ASCII characters. Histograms are built with 100 equal-width bins. The median toxicity of the two distributions are respectively 0.08 and 0.07.




Figure 3: Top 10 most frequent Perspective API scores rounded at 8 decimal digits for tweets shared in German-speaking countries (Austria, Switzerland and Germany) and those in other EU countries, from Dataset 1.

2021). Lastly, we analyze biases in Perspective API scores when considering a random set of Wikipedia summaries in their corresponding English and German language versions (Dataset 3).

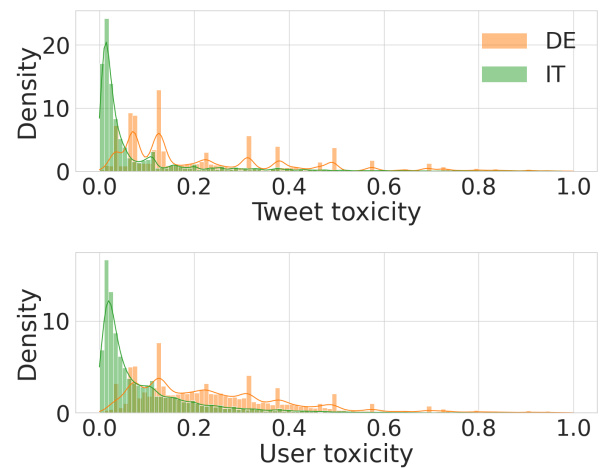


Figure 4: Distribution of toxicity scores for COVID-19 vaccine-related tweets (top) and users (bottom) for German and Italian language, from Dataset 2. Histograms are built with 100 equal-width bins. Median values for tweet toxicity are DE = 0.132, IT = 0.026; for user toxicity DE = 0.212, IT = 0.048.

German Tweets Are Labeled as More Toxic Than Those From Other Languages

We first study the distribution of toxicity scores for tweets shared in German-speaking countries – namely Austria, Germany, and Switzerland – compared to other countries. As shown in Figure 1, tweets from German-speaking countries (which are mostly in the German language) are significantly more toxic (Kruskal-Wallis $P < .001$) than those from other countries, with a median toxicity of 0.075 versus 0.023 when considering tweets from the two groups of countries altogether. The latter distributions closely resemble smooth exponential distributions, which is typical of a range of phenomena on social media (Adamic et al. 2001; Kwak et al. 2010), whereas the former are spiky – suggesting the presence of some type of error or classification artifact.

Non-ASCII Characters Are Not Related to Higher Toxicity of Tweets To investigate whether the presence of non-Latin characters (such as ä, ö, ü) might cause this abnormal behavior, we also plot in Figure 1 the distribution of tweets shared in Sweden separately from other countries (see middle panel), finding no evidence of spikes. We further compare the toxicity of texts in German that contain at least one non-Latin character (i.e., with an ASCII code greater than 127) from those that only contain Latin characters, finding very similar spiky distributions in the two classes (see Figure 2). We observe similar patterns also when separating tweets on the number of non-ASCII characters they contain in the range [1, 12] (i.e., up to the 99th percentile of the number of non-ASCII characters in tweets from German-speaking countries). We provide the figure in the Appendix (Figure 10).

Abnormal Peaks of Toxicity Score in German Tweets

To explore abnormal peaks of toxicity, we look into the top 10 most frequent toxicity scores returned by the API after rounding them at 8 decimal digits. As shown in Figure 3, tweets in German exhibit an unusual pattern, as some very specific scores are extremely common. While a smooth distribution would imply that each score should account for a negligible proportion of all scores, as in the case of the distribution for all other languages (in grey in Figure 3), a few scores are disproportionately frequent for German texts – in line with abnormal spikes highlighted in Figures 1 and 2, accounting for more than 10% of the total toxicity scores returned by the Perspective API. This offers further evidence of the presence of some internal issues in the toxicity model leading to erroneous labeling of German language text.

No Evidence of Specific Words Associated With Abnormal Toxicity Scores

Notwithstanding the fact that Perspective API is a black-box machine learning model, we perform an exploratory analysis to understand whether specific words might cause the artifacts in the toxicity score. Specifically, we first assign tweets to two classes: *Abnormal* tweets if they are assigned one of the scores described in the previous section, *Normal* otherwise. Then, we compute TF-IDF embeddings of such tweets (excluding words that appear less than 1% of the time), and we train and test two off-the-shelf machine learning classifiers (Random Forest and Gaussian Naive Bayes) using the *scikit-learn* (Pedregosa et al. 2011) library in the binary classification task of predicting whether a tweet belong to one of the two classes. We observe that they perform no better than a random baseline across several performance metrics, as shown in Figure 11 of the Appendix. We also leverage the *Fightin' Words* approach first described in (Monroe, Colaresi, and Quinn 2008) and recently employed by (Gligoric et al. 2024) to measure statistically significant differences in words between the two classes, which might be causing the spikes in the distribution. By looking at the top-10 differentiating words for the Abnormal class, as shown in Figure 12 of the Appendix, we observe that they are associated with a slightly larger toxicity score and correspond to very common terms (e.g., *nicht*, *die* and *sie*). However, tweets that do not contain such words still exhibit the abnormal peaks.

German Vaccine Tweets Are More Prone to Removal

To examine the implications of the identified biases in toxicity scores for realistic real-world scenarios of content moderation, we leverage the dataset on COVID-19 vaccine-related tweets. COVID-19 information has frequently been raised as an example in which content moderation is particularly crucial (Ferrara, Cresci, and Luceri 2020; Krishnan et al. 2021).

Toxicity of Tweets and Users We first analyze the distribution of toxicity scores for tweets shared in German and Italian. As shown in the top panel of Figure 4, German tweets are significantly more toxic than Italian tweets. The bottom panel of Figure 4 highlights that the same result also holds when considering user-level toxicity, which we computed as the average of all tweets shared by each user. The




Figure 5: Distribution of toxicity scores for German COVID-19 vaccine-related tweets and for their English translation, from Dataset 2. Histograms are built with 100 equal-width bins. Median values are: DE = 0.132, EN = 0.012.




Figure 6: Distribution of the differences in toxicity scores for German and Italian tweets compared to their English translation, using a random sample of 10 000 tweets from Dataset 2. Histograms are built with 100 equal-width bins. Median values are DE = +0.085, IT = -0.006.

Italian and German distributions are statistically different in both cases according to a Kruskal-Wallis test ($P < .001$).

Toxicity of Tweets Translated to English To rule out the possibility that the higher toxicity scores assigned to German texts are due to more toxic content, we translate both German and Italian tweets into English and we subsequently use Perspective API to compute toxicity scores for the English texts. We leverage Argos Translate⁴ to obtain English translations of our tweets. Figure 5 shows the distribution of toxicity scores for all German tweets (solid line) and their English translation equivalents (dashed line). We observe that the tweets translated in English are much less toxic, with their distribution being similar to that of generic tweets in non-German languages, presented in Figure 1.

However, it may be that the machine-translation will tend to reduce the toxicity of messages, for instance by avoiding to translate certain swearwords or slang. To examine whether this is the case, we also compute the difference in toxicity between a random sample of 10 000 tweets in German compared to their respective English and Italian translation. Figure 6 shows the distribution of these differences.

⁴<http://www.argosopentech.com/>




Figure 7: Distributions of difference in toxicity between German tweets from Dataset 2 and their English translation, divided by 20 bins of toxicity computed on the German version of each tweet. Outliers are not shown and triangles indicate the mean of the distribution.

We observe that, on average, German messages are 8.5% more toxic (i.e., median difference = +0.085) than the corresponding English messages, as also shown by the skewed orange-colored distribution. Contrarily, the differences between Italian and English-translated messages appear to follow a normal distribution centered around 0 and, in fact, Italian messages only show negligible differences (median difference = 0.01) when compared to their English translations. These results unequivocally demonstrate that the larger toxicity scores for German texts do not depend on the content of such texts, but are rather the consequence of an inherent language bias in Perspective API. To rule out potential biases in the translation that might modify the toxicity, we perform the same experiments by starting from a random sample of 10 000 Italian tweets and translating them to German and English to compute differences. Also in this case, we obtained similar results, but omit the figure for the sake of brevity.

Correlation Between Toxicity and Difference With English Translation

We study how Perspective bias affects German tweets in the range of toxicity [0, 1] by dividing tweets of Dataset 2 in bins based on their German toxicity score, and computing the difference with the toxicity score assigned to their English translation, for each bin. Results are shown in Figure 7 and highlight that large differences are found in tweets with high German toxicity scores. We also find a significant positive Pearson correlation ($R=0.9$, $P < .001$) between difference and toxicity (not binned). These results highlight that the bias of Perspective API is significantly higher for more toxic tweets. In other words, German tweets with a larger toxicity score are significantly over-estimated as toxic.

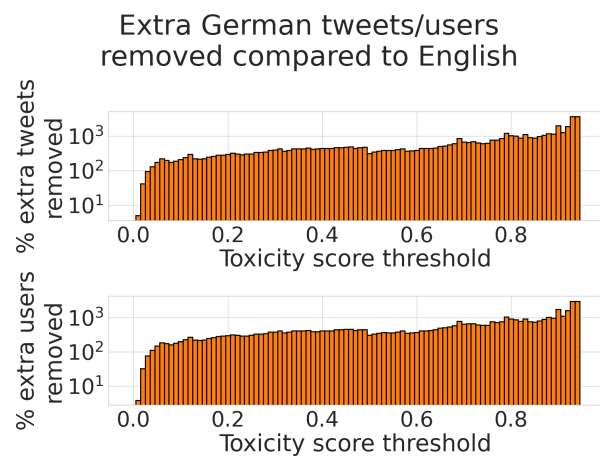


Figure 8: Proportion of extra tweets (top) and users (bottom) removed in German compared to English for different choices of toxicity score as a threshold, i.e., tweets/users above a given score would be removed. The median increase in percentage points is +429.94 for tweets and +409.84 for users, meaning that on average more than 4 times German tweets and users would be removed compared to English.

Implications for Content Moderation To assess the implications of this bias, we analyze how moderation strategies aimed at minimizing online toxicity might affect decision-making processes, especially when a moderator or an automatic system relies on toxicity scores from the Perspective API to oversee and regulate discussions on a social media platform (Rieder and Skop 2021).

Specifically, we focus on how different toxicity thresholds used for removing tweets or users might affect COVID-19 vaccine conversations in German when compared to their translation in English, using the random sample of 10 000 tweets employed in the previous analysis. In the top panel of Figure 8, we simulate removing all tweets whose toxicity score exceeds a given threshold, and we show the percentage of extra tweets removed when considering German scores compared to English scores. For instance, by adopting the threshold value of 0.7 recommended by Perspective API⁵, 10 times more German tweets would be removed, and similarly for the number of German users (as shown in the bottom panel of the figure). The median increase in percentage points is +429.94 for tweets and +409.84 for users, meaning that on average more than four times German tweets and users would be removed compared to English.

Wikipedia Texts in German Are More Toxic Than in English

To verify that the identified biases do not merely apply to Twitter messages, we also examine the distribution of toxicity scores when classifying a random sample of summaries from Wikipedia. As shown in the top panel of Figure 9, also

⁵https://developers.perspectiveapi.com/s/about-the-api-score?language=en_US

for these texts, German exhibits significantly larger toxicity scores than those in English (Kruskal-Wallis $P < .001$). We also observe peaks in the distribution of German scores, which resemble those of previous distributions, such as Figure 1. We perform an additional robustness test by analyzing Wikipedia summaries in three non-European languages supported by Perspective API (Arabic, Chinese and Japanese) and computing their toxicity scores. As shown in the bottom panel of Figure 9, they do not exhibit abnormal peaks as for the German language. An analysis of the most frequent scores confirms the findings provided in Figure 3, with German texts exhibiting scores that are disproportionately more frequent than other languages.

Discussion

Contributions

Perspective API has emerged as the technological foundation for a broad spectrum of applications, stretching from academic research to practical implementations in content moderation across various online platforms. In the realm of research, scholars and data scientists utilize Perspective API to analyze online discourse, study patterns of hate speech and harassment, and explore the dynamics of digital communication. Beyond the academic sphere, Perspective API plays a pivotal role in content moderation for websites, forums, and large-scale social media platforms. By automatically flagging potentially harmful or abusive comments, it aids moderators in maintaining a constructive and respectful environment, enabling large-scale platforms to manage large volumes of user-generated content that would make manual moderation impractical. Perspective API is often integrated into comment systems and discussion forums, enabling preemptive identification and removal of messages that are classified as toxic. For these reasons, automated tools such as Perspective API apparently contribute to overcoming the challenges in content moderation and to respecting the regulatory obligations that large online platforms must abide by (Trujillo, Fagni, and Cresci 2023). However, while the widespread application of such tools means that these systems provide important support for researchers and platforms alike, it also highlights the potential negative impact of any biases or inaccuracies in the data and algorithms.

This paper identified a significant bias within the Perspective API model when it comes to processing German-language messages: our analysis reveals that the algorithm tends to assign disproportionately higher toxicity scores to texts written in German, compared to similar content in other languages. Specifically, our examination of Twitter posts revealed that the same content reported in German was more toxic than its English-translated version. This pattern was further validated by analyzing random texts from Wikipedia summaries in both languages. The distribution of toxicity returned by the API for German appears spiky, rather than following the smooth distribution of other languages – suggesting some errors or artifacts stemming from the underlying machine learning model. Due to the black box and proprietary nature of the model, we unfortunately cannot examine the underlying reasons for these results,




Figure 9: Distribution of toxicity for over 10 000 random Wikipedia page summaries in Arabic (ar), Chinese (zh), English (en), German (de) and Japanese (ja). Histograms are built with 100 equal-width bins. We limit the x-axis to 0.3 for clarity reasons, as the distribution is highly skewed towards 0.

which calls for additional research efforts. However, an exploratory analysis showed that no specific characters and words are associated with such spiky behaviour, and that no abnormal patterns are observed for other non-European languages such as Arabic, Chinese and Japanese.

This bias raises profound concerns, as it not only skews research findings but also impacts the efficacy of content moderation on digital platforms. In the academic context, such a bias can lead to misleading conclusions about the nature and prevalence of toxicity within German-speaking online communities. For content moderators and platform administrators, this skewed assessment could result in unjust censorship or the inadvertent promotion of a hostile online environment for German-speaking users. For academic and journalistic research relying on the API, the implications are far-reaching, potentially affecting the perception and treatment of German digital discourse globally.

The problems identified in this paper highlight the inherent risks as researchers become increasingly dependent on proprietary models for data analysis and AI. The closed and proprietary APIs offered by major tech companies provide researchers with easy access to vast amounts of data and sophisticated analysis methods, which can otherwise be challenging to access. This ready availability is invaluable for studying complex social phenomena, trends, and behaviors in the digital age. However, the growing dependency on closed models also raises critical concerns. The opaque nature of these proprietary systems means that researchers are unable to scrutinize or understand the inner workings of the algorithms, and identify potential biases and limitations. The result is a substantial risk of systemic biases being embedded and perpetuated through their employment. This lack of transparency can furthermore inadvertently influence the di-

rection and conclusions of studies, as researchers must operate within the constraints and methodologies predefined by these tools. While proprietary APIs have opened new avenues for exploring social sciences, they also necessitate a cautious and critical approach to ensure the integrity of research and its reproducibility (Pozzobon et al. 2023) – rather than taking the technology firms by their word.

Limitations

There are some limitations in our work. We based our investigation on data from Twitter, which is not the most widely adopted social platform in many countries, and therefore not fully representative of the population. However, our robustness analyses with Wikipedia data suggest that the highlighted issues do not depend on user demographics or the specific context/topic of conversation. Besides, we only analyzed the Toxicity attribute of the API, and leave an investigation of other attributes (such as Insult or Profanity) for future research. Lastly, We did not evaluate whether this bias manifests in other LLM-based models for toxicity detection.

Data Availability

To ensure reproducibility, we provide access to the code and datasets employed in our analysis in a public repository available at: <https://github.com/gi-ux/Toxic-Bias-Perspective-API-Data>. For what concerns Dataset 1 and Dataset 2, to comply with Twitter’s Terms of Service, we can only provide tweet IDs, which can be employed (*hydrated*) to retrieve the data by means of X/Twitter API. We acknowledge that access to the `GET /2/tweets` endpoint for researchers is not as easy as in the past, and we encourage interested researchers to reach out to the authors in case they encounter difficulties in replicating the analyses. For what concerns Dataset 3, we provide the full data, as it is already publicly accessible. We notice that we do not attempt to de-identify or disclose any personal or sensitive data.

References

Adamic, L. A.; Lukose, R. M.; Puniyani, A. R.; and Huberman, B. A. 2001. Search in power-law networks. *Physical review E*, 64(4): 046135.

Czymara, C. S.; Dochow-Sondershaus, S.; Drouhot, L. G.; Simsek, M.; and Spörlein, C. 2023. Catalyst of hate? Ethnic insulting on YouTube in the aftermath of terror attacks in France, Germany and the United Kingdom 2014–2017. *Journal of Ethnic and Migration Studies*, 49(2): 535–553.

Ferrara, E.; Cresci, S.; and Luceri, L. 2020. Misinformation, manipulation, and abuse on social media in the era of COVID-19. *Journal of Computational Social Science*, 3: 271–277.

Friedl, P. 2023. Dis/similarities in the design and development of legal and algorithmic normative systems: The case of Perspective API. *Law, Innovation and Technology*, 15(1): 25–59.

Gargee, S.; Gopinath, P. B.; Kancharla, S. R. S.; Anand, C.; and Babu, A. S. 2022. Analyzing and addressing the difference in toxicity prediction between different comments

with same semantic meaning in Google’s Perspective API. In *The 7th International Conference on ICT for Sustainable Development (ICT4SD’22)*, 455–464.

Gehman, S.; Gururangan, S.; Sap, M.; Choi, Y.; and Smith, N. A. 2020. RealToxicityPrompts: Evaluating neural toxic degeneration in language models. In *The 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP’20)*, 3356–3369.

Ghosh Roy, S.; Narayan, U.; Raha, T.; Abid, Z.; and Varma, V. 2020. Leveraging multilingual transformers for hate speech detection. In *The 12th meeting of Forum for Information Retrieval Evaluation (FIRE’20)*.

Giovanni, M. D.; Pierri, F.; Torres-Lugo, C.; and Brambilla, M. 2022. VaccinEU: COVID-19 vaccine conversations on Twitter in French, German and Italian. *Proceedings of the International AAAI Conference on Web and Social Media (ICWSM’22)*.

Gligoric, K.; Cheng, M.; Zheng, L.; Durmus, E.; and Jurafsky, D. 2024. NLP Systems That Can’t Tell Use from Mention Censor Counterspeech, but Teaching the Distinction Helps. *NAACL 2024*.

Hoseini, M.; Melo, P.; Benevenuto, F.; Feldmann, A.; and Zannettou, S. 2023. On the globalization of the QAnon conspiracy theory through Telegram. In *Proceedings of the 15th ACM Web Science Conference (WebSci’23)*, 75–85.

Hua, Y.; Naaman, M.; and Ristenpart, T. 2020. Characterizing Twitter users who engage in adversarial interactions against political candidates. In *The 2020 CHI Conference on Human Factors in Computing Systems (CHI’20)*, 1–13.

Hua, Y.; Ristenpart, T.; and Naaman, M. 2020. Towards measuring adversarial Twitter interactions against candidates in the US midterm elections. In *Proceedings of the International AAAI Conference on Web and Social Media (ICWSM’20)*, volume 14, 272–282.

Krishnan, N.; Gu, J.; Tromble, R.; and Abroms, L. C. 2021. Research note: Examining how various social media platforms have responded to COVID-19 misinformation. *Harvard Kennedy School Misinformation Review*, 2(6): 1–25.

Kwak, H.; Lee, C.; Park, H.; and Moon, S. 2010. What is Twitter, a social network or a news media? In *Proceedings of the 19th international conference on World wide web*, 591–600.

Lasser, J.; Herderich, A.; Garland, J.; Aroyehun, S. T.; Garcia, D.; and Galesic, M. 2023. Collective moderation of hate, toxicity, and extremity in online discussions. *arXiv preprint:2303.00357*.

Lees, A.; Tran, V. Q.; Tay, Y.; Sorensen, J.; Gupta, J.; Metzler, D.; and Vasserman, L. 2022. A new generation of Perspective API: Efficient multilingual character-level transformers. In *The 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD’22)*, 3197–3207.

Mihaljević, H.; and Steffen, E. 2023. How toxic is antisemitism? Potentials and limitations of automated toxicity scoring for antisemitic online content. *arXiv preprint:2310.04465*.

Monroe, B. L.; Colaresi, M. P.; and Quinn, K. M. 2008. Fightin' words: Lexical feature selection and evaluation for identifying the content of political conflict. *Political Analysis*, 16(4): 372–403.

Pavlopoulos, J.; Thain, N.; Dixon, L.; and Androutsopoulos, I. 2019. ConvAI at SemEval-2019 Task 6: Offensive language identification and categorization with Perspective and BERT. In *Proceedings of the 13th International Workshop on Semantic Evaluation (SemEval'19)*, 571–576.

Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; et al. 2011. Scikit-learn: Machine learning in Python. *the Journal of machine Learning research*, 12: 2825–2830.

Pierrri, F. 2020. The diffusion of mainstream and disinformation news on Twitter: the case of Italy and France. In *Companion proceedings of the web conference 2020*, 617–622.

Pozzobon, L. A.; Ermis, B.; Lewis, P.; and Hooker, S. 2023. On the Challenges of Using Black-Box APIs for Toxicity Evaluation in Research. In *ICLR 2023 Workshop on Trustworthy and Reliable Large-Scale Machine Learning Models*.

Rieder, B.; and Skop, Y. 2021. The fabrics of machine moderation: Studying the technical, normative, and organizational structure of Perspective API. *Big Data & Society*, 8(2).

Rye, E.; Blackburn, J.; and Beverly, R. 2020. Reading in-between the lines: An analysis of Dissenter. In *Proceedings of the ACM Internet Measurement Conference (IMC'20)*, 133–146.

Smirnov, I.; Oprea, C.; and Strohmaier, M. 2023. Toxic comments are associated with reduced activity of volunteer editors on Wikipedia. *PNAS nexus*, 2.

Trujillo, A.; Fagni, T.; and Cresci, S. 2023. The DSA Transparency Database: Auditing Self-reported Moderation Actions by Social Media. *arXiv:2312.10269*.

Vogels, E. A. 2021. The state of online harassment. *Pew Research Center*, 13: 625.

Welbl, J.; Glaese, A.; Uesato, J.; Dathathri, S.; Mellor, J.; Hendricks, L. A.; Anderson, K.; Kohli, P.; Coppin, B.; and Huang, P.-S. 2021. Challenges in detoxifying language models. In *The 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP'21)*, 2447–2469.

Wich, M.; Gorniak, A.; Eder, T.; Bartmann, D.; Cakici, B. E.; and Groh, G. 2022. Introducing an abusive language classification framework for Telegram to investigate the German hater community. In *The 2022 International AAAI Conference on Web and Social Media (ICWSM'22)*, 1133–1144.

Paper Checklist

1. For most authors...

- (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? **Yes, because we use data**

and tools publicly available to researchers, and we do not attempt to de-identify or disclose private data.

- (b) Do your main claims in the abstract and introduction accurately reflect the paper's contributions and scope? **Yes.**
- (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? **Yes.**
- (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? **Yes.**
- (e) Did you describe the limitations of your work? **Yes.**
- (f) Did you discuss any potential negative societal impacts of your work? **No, because we envision positive implications of our findings as detailed in the last part of the Discussion.**
- (g) Did you discuss any potential misuse of your work? **NA**
- (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? **Yes.**
- (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? **Yes, and we also discuss ethical implications in a dedicated subsection of the Discussion.**

2. Additionally, if your study involves hypotheses testing...

- (a) Did you clearly state the assumptions underlying all theoretical results? **NA.**
- (b) Have you provided justifications for all theoretical results? **NA.**
- (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? **NA.**
- (d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? **NA.**
- (e) Did you address potential biases or limitations in your theoretical framework? **NA.**
- (f) Have you related your theoretical results to the existing literature in social science? **NA.**
- (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? **NA.**

3. Additionally, if you are including theoretical proofs...

- (a) Did you state the full set of assumptions of all theoretical results? **NA.**
- (b) Did you include complete proofs of all theoretical results? **NA.**

4. Additionally, if you ran machine learning experiments...

- (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **No.**
- (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **No.**

- (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **NA**.
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **No**.
 - (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? **No**.
 - (f) Do you discuss what is “the cost“ of misclassification and fault (in)tolerance? **No**.
5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets, **without compromising anonymity**...
- (a) If your work uses existing assets, did you cite the creators? **Yes**.
 - (b) Did you mention the license of the assets? **NA**.
 - (c) Did you include any new assets in the supplemental material or as a URL? **NA**.
 - (d) Did you discuss whether and how consent was obtained from people whose data you’re using/curating? **No**.
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **No**.
 - (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR (see ?)? **NA**.
 - (g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset (see ?)? **NA**.
6. Additionally, if you used crowdsourcing or conducted research with human subjects, **without compromising anonymity**...
- (a) Did you include the full text of instructions given to participants and screenshots? **NA**.
 - (b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? **NA**.
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? **NA**.
 - (d) Did you discuss how data is stored, shared, and de-identified? **NA**.

Ethical Statement

We envision positive ethical implications for this work, particularly concerning future research on online toxicity as well as on the policy design of online platforms. Our findings underscore the potential consequences of relying on black-box models and APIs, without a thorough understanding of their language biases. The identification of intrinsic biases within the multilingual Perspective API highlights the risk of inadvertently amplifying or suppressing certain voices in online spaces. The observed higher levels of toxicity assigned to German content compared to other languages raise concerns about the fairness and equity of ongoing content moderation processes, which could potentially lead to

disproportionate censorship or scrutiny of specific linguistic communities. This study urges a reevaluation of the ethical dimensions surrounding the deployment of such tools, emphasizing the need for transparency, accountability, and ongoing refinement to ensure that AI-driven content moderation aligns with principles of impartiality and cultural sensitivity, thereby mitigating the inadvertent amplification of biases in online discourse.

Appendix

Here we provide supplementary figures and results referenced and commented in the main text.

Figure 10 provides the distribution of toxicity scores for tweets shared in German-speaking countries from Dataset 1 separating them by the number of non-ASCII characters they contain.

Figure 11 provides the classification performance, computed with a 5-fold cross-validation, of three off-the-shelf classifiers in the binary task of predicting whether a tweet belongs to the Abnormal class described in the main text.

Figure 12 provides the distributions of toxicity scores for tweets containing (or not) one of the top 10 words differentiating the Abnormal class (in terms of Z-score). They are translated into English as follows:

- nicht: not
- die: the
- sie: they/she
- man: one
- ist: is
- zu: to
- sich: oneself
- wenn: when
- nur: only
- dass: that

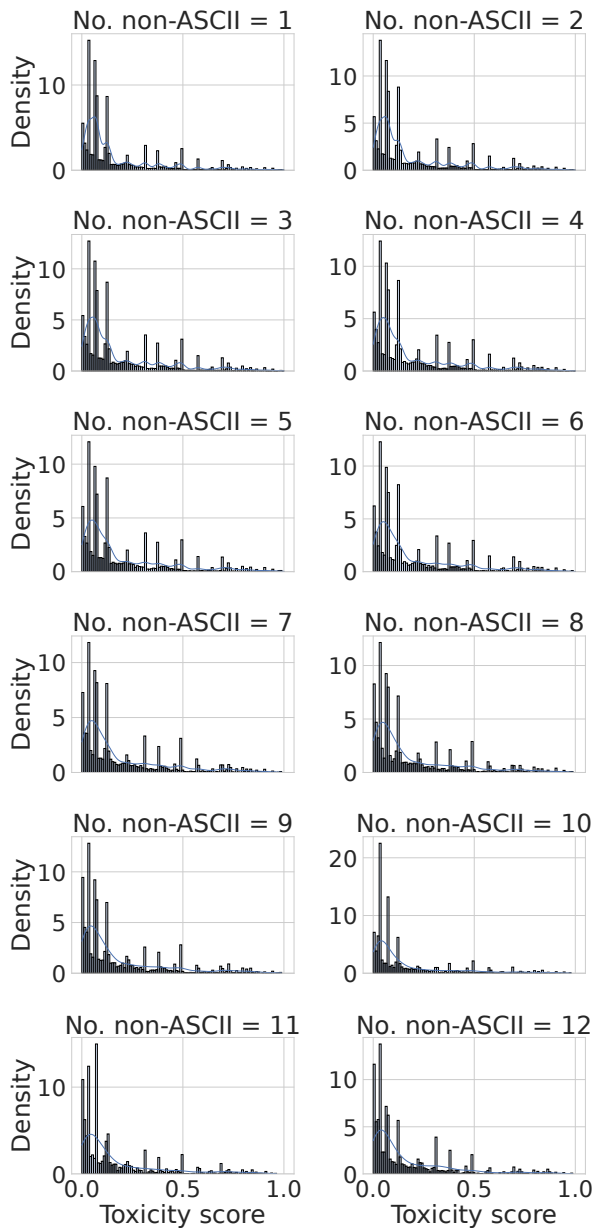


Figure 10: Distributions of toxicity scores for tweets from German-speaking countries in Dataset 1 based on the number of non-ASCII characters they contain.

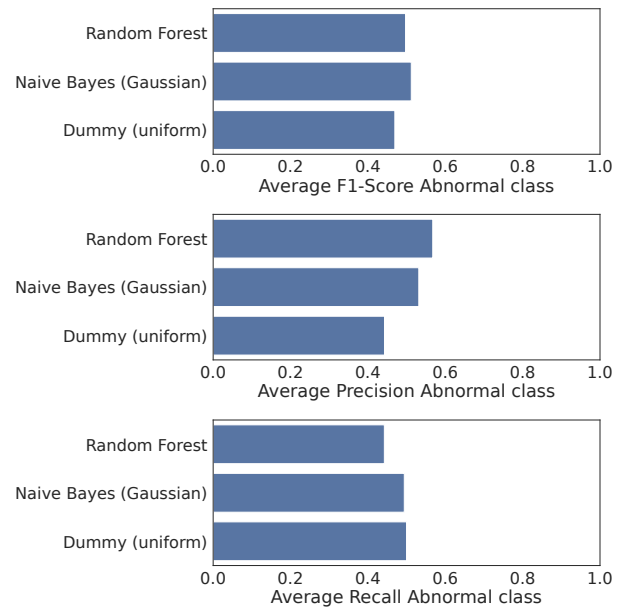


Figure 11: Performance evaluation of three classifiers in the binary task of predicting whether a tweet belongs to the class of Abnormal toxicity scores described in the main text. We show the mean value of each metric in a 5-fold cross-validation.

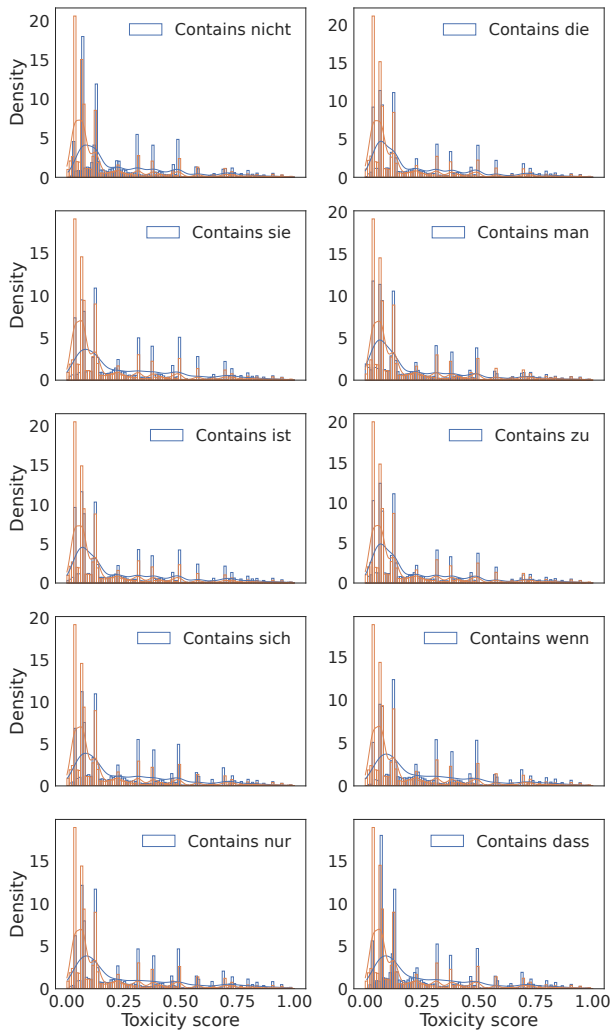


Figure 12: Distributions of toxicity scores for tweets from German-speaking countries in Dataset 1 that contain (or not) one of the top-10 words differentiating the Abnormal class.