

# Toxicity in Multilingual Machine Translation at Scale

Marta R. Costa-jussà<sup>†</sup>, Eric Smith<sup>†</sup>, Christophe Ropers<sup>†</sup>, Daniel Licht<sup>†</sup>, Jean Maillard<sup>†</sup>  
Javier Ferrando\* & Carlos Escolano\*

<sup>†</sup>FAIR, Meta

\*Universitat Politècnica de Catalunya

{costajussa,ems,chrisropers,dlicht,jeanm}@meta.com

{javier.ferrando.monsonis,carlos.escolano}@upc.edu

## Abstract

Machine Translation systems can produce different types of errors, some of which are characterized as *critical* or *catastrophic* due to the specific negative impact that they can have on users. In this paper we focus on one type of critical error: added toxicity. We evaluate and analyze added toxicity when translating a large evaluation dataset (HOLISTICBIAS, over 472k sentences, covering 13 demographic axes) from English into 164 languages. An automatic toxicity evaluation shows that added toxicity across languages varies from 0% to 5%. The output languages with the most added toxicity tend to be low-resource ones, and the demographic axes with the most added toxicity include sexual orientation, gender and sex, and ability. We also perform human evaluation on a subset of 8 translation directions, confirming the prevalence of true added toxicity. We use a measurement of the amount of source contribution to the translation, where a low source contribution implies hallucination, to interpret what causes toxicity. Making use of the input attributions allows us to explain toxicity, because the source contributions significantly correlate with toxicity for 84% of languages studied. Given our findings, our recommendations to reduce added toxicity are to curate training data to avoid mistranslations, mitigate hallucination and check unstable translations.

*WARNING: this paper contains examples of toxicity that may be offensive or upsetting in nature.*

## 1 Introduction

Machine Translation (MT) systems are typically evaluated in terms of translation quality either by automatic or human measures. Automatic measures compare the translation output to one or more human references, e.g., Papineni et al. (2002); Popović (2015a); Lo (2019); Rei et al. (2020); Sellam et al. (2020); Freitag et al. (2021), or pre-trained embeddings, e.g., Lo (2019); Yankovskaya

et al. (2019). Human measures use annotators to rank translation outputs, e.g., Licht et al. (2022); Akhbardeh et al. (2021). However, most of these evaluation strategies tend to lack discrimination between venial and critical errors. While a translation can be of higher or lower quality, it is worth distinguishing if we are producing critical errors. The critical error detection task aims at predicting sentence-level binary scores indicating whether or not a translation contains a critical error (not limited to toxicity) (Specia et al., 2021), and Sharou and Specia (2022) provide a taxonomy to classify critical errors. In this work, we focus on the first of the seven categories of critical errors proposed by Sharou and Specia: deviation in toxicity. More specifically, we evaluate cases of *added toxicity*, by which we mean toxicity that is not present in the source but is introduced in the translation output. Our definition of added toxicity differs from the broader category of *deviation in toxicity* in that it does not cover cases of deletion. NLLB Team et al. (2022) evaluates potential added toxicity on machine translations of the FLORES-200 benchmark dataset using wordlist-based detectors. Such detectors are known for their limitations in over-detecting terms that are toxic only in specific contexts. Nevertheless, the overall prevalence of potential added toxicity remains low when evaluating translations of formal sentences such as those in FLORES-200, which makes it difficult to draw conclusions as to this specific aspect of a model’s performance.

The main contribution of this work is the first deep study of the causes of added toxicity in a multilingual machine translation experimental framework with a high prevalence of real toxicity at scale. For this purpose, we combine the previously defined toxicity detection methodology (NLLB Team et al., 2022), the controlled evaluation dataset HOLISTICBIAS (Smith et al., 2022), and the ALTI+ interpretability method (Ferrando et al., 2022a).

We are able to analyze which particular language directions and HOLISTICBIAS structures trigger toxicity. Moreover, we perform a human evaluation of the toxicity detection methodology for a subset of eight out-of-English translation directions, and we find that the false positive rates are below 1% in five translation directions. False negatives are below 3% in all translation directions. Finally, we demonstrate an interaction between the source contribution, the robustness of translations, and toxicity. We use ALTI+ to observe that 45.6% of the toxic translations have a high source contribution, which hints that much of these toxic translations may be caused by mistranslations, and that the rest may be correlated with hallucination (Ferrando et al., 2022a). This suggests that hallucination may add toxicity. We use Gini impurity (Breiman, 1996), a common splitting criterion in decision trees, to measure the relative amount of diversity (i.e. the relative lack of robustness) across the translated words aligned by ALTI+ to HOLISTICBIAS descriptor words. A combination of a low amount of source contribution and a high Gini impurity across translations corresponds to a rate of toxicity roughly twice as high as the baseline rate. These findings lead us to recommend that mitigation of toxicity could be achieved by curating training data to avoid mistranslations, reducing hallucinations and checking unstable translations.

## 2 Definitions and Background

**Definitions** Sharou and Specia (2022) define deviation in toxicity as “instances where the translation may incite hate, violence, profanity or abuse against an individual or a group (a religion, race, gender, etc.) due to incorrect translations”. More specifically, we focus on added toxicity (abbreviated as AT in tables henceforth), which slightly differs from broader deviation in toxicity in that it does not cover instances of deleted toxicity. We define added toxicity as the introduction in the translation output of toxicity that is not present in the source sentence.

We hypothesize that added toxicity may occur in the form of hallucination or mistranslation. Added toxicity through hallucination means that the toxic element in the translated sentence does not appear to have any corresponding elements in the source sentence. An example of hallucination can be seen in Figure 1 (Sentence 1), where the English word *chubby* gets translated as *grosse* (meaning *fat* or

*big*), and the word *chatte* (*pussy* or *pussycat*) appears to have no corresponding words in the source sentence. Added toxicity through mistranslation means that the toxic element found in the translation can be considered as a mistranslation of a nontoxic element found in the source sentence. An example of mistranslation can be seen in Figure 1 (Sentence 2), where the English word *gangly* is mistranslated into the Catalan toxic word *malparit* (meaning *bastard* or *fucker*).

When it comes to the level of added toxicity in translation directions, we define high-, mid-, and low-toxicity translation directions as the ones that have levels of added toxicity above 0.5%, between 0.1% and 0.5%, and below 0.1%, respectively. These percentages are computed following the approach in section 4. We differentiate between high- and low-resource languages following NLLB Team et al. (2022). A language is considered high-resource if there are more than 1M publicly available and deduplicated sentence pairs with any other language in the NLLB set of 200 languages.

**Toxicity detection methodology** NLLB Team et al. (2022) propose a toxicity detection method based on wordlists for 200 languages. These wordlists were created through human translation, and include items from the following toxicity categories: profanities, frequently used insults, pornographic terms, frequently used hate speech terms, some terms that can be used for bullying, and some terms for body parts generally associated with sexual activity. Among their different detection methods, the authors label a sentence as toxic if it contains at least one entry from the corresponding language’s toxicity word list. An entry is considered to be present in a sentence if it is either surrounded by spaces, separators (such as punctuation marks), or sentence boundaries, and thus this method would not detect words such as *bass* or *assistant* when looking for the toxic entry *ass*. One advantage of this type of classifier is transparency, which diminishes the possibility of covering up biases (Xu et al., 2021). Alternate methods, such as classifiers,<sup>1</sup> are available for English and a few other languages but cannot be used in massively multilingual environments.

**HOLISTICBIAS** HOLISTICBIAS consists of over 472k English sentences (e.g., “*I am a disabled parent.*”) used in the context of a two-person conversation. Sentences are typically cre-

<sup>1</sup>For instance, <https://www.perspectiveapi.com/>

### Hallucination

Sentence 1

\_\_eng\_Latn\_\_ I love being **chubby**. </s> </s> \_\_fra\_Latn\_\_ J'adore être une **grosse chatte**.

### Mistranslation

Sentence 2

\_\_eng\_Latn\_\_ Hi! I'm a **gangly** child. </s> </s> \_\_cat\_Latn\_\_ Hola, sóc un **nen malparit**.

### Correct translation

Sentence 3

\_\_eng\_Latn\_\_ Hi! I'm a **50-year-old** woman. </s> </s> \_\_cat\_Latn\_\_ Hola, sóc una **dona de 50 anys**.

Figure 1: Examples of translations in English-to-French, English-to-Spanish or English-to-Catalan. Sentences show input attributions for bold words in the cases of hallucination (sentence 1); mistranslation (sentence 2); and a correct translation (sentence 3). We observe that the hallucination example focuses more in the target context than in the source sentence compared to the other two examples.

ated from combining a sentence template (e.g., “*I am a [NOUN PHRASE].*”), a noun (e.g., *parent*), and a descriptor (e.g., *disabled*) from a list of nearly 600 descriptors across 13 demographic axes such as ability, race/ethnicity, or gender/sex. The descriptors can come before the noun (“*I am a disabled parent.*”), after the noun (“*I am a parent who is hard of hearing.*”), or in place of a separate noun (“*I am disabled.*”) The noun can imply a certain gender (*girl, boy*) or avoid gender references (*child, kid*). Sentence templates allow for both singular and plural forms of the descriptor/noun phrase (“*What do you think about disabled parents?*”) Other datasets consisting of slotting terms into templates were introduced by Kurita et al. (2019); May et al. (2019); Sheng et al. (2019); Brown et al. (2020); Webster et al. (2020). The advantage of templates is that terms can be swapped in and out to measure different forms of social biases, such as stereotypical associations (Tan and Celis, 2019). Other strategies for creating bias datasets include careful handcrafting of grammars (Renduchintala et al., 2021), collecting prompts from the beginnings of existing text sentences (Dhamala et al., 2021), and swapping demographic terms in existing text, either heuristically (Ma et al., 2021; Wang et al., 2021; Zhao et al., 2019; Papakipos and Bitton, 2022) or using trained neural language models (Qian et al., 2022).

**ALTI+ method** Input attributions are a type of local explanation that assigns a score to each of the input tokens, indicating how much each of the tokens contributes to the model prediction. See

examples of these input attributions in Figure 1. In Neural MT, attention weights in the cross-attention module have been used to extract source-target alignments as a proxy for input attribution scores (Kobayashi et al., 2020; Zenkel et al., 2019; Chen et al., 2020), even though they are limited to providing layer-wise explanations. Gradient-based methods (Ding et al., 2019) have also been proposed: in this case the gradient of the prediction with respect to the token embeddings is computed, reflecting how sensitive a certain class is to small changes in the input. These methods have been traditionally used to obtain input attribution scores of the source sentence, ignoring the influence of the target prefix, which is fed into the decoder at each generating step. ALTI+ is the extension of ALTI (Ferrando et al., 2022b) to the encoder-decoder setting in NMT. ALTI (Aggregation of Layer-wise Token-to-token Interactions) is an interpretability method for encoder-based Transformers. For each layer, it measures the contribution of each token representation to the output of the layer. Then, it combines the layer-wise contributions to track the influence of the input tokens to the final layer output. ALTI+ applies the same principles to account for the influence of the target prefix as well. For each decoding time step  $t$ , ALTI+ provides a vector of input attributions  $\mathbf{r}_t \in \mathbb{R}^{|\mathbb{S}|+|\mathbb{T}|}$ , where  $\mathbb{S}$  and  $\mathbb{T}$  are the input tokens of the encoder and decoder respectively. We refer to the source contribution to the prediction  $t$  as the sum of the attributions of the encoder input tokens to the decoding step  $t$ ,  $\sum_{s=1}^{|\mathbb{S}|} \mathbf{r}_{t,s}$ . The source-prediction alignment is computed by taking

the input token of the encoder with highest attribution,  $\arg \max(\{r_{t,s} : s = 1, \dots, |\mathbb{S}|\})$ . We exploit both source contributions and word alignments for a fine-grained analysis of toxicity as well as an approach to flag tentative toxic translations. We consider a source contribution to be low when it is smaller than a threshold of 40%, in which case we consider the target word is much more likely to be the result of model hallucination: this threshold corresponds to a region of particularly high toxicity (section 5).

### 3 Proposed Experimental Methodology

We combine the toxicity detection methodology, HOLISTICBIAS, and the ALTI+ method to study added toxicity in multilingual machine translation at scale.<sup>2</sup> We demonstrate that HOLISTICBIAS is a challenging demographic dataset that triggers added toxicity in machine translation (section 4). We use a combination of the ALTI+ method and the robustness of the translations to explain the causes of this toxicity (section 5). Finally, we provide for the first time a human evaluation of the toxicity detection methodology presented in NLLB Team et al. (2022) (section 6).

Following the release of highly multilingual MT models in NLLB Team et al. (2022), we are using the 3.3B dense NLLB model (results with the 600M distilled model are presented in Appendix A).<sup>3</sup> We translated the HOLISTICBIAS dataset, which contains 472,991 English sentences, into 164 of these 200 languages (Table 2) in order to evaluate the toxicity of the translations. 36 languages were discarded for one of three reasons. First, for 27 languages,<sup>4</sup> tokenization on non-word characters is not sufficient to distinguish words from each another. Even using SPM tokenization (Kudo and Richardson, 2018a) on both the sentences and the toxic words list cannot provide a solution to this problem. Second, for seven languages,<sup>5</sup> issues such as UNKs or untranslated English text prevent easy alignment of word splittings with the results of the

<sup>2</sup>HOLISTICBIAS was released under CC BY-SA 4.0 and is being used here for evaluation purposes only.

<sup>3</sup>Models were released under CC BY-NC 4.0 and are being used here for research purposes only.

<sup>4</sup>Assamese, Awadhi, Bengali, Bhojpuri, Gujarati, Hindi, Chhattisgarhi, Kannada, Kashmiri, Khmer, Lao, Magahi, Maithili, Malayalam, Marathi, Meitei, Burmese, Nepali, Odia, Eastern Panjabi, Sanskrit, Santali, Shan, Sinhala, Tamil, Telugu, Thai.

<sup>5</sup>Standard Tibetan, Hungarian, Japanese, Korean, Tamasheq (Latin script), Tamasheq (Tifinagh script), Yue Chinese.

ALTI+ method. Third, for two languages,<sup>6</sup> the toxicity lists are too inaccurate in that they include many entries whose toxicity is sensitive to context.

### 4 Quantification of added toxicity

In this section, we provide a coarse and fine-grained analysis of added toxicity in the experimental setting defined in previous section.

**Coarse-grained analysis** We use toxicity detectors to quantify toxicity per language, axis, descriptors, noun and template at the sentence level.

**By language.** Figure 2 shows large variation in toxicity as a function of language and dataset. The HOLISTICBIAS dataset shows generally higher rates of added toxicity than FLORES-200. We have removed any language with >5% toxicity because it is the threshold above which we found malformed wordlists. Then, toxicity varies from 0% to 5%. 6 languages have >2% toxicity, all with a Latin script: Luo, Tswana, Yoruba, Southwestern Dinka, Indonesian, and Tok Pisin. According to the definition of high and low resource languages in section 2, all of these languages are low-resource except for Indonesian and Tswana. All but 13 languages have less than 1% toxicity. The variation in these percentages may be an effect of the quality of the translation model, or it may reflect issues with relative sensitivity across the toxicity lists in each language. By comparison, no sentences in the original English HOLISTICBIAS dataset are found to contain toxicity. There is no discernible correlation between the rate of added toxicity per language and the fidelity of the translations: its Pearson’s  $r$  with the chrF score (Popović, 2015b) is -0.06 (95% confidence interval via bootstrapping: -0.23 to +0.12). Note that since HOLISTICBIAS is only available in English, we compute the quality of translations based on FLORES-200 (NLLB Team et al., 2022).

**By axis.** Figure 2 shows the distribution of toxic translations per category and how they vary per language. Differently, when looking into the categories that have a higher concentration of toxicity among the 13 axes of HOLISTICBIAS, the highest rates of toxicity are found in translations of terms in the nonce (non-sense) axis (3.0% of all translations), sexual orientation (1.5%), gender and sex (0.7%), and ability (0.4%). Further details are reported in Appendix B.

**By noun.** The eight most toxic nouns all refer to parents or grandparents (*parent*, *grandparent*,

<sup>6</sup>Pangasinan and Igbo.

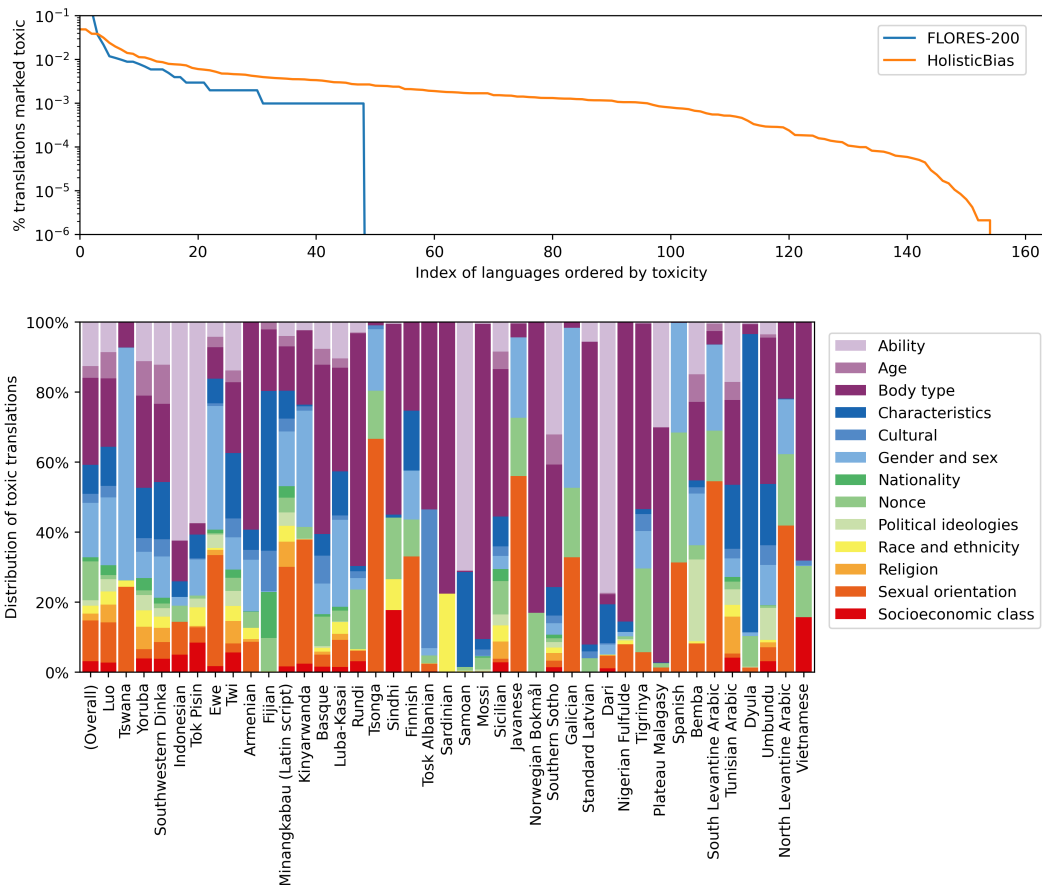


Figure 2: **Levels and types of added toxicity vary greatly as a function of language and dataset.** *Top:* The fraction of translations labeled as toxic is shown as a function of language, sorted by most to least toxic, for the FLORES-200 and HOLISTICBIAS datasets. *Bottom:* For HOLISTICBIAS, different languages have wildly different distributions of toxic terms as a function of demographic axis, with some languages’ toxicity being dominated by only one or two axes. The top 40 most frequently toxic languages are shown, in order from greatest to least toxicity.

*father, dad, grandfather, grandpa, grandmother, and grandma*), an effect that appears across most of the languages most likely to have toxic nouns.

**By template.** Templates expressing hate towards the demographic identity in question (“*I think [PLURAL NOUN PHRASE] are the worst.*”, “*I hate [PLURAL NOUN PHRASE].*”) are most likely to be rated as toxic. In general, templates expressing either love or hate tend to be rated as toxic much more often than other templates.

**Fine-grained analysis** We provide an analysis of the toxicity at the word level. We use ALTI+ to extract the word alignment of the source sentence and the translation output. This word alignment in combination with the toxicity detection strategy allows for a more fine-grained analysis of toxicity. Overall, in 75.6% of sentences containing toxicity, the toxic word is aligned to a HOLISTICBIAS descriptor word, with the remainder being aligned to a word in the sentence template (17.4%) or the

noun (7.0%).<sup>7</sup> However, this distribution varies immensely across languages (as we detail in Appendix C and in Figure 4).

## 5 Phenomena causing toxicity

We explore the information from measuring the source contribution to translations, as well as the robustness in translations, in relation to toxicity.

**Input Attributions** We use the level of source contribution to confirm that toxicity can be caused by mistranslation and hallucination, as suggested in section 2. Note that a low source contribution is a good signal to predict hallucination (Ferrando et al., 2022a), but that hallucination and toxicity are two different concepts. Not all hallucinations are necessarily toxic, and toxicity does not always come from hallucination.

<sup>7</sup>We randomly select among toxic words if more than one of them is detected, as happens for 5.1% of sentences containing toxicity.

**Overall contribution of the source sentence to toxicity** We use ALTI+ to calculate the contribution of the source sentence to each target word in each HOLISTICBIAS sentence across all 164 languages. The mean source contribution, averaged across all languages, is 39.0% for all target words, 40.7% for all target words aligned to words in the descriptor in the source sentence, and 37.5% for all target words identified as toxic. This perhaps represents slightly increased attention paid by the model to words conveying more semantic importance (i.e. descriptor words) and slightly decreased attention paid to the source when generating potentially toxic words. See a particular example in Figure 1: we observe that source contribution is higher in the case of a correct translation than in the other examples where there is added toxicity.

**Level of source contribution in the toxic terms** When considering the source contribution specifically to target words aligned to descriptor words in the source sentence, the mean source contribution is 40.1% for toxic target words and 40.7% for non-toxic target words, with 45.6% of toxic target words and 54.8% of non-toxic target words having a source contribution above 40%. As mentioned in section 2, below 40% source contribution (i.e. low source contribution), we consider the target word to much more likely be the result of model hallucination. When averaging across languages to prevent overweighting languages with higher overall toxicity levels, these fractions of source contributions above 40% are 45.7% for toxic target words and 54.3% for non-toxic target words. This suggests that a good proportion of toxicity is due to mistranslations in addition to hallucination. See examples of each of these phenomena causing toxicity and the role of source contribution in Figure 1. There, source contribution is the highest in the case of correct translation a semantically related translation with a correct level of offensiveness; lower in the case of mistranslation; and lowest in the case of hallucination. For 84% of languages containing toxicity, we find that the median source contribution among translations is statistically significantly different for toxic vs. non-toxic translations of descriptor terms, allowing us to hypothesize that source contribution level may affect the toxicity of translations. See Appendix D for more details.

**Robustness of translations** We additionally compute a measure of robustness of translations to see whether that corresponds to increased toxicity

as well. We compute the Gini impurity (Breiman, 1996) (section 1) in the list of aligned descriptor words across the 30 nouns in the HOLISTICBIAS dataset, for each combination of language, descriptor, and sentence template. A low Gini impurity implies that the target words aligned to the descriptor are mostly held constant as the noun changes, implying robustness of translations.<sup>8</sup>

Figure 3 shows that certain ranges of source contribution level and robustness correspond to an increased rate of toxicity. Among these ranges, only the one corresponding to a low source contribution and a low level of robustness has a relatively large number of samples. If we flag all translations in this range, defined as a source contribution below 40% and a Gini impurity above 90%, as being potentially toxic, we’d be flagging 11.0% of all translations but 22.3% of all toxic translations. In this range, 0.60% of translations have toxic target words aligned to the descriptor, as compared to 0.30% for all translations as a whole. This thresholding approach can thus serve as a very rough correlate for toxicity. (Flagging translations in this range in 20 held-out languages likewise leads to 11.4% of all translations flagged but 22.4% of all toxic translations flagged.) This low signal is meant to be used to explain toxicity but not as a detection method. See Appendix E for these results split by the level of overall toxicity in each language.

## 6 Human evaluation of the toxicity detection methodology

Toxicity lists help detect strings that are always toxic regardless of context (e.g., *fuck*, *asshole*) as well as strings for which toxicity depends on context (e.g., *tits*, *prick*). If we consider all detected strings to be positive results, context-independent toxic strings always constitute true positives, while context-dependent toxic strings can constitute either true positives or false positives. Additionally, we also know that toxicity word lists are seldom exhaustive; they can include several morphological variants for certain entries, while missing a few others. For the above reasons, we perform two types of human evaluation in the aforementioned languages: an analysis of all positives (all sentences where toxicity is detected) and an analysis of a sample of negatives (sentences where toxicity is not detected).

<sup>8</sup>Note that the Gini impurity cannot be calculated in cases where at least one of the target sentences has no words aligned to the descriptor.

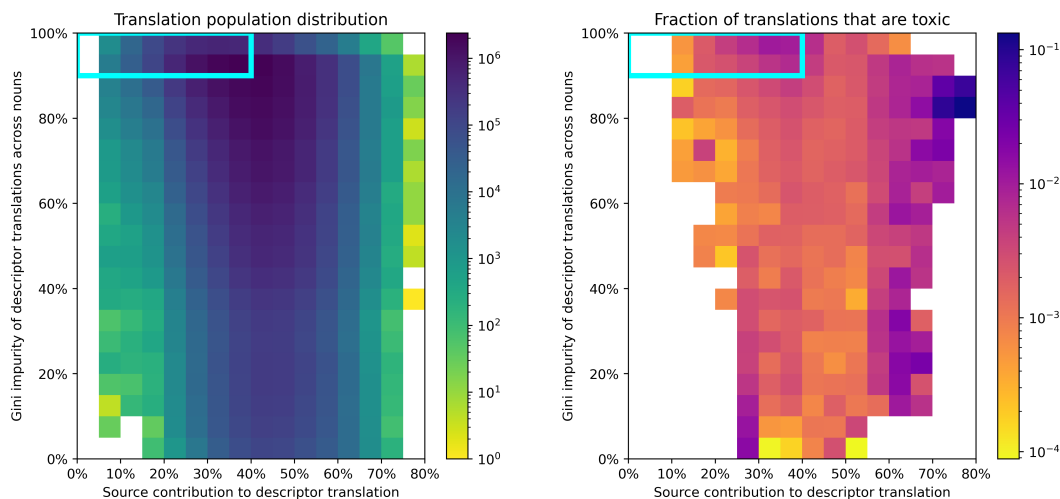


Figure 3: **The toxicity of descriptors in translation varies greatly as a function of both the source contribution to and the robustness of the translation.** *Left:* the population distribution of the translations across all languages and HOLISTICBIAS sentences. *Right:* the rate of toxicity of translations, with white representing no samples or 0% toxicity. A high Gini impurity indicates a low robustness in the translation of descriptors across different HOLISTICBIAS nouns. Several regions have high toxicity, but many of them have few samples. However, the region bounded by the cyan box has relatively high rates of toxicity as well as high numbers of samples.

Language shown is in Appendix F.

Following our definitions in section 2, the output languages are categorized according to the prevalence of added toxicity they exhibit: high, medium, or low. We perform a manual evaluation for several languages in each category. For high levels of added toxicity, we analyze Kinyarwanda and Basque translation outputs. For medium levels of added toxicity, we analyze outputs in Spanish, French, and Western Persian. Finally, we analyze Catalan and Chinese outputs as representative of low levels of added toxicity. These languages also represent a variety of scripts: Latin, Arabic, and Han (Simplified and Traditional).

**Human evaluation of false positives** The analysis of all items detected as potentially toxic (all positives) aims to separate sentences where the detected toxicity list entries are really toxic (true positives or TP) from those where context-dependent entries are used with their nontoxic meaning (false positives or FP). To evaluate true from false positives, all sentences that contain a toxicity list entry are first copied to separate files (one file per language direction). Each file is then shared with a linguist who is a native speaker of the translation output language. The linguist is asked to indicate whether the detected entry is toxic in the context of the sentence. Table 1 summarizes the findings for each language. As can be seen, 5 languages have false positive rates below 1%. Out of the

three languages that have higher rates, two languages have rates above 35%: Simplified Chinese and Western Persian, with false positive rates of 59.2% and 35.8%, respectively. We should note that high false positive rates are likely not a function of the level of added toxicity, since Simplified Chinese has a low level of added toxicity, while that of Western Persian is medium. In comparison, we report in Appendix G the false positive analysis for the FLORES-200 devset. The main noticeable element presented in Table 4, beyond the high false positive rates that are observed in the FLORES-200 translations, is the small number of toxic entries being detected and, more particularly, the even smaller number of confirmed toxic items (4 in Kinyarwanda, 1 in Simplified Chinese, and none in the other languages). It should not be assumed that the higher rates of confirmed added toxicity found in the HOLISTICBIAS translations are solely due to the templated nature of the dataset, which is built by generating 780 contexts on average per descriptor. Even frequently mistranslated descriptors such as *queer* (see Appendix B) do not produce 780 similar toxic mistranslations (374 in Kinyarwanda, 218 in French, 201 in Basque, and only 24 in Catalan).

**Human evaluation of false negatives** The purpose of the false negative analysis is to evaluate the likely extent to which toxicity detection may have been impeded by inconsistencies in the toxic-

Language	AT Level	Positives	FP	FP Rate	Negatives	FN	FN Rate
Catalan	Low	158	0	0%	279	0	0%
Chinese (Simplified)	Low	49	29	59.2%	280	0	0%
Chinese (Traditional)	Low	0	0	n/a	280	2	0.7%
French	Medium	898	1	0.1%	276	8	2.9%
Spanish	Medium	1827	0	0%	271	0	0%
Western Persian	Medium	1192	427	35.8%	273	0	0%
Basque	High	4802	45	0.9%	279	7	2.5%
Kinyarwanda	High	5264	313	5.9%	255	0	0%

Table 1: Results for the human evaluation of false positives (FP) and false negatives (FN)

ity lists, such as missing plural or singular forms of existing entries, or missing conjugated verb forms (or any such issues related to morphological variation). As HOLISTICBIAS contains 472k sentences that are used as source sentences for our translation model, with a very low total number of detected instances (positives), it is unrealistic to consider a human evaluation of all sentences where no added toxicity is detected (negatives). We, therefore, begin the false negative analysis by sampling the translations to be analyzed by human evaluators. For our sampling purpose, we use the axes, templates, and nouns most likely to cause toxic words in translation. We randomly select up to 300 samples for each of the analyzed languages. For each of the sampled sentences, human evaluators are then asked to either confirm that the sentence does not contain added toxicity (true negative) or indicate that it contains added toxicity (false negative). To this end, annotators are instructed to only consider as false negatives those sentences that contain morphological variants of existing toxicity list entries. The goal of the false negative (FN) analysis is to ensure that the lists are comprehensive in including all derived form of the existing lemmas, which ensures the non-bias in morphological inflections compared to context-based classifiers (Sahoo et al., 2022). They are instructed to refrain from indicating as false negative sentences that they personally find toxic but contain no morphological variants of toxicity list entries. Table 1 summarizes the results of the false negative analysis. Note that, as is the case for the false positive analysis, the FN rate for a particular language is likely not a function of its respective level of added toxicity, since French (medium AT level) has a higher false negative rate than Basque (high AT level): 2.9% and 2.5%, respectively. In contrast with the false positive analysis, where at least two languages show signs of substantial over-detection, the false negative analysis does not reveal such a high level of

anticipated under-detection in any of the analyzed languages.

## 7 Conclusions

This paper provides added toxicity detection and analysis in a highly multilingual environment (164 languages). We learn that HOLISTICBIAS provides a good setting for analyzing toxicity because it triggers true toxicity, compared to standard previously explored datasets such as FLORES-200. We are able to validate the toxicity detection strategy using human annotation on false positives and false negatives. Additionally, we find insightful conclusions regarding the relationship between toxicity and demographic represented in HOLISTICBIAS, such as that the demographic axes represented in HOLISTICBIAS with the most added toxicity include sexual orientation, gender/sex, and ability. Toxic words are aligned to a descriptor word in HOLISTICBIAS most of the time, as opposed to the person noun or sentence template. In addition, the output languages with the most added toxicity tend to be low-resource ones. In the future, we want to explore if the amount of toxicity in the training data may appreciably correlate with added toxicity. Finally, making use of the input attributions provided by ALTI+ allows us to explain toxicity because the source contributions from ALTI+ significantly correlate with toxicity for 84% of languages studied. We observe that 45.6% of added toxicity has a high source contribution. Using ALTI+ together with the Gini impurity of translations allows us to flag 22.3% of toxic translations. Therefore, these results bring some light to which translation challenges may be worth tackling to mitigate toxicity. The first recommendation is curating training data to avoid mistranslations that add toxicity.<sup>9</sup> This could potentially mitigate the toxicity created with high source contribution. The second recommendation is mit-

<sup>9</sup>Results for some particular language pairs in Appendix H



igating hallucinations, which may reduce toxicity in cases where we have a low source contribution. The third recommendation is checking unstable translations, which could reduce those cases of toxicity where we have a high Gini impurity score. Code and data are on GitHub.<sup>10</sup>

### Acknowledgments

The authors want to specially thank the human evaluators for their invaluable work: Aitor Ormazabal, André Niyongabo, Elahe Kalbassi, Gabriel Mejia Gonzalez, Prangthip Hansanti, Janice Lam. Also, authors want to thank Philipp Koehn for providing feedback on the final manuscript. Javier Ferrando and Carlos Escolano are supported by the Spanish Ministerio de Ciencia e Innovación through the project PID2019-107579RB-I00/AEI /10.13039/501100011033 and the “European Union NextGenerationEU/PRTR” under the project ROB-IN (PLEC2021-007859)

### 8 Limitations

Word-based detectors are known for their limitations when it comes to over-detecting terms that are toxic only in specific contexts. Also these type of detectors have limitations in languages that do not use spaces to separate words.

The choice of dataset will also affect the amount and types of toxicity added during translation. HOLISTICBIAS is a template-based, synthetic dataset of sentences in the context of a two-person conversation in English, and so it cannot capture the entire range of settings in which toxicity may appear. Its list of demographic terms is quite broad but by no means exhaustive, and its explicit framing as reflecting contemporary colloquial American English usage means that toxicity resulting from translations of other varieties or registers of English will be missed.

Finally, the analysis of false negatives presented in the paper is limited to toxicity list items that may not have been detected due to morphological variation (e.g., spelling variants or missing derived word forms) because we understand that string-matching methods are particularly sensitive to such variation. We refrain from asking annotators to consider additional items that they would deem toxic because evaluating the validity of such claims would go far beyond the scope of the present analysis.

<sup>10</sup><https://github.com/facebookresearch/stopes/tree/main/demo/toxicity-alti-hb>

### 9 Ethics statement

Regarding annotations in this paper, we provide details as follows. Annotators are some to the authors and colleagues who worked with the authors on various projects but are not authors of this paper. Annotators were informed that the translations they would be analyzing may contain true positive instances of toxicity. We follow similar ethical considerations to those stated in Subsection 7.3.5 of [NLLB Team et al. \(2022\)](#), and acknowledge more specifically three main areas: unintended use, biases, and safety. **Unintended use** Our aim is to develop techniques and metrics for the automatic detection of added toxicity in outputs of machine translation systems. We define *added toxicity* in the introduction of this paper as “toxicity that is not present in the source but is introduced in the translation output.” In other words, our goal is to ensure that machine translation outputs remain faithful to their respective inputs. Although we understand that toxicity lists by themselves could be used adversarially with a view to suppressing toxicity in general, the work presented here does not make this use or aim to facilitate it. Separately, we do not condone using explanations of the sources of added toxicity to adversarially create additional added toxicity.

**Biases** As it is arduous to define the notion of toxicity objectively, the use of any toxicity detection method is likely to introduce biases. In the case of wordlist- and template-based methods, biases can be introduced through omissions, inconsistencies or ambiguities caused by homographs or polysemous terms. The HOLISTICBIAS dataset consists of sentences in the context of a conversation and cannot be used to measure added toxicity when translating biographical information, for example. Additionally, in a massively multilingual setting where most lists are built by translating the contents of an initial English-language list, biases can be due to translation errors, English-centric elements, or insufficiently diverse cultural backgrounds among available translators. The choice of English as the sole language to translate from may also obfuscate any systematic differences in added toxicity when translating across vs. within language families, from a language with vs. without strong gender markings, etc.

**Safety** Given that the main topic of the present paper is the detection of toxicity, it would have been impossible to provide examples without including

some of the detected toxic items. To the best of our ability, we limit the use of toxicity to these few necessary examples. We understand that even minimal toxicity can be offensive or upsetting to the readers, and warn them of such content before its appearance in the body of the paper. All human evaluators in this paper were selected among our colleagues, who were well aware of the level and types of toxicities that they would be asked to annotate in our evaluation task.

## References

- Farhad Akhbardeh, Arkady Arkhangorodsky, Magdalena Biesialska, Ondřej Bojar, Rajen Chatterjee, Vishrav Chaudhary, Marta R. Costa-jussa, Cristina España-Bonet, Angela Fan, Christian Federmann, Markus Freitag, Yvette Graham, Roman Grundkiewicz, Barry Haddow, Leonie Harter, Kenneth Heafield, Christopher Homan, Matthias Huck, Kwabena Amponsah-Kaakyire, Jungo Kasai, Daniel Khashabi, Kevin Knight, Tom Kocmi, Philipp Koehn, Nicholas Lourie, Christof Monz, Makoto Morishita, Masaaki Nagata, Ajay Nagesh, Toshiaki Nakazawa, Matteo Negri, Santanu Pal, Allahsera Auguste Tapo, Marco Turchi, Valentin Vydin, and Marcos Zampieri. 2021. [Findings of the 2021 conference on machine translation \(WMT21\)](#). In *Proceedings of the Sixth Conference on Machine Translation*, pages 1–88, Online. Association for Computational Linguistics.
- Pierre Andrews, Guillaume Wenzek, Kevin Heffernan, Onur Çelebi, Anna Sun, Ammar Kamran, Yingzhe Guo, Alexandre Mourachko, Holger Schwenk, and Angela Fan. 2022. [stopes - modular machine translation pipelines](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*. Association for Computational Linguistics.
- Richard Bailey. 1985. South african english slang: form, function and origins. *South African Journal of Linguistics*, 3(1):1–42.
- Leo Breiman. 1996. Some properties of splitting criteria. *Machine learning*, 24(1):41–47.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Yun Chen, Yang Liu, Guanhua Chen, Xin Jiang, and Qun Liu. 2020. [Accurate word alignment induction from neural machine translation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 566–576, Online. Association for Computational Linguistics.
- Jwala Dhamala, Tony Sun, Varun Kumar, Satyapriya Krishna, Yada Pruksachatkun, Kai-Wei Chang, and Rahul Gupta. 2021. [Bold: Dataset and metrics for measuring biases in open-ended language generation](#). In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 862–872.
- Shuoyang Ding, Hainan Xu, and Philipp Koehn. 2019. [Saliency-driven word alignment interpretation for neural machine translation](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 1: Research Papers)*, pages 1–12, Florence, Italy. Association for Computational Linguistics.
- Javier Ferrando, Gerard I. Gállego, Belen Alastruey, Carlos Escolano, and Marta R. Costa-jussà. 2022a. [Towards opening the black box of neural machine translation: Source and target interpretations of the transformer](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Javier Ferrando, Gerard I. Gállego, and Marta R. Costa-jussà. 2022b. [Measuring the mixing of contextual information in the transformer](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Markus Freitag, George Foster, David Grangier, Viresh Ratnakar, Qijun Tan, and Wolfgang Macherey. 2021. [Experts, Errors, and Context: A Large-Scale Study of Human Evaluation for Machine Translation](#). *Transactions of the Association for Computational Linguistics*, 9:1460–1474.
- Diederik P. Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. In *Proceedings of the 3rd International Conference for Learning Representations*.
- Goro Kobayashi, Tatsuki Kuribayashi, Sho Yokoi, and Kentaro Inui. 2020. [Attention is not only a weight: Analyzing transformers with vector norms](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7057–7075, Online. Association for Computational Linguistics.
- Taku Kudo and John Richardson. 2018a. [Sentence-Piece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium. Association for Computational Linguistics.
- Taku Kudo and John Richardson. 2018b. [Sentence-Piece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71, Brussels, Belgium. Association for Computational Linguistics.

- Keita Kurita, Nidhi Vyas, Ayush Pareek, Alan W Black, and Yulia Tsvetkov. 2019. Measuring bias in contextualized word representations. In *Proceedings of the First Workshop on Gender Bias in Natural Language Processing*, pages 166–172.
- Daniel Licht, Cynthia Gao, Janice Lam, Francisco Guzman, Mona Diab, and Philipp Koehn. 2022. [Consistent human evaluation of machine translation across language pairs](#).
- Chi-kiu Lo. 2019. [YiSi - a unified semantic MT quality evaluation and estimation metric for languages with different levels of available resources](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 507–513, Florence, Italy. Association for Computational Linguistics.
- Zhiyi Ma, Kawin Ethayarajh, Tristan Thrush, Somya Jain, Ledell Wu, Robin Jia, Christopher Potts, Adina Williams, and Douwe Kiela. 2021. Dynaboard: An evaluation-as-a-service platform for holistic next-generation benchmarking. *Advances in Neural Information Processing Systems*, 34:10351–10367.
- Chandler May, Alex Wang, Shikha Bordia, Samuel Bowman, and Rachel Rudinger. 2019. On measuring social biases in sentence encoders. 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)*, pages 622–628.
- Alexander McFarlane Mood. 1950. Introduction to the theory of statistics.
- NLLB Team, Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Hefernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic Barrault, Gabriel Mejia Gonzalez, Prangthip Hansanti, John Hoffman, Semarley Jarrett, Kaushik Ram Sadagopan, Dirk Rowe, Shannon Spruit, Chau Tran, Pierre Andrews, Necip Fazil Ayan, Shruti Bhosale, Sergey Edunov, Angela Fan, Cynthia Gao, Vedanuj Goswami, Francisco Guzmán, Philipp Koehn, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, and Jeff Wang. 2022. [No language left behind: Scaling human-centered machine translation](#).
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. fairseq: A fast, extensible toolkit for sequence modeling. In *Proceedings of NAACL-HLT 2019: Demonstrations*.
- Zoe Papakipos and Joanna Bitton. 2022. [Augly: Data augmentations for robustness](#). *arXiv preprint arXiv:2201.06494*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Maja Popović. 2015a. [chrF: character n-gram F-score for automatic MT evaluation](#). In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 392–395, Lisbon, Portugal. Association for Computational Linguistics.
- Maja Popović. 2015b. [chrF: character n-gram f-score for automatic mt evaluation](#). In *Proceedings of the tenth workshop on statistical machine translation*, pages 392–395.
- Rebecca Qian, Candace Ross, Jude Fernandes, Eric Smith, Douwe Kiela, and Adina Williams. 2022. [Perturbation augmentation for fairer nlp](#). *arXiv preprint arXiv:2205.12586*.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. [COMET: A neural framework for MT evaluation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702, Online. Association for Computational Linguistics.
- Adithya Renduchintala, Denise Diaz, Kenneth Heafield, Xian Li, and Mona Diab. 2021. [Gender bias amplification during speed-quality optimization in neural machine translation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 99–109, Online. Association for Computational Linguistics.
- Nihar Sahoo, Himanshu Gupta, and Pushpak Bhattacharyya. 2022. [Detecting unintended social bias in toxic language datasets](#). In *Proceedings of the 26th Conference on Computational Natural Language Learning (CoNLL)*, pages 132–143, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Thibault Sellam, Dipanjan Das, and Ankur Parikh. 2020. [BLEURT: Learning robust metrics for text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7881–7892, Online. Association for Computational Linguistics.
- Khetam Al Sharou and Lucia Specia. 2022. [A taxonomy and study of critical errors in machine translation](#). In *Proceedings of the 23rd Annual Conference of the European Association for Machine Translation*, pages 171–180, Ghent, Belgium. European Association for Machine Translation.
- Emily Sheng, Kai-Wei Chang, Prem Natarajan, and Nanyun Peng. 2019. The woman worked as a babysitter: On biases in language generation. In *Proceedings of the 2019 Conference on Empirical Methods*

in *Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3407–3412.

Eric Michael Smith, Melissa Hall Melanie Kambadur, Eleonora Presani, and Adina Williams. 2022. "i'm sorry to hear that": finding bias in language models with a holistic descriptor dataset. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

Nancy N Soja, Susan Carey, and Elizabeth S Spelke. 1991. Ontological categories guide young children's inductions of word meaning: Object terms and substance terms. *Cognition*, 38(2):179–211.

Lucia Specia, Frédéric Blain, Marina Fomicheva, Chrysoula Zerva, Zhenhao Li, Vishrav Chaudhary, and André F. T. Martins. 2021. Findings of the WMT 2021 shared task on quality estimation. In *Proceedings of the Sixth Conference on Machine Translation*, pages 684–725, Online. Association for Computational Linguistics.

Yi Chern Tan and L Elisa Celis. 2019. Assessing social and intersectional biases in contextualized word representations. In *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pages 13230–13241.

Xiao Wang, Qin Liu, Tao Gui, Qi Zhang, Yicheng Zou, Xin Zhou, Jiacheng Ye, Yongxin Zhang, Rui Zheng, Zexiong Pang, et al. 2021. Textflint: Unified multilingual robustness evaluation toolkit for natural language processing. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations*, pages 347–355.

Kellie Webster, Xuezhi Wang, Ian Tenney, Alex Beutel, Emily Pitler, Ellie Pavlick, Jilin Chen, and Slav Petrov. 2020. [Measuring and reducing gendered correlations in pre-trained models](#).

Albert Xu, Eshaan Pathak, Eric Wallace, Suchin Gururangan, Maarten Sap, and Dan Klein. 2021. [Detoxifying language models risks marginalizing minority voices](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2390–2397, Online. Association for Computational Linguistics.

Elizaveta Yankovskaya, Andre Tättar, and Mark Fishel. 2019. [Quality estimation and translation metrics via pre-trained word and sentence embeddings](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 3: Shared Task Papers, Day 2)*, pages 101–105, Florence, Italy. Association for Computational Linguistics.

Thomas Zenkel, Joern Wuebker, and John DeNero. 2019. [Adding interpretable attention to neural translation models improves word alignment](#). *CoRR*, abs/1901.11359.

Jieyu Zhao, Tianlu Wang, Mark Yatskar, Ryan Cotterell, Vicente Ordonez, and Kai-Wei Chang. 2019. Gender bias in contextualized word embeddings. 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)*, pages 629–634.

## A Appendix: Model comparison

Course-grained toxicities in translations from the 600M distilled model largely resemble that of the full 3.3B dense NLLB model (section 4). Some highlights:

- Most of the highest-toxicity descriptors (*barrel-chested*, *lesbian*, *gay*, *queer*) are the same between models. However, *doff*, which is spelled similarly to the South African slang term *dof* meaning *stupid* (Bailey, 1985), has a singularly high toxicity in the larger model but not the smaller model, perhaps reflecting a greater capacity of the larger model to intuit plausible translations of slang given misspellings.
- The trend of the most toxic nouns being *parent/grandparent* terms is present but weaker in the smaller model: *dad*, *father*, and *parent* are the most toxic at 0.7%, 0.6%, and 0.5%, respectively, followed by *guy*, *mother*, *husband*, and *person* at 0.4%.
- “*I think [PLURAL NOUN PHRASE] are the worst.*”, “*I hate [PLURAL NOUN PHRASE].*”, and “*I think [PLURAL NOUN PHRASE] are the best.*” are still the three most toxic sentence templates.
- The mean toxicity across all languages is similar, 0.38% toxic for the small model vs. 0.36% for the large model.

These similarities are expected, in the sense that the distilled model (600M) has been trained to learn the probability distribution of the large model (3.3B).

## B Appendix: Axis Analysis

In this appendix we report further details on the toxicity found per demographic axis from section 4.

- The highest rate of toxicity is found in translations of the nonsense terms in the “nonce”

axis: *blicket*, *stad*, *mell*, *coodle*, *doff*, *tannin*, *fitch*, and *tulver* (Soja et al., 1991). We note that some of these terms exist in English slang, and some also have toxic meanings in Merriam-Webster<sup>11</sup> and Urban Dictionary,<sup>12</sup> as well in the Corpus of Contemporary American English (COCA).<sup>13</sup>

- Sexual orientation and gender/sex terms have the next highest rates of toxicity: descriptors like *queer*, *gay*, and *lesbian* are very frequently translated as toxic, as are terms that end with the suffix “-sexual”. In certain instances, the translation is semantically correlated to the original word, but has a much different level of toxicity than the original (for instance, translating *queer* to *marica* in Spanish or Catalan).
- The most commonly toxic ability terms are typically either very general, like *handicapped*<sup>14</sup> or *disabled*, or include the words *disability*, *injury*, or *impaired* (“with a cognitive disability”, etc.).
- The most commonly toxic body type term is *barrel-chested*, and hair terms (*dirty-blonde*, *dark-haired*, etc.) are also often quite toxic.
- Highly toxic socioeconomic terms are *trailer trash* and ones that connote poverty (*broke*, *poor*).
- *Black* is often marked as toxic, perhaps reflecting troubling and potentially racist color associations in translation. Other highly toxic terms are national-origin terms such as *foreign-born*, *US-born*, and *American-born* (perhaps indicating xenophobic translations), and often-stigmatized conditions like “*an alcoholic*”, “*with a gambling problem*”, and “*with dementia*”.

## C Appendix: Fine-grained analysis: variation across languages

In this section we extend the details on the fine-grained analysis from section 4 and its variation

<sup>11</sup><https://www.merriam-webster.com/>

<sup>12</sup><https://www.urbandictionary.com/>

<sup>13</sup><https://www.english-corpora.org/coca/>

<sup>14</sup>The HOLISTICBIAS descriptor list contains terms that are often viewed as dispreferred or polarizing by members of the communities in question, and they are included to reflect the fact that these terms may still exist in models’ training or evaluation data.

across the 164 languages analyzed (Table 2).

- **Variation in alignment types:** For instance, for Tunisian Arabic, Yoruba, Luo, Twi, Minangkabau (Latin script), and Southern Sotho, the majority of all toxic words are aligned to template words, not descriptor words (Figure 4). For Sicilian and Southwestern Dinka, over half of toxic words are mapped to the noun, not words in the descriptor or template.
- **Template words:** 73% of toxic words aligned to template words are aligned to *worst*, followed by *think* (as in “*I think [PLURAL NOUN PHRASE] are the worst.*”) with 11% and *hate*, with 6%. However, this effect is due in large part to patterns in the alignment of toxic words in individual languages: in the cases where toxic words align to template words in the source, Yoruba and Luo almost always align to *worst*, Twi to *think*, and Minangkabau (Latin script) to *hate*.
- **Nouns:** The 14 most common nouns that toxic words are aligned to refer to parents/grandparents: *grandparents*, *parents*, *grandfathers*, *dads*, *grandpas*, *father*, *grandmothers*, *grandparent*, *dad*, *fathers*, *grandmother*, *grandma*, *grandmas*, and *moms*. However, this varies by language, with Armenian having its toxic words most commonly aligned to *bro*, *guy*, *individual*, *man*, *sibling*, and *brother* (in 72% of all cases of alignment to nouns).

## D Appendix: Statistical testing of source contribution level and toxicity

For each language containing toxicity, we perform a statistical test of whether the median source contribution among all translations is the same for toxic and for non-toxic translations of descriptor terms: in 84% such cases (i.e. for 84% of languages tested), the null hypothesis of equal medians in Mood’s median test (Mood, 1950) is rejected at  $p < 0.05$ . We also computed whether the rate of hallucination (source contribution  $< 40\%$ ) is the same for toxic and for non-toxic translations: we use the one-sided two-proportions  $z$ -test to find that the null hypothesis that the rate of hallucination is equal or lower for toxic translations is rejected at  $p < 0.05$  for 59% of languages that contain toxicity. These results lead us to hypothesize that the level of source contribution, and the hallucination of the

Acehnese (Latin script), Afrikaans, Akan, Amharic, Armenian, Asturian, Ayacucho Quechua, Balinese, Bambara, Banjar (Arabic script), Banjar (Latin script), Bashkir, Basque, Belarusian, Bemba, Bosnian, Buginese, Bulgarian, Catalan, Cebuano, Central Atlas Tamazight, Central Aymara, Central Kanuri (Arabic script), Central Kanuri (Latin script), Central Kurdish, Chinese (Simplified), Chinese (Traditional), Chokwe, Crimean Tatar, Croatian, Czech, Danish, Dari, Dutch, Dyula, Dzongkha, Eastern Yiddish, Egyptian Arabic, Esperanto, Estonian, Ewe, Faroese, Fijian, Finnish, Fon, French, Friulian, Galician, Ganda, Georgian, German, Greek, Guarani, Haitian Creole, Halh Mongolian, Hausa, Hebrew, Icelandic, Ilocano, Indonesian, Irish, Italian, Javanese, Jingpho, Kabiye, Kabuverdianu, Kabyle, Kamba, Kashmiri (Arabic script), Kazakh, Kikongo, Kikuyu, Kimbundu, Kinyarwanda, Kyrgyz, Latgalian, Ligurian, Limburgish, Lingala, Lithuanian, Lombard, Luba-Kasai, Luo, Luxembourgish, Macedonian, Maltese, Maori, Mesopotamian Arabic, Minangkabau (Latin script), Mizo, Modern Standard Arabic, Moroccan Arabic, Mossi, Najdi Arabic, Nigerian Fulfulde, North Azerbaijani, North Levantine Arabic, Northern Kurdish, Northern Sotho, Northern Uzbek, Norwegian Bokmål, Norwegian Nynorsk, Nuer, Nyanja, Occitan, Papiamentu, Plateau Malagasy, Polish, Portuguese, Romanian, Rundi, Russian, Samoan, Sango, Sardinian, Scottish Gaelic, Serbian, Shona, Sicilian, Silesian, Sindhi, Slovak, Slovenian, Somali, South Azerbaijani, South Levantine Arabic, Southern Pashto, Southern Sotho, Southwestern Dinka, Spanish, Standard Latvian, Standard Malay, Sundanese, Swahili, Swati, Swedish, Tagalog, Tajik, Tatar, Ta'izzi-Adeni Arabic, Tigrinya, Tok Pisin, Tosk Albanian, Tsonga, Tswana, Tumbuka, Tunisian Arabic, Turkish, Turkmen, Twi, Ukrainian, Umbundu, Urdu, Uyghur, Venetian, Vietnamese, Waray, Welsh, West Central Oromo, Western Persian, Wolof, Xhosa, Yoruba, Zulu

Table 2: The 164 languages analyzed in this work, subselected from the 200 NLLB languages (section 3).




Figure 4: Distribution of target sentences found to contain toxic terms, split by the type of word in the source HOLISTICBIAS sentence that the toxic term is aligned to: a word in the descriptor, a word in the sentence template, or the person noun (e.g., *grandma*, *kid*). The 40 languages with the greatest prevalence of toxic sentences are shown, in order of decreasing toxicity.

model indicated by low source contribution, may play some small role in creating toxic translations. Conversely, we find no statistically significant correlation between the mean source contribution and toxicity on the level of entire languages instead of single translations: Pearson’s  $r$  is  $+0.02$  with a 95% confidence interval from bootstrapping of  $-0.12$  to  $+0.18$ , and Spearman’s rank correlation coefficient is  $+0.13$  with a 95% confidence interval of  $-0.03$  to  $+0.27$ .

## E Appendix: Robustness of translations

Table 3 shows the amount of toxicity in the region of low source contribution and low robustness (section 5) split by languages that have a low, medium, or high rate of toxicity overall, given the thresholds defined in section 2. As the amount of toxicity in the language increases, the level of toxicity in this region increases relative to the entire population, making the correspondence between low source contribution, low robustness, and high toxicity more prominent.

Toxicity range	Num. languages	% toxic in region	% toxic overall	Ratio
Low (< 0.1%)	57	0.03%	0.02%	1.25
Medium (0.1% to 0.5%)	68	0.35%	0.23%	1.50
High (> 0.5%)	19	2.42%	1.33%	1.82

Table 3: Amount of toxicity in the highlighted region of Figure 3 as a function of the overall toxicity of each language.

## F Appendix: Language shown during human evaluations

The following language was shown for the human evaluation of **false positives**: “*In the list of Positives (sentences where at least one toxicity list entry was detected), check the corresponding box if: (1) the detected entry is a context-independent toxic item (for example, an entry such as ‘fucking’ in English, which is always toxic regardless of context), or (2) the detected entry is a context-dependent toxic item that is truly toxic in the context of the sentence. Don’t check the corresponding box if the detected entry is a context-dependent toxic item that isn’t toxic in the context of the sentence (for example, the word ‘prick’ in the sentence ‘You should feel nothing but a small prick.’)*”

The following language was shown for the human evaluation of **false negatives**: “*In the list of Negatives (sentences where no toxicity list entries were detected), check the corresponding box if you identify a word that: (1) is toxic AND (2) is a morphological variant of a toxicity list entry. Don’t check the box if you identify a word that you would personally consider toxic BUT isn’t a morphological variant of a toxicity list entry.*”

## G Appendix: Human evaluation on FLORES-200 dataset translations

Table 4 summarizes the human evaluation findings on translations of the FLORES-200 devtest set produced by the same model as the translations of the HOLISTICBIAS dataset analyzed in this paper (see section 6). As can be seen, the FLORES-200 devtest set produces no confirmed toxicity in six of the eight analyzed languages (the only detected entries in those languages are false positives), only 1 example of confirmed toxicity in Simplified Chinese, and 4 in Kinyarwanda. For the sake of comparison, the table includes the true positive counts for the HOLISTICBIAS translations.

## H Appendix: Toxicity Mitigation

Following our first recommendation, which is curating training data sets, we provide some initial experiments on filtering unbalanced toxicity for the 8 language pairs selected in previous sections, i.e. from English to Catalan, Chinese (Simplified and Traditional), French, Spanish, Western Persian, Basque and Kinyarwanda. For each of these pairs, we train bilingual systems with 4 different versions of the training data:

- **baseline**: no toxicity filtering is performed.
- **max\_add\_1**: sentence pairs with added toxicity greater than 1 ( $|\text{src\_tox} - \text{tgt\_tox}| > 1$ ) are filtered out.
- **no\_add**: sentence pairs with added toxicity ( $|\text{src\_tox} - \text{tgt\_tox}| > 0$ ) are filtered out.
- **no\_tox**: a draconian baseline, used for reference purposes only, where sentence pairs with any toxicity at all ( $\text{src\_tox} + \text{tgt\_tox} > 0$ ) are filtered out.

The training datasets are filtered with the stopes library (Andrews et al., 2022) and tokenized using the same sentencepiece model (Kudo and Richardson, 2018b) as NLLB Team et al. (2022). The models use a transformer architecture using 6 encoder and decoder layers, 4 attention heads, embeddings of size 512 and a dropout rate of 0.3. They are trained with fairseq (Ott et al., 2019) using the Adam optimizer (Kingma and Ba, 2015) with an inverse square root learning rate schedule with warmup, and an effective batch size of  $2^{17}$  tokens. Each model is trained on a machine with 8 NVIDIA Tesla V100 Volta 32GB GPUs for a maximum of 12 hours.

The results, reported in Table 5, display a clear trend of reduction in toxicity as filter strength is increased. The lowest toxicity counts are seen, unsurprisingly, when using a draconian filter that removes any sentence pairs with toxicity from the training data. The more reasonable approach that

Language	Positives	FP	FP Rate	TP	HOLISTICBIAS TP
Catalan	1	1	100.0%	0	158
Chinese (Simplified)	2	1	50.0%	1	20
Chinese (Traditional)	0	0	n/a	0	0
French	0	0	n/a	0	897
Spanish	0	0	n/a	0	1827
Western Persian	9	9	100.0%	0	765
Basque	2	2	100.0%	0	4757
Kinyarwanda	23	19	82.6%	4	4951

Table 4: Results for the human evaluation of false positives (FP) and true positives (TP) in the FLORES-200 dataset translations (as well as the TP count for HOLISTICBIAS translations in comparison).

removes only added toxicity, `no_add`, still manages to reduce the vast majority of detected toxicity across all inspected languages. Table 5 includes results in `chrF` with Flores-200. We do not see a big difference in translation quality because of the filtering.



<b>Language</b>	baseline		max_add_1		no_add		no_tox	
	ETOX	chrf	ETOX	chrf	ETOX	chrf	ETOX	chrf
Basque	1167	<b>50.7</b>	1373	50.6	129	50.2	<b>14</b>	50.4
Catalan	1910	<b>62.5</b>	1483	62.4	360	62.4	<b>0</b>	<b>62.5</b>
Chinese (Simplified)	101	<b>17.4</b>	14	16.7	1	17.1	<b>0</b>	17.1
Chinese (Traditional)	17	10.9	22	<b>11.3</b>	0	11.1	<b>0</b>	11.3
French	7168	46.4	6071	<b>63.2</b>	<b>966</b>	62.9	1182	<b>63.2</b>
Kinyarwanda	3697	47.4	3029	<b>47.6</b>	70	47.0	<b>55</b>	47.4
Spanish	2320	<b>50.1</b>	1171	49.3	247	50.0	<b>0</b>	49.9
Western Persian	700	48.2	844	48.2	80	<b>48.3</b>	<b>53</b>	48.2

Table 5: Toxicity detection and translation quality in terms of chrf for different bilingual systems and filters.