

Data-Efficient Strategies for Expanding Hate Speech Detection into Under-Resourced Languages

Paul Röttger¹, Debora Nozza², Federico Bianchi³, and Dirk Hovy²

¹University of Oxford

²Bocconi University

³Stanford University

Abstract

Hate speech is a global phenomenon, but most hate speech datasets so far focus on English-language content. This hinders the development of more effective hate speech detection models in hundreds of languages spoken by billions across the world. More data is needed, but annotating hateful content is expensive, time-consuming and potentially harmful to annotators. To mitigate these issues, we explore data-efficient strategies for expanding hate speech detection into under-resourced languages. In a series of experiments with mono- and multilingual models across five non-English languages, we find that 1) a small amount of target-language fine-tuning data is needed to achieve strong performance, 2) the benefits of using more such data decrease exponentially, and 3) initial fine-tuning on readily-available English data can partially substitute target-language data and improve model generalisability. Based on these findings, we formulate actionable recommendations for hate speech detection in low-resource language settings.

Content warning: This article contains illustrative examples of hateful language.

1 Introduction

Hate speech is a global phenomenon, but most hate speech datasets so far focus on English-language content (Vidgen and Derczynski, 2020; Poletto et al., 2021). This hinders the development of effective models for detecting hate speech in other languages. As a consequence, billions of non-English speakers across the world are less protected against online hate, and even giant social media platforms have clear language gaps in their content moderation systems (Simonite, 2021; Marinescu, 2021).

Zero-shot cross-lingual transfer, where large multilingual language models are fine-tuned on one source language and then applied to another target language, may appear like a potential solution to the issue of language-specific resource

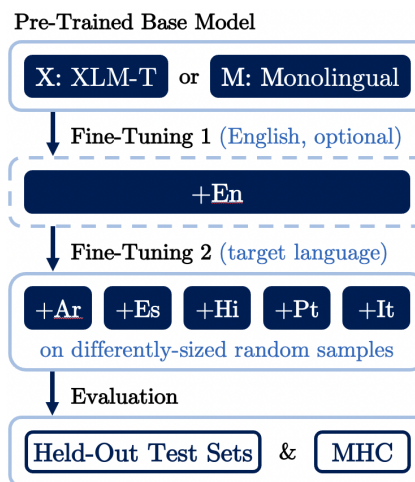


Figure 1: Overview of our experimental setup. We use ISO 639-1 codes to denote the different languages. MHC is Multilingual HateCheck (Röttger et al., 2022).

scarcity. However, while this approach performs well on some tasks (Conneau et al., 2020; Barbieri et al., 2022), it fails on many others (Lauscher et al., 2020; Hu et al., 2020). For hate speech detection in particular, cross-lingual performance in zero-shot settings is lacking (Stappen et al., 2020; Leite et al., 2020). For example, zero-shot cross-lingual transfer cannot account for language-specific taboo expressions that play a key role in classification (Nozza, 2021). Conversely, hate speech detection models trained or fine-tuned directly on the target language, i.e. in few- and many-shot settings, are consistently found to perform best (Aluru et al., 2020; Pelicon et al., 2021).

So, **how do we build hate speech detection models for hundreds more languages?** We need at least some labelled data in the target language to make models effective, but data annotation is difficult, time-consuming and expensive. It requires resources that are often very limited for non-English languages. Annotating *hateful* content in particular also risks exposing annotators to harm in the process (Vidgen et al., 2019; Derczynski et al., 2022).

In this article, we explore strategies for hate speech detection in under-resourced languages that make *efficient* use of labelled data, to build *effective* models while also minimising annotation cost and risk of harm to annotators. For this purpose, we conduct a series of experiments using mono- and multilingual models fine-tuned on differently-sized random samples of labelled hate speech data in English as well as Arabic, Spanish, Hindi, Portuguese and Italian. Our **key findings** are:

1. A small amount of labelled target-language data is needed to achieve strong performance on held-out test sets.
2. The benefits of using more such data decrease exponentially.
3. Initial fine-tuning on readily-available English data can partially substitute target-language data and improve model generalisability.

Based on these findings, we formulate and discuss **five recommendations** for expanding hate speech detection into under-resourced languages:

1. Collect and label target-language data.
2. Start by labelling a small set, then iterate.
3. Use diverse data collection methods to increase marginal benefits of annotation.
4. Use multilingual models for unlocking readily-available data in high-resource languages for initial fine-tuning.
5. Evaluate out-of-domain performance to reveal potential weaknesses in generalisability.

With these recommendations, we hope to facilitate the development of new hate speech detection models for yet-unserved languages.¹

Definition of Hate Speech Definitions of hate speech vary across cultural and legal settings. Following Röttger et al. (2021), we define hate speech as *abuse that is targeted at a protected group or at its members for being a part of that group*. Protected groups are based on characteristics such as gender identity, race or religion, which broadly reflects Western legal consensus, particularly the US 1964 Civil Rights Act, the UK’s 2010 Equality Act and the EU’s Charter of Fundamental Rights. Based on these definitions, we approach hate speech detection as the binary classification of content as either hateful or non-hateful.

¹We make all data and code to reproduce our experiments available on [GitHub](#).

2 Experiments

All our experiments follow the setup described in Figure 1. We start by loading a pre-trained mono- or multilingual transformer model for sequence classification. For multilingual models, there is an optional first phase of fine-tuning on English data. This is to simulate using readily-available data from a high-resource language that is not the target language. For all models, there is then a second phase of fine-tuning on differently-sized random samples of data in the target language. This is to simulate using scarce data from an under-resourced language. Finally, all models are evaluated on the held-out test set corresponding to the target-language dataset they were fine-tuned on as well as the target-language test suite from Multilingual HateCheck (Röttger et al., 2022).

2.1 Data

For all our experiments, we use hate speech datasets from [hatespeechdata.com](#), which was first introduced by Vidgen and Derczynski (2020) and is now the largest public repository of datasets annotated for hate, abuse and offensive language. At the time of our review in May 2022, the site listed 53 English datasets as well as 57 datasets in 24 other languages. From these datasets, we select those for our experiments that a) contain explicit labels for hate, and b) use a definition of hate for data annotation that aligns with our own (§1).

2.1.1 Fine-Tuning 1: English

For the optional first phase of fine-tuning in English, we use one of three English datasets. DYN21_EN by Vidgen et al. (2021) contains 41,255 entries, of which 53.9% are labelled as hateful. The entries were hand-crafted by annotators to be challenging to hate speech detection models, using the Dynabench platform (Kiela et al., 2021). FOU18_EN by Founta et al. (2018) contains 99,996 tweets, of which 4.97% are labelled as hateful. KEN20_EN by Kennedy et al. (2020) contains 39,565 comments from Youtube, Twitter and Reddit, of which 29.31% are labelled as hateful.

From each of these three English datasets, we sample 20,000 entries for the optional first phase of fine-tuning, plus another 500 entries for development and 2,000 for testing. To align the proportion of hate across English datasets, we use random sampling for DYN21_EN, while for KEN20_EN and FOU18_EN we retain all hateful entries and

then sample from non-hateful entries, so that the proportion of hate in the two datasets increases to 50.0% and 22.0%, respectively.

2.1.2 Fine-Tuning 2: Target Language

For fine-tuning in the target language, we use one of five datasets in five different target languages. BAS19_ES compiled by Basile et al. (2019) for SemEval 2019 contains 4,950 Spanish tweets, of which 41.5% are labelled as hateful. FOR19_PT by Fortuna et al. (2019) contains 5,670 Portuguese tweets, of which 31.5% are labelled as hateful. HAS21_HI compiled by Modha et al. (2021) for HASOC 2021 contains 4,594 Hindi tweets, of which 12.3% are labelled as hateful. OUS19_AR by Ousidhoum et al. (2019) contains 3,353 Arabic tweets, of which 22.5% are labelled as hateful. SAN20_IT compiled by Sanguinetti et al. (2020) for EvalIta 2020 contains 8,100 Italian tweets, of which 41.8% are labelled as hateful.

From each of these five target-language datasets, we randomly sample differently-sized subsets for target-language fine-tuning. Like in English, we set aside 500 entries for development and 2,000 for testing.² From the remaining data, we sample subsets in 12 different sizes – 10, 20, 30, 40, 50, 100, 200, 300, 400, 500, 1,000 and 2,000 entries – so that we are able to evaluate the effects of using more or less labelled data within and across different orders of magnitude.³ Zhao et al. (2021) show that there can be large sampling effects when fine-tuning on small amounts of data. To mitigate this issue, we use 10 different random seeds for each sample size, so that in total we have 120 different samples in each language, and 600 samples across the five non-English languages.

2.2 Models

Multilingual Models We fine-tune and evaluate XLM-T (Barbieri et al., 2022), an XLM-R model (Conneau et al., 2020) pre-trained on an additional 198 million Twitter posts in over 30 languages. XLM-R is a widely-used architecture for multilingual language modelling, which has been shown to achieve near state-of-the-art performance on multilingual hate speech detection (Banerjee et al., 2021; Modha et al., 2021). We chose XLM-T because it strongly outperformed XLM-R across our target language test sets in initial experiments.

²Due to limited dataset size, we only set aside 300 dev and 1,000 test entries for OUS19_AR (n=3,353).

³There is at least one hateful entry in every sample.

Monolingual Models For each of the five target languages, we also fine-tune and evaluate a monolingual transformer model from HuggingFace. For Spanish, we use RoBERTuito (Pérez et al., 2021). For Portuguese, we use BERTimbau (Souza et al., 2020). For Hindi, we use Hindi BERT. For Arabic, we use AraBERT v2 (Antoun et al., 2020). For Italian, we use UmBERTo. Details on model training can be found in Appendix A.

Model Notation We denote all models by an additive code. The first part is either M for a monolingual model or X for XLM-T. For XLM-T, the second part of the code is DEN, FEN or KEN, for models fine-tuned on 20,000 entries from DYN21_EN, FOU18_EN or KEN20_EN. For all models, the final part of the code is ES, PT, HI, AR or IT, corresponding to the target language that the model was finetuned on. For example, M+IT denotes the monolingual Italian model, UmBERTo, fine-tuned on SAN20_IT, and X+KEN+AR denotes an XLM-T model fine-tuned first on 20,000 English entries from KEN20_EN and then on OUS19_AR.

2.3 Evaluation Setup

Held-Out Test Sets + MHC We test all models on the held-out test sets corresponding to their target-language fine-tuning data, to evaluate their in-domain performance (§2.4). For example, we test X+KEN+IT, which was fine-tuned on SAN20_IT data, on the SAN20_IT test set. Additionally, we test all models on the matching target-language test suite from Multilingual HateCheck (MHC). MHC is a collection of around 3,000 test cases for different kinds of hate as well as challenging non-hate in each of ten different languages (Röttger et al., 2022). We use MHC to evaluate out-of-domain generalisability (§2.5).

Evaluation Metrics We use macro F1 to evaluate model performance because most of our test sets as well as MHC are imbalanced. To give context for interpreting performance, we show baseline model results in all figures: macro F1 for always predicting the hateful class ("always hate"), for never predicting the hateful class ("never hate") and for predicting both classes with equal probability ("50/50"). We also show bootstrapped 95% confidence intervals around the average macro F1, which is calculated across the 10 random seeds for each sample size. These confidence intervals are expected to be wider for models fine-tuned on less data because of larger sampling effects.

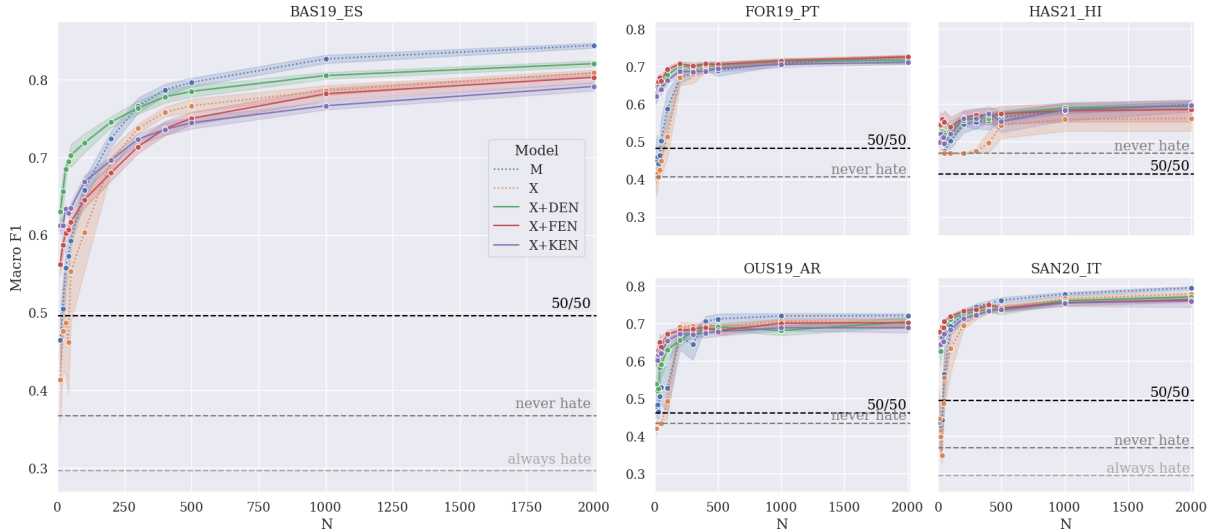


Figure 2: Macro F1 on target-language held-out test sets across models fine-tuned on up to $N=2,000$ target-language entries. Model notation as described in §2.2. Confidence intervals and model baselines as described in §2.4. We provide larger versions of all five graphs in Appendix B.

	BAS19_ES			FOR19_PT			HAS21_HI			OUS19_AR			SAN20_IT		
N	20	200	2,000	20	200	2,000	20	200	2,000	20	200	2,000	20	200	2,000
M	0.50	0.72	0.84	0.46	0.67	0.73	0.47	0.55	0.60	0.48	0.66	0.72	0.40	0.73	0.79
X	0.48	0.70	0.81	0.42	0.67	0.73	0.47	0.47	0.56	0.43	0.69	0.70	0.40	0.70	0.78
X+DEN	0.66	0.75	0.82	0.63	0.70	0.72	0.52	0.56	0.60	0.52	0.66	0.70	0.63	0.73	0.77
X+FEN	0.59	0.68	0.80	0.66	0.71	0.73	0.55	0.56	0.59	0.61	0.68	0.70	0.66	0.73	0.76
X+KEN	0.61	0.70	0.79	0.65	0.69	0.71	0.52	0.56	0.60	0.60	0.67	0.69	0.64	0.71	0.76

Table 1: Macro F1 on respective held-out test sets for models fine-tuned on N target-language entries, averaged across 10 random seeds for each N . Best performance for a given N in **bold**. Results across all N in Appendix B.

2.4 Testing on Held-Out Test Sets

When evaluating our mono- and multilingual models on their corresponding target-language test sets, we find a set of consistent patterns in model performance. We visualise overall performance in Figure 2 and highlight key data points in Table 1.

First, there is an enormous benefit from even very small amounts of target-language fine-tuning data N . Model performance increases sharply across models and held-out test sets up to around $N=200$. For example, the performance of $X+DEN+ES$ increases from around 0.63 at $N=10$ to 0.70 at $N=50$, to 0.75 at $N=200$.

On the other hand, larger amounts of target-language fine-tuning data correspond to much less of an improvement in model performance. Across models and held-out test sets, there is a steep decrease in the marginal benefits of increasing N . $M+IT$, for example, improves by 0.33 macro F1 from $N=20$ to $N=200$, and by just 0.06 from $N=200$ to $N=2,000$. $X+PT$ improves by 0.25 macro F1

from $N=20$ to $N=200$, and by just 0.06 from $N=200$ to $N=2,000$. We analyse these decreasing marginal benefits using linear regression in §2.6.

Further, there is a clear benefit to a first phase of fine-tuning on English data, when there is limited target-language data. Absolute performance differs across test sets, but multilingual models with initial fine-tuning on English data (i.e. $X+DEN$, $X+FEN$ and $X+KEN$) perform substantially better than those without (i.e. M and X), up to around $N=200$. At $N=20$, there is up to 0.26 macro F1 difference between the former and the latter. Conversely, models without initial fine-tuning on English data need substantially more target-language data to achieve the same performance. For example, $X+DEN+PT$ at $N=100$ performs as well as $M+PT$ at $N=300$ on $FOR19_PT$.

Relatedly, there are clear differences in model performance based on which English dataset was used in the first fine-tuning phase, when there is limited target-language data. Among the three En-

glish datasets we evaluated, X+DEN performs best for up to around N=200 on BAS19_ES, whereas X+FEN wins out for the other four test sets.

After the strong initial divergence for low amounts of target-language fine-tuning data, performance tends to align across models, regardless of whether they were first fine-tuned on English data or only fine-tuned on the target language. From around N=200 onwards, for a given test set, all models we evaluate have broadly comparable performance. For example, all models score 0.71 to 0.72 macro F1 on FOR19_PT for N=1,000.

Finally, we find that if there is some divergence in model performance on the held-out test sets for larger amounts of target-language fine-tuning data, then monolingual models tend to outperform multilingual models. For example, despite the disadvantage of monolingual models for $N \leq 200$, we find that M+ES achieves 0.84 macro F1 on BAS19_ES for N=2,000, compared 0.82 macro F1 for X+DEN+ES. This difference is visible across

test sets and significant at 95% confidence for the Spanish, Arabic and Italian test sets.

2.5 Testing on MHC

Performance on the target-language test suites from MHC, which none of the models saw during fine-tuning, differs strongly and systematically from results on the held-out test sets (Figure 3). Key data points are highlighted in Table 2.

First, as for the held-out test sets, there is an initial benefit to a first phase of fine-tuning on English data, up to around N=500 for MHC. However, there are clear performance differences depending on which English dataset was used. Models fine-tuned on DYN21_EN perform best by far across MHC for any N. On Italian HateCheck for example, X+DEN achieves 0.79 macro F1 at N=20 compared to 0.55 for X+FEN+IT and 0.29 for the monolingual model M+IT.

Second, for Spanish, Portuguese and Italian, there is little to no benefit to fine-tuning on more

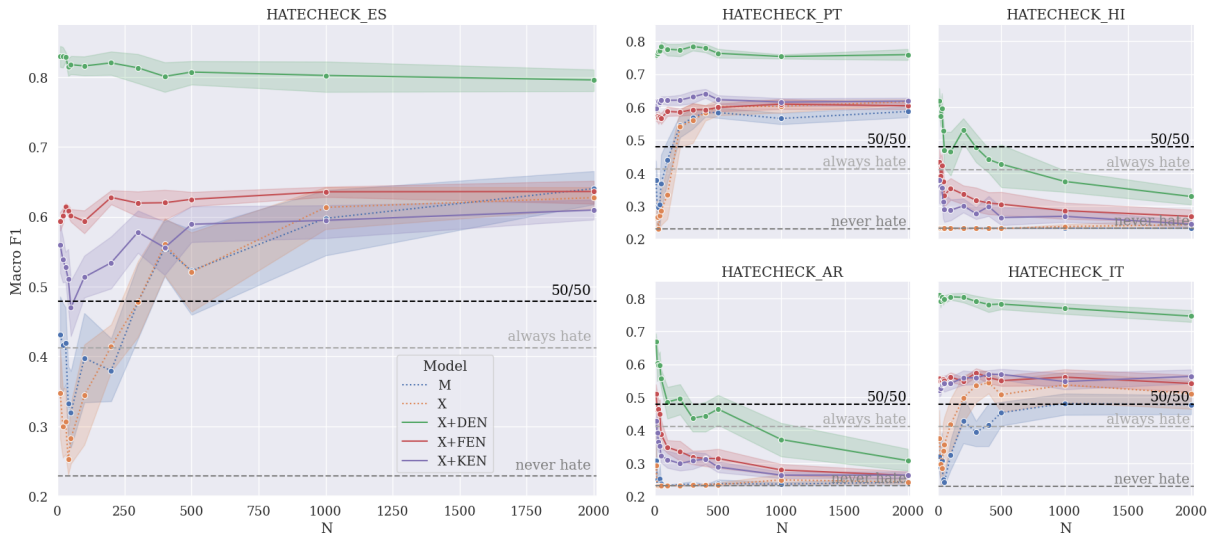


Figure 3: Macro F1 on target-language tests from MHC across models fine-tuned on up to N = 2,000 target-language entries. Model notation as described in §2.2. Confidence intervals and model baselines as described in §2.4. We provide larger versions of all five graphs in Appendix C.

	HATECHECK_ES			HATECHECK_PT			HATECHECK_HI			HATECHECK_AR			HATECHECK_IT		
N	20	200	2,000	20	200	2,000	20	200	2,000	20	200	2,000	20	200	2,000
M	0.42	0.38	0.64	0.30	0.54	0.59	0.23	0.23	0.23	0.25	0.23	0.24	0.29	0.43	0.48
X	0.30	0.41	0.63	0.27	0.54	0.62	0.23	0.23	0.24	0.23	0.23	0.24	0.30	0.50	0.51
X+DEN	0.83	0.82	0.80	0.76	0.77	0.76	0.57	0.53	0.33	0.60	0.50	0.31	0.79	0.80	0.75
X+FEN	0.60	0.63	0.64	0.57	0.58	0.60	0.39	0.34	0.27	0.46	0.34	0.26	0.55	0.55	0.54
X+KEN	0.54	0.53	0.61	0.62	0.62	0.62	0.36	0.30	0.24	0.39	0.30	0.26	0.53	0.56	0.56

Table 2: Macro F1 on target-language HateCheck for models fine-tuned on N target-language entries, averaged across 10 random seeds for each N. Best performance for a given N in **bold**. Results across all N in Appendix C.

target-language data, as measured by performance on MHC. Only those models that were not first fine-tuned on English data (i.e. M and X) are improved by fine-tuning on more target-language data, up to around $N=500$. None of the models benefit from even larger N . On Portuguese HateCheck for example, X+KEN+PT scores 0.62 macro F1 at $N=500$, $N=1,000$ and $N=2,000$.

Third, for Hindi and Arabic, training on more target-language data actively harms performance on MHC. For larger N , all models degrade towards only predicting the non-hateful class. On Arabic HateCheck, for example, X+FEN+AR achieves 0.26 macro F1 at $N=2,000$, which is only slightly better than the “never hate” baseline at 0.23 macro F1. This is despite the exact same model scoring 0.70 macro F1 on OUS19_AR at $N=2,000$. We further discuss this finding in §3.

Lastly model performance on MHC is much more volatile across random seeds compared to performance on the held-out test sets. This can be seen from the wider confidence intervals in Figure 3 compared to Figure 2.

2.6 Regression Analysis

When evaluating models on their corresponding held-out test sets, we saw decreasing benefits to adding more target-language fine-tuning data (§2.4). To quantify this observation and evidence its significance, we fit ordinary least squares (OLS) linear regressions for each model on each test set. OLS is commonly used in economics and the social sciences to analyse empirical relationships between an outcome variable and one or more regressor variables. In our case, the outcome is model performance as measured by macro F1 and the regressor is the amount of target-language fine-tuning data N . Since the relationship between the two variables is clearly non-linear and plausibly exponential, we regress macro F1 on the natural logarithm of N instead of just N , so that our regression equation for a given random seed i takes the form of

$$F_i = \beta_0 + \beta_1 \ln(N_i) + \varepsilon_i$$

where F_i is the observed macro F1 and ε_i the error term. β_0 , the intercept, can be interpreted as the expected zero-shot macro F1, i.e. at $N=0$. β_1 , the slope, can be interpreted as the expected increase in macro F1 from increasing the amount of target-language fine-tuning data N by 1%, due to our log-linear regression setup. In addition to reporting

regression coefficients, we also report adjusted R^2 as a goodness-of-fit measure, which is calculated as the percentage of variance in macro F1 explained by our regression. Lastly, we report the expected percentage-point increase in macro F1 from doubling N , denoted as $2N$, which is calculated as $\beta_1 \ln(2)$. Since we saw that variance in macro F1 changes with N , we use heteroskedasticity-robust HC3 standard errors (Hayes and Cai, 2007).

Model: M

Test Set	β_0	β_1	R^2	$2N$
BAS19_ES	0.2879	0.0792	91.97	5.49
FOR19_PT	0.2628	0.0670	71.23	4.65
HAS21_HI	0.2998	0.0607	69.11	4.21
OUS19_AR	0.2092	0.0866	70.08	6.00
SAN20_IT	0.3784	0.0294	77.68	2.04

Model: X+DEN

Test Set	β_0	β_1	R^2	$2N$
BAS19_ES	0.5543	0.0363	87.75	2.52
FOR19_PT	0.5756	0.0205	63.66	1.42
HAS21_HI	0.4691	0.0163	43.91	1.13
OUS19_AR	0.4320	0.0391	64.99	2.71
SAN20_IT	0.5724	0.0272	69.42	1.89

Table 3: OLS results for M and X+DEN. All regression coefficients β_0 and β_1 are significant at $p < 0.001$. We show results for all models in Appendix D.

We find that the natural logarithm of the amount of target-language fine-tuning data N provides a strikingly good explanation for model performance as measured by macro F1 (Table 3). R^2 for our regressions, where N is the sole explanatory variable, is generally very high, up to 91.97%, which indicates near-perfect goodness-of-fit. This strongly suggests that there is an exponential decrease in the benefits of using more target-language fine-tuning data. Conversely, there is a constant benefit to doubling the amount of target-language fine-tuning data: We expect the same increase in macro F1 whether we increase N ten-fold from 10 to 20 or from 1,000 to 2,000.

Further, we can confirm our earlier findings that a) initial fine-tuning on English data improves performance for low N , while b) all models perform similarly well for larger N . In our regression, this is visible in monolingual M models having much lower estimated zero-shot performance β_0

paired with much larger β_1 and $2N$ effects compared to the X+DEN models. For example, doubling the amount of Spanish fine-tuning data for M ($\beta_0=0.29$) corresponds to an expected increase in macro F1 of 5.49 points on BAS19_ES, whereas for X+DEN ($\beta_0=0.55$) we only expect an increase of 2.52 points.

3 Discussion and Recommendations

Based on the sum of our experimental results, we formulate five recommendations for the data-efficient development of effective hate speech detection models for under-resourced languages.

Recommendation 1

Collect and label target-language data.

We found enormous benefits from fine-tuning models on even small amounts of target-language data. Zero-shot cross-lingual transfer performs worse than naive baselines, but few shots go a long way.

Recommendation 2

Do not annotate all the data you collect right away. Instead, start small and iterate.

We found that the benefits of using more target-language fine-tuning data decrease exponentially (Table 3), which suggests that the data in each dataset is not diverse enough to warrant full annotation. The datasets we used in our experiments contained several thousand entries, but most of the useful signal in them can be learned from just a few dozen or hundred randomly-sampled entries.⁴ Thus, most of the annotation cost and the potential harm to annotators in compiling these datasets created very little performance benefit. Assuming constant cost per additional annotation, improving F1 by just one percentage point grows exponentially more expensive. Our results suggest that dataset creators should at first annotate smaller subsets of the data they collected and iteratively evaluate what marginal benefits additional annotations are likely to bring compared to their costs.

Recommendation 3

Use diverse data collection strategies to minimise redundant data annotation.

⁴Future work could explore more deliberate selection strategies, like active learning, which would likely achieve even better performance with fewer target-language entries (see for example Markov et al., 2022).

We do not recommend against annotating target-language data. Instead, we recommend against annotating too much data that was collected using the same method. The HAS21_HI dataset, for example, was collected using trending hashtags related to Covid and Indian politics (Modha et al., 2021), and Ousidhoum et al. (2019) used topical keywords related to Islamic inter-sect hate for OUS19_AR. As a consequence, these datasets necessarily have a narrow topical focus as well as limited linguistic diversity. This kind of *selection bias* has been confirmed for several other hate speech datasets (Wiegand et al., 2019; Ousidhoum et al., 2020). MHC, on the other hand, contains diverse statements against seven different target groups in each language (Röttger et al., 2022). Our results show that very quickly additional entries from a given dataset carry little new and useful information (Figure 2), and that fine-tuning on particularly narrow datasets like HAS21_HI and OUS19_AR can even harm model generalisability as measured by MHC (Figure 3). Conversely, more diverse strategies for data collection, such as sampling from multiple platforms, sampling from a diverse set of users, or using a wider variety of topical keywords, are likely to result in data that is more worth annotating.

Recommendation 4

Use multilingual models to unlock readily-available data in high-resource languages for initial fine-tuning.

We found that an initial phase of fine-tuning on English data increases model performance on held-out test sets when there is little target-language fine-tuning data (Figure 2). Thus, English data can partly substitute target-language data in few-shot settings. It also substantially improves out-of-domain generalisability to the MHC test suites (Figure 3), with benefits varying by which English dataset is used. Only multilingual models can access these benefits, because they can be fine-tuned on languages other than just the target language. In our experiments, we used only three of the many English hate speech datasets that have been published to date. Of them, we use only one at a time, and only a subset of the full data. Future work could explore the benefits from initial fine-tuning on larger and more diverse datasets, or a combination of datasets in English and other languages that may be more available than the target language.

Recommendation 5

Do not rely on just one held-out test set to evaluate model performance in the target language.

Monolingual models appear strong when evaluated on held-out test sets, and sometimes even outperform multilingual models when there is relatively much target-language fine-tuning data (Figure 2). However, this hides clear weaknesses in their generalisability as measured by MHC (Figure 3). These findings suggest that monolingual models are more at risk of overfitting to dataset-specific features than multilingual models fine-tuned first on data from another language. In higher-resource settings, monolingual models are still likely to outperform multilingual models (Martin et al., 2020; Nozza et al., 2020; Armengol-Estapé et al., 2021).

4 Related Work

Cross-Lingual Transfer Learning Only a very small number of the over 7,000 languages in the world are well-represented in natural language processing resources (Joshi et al., 2020). This has motivated much research into cross-lingual transfer, where large multilingual language models are fine-tuned on one (usually high-resource) source language and then applied to another (usually under-resourced) target language. Cross-lingual transfer has been systematically evaluated for a wide range of NLP tasks and for varying amounts of *shots*, i.e. target-language training data. Evidence on the efficacy of *zero-shot* transfer is mixed, across tasks and for transfer between different language pairs (Artetxe and Schwenk, 2019; Pires et al., 2019; Wu and Dredze, 2019; Lin et al., 2019; Hu et al., 2020; Liang et al., 2020; Ruder et al., 2021, inter alia). In comparison, *few-shot* approaches are generally found to be more effective (Lauscher et al., 2020; Hedderich et al., 2020; Zhao et al., 2021; Hung et al., 2022). In *many-shot* settings, multilingual models are often outperformed by their language-specific monolingual counterparts (Martin et al., 2020; Nozza et al., 2020; Armengol-Estapé et al., 2021). In this paper, we confirmed and expanded on key results from this body of prior work, such as the effectiveness of few-shot learning, specifically for the task of hate speech detection. We also introduced OLS linear regression as an effective method for quantifying data efficiency, i.e. dynamic cost-benefit trade-offs in data annotation.

Cross-Lingual Hate Speech Detection Research on cross-lingual hate speech detection has generally had a more narrow scope than ours. Nozza et al. (2020) demonstrate weaknesses of zero-shot transfer between English, Italian and Spanish. Leite et al. (2020) find similar results for Brazilian Portuguese, Bigoulaeva et al. (2021) for German, and Stappen et al. (2020) for Spanish, while also providing initial evidence for decreasing returns to fine-tuning on more hate speech data (see also Aluru et al., 2020). Pelicon et al. (2021) and Ahn et al. (2020), like Stappen et al. (2020), find that intermediate training on other languages is effective when little target-language data is available. By comparison, we covered a wider range of languages, with two datasets per language and a consistent definition of hate across datasets. We provided a more systematic evaluation of marginal benefits to additional target-language fine-tuning data based on OLS linear regression, as well as a first assessment of out-of-domain performance using the recently-released Multilingual HateCheck test suites. We are also the first to formulate concrete recommendations for expanding hate speech detection into more under-resourced languages.

5 Conclusion

In this paper, we explored strategies for hate speech detection in under-resourced language settings that are effective while also minimising annotation cost and risk of harm to annotators.

We conducted a series of experiments using mono- and multilingual language models fine-tuned on differently-sized random samples of labelled hate speech data in English as well as Arabic, Spanish, Hindi, Portuguese and Italian. We evaluated all models on held-out test sets as well as out-of-domain target-language test suites from Multilingual HateCheck, and used OLS linear regression to quantify the marginal benefits of target-language fine-tuning data. In our experiments, we found that 1) a small amount of target-language fine-tuning data is needed to achieve strong performance on held-out test sets, 2) the benefits of using more such data decrease exponentially, and 3) initial fine-tuning on readily-available English data can partially substitute target-language data and improve model generalisability. Based on our findings, we formulated five actionable recommendations for expanding hate speech detection into under-resourced languages.

Most hate speech research and resources so far have focused on English-language content. However, hate speech is a global phenomenon, and hate speech detection models are needed in hundreds more languages. With our findings and recommendations, we hope that we can facilitate the development of models for yet-unserved languages and thus help to close language gaps that right now leave billions of non-English speakers world less protected against online hate.

Acknowledgments

This article was completed during Paul Röttger’s 2022 research visit to Dirk Hovy’s research group at Bocconi University in Milan. Paul Röttger was funded by the German Academic Scholarship Foundation. Debora Nozza received funding from Fondazione Cariplo (grant No. 2020-4288, MONICA). Federico Bianchi was a member of the Bocconi Institute for Data Science and Analysis (BIDSA) when parts of this work were completed. Dirk Hovy received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No. 949944). Debora Nozza and Dirk Hovy are members of the Data and Marketing Insights (DMI) Unit of the Bocconi Institute for Data Science and Analysis (BIDSA). We thank the MilaNLP Lab and the Pierrehumbert Language Modelling Group for their helpful comments, and all reviewers for their constructive feedback.

Limitations

Our analysis is necessarily constrained by the availability of suitable hate speech datasets, particularly for non-English languages. The datasets we used are idiosyncratic, not just because their language differs but also because they were collected using different methods at different times. This is why we focus on analysing general patterns across datasets in different languages, rather than attempting like-for-like comparisons between individual languages. As new datasets are created, future work could expand our analysis to cover more languages as well as more datasets within each language.

Further, our experiments and the recommendations we formulate based on their results presup-

pose the availability of at least some models tailored to the target language. Multilingual models like XLM-R and XLM-T are respectively pre-trained on 100 and on 30 languages (Conneau et al., 2020; Barbieri et al., 2022), and there are active efforts to expand language model coverage to hundreds of more languages (Wang et al., 2022). However, even the most multilingual models available today cover only a fraction of the world’s languages. Similarly, there is a growing number of monolingual transformer models for different languages (Nozza et al., 2020) but most languages are missing such a model.

Also, we focus on varying amounts of fine-tuning data while assuming the existence of held-out development and test sets to evaluate our experiments. However large or small these sets are, labelling them also creates costs and risks of harm to annotators. This needs to be factored into cost assessments for expanding hate speech detection into an under-resourced language.

For reasons of scope and prioritisation, we did not explore the full range of possibilities for improving cross-lingual model performance. For example, we chose English data for the initial fine-tuning because it is most readily available, but English may not be the optimal choice (Lin et al., 2019). Further, prior work on hate speech detection has found some benefits to data augmentation based on machine translation (Pamungkas et al., 2021; Wang and Banko, 2021) or semi-supervised learning (Zia et al., 2022), which could be explored in future research.

Lastly, our definition of hate speech (§1) is based on a Western legal consensus, in which the presence of hate speech hinges on the target being a protected group. This definition is consistent throughout the datasets we use for our experiments. However, which groups are considered protected may differ across legal and cultural settings. The United Arab Emirates, for example, do not consider sexual orientation a protected characteristic (UAE, 2022). To avoid Eurocentrism, potential differences across legal and cultural settings need to be accounted for when expanding hate speech detection into new languages and when creating new hate speech datasets.

Ethical Considerations

Language Representation Besides English, our experiments cover a limited set of five languages:

Arabic, Spanish, Hindi, Portuguese and Italian. This is due to the very scarcity of suitable non-English datasets for hate speech detection that our work seeks to address. While certainly under-resourced and under-represented in the context of hate speech detection, the five non-English languages we consider are widely spoken, and they have received at least some attention in hate speech research. There is a clear need for creating hate speech datasets in even more under-resourced languages, to facilitate the creation of hate speech detection models in these languages and also to enable even more comprehensive analyses of the kind we conducted here.

Biases in Multilingual Models Zhao et al. (2020) show that multilingual models exhibit similar biases as their monolingual counterparts. Therefore, there is a risk that the models we trained, especially those trained on fewer target-language entries, reflect the biases present in their training and fine-tuning data. Our analyses focused on general model efficacy rather than bias evaluation. Future work could for example make use of Multilingual HateCheck to evaluate model performance across different targeted groups.

Environmental Impact We fine-tuned 3,000 models for our experiments: five model types for each of five target languages with 12 differently-sized samples for 10 random seeds in each language. However, each model fine-tuning step was very computationally inexpensive, especially because we used very small amounts of fine-tuning data. We were able to run all our experimentation on the University of Oxford’s Advanced Research Computing cloud CPUs. Fine-tuning 120 models for a given language took only around four hours for monolingual models and around seven hours for XLM-T. Relative to the concerns raised around the environmental costs of pre-training large language models (Strubell et al., 2019; Henderson et al., 2020; Bender et al., 2021), or even larger-scale fine-tuning with hyperparameter tuning, we therefore consider the environmental costs of our work to be relatively minor. Further, our findings enable the training more effective models with less data and may thus reduce the environmental impact of future work.

References

- Hwijeen Ahn, Jimin Sun, Chan Young Park, and Jungyun Seo. 2020. [NLPDove at SemEval-2020 task 12: Improving offensive language detection with cross-lingual transfer](#). In *Proceedings of the Fourteenth Workshop on Semantic Evaluation*, pages 1576–1586, Barcelona (online). International Committee for Computational Linguistics.
- Sai Saketh Aluru, Binny Mathew, Punyajoy Saha, and Animesh Mukherjee. 2020. [A deep dive into multilingual hate speech classification](#). In *Machine Learning and Knowledge Discovery in Databases. Applied Data Science and Demo Track: European Conference, ECML PKDD 2020, Ghent, Belgium, September 14–18, 2020, Proceedings, Part V*, page 423–439, Berlin, Heidelberg. Springer-Verlag.
- Wissam Antoun, Fady Baly, and Hazem Hajj. 2020. [AraBERT: Transformer-based model for Arabic language understanding](#). In *Proceedings of the 4th Workshop on Open-Source Arabic Corpora and Processing Tools, with a Shared Task on Offensive Language Detection*, pages 9–15, Marseille, France. European Language Resource Association.
- Jordi Armengol-Estapé, Casimiro Pio Carrino, Carlos Rodriguez-Penagos, Ona de Gibert Bonet, Carme Armentano-Oller, Aitor Gonzalez-Agirre, Maite Melero, and Marta Villegas. 2021. [Are multilingual models the best choice for moderately under-resourced languages? A comprehensive assessment for Catalan](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4933–4946, Online. Association for Computational Linguistics.
- Mikel Artetxe and Holger Schwenk. 2019. [Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond](#). *Transactions of the Association for Computational Linguistics*, 7:597–610.
- Somnath Banerjee, Maulindu Sarkar, Nancy Agrawal, Punyajoy Saha, and Mithun Das. 2021. [Exploring transformer based models to identify hate speech and offensive content in English and Indo-Aryan languages](#). *arXiv preprint arXiv:2111.13974*.
- Francesco Barbieri, Luis Espinosa Anke, and Jose Camacho-Collados. 2022. [XLM-T: Multilingual language models in Twitter for sentiment analysis and beyond](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 258–266, Marseille, France. European Language Resources Association.
- Valerio Basile, Cristina Bosco, Elisabetta Fersini, Debora Nozza, Viviana Patti, Francisco Manuel Rangel Pardo, Paolo Rosso, and Manuela Sanguinetti. 2019. [SemEval-2019 task 5: Multilingual detection of hate speech against immigrants and women in Twitter](#). In *Proceedings of the 13th International*

- Workshop on Semantic Evaluation*, pages 54–63, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. [On the dangers of stochastic parrots: Can language models be too big?](#) In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '21, page 610–623, New York, NY, USA. Association for Computing Machinery.
- Irina Bigoulaeva, Viktor Hangya, and Alexander Fraser. 2021. [Cross-lingual transfer learning for hate speech detection](#). In *Proceedings of the First Workshop on Language Technology for Equality, Diversity and Inclusion*, pages 15–25, Kyiv. Association for Computational Linguistics.
- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.
- Leon Derczynski, Hannah Rose Kirk, Abeba Birhane, and Bertie Vidgen. 2022. Handling and presenting harmful text. *arXiv preprint arXiv:2204.14256*.
- Paula Fortuna, João Rocha da Silva, Juan Soler-Company, Leo Wanner, and Sérgio Nunes. 2019. [A hierarchically-labeled Portuguese hate speech dataset](#). In *Proceedings of the Third Workshop on Abusive Language Online*, pages 94–104, Florence, Italy. Association for Computational Linguistics.
- Antigoni Founta, Constantinos Djouvas, Despoina Chatzakou, Ilias Leontiadis, Jeremy Blackburn, Gianluca Stringhini, Athena Vakali, Michael Sirivianos, and Nicolas Kourtellis. 2018. [Large scale crowdsourcing and characterization of Twitter abusive behavior](#). *Proceedings of the International AAAI Conference on Web and Social Media*, 12(1).
- Andrew F Hayes and Li Cai. 2007. [Using heteroskedasticity-consistent standard error estimators in OLS regression: An introduction and software implementation](#). *Behavior research methods*, 39(4):709–722.
- Michael A. Hedderich, David Adelani, Dawei Zhu, Jesusjoba Alabi, Udia Markus, and Dietrich Klakow. 2020. [Transfer learning and distant supervision for multilingual transformer models: A study on African languages](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2580–2591, Online. Association for Computational Linguistics.
- Peter Henderson, Jieru Hu, Joshua Romoff, Emma Brunskill, Dan Jurafsky, and Joelle Pineau. 2020. [Towards the systematic reporting of the energy and carbon footprints of machine learning](#). *Journal of Machine Learning Research*, 21(248):1–43.
- Junjie Hu, Sebastian Ruder, Aditya Siddhant, Graham Neubig, Orhan Firat, and Melvin Johnson. 2020. Xtreme: A massively multilingual multi-task benchmark for evaluating cross-lingual generalisation. In *International Conference on Machine Learning*, pages 4411–4421. PMLR.
- Chia-Chien Hung, Anne Lauscher, Ivan Vulić, Simone Ponzetto, and Goran Glavaš. 2022. [Multi2WOZ: A robust multilingual dataset and conversational pre-training for task-oriented dialog](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3687–3703, Seattle, United States. Association for Computational Linguistics.
- Pratik Joshi, Sebastin Santy, Amar Budhiraja, Kalika Bali, and Monojit Choudhury. 2020. [The state and fate of linguistic diversity and inclusion in the NLP world](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6282–6293, Online. Association for Computational Linguistics.
- Brendan Kennedy, Xisen Jin, Aida Mostafazadeh Davani, Morteza Dehghani, and Xiang Ren. 2020. [Contextualizing hate speech classifiers with post-hoc explanation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5435–5442, Online. Association for Computational Linguistics.
- Douwe Kiela, Max Bartolo, Yixin Nie, Divyansh Kaushik, Atticus Geiger, Zhengxuan Wu, Bertie Vidgen, Grusha Prasad, Amanpreet Singh, Pratik Ringshia, Zhiyi Ma, Tristan Thrush, Sebastian Riedel, Zeerak Waseem, Pontus Stenetorp, Robin Jia, Mohit Bansal, Christopher Potts, and Adina Williams. 2021. [Dynabench: Rethinking benchmarking in NLP](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4110–4124, Online. Association for Computational Linguistics.
- Anne Lauscher, Vinit Ravishankar, Ivan Vulić, and Goran Glavaš. 2020. [From zero to hero: On the limitations of zero-shot language transfer with multilingual Transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4483–4499, Online. Association for Computational Linguistics.
- João Augusto Leite, Diego Silva, Kalina Bontcheva, and Carolina Scarton. 2020. [Toxic language detection in social media for Brazilian Portuguese: New dataset and multilingual analysis](#). In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the*

- 10th International Joint Conference on Natural Language Processing, pages 914–924, Suzhou, China. Association for Computational Linguistics.
- Yaobo Liang, Nan Duan, Yeyun Gong, Ning Wu, Fenfei Guo, Weizhen Qi, Ming Gong, Linjun Shou, Daxin Jiang, Guihong Cao, Xiaodong Fan, Ruofei Zhang, Rahul Agrawal, Edward Cui, Sining Wei, Taroon Bharti, Ying Qiao, Jiun-Hung Chen, Winnie Wu, Shuguang Liu, Fan Yang, Daniel Campos, Rangan Majumder, and Ming Zhou. 2020. [XGLUE: A new benchmark dataset for cross-lingual pre-training, understanding and generation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6008–6018, Online. Association for Computational Linguistics.
- Yu-Hsiang Lin, Chian-Yu Chen, Jean Lee, Zirui Li, Yuyan Zhang, Mengzhou Xia, Shruti Rijhwani, Junxian He, Zhisong Zhang, Xuezhe Ma, Antonios Anastasopoulos, Patrick Littell, and Graham Neubig. 2019. [Choosing transfer languages for cross-lingual learning](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3125–3135, Florence, Italy. Association for Computational Linguistics.
- Delia Marinescu. 2021. Facebook’s content moderation language barrier. *New America*.
- Todor Markov, Chong Zhang, Sandhini Agarwal, Tyna Eloundou, Teddy Lee, Steven Adler, Angela Jiang, and Lilian Weng. 2022. A holistic approach to undesired content detection in the real world. *arXiv preprint arXiv:2208.03274*.
- Louis Martin, Benjamin Muller, Pedro Javier Ortiz Suárez, Yoann Dupont, Laurent Romary, Éric de la Clergerie, Djamel Seddah, and Benoît Sagot. 2020. [CamemBERT: a tasty French language model](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7203–7219, Online. Association for Computational Linguistics.
- Sandip Modha, Thomas Mandl, Gautam Kishore Shahi, Hiren Madhu, Shrey Satapara, Tharindu Ranasinghe, and Marcos Zampieri. 2021. [Overview of the hasoc subtrack at fire 2021: Hate speech and offensive content identification in english and indo-aryan languages and conversational hate speech](#). In *Forum for Information Retrieval Evaluation, FIRE 2021*, page 1–3, New York, NY, USA. Association for Computational Machinery.
- Debora Nozza. 2021. [Exposing the limits of zero-shot cross-lingual hate speech detection](#). 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 907–914, Online. Association for Computational Linguistics.
- Debora Nozza, Federico Bianchi, and Dirk Hovy. 2020. What the [MASK]? making sense of language-specific BERT models. *arXiv preprint arXiv:2003.02912*.
- Nedjma Ousidhoum, Zizheng Lin, Hongming Zhang, Yangqiu Song, and Dit-Yan Yeung. 2019. [Multilingual and multi-aspect hate speech analysis](#). 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 4675–4684, Hong Kong, China. Association for Computational Linguistics.
- Nedjma Ousidhoum, Yangqiu Song, and Dit-Yan Yeung. 2020. [Comparative Evaluation of Label-Agnostic Selection Bias in Multilingual Hate Speech Datasets](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2532–2542. Association for Computational Linguistics.
- Endang Wahyu Pamungkas, Valerio Basile, and Viviana Patti. 2021. [A joint learning approach with knowledge injection for zero-shot cross-lingual hate speech detection](#). *Information Processing & Management*, 58(4):102544.
- Andraž Pelicon, Ravi Shekhar, Blaž Škrlj, Matthew Purver, and Senja Pollak. 2021. [Investigating cross-lingual training for offensive language detection](#). *PeerJ Computer Science*, 7:e559. Publisher: PeerJ Inc.
- Telmo Pires, Eva Schlinger, and Dan Garrette. 2019. [How multilingual is multilingual BERT?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4996–5001, Florence, Italy. Association for Computational Linguistics.
- Fabio Poletto, Valerio Basile, Manuela Sanguinetti, Cristina Bosco, and Viviana Patti. 2021. [Resources and benchmark corpora for hate speech detection: a systematic review](#). *Language Resources and Evaluation*, 55(2):477–523.
- Juan Manuel Pérez, Damián A. Furman, Laura Alonso Alemany, and Franco Luque. 2021. [RoBERTuito: a pre-trained language model for social media text in Spanish](#). *arXiv preprint arXiv:2111.09453*.
- Paul Röttger, Haitham Seelawi, Debora Nozza, Zeerak Talat, and Bertie Vidgen. 2022. [Multilingual HateCheck: Functional tests for multilingual hate speech detection models](#). In *Proceedings of the Sixth Workshop on Online Abuse and Harms (WOAH)*, pages 154–169, Seattle, Washington (Hybrid). Association for Computational Linguistics.
- Paul Röttger, Bertie Vidgen, Dong Nguyen, Zeerak Talat, Helen Margetts, and Janet Pierrehumbert. 2021. [HateCheck: Functional tests for hate speech detection models](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*,

- pages 41–58, Online. Association for Computational Linguistics.
- Sebastian Ruder, Noah Constant, Jan Botha, Aditya Siddhant, Orhan Firat, Jinlan Fu, Pengfei Liu, Junjie Hu, Dan Garrette, Graham Neubig, and Melvin Johnson. 2021. [XTREME-R: Towards more challenging and nuanced multilingual evaluation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10215–10245, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Manuela Sanguinetti, Gloria Comandini, Elisa Di Nuovo, Simona Frenda, Marco Stranisci, Cristina Bosco, Tommaso Caselli, Viviana Patti, and Irene Russo. 2020. [Haspeede 2 @ EVALITA2020: overview of the EVALITA 2020 hate speech detection task](#). In *Proceedings of the Seventh Evaluation Campaign of Natural Language Processing and Speech Tools for Italian. Final Workshop (EVALITA 2020), Online event, December 17th, 2020*, volume 2765 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Tom Simonite. 2021. Facebook is everywhere; its moderation is nowhere close. *Wired*.
- Fábio Souza, Rodrigo Nogueira, and Roberto Lotufo. 2020. [BERTimbau: Pretrained BERT models for Brazilian Portuguese](#). In *Intelligent Systems*, pages 403–417, Cham. Springer International Publishing.
- Lukas Stappen, Fabian Brunn, and Björn W. Schuller. 2020. [Cross-lingual zero- and few-shot hate speech detection utilising frozen transformer language models and AXEL](#). *CoRR*, abs/2004.13850.
- Emma Strubell, Ananya Ganesh, and Andrew McCallum. 2019. [Energy and policy considerations for deep learning in NLP](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3645–3650, Florence, Italy. Association for Computational Linguistics.
- UAE. 2022. [Facebook to pay \\$52m to content moderators over ptsd](#). *u.ae*.
- Bertie Vidgen and Leon Derczynski. 2020. [Directions in abusive language training data, a systematic review: Garbage in, garbage out](#). *PLOS ONE*, 15(12):e0243300.
- Bertie Vidgen, Alex Harris, Dong Nguyen, Rebekah Tromble, Scott Hale, and Helen Margetts. 2019. [Challenges and frontiers in abusive content detection](#). In *Proceedings of the Third Workshop on Abusive Language Online*, pages 80–93, Florence, Italy. Association for Computational Linguistics.
- Bertie Vidgen, Tristan Thrush, Zeerak Talat, and Douwe Kiela. 2021. [Learning from the worst: Dynamically generated datasets to improve online hate detection](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1667–1682, Online. Association for Computational Linguistics.
- Cindy Wang and Michele Banko. 2021. [Practical Transformer-based Multilingual Text Classification](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Industry Papers*, pages 121–129, Online. Association for Computational Linguistics.
- Xinyi Wang, Sebastian Ruder, and Graham Neubig. 2022. [Expanding pretrained models to thousands more languages via lexicon-based adaptation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 863–877, Dublin, Ireland. Association for Computational Linguistics.
- Michael Wiegand, Josef Ruppenhofer, and Thomas Kleinbauer. 2019. [Detection of abusive language: The problem of biased datasets](#). 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 602–608, Minneapolis, Minnesota. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Shijie Wu and Mark Dredze. 2019. [Beto, bentz, becas: The surprising cross-lingual effectiveness of BERT](#). 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 833–844, Hong Kong, China. Association for Computational Linguistics.
- Jieyu Zhao, Subhabrata Mukherjee, Saghar Hosseini, Kai-Wei Chang, and Ahmed Hassan Awadallah. 2020. [Gender bias in multilingual embeddings and cross-lingual transfer](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2896–2907, Online. Association for Computational Linguistics.
- Mengjie Zhao, Yi Zhu, Ehsan Shareghi, Ivan Vulić, Roi Reichart, Anna Korhonen, and Hinrich Schütze. 2021. [A closer look at few-shot crosslingual transfer: The choice of shots matters](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1:*

Long Papers), pages 5751–5767, Online. Association for Computational Linguistics.

Haris Bin Zia, Ignacio Castro, Arkaitz Zubiaga, and Gareth Tyson. 2022. Improving zero-shot cross-lingual hate speech detection with pseudo-label fine-tuning of transformer language models. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 16, pages 1435–1439.

A Details on Model Training

Parameters We implemented XLM-T as well as the five language-specific monolingual models (§2.2) using the HuggingFace transformers library (Wolf et al., 2020). Training batch size was 16. Maximum sequence length was 128 tokens. Otherwise, we used default parameters. The optional first phase of fine-tuning on English data was for three epochs. Fine-tuning on the target language was for five epochs, with subsequent best epoch selection based on macro F1 on the held-out development set.

Computation We ran all experiments on our institution’s 16-core cloud CPUs. Fine-tuning 120 models for a one of the five target languages took around four hours for monolingual models and around seven hours for XLM-T.

B Macro F1 on Held-Out Test Sets

For larger versions of the graphs in Figure 2 from the main body, see Figure 4 on the next page. See Table 5 for a version of Table 1 across all N.

C Macro F1 on MHC

For larger versions of the graphs in Figure 3 from the main body, see Figure 5 below. Similarly, see Table 6 for a version of Table 2 across all N.

D Regression Results for All Models

See Table 4 to the right.

Model: M

Test Set	β_0	β_1	R^2	$2N$
BAS19_ES	0.2879	0.0792	91.97	5.49
FOR19_PT	0.2628	0.0670	71.23	4.65
HAS21_HI	0.2998	0.0607	69.11	4.21
OUS19_ES	0.2092	0.0866	70.08	6.00
SAN20_IT	0.3784	0.0294	77.68	2.04

Model: X

Test Set	β_0	β_1	R^2	$2N$
BAS19_ES	0.2107	0.0856	75.39	5.93
FOR19_PT	0.1727	0.0808	79.17	5.6
HAS21_HI	0.1727	0.0808	79.17	5.6
OUS19_ES	0.1852	0.0869	69.09	6.02
SAN20_IT	0.4004	0.019	40.73	1.32

Model: X+DEN

Test Set	β_0	β_1	R^2	$2N$
BAS19_ES	0.5543	0.0363	87.75	2.52
FOR19_PT	0.5756	0.0205	63.66	1.42
HAS21_HI	0.4691	0.0163	43.91	1.13
OUS19_ES	0.4320	0.0391	64.99	2.71
SAN20_IT	0.5724	0.0272	69.42	1.89

Model: X+FEN

Test Set	β_0	β_1	R^2	$2N$
BAS19_ES	0.4347	0.0489	93.95	3.39
FOR19_PT	0.6167	0.0148	60.96	1.02
HAS21_HI	0.5684	0.0192	45.73	1.33
OUS19_ES	0.6256	0.0191	63.71	1.32
SAN20_IT	0.5227	0.008	7.81	0.56

Model: X+KEN

Test Set	β_0	β_1	R^2	$2N$
BAS19_ES	0.5011	0.0381	93.59	2.64
FOR19_PT	0.5797	0.0182	72.81	1.26
HAS21_HI	0.5545	0.0195	40.66	1.35
OUS19_ES	0.5712	0.026	82.84	1.8
SAN20_IT	0.4444	0.0198	42.5	1.37

Table 4: OLS results for all models in our experiments, expanding on Table 3 from the main body. All regression coefficients β_0 and β_1 are significant at $p < 0.001$.

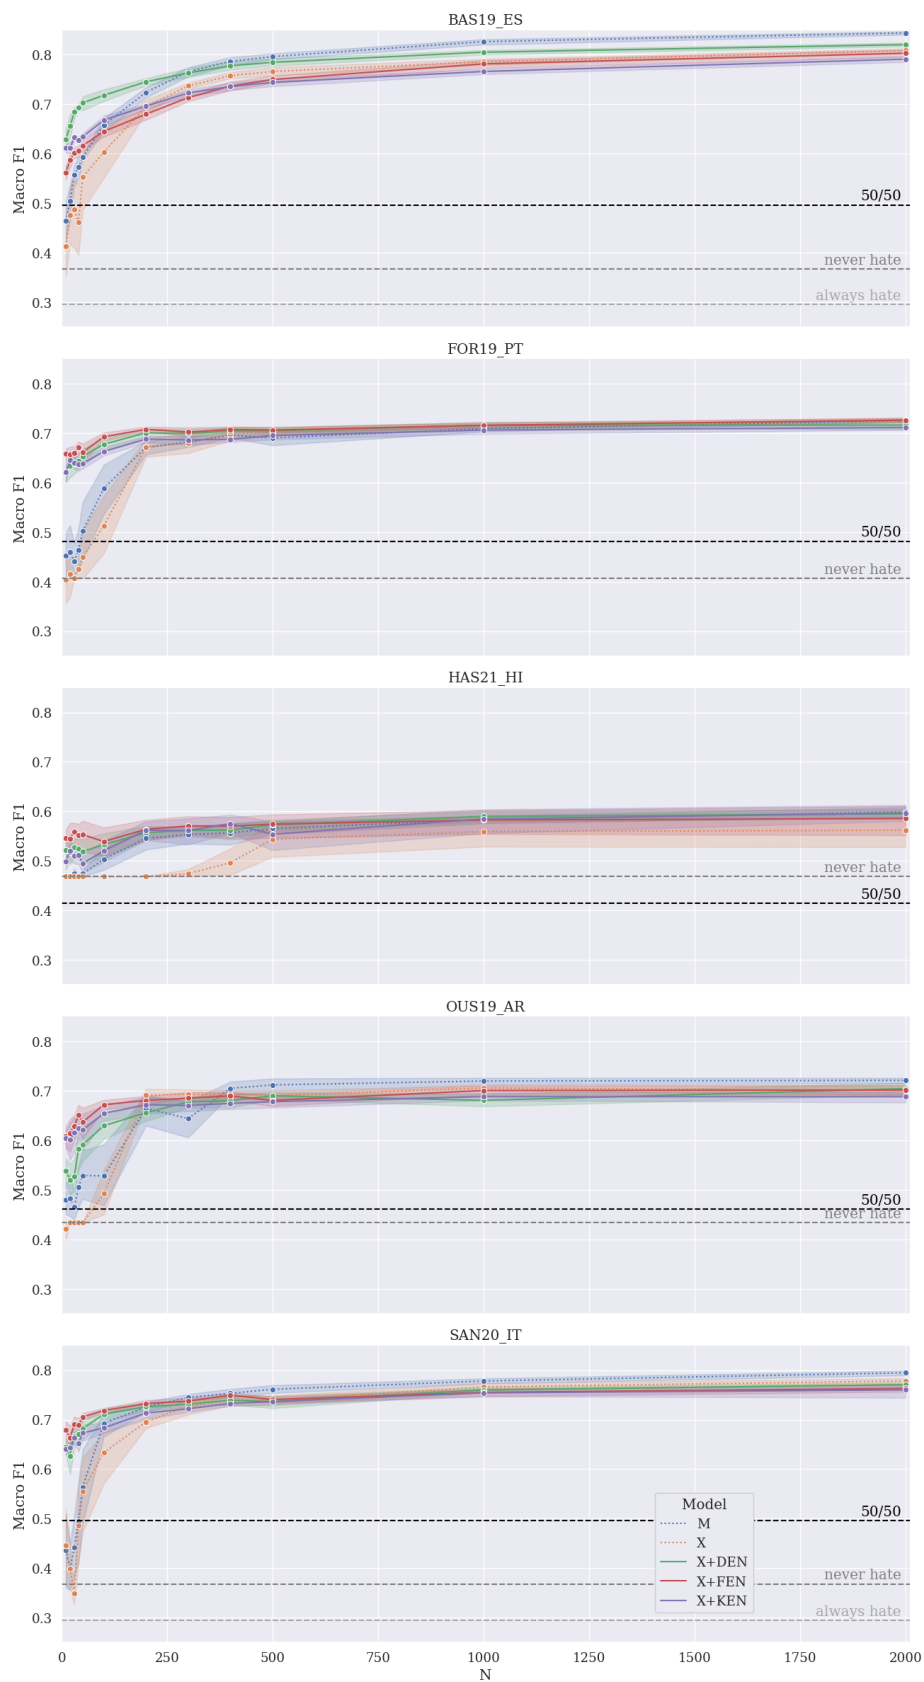


Figure 4: Macro F1 on target-language held-out test sets across models trained on up to $N=2,000$ target-language entries. Model notation as described in §2.2. Confidence intervals and model baselines as described in §2.4.

BAS19_Es												
N	10	20	30	40	50	100	200	300	400	500	1000	2,000
M	0.46	0.50	0.56	0.57	0.59	0.66	0.72	0.77	0.79	0.80	0.83	0.84
X	0.41	0.48	0.49	0.46	0.55	0.60	0.70	0.74	0.76	0.77	0.79	0.81
X+DEN	0.63	0.66	0.68	0.69	0.70	0.72	0.75	0.76	0.78	0.78	0.81	0.82
X+FEN	0.56	0.59	0.60	0.61	0.62	0.65	0.68	0.71	0.74	0.75	0.78	0.80
X+KEN	0.61	0.61	0.63	0.63	0.63	0.67	0.70	0.72	0.74	0.74	0.77	0.79

FOR19_Pt												
N	10	20	30	40	50	100	200	300	400	500	1000	2,000
M	0.45	0.46	0.44	0.46	0.50	0.59	0.67	0.68	0.70	0.69	0.71	0.73
X	0.40	0.42	0.41	0.42	0.45	0.51	0.67	0.68	0.70	0.70	0.71	0.73
X+DEN	0.62	0.63	0.64	0.64	0.65	0.68	0.70	0.70	0.70	0.70	0.72	0.72
X+FEN	0.66	0.66	0.66	0.67	0.66	0.69	0.71	0.70	0.71	0.71	0.72	0.73
X+KEN	0.62	0.65	0.64	0.64	0.64	0.66	0.69	0.69	0.69	0.70	0.71	0.71

HAS21_Hi												
N	10	20	30	40	50	100	200	300	400	500	1000	2,000
M	0.47	0.47	0.47	0.47	0.47	0.50	0.55	0.55	0.56	0.57	0.58	0.60
X	0.47	0.47	0.47	0.47	0.47	0.47	0.47	0.47	0.50	0.54	0.56	0.56
X+DEN	0.52	0.52	0.53	0.53	0.52	0.53	0.56	0.56	0.56	0.57	0.59	0.60
X+FEN	0.55	0.55	0.56	0.55	0.55	0.54	0.56	0.57	0.57	0.57	0.58	0.59
X+KEN	0.50	0.52	0.51	0.51	0.49	0.52	0.56	0.56	0.58	0.55	0.58	0.60

OUS19_AR												
N	10	20	30	40	50	100	200	300	400	500	1000	2,000
M	0.48	0.48	0.46	0.51	0.53	0.53	0.66	0.64	0.70	0.71	0.72	0.72
X	0.42	0.43	0.43	0.43	0.43	0.49	0.69	0.69	0.69	0.69	0.71	0.70
X+DEN	0.54	0.52	0.53	0.58	0.59	0.63	0.66	0.68	0.68	0.69	0.68	0.70
X+FEN	0.61	0.61	0.63	0.65	0.64	0.67	0.68	0.69	0.69	0.68	0.70	0.70
X+KEN	0.60	0.60	0.62	0.62	0.62	0.65	0.67	0.67	0.67	0.68	0.69	0.69

SAN20_IT												
N	10	20	30	40	50	100	200	300	400	500	1000	2,000
M	0.44	0.40	0.44	0.50	0.56	0.69	0.73	0.75	0.75	0.76	0.78	0.79
X	0.45	0.40	0.35	0.49	0.56	0.63	0.70	0.73	0.74	0.74	0.77	0.78
X+DEN	0.64	0.63	0.67	0.67	0.68	0.71	0.73	0.73	0.74	0.74	0.76	0.77
X+FEN	0.68	0.66	0.69	0.69	0.71	0.72	0.73	0.74	0.75	0.74	0.76	0.76
X+KEN	0.64	0.64	0.66	0.65	0.67	0.68	0.71	0.72	0.73	0.74	0.75	0.76

Table 5: Macro F1 on respective held-out test sets for models fine-tuned on N target-language entries, averaged across 10 random seeds for each N. Best performance for a given N in **bold**.

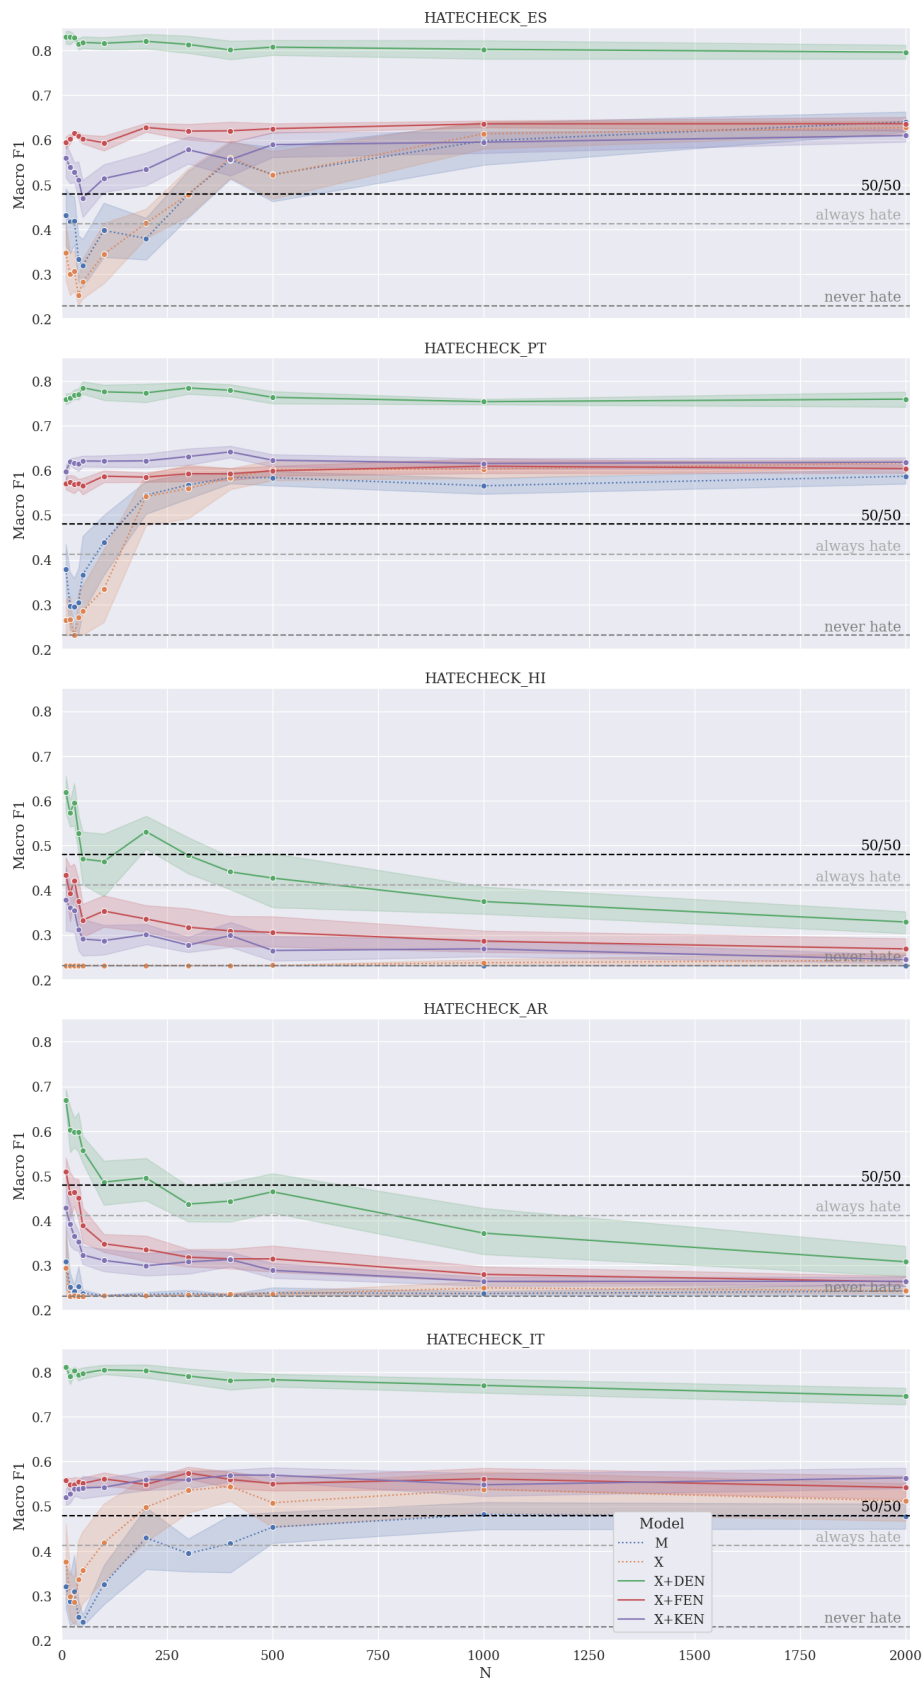


Figure 5: Macro F1 on target-language tests from MHC across models trained on up to $N = 2,000$ target-language entries. Model notation as described in §2.2. Confidence intervals and model baselines as described in §2.4.

HATECHECK_ES												
N	10	20	30	40	50	100	200	300	400	500	1000	2,000
M	0.43	0.42	0.42	0.33	0.32	0.40	0.38	0.48	0.56	0.52	0.60	0.64
X	0.35	0.30	0.31	0.25	0.28	0.34	0.41	0.48	0.56	0.52	0.61	0.63
X+DEN	0.83	0.83	0.83	0.82	0.82	0.82	0.82	0.81	0.80	0.81	0.80	0.80
X+FEN	0.59	0.60	0.61	0.61	0.60	0.59	0.63	0.62	0.62	0.63	0.64	0.64
X+KEN	0.56	0.54	0.53	0.51	0.47	0.51	0.53	0.58	0.56	0.59	0.59	0.61

HATECHECK_PT												
N	10	20	30	40	50	100	200	300	400	500	1000	2,000
M	0.38	0.30	0.30	0.31	0.37	0.44	0.54	0.57	0.58	0.58	0.57	0.59
X	0.27	0.27	0.23	0.27	0.28	0.33	0.54	0.56	0.58	0.60	0.60	0.62
X+DEN	0.76	0.76	0.77	0.77	0.78	0.78	0.77	0.78	0.78	0.76	0.75	0.76
X+FEN	0.57	0.57	0.57	0.57	0.57	0.59	0.58	0.59	0.59	0.60	0.61	0.60
X+KEN	0.60	0.62	0.62	0.62	0.62	0.62	0.62	0.63	0.64	0.62	0.62	0.62

HATECHECK_HI												
N	10	20	30	40	50	100	200	300	400	500	1000	2,000
M	0.23	0.23	0.23	0.23	0.23	0.23	0.23	0.23	0.23	0.23	0.23	0.23
X	0.23	0.23	0.23	0.23	0.23	0.23	0.23	0.23	0.23	0.23	0.24	0.24
X+DEN	0.62	0.57	0.60	0.53	0.47	0.46	0.53	0.48	0.44	0.43	0.37	0.33
X+FEN	0.43	0.39	0.42	0.37	0.33	0.35	0.34	0.32	0.31	0.31	0.29	0.27
X+KEN	0.38	0.36	0.35	0.31	0.29	0.29	0.30	0.28	0.30	0.27	0.27	0.24

HATECHECK_AR												
N	10	20	30	40	50	100	200	300	400	500	1000	2,000
M	0.31	0.25	0.24	0.25	0.24	0.23	0.23	0.24	0.23	0.24	0.24	0.24
X	0.29	0.23	0.23	0.23	0.23	0.23	0.23	0.23	0.24	0.24	0.25	0.24
X+DEN	0.67	0.60	0.60	0.60	0.56	0.49	0.50	0.44	0.44	0.46	0.37	0.31
X+FEN	0.51	0.46	0.46	0.45	0.39	0.35	0.34	0.32	0.31	0.31	0.28	0.26
X+KEN	0.43	0.39	0.37	0.35	0.32	0.31	0.30	0.31	0.31	0.29	0.26	0.26

HATECHECK_IT												
N	10	20	30	40	50	100	200	300	400	500	1000	2,000
M	0.32	0.29	0.31	0.25	0.24	0.33	0.43	0.39	0.42	0.45	0.48	0.48
X	0.38	0.30	0.29	0.34	0.36	0.42	0.50	0.54	0.54	0.51	0.54	0.51
X+DEN	0.81	0.79	0.80	0.79	0.80	0.80	0.80	0.79	0.78	0.78	0.77	0.75
X+FEN	0.56	0.55	0.55	0.55	0.55	0.56	0.55	0.57	0.56	0.55	0.56	0.54
X+KEN	0.52	0.53	0.54	0.54	0.54	0.54	0.56	0.56	0.57	0.57	0.55	0.56

Table 6: Macro F1 on target-language HateCheck for models fine-tuned on N target-language entries, averaged across 10 random seeds for each N. Best performance for a given N in **bold**.