

# Learning from the Worst: Dynamically Generated Datasets to Improve Online Hate Detection

Bertie Vidgen<sup>†</sup>, Tristan Thrush<sup>‡</sup>, Zeerak Waseem<sup>\*</sup>, Douwe Kiela<sup>‡</sup>

<sup>†</sup>The Alan Turing Institute; <sup>\*</sup>University of Sheffield; <sup>‡</sup>Facebook AI Research  
bvidgen@turing.ac.uk

## Abstract

We present a human-and-model-in-the-loop process for dynamically generating datasets and training better performing and more robust hate detection models. We provide a new dataset of  $\sim 40,000$  entries, generated and labelled by trained annotators over four rounds of dynamic data creation. It includes  $\sim 15,000$  challenging perturbations and each hateful entry has fine-grained labels for the type and target of hate. Hateful entries make up 54% of the dataset, which is substantially higher than comparable datasets. We show that model performance is substantially improved using this approach. Models trained on later rounds of data collection perform better on test sets and are harder for annotators to trick. They also have better performance on HATECHECK, a suite of functional tests for online hate detection. We provide the code, dataset and annotation guidelines for other researchers to use.

## 1 Introduction

Accurate detection of online hate speech is important for ensuring that such content can be found and tackled scalably, minimizing the risk that harm will be inflicted on victims and making online spaces more accessible and safe. However, detecting online hate has proven remarkably difficult and concerns have been raised about the performance, robustness, generalisability and fairness of even state-of-the-art models (Waseem et al., 2018; Vidgen et al., 2019a; Caselli et al., 2020; Mishra et al., 2019; Poletto et al., 2020). To address these challenges, we present a human-and-model-in-the-loop process for collecting data and training hate detection models.

Our approach encompasses four rounds of data generation and model training. We first trained a classification model using previously released hate speech datasets. We then tasked annotators with

presenting content that would trick the model and yield *misclassifications*. At the end of the round we trained a new model using the newly presented data. In the next round the process was repeated with the new model in the loop for the annotators to trick. We had four rounds but this approach could, in principle, be continued indefinitely.

Round 1 contains original content created synthetically by annotators. Rounds 2, 3 and 4 are split into half original content and half perturbations. The perturbations are challenging ‘contrast sets’, which manipulate the original text just enough to flip the label (e.g. from ‘Hate’ to ‘Not Hate’) (Kaushik et al., 2019; Gardner et al., 2020). In Rounds 3 and 4 we also tasked annotators with exploring specific types of hate and taking close inspiration from real-world hate sites to make content as adversarial, realistic, and varied as possible.

Models have lower accuracy when evaluated on test sets from later rounds as the content becomes more adversarial. Similarly, the rate at which annotators trick models also decreases as rounds progress (see Table 3). At the same time, models trained on data from later rounds achieve higher accuracy, indicating that their performance improves (see Table 4). We verify improved model performance by evaluating them against the HATECHECK functional tests (Röttger et al., 2020), with accuracy improving from 60% in Round 1 to 95% in Round 4. In this way the models ‘learn from the worst’ because as the rounds progress (a) they become increasingly accurate in detecting hate which means that (b) annotators have to provide more challenging content in order to trick them.

We make three contributions to online hate classification research. First, we present a human-and-model-in-the-loop process for training online hate detection models. Second, we present a dataset of 40,000 entries, of which 54% are hate. It includes fine-grained annotations by trained annotators for

label, type and target (where applicable). Third, we present high quality and robust hate detection models. All data, code and annotation guidelines are available.<sup>1</sup>

## 2 Background

**Benchmark datasets** Several benchmark datasets have been put forward for online hate classification (Waseem and Hovy, 2016; Waseem, 2016; Davidson et al., 2017; Founta et al., 2018; Mandl et al., 2019; Zampieri et al., 2019, 2020; Vidgen et al., 2020, 2021). These datasets offer a comparison point for detection systems and have focused the field’s attention on important subtasks, such as classification across different languages, domains and targets of hate. Performance on some benchmark datasets has increased substantially through the use of more advanced models. For instance, in the original Waseem and Hovy (2016) paper in 2016, the authors achieved an F1 of 0.74. By 2018 this had increased to 0.93 (Pitsilis et al., 2018).

Numerous problems have been identified with hate speech training datasets, such as lacking linguistic variety, being inexpertly annotated and degrading over time (Vidgen et al., 2019a; Poletto et al., 2020). Vidgen and Derczynski (2020) examined 63 open-source abusive language datasets and found that 27 (43%) were sourced from Twitter (Vidgen and Derczynski, 2020). In addition, many datasets are formed with bootstrapped sampling, such as keyword searches, due to the low prevalence of hate speech ‘in the wild’ (Vidgen et al., 2019b). Such bootstrapping can substantially bias the nature and coverage of datasets (Wiegand et al., 2019). Models trained on historical data may also not be effective for present-day hate classification models given how quickly online conversations evolve (Nobata et al., 2016).

**Model limitations** Systems trained on existing datasets have been shown to lack accuracy, robustness and generalisability, creating a range of false positives and false negatives (Schmidt and Wiegand, 2017; Mishra et al., 2019; Vidgen and Derczynski, 2020; Röttger et al., 2020; Mathew et al., 2020). These errors often make models unsuitable for use in downstream tasks, such as moderating online content or measuring online hate.

False positives are non-hateful entries which are incorrectly classified as hateful. Vidgen et al. (2020) report that 29% of errors from a classifier for East Asian prejudice are due to lexical similarities between hateful and non-hateful entries, such as abuse directed towards out-of-scope targets being misclassified as Sinophobic. Other research shows that some identity terms (e.g. ‘gay’) are substantially more likely to appear in toxic content in training datasets, leading models to overfit on them (Dixon et al., 2018; Kennedy et al., 2020). Similarly, many models overfit on the use of slurs and pejorative terms, treating them as hateful irrespective of how they are used (Waseem et al., 2018; Davidson et al., 2017; Kurrek et al., 2020; Palmer et al., 2020). This is problematic when the terms are used as part of counter speech (Wright et al., 2017; Chung et al., 2019) or have been reclaimed by the targeted group (Waseem et al., 2018; Sap et al., 2019). Models can also misclassify interpersonal abuse and incivil language as hateful (Wulczyn et al., 2017a; Zampieri et al., 2019; Palmer et al., 2020).

False negatives are hateful entries which are incorrectly classified as non-hateful. Gröndahl et al. (2018) show that making simple changes such as inserting spelling errors, using leetspeak<sup>2</sup>, changing word boundaries, and appending words can lead to misclassifications of hate. Hosseini et al. (2017) also investigate how detection models can be attacked and report similar findings. In other cases, false negatives can be provoked by changing the ‘sensitive’ attribute of hateful content, such as changing the target from ‘gay’ to ‘black’ people (Garg et al., 2019). This can happen when models are trained on data which only contains hate directed against a limited set of targets (Salminen et al., 2020). Another source of false negatives is when classification systems are applied to out-of-domain settings, such as system trained on Twitter data being applied to data from Gab (Karan and Šnajder, 2018; Pamungkas et al., 2020; Swamy et al., 2019; Basile et al., 2019; Salminen et al., 2020). Subtle and implicit forms of hate speech can also create false negatives (Vidgen and Yasseri, 2019; Palmer et al., 2020; Mathew et al., 2020), as well as more ‘complex’ forms of speech such as sarcasm, irony, adjective nominalization and rhetorical questions (Caselli et al., 2020; Vidgen et al.,

<sup>1</sup><https://github.com/bvidgen/Dynamically-Generated-Hate-Speech-Dataset>

<sup>2</sup>Leetspeak refers to the obfuscation of words by replacing letters with similar looking numbers and symbols.

2019a).

**Dynamic benchmarking and contrast sets** Addressing the numerous flaws of hate detection models is a difficult task. The problem may partly lie in the use of static benchmark datasets and fixed model evaluations. In other areas of Natural Language Processing, several alternative model training and dataset construction paradigms have been presented, involving dynamic and iterative approaches. In a dynamic dataset creation setup, annotators are incentivised to produce high-quality ‘adversarial’ samples which are challenging for baseline models, repeating the process over multiple rounds (Nie et al., 2020). This offers a more targeted way of collecting data. Dinan et al. (2019) ask crowd-workers to ‘break’ a BERT model trained to identify toxic comments and then retrain it using the new examples. Their final model is more robust to complex forms of offensive content, such as entries with figurative language and without profanities.

Another way of addressing the limitations of static datasets is through creating ‘contrast sets’ of perturbations (Kaushik et al., 2019; Gardner et al., 2020). By making minimal label-changing modifications that preserve ‘lexical/syntactic artifacts present in the original example’ (Gardner et al., 2020, p. 1308) the risk of overfitting on spurious correlations is minimized. Perturbations have only received limited attention in the context of hate detection. Samory et al. (2020) create 2,000 ‘hard-to-classify’ not-sexist examples which contrast sexist examples in their dataset. They show that fine-tuning a BERT model with the contrast set produces more robust classification system.

Dynamic benchmarking and contrast sets highlight the effectiveness of developing datasets in a directed and adaptive way, ensuring that models learn from and are evaluated on the most challenging content. However, to date, these approaches remain under-explored for hate speech detection and to the best of our knowledge no prior work in hate speech detection has combined the two approaches within one system.

### 3 Dataset labels

Previous research shows the limitations of using only a binary labelling schema (i.e., ‘Hate’ and ‘Not Hate’). However, there are few established taxonomies and standards in online hate research, and most of the existing datasets have been labelled

with very different schemas. The hierarchical taxonomy we present aims for a balance between granularity versus conceptual distinctiveness and annotation simplicity, following the guidance of Nickserson et al. (2013). All entries are assigned to either ‘Hate’ or ‘Not Hate’. ‘Hate’ is defined as “abusive speech targeting specific group characteristics, such as ethnic origin, religion, gender, or sexual orientation.” (Warner and Hirschberg, 2012). For ‘Hate’, we also annotate secondary labels for the type and target of hate. The taxonomy for the type of hate draws on and extends previous work, including Waseem and Hovy (2016); Vidgen et al. (2019a); Zampieri et al. (2019).

#### 3.1 Types of hate

**Derogation** Content which explicitly attacks, demonizes, demeans or insults a group. This resembles similar definitions from Davidson et al. (2017), who define hate as content that is ‘derogatory’, Waseem and Hovy (2016) who include ‘attacks’ in their definition, and Zampieri et al. (2019) who include ‘insults’.

**Animosity** Content which expresses abuse against a group in an implicit or subtle manner. It is similar to the ‘implicit’ and ‘covert’ categories used in other taxonomies (Waseem et al., 2017; Vidgen and Yasseri, 2019; Kumar et al., 2018).

**Threatening language** Content which expresses intention to, support for, or encourages inflicting harm on a group, or identified members of the group. This category is used in datasets by Hamner (2014), Golbeck et al. (2017) and Anzovino et al. (2018).

**Support for hateful entities** Content which explicitly glorifies, justifies or supports hateful actions, events, organizations, tropes and individuals (collectively, ‘entities’).

**Dehumanization** Content which ‘perceiv[es] or treat[s] people as less than human’ (Haslam and Stratemeyer, 2016). It often involves describing groups as leeches, cockroaches, insects, germs or rats (Mendelsohn et al., 2020).

#### 3.2 Targets of hate

Hate can be targeted against any vulnerable, marginalized or discriminated-against group. We provided annotators with a non-exhaustive list of 29 identities to focus on (e.g., women, black people, Muslims, Jewish people and gay people), as well

as a small number of intersectional variations (e.g., ‘Muslim women’). They are given in Appendix A. Some identities were considered out-of-scope for Hate, including men, white people, and heterosexuals.

## 4 Annotation

Data was annotated using an open-source web platform for dynamic dataset creation and model benchmarking.<sup>3</sup> The platform supports human-and-model-in-the-loop dataset creation for a variety of NLP tasks. Annotation was overseen by two experts in online hate. The annotation process is described in the following section. Annotation guidelines were created at the start of the project and then updated after each round in response to the increased need for detail from annotators. We followed the guidance for protecting and monitoring annotator well-being provided by Vidgen et al. (2019a). 20 annotators were recruited. They received extensive training and feedback during the project. Full details on the annotation team are given in Appendix E. The small pool of annotators was driven by the logistical constraints of hiring and training them to the required standard and protecting their welfare given the sensitivity and complexity of the topic. Nonetheless, it raises the potential for bias. We take steps to address this in our test set construction and provide an annotator ID with each entry in our publicly-released dataset to enable further research into this issue.

## 5 Dataset formation

The dataset was generated over four rounds, each of which involved  $\sim 10,000$  entries. The final dataset comprises 41,255 entries, as shown in Table 1. The ten groups that are targeted most often are given in Table 2. Entries could target multiple groups. After each round, the data was split into training, dev and test splits of 80%, 10% and 10%, respectively. Approximately half of the entries in the test sets are produced by annotators who do not appear in the training and dev sets (between 1 and 4 in each round). This makes the test sets more challenging and minimizes the risk of annotator bias given our relatively small pool of annotators (Geva et al., 2019). The other half of each test set consists of content from annotators who do appear in the training and dev sets.

<sup>3</sup><https://anonymized-url>

Rounds 2, 3 and 4 contain perturbations. In 18 cases the perturbation does not flip the label. This mistake was only identified after completion of the paper and is left in the dataset. These cases can be identified by checking whether original and perturbed entries that have been linked together have the same labels (e.g., whether an original and perturbation are both assigned to ‘Hate’).

**Target model implementation** Every round has a model in the loop, which we call the ‘target model’. The target model is always trained on a combination of data collected in the previous round(s). For instance, **M2** is the target model used in **R2**, and was trained on **R1** and **R0** data. For consistency, we use the same model architecture everywhere, specifically RoBERTa (Liu et al., 2019) with a sequence classification head. We use the implementation from the Transformers (Wolf et al., 2019) library. More details are available in appendix D.

For each new target model, we identify the best sampling ratio of previous rounds’ data using the dev sets. **M1** is trained on **R0** data. **M2** is trained on **R0** data and **R1** upsampled to a factor of five. **M3** is trained on the data used for **M2** and **R2** data upsampled to a factor of one hundred. **M4** is trained on the data used for **M3** and one lot of the **R3** data.

### 5.1 Round 1 (R1)

The target model in **R1** is **M1**, a RoBERTa model trained on **R0** which consists of 11 English language training datasets for hate and toxicity taken from [hatespeechdata.com](https://hatespeechdata.com), as reported in Vidgen and Derczynski (2020). It includes widely-used datasets provided by Waseem (2016), Davidson et al. (2017) and Founta et al. (2018). It comprises 468,928 entries, of which 22% are hateful/toxic. The dataset was anonymized by replacing usernames, indicated by the ‘@’ symbol. URLs were also replaced with a special token. In **R1**, annotators were instructed to enter synthetic content into the model that would trick **M1** using their own creativity and by exploiting any model weaknesses they identified through the real-time feedback.

All entries were validated by one other annotator and entries marked as incorrect were sent for review by expert annotators. This happened with 1,011 entries. 385 entries were excluded for being entirely incorrect. In the other cases, the expert annotator decided the final label and/or made minor



Label	Type	Total	R1	R2	R3	R4
Hate	Not given	7,197	7,197	0	0	0
	Animosity	3,439	0	758	1,206	1,475
	Dehumanization	906	0	261	315	330
	Derogation	9,907	0	3,574	3,206	3,127
	Support	207	0	41	104	62
	Threatening	606	0	376	148	82
	Total	22,262	7,197	5,010	4,979	5,076
Not Hate	/	18,993	3,960	4,986	4,971	5,076
All	TOTAL	41,255	11,157	9,996	9,950	10,152

Table 1: Summary of data collected in each round

Target	Number of entries
Black people	2,278
Women	2,192
Jewish people	1,293
Muslims	1,144
Trans people	972
Gay people	875
Immigrants	823
Disabled people	575
Refugees	533
Arabs	410

Table 2: Most common targets of hate in the dataset

adjustments to the text. The final dataset comprises 11,157 entries of which 7,197 are ‘Hate’ (65%) and 3,960 are ‘Not Hate’ (35%). The type and target of Hate was not recorded by annotators in **R1**.

## 5.2 Round 2 (R2)

A total of 9,996 entries were entered in **R2**. The hateful entries are split between Derogation (3,577, 72%), Dehumanization (255, 5%), Threats (380, 8%), Support for hateful entities (39, 1%) and Animosity (759, 15%). In **R2** we gave annotators adversarial ‘pivots’ to guide their work, which we identified from a review of previous literature (see Section 2). The 10 hateful and 12 not hateful adversarial pivots, with examples and a description, are given in Appendix B. Half of **R2** comprises originally entered content and the other half comprises perturbed contrast sets.

Following Gardner et al. (2020), perturbations were created offline without feedback from a model-in-the-loop. Annotators were given four

main points of guidance: (1) ensure perturbed entries are realistic, (2) firmly meet the criteria of the flipped label and type, (3) maximize diversity within the dataset in terms of type, target and how entries are perturbed and (4) make the least changes possible while meeting (1), (2) and (3). Common strategies for perturbing entries included changing the target (e.g., from ‘black people’ to ‘the local council’), changing the sentiment (e.g. ‘It’s wonderful having gay people round here’), negating an attack (e.g. ‘Muslims are *not* a threat to the UK’) and quoting or commenting on hate.

Of the original entries, those which fooled **M1** were validated by between three and five other annotators. Every perturbation was validated by one other annotator. Annotators could select: (1) correct if they agreed with the label and, for Hate, the type/target, (2) incorrect if the label was wrong or (3) flag if they thought the entry was unrealistic and/or they agreed with the label for hate but disagreed with the type or target. Krippendorff’s alpha is 0.815 for all original entries if all ‘flagged’ entries are treated as ‘incorrect’, indicating extremely high levels of agreement (Hallgren, 2012). All of the original entries identified by at least two validators as incorrect/flagged, and perturbations which were identified by one validator as incorrect/flagged, were sent for review by an expert annotator. This happened in 760 cases in this round.

**Lessons from R2** The validation and review process identified some limitations of the **R2** dataset. First, several ‘template’ statements were entered by annotators. These are entries which have a standardized syntax and/or lexicon, with only the identity changed, such as ‘[Identity] are [negative attribute]’. When there are many cases of each tem-

plate they are easy for the model to correctly classify because they create a simple decision boundary. Discussion sessions showed that annotators used templates (i) to ensure coverage of different identities (an important consideration in making a *generalisable* online hate classifier) and (ii) to maximally exploit model weaknesses to increase their model error rate. We banned the use of templates. Second, in attempting to meet the ‘pivots’ they were assigned, some annotators created unrealistic entries. We updated guidance to emphasize the importance of realism. Third, the pool of 10 trained annotators is large for a project annotating online hate but annotator biases were still produced. Model performance was high in **R2** when evaluated on a training/dev/test split with all annotators stratified. We then held out some annotators’ content and performance dropped substantially. We use this setup for all model evaluations.

### 5.3 Round 3 (R3)

In **R3** annotators were tasked with finding real-world hateful online content to inspire their entries. All real-world content was subject to at least one substantial adjustment prior to being presented to the model. 9,950 entries were entered in **R3**. The hateful entries are split between Derogation (3,205, 64%), Dehumanization (315, 6%), Threats (147, 3%), Support for hateful entities (104, 2%) and Animosity (1,210, 24%). Half of **R3** comprises originally entered content (4,975) and half comprises perturbed contrast sets (4,975).

The same validation procedure was used as with **R2**. Krippendorff’s alpha was 0.55 for all original entries if all ‘flagged’ entries are treated as ‘incorrect’, indicating moderate levels of agreement (Hallgren, 2012). This is lower than **R2**, but still comparable with other hate speech datasets (e.g., Wulczyn et al. (2017b) achieve Krippendorff’s alpha of 0.45). Note that more content is labelled as Animosity compared with **R2** (24% compared with 15%), which tends to have higher levels of disagreement. 981 entries were reviewed by the expert annotators.

### 5.4 Round 4 (R4)

As with **R3**, annotators searched for real-world hateful online content to inspire their entries. In addition, each annotator was given a target identity to focus on (e.g., Muslims, women, Jewish people). The annotators (i) investigated hateful online forums and communities relevant to the target

identity to find the most challenging and nuanced content and (ii) looked for challenging non-hate examples, such as neutral discussions of the identity. 10,152 entries were entered in **R4**, comprising 5,076 ‘Hate’ and 5,076 ‘Not Hate’. The hateful entries are split between Derogation (3,128, 62%), Dehumanization (331, 7%), Threats (82, 2%), Support for hateful entities (61, 1%) and Animosity (1,474, 29%). Half of **R4** comprises originally entered content (5,076) and half comprises perturbed contrast sets (5,076). The same validation procedure was used as in **R2** and **R3**. Krippendorff’s alpha was 0.52 for all original entries if all ‘flagged’ entries are treated as ‘incorrect’, indicating moderate levels of agreement (Hallgren, 2012). This is similar to **R2**. 967 entries were reviewed by the expert annotators following the validation process.

## 6 Model performance

In this section, we examine the performance of models on the collected data, both when used in-the-loop during data collection (measured by the model error rate on new content shown by annotators), as well as when separately evaluated against the test sets in each round’s data. We also examine how models generalize by evaluating them on the out-of-domain suite of diagnostic functional tests in HATECHECK.

### 6.1 Model error rate

The model error rate is the rate at which annotator-generated content tricks the model. It decreases as the rounds progress, as shown in Table 3. **M1**, which was trained on a large set of public hate speech datasets, was the most easily tricked, even though many annotators were learning and had not been given advice on its weaknesses. 54.7% of entries tricked it, including 64.6% of Hate and 49.2% of Not Hate. Only 27.7% of content tricked the final model (**M4**), including 23.7% of Hate and 31.7% of Not Hate. The type of hate affected how frequently entries tricked the model. In general, more explicit and overt forms of hate had the lowest model error rates, with threatening language and dehumanization at 18.2% and 24.8% on average, whereas support for hateful entities and animosity had the highest error (55.4% and 46.4% respectively). The model error rate falls as the rounds progress but nonetheless this metric potentially still underestimates the increasing difficulty of the rounds and the improvement in the models.

Round	Total	Not	Hate	Animosity	Dehuman-ization	Derogation	Support	Threatening
R1	54.7%	64.6%	49.2%	-	-	-	-	-
R2	34.3%	38.9%	29.7%	40.1%	25.5%	28.7%	53.8%	18.4%
R3	27.8%	20.5%	35.1%	53.8%	27.9%	29.2%	59.6%	17.7%
R4	27.7%	23.7%	31.7%	44.5%	21.1%	26.9%	49.2%	18.3%
All	36.6%	35.4%	37.7%	46.4%	24.8%	28.3%	55.4%	18.2%

Table 3: Error rate for target models in each round. Error rate decreases as the rounds progress, indicating that models become harder to trick. Annotators were not given real-time feedback on whether their entries tricked the model when creating perturbations. More information about tuning is available in appendix D

Annotators became more experienced and skilled over the annotation process, and entered progressively more adversarial content. As such the content that annotators enter becomes far harder to classify in the later rounds, which is also reflected in all models’ lower performance on the later round test sets (see Table 4).

## 6.2 Test set performance

Table 4 shows the macro F1 of models trained on different combinations of data, evaluated on the test sets from each round (see Appendix C for dev set performance). The target models achieve lower scores when evaluated on test sets from the later rounds, demonstrating that the dynamic approach to data collection leads to increasingly more challenging data. The highest scores for **R3** and **R4** data are in the mid-70s, compared to the high 70s in **R2** and low 90s in **R1**. Generally, the target models from the later rounds have higher performance across the test sets. For instance, **M4** is the best performing model on **R1**, **R2** and **R4** data. It achieves 75.97 on the **R4** data whereas **M3** achieves 74.83 and **M2** only 60.87. A notable exception is **M1** which outperforms **M2** on the **R3** and **R4** test sets.

Table 4 presents the results for models trained on just the training sets from each round (with no upsampling), indicated by M(RX only). In general the performance is lower than the equivalent target model. For instance, **M4** achieves macro F1 of 75.97 on the **R4** test data. **M(R3 only)** achieves 73.16 on that test set and **M(R4 only)** just 69.6. In other cases, models which are trained on just one round perform well on some rounds but are far worse on others. Overall, building models cumulatively leads to more consistent performance. Table 4 also shows models trained on the cumulative rounds of data with no upsampling, indicated by M(RX+RY). In general, performance is lower with-

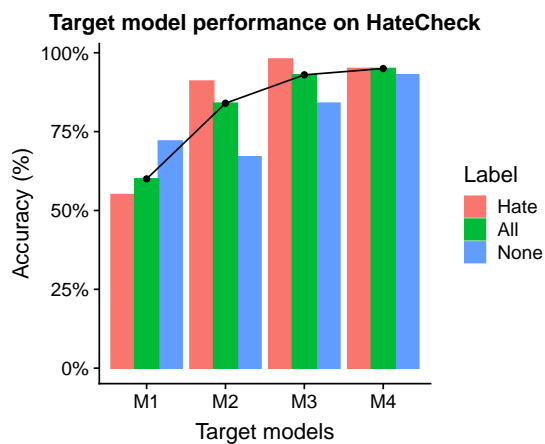


Figure 1: Performance of target models on the HATE-CHECK test suite.

out upsampling; the F1 of **M3** is 2 points higher on the **R3** test set than the equivalent non-upsampled model (**M(R0+R1+R2)**).

## 6.3 HateCheck

To better understand the weaknesses of the target models from each round, we apply them to HATECHECK, as presented by Röttger et al. (2020). HATECHECK is a suite of functional tests for hate speech detection models, based on the testing framework introduced by Ribeiro et al. (2020). It comprises 29 tests, of which 18 correspond to distinct expressions of hate and the other 11 are non-hateful contrasts. The selection of functional tests is motivated by a review of previous literature and interviews with 21 NGO workers. From the 29 tests in the suite, 3,728 labelled entries are generated in the dataset of which 69% are ‘Hate’ and 31% are ‘Not Hate’.

Performance of target models trained on later rounds is substantially higher, increasing from 60% (on both ‘Hate’ and ‘Not Hate’) combined for **M1**

Model	R1	R2	R3	R4
M1 (R1 Target)	44.84±1.1	54.42±0.45	66.07±1.03	60.91±0.4
M2 (R2 Target)	90.17±1.42	66.05±0.67	62.89±1.26	60.87±1.62
M3 (R3 Target)	91.37±1.26	77.14±1.26	<b>76.97±0.49</b>	74.83±0.92
M4 (R4 Target)	<b>92.01±0.6</b>	<b>78.02±0.91</b>	75.89±0.62	<b>75.97±0.96</b>
M(R1 only)	<b>92.20±0.55</b>	62.87±0.63	47.67±1.04	52.37±1.27
M(R2 only)	80.73±0.4	76.52±0.7	<b>77.43±0.51</b>	<b>74.88±0.85</b>
M(R3 only)	72.71±1.05	<b>78.55±0.71</b>	74.14±1.5	73.16±0.58
M(R4 only)	72.26±1.3	76.78±1.65	77.21±0.43	69.6±0.6
M(R0+R1)	88.78±0.89	66.15±0.77	67.15±1.11	63.44±0.26
M(R0+R1+R2)	91.09±0.37	74.73±0.95	74.73±0.46	71.59±0.59
M(R0+R1+R2+R3)	<b>91.17±0.99</b>	77.03±0.72	74.6±0.48	<b>73.94±0.94</b>
M(R0+R1+R2+R3+R4)	90.3±0.96	<b>77.93±0.84</b>	<b>76.79±0.24</b>	72.93±0.56

Table 4: Macro F1 with standard deviation over 5 training rounds, evaluated on each rounds’ test set. Early-stopping is performed on the latest dev set for each round where dev results are obtained at least once per epoch, out of four epochs.

to 95% for **M4**. Performance is better than all four models evaluated by Röttger et al. (2020), of which Perspective’s toxicity classifier<sup>4</sup> is best performing with 77% overall accuracy, including 90% on ‘Hate’ and 48% on ‘Not Hate’. Notably, the performance of **M4** is consistent across both ‘Hate’ and ‘Not Hate’, achieving 95% and 93% respectively. This is in contrast to earlier target models, such as **M2** which achieves 91% on ‘Hate’ but only 67% on ‘Not Hate’ (note that this is actually a *reduction* in performance from **M1** on ‘Not Hate’). Note that HATECHECK only has negative predictive power. These results indicate the *absence* of particular weaknesses in models rather than necessarily characterising generalisable strengths.

A further caveat is that in **R2** the annotators were given adversarial pivots to improve their ability to trick the models (See above). These pivots exploit similar model weaknesses as the functional tests in HATECHECK expose, which creates a risk that this gold standard is not truly independent. We did not identify any exact matches, although after lowering case and removing punctuation there are 21 matches. This is just 0.05% of our dataset but indicates a risk of potential overlap and cross-dataset similarity.

## 7 Discussion

Online hate detection is a complex and nuanced problem, and creating systems that are accurate,

<sup>4</sup>See: <https://www.perspectiveapi.com/#/home>.

robust and generalisable across target, type and domain has proven difficult for AI-based solutions. It requires having datasets which are large, varied, expertly annotated and contain challenging content. Dynamic dataset generation offers a powerful and scalable way of creating these datasets, and training and evaluating more robust and high performing models. Over the four rounds of model training and evaluation we show that the performance of target models improves, as measured by their accuracy on the test sets. The robustness of the target models from later rounds also increases, as shown by their better performance on HATECHECK.

Dynamic data creation systems offer several advantages for training better performing models. First, problems can be addressed as work is conducted – rather than creating the dataset and then discovering any inadvertent design flaws. For instance, we continually worked with annotators to improve their understanding of the guidelines and strategies for tricking the model. We also introduced perturbations to ensure that content was more challenging. Second, annotators can input more challenging content because their work is guided by real-time feedback from the target model. Discussion sessions showed that annotators responded to the models’ feedback in each round, adjusting their content to find better ways to trick it. This process of people trying to find ways to circumvent hate speech models such that their content goes undetected is something that happens often in the real world. Third, dynamic datasets can be



constructed to better meet the requirements of machine learning; our dataset is balanced, comprising ~54% hate. It includes hate targeted against a large number of targets, providing variety for the model to learn from, and many entries were constructed to include known challenging content, such as use of slurs and identity referents.

However, our approach also presents some challenges. First, it requires substantial infrastructure and resources. This project would not have been possible without the use of an online interface and a backend that can serve up state-of-the-art hate speech detection models with relatively low latency. Second, it requires substantial domain expertise from dataset creators as well as annotators, such as knowing where to find real-world hate to inspire synthetic entries. This requires a cross-disciplinary team, combining social science with linguistics and machine learning expertise. Third, evaluating and validating content in a time-constrained dynamic setting can introduce new pressures on the annotation process. The perturbation process also requires additional annotator training, or else might introduce other inadvertent biases.

## 8 Conclusion

We presented a human-and-model-in-the-loop process for training an online hate detection system. It was employed dynamically to collect four rounds of hate speech datasets. The datasets are large and high quality, having been obtained using only expert annotators. They have fine-grained annotations for the type and target of hate, and include perturbations to increase the dataset difficulty. We demonstrated that the models trained on these dynamically generated datasets are much better at the task of hate speech detection, including evaluation on out-of-domain functional test suites.

In future work we aim to expand the size and diversity of the annotator pool for further rounds of dynamic adversarial data collection. We would like to evaluate different models in-the-loop beyond RoBERTa. The datasets also open many new avenues of investigation, including training models on only original entries and evaluating against perturbations (and vice versa) and training multi-label results for type and target of hate. Data collection for future rounds is ongoing.

## Impact Statement & Ethical Considerations

In the Impact Statement we address relevant ethical considerations that were not explicitly discussed in the main body of the paper.

**Data** The entries in the dataset were created by the annotation team and, where needed, reviewed by the expert annotators. In no cases did annotators enter content that they found on online sites. All entries which were closely inspired by real-world content (e.g., data entered during round 4) had substantial adjustments made to them. As such, the data is synthetic.

**Annotator Compensation** We employed a team of twenty annotators to enter content who worked varying hours on a flexible basis over four months. Annotators were compensated at a rate of £16 per hour. The rate was set 50% above the local living wage (£10.85), even though all work was completed remotely. All training time and meetings were paid.

**Intended Use** The approach, dataset and models presented here are intended to support more accurate and robust detection and classification of online hate. We anticipate that the high-quality and fine-grained labels in the dataset will advance research in online hate in other ways, such as enabling multiclass classification of types and targets of online hate.

**Potential Misuse** The dataset and models we present could in principle be used to train a generative hate speech model. Alternatively, the dataset and models could be used to better understand the limitations of current detection tools and then attack them. For instance, if a malicious actor investigated our models then they could better understand what content tricks content moderation tools and then use this knowledge to avoid their content being flagged on social media platforms. However, we believe that these outcomes are unlikely. We do not report any new weaknesses that have not been established in previous research, and the models we present still contain several limitations. Further, it is unlikely that a malicious actor would be able to train a powerful enough generative model from this dataset (given its size and composition) to affect their activities. Overall, the scientific and social benefits of the present research arguably outweighs the small risk of their misuse.

## References

- Maria Anzovino, Elisabetta Fersini, and Paolo Rosso. 2018. [Automatic identification and classification of misogynistic language on Twitter](#). In *NLDB*, pages 57–64.
- Valerio Basile, Cristina Bosco, Elisabetta Fersini, Debora Nozza, Viviana Patti, Francisco Manuel Rangel Pardo, Paolo Rosso, and Manuela Sanguinetti. 2019. [Sem Eval-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.
- Emily M. Bender and Batya Friedman. 2018. [Data statements for natural language processing: Toward mitigating system bias and enabling better science](#). *Transactions of the Association for Computational Linguistics*, 6:587–604.
- Tommaso Caselli, Valerio Basile, Jelena Mitrović, Inga Kartoziya, and Michael Granitzer. 2020. [I Feel Offended, Don't Be Abusive! Implicit/Explicit Messages in Offensive and Abusive Language](#). In *Proceedings of the 12th Conference on Language Resources and Evaluation (LREC)*, pages 6193–6202.
- Yi-Ling Chung, Elizaveta Kuzmenko, Serra Sinem Tekiroglu, and Marco Guerini. 2019. [CONAN - COunter NARratives through Nichesourcing: a Multilingual Dataset of Responses to Fight Online Hate Speech](#). In *Proceedings of the 57th Annual Meeting of the ACL*, pages 2819–2829.
- Thomas Davidson, Dana Warmsley, Michael Macy, and Ingmar Weber. 2017. [Automated Hate Speech Detection and the Problem of Offensive Language](#). In *Proceedings of the 11th ICWSM*, pages 1–4.
- Emily Dinan, Samuel Humeau, Bharath Chintagunta, and Jason Weston. 2019. [Build it break it fix it for dialogue safety: Robustness from adversarial human attack](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*, pages 4537–4546.
- Lucas Dixon, John Li, Jeffrey Sorensen, Nithum Thain, and Lucy Vasserman. 2018. [Measuring and Mitigating Unintended Bias in Text Classification](#). In *AAAI/ACM Conference on AI, Ethics, and Society*, pages 67–73.
- Antigoni-maria 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](#). In *ICWSM*, pages 1–11.
- Matt Gardner, Yoav Artzi, Victoria Basmova, Jonathan Berant, Ben Bogin, Sihao Chen, Pradeep Dasigi, Dheeru Dua, Yanai Elazar, Ananth Gottumukkala, Nitish Gupta, Hanna Hajishirzi, Gabriel Ilharco, Daniel Khoshabi, Kevin Lin, Jiangming Liu, Nelson F. Liu, Phoebe Mulcaire, Qiang Ning, Sameer Singh, Noah A. Smith, Sanjay Subramanian, Reut Tsarfaty, Eric Wallace, Ally Zhang, and Ben Zhou. 2020. [Evaluating NLP Models' Local Decision Boundaries via Contrast Sets](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1307–1323.
- Sahaj Garg, Ankur Taly, Vincent Perot, Ed H. Chi, Nicole Limtiaco, and Alex Beutel. 2019. [Counterfactual fairness in text classification through robustness](#). *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, pages 219–226.
- Mor Geva, Yoav Goldberg, and Jonathan Berant. 2019. [Are we modeling the task or the annotator? an investigation of annotator bias in natural language understanding datasets](#). 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 1161–1166, Hong Kong, China. Association for Computational Linguistics.
- Jennifer Golbeck, Alicia A. Geller, Jayanee Thanki, Shalmali Naik, Kelly M. Hoffman, Derek Michael Wu, Alexandra Berlinger, Priyanka Vengataraman, Shivika Khare, Zahra Ashktorab, Marianna J. Martindale, Gaurav Shahane, Paul Chekalos, Jenny Hottle, Siddharth Bhagwan, Raja Rajan Gnanasekaran, Rajesh Kumar Gnanasekaran, Rashad O. Banjo, Piyush Ramachandran, Lisa Rogers, Kristine M. Rogers, Quint Gergory, Heather L. Nixon, Meghna Sardana Sarin, Zijian Wan, Cody Buntain, Ryan Lau, and Vichita Jienjittler. 2017. [A Large Labeled Corpus for Online Harassment Research](#). In *Proceedings of the ACM Conference on Web Science*, pages 229–233.
- Tommi Gröndahl, Luca Pajola, Mika Juuti, Mauro Conti, and N. Asokan. 2018. [All You Need is "Love": Evading Hate-speech Detection](#). In *Proceedings of the 11th ACM Workshop on Artificial Intelligence and Security*, pages 2–12.
- Kevin A Hallgren. 2012. [Computing Inter-Rater Reliability for Observational Data: An Overview and Tutorial](#). *Tutorials in quantitative methods for psychology*, 8(1):23–34.
- Hugo Lewi Hammer. 2014. [Detecting threats of violence in online discussions using bigrams of important words](#). In *Proceedings - 2014 IEEE Joint Intelligence and Security Informatics Conference, JISIC 2014*, page 319.
- N. Haslam and M. Stratemeyer. 2016. [Recent research on dehumanization](#). *Current Opinion in Psychology*, 11:25–29.
- Hossein Hosseini, Sreeram Kannan, Baosen Zhang, and Radha Poovendran. 2017. [Deceiving Google's Perspective API Built for Detecting Toxic Comments](#). In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 1305–1309.

- Mladen Karan and Jan Šnajder. 2018. [Cross-Domain Detection of Abusive Language Online](#). In *2nd Workshop on Abusive Language Online*, pages 132–137.
- Divyansh Kaushik, Eduard Hovy, and Zachary C Lipton. 2019. Learning the difference that makes a difference with counterfactually-augmented data. *arXiv preprint arXiv:1909.12434*.
- 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.
- Ritesh Kumar, Atul Kr Ojha, Shervin Malmasi, and Marcos Zampieri. 2018. Benchmarking Aggression Identification in Social Media. In *Proceedings of the First Workshop on Trolling, Aggression and Cyberbullying*, 1, pages 1–11.
- Jana Kurrek, Haji Mohammad Saleem, and Derek Ruths. 2020. [Towards a Comprehensive Taxonomy and Large-Scale Annotated Corpus for Online Slur Usage](#). In *Proceedings of the Fourth Workshop on Online Abuse and Harms*, pages 138–149.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv:1907.11692*.
- Thomas Mandl, Sandip Modha, Prasenjit Majumder, Daksh Patel, Mohana Dave, Chintak Mandlia, and Aditya Patel. 2019. [Overview of the HASOC track at FIRE 2019: Hate speech and offensive content identification in Indo-European languages](#). In *FIRE '19: Proceedings of the 11th Forum for Information Retrieval Evaluation*, pages 14–17.
- Binny Mathew, Punyajoy Saha, Seid Muhie Yimam, Chris Biemann, Pawan Goyal, and Animesh Mukherjee. 2020. [Hatexplain: A benchmark dataset for explainable hate speech detection](#).
- Julia Mendelsohn, Yulia Tsvetkov, and Dan Jurafsky. 2020. [A Framework for the Computational Linguistic Analysis of Dehumanization](#). *Frontiers in Artificial Intelligence*, 3(August):1–24.
- Pushkar Mishra, Helen Yannakoudakis, and Ekaterina Shutova. 2019. [Tackling online abuse: A survey of automated abuse detection methods](#). *arXiv:1908.06024v2*, pages 1–17.
- Robert C Nickerson, Upkar Varshney, and Jan Muntermann. 2013. [A method for taxonomy development and its application in information systems](#). *European Journal of Information Systems*, 22(3):336–359.
- Yixin Nie, Adina Williams, Emily Dinan, Mohit Bansal, Jason Weston, and Douwe Kiela. 2020. [Adversarial NLI: A new benchmark for natural language understanding](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4885–4901, Online. Association for Computational Linguistics.
- Chikashi Nobata, Achint Thomas, Yashar Mehdad, Yi Chang, and Joel Tetreault. 2016. [Abusive Language Detection in Online User Content](#). In *World Wide Web Conference*, pages 145–153.
- Alexis Palmer, Christine Carr, Melissa Robinson, and Jordan Sanders. 2020. COLDF: Annotation scheme and evaluation data set for complex offensive language in English. *The Journal for Language Technology and Computational Linguistics*, 34(1):1–28.
- Endang Pamungkas, Valerio Basile, and Viviana Patti. 2020. Misogyny Detection in Twitter: a Multilingual and Cross-Domain Study. *Information Processing and Management*, 57(6).
- Georgios K. Pitsilis, Heri Ramampiaro, and Helge Langseth. 2018. [Detecting Offensive Language in Tweets Using Deep Learning](#). *arXiv:1801.04433v1*, pages 1–17.
- Fabio Poletto, Valerio Basile, Manuela Sanguinetti, Cristina Bosco, and Viviana Patti. 2020. [Resources and benchmark corpora for hate speech detection: a systematic review](#). *Language Resources and Evaluation*, pages 1–47.
- Marco Tulio Ribeiro, Tongshuang Wu, Carlos Guestrin, and Sameer Singh. 2020. [Beyond accuracy: Behavioral testing of NLP models with CheckList](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4902–4912, Online. Association for Computational Linguistics.
- Paul Röttger, Bertram Vidgen, Dong Nguyen, Zeerak Waseem, Helen Margetts, and Janet Pierrehumbert. 2020. [Hatecheck: Functional tests for hate speech detection models](#).
- Joni Salminen, Maximilian Hopf, Shammur A. Chowdhury, Soon gyo Jung, Hind Almerkhi, and Bernard J. Jansen. 2020. [Developing an online hate classifier for multiple social media platforms](#). *Human-centric Computing and Information Sciences*, 10(1):1–34.
- Mattia Samory, Indira Sen, Julian Kohne, Fabian Flock, and Claudia Wagner. 2020. [Unsex me here: Revisiting sexism detection using psychological scales and adversarial samples](#). *arXiv preprint:2004.12764v1*, pages 1–11.
- Maarten Sap, Dallas Card, Saadia Gabriel, Yejin Choi, Noah A Smith, and Paul G Allen. 2019. [The Risk of Racial Bias in Hate Speech Detection](#). In *Proceedings of the 57th Annual Meeting of the ACL*, pages 1668–1678.

- Anna Schmidt and Michael Wiegand. 2017. [A Survey on Hate Speech Detection using Natural Language Processing](#). In *Proceedings of the Fifth International Workshop on Natural Language Processing for Social Media*, pages 1–10. Association for Computational Linguistics.
- Steve Durairaj Swamy, Anupam Jamatia, and Björn Gambäck. 2019. [Studying generalisability across abusive language detection datasets](#). In *CoNLL 2019 - 23rd Conference on Computational Natural Language Learning, Proceedings of the Conference*, pages 940–950.
- Bertie Vidgen, Austin Botelho, David Broniatowski, Ella Guest, Matthew Hall, Helen Margetts, Rebekah Tromble, Zeerak Waseem, and Scott Hale. 2020. [Detecting East Asian Prejudice on Social Media](#). *arXiv:2005.03909v1*, pages 1–12.
- Bertie Vidgen and Leon Derczynski. 2020. [Directions in Abusive Language Training Data: Garbage In, Garbage Out](#). *Plos ONE*, pages 1–26.
- Bertie Vidgen, Alex Harris, Dong Nguyen, Rebekah Tromble, Scott Hale, and Helen Margetts. 2019a. [Challenges and frontiers in abusive content detection](#). In *Proceedings of the Third Workshop on Abusive Language Online (ACL)*, pages 80–93.
- Bertie Vidgen, Helen Margetts, and Alex Harris. 2019b. [How much online abuse is there? A systematic review of evidence for the UK](#). The Alan Turing Institute, London.
- Bertie Vidgen, Dong Nguyen, Helen Margetts, Patricia Rossini, and Rebekah Tromble. 2021. [Introducing CAD: the contextual abuse dataset](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2289–2303, Online. Association for Computational Linguistics.
- Bertie Vidgen and Taha Yasseri. 2019. [Detecting weak and strong Islamophobic hate speech on social media](#). *Journal of Information Technology & Politics*, 17(1):66–78.
- William Warner and Julia Hirschberg. 2012. [Detecting hate speech on the world wide web](#). In *Proceedings of the Second Workshop on Language in Social Media*, pages 19–26.
- Zeerak Waseem. 2016. [Are you a racist or am I seeing things? annotator influence on hate speech detection on twitter](#). In *Proceedings of the First Workshop on NLP and Computational Social Science*, pages 138–142, Austin, Texas. Association for Computational Linguistics.
- Zeerak Waseem, Thomas Davidson, Dana Warmusley, and Ingmar Weber. 2017. [Understanding Abuse: A Typology of Abusive Language Detection Subtasks](#). In *Proceedings of the First Workshop on Abusive Language Online*, pages 78–84.
- Zeerak Waseem and Dirk Hovy. 2016. [Hateful Symbols or Hateful People? Predictive Features for Hate Speech Detection on Twitter](#). In *NAACL-HLT*, pages 88–93.
- Zeerak Waseem, James Thorne, and Joachim Bingel. 2018. [Bridging the gaps: Multi task learning for domain transfer of hate speech detection](#). In Jennifer Golbeck, editor, *Online Harassment*, pages 29–55. Springer International Publishing, Cham.
- Michael Wiegand, Josef Ruppenhofer, and Thomas Kleinbauer. 2019. [Detection of Abusive Language: the Problem of Biased Datasets](#). In *NAACL-HLT*, pages 602–608, Minneapolis. ACL.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, R’emi Louf, Morgan Funtowicz, and Jamie Brew. 2019. [Huggingface’s transformers: State-of-the-art natural language processing](#). *ArXiv*, abs/1910.03771.
- Lucas Wright, Derek Ruths, Kelly P Dillon, Haji Mohammad Saleem, and Susan Benesch. 2017. [Vectors for counterspeech on Twitter](#). In *Proceedings of the First Workshop on Abusive Language Online*, pages 57–62, Vancouver, BC, Canada. Association for Computational Linguistics.
- Ellery Wulczyn, Nithum Thain, and Lucas Dixon. 2017a. [Ex Machina: Personal Attacks Seen at Scale](#). In *Proceedings of the International World Wide Web Conference*, pages 1391–1399.
- Ellery Wulczyn, Nithum Thain, and Lucas Dixon. 2017b. [Ex Machina: Personal Attacks Seen at Scale](#). In *Proceedings of the International World Wide Web Conference*, pages 1391–1399.
- Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Sara Rosenthal, Noura Farra, and Ritesh Kumar. 2019. [Predicting the type and target of offensive posts in social media](#). In *Proceedings of NAACL HLT 2019*, volume 1, pages 1415–1420.
- Marcos Zampieri, Preslav Nakov, Sara Rosenthal, Pepa Atanasova, Georgi Karadzhov, Hamdy Mubarak, Leon Derczynski, Zeses Pitenis, and Çağrı Çöltekin. 2020. [SemEval-2020 Task 12: Multilingual Offensive Language Identification in Social Media \(OffenseEval 2020\)](#). *arXiv preprint*, pages 1–23.



## A List of identities

See Table 5 for a list of identities focused on during data collection.

## B List of pivots in R2

### Hate pivots

**Convoluting statements** Statements which express hate through a convoluted sentence structure, e.g. ‘I’ll make sure that X don’t exist anymore because that’s what I want to do and I don’t care if nobody knows it’.

### Negation of neutrality in the second clause

Statements which initially appear to be neutral (or positive) about an identity, but then undermine it with a caveat, e.g. ‘some people say it’s not fair to call X dumb. . . but 9 times out of 10 it’s true!’.

**Long phrases** Statements which are long, often involving multiple clauses and sentences. In such cases the hate may use more subtle forms of language and/or rely more heavily on tone, e.g. ‘Italy is a beautiful country and I love it so much, from its food to its fashion, from its history to its arts, but I am very worried about the current trends which remind me of the 1930s, So let’s import X to change up the food, fashion, history and arts?’

**Rhetorical questions** Statements which ask a rhetorical question in order to make a hateful point. Their goal is to express hate rather than to elicit an actual answer, e.g. ‘Aren’t all X penny pinching, after all?’

**Misspelling (swapped characters)** Words which are intentionally misspelt so that, although the meaning is clear to a human, the model misclassifies it because the hate is ‘obfuscated’.<sup>5</sup> e.g. *n\*ggga*: ‘N!gg@’, ‘n!ggg’, ‘n1gg4’, *f\*g*: ‘f@g’, ‘f4g’, ‘f@gz’, *k\*ke*: ‘k1ke’, ‘k1k3’.

**Misspelling (elongation)** Words which are misspelt so that they are far longer, with many repeated letters, e.g. ‘I think letting gaaaaaaay people have rights is just soooooo awful’.

**Misspelling (spaces)** Words which are misspelt so that there are spaces between letters. Misspelling of terms includes: ‘T R A N N Y S’, ‘J e w s’, ‘N i g g a’.

<sup>5</sup>Note we use \* for the first vowel in each slur to minimize use of harmful language.

**Synonym (hate expressed with infrequently used terms)** Statements which contain rarely used terms, e.g. rather than saying ‘X are scum’ they use the phrase, ‘X are all sour misanthropes’. Equally, such statements might just direct unusual / less common forms of abuse at the targeted identity. Infrequently used terms to insult a group include: ‘Obnoxious’, ‘Heinous’, ‘Abhorrent’, ‘Unlawful’, ‘Dishonorable’.

**Terms with positive sentiment** Statements which contain positive phrases and terms but still express hate. In particular, phrases which use the term ‘love’ are often identified as not hateful, e.g. ‘I just love it when all of the brown people leave my office so I don’t have to look at them’.

**Random statement** Statements which are ‘random’ and do not contain any features which are likely to trick the model.

### Not Hate pivots

**Counter speech (by another identity)** Statements which challenge, counter or dismiss hate, which are expressed by someone who does *not* identify with the group being attacked. This often involves referring to or quoting hateful content – which models may mistake for genuine hate.

**Counter speech (by target identity)** Statements which challenge, counter or dismiss hate, which are expressed by someone who is from the group being attacked.

**Counter speech (use of negation)** Statements which are neutral and state that a group does not have a (negative or neutral) attribute. This is presented primarily as a contrast set to *Negated non-hate* (in the Hateful pivots), e.g. ‘Jews are not money lovers’.

**Not hateful use of a term that is also a referent for an identity (polysemy)** Statements that mention a referent for an identity (in a way that is not meant to refer to the identity). For instance, using ‘black’ as a colour, e.g. ‘the black cat down the road really ruins the neighbourhood’.

**Use of profanities** Statements which contain a profanity in a not hateful way. They are often used as an adjective or adverb to describe an emotion or to place emphasis on what is being said, e.g. ‘fucking hell today was a lot of bullshit’.

**Negativity against objects** Statements which attack, criticise or express negativity against inanimate objects, such as sofa or a car, e.g. ‘this cup is totally rubbish’.

**Personal abuse (direct)** Statements which are aggressive, insulting or abusive against an individual using a direct personal pronoun (i.e. ‘you’), e.g. ‘you are a complete joke and no-one respects you, loser’.

**Personal abuse (indirect)** Statements which are aggressive, insulting or abusive against an individual who is not part of the conversation and as such is referred to with an indirect personal pronoun (i.e. ‘he’, ‘she’, ‘they’), e.g. ‘he is such a waste of space. I hope he dies’.

**Negativity against concepts** Statements which attack, criticise or express negativity against concepts and ideologies, such as political ideologies, economic ideas and philosophical ideals, e.g. ‘I’ve never trusted capitalism. It’s bullshit and it fucks society over’.

**Negativity against animals** Statements which attack, criticise or express negativity against animals, e.g. ‘dogs are just beasts, kick them if they annoy you’.

**Negativity against institutions** Statements which attack, criticise or express negativity against institutions; such as large organisations, governments and bodies, e.g. ‘the NHS is a badly run and pointless organisation which is the source of so much harm’.

**Negativity against others** Statements which attack, criticise or express negativity against something that is NOT an identity – and the targets are not identified elsewhere in this typology, e.g. ‘the air round here is toxic, it smells like terrible’.

## C Development set performance

Table 6 shows dev set performance numbers.

## D Model, Training, and Evaluation Details

The model architecture was the roberta-base model from Huggingface (<https://huggingface.co/>), with a sequence classification head. This model has approximately 125 million parameters. Training each model took no longer than approximately a

day, on average, with 8 GPUs on the FAIR cluster. All models were trained with a learning rate of  $2e-5$  with the default optimizer that Huggingface’s sequence classification routine uses. Target model hyperparameter search was as follows: the R2 target was trained for 3 epochs on the R1 target training data, plus multiples of the round 1 data from {1, 5, 10, 20, 40, 100} (the best was 5). The R3 target was trained for 3 epochs on the R2 target training data, plus multiples of the round 2 data from {1, 5, 10, 20, 40, 100} (the best was 100). The R4 target was trained on the R3 target training data for 4 epochs, plus multiples of the round 3 data from {1, 5, 10, 20, 40, 100, 200} (the best was 1); early stopping based on loss on the dev set (measured multiple times per epoch) was performed. The dev set we used for tuning target models was the latest dev set we had at each round. We did not perform hyperparameter search on the non-target models, with the exception of training 5 seeds of each and early stopping based on dev set loss throughout 4 training epochs. We recall that model performance typically did not vary by much more than 5% through our hyperparameter searches.

## E Data statement

Following [Bender and Friedman \(2018\)](#) we provide a data statement, which documents the process and provenance of the final dataset.

**A. CURATION RATIONALE** In order to study the potential of dynamically generated datasets for improving online hate detection, we used an online interface to generate a large-scale synthetic dataset of 40,000 entries, collected over 4 rounds, with a ‘model-in-the-loop’ design. Data was not sampled. Instead a team of trained annotators created synthetic content to enter into the interface.

**B. LANGUAGE VARIETY** All of the content is in English. We opted for English language due to the available annotation team, and resources and the project leaders’ expertise. The system that we developed could, in principle, be applied to other languages.

**C. SPEAKER DEMOGRAPHICS** Due to the synthetic nature of the dataset, the speakers are the same as the annotators.

**D. ANNOTATOR DEMOGRAPHICS** Annotator demographics are reported in the paper, and

are reproduced here for fullness. Annotation guidelines were created at the start of the project and then updated after each round. Annotations guidelines will be publicly released with the dataset. We followed the guidance for protecting and monitoring annotator well-being provided by Vidgen et al. (2019a). 20 annotators were recruited. Ten were recruited to work for 12 weeks and ten were recruited for the final four weeks. Annotators completed different amounts of work depending on their availability, which is recorded in the dataset.

All annotators attended a project onboarding session, half day training session, at least one one-to-one session and a daily 'standup' meeting when working. They were given a test assignment and guidelines to review before starting work and received feedback after each round. Annotators could ask the experts questions in real-time over a messaging platform.

Of the 20 annotators, 35% were male and 65% were female. 65% were 18-29 and 35% were 30-39. 10% were educated to high school level, 20% to undergraduate, 45% to taught masters and 25% to research degree (i.e. PhD). 70% were native English speakers and 30% were non-native but fluent. Respondents had a range of nationalities, including British (60%), as well as Polish, Spanish and Iraqi. Most annotators identified as ethnically white (70%), followed by Middle Eastern (20%) and two or more ethnicities (10%). Participants all used social media regularly, including 75% who used it more than once per day. All participants had seen other people targeted by online abuse before, and 55% had been targeted personally.

**E. SPEECH SITUATION** All data was created from 21st September 2020 until 14th January 2021. During the project annotators visited a range of online platforms, with adequate care and supervision from the project leaders, to better understand online hate as it appears online.

**F. TEXT CHARACTERISTICS** The composition of the dataset, including the distribution of the Primary label (Hate and Not) and the type (Derogation, Animosity, Threatening, Support, Dehumanization and None Given) is described in the paper.

Category of identity	Identity
Disability	People with disabilities
Gender	Gender minorities (e.g. non binary)
Gender	Women
Gender	Trans
Immigration status	Immigrants
Immigration status	Foreigner
Immigration status	Refugee
Immigration status	Asylum seeker
Race / Ethnicity	Black people
Race / Ethnicity	Indigenous
Race / Ethnicity	East Asians (e.g. China)
Race / Ethnicity	East Asians (e.g. Korea)
Race / Ethnicity	South East Asians (e.g. India)
Race / Ethnicity	Pakistanis
Race / Ethnicity	Aboriginal people (e.g. Native Americans)
Race / Ethnicity	Mixed race
Race / Ethnicity	Minority groups
Race / Ethnicity	Arabs
Race / Ethnicity	Travellers (e.g. Roma)
Race / Ethnicity	People from Africa
Religion or belief	Muslims
Religion or belief	Jews
Sexual orientation	Gay
Sexual orientation	Lesbian
Sexual orientation	Bisexual
National origin	Polish
Religion or belief	Hindus
Class	Working class
Race / Ethnicity	Hispanic (e.g. Latinx)
Intersectional	Black women
Intersectional	Black men
Intersectional	Indigenous women
Intersectional	Asian women
Intersectional	Muslim women

Table 5: List of high priority identities

Model	R1	R2	R3	R4
M1 (R1 Target)	41.4±0.91	61.06±0.43	58.18±0.69	55.46±0.63
M2 (R2 Target)	<b>95.38±0.25</b>	68.86±0.71	66.46±1.09	63.17±0.8
M3 (R3 Target)	94.55±0.65	85.04±0.63	76.77±0.57	74.4±0.9
M4 (R4 Target)	94.92±0.45	<b>85.32±0.29</b>	<b>77.52±0.68</b>	<b>76.42±0.82</b>
M(R1)	<b>95.69±0.29</b>	61.88±0.98	57.75±0.8	58.54±0.52
M(R2)	81.28±0.2	<b>84.36±0.4</b>	75.8±0.55	<b>74.29±1.05</b>
M(R3)	76.79±1.18	79.6±0.99	75.5±0.48	74.19±1.07
M(R4)	78.05±1.09	80.21±0.31	<b>75.63±0.49</b>	72.54±0.64
M(R0+R1)	<b>93.92±0.3</b>	69.43±1.58	65.48±0.48	63.99±0.74
M(R0+R1+R2)	93.13±0.24	82.82±0.8	73.66±0.75	72.28±0.84
M(R0+R1+R2+R3)	93.43±0.39	84.66±0.6	75.81±0.29	<b>75.85±1.0</b>
M(R0+R1+R2+R3+R4)	92.73±0.82	<b>86.0±0.69</b>	<b>77.0±0.59</b>	75.7±0.69

Table 6: Macro F1 with standard deviation over 5 training rounds, evaluated on each rounds’ dev set. Early-stopping is performed on the latest development set for each round where dev results are obtained at least once per epoch, out of four epochs.