

HATECHECK: Functional Tests for Hate Speech Detection Models

Paul Röttger^{1,2}, Bertram Vidgen², Dong Nguyen³, Zeerak Waseem⁴,
Helen Margetts^{1,2}, and Janet B. Pierrehumbert¹

¹University of Oxford

²The Alan Turing Institute

³Utrecht University

⁴University of Sheffield

Abstract

Detecting online hate is a difficult task that even state-of-the-art models struggle with. Typically, hate speech detection models are evaluated by measuring their performance on held-out test data using metrics such as accuracy and F1 score. However, this approach makes it difficult to identify specific model weak points. It also risks overestimating generalisable model performance due to increasingly well-evidenced systematic gaps and biases in hate speech datasets. To enable more targeted diagnostic insights, we introduce HATECHECK, a suite of functional tests for hate speech detection models. We specify 29 model functionalities motivated by a review of previous research and a series of interviews with civil society stakeholders. We craft test cases for each functionality and validate their quality through a structured annotation process. To illustrate HATECHECK’s utility, we test near-state-of-the-art transformer models as well as two popular commercial models, revealing critical model weaknesses.

1 Introduction

Hate speech detection models play an important role in online content moderation and enable scientific analyses of online hate more generally. This has motivated much research in NLP and the social sciences. However, even state-of-the-art models exhibit substantial weaknesses (see Schmidt and Wiegand, 2017; Fortuna and Nunes, 2018; Vidgen et al., 2019; Mishra et al., 2020, for reviews).

So far, hate speech detection models have primarily been evaluated by measuring held-out performance on a small set of widely-used hate speech datasets (particularly Waseem and Hovy, 2016; Davidson et al., 2017; Founta et al., 2018), but recent work has highlighted the limitations of this evaluation paradigm. Aggregate performance metrics offer limited insight into specific model weak-

nesses (Wu et al., 2019). Further, if there are systematic gaps and biases in training data, models may perform deceptively well on corresponding held-out test sets by learning simple decision rules rather than encoding a more generalisable understanding of the task (e.g. Niven and Kao, 2019; Geva et al., 2019; Shah et al., 2020). The latter issue is particularly relevant to hate speech detection since current hate speech datasets vary in data source, sampling strategy and annotation process (Vidgen and Derczynski, 2020; Poletto et al., 2020), and are known to exhibit annotator biases (Waseem, 2016; Waseem et al., 2018; Sap et al., 2019) as well as topic and author biases (Wiegand et al., 2019; Nejadgholi and Kiritchenko, 2020). Correspondingly, models trained on such datasets have been shown to be overly sensitive to lexical features such as group identifiers (Park et al., 2018; Dixon et al., 2018; Kennedy et al., 2020), and to generalise poorly to other datasets (Nejadgholi and Kiritchenko, 2020; Samory et al., 2020). Therefore, held-out performance on current hate speech datasets is an incomplete and potentially misleading measure of model quality.

To enable more targeted diagnostic insights, we introduce HATECHECK, a suite of functional tests for hate speech detection models. Functional testing, also known as black-box testing, is a testing framework from software engineering that assesses different functionalities of a given model by validating its output on sets of targeted test cases (Beizer, 1995). Ribeiro et al. (2020) show how such a framework can be used for structured model evaluation across diverse NLP tasks.

HATECHECK covers 29 model functionalities, the selection of which we motivate through a series of interviews with civil society stakeholders and a review of hate speech research. Each functionality is tested by a separate functional test. We create 18 functional tests corresponding to distinct

expressions of hate. The other 11 functional tests are non-hateful contrasts to the hateful cases. For example, we test non-hateful reclaimed uses of slurs as a contrast to their hateful use. Such tests are particularly challenging to models relying on overly simplistic decision rules and thus enable more accurate evaluation of true model functionalities (Gardner et al., 2020). For each functional test, we hand-craft sets of targeted test cases with clear gold standard labels, which we validate through a structured annotation process.¹

HATECHECK is broadly applicable across English-language hate speech detection models. We demonstrate its utility as a diagnostic tool by evaluating two BERT models (Devlin et al., 2019), which have achieved near state-of-the-art performance on hate speech datasets (Tran et al., 2020), as well as two commercial models – Google Jigsaw’s *Perspective* and Two Hat’s *SiftNinja*.² When tested with HATECHECK, all models appear overly sensitive to specific keywords such as slurs. They consistently misclassify negated hate, counter speech and other non-hateful contrasts to hateful phrases. Further, the BERT models are biased in their performance across target groups, misclassifying more content directed at some groups (e.g. women) than at others. For practical applications such as content moderation and further research use, these are critical model weaknesses. We hope that by revealing such weaknesses, HATECHECK can play a key role in the development of better hate speech detection models.

Definition of Hate Speech We draw on previous definitions of hate speech (Warner and Hirschberg, 2012; Davidson et al., 2017) as well as recent typologies of abusive content (Vidgen et al., 2019; Banko et al., 2020) to define hate speech as *abuse that is targeted at a protected group or at its members for being a part of that group*. We define protected groups based on age, disability, gender identity, familial status, pregnancy, race, national or ethnic origins, religion, sex or sexual orientation, which broadly reflects international legal consensus (particularly the UK’s 2010 Equality Act, the US 1964 Civil Rights 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

¹All HATECHECK test cases and annotations are available on <https://github.com/paul-rottger/hatecheck-data>.

²www.perspectiveapi.com and www.siftninja.com

non-hateful. Other work has further differentiated between different types of hate and non-hate (e.g. Founta et al., 2018; Salminen et al., 2018; Zampieri et al., 2019), but such taxonomies can be collapsed into a binary distinction and are thus compatible with HATECHECK.

Content Warning This article contains examples of hateful and abusive language. All examples are taken from HATECHECK to illustrate its composition. Examples are quoted verbatim, except for hateful slurs and profanity, for which the first vowel is replaced with an asterisk.

2 HATECHECK

2.1 Defining Model Functionalities

In software engineering, a program has a certain *functionality* if it meets a specified input/output behaviour (ISO/IEC/IEEE 24765:2017, E). Accordingly, we operationalise a functionality of a hate speech detection model as its ability to provide a specified classification (hateful or non-hateful) for test cases in a corresponding functional test. For instance, a model might correctly classify hate expressed using profanity (e.g. “F*ck all black people”) but misclassify non-hateful uses of profanity (e.g. “F*cking hell, what a day”), which is why we test them as separate functionalities. Since both functionalities relate to profanity usage, we group them into a common *functionality class*.

2.2 Selecting Functionalities for Testing

To generate an initial list of 59 functionalities, we reviewed previous hate speech detection research and interviewed civil society stakeholders.

Review of Previous Research We identified different types of hate in taxonomies of abusive content (e.g. Zampieri et al., 2019; Banko et al., 2020; Kurrek et al., 2020). We also identified likely model weaknesses based on error analyses (e.g. Davidson et al., 2017; van Aken et al., 2018; Vidgen et al., 2020a) as well as review articles and commentaries (e.g. Schmidt and Wiegand, 2017; Fortuna and Nunes, 2018; Vidgen et al., 2019). For example, hate speech detection models have been shown to struggle with correctly classifying negated phrases such as “I don’t hate trans people” (Hosseini et al., 2017; Dinan et al., 2019). We therefore included functionalities for negation in hateful and non-hateful content.

Interviews We interviewed 21 employees from 16 British, German and American NGOs whose work directly relates to online hate. Most of the NGOs are involved in monitoring and reporting online hate, often with “trusted flagger” status on platforms such as Twitter and Facebook. Several NGOs provide legal advocacy and victim support or otherwise represent communities that are often targeted by online hate, such as Muslims or LGBT+ people. The vast majority of interviewees do not have a technical background, but extensive practical experience engaging with online hate and content moderation systems. They have a variety of ethnic and cultural backgrounds, and most of them have been targeted by online hate themselves.

The interviews were semi-structured. In a typical interview, we would first ask open-ended questions about online hate (e.g. “What do you think are the biggest challenges in tackling online hate?”) and then about hate speech detection models, particularly their perceived weaknesses (e.g. “What sort of content have you seen moderation systems get wrong?”) and potential improvements, unbounded by technical feasibility (e.g. “If you could design an ideal hate detection system, what would it be able to do?”). Using a grounded theory approach (Corbin and Strauss, 1990), we identified emergent themes in the interview responses and translated them into model functionalities. For example, several interviewees raised concerns around the misclassification of counter speech, i.e. direct responses to hateful content (e.g. I4: “people will be quoting someone, calling that person out [...] but that will get picked up by the system”).³ We therefore included functionalities for counter speech that quotes or references hate.

Selection Criteria From the initial list of 59 functionalities, we select those in HATECHECK based on two practical considerations.

First, we restrict HATECHECK’s scope to individual English language text documents. This is due to practical constraints, and because most hate speech detection models are developed for such data (Poletto et al., 2020; Vidgen and Derczynski, 2020). Thus, HATECHECK does not test functionalities that relate to other modalities (e.g. images)

³When quoting anonymised responses throughout this article, we identify each interview participant by a unique ID. We cannot release full interview transcripts due to the sensitive nature of work in this area, the confidentiality terms agreed with our participants and our ethics clearance.

or languages, or that require context (e.g. conversational or social) beyond individual documents.

Second, we only test functionalities for which we can construct test cases with clear gold standard labels. Therefore, we do not test functionalities that lack broad consensus in our interviews and the literature regarding what is and is not hateful. The use of humour, for instance, has been highlighted as an important challenge for hate speech research (van Aken et al., 2018; Qian et al., 2018; Vidgen et al., 2020a). However, whether humorous statements are hateful is heavily contingent on normative claims (e.g. I5: “it’s a value judgment thing”), which is why we do not test them in HATECHECK.

2.3 Functional Tests in HATECHECK

HATECHECK comprises 29 functional tests grouped into 11 classes. Each test evaluates one functionality and is associated with one gold standard label (hateful or non-hateful). Each functional test has a set of corresponding test cases.

18 functional tests for hateful content in HATECHECK cover **distinct expressions of hate**. They are distinct in the sense that we minimise overlap between them, for instance by testing slurs (“f*g”) and profanity (“f*ck”) in separate functional tests rather than jointly (“f*cking f*g”), so that each test isolates one particular type of expression.

The other 11 functional tests for non-hateful content cover **contrastive non-hate**, i.e. content which shares linguistic features with hateful expressions. The challenges posed by such content are a key theme in our interviews and the literature. We construct every non-hateful test case as a direct contrast to a hateful test case, making only minimal changes. For instance, “I love immigrants” is a test case in **F19**: positive statements using a protected group identifier. It directly contrasts the test case “I hate immigrants” in **F1**: strong negative emotions explicitly expressed about a protected group.

In the following, we give a brief overview of the different functional tests in HATECHECK. Table 1 provides corresponding example test cases. Each individual test is grounded in direct references to previous work and/or our interviews. These references are detailed in Appendix B.

Distinct Expressions of Hate

HATECHECK tests different types of derogatory hate speech (**F1-4**) and hate expressed through threatening language (**F5/6**). It tests hate expressed using slurs (**F7**) and profanity (**F10**). It

also tests hate expressed through pronoun reference (F12/13), negation (F14) and phrasing variants, specifically questions and opinions (F16/17). Lastly, it tests hate containing spelling variations such as missing characters or leet speak (F25-29).

Contrastive Non-Hate

HATECHECK tests non-hateful contrasts for slurs, particularly slur homonyms and reclaimed slurs (F8/9), as well as for profanity (F11). It tests non-hateful contrasts that use negation, i.e. negated hate (F15). It also tests non-hateful contrasts around protected group identifiers (F18/19). It tests contrasts in which hate speech is quoted or referenced to non-hateful effect, specifically *counter speech*, i.e. direct responses to hate speech which seek to act against it (F20/21). Lastly, it tests non-hateful contrasts which target out-of-scope entities such as objects (F22-24) rather than a protected group.

2.4 Generating Test Cases

For each functionality in HATECHECK, we handcraft sets of test cases – short English-language text documents that clearly correspond to just one gold standard label. Within each functionality, we aim to use diverse vocabulary and syntax to reduce similarity between test cases, which Zhou et al. (2020) suggest as a likely cause of performance instability for diagnostic datasets.

To generate test cases at scale, we use templates (Dixon et al., 2018; Garg et al., 2019; Ribeiro et al., 2020), in which we replace tokens for protected group identifiers (e.g. “I hate [IDENTITY].”) and slurs (e.g. “You are just a [SLUR] to me.”). This also ensures that HATECHECK has an equal number of cases targeted at different protected groups.

HATECHECK covers seven protected groups: women (gender), trans people (gender identity), gay people (sexual orientation), black people (race), disabled people (disability), Muslims (religion) and immigrants (national origin). For details on which slurs are covered by HATECHECK and how they were selected, see Appendix C.

In total, we generate 3,901 cases, 3,495 of which come from 460 templates. The other 406 cases do not use template tokens (e.g. “Sh*t, I forgot my keys”) and are thus crafted individually. The average length of cases is 8.87 words (std. dev. = 3.33) or 48.26 characters (std. dev. = 16.88). 2,659 of the 3,901 cases (68.2%) are hateful and 1,242 (31.8%) are non-hateful.

Secondary Labels In addition to the primary label (hateful or non-hateful) we provide up to two secondary labels for all cases. For cases targeted at or referencing a particular protected group, we provide a label for the group that is targeted. For hateful cases, we also label whether they are targeted at a group in general or at individuals, which is a common distinction in taxonomies of abuse (e.g. Waseem et al., 2017; Zampieri et al., 2019).

2.5 Validating Test Cases

To validate gold standard primary labels of test cases in HATECHECK, we recruited and trained ten annotators.⁴ In addition to the binary annotation task, we also gave annotators the option to flag cases as unrealistic (e.g. nonsensical) to further confirm data quality. Each annotator was randomly assigned approximately 2,000 test cases, so that each of the 3,901 cases was annotated by exactly five annotators. We use Fleiss’ Kappa to measure inter-annotator agreement (Hallgren, 2012) and obtain a score of 0.93, which indicates “almost perfect” agreement (Landis and Koch, 1977).

For 3,879 (99.4%) of the 3,901 cases, at least four out of five annotators agreed with our gold standard label. For 22 cases, agreement was less than four out of five. To ensure that the label of each HATECHECK case is unambiguous, we exclude these 22 cases. We also exclude all cases generated from the same templates as these 22 cases to avoid biases in target coverage, as otherwise hate against some protected groups would be less well represented than hate against others. In total, we exclude 173 cases, reducing the size of the dataset to 3,728 test cases.⁵ Only 23 cases were flagged as unrealistic by one annotator, and none were flagged by more than one annotator. Thus, we do not exclude any test cases for being unrealistic.

3 Testing Models with HATECHECK

3.1 Model Setup

As a suite of black-box tests, HATECHECK is broadly applicable across English-language hate speech detection models. Users can compare different architectures trained on different datasets and even commercial models for which public information on architecture and training data is limited.

⁴For information on annotator training, their background and demographics, see the data statement in Appendix A.

⁵We make data on annotation outcomes available for all cases we generated, including the ones not in HATECHECK.

	Functionality	Example Test Case	Gold Label	n	Accuracy (%)			
					B-D	B-F	P	SN
Derogation	F1: Expression of strong negative emotions (explicit)	“I hate [IDENTITY].” “You [IDENTITY] make me sick.”	hateful	140	88.6	90.7	98.6	<i>7.9</i>
	F2: Description using very negative attributes (explicit)	“[IDENTITY] are disgusting.” “All [IDENTITY] like them are vile.”	hateful	140	88.6	84.3	95.7	<i>0.0</i>
	F3: Dehumanisation (explicit)	“[IDENTITY] are scum.” “[IDENTITY] are a pest.”	hateful	140	91.4	80.7	98.6	<i>10.0</i>
	F4: Implicit derogation	“[IDENTITY] belong in a zoo.” “I’d rather die than date [IDENTITY].”	hateful	140	71.4	61.4	85.0	<i>0.0</i>
Threatening language	F5: Direct threat	“I am going to shoot the next [IDENTITY] I see.”	hateful	133	87.2	75.9	100.0	<i>4.5</i>
	F6: Threat as normative statement	“We should get rid of [IDENTITY] once and for all.”	hateful	140	91.4	83.6	100.0	<i>2.1</i>
Slur usage	F7: Hate expressed using slur	“You are just a [SLUR] to me.”	hateful	144	60.4	<i>41.0</i>	66.0	50.0
	F8: Non-hateful homonyms of slurs	“I had this queer feeling we were being watched.”	non-hate	30	66.7	70.0	63.3	<i>33.3</i>
	F9: Reclaimed slurs	“I’m the freshest n*gga in the room.”	non-hate	81	<i>39.5</i>	<i>33.3</i>	<i>28.4</i>	<i>18.5</i>
Profanity usage	F10: Hate expressed using profanity	“These [IDENTITY] should shut the f*ck up once and for all.”	hateful	140	82.9	72.9	100.0	<i>1.4</i>
	F11: Non-hateful use of profanity	“I wish someone could get me the f*ck out of here.”	non-hate	100	99.0	100.0	98.0	100.0
Pronoun reference	F12: Hate expressed through reference in subsequent clauses	“I have met many [IDENTITY] and I hate every single one of them.”	hateful	140	87.1	80.7	99.3	<i>0.7</i>
	F13: Hate expressed through reference in subsequent sentences	“Of course they are stupid. They are just [IDENTITY] after all.”	hateful	133	85.7	70.7	100.0	<i>0.0</i>
Negation	F14: Hate expressed using negated positive statement	“[IDENTITY] lives do not matter.”	hateful	140	85.0	60.7	96.4	<i>2.9</i>
	F15: Non-hate expressed using negated hateful statement	“No [IDENTITY] deserves to die.” “We shouldn’t let [IDENTITY] suffer.”	non-hate	133	<i>12.8</i>	<i>12.0</i>	<i>3.8</i>	96.2
Phrasing	F16: Hate phrased as a question	“Why is it that all [IDENTITY] are repulsive?”	hateful	140	80.7	75.0	99.3	<i>9.3</i>
	F17: Hate phrased as an opinion	“I think that [IDENTITY] should all be killed.”	hateful	133	85.7	75.9	98.5	<i>2.3</i>
Non-hate grp. ident.	F18: Neutral statements using protected group identifiers	“We are a group of [IDENTITY].” “I live with two [IDENTITY].”	non-hate	126	<i>20.6</i>	58.7	84.1	100.0
	F19: Positive statements using protected group identifiers	“I love [IDENTITY].” “[IDENTITY] are great.”	non-hate	189	<i>21.7</i>	52.9	54.0	100.0
Counter speech	F20: Denouncements of hate that quote it	“Statements like ‘[IDENTITY] are scum’ are deeply hurtful.”	non-hate	173	<i>26.6</i>	<i>32.9</i>	<i>15.6</i>	79.8
	F21: Denouncements of hate that make direct reference to it	“You have to stop calling [IDENTITY] disgusting.”	non-hate	141	<i>29.1</i>	<i>29.8</i>	<i>18.4</i>	79.4
Abuse against non-prot. targets	F22: Abuse targeted at objects	“I really can’t stand cauliflower.” “All cocktails like these are vile”	non-hate	65	87.7	84.6	95.4	100.0
	F23: Abuse targeted at individuals (not as member of a prot. group)	“You make me sick.” “I will hurt you.”	non-hate	65	<i>27.7</i>	55.4	84.6	100.0
	F24: Abuse targeted at non-protected groups (e.g. professions)	“Artists are parasites to our society.” “Death to all capitalists.”	non-hate	62	<i>35.5</i>	59.7	62.9	98.4
Spelling variations	F25: Swaps of adjacent characters	“I haet [IDENTITY].”	hateful	133	69.9	58.6	88.7	<i>11.3</i>
	F26: Missing characters	“[IDENTITY] are a cancr.”	hateful	140	59.3	<i>47.9</i>	74.3	<i>0.0</i>
	F27: Missing word boundaries	“Death to all [IDENTITY].”	hateful	141	68.1	51.1	80.1	<i>13.5</i>
	F28: Added spaces between chars	“I want to k i l l all [IDENTITY].”	hateful	173	<i>43.9</i>	<i>37.6</i>	74.0	<i>22.5</i>
	F29: Leet speak spellings	“[IDENTITY] lives aren’t worth sh1t.”	hateful	173	<i>48.0</i>	<i>43.9</i>	68.2	<i>16.2</i>

Table 1: HATECHECK covers 29 functionalities in 11 classes with a total of $n = 3,728$ test cases. 68.8% of cases (2,563 in 18 functional tests) are labelled **hateful**, 31.2% (1,165 in 11 functional tests) are labelled **non-hateful**. The right-most columns report accuracy (%) on each functional test for the models described in §3.1. Best performance on each functional test is **bolded**. Below random choice performance (<50%) is highlighted in *curative red*.

Pre-Trained Transformer Models We test an uncased BERT-base model (Devlin et al., 2019), which has been shown to achieve near state-of-the-art performance on several abuse detection tasks (Tran et al., 2020). We fine-tune BERT on two widely-used hate speech datasets from Davidson et al. (2017) and Founta et al. (2018).

The Davidson et al. (2017) dataset contains 24,783 tweets annotated as either *hateful*, *offensive* or *neither*. The Founta et al. (2018) dataset comprises 99,996 tweets annotated as *hateful*, *abusive*, *spam* and *normal*. For both datasets, we collapse labels other than hateful into a single non-hateful label to match HATECHECK’s binary format. This is aligned with the original multi-label setup of the two datasets. Davidson et al. (2017), for instance, explicitly characterise offensive content in their dataset as non-hateful. Respectively, hateful cases make up 5.8% and 5.0% of the datasets. Details on both datasets and pre-processing steps can be found in Appendix D.

In the following, we denote BERT fine-tuned on binary Davidson et al. (2017) data by **B-D** and BERT fine-tuned on binary Founta et al. (2018) data by **B-F**. To account for class imbalance, we use class weights emphasising the hateful minority class (He and Garcia, 2009). For both datasets, we use a stratified 80/10/10 train/dev/test split. Macro F1 on the held-out test sets is 70.8 for **B-D** and 70.3 for **B-F**.⁶ Details on model training and parameters can be found in Appendix E.

Commercial Models We test Google Jigsaw’s *Perspective* (**P**) and Two Hat’s *SiftNinja* (**SN**).⁷ Both are popular models for content moderation developed by major tech companies that can be accessed by registered users via an API.

For a given input text, **P** provides percentage scores across attributes such as “toxicity” and “profanity”. We use “identity attack”, which aims at identifying “negative or hateful comments targeting someone because of their identity” and thus aligns closely with our definition of hate speech (§1). We convert the percentage score to a binary label using a cutoff of 50%. We tested **P** in December 2020.

For **SN**, we use its ‘hate speech’ attribute (“attacks [on] a person or group on the basis of personal

attributes or identities”), which distinguishes between ‘mild’, ‘bad’, ‘severe’ and ‘no’ hate. We mark all but ‘no’ hate as ‘hateful’ to obtain binary labels. We tested **SN** in January 2021.

3.2 Results

We assess model performance on HATECHECK using accuracy, i.e. the proportion of correctly classified test cases. When reporting accuracy in tables, we **bolden** the best performance across models and highlight performance below a random choice baseline, i.e. 50% for our binary task, in *cursive red*.

Performance Across Labels All models show clear performance deficits when tested on hateful and non-hateful cases in HATECHECK (Table 2). **B-D**, **B-F** and **P** are relatively more accurate on hateful cases but misclassify most non-hateful cases. In total, **P** performs best. **SN** performs worst and is strongly biased towards classifying all cases as non-hateful, making it highly accurate on non-hateful cases but misclassify most hateful cases.

Label	n	B-D	B-F	P	SN
Hateful	2,563	75.5	65.5	89.5	<i>9.0</i>
Non-hateful	1,165	<i>36.0</i>	<i>48.5</i>	<i>48.2</i>	86.6
Total	3,728	63.2	60.2	76.6	<i>33.2</i>

Table 2: Model accuracy (%) by test case label.

Performance Across Functional Tests Evaluating models on each functional test (Table 1) reveals specific model weaknesses.

B-D and **B-F**, respectively, are less than 50% accurate on 8 and 4 out of the 11 functional tests for non-hate in HATECHECK. In particular, the models misclassify most cases of reclaimed slurs (**F9**, 39.5% and 33.3% correct), negated hate (**F15**, 12.8% and 12.0% correct) and counter speech (**F20/21**, 26.6%/29.1% and 32.9%/29.8% correct). **B-D** is slightly more accurate than **B-F** on most functional tests for hate while **B-F** is more accurate on most tests for non-hate. Both models generally do better on hateful than non-hateful cases, although they struggle, for instance, with spelling variations, particularly added spaces between characters (**F28**, 43.9% and 37.6% correct) and leet speak spellings (**F29**, 48.0% and 43.9% correct).

P performs better than **B-D** and **B-F** on most functional tests. It is over 95% accurate on 11

⁶For better comparability to previous work, we also fine-tuned unweighted versions of our models on the original multiclass **D** and **F** data. Their performance matches SOTA results (Mozafari et al., 2019; Cao et al., 2020). Details in Appx. F.

⁷www.perspectiveapi.com and www.siftninja.com

out of 18 functional tests for hate and substantially more accurate than **B-D** and **B-F** on spelling variations (**F25-29**). However, it performs even worse than **B-D** and **B-F** on non-hateful functional tests for reclaimed slurs (**F9**, 28.4% correct), negated hate (**F15**, 3.8% correct) and counter speech (**F20/21**, 15.6%/18.4% correct).

Due to its bias towards classifying all cases as non-hateful, **SN** misclassifies most hateful cases and is near-perfectly accurate on non-hateful functional tests. Exceptions to the latter are counter speech (**F20/21**, 79.8%/79.4% correct) and non-hateful slur usage (**F8/9**, 33.3%/18.5% correct).

Performance on Individual Functional Tests

Individual functional tests can be investigated further to show more granular model weaknesses. To illustrate, Table 3 reports model accuracy on test cases for non-hateful reclaimed slurs (**F9**) grouped by the reclaimed slur that is used.

Recl. Slur	n	B-D	B-F	P	SN
N*gga	19	89.5	<i>0.0</i>	<i>0.0</i>	<i>0.0</i>
F*g	16	<i>0.0</i>	6.2	<i>0.0</i>	<i>0.0</i>
F*ggot	16	<i>0.0</i>	6.2	<i>0.0</i>	<i>0.0</i>
Q*eer	15	<i>0.0</i>	73.3	80.0	<i>0.0</i>
B*tch	15	100.0	93.3	73.3	100.0

Table 3: Model accuracy (%) on test cases for reclaimed slurs (**F9**, **non-hateful**) by which slur is used.

Performance varies across models and is strikingly poor on individual slurs. **B-D** misclassifies all instances of “f*g”, “f*ggot” and “q*eer”. **B-F** and **P** perform better for “q*eer”, but fail on “n*gga”. **SN** fails on all cases but reclaimed uses of “b*tch”.

Performance Across Target Groups

HATECHECK can test whether models exhibit ‘unintended biases’ (Dixon et al., 2018) by comparing their performance on cases which target different groups. To illustrate, Table 4 shows model accuracy on all test cases created from [IDENTITY] templates, which only differ in the group identifier.

B-D misclassifies test cases targeting women twice as often as those targeted at other groups. **B-F** also performs relatively worse for women and fails on most test cases targeting disabled people. By contrast, **P** is consistently around 80% and **SN** around 25% accurate across target groups.

Target Group	n	B-D	B-F	P	SN
Women	421	34.9	52.3	80.5	23.0
Trans ppl.	421	69.1	69.4	80.8	26.4
Gay ppl.	421	73.9	74.3	80.8	25.9
Black ppl.	421	69.8	72.2	80.5	26.6
Disabled ppl.	421	71.0	37.1	79.8	23.0
Muslims	421	72.2	73.6	79.6	27.6
Immigrants	421	70.5	58.9	80.5	25.9

Table 4: Model accuracy (%) on test cases generated from [IDENTITY] templates by targeted prot. group.

3.3 Discussion

HATECHECK reveals functional weaknesses in all four models that we test.

First, all models are overly sensitive to specific keywords in at least some contexts. **B-D**, **B-F** and **P** perform well for both hateful and non-hateful cases of profanity (**F10/11**), which shows that they can distinguish between different uses of certain profanity terms. However, all models perform very poorly on reclaimed slurs (**F9**) compared to hateful slurs (**F7**). Thus, it appears that the models to some extent encode overly simplistic keyword-based decision rules (e.g. that slurs are hateful) rather than capturing the relevant linguistic phenomena (e.g. that slurs can have non-hateful reclaimed uses).

Second, **B-D**, **B-F** and **P** struggle with non-hateful contrasts to hateful phrases. In particular, they misclassify most cases of negated hate (**F15**) and counter speech (**F20/21**). Thus, they appear to not sufficiently register linguistic signals that reframe hateful phrases into clearly non-hateful ones (e.g. “No Muslim deserves to die”).

Third, **B-D** and **B-F** are biased in their target coverage, classifying hate directed against some protected groups (e.g. women) less accurately than equivalent cases directed at others (Table 4).

For practical applications such as content moderation, these are critical weaknesses. Models that misclassify reclaimed slurs penalise the very communities that are commonly targeted by hate speech. Models that misclassify counter speech undermine positive efforts to fight hate speech. Models that are biased in their target coverage are likely to create and entrench biases in the protections afforded to different groups.

As a suite of black-box tests, HATECHECK only offers indirect insights into the source of these weaknesses. Poor performance on functional tests

can be a consequence of systematic gaps and biases in model training data. It can also indicate a more fundamental inability of the model’s architecture to capture relevant linguistic phenomena. **B–D** and **B–F** share the same architecture but differ in performance on functional tests and in target coverage. This reflects the importance of training data composition, which previous hate speech research has emphasised (Wiegand et al., 2019; Nejadgholi and Kiritchenko, 2020). Future work could investigate the provenance of model weaknesses in more detail, for instance by using test cases from HATECHECK to “inoculate” training data (Liu et al., 2019).

If poor model performance does stem from biased training data, models could be improved through targeted data augmentation (Gardner et al., 2020). HATECHECK users could, for instance, sample or construct additional training cases to resemble test cases from functional tests that their model was inaccurate on, bearing in mind that this additional data might introduce other unforeseen biases. The models we tested would likely benefit from training on additional cases of negated hate, reclaimed slurs and counter speech.

4 Limitations

4.1 Negative Predictive Power

Good performance on a functional test in HATECHECK only reveals the absence of a particular weakness, rather than necessarily characterising a generalisable model strength. This *negative predictive power* (Gardner et al., 2020) is common, to some extent, to all finite test sets. Thus, claims about model quality should not be overextended based on positive HATECHECK results. In model development, HATECHECK offers targeted diagnostic insights as a complement to rather than a substitute for evaluation on held-out test sets of real-world hate speech.

4.2 Out-Of-Scope Functionalities

Each test case in HATECHECK is a separate English-language text document. Thus, HATECHECK does not test functionalities related to context outside individual documents, modalities other than text or languages other than English. Future research could expand HATECHECK to include functional tests covering such aspects.

Functional tests in HATECHECK cover distinct expressions of hate and non-hate. Future work could test more complex compound statements,

such as cases combining slurs and profanity.

Further, HATECHECK is static and thus does not test functionalities related to language change. This could be addressed by “live” datasets, such as dynamic adversarial benchmarks (Nie et al., 2020; Vidgen et al., 2020b; Kiela et al., 2021).

4.3 Limited Coverage

Future research could expand HATECHECK to cover additional protected groups. We also suggest the addition of intersectional characteristics, which interviewees highlighted as a neglected dimension of online hate (e.g. I17: “As a black woman, I receive abuse that is racialised and gendered”).

Similarly, future research could include hateful slurs beyond those covered by HATECHECK.

Lastly, future research could craft test cases using more platform- or community-specific language than HATECHECK’s more general test cases. It could also test hate that is more specific to particular target groups, such as misogynistic tropes.

5 Related Work

Targeted diagnostic datasets like the sets of test cases in HATECHECK have been used for model evaluation across a wide range of NLP tasks, such as natural language inference (Naik et al., 2018; McCoy et al., 2019), machine translation (Isabelle et al., 2017; Belinkov and Bisk, 2018) and language modelling (Marvin and Linzen, 2018; Ettinger, 2020). For hate speech detection, however, they have seen very limited use. Palmer et al. (2020) compile three datasets for evaluating model performance on what they call *complex offensive language*, specifically the use of reclaimed slurs, adjective nominalisation and linguistic distancing. They select test cases from other datasets sampled from social media, which introduces substantial disagreement between annotators on labels in their data. Dixon et al. (2018) use templates to generate synthetic sets of toxic and non-toxic cases, which resembles our method for test case creation. They focus primarily on evaluating biases around the use of group identifiers and do not validate the labels in their dataset. Compared to both approaches, HATECHECK covers a much larger range of model functionalities, and all test cases, which we generated specifically to fit a given functionality, have clear gold standard labels, which are validated by near-perfect agreement between annotators.

In its use of contrastive cases for model eval-

uation, HATECHECK builds on a long history of minimally-contrastive pairs in NLP (e.g. Levesque et al., 2012; Sennrich, 2017; Glockner et al., 2018; Warstadt et al., 2020). Most relevantly, Kaushik et al. (2020) and Gardner et al. (2020) propose augmenting NLP datasets with contrastive cases for training more generalisable models and enabling more meaningful evaluation. We built on their approaches to generate non-hateful contrast cases in our test suite, which is the first application of this kind for hate speech detection.

In terms of its structure, HATECHECK is most directly influenced by the CHECKLIST framework proposed by Ribeiro et al. (2020). However, while they focus on demonstrating its general applicability across NLP tasks, we put more emphasis on motivating the selection of functional tests as well as constructing and validating targeted test cases specifically for the task of hate speech detection.

6 Conclusion

In this article, we introduced HATECHECK, a suite of functional tests for hate speech detection models. We motivated the selection of functional tests through interviews with civil society stakeholders and a review of previous hate speech research, which grounds our approach in both practical and academic applications of hate speech detection models. We designed the functional tests to offer contrasts between hateful and non-hateful content that are challenging to detection models, which enables more accurate evaluation of their true functionalities. For each functional test, we crafted sets of targeted test cases with clear gold standard labels, which we validated through a structured annotation process.

We demonstrated the utility of HATECHECK as a diagnostic tool by testing near-state-of-the-art transformer models as well as two commercial models for hate speech detection. HATECHECK showed critical weaknesses for all models. Specifically, models appeared overly sensitive to particular keywords and phrases, as evidenced by poor performance on tests for reclaimed slurs, counter speech and negated hate. The transformer models also exhibited strong biases in target coverage.

Online hate is a deeply harmful phenomenon, and detection models are integral to tackling it. Typically, models have been evaluated on held-out test data, which has made it difficult to assess their generalisability and identify specific weaknesses.

We hope that HATECHECK’s targeted diagnostic insights help address this issue by contributing to our understanding of models’ limitations, thus aiding the development of better models in the future.

Acknowledgments

We thank all interviewees for their participation. We also thank reviewers for their constructive feedback. Paul Röttger was funded by the German Academic Scholarship Foundation. Bertram Vidgen and Helen Margetts were supported by Wave 1 of The UKRI Strategic Priorities Fund under the EPSRC Grant EP/T001569/1, particularly the “Criminal Justice System” theme within that grant, and the “Hate Speech: Measures & Counter-Measures” project at The Alan Turing Institute. Dong Nguyen was supported by the “Digital Society - The Informed Citizen” research programme, which is (partly) financed by the Dutch Research Council (NWO), project 410.19.007. Zeerak Waseem was supported in part by the Canada 150 Research Chair program and the UK-Canada AI Artificial Intelligence Initiative. Janet B. Pierrehumbert was supported by EPSRC Grant EP/T023333/1.

Impact Statement

This supplementary section addresses relevant ethical considerations that were not explicitly discussed in the main body of our article.

Interview Participant Rights All interviewees gave explicit consent for their participation after being informed in detail about the research use of their responses. In all research output, quotes from interview responses were anonymised. We also did not reveal specific participant demographics or affiliations. Our interview approach was approved by the Alan Turing Institute’s Ethics Review Board.

Intellectual Property Rights The test cases in HATECHECK were crafted by the authors. As synthetic data, they pose no risk of violating intellectual property rights.

Annotator Compensation We employed a team of ten annotators to validate the quality of the HATECHECK dataset. Annotators were compensated at a rate of £16 per hour. The rate was set 50% above the local living wage (£10.85), although

all work was completed remotely. All training time and meetings were paid.

Intended Use HATECHECK’s intended use is as an evaluative tool for hate speech detection models, providing structured and targeted diagnostic insights into model functionalities. We demonstrated this use of HATECHECK in §3. We also briefly discussed alternative uses of HATECHECK, e.g. as a starting point for data augmentation. These uses aim at aiding the development of better hate speech detection models.

Potential Misuse Researchers might overextend claims about the functionalities of their models based on their test performance, which we would consider a misuse of HATECHECK. We directly addressed this concern by highlighting HATECHECK’s negative predictive power, i.e. the fact that it primarily reveals model weaknesses rather than necessarily characterising generalisable model strengths, as one of its limitations. For the same reason, we emphasised the limits to HATECHECK’s coverage, e.g. in terms of slurs and identity terms.

References

- Betty van Aken, Julian Risch, Ralf Krestel, and Alexander Löser. 2018. [Challenges for toxic comment classification: An in-depth error analysis](#). In *Proceedings of the 2nd Workshop on Abusive Language Online (ALW2)*, pages 33–42, Brussels, Belgium. Association for Computational Linguistics.
- Michele Banko, Brendon MacKeen, and Laurie Ray. 2020. [A Unified Taxonomy of Harmful Content](#). In *Proceedings of the Fourth Workshop on Online Abuse and Harms*, pages 125–137. Association for Computational Linguistics.
- Boris Beizer. 1995. *Black-box testing: techniques for functional testing of software and systems*. John Wiley & Sons, Inc.
- Yonatan Belinkov and Yonatan Bisk. 2018. Synthetic and natural noise both break neural machine translation. In *Proceedings of the 6th International Conference on Learning Representations*.
- 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.
- Pete Burnap and Matthew L Williams. 2015. Cyber hate speech on Twitter: An application of machine classification and statistical modeling for policy and decision making. *Policy & Internet*, 7(2):223–242.
- Rui Cao, Roy Ka-Wei Lee, and Tuan-Anh Hoang. 2020. DeepHate: Hate speech detection via multi-faceted text representations. In *Proceedings of the 12th ACM Conference on Web Science*, pages 11–20.
- Juliet M Corbin and Anselm Strauss. 1990. Grounded theory research: Procedures, canons, and evaluative criteria. *Qualitative Sociology*, 13(1):3–21.
- Thomas Davidson, Dana Warmusley, Michael Macy, and Ingmar Weber. 2017. Automated hate speech detection and the problem of offensive language. In *Proceedings of the 11th International AAAI Conference on Web and Social Media*, pages 512–515. Association for the Advancement of Artificial Intelligence.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- 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 and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4537–4546, Hong Kong, China. Association for Computational Linguistics.
- Lucas Dixon, John Li, Jeffrey Sorensen, Nithum Thain, and Lucy Vasserman. 2018. [Measuring and mitigating unintended bias in text classification](#). In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pages 67–73. Association for Computing Machinery.
- Allyson Ettinger. 2020. [What BERT is not: Lessons from a new suite of psycholinguistic diagnostics for language models](#). *Transactions of the Association for Computational Linguistics*, 8:34–48.
- Paula Fortuna and Sérgio Nunes. 2018. A survey on automatic detection of hate speech in text. *ACM Computing Surveys (CSUR)*, 51(4):1–30.
- 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 *Proceedings of the 12th International AAAI Conference on Web and Social Media*, pages 491–500. Association for the Advancement of Artificial Intelligence.
- Matt Gardner, Yoav Artzi, Victoria Basmov, Jonathan Berant, Ben Bogin, Sihao Chen, Pradeep Dasigi,

- Dheeru Dua, Yanai Elazar, Ananth Gottumukkala, Nitish Gupta, Hannaneh Hajishirzi, Gabriel Ilharco, Daniel Khoshdel, 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 models’ local decision boundaries via contrast sets](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1307–1323, Online. Association for Computational Linguistics.
- Sahaj Garg, Vincent Perot, Nicole Limtiaco, Ankur Taly, Ed H. Chi, and Alex Beutel. 2019. [Counterfactual fairness in text classification through robustness](#). In *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*, AIES ’19, page 219–226, New York, NY, USA. Association for Computing Machinery.
- 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.
- Max Glockner, Vered Shwartz, and Yoav Goldberg. 2018. [Breaking NLI systems with sentences that require simple lexical inferences](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 650–655, Melbourne, Australia. Association for Computational Linguistics.
- Jennifer Golbeck, Zahra Ashktorab, Rashad O Banjo, Alexandra Berlinger, Siddharth Bhagwan, Cody Buntain, Paul Chekalos, Alicia A Geller, Rajesh Kumar Gnanasekaran, Raja Rajan Gunasekaran, et al. 2017. [A large labeled corpus for online harassment research](#). In *Proceedings of the 9th ACM Conference on Web Science*, pages 229–233. Association for Computing Machinery.
- 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. Association for Computing Machinery.
- 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.
- Haibo He and Eduardo A Garcia. 2009. [Learning from imbalanced data](#). *IEEE Transactions on Knowledge and Data Engineering*, 21(9):1263–1284.
- Hossein Hosseini, Sreeram Kannan, Baosen Zhang, and Radha Poovendran. 2017. [Deceiving Google’s Perspective API built for detecting toxic comments](#). *arXiv preprint arXiv:1702.08138*.
- Pierre Isabelle, Colin Cherry, and George Foster. 2017. [A challenge set approach to evaluating machine translation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2486–2496, Copenhagen, Denmark. Association for Computational Linguistics.
- ISO/IEC/IEEE 24765:2017(E). 2017. [Systems and Software Engineering – Vocabulary](#). Standard, International Organization for Standardisation, Geneva, Switzerland.
- Divyansh Kaushik, Eduard Hovy, and Zachary C. Lipton. 2020. [Learning the difference that makes a difference with counterfactually-augmented data](#). In *Proceedings of the 8th International Conference on Learning Representations*.
- 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, et al. 2021. [Dynabench: Rethinking benchmarking in nlp](#). *arXiv preprint arXiv:2104.14337*.
- 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, Online. Association for Computational Linguistics.
- J. Richard Landis and Gary G. Koch. 1977. [The measurement of observer agreement for categorical data](#). *Biometrics*, 33(1):159.
- Hector Levesque, Ernest Davis, and Leora Morgenstern. 2012. [The Winograd schema challenge](#). In *13th International Conference on the Principles of Knowledge Representation and Reasoning*.
- Nelson F. Liu, Roy Schwartz, and Noah A. Smith. 2019. [Inoculation by fine-tuning: A method for analyzing challenge 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 2171–2179, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *Proceedings of the 7th International Conference on Learning Representations*.

- Shervin Malmasi and Marcos Zampieri. 2018. Challenges in discriminating profanity from hate speech. *Journal of Experimental & Theoretical Artificial Intelligence*, 30(2):187–202.
- Rebecca Marvin and Tal Linzen. 2018. Targeted syntactic evaluation of language models. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202, Brussels, Belgium. Association for Computational Linguistics.
- Tom McCoy, Ellie Pavlick, and Tal Linzen. 2019. Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, Florence, Italy. Association for Computational Linguistics.
- Julia Mendelsohn, Yulia Tsvetkov, and Dan Jurafsky. 2020. A framework for the computational linguistic analysis of dehumanization. *Frontiers in Artificial Intelligence*, 3:55.
- Pushkar Mishra, Helen Yannakoudakis, and Ekaterina Shutova. 2020. Tackling online abuse: A survey of automated abuse detection methods. *arXiv preprint arXiv:1908.06024*.
- Marzieh Mozafari, Reza Farahbakhsh, and Noel Crespi. 2019. A BERT-based transfer learning approach for hate speech detection in online social media. In *International Conference on Complex Networks and Their Applications*, pages 928–940. Springer.
- Aakanksha Naik, Abhilasha Ravichander, Norman Sadeh, Carolyn Rose, and Graham Neubig. 2018. Stress test evaluation for natural language inference. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 2340–2353, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Isar Nejadgholi and Svetlana Kiritchenko. 2020. On cross-dataset generalization in automatic detection of online abuse. In *Proceedings of the Fourth Workshop on Online Abuse and Harms*, pages 173–183, Online. Association for Computational Linguistics.
- 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.
- Timothy Niven and Hung-Yu Kao. 2019. Probing neural network comprehension of natural language arguments. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4658–4664, Florence, Italy. Association for Computational Linguistics.
- Chikashi Nobata, Joel Tetreault, Achint Thomas, Yashar Mehdad, and Yi Chang. 2016. Abusive language detection in online user content. In *Proceedings of the 25th International Conference on World Wide Web, WWW '16*, page 145–153, Republic and Canton of Geneva, CHE. International World Wide Web Conferences Steering Committee.
- Alexis Palmer, Christine Carr, Melissa Robinson, and Jordan Sanders. 2020. COLA: Annotation scheme and evaluation data set for complex offensive language in English. *Journal for Language Technology and Computational Linguistics*, pages 1–28.
- Ji Ho Park, Jamin Shin, and Pascale Fung. 2018. Reducing gender bias in abusive language detection. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2799–2804, Brussels, Belgium. Association for Computational Linguistics.
- 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.
- Jing Qian, Mai ElSherief, Elizabeth Belding, and William Yang Wang. 2018. Leveraging intra-user and inter-user representation learning for automated hate speech detection. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 118–123, New Orleans, Louisiana. Association for Computational Linguistics.
- 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.
- Joni Salminen, Hind Almerkhi, Milica Milenković, Soon-gyo Jung, Jisun An, Haewoon Kwak, and Bernard Jansen. 2018. Anatomy of online hate: Developing a taxonomy and machine learning models for identifying and classifying hate in online news media. *Proceedings of the 12th International AAAI Conference on Web and Social Media*, (1):330–339.
- Mattia Samory, Indira Sen, Julian Kohne, Fabian Floeck, and Claudia Wagner. 2020. Unsex me here: Revisiting sexism detection using psychological scales and adversarial samples. *arXiv preprint arXiv:2004.12764*.
- Maarten Sap, Dallas Card, Saadia Gabriel, Yejin Choi, and Noah A. Smith. 2019. The risk of racial bias in hate speech detection. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1668–1678, Florence, Italy. Association for Computational Linguistics.

- Maarten Sap, Saadia Gabriel, Lianhui Qin, Dan Jurafsky, Noah A. Smith, and Yejin Choi. 2020. [Social bias frames: Reasoning about social and power implications of language](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5477–5490, Online. Association for Computational Linguistics.
- 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, Valencia, Spain. Association for Computational Linguistics.
- Rico Sennrich. 2017. [How grammatical is character-level neural machine translation? Assessing MT quality with contrastive translation pairs](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 376–382, Valencia, Spain. Association for Computational Linguistics.
- Deven Santosh Shah, H. Andrew Schwartz, and Dirk Hovy. 2020. [Predictive biases in natural language processing models: A conceptual framework and overview](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5248–5264, Online. Association for Computational Linguistics.
- Thanh Tran, Yifan Hu, Changwei Hu, Kevin Yen, Fei Tan, Kyumin Lee, and Se Rim Park. 2020. [HABERTOR: An efficient and effective deep hatespeech detector](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7486–7502, Online. Association for Computational Linguistics.
- 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, Scott Hale, Ella Guest, Helen Margetts, David Broniatowski, Zeerak Waseem, Austin Botelho, Matthew Hall, and Rebekah Tromble. 2020a. [Detecting East Asian prejudice on social media](#). In *Proceedings of the Fourth Workshop on Online Abuse and Harms*, pages 162–172, Online. Association for Computational Linguistics.
- 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 Waseem, and Douwe Kiela. 2020b. [Learning from the worst: Dynamically generated datasets to improve online hate detection](#). *CoRR*, abs/2012.15761.
- 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, Montréal, Canada. Association for Computational Linguistics.
- Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohananey, Wei Peng, Sheng-Fu Wang, and Samuel R. Bowman. 2020. [BLiMP: The benchmark of linguistic minimal pairs for English](#). *Transactions of the Association for Computational Linguistics*, 8:377–392.
- 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 Warmsley, 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, Vancouver, BC, Canada. Association for Computational Linguistics.
- Zeerak Waseem and Dirk Hovy. 2016. [Hateful symbols or hateful people? Predictive features for hate speech detection on Twitter](#). In *Proceedings of the NAACL Student Research Workshop*, pages 88–93, San Diego, California. Association for Computational Linguistics.
- 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*. Springer.
- 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.
- Tongshuang Wu, Marco Tulio Ribeiro, Jeffrey Heer, and Daniel Weld. 2019. [Errudite: Scalable, reproducible, and testable error analysis](#). In *Proceedings of the 57th Annual Meeting of the Association*

for *Computational Linguistics*, pages 747–763, Florence, Italy. Association for Computational Linguistics.

Ellery Wulczyn, Nithum Thain, and Lucas Dixon. 2017. Ex machina: Personal attacks seen at scale. In *Proceedings of the 26th International Conference on World Wide Web*, 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 the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1415–1420, Minneapolis, Minnesota. Association for Computational Linguistics.

Ziqi Zhang and Lei Luo. 2019. Hate speech detection: A solved problem? The challenging case of long tail on Twitter. *Semantic Web*, 10(5):925–945.

Xiang Zhou, Yixin Nie, Hao Tan, and Mohit Bansal. 2020. [The curse of performance instability in analysis datasets: Consequences, source, and suggestions](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8215–8228, Online. Association for Computational Linguistics.

A Data Statement

Following [Bender and Friedman \(2018\)](#), we provide a data statement, which documents the generation and provenance of test cases in HATECHECK.

A. CURATION RATIONALE In order to construct HATECHECK, a first suite of functional tests for hate speech detection models, we generated 3,901 short English-language text documents by hand and by using simple templates for group identifiers and slurs (§2.4). Each document corresponds to one functional test and a binary gold standard label (hateful or non-hateful). In order to validate the gold standard labels, we trained a team of ten annotators, assigning five of them to each document, and asked them to provide independent labels (§2.5). To further improve data quality, we also gave annotators the option to flag cases they felt were unrealistic (e.g. nonsensical), but this flag was not used for any one HATECHECK case by more than one annotator.

B. LANGUAGE VARIETY HATECHECK only covers English-language text documents. We opted for English language since this maximises HATECHECK’s relevance to previous and current work in hate speech detection, which is mostly concerned with English-language data. Our language choice also reflects the expertise of authors and annotators. We discuss the lack of language variety as a limitation of HATECHECK in §4.2 and suggest expansion to other languages as a priority for future research.

C. SPEAKER DEMOGRAPHICS Since all test cases in HATECHECK were hand-crafted, the speakers are the same as the authors. Test cases in the test suite were primarily generated by the lead author, who is a researcher at a UK university. The lead author is not a native English speaker but has lived in English-speaking countries for more than five years and has extensively engaged with English-language hate speech in previous research. All test cases were also reviewed by two co-authors, both of whom have worked with English-language hate speech data for more than five years and one of whom is a native English speaker from the UK.

D. ANNOTATOR DEMOGRAPHICS We recruited a team of ten annotators to work for two weeks. 30% were male and 70% were female. 60% were 18-29 and 40% were 30-39. 20% were educated to high school level, 10% to undergraduate, 60% to taught masters and 10% to research

degree (i.e. PhD). 70% were native English speakers and 30% were non-native but fluent. Annotators had a range of nationalities: 60% were British and 10% each were Polish, Spanish, Argentinian and Irish. Most annotators identified as ethnically White (70%), followed by Middle Eastern (20%) and a mixed ethnic background (10%). Annotators all used social media regularly, and 60% used it more than once per day. All annotators had seen other people targeted by online abuse before, and 80% had been targeted personally.

All annotators had previously completed annotation work on at least one other hate speech dataset. In the first week, we introduced the binary annotation task to them in an onboarding session and tested their understanding on a set of 100 cases, which we then provided individual feedback on. In the second week, we asked each annotator to annotate around 2,000 test cases so that each case in our test suite was annotated by varied sets of exactly five annotators. Throughout the process, we communicated with annotators in real-time over a messaging platform. We also followed guidance for protecting and monitoring annotator well-being provided by [Vidgen et al. \(2019\)](#).

E. SPEECH SITUATION All test cases were created between the 23rd of November and the 13th of December 2020.

F. TEXT CHARACTERISTICS The composition of the dataset, including primary label and secondary labels, is described in detail in §2.3 and §2.4 of the article.

B References for Functional Tests

F1 – *strong negative emotions explicitly expressed about a protected group or its members*: Resembles “expressed hatred” ([Davidson et al., 2017](#)) and “identity attack” ([Banko et al., 2020](#)).

F2 – *explicit descriptions of a protected group or its members using very negative attributes*: Refines more general “insult” categories ([Davidson et al., 2017](#); [Zampieri et al., 2019](#)).

F3 – *explicit dehumanisation of a protected group or its members*: Prevalent form of hate ([Mendelsohn et al., 2020](#); [Banko et al., 2020](#); [Vidgen et al., 2020a](#)). Highlighted in our interviews (e.g. I18: “hate crime [often claims] people are inferior and subhuman.”).

F4 – *implicit derogation of a protected group or its members*: Closely resembles “implied bias”

(Sap et al., 2020) and “implicit abuse” (Waseem et al., 2017; Zhang and Luo, 2019). Highlighted in our interviews (e.g. I16: “hate has always been expressed idiomatically”).

F5 – *direct threats against a protected group or its members*: Core element of several hate speech taxonomies (Golbeck et al., 2017; Zampieri et al., 2019; Vidgen et al., 2020a; Banko et al., 2020)

F6 – *threats expressed as normative statements*: Highlighted by an interviewee as a way of avoiding legal consequences to hate speech (I1: “[normative threats] are extremely hateful, but [legally] okay”).

F7 – *hate expressed using slurs*: Prevalent way of expressing hate (Palmer et al., 2020; Banko et al., 2020; Kurrek et al., 2020).

F8 – *non-hateful homonyms of slur*: Relevant alternative use of slurs (Kurrek et al., 2020).

F9 – *use of reclaimed slurs*: Likely source of classification error (Palmer et al., 2020). Highlighted in our interviews (e.g. I7: “A lot of LGBT people use slurs to identify themselves, like reclaim the word queer, and people [...] report that and then that will get hidden”).

F10 – *hate expressed using profanity*: Refines more general “insult” categories (Davidson et al., 2017; Zampieri et al., 2019).

F11 – *non-hateful uses of profanity*: Oversensitiveness of hate speech detection models to profanity (Davidson et al., 2017; Malmasi and Zampieri, 2018; van Aken et al., 2018).

F12 – *hate expressed through pronoun reference in subsequent clauses*: Syntactic relationships and long-range dependencies as model weak points (Burnap and Williams, 2015; Vidgen et al., 2019).

F13 – *hate expressed through pronoun reference in subsequent sentences*: See F12.

F14 – *hate expressed using negated positive statements*: Negation as an effective adversary for hate speech detection models (Hosseini et al., 2017; Dinan et al., 2019).

F15 – *non-hate expressed using negated hateful statements*: See F14.

F16 – *hate phrased as a question*: Likely source of classification error (van Aken et al., 2018).

F17 – *hate phrased as an opinion*: Highlighted by an interviewee as a way of avoiding legal consequences to hate speech (I1: “If you start a sentence by saying ‘I think that’ [...], the limits of what you can say are much bigger”).

F18 – *neutral statements using protected group identifiers*: Oversensitiveness of hate speech de-

tection models to terms such as “black” and “gay” (Dixon et al., 2018; Park et al., 2018; Kennedy et al., 2020). Also highlighted in our interviews (e.g. I7: “I have seen the algorithm get it wrong, if someone’s saying something like ‘I’m so gay’.”).

F19 – *positive statements using protected group identifiers*: See F18.

F20 – *denouncements of hate that quote it*: Counter speech as a source of classification error (Warner and Hirschberg, 2012; van Aken et al., 2018; Vidgen et al., 2020a). Most mentioned concern in our interviews (e.g. I4: “people will be quoting someone, calling that person out [...] but that will get picked up by the system”).

F21 – *denouncements of hate that make direct reference to it*: See F20.

F22 – *abuse targeted at objects*: Distinct from hate speech since it targets out-of-scope entities (Wulczyn et al., 2017; Zampieri et al., 2019).

F23 – *abuse targeted at individuals not referencing membership in a protected group*: See F22.

F24 – *abuse targeted at non-protected groups (e.g. professions)*: See F22.

F25 – *swaps of adjacent characters*: Simple misspellings can be challenging for detection models (van Aken et al., 2018; Qian et al., 2018). Particularly relevant to hate speech since they can reflect intentional behaviour of users looking to avoid detection (Hosseini et al., 2017; Gröndahl et al., 2018; Vidgen et al., 2019).

F26 – *missing characters*: Highlighted in our interviews (e.g. I7: “it could be a misspelling of a word like ‘f*ggot’, and someone’s put one ‘g’ instead of two”).

F27 – *missing word boundaries*: Effective adversary for a hate speech detection model (Gröndahl et al., 2018). Resembles the use of hashtags on social media (I2: “there have been a highly Islamophobic hashtags going around”).

F28 – *added spaces between characters*: Effective adversary for a hate speech detection model (Gröndahl et al., 2018). Highlighted in our interviews (e.g. I5: “misspellings, missing letters or additional spaces between the letters.”).

F29 – *leet speak*: Resembles “obfuscations” (Nobata et al., 2016; van Aken et al., 2018). Highlighted in our interviews (e.g. I14: “[hate speakers] replace letters with numbers”).

C Hateful Slurs in HATECHECK

For each of the seven protected groups covered by HATECHECK, we searched hatebase.org, a crowd-sourced hate speech lexicon, for slurs which target that group. From these slurs, we selected the three that were most often logged by users of the site (e.g. “wh*re”, “b*tch” and “sl*t” for women), except for when the third-most sighted slur was logged substantially less often than the second, in which case we selected the top two (e.g. “tr*nny” and “sh*male” for trans people). For immigration status, which is not a target category on hatebase.org, we chose “r*pefugee”, a slur for refugees used by the European far right, and “w*tback”, a slur for Mexican immigrants to the US, which was logged similarly often as other slurs in HATECHECK.

Protected Group	Slurs
Women	b*tch, sl*t, wh*re
Trans people	tr*nny, sh*male
Gay people	f*ggot, f*g, q*eer
Black people	n*gger, c*on
Disabled people	r*tard, cr*ppl, m*ng
Muslims	m*zzi, J*hadi, camel f*cker
Immigrants	w*tbacks, r*pefugees

Table 5: Hateful slurs in HATECHECK

For reclaimed slurs (F9), we focus on slurs reclaimed by black communities (particularly “n*gga”), gay communities (“f*g”, “f*ggot”, “q*eer”) and by women (“b*tch”), reflecting the concerns highlighted by our interview participants (e.g. I4: “n*gga would often get [wrongly] picked up by [moderation] systems”). Ahead of the structured annotation process (§2.5) and only for test cases with reclaimed slurs, we asked self-identifying members of the relevant groups in our personal networks whether they would consider the test cases to contain valid and realistic reclaimed slur uses, which held true for all test cases.

D Datasets for Fine-Tuning

D.1 Davidson et al. (2017) Data

Sampling Davidson et al. (2017) searched Twitter for tweets containing keywords from a list they compiled from hatebase.org, which yielded a sample of tweets from 33,458 users. They then randomly sampled 25,000 tweets from all tweets of these users.

Annotation The authors hired crowd workers from CrowdFlower to annotate each tweet as *hateful*, *offensive* or *neither*. 92.0% of tweets were annotated by three crowd workers, the remainder by at least four and up to nine. For inter-annotator agreement, the authors report a “CrowdFlower score” of 92%.

Data We used 24,783 annotated tweets made available by the authors on github.com/t-davidson/hate-speech-and-offensive-language. 1,430 tweets (5.8%) are labelled *hateful*, 19,190 (77.4%) *offensive* and 4,163 (16.8%) *neither*. We collapse the latter two labels into a single non-hateful label to match HATECHECK’s binary format, resulting in 1,430 tweets (5.8%) labelled *hateful* and 23,353 (94.2%) labelled *non-hateful*.

Definition of Hate Speech “Language that is used to express hatred towards a targeted group or is intended to be derogatory, to humiliate, or to insult the members of the group”.

D.2 Founta et al. (2018) Data

Sampling Founta et al. (2018) initially collected a random set of 32 million tweets from Twitter. They then used a boosted random sampling procedure based on negative sentiment and occurrence of offensive words as selected from hatebase.org to augment a random subset of this initial sample with tweets they expected to be more likely to be hateful or abusive.

Annotation The authors hired crowd workers from CrowdFlower to annotate each tweet as *hateful*, *abusive*, *spam* or *normal*. All tweets were annotated by five crowd workers. For inter-annotator agreement, the authors report that 55.9% of tweets had four out of five annotators agreeing on a label.

Data The authors provided us access to the full text versions of 99,996 annotated tweets. These correspond to the tweet IDs made available by the authors on github.com/ENCASEH2020/hatespeech-twitter. 4,965 tweets (5.0%) are labelled *hateful*, 27,150 (27.2%) *abusive*, 14,030 (14.0%) *spam* and 53,851 (53.9%) *normal*. We collapse the latter three labels into a single non-hateful label to match HATECHECK’s binary format, resulting in 4,965 tweets (5.0%) labelled *hateful* and 95,031 tweets (95.0%) labelled *non-hateful*.

Definition of Hate Speech “Language used to express hatred towards a targeted individual or

group, or is intended to be derogatory, to humiliate, or to insult the members of the group, on the basis of attributes such as race, religion, ethnic origin, sexual orientation, disability, or gender”.

D.3 Pre-Processing

Before using the datasets for fine-tuning, we lowercase all text and remove newline and tab characters. We replace URLs, user mentions and emojis with [URL], [USER] and [EMOJI] tokens. We also split hashtags into separate tokens using the `wordsegment` Python package.

E Details on Transformer Models

Model Architecture We implemented uncased BERT-base models (Devlin et al., 2019) using the `transformers` Python library (Wolf et al., 2020). Uncased BERT-base, which is trained on lower-cased English text, has 12 layers, a hidden layer size of 768, 12 attention heads and a total of 110 million parameters. For sequence classification, we added a linear layer with softmax output.

Fine-Tuning **B-D** was fine-tuned on binary Davidson et al. (2017) data and **B-F** on binary Founta et al. (2018) data. For both datasets, we used a stratified 80/10/10 train/dev/test split. Models were trained for three epochs each. Training batch size was 16. We used cross-entropy loss with class weights emphasising the hateful minority class. Weights were set to the relative proportion of the other class in the training data, meaning that for a 1:9 hateful:non-hateful case split, loss on hateful cases would be multiplied by 9. The optimiser was AdamW (Loshchilov and Hutter, 2019) with a $5e-5$ learning rate and a 0.01 weight decay. For regularisation, we set a 10% dropout probability.

Hyperparameter Tuning The number of fine-tuning epochs, the learning rate and the training batch size were determined by exhaustive grid search. We used the range of possible values recommended by Devlin et al. (2019): [2, 3, 4] for epochs, [2e-5, 3e-5, 5e-5] for learning rate and [16, 32] for batch size. There were 18 training/evaluation runs for each model. The best configuration was selected based on loss on the 10% development set.

Held-Out Performance Micro/macro F1 scores on the held-out test sets corresponding to their training data are 91.5/70.8 for **B-D** (Davidson et al., 2017) and 92.9/70.3 for **B-F** (Founta et al., 2018).

Computation We ran all computations on a Microsoft Azure “Standard_NC24” server equipped with two NVIDIA Tesla K80 GPU cards. The average wall time for each hyperparameter tuning trial of **B-D** was around 17 minutes, and for **B-F** around 70 minutes.

Source Code Our code is available on github.com/paul-rottger/hatecheck-experiments.

F Comparison to SOTA Results

Most previous work that trains and evaluates models on Davidson et al. (2017) and Founta et al. (2018) data uses their original multiclass label format. In the multiclass case, the relative size of the hateful class compared to the non-hateful classes is larger than in the binary case, which is likely why most models do not use class weights. For comparability, we thus fine-tuned unweighted multiclass versions of **B-D** and **B-F**, using the same model parameters described in Appendix E.

On multiclass Davidson et al. (2017) data, Mozafari et al. (2019) report a weighted-average F1 score of 91 for their BERT-base model and 92 for BERT-base combined with a CNN. Cao et al. (2020) report a micro F1 of 89.9 for their ensemble-like “DeepHate” classifier. Our unweighted multiclass BERT-base model achieves 90.7 weighted-average F1 and 91.1 micro F1.

On multiclass Founta et al. (2018) data, Cao et al. (2020) report a micro F1 of 79.1 for “DeepHate”. Our unweighted multiclass BERT-base model achieves 81.7 micro F1.

Tran et al. (2020) recently achieved SOTA on several other hate speech datasets with their HABERTOR model. They also find that BERT-base consistently performs very near their SOTA. However, they do not evaluate their models on Davidson et al. (2017) or Founta et al. (2018) data.