

Is More Data Better? Re-thinking the Importance of Efficiency in Abusive Language Detection with Transformers-Based Active Learning

Hannah Rose Kirk^{1,2,3‡}, Bertie Vidgen^{1,3}, Scott A. Hale^{1,2,4}

¹University of Oxford, ²The Alan Turing Institute, ³Rewire, ⁴Meedan

[‡]hannah.kirk@oii.ox.ac.uk

Abstract

Annotating abusive language is expensive, logistically complex and creates a risk of psychological harm. However, most machine learning research has prioritized maximizing *effectiveness* (i.e., F1 or accuracy score) rather than data *efficiency* (i.e., minimizing the amount of data that is annotated). In this paper, we use simulated experiments over two datasets at varying percentages of abuse to demonstrate that transformers-based active learning is a promising approach to substantially raise efficiency whilst still maintaining high effectiveness, especially when abusive content is a smaller percentage of the dataset. This approach requires a fraction of labeled data to reach performance equivalent to training over the full dataset.

1 Introduction

Online abuse, such as hate and harassment, can inflict psychological harm on victims (Gelber and McNamara, 2016), disrupt communities (Mohan et al., 2017) and even lead to physical attacks (Williams et al., 2019). Machine learning solutions can be used to automatically detect abusive content at scale, helping to tackle this growing problem (Gillespie, 2020). An *effective* model is one which makes few misclassifications, minimizing the risk of harm from false positives and negatives: false negatives mean that users are not fully protected from abuse while false positives constrain free expression. Most models to automatically detect abuse are trained to maximize effectiveness via “passive” supervised learning over large labeled datasets. However, although collecting large amounts of social media data is relatively cheap and easy, annotating data is expensive, logistically complicated and creates a risk of inflicting psychological harm on annotators through vicarious trauma (Roberts, 2019; Steiger et al., 2021). Thus, an *efficient* model, which achieves a given level of performance with few labeled examples, is highly desirable for abusive content detection.




Figure 1: Transformers-based active learning beats fully-supervised baseline with 1.5% of the 20,000 examples.

Our central objective is to demonstrate how to maximize efficiency and effectiveness when training abuse detection systems, and in this paper, we focus on active learning (AL). AL is an iterative human-in-the-loop approach that selects entries for annotation only if they are ‘informative’ (Lewis and Gale, 1994; Settles, 2009). While AL has shown promise for abusive language dataset creation (Charitidis et al., 2020; Mollas et al., 2020; Rahman et al., 2021; Bashar and Nayak, 2021; Abidin et al., 2021), there are several open questions about the most appropriate configuration and use. In particular, only one paper uses transformers-based AL for abusive language detection (Ein-Dor et al., 2020) to our knowledge, although the benefits of AL for other classification tasks is clear (Schroder et al., 2022; Ein-Dor et al., 2020; Yuan et al., 2020). Pre-trained transformer models have been widely adopted for abuse detection, but while they can be fine-tuned on relatively few examples for specific tasks (Devlin et al., 2018; Qiu et al., 2020), they are still commonly used with large datasets (e.g. Mozafari et al., 2019; Mutanga et al., 2020; Isaksen and Gambäck, 2020; Koufakou et al., 2020). Our first subquestion asks, **RQ1.1:** *What effect do model pre-training and architecture have on efficiency and effectiveness?* To answer RQ1.1, we evaluate transformers- and traditional-based AL

in a simulated setup using two already-labeled abusive language datasets.

One challenge in abusive language detection is class imbalance as, although extremely harmful, abuse comprises a small portion of online content (Vidgen et al., 2019). Prior AL work primarily uses datasets at their given class imbalances and thus has not disentangled how class imbalance versus linguistic features affect the design choices needed for efficient AL. This is a problem given that most abusive language datasets do not reflect the imbalance actually observed in the wild. Our second subquestion addresses this issue, **RQ1.2: What effect does class imbalance have on efficiency and effectiveness?** To answer RQ1.2, we artificially-rebalance the datasets at different percentages of abuse.

In addressing these questions, we find that more data is not always better and can actually be worse, showing that effectiveness and efficiency are not always in tension with one another. With extensive pre-training and greater model complexity, a transformers-based AL approach achieves high performance with only a few hundred examples. Crucially, we show that the value of transformers-based AL (relative to random sampling) is larger for more imbalanced data (i.e., data that more closely reflects the real-world). For 5% abuse, the performance of a transformers-based AL strategy over 3% of a 20k dataset can even surpass the F1 of a model passively trained over the full dataset by 5 percentage points (Fig. 1). In §4 we describe caveats of our findings and implications for future research in abusive language detection.¹

2 Methods

2.1 Active Learning Set-Up

AL typically consists of four components: 1) a classification model, 2) pools of unlabeled data \mathcal{U} and labeled data \mathcal{L} , 3) a query strategy for identifying data to be labeled, and 4) an ‘oracle’ (e.g., human annotators) to label the data. First, seed examples are taken from \mathcal{U} and sent to the oracle(s) for labeling. These examples initialize the classification model. Second, batches of examples are iteratively sampled from the remaining unlabeled pool, using a query strategy to estimate their ‘informativeness’ to the initialized classification model.²

¹Code at [ActiveTransformers-for-AbusiveLanguage](https://github.com/ActiveTransformers-for-AbusiveLanguage).

²Note that *batch-mode active learning* is a common application in both research and industry, given its more practi-

Table 1: Summary of source datasets (in gray) and their artificially-rebalanced versions.

Dataset	Imbalance	Train [†]		Test [*]	
		abuse	non-abuse	abuse	non-abuse
wiki	12%	10,834	81,852	2,756	20,422
wiki50	50%	10,000	10,000	2,500	2,500
wiki10	10%	2,000	18,000	500	4,500
wiki5	5%	1,000	19,000	250	4,750
tweets	32%	28,955	61,041	3,160	6,840
tweets50	50%	10,000	10,000	2,500	2,500
tweets10	10%	2,000	18,000	500	4,500
tweets5	5%	1,000	19,000	250	4,750

Notes: [†] Train is used as the unlabeled pool ($n = 20,000$)

^{*} Test is used for held-out evaluation ($n = 5,000$)

Each queried batch is labeled and added to \mathcal{L} . Finally, the classifier is re-trained over \mathcal{L} .³

2.2 Dataset Selection and Processing

AL is path-dependent—i.e., later decisions are dependent upon earlier ones; so, experimenting in real-world settings is prohibitively costly and risky to annotator well-being. To reproduce the process without labeling new data, we use existing labeled datasets but withhold the labels until the model requests their annotation. We examined a list of publicly available, annotated datasets for abusive language detection⁴ and found two that were sufficiently large and contained enough abusive instances to facilitate our experimental approach. The **wiki** dataset (Wulczyn et al., 2017) contains comments from Wikipedia editors, labeled for whether they contain personal attacks. A test set is pre-defined; we take our test instances from this set. The **tweets** dataset (Founta et al., 2018) contains tweets which have been assigned to one of four classes. We binarize by combining the abusive and hate speech classes (=1) and the normal and spam classes (=0) to allow for cross-dataset comparison (Wiegand et al., 2019; Ein-Dor et al., 2020). A test set is not pre-defined; so, we set aside 10% of the data for testing that is never used for training.

To disentangle the merits of AL across class imbalances, we construct three new datasets for both **wiki** and **tweets** that have different class distributions: 50% abuse, 10% abuse and 5% abuse. This creates 6 datasets in total (see Tab. 1). To control

cal application to annotation workflows and model retraining times (Settles, 2009, p. 35).

³We train from scratch to avoid overfitting to previous iterations (Ein-Dor et al., 2020; Hu et al., 2018).

⁴<https://hatespeechdata.com>

dataset size and ensure we have sufficient positive instances for all imbalances, we assume that each unlabeled pool has 20,000 examples.⁵ We experiment with multiple AL strategies to select 2,000 examples for annotation as early experiments showed further iterations did not affect performance.⁶

2.3 Experimental Setup

We use 2 model architectures, 2 query strategies and 6 artificially-rebalanced datasets, giving 24 experiments each of which we repeat with 3 random seeds. Each experiment uses the same sized unlabeled pool, training budget and test set (see Tab. 1). In figures, we present the mean run (line) and standard deviation (shaded). For transformers-based AL, we use distil-roBERTa (**dBERT**), which performs competitively to larger transformer models (Sanh et al., 2019), also in an AL setting (Schröder et al., 2022). For traditional AL without pre-training, we use a linear support vector machine (**SVM**) as a simple, fast and lightweight baseline.⁷ For active data acquisition, we try three AL strategies; LeastConfidence, which selects items close to the decision boundary (Lewis and Gale, 1994), is presented in the paper while the other strategies are in the Appendix.⁸ For comparison, we randomly sample items from the unlabeled pool at each iteration. Alongside model and query strategy, AL requires an initial seed size, seed acquisition strategy and batch size. We experimentally determined the best values for these parameters: an initial seed of 20 examples selected via a keyword-heuristic (Ein-Dor et al., 2020) and batches of 50 examples.⁹

2.4 Evaluation

As a baseline, we use the passive macro-F1 score over the full dataset of 20,000 entries ($F1_{20k}$). For each AL strategy, we measure efficiency on the held-out test set as the number of examples needed to surpass 90% of $F1_{20k}$, which we call N_{90} .¹⁰ For effectiveness, we use the maximum F1 score achieved by each AL strategy, which we call $F1_{AL}$.

⁵The **wiki** dataset has 10,834 abusive entries; so, at 50% abuse, the upper limit on a rebalanced pool is 21,668.

⁶AL experiments are implemented in the Python `small-text` library (Schröder et al., 2021)

⁷Appendix A presents details of model training.

⁸We also test GreedyCoreSet (Sener and Savarese, 2017) and EmbeddingKMeans (Yuan et al., 2020), but LeastConfidence outperformed them.

⁹We present pilot experiments in Appendix B and C.

¹⁰To fairly compare models, we calculate N_{90} relative to *best* $F1_{20k}$ (achieved by dBERT in all cases).

Table 2: Efficiency and effectiveness of each classifier (transformers vs SVM) with LeastConfidence sampling.

Dataset	Classifier	$F1_{20k}^\dagger$	$F1_{AL}$	N_{90}
wiki50	dBERT	0.920	0.920	170
	SVM	0.875	0.836	1570
wiki10	dBERT	0.859	0.866	170
	SVM	0.809	0.810	320
wiki5	dBERT	0.807	0.855	220
	SVM	0.785	0.780	170
tweets50	dBERT	0.939	0.938	170
	SVM	0.931	0.926	220
tweets10	dBERT	0.904	0.902	220
	SVM	0.893	0.901	170
tweets5	dBERT	0.844	0.856	300
	SVM	0.825	0.830	170

Notes: † global metric from passive training over full, re-balanced dataset

3 Results

Efficiency & Effectiveness For each dataset, we find active strategies that need just 170 examples (0.8% of the full dataset) to reach 90% of passive supervised learning performance (see Tab. 2). When training over the full dataset, dBERT always outperforms SVM, models have worse performance on more imbalanced datasets, and **wiki** is harder to predict than **tweets** (Tab. 2). In all cases, LeastConfidence outperforms the random baseline, and the gain is larger for lower percentages of abuse: for **wiki10** and **wiki5**, N_{90} is lower by 150 and 100 examples, respectively. AL can even outperform passive supervised learning over the full dataset, showing there is no efficiency-effectiveness trade-off. For the majority of datasets, dBERT with LeastConfidence over 2,000 examples matches or surpasses the F1 score of a model trained passively over the whole dataset ($F1_{AL} \geq F1_{20k}$ in Tab. 2). For **wiki5**, it is 5 percentage points (pp) higher (Fig. 1).

The Effect of Pre-Training We find AL has a bigger impact for SVM than dBERT, shown by the larger gap to the random baselines (Fig. 2). With its extensive pre-training, dBERT achieves high performance with few examples, even if randomly selected. Nonetheless, an AL component still enhances dBERT performance above the random baseline especially with imbalanced data (as found by Schröder et al., 2022; Ein-Dor et al., 2020), requiring 150 and 100 fewer examples for N_{90} , and raising F1 score by 2pp and 4pp, for **wiki5** and **wiki10** respectively.

Train Distribution To assess why AL is more impactful with imbalanced data, we evaluate the

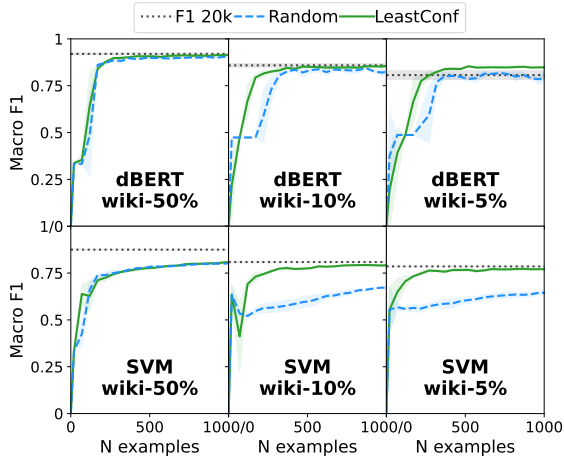


Figure 2: The contribution of pre-training vs active data acquisition.

distribution of the labeled pool at each iteration (Fig. 3). The random baseline tends to the original distribution as expected but the LeastConfidence strategy actively selects abusive examples from the pool and tends toward a balanced distribution.

Out-of-domain Testing The high performance of models trained on few examples raises a risk that they are overfitting and may not generalize. We take the models trained on each of the three class imbalances for **wiki** and test them on their equivalent **tweets** dataset, and vice versa. As with in-domain results, models trained on **wiki** and applied to **tweets** reach $F1_{20k}$ within few iterations. The gap between LeastConfidence and the random baseline is larger for out-of-domain evaluation versus in-domain (Fig. 4). A similar pattern occurs for other imbalances (see Appendix D). This suggests that our results for these two datasets are not overfitting.

4 Discussion

In response to our central research objective, we find strategies which are both effective and efficient, requiring far fewer examples to reach performance equivalent to passive training over the full dataset. These results suggest that passive approaches may be needlessly expensive and place annotators at unnecessary risk of harm. For RQ1.1, we find that coupling pre-trained transformers with AL is a successful approach which leverages the benefits of careful training data selection with the previously demonstrated strong capabilities of pre-trained language models for few-shot learning (Brown et al., 2020; Gao et al., 2021; Schick and Schütze, 2021).

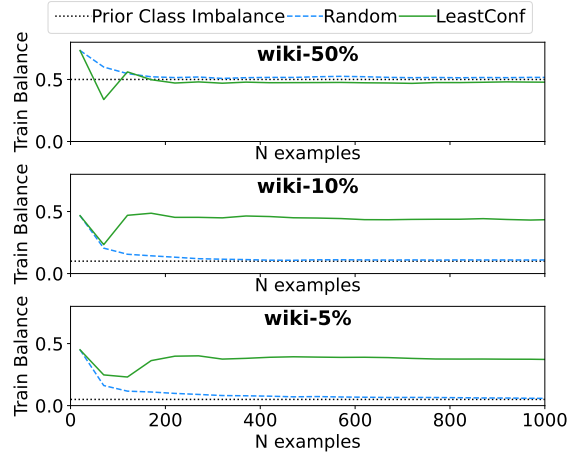


Figure 3: Label imbalance during training (dBERT).

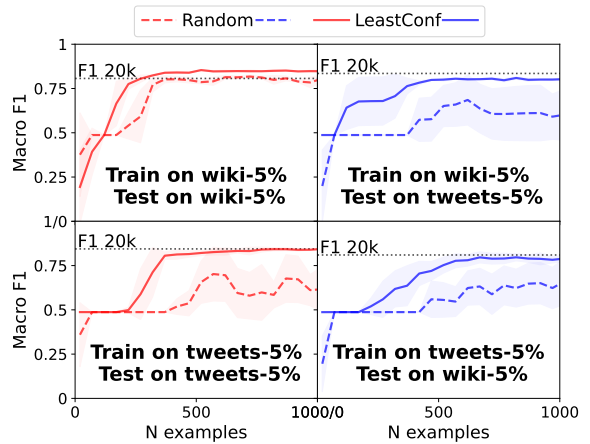


Figure 4: Cross-dataset generalization (dBERT).

However, the compute required to fine-tune a new transformer model in each iteration means AL may have a large environmental footprint (Bender et al., 2021). In some instances, SVMs with AL produce competitive results and have smaller environmental costs. For RQ1.2, we find transformers-based AL is particularly valuable under more extreme class imbalance because it iteratively balances the distribution. Our findings are subject to some limitations, which present avenues for future work.

How does data sampling, class labels and linguistic diversity affect performance? We evaluate against two datasets with pre-existing labels, which we simplify into a binary task. This binarization was required to allow comparison across datasets. The **wiki** dataset samples banned comments and **tweets** samples with keywords and sentiment analysis. While these datasets were the only

publicly-available datasets large enough for this work, [Wiegand et al. \(2019\)](#) shows that they lack diversity, contain numerous biases, and cover abuse which is mostly explicit. This may make it easier for models to learn the task and generalize in fewer examples. Future work should evaluate the success and generalizability of AL for fine-grained labels and implicit abuse.

How does the number of model parameters affect performance? For computational feasibility and environmental concerns, we use distil-BERT but future work could assess if larger transformers models set higher baselines from passive training over the full dataset.

Are certain AL strategies well-suited to abusive language detection? We evaluate three commonly-used AL strategies, finding that Least-Confidence performs best, but none are tailored explicitly to abusive language. Contrastive Active Learning ([Margatina et al., 2021](#)) may be particularly useful: by finding linguistically similar entries on either side of the decision boundary, it may prevent overfitting to certain slurs, profanities or identities.

Do the experimental findings generalize to real-world settings? Our motivation for maximizing efficiency is to reduce financial costs and risk of harm to annotators, which we operationalize in terms of the number of labeled examples they view. In practice, costs are variable because entries which are more ‘uncertain’ to the model may also be more time-consuming, challenging or harmful for humans to label ([Haertel et al., 2015](#)). In a real-world setting, the work of the human annotators must be scaled up and down in response to labeling demands, which may incur additional costs. Crowdsourced annotators can provide labels on demand when a new batch of entries is launched. With an expert annotation team, there may be a cost of paying annotators during re-training. Furthermore, it is important to note that the scope and scale of realized harm depends on both the total number of annotators as well as their identity, positionality and working conditions. While our approach simulates the labeling process with one groundtruth label, we make no assumptions on how this groundtruth is obtained—either via a single annotator or with some aggregation function over multiple annotator votes—so, our method is applicable to any number or constitution of annotators. We only make the

light assumption that less exposure to harm is a good thing—whether that is many people being exposed a little less or few people being exposed a lot less. Future work is needed beyond our simulated set-up to calculate a more realistic cost-benefit ratio of AL, both in terms of financial and psychological costs.

We are exploring these questions in future work but simultaneously encourage the community to consider the need for efficiency in abusive language detection because of the costs, complexities and risk of harm to annotator well-being from inefficient data labeling.

Acknowledgments

This work was supported in part by JADE: Joint Academic Data science Endeavour - 2 under the EPSRC Grant EP/T022205/1, and by the The Alan Turing Institute under EPSRC grant EP/N510129/1. Hannah Rose Kirk was supported by the Economic and Social Research Council grant ES/P000649/1. We thank Paul Röttger and our anonymous reviewers for their comments.

References

- Muhammad Ilham Abidin, Khairil Anwar Notodiputro, and Bagus Sartono. 2021. [Improving Classification Model Performances using an Active Learning Method to Detect Hate Speech in Twitter](#). *Indonesian Journal of Statistics and Its Applications*, 5(1):26–38.
- Md Abul Bashar and Richi Nayak. 2021. [Active Learning for Effectively Fine-Tuning Transfer Learning to Downstream Task](#). *ACM Transactions on Intelligent Systems and Technology (TIST)*, 12(2).
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. [On the dangers of stochastic parrots: Can language models be too big?](#). In *Conference on Fairness, Accountability, and Transparency (FAccT '21)*. ACM, New York, NY, USA.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language Models are Few-Shot Learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.

- Polychronis Charitidis, Stavros Doropoulos, Stavros Vologiannidis, Ioannis Papastergiou, and Sophia Karakeva. 2020. [Towards countering hate speech against journalists on social media](#). *Online Social Networks and Media*, 17:100071.
- 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, ICWSM '17*, pages 512–515.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding](#). *NAACL HLT 2019 - 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies - Proceedings of the Conference*, 1:4171–4186.
- Liat Ein-Dor, Alon Halfon, Ariel Gera, Eyal Shnarch, Lena Dankin, Leshem Choshen, Marina Danilevsky, Ranit Aharonov, Yoav Katz, and Noam Slonim. 2020. [Active Learning for BERT: An Empirical Study](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, pages 7949–7962.
- Mai ElSherief, Vivek Kulkarni, Dana Nguyen, Wang William Y., and Elizabeth Belding. 2018a. Hate Lingo: A Target-based Linguistic Analysis of Hate Speech in Social Media. In *Proceedings of the 12th International AAAI Conference on Web and Social Media, ICWSM '18*.
- Mai ElSherief, Shirin Nilizadeh, Dana Nguyen, Vigna Giovanni, and Elizabeth Belding. 2018b. Peer to Peer Hate: Hate Instigators and Their Targets. In *Proceedings of the 12th International AAAI Conference on Web and Social Media, ICWSM '18*.
- 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](#). *12th International AAAI Conference on Web and Social Media, ICWSM 2018*, pages 491–500.
- Robert J Gabriel. 2018. [Full List of Bad Words and Top Swear Words Banned by Google](#).
- Tianyu Gao, Adam Fisch, and Danqi Chen. 2021. [Making Pre-trained Language Models Better Few-shot Learners](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3816–3830, Online. Association for Computational Linguistics.
- Katharine Gelber and Luke McNamara. 2016. [Evidencing the harms of hate speech](#). *Social Identities*, 22(3):324–341.
- Tarleton Gillespie. 2020. [Content moderation, AI, and the question of scale](#). *Big Data & Society*, 7(2):205395172094323.
- Robbie Haertel, Eric Ringger, Kevin Seppi, and Paul Felt. 2015. [An Analytic and Empirical Evaluation of Return-on-Investment-Based Active Learning](#). In *Proceedings of The 9th Linguistic Annotation Workshop*, pages 11–20. Association for Computational Linguistics.
- Peiyun Hu, Zachary C. Lipton, Anima Anandkumar, and Deva Ramanan. 2018. [Active Learning with Partial Feedback](#). *7th International Conference on Learning Representations, ICLR 2019*.
- Vebjørn Isaksen and Björn Gambäck. 2020. [Using transfer-based language models to detect hateful and offensive language online](#). In *Proceedings of the Fourth Workshop on Online Abuse and Harms*, pages 16–27, Online. Association for Computational Linguistics.
- Anna Koufakou, Endang Wahyu Pamungkas, Valerio Basile, and Viviana Patti. 2020. [HurtBERT: Incorporating lexical features with BERT for the detection of abusive language](#). In *Proceedings of the Fourth Workshop on Online Abuse and Harms*, pages 34–43, Online. Association for Computational Linguistics.
- David D Lewis and William A Gale. 1994. [A Sequential Algorithm for Training Text Classifiers](#). In *SIGIR '94*, pages 3–12. Springer London, London.
- Katerina Margatina, Giorgos Vernikos, Loïc Barrault, and Nikolaos Aletras. 2021. Active learning by acquiring contrastive examples. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 650–663.
- Shruthi Mohan, Apala Guha, Michael Harris, Fred Popowich, Ashley Schuster, and Chris Priebe. 2017. [The Impact of Toxic Language on the Health of Reddit Communities](#). In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 10233 LNAI, pages 51–56. Springer, Cham.
- Ioannis Mollas, Zoe Chrysopoulou, Stamatis Karlos, and Grigorios Tsoumakas. 2020. [ETHOS: an Online Hate Speech Detection Dataset](#).
- 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.
- R Mutanga, Nalindren Naicker, and Oludayo O Olugbara. 2020. Hate speech detection in twitter using transformer methods. *International Journal of Advanced Computer Science and Applications*, 11(01).

- Xi Peng Qiu, Tian Xiang Sun, Yi Ge Xu, Yun Fan Shao, Ning Dai, and Xuan Jing Huang. 2020. [Pre-trained models for natural language processing: A survey](#). *Science China Technological Sciences* 2020 63:10, 63(10):1872–1897.
- Md Mustafizur Rahman, Dinesh Balakrishnan, Dhiraaj Murthy, Mucahid Kutlu, and Matthew Lease. 2021. [An Information Retrieval Approach to Building Datasets for Hate Speech Detection](#). *35th Conference on Neural Information Processing Systems (NeurIPS 2021) Track on Datasets and Benchmarks*.
- Sarah T. Roberts. 2019. *Behind the Screen: Content Moderation in the Shadows of Social Media*. Yale University Press, New Haven, CT.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- Timo Schick and Hinrich Schütze. 2021. [It’s Not Just Size That Matters: Small Language Models Are Also Few-Shot Learners](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2339–2352, Online. Association for Computational Linguistics.
- Christopher Schröder, Lydia Müller, Andreas Niekler, and Martin Potthast. 2021. [Small-text: Active Learning for Text Classification in Python](#).
- Christopher Schröder, Andreas Niekler, and Martin Potthast. 2022. [Revisiting Uncertainty-based Query Strategies for Active Learning with Transformers](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2194–2203, Dublin, Ireland. Association for Computational Linguistics.
- Ozan Sener and Silvio Savarese. 2017. [Active Learning for Convolutional Neural Networks: A Core-Set Approach](#). *6th International Conference on Learning Representations, ICLR 2018 - Conference Track Proceedings*.
- Burr Settles. 2009. [Active Learning Literature Survey](#). Technical Report, University of Wisconsin-Madison Department of Computer Sciences.
- Miriah Steiger, Timir J Bharucha, Sukrit Venkatagiri, Martin J. Riedl, and Matthew Lease. 2021. [The Psychological Well-Being of Content Moderators](#). *CHI ’21: Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–14.
- Bertie Vidgen, Helen Margetts, and Alex Harris. 2019. [How much online abuse is there? A systematic review of evidence for the UK](#). Technical report, The Alan Turing Institute.
- 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. Association for Computational Linguistics.
- Matthew L. Williams, Pete Burnap, Amir Javed, Han Liu, and Sefa Ozalp. 2019. [Hate in the Machine: Anti-Black and Anti-Muslim Social Media Posts as Predictors of Offline Racially and Religiously Aggravated Crime](#). *The British Journal of Criminology*, 60(1):93–117.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi 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 M. Rush. 2019. [HuggingFace’s Transformers: State-of-the-art Natural Language Processing](#).
- 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, Republic and Canton of Geneva, Switzerland. International World Wide Web Conferences Steering Committee.
- Michelle Yuan, Hsuan-Tien Lin, and Jordan Boyd-Graber. 2020. Cold-start active learning through self-supervised language modeling. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7935–7948.

A Details of Dataset Processing and Model Training

We use two English-language datasets which were curated for the task of automated abuse detection (Wulczyn et al., 2017; Founta et al., 2018). The **wiki** dataset can be downloaded from <https://github.com/ewulczyn/wiki-detox> and is licensed under Apache License, Version 2.0. The **tweets** dataset can be downloaded with tweet ids from <https://github.com/ENCASEH2020/hatespeech-twitter>.

These datasets cover two different domains: Wikipedia and Twitter. Each dataset is cleaned by removing extra white space, dropping duplicates and converting usernames, URLs and emoji to special tokens.

We fine-tune distil-roBERTa using the transformers integration with the small-text python package (Wolf et al., 2019; Schröder et al., 2021). distil-roBERTa has six layers, 768 hidden units, and 82M parameters. We encode input texts using the distil-roBERTa tokenizer, with added special tokens for usernames, URLs and emoji. All models were trained for 3 epochs with early stopping based on the cross-validation loss, a learning rate of $2e - 5$ and a weighted Adam optimizer. All other hyperparameters are set to their small-text defaults. In each active learning iteration, we use 10% of each labeled batch for validation. As a baseline to transformers-based AL, we use a support vector machine with no pre-training which we implement with sklearn. To encode a vector representation of input texts, we use a TF-IDF transformation fitted to the training dataset.

All experiments were run on the JADE-2 cluster using one NVIDIA Tesla V100 GPU per experiment. For transformer-models, it took on average 1.5 hours to run each experiment. For SVM, it took less than a minute to run each experiment and these can be easily be run on a CPU. We repeat each experiment three times using three seeds to initialize a pseudo-random number generator.

B Sampling with Keywords

We use a heuristic to weakly label examples from the unlabeled pool to be selected for the initial seed. Keywords are a commonly-used approach (e.g. see Ein-Dor et al., 2020) and searching for text matches is computationally efficient over a large pool of unlabeled examples. However, the keyword heuristic

Table 3: The effect of varied keyword density thresholds on F1, false positive rate (FPR) and false negative rate (FNR).

K	F1	FPR	FNR
wiki			
1.0%	76.0%	2.7%	52.8%
5.0%	69.0%	0.5%	71.8%
10.0%	91.0%	0.1%	87.4%
25.0%	49.0%	0.0%	98.4%
Tweets			
1.0%	85.0%	4.5%	29.6%
5.0%	80.0%	2.9%	42.7%
10.0%	83.0%	0.9%	76.4%
25.0%	75.0%	0.2%	98.5%

only approximates the true label and can introduce biases due to non-abusive use of offense and profanities. In our data, we rely on a keyword density measure (K) which equals the number of keyword matches over the total tokens in a text instance. We then experiment with varied thresholds of $K \in [1\%, 5\%, 10\%, 25\%]$ for a weak label of abusive text. A higher threshold reduces false positives but also decreases true positives. We find a threshold of 5% best balances these directional effects. Making predictions using a keyword heuristic with a 5% cut-off achieves an F1-score relative to the true labels of 69% for wiki and 80% for tweets. Using this threshold, examples are expected to be abusive if the percentage of keywords in total tokens exceeds 5%. We then sample equal numbers of expected abusive and non-abusive examples from the pool, reveal their true labels and initialize the classifier by training over this seed.

C Additional Experimental Analysis

Table 4: The best AL parameters and performance for each classifier (transformers vs SVM).

Dataset	Classifier	Best AL Combinations*				Metrics		
		Seed	Cold	Batch	Query	F1 _{20k} [†]	F1 _{AL}	N ₉₀
wiki50	dBERT	20	Random	50	LC	0.920	0.922	170
	SVM	20	Random	50	LC	0.875	0.838	1520
wiki10	dBERT	20	Heuristic	50	LC	0.859	0.866	170
	SVM	20	Heuristic	50	LC	0.809	0.810	320
wiki5	dBERT	20	Heuristic	50	LC	0.807	0.855	220
	SVM	20	Heuristic	50	LC	0.785	0.780	170
tweets50	dBERT	20	Random	50	LC	0.939	0.938	170
	SVM	20	Random	50	LC	0.931	0.926	220
tweets10	dBERT	20	Heuristic	50	LC	0.904	0.902	220
	SVM	20	Random	50	LC	0.893	0.901	170
tweets5	dBERT	200	Heuristic	50	LC	0.844	0.856	300
	SVM	20	Heuristic	50	LC	0.825	0.830	170

Notes: [†] global metric from passive training over the full dataset

* calculated by averaging the rank performance on F1_{AL}, N₉₀

Tab. 4 shows the best parameters for each dataset and each classifier. In Fig. 6, we present the learning curve and comparisons of each experimental

variable for both datasets and classifiers. In each panel of Fig. 6, we vary one parameter whilst holding all others fixed. This allows us to evaluate the impact of one variable, *ceteris paribus*. Namely, the reference values are those reported in the main paper: seed size of 20 selected by heuristics-based sampling and a batch size of 50 queried by LeastConfidence strategy.

Seed and Batch Size We test two choices for seed size (20, 200), and three choices for batch size (50, 100, 500). We find AL is more efficient with smaller seeds and batch sizes. The F1 score achieved with a seed of 20 and four AL iterations of 50 ($|\mathcal{L}| = 220$) exceeds that reached with a seed of 200 and 0 iterations ($|\mathcal{L}| = 200$) by 55pp for **wiki50**, 4pp for **wiki10**, and 10pp for **wiki5**. Batch sizes of 100 and 500 are less efficient than 50, with 700–1,100 and 150–200 more examples needed for N_{90} , respectively.

Seed Acquisition Strategy (Cold) We evaluate two choices to select the examples for the seed. (1) **Random**: Seed examples are randomly selected. Depending on the class distribution of the unlabeled pool (which, in real world settings, is unknown) only non-abusive content might be identified. For datasets expected to be approximately balanced, a randomly-selected seed has a high probability of including both class labels. (2) **Heuristics**: Seed examples are selected using keywords ($n = 652$), taken from the abusive language literature (Davidson et al., 2017; ElSherief et al., 2018a,b; Gabriel, 2018). For **wiki50**, random- and heuristics-based initialization achieve equivalent N_{90} . However, with a seed of 20, a third of randomly-initialized experiments fail on **wiki10** and all experiments fail for **wiki5**. This shows that when the data is imbalanced, a random seed is suboptimal because both class labels are not observed.

Query Strategy In addition to LeastConfidence (LC), we evaluate two further strategies coupled with dBERT: 1) **GreedyCoreSet** is a data-based diversity strategy which selects items representative of the full set (Sener and Savarese, 2017) and 2) **EmbeddingKMeans** is a data-based diversity strategy which uses a dense embedding representation (such as BERT embeddings) to cluster and sample from the nearest neighbors of the k centroids (Yuan et al., 2020). On our datasets, these two strategies are high performing in terms of the maximum F1 score they achieve over 2,000 exam-

ples, but take longer to learn and are less efficient than LeastConfidence.

D Generalizability of Performance

In the main paper, we present the results of cross-dataset generalization with 5% abuse. In Fig. 5, we demonstrate the equivalent results for all class imbalances and both datasets. In general, tweets is harder to predict than wiki, so we see a larger change in performance when training on tweets and evaluating on wiki. For 50% and 10% abuse, performance is similar across test sets. For 5% abuse, there is a larger difference especially for the random baseline. However, in all cases, the performance of the LeastConfidence strategy generalizes well to out-of-domain testing, at least for these two datasets which are similar in their proportion of explicit abuse (Wiegand et al., 2019).

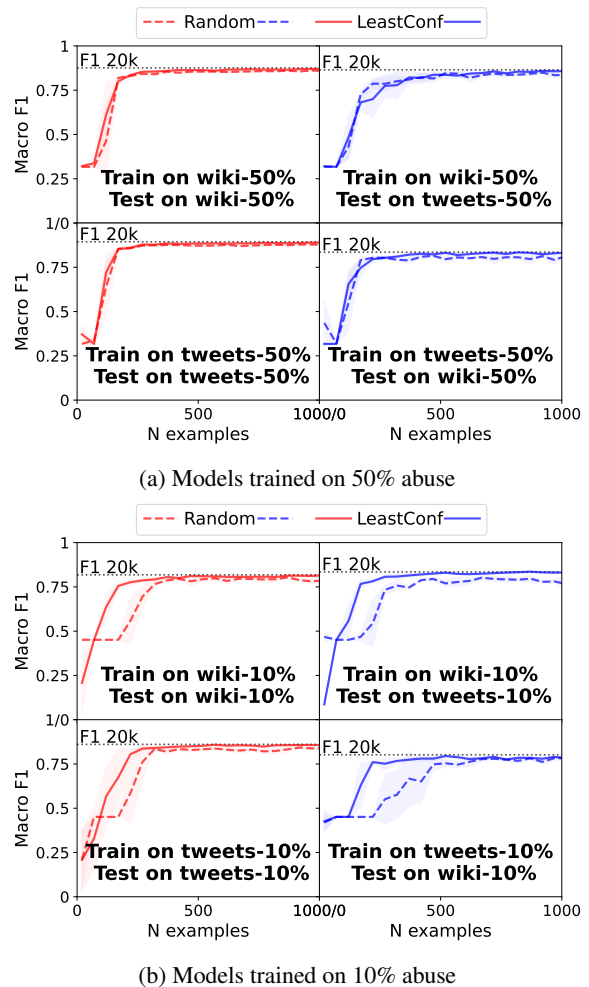
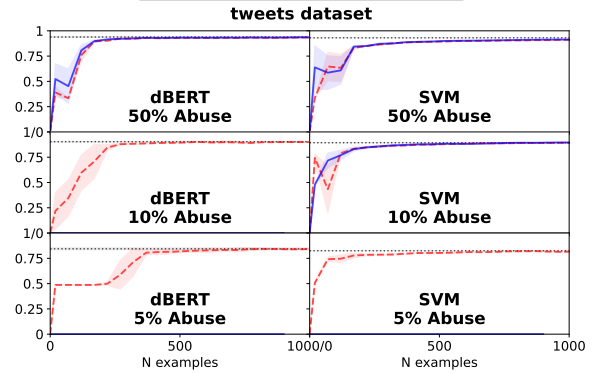

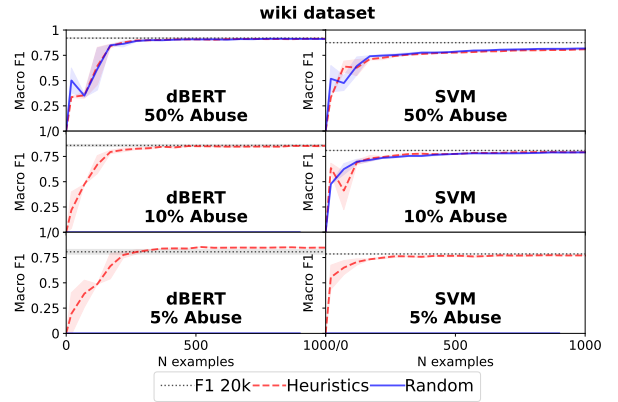
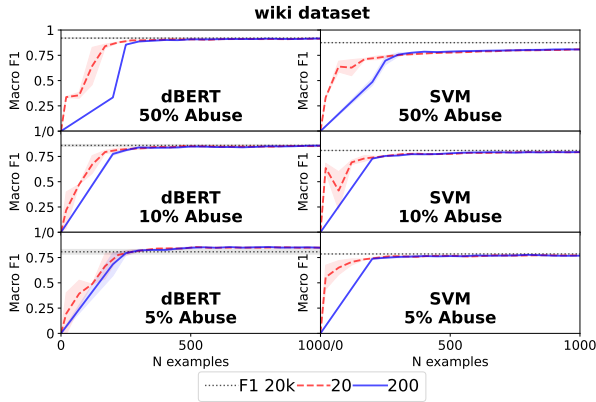
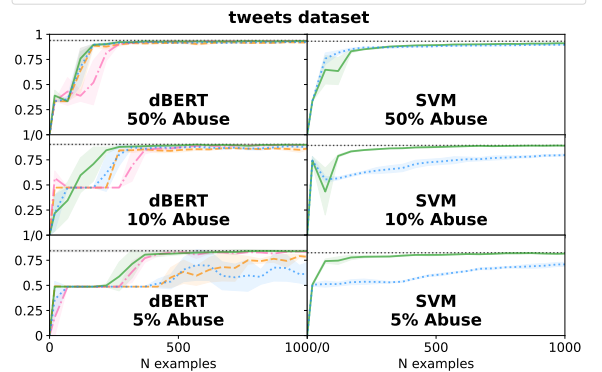
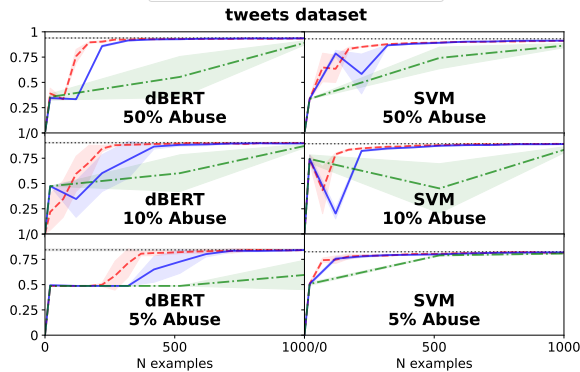
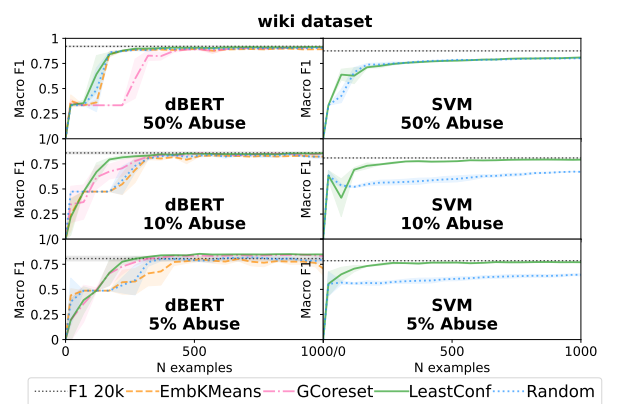
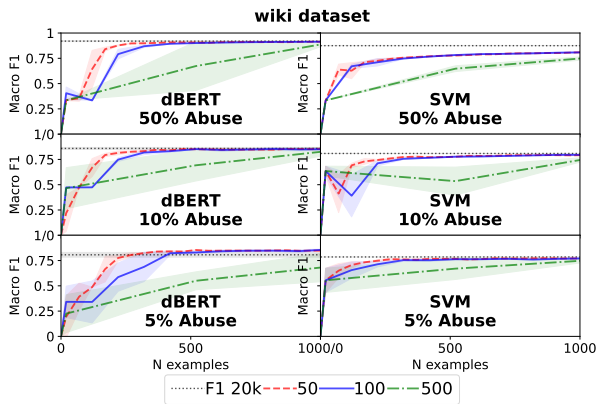


Figure 5: Cross-dataset generalization (dBERT) for 50% and 10% abuse.



(a) Seed Size

(b) Cold Strategy



(c) Batch Size

(d) Query Strategy

Figure 6: Learning curves per dataset-class imbalance pair showing the effect of isolated experimental variables on traditional (SVM) and transformers-based (dBERT) active learning.