

---

# On the Zero-Shot Generalization of Machine-Generated Text Detectors

---

**Xiao Pu**  
Peking University  
puxiao@stu.pku.edu.cn

**Jingyu Zhang**  
Johns Hopkins University  
jzhan237@jhu.edu

**Xiaochuang Han**  
University of Washington  
xhan77@cs.washington.edu

**Yulia Tsvetkov**  
University of Washington  
yuliats@cs.washington.edu

**Tianxing He**  
University of Washington  
goosehe@cs.washington.edu

## Abstract

The rampant proliferation of large language models, fluent enough to generate text indistinguishable from human-written language, gives unprecedented importance to the detection of machine-generated text. This work is motivated by an important research question: How will the detectors of machine-generated text perform on outputs of a new generator, that the detectors were not trained on? We begin by collecting generation data from a wide range of LLMs, and train neural detectors on data from each generator and test its performance on held-out generators. While none of the detectors can generalize to all generators, we observe a consistent and interesting pattern that the detectors trained on data from a medium-size LLM can zero-shot generalize to the larger version. As a concrete application, we demonstrate that robust detectors can be built on an ensemble of training data from medium-sized models.

## 1 Introduction

Thanks to large-scale pretraining and tuning with human feedback [Ouyang et al., 2022], large language models (LLMs) [Chung et al., 2022, Zhang et al., 2022, Touvron et al., 2023] are now able to follow instructions and generate realistic and consistent texts. A prominent example is the recently developed ChatGPT or GPT4 model [OpenAI, 2023], which when instructed, can write documents, create executable code, or answer questions that require world knowledge. In a lot of scenarios, the machine-generated texts have high quality and cannot easily be distinguished from genuine human texts [Dugan et al., 2022, Gehrmann et al., 2019]. These trends give an unprecedented importance to the detection of machine-generated text [Su et al., 2023, Jawahar et al., 2020, Pagnoni et al., 2022a]. A lot of work has been devoted to proposing efficient detection models or algorithms [Mitchell et al., 2023, Kirchenbauer et al., 2023, Zellers et al., 2019]. However, in most studies, the detector is tested on the same generator model that it is trained/tuned on.

This study is motivated by an underexplored research question: How will the detector perform on a different generator that it is not trained on? This question is important due to multiple reasons: (1) LLMs are becoming increasingly large and expensive. Some of the most recent models are either too large to fit into a common GPU (e.g., LLaMA-65B) or require payment from the user [OpenAI, 2023], making the collection of training samples difficult. (2) The number of released LLMs is growing rapidly. In a real application scenario, the detector needs to cover a wide range of LLMs (including the ones the detector is not trained on), instead of only one generator.

In this work, we collect generation data from a wide range of LLMs. We then train neural detectors on data from each generator and test its performance on other generators. Our primary findings include: (1) In many cases, detectors can zero-shot generalize to a held-out generator (Figure 1). In particular, we observe an interesting pattern that the detector for the medium version of an LLM can

---

Code and datasets will be available at <https://github.com/SophiaPx/detectors-generalization>.




Figure 1: The generalization ability of detectors when applied to different generators, measured by Acc-Gap (defined in §2). Detectors for a medium-size generator can zero-shot generalize to the larger-version model (highlighted by dotted green).

generalize to the larger version. (2) None of the detectors generalizes to all generators, implying that an ensemble of detectors/data is necessary for a wider coverage. (3) As a concrete application, we demonstrate that robust detectors can be built on an ensemble of training data from medium-sized models; Excluding large-versions only leads to a minor drop in performance.

## 2 Methodology

We begin by giving an overview of our experiment structure and establish some notations. This study includes detection of a range of popular LLMs (detailed in §3), and we construct train/dev/test sets for each generator. In §4.1, we train neural detectors on data from each generator and test its performance on other generators. In §4.2, we further consider an ensemble setting, where the detector is trained on data composed of multiple generators, and test its generalization ability on held-out generators.

We denote the detector model trained on data from generator model  $M$  as  $D_M$ . Since we will test the accuracy of  $D_M$  on data from different generators, we use  $\text{Acc}_N(D_M)$  to denote the accuracy of  $D_M$  on the test set of generator  $N$ . Finally, we define  $\text{Acc-Gap}_N^{D_M}$  to measure the drop of performance when the detector is trained on generator  $M$  instead of  $N$  itself:

$$\text{Acc-Gap}_N^{D_M} = \text{Acc}_N(D_N) - \text{Acc}_N(D_M). \quad (1)$$

We expect Acc-Gap to be larger than zero in general, and a large Acc-Gap means  $D_M$  has poor generalization on generator  $N$ .

## 3 Experiment Setup

**Generators** We include the detection of a total of 13 popular LMs in our study, including GPT-1 [Radford et al., 2018], GPT-2 models (small, medium, large, and xl) [Radford et al., 2019], GPT-3 (text-davinci-003) [Ouyang et al., 2022], GPT-4 [OpenAI, 2023], three GPT-Neo models (125M, 1.3B and 2.7B) [Black et al., 2021, Gao et al., 2020], GPT-J [Wang, 2021], and LLaMA (7B and 13B) [Touvron et al., 2023].<sup>1</sup>

**Datasets** We consider data from three domains: news, reviews and knowledge. For the news domain, we utilize the RealNews dataset, which is a subset of the c4 dataset [Raffel et al., 2019] named "realnewslike". For the reviews domain, we utilize the IMDBreview dataset [Maas et al., 2011]. As for the knowledge domain, we utilize the Wikipedia dataset [Foundation]<sup>2</sup>. For each dataset, we first randomly sample 5000 real-world human-written samples, with a train/dev/test split ratio of 8:1:1. For all samples, we truncate the first 20 tokens to serve as prompts and feed them into different generators for text continuation, yielding 5000 machine-generated samples. For generation we apply nucleus sampling [Holtzman et al., 2020] with  $p = 0.96$ , following the setting in Pagnoni et al. [2022b]. We truncate each sample so that its length is around 120 tokens. For all training or test sets in this work, we keep the ratio of human and machine text to be 1:1.

**Detectors** For data from each generator, we train a ELECTRA-large model [Clark et al., 2020] as a binary classifier. The detectors were trained for 1 epoch with a learning rate of 5e-6 (training for

<sup>1</sup>There are larger versions of LLaMA, but we find it difficult to fit it into our GPU.

<sup>2</sup>The Wikipedia dataset we used is directly obtained from Hugging Face, data subset "20220301.en" (Page link: <https://huggingface.co/datasets/wikipedia>).

more epochs only gives minimal improvement on the dev set). For the data-mix baseline and pruned models in §4.2, 3 epochs of training is used. We use Adam optimizer [Kingma and Ba, 2014] with  $\beta_1 = 0.9, \beta_2 = 0.999$ . The average accuracy (when tested on the same generator it is trained on) of all detectors in news, review and knowledge domains are 94.1%, 96.2% and 94.9%, separately.

## 4 Experiment Results

### 4.1 On Generalization Ability of Detectors

As explained in §2, we compute Acc-Gap to reflect the generalization ability of detectors trained on each generator. Figure 1 depicts the Acc-Gap of each detector/generator pair. We link from node  $M$  to node  $N$  if  $\text{Acc-Gap}_N^{D_M} < T$  (good generalization), where the threshold  $T$  is set to a small number from  $\{1\%, 2\%, 4\%\}$ . On the other hand, in Figure 2 (Appendix B), node  $M$  is linked to node  $N$  when  $\text{Acc-Gap}_N^{D_M} > 20\%$  (poor generalization). For statistical significance, we utilize bootstrapping [Koehn, 2004] and generate 100 virtual test sets by sampling with replacement from the original test set. We then conduct one-sided t-test and use a  $p$ -value of 0.05.

We observe two interesting patterns shared across the three datasets. First, **the detectors for the medium-version LLMs can generalize to the large-version models**. For example,  $D_{\text{LLaMA7B}}$  generalizes to LLaMA13B, and  $D_{\text{GPT3}}$  generalizes to GPT4.<sup>3</sup> This is somewhat surprising because generations from the large-version generator is commonly considered to have higher quality. Interestingly, the generalization of the reverse direction is weaker on RealNews and IMDBReview. As shown in Table 1, when attempting to generalize from the large-version models to medium ones using ELECTRA detectors, the generalization performance is slightly worse, reflected by a larger Acc-Gap. For the reason behind, we conjecture that comparing to the larger model, the medium generator is making a similar but wider range of artifacts in its generations, leading to a smooth generalization to the detection of the larger model. We also experiment with additional base detectors, e.g. ALBERT Large v2 [Lan et al., 2019] and find that the key observation—that the detectors trained for the medium-size models can generalize to larger-size models—still holds. These results are omitted for brevity.

Second, Figure 2 (Appendix B) shows that **none of the detectors, on its own, can generalize to all generators**. In particular, GPT3 and GPT4 seem “isolated” from other families of generators. This result indicates that if we want an “universal” detector which can cover all generators, an ensemble of detectors/data is necessary. We explore this direction in the next section.

### 4.2 Pruning Out Large-Version LLMs in a Mixed Training Dataset

We now demonstrate a concrete application of our findings, and the following realistic threat scenario is considered: The task is still binary classification but the machine text is composed of generations from a range of models (listed in §3). For simplicity, we use a uniform data ratio for the generators. Following results of the last section, an ensemble of detectors/data is necessary. We begin by comparing two baselines: (1) Model ensemble, where we aggregate predictions from all detectors by majority voting or confidence (probability) average; (2) Data mixing, where we train a new ELECTRA-large detector by mixing up the training data from all generators.<sup>4</sup>

For each baseline detector  $D$ , we report the average accuracy on all generators, and the worst-case accuracy which is  $\min_N \text{Acc}_N(D)$ . Accuracy on the four largest generators is also reported. We conduct experiments on RealNews and IMDBReview datasets, and the results for baselines on are shown in the left part of Table 2. It is shown that the data-mix model outperforms the ensemble approach by a large margin. Therefore, we base our pruning experiments on the data-mix model. Following insights from the last section, we then prune out data from the large-version language

M	N	RealNews		IMDBReview	
		Gap $_N^{D_M}$	Gap $_M^{D_N}$	Gap $_N^{D_M}$	Gap $_M^{D_N}$
<b>GPT3</b>	<b>GPT4</b>	3.64%	5.46%	1.47%	5.48%
<b>LLa7B</b>	<b>LLa13B</b>	-1.11%	1.18%	-1.50%	1.47%
<b>Neo1.3B</b>	<b>Neo2.7B</b>	0.04%	2.16%	-2.31%	3.41%
<b>GPT2lg</b>	<b>GPT2xl</b>	-4.40%	4.94%	-0.02%	0.32%

Table 1: Acc-Gap from medium-version to large-version models, as well as in the opposite direction. The generalization from the medium-version model to the large-version is better than the opposite direction.

<sup>3</sup>Strictly speaking, GPT3 is not a “small version” of GPT4. But they are from the same company, and our experiments consistently show they are strongly related.

<sup>4</sup>We have also tried another baseline where we average parameters from all detectors. However, the performance is worse than the majority-voting baseline, and we omit this result.

	Baseline			Pruned (Proposed)		Pruned (Comparison)	
RealNews Acc(%)	Ensemble		Data-mix	-GPT4 -L13B	-GPT4-GNeo2.7B -L13B-GPT2xl	-GPT3 -GPT4	-L13B -L7B
Average	81.1	81.1	<b>88.0</b>	88.6 (+0.6)	88.2 (+0.2)	79.5 (-8.5)	91.6 (+3.6)
Worst-case	50.6	50.5	<b>84.5</b>	84.6 (+0.1)	83.6 (-0.9)	42.7 (-41.8)	80.8 (-3.7)
GPT4	50.6	50.5	<b>88.2</b>	84.6 (-3.6)	85.9 (-2.3)	42.7 (-45.5)	93.2 (+5.0)
GPT3	52.9	52.8	<b>87.0</b>	87.2 (+0.2)	87.5 (+0.5)	52.3 (-34.7)	91.4 (+4.4)
L13B	58.8	58.4	<b>84.5</b>	85.0 (+0.5)	83.6 (-0.9)	83.5 (-1.0)	80.8 (-3.7)
L7B	61.6	61.3	<b>86.0</b>	86.1 (+0.1)	86.1 (+0.1)	83.7 (-2.3)	82.6 (-3.4)
IMDBReview Acc(%)	Ensemble		Data-mix	-GPT4 -L13B	-GPT4-GNeo2.7B -L13B-GPT2xl	-GPT3 -GPT4	-L13B -L7B
Average	85.2	85.1	<b>94.3</b>	93.2 (-1.1)	94.1 (-0.2)	89.0 (-5.3)	93.7 (-0.6)
Worst-case	52.0	51.8	<b>93.7</b>	92.2 (-1.5)	92.7 (-1.0)	62.6 (-31.1)	89.8 (-3.9)
GPT4	52.0	51.8	<b>94.4</b>	93.4 (-1.0)	94.4 (0)	62.8 (-31.6)	94.4 (0)
GPT3	54.1	54.2	<b>93.7</b>	93.1 (-0.6)	93.5 (-0.2)	62.6 (-31.1)	93.7 (0)
L13B	70.2	70.1	<b>94.0</b>	92.2 (-1.8)	92.7 (-1.3)	93.5 (-0.5)	89.8 (-4.2)
L7B	72.1	71.9	<b>94.1</b>	93.0 (-1.1)	93.9 (-0.2)	93.7 (-0.4)	91.7 (-2.4)

Table 2: Accuracy of the baseline detectors and detectors trained on pruned data. “L13B/7B” refers to the LLaMA 13B/7B generator. We highlight the results for the data-mix model because it serves as the base for the pruned models. It is shown that pruning out the large-version LLMs only induce minimal accuracy loss.

models (i.e., GPT4, GPT-Neo2.7B, LLaMA13B and GPT-2xl) and train a detector by mixing up training data from the remaining generators.<sup>5</sup> The degree of drop on the worst-case accuracy reflects the zero-shot generalization ability of the proposed detector.

Also shown in Table 2, the accuracy of the proposed detectors (both average and worst-case) remains similar to or only slightly decreases compared to the data-mix baseline. Figure 3 in appendix provides detailed information on the changes in accuracy after pruning out four large-version models. The accuracy of the proposed detector only experiences a slight decrease (<3%) for GPT4 and LLaMA13B. **Our results show that in the case of limited budget or computing, data from the medium-version LM can decently approximate the large-version in an ensembled data collection.** On the right part of Table 2, we conduct comparison experiments where both medium and large versions are pruned out. As expected, this results in a worse performance on detection of the pruned generators, reflected by the worst-case accuracy. Especially, the comparison experiment of pruning out both GPT3 and GPT4 is quite alarming: The detector trained from combined data of all other generators only has accuracy around 42% (RealNews) or 62% (IMDBReview). This implies that if OpenAI did not give public access to generations of the two models, existing detectors would fail.

## 5 Related Work

We now discuss the literature most related to our work, and defer a more complete review to Appendix §A. Pagnoni et al. [2022a] demonstrates the degraded performance of trained detectors under different threat scenarios, while the range of generator models is not as wide or up-to-date as our work. Liang et al. [2023] studies the bias of detectors for LLMs in the case of non-native English writers. In a concurrent work, Miresghallah et al. [2023] studies the generalization of detectors under the DetectGPT [Mitchell et al., 2023] algorithm, which is also shown to be far from perfect. Comparing to a trained detector, DetectGPT relies on access to the generator LLM, which might be expensive.

## 6 Conclusion and Discussion

In this work, we observe a generalization relationship among detectors trained on different generators in three domains, where detectors for medium-version models demonstrate the ability to effectively generalize to the larger-version. Building upon this finding, we prune out data from large-version generators in an ensembled training dataset and demonstrate that the performance loss is minimal. Our results indicate that practitioners with limited budget or computing resources can use data from medium-size LLMs as a good approximation for the large version. With the rapid release of various LLMs and generation APIs, a detector needs to cover a wide range of generators. While our work makes some initial progress, our experiments show that the detection of an unseen (or non-public) generator is still a difficult and open question. We hope our work could motivate more research devoted to this important direction.

<sup>5</sup>Our data collection for GPT4 costs around \$450.

## References

- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Gray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=TG8KACxEON>.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Alex Castro-Ros, Marie Pellat, Kevin Robinson, Dasha Valter, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. Scaling instruction-finetuned language models, 2022.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, Todor Mihaylov, Myle Ott, Sam Shleifer, Kurt Shuster, Daniel Simig, Punit Singh Koura, Anjali Sridhar, Tianlu Wang, and Luke Zettlemoyer. Opt: Open pre-trained transformer language models, 2022.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- OpenAI. Gpt-4 technical report, 2023.
- Liam Dugan, Daphne Ippolito, Arun Kirubarajan, Sherry Shi, and Chris Callison-Burch. Real or fake text?: Investigating human ability to detect boundaries between human-written and machine-generated text. *arXiv*, 2022. doi: 10.48550/ARXIV.2212.12672. URL <https://arxiv.org/abs/2212.12672>.
- Sebastian Gehrmann, Hendrik Strobelt, and Alexander Rush. GLTR: Statistical detection and visualization of generated text. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 111–116, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-3019. URL <https://aclanthology.org/P19-3019>.
- Jinyan Su, Terry Yue Zhuo, Di Wang, and Preslav Nakov. Detectllm: Leveraging log rank information for zero-shot detection of machine-generated text, 2023.
- Ganesh Jawahar, Muhammad Abdul-Mageed, and Laks Lakshmanan, V.S. Automatic detection of machine generated text: A critical survey. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 2296–2309, Barcelona, Spain (Online), December 2020. International Committee on Computational Linguistics. doi: 10.18653/v1/2020.coling-main.208. URL <https://aclanthology.org/2020.coling-main.208>.
- Artidoro Pagnoni, Martin Graciarena, and Yulia Tsvetkov. Threat scenarios and best practices to detect neural fake news. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 1233–1249, Gyeongju, Republic of Korea, October 2022a. International Committee on Computational Linguistics. URL <https://aclanthology.org/2022.coling-1.106>.
- Eric Mitchell, Yoonho Lee, Alexander Khazatsky, Christopher D Manning, and Chelsea Finn. Detectgpt: Zero-shot machine-generated text detection using probability curvature. *arXiv preprint arXiv:2301.11305*, 2023.
- John Kirchenbauer, Jonas Geiping, Yuxin Wen, Jonathan Katz, Ian Miers, and Tom Goldstein. A watermark for large language models, 2023.
- Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi. Defending against neural fake news. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 9054–9065. Curran Associates, Inc., 2019. URL <http://papers.nips.cc/paper/9106-defending-against-neural-fake-news.pdf>.

- Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. 2018.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Sid Black, Gao Leo, Phil Wang, Connor Leahy, and Stella Biderman. GPT-Neo: Large Scale Autoregressive Language Modeling with Mesh-Tensorflow, March 2021. URL <https://doi.org/10.5281/zenodo.5297715>.
- Leo Gao, Stella Biderman, Sid Black, Laurence Golding, Travis Hoppe, Charles Foster, Jason Phang, Horace He, Anish Thite, Noa Nabeshima, et al. The pile: An 800gb dataset of diverse text for language modeling. *arXiv preprint arXiv:2101.00027*, 2020.
- Ben Wang. Mesh-Transformer-JAX: Model-Parallel Implementation of Transformer Language Model with JAX. <https://github.com/kingoflolz/mesh-transformer-jax>, May 2021.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv e-prints*, 2019.
- Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 142–150, Portland, Oregon, USA, June 2011. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/P11-1015>.
- Wikimedia Foundation. Wikimedia downloads. URL <https://dumps.wikimedia.org>.
- Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text degeneration. In *International Conference on Learning Representations*, 2020. URL <https://openreview.net/forum?id=rygGQyrFvH>.
- Artidoro Pagnoni, Martin Graciarena, and Yulia Tsvetkov. Threat scenarios and best practices to detect neural fake news. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 1233–1249, 2022b.
- Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. Electra: Pre-training text encoders as discriminators rather than generators. *arXiv preprint arXiv:2003.10555*, 2020.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Philipp Koehn. Statistical significance tests for machine translation evaluation. In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*, pages 388–395, Barcelona, Spain, July 2004. Association for Computational Linguistics. URL <https://aclanthology.org/W04-3250>.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*, 2019.
- Weixin Liang, Mert Yuksekgonul, Yining Mao, Eric Wu, and James Zou. Gpt detectors are biased against non-native english writers, 2023.
- Fatemehsadat Miresghallah, Justus Matterern, Sicun Gao, Reza Shokri, and Taylor Berg-Kirkpatrick. Smaller language models are better black-box machine-generated text detectors, 2023.
- Anton Bakhtin, Sam Gross, Myle Ott, Yuntian Deng, Marc’ Aurelio Ranzato, and Arthur Szlam. Real or fake? learning to discriminate machine from human generated text, 2019.
- Rowan Zellers, Ari Holtzman, Hannah Rashkin, Yonatan Bisk, Ali Farhadi, Franziska Roesner, and Yejin Choi. Defending against neural fake news, 2020.

Irene Solaiman, Miles Brundage, Jack Clark, Amanda Askill, Ariel Herbert-Voss, Jeff Wu, Alec Radford, Gretchen Krueger, Jong Wook Kim, Sarah Kreps, Miles McCain, Alex Newhouse, Jason Blazakis, Kris McGuffie, and Jasmine Wang. Release strategies and the social impacts of language models, 2019.

Daphne Ippolito, Daniel Duckworth, Chris Callison-Burch, and Douglas Eck. Automatic detection of generated text is easiest when humans are fooled, 2020.

Xiaoming Liu, Zhaohan Zhang, Yichen Wang, Yu Lan, and Chao Shen. Coco: Coherence-enhanced machine-generated text detection under data limitation with contrastive learning. *arXiv preprint arXiv:2212.10341*, 2022.

Sahar Abdelnabi and Mario Fritz. Adversarial watermarking transformer: Towards tracing text provenance with data hiding. In *2021 IEEE Symposium on Security and Privacy (SP)*, pages 121–140. IEEE, 2021.

## Supplemental Materials

### A Related Work

Research on detecting machine-generated text can be roughly divided into two categories: supervised training and zero-shot detection (To clarify, in the literature “zero-shot” usually means that the approach does not require training data, while our work focus on zero-shot generalization to the detection of a held-out generator).

In the case of supervised methods, Bakhtin et al. [2019] train an energy-based model to identify machine-generated text. Zellers et al. [2020] train a GROVER detector and finds that models exhibiting superior performance in generating neural disinformation are also highly effective in detecting their own generated content. Both Solaiman et al. [2019] and Ippolito et al. [2020] propose zero-shot approaches to detect machine-generated text and evaluate the capability of pretrained models. Liu et al. [2022] present a coherence-based contrastive learning model to detect the machine-generated text under low-resource scenario. Kirchenbauer et al. [2023] propose a watermarking method [Abdelnabi and Fritz, 2021] which introduces designed noise which is imperceptible to human readers. Mitchell et al. [2023] propose DetectGPT, a zero-shot method that utilizes a novel curvature-based criterion to determine whether a text is generated by a specific model. This approach has demonstrated superior detection capabilities compared to other existing zero-shot methods. While DetectGPT does not require training a separate detector, it relies on access to the generator LLM, which can be costly. Recently, Su et al. [2023] follow up the work of DetectGPT and introduce two new zero-shot methods: DetectLLM-LRR and DetectLLM-NPR.

### B Auxiliary Results

In Figure 3, we present detailed information on the changes in accuracy after pruning out four large-version models. In Figure 2, we plot detector-generator pairs with large (>20%) Acc-Gap on the three datasets. It shows that none of detectors is able to generalize to all generators. For example, all detectors except  $D_{GPT4}$  has large accuracy gap for GPT3.

In Figure 4, 5 and 6, we give detailed heatmaps of Acc-Gap for every detector/generator pair on the three datasets. The reported numbers are calculated as the averages of Acc-Gap obtained by bootstrapping 100 times.




Figure 3: Accuracy comparison on each generator before and after pruning out four large-version LLMs on the RealNews dataset.






Figure 2: Detector-generator pairs with large (>20%) accuracy gap.




Figure 4: Heatmap of Acc-Gap for detector/generator pairs on the RealNews dataset.




Figure 5: Heatmap of Acc-Gap for detector/generator pairs on the IMDBReview dataset.




Figure 6: Heatmap of Acc-Gap for detector/generator pairs on the Wikipedia dataset.