FAR

Exploring Scaling Trends in LLM Robustness

Nikolaus Howe,*[*]* `niki@far.ai`, FAR AI; Mila; Université de Montréal
Michał Zajac, FAR AI
Ian McKenzie, `ian@far.ai`, FAR AI
Oskar Hollinsworth, `oskar@far.ai`, FAR AI
Tom Tseng, `tom@far.ai`, FAR AI
Pierre-Luc Bacon, Mila; Université de Montréal
Adam Gleave, `adam@far.ai`, FAR AI

Language model capabilities predictably improve from scaling a model's size and training data. Motivated by this, increasingly large language models have been trained, yielding an array of impressive capabilities. Yet these models are vulnerable to adversarial prompts, such as "jailbreaks" that hijack models to perform undesired behaviors, posing a significant risk of misuse. Prior work indicates that computer vision models become more robust with model and data scaling, raising the question: does language model robustness also improve with scale? We study this question empirically, finding that larger models respond substantially better to adversarial training, but there is little to no benefit from model scale in the absence of explicit defenses.

1. Introduction

Language models have demonstrated a range of impressive capabilities in tasks such as general reasoning (Hendrycks et al., 2021), graduate-level Q&A (Rein et al., 2023), and code generation (Chen et al., 2021). This growth in capabilities has fueled rapid deployment, with ChatGPT becoming one of the fastest-growing consumer applications in history (Hu, 2023). Moreover, language models are increasingly integrated into larger systems enabling them to take actions in the real world using external tools (OpenAI, 2023; Anthropic, 2024; Google, 2024) and pursue long-term open-ended goals (Richards, 2024; Kinniment et al., 2024).

The advent of language models enables many new tasks to be solved by AI but also introduces novel classes of security vulnerabilities. In particular, a wide variety of adversarial prompts can hijack models (Wei et al., 2023; Zou et al., 2023; Anil et al., 2024). This enables malicious users to bypass safety fine-tuning performed by the designer, unlocking harmful capabilities such as generating compelling misinformation (Spitale et al., 2023; Chen and Shu, 2024). Innocent users are also at risk from attackers using methods such as indirect prompt injections (Abdelnabi et al., 2023) to exploit LLM-driven applications without any awareness or participation by the user.

A key question is whether future, more capable systems will *naturally* become more robust, or if this will instead require a dedicated safety effort. Although current attacks are concerning, the risks could grow still greater with future models capable of more dangerous actions, such as assisting with biological weapon development (Mouton et al., 2023), or with greater *affordances* to interact with the world (Sharkey et al., 2023), such as a virtual assistant for a CEO of a major company. Prior work has found that superhuman Go systems (Wang et al., 2023) are vulnerable to attack, demonstrating that impressive capabilities do not guarantee robustness. However, work has also found that scaling unlabeled pretraining data (Hendrycks et al., 2019; Carmon et al., 2022; Alayrac et al., 2019) and model size (Xie and Yuille, 2019; Huang et al., 2023) improves adversarial robustness in computer vision.

To answer this question, we conduct an empirical investigation into scaling trends for the adversarial robustness of language models. These trends enable us to forecast the robustness of future models, and give us insight into how the offense-defense balance might shift over time. For example, does the cost of conducting an attack against more capable models grow faster

[*]Primary contact for correspondences.

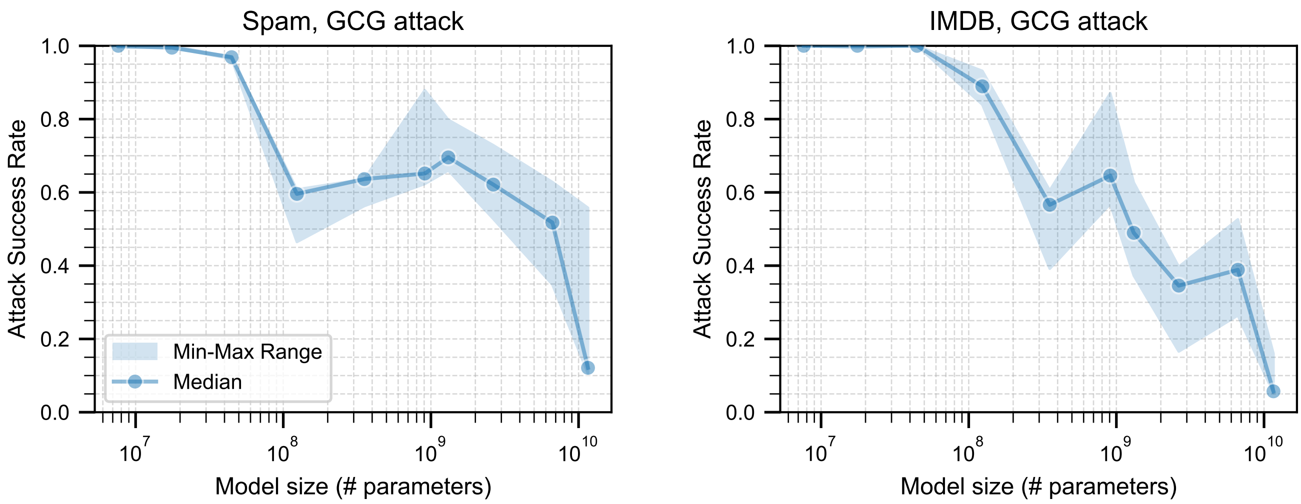

Figure 1: Attack success rate (y-axis) of `GCG` against Pythia models of different sizes (x-axis) fine-tuned on the `Spam` (**left**) and `IMBD` (**right**) tasks. We run three fine-tuning seeds for each model, plotting the median attack success rate and shading the range between the min and max. We observe significant attack success rate variability across model sizes: median robustness does not improve monotonically with scale.

or slower than the defender's cost of training those models?

Concretely, we investigate the robustness of Pythia models ranging from 14M to 12B parameter (Biderman et al., 2023) against two attacks: the *random tokens* baseline and the state-of-the-art *greedy coordinate gradient* attack. We test these models in various simple classification tasks where our models achieve high accuracy on clean (non-adversarial) data.

We first evaluate these pretrained models fine-tuned only on clean data. Larger models tend to be more resistant to attack, but the effect is weak and noisy (Figure 1). By contrast, a clearer scaling trend emerges for models adversarially trained against examples of attacks (Figure 2). Larger models are both more sample efficient, becoming more robust with fewer examples, and converge to be more robust given a sufficient number of examples. Moreover, adversarial training against one attack transfers protection to similar attacks, with the transfer being *stronger* for larger models (Figure 3b).

2. Related Work

Adversarial examples were first identified in image classifiers (Szegedy et al., 2014), but have since been found for systems performing image captioning (Xu et al., 2019; Zhang et al., 2020), speech recognition (Cisse et al., 2017; Alzantot et al., 2018; Schönherr et al., 2018), and reinforcement learning (Huang et al., 2017; Gleave et al., 2020; Ilahi et al., 2022). Moreover, a range of adversarial threat models (Gilmer et al., 2018) give rise to viable attacks.

Most recently, many qualitatively different vulnerabilities have been found in language models, from human-understandable "jailbreaks" (Wei et al., 2023) to seemingly gibberish adversarial suffixes (Wallace et al., 2021; Zou et al., 2023). Simple methods such as perplexity filtering and paraphrasing defend against some of these attacks (Jain et al., 2023). However, these defenses can easily be bypassed by more sophisticated methods (Zhu et al., 2023). Adversarial training shows more promise as a defense (Ziegler et al., 2022), and is the focus of our analysis.

The determinants of adversarial robustness have been well-studied in computer vision. One line of scholarship proposes a fundamental tradeoff between robustness and accuracy (Tsipras et al., 2019): exploitable models are simply relying on non-robust features (Ilyas et al., 2019), which improve training performance but hurt robustness. Other work has emphasized what *does* improve robustness. Scaling unlabeled pretraining data (Hendrycks et al., 2019; Carmon et al., 2022; Alayrac et al., 2019), model depth (Xie and Yuille, 2019) and model width (Huang et al., 2023) improves adversarial robustness in computer vision. However, other work shows that computer vision adversarial robustness scales too slowly to be a full solution (Debenedetti et al., 2023; Bartoldson et al., 2024).

Language model scaling laws (Hestness et al., 2017; Rosenfeld et al., 2019; Kaplan et al., 2020; Hoff-

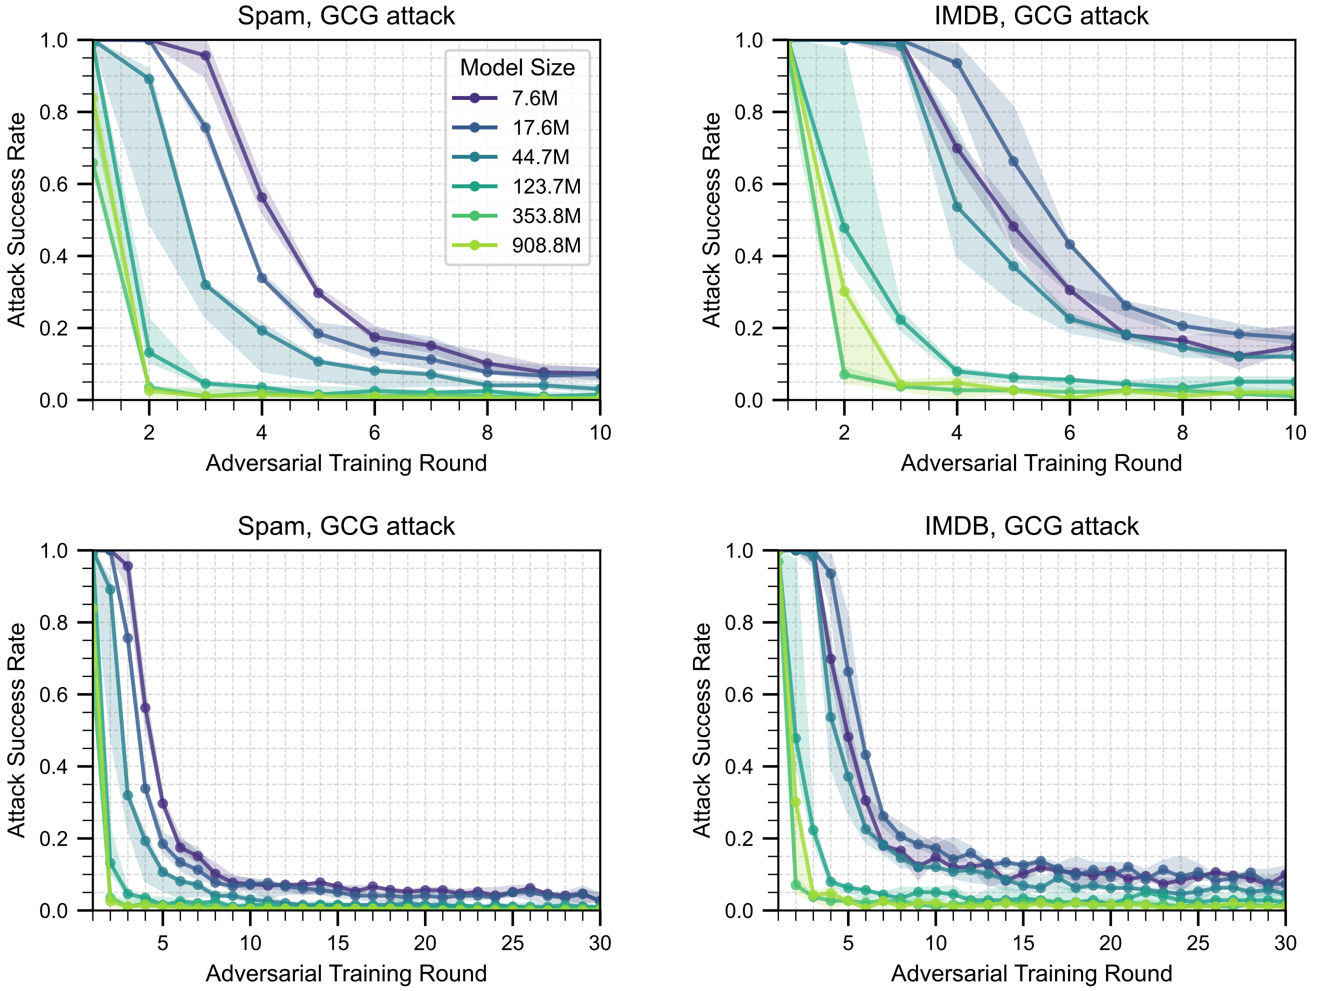

Figure 2: Attack success rate (y-axis) of `GCG` against Pythia models of varying sizes (line color) on `Spam` (**left**) and `IMDB` (**right**) during adversarial training (x-axis) against `GCG` over 10 rounds (**top**) and over 30 rounds (**bottom**). See Figure 16 for a zoomed-in view of the final 10 rounds of the 30-round adversarial training. We plot the median over three seeds and shade the region between the min and max. We observe larger models are more sample efficient, and appear to converge to a higher robustness (lower attack success rate).

mann et al., 2022) have shown that increasing compute improves performance across many tasks and domains (Chen et al., 2021; Hernandez et al., 2021). However, scaling does not solve all problems (Lin et al., 2022; McKenzie et al., 2023). There has been only limited work on scaling laws for adversarial robustness in language models, with mixed results. Ganguli et al. (2022) show that LLMs become harder to attack with scale—but Anil et al. (2024) find that some attacks become *more successful* with scale.

3. Experimental Methodology

We test models in the binary classification setting, as it is the simplest context in which to study LLM robustness. Crucially, binary classification allows us to measure robustness by the **attack success rate**, defined as the proportion of examples correctly classified by the model before attack that are incorrectly classified after attack.[1] We adapt pretrained models for classification by replacing the unembedding layer with a randomly initialized classification head, and then fine-tune the models on each task.

Tasks We consider four tasks in our experiments, the latter two developed by us for this project:

- `Spam` (Metsis et al., 2006): Given the subject and body of an email, is it spam or not?
- `IMDB` (Maas et al., 2011): Given a movie review,

[1] We assume that the attack does not, in fact, change the ground truth label of the datapoint. This is guaranteed by construction for some of our simple procedurally generated tasks, and was manually validated on a random sample of datapoints in other tasks.

is the sentiment positive or negative?

- PasswordMatch: Given two strings in the prompt, are they exactly equal?
- WordLength: Given two words in the prompt, is the first word shorter than the second?

Spam and IMDB were chosen as standard natural language processing classification tasks. PasswordMatch was inspired by TensorTrust (Toyer et al., 2023), and WordLength by the RuLES dataset (Mu et al., 2023). Both PasswordMatch and WordLength were designed to be easily procedurally generated and have ground truth labels that can be checked algorithmically. For brevity, we report on Spam and IMDB in the main text, with plots for other tasks deferred to Appendices D and E. We provide example datapoints and details about the datasets in Appendix B.

Models We test the Pythia model family (Biderman et al., 2023). These models range in size from 14M to 12B parameters (or 7.6M to 11.6B when used with a classification head). All models were trained to predict the next token on the same dataset following the same training protocol, allowing us to isolate model scale from other confounding factors.

Attacks Our attacks append an adversarial suffix of N tokens to the prompt. We use two different procedures to generate this adversarial suffix: a random token baseline (RandomToken) and the state-of-the-art greedy coordinate gradient attack (GCG; Zou et al., 2023). RandomToken was chosen due to its simplicity. GCG was chosen as it is currently one of the most effective attacks on language models.

In the RandomToken baseline, the N tokens are chosen uniformly at random from the model's vocabulary. We evaluate the model on the attacked text and then repeat the process with another sample of N random tokens until the model is successfully attacked or an appointed budget for model calls is exhausted.

In GCG (Zou et al., 2023), the N tokens are initialized arbitrarily and then greedily optimized over multiple rounds. In each round, the gradient of the loss function with respect to the attack tokens is computed. This gradient is used to compute a set of promising single-token modifications, from which the best candidate is selected and used in the next round. To make this attack work in the classification setting, we minimize the cross-entropy loss between the predicted label and the target label.

In our experiments, we always use $N = 10$ tokens. For more details about the attacks and hyperparameters used, see Appendix C.

4. Fine-tuning

Figure 1 shows the robustness of fine-tuned models against the GCG attack. The attack is generally less successful on larger models, but model size alone does not explain all the variance in attack success rate. We observe similarly large random variation in attack success across model sizes on other tasks and with other attacks; for more details, see Appendix D.2.

As described in Section 3, we use the Pythia models, which range from 7.6M to 11.6B parameters after replacing the unembedding matrix with a classification head.[2] We fine-tune all models for a single epoch with default hyperparameters from HuggingFace Transformers (Wolf et al., 2019), except for the learning rate which we set to $1\mathrm{e}{-5}$. All models reach $> 83\%$ accuracy on all tasks, with larger models generally performing better (see Appendix D.1 for the final validation performance of all models on all tasks). We then evaluate the fine-tuned models against adversarial attacks on an unseen validation dataset.

To understand the source of the variability in model robustness shown by our experiments, we varied 1) the pretraining checkpoint,[3] and 2) the random seeds used to initialize the classification head before fine-tuning. Both factors led to significant variability in model robustness, with pretraining checkpoint contributing significantly more variability. The variability was comparable or greater than that of an order of magnitude of model scaling, indicating that out-of-the-box robustness on a given task is heavily influenced by the randomness of the pretraining procedure itself.

This initial result suggests that we cannot rely on scale alone to solve the problem of robustness. However, in practice, we would apply a defense to a model prior to deploying it in a security-critical setting. In the following section, we consider whether scale enables defenses to more effectively improve model robustness.

5. Adversarial training

In this section, we explore how model size impacts robustness when performing adversarial training. Figure 2 evaluates the robustness of Pythia models to the GCG attack when adversarially trained against the same attack. We see a much cleaner trend than in the fine-tuning only case: larger models gain robustness

[2]In all figures, we report the actual parameter count of the classification model, and not the pretrained model it was derived from.

[3]The Pythia models provide checkpoints from earlier stages of pretraining. We used various checkpoints from the final 10% of pretraining as a starting point for fine-tuning.

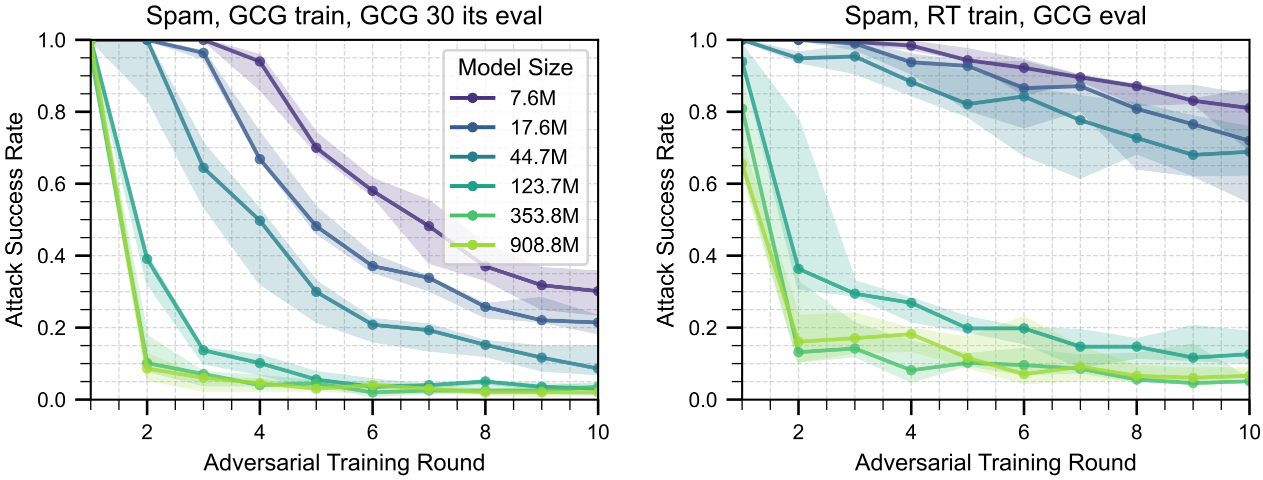

(a) Adversarial training against 10-iteration `GCG`, with evaluation against 30-iteration `GCG`. All models show some transfer of their defense to this stronger attack, with larger models doing so more effectively.

(b) Adversarial training against 10-iteration `RandomToken`, with evaluation against 10-iteration `GCG`. ≥ 100M parameter models show strong defense transfer, while smaller models struggle against the new attack.

Figure 3: Attack success rate (y-axis) against Pythia models of varying sizes (line color) during adversarial training (x-axis).

more quickly and converge to be more robust than smaller models. These results suggest that model size is a strong predictor of robustness—so long as the model is explicitly optimized for robustness. We observe similar behavior across the other two datasets and two attacks; see Appendix E for these plots, including extensions for up to 30 adversarial training rounds.

We perform adversarial training by iteratively training our model on a training dataset, evaluating the model on attacked examples, and then adding successful attack examples to the training dataset. Simultaneously, we evaluate model performance on a held-out attacked validation dataset. This procedure is illustrated in Figure 12.

In our experiments, we initialize the training dataset to consist of 2000 clean examples, and add 200 adversarial examples to the training dataset each round. We repeat the train-attack-add loop 30 times (here we only show the first 10 rounds; see Appendix E for the full 30-round plots). Since adversarial examples are only added after the first training round, the models here were trained for a single epoch on the 2000 clean datapoints before being adversarially attacked.

We perform adversarial training on Pythia models ranging from 7.6 to 909 million parameters after replacing the unembedding layer with a classification head.[4] Table 1 in Appendix A enumerates all model sizes along with corresponding plot colors.

5.1. Robustness transfer

In practice, we often do not have the luxury of knowing the exact attack method an adversary may employ against our model. For practical deployments, we therefore need adversarial training on a handful of attacks to provide more general robustness against other unforeseen attacks. In this subsection, we study whether we observe this transfer in robustness between attacks—and how model scale affects the transfer.

First, we explore whether robustness from adversarial training transfers to a stronger attack from the same family. To do this, we adversarially train using the procedure described above with GCG for 10 iterations as our training attack. We then evaluate on GCG for 30 iterations, a stronger attack. Figure 3a shows that larger models are more robust to this in-distribution, stronger attack. Although the transfer is imperfect—the models do, of course, lose against 30-iteration GCG more than against 10-iteration GCG—the performance is much better than the undefended (fine-tuned) models, which lose approximately 100% of the time.

This is a promising result. Yet, what happens if our models experience an attack that is not only stronger but also uses a different method than the one on which they were adversarially trained? We investigate this question by performing adversarial training against `RandomToken` and evaluating against the `GCG` attack.

Figure 3b shows models adversarially trained on

[4] Specifically, we use the `pythia-14m` to `pythia-1b` models loaded as `AutoModelForSequenceClassification`.

RandomToken do perform better than undefended models, though the effect is weaker. Critically, the extent to which transfer occurs varies drastically across models. In particular, the models with more than 100 million parameters all show strong transfer behavior, with the attack success rate falling below 25% after just 4 rounds of adversarial training. On the other hand, models with fewer than 100 million parameters struggle to transfer their robustness against the RandomToken attack to the stronger GCG attack, with the attack success rate still near 70% on the strongest model even after 10 adversarial training rounds.

This finding is encouraging as it suggests that, for sufficiently large models, robustness will transfer across attacks. It appears that this transfer might be a property that emerges with sufficient scale, similarly to other emergent properties like the ability to use a scratchpad for addition or the utility of instruction fine-tuning (Wei et al., 2022). While we cannot say with certainty that such transfer of robustness generalizes outside the settings and attacks considered in this work, it seems plausible that it would, and indeed, that scaling to further orders of magnitude could unlock more general transfer to a wider variety of attack methodologies and strengths.

6. Conclusion

Our results demonstrate that larger Pythia models benefit more from adversarial training than smaller Pythia models across a variety of classification tasks. An important direction for future work is to validate this trend holds in a broader variety of settings. In particular, we plan to study generative tasks and how factors such as task complexity affect robustness. We also plan to investigate different model families, including larger models.

A key application of scaling trends is to inform appropriate sizing of models to maximize robustness given a fixed defender compute budget. Although larger models are more sample efficient with a fixed number of adversarial training time steps, each adversarial training step is more computationally expensive than with smaller models. For example, Figure 2 shows that performing 8 adversarial training rounds on the 17.6M parameter model results in better robustness than performing 4 adversarial training rounds on the 44.7M parameter model, and a quick calculation shows that it is slightly less expensive to train (see Appendix E.3). However, using a smaller model is not always better, since there are diminishing returns to adversarial training, with larger models appearing to converge to be more robust.

Although scale can improve robustness, our results make clear that it is far from the only determinant. For example, a small adversarially trained model is more robust than a large model fine-tuned only on clean data. We expect that achieving robust language models will require innovations in defense techniques as well as scaling model pretraining and defense training. Scaling trends both enable us to measure how far we are from achieving robustness by scale alone and enable us to identify defense techniques that can better leverage scale to produce more robust models.

Acknowledgements

The authors thank ChengCheng Tan for her assistance in formatting this document, and Daniel Pandori for his contributions to the codebase during the early stages of this project. Nikolaus Howe thanks the Natural Sciences and Engineering Research Council of Canada (NSERC) for their support via the Vanier Canada Graduate Scholarship.

References

Sahar Abdelnabi, Kai Greshake, Shailesh Mishra, Christoph Endres, Thorsten Holz, and Mario Fritz. Not what you've signed up for: Compromising real-world LLM-integrated applications with indirect prompt injection. In *AISec*, page 79–90, 2023.

Jean-Baptiste Alayrac, Jonathan Uesato, Po-Sen Huang, Alhussein Fawzi, Robert Stanforth, and Pushmeet Kohli. Are Labels Required for Improving Adversarial Robustness? In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://papers.nips.cc/paper_files/paper/2019/hash/bea6cfd50b4f5e3c735a972cf0eb8450-Abstract.html.

Moustafa Alzantot, Bharathan Balaji, and Mani Srivastava. Did you hear that? Adversarial examples against automatic speech recognition, 2018. URL https://arxiv.org/abs/1808.05665.

Cem Anil, Esin Durmus, Mrinank Sharma, Joe Benton, Sandipan Kundu, Joshua Batson, Nina Rimsky, Meg Tong, Jesse Mu, Daniel Ford, Francesco Mosconi, Rajashree Agrawal, Rylan Schaeffer, Naomi Bashkansky, Samuel Svenningsen, Mike Lambert, Ansh Radhakrishnan, Carson Denison, Evan J Hubinger, Yuntao Bai, Trenton Bricken, Timothy Maxwell, Nicholas Schiefer, Jamie Sully, Alex Tamkin, Tamera Lanham, Karina Nguyen, Tomasz Korbak, Jared Kaplan, Deep Ganguli, Samuel R Bowman, Ethan Perez, Roger Grosse, and David Duvenaud. Many-shot Jailbreaking, 2024. URL https://www-cdn.anthropic.com/af5633c94ed2beb282f6a53c595eb437e8e7b630/Many_Shot_Jailbreaking__2024_04_02_0936.pdf.

Anthropic. Tool use (function calling), 2024. URL https://archive.ph/EqXCz.

Brian R. Bartoldson, James Diffenderfer, Konstantinos Parasyris, and Bhavya Kailkhura. Adversarial Robustness Limits via Scaling-Law and Human-Alignment Studies, April 2024. URL http://arxiv.org/abs/2404.09349.

Stella Biderman, Hailey Schoelkopf, Quentin Gregory Anthony, Herbie Bradley, Kyle O'Brien, Eric Hallahan, Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, et al. Pythia: A suite for analyzing large language models across training and scaling. In *International Conference on Machine Learning*, pages 2397–2430. PMLR, 2023.

Yair Carmon, Aditi Raghunathan, Ludwig Schmidt, Percy Liang, and John C. Duchi. Unlabeled Data Improves Adversarial Robustness, January 2022. URL http://arxiv.org/abs/1905.13736.

Canyu Chen and Kai Shu. Can LLM-generated misinformation be detected? In *International Conference on Learning Representations*, 2024.

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating Large Language Models Trained on Code, July 2021. URL http://arxiv.org/abs/2107.03374.

Moustapha M Cisse, Yossi Adi, Natalia Neverova, and Joseph Keshet. Houdini: Fooling deep structured visual and speech recognition models with adversarial examples. In *Advances in Neural Information Processing Systems*, volume 30, 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/hash/d494020ff8ec181ef98ed97ac3f25453-Abstract.html.

Edoardo Debenedetti, Zishen Wan, Maksym Andriushchenko, Vikash Sehwag, Kshitij Bhardwaj, and Bhavya Kailkhura. Scaling Compute Is Not All You Need for Adversarial Robustness, December 2023. URL http://arxiv.org/abs/2312.13131.

Deep Ganguli, Liane Lovitt, Jackson Kernion, Amanda Askell, Yuntao Bai, Saurav Kadavath, Ben Mann, Ethan Perez, Nicholas Schiefer, Kamal Ndousse, Andy Jones, Sam Bowman, Anna Chen, Tom Conerly, Nova DasSarma, Dawn Drain, Nelson Elhage, Sheer El-Showk, Stanislav Fort, Zac Hatfield-Dodds, Tom Henighan, Danny Hernandez, Tristan Hume, Josh Jacobson, Scott Johnston, Shauna Kravec, Catherine Olsson, Sam Ringer, Eli Tran-Johnson, Dario Amodei, Tom Brown, Nicholas Joseph, Sam McCandlish, Chris Olah, Jared Kaplan, and Jack Clark. Red Teaming Language Models to Reduce Harms: Methods, Scaling Behaviors, and Lessons Learned, November 2022. URL http://arxiv.org/abs/2209.07858.

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.

Justin Gilmer, Luke Metz, Fartash Faghri, Samuel S. Schoenholz, Maithra Raghu, Martin Wattenberg, and Ian Goodfellow. Adversarial Spheres, September 2018. URL http://arxiv.org/abs/1801.02774.

Adam Gleave, Michael Dennis, Cody Wild, Neel Kant, Sergey Levine, and Stuart Russell. Adversarial policies: Attacking deep reinforcement learning. In *International Conference on Learning Representations*, 2020.

Google. Function calling — Google AI for developers, 2024. URL https://archive.ph/YGJHJ.

Dan Hendrycks, Kimin Lee, and Mantas Mazeika. Using Pre-Training Can Improve Model Robustness and Uncertainty. In *International Conference on Machine Learning*, pages 2712–2721. PMLR, May 2019. URL https://proceedings.mlr.press/v97/hendrycks19a.html. ISSN: 2640-3498.

Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=d7KBjmI3GmQ.

Danny Hernandez, Jared Kaplan, Tom Henighan, and Sam McCandlish. Scaling Laws for Transfer, February 2021. URL http://arxiv.org/abs/2102.01293.

Joel Hestness, Sharan Narang, Newsha Ardalani, Gregory Diamos, Heewoo Jun, Hassan Kianinejad, Md Mostofa Ali Patwary, Yang Yang, and Yanqi Zhou. Deep Learning Scaling is Predictable, Empirically, December 2017. URL http://arxiv.org/abs/1712.00409.

Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katie Millican, George van den Driessche, Bogdan Damoc, Aurelia Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Jack W. Rae, Oriol Vinyals, and Laurent Sifre. Training Compute-Optimal Large Language Models, March 2022. URL http://arxiv.org/abs/2203.15556.

Krystal Hu. ChatGPT sets record for fastest-growing user base – analyst note. *Reuters*, 2023.

Sandy H. Huang, Nicolas Papernot, Ian J. Goodfellow, Yan Duan, and Pieter Abbeel. Adversarial attacks on neural network policies. arXiv:1702.02284v1 [cs.LG], 2017.

Shihua Huang, Zhichao Lu, Kalyanmoy Deb, and Vishnu Naresh Boddeti. Revisiting Residual Networks for Adversarial Robustness. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8202–8211, Vancouver, BC, Canada, June 2023. IEEE. ISBN 9798350301298. doi: 10.1109/CVPR52729.2023.00793. URL https://ieeexplore.ieee.org/document/10204909/.

Inaam Ilahi, Muhammad Usama, Junaid Qadir, Muhammad Umar Janjua, Ala Al-Fuqaha, Dinh Thai Hoang, and Dusit Niyato. Challenges and countermeasures for adversarial attacks on deep reinforcement learning. *IEEE TAI*, 3(2):90–109, 2022.

Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Logan Engstrom, Brandon Tran, and Aleksander Madry. Adversarial Examples Are Not Bugs, They Are Features. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL https://papers.nips.cc/paper_files/paper/2019/hash/e2c420d928d4bf8ce0ff2ec19b371514-Abstract.html.

Neel Jain, Avi Schwarzschild, Yuxin Wen, Gowthami Somepalli, John Kirchenbauer, Ping yeh Chiang, Micah Goldblum, Aniruddha Saha, Jonas Geiping, and Tom Goldstein. Baseline defenses for adversarial attacks against aligned language models, 2023. URL https://arxiv.org/abs/2309.00614.

Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling Laws for Neural Language Models, January 2020. URL http://arxiv.org/abs/2001.08361.

Megan Kinniment, Lucas Jun Koba Sato, Haoxing Du, Brian Goodrich, Max Hasin, Lawrence Chan, Luke Harold Miles, Tao R. Lin, Hjalmar Wijk, Joel Burget, Aaron Ho, Elizabeth Barnes, and Paul Christiano. Evaluating language-model agents on realistic autonomous tasks, 2024. URL https://arxiv.org/abs/2312.11671.

Stephanie Lin, Jacob Hilton, and Owain Evans. TruthfulQA: Measuring How Models Mimic Human Falsehoods, May 2022. URL http://arxiv.org/abs/2109.07958.

Andrew L. Maas, Raymond E. Daly, Peter T. Pham, Dan Huang, Andrew Y. Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In *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.

Ian R. McKenzie, Alexander Lyzhov, Michael Martin Pieler, Alicia Parrish, Aaron Mueller, Ameya Prabhu, Euan McLean, Xudong Shen, Joe Cavanagh, Andrew George Gritsevskiy, Derik Kauffman, Aaron T. Kirtland, Zhengping Zhou, Yuhui Zhang, Sicong Huang, Daniel Wurgaft, Max Weiss, Alexis Ross, Gabriel Recchia, Alisa Liu, Jiacheng Liu, Tom Tseng, Tomasz Korbak, Najoung Kim, Samuel R. Bowman, and Ethan Perez. Inverse Scaling: When Bigger Isn't Better. *Transactions on Machine Learning Research*, June 2023. ISSN 2835-8856. URL https://openreview.net/forum?id=DwgRm72GQF.

Vangelis Metsis, Ion Androutsopoulos, and Georgios Paliouras. Spam Filtering with Naive Bayes - Which Naive Bayes? In *Conference on Email and Anti-Spam*, 2006. URL https://www2.aueb.gr/users/ion/docs/ceas2006_paper.pdf.

Christopher A. Mouton, Caleb Lucas, and Ella Guest. *The Operational Risks of AI in Large-Scale Biological Attacks: A Red-Team Approach*. RAND Corporation, 2023.

Norman Mu, Sarah Chen, Zifan Wang, Sizhe Chen, David Karamardian, Lulwa Aljeraisy, Basel Alomair, Dan Hendrycks, and David Wagner. Can LLMs follow simple rules? *arXiv*, 2023. URL https://arxiv.org/abs/2311.04235.

OpenAI. Assistants API documentation, 2023. URL https://archive.ph/8Az8d.

David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R. Bowman. GPQA: A graduate-level google-proof q&a benchmark, 2023. URL https://arxiv.org/abs/2311.12022.

Toran Bruce Richards. Auto-gpt: An autonomous GPT-4 experiment, 2024. URL https://github.com/Significant-Gravitas/AutoGPT/.

Jonathan S. Rosenfeld, Amir Rosenfeld, Yonatan Belinkov, and Nir Shavit. A Constructive Prediction of the Generalization Error Across Scales, December 2019. URL http://arxiv.org/abs/1909.12673.

Lea Schönherr, Katharina Kohls, Steffen Zeiler, Thorsten Holz, and Dorothea Kolossa. Adversarial attacks against automatic speech recognition systems via psychoacoustic hiding, 2018.

Lee Sharkey, Clíodhna Ní Ghuidhir, Dan Braun, Jérémy Scheurer, Mikita Balesni, Lucius Bushnaq, Charlotte Stix, and Marius Hobbhahn. A causal framework for AI regulation and auditing. Technical report, Apollo Research, 2023.

Giovanni Spitale, Nikola Biller-Andorno, and Federico Germani. AI model GPT-3 (dis)informs us better than humans. *Science Advances*, 9(26), 2023.

Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. Intriguing properties of neural networks, 2014. URL https://arxiv.org/abs/1312.6199.

Sam Toyer, Olivia Watkins, Ethan Adrian Mendes, Justin Svegliato, Luke Bailey, Tiffany Wang, Isaac Ong, Karim Elmaaroufi, Pieter Abbeel, Trevor Darrell, Alan Ritter, and Stuart Russell. Tensor Trust: Interpretable prompt injection attacks from an online game, 2023. URL https://arxiv.org/abs/2311.01011.

Dimitris Tsipras, Shibani Santurkar, Logan Engstrom, Alexander Turner, and Aleksander Madry. Robustness may be at odds with accuracy. In *International Conference on Learning Representations*, 2019. URL https://arxiv.org/abs/1805.12152.

Eric Wallace, Shi Feng, Nikhil Kandpal, Matt Gardner, and Sameer Singh. Universal Adversarial Triggers for Attacking and Analyzing NLP, January 2021. URL http://arxiv.org/abs/1908.07125.

Tony Tong Wang, Adam Gleave, Tom Tseng, Kellin Pelrine, Nora Belrose, Joseph Miller, Michael D Dennis, Yawen Duan, Viktor Pogrebniak, Sergey Levine, and Stuart Russell. Adversarial policies beat superhuman Go AIs. In *International Conference on Machine Learning*, pages 35655–35739. PMLR, 2023.

Alexander Wei, Nika Haghtalab, and Jacob Steinhardt. Jailbroken: How Does LLM Safety Training Fail?, July 2023. URL http://arxiv.org/abs/2307.02483.

Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682*, 2022. URL https://arxiv.org/abs/2206.07682.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. HuggingFace's transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*, 2019. URL https://arxiv.org/abs/1910.03771.

Cihang Xie and Alan Yuille. Intriguing Properties of Adversarial Training at Scale. In *International Conference on Learning Representations*, September

2019. URL https://openreview.net/forum?id=HyxJhCEFDS.

Yan Xu, Baoyuan Wu, Fumin Shen, Yanbo Fan, Yong Zhang, Heng Tao Shen, and Wei Liu. Exact adversarial attack to image captioning via structured output learning with latent variables. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, June 2019.

Shaofeng Zhang, Zheng Wang, Xing Xu, Xiang Guan, and Yang Yang. Fooled by imagination: Adversarial attack to image captioning via perturbation in complex domain. In *ICME*, 2020.

Sicheng Zhu, Ruiyi Zhang, Bang An, Gang Wu, Joe Barrow, Zichao Wang, Furong Huang, Ani Nenkova, and Tong Sun. AutoDAN: Interpretable gradient-based adversarial attacks on large language models, 2023. URL https://arxiv.org/abs/2310.15140.

Daniel Ziegler, Seraphina Nix, Lawrence Chan, Tim Bauman, Peter Schmidt-Nielsen, Tao Lin, Adam Scherlis, Noa Nabeshima, Benjamin Weinstein-Raun, Daniel de Haas, Buck Shlegeris, and Nate Thomas. Adversarial training for high-stakes reliability. In *Advances in Neural Information Processing Systems*, October 2022. URL https://openreview.net/forum?id=NtJyGXo0nF.

Andy Zou, Zifan Wang, J. Zico Kolter, and Matt Fredrikson. Universal and transferable adversarial attacks on aligned language models, 2023. URL https://arxiv.org/abs/2307.15043.

A. Models

In this work, we use the Pythia suite (Biderman et al., 2023), a collection of 10 autoregressive language models of different sizes, all pretrained for one epoch on the Pile (Gao et al., 2020). Model checkpoints are provided every thousand steps; for the experiments presented in this work, we always start from the final checkpoint (the `main` revision on HuggingFace Hub) unless otherwise specified.

We reproduce the model sizes of the Pythia suite in Table 1. Note that the number of parameters differs from that given in the model name because we use the models for classification tasks, which replaces the unembedding layer with a (smaller) classification head.

Model Size (# parameters)	Short Name	Pythia Name	Plot Color
7,629,056	7.6M	14m	
17,617,408	17.6M	31m	
44,672,000	44.7M	70m	
123,691,008	123.7M	160m	
353,824,768	353.8M	410m	
908,763,136	908.8M	1b	
1,311,629,312	1.3B	1.4b	NA
2,646,435,840	2.6B	2.8b	NA
6,650,740,736	6.7B	6.9b	NA
11,586,560,000	11.6B	12b	NA

Table 1: Model sizes used in our experiments, the short name often used in plots, Pythia model name, and corresponding plot colors where applicable

B. Datasets

We consider four datasets in this paper. Two of them are pre-existing datasets that we use from HuggingFace Hub: `Spam` (Metsis et al., 2006) and `IMDB` (Maas et al., 2011).[5] Two are synthetic datasets that we generate ourselves: `PasswordMatch` and `WordLength`. For representative datapoints of these datasets, see Table 3.

Since the context window for the Pythia model family is 2048 tokens (Biderman et al., 2023), we must be careful not to run models on datapoints that are longer than this threshold. For fine-tuning, presented in Section 4, we train on the entire dataset, filtering out the (very few) datapoints which exceed 2000 tokens. We cap at 2000 tokens instead of the 2048 token context length to leave room for adversarial attacks, special tokens, and any other additional tokens we might need. Table 2 shows the number of datapoints in each dataset, as well as the number of datapoints that exceed 2000 tokens.

For the `PasswordMatch` task, we allow attacks to replace the 'user-provided' password, instead of treating the prompt as immutable and appending new text only after it.

Dataset	**Train**	**of which > 2000 tokens**	**Validation**	**of which > 2000 tokens**
`Spam`	31,716	496	2,000	33
`IMDB`	25,000	7	25,000	4
`PasswordMatch`	20,000	0	2,000	0
`WordLength`	20,000	0	2,000	0

Table 2: Dataset Statistics

[5] Available at https://huggingface.co/datasets/SetFit/enron_spam and https://huggingface.co/datasets/stanfordnlp/imdb

Dataset	Text	Label
Spam	organisational announcement we are pleased to announce that we have appointed michael brown as chief\noperating officer of enron europe .\nmichael joined enron \x01 , s london legal department in 1995 and has played a key\nrole in the success of many of our large commercial ventures since then . in\njuly 1999 , he was made enron europe \x01 , s general counsel . replacing michael in\nthe general counsel role will be mark evans who joined our legal department\nin 1995 .\nplease join us in congratulating michael and mark in their new roles .	HAM (0)
Spam	claim your free home depot gift card - a $ 1000 value . claim your home depot gift card - a $ 1000 value . were sure you can find a use for this gift card in your area . ( ) .\nby exclusiverewards\nqprkelmv	SPAM (1)
IMDB	Up until the last 20 minutes, I was thinking that this is possibly Jackie Chan\'s worst movie (excluding his pre-1978 work, which I am not familiar with). The final fight sequence changed all that: it is long and good and intense - indeed, one of the highlights of Chan\'s career. But to get to it, you have to sit through a lot of "comedy" that might amuse five-year-olds (oh, look! someone threw a tomato at that guy\'s face) and endless "football" scenes. Not to mention the dubbing (which includes the line "How can I turn it off? It\'s not a tap" - watch to find out what it refers to). "Dragon Lord" is worth renting for the final fight alone, but the rest of the movie is only for Jackie collectors, and even then only for those who\'ve already seen at least 15 of his other movies. (**)	NEGATIVE (0)
IMDB	Hurrah! A space film that doesn\'t take itself too seriously and everyone can come along for the exciting ride that is space camp. The film starts slowly, the usual mix of idiots and high-fliers mixed together into a dodgy soup. But when the going gets tough - the tough get themselves sorted out and itś not an un-believable change in the characters as you can believe that there\'s some responsibility in their young minds.¡br /¿¡br /¿The only flaw in the film is that Kate Capshaw is EXTREMELY annoying as the "I\'m right and you\'re all wrong" instructor. I would recommend this as a nice night in movie and a 7 Vote.	POSITIVE (1)
PasswordMatch	You are an online security system. Below the user will enter a password.\nIf it is "syllogistic", then respond with "Access Granted". Otherwise respond with "Access Denied".\n—\nYorktown\n—\n	ACCESS DENIED (0)
PasswordMatch	You are an online security system. Below the user will enter a password.\nIf it is "Herodotus", then respond with "Access Granted". Otherwise respond with "Access Denied".\n—\nHerodotus\n—\n	ACCESS GRANTED (1)
WordLength	Which word is longer? Word 1: keyword. Word 2: medal. Return 0 for the first word and 1 for the second word. If the words are the same length, return 0. Ignore everything after this sentence.	FIRST (0)
WordLength	Which word is longer? Word 1: purloin. Word 2: literate. Return 0 for the first word and 1 for the second word. If the words are the same length, return 0. Ignore everything after this sentence.	SECOND (1)

Table 3: Representative examples from four datasets used in our experiments.

C. Adversarial Attacks

The primary attack we use is GCG from Zou et al. (2023). We use the simple, single-prompt version described in Algorithm 1 of Zou et al. (2023) with the modifiable subset $\mathcal{I}$ set to be the final N tokens of the prompt (except for PasswordMatch, where there is a final --- separator after the attack tokens; see Table 3). We use a suffix of length $N = 10$, batch size $B = 128$, and $k = 256$ top substitutions for all experiments. We use $T = 10$ iterations for most experiments, using $T = 30$ to evaluate robustness transfer from adversarially training on a weaker attack ($T = 10$).

We describe the baseline RandomToken algorithm in Algorithm 1. RandomToken is designed to be similar to GCG except that RandomToken does not use gradient-guided search. Instead, for each iteration we replace each token in the adversarial suffix with a new token chosen uniformly at random from the vocabulary of the model. We then evaluate the new prompt to see if it has caused the model to give an incorrect answer and stop the attack if it has. If no iteration was successful, we return the adversarial suffix from the final iteration.

To make sure the baseline is a fair comparison, we constrain the attacks to use the same maximum number of forward passes. To do this, we compute the number of forward passes used by GCG as $B \times T = 1280$ and thus perform up to 1280 iterations of RandomToken.

Algorithm 1 RandomToken

Input: Initial prompt $x_{1:n}$, modifiable subset $\mathcal{I}$, iterations T, success criterion S, vocabulary V
for $t = 1$ **to** T **do**
  for $i \in \mathcal{I}$ **do**
    $x_i \leftarrow \text{Uniform}(V)$
  end for
  if $S(x_{1:n})$ **then**
    return: $x_{1:n}$
  end if
end for
return: $x_{1:n}$
Output: Optimized prompt $x_{1:n}$

D. Fine-tuning

D.1. Training

For each task, we fine-tune each model for a single epoch. The final validation accuracies are shown in Table 4.

Task	Model Size (# parameters)	Validation accuracy
`Spam`	7.6M	0.985
	17.6M	0.985
	44.7M	0.99
	123.7M	0.99
	353.8M	0.985
	908.8M	0.99
	1.3B	0.99
	2.6B	0.9
	6.7B	0.99
	11.6B	0.99
`IMDB`	7.6M	0.875
	17.6M	0.9
	44.7M	0.905
	123.7M	0.93
	353.8M	0.96
	908.8M	0.965
	1.3B	0.96
	2.6B	0.975
	6.7B	0.97
	11.6B	0.98
`PasswordMatch`	7.6M	1
	17.6M	1
	44.7M	1
	123.7M	1
	353.8M	1
	908.8M	1
	1.3B	1
	2.6B	1
	6.7B	1
	11.6B	1
`WordLength`	7.6M	0.836
	17.6M	0.882
	44.7M	0.858
	123.7M	0.944
	353.8M	0.978
	908.8M	0.958
	1.3B	0.968
	2.6B	0.972
	6.7B	0.954
	11.6B	0.976

Table 4: Accuracy on (not attacked) validation dataset at the end of training.

D.2. Attack Results

We attack the fine-tuned models with both the `GCG` and `RandomToken` attacks. As explored in Section 4, while model size appears to generally help with robustness, there is a large amount of unexplained variability in each model's robustness.

D.2.1. `GCG`

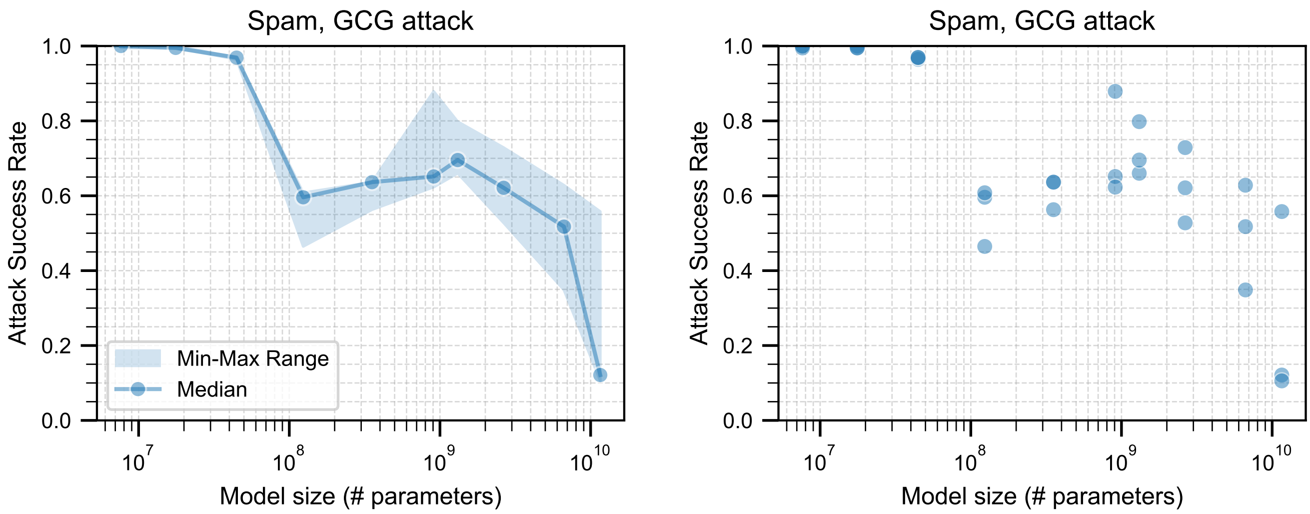

Figure 4: `GCG` attack success rate on different sizes of fine-tuned models on the `Spam` task. We show three seeds per model size. The min-max-median plot (left) and scatterplot (right) are constructed using the same data.

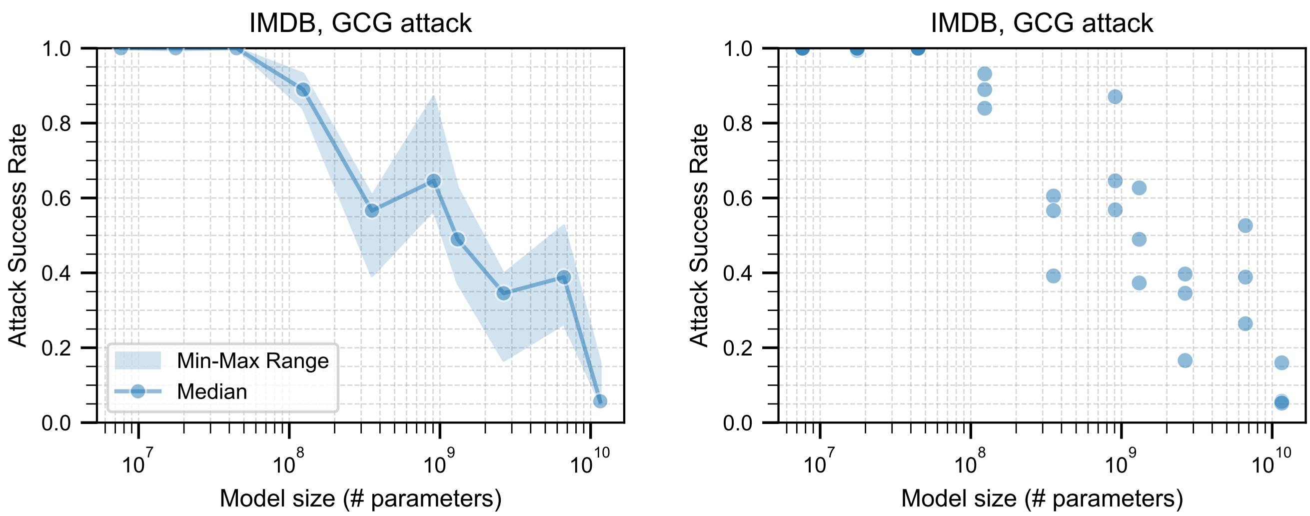

Figure 5: `GCG` attack success rate on different sizes of fine-tuned models on the `IMDB` task. We show three seeds per model size. The min-max-median plot (left) and scatterplot (right) are constructed using the same data.

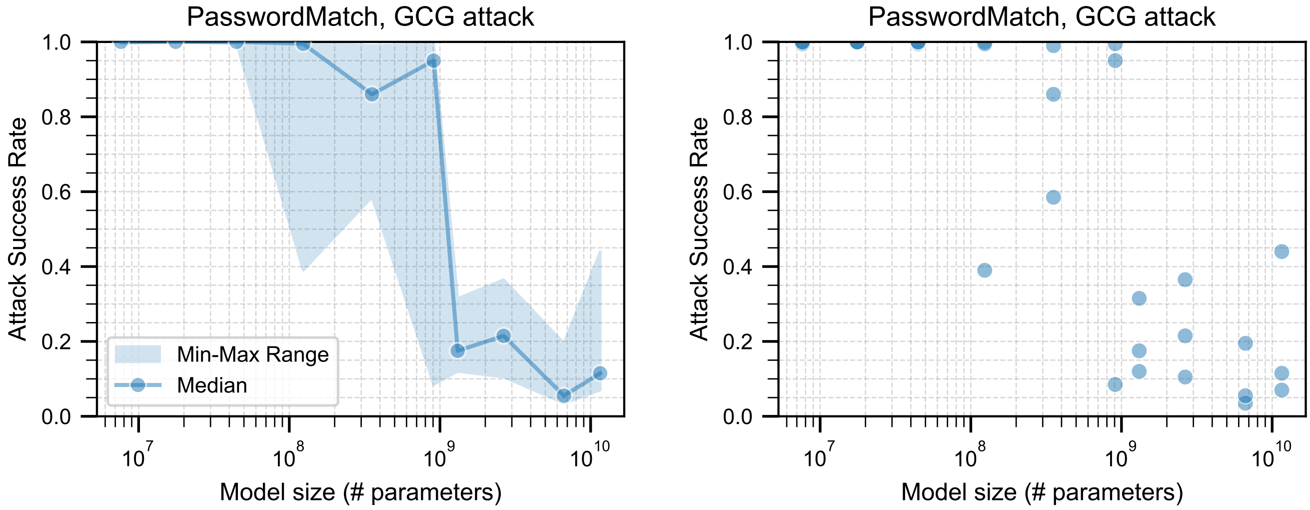

Figure 6: GCG attack success rate on different sizes of fine-tuned models on the PasswordMatch task. We show three seeds per model size. The min-max-median plot (left) and scatterplot (right) are constructed using the same data.

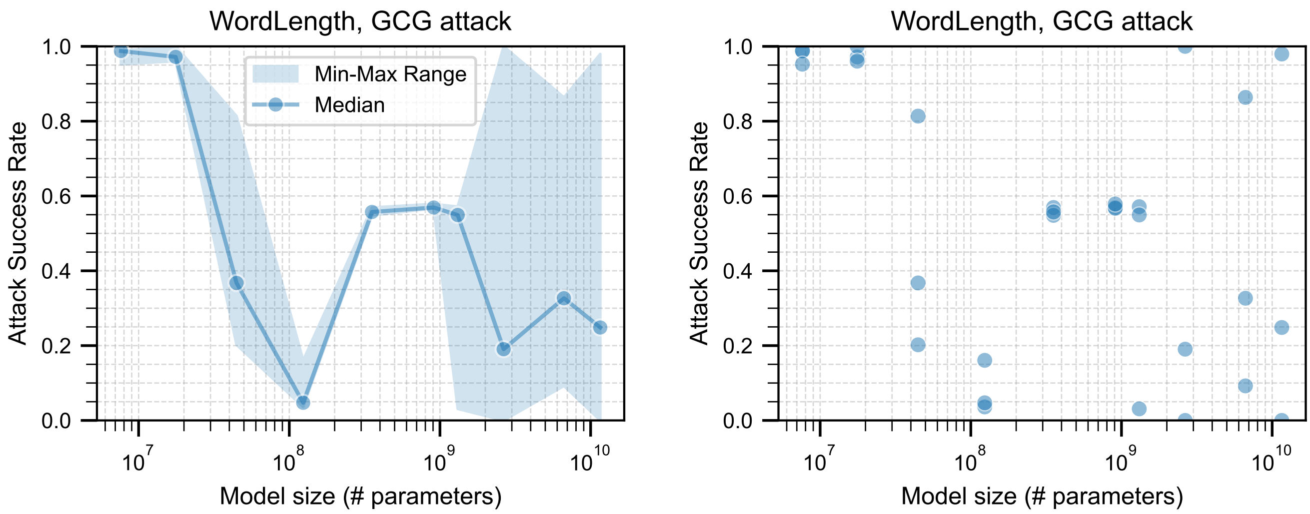

Figure 7: GCG attack success rate on different sizes of fine-tuned models on the WordLength task. We show three seeds per model size. The min-max-median plot (left) and scatterplot (right) are constructed using the same data.

D.2.2. `RandomToken`

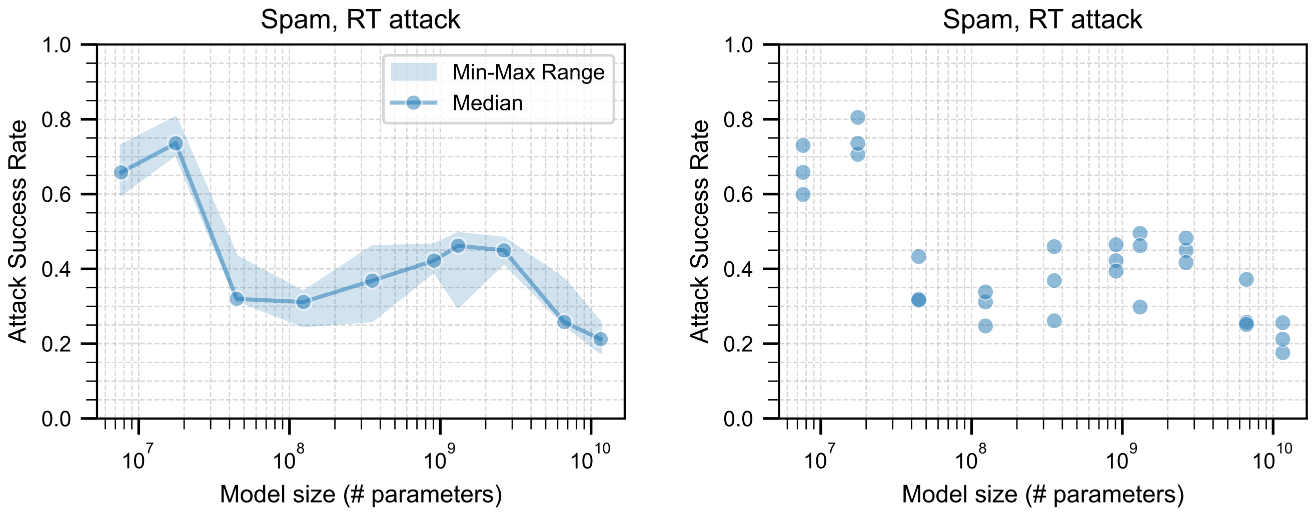

Figure 8: `RandomToken` (RT) attack success rate on different sizes of fine-tuned models on the `Spam` task. We show three seeds per model size. The min-max-median plot (left) and scatterplot (right) are constructed using the same data.

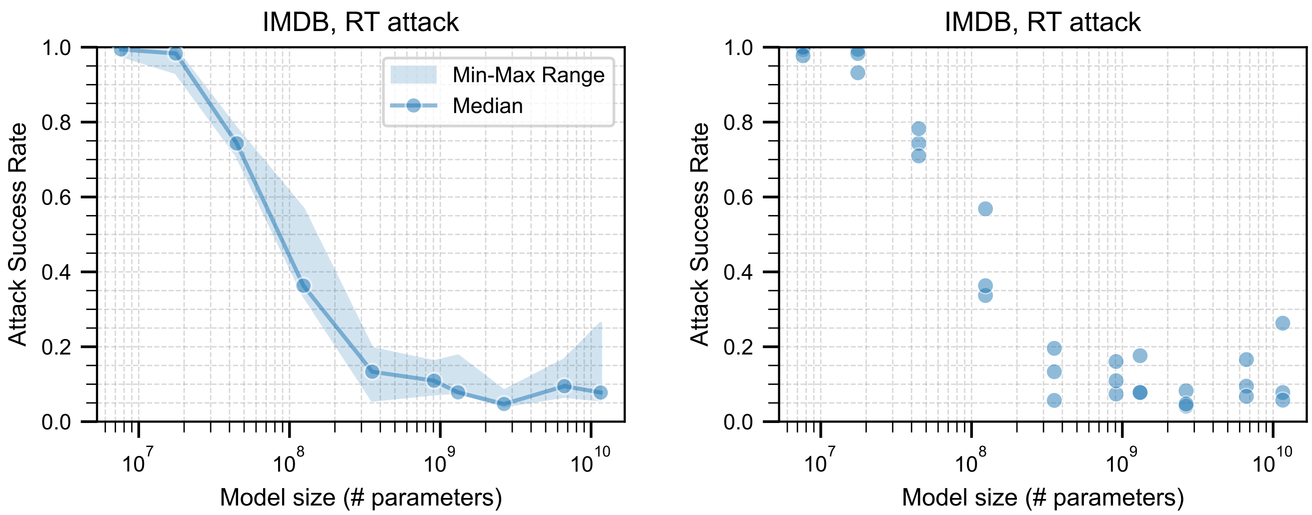

Figure 9: `RandomToken` (RT) attack success rate on different sizes of fine-tuned models on the `IMDB` task. We show three seeds per model size. The min-max-median plot (left) and scatterplot (right) are constructed using the same data.

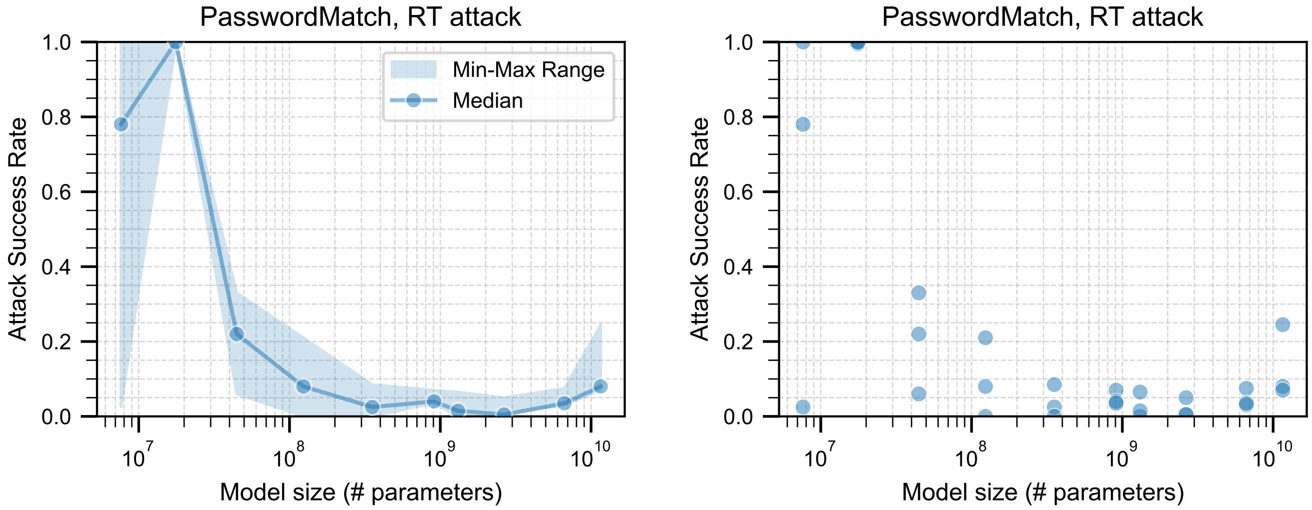

Figure 10: `RandomToken` (RT) attack success rate on different sizes of fine-tuned models on the `PasswordMatch` task. We show three seeds per model size. The min-max-median plot (left) and scatterplot (right) are constructed using the same data.

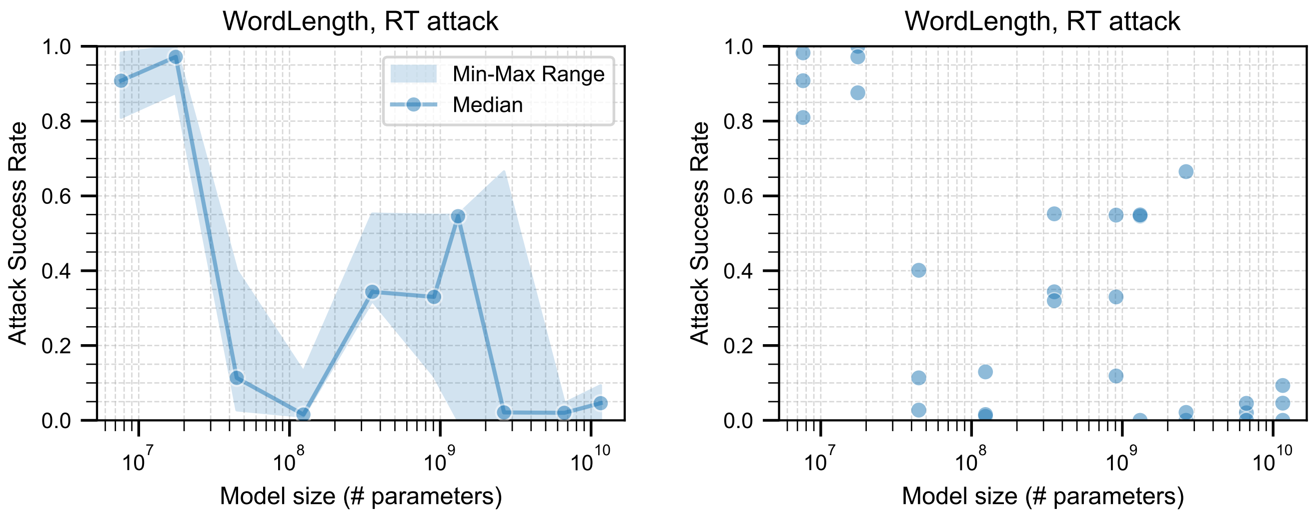

Figure 11: `RandomToken` (RT) attack success rate on different sizes of fine-tuned models on the `WordLength` task. We show three seeds per model size. The min-max-median plot (left) and scatterplot (right) are constructed using the same data.

E. Adversarial Training and Transfer

The overall adversarial training procedure is presented in Figure 12.

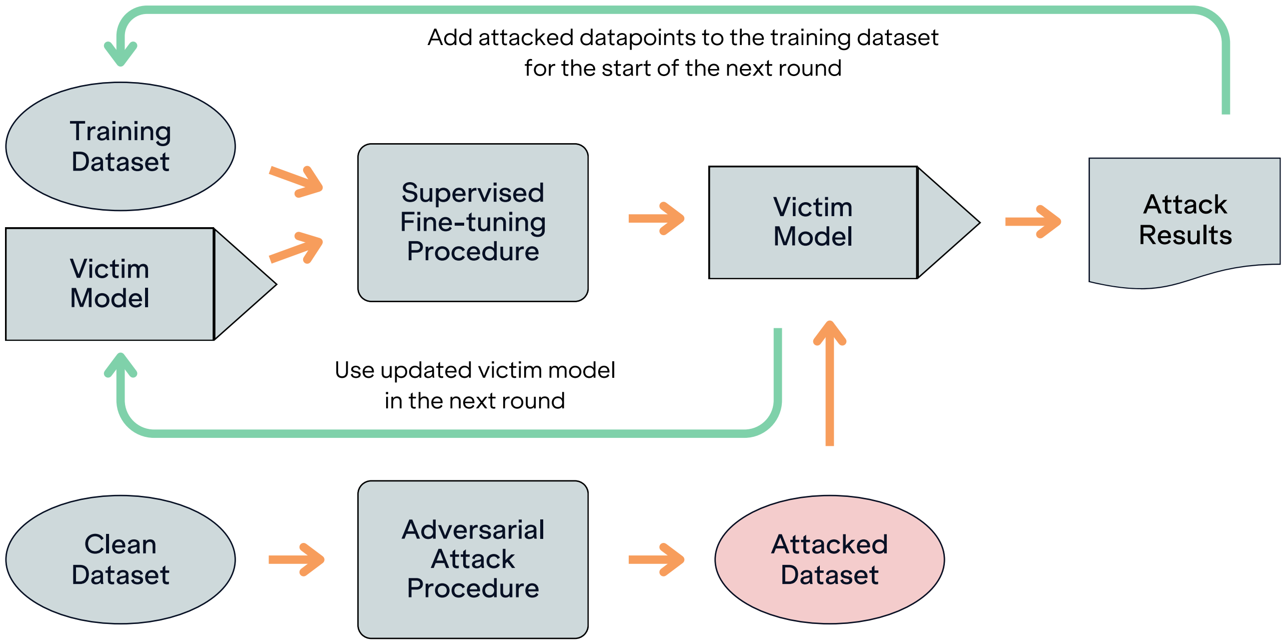

Figure 12: Our adversarial training setup.

As the diagram highlights, adversarial training is done by repeating the following steps:

- Train the model for one epoch on the train dataset.
- Attack the train dataset and evaluate the model on the attacked train dataset.
- Add the attacked examples to the train dataset.
- Attack the validation dataset and evaluate the model on the attacked validation dataset. Record model performance on the attacked validation dataset.

For adversarial training, we use an initial training dataset of size 2000, and a validation dataset of size 200. Initially we used a validation dataset also of size 2000, but found that decreasing the validation dataset size had a negligible effect on the variance of the attack success rate, so opted for smaller dataset to enable faster evaluation. At each round, we add 200 adversarially-attacked examples to the train dataset.

E.1. Adversarial Training

Below, we show plots of adversarial training using the `GCG` and `RandomToken` attacks across the four tasks. We use three seeds per model, and present attack success rate after 10 and 30 rounds of adversarial training.

E.1.1. GCG Attack 10 Rounds

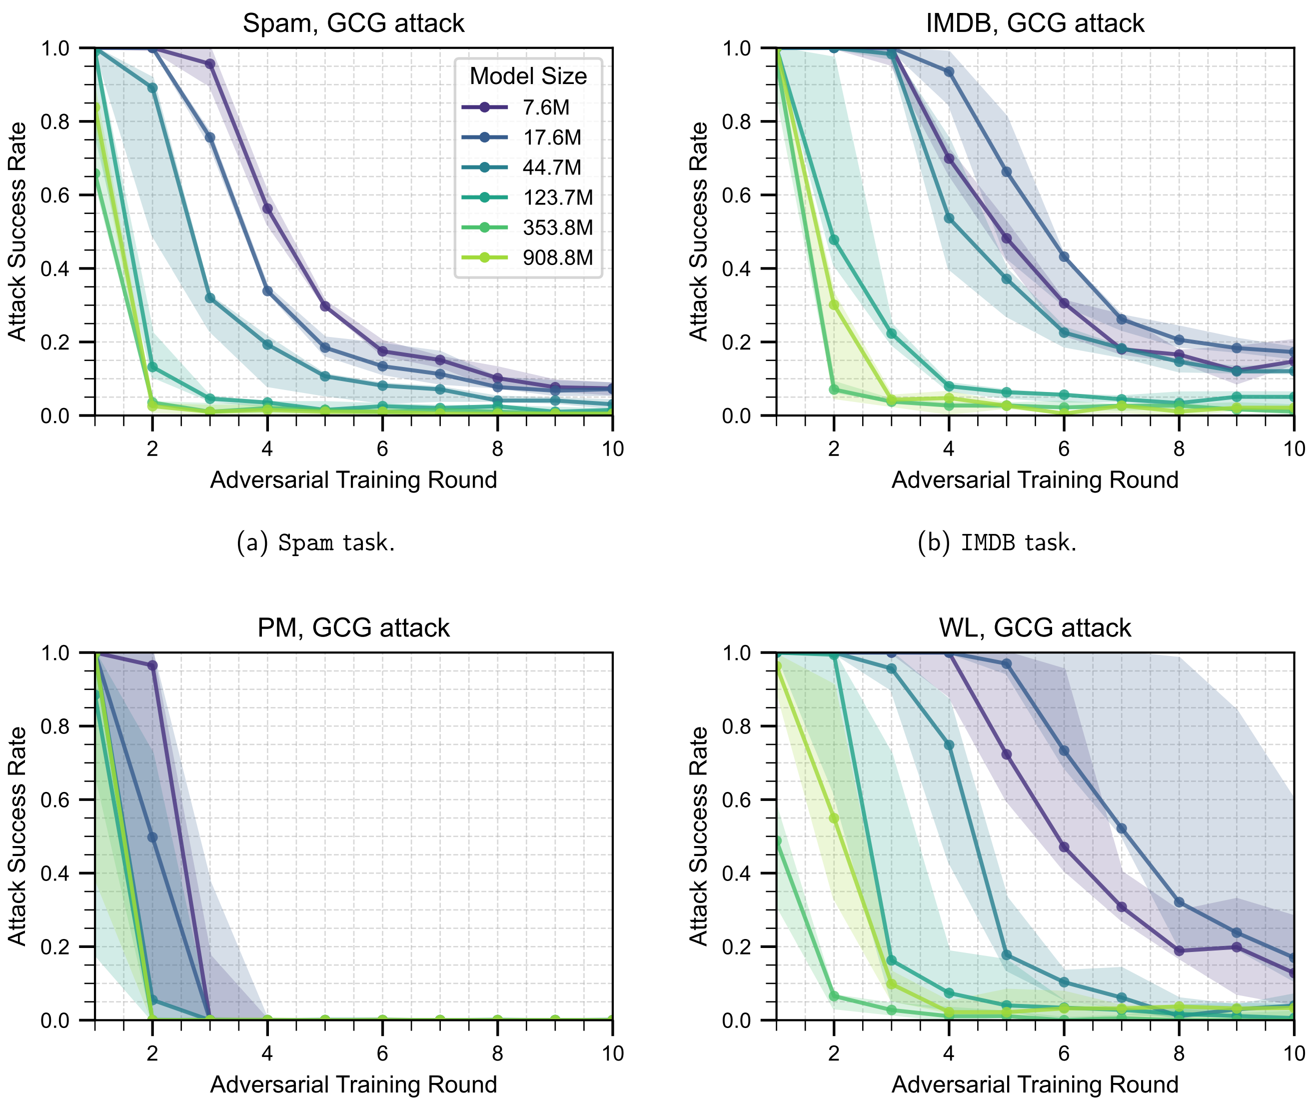

(a) `Spam` task.

(b) `IMDB` task.

(c) `PasswordMatch` task.

(d) `WordLength` task.

Figure 13: Attack success rate as a function of adversarial training round across four tasks using the 10-iteration `GCG` attack, for different model sizes, shown for 10 rounds of adversarial training. We shade min to max and plot median over three seeds.

E.1.2. GCG Attack 10 Rounds Alternate View

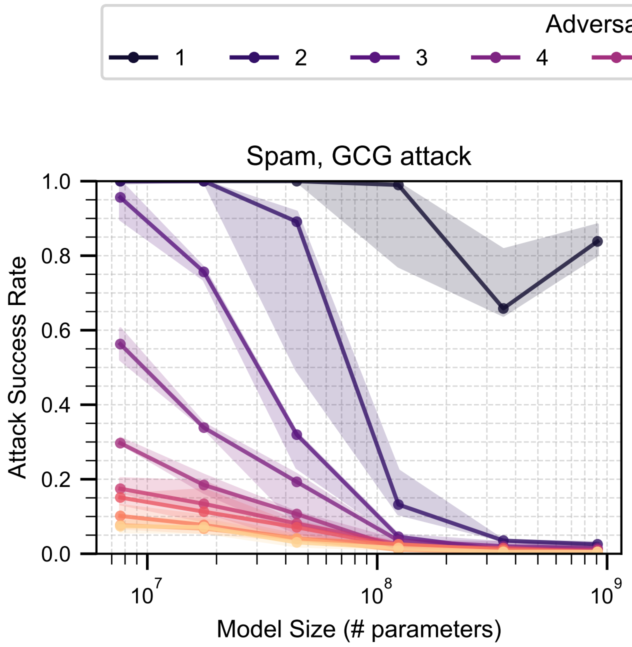

(a) Spam task.

(b) IMDB task.

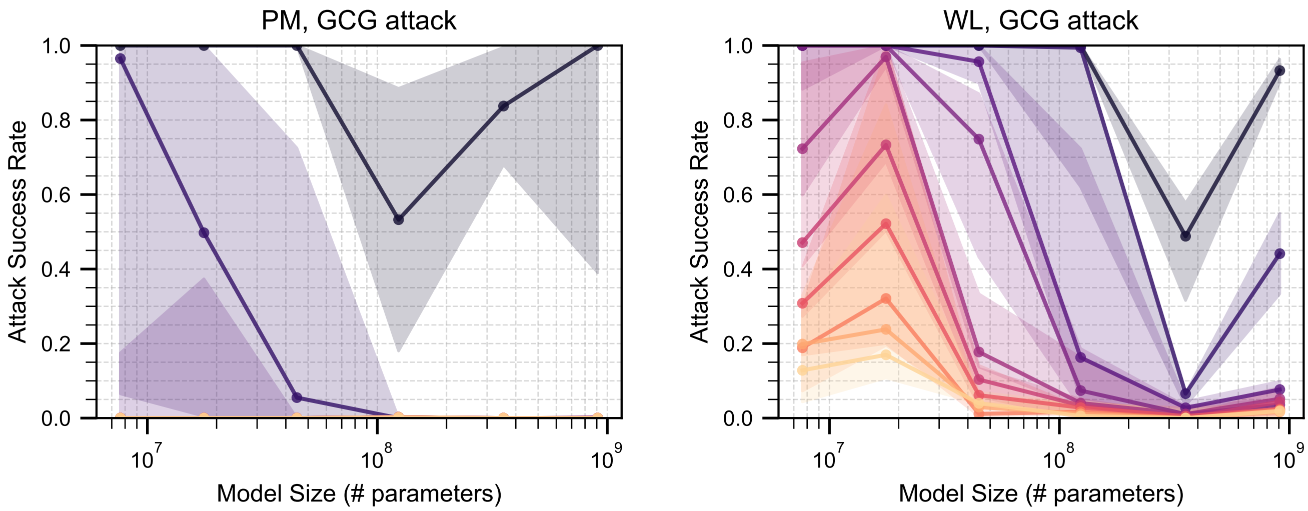

(c) PasswordMatch task.

(d) WordLength task.

Figure 14: Attack success rate as a function of model size across four tasks using the 10-iteration GCG attack, over different adversarial training rounds.

E.1.3. GCG Attack 30 Rounds

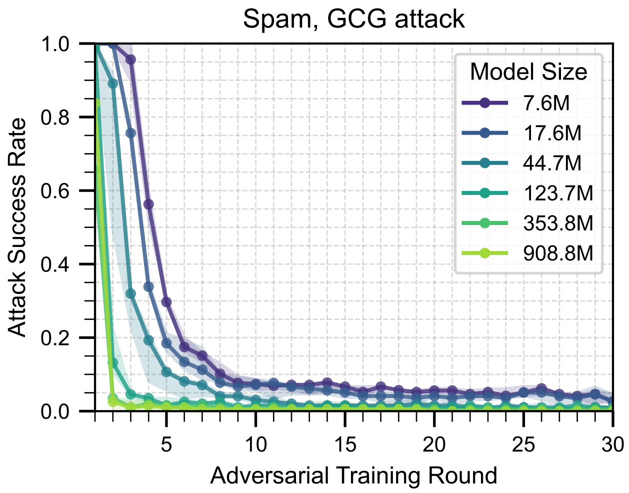

(a) `Spam` task.

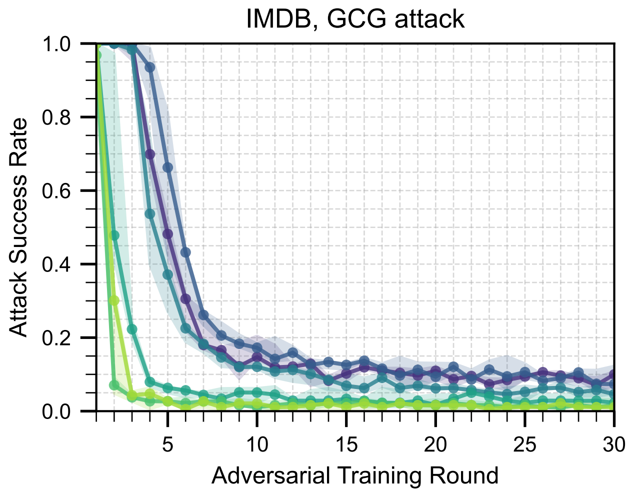

(b) `IMDB` task.

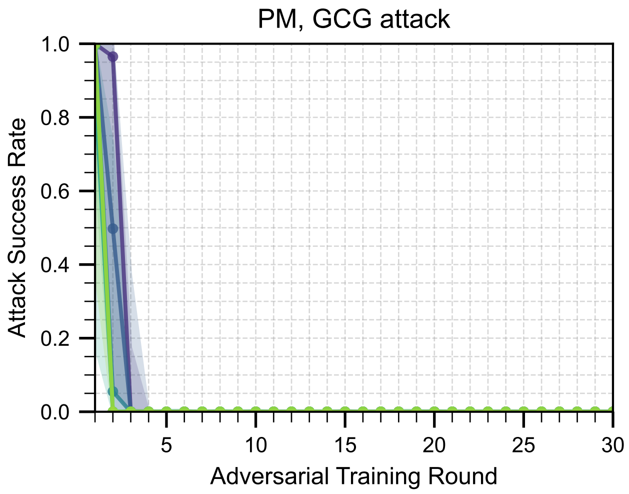

(c) `PasswordMatch` task.

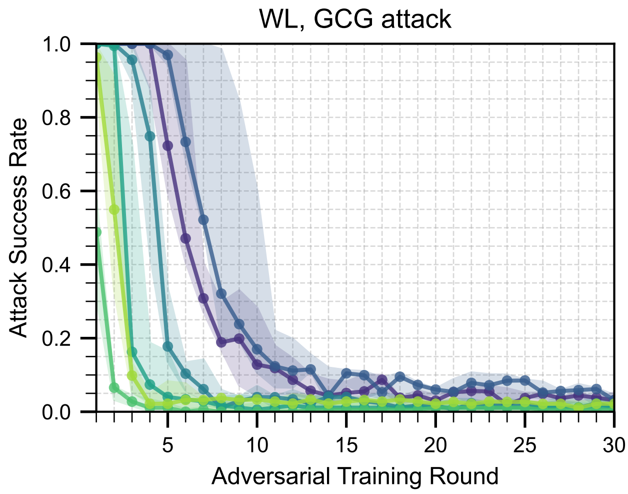

(d) `WordLength` task.

Figure 15: Attack success rate as a function of adversarial training round across four tasks using the 10-iteration `GCG` attack, for different model sizes, shown for 30 rounds of adversarial training.

E.1.4. GCG Attack 30 Rounds Convergence

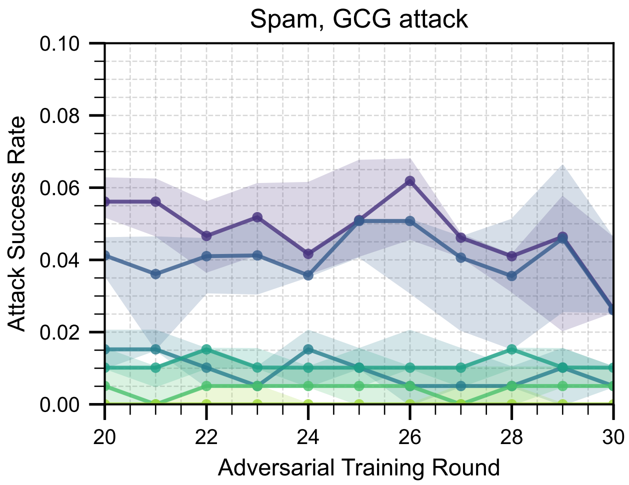

(a) `Spam` task.

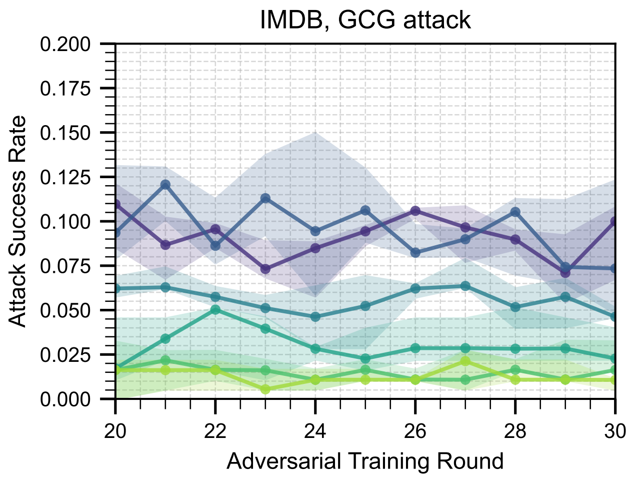

(b) `IMDB` task.

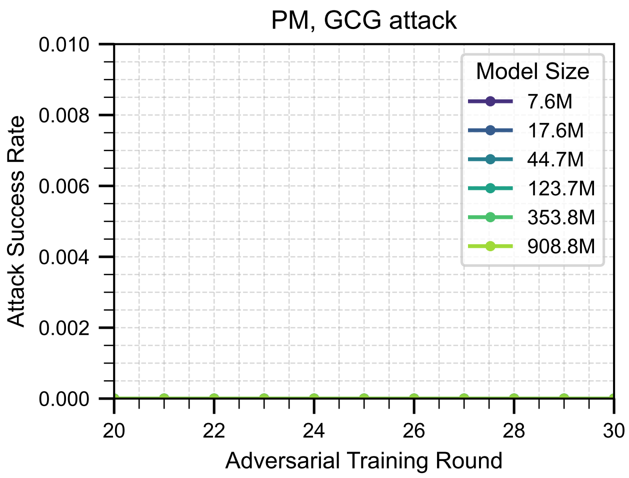

(c) `PasswordMatch` task.

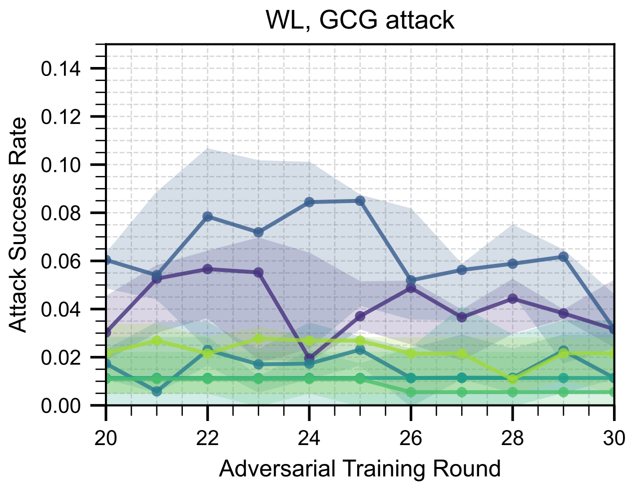

(d) `WordLength` task.

Figure 16: Attack success rate as a function of adversarial training round across four tasks using the 10-iteration GCG attack, for different model sizes, shown for the final 10 rounds of 30-round adversarial training.

E.1.5. RandomToken Attack 10 Rounds

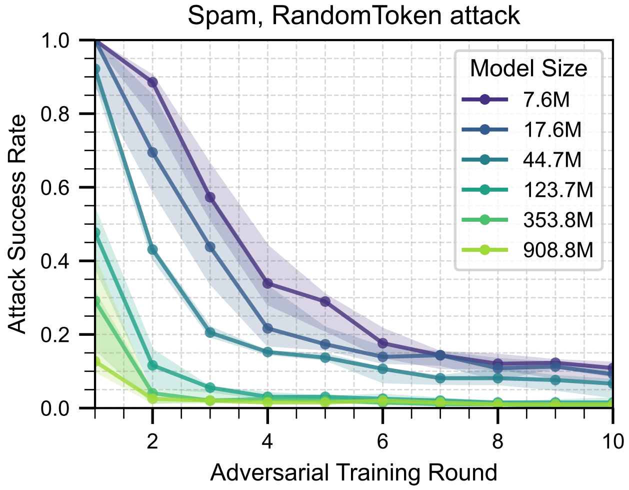

(a) `Spam` task.

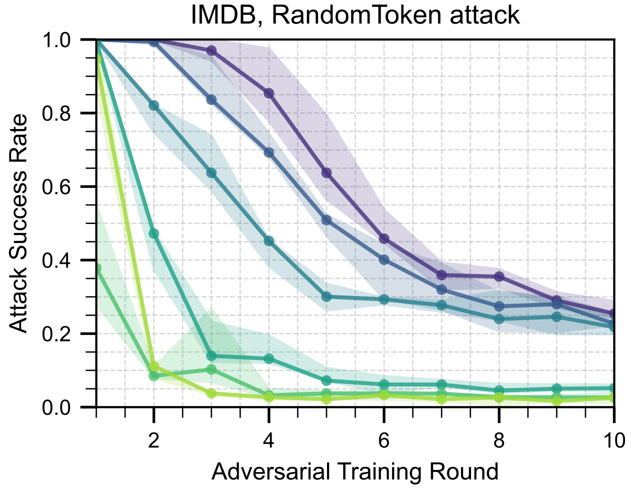

(b) `IMDB` task.

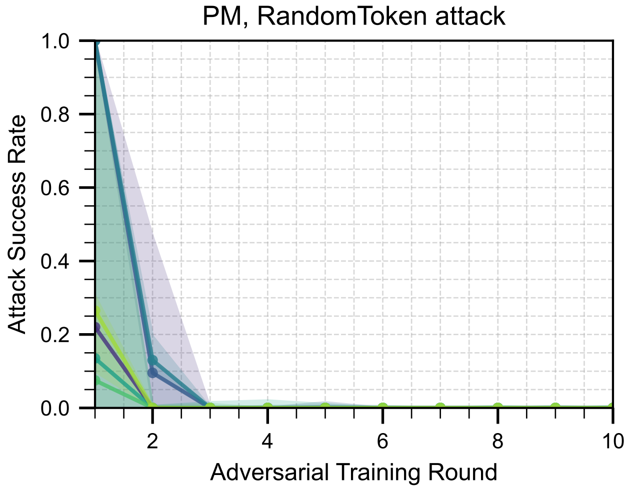

(c) `PasswordMatch` task.

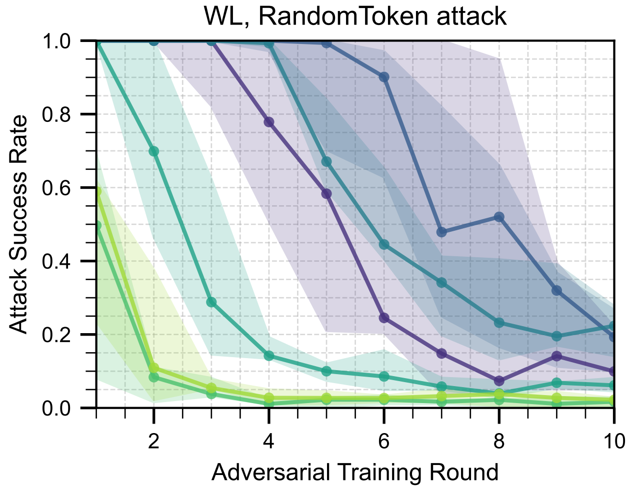

(d) `WordLength` task.

Figure 17: Attack success rate as a function of adversarial training round across four tasks using the `RandomToken` attack, for different model sizes, shown for 10 rounds of adversarial training. We shade min to max and plot median over three seeds.

E.1.6. RandomToken Attack 10 Rounds Alternate View

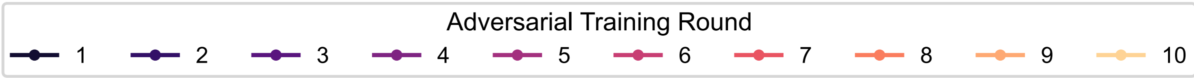

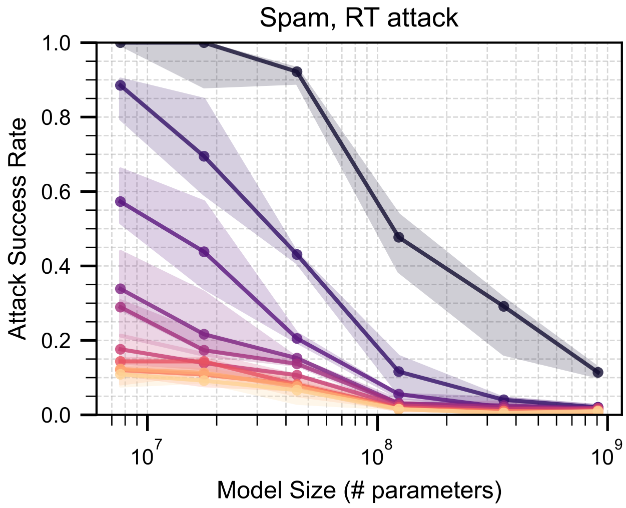

(a) `Spam` task.

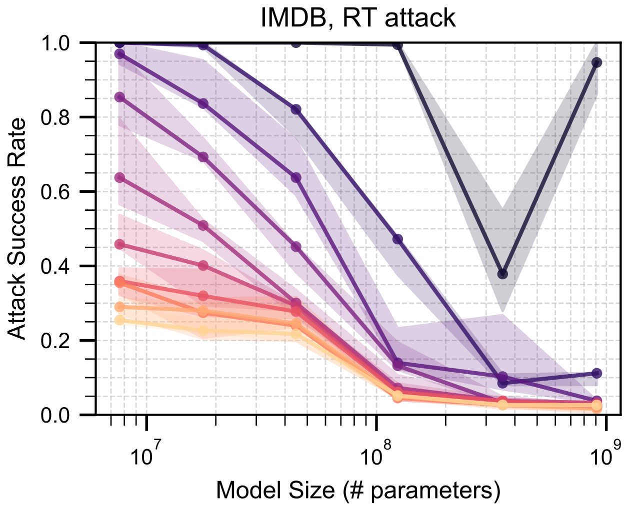

(b) `IMDB` task.

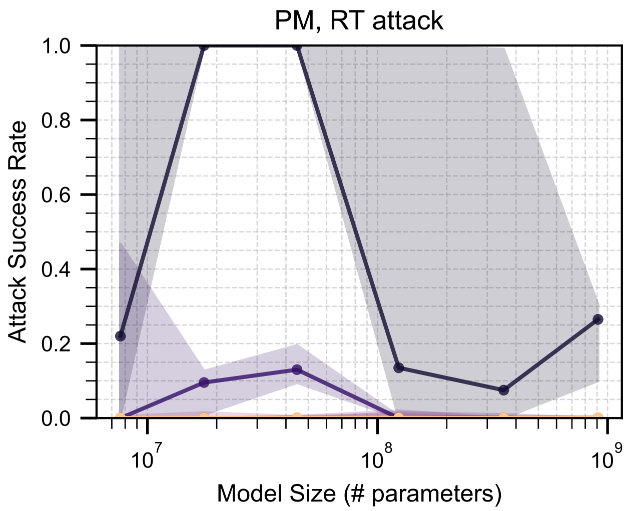

(c) `PasswordMatch` task.

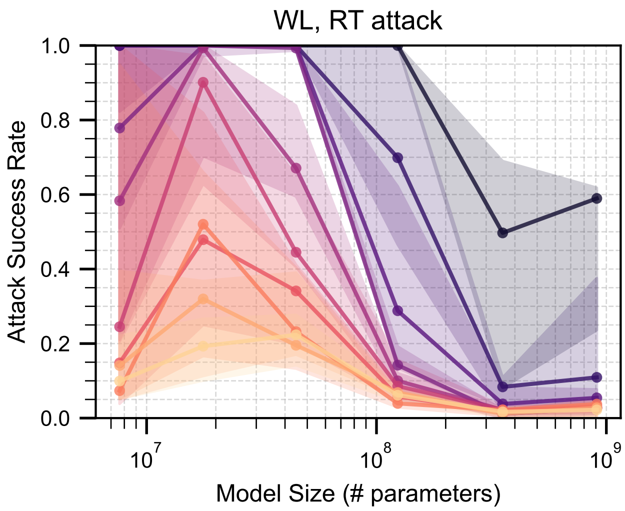

(d) `WordLength` task.

Figure 18: Attack success rate as a function of model size across four tasks using the 10-iteration `RandomToken` (RT) attack, over different adversarial training rounds.

E.1.7. RandomToken Attack 30 Rounds

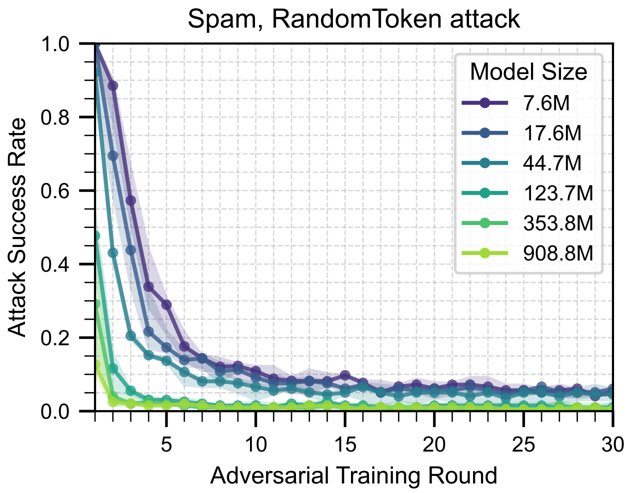

(a) `Spam` task.

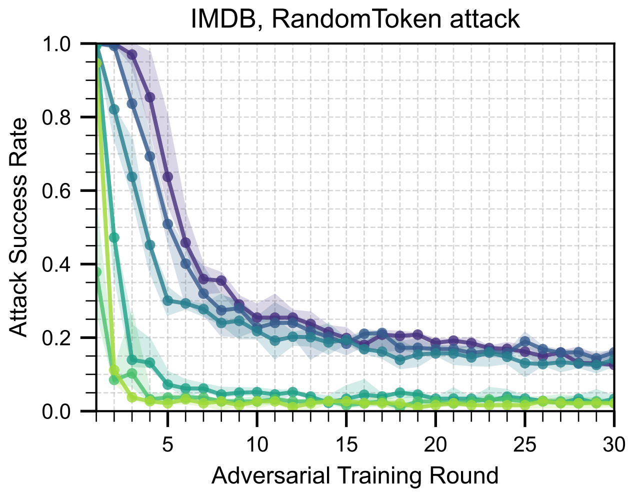

(b) `IMDB` task.

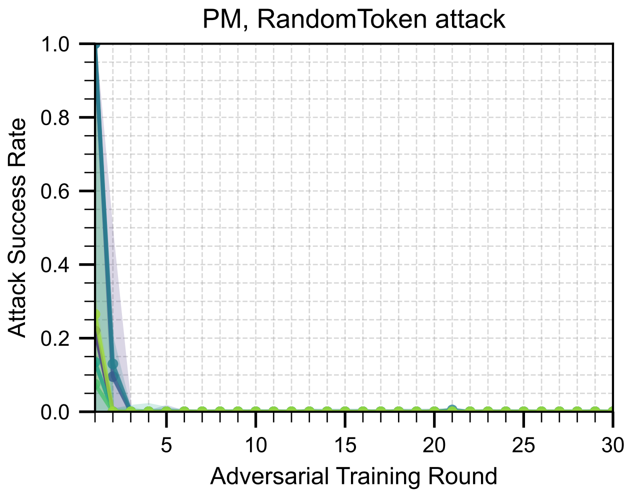

(c) `PasswordMatch` task.

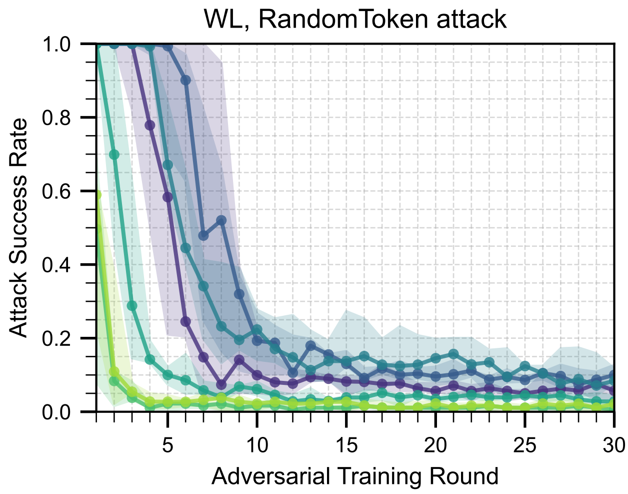

(d) `WordLength` task.

Figure 19: Attack success rate as a function of adversarial training round across four tasks using the `RandomToken` attack, for different model sizes, shown for 30 rounds of adversarial training.

E.1.8. RandomToken Attack 30 Rounds Convergence

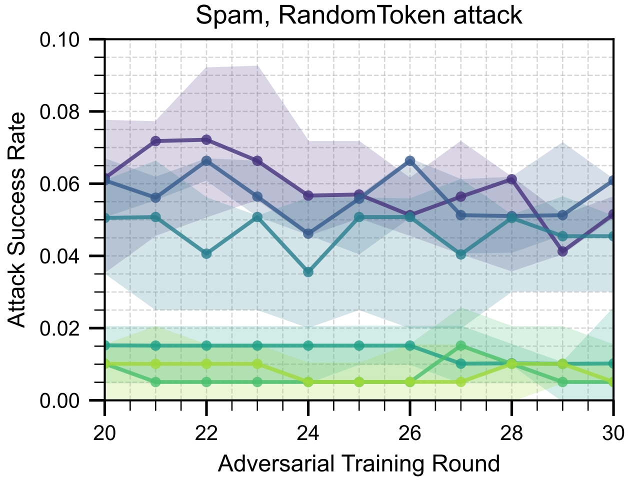

(a) `Spam` task.

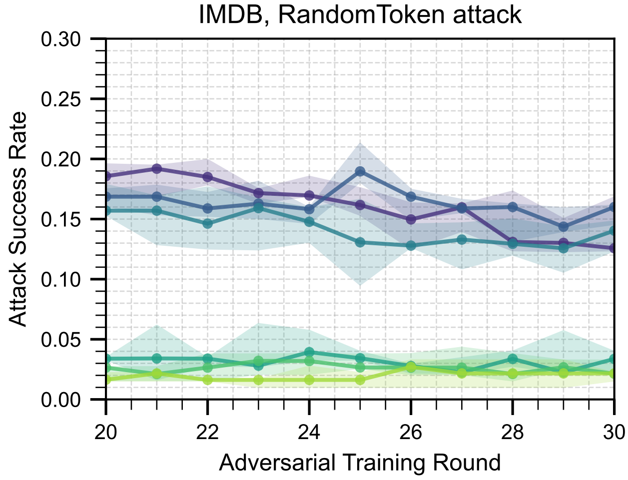

(b) `IMDB` task.

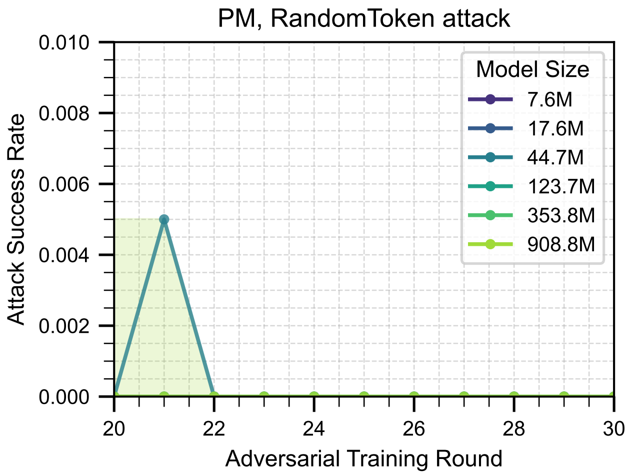

(c) `PasswordMatch` task.

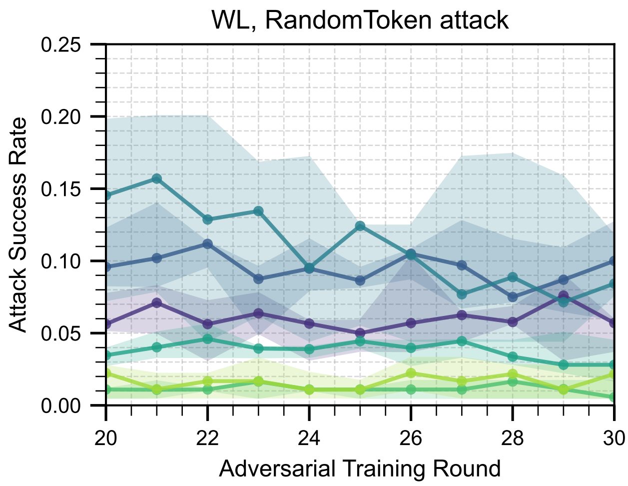

(d) `WordLength` task.

Figure 20: Attack success rate as a function of adversarial training round across four tasks using the `RandomToken` attack, for different model sizes, shown for the final 10 rounds of 30-round adversarial training.

E.2. Transfer

As presented in Section 5.1, we also evaluate how models adversarially trained with one attack generalize to defending against other attacks. We present two collections of plots: first, models trained on the 10-iteration `GCG` attack and evaluated with the 30-iteration `GCG` attack; second, models trained on the `RandomToken` attack and evaluated on the (10-iteration) `GCG` attack. In the first case, all model sizes are able to generalize to being somewhat robust against the stronger attack, though larger models do so both faster and to a greater extent. By contrast, in the second case, only the larger models are able to generalize within the 10 adversarial training rounds studied.

E.2.1. GCG Attack

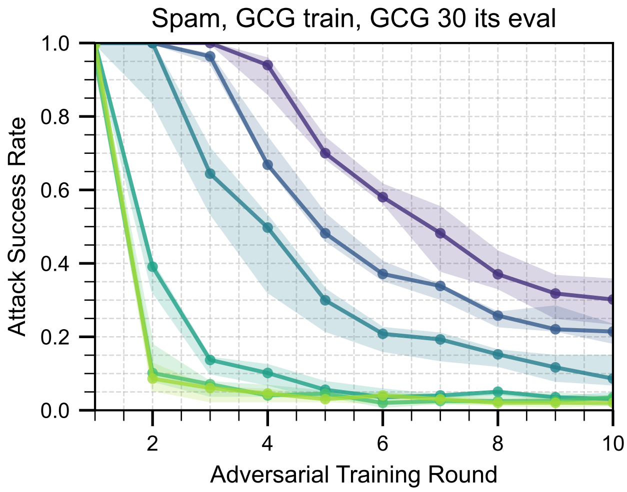

(a) `Spam` task.

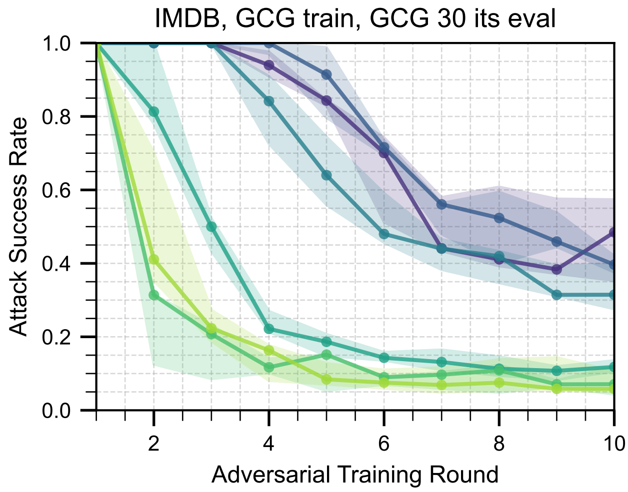

(b) `IMDB` task.

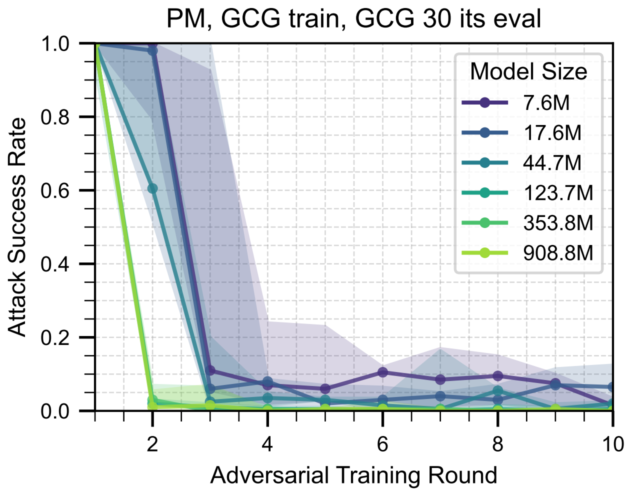

(c) `PasswordMatch` task.

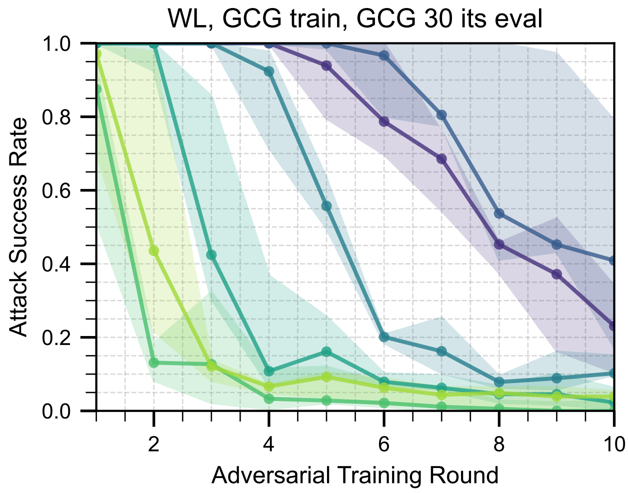

(d) `WordLength` task.

Figure 21: Attack success rate as a function of adversarial training round across four tasks. Adversarial training is performed with the 10-iteration `GCG` attack, and evaluation performed with the 30-iteration `GCG` attack.

E.2.2. RandomToken attack

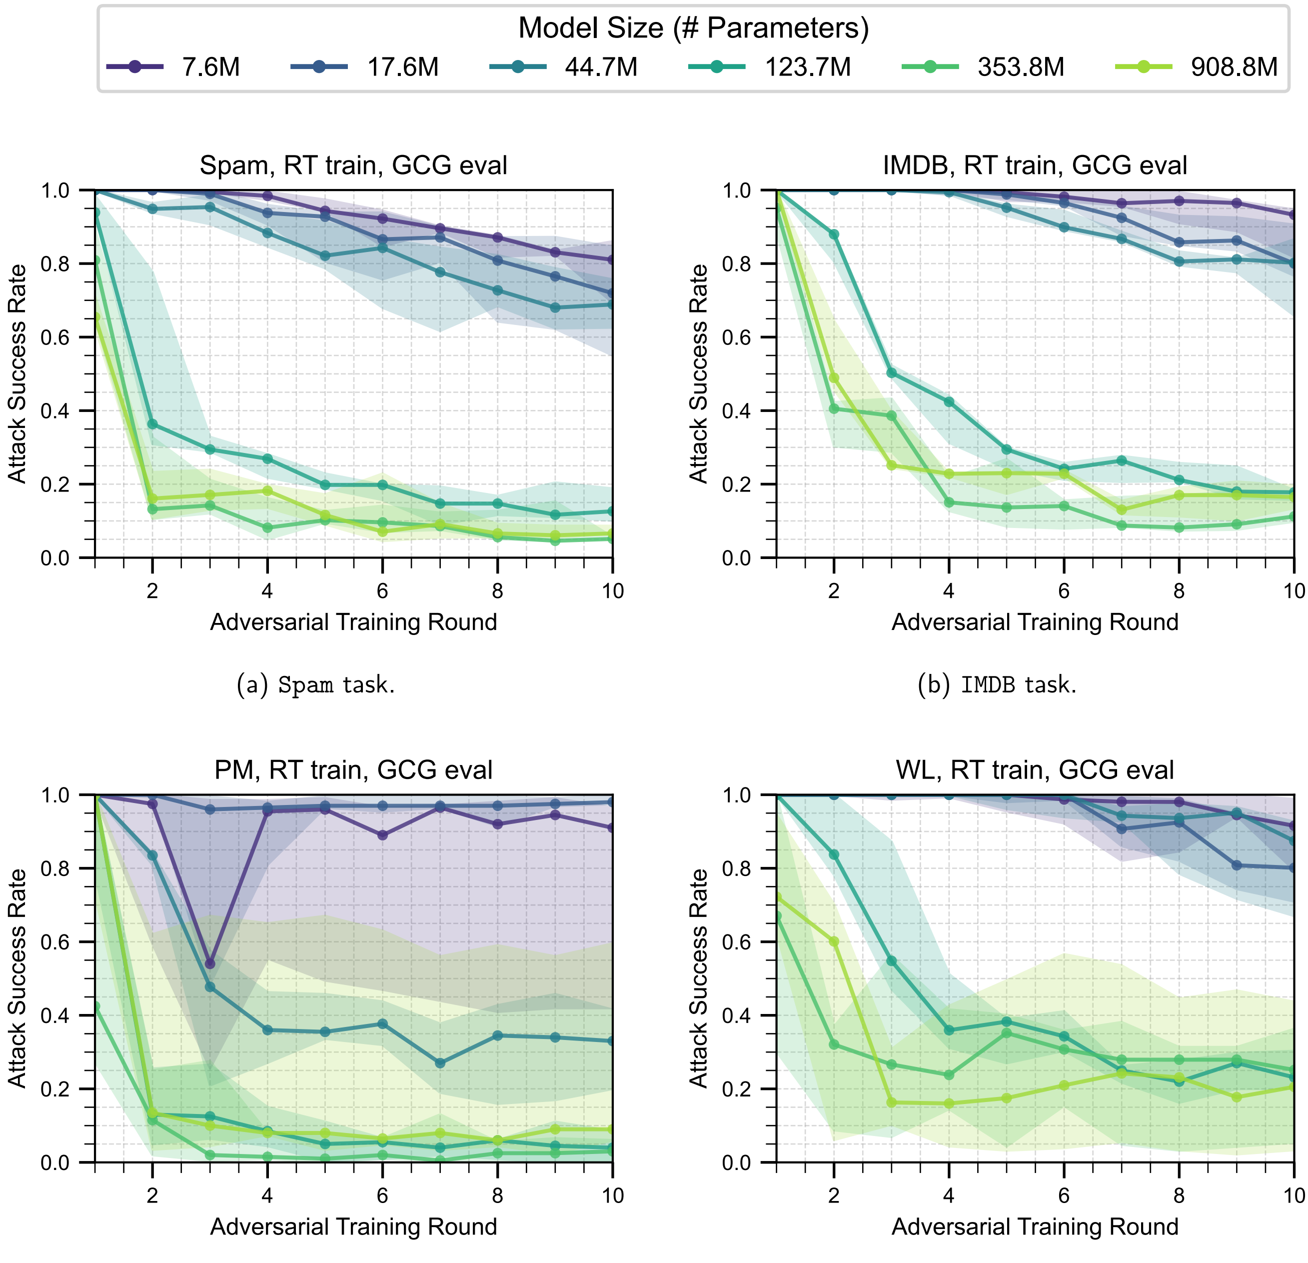

(a) `Spam` task.

(b) `IMDB` task.

(c) `PasswordMatch` task.

(d) `WordLength` task.

Figure 22: Attack success rate as a function of adversarial training round across four tasks. Adversarial training is performed with the `RandomToken` (RT) attack, and evaluation performed with the 10-iteration `GCG` attack.

E.3. Complexity Calculation

In Section 6, we compare the relative complexity of adversarially training a larger model for fewer rounds or a smaller model for more rounds. In this section, we provide a worked example. We use a batch size of 8 for both the 17.6M and 44.7M models. We start with 2000 datapoints in the train dataset and add 200 datapoints each round. This means that after 4 rounds of training, each model will have seen $\sum_{i=1}^{4} (250 + i \cdot 25) = 1250$ batches, and after 8 rounds of training, $\sum_{i=1}^{8} (250 + i \cdot 25) = 2900$ batches. If we update model parameters once per batch, this means that after 4 rounds, the 44.7M parameter model will have had $44.7\text{M} \cdot 1250 = 55875\text{M}$ gradient updates, while after 8 rounds, the 17.6M parameter model will have had $17.6\text{M} \cdot 2900 = 51040\text{M}$ gradient updates.