Exploring Adversarial Robustness in Classification tasks using DNA Language Models

Hyunwoo Yoo [*1] **Haebin Shin** [*2] **Kaidi Xu** [*3] **Gail L. Rosen** [*1]

Abstract

DNA Language Models, such as GROVER, DNABERT2 and the Nucleotide Transformer, operate on DNA sequences that inherently contain sequencing errors, mutations, and laboratory-induced noise, which may significantly impact model performance. Despite the importance of this issue, the robustness of DNA language models remains largely underexplored. In this paper, we comprehensively investigate their robustness in DNA classification by applying various adversarial attack strategies: the character (nucleotide substitutions), word (codon modifications), and sentence levels (back-translation-based transformations) to systematically analyze model vulnerabilities. Our results demonstrate that DNA language models are highly susceptible to adversarial attacks, leading to significant performance degradation. Furthermore, we explore adversarial training method as a defense mechanism, which enhances both robustness and classification accuracy. This study highlights the limitations of DNA language models and underscores the necessity of robustness in bioinformatics.

1. Introduction

Transformer-based language models are increasingly being adopted in bioinformatics, leveraging NLP techniques to tackle sequence classification and functional prediction tasks, which have traditionally relied on alignment-based methods(Steinegger & Söding, 2017; Buchfink et al., 2021). Notably, language models such as DNABERT2 (Zhou et al., 2023), Nucleotide Transformer (Dalla-Torre et al., 2023), and GROVER (Sanabria et al., 2024) leverage large-scale genomic sequence data as textual data and are specialized for specific bioinformatics downstream tasks. By treating DNA sequences not merely as strings but as sequence data with contextual information, these models introduce a novel approach to solving bioinformatics problems.

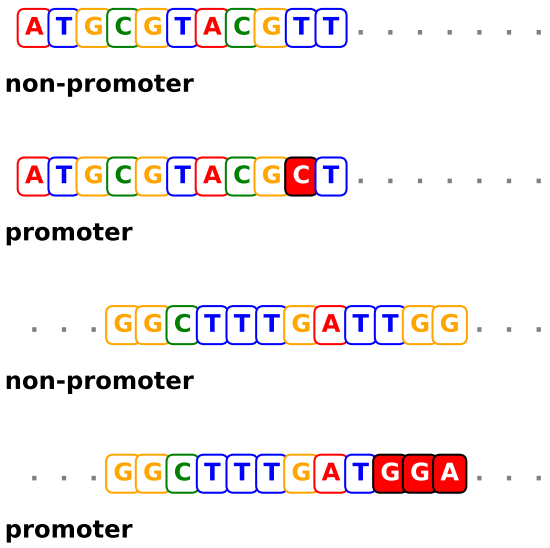

Figure 1. Adversarial examples with nucleotides and codon change, which will be misclassified by a non-promoter to a promoter.

However, despite these advancements, the robustness of DNA language models remains underexplored. In real-world biological environments, DNA sequences are susceptible to sequencing errors, mutations, and data noise introduced during the extraction process in laboratories (Ono et al., 2021; Ma et al., 2019). However, there is a lack of systematic research analyzing the impact of such variations on model performance. While robustness studies on language models have been active in text classification, research on DNA sequence classification models remains limited. Given the increasing adoption of DNA classification models in clinical and biotechnological applications, it is crucial to assess

[*]Equal contribution [1]Department of Electrical and Computer Engineering, Drexel University, Philadelphia, United States [2]KAIST AI, Seoul, Korea [3]Department of Computer Science, Drexel University, Philadelphia, United States. Correspondence to: Hyunwoo Yoo <hty23@drexel.edu>, Haebin Shin <haebin.shin@kaist.ac.kr>, Kaidi Xu <kx46@drexel.edu>, Gail L. Rosen <glr26@drexel.edu>.

Proceedings of the Workshop on Generative AI for Biology at the 42nd International Conference on Machine Learning, Vancouver, Canada. PMLR 267, 2025.

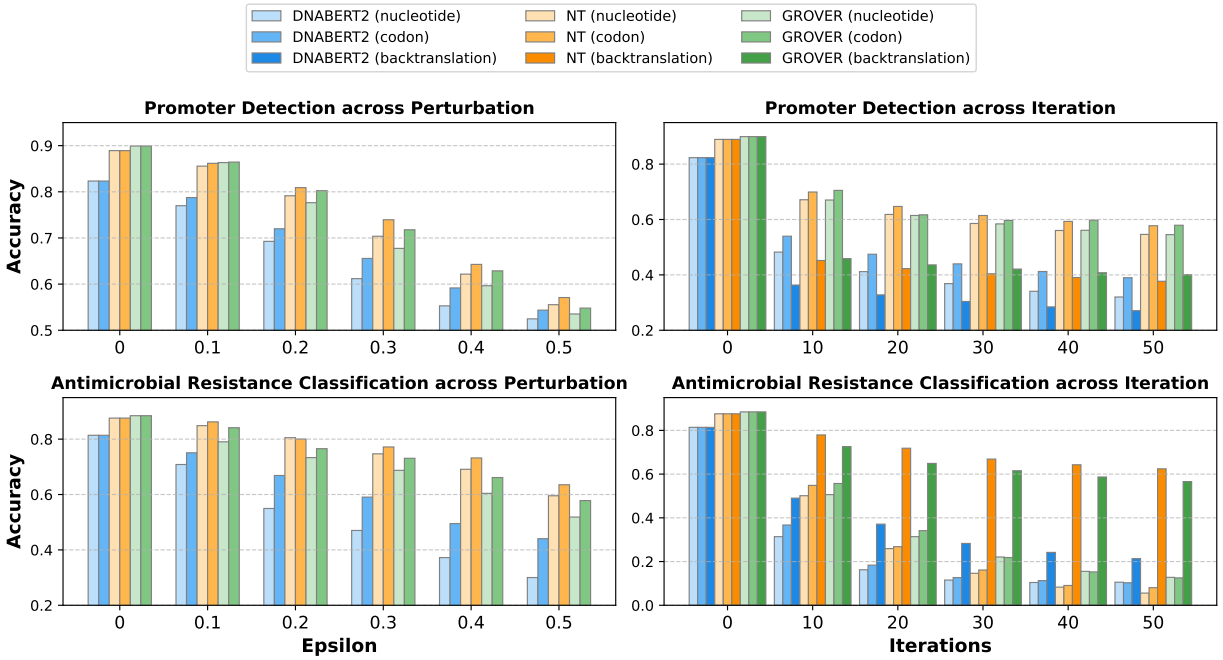

Figure 2. Anti-Microbial Resistance (AMR) gene Classification and Promoter Detection Results

their reliability in real-world scenarios.

In this paper, we explore the adversarial robustness of DNA language models in classification tasks. In addition to existing character-level attacks (nucleotide-level attacks (Kuleshov et al., 2021)), we introduce word-level (codon-level) and sentence-level (backtranslation) attacks. We demonstrate that these attacks lead to significant performance degradation in Anti-Microbial Resistance (AMR) gene classification (Yoo et al., 2024) and promoter detection (Zhou et al., 2023).

Furthermore, we investigate whether a simple adversarial training strategy can mitigate the vulnerabilities of DNA language models. The results show that while its effectiveness varies by attack type and dataset, it enhances robustness and can occasionally improve overall model performance.

Our main contributions can be summarized into three folds:

- We comprehensively explore the robustness of DNA language models through multi-granularity attack levels: **character-level (nucleotide-level)**, **word-level (codon-level)**, and **sentence-level (back-translation)**. Notably, we introduce **codon-level** and **back-translation** attacks, extending the scope beyond conventional character-level perturbations for a more thorough robustness evaluation.

- We demonstrate that increasing perturbation strength causes the most severe performance degradation in nucleotide-level perturbations, though at the risk of disrupting biological context. In the AMR gene classification with drug labels task, our proposed back-translation-based perturbations effectively degrade classification performance while preserving semantic meaning.

- We show that even simple adversarial training effectively mitigates vulnerabilities and improves classification accuracy over the original model, highlighting its potential to enhance the real-world reliability of DNA classification models.

2. Related Work

Adversarial Attacks in Text Classification Adversarial attacks in text classification can be broadly categorized into character-level, word-level, and sentence-level attacks (Ebrahimi et al., 2018). Character-level attacks involve random substitutions at the character level, or modifications at the word embedding or language model level (Jia & Liang, 2017; Alzantot et al., 2018). Word-level attacks follow a similar approach. Sentence-level attacks include paraphrasing, backtranslation (Ribeiro et al., 2018), and sentence embedding-level perturbations. (Zhang et al., 2019).

Adversarial Attacks in DNA Sequence Classification Kuleshov et al. (2021) and Montserrat & Ioannidis (2022) demonstrated that nucleotide-level adversarial attacks, including perturbations on single nucleotide polymorphism

(SNP)-based ancestry classification models, significantly degrade the performance of DNA sequence classifiers. In this paper, inspired by word-level and sentence-level perturbations, we further explore codon-level attacks and backtranslation-based DNA attacks. Recent jailbreak-style attacks on DNA language models, such as GeneBreaker (Zhang et al., 2025) and GenoArmory (Luo et al., 2025), provide complementary perspectives focused on targeted adversarial behaviors.

Adversarial Training Several studies have shown that adversarial training enhances model robustness (Madry et al., 2018; Zhang et al., 2022; Liu & Sun, 2023), including in DNA sequence classification, where character-level adversarial examples are used to improve resilience (Kuleshov et al., 2021; Montserrat & Ioannidis, 2022).

3. Adversarial Attack on DNA Language Models

We examine the impact of adversarial attacks on DNABERT2 (Zhou et al., 2023), Nucleotide Transformer (Dalla-Torre et al., 2023), and GROVER (Sanabria et al., 2024) across three levels attack strategies: character-level (nucleotide substitutions (Kuleshov et al., 2021)), word-level (codon modifications), and sentence-level (back-translation-based transformations). Our analysis focuses on DNA classification task using AMR gene (Yoo et al., 2024) and promoter detection (Zhou et al., 2023). These models differ in tokenizer granularity and pretraining scale, allowing us to compare adversarial robustness across a representative spectrum of DNA language model configurations. Further implementation details and benchmark details are provided in Appendix A and Appendix B.

Nucleotide-Level Attack Nucleotides are the fundamental building blocks of nucleic acids, such as DNA and RNA, crucial for storing and transmitting genetic information. In DNA sequences, nucleotides are represented by single characters, making nucleotide-level attacks analogous to character-level attacks in text. We conducted nucleotide-level attacks using an enumeration approach combined with trial and error search methods (Iyyer et al., 2018; Ribeiro et al., 2018; Belinkov & Bisk, 2018).

Codon-Level Attack A codon is a sequence of three nucleotides in DNA or RNA that specifies a particular amino acid. Codons are essential for protein synthesis. Codon-level attack is similar to word-level attacks and is also performed using the enumeration search method including trial and error.

Backtranslation Attack We introduce backtranslation based attack method for DNA sequences. In biological translation, translation refers to converting mRNA sequences into protein amino acid sequences, and reverse translation refers to generating possible nucleotide sequences based on protein amino acid sequences.

These three perturbation strategies differ in both granularity and biological semantic preservation. Nucleotide-level attacks modify single bases and often introduce biologically implausible variations. Codon-level attacks perturb three-base units that may affect amino acid interpretations. Backtranslation-based attacks, in contrast, generate synonymous nucleotide sequences by substituting codons with biologically equivalent alternatives, thereby preserving the encoded amino acid sequence while altering the surface form of the DNA.

4. Results

4.1. Attack Effectiveness and Context Preservation

In Figure 2, nucleotide-level attacks (Kuleshov et al., 2021) result in the most significant performance degradation. While these attacks are highly effective. Conversely, codon-level attacks are better at preserving sequence meaning since they group sequences based on meaningful units, such as codons. Backtranslation attacks preserve context the best but are relatively less effective in reducing accuracy. This trend is evident in the first graph, where there is a sharp decline in model accuracy with increasing epsilon values for nucleotide-level attacks, particularly affecting DNABERT2 (Zhou et al., 2023). In contrast, the Nucleotide Transformer maintains relatively high accuracy even as epsilon increased.

4.2. Model Robustness Comparison

The Nucleotide Transformer (Dalla-Torre et al., 2023) and GROVER (Sanabria et al., 2024) demonstrate superior robustness compared to DNABERT2 (Zhou et al., 2023). Due to its larger model capacity and ability to handle longer sequences, the Nucleotide Transformer exhibits increased resistance to adversarial attacks. Compared to DNABERT2, GROVER employs an extensive BPE optimization process involving hundreds of iterations enhancing its ability to model long-range dependencies in DNA sequences. These architectural refinements are likely to improve its robustness against adversarial attacks. This robustness is particularly notable in experiments involving backtranslation attacks. The second graph in Figure 2 illustrates that DNABERT2's accuracy significantly drops with increasing iterations, whereas the Nucleotide Transformer and GROVER maintain high accuracy, demonstrating its robustness due to its larger model capacity or longer sequence processing capabilities.

Moreover, DNABERT2 promoter detection model demon-

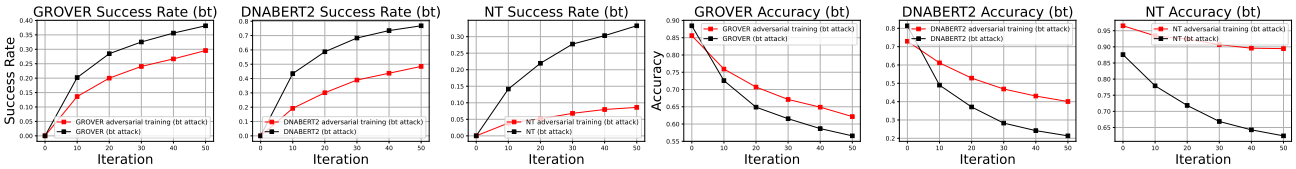

Figure 3. Comparison of success rates and accuracies between adversarial training and standard training in AMR gene Classification across increasing iterations of backtranslation attacks

strates greater robustness compared to DNABERT2 AMR drug classification models, potentially due to the larger volume of training data utilized. As illustrated in the third graph, with increasing epsilon values, the accuracy of DNABERT2 promoter detection model declines, though to a lesser extent than in AMR detection models. However, DNABERT2 still performs worse than the Nucleotide Transformer and GROVER under nucleotide-level attacks. The fourth graph further shows that, in AMR drug classification tasks, the Nucleotide Transformer and GROVER maintain the highest accuracy against backtranslation attacks, whereas DNABERT2 appears to be relatively more vulnerable to backtranslation attacks. We further address the robustness of a general language model in Appendix C.

4.3. Comparative Analysis of Classification Tasks

The experimental results indicate that the Nucleotide Transformer and GROVER possess higher robustness against adversarial attacks compared to DNABERT2, especially in the context of backtranslation attacks. Additionally, DNABERT2 trained with larger datasets, such as those used for promoter detection, demonstrates enhanced robustness. This comparative analysis highlights the importance of model capacity, training data size and The ability to process long sequences in improving the resilience of DNA sequence classification models against various adversarial attack methods.

4.4. Performance Against Defense Method

We further investigate adversarial training as a defense mechanism against adversarial attacks, incorporating adversarial examples into the training process to improve model robustness.

Figure 3 illustrates the changes in success rate and accuracy over iterations when applying the Back-translate attack in the AMR drug classification task. The results indicate that the success rate of adversarially trained models (red) increases at a slower rate compared to standard training models (black). This indicates that models trained with adversarial examples are more resistant to such attacks, as shown by the slower increase in attack success rate. Furthermore, the accuracy graphs show that adversarially trained models (red) exhibit a smaller decline in accuracy compared to standard training (black). This observation implies that adversarial training improves the model's generalization performance and mitigates performance degradation when exposed to adversarial attacks. For GROVER and DNABERT2, the initial accuracy of adversarially trained models is lower than that of standard training models in Figure 6 and Figure 8. However, as iterations increase, adversarially trained models demonstrate relatively higher accuracy, indicating that repeated exposure to attacks reinforce the model's resilience. In contrast, Nucleotide Transformer exhibits superior robustness, with adversarially trained models outperforming standard training models in accuracy from the outset. Additionally, the accuracy degradation over iterations is the least pronounced among the three models. This suggests that the larger model capacity and ability to handle longer sequences of Nucleotide Transformer further amplify the benefits of adversarial training, making it more resistant to adversarial perturbations. Further details are described in Appendix E

5. Conclusion

In this study, we utilize three methods to generate DNA adversarial examples, proposing techniques for generating adversarial examples using nucleotide based attack, codon-based attacks and backtranslation. These methods successfully degrade the performance of AMR drug classification and promoter detection models. Nucleotide-level attacks are the most effective, although they risk disrupting the DNA context. Conversely, backtranslation better preserves meaning and context. Notably, adversarial training as a defense strategy proves effective in improving model robustness, and in some cases, even surpassing the performance of standard training methods. By leveraging adversarial training, models become more resilient against such attacks. These findings underscore the importance of building DNA language models that are robust not only to worst-case adversarial manipulations, but also to natural sources of variation. Such robustness is especially critical in practical settings where sequencing noise, synthetic design variation, or cross-lab inconsistencies may lead to unreliable predictions.

6. Limitations

This study has several limitations. First, the datasets used are limited to AMR gene classification with drug classes which may restrict the generalization of the results. Expanding experiments to include datasets such as AMR gene families or AMR mechanisms would provide a more comprehensive evaluation of adversarial robustness.

The study evaluates adversarial attack methods on three distinct DNA language models such as DNABERT2, GROVER, and Nucleotide Transformer. While these models cover a range of architectures and sizes, extending the evaluation to additional DNA classification models could provide further insights into model-specific robustness.

Acknowledgements

This work was supported in part by the National Science Foundation (NSF) under Grant Number 2107108.

References

Alcock, B. P., Raphenya, A. R., Lau, T. T. Y., Tsang, K. K., Bouchard, M., Edalatmand, A., Huynh, W., Nguyen, A.-L. V., Cheng, A. A., Liu, S., Min, S. Y., Miroshnichenko, A., Tran, H.-K., Werfalli, R. E., Nasir, J. A., Oloni, M., Speicher, D. J., Florescu, A., Singh, B., Faltyn, M., Hernandez-Koutoucheva, A., Sharma, A. N., Bordeleau, E., Pawlowski, A. C., Zubyk, H. L., Dooley, D., Griffiths, E., Maguire, F., Winsor, G. L., Beiko, R. G., Brinkman, F. S. L., Hsiao, W. W. L., Domselaar, G. V., and McArthur, A. G. CARD 2020: antibiotic resistome surveillance with the comprehensive antibiotic resistance database. *Nucleic Acids Research*, 48(D1):D517–D525, Jan 2020. doi: 10.1093/nar/gkz935. URL https://doi.org/10.1093/nar/gkz935. Published online 2019 Oct 29.

Alzantot, M., Sharma, Y., Elgohary, A., Ho, B.-J., Srivastava, M., and Chang, K.-W. Generating natural language adversarial examples. In Riloff, E., Chiang, D., Hockenmaier, J., and Tsujii, J. (eds.), *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 2890–2896, Brussels, Belgium, October-November 2018. Association for Computational Linguistics.

Belinkov, Y. and Bisk, Y. Synthetic and natural noise both break neural machine translation. In *International Conference on Learning Representations*, 2018. URL https://openreview.net/forum?id=BJ8vJebC-.

Bonin, N., Doster, E., Worley, H., Pinnell, L. J., Bravo, J. E., Ferm, P., Marini, S., Prosperi, M., Noyes, N., Morley, P. S., and Boucher, C. MEGARes and AMR++, v3.0: an updated comprehensive database of antimicrobial resistance determinants and an improved software pipeline for classification using high-throughput sequencing. *Nucleic Acids Research*, 51(D1):D744–D752, Jan 2023. doi: 10.1093/nar/gkac1047.

Buchfink, B., Reuter, K., and Drost, H.-G. Sensitive protein alignments at tree-of-life scale using diamond. *Nature Methods*, 18:366–368, 2021.

Cook, C. E., Bergman, M. T., Finn, R. D., Cochrane, G., Birney, E., and Apweiler, R. The european bioinformatics institute in 2016: Data growth and integration. *Nucleic Acids Research*, 44(D1):D20–D26, 2016. doi: 10.1093/nar/gkv1352.

Dalla-Torre, H., Gonzalez, L., Mendoza Revilla, J., Lopez Carranza, N., Grzywaczewski, A. H., Oteri, F., Dallago, C., Trop, E., Sirelkhatim, H., Richard, G., Skwark, M., Beguir, K., Lopez, M., and Pierrot, T. The nucleotide transformer: Building and evaluating robust foun-

dation models for human genomics. *bioRxiv*, 2023. doi: 10.1101/2023.01.11.523679.

Ebrahimi, J., Rao, A., Lowd, D., and Dou, D. HotFlip: White-box adversarial examples for text classification. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL 2018)*, 2018. URL https://aclanthology.org/P18-2006/.

Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., and Chen, W. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=nZeVKeeFYf9.

Iyyer, M., Wieting, J., Gimpel, K., and Zettlemoyer, L. Adversarial example generation with syntactically controlled paraphrase networks. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, June 2018. URL https://aclanthology.org/N18-1170/.

Jia, R. and Liang, P. Adversarial examples for evaluating reading comprehension systems. In Palmer, M., Hwa, R., and Riedel, S. (eds.), *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pp. 2021–2031, Copenhagen, Denmark, September 2017. Association for Computational Linguistics.

Kuleshov, V., Nikishin, E., Thakoor, S., Lau, T., and Ermon, S. Quantifying and understanding adversarial examples in discrete input spaces. *arXiv preprint arXiv:2112.06276*, 2021. doi: 10.48550/arXiv.2112.06276.

Liu, J., Yang, M., Yu, Y., Xu, H., Li, K., and Zhou, X. Large language models in bioinformatics: Applications and perspectives. *arXiv*, Preprint(arXiv:2401.04155v1), Jan 2024. URL https://arxiv.org/abs/2401.04155v1.

Liu, T. and Sun, Y. End-to-end adversarial sample generation for data augmentation. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pp. 11359–11368, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-emnlp.760. URL https://aclanthology.org/2023.findings-emnlp.760/.

Luo, H., Qiu, C., Wang, Y., Wu, S., Yu, J., Liu, H., Wang, B., and Chen, Y. Genoarmory: A unified evaluation framework for adversarial attacks on genomic foundation models. *arXiv preprint arXiv:2505.10983*, 2025. doi: 10.48550/arXiv.2505.10983.

Ma, X., Shao, Y., Tian, L., Flasch, D. A., Mulder, H. L., Edmonson, M. N., Liu, Y., Chen, X., Newman, S., Nakitandwe, J., Li, Y., Li, B., Shen, S., Wang, Z., Shurtleff, S., Robison, L. L., Levy, S., Easton, J., and Zhang, J. Analysis of error profiles in deep next-generation sequencing data. *Genome Biology*, 20(50), March 2019. doi: 10.1186/s13059-019-1659-6. URL https://genomebiology.biomedcentral.com/articles/10.1186/s13059-019-1659-6.

Madry, A., Makelov, A., Schmidt, L., Tsipras, D., and Vladu, A. Towards deep learning models resistant to adversarial attacks. In *International Conference on Learning Representations*, 2018. URL https://openreview.net/forum?id=rJzIBfZAb.

Montserrat, D. M. and Ioannidis, A. G. Adversarial attacks on genotype sequences. *bioRxiv*, 2022. doi: 10.1101/2022.11.07.515527.

Ono, Y., Asai, K., and Hamada, M. PBSIM2: a simulator for long-read sequencers with a novel generative model of quality scores. *Bioinformatics*, 37(5): 589–595, March 2021. doi: 10.1093/bioinformatics/btaa835. URL https://doi.org/10.1093/bioinformatics/btaa835. Published: 25 September 2020.

Ribeiro, M. T., Singh, S., and Guestrin, C. Semantically equivalent adversarial rules for debugging nlp models. In Gurevych, I. and Miyao, Y. (eds.), *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 856–865, Melbourne, Australia, July 2018. Association for Computational Linguistics.

Sanabria, M., Hirsch, J., Joubert, P. M., and Poetsch, A. R. Dna language model grover learns sequence context in the human genome. *Nature Machine Intelligence*, 6:911–923, August 2024. doi: 10.1038/s42256-024-00872-0.

Sarumi, O. A. and Heider, D. Large language models and their applications in bioinformatics. *Computational and Structural Biotechnology Journal*, 23:3498–3505, December 2024.

Steinegger, M. and Söding, J. Mmseqs2 enables sensitive protein sequence searching for the analysis of massive data sets. *Nature Biotechnology*, 35:1026–1028, 2017.

Yoo, H., Sokhansanj, B., Brown, J. R., and Rosen, G. Predicting anti-microbial resistance using large language models. *arXiv preprint arXiv:2401.00642*, 2024. doi: 10.48550/arXiv.2401.00642.

Zhang, C., Zhou, X., Wan, Y., Zheng, X., Chang, K.-W., and Hsieh, C.-J. Improving the adversarial robustness of

nlp models by information bottleneck. In *Findings of the Association for Computational Linguistics: ACL 2022*, pp. 3588–3598, Dublin, Ireland, May 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.findings-acl.284. URL https://aclanthology.org/2022.findings-acl.284/.

Zhang, H., Chen, H., Song, Z., Boning, D., Dhillon, I. S., and Hsieh, C.-J. The limitations of adversarial training and the blind-spot attack. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=H1g0piA9tQ. Huan Zhang and Hongge Chen contributed equally.

Zhang, Z., Zhou, Z., Jin, R., Cong, L., and Wang, M. Genebreaker: Jailbreak attacks against dna language models with pathogenicity guidance. *arXiv preprint arXiv:2505.23839*, 2025. URL https://doi.org/10.48550/arXiv.2505.23839.

Zhou, Z., Ji, Y., Li, W., Dutta, P., Davuluri, R., and Liu, H. DNABERT-2: Efficient foundation model and benchmark for multi-species genome. *arXiv preprint arXiv:2306.15006*, 2023. doi: 10.48550/arXiv.2306.15006.

A. Implementation Details

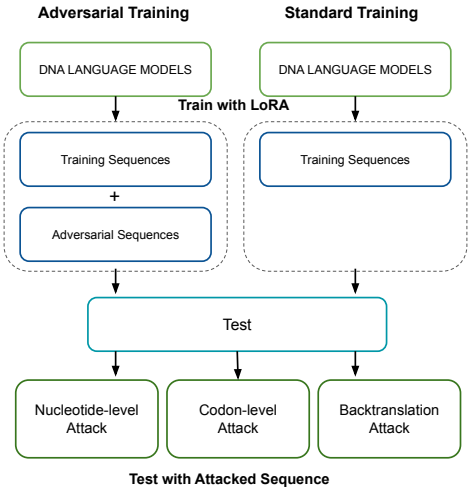

Figure 4. Overview of Adversarial Training(Fine-tuning) for DNA Language Models and Testing Adversarial Attacks

We utilize LoRA (Hu et al., 2022) to finetune the DNA language models for each benchmarks. The learning rate is set to 0.0005, the number of epochs is 2, the batch size is 64. Single A100 GPU is used, and the entire process takes approximately one hour. The number of parameters for each model is as follows: DNABERT-2 (117M), Nucleotide Transformer (2.5B), and GROVER (90M).

B. Details of Benchmarks

Antimicrobial Resistance Classification We employ the methodology outlined in Yoo et al. (2024) to integrate datasets from CARD v2 (Alcock et al., 2020) and MEGARes v3 (Bonin et al., 2023) for investigating antibiotic resistance. Classes with fewer than 15 instances are excluded. The remaining data is divided into training, testing, and validation sets, maintaining the same proportions as the study. Data integration is performed using the European Bioinformatics Institute Antibiotic Resistance Ontology (EBI ARO; Cook et al., 2016), with irrelevant classes excluded. Overall, the drug class classification dataset for AMR gene classification is split into 75% training data, 20% test data, and 5% validation data. There are a total of 9 drug classes.

Promoter Detection The promoter detection data from the Genome Understanding Evaluation (GUE) benchmark dataset introduced in the DNABERT-2 is used to detect promoters in gene sequences (Zhou et al., 2023). Promoters are DNA sequences that regulate the initiation of gene transcription and play a crucial role in gene expression regulation and understanding biological phenomena.

The data is extracted from gene sequences of various species, providing a broad context for biological research. Each data point includes both sequences that contain promoters and those that do not, enabling the model to learn how to distinguish promoters. This dataset includes complex promoter sequence elements known as MIX elements, which are found in specific genes. We utilize the prom 300 all dataset from the GUE benchmark for promoter detection. The dataset is split into approximately 80% training, 10% development, and 10% test data, with 2 classes for binary classification.

C. Adversarial Evaluation of GPT-4o on DNA AMR Classification

We present additional analysis on the adversarial robustness of `gpt-4o-mini-2024-07-18` fine-tuned for DNA AMR classification. This study has demonstrated that DNA Language Models are vulnerable to adversarial attacks, and we extend this investigation to examine whether general-purpose LLMs, when adapted for biological sequence classification, exhibit similar weaknesses.

By applying codon-level and nucleotide-level perturbations, we observe a sharp decline in model accuracy as perturbation

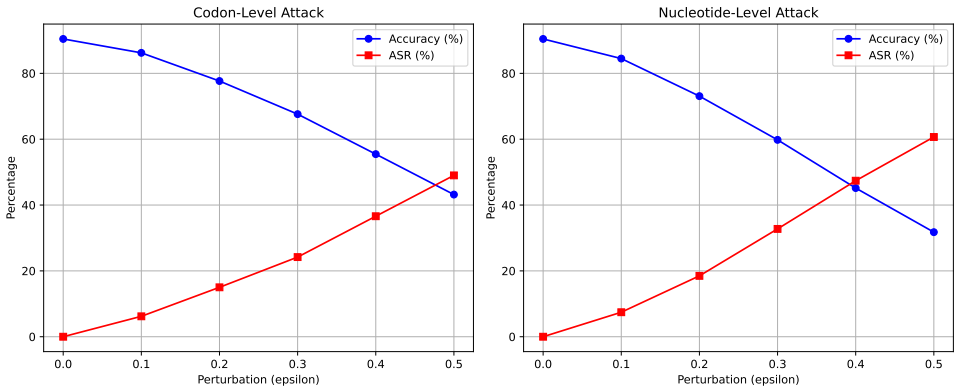

Figure 5. Accuracy and attack success rate (ASR) under codon-level and nucleotide-level perturbations in Antimicrobial Resistance Drug Classification Across Different Perturbation Levels

intensity increased (dropping to 31.76% at the highest level), while the attack success rate (ASR) rises to a maximum of 60.67% in Figure 5.

With the growing interest in applying LLMs to DNA sequence analysis (Sarumi & Heider, 2024; Liu et al., 2024), ensuring robustness remains a critical challenge. Adversarial training and domain-specific robustness strategies may be necessary to enhance the reliability of such models in biological sequence tasks.

D. Effect of Realistic Sequencing Errors on Robustness

To evaluate model robustness under more realistic sequencing scenarios, we simulate sequencing errors using PBSIM2 (Ono et al., 2021), incorporating real-world error patterns observed in PacBio and Oxford Nanopore sequencing. Sequencing Error Simulation is a technique used to model errors occurring during the genome sequencing process, playing a crucial role in genome analysis, algorithm evaluation, and robustness assessment in bioinformatics. PBSIM2, widely used for this purpose, simulates insertion, deletion, and substitution errors, as well as the length distribution observed in long-read sequencing technologies such as PacBio and Oxford Nanopore, enabling the generation of reads that closely resemble real sequencing data. The simulated sequencing error data can be utilized to test model robustness by incorporating noise similar to that encountered in real sequencing experiments. This holds significant importance from the perspective of adversarial training and robustness assessment, as it helps deep learning-based genomic analysis models maintain predictive accuracy while improving their resilience to sequencing errors. Our experiments show that the model's performance degrades under these conditions, indicating that realistic sequencing errors can pose a challenge to robustness. In the promoter detection task with Nucleotide Transformer, accuracy dropped from 0.8892 to 0.8608.

E. Comparisons of Adversarial Training and Standard Training

Figure 6, Figure 7, Figure 8, and Figure 9 compare the performance differences between adversarially trained models and conventionally trained models (black) as the number of attack iterations and perturbation magnitude (epsilon) increase. Epsilon represents the rate of change in the sequence for the attack, indicating the proportion of the sequence that is altered. Iterations indicate the number of times the attack is performed. While all models experience a decrease in accuracy as the number of iterations and perturbation levels increase, adversarially trained models generally maintain higher accuracy and exhibit a more gradual decline in performance. In terms of attack success rate, conventionally trained models show a sharp increase in attack success as the number of iterations and perturbation magnitude grow, whereas adversarially trained models exhibit a more moderate increase, demonstrating greater resistance to adversarial attacks. These findings suggest that adversarial training effectively enhances model robustness against iterative attacks and higher perturbation magnitudes while mitigating performance degradation. Figures present accuracy and success rate trends under different perturbation settings. However, all results are based on a single run.

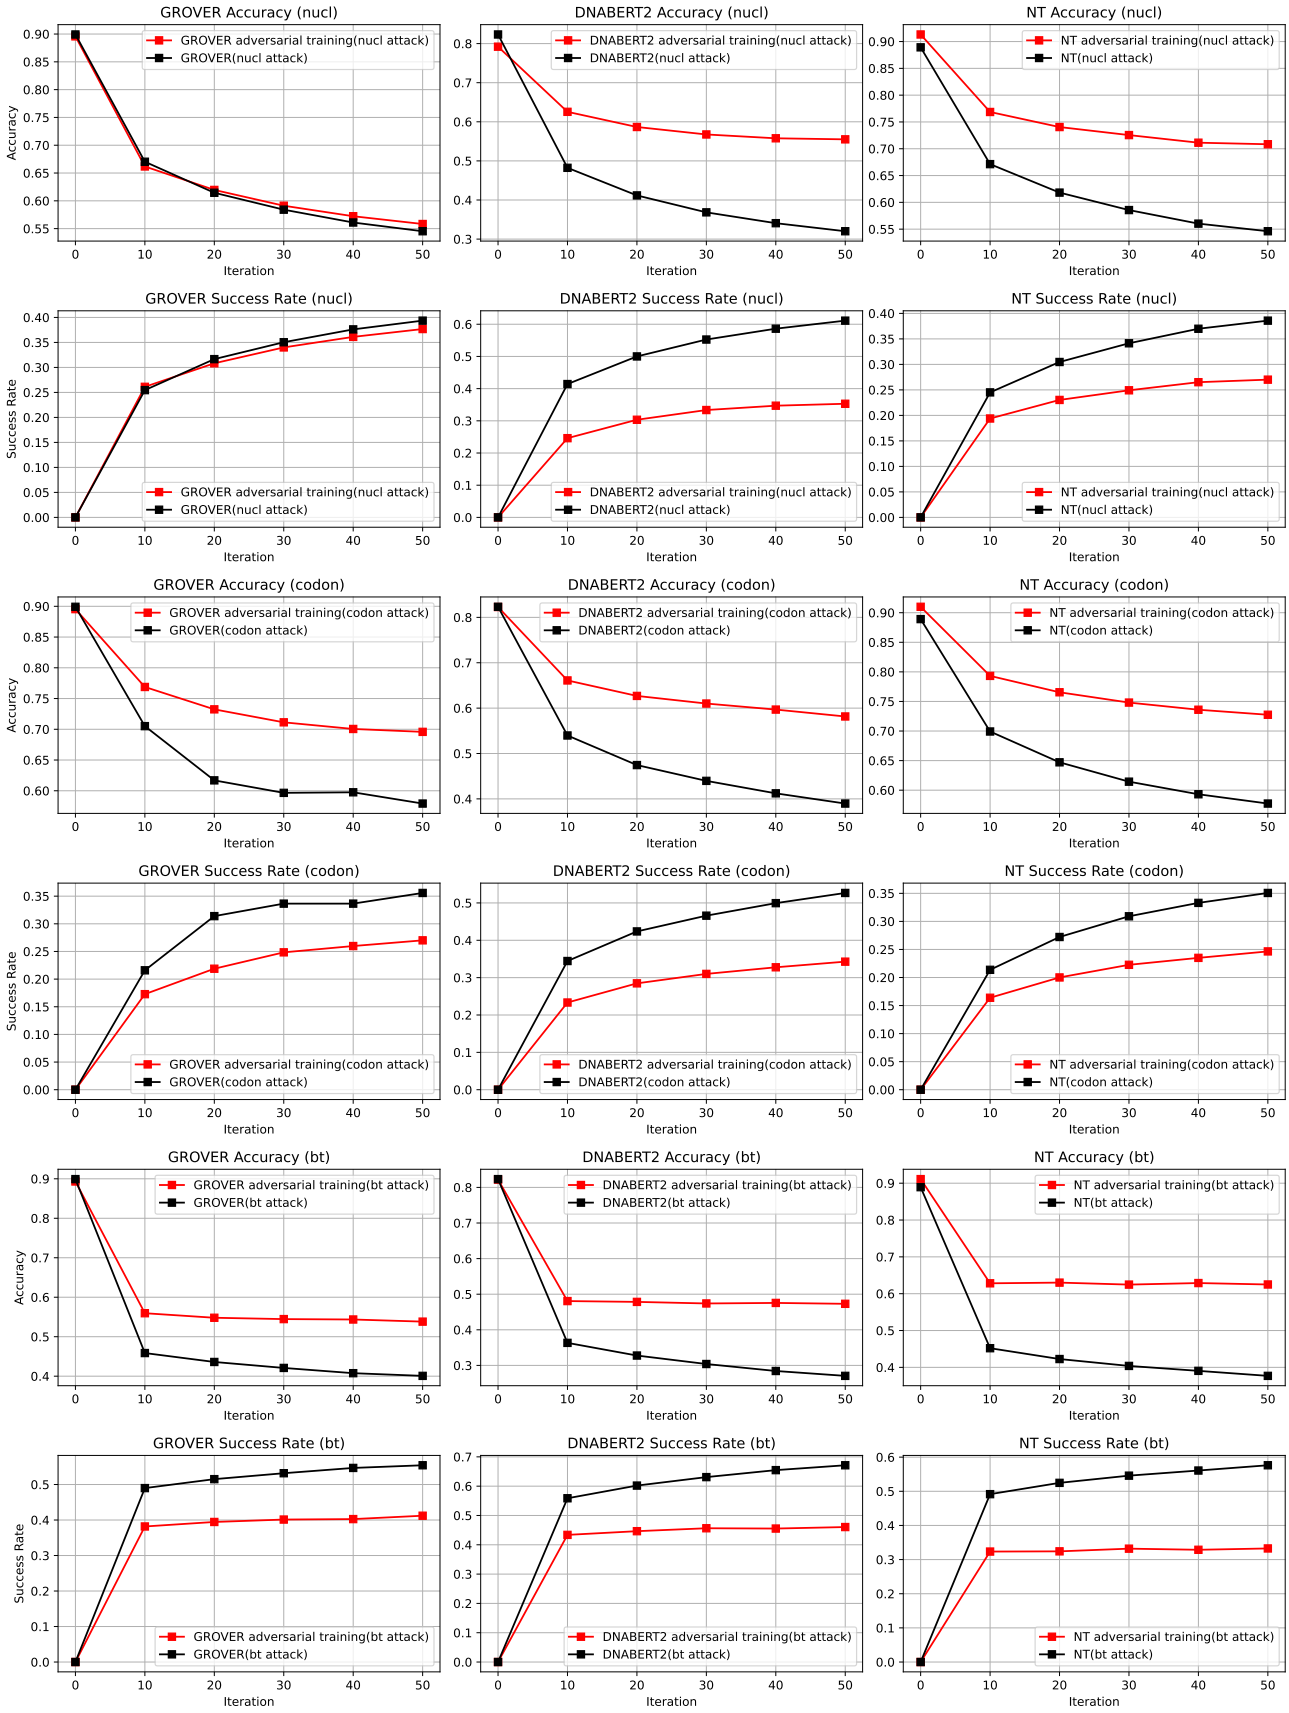

Figure 6. Comparison of Adversarial Training and Standard Training in Promoter Detection with Increasing Iterations

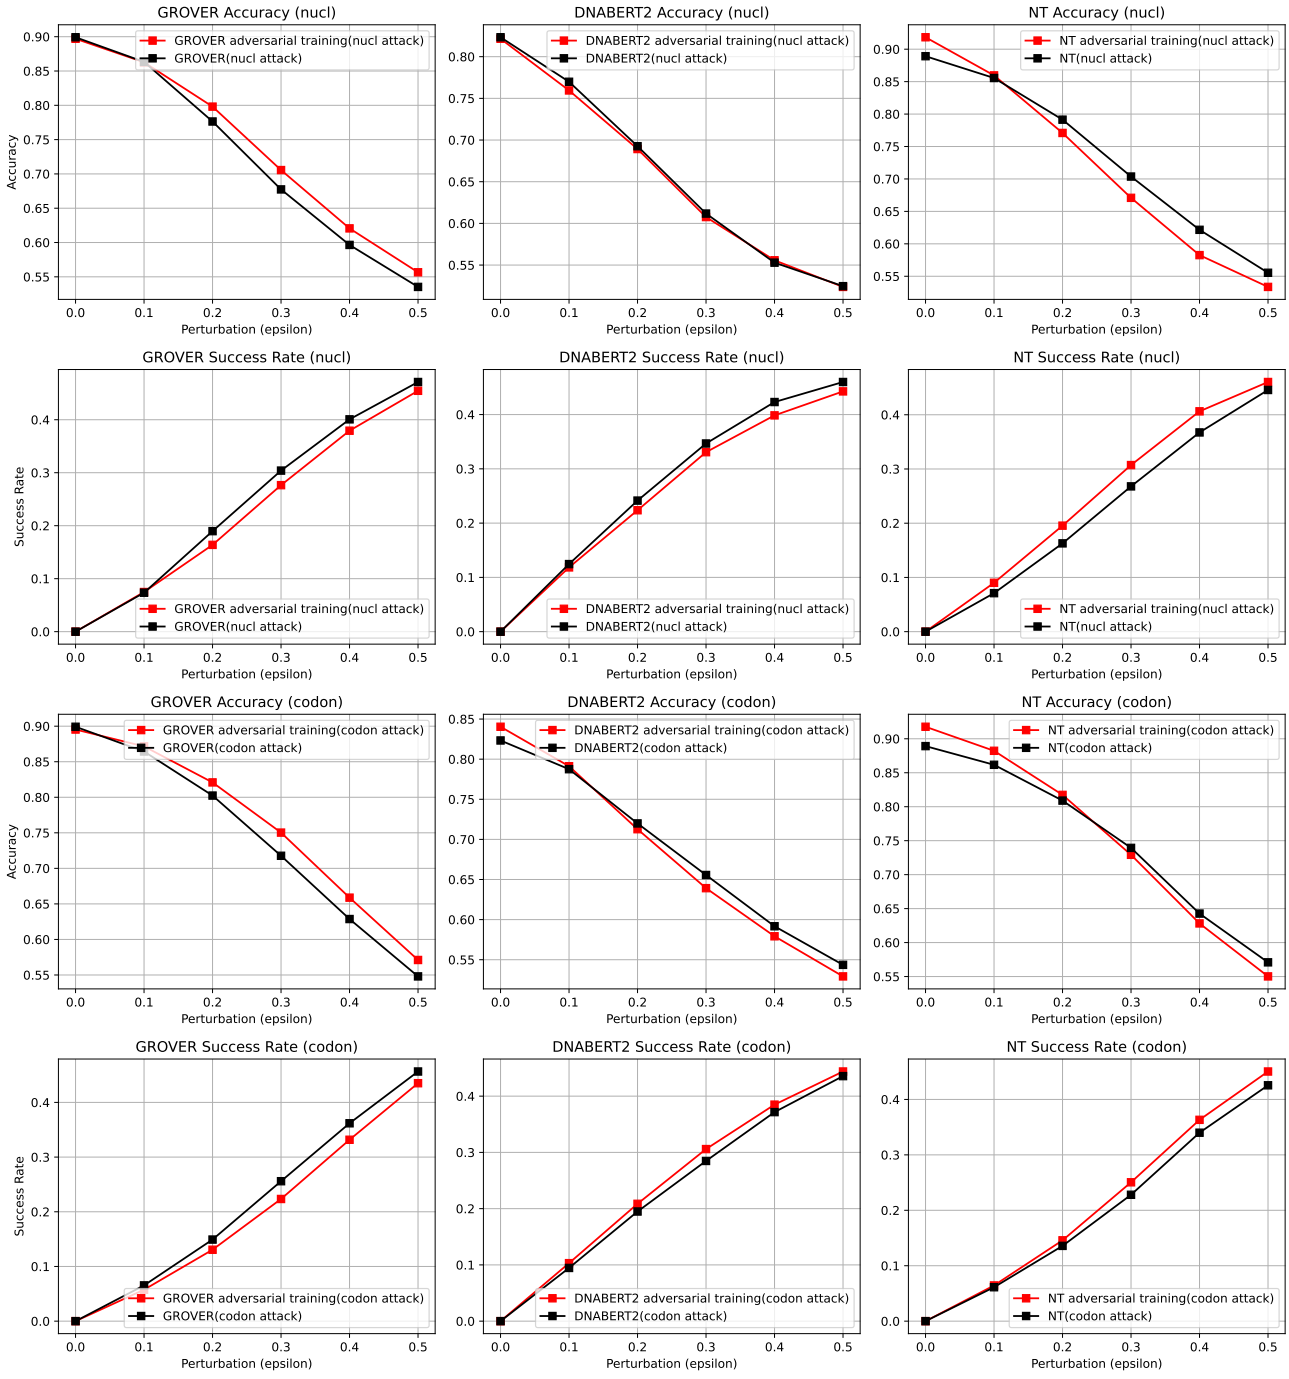

Figure 7. Comparison of Adversarial Training and Standard Training in Promoter Detection Across Different Perturbation Levels

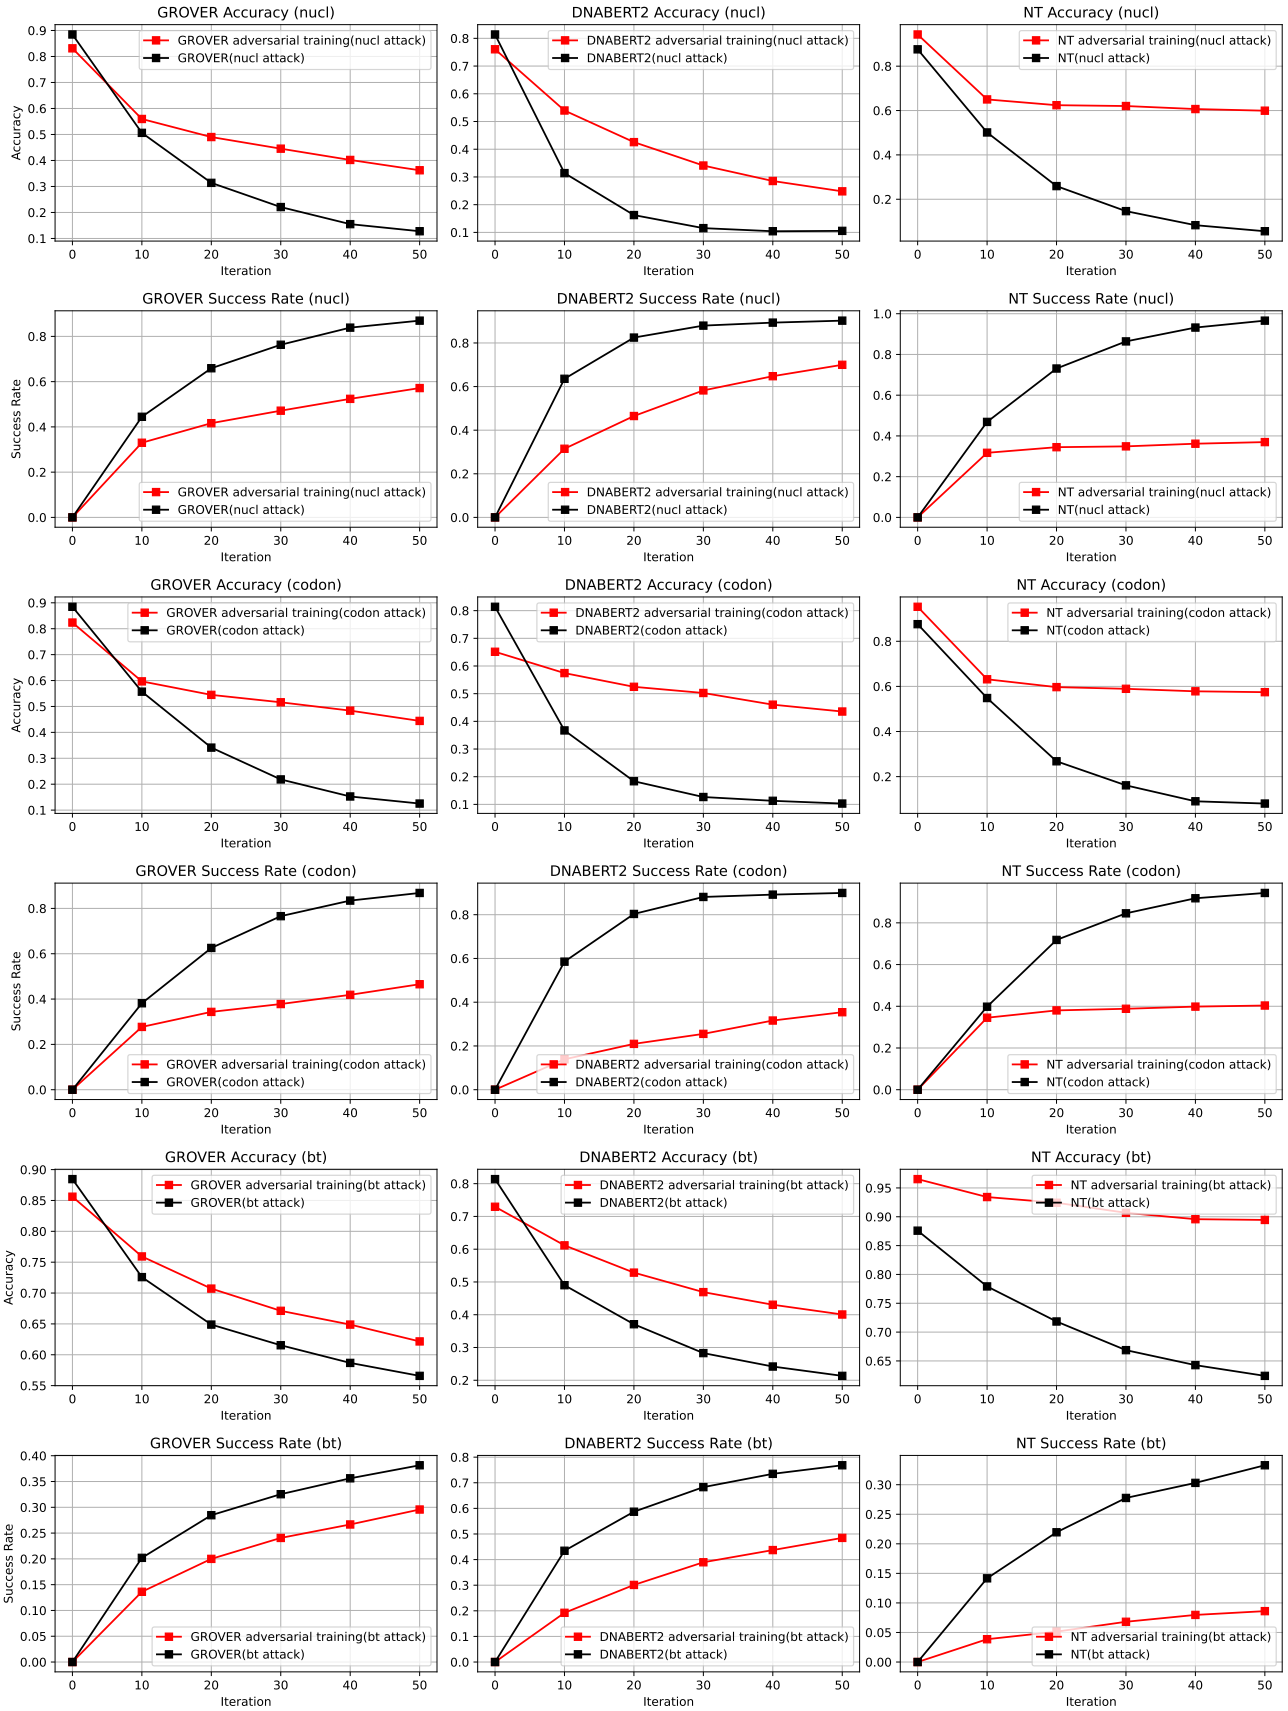

Figure 8. Comparison of Adversarial Training and Standard Training in Antimicrobial Resistance Drug Classification with Increasing Iterations

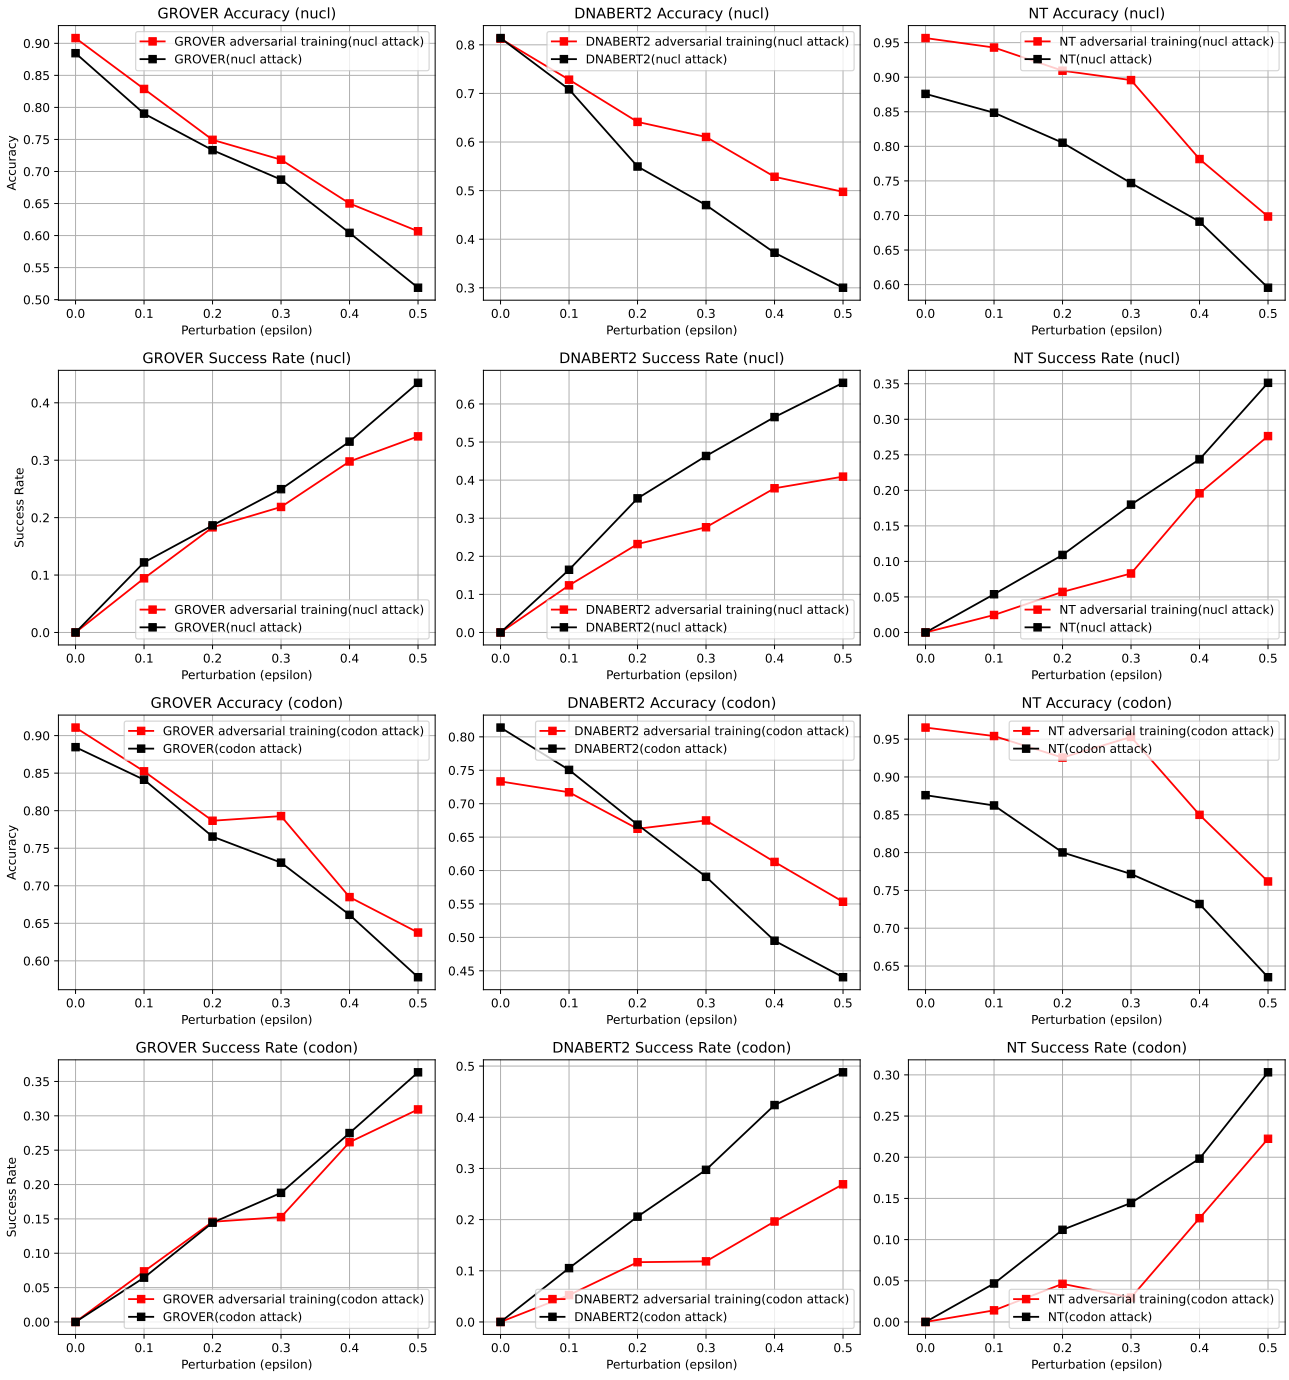

Figure 9. Comparison of Adversarial Training and Standard Training in Antimicrobial Resistance Drug Classification Across Different Perturbation Levels