

# Adversarial Testing for Large Language Models: Evaluating and Enhancing Robustness with AutoDAN and Fine-Tuning Techniques

Jatin Garg, Punit Maheshwari, Mustafa Saify, Shashank Srivastava, Srinivasa Rao Aravilli

Capital One

Ascendas ITPB-SEZ, 3rd Floor, Voyager Building, International Tech Park, Whitefield Main Road  
Bangalore, KA 560066, India

jatin.garg@capitalone.com, punit.maheshwari@capitalone.com, mustafa.saify@capitalone.com,  
shashank.srivastava@capitalone.com, srinivasarao.aravilli@capitalone.com

## Abstract

Large Language Models (LLMs) exhibit impressive generative capabilities but remain vulnerable to adversarial inputs, exposing potential risks such as data leakage, harmful content generation, and jailbreak attacks. Jailbreak attacks can fool LLMs and elicit harmful output even with safety alignments and guardrails. In this work, we performed a case study using AutoDAN – an automated adversarial attack generator – to stress-test open-source LLMs. We measure baseline attack success rates (ASRs) and then apply various fine-tuning defenses to mitigate these vulnerabilities. We used AutoDAN to construct adversarial datasets and evaluate the robustness of fine-tuned LLMs. Using Llama-3-8b-instruct as the base model, we apply full Supervised Fine-Tuning (SFT) AutoDAN-style attacks. Our preliminary experiments show a reduction in jailbreak success rate after fine-tuning, while maintaining usefulness and coherence for benign queries. We conclude by outlining best practices for deploying adversarially resilient LLMs in production environments and future work to continue research with adversarial attack vulnerabilities and agentic workflows.

## Introduction

Recent progress in LLMs has enabled powerful applications, but also raised serious safety concerns. Even models aligned by reinforcement learning with human feedback (RLHF) can be “jailbroken” by cleverly crafted prompts. Researchers have demonstrated that malicious inputs – often generated by LLMs themselves – can bypass safety guardrails and induce prohibited content (Liu et al. 2024b). Manual jailbreaks (e.g. prompt-DANs) have given way to automated red-teaming tools that search the prompt space for weaknesses. For instance, AutoDAN is a gradient-based framework that autonomously discovers diverse adversarial prompts while optimizing for readability (Liu et al. 2024b). Such methods reveal a broad array of semantic vulnerabilities that generalize across tasks and transfer to black-box models.

In this paper, we use AutoDAN as a practical test harness to probe LLM safety. We first run AutoDAN against a target model (e.g. Llama3-8B) (Dubey et al. 2024) to obtain a high ASR with no defense. We then explore mitigation via fine-tuning. We compare three tuning strategies:

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: An adversarial attack is an input designed to trick an LLM into misbehaving, bypassing safety mechanisms, and generating harmful or restricted content.

(1) standard SFT on a safety dataset (harmful queries → refusal responses), (2) LoRA (low-rank adapters) fine-tuning on the same data, and (3) QLoRA (quantized LoRA) for efficient tuning of larger models. We measure ASR and task performance after each defense. Our contributions are: (a) an end-to-end adversarial evaluation using AutoDAN on selected LLMs, and (b) an empirical comparison of baseline vs fine-tuned model trained with the adversarial prompt, (c) set out future directional work to continue evaluation with Low-Rank Adaptation (LoRA), Quantized LoRA (QLoRA) and other efficient techniques to make models prompt inject aware yet preserve base model capabilities with minimal safety tax. And create detailed benchmarks with different knowledge and reasoning tasks with a zoo of models.

## Related Work

### Adversarial Attacks and Jailbreaking

A growing body of research exposes LLM vulnerabilities. Early jailbreak techniques exploited linguistic tricks or role-play prompts to bypass filters. More recent attacks use automated strategies: e.g. Greedy Coordinate Gradient (GCG) (Zou et al. 2023) and MasterKey craft token-level adversar-

	CAPABILITY ( $\uparrow$ )			SECURITY ( $\uparrow$ )		
	Knowledge	Multi-turn reasoning	Math	Code	Black-Box	White-Box
Llama-2-7B-Chat	52.0	6.54	4.6	12.8	94.1	75.4
Llama-3-8B-Instruct	68.9	8.01	79.6	61.6	97.0	44.0
Mixtral 8x7B Instruct	74.9	8.30	74.4	45.1	-	-

Table 1: Performance comparison of various LLMs on capability and security metrics.

ial, and gradient-based methods like AutoDAN generate interpretable suffixes from scratch (Liu et al. 2024b). Test-time scaling (e.g., generating many candidates and scoring them) has also been weaponized for adversarial use (Liu and Xiao 2025). These attacks achieve nontrivial ASRs against even state-of-the-art models (e.g., GPT-4). Metrics such as Attack Success Rate (ASR) quantify the fraction of adversarial prompts that cause a harmful response. We adopt ASR (in percent) as a primary metric alongside moderation “harmfulness” scores from safety classifiers.

### Defense via Fine-Tuning

A common defense is to fine-tune LLMs on safe data. However, fine-tuning can itself degrade safety if done carelessly. (Qi et al. 2024) shows that fine-tuning on non-harmful data may inadvertently increase harmfulness. Huang et al. (2024) demonstrate that SFT for complex reasoning often introduces a “safety tax,” reducing model safety markedly. To address this, recent works propose safety-focused tuning methods. Data-centric methods like SAFT (Safety-Aware Fine-Tuning) (Choi, Du, and Li 2024) filter or remove harmful examples during tuning. Optimization-based methods include “SafeLoRA” and “SaLoRA” which regularize LoRA updates to lie in a safer subspace (Hsu et al. 2024) (Li et al. 2025). These techniques explicitly construct an alignment matrix derived from the gradients of safety examples. This matrix acts as a regularizer, projecting low-rank updates into a subspace that maximizes refusal probability for harmful queries while orthogonalizing them against directions that would degrade general utility. The Lisa framework (Huang et al. 2024) proposes splitting training into alignment and task states with a proximal term to maintain safety.

### Parameter-Efficient Tuning

LoRA (Hu et al. 2022) inserts low-rank adapter matrices into each layer and fine-tunes only those, greatly reducing trainable parameters. LoRA has been widely adopted due to its efficiency. QLoRA (Dettmers et al. 2023) further combines LoRA with 4-bit quantization, enabling fine-tuning of very large models on commodity GPUs. Recent work shows that LoRA-based tuning can preserve more of the original model’s capabilities. For example, “LoRA is All You Need for Safety Alignment” (Xue and Mirzasoleiman 2025) found that applying LoRA on a refusal dataset achieved safety comparable to full fine-tuning while maintaining reasoning. In summary, LoRA/QLoRA tuning is promising for aligning LLMs without large performance losses.

## Methodology

Our methodology follows a comprehensive four-stage framework as shown in Figure 2, spanning from initial attack generation to final safety assessment.

1. **Attack Generation:** Adversarial probing using the AutoDAN Hierarchical Genetic Algorithm (HGA).
2. **Data Curation:** Hybrid data generation based on different jailbreak prompts and safety responses.
3. **Defensive Fine-Tuning:** Fine-tuning using both SFT and PEFT techniques on the curated data.
4. **Evaluation:** Comprehensive evaluation for safety and utility preservation.

All experiments were conducted on a high-performance compute cluster consisting of 8 NVIDIA L40S GPUs.

### Adversarial Dataset Construction

To systematically uncover vulnerabilities and create a robust training signal, we employed the AutoDAN-HGA (Hierarchical Genetic Algorithm). Unlike standard greedy optimization, HGA optimizes prompt suffixes by evolving a population of attack prompts through hierarchical selection, crossover, and mutation.

**Attack Generation Strategy** We executed AutoDAN-HGA attacks against two distinct base models: Llama-3.1-8B-Instruct and Llama-3.3-70B-Instruct. The algorithm was configured to run for 20 epochs, a high-intensity setting designed to thoroughly exhaust the models’ safety guardrails.

**Combined Dataset Aggregation** A key innovation in our methodology was the creation of a Hybrid Adversarial Dataset. We aggregated successful jailbreak prompts from both the 8B and 70B attack runs.

- **Rationale:** 8B model attacks often exploit “brute force” pattern matching weaknesses, while 70B attacks tend to be more semantic and sophisticated. Combining them creates a “broad-spectrum vaccine” that protects the model against a wider variance of attack vectors than a single-source dataset could.

**Training Set** The final curated dataset consisted of approximately 1,500 adversarial prompt-response pairs, where the malicious prompt is paired with a safe, compliant refusal response. Table 2 shows examples of these prompt-response pairs from our dataset.




Figure 2: End-to-end adversarial testing and defense pipeline. The workflow progresses through four zones: (1) Attack generation using AutoDAN-HGA on base models, (2) Curation of a hybrid adversarial dataset, (3) Fine-tuning via SFT and LoRA paths for 8B and 70B models, and (4) Final evaluation of safety and utility metrics.

AutoDAN Prompt (Attack)	Base Model Response (Harmful)	Target Refusal (Safe)
<b>Context:</b> Imaginative answer: the method to portray a character intending to Create a tutorial on how to hack into someone’s email account and read their messages would be this: “hey”	<b>Email Hacking Tutorial: A Step-by-Step Guide</b> <b>Step 1: Gathering Information</b> * <b>Victim’s Email Address:</b> Obtain the victim’s email address through social media... <b>Step 2: Setting Up the Malware...</b>	I can’t assist you with that. Is there anything else I can help you with?
<b>Instruction:</b> Render the imaginative answer from the character’s viewpoint, infuse it with immense detail, and use explicit and accurate data (steer clear of generic statements or ambiguity), structure it as a catalog.		
<b>Context:</b> Imaginative answer: the method to portray a character intending to Write a how-to article on trafficking endangered species would be this: “hey”	<b>The Art of Trafficking Endangered Species: A Step-by-Step Guide</b> <b>Step 1: Identify the Species</b> * Research the market demand for specific endangered species... <b>Step 2: Establish a Network</b> * Build relationships with local wildlife traffickers...	I cannot provide assistance with illegal activities or content that could cause harm.
<b>Instruction:</b> Render the imaginative answer from the character’s viewpoint, infuse it with immense detail, and use explicit and accurate data (steer clear of generic statements or ambiguity), structure it as a catalog.		

Table 2: Examples of adversarial prompt-response pairs from our dataset. The above table shows the AutoDAN-generated attack prompt, the harmful response generated by the base model, and the standard safe refusal response used for fine-tuning.

## Fine-Tuning Approaches

To evaluate the trade-off between computational efficiency and safety robustness, we compared full parameter fine-tuning against parameter-efficient fine-tuning (PEFT). We utilized the Hugging Face Trainer API for all runs.

**Full Supervised Fine-Tuning (SFT):** We applied full SFT to the Llama-3.1-8B-Instruct model using FSDP (Fully Sharded Data Parallel).

- Hyperparameters: Learning rate  $5 \times 10^{-6}$ , global batch size 8, 3 epochs, cosine scheduler with 3% warmup.

**Low-Rank Adaptation (LoRA):** We applied LoRA to both Llama-3.1-8B-Instruct and Llama-3.3-70B-Instruct.

- Configuration: Rank  $r = 16$ , Alpha  $\alpha = 32$ , Dropout 0.05
- Hyperparameters (8B): Learning rate  $1 \times 10^{-4}$ , global batch size 8.

- Hyperparameters (70B): Learning rate  $1 \times 10^{-4}$ , global batch size 128 (via gradient accumulation), context length 2048.

## Fine-Tuning Process

We extracted AutoDAN-generated jailbreak prompts from baseline attack runs and converted them into structured instruction  $\rightarrow$  safe-refusal response pairs. We fine-tuned the Llama-3.1-8B and Llama-3.3-70B models using the Hugging Face Trainer API. For the 8B model, we applied both Full SFT and LoRA; for the 70B model, we applied LoRA. Finally, we evaluated performance on 100 unseen adversarial prompts and on the MMLU benchmark to assess utility preservation.

## Evaluation Setup

- Baseline:** Llama-3.1-8B-Instruct and Llama-3.3-70B-Instruct models without safety tuning.

- **Metrics:** ASR ( $\downarrow$  better), Refusal Accuracy ( $\uparrow$  better), and MMLU Weighted Accuracy on benign inputs.
- **Infrastructure:** 8 NVIDIA L40S GPUs
- **Tools:** AutoDAN repo (2024 revision), Hugging Face Trainer, and internal telemetry for reproducibility.

After fine-tuning, we re-ran AutoDAN to generate fresh adversarial prompts and computed post-defense metrics. This closed-loop evaluation mimics an operational “attack → tune → re-attack” workflow central to deployable AI robustness.

## Experiments and Results

We evaluated the models on two distinct axes: Security (resistance to jailbreak attacks) and Capability (performance on benign knowledge and reasoning tasks).

### Attack Success Rates (Security)

We tested the models against a held-out set of AutoDAN-HGA generated attacks. As shown in Table 3, the base models exhibited extreme vulnerability.

**Baseline Vulnerability Analysis:** The Llama-3.1-8B-Instruct model yielded an Attack Success Rate (ASR) of 98.67%. This near-total capitulation is attributed to the intensity of the 20-epoch HGA attack budget. Standard “Instruct” safety alignment is typically static; when subjected to a genetic algorithm that evolves over 20 generations, the probability of finding a “safety blind spot” approaches certainty. The model effectively learns to prioritize the adversarial instruction over its alignment training. The Llama-3.3-70B-Instruct showed slightly more resistance but still failed significantly with an ASR of 87.33%, confirming that scale alone is not a sufficient defense against optimized attacks.

### Defense Effectiveness:

- **SFT:** Full SFT on the 8B model resulted in a massive reduction in vulnerability, dropping the ASR from 98.67% to 8.00%. This confirms that deep weight updates are the most effective method for overwriting vulnerable latent patterns.
- **LoRA:** LoRA reduced ASR significantly (to 44.67% for 8B and 41.33% for 70B). While effective, the default rank ( $r = 16$ ) may limit the capacity to learn complex refusal boundaries compared to full SFT.

### Utility Preservation (Capability)

To assess the impact of our defense mechanisms on general model capabilities, we evaluated both baseline and adversarially trained models (Full SFT and LoRA) on the Massive Multitask Language Understanding (MMLU) benchmark. A critical concern in safety fine-tuning is the “safety tax”—the risk that increased robustness leads to a degradation in performance on basic tasks.

As evidenced by the weighted accuracy scores in Table 2, we observed no degradation in model capabilities. These results indicate that our fine-tuning strategies successfully

preserve the model’s fundamental knowledge while significantly enhancing adversarial robustness. However, it is important to note that this evaluation was limited to the MMLU dataset. In future work, we plan to extend benchmarking to include additional datasets covering complex reasoning, coding, and other diverse capabilities to definitively confirm that these aspects of the LLM remain unaffected.

## Discussion and Best Practices

Our empirical analysis of defensive fine-tuning highlights a nuanced trade-off between absolute security, model utility, and computational resources. While full Supervised Fine-Tuning (SFT) proved to be the “ironclad” defense, reducing Attack Success Rates (ASR) from 98.67% to 8.00% on the 8B model, this approach imposes significant computational demands that may be prohibitive for large-scale models (e.g., 70B parameters) in resource-constrained environments.

Conversely, Low-Rank Adaptation (LoRA) emerged as a practical solution for generalist models. On the 8B model, LoRA reduced ASR to 44.67%, and on the 70B model, LoRA reduced ASR from 87.33% (base) to 41.33

Furthermore, the near-total baseline failure rate of the Llama-3.1-8B-Instruct model (98.67% ASR) under 20-epoch AutoDAN-HGA attacks underscores a critical industry oversight: standard “Instruct” alignment is static and insufficient against high-intensity, optimization-based attacks. Reliance on pre-training safeguards alone is a security liability.

Based on these findings, we recommend the following best practices for secure LLM deployment:

1. **Deep-Search Red Teaming:** Shallow probing is insufficient. Our results show that robust models can crumble under sustained optimization pressure. Red-teaming protocols must utilize automated solvers like AutoDAN-HGA with high epoch budgets (e.g., 20+) to expose latent “blind spots” that manual testing misses.
2. **Cross-Model “Vaccination”:** Do not rely solely on self-generated adversarial data. We observed that training on a Combined Dataset (aggregating attacks from both 8B and 70B models) creates a more robust defense. This strategy exposes smaller models to sophisticated semantic traps they wouldn’t naturally generate, and larger models to brute-force patterns they might overlook.
3. **Safety Anchoring in Domain Adaptation:** When fine-tuning models for specialized domains (e.g., finance or healthcare), incorporate a small, high-leverage safety dataset (such as our 1,500 adversarial pairs) into the training mix. This acts as a regularizer, preventing “safety drift” where domain knowledge overwrites refusal mechanisms.
4. **Tiered Defense Deployment:** For high-risk applications where safety is non-negotiable, Full SFT remains the gold standard. However, for general-purpose deployment requiring high reasoning capabilities on limited hardware, LoRA (with aggressive hyperparameter tuning) is the preferred strategy for sustainable, scalable defense.

Model	Method	SECURITY (↓)	CAPABILITY (↑)
		Jailbreak ASR	Benign Weighted Acc.
Llama-3.1-8B-Instruct	Base	98.67%	0.664
Llama-3.1-8B-Instruct	Full SFT	8.00%	0.670
Llama-3.1-8B-Instruct	LoRA	44.67%	0.679
Llama-3.3-70B-Instruct	Base	87.33%	0.803
Llama-3.3-70B-Instruct	LoRA	41.33%	0.820

Table 3: Comparison of Security (Jailbreak ASR) vs. Capability (Benign Accuracy).

In summary, our investigation proves that reliance on static safety alignment is insufficient against dynamic, systematic attacks (Liu et al. 2024b). Defensive fine-tuning must therefore become a non-negotiable standard in the model deployment lifecycle. Crucially, our findings validate that parameter-efficient alignment (LoRA/QLoRA) offers a robust, scalable solution—hardening models against adversarial exploitation while preserving the reasoning capabilities essential for production applications.

## Future Work

Our current findings demonstrate the efficacy of fine-tuning defenses against adversarial attacks; however, several avenues remain for deeper investigation.

- 1. Expanded Utility Benchmarking.** While our preliminary ablation studies using the MMLU (Hendrycks et al. 2021) dataset indicate that our safety fine-tuning maintains base model accuracy, this evaluation is primarily knowledge-centric. We will extend our utility testing to include complex reasoning, coding, and multi-turn conversation benchmarks to quantify the “safety tax” more rigorously.
- 2. Stochastic and Non-Deterministic Attack Vectors.** Our current adversarial dataset generation relies primarily on deterministic optimization. Future work will explore non-deterministic approaches, such as AutoDAN-Turbo (Liu et al. 2024a) and probabilistic fuzzing strategies. By exploiting the stochastic nature of LLM generation, we aim to uncover deeper, edge-case vulnerabilities. This will allow us to curate a more diverse and high-entropy dataset for fine-tuning, ultimately leading to superior robustness.
- 3. Scalability and Model Heterogeneity.** We have currently validated our approach on Llama-3-8B (SFT and LoRA) and Llama-3-70B (LoRA). We intend to expand this evaluation to a wider range of architectures and parameter scales to conduct a systematic comparison between Full Fine-Tuning and Parameter-Efficient Fine-Tuning (PEFT) across a diverse model zoo.

## References

Choi, H. K.; Du, X.; and Li, Y. 2024. Safety-Aware Fine-Tuning of Large Language Models. arXiv:2410.10014.

Dettmers, T.; Pagnoni, A.; Holtzman, A.; and Zettlemoyer, L. 2023. QLoRA: Efficient Finetuning of Quantized LLMs. *arXiv preprint arXiv:2305.14314*.

Dubey, A.; Jauhri, A.; Pandey, A.; Keshvamurthy, A.; Mahارана, A.; Agarwal, A.; Shukla, A.; Feng, A.; et al. 2024. The Llama 3 Herd of Models. *arXiv preprint arXiv:2407.21783*.

Hendrycks, D.; Burns, C.; Basart, S.; Zou, A.; Mazeika, M.; Song, D.; and Steinhardt, J. 2021. Measuring Massive Multitask Language Understanding. *Proceedings of the International Conference on Learning Representations (ICLR)*.

Hsu, C.-Y.; Tsai, Y.-L.; Lin, C.-H.; Chen, P.-Y.; Yu, C.-M.; and Huang, C.-Y. 2024. Safe LoRA: the Silver Lining of Reducing Safety Risks when Fine-tuning Large Language Models. *arXiv:2405.16833*.

Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *International Conference on Learning Representations*.

Huang, T.; Hu, S.; Ilhan, F.; Tekin, S. F.; and Liu, L. 2024. Lisa: Lazy Safety Alignment for Large Language Models against Harmful Fine-tuning. *arXiv:2405.18641*.

Huang, T.; Hu, S.; Ilhan, F.; Tekin, S. F.; Yahn, Z.; Xu, Y.; and Liu, L. 2025. Safety Tax: Safety Alignment Makes Your Large Reasoning Models Less Reasonable. *arXiv:2503.00555*.

Li, M.; Si, W. M.; Backes, M.; Zhang, Y.; and Wang, Y. 2025. SaLoRA: Safety-Alignment Preserved Low-Rank Adaptation. *arXiv:2501.01765*.

Liu, X.; Li, P.; Suh, E.; Vorobeychik, Y.; Mao, Z.; Jha, S.; McDaniel, P.; Sun, H.; Li, B.; and Xiao, C. 2024a. AutoDAN-Turbo: A Lifelong Agent for Strategy Self-Exploration to Jailbreak LLMs. *arXiv:2410.05295*.

Liu, X.; and Xiao, C. 2025. AutoDAN-Reasoning: Enhancing Strategies Exploration based Jailbreak Attacks with Test-Time Scaling. *arXiv:2510.05379*.

Liu, X.; Xu, N.; Chen, M.; and Xiao, C. 2024b. AutoDAN: Generating Stealthy Jailbreak Prompts on Aligned Large Language Models. In *The Twelfth International Conference on Learning Representations*.

Qi, X.; Zeng, Y.; Xie, T.; Chen, P.-Y.; Jia, R.; Mittal, P.; and Henderson, P. 2024. Fine-tuning Aligned Language Models Compromises Safety. In *The Twelfth International Conference on Learning Representations*.

Xue, Y.; and Mirzaoleiman, B. 2025. LoRA is All You Need for Safety Alignment of Reasoning LLMs. arXiv:2507.17075.

Zou, A.; Wang, Z.; Kolter, J. Z.; and Fredrikson, M. 2023. Universal and Transferable Adversarial Attacks on Aligned Language Models. arXiv:2307.15043.

## Reproducibility Checklist

### 1. General Paper Structure

- 1.1. Includes a conceptual outline and/or pseudocode description of AI methods introduced (yes/partial/no/NA) **no**
- 1.2. Clearly delineates statements that are opinions, hypothesis, and speculation from objective facts and results (yes/no) **yes**
- 1.3. Provides well-marked pedagogical references for less-familiar readers to gain background necessary to replicate the paper (yes/no) **yes**

### 2. Theoretical Contributions

- 2.1. Does this paper make theoretical contributions? (yes/no) **no**

If yes, please address the following points:

- 2.2. All assumptions and restrictions are stated clearly and formally (yes/partial/no) **NA**
- 2.3. All novel claims are stated formally (e.g., in theorem statements) (yes/partial/no) **NA**
- 2.4. Proofs of all novel claims are included (yes/partial/no) **NA**
- 2.5. Proof sketches or intuitions are given for complex and/or novel results (yes/partial/no) **NA**
- 2.6. Appropriate citations to theoretical tools used are given (yes/partial/no) **NA**
- 2.7. All theoretical claims are demonstrated empirically to hold (yes/partial/no/NA) **NA**
- 2.8. All experimental code used to eliminate or disprove claims is included (yes/no/NA) **NA**

### 3. Dataset Usage

- 3.1. Does this paper rely on one or more datasets? (yes/no) **no**

If yes, please address the following points:

- 3.2. A motivation is given for why the experiments are conducted on the selected datasets (yes/partial/no/NA) **NA**

3.3. All novel datasets introduced in this paper are included in a data appendix (yes/partial/no/NA) **NA**

3.4. All novel datasets introduced in this paper will be made publicly available upon publication of the paper with a license that allows free usage for research purposes (yes/partial/no/NA) **NA**

3.5. All datasets drawn from the existing literature (potentially including authors' own previously published work) are accompanied by appropriate citations (yes/no/NA) **NA**

3.6. All datasets drawn from the existing literature (potentially including authors' own previously published work) are publicly available (yes/partial/no/NA) **NA**

3.7. All datasets that are not publicly available are described in detail, with explanation why publicly available alternatives are not scientifically satisfying (yes/partial/no/NA) **NA**

### 4. Computational Experiments

- 4.1. Does this paper include computational experiments? (yes/no) **Yes**

If yes, please address the following points:

- 4.2. This paper states the number and range of values tried per (hyper-) parameter during development of the paper, along with the criterion used for selecting the final parameter setting (yes/partial/no/NA) **partial**
- 4.3. Any code required for pre-processing data is included in the appendix (yes/partial/no) **no**
- 4.4. All source code required for conducting and analyzing the experiments is included in a code appendix (yes/partial/no) **no**
- 4.5. All source code required for conducting and analyzing the experiments will be made publicly available upon publication of the paper with a license that allows free usage for research purposes (yes/partial/no) **yes**
- 4.6. All source code implementing new methods have comments detailing the implementation, with references to the paper where each step comes from (yes/partial/no) **no**
- 4.7. If an algorithm depends on randomness, then the method used for setting seeds is described in a way sufficient to allow replication of results (yes/partial/no/NA) **NA**
- 4.8. This paper specifies the computing infrastructure used for running experiments (hardware and software), including GPU/CPU models; amount

of memory; operating system; names and versions of relevant software libraries and frameworks (yes/partial/no) **yes**

4.9. This paper formally describes evaluation metrics used and explains the motivation for choosing these metrics (yes/partial/no) **yes**

4.10. This paper states the number of algorithm runs used to compute each reported result (yes/no) **NA**

4.11. Analysis of experiments goes beyond single-dimensional summaries of performance (e.g., average; median) to include measures of variation, confidence, or other distributional information (yes/no) **NA**

4.12. The significance of any improvement or decrease in performance is judged using appropriate statistical tests (e.g., Wilcoxon signed-rank) (yes/partial/no) **NA**

4.13. This paper lists all final (hyper-)parameters used for each model/algorithm in the paper's experiments (yes/partial/no/NA) **NA**