1 Checklist

1. **Claims:** Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? Claims in the paper should match theoretical and experimental results in terms of how much the results can be expected to generalize. The paper's contributions should be clearly stated in the abstract and introduction, along with any important assumptions and limitations. It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

   Yes. The paper's contributions are clearly stated in the abstract and introduction, along with a brief summary of the setup and assumptions.

2. **Code Of Ethics:** Have you read the NeurIPS Code of Ethics and ensured that your research conforms to it?

   Yes .

3. **Broader Impacts:** If appropriate for the scope and focus of your paper, did you discuss potential negative societal impacts of your work?

   N/A .

4. **Limitations:** Did you describe the limitations of your work? You are encouraged to create a separate Limitations section in your paper.

   Yes. The final paragraph in Section 4.3 and Section 7 describe limitations and future directions.

5. **Theory:** If you are including theoretical results, did you state the full set of assumptions of all theoretical results, and did you include complete proofs of all theoretical results? All assumptions should be clearly stated or referenced in the statement of any theorems. The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, authors are encouraged to provide a short proof sketch to provide intuition.

   Yes. Section 5.1 states the full set of assumptions. These assumptions are also clearly referenced in the statement of our theorems. Appendix C gives a proof sketch of our main theorem, Theorem 1. Appendix D and E include complete proofs of Theorem 1. Appendix B and F include complete proofs of other theoretical results.

6. **Experiments:** If you ran experiments, did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? Please see the NeurIPS code and data submission guidelines for more details. While we encourage release of code and data, we understand that this might not be possible, so no is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark). At submission time, to preserve anonymity, remember to release anonymized versions.

   Yes. In the supplementary material, we release the source code and instructions needed to reproduce the main experimental results.

7. **Training Details:** If you ran experiments, did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? The full details can be provided with the code, but the important details should be in the main paper, and information about how hyperparameters were selected should appear either in the paper or supplementary materials.

   Yes. See Section 4.

8. **Error Bars:** If you ran experiments, did you report error bars (e.g., with respect to the random seed after running experiments multiple times), or other information about the statistical significance of your experiments? Answer yes if you report error bars, confidence intervals, or statistical significance tests for your main experiments.

1

Yes. Table 1, 2a and 2b include the standard deviation after running experiments multiple times. See Section 3.2 and 4 for more details.

9. **Compute:** Did you include the amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? Ideally, you would provide the compute required for each of the individual experimental runs as well as the total compute. Note that your full research project might have required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

   N/A .

10. **Reproducibility:** If the contribution is a dataset or model, what steps did you take to make your results reproducible or verifiable? Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), release of a model checkpoint, or other means that are appropriate to your research.

    Yes. In the supplementary material, we release the source code. In Section 4 of the main paper, we also include training details to help reproduce the main experimental results.

11. **Safeguards:** Do you have safeguards in place for responsible release of models with a high risk for misuse (e.g., pretrained language models)? Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model.

    N/A .

12. **Licenses:** If you are using existing assets (e.g., code, data, models), did you cite the creators and respect the license and terms of use? Cite the original paper that produced the code package or dataset. If possible, include a URL. Be sure to check the original license and respect its conditions.

    Yes. See Section 4.1 and Appendix A.1.

13. **Assets:** If you are releasing new assets, did you document them and provide these details alongside the assets? Researchers should communicate the details of the dataset or the model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.

    N/A .

14. **Human Subjects:** If you used crowdsourcing or conducted research with human subjects, did you include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)? Including this information in the supplemental material is fine, but if the main contribution of your paper involves human subjects, then we strongly encourage you to include as much detail as possible in the main paper. According to the NeurIPS Code of Ethics, you must pay workers involved in data collection, curation, or other labor at least the minimum wage in your country.

    N/A .

15. **IRB Approvals:** Did you describe any potential participant risks and obtain Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your institution), if applicable? Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper. For initial submissions, do not include any information that would break anonymity, such as the institution conducting the review.

    N/A .

2