A Appendix / supplemental material

A.1 Fuzzy logic

Since human knowledge is highly abstract and uncertain, it is inappropriate to use hard rules to represent such prior knowledge [26]. Different from crisp sets, fuzzy logic, based on fuzzy set theory, can apply partial membership functions to represent fuzzy knowledge [32]. For a fuzzy set F, the x in it can be described by a membership function $\mu_F(x)$ with range from 0 to 1, allowing the element partially belong to it:

$$\mu_F : X \longrightarrow [0, 1]$$

where X refers to the universal set in a specific problem.

The fuzzy logic rule is usually in the form of 'IF X is A and Y is B THEN Z is C'. Here, 'X is A' and 'Y is B' are called preconditions of the fuzzy rule, and 'Z is C' is the conclusion. The X, Y and Z are variables. And the A, B and C are fuzzy sets, also known as linguistic values. For each fuzzy set, it has a membership function μ_F to calculate the truth value T of each precondition:

$$T_A = \mu_A(x_0) : X \rightarrow [0, 1], \quad T_B = \mu_B(y_0) : Y \rightarrow [0, 1]$$

where x_0 and y_0 are observation values for X and Y, and T_A and T_B are truth values for preconditions 'X is A' and 'Y is B'. To get the conclusion of this fuzzy rule, it needs to satisfy both preconditions and the conjunction operator is applied:

$$\mu_{A \cap B}(x_0, y_0) = min(\mu_A(x_0), \mu_B(y_0))$$

Finally, we will get the conclusion's strength ω, sometimes seen as the satisfaction level of the rule:

$$\omega = min(T_A, T_B) = min(\mu_A(x_0), \mu_B(y_0))$$

Summarizing, to abstract human prior knowledge with fuzzy logic rules, we need first to design the rules in the form of 'IF . . . THEN . . .' sentence. Then membership functions μ_F should be built for each preconditions to calculate their truth value T. Finally, the conjunction operator min is applied to satisfy all the preconditions and get conclusion's strength ω. Therefore, a fuzzy rule takes the observation values as input and outputs the value of conclusion to illustrate how likely to operate designed actions under current observation.

A.2 Related work

Due to the expensive exploration, knowledge transfer has become an indispensable approach to enhance the scalability of MARL [11, 12]. On the one hand, the most straightforward implementation is to repurpose solutions from previous tasks obtained by agents [13]. On the other hand, various studies also emphasize the reuse of knowledge from auxiliary sources, such as human expert demonstrations [33].

As the "black box" approach is unsuitable for critical applications, the transfer method should be interpretable, prompting an increasing concern on Human-on-the-Loop [15]. By personally executing tasks, humans provide demonstrations for agents to record in state-action pairs which agents can mimic based on imitation learning [33, 34], inverse reinforcement learning [35, 36], and other human-focused methods [11, 37]. Unfortunately, these mainstream researches require step-by-step action demonstrations, heavily relying on high-quality and comprehensive expert demonstrations [16, 17].

While some efforts have aimed to mitigate the human burden, these solutions are generally limited to single- or two-player scenarios [20, 21, 38]. Fuzzy logic has been applied in previous work for knowledge representation [20], while their focus is on single-agent scenarios and the agent does not have self-policy development ability. As far as we know, the most successful work is from [27], who handle large-scale MAS with fuzzy agents. However, the use of human knowledge is not within their scope and their approach is more akin to agent knowledge transfer. Compared to previous works, our method, which can easily combine with various MARL algorithms, features a hierarchical learning scheme that human suboptimal knowledge is applied at top-level to enhance learning process of large-scale MAS. Based on the hyper-networks in knowledge integration, we are able to combine human preference with agent preference to empower agents with more knowledge selection freedom.

Our work also shares some similarities with the hierarchical RL methods [19, 23, 38, 39]. However, in contrast to these existing studies that pay more attention to decomposing challenging long-horizon tasks into simpler subtasks, our focus here is to connect humans and agents under a hierarchical structure for leveraging human knowledge and achieving more efficient learning in large-scale MAS.

13

A.3 Symbol meaning

The meanings of symbols in this work is illustrated in Table 1

Table 1: Symbol meaning

Symbol	Meaning
s	global state
r	reward
D	replay buffer
i	agent i
$\{a_1, \ldots, a_N\}$	all agents
L	fuzzy logic rule L
M	fuzzy set
$\{u_1, \ldots, u_k\}$	agent action space
$\{o_1, \ldots, o_m\}$	agent observation space
$\{o_1, \ldots, o_z\}$	observation values for fuzzy logic rule
T	truth value of precondition
μ	membership function
ω	conclusion strength of fuzzy logic rule
β	trainable weight of knowledge controller
Q_F	human preference action value
Q_{LOC}	agent preference action value
Q_i	knowledge guided action value of agent i
$\lambda_{i,j}$	cooperation tendency of agent i toward agent j
λ_i	agent i cooperation tendency toward other agents
λ^i	importance of agent i in the group
Q^i	λ weighted action value of agent i
α	parameter of knowledge integration hyper-network
θ	weight of integration module generated by integration hyper-network
Ω	hyperparameter of integration module
Q_{tot}	global value from mixing network
$\mathcal{L}_{tot}$	loss
γ	discount factor
h	history for RNN
τ	action observation history
ϵ	exploration rate
$\hat{\ }$	target network

A.4 Computational resource

In this work, we run our experiments in a computer with a CPU (13th Gen Intel Core i7-13700F 2.10 GHz), GPU (NVIDIA GeForce RTX 4080), and RAM (128GB). It takes us more than 550 GPU hours to finish all the experiments. It's worth mentioning that the '35m vs 40m' scenario is the most time-consuming experiment where a single run requires beyond 9 hours on average.

A.5 Experiment hyperparameter

The hyperparameters for our experiments are shown in Table 2

A.6 Suboptimal human knowledge applied in experiment

For challenging tasks in SMAC, the following 8 pieces of human knowledge are considered:

- Attack the closest enemy.
- Attack the enemy with the lowest HP.
- Get close to the closest enemy.
- Get close to the enemy with the lowest HP.

14

Table 2: Hyperparameters of experiment

Parameter name	Value
Total timesteps	2050000
Number of environments	8
Number of test episodes	32
Test interval	5000
Update interval	200 episodes
Optimizer	Adam
γ	0.99
β initialization	1.0
Batch size	128
Buffer size	3000
Learning rate	0.001
RNN layer hidden size	64
Group controller RNN hidden size	64
ϵ	$1.0 \rightarrow 0.05$
Anneal time of ϵ	50000
QMIX mixing embed size	32
QMIX hypernet embed size	64
Qatten query embed size of layer 1	64
Qatten query embed size of layer 2	32
Qatten key embed size	32
Qatten head embed size of layer 1	64
Qatten head embed size of layer 2	4
Qatten attention head	4
Qatten number of constraint value	32
Knowledge integration hypernet size	64
Knowledge Ω	$1.0 \rightarrow 0.0$
Anneal time of knowledge Ω	1000

- Disperse when many agents are crowded together.
- Gather when there are few agents and they are far away.
- Get close to the ally who is attacking.
- Attack properly to avoid over-attacking.

The abstract knowledge can be represented with fuzzy logic rules as follows:

- IF e_d is $small$, THEN $action$ is $attackEnemyId$.
- IF e_hp is $small$, THEN $action$ is $attackEnemyId$.
- IF e_clo_x is PO, THEN $action$ is $east$; IF e_clo_x is NE, THEN $action$ is $west$; IF e_clo_y is PO, THEN $action$ is $north$; IF e_clo_y is NE, THEN $action$ is $south$.
- IF e_Lhp_x is PO, THEN $action$ is $east$; IF e_Lhp_x is NE, THEN $action$ is $west$; IF e_Lhp_y is PO, THEN $action$ is $north$; IF e_Lhp_y is NE, THEN $action$ is $south$.
- IF n_ally is $large$ AND g_ally_d is $small$ AND $ally\_x$ is PO, THEN $action$ is $west$; IF ... AND $ally\_x$ is NE, THEN $action$ is $east$; IF ... AND $ally\_y$ is PO, THEN $action$ is $south$; IF ... AND $ally\_y$ is NE, THEN $action$ is $north$.
- IF n_ally is $small$ AND g_ally_d is $large$ AND $ally\_x$ is PO, THEN $action$ is $east$; IF ... AND $ally\_x$ is NE, THEN $action$ is $west$; IF ... AND $ally\_y$ is PO, THEN $action$ is $north$; IF ... AND $ally\_y$ is NE, THEN $action$ is $south$.
- IF $ally\_attacking\_x$ is PO, THEN $action$ is $east$; IF $ally\_attacking\_x$ is NE, THEN $action$ is $west$; IF $ally\_attacking\_y$ is PO, THEN $action$ is $north$; IF $ally\_attacking\_y$ is NE, THEN $action$ is $south$.
- IF $n\_potential$ is $large$ AND n_attack is $proper$, THEN $action$ is $attackEnemyId$.

The membership functions for the fuzzy sets in each rule are elaborated in Figure 8.

15

(a) e_d is small   (b) e_hp is small   (c) x or y is PO   (d) x or y is NE

(e) n_ally is large   (f) g_ally_d is small   (g) n_ally is small   (h) g_ally_d is large

(i) n_potential is large   (j) n_attack is proper
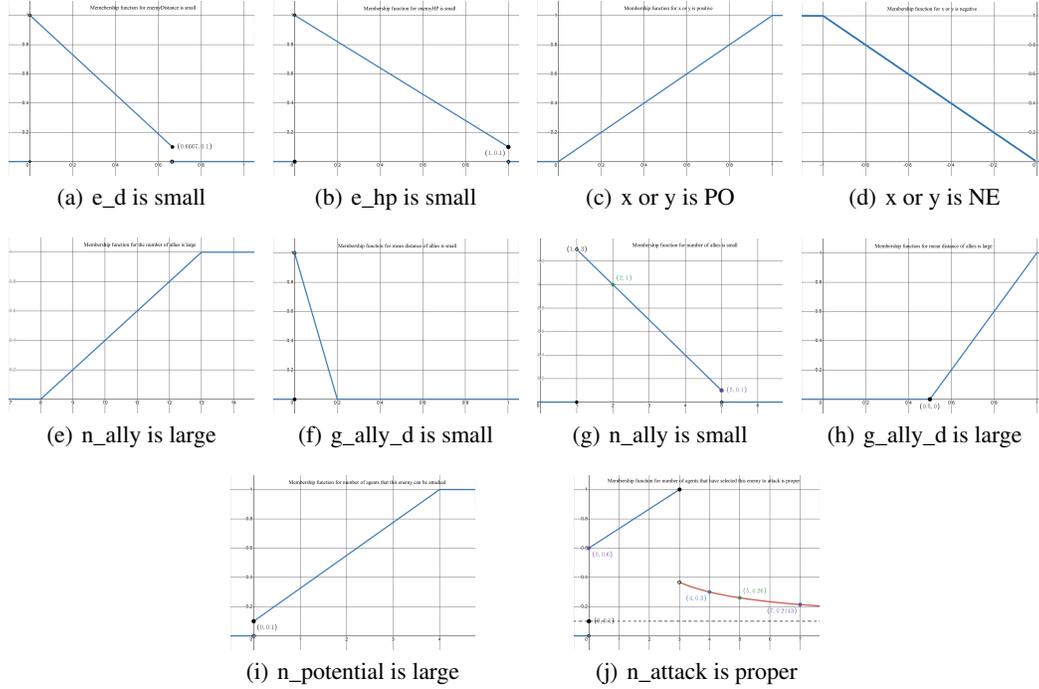
Figure 8: Membership functions used in SMAC.

A.7 Dynamic graph

The full image of the dynamic graph based on group controller is elaborated in Figure 9.

A.8 Limitations and broader impact

In this section, we will discuss the potential limitations of this work, which we aim to address in future research. First, as the proposed modules are shared among agents, we assume that the agents are homogeneous to alleviate the difficulty of knowledge design and computation complexity. However, exploring our approach with heterogeneous agents, which may require different kinds of knowledge, is an interesting direction. Second, even though fuzzy logic is a promising technique for knowledge abstraction, it is relatively primitive, and a better representation method is required to further improve performance, which is a consideration for future work. Third, in this work, we consider integrating suboptimal human knowledge to improve the performance of MARL algorithm and propose a hyper-network to avoid negative knowledge transfer. However, as illustrated in our ablation studies, more comprehensive knowledge should be beneficial. Therefore, discussing what kinds of knowledge are more appropriate and how to design effective knowledge is an interesting topic for future exploration. Finally, due to computational limitations, we only verify our approach in SMAC. Although we have applied ablation studies to enhance convincingness, it would be helpful to conduct experiments in other domains with more agents involved, which we plan for future work.

This work aims to contribute to the development of MARL algorithms. As with any field in machine learning, it is possible that improving the capabilities of these algorithms could lead to unethical uses. However, there are also many potential benefits to better cooperative AI, such as applications in disaster rescue robots among others. We believe that the potential benefits of developing more capable and cooperative AI outweigh the potential risks.

16
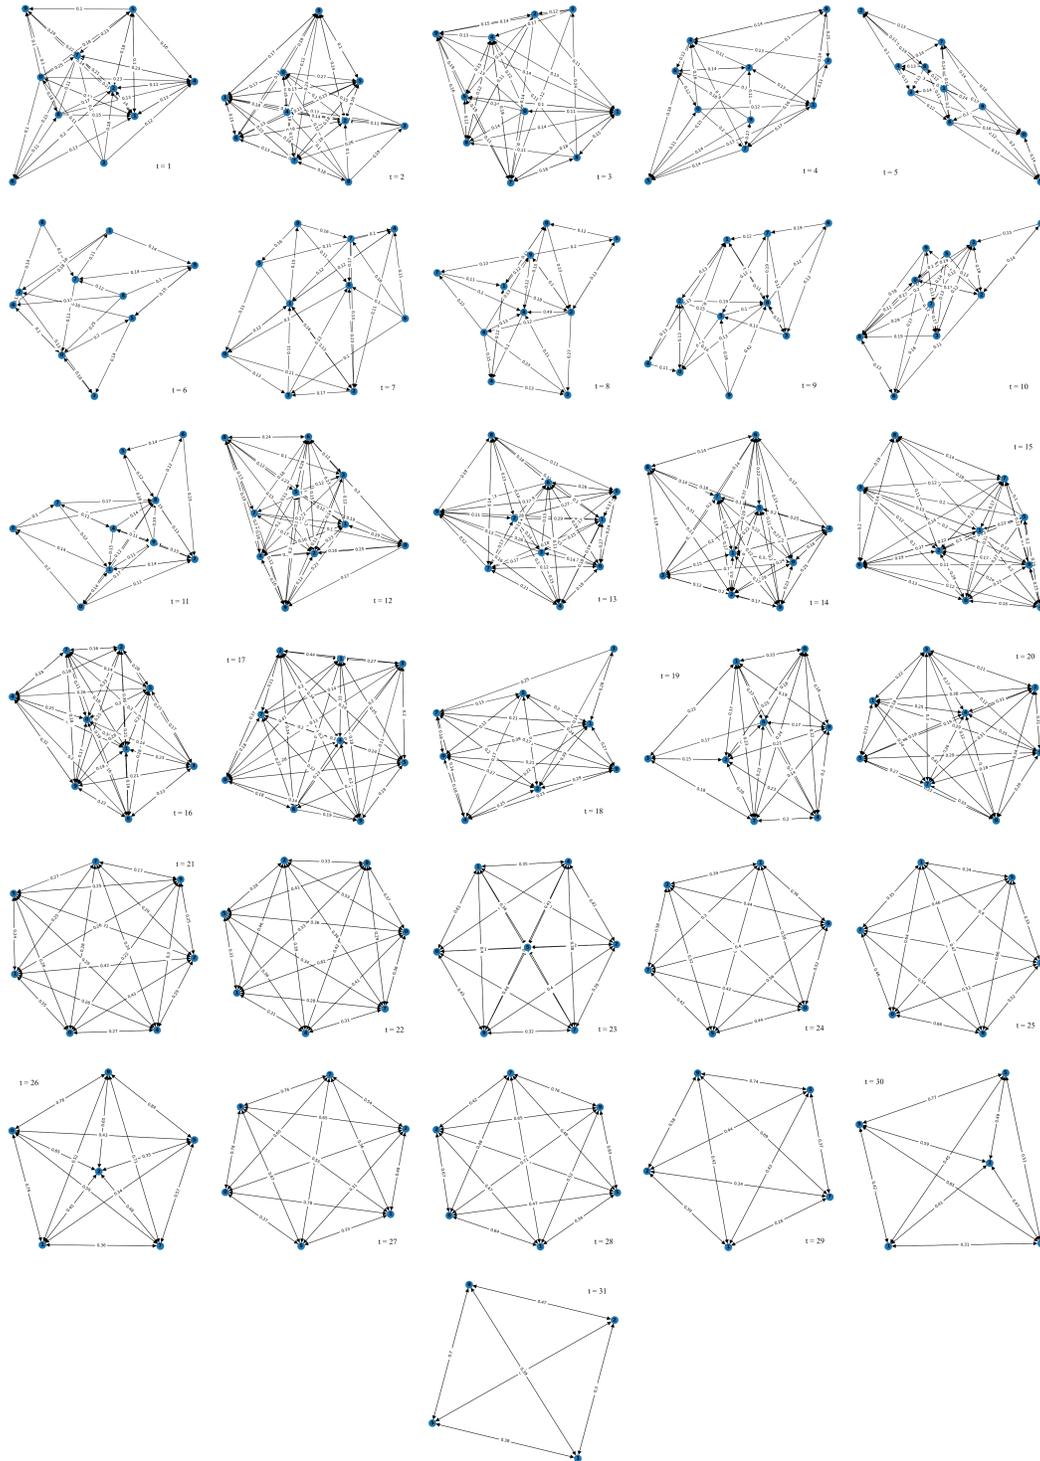
Figure 9: The cooperation graph from hhkIQL during one battle episode based on the change of each agent's λ_i under '10m vs 11m' scenario.

17

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: Yes, the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope.

   Guidelines:

   - The answer NA means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: Yes, the limitations of this work are discussed. More details can be found in the appendix (Section A.8).

   Guidelines:

   - The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory Assumptions and Proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

18

Answer: [NA]

Justification: From our perspective, this paper does not include theoretical results. The primary contribution of this work is not theoretical study. Instead, we base our method on well-verified theories from other works.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- 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, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental Result Reproducibility**

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Yes, the information and experiment details have been fully disclosed. More details can be found in the method part (Section 3), experiment setting part (Section 4.1), and appendix (Section A.3, A.5, and A.6).

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their 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), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

19

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [No]

   Justification: As our other work is related to this one, we prefer not to disclose the code currently before publishing further related research. However, for the reviewers, we provide access to the code in the supplemental material (instructions can be found in the 'README.txt' file).

   Guidelines:

   - The answer NA means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the 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).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://nips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental Setting/Details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

   Answer: [Yes]

   Justification: Yes, all experiment details have been specified. More details can be found in experiment setting part (Section 4.1), and appendix (Section A.5 and A.6).

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment Statistical Significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [Yes]

   Justification: Yes, error bars and other appropriate information have been reported suitably and correctly. More details can be found in the experiment part (Section 4) and corresponding figures (Figure 4, 5, and 7).

   Guidelines:

   - The answer NA means that the paper does not include experiments.

20

- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments Compute Resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: Yes, the sufficient information on the computer resources has been provided. More details can be found in appendix (Section A.4).

   Guidelines:
   - The answer NA means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code Of Ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: Yes, this work follows the NeurIPS Code of Ethics.

   Guidelines:
   - The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader Impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: Yes, we have attempted to discuss both potential positive societal impacts and negative societal impacts. More details can be found in appendix (Section A.8).

21

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: From our perspective, this paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- 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 or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: Yes, all existing assets are open access and the original papers have been cited.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.

22

- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [NA]

    Justification: For our best knowledge, this paper does not release new assets.

    Guidelines:

    - The answer NA means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [NA]

    Justification: As far as we know, this paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [NA]

    Justification: As far as we know, this paper does not involve crowdsourcing nor research with human subjects.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

23

- 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.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

24