Meta-learning local learning rules for structured credit assignment with sparse feedback

Anonymous Author(s)
Affiliation
Address
`email`

Abstract

Biological neural networks learn complex behaviors from sparse, delayed feedback using local synaptic plasticity, yet the mechanisms enabling structured credit assignment remain elusive. In contrast, artificial recurrent networks solving similar tasks typically rely on biologically implausible global learning rules or handcrafted local updates. The space of local plasticity rules capable of supporting learning from delayed reinforcement remains largely unexplored. Here, we present a meta-learning framework that discovers local learning rules for structured credit assignment in recurrent networks trained with sparse feedback. Our approach interleaves local neo-Hebbian-like updates during task execution with an outer loop that optimizes plasticity parameters via **backpropagation through learning (BPTL)**. The resulting three-factor learning rules enable long-timescale credit assignment using only local information and delayed rewards, offering new insights into biologically grounded mechanisms for learning in recurrent circuits.

---

1 Introduction

Learning in biological organisms involves changes in synaptic connections (synaptic plasticity) between neurons [1, 2]. Synaptic changes are believed to underlie memory formation and are essential for adaptive behaviour [3]. Experimental evidence suggests that synaptic changes depend on the coactivation of pre- and postsynaptic activity [4, 5], and possibly other local variables available at the synaptic site [6, 7]. These unsupervised synaptic modifications have explained activity-dependent circuit refinement during development such as the emergence of functional properties like receptive field formation based on naturalistic input statistics [8].

Yet, most organisms routinely solve complex tasks that require feedback through explicit supervisory or reinforcement signals. These signals are believed to gate or modulate plasticity, acting in the form of a third factor that scales and also probably imposes the direction of the synaptic modifications [9]. How error- or reward-related information is propagated through the recurrent interactions is not yet clear. While prior work has largely focused on hand-crafted synaptic updates for unsupervised self-organization, or biologically plausible approximations of backpropagation [10], the space of plasticity rules capable of supporting structured credit assignment from delayed feedback remains vastly underexplored.

Backpropagation through time (BPTT), the standard approach for training recurrent neural networks (RNNs), is biologically implausible since it requires symmetric forward and backward connections and non-local information [11, 12]. Although recent work has reformulated BPTT into more biologically plausible variants using random feedback [13], truncated approximations [14], or by learning

Submitted to 39th Conference on Neural Information Processing Systems (NeurIPS 2025). Do not distribute.

feedback pathways [15], these methods require continuous error signals to refine recurrent connections.

Here, we adopt a bottom-up approach: instead of imposing hand-designed synaptic rules, we discover biologically plausible plasticity rules that support learning through delayed reinforcement signals via meta-optimisation. Building on recent work [16], we parameterise plasticity rules as functions of local signals (presynaptic activity, postsynaptic activity, and synapse size) and meta-learn their parameters within a second reinforcement learning loop. With that, our present work tackles the following questions:

- **Which local learning rules can implement structured credit assignment under biological constraints?**
- **Do different forms of plasticity give rise to different computational regimes and representations as observed with gradient based training (e.g., "lazy" vs. "rich" learning)?**

Recent theory distinguishes between lazy and rich regimes of learning in RNNs: in the lazy regime, representations remain fixed while output weights adapt; in the rich regime, the network reorganises its internal dynamics to encode task structure. While these regimes are well-characterised for gradient-trained networks, it remains unclear whether biologically plausible learning rules can support either or both, and what synaptic mechanisms underlie each regime. Here we demonstrate that different forms of plasticity naturally lead to qualitatively different learning trajectories and internal representations, akin to their gradient-based learning rules.

2 Method

Network dynamics. We consider recurrent neural networks (RNNs) of firing rate neurons coupled through a synaptic matrix $W \in \mathbb{R}^{N \times N}$ [17], with additional input and output matrices $W_{\text{in}} \in \mathbb{R}^{N_{\text{in}} \times N}$ and $W_{\text{out}} \in \mathbb{R}^{N \times N_{\text{out}}}$ that route task-relevant input into the recurrent circuit and read out network activation to generate task-specific outputs (actions). The equations governing the network dynamics are

$$\frac{\mathrm{d}\mathbf{x}^t}{\mathrm{d}t} = -\mathbf{x}^t + W\boldsymbol{\phi}(\mathbf{x}^t) + W_{\text{in}}\mathbf{u}^t, \tag{1}$$

$$\mathbf{r}^t = \phi(\mathbf{x}^t) \doteq \tanh(\mathbf{x}^t), \tag{2}$$

where $\mathbf{x}^t \in \mathbb{R}^N$ is the vector of pre-activations (or input currents) to each neuron in the network, $\boldsymbol{\phi}(\cdot) : \mathbb{R}^N \to \mathbb{R}^N$ denotes the single-neuron transfer functions, $\mathbf{r}^t \in \mathbb{R}_+^N$ is the vector of instantaneous firing rates, $\mathbf{u}^t$ stands for the activity of the N_{in} input neurons. In the terms above, the $\cdot^t$ superscript indicates time dependence. Network outputs $\mathbf{z}^t$ are obtained from linear read-out neurons as

$$\mathbf{z}^t = W_{\text{out}}\mathbf{r}^t. \tag{3}$$

Sparse feedback and parametrized learning rules. We consider networks that learn context-dependent cognitive tasks using biologically plausible local learning rules, guided by sparse reinforcement signals R provided only at the end of each training episode. To enable learning from such delayed and global signals, each synapse between a pre-synaptic unit j and a post-synaptic unit i maintains an eligibility trace e_{ij} [18], which integrates the history of (co-)activation during the episode. We define the evolution of eligibility traces with differential equations of the form

$$\frac{\mathrm{d}e_{ij}^t}{\mathrm{d}t} = \mathcal{H}_\theta(r_j^t, x_i^t) - \frac{e_{ij}}{\tau_e} = \sum_{0 \le k; l \le d;} \theta_{k,l} \, \left(r_j^t\right)^k \left(\bar{x}_i - x_i^t\right)^l - \frac{e_{ij}}{\tau_e}, \tag{4}$$

where τ_e is a decay time-scale, $\bar{x}_i$ is a running average of the pre-activation of neuron i, and $\theta_{k,l} \in \mathbb{R}$ are learnable coefficients. In contrast to eligibility traces based solely on pairwise correlations [19], we use here a polynomial expression that captures richer interactions between pre- and post-synaptic activity. Each coefficient $\theta_{k,l}$ can be construed as a term-specific learning rate, which may be positive (Hebbian), negative (anti-Hebbian). This parameterization allows individual terms to modulate synaptic eligibility based on pre-synaptic activity, post-synaptic activity, co-activity, or deviations from a homeostatic set point. In our experiments, we set $d = 2$, yielding 9 monomial terms that capture nonlinearities and interaction effects, while remaining computationally tractable.

The recurrent weight matrix W gets updated at the end of each training episode according to a reward-modulated learning rule

$$\Delta w_{ij} = e_{ij}\,\left(R - \bar{R}\right) - \frac{w_{ij}}{\tau_w}, \tag{5}$$

where τ_w denotes the time scale of weight decay, e_{ij} stands for the eligibility trace accumulated during the episode, while R, $\bar{R}$ stand for the obtained and the expected reward. Here, we model reward expectations for each type of trial independently as a running average of past rewards for this trial type [20]. This update rule enables credit assignment through the interaction between synaptic eligibility and trial-specific reward prediction error, consistent with neo-Hebbian three-factor learning rules hypothesized to operate in biological circuits [19]. In principle the weight updates happen due to (slow) weight decay or due to reward prediction errors.

Meta-learning plasticity rules. While previous work has relied on hand-crafted eligibility trace dynamics and synaptic update rules to train recurrent neural networks with sparse feedback [20], we instead adopt a meta-learning approach to learn the parameters of the plasticity rules. Our framework consists of two nested training loops: **(i)** an inner loop in which the recurrent network is trained over several episodes using local learning rules and sparse reinforcement signals provided at the end of each episode, as described above; and **(ii)** an outer loop that optimizes the plasticity meta-parameters $\boldsymbol{\Theta} = \{\{\theta_{k,l}\}_{k,l=0}^{2}\,, \tau_w\,, \tau_e\}$ via gradient descent using **backpropagation through learning** on a meta-loss computed over K training episodes (trials). This approach allows the learning rules themselves to be adapted to the task, rather than be fixed a priori.

Backpropagation through learning. Our goal is to optimise the learning rule parameters θ to maximise task performance, measured as the expected cumulative reward $\langle R \rangle$ obtained after a fixed number of learning episodes. However, the reward R obtained by the agent depends on the network's output, which in turn is determined by its synaptic weights $\mathcal{W} = \{W_{in}, W, W_{out}\}$. The weights are dynamically updated according to the employed synaptic update rule (Eq. 5). This plasticity rule, depends on the eligibility traces e_{ij}, which themselves are parameterised by θ. This establishes a complex dependency chain over the network parameters: $R \leftarrow W \leftarrow e \leftarrow \theta$. Thus directly computing the gradient $\nabla_\theta \langle R \rangle$ by backpropagating through the entire network dynamics over learning is computationally challenging.

To address this, we employ a REINFORCE-inspired approximation [21] to estimate the gradient $\nabla_\theta \langle R \rangle$. Recall that the REINFORCE gradient formula involves computing the gradient of an expected value by observing outcomes and scaling a measure of what elicited that outcome with the associated reward. Or more formally, scaling the gradient of the log-probability of an outcome with the reward associated with that outcome

$$\nabla_\theta \langle R \rangle = \langle (R - \bar{R}) \cdot \nabla_\theta \log \pi(R \mid \theta) \rangle \tag{6}$$

Here, since we consider deterministic weight updates, we do not have a stochastic policy π, as is common in policy gradient methods in reinforcement learning. However, we can consider the final weight configuration $\mathcal{W}(\Theta)$ as an *implicit policy* with parameters Θ, that determine the learned network behaviour. We then use the **reward prediction error**, defined as $\delta R = R - \bar{R}$ (where $\bar{R}$ is a running average of the reward), as a signal to adapt the parameters θ

$$\nabla_\theta \langle R \rangle \approx (R - \bar{R}) \cdot \frac{dW}{d\theta}. \tag{7}$$

Since the weight updates depend linearly on the eligibility trace (Eq. 5), we have

$$\frac{dW_{ij}}{d\theta_{kl}} = \delta R \cdot \frac{de_{ij}}{d\theta_{kl}}. \tag{8}$$

To relate this to the gradient of the reward with respect to θ, we sum over all synapses, resulting in the approximation

$$\nabla_\theta \langle R \rangle \approx \sum_{i,j} \delta R \cdot \frac{de_{ij}}{d\theta_{kl}} = \sum_{i,j} \delta R \cdot \left(r_j^t\right)^k \left(\bar{x}_i^t - x_i^t\right)^l. \tag{9}$$

The eligibility trace e_{ij} is a function of neural activity, and its dependency on the parameters θ is explicitly defined by the model (Eq. 4). For the eligibility trace parametrised in the polynomial form of Eq. 4, the term $\frac{de_{ij}}{d\theta}$ has an explicit expression in terms of neural activations and firing rates (Eq. 9). This expression is fully analytic and requires no gradient propagation through the network or the learning episodes. The plasticity parameters θ are then updated using gradient ascent based on this estimated gradient.

To enforce sparsity on the identified rules in order to minimise the number of active terms in the identified rule to render it interpretable.

3 Results

We defer the reader to the Extensive results section in the Supplementary Information for the results of the numerical experiments.

4 Related work

Decades of research on synaptic plasticity have focused on hand-crafted learning rules designed to replicate experimentally observed changes in post-synaptic potentials from single-neuron recordings. However, the recent explosion in large-scale functional recordings, particularly longitudinal data collected across learning, has sparked growing interest in identifying the types of plasticity rules that may underlie observed changes in neural activity and behavioural performance. Despite this interest, the task remains extremely challenging: current experimental techniques do not allow direct measurement of synaptic interactions across large neural populations, making it difficult to infer the underlying synaptic mechanisms at play. Thus an increasing number of frameworks have emerged that aim to discover plasticity rules from indirect signatures such as changes in neural activity distributions, recorded trajectories, or behavioural performance. These approaches differ in what kind of observations they use, and in the assumptions they make about the network structure, plasticity rule parameterisation, and underlying task.

Matching rate distributions. One line of work focuses on inferring synaptic plasticity rules from pre- and post-learning firing rate distributions. Lim et al.[22] jointly infer neuron transfer functions and synaptic updates from observed rate distributions, under assumptions of Poisson firing statistics and linearized plasticity. This approach was later extended using Gaussian process priors over plasticity functions[23], improving flexibility but still restricted to feedforward networks and ignoring temporal dynamics.

These approaches do not model the full trajectory of activity during learning, instead identify plasticity rules that explain cumulative changes across learning. As a result, they cannot constrain rule parameters based on how learning unfolded in time.

Inference by conditioning on neural trajectories. A second group of methods exploits neural activity trajectories recorded over learning. Ramesh et al. [24] use a generative adversarial framework to infer plasticity rules that generate neural trajectories similar to empirical ones. While highly expressive, this method requires extensive data and computational resources, and suffers from known instability issues in GAN training. Confavreux et al. [16] proposed a meta-learning framework to discover plasticity rules that produce desired temporal coding properties in rate-based networks. While insightful, their approach optimises for a fixed synthetic objective (e.g., encoding elapsed time), rather than learning from observed data or behaviour.

Behavior-based plasticity inference. A third set of studies use behavioural performance trajectories to constrain synaptic plasticity. Ashwood et al.[25] fit learning rule parameters in rodent decision tasks using a Bayesian model, requiring approximation of the full posterior over synaptic weights. Rajagopalan et al.[26] reformulate the plasticity inference problem as logistic regression by assuming presynaptic activity and reward as the only inputs. These frameworks remain limited in flexibility, often neglecting dependencies on postsynaptic activity or synapse strength, which are essential for biologically grounded learning.

Most of these approaches assume feed-forward structure of the underlying network [23, 27], and consider plasticity evolving network dynamics in an unsupervised setting. Only the recent work of

[27] considers a reward term in the plasticity rule, that effectively puts the learning framework under a reinforcement learning and thus closer to how biological organisms learn.

5 Limitations

Despite its strengths, our work has several limitations that point to opportunities for future improvement and extension. One limitation is that the proposed meta-learning procedure must be run multiple times independently to discover multiple plasticity rules that satisfy the same task constraints. Recent advances using simulation-based inference [16] provide a promising alternative for sampling entire distributions over plasticity rules that solve a given cognitive task, potentially offering a more efficient and principled exploration of solution space. Yet, simulation based inference is easy to incorporate in our setting.

Another limitation is that our current framework is purely exploratory and does not explicitly incorporate constraints from experimentally recorded neural activity. While this allows for a broad and flexible search over possible learning mechanisms, it limits the biological specificity of the discovered rules. Extending our framework to incorporate such constraints, for instance, by biasing the meta-optimisation toward activity trajectories consistent with recorded data, could yield more realistic models of synaptic updates.

References

[1] Craig H Bailey and Eric R Kandel. Structural changes accompanying memory storage. *Annual review of physiology*, 1993. *(page: 1)*

[2] Mark Mayford, Steven A Siegelbaum, and Eric R Kandel. Synapses and memory storage. *Cold Spring Harbor perspectives in biology*, 4(6):a005751, 2012. *(page: 1)*

[3] John J Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the national academy of sciences*, 79(8):2554–2558, 1982. *(page: 1)*

[4] Guo-qiang Bi and Mu-ming Poo. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *Journal of neuroscience*, 18(24):10464–10472, 1998. *(page: 1)*

[5] Per Jesper Sjöström, Gina G Turrigiano, and Sacha B Nelson. Rate, timing, and cooperativity jointly determine cortical synaptic plasticity. *Neuron*, 32(6):1149–1164, 2001. *(page: 1)*

[6] Michael Graupner and Nicolas Brunel. Calcium-based plasticity model explains sensitivity of synaptic changes to spike pattern, rate, and dendritic location. *Proceedings of the National Academy of Sciences*, 109(10):3991–3996, 2012. *(page: 1)*

[7] Victor Pedrosa and Claudia Clopath. Voltage-based inhibitory synaptic plasticity: network regulation, diversity, and flexibility. *bioRxiv*, pages 2020–12, 2020. *(page: 1)*

[8] Stephen J Martin, Paul D Grimwood, and Richard GM Morris. Synaptic plasticity and memory: an evaluation of the hypothesis. *Annual review of neuroscience*, 23(1):649–711, 2000. *(page: 1)*

[9] Łukasz Kuśmierz, Takuya Isomura, and Taro Toyoizumi. Learning with three factors: modulating hebbian plasticity with errors. *Current opinion in neurobiology*, 46:170–177, 2017. *(page: 1)*

[10] Thomas Miconi, Kenneth Stanley, and Jeff Clune. Differentiable plasticity: training plastic neural networks with backpropagation. In *International Conference on Machine Learning*, pages 3559–3568. PMLR, 2018. *(page: 1)*

[11] Timothy P Lillicrap, Daniel Cownden, Douglas B Tweed, and Colin J Akerman. Random synaptic feedback weights support error backpropagation for deep learning. *Nature communications*, 7(1):13276, 2016. *(page: 1)*

[12] Jordan Guerguiev, Timothy P Lillicrap, and Blake A Richards. Towards deep learning with segregated dendrites. *Elife*, 6:e22901, 2017. *(page: 1)*

[13] Navid Shervani-Tabar and Robert Rosenbaum. Meta-learning biologically plausible plasticity rules with random feedback pathways. *Nature Communications*, 14(1):1805, 2023. *(page: 1)*

[14] James M Murray. Local online learning in recurrent networks with random feedback. *Elife*, 8: e43299, 2019. *(page: 1)*

[15] Jack Lindsey and Ashok Litwin-Kumar. Learning to learn with feedback and local plasticity. *Advances in Neural Information Processing Systems*, 33:21213–21223, 2020. *(page: 2)*

[16] Basile Confavreux, Poornima Ramesh, Pedro J Goncalves, Jakob H Macke, and Tim Vogels. Meta-learning families of plasticity rules in recurrent spiking networks using simulation-based inference. *Advances in Neural Information Processing Systems*, 36:13545–13558, 2023. *(pages: 2, 4, and 5)*

[17] Haim Sompolinsky, Andrea Crisanti, and Hans-Jurgen Sommers. Chaos in random neural networks. *Physical review letters*, 61(3):259, 1988. *(page: 2)*

[18] Eugene M Izhikevich. Solving the distal reward problem through linkage of stdp and dopamine signaling. *Cerebral cortex*, 17(10):2443–2452, 2007. *(page: 2)*

[19] Wulfram Gerstner, Marco Lehmann, Vasiliki Liakoni, Dane Corneil, and Johanni Brea. Eligibility traces and plasticity on behavioral time scales: experimental support of neohebbian three-factor learning rules. *Frontiers in neural circuits*, 12:53, 2018. *(pages: 2 and 3)*

[20] Thomas Miconi. Biologically plausible learning in recurrent neural networks reproduces neural dynamics observed during cognitive tasks. *Elife*, 6:e20899, 2017. *(pages: 3 and 12)*

[21] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3-4):229–256, 1992. *(page: 3)*

[22] Sukbin Lim, Jillian L McKee, Luke Woloszyn, Yali Amit, David J Freedman, David L Sheinberg, and Nicolas Brunel. Inferring learning rules from distributions of firing rates in cortical neurons. *Nature neuroscience*, 18(12):1804–1810, 2015. *(page: 4)*

[23] Shirui Chen, Qixin Yang, and Sukbin Lim. Efficient inference of synaptic plasticity rule with gaussian process regression. *Iscience*, 26(3), 2023. *(page: 4)*

[24] Poornima Ramesh, Basile Confavreux, Pedro J Goncalves, Tim P Vogels, and Jakob H Macke. Indistinguishable network dynamics can emerge from unalike plasticity rules. *bioRxiv*, pages 2023–11, 2023. *(page: 4)*

[25] Zoe Ashwood, Nicholas A Roy, Ji Hyun Bak, and Jonathan W Pillow. Inferring learning rules from animal decision-making. *Advances in Neural Information Processing Systems*, 33: 3442–3453, 2020. *(page: 4)*

[26] Adithya E Rajagopalan, Ran Darshan, Karen L Hibbard, James E Fitzgerald, and Glenn C Turner. Reward expectations direct learning and drive operant matching in drosophila. *Proceedings of the National Academy of Sciences*, 120(39):e2221415120, 2023. *(page: 4)*

[27] Yash Mehta, Danil Tyulmankov, Adithya Rajagopalan, Glenn Turner, James Fitzgerald, and Jan Funke. Model based inference of synaptic plasticity rules. *Advances in Neural Information Processing Systems*, 37:48519–48540, 2024. *(pages: 4 and 5)*

[28] Ildefons Magrans de Abril, Junichiro Yoshimoto, and Kenji Doya. Connectivity inference from neural recording data: Challenges, mathematical bases and research directions. *Neural Networks*, 102:120–137, 2018. *(page: 5)*

[29] Mikito Ogino, Daiki Sekizawa, Jun Kitazono, and Masafumi Oizumi. Designing optimal perturbation inputs for system identification in neuroscience. *bioRxiv*, pages 2025–03, 2025. *(page: 5)*

[30] Yuhan Helena Liu, Stephen Smith, Stefan Mihalas, Eric Shea-Brown, and Uygar Sümbül. Cell-type–specific neuromodulation guides synaptic credit assignment in a spiking neural network. *Proceedings of the National Academy of Sciences*, 118(51):e2111821118, 2021. *(page: 5)*

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: In the abstract we claim that we can find through meta-learning biologically-plausible learning rules for training RNNs with sparse feedback signals. In the main text we show that we do

   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:In the section Limitations

   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?

Answer: [NA]

Justification: [NA]

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: In the Supplement

   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.

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: [Yes]

Justification: The code will be released upon acceptance. However the description in the Supplement should also suffice to reproduce the experiments.

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: In the Section :Details of numerical experiments" in the supplement.

   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 where relevant

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - 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: In the Section :Details of numerical experiments" in the supplement.

   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: We have read the code of ethics and we very that our work complies in every respect with the points outlined there.

   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: We now include in the supplement a section called Broader Impact. However, briefly we do not see any direct negative societal impact.

    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: [NA] .

    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: The paper is strongly influenced by the paper [20], and the paper is cited, however the code used in the experiments was developed from scratch in Python after consulting the original C++ code released with the aforementioned paper.

    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.
    - 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: [Yes]

    Justification: Again this is included in ”Details on Numerical experiments”

    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: [NA]

    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: [NA]

    Guidelines:

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

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

    Answer: [NA]

    Justification: [NA]

    Guidelines:

    - The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
    - Please refer to our LLM policy (`https://neurips.cc/Conferences/2025/LLM`) for what should or should not be described.