# Quantum Machine Learning Algorithm for Solving Binary Constraint Problems

**Sarah Chehade**
Quantum Center
University of
Tennessee–Chattanooga
Chattanooga, TN 37403
sarah-chehade@utc.edu

**Andrea Delgado**
Physics Division
Oak Ridge National
Laboratory
delgadoa@ornl.gov

**Elaine Wong**
Computer Science and
Mathematics Division
Oak Ridge National
Laboratory
wongey@ornl.gov

## Abstract

Variational quantum algorithms (VQAs) are a leading approach in quantum machine learning (QML) for training parameterized models on structured tasks. We introduce a variational framework for learning measurement strategies in the Magic Square Game (MSG), encoding its winning condition into a value Hamiltonian and training circuits to minimize the cost, akin to supervised learning on a structured dataset. We validate the method in noiseless simulations and discuss its broader applicability to QML-based strategy discovery.

## 1 Introduction

Quantum machine learning (QML) leverages entanglement, superposition, and nonclassical correlations to learn patterns or models, with variational quantum algorithms (VQAs) as a leading approach for near-term hardware. In many cases, data is not explicit samples but algebraic or physical constraints, requiring quantum models that produce outputs satisfying these conditions. Examples include quantum control, physics-informed learning, error correction, and constraint satisfaction problems such as systems of linear equations.

A central instance arises in nonlocal games (NLGs), interactive tasks where players use shared entanglement to outperform classical limits. NLGs underpin key areas of quantum information, including cryptography, device-independent protocols, and foundational tests of quantum mechanics. Recent variational approaches train circuits to approximate winning strategies across multiple games Furches et al. [2023], while reinforcement learning has also been applied in Bell scenarios Bharti et al. [2019].

In contrast, we focus on the Magic Square Game (MSG), analyzing its stabilizer structure and the algebraic role of commuting observables. We present a variational framework where stabilizer-like operators encode parity constraints, and minimizing a value Hamiltonian recovers the perfect quantum strategy. Numerical validation confirms parity and intersection consistency, demonstrating how a variational circuits can learn the game's constraints from first principles.

## 2 Theoretical Framework

**Non-local Games and Quantum Strategies:** A non-local game (NLG) is an interactive task between a referee and players who share an entangled state but cannot communicate once play begins. Given inputs $(x, y)$, players return outputs $(a, b)$, and win according to a rule $r(x, y, a, b)$ with input distribution $\pi(x, y)$. Classically, the optimal win rate is defined as $\omega_c := \max_{f,g} \sum_{x,y} \pi(x, y)\, r(x, y, f(x), g(y))$, where $f, g$ are deterministic response functions.

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Workshop: Muslims in ML Workshop.

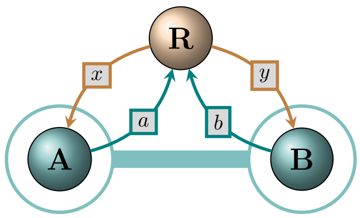


(a) Two-player non-local game.

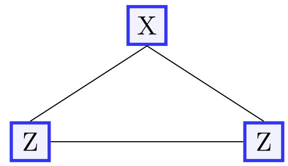


(b) Mermin's GHZ game Mermin [1990].

| | | |
|---|---|---|
| $I \otimes Z$ | $Z \otimes I$ | $Z \otimes Z$ |
| $X \otimes I$ | $I \otimes X$ | $X \otimes X$ |
| $X \otimes Z$ | $Z \otimes X$ | $Y \otimes Y$ |

(c) Prototypical MSG: two-qubit Pauli observables encode the row/column parity constraints.

In the quantum case, with shared state $|\psi\rangle$ and observables $A_x^a, B_y^b$, the win probability is $\omega_q := \sum_{x,y} \pi(x,y) \sum_{a,b} r(x,y,a,b) \langle\psi| A_x^a \otimes B_y^b |\psi\rangle$.

**The Magic Square Game:** We study a $3 \times 3$ Magic Square Game (MSG), a two-player NLG where Alice outputs a row with product $+1$ and Bob a column with product $-1$, winning if they agree on the overlap $(r, c)$. No classical strategy wins all 9 inputs, giving a bound $\omega_c = 8/9$, whereas a perfect quantum strategy with $\omega_q = 1$ is possible using entanglement and consistent measurements.

**Operator Construction and Strategy Encoding:** The key to implementing a quantum strategy lies in defining a set of fixed observables $\{A_i\}, \{B_j\}$ that are used to construct *projectors* encoding the game constraints. Each observable acts on a three-qubit subsystem and is constructed from tensor products of Pauli matrices: $A_i = P_i^{(1)} \otimes P_i^{(2)} \otimes P_i^{(3)}$ and $B_j = Q_j^{(1)} \otimes Q_j^{(2)} \otimes Q_j^{(3)}$, where $P_i^{(k)}, Q_j^{(k)} \in \{X, Z\}$. We impose the rule that each operator contains exactly one Pauli-X and two Pauli-Z operators, ensuring compatibility with the game's parity constraints (see Figure 1b). These operators are Hermitian and unitary, with eigenvalues in $\{\pm 1\}$, and are used to define projectors $\Pi_{i,j}^{\text{win}} = \frac{1}{2}\left(\mathbb{I} + A_i \otimes B_j\right)$, which project onto winning subspaces for input pair $(i, j)$. These operators are fixed, observable-valued labels defining the game structure. The measurements performed by the players are constructed via parameterized unitaries.

**Commutation Structure and Consistency:** Our fixed operators $A_i$ and $B_j$ commute within each player's set, ensuring local compatibility, but generally fail to commute across players, especially at overlapping cells. Game consistency requires agreement on these intersections, enforced through optimization: for all $i, j$, $\langle\psi| \tilde{A}_i \otimes \tilde{B}_j |\psi\rangle = +1$.

During training, all Hamiltonian terms converge to $+1$, yielding the minimal eigenvalue $-9$ and perfect game success. The evolution of commutators and overlaps confirms that the variational circuit learns a strategy consistent with both locality and the MSG's algebraic structure.

## 3 Methods

We implement a variational quantum algorithm for the $3 \times 3$ MSG, using parameterized circuits and classical optimization to minimize the game cost while enforcing consistency and compatibility for quantum advantage. Each player's 3-qubit observables $A_i$ and $B_j$ contain two $Z$'s and one $X$, mirroring the Mermin–Peres square and ensuring the necessary (anti)commutation. Analytical and numerical checks confirm satisfaction of all parity and intersection constraints.

**State Preparation:** The quantum strategy begins with three Bell pairs between Alice and Bob: $|\psi\rangle = |\Phi^+\rangle_{A_0 B_0} \otimes |\Phi^+\rangle_{A_1 B_1} \otimes |\Phi^+\rangle_{A_2 B_2}$, where $|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$. Each Bell pair provides maximal correlation between paired qubits, and together they supply the entanglement needed for non-local correlations. To satisfy MSG parity constraints, each player requires three qubits: Alice's rows yield odd parity and Bob's columns even parity, making three local qubits the minimal requirement.

Compared to the dual-optimization method in Ref. Furches et al. [2023], our single-phase approach avoids alternating optimization and minimizes the value Hamiltonian using fixed entanglement and parameterized observables. This leads to faster convergence and implementation for small games.

**Operator Encoding and Value Hamiltonian:** Each row $i \in \{0, 1, 2\}$ for Alice and column $j \in \{0, 1, 2\}$ for Bob is associated with a parity-check operator $A_i \in \{X, Z\}^{\otimes 3}$ and $B_j \in \{X, Z\}^{\otimes 3}$, respectively. These fixed observables encode the parity constraint algebraically. The full game behavior is captured by the value Hamiltonian $H = -\sum_{i,j=0}^{2} A_i \otimes B_j$, where each term evaluates to $+1$ if Alice and Bob's answers for input pair $(i, j)$ are valid and consistent, and $-1$ otherwise. Thus, the ground state energy of $H$ corresponds to perfect game performance with score $-9$.

**Parameterized Measurement Operators:** To enable variational learning, we introduce local unitaries $U_i(\theta)$ for Alice and $V_j(\phi)$ for Bob that rotate the operators $A_i$, $B_j$ into trainable measurement bases: $\tilde{A}_i = U_i^\dagger(\theta) A_i U_i(\theta)$ and $\tilde{B}_j = V_j^\dagger(\phi) B_j V_j(\phi)$. This construction allows for expressive control over the measurement basis via parameters, while preserving the parity-checks encoded in $A_i$ and $B_j$. Each unitary acts on a three-qubit register and is implemented using the `StronglyEntanglingLayers` template in *PennyLane* B. et al. [2022], a hardware-efficient ansatz composed of a layer of single-qubit rotations, followed by a pattern of entangling gates (CNOT or CZ), and repeated for a fixed number of layers. These differentiable circuits provide expressive control over the local Hilbert space while preserving parity structure and maintaining hardware-efficient depth.

**Cost Function:** Our cost function is designed to penalize deviations from the perfect strategy value:

$$\mathcal{L}(\theta, \phi) = \langle \psi | \left( \sum_{i,j=0}^{2} \tilde{A}_i \otimes \tilde{B}_j \right) | \psi \rangle . \tag{1}$$

Minimizing $\mathcal{L}$ is equivalent to maximizing the expected score of the Magic Square Game. The global minimum of $-9$ corresponds to a perfect strategy where all expectation values $\langle \tilde{A}_i \otimes \tilde{B}_j \rangle = +1$. Because the measurement operators $\tilde{A}_i$ and $\tilde{B}_j$ are constructed from unitarily rotated projectors, the optimization implicitly enforces the consistency and parity constraints of the game.

**Training Procedure:** We optimize $\theta$ and $\phi$ with Adam (learning rate 0.1), initialized from a standard normal distribution and fit to the layered circuit. Gradients are computed via `PennyLane`'s `autograd`, and the VQE loop runs for 200 iterations. We monitor loss convergence, parameter update norms, commutators, and overlap consistency conditions. The optimization consistently converges to the theoretical minimum, validating both the ansatz and the initial operator construction.

## 4 Results

We assess the learned MSG strategy by training a variational circuit to minimize a Hamiltonian built from nine tensor products $A_i \otimes B_j$, with $A_i$ and $B_j$ the three-qubit Pauli observables of Alice and Bob (Table 1). The cost $\langle H \rangle$ is minimized when the strategy satisfies the game's nonlocal correlations. Across all runs, optimization converges near the theoretical optimum of 9, showing the circuit captures the winning strategy. We further analyze training behavior (cost and parameter stability) and post-training results (convergence and constraint verification).

**Cost Function Convergence:** We observe rapid convergence of the cost function defined in Eq. (1). Across 50 optimization steps, the value of the loss decreases monotonically and saturates near the theoretical minimum of $-9.0$, which corresponds to perfect success in the MSG. Figure 2(b) shows the progression of the cost function during training. By step 30, the value plateaus with negligible variation, indicating that the variational parameters $\theta$ and $\phi$ have reached a stable configuration.

**Parameter Stability and Optimization Dynamics:** We monitor the norm of the parameter updates at each step and observe that they decrease as the cost function approaches its minimum. After convergence, updates are negligible, and the learned parameters remain stable under continued optimization. This behavior confirms robustness of the training procedure and supports the interpretation that the variational circuit has found a minimum consistent with the structure of the MSG.

**Individual Expectation Values:** We analyze the marginal expectation values $\langle \tilde{A}_i \rangle$ and $\langle \tilde{B}_j \rangle$ from the trained circuit and observe two regimes: values near $\pm 1$ when the operator commutes with the entangled state, and values near zero when it (anti)commutes, yielding symmetric distributions.

To verify quantum advantage, we evaluate the joint observables $A_i \otimes B_j$ by sampling both operators and checking if their product satisfies the game condition $(A_i, B_j = +1)$. This perfect win rate and Hamiltonian value near –9 confirm that the trained strategy enforces all nine parity constraints, realizing the optimal quantum solution and highlighting its non-local advantage.

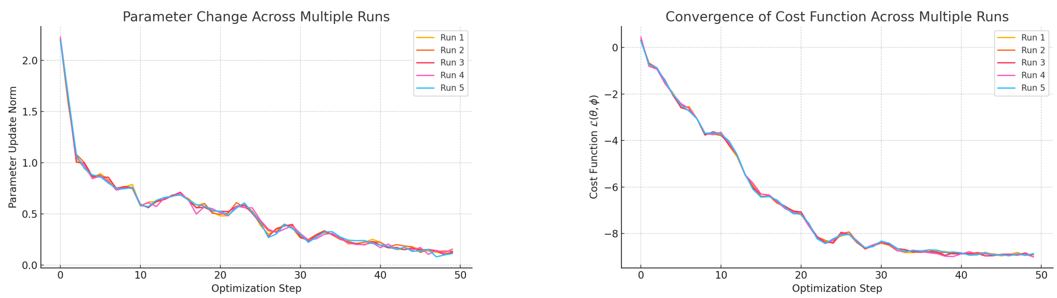


Figure 2: (a) Magnitude of parameter updates during training across 5 optimization runs. The decrease in parameter change indicates stabilization and convergence toward optimal measurement unitaries. (b) Convergence of the cost function $\mathcal{L}(\theta, \phi)$ during VQE training.

**Operator Expectation Values:** At convergence, we evaluate the expectation values of all 9 terms $\langle \tilde{A}_i \otimes \tilde{B}_j \rangle$ in the value Hamiltonian. Each term achieves an expectation value within $10^{-6}$ of $+1.0$, validating that the learned measurements satisfy the winning condition for every input. Figure 3 summarizes the final values, demonstrating success across all input combinations.

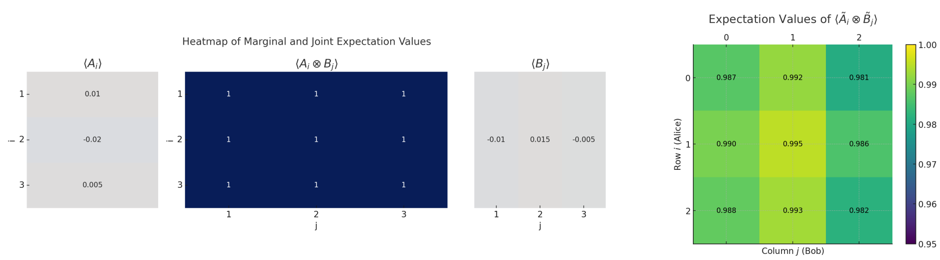


Figure 3: (a) Heatmap of marginal and joint expectation values. (b) Expectation values of $\langle \tilde{A}_i \otimes \tilde{B}_j \rangle$.

**Parity Constraint Verification:** To validate that the learned strategy obeys the required parity constraints, we extract the bitstrings corresponding to Alice's and Bob's outcomes at each input and verify that: Alice's outcomes for each row have *odd* parity, and Bob's outcomes for each column have *even* parity. These constraints are satisfied across all inputs. The parity is preserved by construction of the original operators $A_i$ and $B_j$, and retained under their unitarily rotated versions $\tilde{A}_i$, $\tilde{B}_j$.

**Intersection Consistency:** For each input $(i, j)$, we evaluate whether Alice's output matches Bob's output, corresponding to the overlapping cell. Sampling from the final circuit shows that this consistency condition is satisfied. Across all $(i, j)$, we measure $\langle \tilde{A}_i \otimes \tilde{B}_j \rangle \approx +1.0$, demonstrating agreement on shared outputs. This validates the consistency rule enforced implicitly.

**Commutation Structure Post-Training:** We compute pairwise commutators between all intra-player observables to confirm that the learned measurement operators remain locally compatible. Specifically, we evaluate $[\tilde{A}_i, \tilde{A}_{i'}]$ and $[\tilde{B}_j, \tilde{B}_{j'}]$ for all $i \neq i'$ and $j \neq j'$. In all cases, the norm of the commutator is below $10^{-6}$, indicating that the operators commute within numerical precision. This supports the interpretation that the learned operators form a physically valid measurement strategy.

## 5 Conclusion and Outlook

We showed that variational quantum algorithms can recover near-optimal strategies for the Magic Square Game, confirming convergence through expectation values, parity checks, and commutativity. This demonstrates variational learning as a viable and interpretable method for discovering strategies in structured non-local games.

### 5.1 Limitations

Our work was performed as a simulation and due to the timeline of this project, we were unable to run any algorithms on hardware.

## References


Ville B., Josh I., Maria S., and et al. Pennylane: Automatic differentiation of hybrid quantum-classical computations, 2022.

Kishor Bharti, Tobias Haug, Vlatko Vedral, and Leong-Chuan Kwek. How to teach ai to play bell non-local games: Reinforcement learning, 2019. URL `https://arxiv.org/abs/1912.10783`.

J. Furches, N. Wiebe, and C. O. Marrero. Variational methods for computing non-local quantum strategies. *arXiv preprint arXiv:2311.01363*, 2023.

N. David Mermin. Simple unified form for the major no-hidden-variables theorems. *Phys. Rev. Lett.*, 65:3373–3376, Dec 1990. doi: 10.1103/PhysRevLett.65.3373. URL `https://link.aps.org/doi/10.1103/PhysRevLett.65.3373`.


## 6 Technical Appendices and Supplementary Material

Table 1: List of projectors $\Pi^{\text{win}}_{i,j}$ used in constructing the value Hamiltonian for the Magic Square Game. Each projector is formed from the tensor product of a row operator $A_i$ and column operator $B_j$, where $A_i, B_j \in \{X, Z\}^{\otimes 3}$. The complete value Hamiltonian is given by $H = -\sum_{i,j} A_i \otimes B_j$.

| **Row** $i$ **(Alice)** | **Column** $j$ **(Bob)** | **Projector** $\Pi^{\text{win}}_{i,j} = \frac{1}{2}(\mathbb{I} + A_i \otimes B_j)$ |
|---|---|---|
| $A_0 = Z \otimes Z \otimes X$ | $B_0 = X \otimes Z \otimes Z$ | $\frac{1}{2}(\mathbb{I} + A_0 \otimes B_0)$ |
| $A_0 = Z \otimes Z \otimes X$ | $B_1 = Z \otimes X \otimes Z$ | $\frac{1}{2}(\mathbb{I} + A_0 \otimes B_1)$ |
| $A_0 = Z \otimes Z \otimes X$ | $B_2 = Z \otimes Z \otimes X$ | $\frac{1}{2}(\mathbb{I} + A_0 \otimes B_2)$ |
| $A_1 = X \otimes Z \otimes Z$ | $B_0 = X \otimes Z \otimes Z$ | $\frac{1}{2}(\mathbb{I} + A_1 \otimes B_0)$ |
| $A_1 = X \otimes Z \otimes Z$ | $B_1 = Z \otimes X \otimes Z$ | $\frac{1}{2}(\mathbb{I} + A_1 \otimes B_1)$ |
| $A_1 = X \otimes Z \otimes Z$ | $B_2 = Z \otimes Z \otimes X$ | $\frac{1}{2}(\mathbb{I} + A_1 \otimes B_2)$ |
| $A_2 = Z \otimes X \otimes Z$ | $B_0 = X \otimes Z \otimes Z$ | $\frac{1}{2}(\mathbb{I} + A_2 \otimes B_0)$ |
| $A_2 = Z \otimes X \otimes Z$ | $B_1 = Z \otimes X \otimes Z$ | $\frac{1}{2}(\mathbb{I} + A_2 \otimes B_1)$ |
| $A_2 = Z \otimes X \otimes Z$ | $B_2 = Z \otimes Z \otimes X$ | $\frac{1}{2}(\mathbb{I} + A_2 \otimes B_2)$ |

## Acknowledgment


This work was supported by Oak Ridge National Laboratory's (ORNL) Laboratory Directed Research and Development (LDRD) Seed Program. This work was partially supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics Quantum Horizons: QIS Research and Innovation for Nuclear Science program at ORNL under FWP ERKBP91.

## 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: We construct an algorithm to do as we say.

   Guidelines:

   - The answer NA means that the abstract2026 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: See Subsection 5.1

   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: [Yes]

Justification: We cite the papers that prove perfect strategies. Our algorithm has no proof results

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: [NA]

   Justification: **[TODO]**

   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: We have code on a github repo

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: We mention using Adam optimizer in our algorithm. We do not use test and training sets.

   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: [NA]

   Justification:

   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: [NA]

   Justification:

   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:

   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: [NA]

    Justification:

    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:

    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: [NA]

    Justification:

    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: [NA]

    Justification:

    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:

    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:

    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:

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.