
Localist Topographic Expert Routing: A Barrel Cortex-Inspired Modular Network for Sensorimotor Processing

Tianfang Zhu¹, Dongli Hu¹, Jiandong Zhou¹, Kai Du^{2*}, Anan Li^{1,3 *}

¹Wuhan National Laboratory for Optoelectronics,
Huazhong University of Science and Technology

²Psychological and cognitive sciences, Tsinghua University

³School of Biomedical Engineering, Hainan University

¹{funfunfun, hudongli, jiandongzhou, aali}@hust.edu.cn ²kai_du@tsinghua.edu.cn

Abstract

Biological sensorimotor systems process information through spatially organized, functionally specialized modules. A canonical example is the rodent barrel cortex, in which each vibrissa (whisker) projects to a dedicated cortical column, forming a precise somatotopic map. This anatomical organization stands in stark contrast to the architectures of most artificial neural networks, which are typically monolithic or rely on expert-isolated mixture-of-experts (MoE) mechanisms. In this work, we introduce a brain-inspired modular architecture that treats the barrel cortex as a biologically constrained instantiation of an expert system. Each module (or “expert”) corresponds to a cortical column composed of multiple neuron subtypes spanning vertical cortical layers. Sensory signals are routed exclusively to their corresponding columns, with inter-column communication restricted to local neighbors via a sparse gating mechanism. Despite these anatomical constraints, our model achieves competitive, state-of-the-art performance on challenging 3D tactile object classification benchmarks. Columnar parameter sharing provides inherent regularization, enabling 97% parameter reduction with improved training stability. Notably, constrained localist routing suppresses inter-module activity correlations, mirroring the barrel cortex’s lateral inhibition for sensory differentiation, while suggesting MoE’s potential to reduce expert redundancy through collaborative constraints. These results suggest how cortical principles of localist expert routing and topographic organization could potentially be translated into machine learning architectures. Code is available at <https://github.com/fun0515/MultiBarrelModel>.

1 Introduction

One of the hallmarks of biological intelligence is its use of modular, topographically structured systems to process sensorimotor information [22, 46, 21]. In rodents, the barrel cortex [38, 45] epitomizes this principle: each whisker maps one-to-one onto a dedicated cortical column, enabling localized, efficient processing of tactile input (Fig. 1A). This precise sensor-to-column mapping supports robust spatial discrimination and energy-efficient computation. Extensive neuroscientific studies have characterized both the vertical microcircuitry [12, 37, 49, 15] within each column—organized across layer 2 (L2) through L6—and the horizontal connections that support lateral integration between neighboring columns [33, 13, 24]. However, the implications of this biologically modular and anatomically constrained design for artificial intelligence remain underexplored.

*Corresponding authors.




Figure 1: Outline of our barrel cortex-constrained localist expert model. **(A)** Neural pathway of “whisker-to-barrel” processing in rodent. Top: whisker signals are transmitted via (1) brainstem neurons and relayed through (2) thalamic neurons to the barrel cortex. Each processing stage (brainstem, thalamus, and barrel cortex) contains somatotopically organized units corresponding to individual whiskers, with the barrel cortex exhibiting hierarchical cortical columns (modified from [38]). Bottom: microscopic image of the barrel cortex, clearly showing horizontally arranged columns and vertically stratified layers. **(B)** Computational architecture of our multi-barrel model. Each barrel column incorporates 37 synaptic connections across eight neuronal subtypes, derived from established neuroscience studies, strictly adhering to the one-to-one “whisker-barrel” mapping by processing only its assigned tactile sensor’s data. Neuronal subtype labels, from layer 2 (L2) to L6, correspond to their neuroscientifically defined nomenclature (e.g., “SSP” denotes L4 spiny stellate pyramidal neurons [9, 44]). Neuronal subtype details in supplementary materials. Adjacent barrels integrate inter-barrel currents through a dynamic gating network. The model employs temporally dilated 1D convolutions [18] to simulate brainstem and thalamic preprocessing without inter-sensor leakage, utilizes 2D convolutions to integrate L5/6 neuronal states across barrels mimicking cortical-subcortical projections, and applies a multilayer perceptron (MLP) for final predictions.

By contrast, contemporary artificial neural networks generally do not impose local specialization. Architectures such as convolutional neural networks (CNNs) [27, 47] and Transformers [50, 8] are predominantly global in their computational flow, lacking explicit modular substructures tied to specific sensory channels. The recent success of Mixture-of-Experts (MoE) models [11, 35, 7, 30] has revived interest in modular computation in artificial intelligence. These models use a gating mechanism to activate a sparse subset of sub-networks (“experts”) per input, improving both efficiency and scalability. Nevertheless, the routing strategy in artificial MoEs fundamentally differs from that of the brain. Artificial MoEs focus on task allocation—not enforced inter-expert collaboration and communication. In contrast, the brain utilizes topographic and localist routing, where each expert (e.g., cortical column) processes only its corresponding sensory signal and communicates primarily with spatially adjacent modules. This contrast raises a key question: can we construct a MoEs-like model based on the biologically grounded routing constraints observed in the brain?

In this work, we propose a barrel cortex-inspired modular expert network for sensorimotor learning. The architecture consists of 39 interacting modules, each corresponding to a tactile sensor (whisker), faithfully replicating the somatotopic organization of the rodent barrel cortex. Each module processes its designated sensory stream using an internal structure composed of multiple biologically inspired neuron subtypes, mirroring cortical L2 through L6 (Fig. 1B). All modules share a common architecture




Figure 2: Horizontal spread of neural activity in the multi-barrel model. **(A)** Schematic of vertical connectivity generalized from published neuroscience studies. Red lines indicate synaptic connections between distinct neuronal subtypes. **(B)** Temporal evolution of barrel states ($\frac{\text{spikes}}{\Delta T}$) in the initial model without inter-barrel currents, where only the central barrel receives sustained current input from its corresponding sensor. Barrel layout reflects the spatial arrangement of the sensor array. **(C)** Same as in (B), but with inter-barrel currents enabled in the model. **(D)** Experimentally observed activation spread across barrels in real optogenetic recordings (adapted from [24]).

and parameter set, reflecting the repeated structure observed across cortical columns, and enabling substantial reductions in parameter count. To enable local integration while preserving spatial structure, we introduce a topographic gating mechanism: each module exchanges information only with its immediate neighbors, reproducing the localized lateral connectivity found in the barrel cortex. This results in sparse, spatially constrained expert activation, in sharp contrast to the global routing strategies used in classical MoE frameworks. Our key contributions are summarized as follows:

- **Biologically constrained modular architecture:** We introduce a modular neural network whose submodules (barrels) replicate the canonical microcircuitry and sensor-mapping of the rodent barrel cortex. Each module integrates multiple neuron subtypes and strictly processes signals from its designated sensor, enabling biologically faithful modular computation.
- **Localist expert routing:** We implement a spatially grounded gating mechanism that enables communication between adjacent modules only, contrasting with current MoEs that neglect inter-expert routing. Experimental results demonstrate that localist connection suppress activity correlations between modules while reducing functional connection distances.
- **Empirical validation on tactile tasks:** Our biologically inspired model achieves state-of-the-art performance on challenging 3D tactile classification datasets [17]. The model reduces parameter usage by 97% through shared columnar weights, while improving training stability. These results demonstrate the viability of embedding strong biological priors into scalable machine learning systems.

Our work positions the barrel cortex as a neurobiologically constrained realization of an expert system, providing a conceptual and empirical bridge between cortical computation and modular artificial intelligence. By integrating neuroscience-inspired modularity with sparse local interactions, we offer a new perspective on the design of efficient, interpretable, and high-performing artificial systems. We argue that principles such as topographic routing, localist specialization, and structural sparsity—long favored by evolution—can inform the development of next-generation machine learning architectures with more brain-like sensorimotor intelligence.

2 Related work

Barrel cortex model: Due to its highly specialized structure and unique tactile functions, the rodent barrel cortex has become a preferred model system for neuroscientists studying sensory perception. Through techniques such as immunohistochemistry, microscopic imaging, and neuronal reconstruction, considerable knowledge has been accumulated on its cellular distribution patterns [37, 49, 43] and vertical [45, 12, 32]/horizontal [38, 24, 14, 3] circuit organization. Based on these anatomical data, computational neuroscientists have developed local circuit models [2, 23, 26] and detailed dendritic models [29] to investigate network dynamics. However, these brain simulation models lack learning capabilities, and few incorporate multi-barrel column architectures. Recently, a trainable superficial layer model that maintains biological plausibility while demonstrating tactile processing functions was proposed [60], though it remains limited to single-barrel representation without incorporating “whisker-barrel” topographic mapping.

Columnar machine learning models: The uniform yet distinctly layered structure of the neocortex, widely reported in neuroscience, has also drawn attention in machine learning. Several studies [20, 19, 41] have mimicked the columnar organization of the neocortex by designing analogous basic units that exhibit hierarchical feature abstraction and local-to-global integration. However, their information pathways are governed by simplistic rules, resulting in limited performance. Recent columnar models [56, 25] have prioritized hardware-friendly physical implementations to improve computational efficiency. However, these models overlook the fact that cortical column functionality is based on complex and diverse neural circuits [22], and still fall short of matching the performance of contemporary machine learning systems. Meanwhile, loosely modular architectures like MoE [11, 7, 35, 30] in large-scale models are gaining prominence. A promising future direction lies in the simulation of the collaborative paradigm of neocortical regions to construct versatile expert systems capable of multifunctional integration.

3 Methodology

3.1 Biologically constrained columnar modules

To enable scalable yet biologically plausible modular replication, we model barrel columns at the neural pathway-level. While prior single column models [4, 6, 34] achieved one-to-one biological fidelity through massive neuron counts (requiring cluster computing), our architecture strategically incorporates 8 excitatory neuron subtypes and 37 documented projection pathways from barrel cortex studies [38, 45, 12, 15, 32, 40, 28, 59], including projections from Layer 4 (L4) to L2/3 neurons and feedback connections from L5/6 to L4 neurons, among others (Fig. 1B). To enhance scalability, we: (1) represent inhibitory effects through negative weights rather than explicit interneurons, and (2) implement each subtype with 32 cells (256 neurons/barrel), a configuration that maintains biological plausibility while enabling efficient scaling to multiple barrels. Detailed information on the eight summarized neuronal subtypes is provided in the supplementary materials.

Neuronal dynamics follow an adaptive Leaky Integrate-and-Fire (aLIF) model [55] with background currents. The membrane potential V update equation is defined as follows:

$$\begin{aligned} V^{(t)} &= e^{-\frac{1}{\tau_m}} \cdot V^{(t-1)} + (1 - e^{-\frac{1}{\tau_m}}) \cdot R_m \cdot I^{(t)} - S^{(t-1)} \cdot \theta^{(t-1)} \\ I^{(t)} &= I_e^{(t)} + I_r^{(t)} + I_{agg}^{(t)} + \epsilon^{(t)}, \quad \epsilon^{(t)} \sim \mathcal{N}(0, 1) \\ S^{(t)} &= \Theta(V^{(t)} - \theta^{(t)}), \quad \Theta(x) = \begin{cases} 1 & x > 0 \\ 0 & \text{otherwise} \end{cases} \end{aligned} \quad (1)$$

, where R_m and τ_m represent the membrane resistance and time constant, respectively. The input currents I are categorized into four distinct components: I_e denotes external input currents, I_r represents intra-barrel synaptic currents, I_{agg} corresponds to aggregated inter-barrel currents from neighboring columns, and ϵ denotes the background noise sampled from a standard normal distribution. When the membrane potential V exceeds the firing threshold θ , the neuron emits a spike S . The neuronal firing threshold undergoes an adaptive elevation through spike-triggered accumulation, governed by the following dynamics:

$$\begin{aligned} \theta^{(t)} &= \theta_{init} + \beta \cdot \eta^{(t)} \\ \eta^{(t)} &= e^{-\frac{1}{\tau_{adp}}} \cdot \eta^{(t-1)} + (1 - e^{-\frac{1}{\tau_{adp}}}) \cdot S^{(t-1)} \end{aligned} \quad (2)$$

Table 1: Comparison across methods on two tactile datasets [17]. *Independent* and *Shared* denote our model’s 39-barrel configurations with independent and shared parameters, respectively. *Single* refers to a single barrel model with equivalent neuron size (39×256).

Method	EvTouch-Objects (%)	EvTouch-Containers (%)
TactileSGNet [17]	89.44	64.17
Grid-based CNN [17]	88.40	60.17
GCN [17]	85.14	58.83
Method in [52]	90.28	-
SnnTdlc [57]	91.04	67.33
AM-SGCN [53]	91.32	-
GGT-SNN [54]	92.36	75.00
Single barrel model	88.89	70.00
Independent multi-barrel model	92.36	85.00
Shared multi-barrel model	94.44	86.67

, where θ_{init} and τ_{adp} denote the resting firing threshold and adaptation time constant, respectively, with β being a constant parameter set to 1.8. The recurrent current I_r received by a postsynaptic neuronal population from other upstream populations within the home column can be computed as:

$$I_r^{(t)} = \sum_{j \in S_i} W_j \cdot S_j, \quad W_j \in \mathbb{R}^{32 \times 32} \quad (3)$$

, where S_i denote the presynaptic subtype ensemble and the corresponding trainable connection weights, respectively. The computation processes for I_{agg} and I_e in Eq. 1 are elaborated in the following two sections.

3.2 Localist inter-barrel routing

Beyond vertical intra-barrel pathways, the barrel cortex employs horizontal inter-barrel connections for sensory integration. Neuroanatomical evidence indicates that lateral connections follow an adjacency-priority principle [33, 13], with inter-barrel signaling predominantly originating from spatially adjacent columns and minimal contributions from distant ones. We implemented this biologically observed connectivity through K-nearest neighbor (KNN) spatial mapping coupled with sparse gating mechanisms.

For a 32-neuron subtype in our model, input signals from neighboring barrels follow these governing equations. First, the most relevant neighboring subtypes indices $\mathcal{T}$ are dynamically selected based on gating weights:

$$\begin{aligned} \mathcal{T}^{(t)} &= TopK(G^{(t)}, \Gamma), \Gamma = \lfloor K \cdot 8 \cdot \gamma \rfloor \\ G^{(t)} &= MLP_\phi(I_e^{(t)}) \in \mathbb{R}^{K \times 8} \end{aligned} \quad (4)$$

, where K denotes the number of spatially adjacent barrels (nearest neighbors). γ represents the fraction of total ($K \cdot 8$) neuronal subtypes to be routed, giving Γ selected subtypes. $I_e^{(t)}$ and ϕ correspond to external input currents and trainable parameters, respectively. $TopK(\cdot)$ returns a binary mask $\in \{0, 1\}^{K \times 8}$ with exactly Γ nonzeros. The spatial distance between barrels is computed based on the sensor coordinates provided in the tactile dataset [17]. K and γ were assigned values of 4 and 0.2, respectively.

Then, the final aggregated neighborhood current $I_{agg}^{(t)}$ is computed by applying the selected subtypes indices from $\mathcal{T}^{(t)}$ (mask) to the spiking states $S_g^{(t)}$, flattening the results to a vector in $\mathbb{R}^{\Gamma \cdot 32}$, and performing a trainable linear transformation:

$$I_{agg}^{(t)} = W_{agg} \cdot Flatten(\mathcal{T}^{(t)} \odot S_g^{(t)}) \quad (5)$$

, where $W_{agg} \in \mathbb{R}^{32 \times (\Gamma \cdot 32)}$ denotes the trainable weight matrix mapping the flattened neighborhood activity vector to the 32 neurons of the target barrel, and $S_g^{(t)} = [s_1^{(t)}, s_2^{(t)}, \dots, s_{K \times 8}^{(t)}] \in \mathbb{R}^{(K \cdot 8) \times 32}$ represents the tensor composed of spiking states from all neuronal subtypes in K neighboring barrels at timestep t .




Figure 3: Performance comparison of three model variants. **(A)** Scores from 10 repeated random training runs for each variant. **(B)** Loss landscapes of independent-parameter and shared-parameter models. Perturbations were applied along two orthogonal directions to trained parameters: $\theta' = \theta + \alpha\eta_1 + \beta\eta_2$, where θ and θ' denote the original and perturbed parameters. α and β are perturbation magnitudes. Top to bottom: results on the EvTouch-Objects and EvTouch-Containers datasets.

Fig. 2B-C illustrates the temporal state evolution in our multi-barrel model under optogenetic-like stimulation [1, 16], achieved through constant current activation of the central sensorimotor barrel. The results reveal that incorporating inter-barrel current coupling induces propagated activation from the stimulated barrel to neighboring regions, progressively diffusing across the entire array. This spatiotemporal propagation pattern aligns with cortical barrel dynamics reported in empirical optogenetic studies of the “whisker-barrel” somatosensory system [24, 14, 3].

3.3 Empirical validation on tactile datasets

We aim to both develop a barrel cortex-inspired architecture and validate the “whisker-barrel” system as a localist expert processor. Below, we describe the public datasets used and our differentiable readin-readout implementation, which bridges biological computation with machine learning optimization through effective gradient propagation.

Two tactile datasets similar to whisker systems, EvTouch-Objects and EvTouch-Containers [17], were employed to benchmark our model against artificial neural networks. These datasets feature temporal signals recorded from independent NeuTouch sensor arrays [48] during interactions with diverse 3D objects, requiring the model to predict object categories based on dynamic tactile inputs. Each data sample has a shape of $[39, 2, T]$, representing two channel signals recorded from 39 sensors over T timesteps. Appendix A.1 provides complete tactile dataset specifications.

We treat each sensor as a rodent whisker (Fig. 1B), where brainstem-thalamic signal preprocessing [42, 5] is simulated via 2 one-dimensional dilated convolutional layers [18]. A linear layer then serves as the thalamic signal relay. The external current I_e received by a neuronal subtype in barrel cortex is computed as:

$$I_e^{(t)} = W_{th}^T \cdot (W_{c2} \otimes \sigma(W_{c1} \otimes X^{(t)})), \quad W_{th} \in \mathbb{R}^{64 \times 32} \quad (6)$$

, where X and $\otimes$ denote the input data and dilated convolution operation, respectively. W_{c1} and W_{c2} represent the first- and second-layer convolution kernels. σ is the activation function. The convolution operates exclusively along the temporal dimension, preventing inter-sensor information leakage. The output tensor dimension becomes $[39, 64, T']$, where $T' = \lfloor \frac{T-d(k-1)-1}{s} \rfloor + 1$ with d , k and s being the dilation rate, kernel size and stride, respectively.

Given the dominant role of L5/6 neurons in driving subcortical projections within barrel cortex [13, 45], we read-out the state of our model’s L5/6 neuronal populations through: (1) spatiotemporal integration via a two-dimensional convolution layer, followed by (2) final classification prediction




Figure 4: Ablation analysis of synaptic connectivity on EvTouch-Containers. **(A)** Model performance after isolated removal of vertical connections and horizontal inter-barrel currents, ranked by descending accuracy. **(B)** Quantitative importance of 37 intra-barrel connections, measured by performance degradation magnitude. **(C)** Four highest-impact signaling pathways identified from (B).

through a multilayer perceptron (MLP). The formula is expressed as follows:

$$\hat{y} = \text{Softmax}(W_{mlp} \cdot \text{Flatten}(W_{c3} \odot S_{L5/6})), \quad S_{L5/6} \in \mathbb{R}^{T' \times 39 \times (32 \cdot 4)} \quad (7)$$

, where $\odot$ denotes the two-dimensional convolution operation. W_{c3} and W_{mlp} represent the convolution kernel and MLP parameters, respectively. $S_{L5/6}$ denotes the spiking states of four neuronal ensembles (32 neurons each) in L5/6 across 39 barrels over T' timesteps. The model employs standard cross-entropy loss, with spiking neuron gradients computed via the Gaussian surrogate function: $\mathcal{N}(V^{(t)}|\theta^{(t)}, \sigma^2)$, where σ is set to 0.5.

Given that our multi-barrel model is extended from a single barrel as introduced in Sec. 3.1, a natural question arises: whether these barrels **share training parameters** or maintain their own independent parameters. Excluding the shared readin and readout pathways, in our 39-barrel configuration, the parameter count of the column modules with shared parameters is reduced by approximately 97% compared to those with independent parameters. In practice, their parameter counts are 59,104 and 2,305,056, respectively, a reduction of two orders of magnitude. This significant decrease in parameters is crucial not only for computational efficiency, but also serves as a form of regularization. In the experimental section, we investigate the tactile task performance of both model variants.

4 Experimental results

4.1 Implementation details

Our multi-barrel model was trained for 200 epochs on both EvTouch-Objects and EvTouch-Containers datasets [17] using an AdamW optimizer [31] with 0.1 weight decay. The initial learning rate of 0.0008 decayed by a factor of 0.8 every 10 epochs. The loss function employed solely the standard cross-entropy criterion. All experiments were conducted on an 80GB NVIDIA A100 GPU.

4.2 Performance comparison with baseline methods

First, we evaluated our multi-barrel model against reported baselines on both EvTouch-Objects and EvTouch-Containers datasets, achieving state-of-the-art (SOTA) classification accuracy (Tab. 1). Our model outperforms the previous best-performing method GGT-SNN [54], achieving accuracy improvements of 2.1% on EvTouch-Objects and 11.7% on EvTouch-Containers. While prior work employed graph neural networks [17, 53, 54] to process sensor topology, our biologically constrained architecture attains comparable performance through simple “whisker-to-barrel” mapping. Notably, collapsing the columnar organization into a single barrel reduces accuracy by 5.6% (EvTouch-Objects) and 16.7% (EvTouch-Containers), demonstrating the functional importance of modular architecture.




Figure 5: Statistic of inter-barrel correlations. **(A)** Barrel-wise localist routing. Arrows indicate the most frequently selected source neighbor barrels, as determined by cumulative counts across samples and timesteps. **(B)** Neural activity of 39 barrels during a single trial. A sliding window (length = 10, stride = 2) was applied to compute pairwise activity correlations. The three most strongly correlated barrel pairs per window were selected as high-correlation pairs. **(C)** Top-5 most frequent strongly correlated barrel pairs on EvTouch-Objects dataset. From left to right are models with blocked and intact horizontal inter-barrel currents. **(D)** Same as (C), but for EvTouch-Containers dataset. **(E)** Statistical metrics between blocked and intact models. From left to right are whole-time mean barrel correlations, windowed mean correlations, mean distances of strongly correlated barrel pairs, and mean distances of Top-5 pairs.

Furthermore, we compared two model variants from Sec. 3.3: parameter-shared (39 barrels) versus independent-parameter configurations. The shared-parameter variant showed superior performance (2.1% higher on EvTouch-Objects, 1.7% on Containers) with greater training stability, whereas independent-parameter model occasionally failed to train (Fig. 3A). Loss landscape analysis reveals this dichotomy: shared parameters create smooth basins enabling robust convergence, whereas independent parameters yield rugged terrain with local fluctuations (Fig. 3B). This likely occurs because independent parameters per barrel generate conflicting gradient updates during training, thereby complicating optimization. In contrast, shared parameters naturally enforce cross-barrel regularization, effectively mitigating stochastic gradient variations. These results validate the functional advantages of biologically uniform organization and modular replication in artificial expert systems.

4.3 Ablation analysis of synaptic connectivity

Then, we systematically assess synaptic importance in our multi-barrel model through ablation experiments, isolating and blocking each connection while quantifying performance degradation. Taking the EvTouch-Containers dataset as an example, inter-barrel aggregated currents show the greatest impact, reducing classification accuracy by approximately 11.7% (Fig. 4A), while some vertical connections exhibit minimal impact.

Fig. 4B demonstrates laminar-specific effects through directional analysis of 37 intra-barrel vertical projections. The four L4 to L2/3 connections collectively contributed 13.3% to performance, whereas two reciprocal L2/3 to L4 connections showed weaker impacts but stronger L5 modulation. We identified four dominant pathways: L4-L2/3, L4-L6, L2/3-L5, and L6-L4 (Fig. 4C). These pathways align with established neuroanatomical principles: the L4-L2/3-L5 microcircuit represents a canonical barrel cortex circuit [38, 32], and L6 cortico-thalamic pyramidal (CTP) cells receive dendritic inputs from L4 pyramidal neurons [40, 45] and project reciprocally to L4 boundaries [28, 59]. This alignment may suggest that cortical microcircuits represent an evolutionarily optimized architecture.

Besides, deactivating specific synaptic connections (e.g., STP to L2P) does not compromise model performance, suggesting the existence of redundant pathways. This functional resilience likely results from substitution by connections with similar directional properties. Consequently, more precise characterization is required to delineate the distinct contributions of specific neuronal subtypes.

4.4 Analysis of localist routing patterns

Next, we examined the localist routing behavior of our multi-barrel model. In contrast to the decentralized, globally routed structure of conventional MoE systems, each expert (i.e., barrel column) in our system communicates only with its immediate neighbors. As illustrated in Fig. 5A, the KNN gating constraints introduced in Sec. 3.2 yields predominantly short-range horizontal connections in both the initial and the trained models; long-range projections spanning multiple barrels are entirely absent.

To quantify the functional impact of these horizontal connections, we analyzed the correlations of barrel-column activity on the testsets, considering both full-trial and sliding-window time scales (Fig. 5B). Introducing inter-barrel currents decreased global mean correlations by 0.07 on EvTouch-Objects and 0.01 on EvTouch-Containers; windowed (local) correlations fell by 0.03 and 0.01, respectively (Fig. 5E). Reduced synchrony suggests that horizontal currents promote richer, more specialized activity patterns across columns, thereby improving the discriminability of complex inputs. This effect echoes the lateral-inhibition mechanism reported in the biological barrel cortex, where activated columns transiently suppress their neighbours to sharpen sensory contrast [39, 36, 10].

The spatial signature of these effects is equally striking. Inter-barrel currents shortened the average distance between strongly correlated column pairs by 0.06 grid units on EvTouch-Objects and 0.17 on EvTouch-Containers; for the top-5 most-correlated pairs, the reductions were even larger—1.8 and 3.48 units, respectively (Fig. 5E). As depicted in Fig. 5C-D, the top-correlation pairs overwhelmingly involved neighbouring barrels once horizontal gating was enabled. Together, these findings show that localist inter-barrel currents simultaneously strengthen short-range functional connectivity and amplify activity differentiation across columns. The results are consistent with long-standing neurophysiological principles [33, 13] and may inform the design of future MoE systems with enforced inter-expert communication to reduce information redundancy [51, 58].

5 Conclusion

In this work, we focus on the rodent barrel cortex as a localist expert system in the brain, developing a multi-barrel model that strictly adheres to the one-to-one “whisker-barrel” somatotopic mapping. Our architecture faithfully replicates the laminar and columnar organization of barrel cortex, achieving state-of-the-art performance on two tactile datasets while successfully balancing biological plausibility and behavioral performance. Experimental results demonstrate that the cortex’s uniform modular architecture facilitates parameter sharing to enhance training stability, while localist routing reduces functional connection distances and suppresses inter-barrel activity correlations. Inter-barrel currents may sharpen perception and reduce redundancy. This work reveals the potential of brain’s native expert systems to inspire next-generation machine learning architectures.

Limitations and future work: On one hand, the redundant synaptic connections observed in Fig. 4B necessitate more refined neuronal subtype differentiation in modeling, which requires deeper anatomical knowledge integration. On the other hand, our multi-barrel model remains distinct from standard MoE architectures and has not yet been validated on mainstream machine learning benchmarks. Developing brain-inspired MoE architectures that emulate cortical coordination mechanisms presents an exciting research direction. Another promising future direction involves embedding our biologically constrained barrel cortex model into realistic real-time interactive environments. This framework would leverage biologically constrained models to explore performance in complex sensory tasks (e.g., navigation), thereby significantly advancing both our understanding of brain operating principles and potential applications.

Acknowledgments

This work was financially supported by the STI 2030-Major Projects (2021ZD0201002), and National Natural Science Foundation of China grants (T2122015, 32471149).

References

- [1] Nicolò Accanto, François GC Blot, Antonio Lorca-Cámara, Valeria Zampini, Florence Bui, Christophe Tourain, Noam Badt, Ori Katz, and Valentina Emiliani. A flexible two-photon fiberscope for fast activity imaging and precise optogenetic photostimulation of neurons in freely moving mice. *Neuron*, 111(2):176–189, 2023.
- [2] Tommer Argaman and David Golomb. Does layer 4 in the barrel cortex function as a balanced circuit when responding to whisker movements? *Neuroscience*, 368:29–45, 2018.
- [3] Rachel Aronoff, Ferenc Matyas, Celine Mateo, Carine Ciron, Bernard Schneider, and Carl CH Petersen. Long-range connectivity of mouse primary somatosensory barrel cortex. *European Journal of Neuroscience*, 31(12):2221–2233, 2010.
- [4] Yazan N Billeh, Binghuang Cai, Sergey L Gratiy, Kael Dai, Ramakrishnan Iyer, Nathan W Gouwens, Reza Abbasi-Asl, Xiaoxuan Jia, Joshua H Siegle, Shawn R Olsen, et al. Systematic integration of structural and functional data into multi-scale models of mouse primary visual cortex. *Neuron*, 106(3):388–403, 2020.
- [5] Manuel A Castro-Alamancos and Morgana Favero. Whisker-related afferents in superior colliculus. *Journal of Neurophysiology*, 115(5):2265–2279, 2016.
- [6] Guozhang Chen, Franz Scherr, and Wolfgang Maass. A data-based large-scale model for primary visual cortex enables brain-like robust and versatile visual processing. *science advances*, 8(44):eabq7592, 2022.
- [7] Damai Dai, Chengqi Deng, Chenggang Zhao, R. X. Xu, Huazuo Gao, Deli Chen, Jiashi Li, Wangding Zeng, Xingkai Yu, Y. Wu, Zhenda Xie, Y. K. Li, Panpan Huang, Fuli Luo, Chong Ruan, Zhifang Sui, and Wenfeng Liang. Deepseekmoe: Towards ultimate expert specialization in mixture-of-experts language models, 2024.
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [9] Veronica Egger, Thomas Nevian, and Randy M Bruno. Subcolumnar dendritic and axonal organization of spiny stellate and star pyramid neurons within a barrel in rat somatosensory cortex. *Cerebral cortex*, 18(4):876–889, 2008.
- [10] Linlin Z Fan, Simon Kheifets, Urs L Böhm, Hao Wu, Kiryl D Piatkevich, Michael E Xie, Vicente Parot, Yooree Ha, Kathryn E Evans, Edward S Boyden, et al. All-optical electrophysiology reveals the role of lateral inhibition in sensory processing in cortical layer 1. *Cell*, 180(3):521–535, 2020.

- [11] William Fedus, Barret Zoph, and Noam Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *Journal of Machine Learning Research*, 23(120):1–39, 2022.
- [12] Dirk Feldmeyer. Excitatory neuronal connectivity in the barrel cortex. *Frontiers in neuroanatomy*, 6:24, 2012.
- [13] Dirk Feldmeyer. Excitatory neuronal connectivity in the barrel cortex. *Frontiers in neuroanatomy*, 6:24, 2012.
- [14] Isabelle Ferezou, Florent Haiss, Luc J Gentet, Rachel Aronoff, Bruno Weber, and Carl CH Petersen. Spatiotemporal dynamics of cortical sensorimotor integration in behaving mice. *Neuron*, 56(5):907–923, 2007.
- [15] Pierre-Marie Gardères, Sébastien Le Gal, Charly Rousseau, Alexandre Mamane, Dan Alin Ganea, and Florent Haiss. Coexistence of state, choice, and sensory integration coding in barrel cortex lii/iii. *Nature Communications*, 15(1):4782, 2024.
- [16] Oliver M Gauld, Adam M Packer, Lloyd E Russell, Henry WP Dalglish, Maya Iuga, Francisco Sacadura, Arnd Roth, Beverley A Clark, and Michael Häusser. A latent pool of neurons silenced by sensory-evoked inhibition can be recruited to enhance perception. *Neuron*, 112(14):2386–2403, 2024.
- [17] Fuqiang Gu, Weicong Sng, Tasbolat Taunyazov, and Harold Soh. Tactilesgnet: A spiking graph neural network for event-based tactile object recognition. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 9876–9882. IEEE, 2020.
- [18] Ilyass Hammouamri, Ismail Khalfaoui-Hassani, and Timothée Masquelier. Learning delays in spiking neural networks using dilated convolutions with learnable spacings. In *The Twelfth International Conference on Learning Representations*, 2024.
- [19] Atif Hashmi and Mikko Lipasti. A cortically inspired learning model. In *International Joint Conference on Computational Intelligence*, pages 373–388. Springer, 2010.
- [20] Atif G Hashmi and Mikko H Lipasti. Cortical columns: Building blocks for intelligent systems. In *2009 IEEE Symposium on Computational Intelligence for Multimedia Signal and Vision Processing*, pages 21–28. IEEE, 2009.
- [21] Jeff Hawkins. *A thousand brains: a new theory of intelligence*. Basic Books, 2021.
- [22] Jeff Hawkins, Subutai Ahmad, and Yuwei Cui. A theory of how columns in the neocortex enable learning the structure of the world. *Frontiers in neural circuits*, 11:295079, 2017.
- [23] Chao Huang, Fleur Zeldenrust, and Tansu Celikel. Cortical representation of touch in silico. *Neuroinformatics*, 20(4):1013–1039, 2022.
- [24] BA Johnson and RD Frostig. Long, intrinsic horizontal axons radiating through and beyond rat barrel cortex have spatial distributions similar to horizontal spreads of activity evoked by whisker stimulation. *Brain Structure and Function*, 221:3617–3639, 2016.
- [25] Mikhail Kiselev. Colanet—a spiking neural network with columnar layered architecture for classification. *arXiv preprint arXiv:2409.01230*, 2024.
- [26] Yves Kremer, Jean-François Léger, Dan Goodman, Romain Brette, and Laurent Bourdieu. Late emergence of the vibrissa direction selectivity map in the rat barrel cortex. *The Journal of neuroscience : the official journal of the Society for Neuroscience*, 31(29):10689—10700, July 2011.
- [27] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017.
- [28] Pratap Kumar and Ora Ohana. Inter-and intralaminar subcircuits of excitatory and inhibitory neurons in layer 6a of the rat barrel cortex. *Journal of Neurophysiology*, 100(4):1909–1922, 2008.

- [29] Maria Lavzin, Sophia Rapoport, Alon Polsky, Liora Garion, and Jackie Schiller. Nonlinear dendritic processing determines angular tuning of barrel cortex neurons in vivo. *Nature*, 490(7420):397–401, 2012.
- [30] Bin Lin, Zhenyu Tang, Yang Ye, Jiayi Cui, Bin Zhu, Peng Jin, Jinfa Huang, Junwu Zhang, Yatian Pang, Munan Ning, et al. Moe-llava: Mixture of experts for large vision-language models. *arXiv preprint arXiv:2401.15947*, 2024.
- [31] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019.
- [32] Joachim Lübke and Dirk Feldmeyer. Excitatory signal flow and connectivity in a cortical column: focus on barrel cortex. *Brain Structure and Function*, 212(1):3–17, 2007.
- [33] Joachim Lübke, Arnd Roth, Dirk Feldmeyer, and Bert Sakmann. Morphometric analysis of the columnar innervation domain of neurons connecting layer 4 and layer 2/3 of juvenile rat barrel cortex. *Cerebral Cortex*, 13(10):1051–1063, 2003.
- [34] Henry Markram, Eilif Muller, Srikanth Ramaswamy, Michael W Reimann, Marwan Abdellah, Carlos Aguado Sanchez, Anastasia Ailamaki, Lidia Alonso-Nanclares, Nicolas Antille, Selim Arsever, et al. Reconstruction and simulation of neocortical microcircuitry. *Cell*, 163(2):456–492, 2015.
- [35] Niklas Muennighoff, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Jacob Morrison, Sewon Min, Weijia Shi, Evan Pete Walsh, Oyvind Tafjord, Nathan Lambert, Yuling Gu, Shane Arora, Akshita Bhagia, Dustin Schwenk, David Wadden, Alexander Wettig, Binyuan Hui, Tim Dettmers, Douwe Kiela, Ali Farhadi, Noah A. Smith, Pang Wei Koh, Amanpreet Singh, and Hannaneh Hajishirzi. OLMoe: Open mixture-of-experts language models. In *The Thirteenth International Conference on Learning Representations*, 2025.
- [36] William Muñoz, Robin Tremblay, Daniel Levenstein, and Bernardo Rudy. Layer-specific modulation of neocortical dendritic inhibition during active wakefulness. *Science*, 355(6328):954–959, 2017.
- [37] Simon P Peron, Jeremy Freeman, Vijay Iyer, Caiying Guo, and Karel Svoboda. A cellular resolution map of barrel cortex activity during tactile behavior. *Neuron*, 86(3):783–799, 2015.
- [38] Carl CH Petersen. Sensorimotor processing in the rodent barrel cortex. *Nature Reviews Neuroscience*, 20(9):533–546, 2019.
- [39] Carl CH Petersen and Sylvain Crochet. Synaptic computation and sensory processing in neocortical layer 2/3. *Neuron*, 78(1):28–48, 2013.
- [40] Guanxiao Qi and Dirk Feldmeyer. Dendritic target region-specific formation of synapses between excitatory layer 4 neurons and layer 6 pyramidal cells. *Cerebral cortex*, 26(4):1569–1579, 2016.
- [41] Sara Sabour, Nicholas Frosst, and Geoffrey E Hinton. Dynamic routing between capsules. *Advances in neural information processing systems*, 30, 2017.
- [42] Keisuke Sehara and Hiroshi Kawasaki. Neuronal circuits with whisker-related patterns. *Molecular neurobiology*, 43:155–162, 2011.
- [43] B Semihcan Sermet, Pavel Truschow, Michael Feyerabend, Johannes M Mayrhofer, Tess B Oram, Ofer Yizhar, Jochen F Staiger, and Carl CH Petersen. Pathway-, layer-and cell-type-specific thalamic input to mouse barrel cortex. *elife*, 8:e52665, 2019.
- [44] Jochen F Staiger, Iris Flagmeyer, Dirk Schubert, Karl Zilles, Rolf Kötter, and Heiko J Luhmann. Functional diversity of layer iv spiny neurons in rat somatosensory cortex: quantitative morphology of electrophysiologically characterized and biocytin labeled cells. *Cerebral Cortex*, 14(6):690–701, 2004.
- [45] Jochen F Staiger and Carl CH Petersen. Neuronal circuits in barrel cortex for whisker sensory perception. *Physiological reviews*, 101(1):353–415, 2021.

- [46] Mototaka Suzuki, Cyriel MA Pennartz, and Jaan Aru. How deep is the brain? the shallow brain hypothesis. *Nature Reviews Neuroscience*, 24(12):778–791, 2023.
- [47] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR, 2019.
- [48] Tasbolat Taunyazov, Weicong Sng, Hian Hian See, Brian Lim, Jethro Kuan, Abdul Fatir Ansari, Benjamin CK Tee, and Harold Soh. Event-driven visual-tactile sensing and learning for robots. *arXiv preprint arXiv:2009.07083*, 2020.
- [49] Vassiliy Tsytsarev, Sung E Kwon, Celine Plachez, Shuxin Zhao, Daniel H O’Connor, and Reha S Erzurumlu. Layers 3 and 4 neurons of the bilateral whisker-barrel cortex. *Neuroscience*, 494:140–151, 2022.
- [50] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [51] Ziyang Xiao, Dongxiang Zhang, Yangjun Wu, Lilin Xu, Yuan Jessica Wang, Xiongwei Han, Xiaojin Fu, Tao Zhong, Jia Zeng, Mingli Song, and Gang Chen. Chain-of-experts: When LLMs meet complex operations research problems. In *The Twelfth International Conference on Learning Representations*, 2024.
- [52] Jing Yang, Xiaoyang Ji, Shaobo Li, Hao Dong, Tingqing Liu, Xu Zhou, and Shuaizhen Yu. Robot tactile data classification method using spiking neural network. In *2021 China Automation Congress (CAC)*, pages 5274–5279. IEEE, 2021.
- [53] Jing Yang, Tingqing Liu, Yaping Ren, Qing Hou, Shaobo Li, and Jianjun Hu. Am-sgcnn: Tactile object recognition for adaptive multichannel spiking graph convolutional neural networks. *IEEE Sensors Journal*, 23(24):30805–30820, 2023.
- [54] Jing Yang, Zukun Yu, Shaobo Li, Yang Cao, JianJun Hu, and Ji Xu. Ggt-snn: Graph learning and gaussian prior integrated spiking graph neural network for event-driven tactile object recognition. *Information Sciences*, 677:120998, 2024.
- [55] Bojian Yin, Federico Corradi, and Sander M. Bohtë. Effective and efficient computation with multiple-timescale spiking recurrent neural networks. *CoRR*, abs/2005.11633, 2020.
- [56] Sangmin Yoo, Yongmo Park, Ziyu Wang, Yuting Wu, Saaketh Medepalli, Wesley Thio, and Wei D Lu. Columnar learning networks for multisensory spatiotemporal learning. *Advanced Intelligent Systems*, 4(11):2200179, 2022.
- [57] Gexiang Zhang, Xihai Zhang, Haina Rong, Prithwineel Paul, Ming Zhu, Ferrante Neri, and Yew-Soon Ong. A layered spiking neural system for classification problems. *International journal of neural systems*, 32(08):2250023, 2022.
- [58] Mohan Zhang, Pingzhi Li, Jie Peng, Mufan Qiu, and Tianlong Chen. Advancing moe efficiency: A collaboration-constrained routing (c2r) strategy for better expert parallelism design. *arXiv preprint arXiv:2504.01337*, 2025.
- [59] Zhong-Wei Zhang and Martin Deschênes. Intracortical axonal projections of lamina vi cells of the primary somatosensory cortex in the rat: a single-cell labeling study. *Journal of Neuroscience*, 17(16):6365–6379, 1997.
- [60] Tianfang Zhu, Dongli Hu, Jiandong Zhou, Kai Du, and Anan LI. Biologically constrained barrel cortex model integrates whisker inputs and replicates key brain network dynamics. In *The Thirteenth International Conference on Learning Representations*, 2025.




Figure 6: Tactile perception task overview. **(A)** Object categorization via several-second tactile scanning using NeuTouch sensor array (datasets: EvTouch-Objects/Containers) [17]. **(B)** NeuTouch’s 39-sensor spatial configuration (adapted from [48]). **(C)** Each sensor modeled as a whisker, with 1D dilated convolutions simulating brainstem-thalamic delay/integration [18, 42, 5].

A Technical Appendices and Supplementary Material

A.1 Tactile task introduction

This section details our tactile perception task. Both employed datasets, EvTouch-Objects and EvTouch-Containers [17], were acquired through an identical protocol: the NeuTouch [48] sensor array engaged in several-second tactile interactions with 3D objects, followed by object category prediction based on the recorded sensor signals (Fig. 6A). The NeuTouch system integrates 39 uniformly configured sensor units, each generating two-channel spiking signals, with their spatial arrangement depicted in Fig. 6B. The datasets were categorized into EvTouch-Objects (720 samples, 36 categories) and EvTouch-Containers (300 samples, 20 categories) based on object taxonomy, with both datasets split into training and test sets at an 8:2 ratio for standardized evaluation.

As described in Sec. 3.3, we modeled each of the 39 tactile sensors as a rodent whisker, with separate barrel modules processing each sensor’s signals independently. Each raw sample in the tactile dataset has a shape of $[39, 2, T]$, representing two-channel signals from 39 sensors over T timesteps, where T equals 250 for EvTouch-Objects and 325 for EvTouch-Containers. We first applied two 1D dilated convolutions along the temporal dimension to simulate brainstem and thalamic preprocessing of whisker signals [42, 5], including delay and integration effects [18]. Dilated convolutions were exclusively applied along the time axis to prevent cross-whisker signal leakage (Fig. 6C). The processed timesteps computed as $T' = \lfloor \frac{T-d(k-1)-1}{s} \rfloor + 1$, where d , k and s represent the dilation rate, kernel size, and stride respectively. Then, the preprocessed whisker signals were independently propagated to their corresponding barrel columns. Finally, a standard 2D convolution operation integrated the barrel-wise L5/6 neuronal states across the entire array, followed by an MLP to generate the final prediction output.

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: The abstract and introduction effectively summarize the paper's contributions, distilled into three key points.

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: The conclusion explicitly discusses the study's limitations, highlighted in bold text.

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: This work does not involve theoretical derivation or proof.

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: This paper details the methodology and datasets, with comprehensive task specifications provided in the appendix.

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 source code is provided in supplementary materials and will be open-sourced upon acceptance. All datasets used are publicly available.

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: Training specifications are comprehensively documented in the implementation details section.

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: This study presents multi-run training results, with all quantitative metrics evaluated on the complete test dataset.

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: All experiments were conducted on a single 80GB A100 GPU, as detailed in the implementation details section.

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 confirm that this research fully complies with 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: [NA]

Justification: As a basic science research, this work has no measurable societal implications or ethical concerns

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: This study involves no data misuse or ethical violations regarding data handling.

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: We have properly and explicitly cited all relevant prior work with accurate authorship attribution.

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: Code is accompanied by clear documentation following community standards.

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: This study does not involve 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: This study does not involve human subjects.

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: LLMs were utilized solely for language editing in this research.

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.