The geometry of primary visual cortex representations is dynamically adapted to task performance

Leyla R. Caglar 1*, Julien Corbo 2*, O. Batuhan Erkat 2,3, Pierre-Olivier Polack 2
1- Windreich Department of AI & Human Health, Icahn School of Medicine at Mount Sinai, New York, NY, USA
2- Center for Molecular and Behavioral Neuroscience, Rutgers University—Newark, Newark, NJ, USA
3- Graduate Program in Neuroscience, Rutgers University—Newark, Newark, NJ, USA
* Contributed equally

Abstract

Perceptual learning optimizes perception by reshaping sensory representations to enhance discrimination and generalization. Previous work has shown that learning a visual orientation discrimination task reshapes the population feature representations in the primary visual cortex (V1) via suppressive mechanisms. Although the computational importance of these changes has not yet been elucidated, it has been proposed that they optimize the geometry of the representation to be readout. Are these feature-encoding changes paired with changes to the representational geometry? To answer this, we investigated the relationship between V1 feature representation, behavioral performance, and neural manifold geometry in trained and naïve mice. Response dimensionality showed increases with task difficulty but was lower in trained animals, suggesting that successful learning reduces dimensionality. Based on manifold capacity, dimensionality, and radius, we further found that representational separability is a stronger predictor of individual behavioral performance. These results confirm that learning alters the geometric properties as early as early sensory representations, optimizing them for linear readout and improving perceptual decision-making.

Keywords: manifold capacity; orientation discrimination; representational geometry; sensory processing; visual perception

Introduction

When learning a task, the representations of the relevant stimuli and their reading out must undergo modifications to optimize detection, discrimination, categorization and generalization. However, these mechanisms' implementation remains elusive. In brains, perceptual or associative learning has been shown to change the neuronal feature representations of the relevant stimuli, even at early processing stages such as the primary sensory cortices. Training for a visual discrimination task alters V1 population activity by enhancing target encoding (Henschke et al., 2020; Jurjut et al., 2017; Poort et al., 2015), sparsening cue-evoked responses (Corbo et al., 2022; Failor et al., 2025), and increasing their representational separability (Poort et al., 2015, Corbo et al., 2025; Failor et al., 2025). How do these modifications of the evoked population activity support learning and performance in a discrimination task? In spite of the variability of these effects, they are associated with similar changes in representational geometry, suggesting that they take part in the same computational mechanisms (Cheng et al., 2025).

For example, the manifold capacity theory framework (MCT; Chung et al., 2016, 2018; Chung & Abbott, 2021; Cohen et al., 2020), examines a system's capacity to linearly separate neural representational manifolds, and has shown that changes in the manifold's radius, dimensionality, and alignment predict the classification capacity of a system and its ability to generalize (Chou et al., 2025; Li et al., 2024). Both artificial (Cohen et al., 2020) and biological systems show that successful learning of a decision variable alters manifold geometry in a way that increases the system's capacity (Chou et al., 2025) and favors their linear separability (Yao et al., 2023). However, given that changes in manifold capacity have predominantly been observed downstream of the hierarchical process, it is unclear whether the mouse V1 implements similar geometric changes in orientation representations to facilitate perceptual discrimination and generalization.

Methodology

We applied the MCT framework to calcium imaging datasets recorded from the mouse visual cortex V1 (Corbo et al., 2024, 2025). The mice were either trained for a Go/NoGo orientation discrimination task ("trained"; N=10) with progressively increasing difficulty (from 0 to 90 degrees (Fig. 1a), or passively viewed the same stimuli ("naive"; N=7). The effects of training on feature encoding are reported in Fig. 1b.

Neural responses' dimensionality was estimated using principal components analysis (PCA) and by visualizing the data's lower dimensional embedding space with non-metric multidimensional scaling (MDS), t-SNE, and ISOMAP. Changes in the system's representational geometry were quantified via direct estimation of classification capacity and numerical estimation of object manifold geometry (radius and dimension; adapted from Cohen et al. (2020)) .

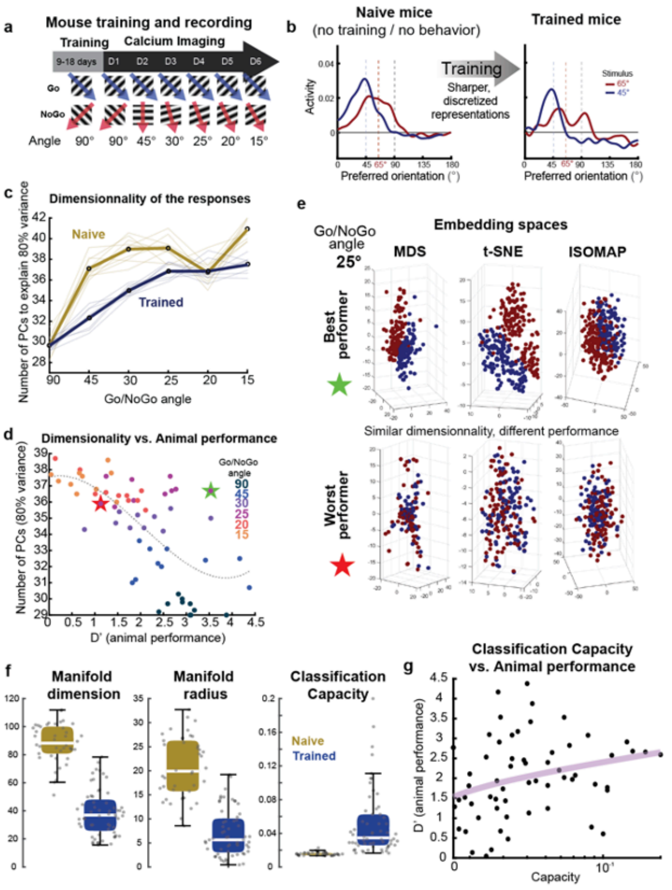

Fig. 1. a: Experimental design. **b**: Results from Corbo et al. (2025) show that after training, the representations in feature encoding space are sharper and the activity is distributed into two discrete domains (e.g. 65° evokes activity at ~45° and 90°) **c**: Dimensionality (number of PCs to explain 80% of the variance) of the V1 evoked responses in naive and trained animals as a function of Go/NoGo angle. **d:** Dimensionality against animal discrimination performance (D'). **e:** The best and worst performers' 25° of Go/NoGo angle representations modeled in MDS, t-SNE, and ISOMAP space. **f:** Manifold dimension, radius and classification capacity for naive and trained mice. **g:** Scatterplot of classification capacity vs. animal performance.

Results

First, to examine the dimensionality of the population representations, we determined the number of principal components necessary to explain 80% of the data's variance. We found that the dimensionality of the responses increased with the similarity of the pair of stimuli in both trained and naive animals (Fig. 1c). Strikingly, the dimensionality was lower in animals performing the task, suggesting that the biological implementation of the task relies on reducing the representations' dimensionality (Fig. 1c).

However, while the dimensionality of responses to visual cues predicted task performance across difficulty levels, differences in dimensionality at the same difficulty level did not explain performance variability (Fig. 1d). Visualizing the best and worst performing mice's go/no-go task data in embedding space showed a clear separation between the representational manifolds for the good but not the bad performers (Fig. 1e), suggesting differences in manifold separability despite similar dimensionality. This separability of manifolds was further evidenced by an increase of the neural manifolds' capacity and a decrease of manifold dimension and radius (Fig. 1f) in trained mice. The increase in capacity was correlated to the behavioral performance of the animals, suggesting that it supports the learning process (Fig. 1g).

Conclusion

Taken together, our results show a relationship between behavioral task performance, representational dimensionality, and manifold separability in the early visual cortex of mice. As the responses' dimensionality decreases, the system's manifold capacity increases, leading to heightened representational separability as measured by numerical estimation of the manifold's geometry. These results confirm that learning alters the geometric properties of early sensory representations as early as in V1, optimizing them for linear readout and improving perceptual decision-making.

References

Cheng, Y.-A., Sanayei, M., Chen, X., Jia, K., Li, S., Fang, F., Watanabe, T., Thiele, A., & Zhang, R.-Y. (2025). A neural geometry approach comprehensively explains apparently conflicting models of visual perceptual learning. *Nature Human Behaviour*, 1–18. https://doi.org/10.1038/s41562-025-02149-x

Chou, C.-N., Kim, R., Arend, L. A., Yang, Y.-Y., Mensh, B. D., Shim, W. M., Perich, M. G., & Chung, S. (2025). *Geometry Linked to Untangling Efficiency Reveals Structure and Computation in Neural Populations* (p.

2024.02.26.582157). bioRxiv. https://doi.org/10.1101/2024.02.26.582157

Chung, S., & Abbott, L. F. (2021). Neural population geometry: An approach for understanding biological and artificial neural networks. *Current Opinion in Neurobiology*, *70*, 137–144. https://doi.org/10.1016/j.conb.2021.10.010

Chung, S., Lee, D. D., & Sompolinsky, H. (2016). Linear readout of object manifolds. *Physical Review E*, *93*(6), 060301. https://doi.org/10.1103/PhysRevE.93.060301

Chung, S., Lee, D. D., & Sompolinsky, H. (2018). Classification and Geometry of General Perceptual Manifolds. *Physical Review X*, *8*(3), 031003. https://doi.org/10.1103/PhysRevX.8.031003

Cohen, U., Chung, S., Lee, D. D., & Sompolinsky, H. (2020). Separability and geometry of object manifolds in deep neural networks. *Nature Communications*, *11*(1), 746. https://doi.org/10.1038/s41467-020-14578-5

Corbo, J., Erkat, O. B., McClure, J., Khdour, H., & Polack, P.-O. (2025). Discretized representations in V1 predict suboptimal orientation discrimination. *Nature Communications*, *16*(1), 41. https://doi.org/10.1038/s41467-024-55409-1

Corbo, J., Erkat, O. B., McClure, J. P., Khdour, H., & Polack, P.-O. (2024). *Discretized representations in V1 predict suboptimal orientation discrimination* (p. 2022.05.13.491867). bioRxiv. https://doi.org/10.1101/2022.05.13.491867

Corbo, J., McClure, J. P., Erkat, O. B., & Polack, P.-O. (2022). Dynamic Distortion of Orientation Representation after Learning in the Mouse Primary Visual Cortex. *Journal of Neuroscience*, *42*(21), 4311–4325. https://doi.org/10.1523/JNEUROSCI.2272-21.2022

Failor, S. W., Carandini, M., & Harris, K. D. (2025). Visual experience orthogonalizes visual cortical stimulus responses via population code transformation. *Cell Reports*, *44*(2). https://doi.org/10.1016/j.celrep.2025.115235

Henschke, J. U., Dylda, E., Katsanevaki, D., Dupuy, N., Currie, S. P., Amvrosiadis, T., Pakan, J. M. P., & Rochefort, N. L. (2020). Reward Association Enhances Stimulus-Specific Representations in Primary Visual Cortex. *Current Biology*, *30*(10), 1866-1880.e5. https://doi.org/10.1016/j.cub.2020.03.018

Jurjut, O., Georgieva, P., Busse, L., & Katzner, S. (2017). Learning Enhances Sensory Processing in Mouse V1 before Improving Behavior. *Journal of Neuroscience*, *37*(27), 6460–6474. https://doi.org/10.1523/JNEUROSCI.3485-16.2017

Li, Q., Sorscher, B., & Sompolinsky, H. (2024). Representations and generalization in artificial and brain neural networks. *Proceedings of the National Academy of Sciences*, *121*(27), e2311805121. https://doi.org/10.1073/pnas.2311805121

Poort, J., Khan, A. G., Pachitariu, M., Nemri, A., Orsolic, I., Krupic, J., Bauza, M., Sahani, M., Keller, G. B., Mrsic-Flogel, T. D., & Hofer, S. B. (2015). Learning Enhances Sensory and Multiple Non-sensory Representations in Primary Visual Cortex. *Neuron*, *86*(6), 1478–1490. https://doi.org/10.1016/j.neuron.2015.05.037

Yao, J. D., Zemlianova, K. O., Hocker, D. L., Savin, C., Constantinople, C. M., Chung, S., & Sanes, D. H. (2023). Transformation of acoustic information to sensory decision variables in the parietal cortex. *Proceedings of the National Academy of Sciences*, *120*(2), e2212120120. https://doi.org/10.1073/pnas.2212120120