Noise Correlations for Feature Learning

Xufeng Caesar Dai (xufeng_dai@brown.edu)
Department of Neuroscience, Brown University, 164 Angell St
Providence, RI 02906 USA

Joonhwa Kim (joonhwa_kim@brown.edu)
Department of Neuroscience, Brown University, 164 Angell St
Providence, RI 02906 USA

Apoorva Bhandari* (apoorva_bhandari@brown.edu)
Department of Cognitive, Linguistic & Psychological Sciences, Brown University, 190 Thayer St
Providence, RI 02912 USA

Matt Nassar* (matthew_nassar@brown.edu)
Department of Neuroscience, Brown University, 164 Angell St
Providence, RI 02906 USA

***co-senior authors**

Abstract:

In real-word learning, individuals continually encounter complex arrays of features, only some of which are crucial to the outcomes they experience. How do they manage to discern which combinations of features are relevant for learning? This study explores how dynamic noise correlations – contextually enhanced correlations in neuronal firing – can focus learning on the most relevant feature dimensions in the current context by leveraging prior experience with these features. Participants were tasked with discriminating multi-dimensional perceptual stimuli under various task conditions that specifically incentivized learning about distinct, combined feature dimensions. We found that people learned preferentially in relevant feature dimensions, but to a degree that differed across individuals. These results motivate ongoing work modeling human subject behavior with neural networks and probing noise correlations in feature representations with fMRI. Our approach provides a window into how adaptive neural mechanisms can enhance the efficiency of learning in complex environments.

Keywords: learning; perceptual learning; feature dimensions; neural networks; noise correlations

Introduction

Learning in high-dimensional environments is inherently challenging, yet people can learn quickly and efficiently, often from minimal data. This raises a critical question: how does the brain adjust the weights connecting the roughly 100 billion neurons in the human brain to learn the right thing from each situation? Human perceptual learning is deeply entwined with the processes by which information is prioritized and encoded by the neuronal population.

While noise correlations can limit the capacity of neural codes by reducing extractable information, they can play a multifaceted role when tuned to varying task structures (Averbeck, Latham, & Pouget, 2006). These correlations vary in behavioral context suggesting that they are adaptive rather than detrimental and can potentially reflect task-driven changes in circuitry (Cohen & Newsome, 2008). More recently, it has been shown that, by focusing neural computation onto task-relevant dimensions, noise correlation can improve learning speed and accuracy (Nassar, Scott, & Bhandari, 2021).

Our study aims to test the noise correlation learning theory by developing a multi-dimensional learning task specifically designed to incentivize and measure learning prioritization across various combined feature dimensions. Here, we report initial results that focus on these first two objectives, setting the stage for our ongoing investigation into how the geometry of neural variability, as measured by fMRI, coordinates learning across relevant feature dimensions.

Method

We employed a dynamic perceptual discrimination paradigm where participants learned to discriminate between multi-dimensional perceptual stimuli: random dot patterns that varied in two dimensions – color (proportion of dots that are orange or purple) and direction of motion (motion coherence upward or downward), as in Figure 1. The study featured two task conditions, each of which required the integration of information from both stimulus dimensions. In each condition, participants viewed a stimulus containing motion and color information and were required to specify one of two possible responses. Within each condition, rules changed occasionally, but always by changing on a fixed feature dimension (ie. rightward/purple). These uncued intra-dimensional shifts involved translational shifts in the learning boundary, requiring them to adapt their decision making within a familiar dimension. These shifts compelled participants to continuously adjust their learning strategies by focusing on the most relevant feature dimension.

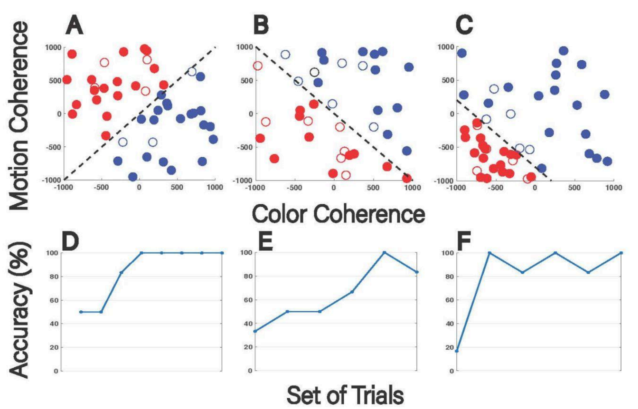

Figure 1: Panel A-C illustrates a participant's responses to three distinct decision boundaries (dashed lines) with solid circles marking correct trials and hollow circles marking incorrect trials. Note that boundary A was encountered in a separate task condition from boundaries in B&C. Panel D-F displays average participant learning curves for these decision boundaries. Notably, the participant's accuracy improved across trials, indicating successful learning.

Results

Our primary goal was to design a task that could measure the degree to which individuals focused learning onto specific task relevant combined feature dimensions. To test our success on this measure we examined behavior with logistic regression that allowed us to assess how color and motion signed coherence values affected choice behavior. We ran this model in sliding windows over time and then examined the change in coefficients from one time point to the next as a proxy for learning. In particular, we interpret the change in coefficients as a learning gradient, providing info not only about whether participants are using a given feature more or less, but also about how participants tended to jointly adjust feature weights.

The left panel in Figure 2 portrays these learning gradients for an individual participant who focused learnings on the relevant features for each condition (blue and red arrows). Note that correlation between the endpoints of the red arrows is positive, whereas the correlation between the endpoints of the blue arrows is negative, signifying differential focus of learning onto feature dimensions corresponding to those shown in figure 1B and 1A respectively. Figure 2B quantifies these correlations for 15 participants (blue points) as well as what would be expected from an agent who performed the task perfectly (red points). Based on the demands of our two learning conditions (Fig 1A&B), we would expect that learning should prioritize the positively correlated color-motion dimension in one condition (y-axis) and the negatively correlated dimension in the other (x-axis). The clustering of participants in the upper-left quadrant suggests that participants are amplifying learning in the dimensions appropriate for each condition.

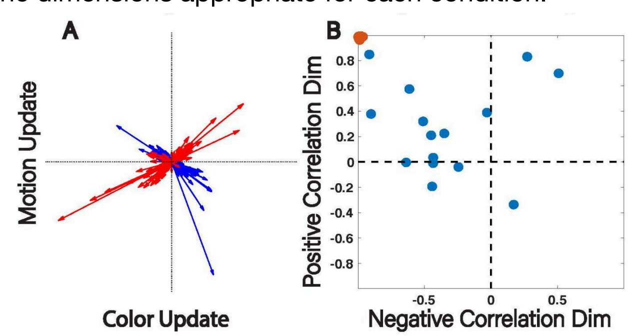

Figure 3: Color and motion learning coefficients were derived using a sliding window with window size of thirty and step size of one. The learning gradient – difference between coefficients across successive windows. Panel A Illustrates how a participant's learning gradients evolve over time. Panel B aggregates all participants' learning gradient correlations.

Conclusion

In conclusion, participants' demonstrated clear learning prioritization and adaptation have highlighted the effectiveness of our approach. Moving forward, we test whether neural variability, as measured by fMRI, drives learning onto specific dimensions.

References

Averbeck BB, Latham PE, Pouget A (2006) Neural correlations, population coding and computation. Nat Rev Neurosci 7:358–366.

Cohen, M. R., & Newsome, W. T. (2008). Context-dependent changes in functional circuitry in visual area MT. *Neuron*, *60*(1), 162–173.

Nassar MR, Scott D, Bhandari A. Noise Correlations for Faster and More Robust Learning. J Neurosci. 2021 Aug 4;41(31):6740-6752. doi: 10.1523/JNEUROSCI.3045-20.2021. Epub 2021 Jun 30. PMID: 34193556; PMCID: PMC8336712.