Research Statement

The goal of my research is to understand how human observers extract meaningful, task-relevant information from complex visual inputs. We are aided in this process by inherent structure in the visual world, where visual features tend to be strongly correlated with the meaning of stimuli. For example, many object categories have stable visual properties (shape, texture, color) that help us infer their identities. At the same time, however, the behavioral relevance of stimuli can change dynamically: how we should interact with each object depends on our internal goals and present task. As a result, operating within the world requires both making inferences based on stable regularities, and allowing for flexibility that supports optimal behavior. My research is motivated by the hypothesis that visual perception is jointly constrained by the statistical structure of the environment and our need for flexible, adaptive behavior.

To address these foundational issues, I use a combination of experimental data collected from humans (functional MRI and behavioral studies) and computational analysis tools such as decoding and encoding models. I also use artificial neural network simulations that complement my work with human data. My research has broad relevance for understanding human visual function and its neural basis in health and disease, as well as for the development of computer vision systems that incorporate principles of human vision. Below, I outline my past work done as a graduate student (UC San Diego; UCSD) and as a postdoc (Carnegie Mellon University; CMU), then discuss future directions for my research group.

How does neural feature selectivity reflect the correlational structure of the environment?

Within visual cortex, a single population of neurons can exhibit selectivity for both low-level visual features (spatial positions, orientations), and high-level semantic properties (object and scene categories). What is the relationship between representations at these distinct levels of complexity? A guiding hypothesis for my research is that the relationship between low- and high-level representations in visual cortex is constrained by natural image statistics. Specifically, high-level categories are characterized by diagnostic sets of low-level features, and these co-occurrence statistics may also be reflected within visual cortex organization. As a postdoc at CMU, my research has worked toward evaluating this hypothesis by better characterizing the fine-grained selectivity of neural populations (i.e., fMRI voxels) within visual cortex for features at multiple levels of complexity. Much of this work uses the computational framework of "encoding models": predictive models that take an arbitrary image viewed by an observer, extract image features using a computational model of interest, and predict the response of a single fMRI voxel as a weighted sum of the model features (see figure, left). A key benefit of this framework is that it allows us to measure feature tuning using responses to complex natural scene images.

In recent work (Henderson, Tarr, & Wehbe, 2023; *J. Neurosci.*), I used this modeling framework to construct encoding models based on **mid-level texture statistics features**. By fitting these models on a large scale naturalistic fMRI dataset (Allen et al., 2022; *Nat. Neurosci.*), I was able to characterize differences in the types of mid-level features to which voxels throughout the visual cortex are sensitive, with feature complexity increasing from early to higher visual areas (see figure, right). We also showed that patterns of sensitivity to texture model features are aligned with broad axes of cortical organization that may reflect semantic scene content. These results have provided a critical step forward in modeling mid-level feature selectivity and its relationship to the higher-level organization of visual cortex.

General modeling framework

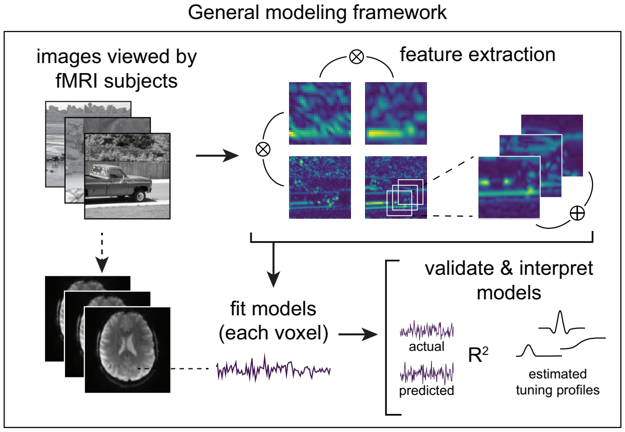

Isolating feature sub-types in the texture statistics model

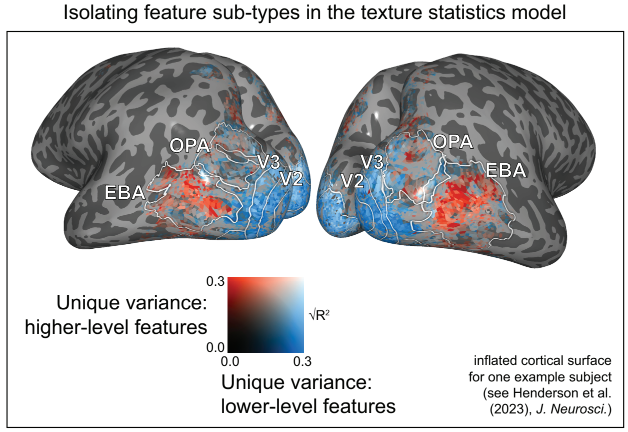

Using this same framework, I've also constructed models based on Gabor features (Henderson, Tarr, & Wehbe, 2023; *J. Vision)*, allowing for mapping of orientation and spatial frequency selectivity across higher visual cortex. As part of this work, I separately performed a large-scale analysis of the statistical correspondence between Gabor features and various high-level semantic categories across a bank of natural scene images. These analyses showed that voxels within category selective visual areas (e.g., face-selective and scene-selective areas) exhibit low-level feature and spatial tuning properties that match the actual **image statistics of their preferred high-level categories**. These results support the hypothesis that natural image statistics may constrain the organization of higher visual cortex.

In addition to modeling low and mid-level visual features, other work from my postdoc has aimed at better understanding selectivity for high-level properties of objects and scenes, with a focus on how such high-level representations may be related to the low-level visual properties of images. In this area, I've worked on evaluating how **real-world object size** is computed within the brain from natural images (Luo, Wehbe, Tarr, & Henderson, 2023; *bioRxiv)*, and have been involved in other work establishing **selectivity for food** images within ventral visual cortex (Jain et al., 2023, *Communications Biology*). Finally, in recent work led by a CMU graduate student, we have developed novel computational tools that can be used to **synthesize an optimal stimulus** for a group of voxels (BrainDiVE; Luo, Henderson, Wehbe, & Tarr, 2023; *NeurIPS*), and generate **natural language captions** that describe the selectivity of visual cortex voxels (BrainSCUBA; Luo, Henderson, Tarr, & Wehbe, 2023; *arXiv*). These approaches provide a powerful new way of characterizing functional selectivity, and will be built upon in ongoing work.

How do task goals impact neural representations of objects and features?

In addition to characterizing visual representations at various levels of complexity, another arm of my research focuses on how these representations are modified to support task goals. In everyday tasks, we are often required to selectively attend to relevant features of objects while ignoring irrelevant features. What neural mechanisms allow us to flexibly represent the most relevant sensory information? During my graduate work at UCSD, I addressed this question using fMRI along with computational analyses to measure visual representations while human subjects performed demanding cognitive tasks.

In one experiment, we explored how performing a **dynamic categorization task** that required switching between orthogonal category boundaries impacts visual representations of two-dimensional shapes (Henderson, Serences, & Rungratsameetaweemana, 2023, *bioRxiv*). We established that shape representations in early visual cortex can be warped during categorization, such that shapes become more discriminable across the currently active boundary. Moreover, these effects were associated with behavioral performance, suggesting a role for early visual cortex in the flexible categorization of stimuli.

In another experiment, I developed a set of three-dimensional novel objects to investigate how task goals impact the representation of **high-level object properties** like identity and viewpoint (Henderson & Serences, 2019, *J. Neurophys.*). Using multivariate decoding, I identified neural populations in frontal and parietal cortex that encoded information related to relevant object dimensions.

Flexible coding of visual information may also impact the mechanisms of information storage in short-term memory. For example, memory tasks requiring precise visual detail may engage early visual cortex, while more abstract tasks may rely more on frontal or parietal regions. I tested this hypothesis using a **spatial short-term memory** task where response predictability was manipulated on each trial (Henderson, Rademaker, & Serences, 2022; *eLife*). We demonstrated that across task conditions, neural representations of a remembered spatial position in early visual cortex appeared to "trade off" with representations of a planned upcoming response in motor-related areas. This indicates that multiple neural coding schemes can be adaptively engaged to support short-term memory.

What role do low-level image statistics play in learning and representations?

Another goal of my research is to understand how low-level visual features interact with learning of high-level information. How do environmental statistics impact how and what we learn? To assess these questions, I use large-scale simulation studies performed using artificial neural networks (ANNs). Rather than taking these models to be perfect analogues for human perception, my approach uses ANNs to explore theoretical ideas about the relationship between stimulus statistics, categories, and representations, which can be used to generate new hypotheses for experimental studies.

While at UCSD, I used this approach to examine the relationship between **orientation tuning biases** (i.e., over-representation of vertical and horizontal orientations) in the visual system and the distribution of orientations in natural images (Henderson & Serences, 2021; *J. Vision*). By training ANNs

(VGG-16) on datasets of images that had been rotated by varying amounts, I demonstrated a link between orientation biases in a model's training set and the representation of orientation that was learned. This work provides support in favor of an efficient coding theory of orientation representation, whereby the brain represents frequently-encountered features with higher fidelity than infrequent features.

I've also used this ANN-based approach to explore how low-level image statistics impact the ease of learning a categorization task. While at CMU, I mentored an undergraduate who was investigating whether poor spatial acuity at the earliest stages of human development might benefit early category learning. By training ANNs (ResNet-50) on low-pass filtered (blurry) and intact images, she demonstrated that initial experience with **blurry images improved basic-level categorization** performance (Jinsi*, Henderson*, & Tarr, 2023; *PLOS ONE*). These results will be built upon in future experimental work.

Ongoing work: How do low-level features contribute to categorization task performance?

If low-level visual features are statistically associated with high-level categories, then observers may be able to infer the category of objects based on their low-level visual properties alone. In ongoing work, I am experimentally testing this idea using a computational texture synthesis algorithm which allows the generation of scrambled images with varying feature complexity (Gatys et al., 2015; *NeurIPS*). Using behavioral and imaging experiments, I am assessing how category representations are formed across the visual hierarchy, and how the formation of category representations depends on the type of task an observer is performing (coarse vs. fine-grained object categorization).

Future work

As an independent investigator, I will build on each of the themes I have explored in previous work, continuing to utilize my core computational and experimental skill set: collection and modeling of human fMRI and behavioral data, as well as simulations using ANNs. The long-term projects that my lab will pursue are organized around three key research aims.

(1) Understand how visual cortex tuning properties reflect the statistics of natural images

In my postdoctoral work, I have made progress on understanding how the correlational structure of the environment is reflected in visual cortex tuning properties. Future projects will build on the computational methods that I developed in this work. As one example, I plan to incorporate additional mid-level features such as color and contour into my computational modeling framework, and identify how selectivity for these features in the brain is related to the distribution of colors and/or contours in natural images.

(2) Characterize how natural image statistics guide cognitive task performance

Much of what is known about attentional guidance is in the context of simplified laboratory tasks, where participants are given explicit instructions about the relevance of a feature or spatial position. In the real world, however, our attention is more often guided by abstract goals, like searching for a high-level category of object or assessing the layout of a scene. For such abstract tasks, attention must be guided by implicit knowledge of what properties are most relevant, based on our experience with the structure of natural images. How does this kind of naturalistic attentional selection impact sensory representations in the visual system? Do attentional modulations accurately reflect the structure of the environment, or are there limits to this correspondence? To examine these and related questions, my lab will use task-driven functional MRI experiments and computational modeling.

(3) Determine how low-level image statistics impact learning of high-level information

To understand how environmental statistics impact learning and representations in the human visual system, my lab will use a combination of ANN simulation studies and human behavioral experiments. One long-term approach will be to develop new datasets that disrupt the typical correspondence between low-level and high-level properties within natural images. We will use these image sets in both experimental and modeling studies, and make them openly available for the broader research community.

Potential for inter-disciplinary collaborations*.* My research is motivated by broad theoretical ideas on the relationship between sensory and cognitive processing in the human brain, and incorporates multi-disciplinary methods. Thus, there are avenues for collaboration with researchers from multiple subfields of psychology, neuroscience, and cognitive science, as well as machine learning and statistics.

Open science and data sharing. I am committed to making my experimental data and code openly available -- all datasets and research code from my previous projects are available on Open Science Framework (https://osf.io/v8b2r/) and GitHub (https://github.com/mmhenderson).