

ARTICLE

An In-depth Exploration of the Interplay between fMRI Methods and Theory in Cognitive Neuroscience**Derek J. Huffman***Department of Psychology, Colby College, Waterville, ME 04901.*<https://doi.org/10.59390/ZABM1739>

Functional magnetic resonance imaging (fMRI) has been a cornerstone of cognitive neuroscience since its invention in the 1990s. The methods that we use for fMRI data analysis allow us to test different theories of the brain, thus different analyses can lead us to different conclusions about how the brain produces cognition. There has been a centuries-long debate about the nature of neural processing, with some theories arguing for functional specialization or localization (e.g., face and scene processing) while other theories suggest that cognition is implemented in distributed representations across many neurons and brain regions. Importantly, these theories have received support via different types of analyses; therefore, having students implement hands-on data analysis to explore the results of different fMRI analyses can allow them to take a firsthand approach to thinking about highly influential theories in cognitive neuroscience. Moreover, these explorations allow students to see that there are not clearcut “right” or “wrong”

answers in cognitive neuroscience, rather we effectively instantiate assumptions within our analytical approaches that can lead us to different conclusions. Here, I provide Python code that uses freely available software and data to teach students how to analyze fMRI data using traditional activation analysis and machine-learning-based multivariate pattern analysis (MVPA). Altogether, these resources help teach students about the paramount importance of methodology in shaping our theories of the brain, and I believe they will be helpful for introductory undergraduate courses, graduate-level courses, and as a first analysis for people working in labs that use fMRI.

Key words: functional magnetic resonance imaging (fMRI); functional specialization; functional localization, distributed representations; multivariate pattern analysis (MVPA); machine learning; free open-source software (FOSS); high-level vision; object recognition; tutorial

The advent of noninvasive brain recording techniques revolutionized the study of the human brain. For example, noninvasive methods allow researchers to study large samples of participants as they dynamically engage in behavioral tasks, which stands in contrast to the relatively static methods of studying patients with damage to their brains, which was the approach in classical cognitive neuropsychology. Initially, many of the early papers using methods such as positron emission tomography (PET) and functional magnetic resonance imaging (fMRI) sought to bolster the findings from patients with localized brain damage. For example, researchers aimed to ascertain whether the specific behavioral changes that were observed in patients with localized lesions were similarly localized in neuroimaging approaches (e.g., language and Broca’s area [left inferior frontal gyrus]) or whether a more distributed network of regions played a role. Interestingly, even early studies that used methods that would become the precursor to PET imaging revealed that a broader network of regions seemed to play a role in cognitive functions, such as language (e.g., the right hemisphere and regions outside of the regions defined by neuropsychology: Ingvar and Schwartz, 1974). Furthermore, early studies also revealed that localized brain damage could give rise to broader network-level changes in the brain, thus suggesting that lesions cause widespread network changes rather than solely affecting the lesioned location (e.g., Cronqvist et al., 1965; for more recent and similar findings with fMRI see:

Gratton et al., 2012; Henson et al., 2016; for recent computational modeling to support these empirical findings see: Alstott et al., 2009). On the other hand, convincing evidence emerged that certain classes of stimuli tended to activate specific brain regions (e.g., faces: Kanwisher et al., 1997; e.g., scenes: Epstein and Kanwisher, 1998). Therefore, a decades-long debate has emerged with one set of prominent theories suggesting that the brain contains functionally specialized modules (Fodor, 1983; Epstein, 2005; Reddy and Kanwisher, 2006; Kanwisher, 2010, 2017) while an alternative set of prominent theories suggests that the brain uses distributed representations in which groups of neurons and brain regions work together in the service of cognition (e.g., distributed representations: McClelland, 1986; Rumelhart et al., 1986; Rumelhart and Todd, 1993; McClelland and Rogers, 2003; Haxby et al., 2001; Haxby, 2012; e.g., the brain-as-a-network theory: Sporns, 2010; Bullmore and Bassett, 2011). If students implement hands-on data analysis to see the contradictory results between these frameworks on the same fMRI data, it will allow them to see that our methods are effectively an instantiation of our assumptions of how the brain works, thus clearly demonstrating that “*what we observe is not nature itself, but nature exposed to our method of questioning*” (Heisenberg, 2007). Moreover, such inquiry-based learning may increase student engagement, retention, and learning, as in other areas of science (Lopatto, 2007; Russell et al., 2007; Rodenbusch et al., 2016; Brabec et al., 2018).

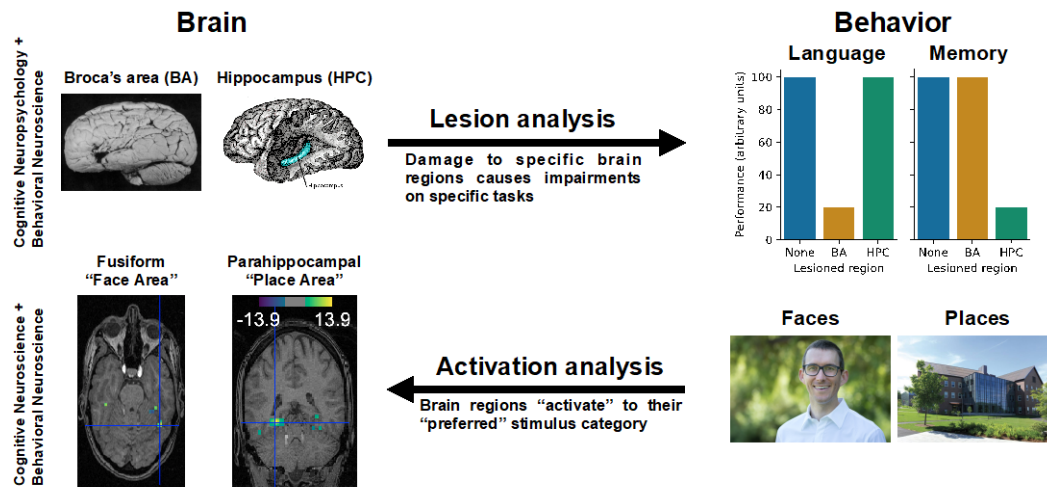


Figure 1. Classic approaches to neuroscience include various methods of attempting to find evidence of localization of function via double dissociation. In human studies these approaches have included neuropsychology and activation analysis, thus these methods can be seen as two sides of the same coin. The classic method of neuropsychology attempts to find specific behavioral changes following localized brain damage (i.e., the arrow goes from the brain to behavior; note that these methods are common throughout neuroscience, e.g., lesion and optogenetic approaches in nonhuman animals). The classic neuroimaging method of activation analysis attempts to find differences in the brain regions that are activated by different behavioral tasks (i.e., the arrow goes from behavior to the brain; note that these methods are common throughout neuroscience, e.g., looking for increases in neural activity within specific brain regions using single unit electrophysiology in nonhuman animals).

Searching for Evidence of functional specialization

Studies in the late 1990s tested the theory of functional specialization for high-level object categories using **activation analysis**. Using this approach, which is essentially an extension of classical neuropsychology, researchers aimed to determine if specific regions of the brain were more active in response to one category or stimulus feature than others (Figure 1). When it is applied to fMRI data, activation analysis typically assesses whether the amount of blood-oxygenation-level dependent (BOLD) activity is greater for one task condition than another, and these techniques typically employ assumptions of spatial clustering by looking for clusters of voxels in nearby brain regions. These techniques are typically referred to as mass univariate approaches (see Chapter 5 of Poldrack et al., 2011) because researchers employ a univariate analysis (e.g., t-test, ANOVA) on all of the voxels within a predefined region of interest (e.g., the whole brain or a subset of brain areas of interest).

As one example of this approach, the discovery of the Fusiform Face Area (FFA; Kanwisher et al., 1997; McCarthy et al., 1997) and the Parahippocampal Place Area (PPA; Epstein and Kanwisher, 1998) provided strong evidence to support the notion that the brain contains functionally specialized regions for processing high-level visual information (i.e., these studies suggest that different brain regions process different categories of visual stimuli). Moreover, the findings from these studies dovetailed nicely with studies of patients with localized damage to their brain. For example, patients with prosopagnosia have sometimes been shown to have a relatively selective impairment in their ability to recognize faces with a relatively intact ability to recognize other object categories (e.g., a patient with

prosopagnosia that could name the make and model of several of a collection of toy cars: Sergent and Signoret, 1992; but note that other accounts have shown a more domain general deficit in face and object recognition: Barton and Corrow, 2016; Geskin and Behrmann, 2018). Therefore, taken together, one set of prominent theories (Fodor, 1983; Epstein, 2005; Reddy and Kanwisher, 2006; Kanwisher, 2010, 2017) argue that the human brain contains specialized regions for processing specific high-level categories and these theories continue to be advanced based on the results of activation analysis.

Searching for Evidence of Distributed Representations

In the late 1990s and early 2000s, fMRI researchers began to test an alternative theory of neural processing: the theory of **distributed representations**. These studies grew out of a rich literature in cognitive science, including neural network modeling in which researchers theorized that neural information is contained within coarse-coded distributed representations (e.g., McClelland, 1986; Rumelhart et al., 1986; Rumelhart and Todd, 1993). Specifically, rather than having specialized regions that process certain categories of stimuli, neurons can be coarsely tuned to represent various stimulus features and the brain or organism can make sense of the neural information by assessing the similarity of the patterns of activity between different classes of stimuli or events. For example, in a distributed representation, the pattern of activity in response to two images of scenes should be more similar than between a scene and a face (Figure 2), however the units (e.g., neurons or in the case of fMRI, voxels [short for volume element], the smallest unit of fMRI data) need not be clustered together (e.g., they can be spatially distributed)

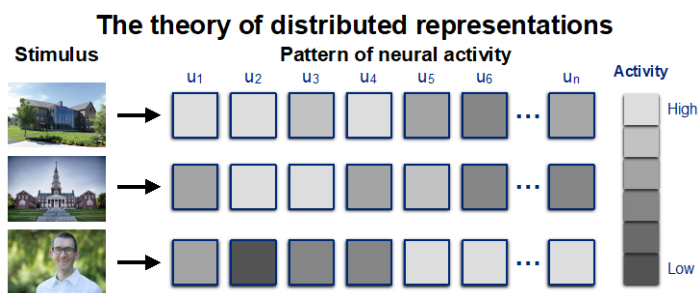


Figure 2. Multivariate pattern analysis (MVPA) enables researchers to test the theory of distributed representations that posits that, rather than being comprised of specialized modules, units (e.g., neurons, voxels) can be coarsely tuned to carry information about several categories. MVPA has been used extensively in neuroscience lately (e.g., for fMRI and EEG analysis in cognitive neuroscience and “ensemble analysis” in multi-unit electrophysiology in nonhuman animals: Eichenbaum and Davis, 1998). The example here depicts that stimuli within a category would elicit a similar pattern of activity across units but relatively distinct patterns of activity in response to stimuli in different categories. Note about abbreviation: u =unit (e.g., neuron, voxel).

nor solely tuned for a specific stimulus feature (e.g., units can respond to different categories).

To test the theory of distributed representations, researchers employed multivariate pattern analysis (MVPA), which allowed them to study patterns of activity across units (voxels for fMRI) rather than looking for localized changes in activity (Edelman et al., 1998; Haxby et al., 2001). Results of these experiments provided evidence that high-level visual cognition is supported by distributed representations through two observations. First, much of the ventral visual stream showed changes in activation (e.g., increases or decreases in BOLD activity) in response to various image categories (Haxby et al., 2001). Second, when Haxby et al. (2001) applied MVPA to the regions that responded maximally to one category (e.g., the putative FFA or PPA), they found that patterns of activity in these voxels could be used to distinguish between the other categories of images, thus suggesting that these regions are not specialized to exclusively carry information about the category of stimulus to which they respond maximally. Similarly, when Haxby et al. (2001) excluded the voxels that responded maximally to any given category, their overall results were the same, thus suggesting their results were not driven by localized regions that preferentially represent a single category. Therefore, these results provided early evidence that high-level categories are coded in distributed representations across the ventral temporal cortex. Altogether, MVPA provided a principled method for testing the theory of distributed representations (McClelland, 1986; Rumelhart et al., 1986; Haxby et al., 2001; Haxby, 2012; Rissman and Wagner, 2012).

The Importance of Learning Programming For Students In Neuroscience

Cognitive neuroscience is a highly interdisciplinary field that has come to rely heavily on computational approaches, however the home departments for cognitive neuroscience classes do not typically emphasize coursework in computer

science (e.g., computer science courses are not required in my home department of psychology; Juavinett, 2020, 2022; Ho et al., 2021). Therefore, there is a huge divide between the skills that students need to know and what they are actually learning. Specifically, I do not know any cognitive neuroscience labs that do not make active use of computer science in all aspects of their projects, such as task development (e.g., showing participants images or movies with toolboxes such as `PsychoPy` in Python: Peirce, 2007; Peirce et al., 2019; or `PsychToolbox` in MATLAB: Brainard, 1997), the creation of more immersive VR tasks with video game engines (e.g., we have worked to develop the Landmarks package for the Unity game engine: Starrett et al., 2020), data collection (e.g., writing scripts to interact with fMRI or EEG hardware; e.g., Lab Streaming Layer [liblsl]: Stenner et al., 2023), data preprocessing (e.g., cleaning fMRI or EEG data; e.g., `Nipype`: Gorgolewski et al., 2011), data analysis (e.g., the Python package `Nilearn` for fMRI data: Abraham et al., 2014; `MNE-Python` for EEG data: Gramfort et al., 2013; the Python package `pandas`: McKinney, 2010) and figure generation (e.g., the Python packages `matplotlib`: Hunter, 2007; and `seaborn`: Waskom, 2021; see Table 1). If we want to teach students about what it is like to be a cognitive neuroscientist, we should expose them to tools that we use in our field, similar to students in art getting the opportunity to practice picking up a paintbrush if they want to become studio artists working on paintings.

There is currently a missed opportunity for providing a more inclusive training opportunity in computational approaches for students across the curriculum (e.g., the Computing in Undergraduate Education [CUE] initiative of the National Science Foundation [NSF], e.g., NSF Award 1935099). For example, students might not even know that they are interested in or capable of becoming a programmer until they have learned about it and seen it in practice within a specific discipline in which they are interested. Like many other practicing neuroscientists, I did not learn how to program until I was a Ph.D. student. Thus, we can increase the diversity of students that go on to use programming by providing opportunities for them to meaningfully engage with programming within a topic that they are interested in studying (e.g., if a student signs up for our course in cognitive neuroscience, then it means that they hopefully have at least a surface-level interest in the topic; Juavinett, 2020). Accordingly, implementing these approaches early in the curriculum provides an opportunity for students to change or extend their course of study for the remainder of their time in college and beyond. Furthermore, early exposure within the undergraduate curriculum provides an opportunity for students to get a more in-depth understanding and proficiency of concepts in higher-level cognitive neuroscience courses and it may increase student engagement and retention, as in research experiences in other areas of science (e.g., Rodenbusch et al., 2016).

Python has firmly established itself as one of the most popular programming languages over the past decade and it offers several features that are attractive for cognitive neuroscience. First, Python is part of the free and open-source software (FOSS) community, thus making it ideal for

Package	Primary use
BrainIAK (Kumar et al., 2022)	fMRI analysis
Matplotlib (Hunter, 2007)	Figure creation
MNE-Python (Gramfort, 2013)	EEG analysis
Nilearn (Abraham et al., 2014)	fMRI analysis
Nipype (Gorgolewski et al., 2011)	fMRI analysis
NumPy (Harris et al., 2020)	Data container
pandas (McKinney, 2010)	Data container
PsychoPy (Peirce, 2007; Peirce et al., 2019)	Stimulus presentation
scikit-learn (Pedregosa et al., 2011)	Data analysis
SciPy (Virtanen et al., 2020)	Data analysis
Seaborn (Waskom, 2021)	Figure creation
statsmodels (Seabold and Perktold, 2010)	Data analysis

Table 1. An overview of valuable Python toolboxes that showcase the benefit of teaching students Python for learning skills for cognitive neuroscience. Teaching students even a little bit of Python programming can go a long way toward helping them learn how to implement all aspects of an experiment, from stimulus presentation (e.g., PsychoPy) to cutting-edge analyses within modern Python toolboxes (e.g., Nilearn, MNE-Python).

using in the classroom, where you no longer have to worry about licenses and expensive software. Moreover, the FOSS ecosystem makes it possible to provide more open and transparent practices in cognitive neuroscience (e.g., sharing our code for data analysis, writing pipelines that will generate all of our figures for a paper) (Nichols et al., 2017). Second, Python has a relatively simple syntax that is an ideal choice for beginner programmers. Third, Python has an extensive set of add-on toolboxes, many of which are specifically important for cognitive neuroscience (Table 1). Specifically, `scikit-learn` (Pedregosa et al., 2011) is one of the leading machine learning toolboxes and Nilearn (Abraham et al., 2014) provides a wrapper for fMRI analysis with many of the `scikit-learn` packages and it provides a collection of functions that are useful for both activation analysis and MVPA. Thus, by teaching students about two of the broad classes of fMRI analysis in these notebooks, educators can also provide a framework for them to understand the more general syntax of running machine learning analysis with `scikit-learn` and Python. Finally, Python is a very marketable skill and thus teaching students about it in your class can make them more marketable for future research or job opportunities (see DISCUSSION).

In the remainder of this paper, I describe a collection of resources that I developed for teaching students how to conduct two of the main analytic techniques for fMRI—activation analysis and multivariate pattern analysis—using Python. Following these Python notebooks, I assign a writing assignment in which I ask students to critically reflect on the disparate results across these two analyses to discuss how methods can guide our theories of the relationship between the brain and high-level cognition. Below, I also provide suggestions for incorporating these Python notebooks into your curriculum.

MATERIALS AND METHODS

Course Overview

I developed these activities for my introductory cognitive

neuroscience course at Colby College, a small liberal arts college. Our semesters are 13 weeks of instruction plus an additional week for final exams, resulting in 25 class meetings of 1.25 hours each. I have used the assignments that I describe here for 3 semesters. The course cap is 35 students and I generally have a full enrollment. The only prerequisite for this class is our Introduction to Psychology course, where students receive a brief overview of the relationship between the brain and behavior, but there is no specific cognitive neuroscience prerequisite. Additionally, and importantly for the purpose of the materials that I describe here, I do not expect students to have any experience with programming or statistics, and the majority of students in my class have no prior programming experience. Thus, I believe you, like me, can implement these assignments at the level of an introductory neuroscience course but I also think they will be useful in upper-division courses, graduate-level courses, and as a first analysis for people working in labs that use fMRI. For more information, please see **Implementing this Module into your Curriculum**.

The learning goals for the course mirror the learning goals for this lesson plan:

1. Describe the scientific method through the lens of cognitive neuroscience.
2. Apply the critical thinking skills of science to new areas of research and analysis.
3. Discuss and critique methodological approaches to cognitive neuroscience, including the strengths and limitations of each method.
4. Evaluate our knowledge (and the limits of our knowledge) of the connection between the brain and high-level cognitive functions.

Thus, our primary goal in this course is to learn *how to think like a cognitive neuroscientist*.

I include both aggregate and specific comments from students in the RESULTS section. The inclusion of the data (from a homework assignment and anonymous course evaluations) was approved by the Institutional Review Board at Colby College.

A Brief Introduction to the Dataset

I use the Haxby et al. (2001) dataset, which is readily available via Nilearn (Abraham et al., 2014). The experiment consisted of showing participants images from 8 categories: faces, houses (i.e., “places”), cats, bottles, scissors, shoes, chairs, and phase-scrambled images of the objects. The object categories were presented in a block design that alternated between 12 seconds of “rest” and 24 seconds of viewing images from a given category. Participants viewed images from all 8 categories within each run and the order of categories was randomized within each run. Each participant completed 12 runs.

I focus the analysis on a single participant because the papers that discovered FFA (Kanwisher, 2010) and PPA (Epstein and Kanwisher, 1998) analyzed data within each participant separately (i.e., having students inspect the resultant activation maps from a contrast of faces vs. the other seven categories as well as “places” vs. the other

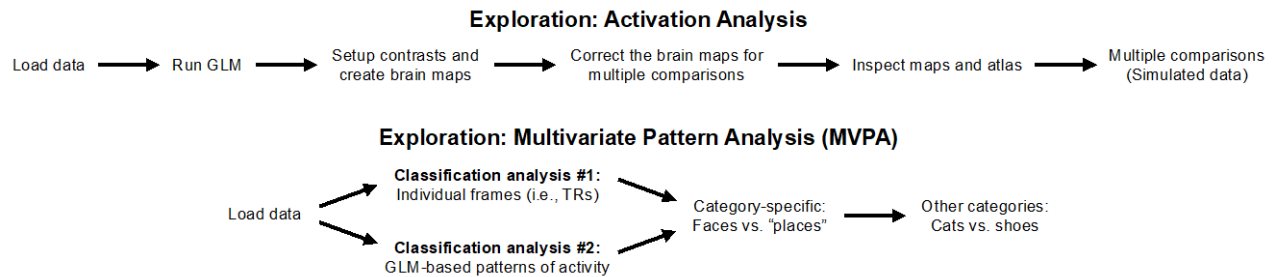


Figure 3. Analysis flow diagram for the two Python notebooks. In the Exploration: Activation Analysis, students learn about how to conduct an activation analysis for a single participant to test whether specialized regions process information about faces and places. They also learn about the importance of correction for multiple comparisons within cognitive neuroscience. In the Exploration: Multivariate Pattern Analysis (MVPA), students test whether voxels that respond maximally to faces or places (i.e., the putative fusiform face area and parahippocampal place area, respectively) also process information about non-preferred categories (here, cats and shoes). Please see the MATERIALS AND METHODS section for more details. Notes about abbreviations: GLM=general linear model, TR=repetition time.

seven categories allows them to take the same approach that the researchers did when they discovered these regions). Additionally, for the purpose of my class, focusing the analysis on a single participant allows me to teach students the general approach for both activation analysis and MVPA. It would be easy to extend these notebooks to include more participants (e.g., by changing the participant that we analyze or by looping the analysis over all of the participants) or to run group-level statistical analysis (e.g., using `Nilearn` functions such as `SecondLevelModel` [for activation analysis] or custom-written code [e.g., to average classification accuracy and compare it to chance, via t-tests or permutation tests, for multivariate pattern analysis]). These additional analyses would be well suited for an upper-division undergraduate course or a graduate-level course.

The Python Notebooks

In these Python notebooks, I use `Nilearn` (Abraham et al., 2014), `pandas` (McKinney, 2010), `NumPy` (Harris et al., 2020), `matplotlib` (Hunter, 2007), and the freely available Haxby dataset (Haxby et al., 2001) to implement two of the main analysis frameworks for fMRI, activation analysis and MVPA, to explore the classic findings that were used to advance the theory that the brain supports high-level visual cognition via functionally specialized regions (Kanwisher et al., 1997; Epstein and Kanwisher, 1998) as well as the opposing findings that support the theory of distributed representations (Haxby et al., 2001). I use `Nilearn` tutorials as starting points for writing these notebooks (Nilearn Contributors et al., 2023), but I heavily modified and extended these examples to allow undergraduate students to recreate the classic discovery of the FFA (Kanwisher et al., 1997), the PPA (Epstein and Kanwisher, 1998), and the initial challenge to these findings that suggests that high-level category representation is instead supported by distributed representations (Haxby et al., 2001). Thus, while there are other resources for learning fMRI analysis (e.g., Nilearn Contributors et al., 2023; Kumar et al., 2020; Jahn et al., 2022), the difference with the tutorials here is that I focus on showcasing the dynamic interplay between our methods and theories of the brain (i.e., rather than focusing on learning methods per se). Moreover, these notebooks

can be implemented in a short period of time, thus allowing integration in an introductory course.

The Python notebooks (i.e., `Jupyter Notebooks`) can be downloaded from my GitHub repository (also see the wiki) and run either within `JupyterHub` (e.g., via your institution), a local installation of `Jupyter Notebooks` (e.g., Anaconda), or without the installation of any software via Google Colaboratory. If you decide to run these notebooks locally via Anaconda, then I recommend that you create a new virtual environment and then you can install `Nilearn` within your new environment (please see this page of the GitHub repository for more details on this approach). Running things via Google Colaboratory is the easiest option in terms of software installation (i.e., there is no installation required); however, this will be the slowest approach for your students. Thus, if possible, I would recommend the other two approaches (e.g., you could use an on-campus computer lab in which you install Python and `Nilearn` for students to run these exercises), but the Colab approach could work well in many cases (e.g., if you have a large class and you have minimal resources via your IT department). All three options are completely free and I have verified that all they all work across multiple operating systems, etc.

In addition to our in-class discussions and the Python notebooks, I provide video walkthroughs of both of the Python notebooks. I find these to be a helpful resource for students because they can pause the videos, rewatch parts that they find difficult, etc. In a nutshell, the videos allow students to work on the notebooks at their own pace and I provide detailed information about the logic of the analysis as well as more background and explanation about each line of code in the notebook. I have also made these videos freely available, and you can find more information about the videos within the GitHub repository.

Using Activation Analysis to Test the Theory of Functional Specialization

In the first notebook, Exploration: Activation Analysis, I teach students how to employ activation analysis to show how it can be used to test the theory of functional specialization. This notebook focuses on having students implement an

analysis to find evidence for the FFA and the PPA. Here, I provide a brief overview of the main steps that I implement in the notebook and I provide some rationale for why I implement each part of the analysis.

The first step of the activation analysis is to load all of the relevant fMRI and behavioral data (i.e., regarding stimulus timing) so that we can have all of the requisite data for analysis. I make use of `Nilearn` functions for loading the data and `pandas` for loading the stimulus timing information (i.e., a `pandas DataFrame` that contains the onset and duration of each stimulus within each run).

The second step is to run a general linear model (GLM) analysis to elucidate the BOLD response to each condition. Here, I use the `FirstLevelModel` function from `Nilearn` to run the analysis. This function also allows you to specify nuisance regressors (i.e., terms in the regression analysis that attempt to deal with noise; e.g., as a form of preprocessing). Here, students will implement a 3rd-order polynomial drift (i.e., this includes three terms: a linear, quadratic, and cubic function) and a high-pass filter of 0.01 Hz, both of which attempt to attenuate low-frequency noise (e.g., caused by scanner drift over the course of each run). You also pass other key information to this function, including the repetition time (TR; i.e., the time between frames; for this dataset, the TR was 2.5 seconds), the type of hemodynamic response function model that you want to run (here, I used the `'spm'` option), as well as other information (which you can see in the notebook). Following the typical syntax of `scikit-learn` (Pedregosa et al., 2011), in `Nilearn` you first set up the parameters of your model and then you can fit that model via object-oriented programming (e.g., here we set up a variable called `fmri_glm` to be the output of our call to the `FirstLevelModel` function, and then we can run the model by calling `fmri_glm.fit()` along with the input of our fMRI data and the stimulus timing information). After fitting the model, I ask students to inspect and interpret the design matrices (i.e., the columns of the design matrix are all of the variables that we are modeling [the task conditions as well as nuisance regressors] and the rows indicate the times [i.e., the frame number]) by viewing them via `matplotlib`, which is a key step in fMRI data analysis.

The third step is to set up contrasts to determine which voxels are more active in response to faces vs. the other 7 categories (i.e., to look for the FFA) as well as to houses vs. the other 7 categories (i.e., to look for the PPA), and the idea is to set up contrasts that sum to zero. Contrasts are common in statistical analyses, and they create a straightforward method for determining if the mean of one condition (or group of conditions) differs from the mean of another condition (or group of conditions). Here, we are specifically interested in determining whether the mean BOLD activity is greater for faces vs. the other visual categories as well as for houses vs. the other 7 categories. We can create these two contrasts by setting up 1-D `NumPy` arrays with the column of interest to a value of 7 and the columns for the other 7 categories to a value of -1. We also set the 3 drift and the constant columns to be 0 for the contrast (i.e., to effectively ignore these from the analysis).

For example, here is the contrast for faces vs. the 7 other categories:

Contrast of Faces vs. other categories

	<code>[-1, -1, -1, +7, -1, -1, -1, -1, 0, 0, 0, 0]</code>
Bottles	
Cats	
Chairs	
Faces	
Houses	
Scissors	
Scrambled	
Shoes	
Drift 1	
Drift 2	
Drift 3	
Constant	

The sum of the contrast is zero ($7 + (-1 * 7) + (0 * 4) = 0$). We code these contrasts by inspecting the design matrix columns to determine which columns correspond to faces (for the faces contrast) or houses (i.e., for the “places” contrast). We compute the contrast using the `Nilearn` method `compute_contrast` (we can implement using our existing `fmri_glm` object via object-oriented programming: `fmri_glm.compute_contrast()` with the input of the `contrast`) and I specify the output to be a z-score for follow-up statistical analysis. Next, I have students generate brain maps by thresholding the z-score map to see which voxels had a contrast that was significantly different than zero. First, we create uncorrected maps by setting a voxel-wise alpha threshold of 0.001 via `Nilearn`'s `threshold_stats_img` (i.e., the threshold will only display voxels that are significantly different than zero, based on a voxel-wise $p < 0.001$). I then have students save the image to an html file so that they can view it later (see step 5) in a 3-D viewer within a web browser to compare it to an atlas of Brodmann areas.

The fourth step is for students to learn about how to apply correction for multiple comparisons to the brain maps. While there are many approaches to correcting for multiple comparisons with fMRI data (e.g., cluster-level thresholds; e.g., see Chapter 7 of Poldrack et al., 2011), for the purpose of this notebook and for my class (e.g., simplicity), I use the false-discovery rate (FDR) method by again using the `Nilearn` function `threshold_stats_img` with an additional input of `'fdr'` for `height_control`.

The fifth step is for students to look at an atlas to determine the names of regions with activated voxels from the faces contrast as well as the “places” contrast. Here, I have students use the `BiolImage Suite mni2tal` webpage (Lacadie et al., 2008). For instructors, please note that the data are in the participant's native space, thus the coordinates within the `BiolImage Suite` page do not map on to the coordinates of the brain maps that we generate in the notebooks. Rather, I ask students to look for the general anatomical features between scans to try to identify the locations within the atlas (e.g., using landmarks such as the lateral ventricles, hippocampus, thalamus, and cerebellum can be useful as they explore the brain maps and the atlas).

In the final part of the notebook, I walk students through the importance of the correction for multiple comparisons by simulating the number of false positives that we would observe with totally random data. Here, the idea is to show students that false positives are an intended consequence of the assumptions of frequentist statistics (e.g., as employed in a t-test). In addition to elucidating the importance of correction for multiple comparisons, I hope

that this section teaches students more about what a p-value actually indicates in the first place (i.e., the proportion of false positives that we are willing to accept), which I feel is a somewhat complicated concept for undergraduates. Please see Figure 3 for a flow diagram of the activation analysis.

Using Multivariate Pattern Analysis to Test the Theory of Distributed Representations

In the second notebook, Exploration: Multivariate Pattern Analysis, I teach students how to implement multivariate pattern analysis to show how it can be used to test and support the theory of distributed representations. The analysis in this notebook focuses on replicating one of the initial classic challenges to the theory of functional specialization by showing that the putative FFA and PPA carry information about categories other than those to which they respond maximally (Haxby et al., 2001). I specifically implement machine-learning analysis using a linear support vector machine (SVM) (Cortes and Vapnik, 1995), which is one of the most commonly applied type of classifier for fMRI analysis (Norman et al., 2006; Grootswagers et al., 2017). One benefit of the linear SVM is that it can handle situations in which we have many features (e.g., voxels) relative to the number of trials (Grootswagers et al., 2017). Here, I will provide a brief overview of the main steps that I implement in the notebook as well as provide some rationale for why I implement each part of the analysis.

The first step of the multivariate pattern analysis, as in the activation analysis, is to load all of the relevant fMRI and behavioral data (i.e., regarding stimulus timing) so that we can have all of the requisite data for analysis. I again make use of `Nilearn` functions for loading the data and `pandas` for loading the stimulus timing information (here, the labels of task events for each frame of data as well as the stimulus onset and duration information for the frame-based analysis as well as the GLM-based analysis, see steps 2 and 3 as well as Figure 3).

The second step is to run multivariate pattern analysis using a linear SVM on the raw data from each individual frame of data (i.e., each TR). In the first analysis, I show students how to analyze the classification accuracy for faces vs. houses (i.e., “places”) within the entire ventral temporal cortex. Here, I use the `Nilearn` functions `Decoder` and `LeaveOneGroupOut`, and I specify the linear SVM by passing the `'svc'` option to the `estimator` parameter for `Decoder`. Moreover, we z-score the data to have a mean of zero and unit variance by setting the `standardize` parameter to `True`, which is a common procedure for working with machine-learning classifiers. We run the analysis using leave-one-run-out cross-validation via the `Nilearn` function `LeaveOneGroupOut`. I then show students how I created functions to run these lines of code with different inputs so that we can look at different category comparisons as well as different regions of interest. Next, we follow up the analysis of the entire ventral temporal cortex by looking at the classification accuracy for faces vs. “places” within voxels that responded maximally to “places” (i.e., these voxels would correspond to the putative PPA). Next, we follow up this analysis by looking at the classification accuracy for the comparison of non-place

stimuli, cats vs. shoes, within the voxels that responded maximally to “places”, which replicates one of the key analyses from Haxby et al. (2001). I then ask students to reflect on the results of these analyses.

The third step is to use a general linear model to create cleaner estimates of the patterns of activity in response to each category for each run for running multivariate pattern analysis. Here, the idea is that the individual frames of data (i.e., TRs) will contain noise, but we can create cleaner patterns of activity by effectively averaging over all of the frames within each run that correspond to a given category. Arguably, the best approach for implementing such a cleaning procedure would be to extract the statistical estimates from a general linear model (e.g., the z-statistics) that is run on each run and then to use these statistical estimates as the input to the classification analysis (e.g., Misaki et al., 2010; Chapter 9 of Poldrack et al., 2011). Thus, here, I modified my code from the activation analysis notebook to enable us to run the GLM procedure and extract the statistical maps for each category for each run. Then, I wrote some custom functions to allow us to run this procedure for various parameters for the region of interest and the comparison of different categories. First, we run the combined GLM/classification analysis for the comparison of faces vs. “places” within the entire ventral temporal cortex. Next, we run the classification of cats vs. shoes within the entire ventral temporal cortex. Then, we repeat these two analyses (i.e., faces vs. “places” and cats vs. shoes) within voxels that respond maximally to faces (i.e., the putative FFA) and then within voxels that respond maximally to “places” (i.e., the putative PPA), which replicates a key comparison of testing how selective the information content is within FFA and PPA (Haxby et al., 2001). Please see Figure 3 for a flow diagram of the multivariate pattern analysis.

Implementing this Module into your Curriculum

Programming Instruction Prior to These Python Notebooks
As I mentioned in the **Course Information** section, I do not require prior programming or statistics background for this course, and the majority of students have no programming background coming into my class. Thus, before completing the Python notebooks, I ask students to read pages from the SciPy Lecture Notes (Varoquaux et al., 2017), watch videos (<https://nsf-cue-frameworks.github.io/www/videos.html>; [part of NSF award 1935099](https://nsf-cue-frameworks.github.io/www/videos.html)), and complete simple assignments to learn about basic Python commands (Table 2). The programming assignments ask students to answer questions from the readings and videos and to get practice with Python (e.g., typing the commands from the reading, creating their own functions to implement simple computations). Altogether, I feel that these two weeks of Python basics help set the stage for the Python analysis notebooks. I would be happy to provide the specific readings and assignments upon request (Tables 2, 5).

How I Teach the Section with These Python Notebooks

For this section of the course, we read Chapters 1-4 and 7 (Ward, 2020) (Table 2). In the first 1.5 weeks of the course, we cover neuroscience fundamentals, including the mind-

Week	Topic	Reading	Programming
1	The mind-body problem	Chapter 1 (Ward)	N/A
2	Overview: from neurons (parts of a neuron, action potentials, neurotransmitters) to systems (broad-scale overview of neuroanatomy)	Chapter 2 (Ward); pg. 12-18 (Varoquaux et al., 2017)	<i>Video:</i> "Assignments & naming – video 1" <i>Assignment:</i> practicing with variables, functions, math operations, lists
3	Electrophysiology (single neurons and EEG); representations (localist, partially distributed representations, fully distributed representations)	Chapter 3 (Ward); pg. 20-26 (Varoquaux et al., 2017)	<i>Videos:</i> "Introduction to Functions", "Calling Built-in Functions", and "Calling Functions in Python" <i>Assignment:</i> dictionaries, tuples, objects the assignment operator, <code>if</code> statements, <code>for</code> loops, list comprehensions, creating and calling functions
4	Brain imaging (fMRI)	Chapter 4 (Ward)	Exploration: Activation analysis
5	Visual cognition	Chapter 7 (Ward)	Exploration: MVPA
6	In-class workshop	N/A	AWA1

Table 2. An overview of the first section of the course in which I use the teaching materials that I describe here. In addition to the assignments listed here, each week students responded to questions that I wrote to get students to think about the big-picture importance of each reading, thus these questions served as a scaffold for the bigger writing assignment at the end of this section. The videos are part of an NSF award for teaching programming: <https://nsf-cue-frameworks.github.io/www/videos.html> I use the 2020.2 edition of the SciPy Lecture Notes (Varoquaux et al., 2017) and omit some sections for concision. Abbreviations: MVPA=multivariate pattern analysis, AWA1=Analytical and Writing Assignment #1 (the combination of the Explorations [i.e., the analytical part] and responses to questions about the interpretation of the results and other materials from discussions and readings [i.e., the writing part]).

body problem, the action potential, neurotransmitters and a broad overview of neuroanatomy in which we discuss the importance of the inputs and outputs of a region in constraining its possible functions (e.g., by discussing primary cortical areas vs. secondary cortical areas vs. association cortex; Week 1 and 2 in Table 2).

In the third week of the course, we discuss the concepts of neural representations, single-cell electrophysiology, scalp electroencephalography (EEG), and the notion of rate codes vs. temporal codes (e.g., oscillations). We first discuss the idea that the neural responses of the primary visual cortex would not give rise to invariant object representation (here, giving an example of how moving

stimuli around in the visual field would elicit totally unique neural responses in early visual cortex). Then, we continue this conversation by talking about the concepts of local representations (i.e., the "grandmother cell" concept) vs. partially distributed representations vs. fully distributed representations (Chapter 3 from Ward, 2020 is a helpful primer for these conversations). We then discuss findings from a study in nonhuman primates (Baylis et al., 1985) that suggests that the brain uses sparse distributed representations for coding of categories such as faces. Then, we discuss evidence for more localized representations in humans (Quiroga et al., 2005) in response to famous faces and places (e.g., the so-called "Halle Berry"/"Jennifer Aniston" neuron paper). Even here, however, we talk about the notion that these results could also be interpreted as evidence for sparse distributed representations (e.g., there was some mixed selectivity of neurons coding for associations between stimuli). Next, we talk about how the properties of synaptic transmission can enable us to measure electrical potentials at the scalp with EEG. We then discuss the concept of event-related potentials, and we focus on the N170, which is an event-related potential that is larger in response to faces than other categories of stimuli (e.g., cars, other objects: Bentin et al., 1996), and it has been interpreted as evidence that face processing recruits specialized neural modules (i.e., evidence of functional specificity). Then, we discuss a study that suggests that the N170 might reflect expertise (i.e., because we have a lot of experience processing faces), rather than something fundamental about faces per se (Tanaka and Curran, 2001). Specifically, Tanaka and Curran (2001) showed a double dissociation between the N170 amplitude in response to birds vs. dogs in bird and dog experts. Thus, I would like to highlight here that I attempt to discuss the concept of visual cognition and the evidence for and against the theory that face processing recruits a specialized neural module, thus students can use these discussions in their writing assignment and it helps set the stage for our discussion of fMRI in the next week and with the Python notebooks. We conclude our discussion this week by talking about the concepts of rate coding vs. temporal coding (e.g., neural coherence, spike-timing-dependent plasticity, and neural oscillations) (Fries, 2005).

In the fourth week of the course, we discuss the concepts of MRI physics, the hemodynamic response, the general linear model (GLM), and the logic of applying the GLM to measure the BOLD activity. When we discuss the GLM, I found it to be helpful to give a real-world example within the realm of plant growth. Here, I ask students to imagine that we begin 3 different interventions or treatments to plants and then we measure their growth over time. The idea is to help students think about the concept that treatments that have a larger effect will result in a faster rate of growth, which would translate into a steeper slope (i.e., a greater resultant beta weight). Similarly, for fMRI activation analyses, voxels that have a larger beta weight for a given condition (e.g., faces) relative to other conditions would suggest that these voxels are influenced more strongly by faces compared to the other conditions. I then proceed to show students examples of expected and simulated fMRI timecourses and I ask them to

visually and mentally calculate beta weights for different examples (in think-pair-share style activities), thus allowing them to get a better understanding of the logic of the GLM-based fMRI analysis prior to completing the notebooks (you can access the slides for the GLM introduction and activities: <https://doi.org/10.17605/OSF.IO/UJWK6>). We also discuss how we can implement machine learning analysis (e.g., linear SVM) by covering a simple example that I developed: deciding whether or not to eat at a restaurant based on both the rating and the cost of the food and then showing how we can apply this logic to fMRI analysis to study distributed representations (you can access the slides for this discussion here: <https://doi.org/10.17605/OSF.IO/UJWK6>).

In the fifth week, we discuss vision, beginning with the primary visual cortex and ending with a discussion of high-level vision and an overview of our findings from the Python notebooks. I begin this section by discussing the topographic organization of the primary somatosensory cortex and primary motor cortex, the tonotopic organization of the primary auditory cortex, and the retinotopic organization of the primary visual cortex. Then, we go into detail about how researchers have used fMRI to further our understanding of early visual cortex in humans. For example, we talk about how researchers have used ring and wedge stimuli to study retinotopic maps in human early visual cortex (Dumoulin and Wandell, 2008; Wandell and Winawer, 2011). Then, I ask students to make predictions about the impact of lesions in various locations of early visual cortex (i.e., to connect fMRI with the study of patients with localized brain damage; here, scotoma, quadrantanopia, and hemianopia). We then talk about the single-cell electrophysiological studies that discovered the nature of the receptive fields of lateral geniculate nucleus and the primary visual cortex, and we focus on “simple cells” in the primary visual cortex. Then, we talk about how fMRI has been used to show that the human primary visual cortex contains similar representations. For example, we discuss how researchers have applied voxelwise encoding models with a Gabor model, which approximates the “simple cell” receptive fields of the primary visual cortex (V1), to show that it provides the best fit for V1 and then is progressively worse higher in the visual hierarchy (V2, V3, and V4) (Kay et al., 2008; Kay, 2011).

We next discuss the coding properties of higher-level visual cortex. First, we discuss a paper that compared voxelwise encoding models of the Gabor model vs. a semantic model and found that the Gabor model does a better job of accounting for responses in early visual areas (e.g., V1 as well as V2 and V3) while the semantic model provided a better fit of higher-level visual areas (e.g., the anterior portion of lateral occipital cortex and anterior occipital cortex; Naselaris et al., 2009). Then, we talk about patients with visual agnosia as an introduction to concept of object recognition. We then discuss a few more examples of seemingly functionally specialized regions for processing movement (typically called V5 in humans or MT in nonhuman primates; e.g., patient studies of akinetopsia: Zihl et al., 1983; Zeki, 1991; as well as the finding that this region responds to the Enigma illusion [i.e., illusory motion]: Zeki et al., 1993), hemineglect following damage to the right parietal

lobe, and the finding that V2 neurons respond to the illusory triangle (Von Der Heydt et al., 1984; for more information see Chapter 7 of Ward, 2020). We also discuss the dorsal vs. ventral visual streams hypothesis for processing “where”/“how” information (i.e., spatial processing) and the ventral visual stream in processing “what” information (i.e., object processing) throughout the course (e.g., when we discuss neuroanatomy and here; e.g., Mishkin et al., 1983). Then, we go into detail discussing the results of our Exploration: Activation Analysis notebook, and I show students how similar our results are to the original studies that discovered the FFA (Kanwisher et al., 1997) and PPA (Epstein and Kanwisher, 1998; I enjoy showing figures of our results next to a paper that was published in Nature).

We conclude this section by discussing the evidence for and against the hypothesis of category specificity for face processing. For example, in support of the theory of functional specificity for face processing, we discuss 1) the findings from our activation analysis, 2) patients with prosopagnosia (e.g., the finding that patients could have relatively selective impairment for processing faces: Sergent and Signoret, 1992; but note that other accounts have shown a more domain general deficit in face and object recognition: Barton and Corrow, 2016; Geskin and Behrmann, 2018), 3) stimulation of putative FFA causes changes in face processing (Parvizi et al., 2012), and 4) we revisit the finding of the greater N170 for faces vs. other categories (Bentin et al., 1996). We then discuss some initial challenges for the hypothesis of category specificity in face processing, and we focus on the expertise hypothesis. For example, I remind students of the finding of differential N170 findings in category experts (i.e., bird and dog experts: Tanaka and Curran, 2001) and I discuss how these findings extend to fMRI, where researchers have shown that the FFA shows differential activity in category experts (e.g., car and bird experts: Gauthier et al., 2000) and the finding that new learning of complex object discrimination caused increased activation of FFA (e.g., Greebles: Gauthier et al., 1999). I conclude by reiterating to students that this week’s homework assignment (Exploration: Multivariate Pattern Analysis) allows students to test the theory of distributed representations.

At the end of this section, students respond to several prompts that get them to think about the results of both Python notebooks as well as how the concept of converging operations (e.g., McNamara, 1991) can provide stronger evidence for our theories (I provide the writing prompts here: <https://doi.org/10.17605/OSF.IO/UJWK6>). Note, as I discussed above, I like the fact that the results of our notebooks as well as an examination of the literature does not support the notion of clearcut “right” or “wrong” answers for how visual cognition works, rather students get a chance to evaluate the data to come to their own conclusions, similar to what we do in our research programs. The week that the writing assignment is due, I hold an in-class workshop in which I provide opportunities for students to ask questions about the course material and the Python notebooks.

RESULTS

Here, I briefly describe the results of the Python analysis




Figure 4. The results of the activation analysis contrast of faces vs. all other categories revealed clusters of voxels in the fusiform gyrus (i.e., putative fusiform face area), as generated in Notebook #1. Here, the map shows the results of FDR-corrected voxelwise threshold of $p < 0.001$. These results can be compared to the original discovery of the fusiform face area (see Figures 1 and 2 from Kanwisher et al., 1997; please note that the left/right information from the Haxby et al. (2001) data was lost so we cannot infer anything about laterality in our data here).

notebooks and I discuss how students responded to these assignments. Please see the MATERIALS AND METHODS and DISCUSSION for suggestions for implementing this module into your curriculum.

Using Activation Analysis to Test the Theory of Functional Specialization

In the first Python notebook (Exploration: Activation Analysis), I show students how to run a contrast to compare the activation of faces vs. the seven other categories (houses [i.e., “places”], cats, bottles, scissors, shoes, chairs, and phase-scrambled images of the objects) to test the theory that face processing is implemented in specialized neural machinery within the FFA (Kanwisher et al., 1997) as well as “places” vs. the seven other categories to test the theory that “place” processing is implemented in specialized neural machinery within the PPA (Epstein and Kanwisher, 1998). In the first analysis, students create brain maps for these two contrasts that are not corrected for multiple comparisons. In the second analysis, students apply an FDR correction to the brain maps for the faces vs. other categories contrast and the “places” vs. other categories contrast. I also show students the importance of correcting for multiple comparisons by exploring simulations of random data to see the proportion of false positives that they observe at various alpha levels as well as various sizes of regions of interest (i.e., number of voxels). Here, they discover that the proportion of false positives corresponds to the alpha level, thus highlighting the importance of the correction for multiple comparisons and hopefully teaching students more about the actual assumptions of frequentist statistics (i.e., these tests control the false positive rate). The results of the activation analysis replicate the classic findings; specifically, we find clusters of activated voxels within the fusiform gyrus for the faces contrast (i.e., the putative FFA; Figure 4) and at the border of the parahippocampal gyrus and the visual association area for the “places” contrast (i.e., the putative PPA; Figure 5).




Figure 5. The results of the activation analysis contrast of “places” vs. all other categories revealed clusters of voxels along the junction of the posterior portion of the parahippocampal gyrus and the anterior portion of the visual association cortex (i.e., putative parahippocampal place area). Please compare these results to the original discovery of the parahippocampal place area (see Figure 2 from Epstein and Kanwisher, 1998; please note that the left/right information from the Haxby et al. (2001) data was lost so we cannot infer anything about laterality in our data here).

Using Multivariate Pattern Analysis to Test the Theory of Distributed Representations

In the second Python notebook (Exploration: Multivariate Pattern Analysis), students can implement machine-learning-based MVPA to test the theory that high-level visual processing is implemented in distributed representations throughout the ventral temporal cortex (Haxby et al., 2001). Students perform two main classes of MVPA: analysis of individual frames of data (i.e., TRs) and analysis of a GLM-based extraction of patterns of activity for each category for each run. Both of these methods are commonly used in fMRI research, and the GLM-based patterns of activity provide cleaner data and results. Moreover, both methods provide the same overall pattern of results; therefore, here, I will focus my discussion on the classification analysis of the GLM-based patterns of activity, but the full results can be observed within the Python notebook.

Students will find that a linear SVM can classify the comparison of faces vs. houses (i.e., “places”) for the GLM-based patterns of activity within the entire ventral temporal cortex (average classification accuracy: 100%), the voxels that responded maximally to faces (i.e., the putative FFA; average classification accuracy: 91.67%), and the voxels that responded maximally to “places” (i.e., the putative PPA; average classification accuracy: 100%; see the left panel of Figure 6). Next, we find that a linear SVM can classify the comparison of other non-face and non-place stimuli—specifically, cats vs. shoes—within the entire ventral temporal cortex (average classification accuracy: 100%), the voxels that responded maximally to faces (i.e., the putative FFA; average classification accuracy: 100%), and the voxels that responded maximally to “places” (i.e., the putative PPA; average classification accuracy: 100%; see the right panel of Figure 6). These results replicate the classic finding that the putative FFA and PPA both contain information about stimulus categories other than the one to which they respond maximally (Haxby et al., 2001; e.g., one

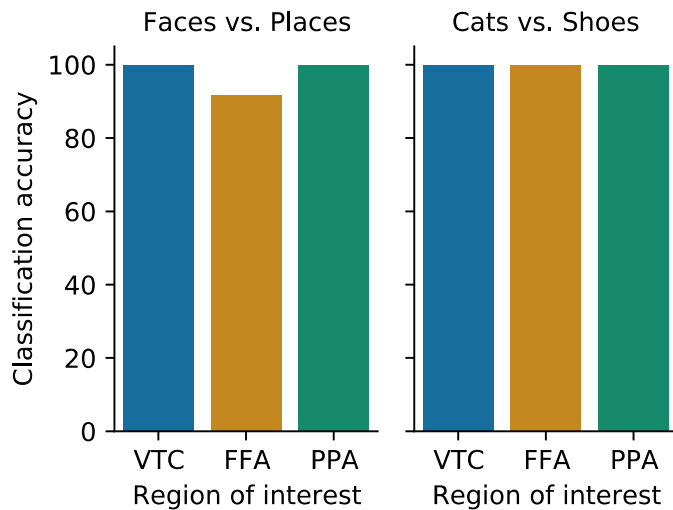


Figure 6. The results of the multivariate pattern analysis revealed very good classification accuracy of faces vs. “places” (all accuracies $\geq 91.67\%$; see left panel) and cats vs. shoes (all accuracies 100%; see right panel) in the ventral temporal cortex as well as in voxels that responded maximally to faces (i.e., putative fusiform face area) and “places” (i.e., putative parahippocampal place area). Altogether, these results support the classic finding that the putative fusiform face area and parahippocampal place area contain information in addition to the categories to which they respond maximally. Note about abbreviations: VTC=ventral temporal cortex, FFA=putative fusiform face area (voxels that responded maximally to faces), PPA=putative parahippocampal place area (voxels that responded maximally to “places”).

would be hard pressed to argue that a cat or a shoe is a “place”-like stimulus).

Altogether, the results of both Python notebooks replicate the classic discovery of FFA and PPA, while also showing that the choice of analysis plays a key role in shaping the kinds of conclusions that we can draw from our data. Therefore, these results show students that there are not clearcut “right” or “wrong” answers in cognitive neuroscience, rather we effectively instantiate our assumptions within our analytical techniques, which can lead us to different conclusions about the relationship between cognition and the brain.

Students Responded Positively to These Assignments

As I mentioned in the MATERIALS AND METHODS, I have used these assignments in three semesters of teaching my introductory cognitive neuroscience course. For the past two semesters, I collected information about students’ perceptions of these assignments via reflection assignments and additional questions in my end-of-the-semester course evaluations. In the reflection assignment, I ask students to respond to several questions, including a question in which I ask them how the assignments meet our course goals. The last time that I taught this course (Fall 2022), all 35 students said that the assignments meet the course learning goals. As one example of student responses to this question, a student wrote: “I think [the assignments] fit well with the goal of learning to think like a cognitive neuroscientist because we got to actually do cognitive neuroscience and interpret

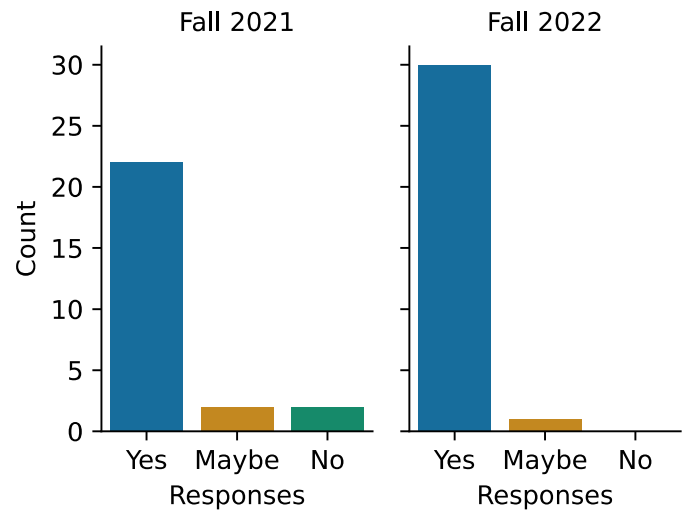


Figure 7. Students’ reflections on the assignments with the Python notebooks and subsequent writing assignments were generally positive. I asked students the following questions: “Do you feel that the Analytical and Writing Assignments supported your learning in this course? For example, how would these assignments compare (in terms of how much you feel you learned) to traditional multiple-choice exams, etc. that might be typical in 200-level courses?” I categorized the students’ responses: “Yes” = they responded positively with detailed comments about why they liked these assignments, “Maybe” = they responded with a mixed response (e.g., they saw value in the assignments but they also saw value in more traditional assessments), “No” = they noted that they thought that traditional assessments would have better supported their learning. Total responses: 2021=26, 2022=31. Also see Table 3.

the results rather than just read about them in a textbook.”

Arguably, anonymous course evaluations give a more realistic view into students’ perception of these assignments, thus I have begun asking students to respond to the following question: “Do you feel that the Analytical and Writing Assignments supported your learning in this course? For example, how would these assignments compare (in terms of how much you feel you learned) to traditional multiple-choice exams, etc. that might be typical in 200-level courses?” Note that the Analytical and Writing Assignments (AWAs) are assignments in which I ask students to complete hands-on analysis and exploration of data (i.e., the analytical portion of the first Analytical and Writing Assignment consists of the Python notebooks that I describe here) and then to critically reflect on these analyses in a writing assignment. I implement two such assignments throughout the semester and the assignments described here account for one of these assignments (see the DISCUSSION for information about the second Analytical and Writing Assignment). Thus, while the responses to the questions below reflect more than solely the notebooks, I feel their responses are representative of students’ feelings of the lesson plan I presented here. I categorized the responses into 3 broad categories: “yes” = they responded positively with detailed comments about why they liked these assignments, “maybe” = they responded with a mixed response (e.g., they saw value in the assignments but they also saw value in more traditional assessments), “no” = they

Student responses
"YES! I think studying and cramming a bunch of info into your head that you're going to purge at the end of the semester is a waste of time, especially when we have the world at our fingertips all the time. I learned how to THINK like a cognitive neuroscientist, learning how to think is a much more valuable lesson than just learning what to think. I think this course successfully accomplished its goal... I came to a liberal arts college to expand my analytical thinking ability, and this course certainly helped assist that goal..."
"Yes, they supported my learning and gave me an opportunity to apply the concepts & theories I learned to analyzing scientific/experimental data. They were also a test on my writing skills as I had to connect different ideas together effectively..."
"I feel like I learned a different kind of information than I would have with exams. I think the analytical and writing assignments gave me the chance to apply and discuss my learning. I really like being able to apply actual techniques used in research..."
"I feel very strongly that the AWAs were a better learning option for me than traditional exams. I actually am a student that does not typically mind exams, but I think some of this material was a bit too complex for me to have the sort of unassisted understanding that would be required for exams. The writing assignments gave prompts that allowed us to walk ourselves through the ideas and REALLY furthered my understanding of the concepts. I think these assignments were very, very well written and well done. These assignments are where my understanding of the course material was solidified."

Table 3. Students' reflections on the assignments with the Python notebooks and subsequent writing responses suggests that the assignments helped us meet our course goals. This table shows the more detailed comments that are summarized categorically in Figure 7.

noted that they thought that traditional assessments (e.g., multiple-choice tests) would have better supported their learning. As you can see in Figure 7, altogether across two semesters, 52 students responded "yes", 3 students responded "maybe" (e.g., some positive and some negative comments), and 2 students responded "no". In summary, the response to these questions were very positive overall. In addition to these data, students that responded "yes" had a variety of enthusiastic comments in my anonymous evaluations (Table 3).

In addition to specific comments about the nature of these assignments, I have heard positive things from students regarding teaching them Python as a skill they can use for applications for internships and jobs. For example, one student wrote in my evaluations, "We learned Python this semester and many internships want their applicants to have some familiarity with it and now I can say that I have that." Moreover, I have had several students change their major to Computational Psychology or take more courses in Computer Science after taking this class.

DISCUSSION

I recommend incorporating these Python notebooks into a section of your course that is focused on the combination of methods and visual cognition (also see MATERIALS AND METHODS for more information). I found that the best way to help students to understand the importance of these modules is to 1) frequently explain the primary importance

of programming for theory-testing via analysis in cognitive neuroscience, 2) integrate the analysis notebooks with an in-depth discussion of visual cognition and the major theories of neural representation (e.g., functional specialization, distributed representations, networks of the brain), 3) incorporate the notebooks into class goals and refer students back to the class goals during these activities and the writing assignment. The first time that I taught this section, I did not make the critical integration of these three aspects into my course. Specifically, I focused this section of the class on the methods of cognitive neuroscience without the in-depth discussion of visual cognition or the connection to the learning goals, which did not have as much of a positive impact on students.

How I Teach the Remainder of the Course: A Continued Discussion of the Concept of Neural Representations

After completing the first section of the course, we move on to complete a section on memory, with a focus on semantic cognition (McClelland and Rogers, 2003; Patterson et al., 2007), episodic memory (O'Reilly et al., 2012; Barnes and Underwood, 1959), and spatial attention and navigation (Chapter 3 from Ekstrom et al., 2018; e.g., place cells: O'Keefe and Dostrovsky, 1971; O'Keefe and Nadel, 1978; head-direction cells: Taube et al., 1990; e.g., grid cells: Hafting et al., 2005; e.g., interacting networks: Ekstrom et al., 2017). Here, we also discuss how models from computational cognitive neuroscience have informed theories of human memory, with a focus on the utility of distributed representations for semantic cognition (i.e., building on our fMRI explorations; we also do a close-read and in-depth discussion of neural networks and semantic cognition via McClelland and Rogers, 2003) vs. the use of pattern-separated representations (i.e., local or sparse distributed representations) for episodic memory (including hands-on exploration of computational models; Chapter 8 from O'Reilly et al., 2012). We also discuss network theories of the brain for supporting spatial navigation (Ekstrom et al., 2017). Therefore, we continue our overarching discussion of theories of neural representations (i.e., local representations vs. distributed representations vs. network interactions). We conclude the second section of the class with a writing assignment based on the students' close-read and in-class discussion and activities of the paper (McClelland and Rogers, 2003), the hands-on exploration of the models of the neocortex and hippocampus (Chapter 8 of O'Reilly et al., 2012), and our discussion of the neuroscience of navigation. In the final section of the course, we explore high-level cognition, where students choose a topic that is the most interesting to them and they present their findings to the class, including creating in-class discussion questions (e.g., think-pair-share-style activities; Table 4).

On the Merits of Implementing Python Early for More Advanced Study Later in the Curriculum

As I discussed above, I use these activities as part of my introductory cognitive neuroscience course because I want to expose students to Python programming as early as possible in our curriculum. Moreover, by teaching students about Python programming and machine learning analysis

Week	Topic	Reading	Assignment
7	Memory I: Studying the brain	Chapter 11 (Ward)	
8	Semantic cognition and neural networks	Pg. 307-315 (Ward); (McClelland and Rogers, 2003)	Discussion and activities re: neural networks and semantic cognition (McClelland and Rogers, 2003)
9	Memory II: A computational cognitive neuroscience approach	Chapter 8 (O'Reilly et al., 2012)	Computer lab: Exploration of memory via models of the 1) neocortex, 2) hippocampus (O'Reilly et al., 2012)
10	Spatial attention and navigation	Pg. 203-216, 224-230 (Ward); Chapter 3 (Ekstrom et al., 2018)	
11	In-class workshop	N/A	AWA2
12	Discussions	Variable	Prepare slides
13	Discussions	Variable	Present; Reflect

Table 4. An overview of the second and third sections of the course in which I use the teaching materials that I describe here. Note that for Week 9 I implement a flipped classroom with pre-lab video lectures and then a full week of hands-on exploration of computational models in a computer lab. In addition to the assignments listed here, each week students responded to questions that I wrote to get students to think about the big-picture importance of each reading, thus these questions served as a scaffold for the bigger writing assignment at the end of Section 2 (in green). The Discussions at the end of the class were student-led presentations on their chosen area of remaining topics. Abbreviations: AWA2=Analytical and Writing Assignment #2 (the combination of the explorations with computational models [i.e., the analytical part] and responses to questions about the interpretation of the results and other materials from discussions and readings [i.e., the writing part]).

of data from the human brain in an introductory course, it opens up the possibility for running more advanced analyses and projects in future semesters. For example, I also teach a combined Seminar in Cognitive Neuroscience and Collaborative Research in Cognitive Neuroscience course pair, where I teach students how to create their own tasks in `PsychoPy` (Peirce, 2007; Peirce et al., 2019), collect their own behavioral and EEG data, and preprocess and run advanced time-frequency and machine-learning based analyses with `MNE-Python` (Gramfort et al., 2013). In Spring 2022, a group of students in the collaborative research course worked on a project looking at spatial memory, and a student from that group went on to conduct her Senior Honors Thesis project on a related but slightly more elaborate experimental design. We presented our results of this project at the annual meeting of the Society for Neuroscience and we plan to finalize data collection and analysis and to write a paper on this project this academic year. Therefore, I believe that teaching Python throughout the cognitive neuroscience curriculum opens significant

opportunities for students and instructors alike, which is especially beneficial at primarily undergraduate universities and colleges that value both teaching and research.

Conclusion

The field of cognitive neuroscience has seen a fierce, decades-long debate regarding the nature of neural representations that support high-level vision and category representation, with one set of prominent theories arguing for functional specificity and another set of prominent theories arguing for distributed representations. Importantly, these theories have received support via different methods, thus the methods that we employ make key assumptions that allow us to test different theories. Here, I designed Python notebooks to teach students how to replicate classic findings that discovered the FFA (Kanwisher et al., 1997; Kanwisher, 2017) and PPA (Epstein and Kanwisher, 1998) as well as important challenges to these theories that suggest the brain instead implements high-level vision and category representation via distributed representations (Haxby et al., 2001). I found that these assignments allow students to gain a deep understanding of the main theories and findings within cognitive neuroscience (based on their responses to a writing assignment in which I asked students to synthesize the findings of the Python notebooks, our in-class discussions, and the readings) and students responded positively to these explorations. Moreover, these assignments form a critical building block toward the rest of my course and upper-division courses within our neuroscience curriculum. I am providing these resources with the aim that we can increase the involvement of students in high-level data analyses early in the curriculum

Links to the resources from this paper
GitHub repository for accessing the Jupyter Notebooks: https://github.com/huffman-spatial-cognition-lab/exploration_of_fMRI_methods_and_theory
Please also see the repository's wiki page: https://github.com/huffman-spatial-cognition-lab/exploration_of_fMRI_methods_and_theory/wiki
Video walkthrough for Exploration: Activation Analysis https://github.com/huffman-spatial-cognition-lab/exploration_of_fMRI_methods_and_theory/wiki/Video-walkthrough:-Activation-analysis
Video walkthrough for Exploration: MVPA https://github.com/huffman-spatial-cognition-lab/exploration_of_fMRI_methods_and_theory/wiki/Video-walkthrough:-MVPA
Additional resources, including slides for teaching GLM and SVM concepts for fMRI and prompts for Analytical and Writing Assignment #1 are available at the following OSF repository: https://doi.org/10.17605/OSF.IO/UJWK6
Videos for teaching basics of Python programming: https://nsf-cue-frameworks.github.io/www/videos.html

Table 5. Links to the resources from this paper, including the Jupyter notebooks, the wiki page for getting the notebooks and environments set up, video walkthroughs, slides, prompts for the writing assignment, and videos for teaching Python basics.

so that they can go on to have successful careers in many areas and we can increase the diversity of students that take an interest in programming.

REFERENCES

- Abraham A, Pedregosa F, Eickenberg M, Gervais P, Mueller A, Kossaifi J, Gramfort A, Thirion B, Varoquaux G (2014) Machine learning for neuroimaging with scikit-learn. *Front Neuroinformatics* 8 Available at <http://journal.frontiersin.org/article/10.3389/fninf.2014.00014/abstract>.
- Alstott J, Breakspear M, Hagmann P, Cammoun L, Sporns O (2009) Modeling the Impact of Lesions in the Human Brain. Friston KJ, ed. *PLoS Comput Biol* 5:e1000408.
- Barnes JM, Underwood BJ (1959) Fate of first-list associations in transfer theory. *J Exp Psychol* 58:97–105.
- Barton JJS, Corrow SL (2016) Selectivity in acquired prosopagnosia: The segregation of divergent and convergent operations. *Neuropsychologia* 83:76–87.
- Baylis GC, Rolls ET, Leonard CM (1985) Selectivity between faces in the responses of a population of neurons in the cortex in the superior temporal sulcus of the monkey. *Brain Res* 342:91–102.
- Bentin S, Allison T, Puce A, Perez E, McCarthy G (1996) Electrophysiological Studies of Face Perception in Humans. *J Cogn Neurosci* 8:551–565.
- Brabec JL, Vos MR, Staab TA, Chan JP (2018) Analysis of Student Attitudes of a Neurobiology Themed Inquiry Based Research Experience in First Year Biology Labs. *J Undergrad Neurosci Educ JUNE Publ FUN Fac Undergrad Neurosci* 17:A1–A9.
- Brainard DH (1997) The Psychophysics Toolbox. *Spat Vis* 10:433–436.
- Bullmore ET, Bassett DS (2011) Brain Graphs: Graphical Models of the Human Brain Connectome. *Annu Rev Clin Psychol* 7:113–140.
- Cortes C, Vapnik V (1995) Support-vector networks. *Mach Learn* 20:273–297.
- Cronqvist S, Ekberg R, Ingvar DH (1965) Regional cerebral blood flow related to neuroradiological findings. *Acta Neurol Scand* 41:176–178.
- Dumoulin SO, Wandell BA (2008) Population receptive field estimates in human visual cortex. *NeuroImage* 39:647–660.
- Edelman S, Grill-Spector K, Kushnir T, Malach R (1998) Toward direct visualization of the internal shape representation space by fMRI. *Psychobiology* 26:309–321.
- Eichenbaum HB, Davis JL eds. (1998) *Neuronal ensembles: strategies for recording and decoding*. Chichester, New York: Wiley.
- Ekstrom AD, Huffman DJ, Starrett M (2017) Interacting networks of brain regions underlie human spatial navigation: a review and novel synthesis of the literature. *J Neurophysiol* 118:3328–3344.
- Ekstrom AD, Spiers HJ, Bohbot VD, Rosenbaum RS (2018) *Human spatial navigation*. Princeton, New Jersey: Princeton University Press.
- Epstein R (2005) The cortical basis of visual scene processing. *Vis Cogn* 12:954–978.
- Epstein R, Kanwisher N (1998) A cortical representation of the local visual environment. *Nature* 392:598–601.
- Fodor JA (1983) *The Modularity of Mind*. The MIT Press. Available at <https://direct.mit.edu/books/book/3985/the-modularity-of-mind>.
- Fries P (2005) A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends Cogn Sci* 9:474–480.
- Gauthier I, Skudlarski P, Gore JC, Anderson AW (2000) Expertise for cars and birds recruits brain areas involved in face recognition. *Nat Neurosci* 3:191–197.
- Gauthier I, Tarr MJ, Anderson AW, Skudlarski P, Gore JC (1999) Activation of the middle fusiform “face area” increases with expertise in recognizing novel objects. *Nat Neurosci* 2:568–573.
- Geskin J, Behrmann M (2018) Congenital prosopagnosia without object agnosia? A literature review. *Cogn Neuropsychol* 35:4–54.
- Gorgolewski K, Burns CD, Madison C, Clark D, Halchenko YO, Waskom ML, Ghosh SS (2011) Nipype: A Flexible, Lightweight and Extensible Neuroimaging Data Processing Framework in Python. *Front Neuroinformatics* 5. Available at <http://journal.frontiersin.org/article/10.3389/fninf.2011.00013/abstract>.
- Gramfort A, Luessi M, Larson E, Engemann DA, Strohmeier D, Brodbeck C, Goj R, Jas M, Brooks T, Parkkonen L, Hämäläinen M (2013) MEG and EEG data analysis with MNE-Python. *Frontiers in Neuroscience* 7. Available at <https://www.frontiersin.org/journals/neuroscience/articles/10.3389/fnins.2013.00267>.
- Gratton C, Nomura EM, Pérez F, D’Esposito M (2012) Focal Brain Lesions to Critical Locations Cause Widespread Disruption of the Modular Organization of the Brain. *J Cogn Neurosci* 24:1275–1285.
- Groetswagers T, Wardle SG, Carlson TA (2017) Decoding Dynamic Brain Patterns from Evoked Responses: A Tutorial on Multivariate Pattern Analysis Applied to Time Series Neuroimaging Data. *J Cogn Neurosci* 29:677–697.
- Hafting T, Fyhn M, Molden S, Moser M-B, Moser EI (2005) Microstructure of a spatial map in the entorhinal cortex. *Nature* 436:801–806.
- Harris CR et al. (2020) Array programming with NumPy. *Nature* 585:357–362.
- Haxby JV (2012) Multivariate pattern analysis of fMRI: The early beginnings. *NeuroImage* 62:852–855.
- Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, Pietrini P (2001) Distributed and Overlapping Representations of Faces and Objects in Ventral Temporal Cortex. *Science* 293:2425–2430.
- Heisenberg W (2007) *Physics & philosophy: the revolution in modern science*, 1st Harper Perennial Modern Classics ed. New York: HarperPerennial.
- Henson RN, Greve A, Cooper E, Gregori M, Simons JS, Geerligs L, Erzinçlioğlu S, Kapur N, Browne G (2016) The effects of hippocampal lesions on MRI measures of structural and functional connectivity: Connectivity Change after Hippocampal Lesion. *Hippocampus* 26:1447–1463.
- Ho Y-Y, Roeser A, Law G, Johnson BR (2021) *Pandemic Teaching: Using the Allen Cell Types Database for Final Semester Projects in an Undergraduate Neurophysiology Lab Course*. *J Undergrad Neurosci Educ JUNE Publ FUN Fac Undergrad Neurosci* 20:A100–A110.
- Hunter JD (2007) *Matplotlib: A 2D Graphics Environment*. *Comput Sci Eng* 9:90–95.
- Ingvar DH, Schwartz MS (1974) Blood flow patterns induced in the dominant hemisphere by speech and reading. *Brain* 97:273–288.
- Jahn A, Levitas D, Holscher E, Johnson JT, Sayal A, Jstaph, JohannesWiesner, Clucas J, Tinashe Michael Tapera, Justbennet (2022) *andrewjahn/AndysBrainBook*: Available at <https://zenodo.org/record/5879293>.
- Juavinett A (2020) *Learning How to Code While Analyzing an Open Access Electrophysiology Dataset*. *J Undergrad Neurosci Educ JUNE Publ FUN Fac Undergrad Neurosci* 19:A94–A104.
- Juavinett AL (2022) The next generation of neuroscientists needs to learn how to code, and we need new ways to teach them. *Neuron* 110:576–578.
- Kanwisher N (2010) Functional specificity in the human brain: A window into the functional architecture of the mind. *Proc Natl*

- Acad Sci 107:11163–11170.
- Kanwisher N (2017) The Quest for the FFA and Where It Led. *J Neurosci* 37:1056–1061.
- Kanwisher N, McDermott J, Chun MM (1997) The Fusiform Face Area: A Module in Human Extrastriate Cortex Specialized for Face Perception. *J Neurosci* 17:4302–4311.
- Kay KN (2011) Understanding Visual Representation by Developing Receptive-Field Models. In: *Visual Population Codes* (Kriegeskorte N, Kreiman G, eds), pp 133–162. The MIT Press. Available at <https://direct.mit.edu/books/book/2176/chapter/57847/Understanding-Visual-Representation-by-Developing>.
- Kay KN, Naselaris T, Prenger RJ, Gallant JL (2008) Identifying natural images from human brain activity. *Nature* 452:352–355.
- Kumar M et al. (2022) BrainIAK: The Brain Imaging Analysis Kit. *Aperture Neuro* 2021:42.
- Kumar M, Ellis CT, Lu Q, Zhang H, Capotà M, Willke TL, Ramadge PJ, Turk-Browne NB, Norman KA (2020) BrainIAK tutorials: User-friendly learning materials for advanced fMRI analysis. *PLOS Comput Biol* 16:e1007549.
- Lacadie CM, Fulbright RK, Rajeevan N, Constable RT, Papademetris X (2008) More accurate Talairach coordinates for neuroimaging using non-linear registration. *NeuroImage* 42:717–725.
- Lopatto D (2007) Undergraduate Research Experiences Support Science Career Decisions and Active Learning. *CBE—Life Sci Educ* 6:297–306.
- McCarthy G, Puce A, Gore JC, Allison T (1997) Face-Specific Processing in the Human Fusiform Gyrus. *J Cogn Neurosci* 9:605–610.
- McClelland JL ed. (1986) *Parallel distributed processing. 2: Psychological and biological models*. Cambridge, Mass: The MIT Press.
- McClelland JL, Rogers TT (2003) The parallel distributed processing approach to semantic cognition. *Nat Rev Neurosci* 4:310–322.
- McKinney W (2010) Data Structures for Statistical Computing in Python. In, pp 56–61. Austin, Texas. Available at <https://conference.scipy.org/proceedings/scipy2010/mckinney.html>.
- McNamara TP (1991) Memory's View of Space. In: *Psychology of Learning and Motivation*, pp 147–186. Elsevier. Available at <https://linkinghub.elsevier.com/retrieve/pii/S007974210860123>.
- Misaki M, Kim Y, Bandettini PA, Kriegeskorte N (2010) Comparison of multivariate classifiers and response normalizations for pattern-information fMRI. *NeuroImage* 53:103–118.
- Mishkin M, Ungerleider LG, Macko KA (1983) Object vision and spatial vision: two cortical pathways. *Trends Neurosci* 6:414–417.
- Naselaris T, Prenger RJ, Kay KN, Oliver M, Gallant JL (2009) Bayesian Reconstruction of Natural Images from Human Brain Activity. *Neuron* 63:902–915.
- Nichols TE, Das S, Eickhoff SB, Evans AC, Glatard T, Hanke M, Kriegeskorte N, Milham MP, Poldrack RA, Poline J-B, Proal E, Thirion B, Van Essen DC, White T, Yeo BTT (2017) Best practices in data analysis and sharing in neuroimaging using MRI. *Nat Neurosci* 20:299–303.
- Nilearn Contributors et al. (2023) Nilearn. Available at <https://zenodo.org/record/8397156>.
- Norman KA, Polyn SM, Detre GJ, Haxby JV (2006) Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends Cogn Sci* 10:424–430.
- O'Keefe J, Dostrovsky J (1971) The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res* 34:171–175.
- O'Keefe J, Nadel L (1978) *The hippocampus as a cognitive map*. Oxford : New York: Clarendon Press ; Oxford University Press.
- O'Reilly RC, Munakata Y, Frank MJ, Hazy TE, Contributors (2012) *Computational Cognitive Neuroscience*. Online Book, 4th Edition. Available at <https://github.com/CompCogNeuro/ed4>.
- Parvizi J, Jacques C, Foster BL, Withoft N, Rangarajan V, Weiner KS, Grill-Spector K (2012) Electrical Stimulation of Human Fusiform Face-Selective Regions Distorts Face Perception. *J Neurosci* 32:14915–14920.
- Patterson K, Nestor PJ, Rogers TT (2007) Where do you know what you know? The representation of semantic knowledge in the human brain. *Nat Rev Neurosci* 8:976–987.
- Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Prettenhofer P, Weiss R, Dubourg V, Vanderplas J, Passos A, Cournapeau D, Brucher M, Perrot M, Duchesnay É (2011) Scikit-Learn: Machine Learning in Python. *J Mach Learn Res* 12:2825–2830.
- Peirce J, Gray JR, Simpson S, MacAskill M, Höchenberger R, Sogo H, Kastman E, Lindeløv JK (2019) PsychoPy2: Experiments in behavior made easy. *Behav Res Methods* 51:195–203.
- Peirce JW (2007) PsychoPy—Psychophysics software in Python. *J Neurosci Methods* 162:8–13.
- Poldrack RA, Mumford JA, Nichols TE (2011) *Handbook of Functional MRI Data Analysis, 1st ed*. Cambridge University Press. Available at <https://www.cambridge.org/core/product/identifier/9780511895029/type/book>.
- Quiroga RQ, Reddy L, Kreiman G, Koch C, Fried I (2005) Invariant visual representation by single neurons in the human brain. *Nature* 435:1102–1107.
- Reddy L, Kanwisher N (2006) Coding of visual objects in the ventral stream. *Curr Opin Neurobiol* 16:408–414.
- Rissman J, Wagner AD (2012) Distributed Representations in Memory: Insights from Functional Brain Imaging. *Annu Rev Psychol* 63:101–128.
- Rodenbusch SE, Hernandez PR, Simmons SL, Dolan EL (2016) Early Engagement in Course-Based Research Increases Graduation Rates and Completion of Science, Engineering, and Mathematics Degrees. *CBE—Life Sci Educ* 15:ar20.
- Rumelhart DE, McClelland JL, AU (1986) *Parallel Distributed Processing: Explorations in the Microstructure of Cognition: Foundations*. The MIT Press. Available at <https://direct.mit.edu/books/book/4424/Parallel-Distributed-ProcessingExplorations-in-the>.
- Rumelhart DE, Todd PM (1993) Learning and connectionist representations. In: *Attention and performance 14: Synergies in experimental psychology, artificial intelligence, and cognitive neuroscience.*, pp 3–30. Cambridge, MA, US: The MIT Press.
- Russell SH, Hancock MP, McCullough J (2007) Benefits of Undergraduate Research Experiences. *Science* 316:548–549.
- Seabold S, Perktold J (2010) statsmodels: Econometric and statistical modeling with python. In: *Proceedings of 9th Python in Science Conference, Austin, 28 June-3 July*, (van der Walt S and Millman J, eds) pp 57-61. doi: 10.25080/Majora-92bf1922-011
- Sergent J, Signoret J-L (1992) Varieties of Functional Deficits in Prosopagnosia. *Cereb Cortex* 2:375–388.
- Sporns O (2010) *Networks of the Brain*. The MIT Press. Available at <https://direct.mit.edu/books/book/2149/networks-of-the-brain>.
- Starrett MJ, McAvan AS, Huffman DJ, Stokes JD, Kyle CT, Smuda DN, Kolarik BS, Laczko J, Ekstrom AD (2020) Landmarks: A solution for spatial navigation and memory experiments in virtual reality. *Behav Res Methods* Available at <http://link.springer.com/10.3758/s13428-020-01481-6>.
- Stenner T, Boulay C, Grivich M, Medine D, Kothe C, Tobiasherkze, Chausner, Grimm G, Xloem, Biancarelli A, Mansencal B, Maanen P, Frey J, Jidong Chen, Kyucrane, Powell S, Clisson P, Phfix (2023) sccn/liblsl: v1.16.2. Available at

- <https://zenodo.org/record/7978343>.
- Tanaka JW, Curran T (2001) A Neural Basis for Expert Object Recognition. *Psychol Sci* 12:43–47.
- Taube JS, Muller R, Ranck J (1990) Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *J Neurosci* 10:420–435.
- Varoquaux G et al. (2017) scipy-lectures/scipy-lecture-notes: Release 2017.1. Available at <https://zenodo.org/record/3894791>.
- Virtanen P et al. (2020) SciPy 1.0: fundamental algorithms for scientific computing in Python. *Nat Methods* 17:261–272.
- Von Der Heydt R, Peterhans E, Baumgartner G (1984) Illusory Contours and Cortical Neuron Responses. *Science* 224:1260–1262.
- Wandell BA, Winawer J (2011) Imaging retinotopic maps in the human brain. *Vision Res* 51:718–737.
- Ward J (2020) *The student's guide to cognitive neuroscience*, 4th edition. New York, NY: Routledge, Taylor & Francis Group.
- Waskom M (2021) seaborn: statistical data visualization. *J Open Source Softw* 6:3021.
- Zeki S (1991) Cerebral akinetopsia (visual motion blindness): a review. *Brain* 114:811–824.
- Zeki S, Watson JDG, Frackowiak RSJ (1993) Going beyond the information given: the relation of illusory visual motion to brain activity. *Proc R Soc Lond B Biol Sci* 252:215–222.
- Zihl J, Von Cramon D, Mai N (1983) Selective disturbance of movement vision after bilateral brain damage. *Brain* 106:313–340.

Received January 18, 2024; revised June 8, 2024; accepted June 20, 2024.

I have no known conflicts of interest to disclose. This work was supported by the Department of Psychology and Colby College. I was also supported in the development of these resources as a faculty affiliate of the Computing in Undergraduate Education (CUE) grant from the National Science Foundation: NSF1935099. I thank the principal investigators and the faculty affiliates of the CUE Grant for helpful discussions and resources. I also thank the students in my first three iterations of PS244 Cognitive Neuroscience for completing these explorations and providing valuable feedback about the efficacy of these resources in achieving our broader course learning goals. I thank Randall Downer for setting up JupyterHub for running these notebooks at Colby.

The images for Broca's area and the hippocampus in Figure 1 are from Wikipedia. The images of the "places" are from Colby College's website, the image of the "face" was courtesy of photographer Brian Huffman. All other images and figures were created by Derek J. Huffman using Python, matplotlib, seaborn, Nilearn, and Keynote.

Please reach out to me via email if you have any questions, concerns, or suggestions about these materials or anything in this paper.

Address correspondence to: Dr. Derek J. Huffman, Department of Psychology, Colby College, 5568 Mayflower Hill Drive, Waterville, ME 04901, Email: derek.huffman@colby.edu

Copyright © 2024 Faculty for Undergraduate Neuroscience
www.funjournal.org