Differential connectivity within the Parahippocampal Place Area

Christopher Baldassano [a,*], Diane M. Beck [b], Li Fei-Fei [a]

[a] Department of Computer Science, Stanford University, Stanford, CA, USA
[b] Beckman Institute and Department of Psychology, University of Illinois at Urbana-Champaign, Urbana, IL, USA

ABSTRACT

The Parahippocampal Place Area (PPA) has traditionally been considered a homogeneous region of interest, but recent evidence from both human studies and animal models has suggested that PPA may be composed of functionally distinct subunits. To investigate this hypothesis, we utilize a functional connectivity measure for fMRI that can estimate connectivity differences at the voxel level. Applying this method to whole-brain data from two experiments, we provide the first direct evidence that anterior and posterior PPA exhibit distinct connectivity patterns, with anterior PPA more strongly connected to regions in the default mode network (including the parieto-medial temporal pathway) and posterior PPA more strongly connected to occipital visual regions. We show that object sensitivity in PPA also has an anterior–posterior gradient, with stronger responses to abstract objects in posterior PPA. These findings cast doubt on the traditional view of PPA as a single coherent region, and suggest that PPA is composed of one subregion specialized for the processing of low-level visual features and object shape, and a separate subregion more involved in memory and scene context.

© 2013 Elsevier Inc. All rights reserved.

Introduction

Over the past two decades, functional magnetic resonance imaging (fMRI) has identified a number of category-selective regions involved in visual processing. Most of these regions have been defined based on differential activation to one category of stimuli over another, but this hypothesis-driven approach to mapping brain regions has significant drawbacks. Adjacent areas that have similar response profiles to the presented stimuli, but different functions, may be mistakenly conflated; for example, functionally distinct subregions have been identified in both object-sensitive lateral occipital complex (LOC) (Grill-Spector et al., 1999) and the extrastriate body area (Weiner and Grill-Spector, 2011).

Another visual region that has been proposed as a candidate for subdivision is the Parahippocampal Place Area (PPA) (Epstein and Kanwisher, 1998). This scene-sensitive area has been heavily implicated in visual scene perception, though the precise nature of the representation in this area has been controversial. Leading models have argued that PPA represents local scene geometry (Epstein et al., 2003), spatial expanse (Kravitz et al., 2011a; Park et al., 2011), space-defining objects (Mullally and Maguire, 2011), or contextual relationships (Bar, 2004). All of these models have implicitly assumed that PPA is a homogeneous unit performing a single functional role, but this view has recently been called into question. In the last several years, a number of researchers have suggested that PPA could have multiple functional components. Differences in spatial frequency response (Rajimehr et al., 2011),

varying deficits resulting from PPA lesions (Epstein, 2008), PPA's overlap with multiple visual field maps (Arcaro et al., 2009), and a clustering meta-analysis (Sewards, 2011) all hint at the possibility that PPA may be comprised of at least two functionally distinct subunits along its posterior–anterior axis. However, studies explicitly searching for a distinction between posterior and anterior PPA have failed to identify major differences (Cant and Xu, 2012; Epstein and Morgan, 2012).

Anatomical data from a proposed macaque homologue of PPA presents an interesting possibility for identifying subregions of human PPA. Although the definition of macaque PPA is still a matter of ongoing research (Nasr et al., 2011; Rajimehr et al., 2011; Sewards, 2011), a possible candidate spans cytoarchitectonically defined parahippocampal areas TH, TF, and TFO (Kravitz et al., 2011b). The most anterior area, TH, is primarily connected to retrosplenial cortex (RSC) (Kravitz et al., 2011b; Suzuki, 2009) and is also connected to the caudal inferior parietal lobule (cIPL) through a parieto-medial temporal pathway (Cavada and Goldman-Rakic, 1989; Kravitz et al., 2011b). The more posterior TF is connected to a similar set of regions, but receives stronger input from ventral visual areas V4 and TEO (Suzuki, 2009). The specific connectivity properties of the most posterior area (TFO) are not yet known, but it has been shown that TFO has a neuronal architecture highly similar to that of ventral visual regions (Saleem et al., 2007). In short, these macaque parahippocampal regions exhibit an anterior–posterior gradient, with the anterior side most related to RSC and cIPL and the posterior side most related to ventral visual areas.

Connectivity results in humans, using both diffusion tensor imaging (DTI) and fMRI, have shown that the parahippocampal region is connected to occipital visual cortex (Kim and Kim, 2005; Libby et al., 2012; Rushworth et al., 2006) as well as RSC and posterior parietal

* Corresponding author at: Department of Computer Science, Stanford University, 353 Serra Mall, Stanford, CA 94305, USA. Fax: +1 650 725 7411.
E-mail address: chrisb33@cs.stanford.edu (C. Baldassano).

cortex (Caspers et al., 2011; Kahn et al., 2008; Libby et al., 2012; Rushworth et al., 2006; Uddin et al., 2010), and PPA is known to combine both spatial and object identity information (Harel et al., 2012). However, it is not known whether the posterior and anterior parts of the PPA connect differentially to these two networks. If human PPA corresponds to some or all of the macaque areas TH/TF/TFO, it should be possible to identify an anterior–posterior gradient in the functional connectivity properties of PPA. Such a finding would not only reinforce the proposed link between PPA and these macaque parahippocampal regions, but also demonstrate that PPA is actually composed of at least two regions operating on different types of visual information, shedding new light on the controversy over its functional properties.

To test whether voxels within PPA have differing connectivity properties, we apply our recent method for learning voxel-level connectivity maps (Baldassano et al., 2012). Unlike standard functional connectivity measures that examine each voxel independently, our method considers all PPA voxels simultaneously to identify subtle differences in connectivity between voxels. After examining how several predefined ROIs connect to PPA, we perform a whole-brain searchlight analysis to identify the distinct cortical networks that connect preferentially to anterior or posterior PPA. We then demonstrate that these connectivity gradients are paired with gradients in functional selectivity, by evaluating the response to scenes and objects across PPA. Finally, we show that the connectivity gradients within PPA extend beyond PPA's borders, placing PPA in the context of ventral occipital and parahippocampal regions.

Materials and methods

Regularized connectivity method

Investigating our hypothesis requires a method that characterizes functional connectivity patterns within a region of interest (ROI), at the voxel level. A number of studies have used fMRI functional connectivity measures to investigate structure within ROIs (Chai et al., 2009; Cohen et al., 2008; Kim et al., 2010; Margulies et al., 2007, 2009; Roy et al., 2009; Zhang et al., 2008), but most previous approaches either do not measure connectivity at the voxel level (requiring spatial downsampling to a small number of subregions) and/or learn connectivity weights separately for each voxel (decreasing sensitivity and making comparisons between voxels more difficult). In our datasets, the PPA connectivity effects are too subtle to be detected by learning weights separately for each voxel (see Supplementary Fig. 1), and require the use of a method which can learn voxel-level connectivity maps that consider all voxels simultaneously. Support vector regression can learn these type of voxel-level connectivity maps (Heinzle et al., 2011), but does not utilize information about the spatial arrangement of the voxels and therefore requires a relatively large amount of data. To address this issue, we developed a method for examining connectivity differences within ROIs that is specifically tailored to small training sets typical in the fMRI setting. This method has been shown to recover voxel-level connectivity properties more accurately and efficiently than previous approaches (Baldassano et al., 2012).

The most common type of analysis for computing functional connectivity between two regions A^1 and A^2 measures how well the mean of all voxel timecourses in A^1 predicts the mean timecourse in A^2. We generalize this approach to identify voxel-level connectivity differences, by learning a *weighted* mean over the voxel timecourses in A^1 that best predicts the mean timecourse in A^2. The learned weights of the voxels in A^1 will then indicate the strength of the functional connection between each voxel and region A^2. Simply allowing each voxel weight to be learned independently leads to severe overfitting on typical fMRI datasets, but fMRI data naturally satisfies some regularity assumptions that can constrain our model. In particular, voxel connectivity properties are likely to be spatially correlated, with nearby voxels typically having more similar connectivity properties than spatially distant

voxels. This reflects a common view of cortical organization, and is especially applicable to blood-oxygen-level dependent (BOLD) signals such as fMRI, since the hemodynamic response is spatially smooth. To incorporate this assumption, we add a spatial regularization term to our model, which encourages each voxel in A^1 to have a connectivity weight similar to its spatially adjacent neighbors.

The learned connectivity maps are therefore a compromise between two objectives. Our first goal is to match the weighted average of the A^1 timecourses to the mean A^2 timecourse, by adjusting the weights. Our second goal is to make the weights spatially smooth, to prevent overfitting and allow our weights to generalize to independent data runs. The relative importance of this second goal is controlled by a hyperparameter λ, allowing us to trade off between having all weights be learned independently ($\lambda = 0$) and having all weights be identical ($\lambda = \infty$).

Mathematically, the connectivity weights are learned by minimizing the convex optimization objective

$$\underset{a,b}{\text{Minimize}} \quad ||(a^T \cdot A^1 + b) - \text{mean}_v(A^2)||_2^2 + \lambda ||D \cdot a||_2^2$$

where a is the connectivity weight map, b is a constant offset, A^1 and A^2 are the (# voxels × # timepoints) data matrices from two ROIs, and mean_v denotes an average across voxels. D is the voxel connectivity matrix, which we design to penalize the mean squared difference between the weight a_i of voxel i, and the weights of voxel i's neighbors: $||D \cdot a||_2^2 = \sum_{i=1}^{N} \frac{1}{|n_i|} \sum_{j \in n_i} (a_i - a_j)^2$ where N is the number of voxels in A^1 and n_i is the set of i's neighbors. The optimal a (for a given choice of λ) can be found efficiently by using a convex optimization package such as CVX (Grant and Boyd, 2011). For further details and validation experiments, see Baldassano et al. (2012).

The following sections describe the collection of the datasets used to learn the connectivity weights a. As will be shown in the Results, PPA's functional connectivity properties are not sensitive to the choice of experimental dataset; the specific details of the stimuli and tasks in these experiments are provided only for reference purposes.

Localizer and object-in-scene experiments

Participants

10 subjects (3 female) with normal or corrected-to-normal vision participated in the object-in-scene and localizer fMRI experiment. The study protocol was approved by the Stanford University Institutional Review Board, and all subjects gave their written informed consent.

Scanning parameters

Imaging data were acquired with a 3 T G.E. Healthcare scanner. A gradient echo, echo-planar sequence was used to obtain functional images [volume repetition time (TR), 2 s; echo time (TE), 30 ms; flip angle, 80°; matrix, 128 × 128 voxels; FOV, 20 cm; 29 oblique 3 mm slices with 1 mm gap; in-plane resolution, 1.56 × 1.56 mm]. The functional data was motion-corrected and each voxel's mean value was scaled to equal 100 (no spatial smoothing was applied). We collected a high-resolution (1 × 1 × 1 mm voxels) structural scan (SPGR; TR, 5.9 ms; TE, 2.0 ms, flip angle, 11°) in each scanning session. The structural scan was used to calculate a transformation between each subject's brain and the Talairach atlas.

Localizer stimuli and procedure

For the localizer experiment, subjects performed 2 runs, each with 12 blocks drawn equally from six categories: child faces, adult faces, indoor scenes, outdoor scenes, objects (abstract sculptures with no semantic meaning), and scrambled objects (these stimuli have been used in previous studies such as Golarai et al., 2007). Images (240 × 240 pixels; subtending 12.8 × 12.8° of visual angle) were presented at fixation.