Acknowledgements

The process of writing this thesis has been very rewarding on both an academic and a personal level, and there are people who deserve recognition, both for assisting me in the work itself and for making the process more enjoyable.

First and foremost, I would like to thank my supervisor Hallvard for his invaluable input and assistance during all stages of the work on the thesis, but I have to especially thank him for his tireless efforts in guiding and helping me through the final writing process.

I would also like to thank the entire Trondheim fMRI group for being very accommodating in helping me and for providing a great social and professional environment.

Table of contents

List of figures

List of tables

1 INTRODUCTION

1.1 Spatial function in the medial temporal lobe

The medial temporal lobe, on the ventromedial aspect of the temporal lobe, consists of the hippocampal formation, amygdaloid complex and the parahippocampal region (Franko et al., 2014). While there is some generalized function shared by all structures in the medial temporal lobe, prominently their importance for long-term declarative memory (Squire and Zola, 1996) and higher-order association feedback for memory storage (Squire and Alvarez, 1995), there is some notable functional specialization found within its structures. I will concentrate on functional specialization for spatial memory and processing, in particular along the anterior-posterior axis.

There is evidence for functional difference along the anterior-posterior medial temporal axis within the spatial domain (Poppenk et al., 2013). The anterior medial temporal lobe, containing head and body of the hippocampus, is important for the initial phase of navigation, including the global reinstatement of the environment and planning a course of action, while the posterior medial temporal lobe is more active throughout the entire course of navigation, possibly by giving more local and detailed representations of the navigated environment (Xu et al., 2010).

Here, I will specifically focus on the parahippocampal cortex and the hippocampus, while most other medial temporal lobe structures will be only briefly mentioned. For example, entorhinal cortex of rats contains grid cells, whos firing fields tile the environment in a hexagonal pattern and provide an allocentric map-like representation (Hafting et al., 2005), perirhinal cortex is centrally involved in visual memory (Malkova & Mishkin, 2003), which is important in most spatial tasks, and processing of scene information seems to be distributed throughout the MTL (Chadwick et al., 2012; Chadwick et al., 2010). However, these functional implications are outside the main scope of the current thesis: the former is specifically focused on a tangential (although important) aspect of spatial function, while the latter is too broad and would be more relevant in a thorough review of spatial function in general.

1.1.1 Hippocampus

Hippocampus is a structure of three-layered archicortex within the medial temporal lobe, and is functionally at the top of the cortical hierarchy (Lavenex and Amaral, 2000). It is

divided into three subregions along its anterior-posterior axis (Lorente De No, 1934), and based on its connectivity and cytoarchitecture it has been divided into dentate gyrus (DG), subiculum and cornu ammonis (CA) 1, 2, 3 and 4 (Andersen et al., 1971; Blaabjerg and Zimmer, 2007).

The primary source of general input projections to the hippocampus comes from the entorhinal cortex, which in turn receives its main input from the perirhinal and parahippocampal cortices (Squire et al., 2004). Specifically environmental input also comes through this route (Lavenex and Amaral, 2000), and these projections are organized in a graded manner by number of projections, with more projections to the posterior hippocampus and fewer to anterior hippocampus (Suzuki and Amaral, 1994; Witter et al., 1989).

The hippocampus has long been known to be important for memory processing (Scoville and Milner, 1957), as well as a slew of other more or less well-defined functions, including spatial memory (Poppenk et al., 2013; Strange et al., 2014). Since the discovery of hippocampal place cells in rodents (O'Keefe and Dostrovsky, 1971), it has been known that the hippocampus supports a cognitive map-like representation of physical space (O'Keefe and Dostrovsky, 1971; O'Keefe and Nadel, 1978). The cognitive map theory for the hippocampus was based on the discovery of hippocampal place cells that fire when the animal is in a particular environmental location, called the cells' place fields (O'Keefe and Dostrovsky, 1971). The firing of these place cell ensembles combine to form a distinct neural code for that specific environment (O'Keefe and Nadel, 1978). The storage capacity of these neural networks of hippocampal place cells is quite large, containing at least enough patterns to be able to store unique representations for eleven environments without pattern overlap (Alme et al., 2014). Entorhinal grid cells, one synapse upstream from the hippocampal place cells (Hafting et al., 2005) contain regularly spaced firing fields and are active across any given environment (Jeffery, 2011). Both place cells and grid cells utilize environmental cues to calibrate their firing to the outside world (Barry et al., 2007; Hafting et al., 2005), and remapping of these hippocampal maps, essentially reorganizing them in response to different environments, may be one of the fundamental mechanisms for navigation and (spatial) memory (Colgin et al., 2008). Whether or not these spatial maps are a smaller piece in the larger role of the hippocampus in memory processing is not yet known (Kyle et al., 2015).

An anterior-posterior specialization has been found in the hippocampus, where the posterior hippocampus (tail) supports fine grained representations of the local environment, the anterior (head) a coarse global representationand the middle (body) a medium-grained representation of the environment (Evensmoen et al., 2015; Evensmoen et al., 2013). This

division was also reflected in the study of London Taxi drivers, where the addition of more fine-grained local representations to their internal maps of London led to an increase in posterior hippocampal volume (Woollett and Maguire, 2011). In rats, when going from the posterior to the anterior hippocampus, the place cells' place fields increase in granularity and become less discriminating (Jung et al., 1994; Kjelstrup et al., 2008). Lesions to the posterior hippocampus lead to a reduced ability to separate close but not distant spatial locations (McTighe et al., 2009).

1.1.2 Parahippocampal cortex and parahippocampal place area

The parahippocampal cortex lies caudally adjacent to the perirhinal cortex (Lavenex and Amaral, 2000). Besides being an important input structure to the hippocampus (Squire et al., 2004), it is, together with entorhinal and perirhinal cortices, involved in an association network hierarchy within the medial temporal lobe (Lavenex and Amaral, 2000). Parahippocampal cortex receives input from areas that are involved in the visual processing of objects as well as from areas processing spatial localization information (Ranganath, 2010). It has been shown that parahippocampal cortex is important for spatial memory (Malkova & Mishkin, 2003; (Squire et al., 2004).

Functionally, a difference in spatial and navigational function has been found along the anterior-posterior axis of the parahippocampal cortex. During navigation, activation in activation in the posterior parahippocampal cortex has been observed during the entire navigation period when subjects were trying to reach specified landmarks, while anteriorly activation was only observed initially when planning how to reach the target landmark within the environment (Spiers and Maguire, 2006, 2008; Xu et al., 2010). It was argued that the posterior parahippocampal cortex is associated with local environmental details, while the anterior parahippocampal cortex is associated with the initial reinstatement of a coarse global representation of the environment (Xu et al., 2010). This reinstatement could be of importance when trying to keep a general environmental overview during retrieval of local information (Oliva & Torralba, 2006).

The parahippocampal place area is located in the posterior part of the parahippocampal cortex (Epstein and Kanwisher, 1998), a visuospatial parahippocampal region which tends to respond more strongly to scenes than faces, objects or other stimuli (Epstein et al., 2003). In addition to scene category separation, the parahippocampal place area has been implicated in several other functions (Baldassano et al., 2013), including local scene geometry, in which it

responds more strongly to place changes than object changes in a scene (Epstein et al., 2003). Lesions to the parahippocampal place area lead to difficulties navigating through both familiar and unfamiliar environments (Epstein et al., 2001), and is generally suggested to be involved with representations of local visual scenes (Epstein, 2008). Direct electrical brain stimulation of the area has also been shown to evoke visual hallicinations of topographic features (Megevand et al., 2014). A functional subdivision along the anterior-posterior axis of the parahippocampal place area has been proposed, based in part on its functional diversity (Aminoff et al., 2013; Baldassano et al., 2013; Rajimehr et al., 2011), the main hypothesis being that the posterior and anterior parahippocampal cortex is important for spatial- and non-spatial context, respectively (Aminoff et al., 2013). However, the search for this hypothesized subdivision has been largely unsuccesful (Baldassano et al., 2013).

1.2 Multi-voxel-pattern-analysis of BOLD fMRI data

1.2.1 BOLD fMRI

Neural and synaptic activity leads to increased neuronal and astrocytic energy consumption through ion pumping activity (Attwell & Laughlin 2001), metabolic demands (Iadecola and Nedergaard 2007) and feedforward neuronal pathways (Attwell & Iadecola 2002). This causes a local increase in O_2 metabolism. At the same time, neural activity causes a local increase in cerebral blood flow that is larger than the increase in local O_2 metabolism, and the net result is therefore a decrease in local deoxygenated blood hemoglobin (Kim and Ogawa, 2012). In blood oxygen level dependent (BOLD) functional magnetic resonance imaging (fMRI) (Ogawa and Lee, 1990; Ogawa et al., 1990), these local decreases in the relative amount of deoxygenated hemoglobin can be seen as an increase in the MR signal and BOLD response, because deoxygenated hemoglobin and its free iron causes local distortions in the magnetic field (Kim and Ogawa, 2012).

1.2.2 Univariate analysis: massively univariate

The historical standard analysis technique for analyzing BOLD fMRI data is the general linear model (GLM). Here, each voxel is analyzed individually, and the statistical model compares each single voxel's MR signal intensity over time to the experimental time series (Friston, Holmes, Worsley & Poline 1995). After accounting for the delay of the hemodynamic

response to brain activity, these two time series can be compared to evaluate how the experimental manipulations correlate with changes in brain activity. This is called the massively univariate approach (Luo & Nichols 2003), since each voxel is analyzed individually as a single dependent regression variable, on the basis of one or multiple predictor variables, i.e. experimental stimuli as well as known nuisance variables like movement and scanner drift.

1.2.3 What univariate analyses miss

Since the goal of univariate fMRI analysis often is to investigate the overall activations in a brain region of interest (ROI), the data is spatially smoothed in order to increase the signal to noise ratio at the subject level and focus the sensitivity on the gross activations (Hopfinger et al., 2000). The effect of this smoothing on the underlying voxel activity distribution is complicated, and the main discussion revolves around whether or not an underlying neural/voxel information pattern is conserved or lost through the effects of spatial smoothing (Op de Beeck 2010;Kriegeskorte et al. 2010; Kamitani & Sawahata 2010). No matter what the result of the smoothing is, in a univariate analysis the information content of voxel patterns are not directly assessed as the model treats each voxel separately. Since it can be argued that the brain is inherently multivariate in nature, with both microscopic (neurons and neuronal populations) and macroscopic (cortical and subcortical regions) working together in distributed networks, the univariate approach may not be enough on its own or even the optimal choice at all in a given investigative context (Raizada & Kriegeskorte 2010). To account for more distributed activity patterns in the brain, we may need to directly measure and contrast the changes in voxel activation patterns across experimental conditions, brain regions and/or people.

1.2.4 Multivariate/multivoxel pattern analyses

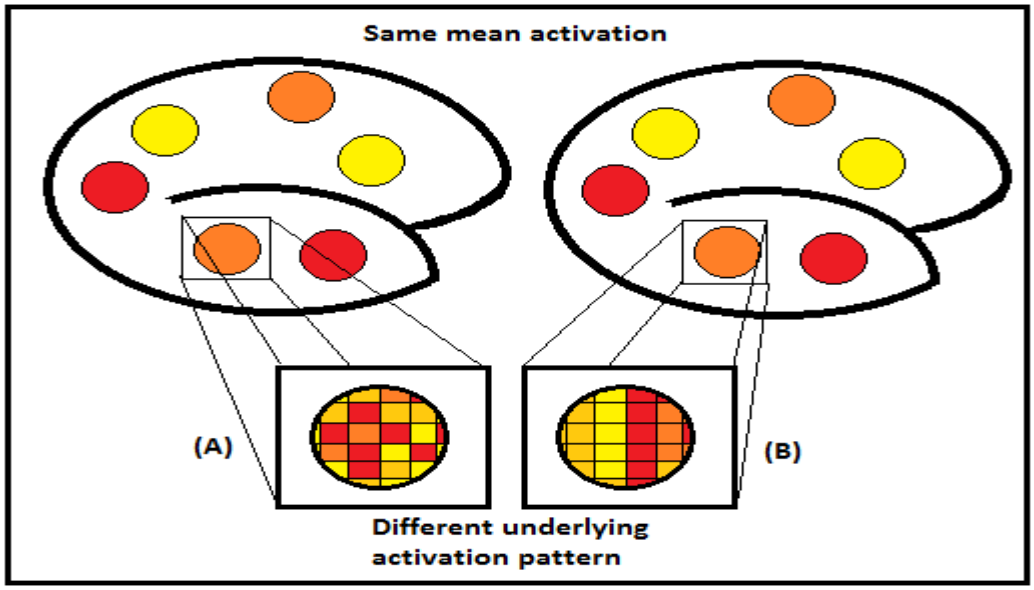

Figure 1-1 Univariate vs Multivariate Activity. The top figures are side views of two brains. Colored circles are mean BOLD fMRI activation across multiple voxels in a region. In a univariate analysis (top), mean regional voxel activations are extracted (orange --> red --> yellow = higher activation compared to baseline). However, the underlying distribution of voxel activations (bottom) can be different even if their shared mean activity is the same. Here, (A) reveals a substantially different underlying pattern compared to (B), but their mean activations over a larger area (eg. ROI) are equivalent. This reveals a key difference between univariate analysis (top) and multivariate pattern analysis (bottom).

Multivariate pattern analysis (MVPA) of fMRI data, also often called multivoxel pattern analysis, since it's a multivariate analysis using the MRI voxels as its dependent variables, is qualitatively separate from the standard univariate analysis. The main difference is that MVPA looks at the pattern of BOLD fMRI signal across many voxels instead of focusing on the individual voxel-wise intensity changes (Norman et al. 2006). Figure 1-1 shows a simple illustration of the main differences between what a univariate and multivariate analysis will be able to detect. If the goal, or one of the goals, of an analysis is to account for the informational content in across-voxels patterns, one of the several MVPA approaches may then be preferable. Importantly, MVPA has been able to detect stable patterns of activity across voxels even in the absence of mean activation (Mur, Bandettini & Kriegeskorte 2009), for example in the orientation sensitivity of human primary visual cortex (Kamitani & Tong 2005; Haynes & Rees 2005). Subsequent work has shown that MVPA can extract contents of higher-order functions, for example intentions (Haynes et al. 2007), working memory (Harrison & Tong 2009) and episodic memory (Collin, Milivojevic & Doeller 2015). It is important to note however, that

this does not necessarily reflect a superiority of multivariate over univariate analyses in fMRI studies, but it does mean that it may be able to detect different aspects of the brain's response to a stimulus (Jimura and Poldrack, 2012). While two single locations analyzed in univariate analyses may not individually carry critical information about a cognitive, sensory or emotional state, analyzing both with a multivariate analysis may show that they convey information in their collective pattern, increasing the information that can be extracted from the measurement (Haynes and Rees, 2006). Generally, univariate analyses may be sensitive to activation of basic processing functions while MVPA detects representations of the specific contents being processed (Raizada et al., 2010).

MVPA has been shown to have the potential to be more sensitive to distributed coding of information, while univariate analyses tend to be more sensitive to the global engagement of the brain during a task/stimulus. MVPA has the potential to reveal activity in additional regions not discovered by a univariate analysis (Jimura & Poldrack 2011). In general, it seems that the MVPA approach yields not just a more sensitive answer to the same problems that we address when using univariate analysis, but that it can answer fundamentally different questions. For example, when looking at pattern information, we can look at spatially overlapping neural representations, which are not amenable to univariate investigations due to the localized pattern distinctions being smoothed away (Raizada & Kriegeskorte 2009) and the single-voxel focus of the approach (Haynes and Rees, 2006). Figure 1-2 shows the difference clearly: while looking only at the total amount of signal in an area (represented here by the total amount of digital ink) would yield the same result for both the letter 'M' and the arrow, looking at the pattern of activations across the area reveals the actual difference.

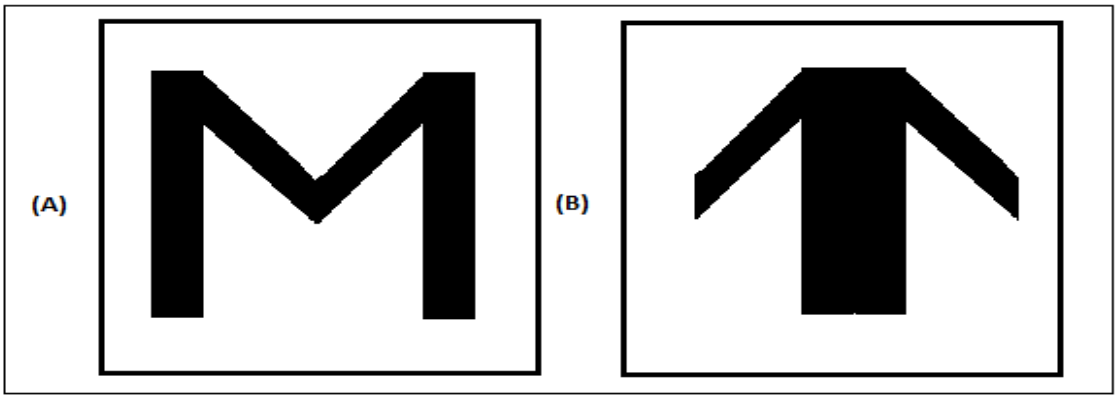

Figure 1-2 Univariate and multivariate data example. The two squares contain the exact same amount of 'signal' (represented by amount of "ink"). However, their spatial distributions are strikingly dissimilar. This illustrates the need to sometimes look deeper than the mean signal intensity in an area and also take into account the possibility that the distribution of activity across the voxels in the area might differ in an important way.

1.2.5 Representational spaces and representational geometries

In the context of neuroscience, neuronal activity (and by extension voxel activity in fMRI) is interpreted as representing some mental content, like a memory, a visually observed image or a motor action, and the researchers' interpretations of these representations can then be used to link the neuronal activity to cognition and more abstract non-physical aspects of brain activity (Kriegeskorte and Kievit, 2013). In order to quantitatively analyze and compare these representations, we need a way to look at them that lets us compute differences between them in a meaningful way. Two ways of constructing and analyzing such representational spaces are: Representational geometries and difference measures in representational similarity analysis (Kriegeskorte and Kievit, 2013; Kriegeskorte et al., 2008) and multi-subject alignment in high-dimensional spaces through hyperalignment (Haxby et al., 2011).

What these have in common is the idea that the representations can be quantified as vectors in a high-dimensional space defined by the dimensions (x,-y,-z-, etc. axes) of the measurement, e.g. voxel activation in fMRI (Haxby et al., 2011; Kriegeskorte et al., 2008). In fMRI, this representational space is a high-dimensional space where the axes (dimensions) are separate voxels, each value on an axis is the voxel intensity for that voxel, and all stimuli are represented as separate vectors defined by their values for each axis/voxel (Haxby et al., 2014) (figure 1-3 A and 1-3 B). So, using only 2 or 3 voxels, we get an intuitively simple picture in two and three dimensions. Figure 1-3B and 1-3C shows the same thing: 1-3C shows the voxel activity for the blue and red stimuli across the three measured voxels, while 1-3B shows how this pattern will be reflected in a high-dimensional space. However, an fMRI experiment will often use several thousand voxels, and this number of dimensions is impossible to convey or even imagine in a simple way, since each additional voxel gives one more axis and thus builds more than three dimensions. Still, the concept is essentially the same if the dimensions are three or twenty thousand. Since the points (or vectors) are defined by their voxel response intensities, we can think of the representational space as a voxel response pattern space.

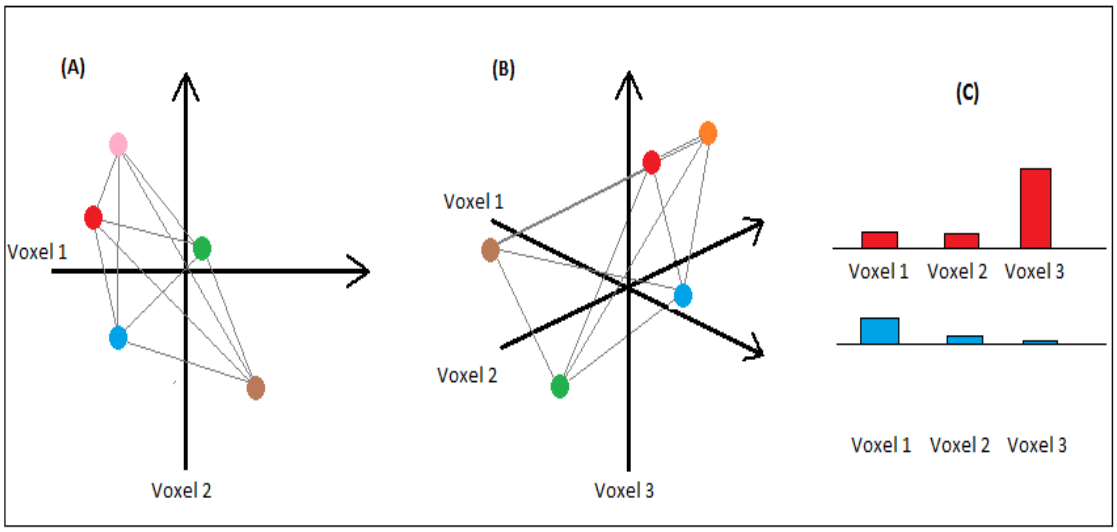

Figure 1-3 Representational spaces . The axes indicate the voxel intensities, with one axis per voxel: A) two voxels gives the space two dimensions for the stimuli (responses) to exist in. (B) three voxels gives three dimensions to the space. Note: each point illustrates the end point coordinate of a vector starting at the origin of the coordinate system, and each component value of the vector corresponds to the intensity score of a voxel (eg. the leftmost (red) point in (A) may represent the vector [-3, 1] for voxels 1 and 2, while the blue point in (B) may represent the vector [3, 0.5, 0], for intensities of voxels 1, 2 and 3 respectively). (C) shows how each point is actually a spatial pattern across voxels; for the red and blue stimulus points in (B), the relative voxel intensities are different, placing them at different coordinates within the three dimensional representational space.

1.2.6 Representational similarity analysis and representational dissimilarity matrices

The activation patterns can also be analyzed by looking at the representational geometry of stimulus responses across different stimuli. The representational geometries for a stimulus set can then be compared across brain areas, people or even between species or computer models. Since the responses are represented in a space, we can perform these comparisons by calculating the distances between the stimulus responses as the similarities (or dissimilarities[1]) between the different patterns of stimulus responses (Kriegeskorte and Kievit, 2013).

Representational similarity analysis (RSA) is the most common method for comparing representational similarities in brain activation patterns (Mumford, 2013). In the context of fMRI, RSA correlates the activity patterns across spatially extended voxels with the experimental stimuli,by taking correlations between all task pairs, in order to quantify the similarity between these activity patterns. The differences between the stimulus responses can

[1] The use of dissimilarities instead of similarities comes from the relative intuitive ease of imagining dissimilarity as an actual distance in real space (more dissimilar = further apart) (Kriegeskorte et al. 2011).

then be represented as their dissimilarities across the voxel space (the grey lines in figure 1-3) and quantified as 1 - correlation across space (1 - r). These dissimilarity scores can then be plotted as a square matrix with n rows and n columns (n = number of stimuli), as in figure 1-4. This representational dissimilarity matrix (RDM) lets us characterize the representational geometries (Kriegeskorte et al., 2008). The resulting RSA values can then be compared in order to show the representational differences between eg. brain regions, subjects, species or computational models (Kriegeskorte et al. 2008; Kriegeskorte 2011). Particularly relevant here, RSA has been used successfully to investigate a scaling of representations along the hippocampal long axis (Collin, Milivojevic & Doeller, 2015). Figure 1-4 shows how the dissimilarities in voxel patterns associated with stimulus pairs can be computed as differences in spatial correlations across voxels.

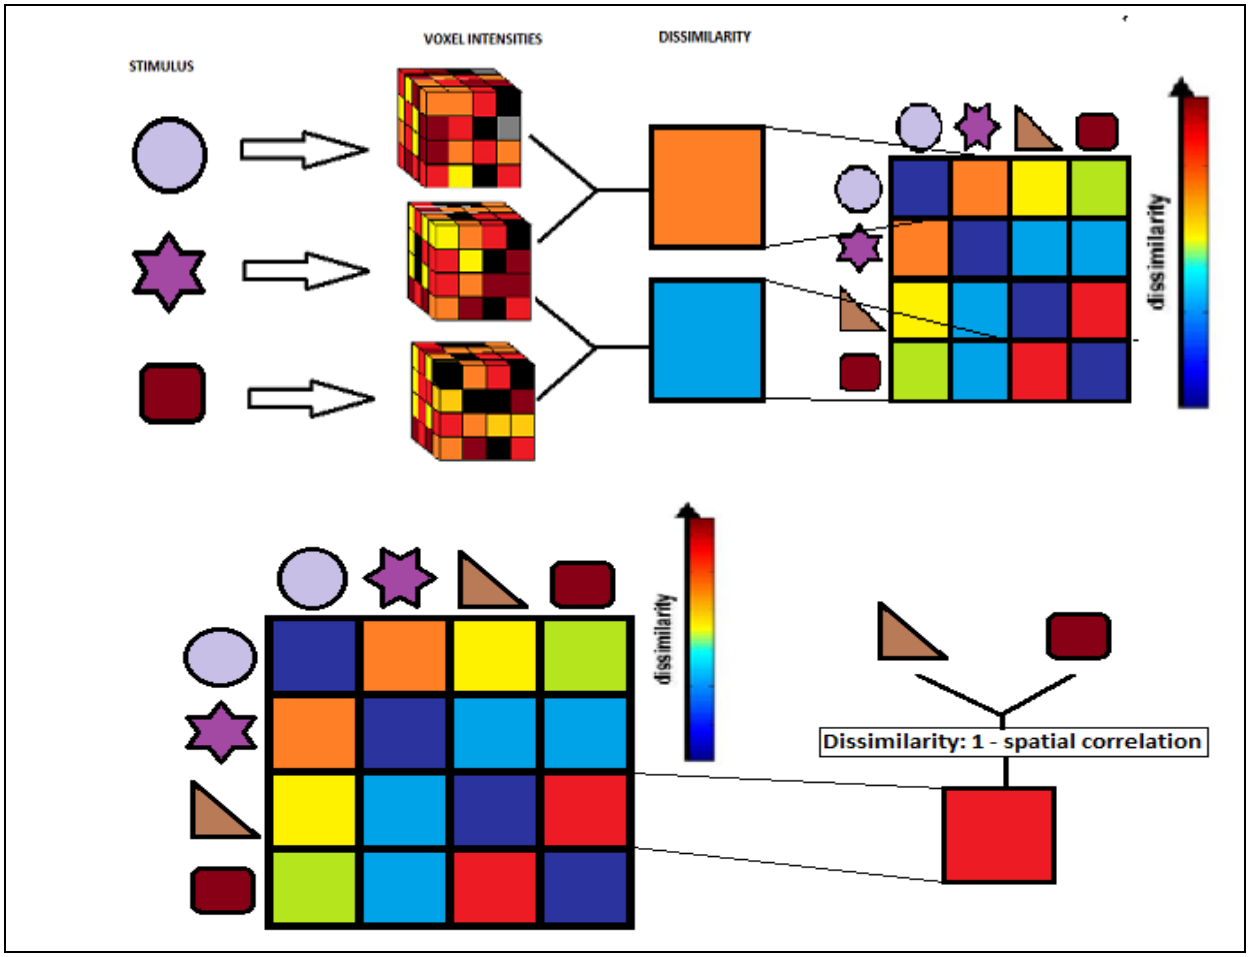

Figure 1-4 Dissimilarity matrix. Top left: Different stimuli lead to different patterns of voxel activiations in an area of the brain (top middle). Top right and bottom left: The colored shapes represent four stimuli, or experimental conditions. For each stimulus pair, the voxel activity pattern associated with those stimuli are compared through spatial correlation. The colors in the square matrix represents a dissimilarity measure between the stimuli of 1 – correlation, so that 0 represents a perfect correlation, 1 no correlation and 2 perfect anticorrelation (note that the matrix is symmetric about the zero diagonal, since all stimuli are compared twice and have perfect (correlations (dark blue) with themselves. Bottom right: how each square is determined. Bottom middle: color coding scheme for the dissimilarity; red is most dissimilar (color bar from Kriegeskorte et al. 2011).

1.2.7 Classifiers and (linear) vector support machines

In order to investigate the correspondence between neural activity patterns and mental representations, different classifiers can be employed. For example, it has become common practice to use machine learning algorithms to calculate the similarity measures used in searchlight analyses and for separation classification of other data structures (Etzel et al., 2013). As discussed in the section on representational spaces (1.2.5), multivoxel activation patterns can be thought of as points in an n-dimensional space (where n equals number of voxels used in the analysis). A linear classifier will then use a linear decision boundary (hyperplane) to separate the response points (multivoxel activity vector endpoints in the space) (figure 1-5). One such classifier is the linear support vector machine (linear SVM), where the 'linear' just means that it uses a linear hyperplane as a separator (Mur et al., 2009).

Linear SVMs are especially popular for these types of classifications in fMRI (Etzel et al., 2013). Several linear classifiers exist, but while they generally seem to perform equally well on fMRI data, the linear SVM is very suitable because it does not assume that the data are multivariate normal and does not use pattern correlation, which may make a classifier insensitive to regional-average differences (Mur et al., 2009). The linear SVM works by margin maximization. It classifies the data of two classes by finding a linear classification boundary hyperplane that separates the classes by assuming an orientation and position that gives it the maximal possible distance from the nearest data point. The classifier is then defined by the points on the margin edge, called the support vectors, and if the classification performs better than chance, based on the activity patterns, this can be taken to signify that the pattern carries important information about the experimental task or condition (Mur et al., 2009). Importantly, this is done on a subset of the data (called the 'training set' as it is what the SVM is using to find the best separating hyperplane) before it is then applied to the rest of the data ('testing set') to see if the classification holds. This grouping into training and testing sets can be achieved by splitting the data randomly into two parts or collecting separate training and testing data (Etzel et al., 2013).

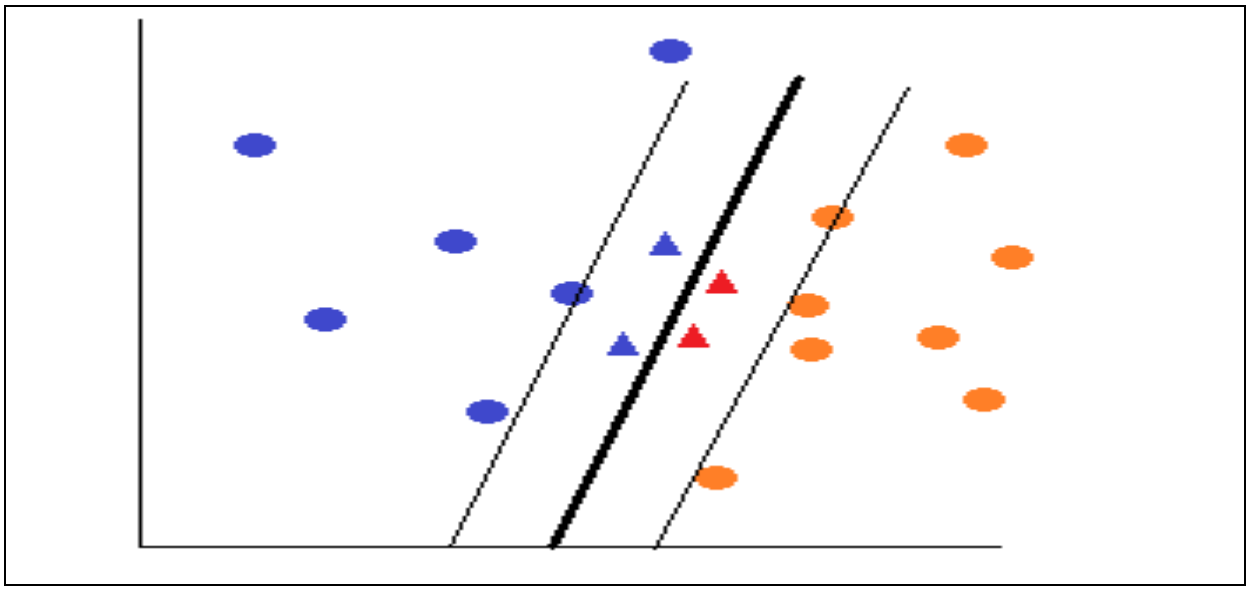

Figure 1-5 Linear support vector machine. The linear support vector machine (SVM) is a classifiser that finds the optimal separating hyperplane between two classes of patterns by first training itself to separate one part of the data and then using the test classifier on another part of the data. The test data classifier thereby determines how succesful the classifier is at separating the data correctly into the two real categories. Thick line = optimal hyperplane. Thin lines = maximum margins on either side, defined only by the closes data points on each side (i.e. each class). Triangles = testing patterns. Circles = training patterns.

1.2.8 Multivoxel pattern analyses of spatial functions in the hippocampus and parahippocampal cortex

Hippocampal activation patterns represent individual positions. Encoding of individual positions refers to how specific locations are encoded without reference to any other object or position, whether the position refers to the position of an object, a goal position or the subject's own current location. For example subjects' locations while navigating between four cornerns in a virtual environment has been succesfully decoded through hippocampal voxel patterns using MVPA, even when the locations were visually matched (Hassabis et al., 2009). Voxel pattern classification has also showed that the hippocampus can separate between three individual goal locations in a circular environment, when navigating to them (Rodriguez, 2010). Neural patterns associated with people's senses of their location within the scanner showed a decoding peak in a voxel cluster extending from the posterior hippocampus into the border between the hippocampus and the parahippocampal gyrus.

Hippocampal and parahippocampal activation patterns also seem to represent the internal relationship between positions independent of our own positions, referred to as an allocentric representation (Ekstrom et al., 2014). A cluster in the entorhinal subiculum showed

a highly similar activation pattern for geocentric direction trials,in which they had to remember the internal relationship between at least two background images and an object, but not egocentric direction (Chadwick et al., 2015). Additonally, voxel activation patterns in parahippocampal cortex and posterior hippocampus have been shown to be more similar the more locations four objects in a virtual scene had in common, while pattern similarity in the anterior hippocampus, for the encoding period compared to the delay period, was associated with object-location encoding (Libby et al., 2014). This suggests that the internal relationship between positions, at least on some level, is embedded in the activation patterns in the hippocampus, parahippocampal cortex and entorhinal cortex. Allocentric encoding is seen as important for creating a cognitive map-like representation of the environment (Ekstrom et al., 2014). Worse map-drawing ability for previously experienced environments has been shown to correlate with increased pattern similarity between environments with a regulary shaped and morphed outer wall, in a hippocampal-parahippocampal cortical cluster (Stokes et al., 2015). A positive correlation was also observed between the ability of a hippocampal-parahippocampal cortical cluster's ability to separeate virtual cities and the ability to retrieve the relative distance of the stores within the cities (Kyle et al., 2015). Taken together, this indicates that activation patterns in the hippocampus and parahippocampal cortex supports representation of the internal relationship between objects, possibly in an allocentric framework.

Hippocampal and parahippocampal activation patterns also seem to involve more coherent environmental representations. When subjects navigated in four small virtual environments, it was found that voxel patterns in the parahippocampal cortex separated between these environments (Hassabis et al., 2009). The hippocampus has also been found to contain neural codes in its voxel activation patterns able to correctly classifiy different virtual environments (Kyle et al., 2015). Finally, environments with similar outer borders are associated with more similar activation patterns in both parahippocampal cortex and hippocampus (Stokes et al., 2015). In total, this suggests that both the hippocampus and parahippocampal cortex, through their activation patterns, are capable of separating between environments.

1.3 Hypotheses and aims

Activity in the hippocampus and parahippocampal cortex, including the parahippocampal place area, have an important role in localization and positional encoding within our spatial memories. Previous studies have investigated how location and position are encoded in the patterns of distributed activity within these regions. Still, a study that investigates encoding of the different parts that an environmental representation consists of, by having the participants reconstruct the environment after encoding, is missing. This is crucially important, because it makes it possible for example to look at encoding of allocentric positional representations per se while controlling for encoding of position-object, position-room, object-room and object representations.

The overlying goal of this thesis was to investigate dissimilarities in patterns of voxel activity, and the classification performance based on these patterns, of the different elements that our environmental representations consist of along the anterior-posterior axis of primarily the hippocampus and parahippocampal cortex. To test this hypothesis we analyzed BOLD fMRI activation patterns from virtual environmental learning, involving both a stimulus presentation and a poststimulus encoding period, using a correlation-based similarity measure and linear support vector machine learning for classification. Each of the 35 small virtual environments used consisted of an outer wall and five objects. For each environment, the participants were tested on their ability to position the objects from a map-like perspective (position test), to associate the objects from the same environment together (objects test), and to associate the objects with the outer wall (objects-room test). The position test gave both a measure of the ability to reconstruct the objects' positional pattern, to position the objects correctly within the pattern, and to place the objects correctly relative to the outer wall.

Univariate analyses have shown that within the hippocampus, and to some extent parahippocampal cortex, anterior parts involve global environmental representations and posterior parts more local environmental representations. Additionally, multivoxel fMRI analyses have been able to detect a gradient of memory representations along the hippocampal long axis, with more similar activation patterns for large-scale memory representations anteriorly, medium-scaled intermediately and small-scale posteriorly (Collin et al., 2015). We therefore hypothesized that anterior hippocampus and parahippocampal cortex would encode individual environments more separate, including their outer wall and objects, through more dissimilar activation patterns than posterior hippocampus and parahippocampal cortex.

Further, univariate analyses have found a positional gradient along the anterior-posterior axis in these data, with fine-grained, medium-grained and coarse-grained positional representations in the posterior, intermediate and anterior hippocampus respectively. We therefore hypothesized that activation patterns in more posterior parts of the hippocampus and parahippocampal cortex would be especially good at correctly classifying fine-grained positional representations, intermediate parts medium-grained representations and anterior parts coarse-grained representations.

Finally, we hypothesized that the posterior part of the parahippocampal cortex, including the parahippocampal place area, will be more stimulus driven, showing increased ability to separate the stimulus presentation- from the poststimulus encoding period.

2 METHODS

2.1 Subjects

Thirty-one right-handed men (18 - 27 years**,** mean 21 years) were recruited. None of them reported any history of neurological disorders, head trauma, or previous or current DSM-IV axis I diagnosis of psychiatric illness, including substance abuse.

2.2 Virtual Reality Environment

The VR-environments were developed in collaboration with Terra Vision AS (Terra Vision, Trondheim, Norway) using the Torque game engine (Garage Games, Eugene, Oregon, US). There were 35 environments, each between 50 and 90 m^2, and each had five unique unrelated objects (figure 2-1 A and B). Every environment had a unique positional pattern of objects. Each object was positioned in one of 16 possible tiles on the floor, and additionally assigned to one of five possible sub-squares within each tile (figure 2-1 B). Furthermore, the outer wall of each environment had one of 10 geometric shapes (figure 2-1 C). The participant moved freely in the environment, but the moving speed was set to 2 m/s (figure 2-1 A).

2.3 fMRI paradigm and test procedures

Using a block design fMRI paradigm, all subjects had to learn 35 unique environments. The paradigm included a 30 second stimulus presentation, a 15 second poststimulus encoding period and a 15 second period where they had to make unrelated odd-even judgments.

When the stimulus presentation period started, subjects were positioned at the door in the environment, and then had to explore the virtual environments freely from a first person perspective (figure 2-1A, left image). During the subsequent poststimulus encoding period, subjects were told to fixate on a cross while continuing to encode the features in the environment. Then, they were given an odd-even task where they were shown random numbers (<100) and told to push the right joystick button if the number was even and the left button if it was odd. This period was used as an implicit baseline.

The order of the 35 environments was randomized both within and between runs, with five new environments in each run over a total of seven runs per subject.

After each run, subjects were given tests to assess the accuracy and success in their recall of the different spatial and non-spatial representations in each of the five learned

environments in the preceding run (figure 2-1). These tests were administered by having the subjects look at a computer screen while they were still in the scanner, and they answered by clicking and dragging objects on the screen using a joystick. All tests were programmed using HTML and javascript (Hansen et al., 2015).

First, they were given the objects test, which assessed their ability to successfully recall which objects were presented together in an environment. Then, the objects-room association test assessed their ability to successfully assign the correct group of objects to the correct room geometry. The final test was an object-positioning test, where they were given several scores based on different types of spatial associations (further details below and in (Evensmoen et al., 2013).

2.4 MRI scanning

Scanning was performed using a 32-channel Head Matrix Coil for all subjects. For the first 24 subjects, a 3T Siemens Trio scanner was used, while the final seven were scanned in a 3T Siemens Skyra scanner (Siemens AG, Erlangen, Germany). To minimize head motion, the subjects were fitted with foam pads. Subjects moved through the virtual environment using an MRI compatible joystick (Current Designs, Philadelphia, US) while the stimuli were presented on an LCD monitor with 1280 x 1024 resolution (Current Designs, Philadelphia, US).

First, subjects familiarized themselves with the equipment, including the monitor and the joystick, and the paradigm by completing practice trials of the different experimental conditions. When the subject showed complete task compliance, the MR scanning started.

2.5 Imaging parameters

T2* weighted, blood-oxygen-level-dependent (BOLD) sensitive images were acquired during the spatial encoding task with a 2D echo-planar imaging pulse sequence. Since a different scanner was used for the first 24 and the final 7 subjects, the 2D echo-planar imaging parameters differed slightly between the two groups.

For the first 24 subjects, the parameters were as follows: TR = 2110.8 ms, TE = 28 ms, FOV = 220 mm x 220 mm, slice thickness = 1.9 mm (no gap), number of slices= 40, matrix = 116x116 producing a voxel size of 1.9x1.9x1.9mm^3, flip angle = 90°.

For the final seven subjects, the corresponding parameters were: TR = 2253.2 ms, TE = 28 ms, FOV = 220 mm x 220 mm, slice thickness = 2.0 mm (no gap), slice number = 40, matrix = 116x116 giving a voxel size of 1.9x1.9x2.0mm^3 , flip angle = 90°.

Imaging slices were positioned as close to 90° on the anterior-posterior direction of the hippocampus as possible, and GRAPPA acceleration was used (factor four). Each functional run consisted of 143 volumes for the first 24 subjects, and 134 volumes for the last seven subjects.

A T1 weighted 3D volume was acquired for anatomical reference, using an MPRage sequence with parameters: TR = 2300 ms, TE = 2.94 ms, FOV = 256 mmx 256 mm x 192 mm, matrix 256x256x192 giving an resolution of 1.0x1.0x1.0 mm^3 , flip angle = 8°.

2.6 Tests of non-positional representations

In the objects test, subjects were shown five groups of objects from the five environments they had just learned. In each group, only three out of the five objects were pre-assigned, and from a group of ten other objects, they had to drag the two missing objects in each group over to their correct group (2-1B). They were calculated according to the total number of correct trials (all five groups completed correctly), the number of groups where only one of the two extra objects were assigned correctly, and failed trials (groups where both added objects were wrong).

In the objects-room test, subjects dragged object groups to a 2D overview of the environment. Five of these overviews were lures (an environment they not had been shown in the recent run). For each room they associated with the correct object group, they got one point. The score was calculated by summing the total number of correct trials, and a score for total failed trials was also calculated by summing these.

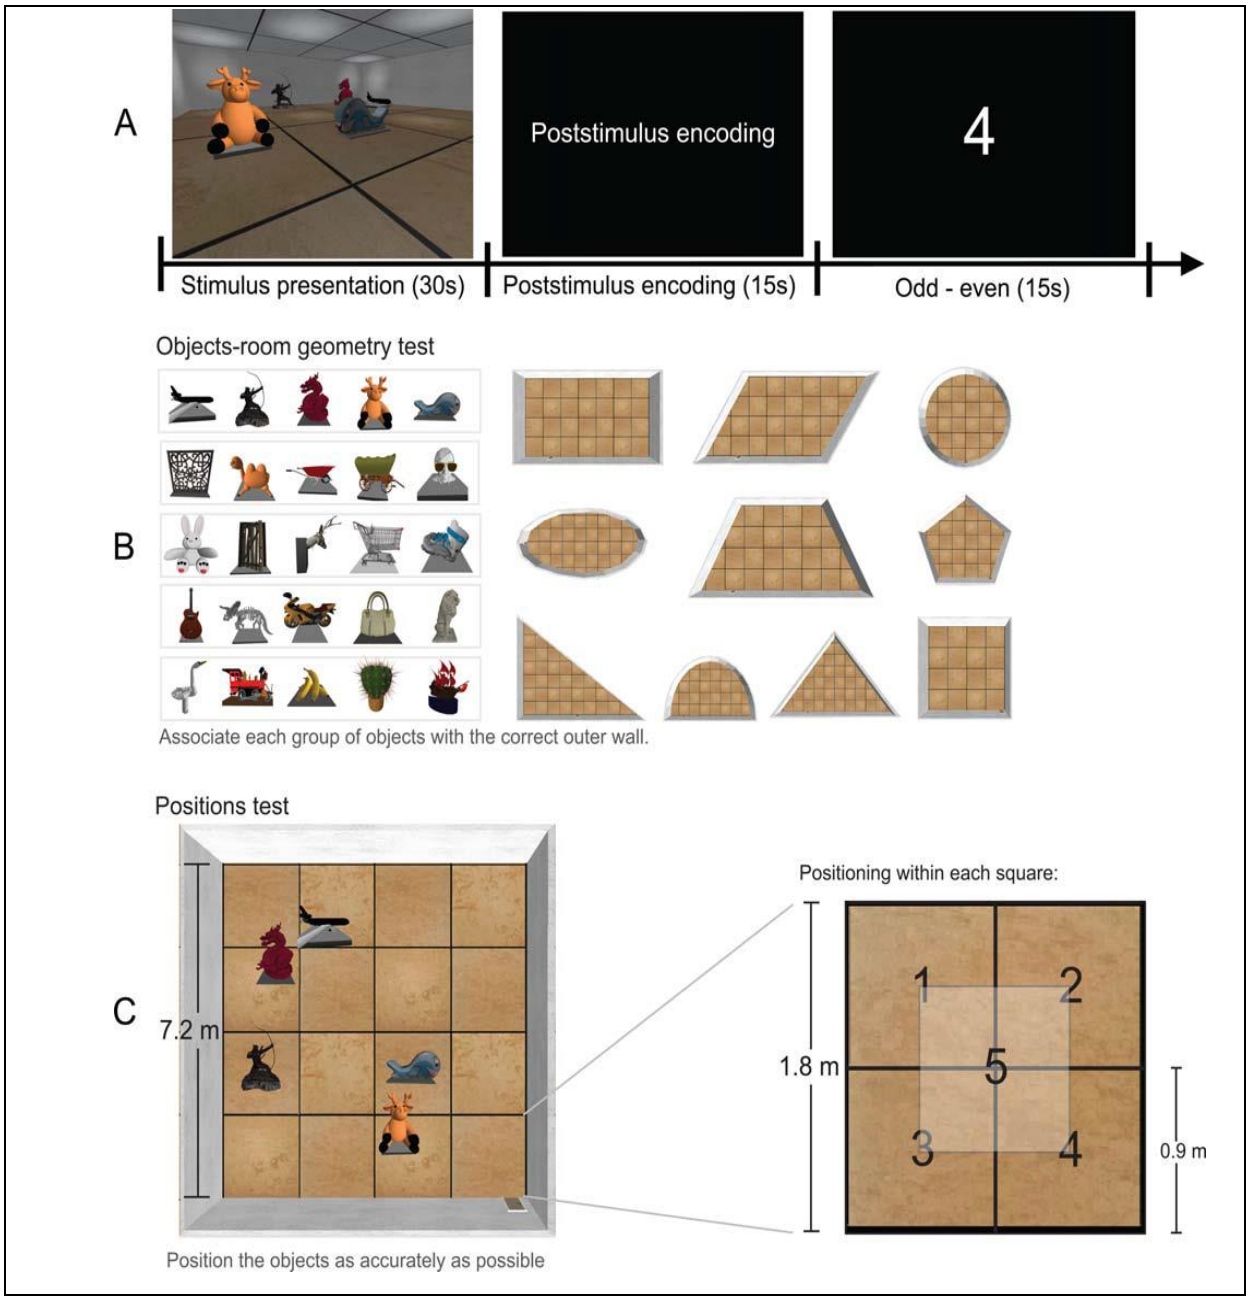

Figure 2-1 Experimental setup. A: fMRI paradigm, with 30 s stimulus presentation period (free exploration), 15 s poststimulus encoding period (memorization of environment) and 15 s where subjects completed odd-even judgments. After five environments were learned, the Objects-room geometry test (B), where they associated a group of objects with the associated geometric shape of the environment, and the Positions test (C), where they positioned the five object as accurately as possible within the 16 squares in the environment, each consisting of 5 subsquares. From this test, a positional accuracy score was calculated. Image reproduced from original study (Evensmoen et al., 2015).

2.7 Tests of positional representations

In the object-position test, subjects had to drag and place all 5 objects associated with an environment as accurately as possible onto the 2D overview of the environment (figure 2-1C). Several different measures were derived from performance on this task:

The positional granularity measure represents each subject's ability to accurately reconstruct the positional pattern of the objects in the implicit coordinate system of the environment, irrespective of object identity and room geometry (translational, rotational and scaling effects between the objects and the environment were removed). This shows how well each subject recreated the positional pattern of objects in an environment, and it represents an allocentric measure, meaning that it reflects the relationships between the objects independently of their position (Ekstrom et al., 2014; O'Keefe and Nadel, 1978; Tolman, 1948). Accuracy scores, measuring how far subjects' reconstruction of the positional pattern of the objects was away from the correct positions, were calculated by summing the total sum of squares from the correct positions. This measure was used as both a continuous and discrete measure (divided into fine-, medium-, coarse-grained, and failed trials). In fine-grained representations, the objects' positioning was placed within an average of one-fourth of a square (figure 2-1C, right image), medium-grained representations were within one square, and coarse-grained representations within one and a half squares. Representations were classified as failed if the score was more than one and a half squares away from the original position.

In the positions-room granularity analysis, scores above 16000 were discarded from the analysis to make sure that a rough idea of the positional pattern was present in all included trials. Object identity was also disregarded here. The granularity measure was estimated by the amount of translation, rotation and scaling that had to be applied to each subject's positional pattern to get the pattern that most closely resembled the original, using least squares. Fine-, medium- and coarse-grained representations were made based on the same criteria as for the positional granularity measure.

2.8 Data preparation and preprocessing

The data were preprocessed using FSL 5.0.9 (Analysis Group, FMRIB, Oxford, UK). First, the seven fMRI runs were motion corrected using MCFLIRT, with the median volume of run 3 as reference, before they were merged using fslmerge. Finally, to exclude voxels in the fMRI data that were most certainly not part of the brain, BET 2 with robust center estimation (Brain Extraction Tool, FMRIB, Oxford, UK) was run on the merged fMRI data file, with a liberal fractional intensity threshold of 0.25.

2.9 Anatomical ROIs

The parahippocampal cortex (PHC) was defined bilaterally combining probabilistic maps of the Harvard Oxford Structural Atlases (part of FSL; http://fsl.fmrib.ox.ac.uk/fsl/fslwiki/Atlases) with defined anatomical boundaries (Franko et al., 2014; Insausti et al., 1998; Pruessner et al., 2002) (figure 3-1). The posterior PHC border was defined as the final coronal slice with visible hippocampal gray matter on the MNI template , i.e. MNI: -42 (Pruessner et al., 2002), and the anterior PHC border 6 mm posteriorly to the end of the gyrus intralimbicus (Franko et al., 2014). The PHC was then divided into three anterior-posterior subregions, i.e. the posterior PHC (y = -42 to y = -38, 1389 voxels), intermediate PHC (y = -37 to y = -33, 1290 voxels), and anterior PHC (y = -32 to y = -27, 1315 voxels). In order to investigate the functional borders of the PHC more closely, a posterior, i.e. lingual, ROI (y = -50 to y = -43, 1850 voxels) and an anterior, i.e. perirhinal/entorhinal, ROI (y = -26 to y = -22, 827 voxels) were defined.

The hippocampus was divided into three anatomical ROIs; posterior (1098 voxels), intermediate (3169 voxels), and anterior (4385 voxels) (Evensmoen et al., 2015; Poppenk et al., 2013).

The caudate nucleus was defined bilaterally (probabilistic maps of the Harvard Oxford Structural Atlases (part of FSL; http://fsl.fmrib.ox.ac.uk/fsl/fslwiki/Atlases), using a probability threshold of 50. The caudate nucleus was divided into three anterior-posterior subregions, i.e. the posterior caudate (y =-21 to y = -4, 619 voxels), intermediate caudate (y = -3 to y = 10, 2810 voxels), and anterior caudate (y = 11 to y = 25, 4033 voxels). The posterior caudate involved more slices because of the relative small size of the posterior slices. Putamen was divided into three anterior-posterior subregions, i.e. the posterior putamen (y = -22 to y = -8, 2609 voxels), intermediate putamen (y = -7 to y = 5, 4994 voxels), and anterior putamen (y = 6 to y = 20, 4688 voxels). The intermediate putamen involved less slices because of the relative large size of the intermediate slices.

The cingulate cortex was also defined bilaterally (probabilistic maps of the Harvard Oxford Structural Atlases (part of FSL; http://fsl.fmrib.ox.ac.uk/fsl/fslwiki/Atlases), using a probability threshold of 50. The cingulate cortex was divided into three anterior-posterior subregions (figure 3-1).

The ventricle ROI was created by created a spherical shaped ROI in the anterior horn of each of the lateral ventricles (370 voxels).

The fusiform ROI was defined as the temporal occipital fusiform cortex bilateral (probabilistic maps of the Harvard Oxford Structural Atlases (part of FSL; http://fsl.fmrib.ox.ac.uk/fsl/fslwiki/Atlases), using a probability threshold of 50 (5066 voxels).

2.10 Dissimilarity matrices

In order to investigate the similarity in the ROI activation pattern across different environments and different levels of encoding, representational dissimilarity matrices (RDMs) were created using the pyMVPA toolbox (Hanke et al., 2009). After minimal detrending, all samples from each condition (see each individual RDM compilation for description of conditions used for them) were averaged into a single exemplar. Correlation distances were computed using a Pearson correlation measure, and the resulting per-subject matrices were color graded by dissimilarity (1 minus Pearson's r). In specified cases, the span of dissimilarities were truncated to show only the negative or positive correlations.

2.11 Classification analyses

Within-subject classification for levels of granularity (fine- vs. coarse-grained, fine-vs. medium-grained and coarse- vs. medium-grained) during the poststimulus encoding period, main experimental periods (poststimulus vs. stimulus, poststimulus vs. odd-even task, and stimulus vs. odd-even) was carried out for each ROI in the python-based pyMVPA toolbox (Hanke et al., 2009), using a linear c-svm classifier (Classification SVM Type 1) with a linear kernel, which transformed input data to a feature space, giving the coefficient of a separating hyperplane. A sensitivity-based feature selector was used to remove non-informative data and reduce noise, using an ANOVA measure to select features with the highest F-scores. Then, the dataset was resampled into multiple instances for leave-one-out folding cross-validation, and when the classifier had been trained on one run of the data, it was then applied to the remaining data. Both single-subject and subject-averaged within-subject classifications were extracted.

2.12 Behavioral-classification analyses

Behavior-classification correlation analyses were carried out in SPSS 21 (IBM 176 Corporation, NY, US). Individual scatterplots of behavioral performance measures and relevant ROI classification accuracies were produced and used to evaluate linear and monotonic

relationships in the data. Since we could not assume nor confirm a linear relationship, but found monotonic relationships, Spearman's rank order correlation coefficient (rho) was used for correlation analyses between object positioning performance and classification performance from the postsimulus encoding, for all relevant ROIs, and to find correlations between granularity trial classification performance (coarse-vs. fine-grained) and positional granularity, positions-room granularity, objects-room test performance, objects test performance and object-position test performance, individually.

3 RESULTS

3.1 Voxel activity pattern dissimilarities

The dissimilarity in the stimulus and poststimulus activation pattern for anterior-posterior and intermediate ROIs in the hippocampus, entorhinal cortex, parahippocampal cortex, caudate, putamen and anterior cingulate cortex were investigated across different positional granularities, room geometries, and individual rooms. Dissimilarities are illustrated as raw *1 minus Pearson's r* representational dissimilarity matrices (RDMs) (figure 3-1.).

3.1.1 Arrangement of ROIs and subjects in representational dissimilarity matrices

All RDM figures (figures 3-2, 3-3, 3-5, 3-6, 7-1, 7-2, 7-3 and 7-4) are arranged in the same general manner: The RDMs for each subject are shown individually along a single row, while each column is a region of interest (ROI) in the brain. This arrangement is shown in figure 3-1, which depicts the first two rows (subjects) from the granularity RDMs in figure 3-3.

Everything in figure 3-1 is applicable to all instances of figures showing RDMs for all subjects across all ROIs, except for the specific RDM used as an example on the bottom. This is only applicable to the granularity RDMs in figure 3-2 and 3-3 and appendix figures 7-1 and 7-2.

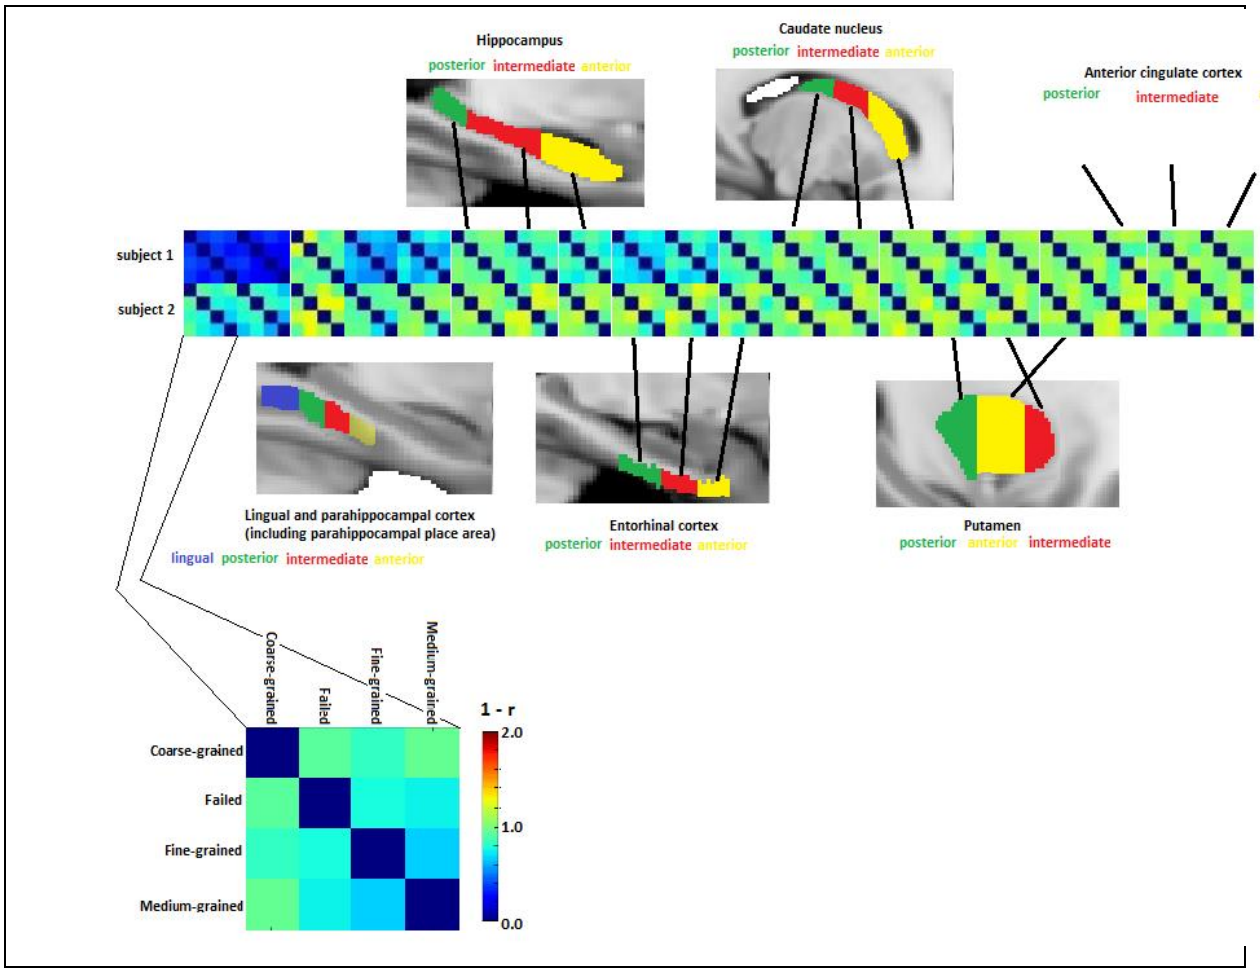

Figure 3-1 Presentation of RMDs for all participants The representational dissimilarity matrices are arranged according to subjects (one subject per row) and ROIs (one ROI per column). The ROIs are indicated in the figure and follow the same order across all RDM presentations. The RDM on the bottom is the granularity RDM. Importantly, any RDM is symmetrical across the diagonal (blue color, indicating perfect correlation). A dissimilarity color bar is shown next to the illustration RDM: the dissimilarity measure here is 1 minus Peason's r, with blue being perfect correlation (1-1 =0) and red perfect anticorrelation [1-(-1) =2].

3.1.2 Granularity-dependent voxel activity pattern dissimilarities

Dissimilarities in voxel activity patterns from the stimulus presentation period for fine-, medium-,coarse-grained, and failed positional encoding in each subject individually across all ROIs in the brain were computed (figure 3-2). The RDMs are arranged as shown in figure 3-1, including the RDM categories on the bottom.

Visual inspection indicates a striking trend, where the lingual ROI and posterior parahippocampal cortex (two leftmost columns) trend clearly toward a lower (bluer) dissimilarity. In other words, the underlying voxel activity pattern in these brain regions is more similar across the three different granularities and the failed trials (figure 3-2). This indicates a general response pattern in these areas that is highly similar across all trials in all or most subjects.

Figure 3-3 shows the dissimilarity in activity patterns in the subsequent poststimulus encoding period for fine-, medium-,coarse-grained, and failed positional encoding. Visually comparing the RDMs in figure 3-2 and 3-3 shows that the most obvious difference from the stimulus presentation period to the poststimulus encoding period is a trend in the two leftmost ROIs (lingual and posterior parahippocampal cortices) towards more dissimilarity in the encoding period. This suggests that the underlying voxel patterns in these brain areas during trials defined as fine-, medium-, course-grained, and failed, are more similar during the stimulus presentation period than during the poststimulus encoding period. Both for the stimulus presentation and poststimulus encoding period, there is very little apparent anticorrelation. This is verified by figures 6-1 and 6-2 in the appendix, which only shows the range between 1 (no correlation) and 2 (perfect anticorrelation).

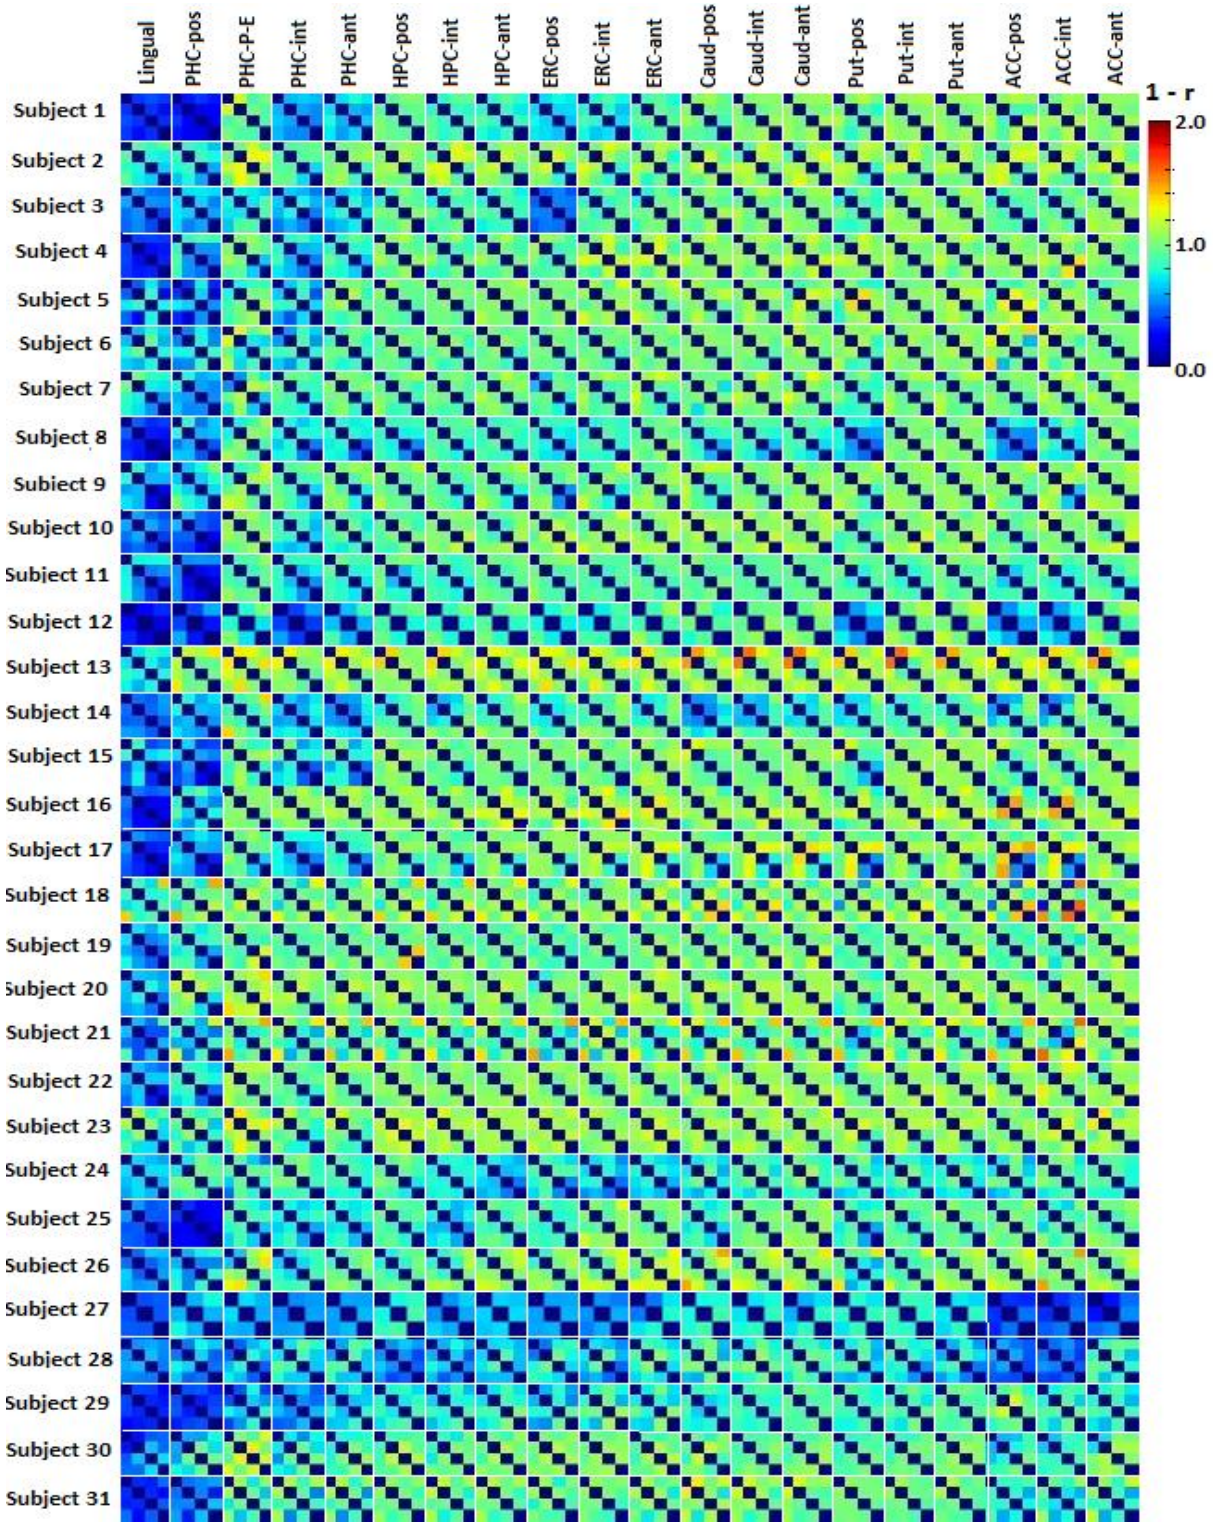

Figure 3-2 Activation pattern dissimilarity for positonal granularities in the stimulus presentation period The granularities in each matrix involved fine-grained. medium-grained, coarse-grained and failed (see also figure 3-1). The figure includes the entire span of dissimilarities, calculated as 1 – Pearson's r, so the range of possible values goes from 0 (perfect correlation) to 2 (perfect anticorrelation). Each row represents a single subject, while the columns represent regions of interest (ROIs). The color bar on the top right shows the color coding along the dissimilarity spectrum. ROI abbreviations: pos=posterior, int=intermediate, ant=anterior; ACC=anterior cingulate cortex, Put=putamen,

Caud=caudate, ERC=entorhinal cortex, HPC=hippocampus, PHC=parahippocampal cortex

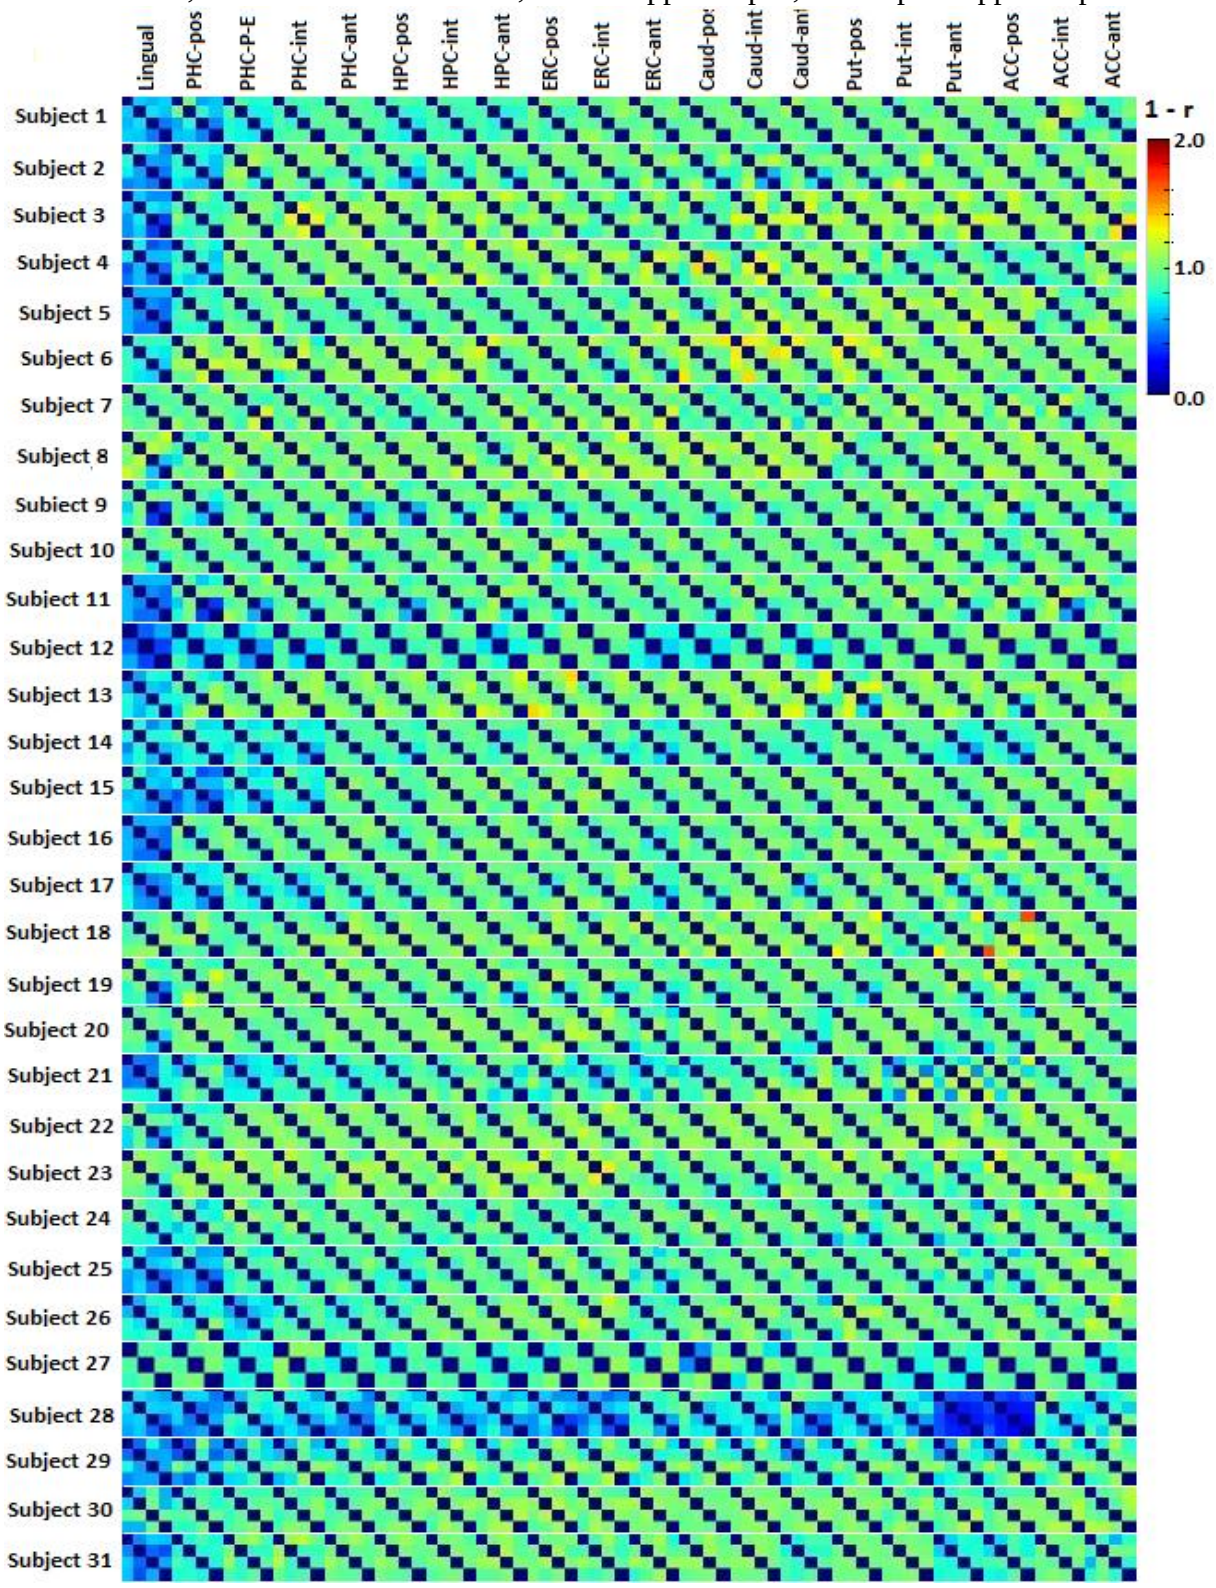

Figure 3-3 Activation pattern dissimilarity for positonal granularities in the poststimulus encoding period The granularities in each matrix involved fine-grained. medium-grained, coarse-grained and failed (see also figure 3-1 The figure includes the entire span of dissimilarities, calculated as 1 – Pearson's r, so the range of possible values goes from 0 (perfect correlation) to 2 (perfect anticorrelation). Each row represents a single subject, while the columns represent regions of interest (ROIs). The color bar on the top right shows the color coding along the dissimilarity spectrum. ROI abbreviations: pos=posterior, int=intermediate, ant=anterior; ACC=anterior cingulate cortex, Put=putamen,

Caud=caudate, ERC=entorhinal cortex, HPC=hippocampus, PHC=parahippocampal cortex

Voxel activity pattern dissimilarities associated with room geometry

Figure 3-4 shows the RDM category setup for the room geometries. This arrangement is kept across all subjects, in all ROIs.

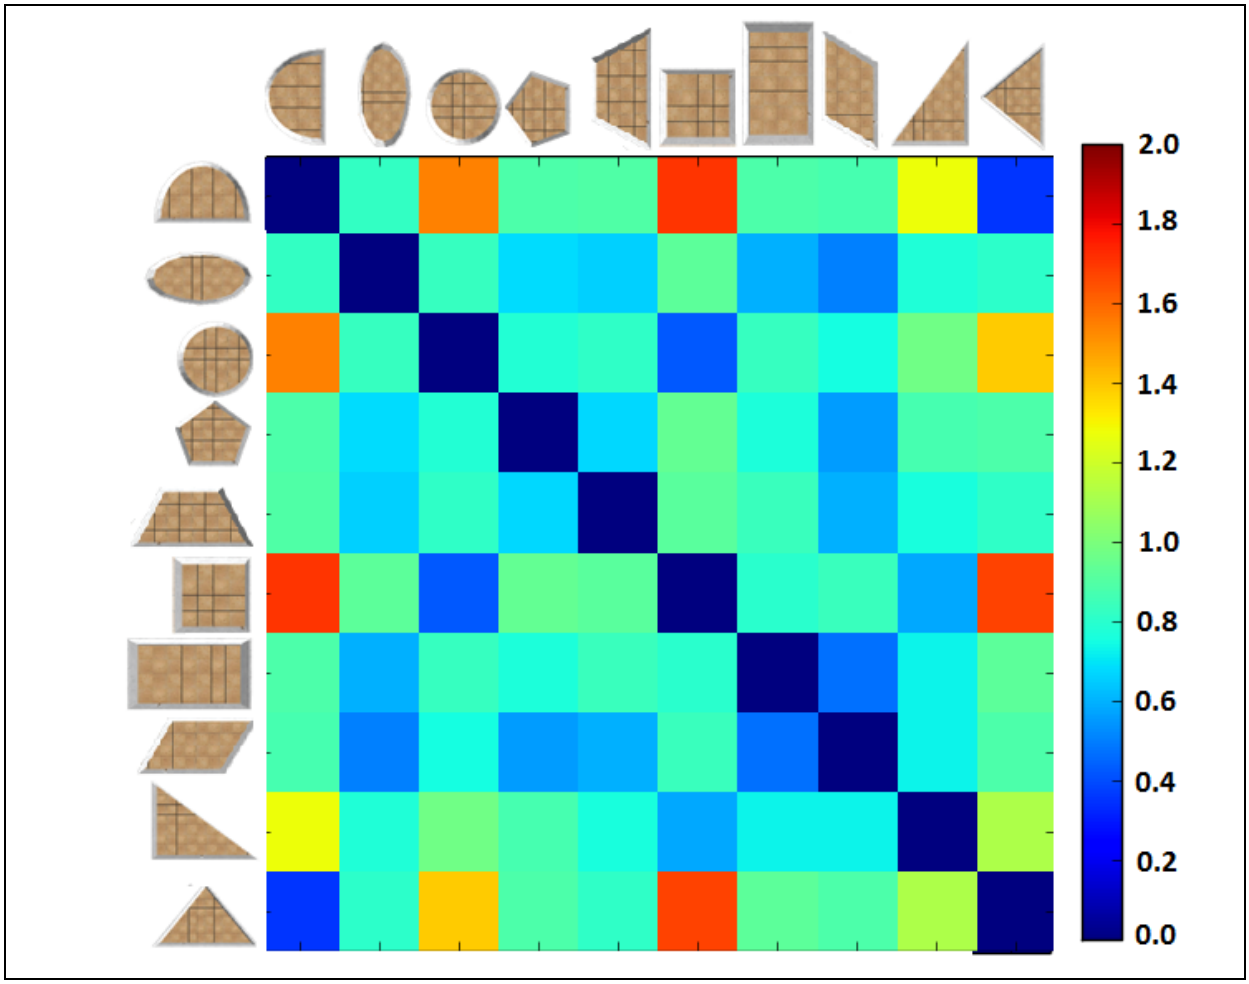

Figure 3-4 Single-RDM arrangement, objects-room geometry binding Importantly, all RDMs are symmetric across the diagonal (dark blue = perfect correlation). Each category/stimulus is a different outer border geometrical shape of the environment the subjects navigated through. The dissimilarity measure is 1-Pearson's r and indicates the dissimilarity between the voxel activity patterns associated with the different room shapes. The full range of dissimilarities (0 to 2) is shown and color coded according to the colorbar on the right.

Dissimilarities in voxel activity patterns from the stimulus presentation period for the ten different outer wall geometries of environments in each subject individually across all ROIs in the brain were computed (figure 3-1). The RDMs are arranged as shown in figure 3-4.

There seems to be a strong trend where the lingual ROI and posterior parahippocampal cortex also here have a lower dissimilarity, meaning that the underlying voxel activity pattern in these brain regions are more similar across the different outer geometries in all or most subjects (figure 3-5). Figure 3-6 shows the dissimilarity in activity patterns in the poststimulus

encoding period for the same outer wall geometries. Comparing the RDMs in figure 3-5 and 3-6, it seems that there is a trend that underlying voxel pattern activity in the lingual and posterior parahippocampal cortices are less similar during the poststimulus encoding period than during the stimulus presentation period. There does not seem to be much anticorrelation in any of the two periods, and this is shown more clearly by figures 7-3 and 7-4 in the appendix, which only shows the range between no correlation and perfect anticorrelation.

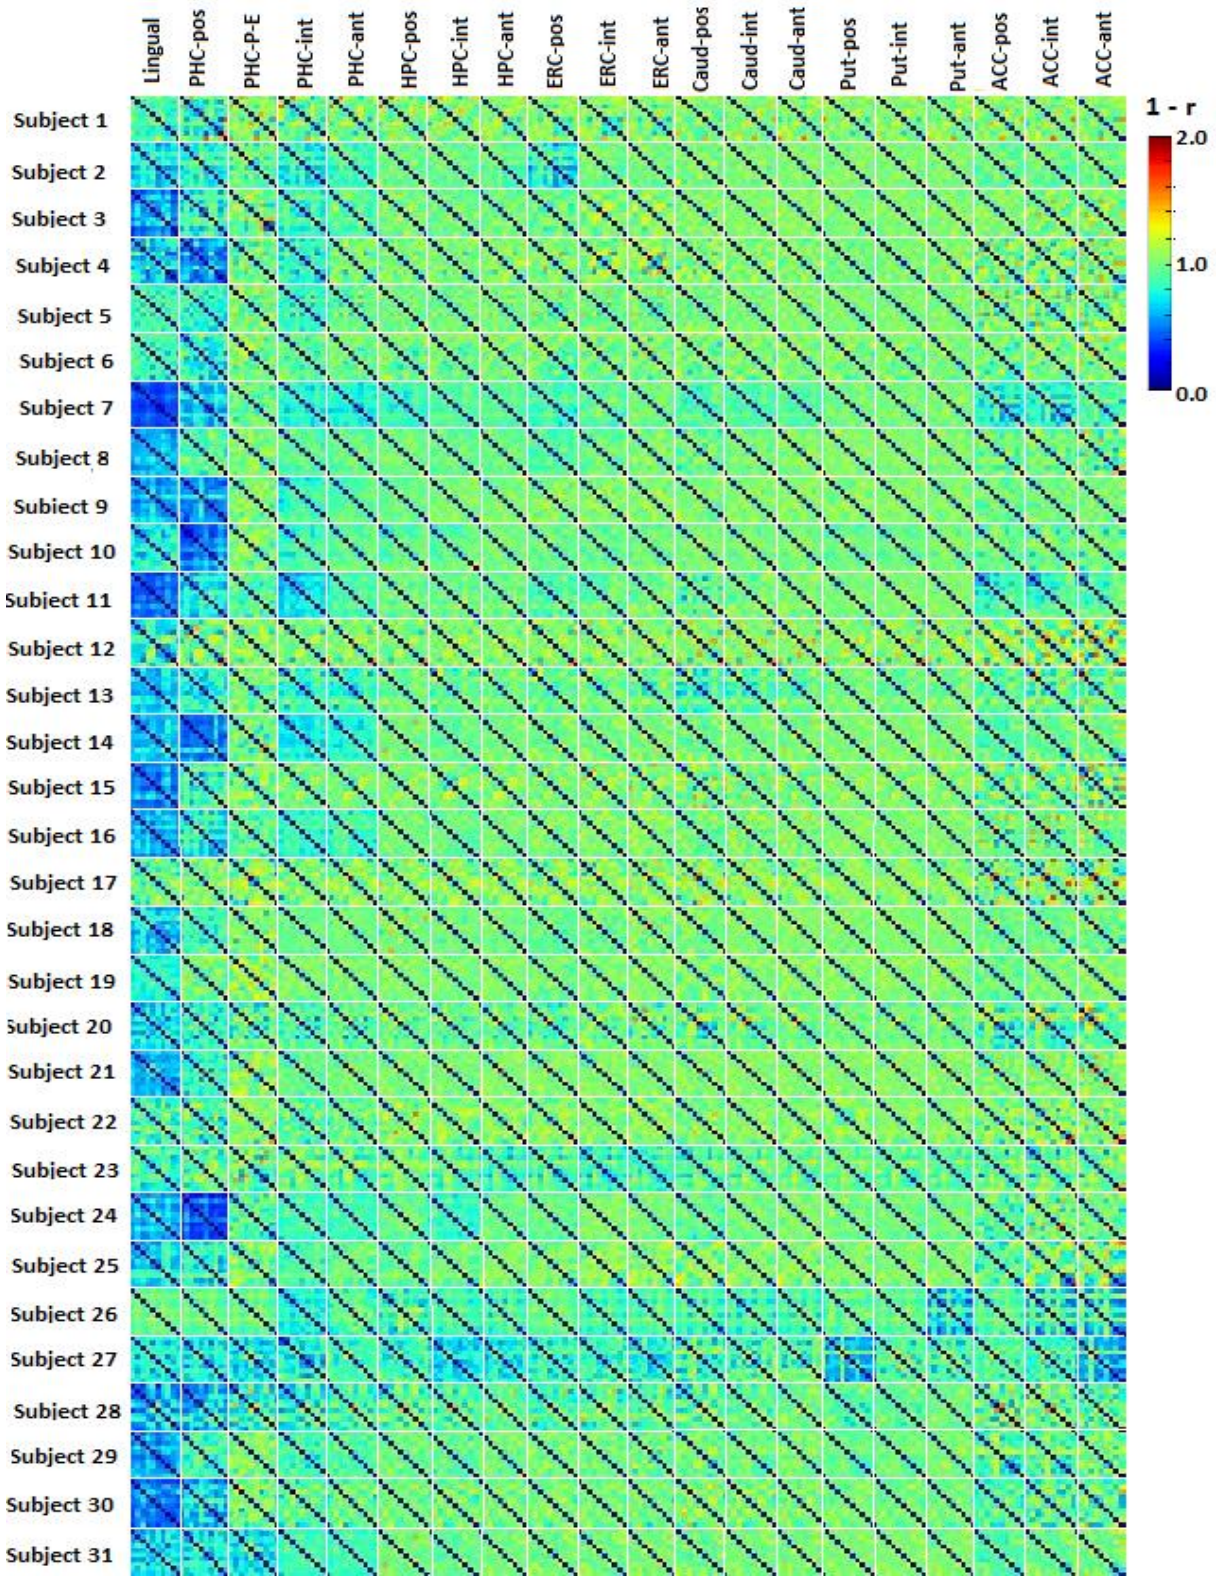

Figure 3-5 Activation pattern dissimilarity for outer border geometrical shapes in the stimulus presentation period The categories in each matrix involved ten different room shapes (see also figure 3-4). The figure includes the entire span of dissimilarities, calculated as 1 – Pearson's r, so the range of possible values goes from 0 (perfect correlation) to 2 (perfect anticorrelation). Each row represents a single subject, while the columns represent regions of interest (ROIs). The color bar on the top right shows the color coding along the dissimilarity spectrum. ROI abbreviations: pos=posterior, int=intermediate, ant=anterior; ACC=anterior cingulate cortex, Put=putamen, Caud=caudate, ERC=entorhinal cortex, HPC=hippocampus, PHC=parahippocampal cortex

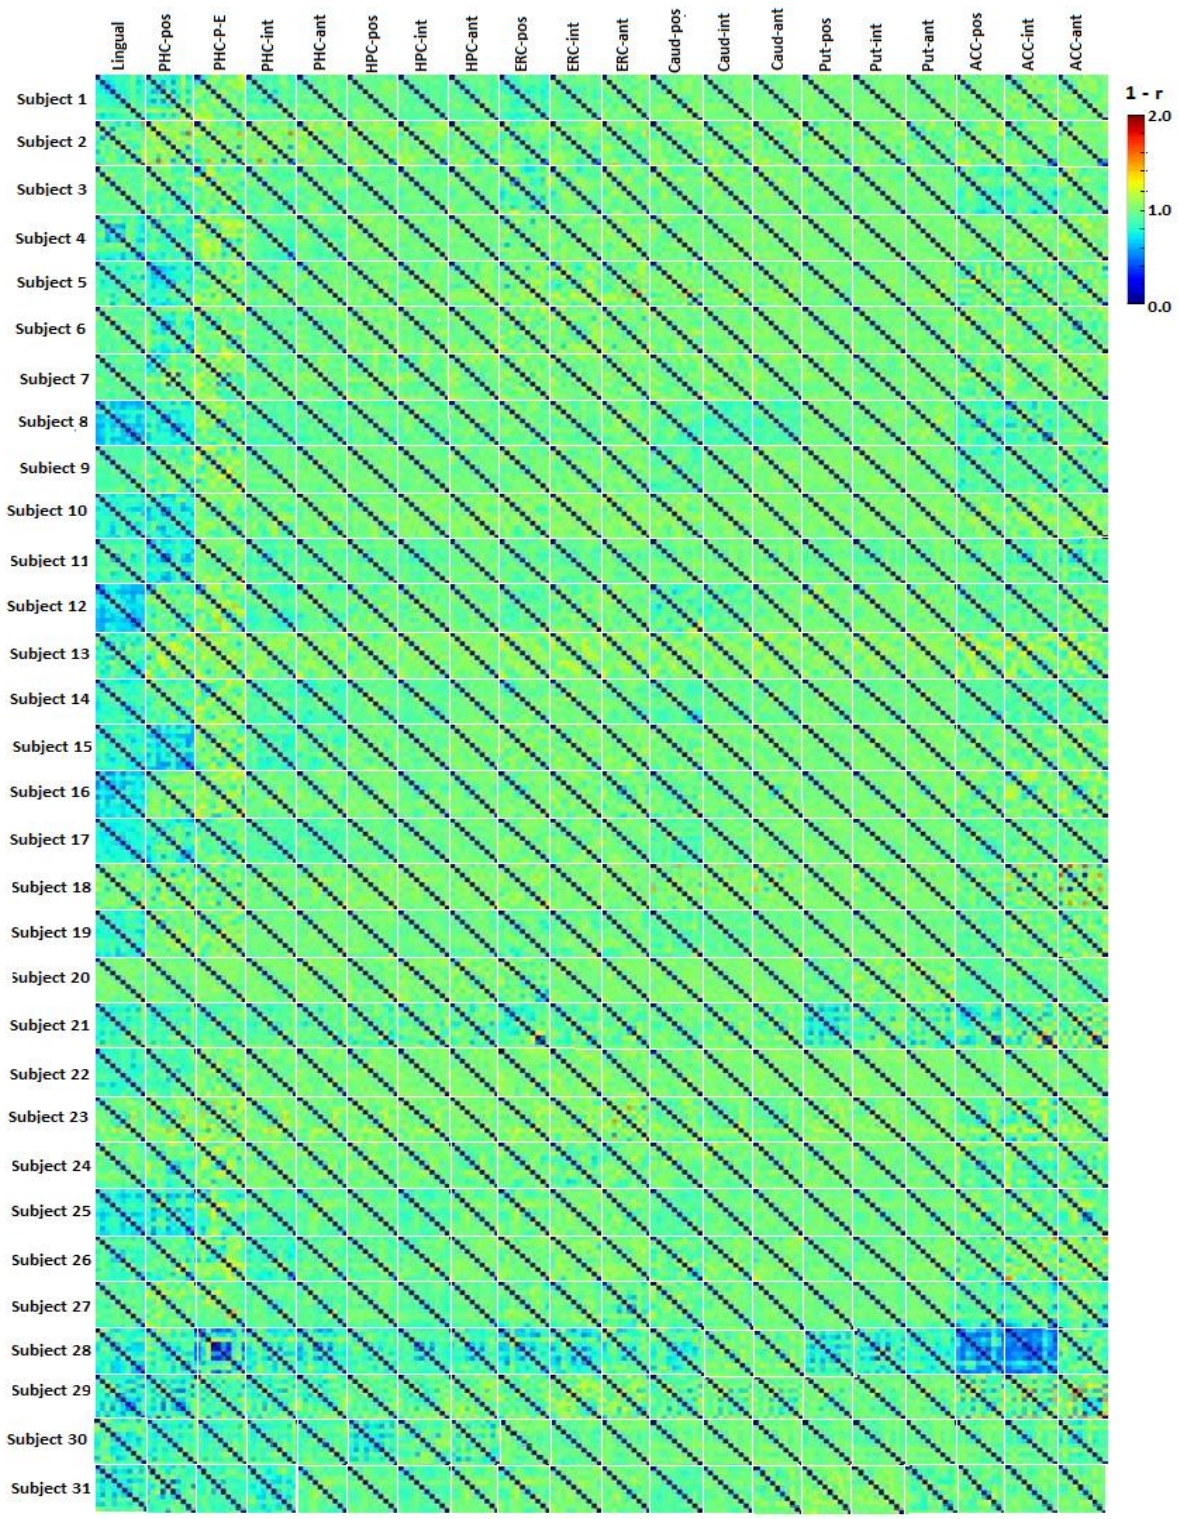

Figure 3-6 Activation pattern dissimilarity for outer border geometrical shapes in the poststimulus encoding period The categories in each matrix involved ten different room shapes (see also figure 3-4). The figure includes the entire span of dissimilarities, calculated as 1 – Pearson's r, so the range of possible values goes from 0 (perfect correlation) to 2 (perfect anticorrelation). Each row represents a single subject, while the columns represent regions of interest (ROIs). The color bar on the top right shows the color coding along the dissimilarity spectrum. ROI abbreviations: pos=posterior, int=intermediate, ant=anterior; ACC=anterior cingulate cortex, Put=putamen, Caud=caudate, ERC=entorhinal cortex, HPC=hippocampus, PHC=parahippocampal cortex

3.1.4 Voxel activity pattern dissimilarities associated with 35 different environments

Figure 3-7 shows the RDM category setup for the 35 different environments. This arrangement is kept across all subjects, in all ROIs.

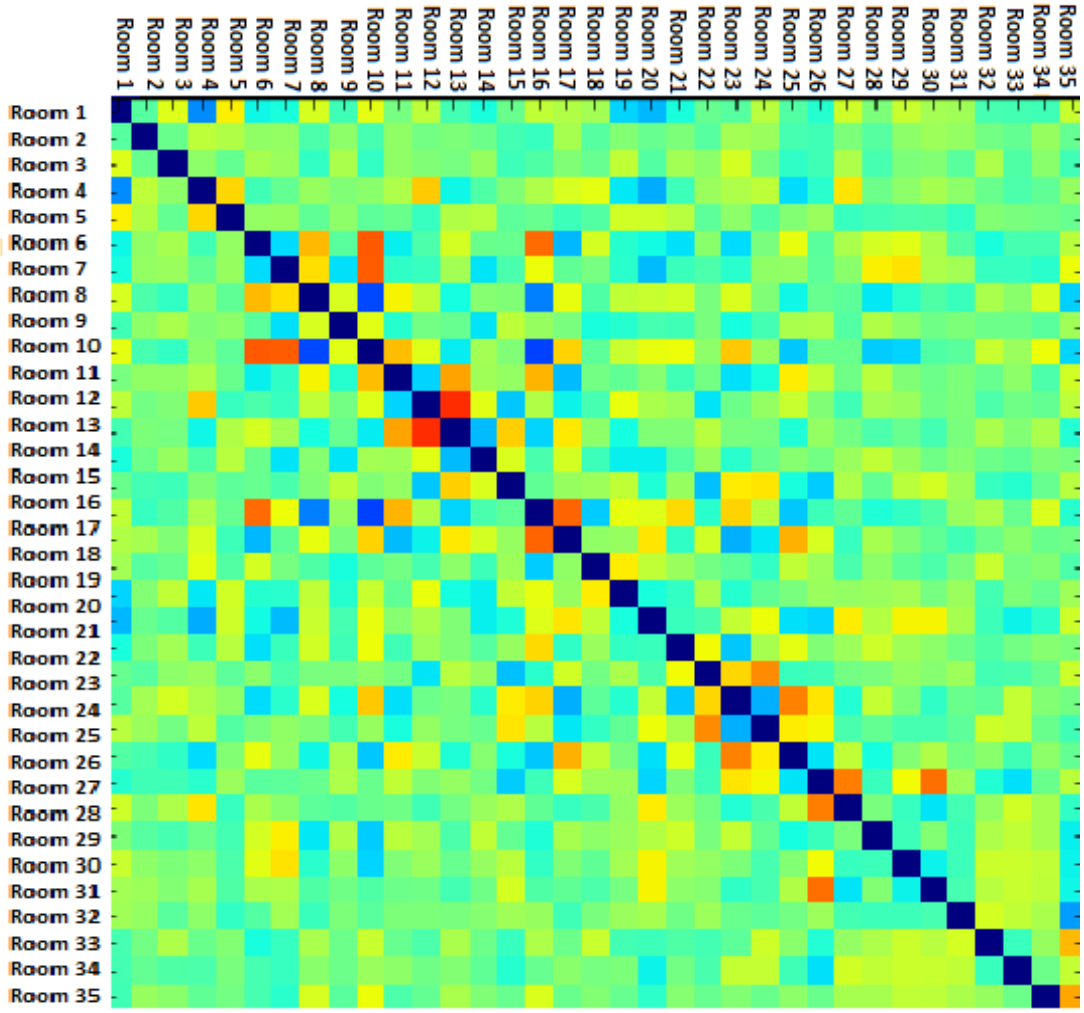

Figure 3-7 Single-RDM arrangement, 35 different environments Importantly, all RDMs are symmetric across the diagonal (blue = perfect correlation). Each category/stimulus is a different outer border geometrical shape of the environment the subjects navigated through. The dissimilarities indicate the 1-Pearson's r dissimilarity between the voxel activity patterns associated with the different room shapes. The full range of dissimilarities (0 to 2) is shown and color coded according to the colorbar on the right

This analysis showed that the activation pattern across all 35 individual environments seemed somewhat more similar in posterior parahippocampal and lingual ROIs during the stimulus presentation period, and that the similarity seemed to drop slightly from the stimulus to the poststimulus period (figures 3.8 and 3.9).

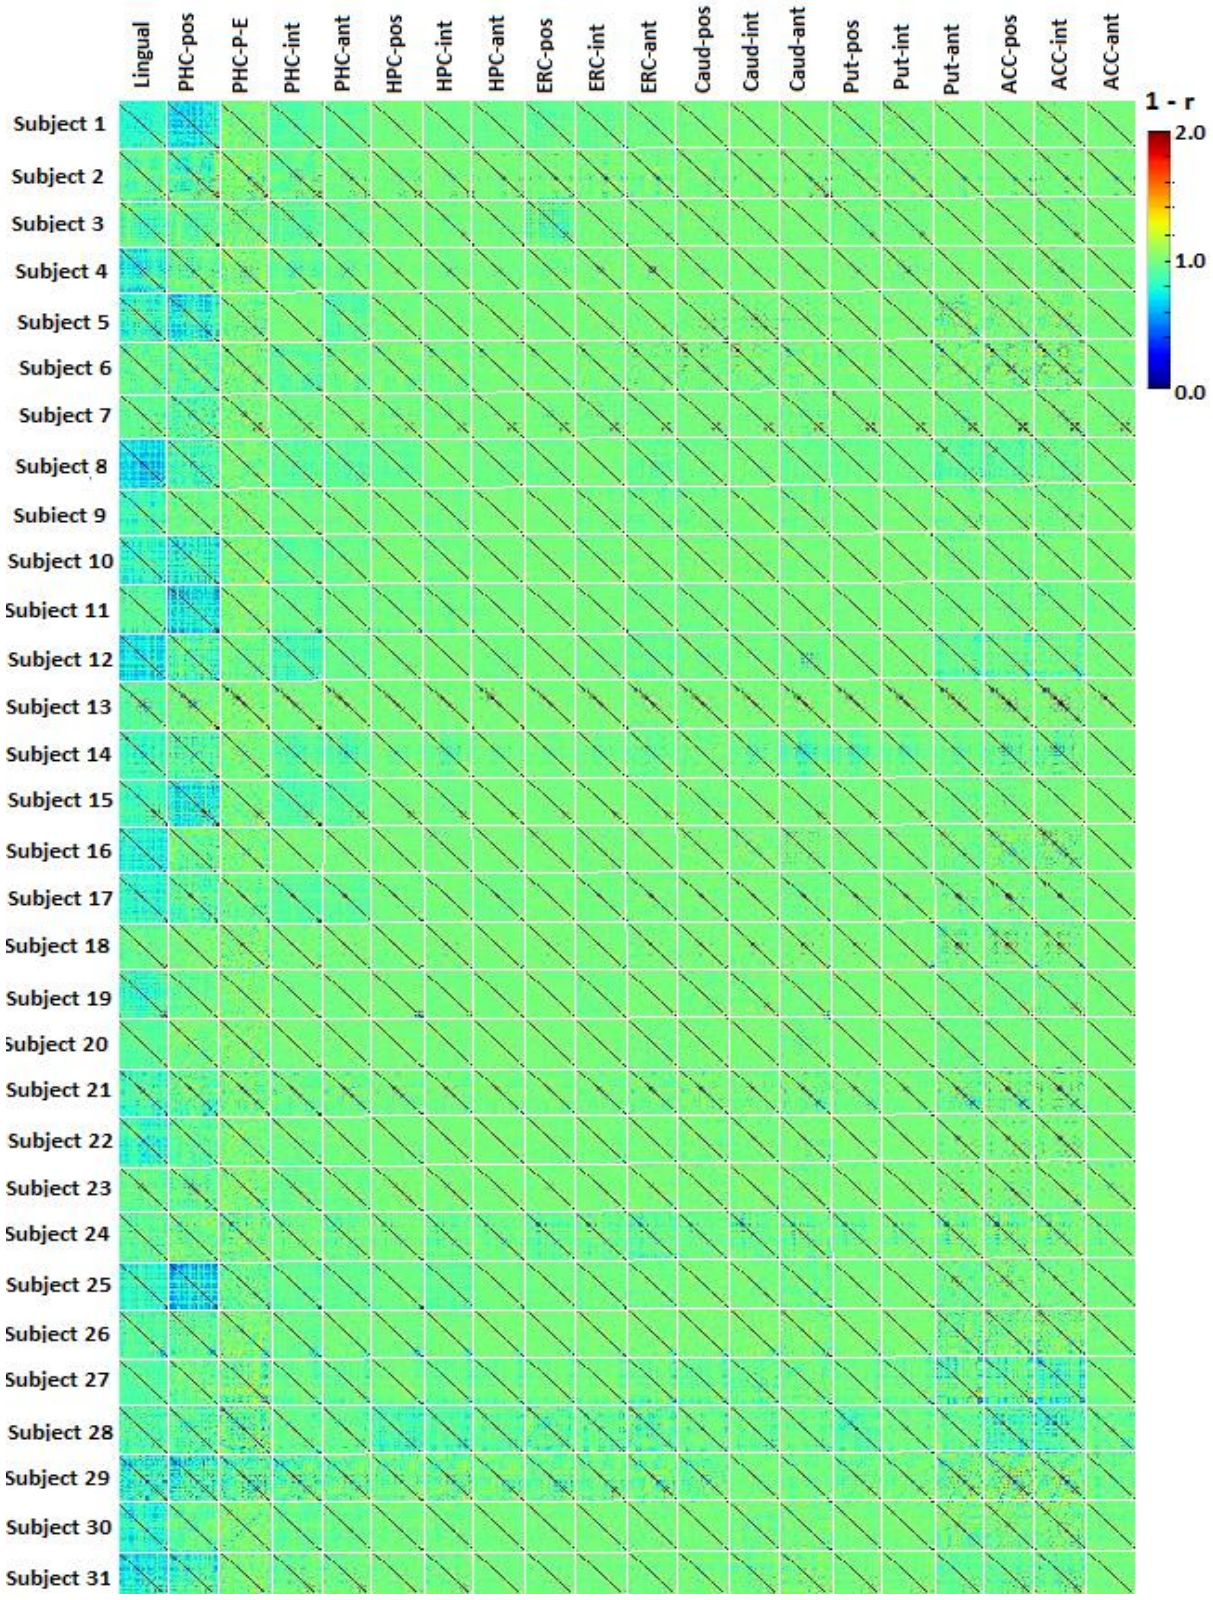

Figure 3-8 Activation pattern dissimilarity for 35 different environments in the stimulus presentation period The categories in each matrix involved 35 different environments (see also figure 3-7). The figure includes the entire span of dissimilarities, calculated as 1 – Pearson's r, so the range of possible values goes from 0 (perfect correlation) to 2 (perfect anticorrelation). Each row represents a single subject, while the columns represent regions of interest (ROIs). The color bar on the top right shows the color coding along the dissimilarity spectrum. ROI abbreviations: pos=posterior, int=intermediate, ant=anterior; ACC=anterior cingulate cortex, Put=putamen, Caud=caudate, ERC=entorhinal cortex, HPC=hippocampus, PHC=parahippocampal cortex

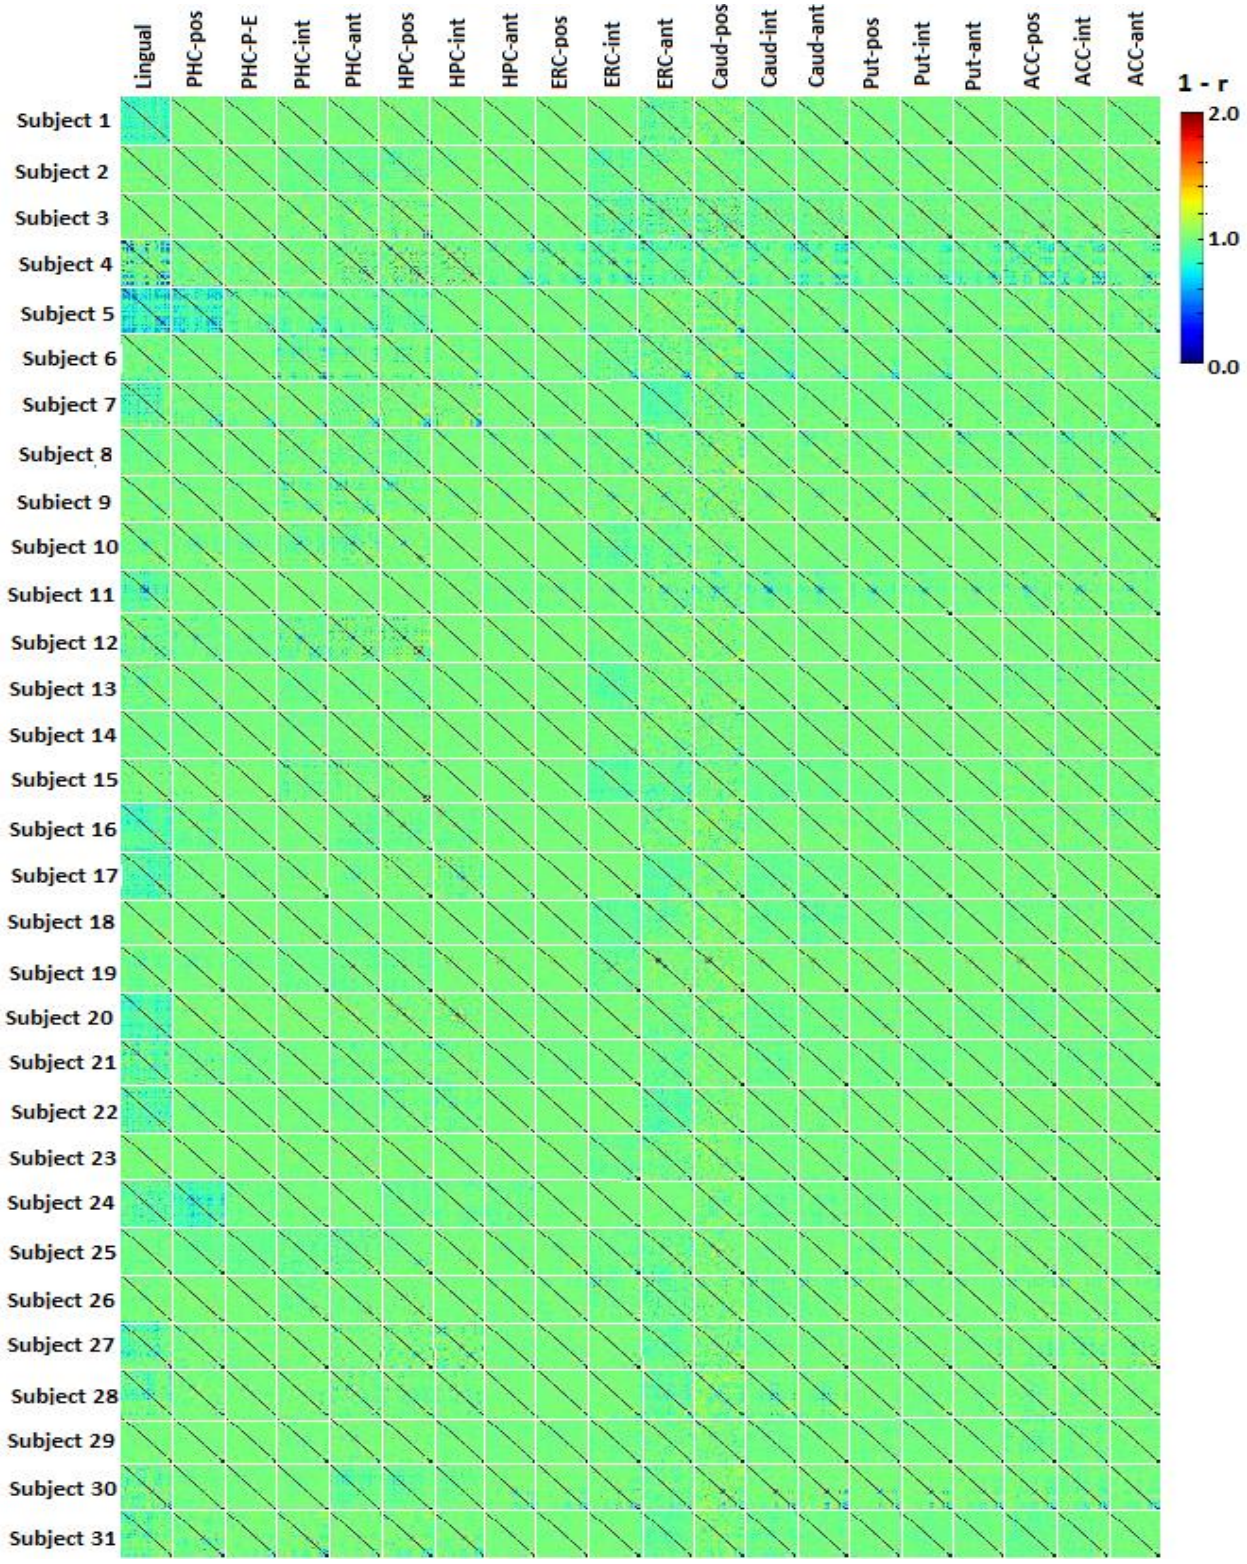

Figure 3-9 Activation pattern dissimilarity for 35 different environments in the poststimulus encoding period The categories in each matrix involved 35 different environments (see also figure 3-7). The figure includes the entire span of dissimilarities, calculated as 1 – Pearson's r, so the range of possible values goes from 0 (perfect correlation) to 2 (perfect anticorrelation). Each row represents a single subject, while the columns represent regions of interest (ROIs). The color bar on the top right shows the color coding along the dissimilarity spectrum. ROI abbreviations: pos=posterior, int=intermediate, ant=anterior; ACC=anterior cingulate cortex, Put=putamen, Caud=caudate, ERC=entorhinal cortex, HPC=hippocampus, PHC=parahippocampal cortex

3.2 Within-subject classification (linear support vector machine)

3.2.1 Main experimental periods

As shown in table 3-1, in all three experimental period comparisons (Stimulus presentation vs Poststimulus encoding, Stimulus presentation vs Odd-even and Poststimulus encoding vs Odd-even), data from the lingual cortex ROI had the highest average within-subject classification performance. For most ROIs, classification accuracy was similar for Poststimulus encoding vs. Stimulus presentation periods and Stimulus presentation vs. Odd-even periods, but there seemed to be a substantial increase in lingual and parahippocampal cortices and posterior hippocampus for Stimulus presentation vs. Odd-even, with classification accuracy increasing with .07 to .12. There seemed to be a large decrease in classification accuracy across all ROIs from Stimulus presentation vs Odd-even to Poststimulus encoding vs Odd-even.

For the fusiform ROI, lingual ROI, parahippocampal ROIs and posterior hippocampus, classification accuracy also seemed to decrease from Stimulus presentation vs Odd-even to Stimulus presentation vs Poststimulus encoding. The ventricular ROI, which would be expected to be functionally uncorrelated to any measured brain function, presupposing that it only contained only non-neural tissue, shows chance-level classification performance across all classification tasks.

Table 3-1 Within-subject classification results for main experimental periods

ROI	Subjects	Stimulus presentation vs. poststimulus encoding	Poststimulus encoding vs. odd-even	Stimulus presentation vs. odd-even
Lingual cortex	31	.76 +/- .01	.65 +/- .01	.88 +/- .01
Posterior parahippocampal cortex	31	.71 +/- .01	.60 +/- .01	.77 +/- .01
Intermed. parahippocampal cortex	31	.70 +/- .01	.63 +/- .01	.79 +/- .02
Anterior parahippocampal cortex	31	.67 +/- .01	.61 +/- .02	.74 +/- .01
Peri/ento/parahippocampal cortex	31	.64 +/- .01	.53 +/- .01	.66 +/- .01
Posterior hippocampus	31	.60 +/- .01	.57 +/- .01	.70 +/- .01
Intermediate hippocampus	31	.69 +/- .01	.59 +/- .02	.71 +/- .02
Anterior hippocampus	31	.67 +/- .01	.56 +/- .02	.68 +/- .01
Posterior entorhinal cortex	31	.65 +/- .01	.55 +/- .01	.65 +/- .01
Intermediate entorhinal cortex	31	.64 +/- .01	.54 +/- .01	.65 +/- .01
Anterior entorhinal cortex	31	.64 +/- .01	.56 +/- .02	.67 +/- .02
Posterior caudate	31	.65 +/- .01	.59 +/- .01	.63 +/- .01
Intermediate caudate	31	.65 +/- .01	.59 +/- .01	.66 +/- .01
Anterior caudate	31	.64 +/- .01	.57 +/- .01	.65 +/- .01
Posterior putamen	31	.64 +/- .01	.56 +/- .01	.62 +/- .01
Intermediate putamen	31	.63 +/- .01	.56 +/- .02	.64 +/- .01
Anterior putamen	31	.65 +/- .01	.56 +/- .01	.65 +/- .01
Posterior anterior cingulate cortex	31	.65 +/- .01	.55 +/- .01	.63 +/- .02
Intermediate ant. cingulate cortex	31	.63 +/- .01	.53 +/- .01	.68 +/- .01
Anterior anterior cingulate cortex	31	.64 +/- .01	.56 +/- .02	.66 +/- .01
Fusiform area	31	.76 +/- .01	.67 +/- .01	.84 +/- .01
Anterior horn of lateral ventricle	31	.58 +/- .02	.51 +/- .02	.54 +/- .02

The table shows average within-subjects classification performance by a linear support vector machine (SVM) on voxel pattern activation data in different regions of interest (ROIs) during the main experimental periods (Stimulus presentation, potstimulus encoding, and odd-even task period) for 31 subjects.

3.2.2 Positional granularity scores

The high classification performance in lingual and parahippocampal cortices for experimental periods continued with a relatively high classification performance for Coarse- versus Fine-grained positional representations of the environments, seemingly marginally better than other ROIs (Table 3-2). The fusiform ROI seems to show a large decrease in classification performance from experimental periods to Coarse-grained vs- Fine-grained

granularity representations, diverging from the other ROIs. The classification for all ROIs seemed to drop when the encoded positional granularity to separate between became more similar, i.e. from Coarse- versus Fine-grained classification to Coarse- versus Medium-grained or Medium- vs Fine-grained classification. .

Table 3-2 Within-subject classification results for degrees of granularity

ROI	Subjects	Coarse-grained vs. fine-grained	Coarse-grained vs. medium-grained	Medium-grained vs. fine-grained
Lingual cortex	31	.70 +/- .06	.61 +/- .04	.55 +/- .04
Posterior parahippocampal cortex	31	.70 +/- .06	.57+/- .04	.56 +/- .04
Intermed. parahippocampal cortex	31	.67 +/- .06	.55 +/- .03	.56 +/- .04
Anterior parahippocampal cortex	31	.68 +/- .06	.57 +/- .04	0.56 +/- .04
Peri/ento/parahippocampal cortex	31	.69 +/- .06	.55 +/- .04	.56 +/- .04
Posterior hippocampus	31	.66 +/- .06	.57 +/- .04	.55 +/- .04
Intermediate hippocampus	31	.68 +/- .05	.56 +/- .03	.54 +/- .03
Anterior hippocampus	31	.67 +/- .06	.54 +/- .04	.55 +/- .04
Posterior entorhinal cortex	31	.66 +/- .06	.54 +/- .04	.55 +/- .04
Intermediate entorhinal cortex	31	.68 +/- .06	.55 +/- .04	.56 +/- .04
Anterior entorhinal cortex	31	.64 +/- .06	.54 +/- .04	.55 +/- .04
Posterior caudate	31	.67 +/- .06	.53 +/- .04	.54 +/- .03
Intermediate caudate	31	.68 +/- .06	.56 +/- .03	.57 +/- .03
Anterior caudate	31	.65 +/- .06	.57 +/- .04	.56 +/- .04
Posterior putamen	31	.68 +/- .06	.54 +/- .03	.56 +/- .04
Intermediate putamen	31	.69 +/- .06	.53 +/- .03	.56 +/- .04
Anterior putamen	31	.66 +/- .06	.55 +/- .03	.55 +/- .03
Posterior anterior cingulate cortex	31	.64 +/- .06	.55 +/- .04	.55 +/- .04
Intermediate ant. cingulate cortex	31	.67 +/- .06	.54 +/- .04	.56 +/- .04
Anterior anterior cingulate cortex	31	.63 +/- .06	.56 +/- .04	.54 +/- .03
Fusiform area	31	.58 +/- .06	.56 +/- .03	.55 +/- .06
Anterior horn of lateral ventricle	31	.57 +/- .05	.54 +/- .05	.55 +/- .04

The table shows average within-subjects classification performance by a linear support vector machine (SVM) on voxel pattern activation data in different regions of interest (ROIs) from the poststimulus encoding period for coarse-grained, medium-grained and fine grained positional representations, for 31 subjects.

3.3 Correlations between performance and classification accuracies

3.3.1 Objects-room binding

Table 3-3 shows that object positioning performance was significantly correlated with lingual ROI classification performance only for both Stimulus presentation vs. Odd-even periods (rho = .515, p = .004) and Stimulus presentation vs. Poststimulus encoding periods (rho = .367, p = .046).

Table 3-3 Correlations between positional granularity and experimental period classification accuracy

ROI	Poststimulus encoding vs. odd-even	Stimulus presentation vs. odd-even	Stimulus presentation vs. poststimulus encoding
Lingual cortex	.075, p = .694	.515, p = .004 **	.367, p = .046*
Posterior parahippocampal cortex	-.032, p = .867	.195, p = .301	.198, p = .294
Intermed. parahippocampal cortex	.071, p = .707	.294, p = .114	.004, p = .981
Anterior parahippocampal cortex	-.216, p = .251	.031, p = .873	.049, p = .795
Peri-/ento-/parahippocampal cortex	-.107, p = .575	.334, p = .071	-.074, p = .698
Posterior hippocampus	.117, p = .538	.175, p = .352	.052, p = .787
Intermediate hippocampus	-.041, p = .828	-.039, p = .838	-.062, p = .746
Anterior hippocampus	-.181, p = .338	.113, p = .551	-.043, p = .822
Posterior entorhinal cortex	.034, p = .859	.103, p = .587	.006, p = .975
Intermediate entorhinal cortex	-.281, p = .133	.089, p = .860	.073, p = .702
Anterior entorhinal cortex	-.077, p = .685	.077, p = .786	-.046, p = .808
Anterior horn of lateral ventricle	-.055, p = .798	.054. p = .721	-.015, p = .854

The table shows the correlation between average within-subject classification performance by a linear support vector machine (SVM) for the voxel patterns in different regions of interest (ROIs) during the three experimental periods (stimulus presentation, poststimulus encoding, and odd-even task period) and object positioning performance (objects-room binding), for 31 subjects. The correlation test statistic is Spearman's rank order coefficient (rho). P-values are 2-sided, unadjusted. * = significant at p = .05, ** = significant at p = .01, *** = significant at p = .005

3.3.2 Positional granularity

Table 3-4 shows that lingual ROI classification performance for voxel activity patterns during Coarse- vs. Fine-grained representations was significantly negatively correlated with both Positional granularity (-0.359, p = .033) and Objects-room test performance (-.350, p = .045). Importantly, a lower positional granularity score indicated better performance. No other correlations were significant, but it's worth noting that Positional granularity, Positions-room

granularity, Objects-room test performance and Objects test performance were all negatively correlated, and the Object-position test score positively correlated, with granularity classification accuracy across all ROIs

Table 3-4 Correlation between performance measures and granularity classification accuracy

ROIs	Positional granularity	Postions-room granularity	Object-position test	Objects-room test	Objects test
Lingual cortex	-.359, p = .033*	-.321, p = .084	.328, p = .150	-.350, p = .045*	-.314, p = .091
Posterior parahippocampal cortex	-.288, p = .123	-.261, p = .164	.268, p = .153	-.333, p = .072	-.336, p = .069
Intermed. parahippocampal cortex	.287, p = .125	-.260, p = .165	.219, p = .246	-.284, p = .128	-.297, p = .135
Anterior parahippocampal cortex	-.293, p = .117	-.267, p = .153	.240, p = .202	-.300, p = .107	-.320, p = .085
Peri/ento/ parahippocampal cortex	-.291, p = .120	-.259, p = .167	.224, p = .234	-.328, p = .083	-.303, p = .104
Posterior hippocampus	-.323, p = .081	-.305, p = .101	.238, p = .205	-.322, p = .109	-.312, p = .093
Intermediate hippocampus	-.289, p = .122	-.253, p = .177	.230, p = .221	-.317, p = .135	-.288, p = .123
Anterior hippocampus	-.263, p = .161	-.200, p = .265	.246, p = .189	-.298, p = .210	-.281, p = .132
Posterior entorhinal cortex	-.276, p = .141	-.258, p = .187	.272, p = .146	-.235, p = .088	-.254, p = .176
Intermediate entorhinal cortex	-.282, p = .131	-.265, p = .223	.254, p = .148	-.317, p = .073	-.297, p = .111
Anterior entorhinal cortex	-.308, p = .098	-.200, p = .147	.232, p = .212	-.326, p = .125	-.316, p = .088
Ventricle	-.211, p = .135	-.155, p = .198	.155, p = .305	-.122, p = .197	-.210, p = .155

The table shows the correlation between average within-subject classification performance by a linear support vector machine (SVM) for the voxel patterns in different regions of interest (ROIs) for coarse-grained vs. fine-grained representations and different performance measures, for 31 subjects. Performance measures are: Positional granularity, positions-room granularity, object-position test, objects-room test, and objects test. The correlation test statistic is Spearman's rank order coefficient (rho). P-values are 2-sided, unadjusted. * = significant at p = .05

4 DISCUSSION

We found that voxel activation patterns only in the posterior parahippocampal cortex and lingual ROI, both being within the borders of the parahippocampal place area, were similar across all virtual environments and levels of granularity encoded, especially during the stimulus presentation period. Further, the lingual ROI was the only region where activation pattern classification correlated significantly with performance measures (i.e. better object positioning, more accurate positional granularity and worse objects-room performance). Classifications based on voxel activation patterns were similar along all ROIs in the anterior-posterior axis, although they often tended to be higher in posterior parahippocampal and lingual ROIs.

4.1 Multivariate encoding along the anterior-posterior axis of the parahippocampal cortex and hippocampus

No functional segregation for activation patterns was observed along the anterior-posterior axis of the hippocampus. All ROIs along the anterior-posterior axis of the hippocampus showed similar within-subject classification accuracy for all three granularity contrasts. Further, the same dissimilarity in activation pattern was observed along the anterior-posterior hippocampal axis using the correlation-dissimilarity measure. This indicates that no functional segregation exists along the anterior-posterior axis of the hippocampus. Previous studies have found that voxel activation patterns contained different information along this axis for sense of self-location and later positional recall (Guterstam et al., 2015), object and scene similarity and object-location encoding (Libby et al., 2014), and episodic memory (Collin et al., 2015). However, the dataset our results are based on is the first, to our knowledge, which looks specifically at encoding of positional representation while controlling for related aspects of environmental processing, and it simply may not be the case that this function specifically is represented in the neural patterns along the hippocampal anterior-posterior axis.

This result also contradicts the univariate analyses of the data, where a granularity gradient was found along the hippocampal anterior-posterior axis, with fine-, medium-, and coarse-grained representations associated with the posterior, intermediate and posterior hippocampus, respectively (Evensmoen et al., 2015). Paradigms that work well for studying and comparing (massively) univariate activation magnitudes do not necessarily transfer as well to studies of the differences between across-voxel patterns, like representational similarity analysis (Mumford, 2013) and linear classification (Etzel et al., 2013). The literature also

contains examples where univariate and multivariate results came to different conclusions regarding representational gradients in this area (Copara et al., 2014). It is, however, not certain how these analyses can be best combined to form a more complete picture of the underlying brain functions, or how and when univariate and multivariate results would be expected to give the same results for positional granularity. Even though a multivoxel study indicated that a gradient within our memories exists, from the association between a limited number of events posteriorly to more complete associational representations in the anterior hippocampus (Collin et al., 2015), same functional region might not represent eg. granularities of episodic memories in the same way it represents positional granularity. Additionally, differences in the multiple choices made for distance measures between studies can influence the results (Walther et al., 2015). Since knowledge about what exactly drives the gradient response in any brain area is not fully understood, it is uncertain what the difference between multivariate and univariate results along these axes means, although the current results support the notion that the hippocampus is primarily driven by the total response amplitude found by the univariate analysis (Evensmoen et al., 2015) and not the pattern of activity investigated here.

The posterior parahippocampal cortex and lingual ROI, the parahippocampal place area, had a unique and stimulus driven activation pattern that was similar throughout and especially important for encoding of local environmental details. The correlation-dissimilarity measure showed a high similarity in the voxel activation pattern in the parahippocampal place area across all conditions and memory categories investigated, i.e. granularity encoded, individual environments, different room geometry, and experimental conditions. Another study found that the parahippocampal cortex showed increased activational pattern similarity across four virtual environments (Kyle et al., 2015). The more similar patterns of activity observed in the parahippocampal place area compared to those seen in other ROIs could reflect a difference in visual involvement, since these areas are functionally involved in visual processing and in humans contain retinotopically mapped information (Orban et al., 2014). Functional connectivity analyses also show that the parahippocampal place area has increased connectivity to visual cortices (Baldassano et al., 2013). Similarity analyses have previously shown that other visual stream areas store category and exemplar information through more similar voxel patterns (Mur et al., 2013). Taken together, this could indicate that in our study the parahippocampal place area processes all exemplars within our environmental geometries, our 35 environments, and the different positional granularities, as exemplars within the same

overarching visuospatial categories. However, this hypothesis will have to be tested for specifically and cannot be answered by our results.

The classification accuracy seemed to be higher for the activation patterns in the posterior parahippocampal cortex and lingual ROI for the separation of the Stimulus presentation periods from both Poststimulus encoding and Odd-even. Increased activation in this part of the brain can be obseserved for spatial scenes (Epstein et al., 1999). In a previous study it was observed increased activation in the posterior parahippocampal cortex throughout the navigation period, argued to be related to more local processing of the environment (Xu et al., 2010). Similar to this, others have argued that the parahippocampal area has a more fundamental representation of local space (Mullally and Maguire, 2011).

Supporting this, increased classification accuracy only for the Stimulus presentation period, vs Odd-even and Poststimulus encoding periods, and not Poststimulus encoding period vs Odd-even, led to less overall binding between the objects and the room and more accurate encoding of the objects positional pattern in the lingual ROI, between subjects. This further indicates that more posterior parts of the extended parahippocampal cortex is especially important for accurate positional encoding. Supporting this, the classification accuracy seemed to be higher for the posterior parahippocampal cortex and lingual ROI than the anterior parahippocampal cortex for Coarse- vs Fine-grained positional encoding. Further, increased classification accuracy for Coarse- vs Fine-grained positional encoding led to more accurate encoding of the objects' positional pattern only in the lingual ROI, across subjects. For an object-grid task, the posterior and not anterior parahippocampal cortex was found to be correlated with recall of exact object binding (Sommer et al., 2005). In a recent review, the posterior part of the parahippocampal cortex was suggested to be more important for spatial associations and the anterior part for non-spatial associations (Aminoff et al., 2013). Taken together, the posterior parahippocampal cortex and lingual ROI, also defined as the parahippocampal place area, seems to be especially important for more fine-grained spatial representations more closely linked to the stimuli.

4.2 Statistical and design considerations regarding the observed results and future directions

The Pearson correlation used as a measure of activation pattern similarity in this study is problematic, especially for the parahippocampal place area, or the posterior parahippocampal

cortex and lingual ROI. The pearson correlation will be influenced by baseline shift and the normalization conducted in this study (Walther et al., 2015). This could explain why we see increased pattern similarity in this regions, especially for the stimulus period, for which we know from a pilot study that the univariate activation in the parahippocampal place area is high even across memory categories. Still, for the parahippocampal place area, increased activation pattern similarity was also seen for the poststimulus period for which no such increase in univariate activation is present. In future analyses, we will implement Euclidean distance as a measure of activation pattern similarity as well, which is not dependent on a reliable baseline and not influenced by mean-pattern substraction (Walther et al., 2015).

Dissimilarity measures that can handle non-linearities are needed. Comparing the raw dissimilarity values between brains and models comes with an assumption that the model can linearly predict these dissimilarities, but we know that nonlinearities in the measurement can render this assumption void (Kriegeskorte et al., 2008; Nili et al., 2014). For this reason, we will use rank-correlations**,** which can handle these non-linearities**,** when comparing activation pattern similarities in future analyses(Kriegeskorte et al., 2008). Still, with rank-correlation, important activation similarity information might be lost (Kriegeskorte and Kievit, 2013).

Noise influenced the representational similarity activation pattern results in this thesis, as we did only a limited amount of noise correction. The only noise correction that was done was detrending, removing temporal drifts. Additionally, the data were motion corrected. In our data we know that for example the anterior parts of the medial temporal lobe, in close proximity to the ventricles, will have a noisier BOLD signal due to physiological susceptibility artifacts (Olman et al., 2009). Distance estimates without a proper noise-correction tend to be positively biased (Diedrichsen et al., 2011; Nili et al., 2014). This means that the activation patterns from especially the anterior medial temporal lobe in this study most likely is less dissimilar than they appeared from our results. In fact, this could explain the dissimilarity observed for most of the ROIs investigated. Still, for the parahippocampal place area, removing noise from our data would most likely make the observed activation pattern similarities even larger. For future analyses, we will use a crossvalidated estimate of the multivariately noise-normalized Mahalanobis distance (Nili et al., 2014; Walther et al., 2015). This unbiased and more reliable estimate will not only reduce the increased dissimiliarty due to noise, but also give a meaningful zero point through crossvalidation, which will allow us to compare two activation patterns for significant differences.

The multivariate pattern analysis will also be performed independent of the predefined anatomical ROIs used here, through searchlight analysis. Searchlight analysis measures the informational content (distributed voxel activation patterns) in a sphere of voxels surrounding a center voxel. As the sphere moves around the brain, using each passing voxel as its center (like a searchlight; thereby the name), it gives each voxel a value corresponding to the voxel activity pattern in the searchlight sphere surrounding it (Kriegeskorte et al. 2006). Scripts for a multivariate searchlight analysis were constructed using a modified version of the one provided in the pyMVPA toolbox framework (Hanke et al., 2009). However, we decided to postpone this effort and undertake future searchlight efforts within the RSA toolbox (Nili et al., 2014). In order to increase our understanding of how the environments we exist in are represented in our patterns of brain activation it is important to look at the brain more as a whole, through searchlight analysis, and not be limited to a few predefined anatomical ROIs.

Since every dataset comes with its own idiosyncracies, we will test future intended analyses on one or more simulated data sets with known ground truths (Nili et al., 2014; Walther et al., 2015). The simulated data could include the challenges that have been outlined here (regional noise variation, between-subject variations, etc.) and an underlying pattern information, and perhaps a set with a null ground truth to look for remaining biases. This could provide vital information about the intended analysis' ability to accurately detect the real patterns of interest in the data.

There are some design issues to consider in the context of the similarity analyses of the different environments, most importantly with the environments not being distributed across multiple runs (each environment was used once per subject). Similarly, it should also be noted that this could also be a problem for granularity dissimilarities: The conditions were defined by subject performance, and there is no control for across-run distribution, and with few failed trials it is probable that these were not distributed across runs.

5 Conclusions

Our study yielded results that differed from those found in the univariate analyses of the same dataset, but we have raised important methodological and statistical challenges that need to be addressed going forward. We found seemingly more similar patterns of voxel activation within posterior regions of the medial temporal lobe, specifically posterior parahippocampal cortex and lingual ROI, both part of the parahippocampal place area. However, it is unknown how this may relate to differences in noise effects due to different noise profiles between the anterior and posterior region of the medial temporal lobe and limitations in current methods. Future studies and analyses are needed.

6 REFERENCES

Alme, C.B., Miao, C., Jezek, K., Treves, A., Moser, E.I., and Moser, M.B. (2014). Place cells in the hippocampus: eleven maps for eleven rooms. Proc Natl Acad Sci U S A 111, 18428-18435.

Aminoff, E.M., Kveraga, K., and Bar, M. (2013). The role of the parahippocampal cortex in cognition. Trends Cogn Sci 17, 379-390.

Andersen, P., Bliss, T.V., and Skrede, K.K. (1971). Lamellar organization of hippocampal pathways. Exp Brain Res 13, 222-238.

Baldassano, C., Beck, D.M., and Fei-Fei, L. (2013). Differential connectivity within the Parahippocampal Place Area. Neuroimage 75, 228-237.

Barry, C., Hayman, R., Burgess, N., and Jeffery, K.J. (2007). Experience-dependent rescaling of entorhinal grids. Nat Neurosci 10, 682-684.

Blaabjerg, M., and Zimmer, J. (2007). The dentate mossy fibers: structural organization, development and plasticity. Prog Brain Res 163, 85-107.

Chadwick, M.J., Bonnici, H.M., and Maguire, E.A. (2012). Decoding information in the human hippocampus: a user's guide. Neuropsychologia 50, 3107-3121.

Chadwick, M.J., Hassabis, D., Weiskopf, N., and Maguire, E.A. (2010). Decoding individual episodic memory traces in the human hippocampus. Curr Biol 20, 544-547.

Chadwick, M.J., Jolly, A.E., Amos, D.P., Hassabis, D., and Spiers, H.J. (2015). A goal direction signal in the human entorhinal/subicular region. Curr Biol 25, 87-92.

Colgin, L.L., Moser, E.I., and Moser, M.B. (2008). Understanding memory through hippocampal remapping. Trends Neurosci 31, 469-477.

Collin, S.H., Milivojevic, B., and Doeller, C.F. (2015). Memory hierarchies map onto the hippocampal long axis in humans. Nat Neurosci 18, 1562-1564.

Copara, M.S., Hassan, A.S., Kyle, C.T., Libby, L.A., Ranganath, C., and Ekstrom, A.D. (2014). Complementary roles of human hippocampal subregions during retrieval of spatiotemporal context. J Neurosci 34, 6834-6842.

Diedrichsen, J., Ridgway, G.R., Friston, K.J., and Wiestler, T. (2011). Comparing the similarity and spatial structure of neural representations: a pattern-component model. Neuroimage 55, 1665-1678.

Ekstrom, A.D., Arnold, A.E., and Iaria, G. (2014). A critical review of the allocentric spatial representation and its neural underpinnings: toward a network-based perspective. Front Hum Neurosci 8, 803.

Epstein, R., Deyoe, E.A., Press, D.Z., Rosen, A.C., and Kanwisher, N. (2001). Neuropsychological evidence for a topographical learning mechanism in parahippocampal cortex. Cogn Neuropsychol 18, 481-508.

Epstein, R., Graham, K.S., and Downing, P.E. (2003). Viewpoint-specific scene representations in human parahippocampal cortex. Neuron 37, 865-876.

Epstein, R., Harris, A., Stanley, D., and Kanwisher, N. (1999). The parahippocampal place area: recognition, navigation, or encoding? Neuron 23, 115-125.

Epstein, R., and Kanwisher, N. (1998). The parahippocampal place area: A cortical representation of the local visual environment. J Cogn Neurosci 10, 20-20.

Epstein, R.A. (2008). Parahippocampal and retrosplenial contributions to human spatial navigation. Trends Cogn Sci 12, 388-396.

Etzel, J.A., Zacks, J.M., and Braver, T.S. (2013). Searchlight analysis: promise, pitfalls, and potential. Neuroimage 78, 261-269.

Evensmoen, H.R., Ladstein, J., Hansen, T.I., Moller, J.A., Witter, M.P., Nadel, L., and Haberg, A.K. (2015). From details to large scale: the representation of environmental positions follows a granularity gradient along the human hippocampal and entorhinal anterior-posterior axis. Hippocampus 25, 119-135.

Evensmoen, H.R., Lehn, H., Xu, J., Witter, M.P., Nadel, L., and Haberg, A.K. (2013). The anterior hippocampus supports a coarse, global environmental representation and the posterior hippocampus supports fine-grained, local environmental representations. J Cogn Neurosci 25, 1908-1925.

Franko, E., Insausti, A.M., Artacho-Perula, E., Insausti, R., and Chavoix, C. (2014). Identification of the human medial temporal lobe regions on magnetic resonance images. Hum Brain Mapp 35, 248-256.

Guterstam, A., Bjornsdotter, M., Gentile, G., and Ehrsson, H.H. (2015). Posterior cingulate cortex integrates the senses of self-location and body ownership. Curr Biol 25, 1416-1425.

Hafting, T., Fyhn, M., Molden, S., Moser, M.B., and Moser, E.I. (2005). Microstructure of a spatial map in the entorhinal cortex. Nature 436, 801-806.

Hanke, M., Halchenko, Y.O., Sederberg, P.B., Hanson, S.J., Haxby, J.V., and Pollmann, S. (2009). PyMVPA: A python toolbox for multivariate pattern analysis of fMRI data. Neuroinformatics 7, 37-53.

Hansen, T.I., Haferstrom, E.C., Brunner, J.F., Lehn, H., and Haberg, A.K. (2015). Initial validation of a web-based self-administered neuropsychological test battery for older adults and seniors. J Clin Exp Neuropsychol 37, 581-594.

Hassabis, D., Chu, C., Rees, G., Weiskopf, N., Molyneux, P.D., and Maguire, E.A. (2009). Decoding neuronal ensembles in the human hippocampus. Curr Biol 19, 546-554.

Haxby, J.V., Connolly, A.C., and Guntupalli, J.S. (2014). Decoding neural representational spaces using multivariate pattern analysis. Annu Rev Neurosci 37, 435-456.

Haxby, J.V., Guntupalli, J.S., Connolly, A.C., Halchenko, Y.O., Conroy, B.R., Gobbini, M.I., Hanke, M., and Ramadge, P.J. (2011). A common, high-dimensional model of the representational space in human ventral temporal cortex. Neuron 72, 404-416.

Haynes, J.D., and Rees, G. (2006). Decoding mental states from brain activity in humans. Nat Rev Neurosci 7, 523-534.

Hopfinger, J.B., Buchel, C., Holmes, A.P., and Friston, K.J. (2000). A study of analysis parameters that influence the sensitivity of event-related fMRI analyses. Neuroimage 11, 326-333.

Insausti, R., Juottonen, K., Soininen, H., Insausti, A., Partanen, K., Vainio, P., Laakso, M., and Pitkanen, A. (1998). MR volumetric analysis of the human entorhinal, perirhinal, and temporopolar cortices. AJNR Am J Neuroradiol 19, 659-671.

Jeffery, K.J. (2011). Place cells, grid cells, attractors, and remapping. Neural Plast 2011, 182602.

Jimura, K., and Poldrack, R.A. (2012). Analyses of regional-average activation and multivoxel pattern information tell complementary stories. Neuropsychologia 50, 544-552.

Jung, M.W., Wiener, S.I., and Mcnaughton, B.L. (1994). Comparison of Spatial Firing Characteristics of Units in Dorsal and Ventral Hippocampus of the Rat. Journal of Neuroscience 14, 7347-7356.

Kim, S.G., and Ogawa, S. (2012). Biophysical and physiological origins of blood oxygenation level-dependent fMRI signals. J Cereb Blood Flow Metab 32, 1188-1206.

Kjelstrup, K.B., Solstad, T., Brun, V.H., Hafting, T., Leutgeb, S., Witter, M.P., Moser, E.I., and Moser, M.-B. (2008). Finite Scale of Spatial Representation in the Hippocampus. Science 321, 140-143.

Kriegeskorte, N., and Kievit, R.A. (2013). Representational geometry: integrating cognition, computation, and the brain. Trends Cogn Sci 17, 401-412.

Kriegeskorte, N., Mur, M., and Bandettini, P. (2008). Representational similarity analysis - connecting the branches of systems neuroscience. Front Syst Neurosci 2, 4.

Kyle, C.T., Stokes, J.D., Lieberman, J.S., Hassan, A.S., and Ekstrom, A.D. (2015). Successful retrieval of competing spatial environments in humans involves hippocampal pattern separation mechanisms. Elife 4.

Lavenex, P., and Amaral, D.G. (2000). Hippocampal-neocortical interaction: a hierarchy of associativity. Hippocampus 10, 420-430.

Libby, L.A., Hannula, D.E., and Ranganath, C. (2014). Medial temporal lobe coding of item and spatial information during relational binding in working memory. J Neurosci 34, 14233-14242.

Lorente De No, R. (1934). Studies on the structure of the cerebral cortex. II. Continuation of the study of the ammonic system. Journal fur Psychologie und Neurologie, 113-177.

McTighe, S.M., Mar, A.C., Romberg, C., Bussey, T.J., and Saksida, L.M. (2009). A new touchscreen test of pattern separation: effect of hippocampal lesions. Neuroreport 20, 881-885.

Megevand, P., Groppe, D.M., Goldfinger, M.S., Hwang, S.T., Kingsley, P.B., Davidesco, I., and Mehta, A.D. (2014). Seeing scenes: topographic visual hallucinations evoked by direct electrical stimulation of the parahippocampal place area. J Neurosci 34, 5399-5405.

Mullally, S.L., and Maguire, E.A. (2011). A new role for the parahippocampal cortex in representing space. J Neurosci 31, 7441-7449.

Mumford, J.A. (2013). Considerations When Using Single-Trial Parameter Estimates in Representational Similarity Analyses.

Mur, M., Bandettini, P.A., and Kriegeskorte, N. (2009). Revealing representational content with pattern-information fMRI--an introductory guide. Soc Cogn Affect Neurosci 4, 101-109.

Mur, M., Meys, M., Bodurka, J., Goebel, R., Bandettini, P.A., and Kriegeskorte, N. (2013). Human Object-Similarity Judgments Reflect and Transcend the Primate-IT Object Representation. Front Psychol 4, 128.

Nili, H., Wingfield, C., Walther, A., Su, L., Marslen-Wilson, W., and Kriegeskorte, N. (2014). A toolbox for representational similarity analysis. PLoS Comput Biol 10, e1003553.

O'Keefe, J., and 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.M., and Nadel, L. (1978). The hippocampus as a cognitive map (London: Oxford University Press).

Ogawa, S., and Lee, T.M. (1990). Magnetic resonance imaging of blood vessels at high fields: in vivo and in vitro measurements and image simulation. Magn Reson Med 16, 9-18.

Ogawa, S., Lee, T.M., Nayak, A.S., and Glynn, P. (1990). Oxygenation-sensitive contrast in magnetic resonance image of rodent brain at high magnetic fields. Magn Reson Med 14, 68-78.

Olman, C.A., Davachi, L., and Inati, S. (2009). Distortion and Signal Loss in Medial Temporal Lobe. PLoS ONE 4, e8160.

Orban, G.A., Zhu, Q., and Vanduffel, W. (2014). The transition in the ventral stream from feature to real-world entity representations. Front Psychol 5, 695.

Poppenk, J., Evensmoen, H.R., Moscovitch, M., and Nadel, L. (2013). Long-axis specialization of the human hippocampus. Trends Cogn Sci 17, 230-240.

Pruessner, J.C., Köhler, S., Crane, J., Pruessner, M., Lord, C., Byrne, A., Kabani, N., Collins, D.L., and Evans, A.C. (2002). Volumetry of Temporopolar, Perirhinal, Entorhinal and Parahippocampal Cortex from High-resolution MR Images: Considering the Variability of the Collateral Sulcus. Cerebral Cortex 12, 1342-1353.

Raizada, R.D., Tsao, F.M., Liu, H.M., and Kuhl, P.K. (2010). Quantifying the adequacy of neural representations for a cross-language phonetic discrimination task: prediction of individual differences. Cereb Cortex 20, 1-12.

Rajimehr, R., Devaney, K.J., Bilenko, N.Y., Young, J.C., and Tootell, R.B. (2011). The "parahippocampal place area" responds preferentially to high spatial frequencies in humans and monkeys. PLoS Biol 9, e1000608.

Ranganath, C. (2010). A unified framework for the functional organization of the medial temporal lobes and the phenomenology of episodic memory. Hippocampus 20, 1263-1290.

Rodriguez, P.F. (2010). Neural decoding of goal locations in spatial navigation in humans with fMRI. Hum Brain Mapp 31, 391-397.

Scoville, W.B., and Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. J Neurol Neurosurg Psychiatry 20, 11-21.

Sommer, T., Rose, M., Weiller, C., and Buchel, C. (2005). Contributions of occipital, parietal and parahippocampal cortex to encoding of object-location associations. Neuropsychologia 43, 732-743.

Spiers, H.J., and Maguire, E.A. (2006). Thoughts, behaviour, and brain dynamics during navigation in the real world. Neuroimage 31, 1826-1840.

Spiers, H.J., and Maguire, E.A. (2008). The dynamic nature of cognition during wayfinding. J Environ Psychol 28, 232-249.

Squire, L.R., and Alvarez, P. (1995). Retrograde amnesia and memory consolidation: a neurobiological perspective. Curr Opin Neurobiol 5, 169-177.

Squire, L.R., Stark, C.E., and Clark, R.E. (2004). The medial temporal lobe. Annu Rev Neurosci 27, 279-306.

Squire, L.R., and Zola, S.M. (1996). Structure and function of declarative and nondeclarative memory systems. Proc Natl Acad Sci U S A 93, 13515-13522.

Stokes, J., Kyle, C., and Ekstrom, A.D. (2015). Complementary roles of human hippocampal subfields in differentiation and integration of spatial context. J Cogn Neurosci 27, 546-559.

Strange, B.A., Witter, M.P., Lein, E.S., and Moser, E.I. (2014). Functional organization of the hippocampal longitudinal axis. Nat Rev Neurosci 15, 655-669.

Suzuki, W.A., and Amaral, D.G. (1994). Topographic organization of the reciprocal connections between the monkey entorhinal cortex and the perirhinal and parahippocampal cortices. J Neurosci 14, 1856-1877.

Tolman, E.C. (1948). Cognitive maps in rats and men. Psychol Rev 55, 189-208.

Walther, A., Nili, H., Ejaz, N., Alink, A., Kriegeskorte, N., and Diedrichsen, J. (2015). Reliability of dissimilarity measures for multi-voxel pattern analysis. Neuroimage.

Witter, M.P., Van Hoesen, G.W., and Amaral, D.G. (1989). Topographical organization of the entorhinal projection to the dentate gyrus of the monkey. J Neurosci 9, 216-228.

Woollett, K., and Maguire, E.A. (2011). Acquiring "the Knowledge" of London's layout drives structural brain changes. Curr Biol 21, 2109-2114.

Xu, J., Evensmoen, H.R., Lehn, H., Pintzka, C.W., and Haberg, A.K. (2010). Persistent posterior and transient anterior medial temporal lobe activity during navigation. Neuroimage 52, 1654-1666.

7 APPENDICES

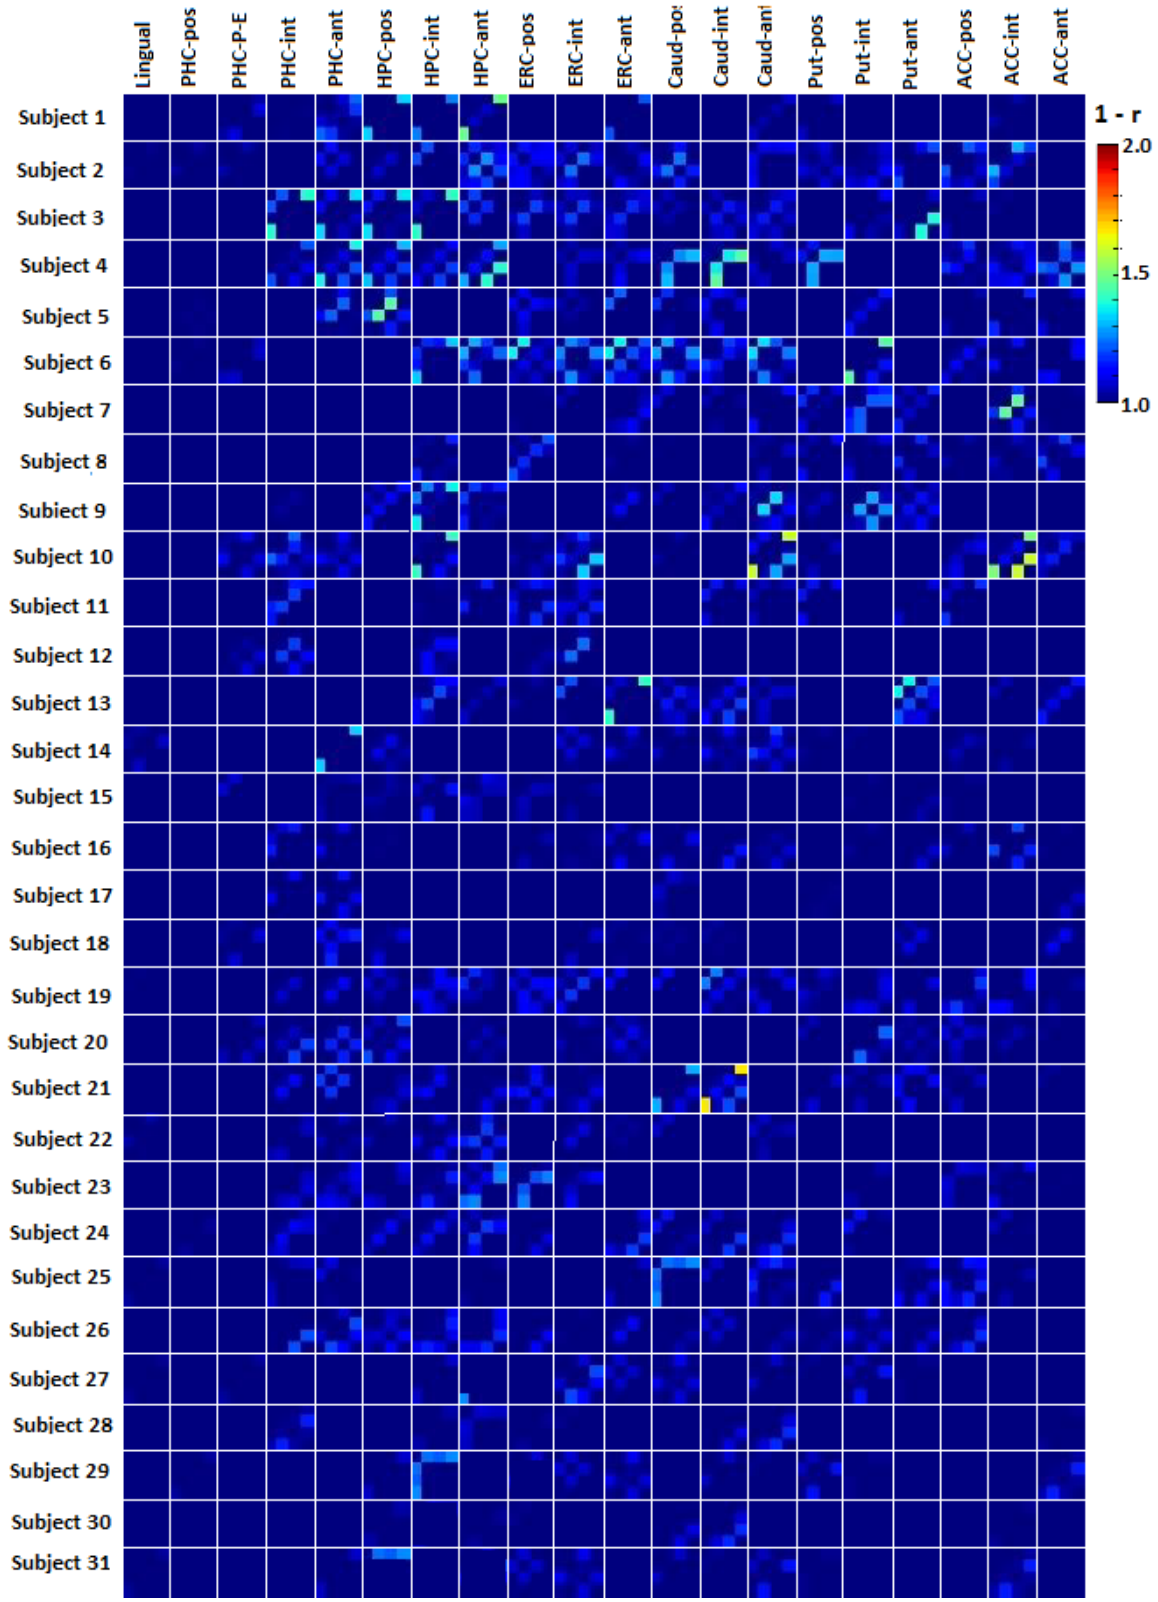

Figure 7-1 Activation pattern dissimilarity for positonal granularities in the stimulus presentation period The granularities in each matrix involved fine-grained. medium-grained, coarse-grained and failed (see also figure 3-1). The figure includes only part of the entire span of dissimilarities, calculated as 1 – Pearson's r, so the range of possible values goes from 1 (zero correlation) to 2 (perfect anticorrelation). Each row represents a single subject, while the columns represent regions of interest (ROIs). The color bar on the top right shows the color coding along the dissimilarity spectrum. ROI abbreviations: pos=posterior, int=intermediate, ant=anterior; ACC=anterior cingulate cortex, Put=putamen, Caud=caudate, ERC=enorhinal cortex, HPC=hippocampus, PHC=parahippocampal cortex.

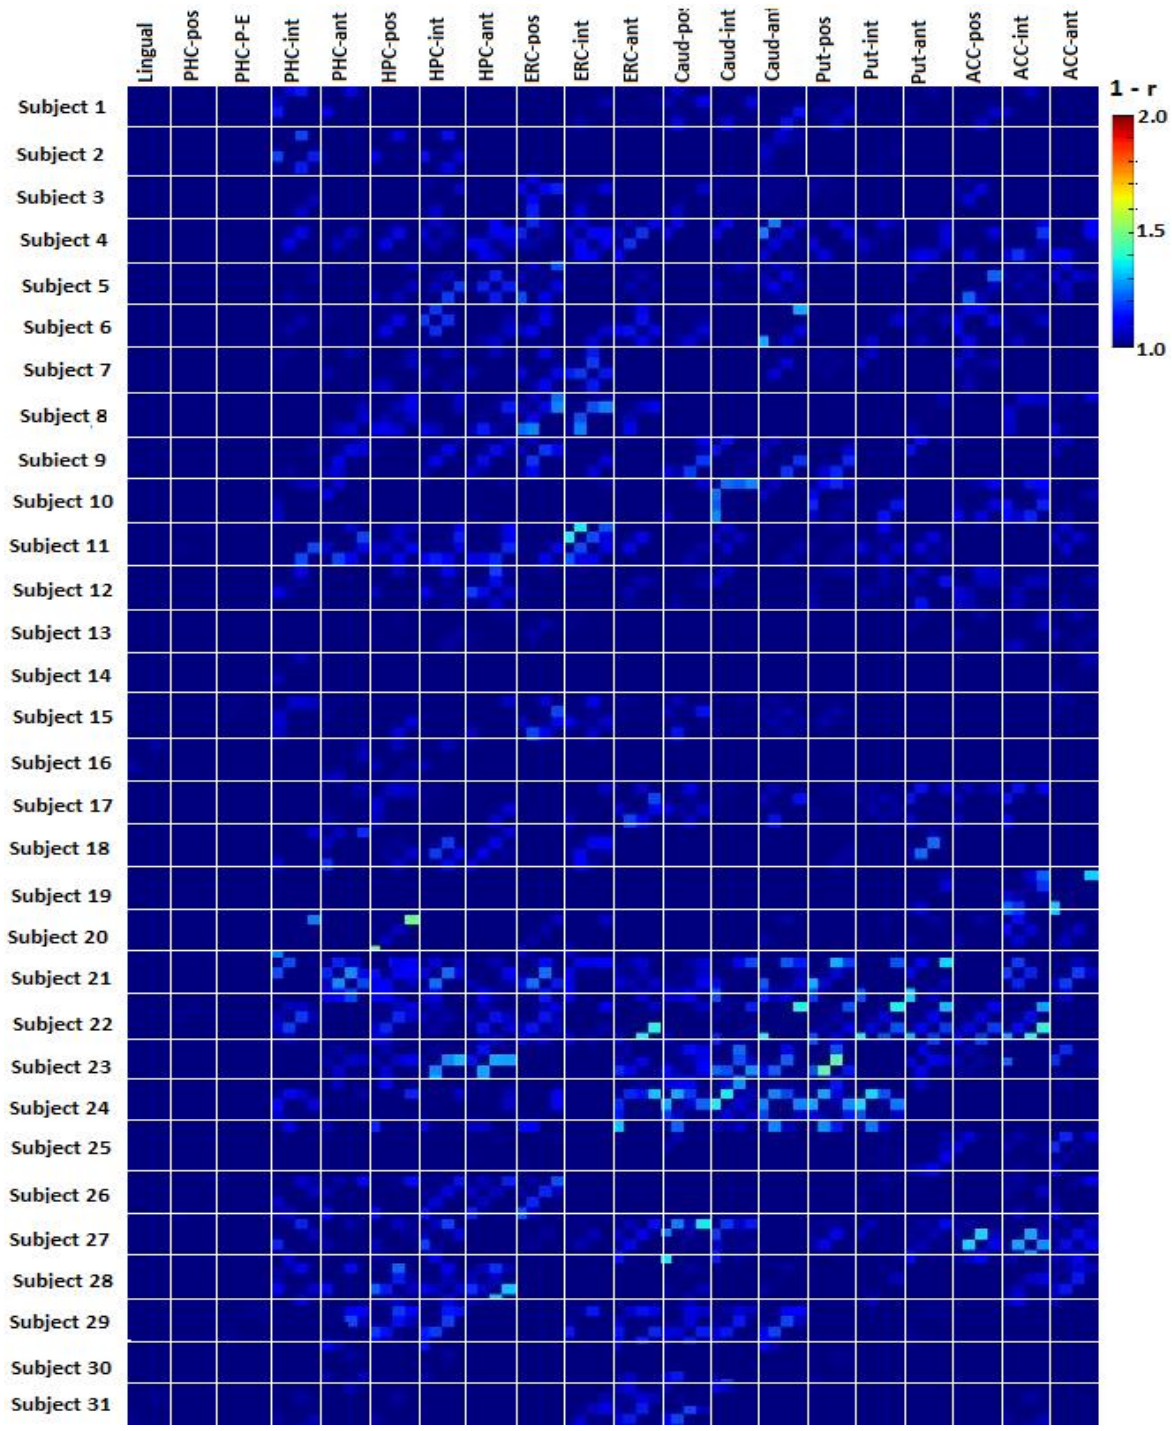

Figure 7-2 Activation pattern dissimilarity for positonal granularities in the poststimulus encoding period The granularities in each matrix involved fine-grained. medium-grained, coarse-grained and failed (see also figure 3-1). The figure includes only part of the entire span of dissimilarities, calculated as 1 – Pearson's r, so the range of possible values goes from 1 (zero correlation) to 2 (perfect anticorrelation). Each row represents a single subject, while the columns represent regions of interest (ROIs).The color bar on the top right shows the color coding along the dissimilarity spectrum. ROI abbreviations: pos=posterior, int=intermediate, ant=anterior; ACC=anterior cingulate cortex, Put=putamen, Caud=caudate, ERC=enorhinal cortex, HPC=hippocampus, PHC=parahippocampal cortex.

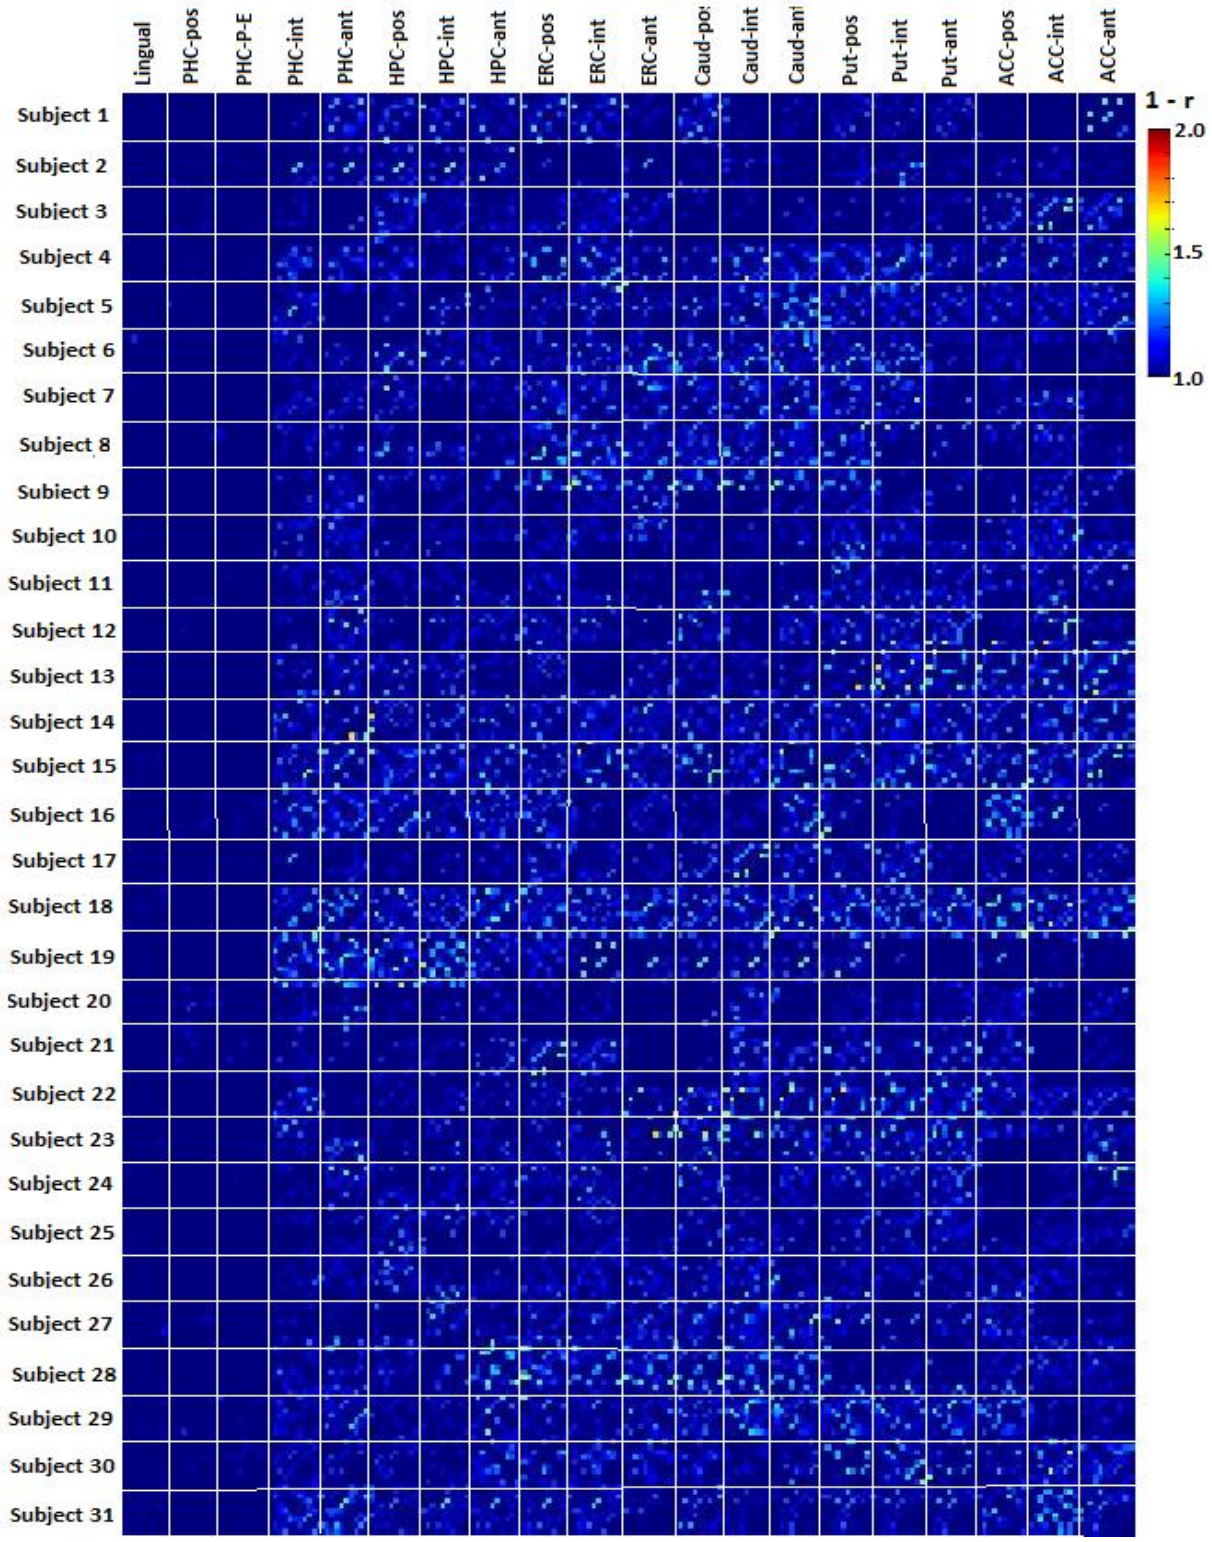

Figure 7-3 Activation pattern dissimilarity for outer border geometrical shapes in the stimulus presentation period The categories in each matrix involved ten different room shapes (see also figure 3-4). The figure includes only part of the entire span of dissimilarities, calculated as 1 – Pearson's r, so the range of possible values goes from 1 (zero correlation) to 2 (perfect anticorrelation). Each row represents a single subject, while the columns represent regions of interest (ROIs).The color bar on the top right shows the color coding along the dissimilarity spectrum. ROI abbreviations: pos=posterior, int=intermediate, ant=anterior; ACC=anterior cingulate cortex, Put=putamen, Caud=caudate, ERC=enorhinal cortex, HPC=hippocampus, PHC=parahippocampal cortex.

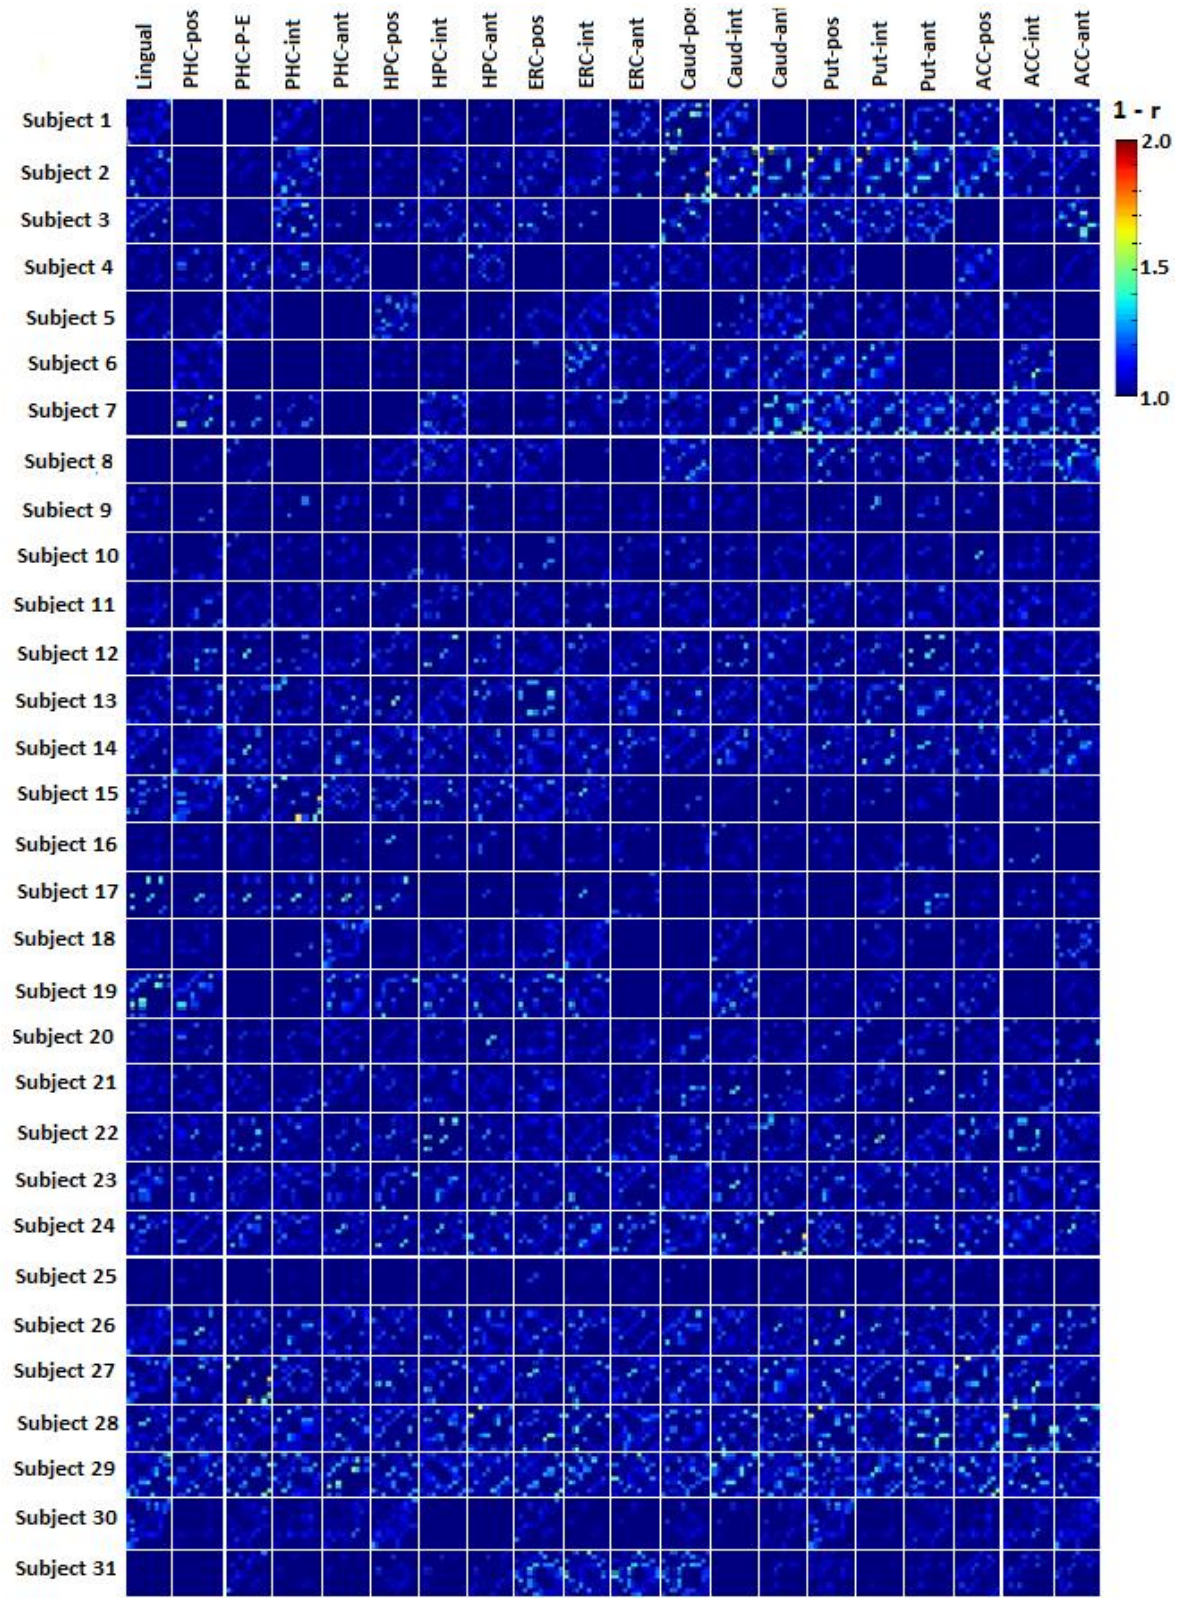

Figure 7-4 Activation pattern dissimilarity for outer border geometrical shapes in the poststimulus encoding period The categories in each matrix involved ten different room shapes (see also figure 3-4). The figure includes only part of the entire span of dissimilarities, calculated as 1 – Pearson's r, so the range of possible values goes from 1 (zero correlation) to 2 (perfect anticorrelation). Each row represents a single subject, while the columns represent regions of interest (ROIs). The color bar on the top right shows the color coding along the dissimilarity spectrum. ROI abbreviations: pos=posterior, int=intermediate, ant=anterior; ACC=anterior cingulate cortex, Put=putamen, Caud=caudate, ERC=enorhinal cortex, HPC=hippocampus, PHC=parahippocampal cortex.