Dynamics of Scene Representations in the Human Brain revealed by MEG and Deep Neural Networks

Radoslaw M. Cichy Aditya Khosla Dimitrios Pantazis Aude Oliva Massachusetts Institute of Technology

{rmcichy, khosla, pantazis, oliva}@mit.edu

Abstract

Human scene recognition is a rapid multistep process evolving over time from single scene image to spatial layout processing. We used multivariate pattern analyses on magnetoencephalography (MEG) data to unravel the time course of this cortical process. Following an early signal for lower-level visual analysis of single scenes at \sim 100ms, we found a marker of real-world scene size, i.e., spatial layout processing, at ~ 250 ms indexing neural representations robust to changes in unrelated scene properties and viewing conditions. For a quantitative explanation that captures the complexity of scene recognition, we compared MEG data to a deep neural network model trained on scene classification. Representations of scene size emerged intrinsically in the model, and resolved emerging neural scene size representation. Together our data provide a first description of an electrophysiological signal for layout processing in humans, and a novel quantitative model of how spatial layout representations may emerge in the human brain. The supplemental materials are available at: http://brainmodels.csail.mit.edu/scene-size

1. Introduction

Perceiving the geometry of space is a core ability shared by all animals, with brain structures for spatial layout perception and navigation preserved across rodents, monkeys and humans [18, 17, 15, 14, 40, 19, 29, 61]. Spatial layout perception, the demarcation of the boundaries and size of real-world visual space, plays a crucial mediating role in spatial cognition [4, 19, 63, 47, 31] between image-specific processing of individual scenes and navigation-related processing. Although the cortical loci of spatial layout perception in humans have been well described [1, 32, 37, 46, 5], the dynamics of spatial cognition remain unexplained, partly because neuronal markers indexing spatial processing remain unknown.

Operationalizing spatial layout as scene size, that is the

size of the space a scene subtends in the real-world [31, 46, 47], we report here an electrophysiological signal of spatial layout perception in the human brain. Using multivariate pattern classification [7, 9, 23] and representational similarity analysis [34, 33, 9] on millisecond-resolved magnetoencephalography data (MEG), we identified a marker of scene size around 250ms, preceded by and distinct from an early signal for lower-level visual analysis of scene images at ~ 100ms. Furthermore, we demonstrated that the scene size marker was independent of both low-level image features (i.e., luminance, contrast, clutter) and semantic properties (the category of the scene, i.e., kitchen, ballroom), thus indexing neural representations robust to changes in viewing conditions as encountered in real-world settings.

To provide a quantitative explanation how space size representations emerge in cortical circuits, we compared brain data to a deep neural network model trained to perform scene categorization [66, 65], termed deep scene network.The deep scene network intrinsically exhibited receptive fields specialized for layout analysis, such as textures and surface layout information, without ever having been explicitly taught any of those features.We showed that the deep scene neural network model predicted the human neural representation of single scenes and scene space size better than a deep object model and standard models of scene and object perception [44, 51]. This demonstrates the ability of the deep scene model to approximate human neural representations at successive levels of processing as they emerge over time.

Together our findings provide a first description of an electrophysiological signal for scene space processing in humans, and offer a novel quantitative and computational model of the dynamics of visual scene space representation in the cortex. Our results suggest that spatial layout representations naturally emerge in cortical circuits learning to differentiate visual environments [44].

2. Materials and Methods

2.1. Participants

Participants were 15 right-handed, healthy volunteers with normal or corrected-to-normal vision (mean age \pm s.d. = 25.87 \pm 5.38 years, 11 female). The Committee on the Use of Humans as Experimental Subjects (COUHES) at MIT approved the experiment and each participant gave written informed consent for participation in the study, for data analysis and publication of study results.

2.2. Stimulus Material and Experimental Design

The image set consisted of 48 scene images differing in four factors with two levels each, namely two scene properties: physical size (small, large) and clutter level (low, high); and two image properties: contrast (low, high) and luminance (low, high) (Figure 1A). There were 3 unique images for every level combination, for example 3 images of small size, low clutter, low contrast and low luminance. The image set was based on behaviorally validated images of scenes differing in size and clutter level, sub-sampling the two highest and lowest levels of factors size and clutter [47]. Small scenes were of size that would typically fit 2-8 people, whereas large scenes would fit hundreds to thousands. Similarly, low clutter level scenes were empty or nearly empty rooms, whereas high clutter scenes contained multiple objects throughout. The contrast and luminance was adjusted to specific values for each image: images of low and high contrast had root mean square values of 34%and 50% respectively; images of low and high luminance had root mean square values of 34% and 51% respectively.

Participants viewed a series of scene images while MEG data was recorded (Figure 1B). Images subtended 8 of visual angle in both width and height and were presented centrally on a gray screen (42.5% luminance) for 0.5s in random order with an inter-stimulus interval (ISI) of 1-1.2s, overlaid with a central red fixation cross. Every 4 trials on average (range 3-5 trials, equally probable) a target image depicting concentric circles was presented prompting participants to press a button and blink their eyes in response. ISI between the concentric-circles and the next trial was 2s to allow time for eye blinks. Target image trials were not included in analysis. Each participant completed 15 runs of 312s each. Every image was presented four times in a run, resulting in 60 trials per image per participant in total.

2.3. MEG Recording

We recorded continuous MEG signals from 306 channels (Elektra Neuromag TRIUX, Elekta, Stockholm) at a sampling rate of 1000Hz. Raw data was band-pass filtered between 0.03 and 330Hz, and pre-processed using spatiotemporal filters (maxfilter software, Elekta, Stockholm). We used Brainstorm [56] to extract peri-stimulus MEG signals from -100 to +900ms with respect to stimulus onset, and then normalized each channel by its baseline (-100 to 0ms) mean and standard deviation, and temporally smoothed the time series with a 20ms sliding window.

2.4. MEG Multivariate Pattern Classification

Single image classification: To determine whether MEG signals can discriminate experimental conditions (scene images), data were subjected to classification analyses using linear support-vector machines (SVM) [41] in the libsym implementation¹ [8] with a fixed regularization parameter C = 1. For each time point t, the processed MEG sensor measurements were concatenated to 306-dimensional pattern vectors, resulting in M = 60 raw pattern vectors per condition (Figure 1B). To reduce computational load and improve signal-to-noise ratio, we subaveraged the M vectors in groups of k = 5 with random assignment, thus obtaining M/k averaged pattern vectors. We then measured the performance of the SVM classifier to discriminate between every pair (i, j) of conditions using a leave-one-out approach: M/k - 1 vectors were randomly assigned to the training test, and 1 vector to the testing set to evaluate the classifier decoding accuracy. The above procedure was repeated 100 times, each with random assignment of the M raw pattern vectors to M/k averaged pattern vectors, and the average decoding accuracy was assigned to the (i, j) element of a 48×48 decoding matrix indexed by condition. The decoding matrix is symmetric with an undefined diagonal. We obtained one decoding matrix (representational dissimilarity matrix or RDM) for each time point t.

Representational clustering analysis for size: Interpreting decoding accuracy as a measure of dissimilarity between patterns, and thus as a distance measure in representational space [33, 9], we partitioned the RDM decoding matrix into within- and between-level segments for the factor scene size (Figure 2A). The average of between-size minus within-size matrix elements produced representational distances (percent decoding accuracy difference) indicative of clustering of visual representations by scene size.

Cross-classification: To assess whether scene size representations were robust to changes of other factors, we used SVM cross-classification assigning different levels of experimental factors to the training and testing set. For example, Figure 2C shows the cross-classification of scene size (small vs. large) across clutter, implemented by limiting the training set to high clutter scenes and the testing set to low clutter scenes. The procedure was repeated with reverse assignment (low clutter for training set and high clutter for testing set) and decoding results were averaged. The training set was 12 times larger (M = 720 raw pattern vectors) than for single-image decoding, as we pooled

¹http://www.csie.ntu.edu.tw/~cjlin/libsvm



Figure 1: Image set and single-image decoding. (A) The stimulus set comprised 48 indoor scene images differing in the size of the space depicted (small vs. large), as well as clutter, contrast, and luminance level; here each experimental factor combination is exemplified by one image. The image set was based on behaviorally validated images of scenes differing in size and clutter level, de-correlating factors size and clutter explicitly by experimental design [47]. Note that size refers to the size of the real-world space depicted on the image, not the stimulus parameters; all images subtended 8 visual angle during the experiment. (B) Time-resolved (1ms steps from -100 to +900ms with respect to stimulus onset) pair-wise support vector machine classification of experimental conditions based on MEG sensor level patterns. Classification independent of other experimental factors. Decoding matrices. (C) Decoding results for single scene classification independent of other experimental factors. Decoding results were averaged across the dark blocks (matrix inset), to control for luminance, contrast, clutter level and scene size differences. Inset shows indexing of matrix by image conditions. Horizontal line below curve indicates significant time points (n = 15, cluster-definition threshold P < 0.05, corrected significance level P < 0.05); gray vertical line indicates image onset.

	onset latency	peak latency		onset latency	neak latency	
aluttan lawal	FC(40, 71)	107(102 - 101)		onset fatency	реак таселеу	
clutter level	50(42 - 71)	107(105 - 191)	size across clutter level	226(134 - 491)	283(101 - 520)	
luminance level	644(68-709)	625(146-725)	Size deloss clutter lever	220 (104 451)	200 (101 020)	
			size across luminance level	183(138 - 244)	217(148 - 277)	
contrast level	53(42 - 128)	74(68-87)	size across fulfillance level	100(100-244)	211(140 - 211)	
			size across contrast level	138(129 - 179)	238(184 - 252)	
			Size deross contrast lever	100 (125 115)	200 (101 202)	
	33 (42 - 126)	14 (00 - 01)	size across contrast level	138(129 - 179)	238(184 - 252)	

(a) Clutter, luminance and contrast level representation time course information.

(b) Time course of cross-classification for scene size.

Table 1: Onset and peak latencies for MEG classification analyses. Onset and peak latency (n = 15, P < 0.05, cluster-level corrected, cluster-definition threshold P < 0.05) with 95% confidence intervals in brackets.

trials across single images that had the same level of clutter and size. We averaged pattern vectors by sub-averaging groups of k = 60 raw pattern vectors before the leave-oneout SVM classification. Cross-classification analysis was performed for the cross-classification of the factors scene size (Figure 2D) and scene clutter (Supplementary Figure 3) with respect to changes across all other factors.

2.5. Low and High-Level Computational Models of Image Statistics

We assessed whether computational models of object and scene recognition predicted scene size from our image material. For this we compared four models: two deep convolutional neural networks that were either trained to perform (1) scene or (2) object classification; (3) the GIST descriptor [44], i.e., a model summarizing the distribution of orientation and spatial frequency in an image that has been



Figure 2: Scene size is discriminated by visual representations. (A) To determine the time course of scene size processing we determined when visual representations clustered by scene size. For this we subtracted mean within-size decoding accuracies (dark gray, -) from between-size decoding accuracies (light gray, +). (B) Scene size was discriminated by visual representations late in time (onset of significance at 141ms (118 - 156ms), peak at 249ms (150 - 274ms). Gray shaded area indicates 95% confidence intervals determined by bootstrapping participants. (C) Cross-classification analysis, exemplified for cross-classification of scene size across clutter level. A classifier was trained to discriminate scene size on high clutter images, and tested on low clutter images. Results were averaged following an opposite assignment of clutter images to training and testing sets. Before enteringcross-classification analysis, MEG trials were grouped by clutter and size level respectively independent of image identity. A similar cross-classification analysis was applied for other image and scene properties. (D) Results of cross-classification analysis indicated robustness of scene size visual representations to changes in other scene and image properties (scene clutter, luminance, and contrast). Horizontal lines indicate significant time points (n = 15, cluster-definition threshold P < 0.05, corrected significance level P < 0.05); gray vertical line indicates image onset. For result curves with 95% confidence intervals see Supplementary Figure 2.

shown to predict scene properties, among them size; and (4) HMAX model [54], a model of object recognition most akin in structure to low-level visual areas V1/V2. We computed the output of each of these models for each image as described below.

Deep neural networks: The deep neural network architecture was implemented following Krizhevsky et al., 2012 [35]. We chose this particular architecture because it was the best performing model in object classification in the ImageNet 2012 competition [52], uses biologically-inspired local operations (convolution, normalization, max-pooling), and has been compared to human and monkey brain activity successfully [22, 27, 28]. The network architecture had 8 layers with the first 5 layers being convolutional and the last 3 fully connected. For an enumeration of units and features for each layer see Table 3. We used the convolution stage of each layer as model output for further analysis.

We constructed two deep neural networks that differed in the visual categorization task and visual material they were trained on. A deep scene model was trained on 216 scene categories from the Places dataset² [66] with 1300 images per category. A deep object model was trained on 683 different objects with 900,000 images from the ImageNet dataset³ [11] with similar number of images per object category (~1300). Both deep neural networks were trained on GPUs using the Caffe toolbox [25]. In detail, the networks were trained for 450,000 iterations, with an initial learning rate of 0.01 and a step multiple of 0.1 every 100,000 iterations. Momentum and weight decay were kept constant at 0.9 and 0.0005 respectively.

To visualize receptive fields (RFs) of model neurons in the deep scene network (Figure 3B) we used a reduction method [65]. In short, for a particular neuron we determined the K images activating the neuron most strongly. To determine the empirical size of the RF, we replicated the K images many times with small random occluders at different positions in the image. We then passed the occluded images into the deep scene network and compared the output to the original image, constructing the discrepancy map that indicates which part of the image drives the neuron. We then recentered discrepancy maps and averaged, generating the final RF. To illustrate the RFs tuning we further plot the image patches corresponding to the top activation regions inside the RFs (Figure 3B).

GIST: For the GIST descriptor [44], each image was filtered by a bank of Gabor filters with 8 orientations and 4 spatial frequencies (32 filters). Filter outputs were averaged in a 4×4 grid, resulting in a 512-dimensional feature vector. The GIST descriptor represents images in terms of spatial frequencies and orientations by position⁴.

HMAX: We used the HMAX model as applied and de-

²http://places.csail.mit.edu/

³http://www.image-net.org/

⁴Code: http://people.csail.mit.edu:/torralba/code/spatialenvelope/

scribed by Serre et al [54], a model inspired by the hierarchical organization of the visual cortex. In short, HMAX consists of two sets of alternating S and C layers, i.e., in total 4 layers. The S-layers convolve the input with predefined filters, and the C layers perform a max operation.

2.6. Linking Computational Models of Vision to Brain data

We used representational dissimilarity analysis to compare the output of computational models to brain data. First, we recorded the output of each model for each of the 48 images of the image set. Then, to compare to human brain data, we calculated the pair-wise dissimilarities between model outputs by 1- Spearman's rank order correlation ρ . This formed 48×48 model dissimilarity matrices (RDMs), one for each layer of each model: 8 for the deep scene and deep object network, 1 for GIST, and 4 for HMAX.

To compare models and brains, we determined whether images that were similarly represented in a computational network were also similarly represented in the brain. This was achieved by computing the similarity (Spearman's ρ) of layer-specific model dissimilarity matrix with the timepoint specific MEG decoding matrix for every subject and time point and averaging results.

We then determined whether the computational models predicted the size of a scene. We formulated an explicit size model, i.e., a 48×48 matrix with entries of 1 where images differed in size and 0 otherwise. Equivalent matrices were produced for scene clutter, contrast and luminance (Supplementary Figure 1). Correlation of the explicit size model with any computational model RDM yielded a measure of how well computational models predicted scene size. Finally, we determined whether the above computational models accounted for neural representations of scene size observed in MEG data. For this, we reformulated the representational clustering analysis in a correlation framework. The two measures are equivalent except that the correlation analysis takes into account the variability of the data, which the clustering analysis does not for the benefit of clear interpretability as percent change in decoding accuracy. The procedure had two steps. First, we calculated the similarity (Spearman's ρ) of the MEG decoding accuracy matrix with the explicit size model for each time point and each participant. Second, we re-calculated the similarity (Spearman's ρ) of the MEG decoding accuracy matrix with the explicit size model after partialling out all of the layer-specific RDMs of a given computational model.

2.7. Statistical Testing

We used permutation tests for cluster-size inference, and bootstrap tests to determine confidence intervals of onset times for maxima, cluster onsets and peak-to-peak latency differences [42, 9, 45].

Sign permutation tests: For the permutation tests, depending on the statistic of interest our null hypothesis was that the MEG decoding time series were equal to 50%chance level, or that the decoding accuracy difference of between- minus within-level segments of the MEG decoding matrix was equal to 0, or that the correlation values were equal to 0. In all cases, under the null hypothesis the sign of the observed effect in the MEG data is randomly permutable, corresponding to a sign-permutation test that randomly multiplies the participant-specific data with +1 or -1. We created 1,000 permutation samples, every time recomputing the statistic of interest. This resulted in an empirical distribution of the data, allowing us to convert our original data, as well as the permutation samples, into Pvalues. We then performed cluster-size inference by setting a P = 0.05 cluster-definition threshold on the original data and permutation samples, and computing a P = 0.05cluster size threshold from the empirical distribution of the resampled data.

Bootstrapping: To calculate confidence intervals (95%) on cluster onset and peak latencies, we bootstrapped the sample of participants 1,000 times with replacement. For each bootstrap sample, we repeated the above permutation analysis yielding distributions of the cluster onset and peak latency, allowing estimation of confidence intervals. In addition, for each bootstrap sample, we determined the peak-to-peak latency difference for scene size clustering and individual scene image classification. This yielded an empirical distribution of peak-to-peak latencies. Setting P < 0.05, we rejected the null hypothesis of a latency difference if the confidence interval did not include 0.

Label permutation tests: For testing the significance of correlation between the computational model RDMs and the scene size model, we relied on a permutation test of image labels. This effectively corresponded to randomly permuting the columns (and accordingly the rows) of the computational model RDMs 1,000 times, and then calculating the correlation between the permuted matrix and the explicit size model matrix. This yielded an empirical distribution of the data, allowing us to convert our statistic into P-values. Effects were reported as significant when passing a P = 0.05 threshold. Results were FDR-corrected for multiple comparisons.

3. Results

Human participants (n = 15) viewed images of 48 realworld indoor scenes that differed in the layout property size, as well as in the level of clutter, contrast and luminance (Figure 1A) while brain activity was recorded with MEG. While often real-world scene size and clutter level correlate, here we de-correlated those stimulus properties explicitly by experimental design, based on independent behavioral validation [47] to allow independent assessment. Im-

	onset latency	peak latency			
GIST	47(45-149)	80(76-159)			
HMAX	48(25-121)	74(61-80)			
deep object network	55(20-61)	97(83-117)			
deep scene network	47(23-59)	83(79-112)			

	onset latency	peak latency
deep scene network minus GIST	58(50-78)	108 (81 - 213)
deep scene network minus HMAX	75(62-86)	108(97 - 122)
deep scene network minus deep object network	-	-

(a) Correlation of models to MEG data.

(b) Comparison of MEG-model correlation for the deep scene network and all other models.

Table 2: Onset and peak latencies for model-MEG representational similarity analysis. Onset and peak latency (n = 15, P < 0.05, cluster-level corrected, cluster-definition threshold P < 0.05) with 95% confidence intervals in brackets.



Figure 3: Predicting emerging neural representations of single scene images by computational models. (A) Architecture of deep convolutional neural network trained on scene categorization (deep scene network). (B) Receptive field (RF) of example deep scene neurons in layers 1, 2, 4, and 5. Each row represents one neuron. The left column indicates size of RF, and the remaining columns indicate image patches most strongly activating these neurons. Lower layers had small RFs with simple Gabor filter-like sensitivity, whereas higher layers had increasingly large RFs sensitive to complex forms. RFs for whole objects, texture, and surface layout information emerged although these features were not explicitly taught to the deep scene model. (C) We used representational dissimilarity analysis to compare visual representations in brains with models. For every time point, we compared subject-specific MEG RDMs (Spearman's ρ) to model RDMs and results were averaged across subjects. (D) All investigated models significantly predicted emerging visual representations in the brain, with superior performance for the deep neural networks compared to HMAX and GIST. Horizontal lines indicate significant time points (n = 15, cluster-definition threshold P < 0.05, corrected significance level P < 0.05); gray vertical line indicates image onset.

ages were presented for 0.5s with an inter-trial interval of 1 - 1.2s (Figure 1B). Participants performed an orthogonal object-detection task on an image of concentric circles appearing every four trials on average. Concentric circle trials were excluded from further analysis.

To determine the timing of cortical scene processing we used a decoding approach: we determined the time course with which experimental conditions (scene images) were discriminated by visual representations in MEG data. For this, we extracted peri-stimulus MEG time series in 1ms resolution from -100 to +900ms with respect to stimulus onset for each subject. For each time point independently we classified scene images pair-wise by MEG sensor patterns(support vector classification, Figure 1C). Timepoint specific classification results (percentage decoding accuracy, 50% chance level) were stored in a 48×48 decoding accuracy matrix, indexed by image conditions in rows and columns (Figure 1C, inset). This matrix is symmetric with

undefined diagonal. Repeating this procedure for every time point yielded a set of decoding matrices (for a movie of decoding accuracy matrices over time, averaged across subjects, see Supplementary Movie 1). Interpreting decoding accuracies as a representational dissimilarity measure, each 48×48 matrix summarized, for a given time point, which conditions were represented similarly (low decoding accuracy) or dissimilarly (high decoding accuracy). The matrix was thus termed MEG representational dissimilarity matrix (RDM) [9, 43].

Throughout, we determined random-effects significance non-parametrically using a cluster-based randomization approach (cluster-definition threshold P < 0.05, corrected significance level P < 0.05) [42, 45, 38]. 95% confidence intervals for mean peak latencies and onsets (reported in parentheses throughout the results) were determined by bootstrapping the participant sample.

layer	conv1	pool/norm1	conv2	pool/norm2	conv3	conv4	conv5	pool5	fc1	fc2	fc3
units	96	96	256	256	384	384	256	256	4096	4096	683 / 216
features	55×55	27×27	27×27	13×13	13×13	13×13	13×13	6×6	1	1	1

Table 3: Number of units and features for each CNN layer. Units and features of the deep neural network architecture were similar as proposed in [35]. All deep neural networks were identical with the exception of the number of nodes in the last layer (output layer) as dictated by the number of training categories, i.e. 683 for the deep object network, 216 for deep scene network. Abbreviations: conv = convolutional layer, pool = pooling layer; norm = normalization layer; fc1 – 3 = fully connected layers. The 8 layers referred to in the manuscript correspond to the convolution stage for layers 1 - 5, and the fc1 – 3 stage for layers 6 - 8 respectively.

3.1. Neural representations of single scene images emerged early in cortical processing

We first investigated the temporal dynamics of imagespecific individual scene information in the brain. To determine the time course with which individual scene images were discriminated by visual representations in MEG data, we averaged the elements of each RDM matrix representing pairwise comparisons with matched experimental factors(luminance, contrast, clutter level and scene size)(Figure 1C). We found that the time course rose sharply after image onset, reaching significance at 50ms(45 - 52ms) and a peak at 97ms (94 - 102ms).This indicates that single scene images were discriminated early by visual representations, similar to single images with other visual content [58, 7, 9, 23], suggesting a common source in early visual areas [9].

3.2. Neural representations of scene size emerged later in time and were robust to changes in viewing conditions and other scene properties

When is the spatial layout property scene size processed by the brain? To investigate, we partitioned the decoding accuracy matrix into two subdivisions: images of different (between subdivision light gray, +) and similar size level (within subdivision, dark gray, -). The difference of mean between-size minus within-size decoding accuracy is a measure of clustering of visual representations by size (Figure). Peaks in this measure indicate time points at which MEG sensor patterns cluster maximally by scene size, suggesting underlying neural visual representations allowing for explicit, linear readout [12] of scene size by the brain. Scene size (Figure 2B) was discriminated first at 141ms (118-156ms) and peaked at 249ms (150-274ms), which was significantly later than the peak in single image classification (P = 0.001, bootstrap test of peak-latency differences).

Equivalent analyses for the experimental factors scene clutter, contrast, and luminance level yielded diverse time courses (Supplementary Figure 1, Table 1(a)). Importantly, representations of low-level image property contrast emerged significantly earlier than scene size (P = 0.004)

and clutter (P = 0.006, bootstrap test of peak-latency differences). For the factor luminance, only a weak effect and thus no significant onset response was observed, suggesting a pre-cortical luminance normalization mechanism.

To be of use in the real world, visual representations of scene size must be robust against changes of other scene properties, such as clutter level (i.e., space filled by different types and amounts of objects) and semantic category (i.e., the label by which we name it), and changes in viewing conditions, such as luminance and contrast. We investigated the robustness of scene size representations to all these factors using cross-classification (Figure 2C; for 95% confidence intervals on curves see Supplementary Figure 2). For this we determined how well a classifier trained to distinguish scenes at one clutter level could distinguish scenes at the other level, while collapsing data across single image conditions of same level in size and clutter. We found that scene size was robust to changes in scene clutter, luminance and contrast (Figure 2D; onsets and peaks in Table 1(b)). Note that by experimental design, the scene category always differed across size level, such that cross- classification also established that scene size was discriminated by visual representations independent of the scene category.

An analogous analysis for clutter level yielded evidence for viewing-condition independent clutter level representations (Supplementary Figure 3), reinforcing the notion of clutter level as a robust and relevant dimension of scene representations in the human brain [47]. Finally, an analysis revealing persistent and transient components of scene representations indicated strong persistent components for scene size and clutter representations, with little or no evidence for contrast and luminance (Supplementary Figure 4). Persistence of scene size and clutter level representations further reinforces the notion of size and clutter level representations being important end products of visual computations kept online by the brain for further processing and behavioral guidance.

In sum, our results constitute evidence for representations of scene size in human brains from non-invasive electrophysiology, apt to describe scene size discrimination under real world changes in viewing conditions.

3.3. Neural representations of single scene images were predicted by deep convolutional neural networks trained on real world scene categorization

Visual scene recognition in cortex is a complex hierarchical multi-step process, whose understanding necessitates a quantitative model that captures this complexity.Here, we evaluated whether an 8-layer deep neural network trained to perform scene classification on 205 different scene categories [66] human scene representations. We refer to this network as deep scene network (Figure 3A).Investigation of the receptive fields (RFs) of model neurons using a reduction method [65] indicated a gradient of increasing complexity from low to high layers, and selectivity to whole objects, texture, and surface layout information (Figure 3B). This suggests that the network might be able to capture information about both single scenes and scene layout properties.

To determine the extent to which visual representations learned by the deep scene model and the human brain are comparable, we used representational similarity analysis [34, 9]. The key idea is that if two images evoke similar responses in the model, they should evoke similar responses in the brain, too.

For the deep neural network, we first estimated image response patterns by computing the output of each model layer to each of the 48 images. We then constructed layer-resolved 48×48 representational dissimilarity matrices (RDMs) by calculating the pairwise dissimilarity (1-Spearman's ρ) across all model response patterns for each layer output.

We then compared (Spearman's ρ) the layer-specific deep scene model RDMs with the time-resolved MEG RDMs and averaged results over layers, yielding a time course indicating how well the deep scene model predicted and thus explained scene representations (Figure 3D). To compare against other models, we performed equivalent analyses to a deep neural network trained on object-categorization (termed deep object network) and standard models of object (HMAX) and scene-recognition (GIST) [44, 53].

We found that the deep object and scene network performed similarly at predicting visual representations over time (Figure 3D, for details see Table 2(a); for layerresolved results see Supplementary Figure 5), and better than the HMAX and GIST models (for direct quantitative comparison see Supplementary Figure 6).

In sum, our results show that brain representations of single scene images were best predicted by deep neural network models trained on real-world categorization tasks, demonstrating the ability of the models to capture the complexity of scene recognition, and their semblance to the human brain representations.

3.4. Representations of scene size emerged in the deep scene model

Beyond prediction of neural representations of single scene images, does the deep scene neural network indicate the spatial layout property scene size? To visualize, we used multidimensional scaling (MDS) on layer-specific model RDMs, and plotted the 48 scene images into the resulting 2D arrangement color-coded for scene size (black= small, gray = large). We found a progression in the representation of scene size in the deep scene network: low layers showed no structure, whereas high layers displayed a progressively clearer representation of scene size(A). A similar, but weaker progression, was visible for the deep object network (Figure 4B). Comparable analysis for HMAX and GIST (Figure 4C,D) found no prominent representation of size.

We quantified this descriptive finding by computing the similarity of model RDMs with an explicit size model (an RDM with entries 0 for images of similar size, 1 for images of dissimilar size; Figure 4E inset). We found a significant effect of size in all models (P < 0.05, FDR-corrected, stars above bars indicate significance). The size effect was larger in the deep neural networks than in GIST and HMAX, it was more pronounced in the high layers, and the deep scene network displayed a significantly stronger effect of scene size than the deep object network in layers 6 - 8 (stars between bars; for all pair-wise layer-specific comparisons see Supplementary Figure 7). A supplementary partial correlation analysis confirmed that the effect of size in the deep scene network was not explained by correlation with the other experimental factors (Supplementary Figure 8).

Together, these results indicate the deep scene network captured scene size better than all other models, and that scene size representations emerge gradually in the deep neural network hierarchy. Thus representations of visual space can emerge intrinsically in neural networks constrained to perform visual scene categorization without being trained to do so directly.

3.5. Neural representations of scene size emerged in the deep scene model

The previous sections demonstrated that representations of scene size emerged in both neural signals (Figure 2) and computational models (Figure 4). To evaluate the overlap between these two representations, we combined representational similarity analysis with partial correlation analysis [10] (Figure 5A).

We first computed the neural representations of scene size by correlating (Spearman's ρ) the MEG RDMs with the explicit size model (black curve). We then repeated the process, but only after partialling out all layer-specific RDMs of a model from the explicit size model (color-coded by model) (Figure 5B). The reasoning is that if neural sig-



Figure 4: **Representation of scene size in computational models of object and scene categorization.** (A–D) Layerspecific RDMs and corresponding 2D multidimensional scaling (MDS) plots for a deep scene network, deep object network, GIST, and HMAX. MDS plots are color-coded by scene size (small = black; large = gray). (E) Quantifying the representation of scene size in computational models. We compared (Spearman's ρ) each model's RDMs with an explicit size model (RDM with entries 0 for images of similar size, 1 for images of dissimilar size). Results are color-coded for each model. (F) Similar to (E) for clutter, contrast and luminance (results shown only for deep scene and object networks). While representations of the abstract scene properties size and clutter emerged with increasing layer number, the low-level image properties contrast and luminance successively abstracted away. Stars above bars indicate statistical significance. Stars between bars indicate significant differences between the corresponding layers of the deep scene vs. object network. Complete layer-wise comparisons available in Supplementary Figure 7. (n = 48; label permutation tests for statistical inference, P < 0.05, FDR-corrected for multiple comparisons).

nals and computational models carry the same scene size information, the scene size effect will vanish in the latter case.

When partialling out the effect of the deep scene network, the scene size effect was considerably reduced and was no longer statistically significant. In all other models, the effect was reduced but was still statistically significant (Figure 5B). Further, the reduction of the size effect was higher for the deep scene network than all other models (Figure 5C). Equivalent analyses for scene clutter, contrast and luminance indicated that the deep scene and object networks abolished all effects, while other models did not (Supplementary Figure 9).

Together, these results show that only the deep scene model captured the neural representation of scene size in the human brain, singling it out as the best of the scene representation models tested here.

4. Discussion

We characterized the emerging representation of scenes in the human brain using multivariate pattern classification methods [7, 9] and representational similarity analysis [34, 33] on combined MEG and computational model data. We found that neural representations of individual scenes and the low-level image property contrast emerged early, followed by the scene layout property scene size at around 250 ms. The neural representation of scene size was robust to changes in viewing conditions and scene properties such as contrast, luminance, clutter level and category. Our results provide novel evidence for an electrophysiological signal of scene processing in humans that remained stable under real-world viewing conditions. To capture the complexity of scene processing in the brain by a computational model, we trained a deep convolutional neural network on scene classification. We found that the deep scene model predicted representations of scenes in the brain and accounted for abstract properties such as scene size and clutter level better than alternative computational models, while abstracting away low-level image properties such as luminance and contrast level.

4.1. A multivariate pattern classification signal for the processing of scene layout property scene size

A large body of evidence from neuropsychology, neuroimaging and invasive work in humans and monkeys has identified locally circumscribed cortical regions of the brain dedicated to the processing of three fundamental visual categories: faces, bodies and scenes [2, 26, 1, 16, 60, 29]. For faces and bodies, respective electrophysiological signals in humans have been identified [2, 3, 24, 36, 55, 57]. However, electrophysiological markers for scene-specific processing have been identified for the auditory modality only [20, 59],

and a visual scene-specific electrophysiological signal had not been described until now.

Our results provide the first evidence for an electrophysiological signal of visual scene size processing in humans. Multivariate pattern classification analysis on MEG data revealed early discrimination of single scene images (peak at 97ms) and the low-level image property contrast (peak at 74ms), whereas the abstract property of space size was discriminated later (peak at 249ms). While early scene-specific information in the MEG likely emerged from low-level visual areas such as V1 [9], the subsequent scene size signal had properties commonly ascribed to higher stages of visual processing in ventral visual cortex: the representation of scene size was tolerant to changes occurring in real world viewing conditions, such as luminance, contrast, clutter level and category. The electrophysiological signal thus reflected scene size representations that could reliably be used for scene recognition in real world settings under changing viewing conditions [48, 12, 13]. This result paves the way to further studies of the representational format of scenes in the brain, e.g. by measuring the modulation of the scene-specific signal by other experimental factors.

The magnitude of the scene size effect, although consistent across subjects and statistically robust to multiple comparison correction, is small with a maximum of $\sim 1\%$. Note however that the size effect, in contrast to single image decoding (peak decodability at $\sim 79\%$), is not a measure of how well single images differing in size can be discriminated, but a difference measure of how much better images of different size can be discriminated rather than images of the same size. Thus, it is a measure of information about scene size over-and-above information distinguishing between any two single scenes. The magnitude of the size effect is comparable to effects reported for abstract visual properties such as animacy (1.9 and 1.1% respectively, [9]).

What might be the exact locus of the observed scene size signal in the brain? Previous research has indicated parametric encoding of scene size in parahippocampal place area (PPA) and retrosplenial cortex [47], corroborating numerous studies showing that spatial properties of scenes such as boundaries and layout are represented in these cortical regions [18, 17, 63]. Both onset and peak latency of the observed scene size signal concurred with reported latencies for parahippocampal cortex [39], suggesting that one or several nodes of the human spatial navigation network might be the source of the scene size signal.

Last, we found that not only scene size representations, but also scene clutter representations were tolerant to changes in viewing conditions, and emerged later than the low-level image contrast representations. These results complement previous findings in object perception research that representations of single objects emerge earlier in time than representations of more abstract properties such as cat-



Figure 5: The deep scene model accounts for more of the MEG size signal than other models. (A) We combined representational similarity with partial correlation analysis to determine which computational models explained emerging representations of scene size in the brain. (B) MEG representations of scene size (termed MEG size signal) before (black) and after (color-coded by model) partialling out the effect of different computational models. Only partialling out the effect of the deep scene network abolished the MEG size signal. (C) Difference in amount of variance partialled out from the size signal: comparing all models to the deep scene network. The deep scene network accounted for more MEG size signal than all other models (n = 15; cluster-definition threshold P < 0.05, significance threshold P < 0.05; results corrected for multiple comparisons by 5 for panel (B) and 3 for panel (C)).

egory membership [7, 9].

4.2. Neural representations of abstract scene properties such as scene size are explained by a deep neural network model trained on scene classification

Scene processing in the brain is a complex process necessitating a formal quantitative model that addresses this complexity. Here, our study of several models of scene and object recognition provided three novel results, each with fundamental theoretical implications.

First, deep neural networks offered the best characterization of neural scene representations compared to other models tested. The superiority of high performing deep neural networks over simpler models indicates that hierarchical architectures might be necessary to capture the structure of single scene representations in the human brain. While previous research has established that deep neural networks capture object representations in human and monkey inferior temporal cortex well, we demonstrated that a deep neural network explained millisecond-resolved dynamics underlying scene recognition from processing of low- to high-level properties, better than other models of object and scene-processing tested.

Concerning high-level abstract scene properties in particular, our results shed lights into the black box of cortical scene processing, providing novel insight both from the perspective of modeling, and of experimental brain science. From a modeling perspective, the near monotonic relationship between the representation of size and clutter level in the deep neural network and the network layer number indicates that scene size is an abstract scene property emerging through complex multi-step processing. From the perspective of experimental brain science, our results provide an advance in understanding neural representations of the processing of abstract scene properties such as spatial layout. Neuronal responses in high-level visual cortex are often sparse and nonlinear, making a full explanation by simple mathematical models in low-dimensional spaces or basic image statistics unlikely [21, 50, 62]. Instead, our result concurs with the finding that complex deep neural networks performing well on visual categorization tasks represent visual stimuli similar to the human brain [6, 64], and extends the claim to abstract properties of visual stimuli.

The second novel finding is that a deep neural network trained specifically on scene categorization had superior representation of scene size compared to a deep neural network trained on objects. Importantly, it also offered the best account of neural representations of scene size in the MEG, indicating that the underlying algorithmic computations matched the neuronal computations in the human brain. This indicates that the constraints imposed by the task the network is trained on, i.e., object or scene categorization, critically influenced the represented features. This makes plausible the notion that spatial representations emerge naturally and intrinsically in neural networks performing scene categorization, such as in the human brain. It further suggests that separate processing streams in the brain for different visual content, such as scenes, objects or faces, might be the result of differential task constraints imposed by classification of the respective visual input [13, 64].

The third novel finding is that representations of abstract scene properties (size, clutter level) emerged with increasing layers in deep neural networks, while low-level image properties (contrast, luminance) were increasingly abstracted away, mirroring the temporal processing sequence in the human brain: representations of low-level image properties emerged first, followed by representations of scene size and clutter level. This suggests common mechanisms in both and further strengthen the idea that deep neural networks are a promising model of the processing hierarchies constituting the human visual system, reinforcing the view of the visual brain as performing increasingly complex feature extraction over time [58, 36, 49, 53, 30, 13].

However, we did not observe a relationship between layer-specific representations in the deep neural networks and temporal dynamics in the human brain. Instead, the MEG signal predominantly reflected representations in low neural network layers (Supplementary Figure 5). One reason for this might be that our particular image set differed strongly in low-level features, thus strongly activating early visual areas that are best modeled by low neural network layers. Activity in low-level visual cortex was thus very strong, potentially masking weaker activity in high-level visual cortex that is invariant to changes in low level features. Another reason might be that while early visual regions are close to the MEG sensors, creating strong MEG signals, scene-processing cortical regions such as PPA are deeply harbored in the brain, creating weaker MEG signals. Future studies using image sets optimized to drive low-and high level visual cortex equally are necessary, to test whether layer-specific representations in deep neural networks can be mapped in both time and in space onto processing stages in the human brain.

4.3. Conclusions

Using a combination of multivariate pattern classification and computational models to study the dynamics in neuronal representation of scenes, we identified a neural marker of spatial layout processing in the human brain, and showed that a deep neural network model of scene categorization explains representations of spatial layout better than other models. Our results pave the way to future studies investigating the temporal dynamics of spatial layout processing, and highlight deep hierarchical architectures as the best models for understanding visual scene representations in the human brain.

Acknowledgements

This work was funded by National Eye Institute grant EY020484 (to A.O.), National Science Foundation grant BCS-1134780 (to D.P.), McGovern Institute Neurotechnology Program (to A.O. and D.P.), a Humboldt Scholarship (to R.M.C), and was conducted at the Athinoula A. Martinos Imaging Center at the McGovern Institute for Brain Research, Massachusetts Institute of Technology. We thank Santani Teng for helpful comments on the manuscript.

References

- G. K. Aguirre, E. Zarahn, and M. Desposito. An area within human ventral cortex sensitive to building stimuli: evidence and implications. *Neuron*, 21(2):373–383, 1998. 1, 10
- [2] T. Allison, H. Ginter, G. McCarthy, A. C. Nobre, A. Puce, M. Luby, and D. D. Spencer. Face recognition in human extrastriate cortex. *Journal of neurophysiology*, 71(2):821– 825, 1994. 10
- [3] S. Bentin, T. Allison, A. Puce, E. Perez, and G. McCarthy. Electrophysiological studies of face perception in humans. *Journal of cognitive neuroscience*, 8(6):551–565, 1996. 10
- [4] C. M. Bird, C. Capponi, J. A. King, C. F. Doeller, and N. Burgess. Establishing the boundaries: the hippocampal contribution to imagining scenes. *The Journal of Neuroscience*, 30(35):11688–11695, 2010. 1
- [5] H. M. Bonnici, D. Kumaran, M. J. Chadwick, N. Weiskopf, D. Hassabis, and E. A. Maguire. Decoding representations of scenes in the medial temporal lobes. *Hippocampus*, 22(5):1143–1153, 2012. 1
- [6] C. F. Cadieu, H. Hong, D. Yamins, N. Pinto, N. J. Majaj, and J. J. DiCarlo. The neural representation benchmark and its evaluation on brain and machine. *arXiv preprint arXiv:1301.3530*, 2013. 11
- [7] T. Carlson, D. A. Tovar, A. Alink, and N. Kriegeskorte. Representational dynamics of object vision: the first 1000 ms. *Journal of vision*, 13(10):1, 2013. 1, 7, 10, 11
- [8] C.-C. Chang and C.-J. Lin. Libsvm: A library for support vector machines. ACM Transactions on Intelligent Systems and Technology (TIST), 2(3):27, 2011. 2
- [9] R. M. Cichy, D. Pantazis, and A. Oliva. Resolving human object recognition in space and time. *Nature neuroscience*, 17(3):455–462, 2014. 1, 2, 5, 6, 7, 8, 10, 11
- [10] A. Clarke and L. K. Tyler. Object-specific semantic coding in human perirhinal cortex. *The Journal of Neuroscience*, 34(14):4766–4775, 2014. 8
- [11] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition*, 2009. CVPR 2009. IEEE Conference on, pages 248–255. IEEE, 2009. 4
- [12] J. J. DiCarlo and D. D. Cox. Untangling invariant object recognition. *Trends in cognitive sciences*, 11(8):333–341, 2007. 7, 10
- [13] J. J. DiCarlo, D. Zoccolan, and N. C. Rust. How does the brain solve visual object recognition? *Neuron*, 73(3):415– 434, 2012. 10, 11, 12
- [14] C. F. Doeller, C. Barry, and N. Burgess. Evidence for grid cells in a human memory network. *Nature*, 463(7281):657– 661, 2010. 1
- [15] C. F. Doeller, J. A. King, and N. Burgess. Parallel striatal and hippocampal systems for landmarks and boundaries in spatial memory. *Proceedings of the National Academy of Sciences*, 105(15):5915–5920, 2008. 1
- [16] P. E. Downing, Y. Jiang, M. Shuman, and N. Kanwisher. A cortical area selective for visual processing of the human body. *Science*, 293(5539):2470–2473, 2001. 10

- [17] R. Epstein, A. Harris, D. Stanley, and N. Kanwisher. The parahippocampal place area: Recognition, navigation, or encoding? *Neuron*, 23(1):115–125, 1999. 1, 10
- [18] R. Epstein and N. Kanwisher. A cortical representation of the local visual environment. *Nature*, 392(6676):598–601, 1998. 1, 10
- [19] R. A. Epstein. Cognitive neuroscience: scene layout from vision and touch. *Current Biology*, 21(11):R437–R438, 2011.
- [20] N. Fujiki, K. A. Riederer, V. JousmaÈki, J. P. MaÈkelaÈ, and R. Hari. Human cortical representation of virtual auditory space: differences between sound azimuth and elevation. *European Journal of Neuroscience*, 16(11):2207–2213, 2002. 10
- [21] I. I. Groen, S. Ghebreab, H. Prins, V. A. Lamme, and H. S. Scholte. From image statistics to scene gist: evoked neural activity reveals transition from low-level natural image structure to scene category. *The Journal of Neuroscience*, 33(48):18814–18824, 2013. 11
- [22] U. Güçlü and M. A. van Gerven. Deep neural networks reveal a gradient in the complexity of neural representations across the brain's ventral visual pathway. *arXiv preprint arXiv:1411.6422*, 2014. 4
- [23] L. Isik, E. M. Meyers, J. Z. Leibo, and T. Poggio. The dynamics of invariant object recognition in the human visual system. *Journal of neurophysiology*, 111(1):91–102, 2014. 1, 7
- [24] D. A. Jeffreys. Evoked potential studies of face and object processing. *Visual Cognition*, 3(1):1–38, 1996. 10
- [25] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. In *Proceedings of the ACM International Conference on Multimedia*, pages 675–678. ACM, 2014. 4
- [26] N. Kanwisher, J. McDermott, and M. M. Chun. The fusiform face area: a module in human extrastriate cortex specialized for face perception. *The Journal of Neuroscience*, 17(11):4302–4311, 1997. 10
- [27] S.-M. Khaligh-Razavi, L. Henriksson, K. Kay, and N. Kriegeskorte. Explaining the hierarchy of visual representational geometries by remixing of features from many computational vision models. *bioRxiv*, page 009936, 2014. 4
- [28] S.-M. Khaligh-Razavi and N. Kriegeskorte. Deep supervised, but not unsupervised, models may explain it cortical representation. *PLOS Computational Biology*, 2014. 4
- [29] S. Kornblith, X. Cheng, S. Ohayon, and D. Y. Tsao. A network for scene processing in the macaque temporal lobe. *Neuron*, 79(4):766–781, 2013. 1, 10
- [30] Z. Kourtzi and C. E. Connor. Neural representations for object perception: structure, category, and adaptive coding. *Annual review of neuroscience*, 34:45–67, 2011. 12
- [31] D. J. Kravitz, C. S. Peng, and C. I. Baker. Real-world scene representations in high-level visual cortex: it's the spaces more than the places. *The Journal of Neuroscience*, 31(20):7322–7333, 2011.

- [32] D. J. Kravitz, K. S. Saleem, C. I. Baker, and M. Mishkin. A new neural framework for visuospatial processing. *Nature Reviews Neuroscience*, 12(4):217–230, 2011.
- [33] N. Kriegeskorte and R. A. Kievit. Representational geometry: integrating cognition, computation, and the brain. *Trends* in cognitive sciences, 17(8):401–412, 2013. 1, 2, 10
- [34] N. Kriegeskorte, M. Mur, and P. Bandettini. Representational similarity analysis–connecting the branches of systems neuroscience. *Frontiers in systems neuroscience*, 2, 2008. 1, 8, 10
- [35] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012. 4, 7
- [36] J. Liu, A. Harris, and N. Kanwisher. Stages of processing in face perception: an meg study. *Nature neuroscience*, 5(9):910–916, 2002. 10, 12
- [37] S. P. MacEvoy and R. A. Epstein. Constructing scenes from objects in human occipitotemporal cortex. *Nature neuro-science*, 14(10):1323–1329, 2011.
- [38] E. Maris and R. Oostenveld. Nonparametric statistical testing of eeg-and meg-data. *Journal of neuroscience methods*, 164(1):177–190, 2007. 6
- [39] F. Mormann, S. Kornblith, R. Q. Quiroga, A. Kraskov, M. Cerf, I. Fried, and C. Koch. Latency and selectivity of single neurons indicate hierarchical processing in the human medial temporal lobe. *The Journal of Neuroscience*, 28(36):8865–8872, 2008. 10
- [40] E. I. Moser, E. Kropff, and M.-B. Moser. Place cells, grid cells, and the brain's spatial representation system. *Annu. Rev. Neurosci.*, 31:69–89, 2008. 1
- [41] K.-R. Müller, S. Mika, G. Rätsch, K. Tsuda, and B. Schölkopf. An introduction to kernel-based learning algorithms. *Neural Networks, IEEE Transactions on*, 12(2):181– 201, 2001. 2
- [42] T. E. Nichols and A. P. Holmes. Nonparametric permutation tests for functional neuroimaging: a primer with examples. *Human brain mapping*, 15(1):1–25, 2002. 5, 6
- [43] H. Nili, C. Wingfield, A. Walther, L. Su, W. Marslen-Wilson, and N. Kriegeskorte. A toolbox for representational similarity analysis. *PLoS Comput. Biol*, 10(4):e1003553, 2014. 6
- [44] A. Oliva and A. Torralba. Modeling the shape of the scene: A holistic representation of the spatial envelope. *International journal of computer vision*, 42(3):145–175, 2001. 1, 3, 4, 8
- [45] D. Pantazis, T. E. Nichols, S. Baillet, and R. M. Leahy. A comparison of random field theory and permutation methods for the statistical analysis of meg data. *Neuroimage*, 25(2):383–394, 2005. 5, 6
- [46] S. Park, T. F. Brady, M. R. Greene, and A. Oliva. Disentangling scene content from spatial boundary: complementary roles for the parahippocampal place area and lateral occipital complex in representing real-world scenes. *The Journal* of Neuroscience, 31(4):1333–1340, 2011. 1
- [47] S. Park, T. Konkle, and A. Oliva. Parametric coding of the size and clutter of natural scenes in the human brain. *Cerebral Cortex*, page bht418, 2014. 1, 2, 3, 5, 7, 10
- [48] T. Poggio and E. Bizzi. Generalization in vision and motor control. *Nature*, 431(7010):768–774, 2004. 10

- [49] L. Reddy and N. Kanwisher. Coding of visual objects in the ventral stream. *Current opinion in neurobiology*, 16(4):408– 414, 2006. 12
- [50] G. E. Rice, D. M. Watson, T. Hartley, and T. J. Andrews. Low-level image properties of visual objects predict patterns of neural response across category-selective regions of the ventral visual pathway. *The Journal of Neuroscience*, 34(26):8837–8844, 2014. 11
- [51] M. Riesenhuber and T. Poggio. Hierarchical models of object recognition in cortex. *Nature neuroscience*, 2(11):1019– 1025, 1999. 1
- [52] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 2014. 4
- [53] T. Serre, A. Oliva, and T. Poggio. A feedforward architecture accounts for rapid categorization. *Proceedings of the National Academy of Sciences*, 104(15):6424–6429, 2007. 8, 12
- [54] T. Serre, L. Wolf, and T. Poggio. Object recognition with features inspired by visual cortex. In *Computer Vision and Pattern Recognition*, 2005. CVPR 2005. IEEE Computer Society Conference on, volume 2, pages 994–1000. IEEE, 2005. 4, 5
- [55] J. J. Stekelenburg and B. de Gelder. The neural correlates of perceiving human bodies: an erp study on the body-inversion effect. *Neuroreport*, 15(5):777–780, 2004. 10
- [56] F. Tadel, S. Baillet, J. C. Mosher, D. Pantazis, and R. M. Leahy. Brainstorm: a user-friendly application for meg/eeg analysis. *Computational intelligence and neuroscience*, 2011:8, 2011. 2
- [57] G. Thierry, A. J. Pegna, C. Dodds, M. Roberts, S. Basan, and P. Downing. An event-related potential component sensitive to images of the human body. *Neuroimage*, 32(2):871–879, 2006. 10
- [58] S. Thorpe, D. Fize, C. Marlot, et al. Speed of processing in the human visual system. *nature*, 381(6582):520–522, 1996.
 7, 12
- [59] H. Tiitinen, N. H. Salminen, K. J. Palomäki, V. T. Mäkinen, P. Alku, and P. J. May. Neuromagnetic recordings reveal the temporal dynamics of auditory spatial processing in the human cortex. *Neuroscience letters*, 396(1):17–22, 2006. 10
- [60] D. Y. Tsao, W. A. Freiwald, R. B. Tootell, and M. S. Livingstone. A cortical region consisting entirely of face-selective cells. *Science*, 311(5761):670–674, 2006. 10
- [61] S. Vaziri, E. T. Carlson, Z. Wang, and C. E. Connor. A channel for 3d environmental shape in anterior inferotemporal cortex. *Neuron*, 84(1):55–62, 2014.
- [62] D. M. Watson, T. Hartley, and T. J. Andrews. Patterns of response to visual scenes are linked to the low-level properties of the image. *Neuroimage*, 99:402–410, 2014. 11
- [63] T. Wolbers, R. L. Klatzky, J. M. Loomis, M. G. Wutte, and N. A. Giudice. Modality-independent coding of spatial layout in the human brain. *Current Biology*, 21(11):984–989, 2011. 1, 10
- [64] D. L. Yamins, H. Hong, C. F. Cadieu, E. A. Solomon, D. Seibert, and J. J. DiCarlo. Performance-optimized hierarchical models predict neural responses in higher visual

cortex. *Proceedings of the National Academy of Sciences*, 111(23):8619–8624, 2014. 11

- [65] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba. Object detectors emerge in deep scene cnns. In *International Conference on Learning Representations (ICLR)*, 2015. 1, 4, 8
- [66] B. Zhou, A. Lapedriza, J. Xiao, A. Torralba, and A. Oliva. Learning deep features for scene recognition using places database. In Advances in Neural Information Processing Systems, pages 487–495, 2014. 1, 4, 8