

MIT Open Access Articles

Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Cichy, Radoslaw Martin, Aditya Khosla, Dimitrios Pantazis, Antonio Torralba, and Aude Oliva. "Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence." *Scientific Reports* 6, Article number:27755 (2016), p.1-12.

As Published: <http://dx.doi.org/10.1038/srep27755>

Publisher: Springer Nature

Persistent URL: <http://hdl.handle.net/1721.1/103585>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of use: Creative Commons Attribution 4.0 International License



SCIENTIFIC REPORTS



OPEN

Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence

Received: 26 January 2016

Accepted: 23 May 2016

Published: 10 June 2016

Radoslaw Martin Cichy^{1,2}, Aditya Khosla¹, Dimitrios Pantazis³, Antonio Torralba¹ & Aude Oliva¹

The complex multi-stage architecture of cortical visual pathways provides the neural basis for efficient visual object recognition in humans. However, the stage-wise computations therein remain poorly understood. Here, we compared temporal (magnetoencephalography) and spatial (functional MRI) visual brain representations with representations in an artificial deep neural network (DNN) tuned to the statistics of real-world visual recognition. We showed that the DNN captured the stages of human visual processing in both time and space from early visual areas towards the dorsal and ventral streams. Further investigation of crucial DNN parameters revealed that while model architecture was important, training on real-world categorization was necessary to enforce spatio-temporal hierarchical relationships with the brain. Together our results provide an algorithmically informed view on the spatio-temporal dynamics of visual object recognition in the human visual brain.

Visual object recognition in humans is mediated by complex multi-stage processing of visual information emerging rapidly in a distributed network of cortical regions^{1–7}. Understanding visual object recognition in cortex thus requires a quantitative model that captures the complexity of the underlying spatio-temporal dynamics^{8–10}.

A major impediment in creating such a model is the highly nonlinear and sparse nature of neural tuning properties in mid- and high-level visual areas^{11–13} that is difficult to capture experimentally, and thus unknown. Previous approaches to modeling object recognition in cortex relied on extrapolation of principles from well understood lower visual areas such as V1^{8,9} and strong manual intervention, achieving only modest task performance compared to humans.

Here we take an alternative route, constructing and comparing against brain signals a visual computational model based on deep neural networks (DNNs)^{14–16}, i.e. computer vision models in which model neuron tuning properties are set by supervised learning without manual intervention^{14,17}. DNNs are the best performing models on computer vision object recognition benchmarks and yield human performance levels on object categorization^{18,19}. We used a tripartite strategy to reveal the spatio-temporal processing cascade underlying human visual object recognition by DNN model comparisons.

First, as object recognition is a process rapidly unfolding over time^{3,20–22}, we compared DNN visual representations to millisecond resolved magnetoencephalography (MEG) brain data. Our results delineate an ordered relationship between the stages of processing in a DNN and the time course with which object representations emerge in the human brain²³.

Second, as object recognition recruits a multitude of distributed brain regions, a full account of object recognition needs to go beyond the analysis of a few pre-defined brain regions^{24–28}, determining the relationship between DNNs and the whole brain. Using a spatially unbiased approach, we revealed a hierarchical relationship between DNNs and the processing cascade of both the ventral and dorsal visual pathway.

¹Computer Science and Artificial Intelligence Laboratory, MIT, Cambridge, MA, USA. ²Department of Education and Psychology, Free University Berlin, Berlin, Germany. ³McGovern Institute for Brain Research, MIT, Cambridge, MA, USA. Correspondence and requests for materials should be addressed to R.M.C. (email: rmcichy@gmail.com)




Figure 1. Deep neural network architecture and properties. (a) The DNN architecture comprised 8 layers. Each of layers 1–5 contained a combination of convolution, max-pooling and normalization stages, whereas the last three layers were fully connected. The DNN takes pixel values as inputs and propagates information feed-forward through the layers, activating model neurons with particular activation values successively at each layer. (b) Visualization of example DNN connections. The thickness of highlighted lines (colored to ease visualization) indicates the weight of the strongest connections going in and out of neurons, starting from a sample neuron in layer 1. Neurons in layer 1 are represented by their filters, and in layers 2–5 by gray dots. For combined visualization of connections between neurons and neuron RF selectivity please visit <http://brainmodels.csail.mit.edu/dnn/drawCNN/>.

Third, interpretation of a DNN-brain comparison depends on the factors shaping the DNN fundamentally: the pre-specified model *architecture*, the *training procedure*, and the *learned task* (e.g. object categorization). By comparing different DNN models to brain data, we demonstrated the influence of each of these factors on the emergence of similarity relations between DNNs and brains in both space and time.

Together, our results provide an algorithmically informed perspective of the spatio-temporal dynamics underlying visual object recognition in the human brain.

Results

Construction of a deep neural network performing at human level in object categorization. To be a plausible model of object recognition in cortex, a computational model must provide high performance on visual object categorization. Latest generations of computer vision models, termed deep neural networks (DNNs), have achieved extraordinary performance, thus raising the question whether their algorithmic representations bear resemblance of the neural computations underlying human vision. To investigate we trained an 8-layer DNN architecture (Fig. 1a) that corresponds to the best-performing model in object classification in the ImageNet Large Scale Visual Recognition Challenge²⁹. Each DNN layer performs simple operations that are implementable in biological circuits, such as convolution, pooling and normalization. We trained the DNN to perform object categorization on everyday object categories (683 categories, with ~1300 images in each category) using back propagation, i.e. the network learned neuronal tuning functions by itself. We termed this neural network *object deep neural network* (*object DNN*). The object DNN performed equally well on object categorization as previous implementations¹⁶ (Suppl. Table 1).

We investigated the coding of visual information in the object DNN by determining the selectivity and size of receptive fields (RF) of the model neurons using a neuroscience-inspired reduction method³⁰. We found that neurons in layer 1 had Gabor filter or color patch-like sensitivity (Fig. 1b, layer 1), while those in increasingly later layers had increasingly larger RFs and sensitivity to complex forms (due to copyright shown without RF visualization; for full visualization of RFs for all neurons in layers 1–5 see <http://brainmodels.csail.mit.edu/dnn/rf/>). This showed that the object DNN had representations in a hierarchy of increasing complexity, akin to representations




Figure 2. Comparison of MEG, fMRI and DNN representations by representational similarity. In each signal space (fMRI, MEG, DNN) we summarized representational structure by calculating the dissimilarity between activation patterns of different pairs of conditions (here exemplified for two objects: bus and orange). This yielded representational dissimilarity matrices (RDMs) indexed in rows and columns by the compared conditions. We calculated millisecond resolved MEG RDMs from -100 ms to $+1,000$ ms with respect to image onset, layer-specific DNN RDMs (layers 1 through 8) and voxel-specific fMRI RDMs in a spatially unbiased cortical surface-based searchlight procedure. RDMs were directly comparable (Spearman's R), facilitating integration across signal spaces. Comparison of DNN with MEG RDMs yielded time courses of similarity between emerging visual representations in the brain and DNN. Comparison of the DNN with fMRI RDMs yielded spatial maps of visual representations common to the human brain and the DNN. Object images shown as exemplars are not examples of the original stimulus set due to copyright; the complete stimulus set is visualized at http://brainmodels.csail.mit.edu/images/stimulus_set.png.

in the primate visual brain hierarchy^{5,7}. To visualize neuronal connectivity (as illustrated by colored lines in Fig. 1b for an example neuron starting in layer 1) and RF selectivity of all neurons in layers 1–5 together we developed an online tool (for full visualization see <http://brainmodels.csail.mit.edu/dnn/drawCNN/>).

Representational similarity analysis was used as the integrative framework for DNN-brain comparison. To compare representations in the object DNN and human brains, we used a 118-image set of natural objects on real-world backgrounds (for visualization please see http://brainmodels.csail.mit.edu/images/stimulus_set.png). Note that these 118 images were not used for training the object DNN to avoid circular inference. We determined how well the network performed on our image set by inspecting its 1 and 5 most confident classification labels for each image (i.e., top-one and top-five classification accuracy). Top-five classification accuracy is a popular and particularly suited success measure for real world image classification, as natural images are likely to contain more than just one object at a time, rendering a top-one classification approach too restrictive. Voting on each of the 118 images is available at <http://brainmodels.csail.mit.edu>. The network classified on the top-one classification task 103/118 (87% success rate) and on the top-five classification task 111/118 images correctly (94% success rate), and thus at a level comparable to humans¹⁸.

We also recorded fMRI and MEG in 15 participants viewing random sequences of the same 118 real-world object image set while conducting an orthogonal task. The experimental design was adapted to the specifics of the measurement technique (Suppl. Fig. 1).

We compared fMRI and MEG brain measurements with the DNN in a common analysis framework with representational similarity analysis³¹ (Fig. 2). The basic idea is that if two images are similarly represented in the brain, they should also be similarly represented in the DNN. To quantify, we first obtained signal measurements in temporally specific MEG sensor activation patterns (1 ms steps from -100 to $+1,000$ ms), in spatially specific fMRI voxel patterns, and in layer-specific model neuron activations of the DNN. To make the different signal spaces (fMRI, MEG, DNN) comparable, we abstracted signals to a similarity space. In detail, for each signal space we computed dissimilarities (1-Spearman's R for DNN and fMRI, percent decoding accuracy in pair-wise classification for MEG) between every pair of conditions (images), as exemplified by images 1 and 2 in Fig. 2. This




Figure 3. Representations in the object DNN correlated with emerging visual representations in the human brain in an ordered fashion. (a) Time courses with which representational similarity in the brain and layers of the deep object network emerged. Color-coded lines above data curves indicate significant time points ($n = 15$, cluster definition threshold $P = 0.05$, cluster threshold $P = 0.05$ Bonferroni-corrected for 8 layers; for onset and peak latencies see Suppl. Table 2). Gray vertical line indicates image onset. (b) Overall peak latency of time courses increased with layer number ($n = 15$, $R = 0.35$, $P = 0.0007$, sign permutation test). Error bars indicate standard error of the mean determined by 10,000 bootstrap samples of the participant pool.

yielded 118×118 representational dissimilarity matrices (RDMs) indexed in rows and columns by the compared conditions. These RDMs were time-resolved for MEG, space-resolved for fMRI, and layer-resolved in DNN. Comparing DNN RDMs with MEG RDMs resulted in time courses highlighting how DNN representations correlated with emerging visual representations. Comparing DNN RDMs with fMRI RDMs resulted in spatial maps indicative of how representations in the object DNN correlated with brain activity.

Representations in the object DNN correlated with emerging visual representations in the human brain. Visual information processing in the brain is a process that rapidly evolves over time^{3,20,21}, and a model of object recognition in cortex should mirror this temporal evolution. While the DNN used here does not model time, it has a clear sequential structure: information flows from one layer to the next in strict order. We thus investigated whether representations in the object DNN correlated with emerging visual representations in the first few hundred milliseconds of vision in sequential order. For this we determined representational similarity between layer-specific DNN representations and MEG data in millisecond steps from -100 to $+1000$ ms with respect to image onset and layer-specific DNN representations. We found that all layers of the object DNN were representationally similar to human brain activity, indicating that the model captures emerging brain visual representations (Fig. 3a, $P < 0.05$ cluster definition threshold, $P < 0.05$ cluster threshold, lines above data curves color-coded same as those indicate significant time points, for details see Suppl. Table 2). We next investigated whether the hierarchy of the layered architecture of the object DNN, as characterized by an increasing size and complexity of model RFs feature selectivity, corresponded to the hierarchy of temporal processing in the brain. That is, we examined whether there was a hierarchical relationship between layer number of the object DNN and the peak latency of the correlation time courses between object DNN and MEG RDMs. Subject-average data are plotted in Fig. 3b. While for the first four DNN layers the relationship was weak and negative, for the last four layers the relationship was strongly positive. When quantified over all layers, the relationship was modestly positive and significant ($n = 15$, Spearman's $R = 0.35$, $P = 0.0007$).

Together these analyses established a correspondence in the sequence of processing steps of a DNN trained for object recognition and the time course with which visual representations emerge in the human brain²³.

Correlation of representations in the visual brain with the object DNN revealed the hierarchical topography in the human ventral and dorsal visual streams. To localize visual representations common to brain and the object DNN, we used a spatially unbiased surface-based searchlight approach. Comparison of representational similarities between fMRI data and object DNN RDMs yielded 8 layer-specific spatial maps identifying the cortical regions where representations in the object DNN correlated with brain activity (Fig. 4, cluster definition threshold $P = 0.05$, cluster-threshold $P = 0.05$ Bonferroni-corrected for multiple comparisons by 16 (8 layers * 2 hemispheres); different viewing angles available in Suppl. Movie 1).

The results indicate a hierarchical correspondence between model network layers and the human visual system. For low DNN layers, similarities of visual representations were confined to the occipital lobe, i.e. low- and mid-level visual regions, and for high DNN layers in more anterior regions in both the ventral and dorsal visual stream. A supplementary volumetric searchlight analysis (Suppl. Text 1, Suppl. Fig. 2; using a false discovery rate correction allowing voxel-wise inference reproduced these findings, yielding corroborative evidence across analysis methodologies. Further, as layer-specific RDMs are correlated with each other (Suppl. Fig. 3a), we conducted a supplementary analysis determining the correlation between layer unique components and brains using partial correlation analysis (Suppl. Fig. 3b). This again yielded evidence for a hierarchical correspondence between model network layers and the human visual system.




Figure 4. Spatial maps of visual representations common to brain and object DNN. There was a correspondence between object DNN hierarchy and the hierarchical topography of visual representations in the human brain. Low layers had significant representational similarities confined to the occipital lobe of the brain, i.e. low- and mid-level visual regions. Higher layers had significant representational similarities with more anterior regions in the temporal and parietal lobe, with layers 7 and 8 reaching far into the inferior temporal cortex and inferior parietal cortex ($n = 15$, cluster definition threshold $P < 0.05$, cluster-threshold $P < 0.05$ Bonferroni-corrected for multiple comparisons by 16 (8 DNN layers * 2 hemispheres)).

These results suggest that hierarchical systems of visual representations emerge in both the human ventral and dorsal visual stream as the result of task constraints of object categorization posed in everyday life, and provide strong evidence for object representations in the dorsal stream independent of attention or motor intention.

Factors determining the correlation between representations in DNN's and cortical visual representations emerging in time. The observation of a positive and hierarchical relationship between the object DNN and brain temporal dynamics poses the fundamental question of the origin of this relationship. Three




Figure 5. Architecture, task, and training procedure influence the correlation between representations in DNNs and temporally emerging brain representations. (a) We created 5 different models: 1) a model trained on object categorization (*object DNN*; Fig. 1); 2) an untrained model initialized with random weights (*untrained DNN*) to determine the effect of architecture alone; 3) a model trained on a different real-world task, scene categorization (*scene DNN*) to investigate the effect of task; and 4,5) a model trained on object categorization with random assignment of image labels (*unecological DNN*), or spatially smoothed noisy images with random assignment of image labels (*noise DNN*), to determine the effect of the training operation independent of task constraints. (b) All DNNs had significant representational similarities to human brains (layer-specific analysis in Suppl. Fig. 4). (c) We contrasted the object DNN against all other models (subtraction of corresponding time series shown in (b)). Representations in the object DNN were more similar to brain representations than any other model except the scene DNN. Lines above data curves significant time points ($n = 15$, cluster definition threshold $P = 0.05$, cluster threshold $P = 0.05$ Bonferroni corrected by 5 (number of models) in (b), and 4 (number of comparisons in (c)); for onset and peak latencies see Suppl. Table 3a,b). Gray vertical lines indicates image onset.

fundamental factors shape DNNs: *architecture, task, and training procedure*. Determining the effect of each is crucial to understanding the emergence of the brain-DNN relationships on the real-world object categorization task. To this goal, we created several different DNN models (Fig. 5a). We reasoned that a comparison of brain with (1) an untrained DNN, i.e. a DNN with the same architecture as the DNN, but random (unlearned) connection weights would reveal the effect of DNN architecture alone, (2) a DNN trained on an alternate categorization task, scene categorization, would reveal the effect of specific task, and (3) a DNN trained on an image set with random unecological assignment of images to category labels, or a DNN trained on noise images, would reveal the effect of the training procedure per se.

To evaluate the hierarchy of temporal and spatial relationships between the human brain and DNNs, we computed layer-specific RDMs for each DNN. To allow direct comparisons across models, we also computed a single summary RDM for each DNN model based on concatenated layer-specific activation vectors.

Concerning the role of architecture, we found that representations in the untrained DNN significantly correlated with emerging brain representations (Fig. 5b). The untrained DNN correlated significantly worse than the object DNN after 166 ms (95% confidence interval 128–198 ms), suggesting late processing stages (Fig. 5c, Suppl. Table 3b). A supplementary layer-specific analysis identified every layer as a significant contributor to this correlation (Suppl. Fig. 4a). Even though the relationship between layer number and the peak latency of brain-DNN similarity time series was hierarchical, it was negative ($R = -0.6$, $P = 0.0003$, Suppl. Fig. 4b) and thus reversed and statistically different from the object DNN ($\Delta R = 0.96$, $P = 0.0003$). This shows that DNN architecture alone, independent of task constraints or training procedures, induces representational similarity to emerging visual representations in the brain, but that constraints imposed by training on a real-world categorization task significantly increases this effect after 160 ms and reverses the direction of the hierarchical relationship.




Figure 6. Architecture, task constraints, and training procedure influence the topographically ordered correlation in representations between DNNs and human brain. (a) Comparison of fMRI representations in V1, IT and IPS1&2 with the layer-specific DNN representations of each model. Error bars indicate standard error of the mean as determined by bootstrapping ($n = 15$). (b) Correlations between layer number and brain-DNN representational similarities for the different models shown in (a). Non-zero correlations indicate hierarchical relationships; positive correlations indicate an increase in brain-DNN similarities towards higher layers, and vice versa for negative correlations. Bars color-coded as DNNs, stars above bars indicate significance (sign-permutation tests, $P < 0.05$, FDR-corrected, for details see Suppl. Table 4a). (c) Comparison of object DNN against all other models (subtraction of corresponding points shown in a). (d) Same as (b), but for the curves shown in (c) (for details see Suppl. Table 4b).

Concerning the role of task, we found that representations in the scene DNN also correlated with emerging brain representations. In the direct comparison between the scene and the object DNN the object DNN performed better with a cluster threshold p -value of 0.04, thus not surviving Bonferroni correction (Fig. 5b,c, Suppl. Fig. 4c). This suggests that task constraints influence the model and possibly also brain in a partly overlapping, and partly dissociable manner. Further, the relationship between layer number and brain-DNN similarity time series was positively hierarchical for the scene DNN ($R = 0.44$, $P = 0.001$, Suppl. Fig. 4d), and not different from the object DNN ($\Delta R = -0.09$, $P = 0.41$), further suggesting overlapping neural mechanisms for object and scene perception.

Concerning the role of the training operation, we found that representations in both the unecological and noise DNNs correlated with brain representations (Fig. 5b, Suppl. Fig. 4e,g), but worse than the object DNN (Fig. 5c). Further, there was no evidence for a hierarchical relationship between layer number and brain-DNN similarity time series for either DNN (unecological DNN: $R = -0.01$; $P = 0.94$; noise DNN: $R = -0.04$, $P = 0.68$; Suppl. Fig. 4f,h), and both had a weaker hierarchical relationship than the object DNN (unecological DNN: $\Delta R = 0.39$, $P = 0.0107$; noise DNN: $\Delta R = 0.36$, $P = 0.0052$). Thus the training operation per se has an effect on the relationship to the brain, but only training on real-world categorization increases brain-DNN similarity and hierarchy.

In summary, we found that although architecture alone was enough for a correlation between DNNs and temporally emerging visual representations, training on real-world categorization was necessary for a hierarchical relationship to emerge. Thus, both architecture and training crucially influence the similarity between DNNs and brains over the first few hundred milliseconds of vision.

Factors determining the topographically ordered correlation between representations in DNN and cortical visual representations in cortex.

The observation of a positive and hierarchical relationship between the object DNN structure and the brain visual pathways motivates an inquiry, akin to the temporal dynamics analysis in the previous section, regarding the role of architecture, task demands and training operation. For this we systematically investigated three regions-of-interest (ROIs): the early visual area V1, and two regions up-stream in the ventral and dorsal stream, the inferior temporal cortex IT and a region encompassing intraparietal sulcus 1 and 2 (IPS1&2), respectively. We examined whether representations in DNN correlated with brain representations in these ROIs (Fig. 6a), and also whether this correlation was hierarchical (Fig. 6, Suppl. Table 4a).

Concerning the role of architecture, we found that representations in the untrained DNN correlated with brain representations better than representations in the object DNN in V1, but worse in IT and IPS1&2 (Fig. 6a,c). Further, the relationship was hierarchical (negative) only in IT ($R = -0.47$, $P = 0.002$) (Fig. 6b; stars above bars). Thus depending on cortical region the DNN architecture alone is enough to induce similarity between a DNN and the brain, but the hierarchy absent (V1, IPS1&2) or reversed (IT) without proper DNN training.

Concerning the role of task, we found the scene DNN had largely similar, albeit weaker, similarity to the brain than the object DNN for all ROIs (Fig. 6a,c), with a significant hierarchical relationship in V1 ($R = -0.68$, $P = 0.002$), but not in IT ($R = 0.26$, $P = 0.155$) or IPS1&2 ($R = 0.30$, $P = 0.08$) (Fig. 6b). In addition, comparing results for the object and scene DNNs directly (Fig. 6c), we found stronger effects for the object DNN in several

layers in all ROIs. Together these results corroborate the conclusions of the MEG analysis, showing that task constraints shape brain representations along both ventral visual streams in a partly overlapping, and partly dissociable manner.

Concerning the role of the training operation, we found that representations in both the unecological and noise DNNs correlated with cortical representations in V1 and IT, but not IPS1&2 (Fig. 6a), and with smaller correlation than the object DNN in all regions (Fig. 6c). A hierarchical relationship was present and negative in V1 and IT, but not IPS1&2 (Fig. 6b, unecological DNN: V1 $R = -0.40$, $P = 0.001$, IT $R = -0.38$, $P = 0.001$, IPS1&2 $R = -0.03$, $P = 0.77$; noise DNN: V1 $R = -0.08$, $P = 0.42$, IT $R = -0.29$, $P = 0.012$, IPS1&2 $R = -0.08$, $P = 0.42$).

Therefore the training on a real-world categorization task, but not the training operation per se, increases the brain-DNN similarity while inducing a hierarchical relationship.

Discussion

Summary. By comparing the spatio-temporal dynamics in the human brain with a deep neural network (DNN) model trained on object categorization, we provided a formal model of object recognition in cortex. We found a correspondence between the object DNN and the brain in both space (fMRI data) and time (MEG data). Both cases demonstrated a hierarchy: in space from low- to high-level visual areas in both ventral and dorsal stream, in time over the visual processing stages in the first few hundred milliseconds of vision. A systematic analysis of the fundamental determinants of this DNN-brain relationship identified that the architecture alone induces similarity, but that training on a real-world categorization task was necessary for a hierarchical relationship to emerge. Our results demonstrate the explanatory and discovery power of the brain-DNN comparison approach to understand the spatio-temporal neural dynamics underlying object recognition. They provide novel evidence for a role of parietal cortex in visual object categorization, and give rise to the idea that the organization of the visual cortex may be influenced by processing constraints imposed by visual categorization the same way that DNN representations were influenced by object categorization tasks.

Representations in the object DNN correlate with brain representations in space and time in a hierarchical fashion. A major impediment in modeling human object recognition in cortex is the lack of principled understanding of exact neuronal tuning in mid- and high-level visual cortex. Previous approaches thus extrapolated principles observed in low-level visual cortex, with limited success in capturing neuronal variability and a much inferior to human behavioral performance^{8,9}.

Our approach allowed us to obviate this limitation by relying on an object recognition model that learns neuronal tuning. By comparing representations between the DNN and the human brain we found a hierarchical correspondence in both space and time: with increasing DNN layer number DNN representations correlated more with cortical representations emerging later in time, and in increasingly higher brain areas in both the dorsal and ventral visual pathway. Our results provide algorithmically informed evidence for the idea of visual processing as a step-wise hierarchical process in time^{3,20,32,33} and along a system of cortical regions^{2,7,34}.

Regarding the temporal correspondence, our results provide evidence for a hierarchical relationship between computer models of vision and the brain. This corroborates and extends previous MEG research showing an ordered correspondence between brain activity and two layers in the HMAX model²³ over time and in source space to 8-layer DNNs. Peak latencies between layers of the object DNN and emerging brain activations ranged between approximately 100 and 160 ms. While in agreement with prior findings about the time necessary for complex object processing³⁵, our results go further by making explicit the step-wise transformations of representational format that may underlie rapid complex object categorization behavior.

In regards to the spatial correspondence, previous studies compared DNNs to the ventral visual stream only, mostly using a spatially limited region-of-interest approach^{26–28}. Here, using a spatially unbiased whole-brain approach³⁶, we discovered a hierarchical correspondence in the dorsal visual pathway. While previous studies have documented object selective responses in dorsal stream in monkeys^{37,38} and humans^{39,40}, it is still debated whether dorsal visual representations are better explained by differential motor action associations or ability to engage attention, rather than category membership or shape representation^{41,42}. Crucially, our results defy explanation by attention or motor-related concepts, as neither played any role in the DNN and thus brain-DNN correspondence. This argues for a stronger role in object recognition than previously appreciated.

Our results thus challenge the classic descriptions of the dorsal pathway as a spatially- or action oriented ‘where’ or ‘how’ pathway^{1,4}, and suggest that current theories describing parietal cortex as related to spatial working memory, visually guided actions and spatial navigation⁶ should be complemented with a role for the dorsal visual stream in object categorization⁴⁰.

Origin and implications of brain-DNN representation similarities. Investigating the influence of crucial parameters determining DNNs, we found an influence of both architecture and task constraints induced by training the DNN on a real-world categorization task. This suggests that similar architectural principles, i.e. convolution, max pooling and normalization govern both model and brains, concurrent with the origin of those principle by observation in the brain⁸. The stronger similarity with early rather than late brain regions might be explained by the fact that neural networks initialized with random weights that involve a convolution, nonlinearity and normalization stage exhibit Gabor-like filters sensitive to oriented edges, and thus similar properties as neurons in early visual areas⁴³.

Although architecture alone induced similarity, training on a real-world categorization task increased similarity – in particular with higher brain regions and processing after ~160 ms – and was necessary for a hierarchical relationship in processing stages between the brain and the DNN to emerge in space and time. This demonstrates that learning constraints imposed by a real-world categorization task crucially shape the representational space of a DNN²⁸, and suggests that the processing hierarchy in the human brain is a the result of computational

constraints imposed by visual object categorization. Such constraints may originate in high-level visual regions such as IT and IPS, be propagated backwards from high-level visual regions through the visual hierarchies through abundantly present feedback connections in the visual stream at all levels⁴⁴ during visual learning⁴⁵, and provide the basis of learning at all stages of the processing in visual brain⁴⁶.

Summary statement. In sum, by comparing deep neural networks to human brains in space and time, we provide a spatio-temporally unbiased algorithmic account of visual object recognition in human cortex.

Methods

Participants. 15 healthy human volunteers (5 female, age: mean \pm s.d. = 26.6 \pm 5.18 years, recruited from a subject pool at Massachusetts Institute of Technology) participated in the experiment. The sample size was based on methodological recommendations in literature for random-effects fMRI and MEG analyses. Written informed consent was obtained from all subjects. The study was approved by the local ethics committee (Institutional Review Board of the Massachusetts Institute of Technology) and conducted according to the principles of the declaration of Helsinki. All methods were carried out in accordance with the approved guidelines.

Visual stimuli. The stimuli presented to humans and computer vision models were 118 color photographs of everyday objects, each from a different category, on natural backgrounds from the ImageNet image database⁴⁷. For visualization please visit http://brainmodels.csail.mit.edu/images/stimulus_set.png.

Experimental design and task. Participants viewed images presented at the center of the screen (4° visual angle) for 0.5 s and overlaid with a light gray fixation cross. The presentation parameters were adapted to the specific requirements of each acquisition technique (Suppl. Fig. 1).

For MEG, participants completed 15 runs of 314 s duration. Each image was presented twice in each MEG run in random order with an inter-trial interval (ITI) of 0.9–1 s. Participants were asked to press a button and blink their eyes in response to a paper clip image shown randomly every 3 to 5 trials (average 4). The paper clip image was not part of the image set, and paper clip trials were excluded from further analysis.

For fMRI, each participant completed two independent sessions of 9–11 runs (486 s duration each) on two separate days. Each run consisted of one presentation of each image in random order, interspersed randomly with 39 null trials (i.e. 25% of all trials) with no stimulus presentation. During the null trials the fixation cross turned darker for 500 ms. Participants reported changes in fixation cross hue with a button press.

MEG acquisition. MEG signals were acquired continuously for the whole session from 306 channels (204 planar gradiometers, 102 magnetometers, Elektra Neuromag TRIUX, Elekta, Stockholm) at a sampling rate of 1,000 Hz, and filtered online between 0.03 and 330 Hz. Head movement parameters were acquired using continuous HPI measurement. To compensate for head movements and to perform noise reduction with spatiotemporal filters^{48,49} we pre-processed data using Maxfilter software (Elekta, Stockholm). We used default parameters (harmonic expansion origin in head frame = [0 0 40] mm; expansion limit for internal multipole base = 8; expansion limit for external multipole base = 3; bad channels automatically excluded from harmonic expansions = 7 s.d. above average; temporal correlation limit = 0.98; buffer length = 10 s). Note that no bad channels were detected. The resulting filtered time series were analyzed with Brainstorm (<http://neuroimage.usc.edu/brainstorm/>). We extracted each trial with a 100 ms baseline and 1,000 ms post-stimulus recordings, removed baseline mean, smoothed data with a 20 ms sliding window, and normalized each channel with its baseline standard deviation. This yielded 30 preprocessed trials per condition and participant.

Note that there was a systematic delay of 26 ms between the stimulus computer and the projector response (as determined previously by photodiode measurements). We accounted for this delay during data acquisition and therefore all reported times are exact.

fMRI acquisition. Magnetic resonance imaging (MRI) was conducted on a 3T Trio scanner (Siemens, Erlangen, Germany) with a 32-channel head coil. We acquired structural images using a standard T1-weighted sequence (192 sagittal slices, FOV = 256 mm², TR = 1,900 ms, TE = 2.52 ms, flip angle = 9°).

For fMRI, we conducted 9–11 runs in which 648 volumes were acquired for each participant (gradient-echo EPI sequence: TR = 750 ms, TE = 30 ms, flip angle = 61°, FOV read = 192 mm, FOV phase = 100% with a partial fraction of 6/8, through-plane acceleration factor 3, bandwidth 1816 Hz/Px, resolution = 3 mm³, slice gap 20%, slices = 33, ascending acquisition). The acquisition volume covered the whole cortex.

Anatomical MRI analysis. We reconstructed the cortical surface of each participant using Freesurfer on the basis of the T1 structural scan⁵⁰. This yielded a discrete triangular mesh representing the cortical surface used for the surface-based two-dimensional (2D) searchlight procedure outlined below.

fMRI analysis. We preprocessed fMRI data using SPM8 (<http://www.fil.ion.ucl.ac.uk/spm/>). For each participant and session separately, fMRI data were realigned and co-registered to the T1 structural scan acquired in the first MRI session. Data was neither normalized nor smoothed. We estimated the fMRI response to the 118 image conditions with a general linear model. Image onsets and duration were entered into the GLM as regressors and convolved with a hemodynamic response function. Movement parameters entered the GLM as nuisance regressors. We then converted each of the 118 estimated GLM parameters into t-values by contrasting each condition estimate against the implicitly modeled baseline. Additionally, we determined the grand-average effect of visual stimulation independent of condition in a separate t-contrast of parameter estimates for all 118 image conditions versus the implicit baseline.

Definition of fMRI regions of interest. We defined three regions-of-interest for each participant: V1 corresponding to the central 4° of the visual field, inferior temporal cortex (IT), and intraparietal sulcus regions 1 and 2 combined (IPS1&2). We defined the V1 ROI based on an anatomical eccentricity template⁵¹. For this, we registered a generic V1 eccentricity template to reconstructed participant-specific cortical surfaces and restricted the template to the central 4° of visual angle. The surface-based ROIs for the left and right hemisphere were resampled to the space of EPI volumes and combined.

To define inferior temporal cortex (IT), we used an anatomical mask of bilateral fusiform and inferior temporal cortex (WFU Pickatlas, IBASPM116 Atlas). To define IPS1&2, we used a combined probabilistic mask of IPS1 and IPS2⁵². Masks in MNI space were reverse-normalized to single-subject functional space. We then restricted the anatomical definition of each ROI for each participant by functional criteria to the 100 most strongly activated voxels in the grand-average contrast of visual stimulation vs. baseline.

fMRI surface-based searchlight construction and analysis. To analyze fMRI data in a spatially unbiased (unrestricted from ROIs) approach, we performed a 2D surface-based searchlight analysis following the approach of Chen *et al.*⁵³. We used a cortical surface-based instead of a volumetric searchlight procedure as the former promises higher spatial specificity. The construction of 2D surface-based searchlights was a two-point procedure. First, we defined 2D searchlight disks on subject-specific reconstructed cortical surfaces by identifying all vertices less than 9 mm away in geodesic space for each vertex *v*. Geodesic distances between vertices were approximated by the length of the shortest path on the surface between two vertices by Dijkstra's algorithm⁵⁰. Second, we extracted fMRI activity patterns in functional space corresponding to the vertices comprising the searchlight disks. Voxels belonging to a searchlight were constrained to appear only once in a searchlight, even if they were nearest neighbor to several vertices. For random effects analysis, i.e. to summarize results across subjects, we estimated a mapping between subject-specific surfaces and an average surface using freesurfer⁵⁰ (fsaverage).

Convolutional neural network architecture and training. We used a deep neural network (DNN) architecture as described by Krizhevsky *et al.*²⁹ (Fig. 1a). We chose this architecture because it was the best-performing neural network in the ImageNet Large Scale Visual Recognition Challenge 2012, it is inspired by biological principles. The network architecture consisted of 8 layers; the first five layers were convolutional; the last three were fully connected. Layers 1 and 2 consisted of three stages: convolution, max pooling and normalization; layers 3–5 consisted of a convolution stage only (enumeration of units and features for each layer in Suppl. Table 5). We used the last processing stage of each layer as model output of each layer for comparison with fMRI and MEG data.

We constructed 5 different DNN models that differed in the categorization task they were trained on (Fig. 5a): (1) *object DNN*, i.e. a model trained on object categorization; (2) *untrained DNN*, i.e. an untrained model initialized with random weights; (3) *scene DNN*, i.e. a model trained on scene categorization; (4) *unecological DNN*, i.e. a model trained on object categorization but with random assignment of label to the training image set; and (5) *noise DNN*, i.e. a model trained to categorize structured noise images. In detail, the object DNN was trained with 900k images of 683 different objects from ImageNet⁴⁷ with roughly equal number of images per object (~1300). The scene DNN, was trained with the recently released PLACES dataset that contains images from different scene categories³⁰. We used 216 scene categories and normalized the total number of images to be equivalent to the number of images used to train the object DNN. For the noise DNN we created an image set consisting of 1000 random categories of 1300 images each. All noise images were sampled independently of each other and had size 256 × 256 with 3 color channels. To generate, each color channel and pixel was sampled independently from a uniform [0, 1] distribution, followed by convolution with a 2D Gaussian filter of size 10 × 10 with standard deviation of 80 pixels. The resulting noise images had small but perceptible spatial gradients.

All DNNs except the untrained DNN were trained on GPUs using the Caffe toolbox (<http://caffe.berkeley-vision.org/>) with the learning parameters set as follows: the networks were trained for 450k iterations, with the initial learning rate set to 0.01 and a step multiple of 0.1 every 100k iterations. The momentum and weight decay were fixed at 0.9 and 0.0005 respectively.

To ascertain that we successfully trained the networks, we determined their performance in predicting the category of images in object and scene databases based on the output of layer 7. As expected, the deep object- and scene networks performed comparably to previous DNNs trained on object and scene categorization, whereas the unecological and noise networks performed at chance level (Suppl. Table 1).

To determine classification accuracy of the object DNN on the 118-image set used to probe the brain here, we determined the 5 most confident classification labels for each image. We then manually verified whether the predicted labels matched the expected object category. Manual verification was required to correctly identify categories that were visually very similar but had different labels e.g., backpack and book bag, or airplane and airliner. Images belonging to categories for which the network was not trained (i.e., person, apple, cattle, sheep) were marked as incorrect. Overall, the network classified 111/118 images correctly, resulting in a 94% success rate, comparable to humans¹⁸ (image-specific voting results available online at <http://brainmodels.csail.mit.edu>).

Visualization of model neuron receptive field properties and DNN connectivity. We used a neuroscience-inspired reduction method to determine the receptive field (RF) properties size and selectivity of model neurons³⁰. In short, for any neuron we determined the $K=25$ most-strongly activating images. 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 DNN and compared the output to the original image, thus constructing a discrepancy map that indicates which portion of the image drives the neuron. Re-centering and averaging discrepancy maps generated the final RF.

To illustrate the selectivity of neuron RFs, we use shaded regions to highlight the image area primarily driving the neuron response (for visualization please see <http://brainmodels.csail.mit.edu/dnn/rf/>). This was obtained by first producing the neuron feature map (the output of a neuron to a given image as it convolves the output of the previous layer), then multiplying the neuron RF with the value of the feature map in each location, summing the contribution across all pixels, and finally thresholding this map at 50% its maximum value.

To illustrate the parameters of the object deep network, we developed a tool (DrawNet; <http://brainmodels.csail.mit.edu/dnn/drawCNN/>) that plots for any chosen neuron in the model 1) the selectivity of the neuron for a particular image, and the strongest connections (weights) between the neurons in the previous and next layer. Only connections with weights that exceed a threshold of 0.75 times the maximum weight for a particular neuron are displayed. DrawNet plots properties for the pooling stage of layers 1, 2 and 5 and for the convolutional stage of layers 3 and 4.

Analysis of fMRI, MEG and computer model data in a common framework. To compare brain imaging data (fMRI, MEG) with the DNN in a common framework we used representational similarity analysis (Fig. 2)^{27,31}. The basic idea is that if two images are similarly represented in the brain, they should be similarly represented in the computer model, too. Pair-wise similarities, or equivalently dissimilarities, between the 118 condition-specific representations can be summarized in a representational dissimilarity matrix (RDM) of size 118×118 , indexed in rows and columns by the compared conditions. Thus representational dissimilarity matrices can be calculated for fMRI (one fMRI RDM for each ROI or searchlight), for MEG (one MEG RDM for each millisecond), and for DNNs (one DNN RDM for each layer). In turn, layer-specific DNN RDMs can be compared to fMRI or MEG RDMs yielding a measure of brain-DNN representational similarity. The specifics of RDM construction for MEG, fMRI and DNNs are given below.

Multivariate analysis of fMRI data yields space-resolved fMRI representational dissimilarity matrices. To compute fMRI RDMs we used a correlation-based approach. The analysis was conducted independently for each subject. First, for each ROI (V1, IT, or IPS1&2) and each of the 118 conditions we extracted condition-specific t-value activation patterns and concatenated them into vectors, forming 118 voxel pattern vectors of length $V = 100$. We then calculated the dissimilarity (1 minus Spearman's R) between t-value patterns for every pair of conditions. This yielded a 118×118 fMRI representational dissimilarity matrix (RDM) indexed in rows and columns by the compared conditions for each ROI. Each fMRI RDM was symmetric across the diagonal, with entries bounded between 0 (no dissimilarity) and 2 (complete dissimilarity).

To analyze fMRI data in a spatially unbiased fashion we used a surface-based searchlight method. Construction of fMRI RDMs was similar to the ROI case above, with the only difference that activation pattern vectors were formed separately for each voxel by using t-values within each corresponding searchlight, thus resulting in voxel-resolved fMRI RDMs.

Construction of DNN layer-resolved and summary DNN representational dissimilarity matrices. To compute DNN RDMs we again used a correlation-based approach. For each layer of the DNN, we extracted condition-specific model neuron activation values and concatenated them into a vector. Then, for each condition pair we computed the dissimilarity (1 minus Spearman's R) between the model activation pattern vectors. This yielded a 118×118 DNN representational dissimilarity matrix (DNN RDM) summarizing the representational dissimilarities for each layer of a network. The DNN RDM is symmetric across the diagonal and bounded between 0 (no dissimilarity) and 2 (complete dissimilarity).

For an analysis of representational dissimilarity at the level of whole DNNs rather than individual layers we modified the aforementioned procedure (Fig. 5b). Layer-specific model neuron activation values were concatenated before entering similarity analysis, yielding a single DNN RDM per model. To balance the contribution of each layer irrespective of the highly different number of neurons per layer, we applied a principal component analysis (PCA) on the condition- and layer-specific activation patterns before concatenation, yielding 117-dimensional summary vectors for each layer and condition. Concatenating the 117-dimensional vector across 8 layers yielded a $117 \times 8 = 936$ dimensional vector per condition that entered similarity analysis.

Multivariate analysis of MEG data yields time-resolved MEG representational dissimilarity matrices. To compute MEG RDMs we used a decoding approach with a linear support vector machine (SVM)^{20,54}. The idea is that if a classifier performs well in predicting condition labels based on MEG data, then the MEG visual representations must be sufficiently dissimilar. Thus, decoding accuracy of a classifier can be interpreted as a dissimilarity measure. The motivation for a classifier-based dissimilarity measure rather than 1 minus Spearman's R (as above) is that a SVM classifier selects MEG sensors that contain discriminative information in noisy data without human intervention. A dissimilarity measure over all sensors might be strongly influenced by noisy channels, and an a-priori sensor selection might introduce a bias, and neglect the fact that different channels contain discriminate information over time.

We extracted MEG sensor level patterns for each millisecond time point (100 ms before to 1,000 ms after image onset) and for each trial. For each time point, MEG sensor level activations were arranged in 306 dimensional vectors (corresponding to the 306 MEG sensors), yielding $M = 30$ pattern vectors per time point and condition). To reduce computational load and improve signal-to-noise ratio, we sub-averaged the M vectors in groups of $k = 5$ with random assignment, thus obtaining $L = M/k$ averaged pattern vectors. For each pair of conditions, we assigned $L-1$ averaged pattern vectors to a training data set used to train a linear support vector machine in the LibSVM implementation (www.csie.ntu.edu.tw/~cjlin/libsvm). The trained SVM was then used to predict the condition labels of the left-out testing data set consisting of the L th averaged pattern vector. We repeated this process 100 times with random assignment of the M raw pattern vectors to L averaged pattern vectors. We assigned

the average decoding accuracy to a decoding accuracy matrix of size 118×118 , with rows and columns indexed by the classified conditions. The matrix was symmetric across the diagonal, with the diagonal undefined. This procedure yielded one 118×118 matrix of decoding accuracies and thus one MEG representational dissimilarity matrix (MEG RDM) for every time point.

Representational similarity analysis compares brain data to DNNs. We used representational similarity analysis to compare layer-specific DNN RDMs to space-resolved fMRI RDMs or time-resolved MEG RDMs (Fig. 2). In particular, fMRI or MEG RDMs were compared to layer-specific DNN RDMs by calculating Spearman's correlation between the lower half of the RDMs excluding the diagonal. All analyses were conducted on single-subject basis.

A comparison of time-resolved MEG RDMs and DNN RDMs (Fig. 2) yielded the time course with which visual representations common to brains and DNNs emerged. For the comparison of fMRI and DNNs RDMs, fMRI searchlight (Fig. 2) and ROI RDMs were compared with DNN RDMs, yielding single ROI values and 2-dimensional brain maps of similarity between human brains and DNNs respectively.

For the searchlight-based fMRI-DNN comparison procedure in detail, we computed the Spearman's R between the DNN RDM of a given layer and the fMRI RDM of a particular voxel in the searchlight approach. The resulting similarity value was assigned to a 2D map at the location of the voxel. Repeating this procedure for each voxel yielded a spatially resolved similarity map indicating common brain-DNN representations. The entire analysis yielded 8 maps, i.e. one for each DNN layer. Subject-specific similarity maps were transformed into a common average cortical surface space before entering random-effects analysis.

Statistical testing. For random-effects inference we used sign permutation tests. In short, we randomly changed the sign of the data points (10,000 permutation samples) for each subject to determine significant effects at a threshold of $P < 0.05$. To correct for multiple comparisons in cases where neighboring tests had a meaningful structure, i.e. neighboring voxels in the searchlight analysis and neighboring time points in the MEG analysis, we used cluster extent (i.e., each data point equally weighted) inference with a cluster extent threshold of $P < 0.05$, Bonferroni-corrected for multiple comparisons. In other cases, we used FDR correction.

To provide estimates of the accuracy of a statistic we bootstrapped the pool of subjects (1,000 bootstraps) and calculated the standard deviation of the sampled bootstrap distribution. This provided the standard error of the statistic.

References

1. Ungerleider, L. G. & Mishkin, M. In *Analysis of Visual Behavior* 549–586 (MIT Press, 1982).
2. Felleman, D. J. & Van Essen, D. C. Distributed Hierarchical Processing in the Primate Cerebral Cortex. *Cereb. Cortex* **1**, 1–47 (1991).
3. Bullier, J. Integrated model of visual processing. *Brain Res. Rev.* **36**, 96–107 (2001).
4. Milner, A. D. & Goodale, M. A. *The visual brain in action*. (Oxford University Press, 2006).
5. Kourtzi, Z. & Connor, C. E. Neural Representations for Object Perception: Structure, Category, and Adaptive Coding. *Annu. Rev. Neurosci.* **34**, 45–67 (2011).
6. Kravitz, D. J., Saleem, K. S., Baker, C. I. & Mishkin, M. A new neural framework for visuospatial processing. *Nat. Rev. Neurosci.* **12**, 217–230 (2011).
7. DiCarlo, J. J., Zoccolan, D. & Rust, N. C. How Does the Brain Solve Visual Object Recognition? *Neuron* **73**, 415–434 (2012).
8. Riesenhuber, M. & Poggio, T. Hierarchical models of object recognition in cortex. *Nat. Neurosci.* **2**, 1019–1025 (1999).
9. Riesenhuber, M. & Poggio, T. Neural mechanisms of object recognition. *Curr. Opin. Neurobiol.* **12**, 162–8 (2002).
10. Naselaris, T., Prenger, R. J., Kay, K. N., Oliver, M. & Gallant, J. L. Bayesian Reconstruction of Natural Images from Human Brain Activity. *Neuron* **63**, 902–915 (2009).
11. David, S. V., Hayden, B. Y. & Gallant, J. L. Spectral Receptive Field Properties Explain Shape Selectivity in Area V4. *J. Neurophysiol.* **96**, 3492–3505 (2006).
12. Wang, G., Tanaka, K. & Tanifuji, M. Optical Imaging of Functional Organization in the Monkey Inferotemporal Cortex. *Science* **272**, 1665–1668 (1996).
13. Yamane, Y., Carlson, E. T., Bowman, K. C., Wang, Z. & Connor, C. E. A neural code for three-dimensional object shape in macaque inferotemporal cortex. *Nat. Neurosci.* **11**, 1352–1360 (2008).
14. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
15. Mnih, V. *et al.* Human-level control through deep reinforcement learning. *Nature* **518**, 529–533 (2015).
16. Zhou, B., Lapedriza, A., Xiao, J., Torralba, A. & Oliva, A. Learning Deep Features for Scene Recognition using Places Database. *Adv. Neural Inf. Process. Syst.* **27** (2014).
17. Rumelhart, D. E., Hinton, G. E. & Williams, R. J. Learning representations by back-propagating errors. *Nature* **323**, 533–536 (1986).
18. Russakovsky, O. *et al.* ImageNet Large Scale Visual Recognition Challenge. *ArXiv14090575 Cs* (2014).
19. He, K., Zhang, X., Ren, S. & Sun, J. Delving Deep into Rectifiers: Surpassing Human-Level Performance on ImageNet Classification. *ArXiv150201852 Cs* (2015).
20. Cichy, R. M., Pantazis, D. & Oliva, A. Resolving human object recognition in space and time. *Nat. Neurosci.* **17**, 455–462 (2014).
21. Schmolesky, M. T. *et al.* Signal Timing Across the Macaque Visual System. *J. Neurophysiol.* **79**, 3272–3278 (1998).
22. Cichy, R., Pantazis, D. & Oliva, A. Similarity-based fusion of MEG and fMRI reveals spatio-temporal dynamics in human cortex during visual object recognition. *bioRxiv* 32656 (2015). doi:10.1101/032656.
23. Clarke, A., Devereux, B. J., Randall, B. & Tyler, L. K. Predicting the Time Course of Individual Objects with MEG. *Cereb. Cortex* **25**, 3602–12 (2015).
24. Agrawal, P., Stansbury, D., Malik, J. & Gallant, J. L. Pixels to Voxels: Modeling Visual Representation in the Human Brain. *ArXiv14075104 Cs Q-Bio* (2014).
25. Cadieu, C. F. *et al.* Deep Neural Networks Rival the Representation of Primate IT Cortex for Core Visual Object Recognition. *PLoS Comput. Biol.* **10**, e1003963 (2014).
26. Güçlü, U. & Gerven, M. A. J. van. Deep Neural Networks Reveal a Gradient in the Complexity of Neural Representations across the Ventral Stream. *J. Neurosci.* **35**, 10005–10014 (2015).
27. Khaligh-Razavi, S.-M. & Kriegeskorte, N. Deep Supervised, but Not Unsupervised, Models May Explain IT Cortical Representation. *PLoS Comput. Biol.* **10**, e1003915 (2014).
28. Yamins, D. L. K. *et al.* Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proc. Natl. Acad. Sci. USA* **111**, 8619–8624 (2014).

29. Krizhevsky, A., Sutskever, I. & Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems* (2012).
30. Zhou, B., Khosla, A., Lapedriza, A., Oliva, A. & Torralba, A. Object Detectors Emerge in Deep Scene CNNs. *Int. Conf. Learn. Represent. ICLR 2015* (2015).
31. Kriegeskorte, N. Representational similarity analysis – connecting the branches of systems neuroscience. *Front. Syst. Neurosci.* **2**, 4 (2008).
32. Mormann, F. *et al.* Latency and Selectivity of Single Neurons Indicate Hierarchical Processing in the Human Medial Temporal Lobe. *J. Neurosci.* **28**, 8865–8872 (2008).
33. Cichy, R. M., Khosla, A., Pantazis, D. & Oliva, A. Dynamics of scene representations in the human brain revealed by magnetoencephalography and deep neural networks. *NeuroImage* doi:10.1016/j.neuroimage.2016.03.063
34. Freiwald, W. A., Tsao, D. Y. & Livingstone, M. S. A face feature space in the macaque temporal lobe. *Nat. Neurosci.* **12**, 1187–1196 (2009).
35. Thorpe, S., Fize, D. & Marlot, C. Speed of processing in the human visual system. *Nature* **381**, 520–522 (1996).
36. Kriegeskorte, N., Goebel, R. & Bandettini, P. Information-based functional brain mapping. *Proc. Natl. Acad. Sci. USA* **103**, 3863–3868 (2006).
37. Janssen, P., Srivastava, S., Ombelet, S. & Orban, G. A. Coding of Shape and Position in Macaque Lateral Intraparietal Area. *J. Neurosci.* **28**, 6679–6690 (2008).
38. Sawamura, H., Georgieva, S., Vogels, R., Vanduffel, W. & Orban, G. A. Using Functional Magnetic Resonance Imaging to Assess Adaptation and Size Invariance of Shape Processing by Humans and Monkeys. *J. Neurosci.* **25**, 4294–4306 (2005).
39. Chao, L. L. & Martin, A. Representation of Manipulable Man-Made Objects in the Dorsal Stream. *NeuroImage* **12**, 478–484 (2000).
40. Konen, C. S. & Kastner, S. Two hierarchically organized neural systems for object information in human visual cortex. *Nat. Neurosci.* **11**, 224–231 (2008).
41. Grill-Spector, K. *et al.* Differential processing of objects under various viewing conditions in the human lateral occipital complex. *Neuron* **24**, 187–203 (1999).
42. Kourtzi, Z. & Kanwisher, N. Cortical Regions Involved in Perceiving Object Shape. *J. Neurosci.* **20**, 3310–3318 (2000).
43. Saxe, A. M. *et al.* On random weights and unsupervised feature learning. in *In NIPS Workshop on Deep Learning and Unsupervised Feature Learning* (2010).
44. DeYoe, E. A., Felleman, D. J., Van Essen, D. C. & McClendon, E. Multiple processing streams in occipitotemporal visual cortex. *Nature* **371**, 151–4 (1994).
45. Ahissar, M. & Hochstein, S. The reverse hierarchy theory of visual perceptual learning. *Trends Cogn. Sci.* **8**, 457–464 (2004).
46. Kourtzi, Z. & DiCarlo, J. J. Learning and neural plasticity in visual object recognition. *Curr. Opin. Neurobiol.* **16**, 152–158 (2006).
47. Deng, J. *et al.* ImageNet: A large-scale hierarchical image database. in *IEEE Conference on Computer Vision and Pattern Recognition, 2009. CVPR 2009 248–255* (2009). doi:10.1109/CVPR.2009.5206848.
48. Taulu, S., Kajola, M. & Simola, J. Suppression of interference and artifacts by the Signal Space Separation Method. *Brain Topogr.* **16**, 269–275 (2004).
49. Taulu, S. & Simola, J. Spatiotemporal signal space separation method for rejecting nearby interference in MEG measurements. *Phys. Med. Biol.* **51**, 1759 (2006).
50. Dale, A. M., Fischl, B. & Sereno, M. I. Cortical Surface-Based Analysis: I. Segmentation and Surface Reconstruction. *Neuroimage* **9**, 179–194 (1999).
51. Benson, N. C. *et al.* The Retinotopic Organization of Striate Cortex Is Well Predicted by Surface Topology. *Curr. Biol.* **22**, 2081–2085 (2012).
52. Wang, L., Mruczek, R. E. B., Arcaro, M. J. & Kastner, S. Probabilistic Maps of Visual Topography in Human Cortex. **25**, 3911–31 (2015).
53. Chen, Y. *et al.* Cortical surface-based searchlight decoding. *NeuroImage* **56**, 582–592 (2011).
54. Cichy, R. M., Ramirez, F. M. & Pantazis, D. Can visual information encoded in cortical columns be decoded from magnetoencephalography data in humans? *Neuroimage* **121**, 193–204 (2015).

Acknowledgements

We thank Chen Yi for assisting in surface-based searchlight analysis. This work was funded by National Eye Institute grant EY020484 (to A.O.), a Google Research Faculty Award (to A.O.), a Feodor Lynen Scholarship of the Humboldt Foundation and the Emmy Noether Program of the Deutsche Forschungsgemeinschaft (CI 241/1-1) (to R.M.C), the McGovern Institute Neurotechnology Program (to A.O. and D.P.), National Science Foundation Award 1532591 (to A.O, A.T and D.P), and was conducted at the Athinoula A. Martinos Imaging Center at the McGovern Institute for Brain Research, Massachusetts Institute of Technology.

Author Contributions

All authors conceived the experiments. R.M.C. and D.P. acquired and analyzed brain data, A.K. trained and analyzed computer models. R.M.C. provided model-brain comparison. R.M.C., A.K., D.P. and A.O. wrote the paper, A.T. provided expertise and feedback. A.O., D.P., A.T. and R.M.C. provided funding.

Additional Information

Supplementary information accompanies this paper at <http://www.nature.com/srep>

Competing financial interests: The authors declare no competing financial interests.

How to cite this article: Cichy, R. M. *et al.* Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence. *Sci. Rep.* **6**, 27755; doi:10.1038/srep27755 (2016).



This work is licensed under a Creative Commons Attribution 4.0 International License. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in the credit line; if the material is not included under the Creative Commons license, users will need to obtain permission from the license holder to reproduce the material. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>