



OPEN Coordinating multiple mental faculties during learning

Xiaoliang Luo¹✉, Robert M. Mok^{2,3}, Brett D. Roads¹ & Bradley C. Love^{1,4}

Complex behavior is supported by the coordination of multiple brain regions. How do brain regions coordinate absent a homunculus? We propose coordination is achieved by a controller-peripheral architecture in which peripherals (e.g., the ventral visual stream) aim to supply needed inputs to their controllers (e.g., the hippocampus and prefrontal cortex) while expending minimal resources. We developed a formal model within this framework to address how multiple brain regions coordinate to support rapid learning from a few example images. The model captured how higher-level activity in the controller shaped lower-level visual representations, affecting their precision and sparsity in a manner that paralleled brain measures. In particular, the peripheral encoded visual information to the extent needed to support the smooth operation of the controller. Alternative models optimized by gradient descent irrespective of architectural constraints could not account for human behavior or brain responses, and, typical of standard deep learning approaches, were unstable trial-by-trial learners. While previous work offered accounts of specific faculties, such as perception, attention, and learning, the controller-peripheral approach is a step toward addressing next generation questions concerning how multiple faculties coordinate.

The ability to master complex tasks requires coordination of multiple perceptual, cognitive, and motor processes subserved by numerous brain regions. Long before the advent of computational neuroscience, philosophers like John Locke and David Hume appreciated that the “faculties” of the mind must coordinate with one another to produce coherent thought^{1,2}. The development of cognitive architectures in the symbolic production system tradition was one attempt to address the coordination challenge^{3–5}. In contrast, deep learning approaches in neuroscience have largely focused on a single faculty (e.g., object recognition) and its supporting circuit (e.g., the ventral visual stream).

We aim to address this gap by developing a general solution to the coordination problem and applying it to the domain of category learning, which requires the coordination of multiple cognitive processes related to attention, learning, object recognition, memory encoding and consolidation, and relies on coordinating multiple brain regions (e.g.,^{6,7}).

Cognitive models of category learning have been developed to capture behavior in rapid learning paradigms, which focus on tasks like category acquisition rather than object recognition. Some of these models have also been linked to neural activity in the brain. For example, SUSTAIN⁸ has been shown to not only account for rapid human learning behavior but also associated neural activity in the hippocampus and ventromedial prefrontal cortex (vmPFC)^{9–14}. Notably, work using SUSTAIN¹² has revealed a process of neural compression in vmPFC, where task-relevant dimensions of information are selectively preserved while irrelevant dimensions are filtered out. However, these models do not address perception (e.g., object recognition) or explain how higher-level notions of attention, which rely on the prefrontal cortex, relate to and potentially influence visual processing along the ventral visual stream. Thus, how perception and cognition coordinate in the brain remains an open question.

Deep Neural Networks (DNNs) address aspects of perception neglected by cognitive models. Although not without their shortcomings^{15–20}, these models can achieve human-level accuracy on object-recognition tasks involving photographs of real-world objects^{21,22}, and exhibit functional correspondence to the primate ventral visual stream^{23–27}. Like cognitive models²⁸, they do not address the coordination problem. Indeed, DNNs’ intended application is restricted to the relatively automatic feed-forward aspects of object recognition referred to as “core object recognition”²⁹. Moreover, DNNs typically learn representations from stationary batches of training data, lacking the ability to account for scenarios where information becomes incrementally available over time (i.e., continual learning; see³⁰ for a review).


¹Department of Experimental Psychology, University College London, 26 Bedford Way, London WC1H 0AP, UK. ²MRC Cognition and Brain Sciences Unit, University of Cambridge, 15 Chaucer Rd, Cambridge CB2 7EF, UK. ³Department of Psychology, Royal Holloway, University of London, Egham TW20 0EX, UK. ⁴The Alan Turing Institute, 96 Euston Rd, London NW1 2DB, UK. ✉email: xiao.luo.17@ucl.ac.uk

Given the complementary roles cognitive and DNN models play in capturing cognition and perception, one obvious path to integration is using the outputs of DNN models as inputs to cognitive models (e.g.,^{31–33}). Although appealing straightforward, this approach does not address how different cognitive processes and their underlying brain regions *interact* to create intelligent behavior. Decades of work in top-down attentional control from neurophysiology³⁴, functional neuroimaging^{35–40}, and neuropsychology^{41,42} suggest that top-down processes from control regions modulate sensory regions. Rather than a simple hand-off from perception to cognition, what is needed is a model that takes images as input, deploys top-down attention, rapidly learns novel categories, and makes decisions in a manner that captures how a multitude of brain regions, including the ventral visual stream, hippocampus, and vmPFC, coordinate to support rich and adaptive behaviors.

To achieve this aim, we propose a general modeling framework that captures mental function through coordinated interactions across multiple brain regions; a *controller-peripheral architecture* (Fig. 1). In this framework, controllers are typically higher-level cognitive regions that directly optimize some external objective tied to behavior whereas peripherals are typically lower-level sensory regions that aim to supply information needed to maintain and support the operation of their controller or controllers. For example, eye movements can be peripheral to higher-level control processes that direct fixations toward information needed for the decision⁴³. By linking controllers and peripherals, accounts of quasi-hierarchical control, multimodal processing, and modularity can be specified. Although we will focus on an account of category learning in which there is only one controller (a cognitive model of the hippocampus and vmPFC) and one peripheral (a perception model of the ventral visual stream), more complex arrangements of controllers and peripherals are possible within the controller-peripheral architecture (Fig. 1A).

Here, we treat the hippocampus and vmPFC as one integrated module (a single controller) because we view them as the primary drivers of overall system in the category learning tasks we consider^{11,12}. Depending on one's scientific aims, one might instead decompose these regions into multiple controllers and/or peripherals¹⁴. Indeed, the hippocampus and vmPFC can serve different purposes when driving behavior, with vmPFC being chiefly responsible for determining goal relevancy reflected in higher-level attention signals and hippocampus being responsible for forming categorical structure over acquired knowledge^{11,12} (also see⁴⁴).

A Possible Controller-Peripheral Arrangements



B A Cross-System Controller-Peripheral Account of Category Learning




Fig. 1. The controller-peripheral architecture provides a general framework for how different brain regions coordinate while performing a task. **(A)** Peripherals aim to supply their controllers with the information they require while expending minimal resources (i.e., costly energy principle). Here, we illustrate a number of possible arrangements of controllers and peripherals. **(ii)** A single controller with multiple peripherals could offer an account of multi-modal integration for the convergence of visual and somatosensory signals in parietal cortex⁵² or semantic hubs in the anterior temporal lobe⁵³. **(iii)** Conversely, multiple controllers with a single peripheral could model eye movements in which multiple controllers related to visual search, obstacle avoidance, social cognition, etc. share this perceptual resource. **(iv)** The controller and peripheral can share reciprocal connections, allowing each to serve both roles interchangeably. This arrangement may be particularly useful for modeling the interaction between the prefrontal cortex (PFC) and premotor cortex in movement execution and recalibration. Controllers and peripherals can be arranged hierarchically as in **(v)**. This arrangement is consistent with hierarchical accounts of the ventral visual stream in object recognition. **(B)** We use the controller-peripheral architecture to develop a model that can learn concepts from a few visual examples. To simplify, we assume a single controller involving the hippocampus and ventromedial prefrontal cortex (vmPFC) and a single peripheral involving the ventral visual stream. The model captures how higher-level goals and outcomes shape activity throughout the ventral visual stream, which aims to provide its controller with needed information while minimizing resource expenditure (i.e., the costly energy principle).

While controllers optimize their states with respect to cognitive demands, peripherals supply sensory input to controllers in a way that preserves controller states and minimizes energy, broadly construed. Peripherals follow a *costly energy principle* that seeks to minimize resource expenditures in accord with a broad range of cognitive effort accounts, including those focused on depletion of blood glucose^{45,46}, opportunity costs^{47–49}, and interference across tasks relying on shared resources^{50,51}. Colloquially, a peripheral is like a worker who gives the boss (the controller) what they want while minimizing unnecessary effort.

In the present contribution, we instantiate a model of category learning within the controller-peripheral framework that offers an account of how vmPFC, the hippocampus, and ventral visual stream coordinate to rapidly learn categories from images. To foreshadow our results, the controller-peripheral approach better accounts for behavior and brain response than other approaches that independently adjust parameters (e.g., weights) to maximize performance as is common in machine learning (ML) systems (e.g., end-to-end optimization by gradient descent learning).

Model overview

The controller follows and extends the principles of a successful model of category learning, SUSTAIN⁸. SUSTAIN provides a good foundation for the controller because SUSTAIN has successfully accounted for both behavior and brain activity (hippocampus and vmPFC) during category learning tasks^{9–12}.

SUSTAIN is a clustering model of category learning that starts off simple and adds clusters in response to surprising errors. SUSTAIN is a cognitive model and, in contrast to deep learning models, takes as input stimulus representations that are hand-coded by the experimenter. Clusters are activated according to how similar they are to the current stimulus. Similarity is a function of attention-weighted distance between the stimulus and cluster. Attention weights are learned and favor stimulus dimensions that are most informative. Clusters compete to respond and inhibit one another. Association weights pass activity from clusters to the output layer (producing possible actions or responses) and are learned to minimize task error. In SUSTAIN, only the most activated cluster has nonzero activation (i.e., a winner-takes-all scheme). The winning cluster moves (i.e., adjusts its receptive field) to be closer to the current stimulus. When a surprising error occurs (e.g., learning that a bat is a mammal and not a bird), a new cluster is recruited and centered on the current stimulus.

The controller module extends SUSTAIN, generalizing several aspects of the original model. In particular, all clusters contribute to varying degrees in the recruitment process and to the output response, which contrasts with SUSTAIN's strict winner-takes-all operation (Eq. 9). In addition, while SUSTAIN's attention mechanism is a localized operation updated based on the winner cluster's relationship with the input, the controller employs a differentiable attention mechanism, which is optimized directly by minimizing task error. These extensions increase the model's applicability and make it fully differentiable, which eases its incorporation into larger systems, such as into our controller-peripheral system considered here.

The peripheral module is based on VGG-16²². VGG-16 is a deep convolutional neural network that was trained to perform object recognition tasks. DNNs, such as VGG-16, are commonly used by neuroscientists to model activity along the ventral visual stream^{23–25,25}. We chose VGG-16 because it is a well known model, has a relatively straightforward architecture, and performs well on benchmarks that assess recognition behavior and agreement with brain responses along the ventral visual stream^{54,55}.

We made a number of changes and extensions to VGG-16 so that it could function as the peripheral. VGG-16 is an object recognition model that takes an image as input and outputs a label from a fixed set of pre-existing categories (e.g., penguin, house, car, etc.) after training on millions of image-label pairings. Instead, we focus on the challenging rapid learning of novel categories from a small set of examples. For this task, we need the peripheral to take an image as input and output a perceptual representation that the controller can take as input (Fig. 1B).

To achieve this, we only preserved the layers of VGG-16 that correspond to regions along the ventral visual stream up to and including lateral occipital cortex (LOC;⁵⁶). We then fine-tuned the peripheral model to output three features in response to a stimulus (an image of a bug; see Supplementary Information, Table S.6). In these studies, the dimensional structure of verbalizable stimuli was conveyed to participants through instructions and examples prior to the experiment (e.g., thin/thick legs, pincer/shovel mouth). To capture this familiarization process in the model, we used this fine-tuning approach to achieve an analog of how human participants' familiarization process prior to category learning. We note that the familiarization process relies on fewer examples compared to fine-tuning in the model. While not perfectly analogous, we consider it a reasonable abstract approximation that captures the abstract process of how people understood the stimuli in the study. We focus on LOC because it is modulated by category learning tasks⁵⁷, including task-related attention modulation⁵⁸, and it is a convergence point for top-down cognitive and bottom-up perceptual processes^{59–61}.

One key aspect of the controller-peripheral account is that the peripheral aims to provide the controller with the information it needs while expending minimal resources (i.e., costly-energy principle). The peripheral altered its operation in response to the controller by adjusting its attention weights. The peripheral's attention mechanism was closely modeled by⁶² in which a nonzero attention weight modulated the output of each DNN filter (i.e., feature). To solve the coordination problem, the peripheral minimizes the reconstruction error of the controller's cluster representations while increasing attention weight sparsity, guided by the costly-energy principle. This principle is implemented as an L_1 penalty on the sum of attention weights, driving some weights toward zero to encourage sparsity. We consider the sparsity of the attention weights (i.e., proportion that are zero) in the DNN module as an indicator for resource expenditure. This perspective aligns with the notion that the visual system may employ sparse coding as an efficient scheme for representing sensory signals (e.g.,^{63–65}). As the peripheral's attention is altered according to the costly-energy principle, how precise stimulus information is encoded along dimensions will vary according to task demands and diverge from the experimenter-defined binary values. Notably, attention mechanisms in the controller and peripheral serve distinct roles. In the

controller, attention is *decisional*, dynamically weighing information to optimize task learning, without direct access to perceptual information. In the peripheral, attention is *perceptual*, optimized to provide the necessary sensory information for the controller's task learning, without direct access to learning errors. These interactions between cognition and perception allow us to predict how encoded perceptual information changes over time (Fig. 4A and Fig. S2) and relate model representations to brain activity (Fig. 4B).

Notice this controller-peripheral approach diverges from standard ML approaches in which all parameters (including peripheral attention weights) are optimized to improve the decisions of the overall system (end-to-end training). Another key difference with ML models is that our controller-peripheral model learns in a trial-by-trial manner consistent with procedures used in human experiments (see Methods). Unlike most DNN models of perception, our peripheral model alters its operation in response to higher-level goals, as reflected by the current state of the controller. Thus, rather than viewing perception and cognition as independent modules in which perception provides the inputs to cognition, we offer an account of how perceptual and cognitive processes interact and coordinate.

Controller-peripheral model optimized to costly energy principle captures complex learning behaviors

We first evaluated if our model can account for learning performance on six learning problems in⁶⁶ and preserve SUSTAIN's strategies in solving these problems.^{8,66} described six category learning tasks and participants showed learning curves that revealed the difficulty order of the category structures (Fig. 2A). Specifically, Type I was the easiest to master, followed by Type II, followed by Types III–V, and Type VI was the hardest.⁶⁶ is a challenging human category learning dataset to fit and has proven difficult for models that take images as inputs.³²

Here, we used images of the insect stimuli from¹¹, who tested participants on⁶⁶ structures during an fMRI scan. The availability of relevant fMRI data enabled us to go beyond previous work to assess cross-area coordination mechanisms in category learning. The eight insects varied along three binary features (thick/thin legs, thick/thin antennae, and pincer/shovel mouths; see Supplementary Information, Table S.6).

The model successfully captured human learning performance (Fig. 2B). Despite the system taking images as inputs, the controller's clustering solutions paralleled SUSTAIN's in terms of the modal number of clusters recruited (2, 4, 6, 6, 6, 8 for Types I–VI respectively; for full results see Supplementary Information, Table S.1). Likewise, the controller's attention weights paralleled SUSTAIN's solution by selectively weighting the relevant stimulus dimensions (Fig. 2C). Operating over images instead of experimenter-defined stimulus representations enables additional predictions. The peripheral module had the most difficulty ascertaining the value of the mandible dimension from the images. In accord with the model, when this dimension was relevant for human learners they made more errors and their response times were longer (see Supplementary Information, Fig. S.5). The model fits of⁶⁶'s six learning problems shown in Fig. 2B, left panel were from⁶⁷'s replication. In that study, simple shape stimuli were used. When applied to these stimuli, our model performs equivalently (Supplementary Information, Fig. S.2). We further tested a visual transformer architecture⁶⁸ as the peripheral module, and the model produced similar learning patterns (Supplementary Information, Fig. S.3).

Although the model fits are impressive, one obvious question is whether the architecture was necessary. One alternative to the controller-peripheral architecture is to simultaneously optimize all aspects of the model to minimize categorization errors, much like how most modern neural networks are trained. Because the model is fully differentiable, this change amounts to the DNN's learning target shifting from serving the controller as a peripheral to adjusting itself to minimize categorization errors. Although the difference is subtle, this model variant proved unstable and could not account for human learning behavior in⁶⁶ or exhibit the task-specific resource expenditure patterns found in⁵⁸. This instability arose because there was no pressure for the DNN to respect the higher-level clustering solutions and the controller module recruited more clusters than necessary. In effect, the different parts of the model lost coordination and became out-of-sync with each other (see full results in Supplementary Information, Fig. S.1, Table S.2, S.3, S.4). Rapid, trial-by-trial learning appears to require the coordination provided by our architecture.

Controller-peripheral framework explains cross-system neural activities

Having established a model that takes images as input and built within the controller-peripheral framework accounts for complex learning behaviors, we evaluate whether the model can capture how brain regions, such as ventral medial prefrontal cortex (vmPFC) and lateral occipital cortex (LOC)^{12,43,58} in the ventral visual stream, coordinate to support these behaviors. Because the previous section established that our controller's clustering solutions matched those of SUSTAIN which have been related to hippocampal activity^{10,11}, we assume the controller's clusters provide a good account of hippocampal activity during these learning tasks. Here, we fitted the model to individuals' behavior and compared model activity to human fMRI data from¹¹ in which participants learned the Type I, II and VI problems. We trained the model on each participant's stimulus sequence in a trial-by-trial manner (see Methods).

Controller attention tracks neural compression in vmPFC

Prefrontal cortex (PFC) is believed to direct attention toward goal-relevant information^{69,70}. In particular, vmPFC may perform information compression by filtering out task-irrelevant information during category learning^{71–74}.

For example,¹² compared patterns of activity in vmPFC to the learned attention weights in SUSTAIN, a cognitive model that is the inspiration for the controller's clustering model, and found that vmPFC performs goal-directed information compression during learning (see Methods 6.5 for information on compression scores which also matched the cognitive model's attention weights). Indeed, vmPFC, mirroring changes in attention




A Six Concept Learning Structures**B Human and Model Learning Behaviors****C Controller Attention Strength**

Fig. 2. The controller-peripheral framework consisting of a clustering module capturing HPC and PFC and a DNN module capturing ventral visual stream, captures human category learning behavior in⁶⁶. **(A)** Six category learning tasks where human participants learnt to classify geometric shapes into one of two categories where stimuli were made up of three binary-valued features (color, shape and size). Typically, models operate over hand-coded three-dimensional vectors. Instead, we trained on actual images, in this case using the insect stimuli from¹¹, a replication of these classic learning problems with image stimulus inputs. The peripheral part of the model, reflecting the ventral visual stream, extracts these three higher-level dimensions for the controller (Fig. 1B). **(B)** The model fits on the left were from⁶⁷'s replication of the six categorization tasks using simple stimuli as shown in Fig. 2A. Our model (right) captured the same difficulty ordering of the six problems using image stimuli from¹¹ with equivalent problem structure. Probability of error is plotted as a function of learning block for each problem type. **(C)** The controller exhibited the same attention strategies as SUSTAIN, solving Type I by attending to one dimension, Type II by attending to two dimensions and Type III–VI by attending to all three dimensions.

weights over learning, had a unique pattern of neural compression, marked by two main effects and their interaction. Neural compression increased over learning blocks, was higher for learning problems with fewer relevant dimensions, and these two factors interacted such that problems with fewer relevant dimensions (i.e., lower complexity) showed greater compression over learning (Fig. 3A).


We evaluated whether a similar relationship exists between the controller's attention weights and vmPFC. Unlike previous models, the attention mechanism in the controller is part of a control system directing the DNN peripheral. We found that compression scores for the controller's attention weights matched the unique signature of vmPFC's compression scores with both main effects and the interaction found (Fig. 3B; Two-way ANOVA main effects: problem complexity, $F(2, 42) = 78.17, p < 0.001$; Learning block, $F(15, 315) = 28.06, p < 0.001$; Interaction: $F(30, 630) = 13.56, p < 0.001$; Supplementary Information, Table S.5), reflecting greater compression over learning for learning problems with fewer relevant features (Fig. 1B).

Peripheral activity aligns with neural representation in ventral visual stream

We found that the controller provides a good account of hippocampal activity in terms of its clustering solutions and of vmPFC compression patterns in terms of its attention weights. Here, we evaluate how the peripheral part of the model adapts itself to provide useful inputs to the controller following the costly energy principle. We focus on the end stage of the peripheral model, which we hypothesize corresponds to LOC.

Category learning modulates activity in LOC, accentuating information that is goal relevant^{43,57,43} found that more highly attended features were better decoded from multi-voxel activity patterns in LOC. Attention was assessed by SUSTAIN's attention weights after fits to individuals' behavior.⁵⁸ found that the intrinsic dimensionality of the BOLD response in LOC was greater when more aspects of the stimulus were relevant to the categorization decision. Both these findings are consistent with LOC activity reflecting higher-level goals.

A vmPFC Neural Compression Tracks Task Demand



B Controller Attention Compression Reflects Task Demand




Fig. 3. (A) A whole-brain voxel-wise linear mixed effects regression was performed in¹², which revealed a vmPFC region that showed a significant interaction between learning block and problem complexity. Neural compression increased over learning blocks and was higher for learning problems with fewer relevant dimensions (each fMRI run consists of four learning blocks; see the original paper for more details). (B) Functional correspondence between the clustering module of the controller-peripheral system and vmPFC in the human brain. The clustering module deploys attention strategies (in terms of attention compression) that track the degree of neural compression in vmPFC across category learning tasks over learning across category structure complexity.




A DNN Peripheral Reflects Feature Relevance**B Peripheral Information Loss Aligns LOC Representations****C Peripheral Attention Reflects LOC Energy Expenditure**

Fig. 4. Performance of the DNN peripheral and its relation to LOC activity during learning. (A) Following the controller's needs and the costly energy principle, task-relevant features (shaded) are more precisely coded than task-irrelevant features (unshaded). (B) The error-rate for a classifier applied to LOC activity to discriminate (decode) between pairs of stimuli mirrored the precision of the peripheral's feature outputs, consistent with our claim that the peripheral's advanced layers correspond to LOC. (C) Following the costly energy principle, the fewer relevant features for a learning problem (VI>II>I), the more zero-valued peripheral attention weights there are.

Such findings are consistent with our controller-peripheral account and costly energy principle. Under this account, LOC should provide needed information to the controller but with minimal resource expenditure. One measure of resource expenditure is the number of nonzero peripheral attention weights. In agreement with⁵⁸, the number of nonzero peripheral attention weights reflected task complexity (Fig. 4C). Specifically, attention layer learning produced the most sparse representation (fewest active units) in Type I, followed by Type II, with the least sparse representation in Type VI ($b = -0.044$, $t(22) = -5.10$, $p < 0.001$).

Unlike other models of the ventral visual stream, we offer an account of how the controller's needs modulate peripheral activity. One prediction is that the precision of information coded by the peripheral should reflect

the needs of the controller. Information sources that are highly attended by the controller should be more precisely coded, whereas those not attended by the controller can also be ignored by the peripheral in accord with the costly energy principle. We observed this pattern in the outputs of the peripheral (Fig. 4A). Information loss was calculated as the cross-entropy error between the output of the DNN peripheral and the true value for the feature (see Methods; Section 6.5). Consistent with neural decoding patterns found in LOC⁴³, the peripheral's output showed low information loss across Types I, II and VI (Fig. 4A) for stimulus dimensions that were highly attended by controller (Fig. 2C).

In contrast to cognitive models like SUSTAIN that take handcrafted features as inputs, our peripheral model processes images to provide a stimulus coding for the controller. Previous work from⁴³ found that more attended features (according to SUSTAIN) were better decoded from the LOC's BOLD response. Removing this featural assumption, we trained linear support vector classifiers to discriminate each pair of stimuli for each task based on LOC activity. We predicted that classifier error should track mean information loss in the peripheral's feature outputs. When the controller demands precise inputs, the peripheral should provide them and, accordingly, we predicted LOC activity will better discriminate between items. As predicted decoding error (i.e., neural information loss) tracked peripheral information loss, which was greatest for Type I (one feature relevant), followed by Type II (two features relevant), followed by Type VI (all three features relevant). The decoding error ($b = -0.008$, $t(22) = -3.20$, $p = 0.002$; Fig. 4B, left) and model information loss decreased as task complexity increased ($b = -0.060$, $t(22) = -3.94$, $p < 0.001$; Fig. 4B, right).

The peripheral's operation is modulated by the needs of the controller, which change over learning. Therefore, as the controller optimizes its high-level attention for the learning task, the peripheral's attention weights should adjust, which in turn should affect the precision of the peripheral's feature outputs. As predicted, we found that as the controller's attention to a feature decreases the information loss in the peripheral's feature output increases and that as the compression of the controller's attention weights increases the sparsity of the peripheral's attention weights increases (Supplementary Information, Fig. S.4).

Discussion

One challenge for neuroscience is explaining how multiple brain regions coordinate to complete a task. Despite many studies indicating control-related modulation in the brain, testable mechanistic theories of inter-region coordination have been lacking. Regions responsible for control are characterized not unlike a homunculus as the field awaits general accounts of how multiple regions coordinate. Here, we proposed a controller-peripheral framework that provides a mechanistic explanation for such coordination. In this framework, peripherals aim to supply required inputs to controllers while minimizing resource expenditure. Working within this framework, we developed a formal model of category learning in which the peripheral corresponded to the ventral visual stream and the controller to the hippocampus and vmPFC. The controller-peripheral framework enabled us to construct a model of category learning that explains the role of and interactions among a number of brain regions, taking images as inputs and generating behavioral outputs (i.e., decisions).

The model detailed how higher-level goals and knowledge state, reflected in the clustering solution and attention weights of the controller, influenced the peripheral's attentional allocation. The controller was implemented by generalizing several aspects of SUSTAIN⁸, a model of human category learning that has successfully captured human behavior and associated brain response in the hippocampus and vmPFC. Our generalization and reformulation is more suitable to neuro-inspired modeling. In particular, multiple clusters can determine how surprising an error is and govern the model's decision. The generalized model is fully differentiable such that the controller's internal representations can be updated by gradient descent to integrate with other controllers and peripherals. In the present work, this formulation allowed the controller to coordinate with a peripheral. Our controller also contrasts with Bayesian clustering models of category learning^{75–77} by including a selective attention mechanism and being able to integrate its error signals with other models, such as deep neural networks (DNNs). In this contribution, that latter ability led to an account of how perceptual and cognitive processes coordinate when learning novel categories from visual images.

While we emphasized the hippocampus, other brain regions may support category learning abilities. One might ask whether the hippocampus is crucial to learning categories and concepts due to evidence in developmental amnesiacs with hippocampal lesions suggesting preserved semantic memory despite impaired episodic memory⁷⁸. However, new evidence indicates that semantic memory is not entirely spared in developmental amnesiacs, who learn new concepts slower⁷⁹ and exhibit atypical semantic knowledge representations⁸⁰.

The peripheral was implemented as a DNN model of object recognition that we augmented with its own attentional mechanism as proposed in^{62,62} characterized attention as the re-purposing of an existing network in response to a goal-directed signal. Attention was implemented as a set of attention weights that modulated activity in an existing network to better align its function with the current task goal. We advanced their proposal by having attention weights shaped by the needs of the controller network, which provides a grounding for what was before an experimenter-defined goal-directed signal. The overall model, consisting of this peripheral and controller, captured a number of findings. The model was able to account for complex learning behaviors that only a subset of successful cognitive models can address.

As predicted, the pattern of attention weights in the controller matched compression patterns in vmPFC¹² while the clustering solution was consistent with representational similarity patterns in the hippocampus. The controller's state influenced that the peripheral, leading to sparser representations when the controller was concerned with fewer aspects of the stimulus⁵⁸. Following the costly energy principle, the precision of stimulus information transmitted to the controller from the peripheral decreased as the controller's need for that information decreased. In accord with our framework, the peripheral preserved resources to the extent that it could while serving the controller. In summary, the model accounted for how people rapidly learn novel concepts from examples, how this behavior is supported by the ventral visual stream, hippocampus, and vmPFC,

and how these regions coordinate with higher-level learning goals shaping the precision and dimensionality of object representations in the ventral visual stream. These contributions required an account of how mental faculties interact and coordinate, which transcends a perception-cognition dichotomy in which perception provides the inputs to cognition.

We hope future developments more closely link our work with related endeavors. For example, while we consider how new concepts are learned, we do not consider how these concepts are consolidated or integrated with existing semantic knowledge^{53,81}. The controller-peripheral architecture also invites comparisons to work in cognitive control that considers how prefrontal cortex (PFC) supports flexible rule-switching⁸². While we focus on medial as opposed to lateral PFC, a controller-peripheral account could integrate these different brain regions and learning processes by including multiple PFC controllers that influence one another.

Our controller-peripheral framework is a domain-agnostic modeling blueprint for theorizing and developing models of coordinated cognitive processes across interactive brain systems. Category learning is one possible application of this framework, which we focused on in this contribution. However, a variety of controller-peripheral arrangements can be tailored to study different brain regions, cognitive processes or systems across multiple modalities (Fig. 1A). For example, an account of multimodal integration for convergence of visual and somatosensory information in parietal cortex could be captured by a framework with a single controller directing multiple peripherals. The flexible modular nature of the controller-peripheral framework can also help to test predictions and offer integrative explanations across levels of mechanism^{83,84}. For example, a controller capturing cognitive constructs tied to behavior can be decomposed into many controllers and/or peripherals that capture lower-level mechanisms such as populations of or neural assemblies across different brain regions¹⁴. While we combined a DNN model with a cognitive model in the current contribution, different computational architectures could also be adopted for different brain regions within the proposed framework. For example, one direction is to incorporate recurrent connections to better understand interconnected brain networks (e.g.,^{85,86}).

Whereas we focused on rapid learning from a few examples, the controller-peripheral architecture could be extended to explain change over longer time-scales. For example, neural pruning (i.e., programmed neuron death), which is essential to brain development⁸⁷, may unfold in accord with the controller-peripheral architecture and costly energy principle. Likewise, wiring patterns over evolutionary time may be explained by our framework. These proposals parallel recent developments in machine learning that find advantages for starting with large models and sparsifying them to reduce resource requirements while maintaining performance⁸⁸.

We hope the controller-peripheral approach will help neuroscientists develop more encompassing accounts of brain function that address how multiple regions coordinate to perform the tasks of interest. The controller-peripheral architecture may strike the proper balance between strict modularity in which regions operate in prescribed ways versus unstructured approaches in which each unit alters its operation to maximize some global objective, such as maximizing reward, accuracy, etc. In our task, comparable alternative models that reflected these two extremes could not account for human performance. For example, models that treated the DNN peripheral as an independent perceptual module did not capture how learning higher-level category representations affects ventral visual stream activity. On the other extreme, alternative models that sought to globally optimize all parameters proved too unstable for trial-by-trial learning because learning updates from different parts of the model were often at odds with one another.

The controller-peripheral architecture itself is not a theory. Instead, it is a framework to facilitate theory building. Here, we use this framework to offer an account of how people learn categories from images. Like other frameworks, such as the “Bayesian Brain”, our framework does not specify an account for every task, nor the function of every brain region. We believe our framework is particularly well suited to addressing how multiple mental faculties coordinate and hope others adopt it for this purpose.

Energy in the controller-peripheral’s costly energy principle refers to some computational resource, which in our model were neuron-like units (also see⁸⁹) which we related to the BOLD response in LOC. For other models constructed in this framework, the resource may take other forms, such as time when conceived as an opportunity cost. However, energy in the current work should not be confused with energy minimized in variational Bayesian inference under the free energy principle⁹⁰. While predictive processing models can make predictions at the level of neural implementation⁹¹, the controller-peripheral architectures also makes information processing or algorithmic claims by specifying how controllers and peripherals are organized, as evidenced by the different behaviors manifested by comparable models in our simulations that did not use the architecture to coordinate processing. Finally, broad frameworks can be compatible with one another. For example, the controller-peripheral focus on coordination may benefit accounts of brain function consisting of differentiable modules⁹² given the instability we observed for such approaches in trial-by-trial learning.

Cross-system coordination is fundamental to how the brain learns. While recent advances in machine learning have accelerated progress in neuroscience, these models can be inconsistent with brain function^{19,20} (but see⁹³). While considering parallels between machine learning models and humans may stimulate substantive research, machine learning models often fail to capture key aspects of human behavior and brain response, which is not surprising given that these models were developed for other purposes. For example, DNNs rely on stationary batches of training data and, unlike humans, they lack the ability to continually and rapidly adapt to changing environments. One use of the controller-peripheral framework is to repurpose engineering models to better suit neuroscience and perhaps in turn offer insights to machine learning. Rather than directly importing or refining models from machine learning, perhaps the controller-peripheral architecture can shift the emphasis in neuroscience to considering how multiple models or modules interrelate to develop more encompassing theories of brain function.

Methods

The peripheral DNN module

DNN architecture and fine-tuning

We instantiated the peripheral of our model under the controller-peripheral framework with a well-known deep convolutional neural network, VGG-16²², a feed-forward architecture consisting of millions of parameters pre-trained on 1.3 million real-world images from the ImageNet database⁹⁴. VGG-16 was originally trained to map real-world images to one-hot vectors across 1,000 pre-defined categories. In our work, this DNN module is to be integrated with a controller (a clustering module), which requires the VGG-16 to be fine-tuned such that the DNN module outputs three-dimensional vectors whose dimensions encode psychological features of the stimulus, such as “0” for thin legs and “1” for thick legs on the first dimension (see full mapping in Supplementary Information Table S.6). The fine-tuning process can be viewed as akin to familiarizing human participants to the experimental stimuli.

To fine-tune the DNN module appropriately, we preserved the layers of VGG-16 that are believed to correspond to regions along the ventral visual stream up to and including LOC⁵⁶. We then replaced layers succeeding “block4_pool” with a three-unit fully-connected layer (301, 056 connection weights; randomly initialized using glorot uniform distribution) whose output units correspond to psychological features of the stimulus. The new layer outputs are gated by a sigmoid function (Eq. 1) which squeezes each unit’s raw activation x_i (unbounded) to a_i .

$$a_i = \frac{1}{1 + \exp(-x_i)} \quad (1)$$

The position to add the new output layer is critical. Intuitively, we should place the new layer deep into the network so we preserve most hierarchical features of the DNN in order to parallel the DNN module to the ventral visual stream. However, not all stimulus features are represented at the same position of the DNN in the same capacity⁹⁵. While DNNs can diverge from important aspects of how humans process visual information, we found a similar pattern in human behaviors that suggest participants are more sensitive to some stimulus features than others (see Supplementary Information Fig. S.5).

In light of this trade-off, we chose the best layer position by evaluating whether the output layer can produce binary-valued stimulus representations under two training procedures. We used binary-valued representations as ground true targets because they represent idealized stimulus features prior to category learning without attentional modulation. We provide the mapping between pixel-level stimulus images and corresponding binary-valued features in Supplementary Information, Table S.6.

To prepare the input data for fine-tuning, we applied data augmentation to the original stimuli. For each stimulus, we applied a random combination of flipping (horizontal or vertical), rotation (0–45 degree) and shear (0–15 degree), determined by a unique random seed. We randomly chose 1,024 different seeds which resulted in 1,024 augmented samples per stimulus. For the first training procedure, we trained the output layer positioned at different locations of VGG-16 on a random 80% of the augmented samples and used the rest 20% for validation (early-stopping). We then tested each candidate DNN module on predicting psychological representations of the eight original stimuli. Each candidate module was scored based on the success rate of predicting binary-valued representations. The second training procedure provided a stricter evaluation for choosing the best layer position. We trained candidate DNN modules on augmented stimuli (same train and validation split ratio applied) but holding out samples of a particular stimulus. We repeated this procedure eight times holding out one type of stimulus at a time. We then tested candidate modules on predicting psychological representation of the corresponding original stimulus that was held out during training. This test is stronger in that it prevented the DNN module from memorizing (i.e., overfitting) the mapping from stimulus images to abstract psychological representations. Each candidate module was scored based on the success rate of predicting binary-valued representations of the held-out stimulus.

We repeated the two procedures over a range of learning rates ($3e-3$, $3e-4$, $3e-5$, $3e-6$) and selected the layer position based on overall performance on the two procedures. If there were two layers with the same performance, we chose the layer that is higher up in the network. All training used the Adam optimizer (with default hyper-parameter settings) and a batch size of 16. We trained each candidate module for 1,000 epochs unless the performance stopped improving on the validation set for 20 consecutive epochs, in which case training would be terminated early. We used the standard cross-entropy error as our loss function and stopping metric. We found “block4_pool” layer achieved the highest metric scores (Full results are in Supplementary Information, Tables S.7 and S.8).

Peripheral attention layer

We inserted a goal-directed attention layer between the “block4_pool” layer and the fine-tuned output layer of the fine-tuned DNN. The peripheral attention layer was implemented in the same fashion as⁶². Additionally, we applied L_1 regularization on the parameters (weights) of the attention layer. In accord with the costly-energy principle, this is to enforce task-driven sparsity over the attention weights. We defined attention modulation as the Hadamard product (filter-wise multiplication) between the preceding layer’s activations and the attention weights. Formally, we denote pre-attention activation for a given stimulus from a DNN layer as x_n , where $x_n \in \mathbb{R}^{H \times W \times F}$ (H and W are the spatial dimensions of the representation and F are the number of filters). We denote the corresponding attention weights as $g \in \mathbb{R}^F$. The attention modulation is then defined as:

$$\mathbf{x}_n^* = \mathbf{x}_n \odot \mathbf{g} \quad (2)$$

where $\mathbf{x}_n^*$ are the post-attention activations that will be passed onto the output layer. The attention layer is trained in conjunction with the clustering module with all attention weights initialized at one. We detailed the training procedure below (Section 6.3).

The controller clustering module

The controller of our proposed model is a clustering module, which follows and extends key principles of a successful category learning model, SUSTAIN⁸. The clustering module is a three-layer feed-forward network. The input layer has a number of nodes each represents one of the psychological features of the stimulus, output from the DNN module (in this case, three nodes). The activation of the i th input node is denoted a_i^{in} . A complete stimulus is denoted $\mathbf{a}^{in} = (a_1^{in}, a_2^{in}, \dots)^T$.

The hidden layer is initialized with no clusters and new clusters are recruited based on the difficulty of the task. For a given stimulus (output from the DNN), each cluster is activated according its psychological similarity to the stimulus captured by the following equation:

$$H_j^{act} = \exp(-c(\sum_i \alpha_i |h_{ji} - a_i^{in}|^r)^{q/r}) \quad (3)$$

where $|h_{ji} - a_i^{in}|^r$ is the dimensional distance between the center of the cluster and the stimulus. In this work, we set $r = 2$ and $q = 1$ (Euclidean distance). The dimensional attention strength α_i , acts as a multiplier on the corresponding dimension. Both the center of the cluster and the attention strength are trainable parameters of the clustering module. Initially, attention strength is equal across dimensions (initialized at $\frac{1}{3}$ for three-dimensional inputs). As learning proceeds, more attention will be allocated to the task relevant dimensions and less to the irrelevant dimensions. Attention weights are always non-negative and normalized to sum to one⁹⁶. Specificity parameter c is fixed (a hyper-parameter).

Clusters compete to respond to input patterns and in turn inhibit one another following

$$H_j^{act} = \frac{\exp(t \cdot H_j^{act})}{\sum_{i=1}^N \exp(t \cdot H_i^{act})} H_j^{act} \quad (4)$$

where t is an inverse temperature parameter. When t is large, inhibition is weak. Contrary to SUSTAIN's winner-take-all (WTA) scheme where only the activation of the most activated cluster is passed onto the output layer, we allow all clusters contribute to model decision subject to normalization:

$$H_j^{out} = \frac{\exp(u \cdot H_j^{act})}{\sum_{i=1}^N \exp(u \cdot H_i^{act})} H_j^{act} \quad (5)$$

where u is a decision parameter. Adjusting u can change how much the module's overall activity (output response) is dependent on a single cluster. When using an extremely large u , model decision reduces to WTA. It is worth noting that while Eqs. 4 and 5 share the same expression, they are intended to capture different processes the brain might implement.

Every cluster has association weights connected to the output layer, hence the activation of output layer unit k is denoted:

$$C_k^{out} = \sum_{j=1}^N w_{kj} H_j^{out} \quad (6)$$

Association weights are trainable parameters of the module, which are initialized from zero. Output activations are further converted to a probability response using

$$p_k = \frac{\exp(d \cdot C_k^{out})}{\sum_{i=1}^K \exp(d \cdot C_i^{out})} \quad (7)$$

where the probability of a given stimulus belonging to category k is the magnitude of output unit k 's activation, scaled by a real-valued decision parameter d relative to the sum of all K output units' activations (exponentiated).

Controller-peripheral learning framework

We train the model within the controller-peripheral learning framework. As the controller, the clustering module is first updated to optimize the global learning objective (i.e., categorization error). As the peripheral, the DNN module is then updated to optimize an intermediate learning objective instructed by the controller (see below). Both modules are iteratively optimized throughout learning. In a given trial, the DNN module receives an image stimulus and outputs a psychological representation of the stimulus for the clustering module. The clustering module receives the stimulus and completes a single learning step that involves *cluster recruitment* and *loss optimization*.

Cluster recruitment

The clustering module is initialized with no clusters and learning always begins with the module creating a new cluster centering on the first trial. In subsequent learning trials, cluster recruitment takes into account all clusters at once to the degree they are activated. Each cluster has a measure of “support” (i.e., consistency) for the correct response which is determined by the direction and magnitude of the association weights (see Eq. 9). Importantly, this measure of support is determined by the ratio of its association weights, as opposed to their absolute magnitude which would advantage older clusters with established association weights.

The *totalSupport* (−1 and 1 inclusive) of the current clustering, which determines whether a new cluster is recruited should *totalSupport* fall below some threshold parameter, is

$$\text{totalSupport} = \frac{\sum_i^N \text{support}_i \cdot H_i^{\text{out}}}{\sum_i^N H_i^{\text{out}}} \quad (8)$$

where H_i^{out} is the output of cluster i and support_i (−1 and 1 inclusive) is the support from cluster i defined as,

$$\text{support}_i = \frac{w_{i,\text{correct}} - w_{i,\text{incorrect}}}{|w_{i,\text{correct}}| + |w_{i,\text{incorrect}}|} \quad (9)$$

where $w_{i,\text{correct}}$ is the association weight from cluster i to the correct output unit and $w_{i,\text{incorrect}}$ is the association weight to the incorrect output (i.e., response) unit. If we only consider the most activated cluster, this recruitment rule reduces to the WTA procedure previously used in SUSTAIN⁸.

Loss optimization

After the cluster recruitment step, parameters of the clustering module, namely the association weights, attention weights and all cluster positions will be updated via gradient descent in order to minimize the global categorization loss as well as a regularization loss:

$$E_n = -\mathbb{E}_{k \in K} [y_k \log(p_k) + (1 - y_k) \log(1 - p_k)] + \gamma \sum_i \alpha_i \log(\alpha_i) \quad (10)$$

The first half of the loss is the cross-entropy error between a stimulus $\mathbf{y}_n = (y_1, y_2, \dots, y_k)^T$ and its prediction by the module $\mathbf{p}_n = (p_1, p_2, \dots, p_k)^T$. The second half of the loss is the entropy of the dimensional attention strength (weighted by a hyper-parameter γ). The entropy term encourages the model to develop selective (non-uniform) attention weights. This is in accordance with eye-tracking results that humans tend to optimize attention to only task diagnostic dimensions when solving Shepard et al.’s problems⁹⁷.

DNN module update

After the clustering module is updated, the DNN module will update to optimize learning objective determined by the clustering module with no direct access to the global categorization error. Specifically, DNN module is directed by the costly-energy principle to both reduce the number of non-zero peripheral attention weights (i.e., increasing sparsity) and to avoid disrupting existing cluster representations in the clustering module (computed by Eq. 3). We define the overall loss as follows:

$$L = \lambda \left[\frac{1}{N} \sum_{n=1}^N (\mathbf{H}_{\text{true};n}^{\text{out}} - \mathbf{H}_{\text{pred};n}^{\text{out}})^2 \right] + \|\mathbf{g}\| \quad (11)$$

where the first half is a reconstruction loss (mean-squared-error; weighted by λ) between the true outputs of all clusters $\mathbf{H}_{\text{true};n}^{\text{out}}$ before attention optimization begins at each trial and the predicted outputs of all clusters $\mathbf{H}_{\text{pred};n}^{\text{out}}$ given the same stimulus after attention optimization. The second half is a L_1 regularization loss on the peripheral attention weights $\mathbf{g}$ that encourages sparsity. The peripheral attention weights are updated using the entire stimulus set to avoid over-fitting to a single stimulus. We set a fixed number of iterations (a hyper-parameter) for this optimization within each trial.

Alternative models

To validate that both the controller-peripheral learning framework and costly-energy principle are necessary for the model to capture human performance and neural responses, we consider three alternative models that lack either or both elements (Fig. S.1A). We compare competing models based on two criteria. First, we evaluate whether the model is able to account for human learning performance in⁶⁶. Second, we evaluate whether the peripheral of the model exhibits patterns of resource expenditure that are in accord with those reported in the prior literature^{58,98,99}. Specifically, we test whether the sparsity of peripheral attention weights – reflecting less energy expenditure – increases with decreasing task difficulty, mirroring results shown in brain imaging studies⁵⁸.

Training procedures are identical across these models but they differ in terms of the optimization objectives. For Model 2, which follows the costly-energy principle but not the controller-peripheral framework, it is optimized to the following loss:

$$L = -\lambda\{\mathbb{E}_{k \in K}[y_k \log(p_k) + (1 - y_k) \log(1 - p_k)]\} + \|g\| \quad (12)$$

The only difference to Model 1 is that in Model 2, both the DNN module and the clustering module are updated to minimize the global categorization error (first half) in addition to the L_1 regularization loss on the peripheral attention weights (second half). For Model 3, which follows the controller-peripheral framework but not the costly-energy principle, it is optimized to:

$$L = \lambda\left[\frac{1}{N} \sum_{n=1}^N (\mathbf{H}_{true;n}^{out} - \mathbf{H}_{pred;n}^{out})^2\right] \quad (13)$$

The only difference to Model 1 is that Model 3 does not have the regularization term on the peripheral attention weights. For Model 4, which follows neither the controller-peripheral framework nor the costly-energy principle, it is optimized to the global categorization loss without the L_1 regularization over peripheral attention weights:

$$L = -\lambda\{\mathbb{E}_{k \in K}[y_k \log(p_k) + (1 - y_k) \log(1 - p_k)]\} \quad (14)$$

Human behavioral and neuroimaging studies

Stimulus set

We applied the model to account for human category learning behavior described in⁶⁶ and replicated in⁶⁷. We used insect stimuli created by¹¹ with the same category structures as the geometric shape stimuli used in the original study (Fig. 2A). The stimulus set consisted of insects with three binary features (thick/thin legs, thick/thin antennae, and pincer/shovel mouths). There are in total eight images representing all combinations of the three binary-valued features (Supplementary Information, Table S.6).

⁶⁶ described six learning tasks where participants learn to classify stimuli into two categories in each task and showed learning curves that revealed the difficulty ordering of the category structures. Type I was the easiest to master, followed by Type II, followed by Types III–V, and Type VI was the hardest. For solving Type I, only one stimulus dimension is relevant whereas two dimensions are relevant for solving Type II (i.e., XOR with an irrelevant dimension). All three dimensions are relevant in Types III–VI. There are a few regularities for Type III–V as they can be classified as rule-plus-exception problems, and solving Type VI requires memorizing all stimuli.

Modeling Shepard et al.,⁶⁶

We set out to evaluate whether our model can capture the classic learning behaviors from⁶⁶ and retain SUSTAIN's strategies in solving the six learning problems⁸. To that end, we trained our model to mirror the learning curves from⁶⁷, which is a replication of⁶⁶. We simulated our model in a trial-by-trial manner consistent with the procedures in the original experiment. Unlike the human results, which are averages of relatively small groups of individuals, we trained the model over 500 independent restarts (analogous to individual participants), each time with a different stimulus sequence. For each restart, the eight stimuli were presented in a randomized order for a total of 32 repetitions. To maintain consistency, we seeded each restart with a specific number, resulted in the same stimulus sequence across problems. We also counterbalanced feature-to-task mappings across restarts. To obtain learning curves to fit to the human data, we computed the probability of error (i.e., 1 – proportion correct) for each repetition over all stimuli and restarts.

Modeling Mack et al.,¹¹

To explore how the computational model's learning mechanism is implemented in the brain during category learning, we set out to fit our model to human category learning behavior and relate the model's internal representations to human fMRI data during category learning.

Building on⁶⁶'s paradigm,¹¹ applied a model-based fMRI approach focusing on how HPC and PFC are involved in concept learning. In their study, participants learned the Type I, II and VI problems whilst during an fMRI scan. All of the participants learned to perform the Type VI problem first, the order of the Types I and II problems was then counterbalanced across participants. Each problem type lasted four scanner runs. There were four blocks within each run and for each behavioral block there were right trials. During scanning, whole-brain images were acquired. Whole brain activation patterns for each stimulus within each run were estimated using an event-specific univariate GLM approach. For details about data preprocessing and GLM modeling,

see Methods (Section 6.5). For details about data acquisition and behavior experiments, we refer readers to the original paper¹¹.

We trained one instance of the model for each participant (with a unique stimulus sequence) in a trial-by-trial basis and fitted their learning curves independently. For each participant, we performed a hierarchical grid search to find the best hyper-parameters. The quality of fits was evaluated based on the average mean-square-error between learning curves of participants and models.

To evaluate whether our model captures learning mechanisms in the brain, we measured the correspondence between model representations to brain regions hypothesized to be involved in different aspects of category learning (vmPFC and LOC).

Model correspondence to vmPFC

To relate high-level clustering module of our model to vmPFC, we evaluated the correspondence between neural compression in vmPFC and attention strategy used by the clustering module during category learning. For each fitted model and a given task, we computed a block-wise attention compression score using the average attention weights within a block of eight trials with unique stimuli based on,

$$entropy = - \sum_{i=1}^3 \alpha_i \log_2 \alpha_i \quad (15)$$

$$compression = 1 + \frac{entropy}{\log_2(\frac{1}{3})} \quad (16)$$

Intuitively, the compression score is formulated as normalized entropy (bounded between 0 and 1 inclusive) indexing the dispersion of attention across stimulus dimensions. If the task complexity is high, requiring attention to multiple features, attention weights will be less selective, which will lead to a low compression score. If the task complexity is low, attention will be allocated to some features more than others, which will lead to a high compression score (see results in Fig. 4A). We tested the significance of the main effects (problem complexity and learning block) and their interaction with a two-way analysis of variance (ANOVA; see full results in Supplementary Information Table S.5). Additionally, we quantified the change of attention compression over time by fitting a linear regression model for each participant model using the compression score as the dependent variable and time (learning block) as the independent variable for each problem type. We further tested the significance of the regression coefficients with an one-sample t-test (see results in Sec. 4).

Model correspondence to LOC

To test whether there is a correspondence between the attention modulated output layer of the DNN and LOC, we evaluated whether stimulus information coding from the model layer is consistent with neural stimulus coding in LOC during category learning.

Stimulus information coding at the output layer of the DNN module is subject to distortion as a result of task-oriented attention in the clustering module, which in turn modulates the psychological representation of the stimulus. We expected that the representation of the irrelevant features will slowly degrade, leading to information loss. Therefore, stimuli with a low information loss would suggest stimulus information is largely preserved in the network which implies that most features are task relevant and can be reconstructed from the network. On the contrary, a high information loss would suggest most features are task irrelevant which can no longer be reconstructed. We quantified the information loss as the cross-entropy error between the stimulus representation before and after category learning. We computed the average information loss over participants per stimulus dimension and per task (see results in Fig. 4A-B)

In the model, information loss can be directly computed using network activities in that we have access to both stimulus representations before and after learning. In the brain, however, such a direct measure is not available. Therefore, we used multivariate pattern analysis (MVPA) to we determined how linearly separable (i.e., confusable) two neural activity patterns are in LOC for every pair of stimuli in each task (decoding error; 1 - decoding accuracy). Intuitively, neural patterns of stimulus pairs that differ by a task-relevant dimension should be more easily separable (less confusable) than stimulus pairs that differ by a task-irrelevant dimension in that information about task relevant dimensions should be better preserved in the brain, akin to less information loss in the model. Information loss should be the lowest for Type VI, as most stimulus pairs differ by at least one relevant dimension (as all features are relevant), and it is desirable to maintain information about relevant features. On the contrary, information loss in Type I should be the highest as many stimulus pairs will differ by an irrelevant dimension (having one relevant and two irrelevant features) and it is not necessary to maintain information irrelevant to the task. We trained linear support vector classifiers on neural activity patterns of stimulus pairs in LOC (SVC; $C = 0.1$; using the Scikit-learn python package;¹⁰⁰). Neural activity patterns for each stimulus within each scanner run were estimated using an event-specific univariate GLM approach (see Sec. 6.5 for details). We fitted the support vector classifiers using a three-fold cross-validation procedure (the first run was excluded from training because learning at the start could be unstable). We computed the average decoding error (1 - classifier accuracy) over stimuli pairs and participants across tasks.

To test the prediction that information loss in both brain and model would linearly scale with the amount of stimulus information required to solve the task (i.e., the number of dimensions relevant to the task), we performed a linear regression analysis on information loss as the dependent variable and the number of relevant dimensions

per task as the independent variables. This was done separately for decoding error in LOC, and information loss in the participant-fitted model. We then performed one-sample t-test over regression coefficients obtained for participants and models separately. A significant downward trend in both cases would support our prediction.

Furthermore, we computed the percentage of zeroed out attention weights on the peripheral attention layer, as a measure of energy efficiency and related it to neural dimensionality estimated in LOC. We assessed the average percentage of zeroed out attention weights over participants and over tasks at the end of learning. We predicted that the dimensionality of attention weights should linearly scale with the amount of stimulus information (i.e., feature dimensions) relevant to the task. To verify such a relationship, we again performed linear regression analysis as above, but using percentage of zero attention weights as the dependent variable and the number of relevant dimensions per task as the independent variables. We performed linear regression on each behaviour-fitted model and performed a one-sample t-test over the regression coefficients to test the significance of the predicted relationship (see results in Sec. 4).

fMRI data processing

Results included in this manuscript come from preprocessing performed using fMRIPrep 21.0.1 (^{101,102}; RRID:SCR_016216), which is based on Nipype;^{103,104}; RRID:SCR_002502).

Anatomical data preprocessing

A total of 22 T1-weighted (T1w) images were found within the input BIDS dataset. The T1-weighted (T1w) image was corrected for intensity non-uniformity (INU) with N4BiasFieldCorrection¹⁰⁵, distributed with ANTs 2.3.3 (¹⁰⁶; RRID:SCR_004757), and used as T1w-reference throughout the workflow. The T1w-reference was then skull-stripped with a Nipype implementation of the antsBrainExtraction.sh workflow (from ANTs), using OASIS30ANTs as target template. Brain tissue segmentation of cerebrospinal fluid (CSF), white-matter (WM) and gray-matter (GM) was performed on the brain-extracted T1w using fast (FSL 6.0.5.1:57b01774, RRID:SCR_002823,¹⁰⁷). Volume-based spatial normalization to one standard space (MNI152NLin2009cAsym) was performed through nonlinear registration with antsRegistration (ANTs 2.3.3), using brain-extracted versions of both T1w reference and the T1w template. The following template was selected for spatial normalization: ICBM 152 Nonlinear Asymmetrical template version 2009c¹⁰⁸; RRID:SCR_008796; TemplateFlow ID: MNI152NLin2009cAsym].

Functional data preprocessing

For each of the 12 BOLD runs found per subject (across all tasks and sessions), the following preprocessing was performed. First, a reference volume and its skull-stripped version were generated using a custom methodology of fMRIPrep. Head-motion parameters with respect to the BOLD reference (transformation matrices, and six corresponding rotation and translation parameters) are estimated before any spatiotemporal filtering using mcflirt (FSL 6.0.5.1:57b01774,¹⁰⁹). The BOLD time-series (including slice-timing correction when applied) were resampled onto their original, native space by applying the transforms to correct for head-motion. These resampled BOLD time-series will be referred to as preprocessed BOLD in original space, or just preprocessed BOLD. The BOLD reference was then co-registered to the T1w reference using mri_coreg (FreeSurfer) followed by flirt (FSL 6.0.5.1:57b01774,¹¹⁰) with the boundary-based registration (Greve and Fischl 2009) cost-function. Co-registration was configured with six degrees of freedom. Several confounding time-series were calculated based on the preprocessed BOLD: framewise displacement (FD), DVARS and three region-wise global signals. FD was computed using two formulations following Power (absolute sum of relative motions,¹¹¹) and Jenkinson (relative root mean square displacement between affines,¹⁰⁹). FD and DVARS are calculated for each functional run, both using their implementations in Nipype (following the definitions by Power et al. 2014). The three global signals are extracted within the CSF, the WM, and the whole-brain masks. Additionally, a set of physiological regressors were extracted to allow for component-based noise correction (CompCor,¹¹²). Principal components are estimated after high-pass filtering the preprocessed BOLD time-series (using a discrete cosine filter with 128s cut-off) for the two CompCor variants: temporal (tCompCor) and anatomical (aCompCor). tCompCor components are then calculated from the top 2% variable voxels within the brain mask. For aCompCor, three probabilistic masks (CSF, WM and combined CSF+WM) are generated in anatomical space. The implementation differs from that of¹¹² in that instead of eroding the masks by 2 pixels on BOLD space, the aCompCor masks are subtracted a mask of pixels that likely contain a volume fraction of GM. This mask is obtained by thresholding the corresponding partial volume map at 0.05, and it ensures components are not extracted from voxels containing a minimal fraction of GM. Finally, these masks are resampled into BOLD space and binarized by thresholding at 0.99 (as in the original implementation). Components are also calculated separately within the WM and CSF masks. For each CompCor decomposition, the k components with the largest singular values are retained, such that the retained components' time series are sufficient to explain 50 percent of variance across the nuisance mask (CSF, WM, combined, or temporal). The remaining components are dropped from consideration. The head-motion estimates calculated in the correction step were also placed within the corresponding confounds file. The confound time series derived from head motion estimates and global signals were expanded with the inclusion of temporal derivatives and quadratic terms for each¹¹³. Frames that exceeded a threshold of 0.5 mm FD or 1.5 standardised DVARS were annotated as motion outliers. All resamplings can be performed with a single interpolation step by composing all the pertinent transformations (i.e. head-motion transform matrices, susceptibility distortion correction when available, and co-registrations to anatomical and output spaces). Gridded (volumetric) resamplings were performed using antsApplyTransforms (ANTs), configured with Lanczos interpolation to minimize the smoothing effects of other kernels¹¹⁴. Non-gridded (surface) resamplings were performed using mri_vol2surf (FreeSurfer).

Many internal operations of fMRIPrep use Nilearn 0.8.1 (¹¹⁵, RRID:SCR_001362), mostly within the functional processing workflow.

fMRI general linear model

We used the general linear model (GLM) in using NiPype¹⁰³, using SPM functions (SPM version 12;¹¹⁶), to obtain estimates of the task-evoked signals for the multivariate pattern analyses (MVPA) where we quantified information loss across category structures. Whole brain activation patterns for each stimulus within each run were estimated using an event-related GLM. For each scan run, we included a GLM with one explanatory variable (EV) for each of the eight stimuli, modelled as 3.5-s boxcar convolved with a canonical hemodynamic response function (HRF) to extract voxel-wise parameter estimates to each of the stimuli. Stimulus EVs for the feedback stimulus (2-s boxcar) and six motion parameters were also included in the GLM (not used in subsequent analyses). This resulted in whole brain activation patterns for each participant for all stimuli across Type I, II and VI learning problems. We conducted a GLM for each run separately for leave-one-run-out cross-validation for MVPA. No spatial smoothing was applied.

Stimulus EVs were 3.5s with stimulus-feedback intervals ranging from 0.5–4.5 (jittered). A feedback was shown for 2s followed by a 4–8s fixation. There were 4 repetitions of each stimulus (eight stimuli) in each scan run (for full details of the task, see¹¹).

Regions of interest

We tested the correspondence between the novel aspects of our model and the brain. Specifically, we hypothesized that activity of the DNN's attention-layer modulated output would correspond to LOC, higher-level visual regions later in the ventral stream. Therefore, we focused the LOC as a region of interest (ROI).

The LOC anatomical masks were taken from¹¹⁷. The masks are provided in T1 structural MRI space (1-mm3), and when transformed into individual functional space (3-mm3), some gray matter voxels are excluded. Therefore, minor smoothing was applied to the T1 mask (Gaussian kernel of 0.2 mm, using fslmaths) for a liberal inclusion of neighboring voxels before transforming to functional space. Left and right masks were smoothed and merged using ANTs.

Data availability

The code for the model and data analysis are publicly available at https://github.com/don-tpanic/brain_data and <https://github.com/don-tpanic/CostlyEnergyPrinciple>. Human participant data (behavioral and fMRI) are available at <https://osf.io/5byhb/>

Received: 10 October 2024; Accepted: 7 February 2025

Published online: 13 February 2025

References

- Locke, J. An essay concerning human understanding. *Read. Hist. Psychol.* **1**, 55–68 (1960).
- Hume, D. An enquiry concerning human understanding. *Essays and treatises on several subjects*, Vol 2: *Containing An enquiry concerning human understanding, A dissertation on the passions, An enquiry concerning the principles of morals, and The natural history of religion*. 3–212 (1779). /record/2008-16196-001.
- Kieras, D. E. An overview of the EPIC architecture for cognition and performance with application to human-computer interaction. *Human-Comput. Interact.* **12**, 391–438 (1997) (/record/1998-00672-004.).
- Young, R. M. & Lewis, R. L. The Soar Cognitive Architecture and Human Working Memory. *Models Work. Memory* 224–256 (1999). /record/1999-02490-006
- Anderson, J. R., Michael, M. & Lebiere, C. ACT-R: A theory of higher level cognition and its relation to visual attention. *Human-Comput. Interact.* **12**, 439–462 (1997).
- Love, B. C. Model-based fMRI analysis of memory. *Curr. Opin. Behav. Sci.* **32**, 88–93 (2020).
- Seger, C. A. & Miller, E. K. Category learning in the brain. *Ann. Rev. Neurosci.* **33**, 203–219 (2010) <https://pubmed.ncbi.nlm.nih.gov/20572771/>.
- Love, B. C., Medin, D. L. & Gureckis, T. M. SUSTAIN: A network model of category learning (2004).
- Love, B. C. & Gureckis, T. M. Models in search of a brain. *Cognit. Affect. Behav. Neurosci.* **7**, 90–108. <https://doi.org/10.3758/CA.BN.7.2.90> (2007).
- Davis, T., Love, B. C. & Preston, A. R. Learning the exception to the rule: Model-based FMRI reveals specialized representations for surprising category members. *Cerebral Cortex (New York, N.Y.: 1991)* **22**, 260–273 (2012). <https://pubmed.ncbi.nlm.nih.gov/21666132/>.
- Mack, M. L., Love, B. C. & Preston, A. R. Dynamic updating of hippocampal object representations reflects new conceptual knowledge. *Proc. Natl. Acad. Sci. United States Am.* **113**, 13203–13208 (2016).
- Mack, M. L., Preston, A. R. & Love, B. C. Ventromedial prefrontal cortex compression during concept learning. *Nat. Commun.* **11**, 46 (2020).
- Mok, R. M. & Love, B. C. A non-spatial account of place and grid cells based on clustering models of concept learning. *Nat. Commun.* **10**, 5685 (2019).
- Mok, R. M. & Love, B. C. A multi-level account of hippocampal function from behaviour to neurons. *bioRxiv* 2022.06.09.495367 (2022). <https://doi.org/10.1101/2022.06.09.495367v1>
- Geirhos, R. et al. Generalisation in humans and deep neural networks. *Adv. Neural. Inf. Process. Syst.* **2018**, 7538–7550 (2018).
- Bracci, S., Ritchie, J. B., Kalfas, I. & Op de Beeck, H. P. The Ventral Visual Pathway Represents Animal Appearance over Animacy, Unlike Human Behavior and Deep Neural Networks. *J. Neurosci.* **39**, 6513–6525 (2019).
- Kietzmann, T. C. et al. Recurrence is required to capture the representational dynamics of the human visual system. *Proc. Natl. Acad. Sci. USA*. **116**, 21854–21863. <https://doi.org/10.1073/pnas.1905544116> (2019).
- Buckner, C. Understanding adversarial examples requires a theory of artefacts for deep learning. *Nat. Mach. Intell.* **2**, 731–736 (2020).
- Saxe, A., Nelli, S. & Summerfield, C. If deep learning is the answer, what is the question?. *Nat. Rev. Neurosci.* **22**, 55–67 (2020).
- Bowers, J. S. et al. Deep problems with neural network models of human vision. *PsyArxiv* (2022). <https://psyarxiv.com/5zf4s/>.

21. Krizhevsky, A., Sutskever, I. & Hinton, G. E. ImageNet classification with deep convolutional neural networks. *Adv. Neural Inf. Process. Syst.* **2**, 1097–1105 (2012).
22. Simonyan, K. & Zisserman, A. Very deep convolutional networks for large-scale image recognition. In Yoshua Bengio and Yann LeCun (ed.) *International Conference on Learning Representations* (2015).
23. Khaligh-Razavi, S.-M. & Kriegeskorte, N. Deep supervised, but not unsupervised, models may explain IT cortical representation. *PLoS Comput. Biol.* **10**, e1003915. <https://doi.org/10.1371/journal.pcbi.1003915> (2014).
24. 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).
25. Güçlü, U. & van Gerven, M. A. J. Deep neural networks reveal a gradient in the complexity of neural representations across the ventral stream. *J. Neurosci.* **35**, 10005–10014 (2015).
26. Eickenberg, M., Gramfort, A., Varoquaux, G. & Thirion, B. Seeing it all: Convolutional network layers map the function of the human visual system. *Neuroimage* **152**, 184–194 (2017).
27. Zeman, A. A., Ritchie, J. B., Bracci, S. & Op de Beeck, H. Orthogonal representations of object shape and category in deep convolutional neural networks and human visual cortex. *Sci. Rep.* **10**, 1–12 (2020).
28. Richler, J. J. & Palmeri, T. J. Visual category learning. *Wiley Interdiscip. Rev.: Cognit. Sci.* **5**, 75–94 (2014).
29. DiCarlo, J. J., Zoccolan, D. & Rust, N. C. How does the brain solve visual object recognition? *Neuron* **73**, 415–434 (2012).
30. Parisi, G. I., Kemker, R., Part, J. L., Kanan, C. & Wermter, S. Continual lifelong learning with neural networks: A review. *Neural Netw.* **113**, 54–71. <https://doi.org/10.1016/j.neunet.2019.01.012> (2018).
31. Guest, O. & Love, B. C. Levels of representation in a deep learning model of categorization. *bioRxiv* 626374 (2019).
32. Tartaglini, A. R., Vong, W. K. & Lake, B. M. Modeling artificial category learning from pixels: Revisiting Shepard, Hovland, and Jenkins (1961) with deep neural networks. *PsyArxiv* (2021). <https://psyarxiv.com/jh3xu/>.
33. Battleday, R. M., Peterson, J. C. & Griffiths, T. L. Capturing human categorization of natural images by combining deep networks and cognitive models. *Nat. Commun.* **11**, 1–14. <https://doi.org/10.1038/s41467-020-18946-z> (2020).
34. Miller, E. K. & Buschman, T. J. Cortical circuits for the control of attention. *Curr. Opin. Neurobiol.* **23**, 216–222 (2013).
35. Corbetta, M. & Shulman, G. L. Control of goal-directed and stimulus-driven attention in the brain. *Nat. Rev. Neurosci.* **3**, 201–215 (2002).
36. Kanwisher, N. & Wojciulik, E. Visual attention: Insights from brain imaging. *Nat. Rev. Neurosci.* **1**, 91–100 (2000).
37. Mangun, G. R. The neuroscience of attention: Attentional control and selection. *Neurosci. Attention: Attentional Control Select.* 1–296 (2012). <https://academic.oup.com/book/3395>.
38. Mok, R. M., O'Donoghue, M. C., Myers, N. E., Drazich, E. H. & Nobre, A. C. Neural markers of category-based selective working memory in aging. *Neuroimage* **194**, 163–173 (2019).
39. Mok, R. M. & Love, B. C. Abstract neural representations of category membership beyond information coding stimulus or response. *J. Cognit. Neurosci.* **34**(10), 1719–1735 (2021).
40. Pessoa, L., Kastner, S. & Ungerleider, L. G. Neuroimaging studies of attention: From modulation of sensory processing to top-down control. *J. Neurosci.* **23**, 3990–3998 (2003).
41. Driver, J. & Vuilleumier, P. Perceptual awareness and its loss in unilateral neglect and extinction. *Cognition* **79**, 39–88 (2001).
42. Mesulam, M. M. A cortical network for directed attention and unilateral neglect. *Ann. Neurol.* **10**, 309–325 (1981).
43. Braunlich, K. & Love, B. C. Occipitotemporal representations reflect individual differences in conceptual knowledge. *J. Exp. Psychol. Gen.* **148**, 1192–1203 (2019).
44. Sučević, J. & Schapiro, A. C. A neural network model of hippocampal contributions to category learning. *eLife* **12**, e77185 (2023).
45. Gailliot, M. T. & Baumeister, R. F. The physiology of willpower: linking blood glucose to self-control. *Personal. Soc. Psychol. Rev.: Off. J. Soc. Personal. Soc. Psychol.* **11**, 303–327 (2007). <https://pubmed.ncbi.nlm.nih.gov/18453466/>.
46. Gailliot, M. T. et al. Self-control relies on glucose as a limited energy source: Willpower is more than a metaphor. *J. Pers. Soc. Psychol.* **92**, 325–336 (2007).
47. Shenhav, A. et al. Toward a rational and mechanistic account of mental effort. *Annu. Rev. Neurosci.* **40**, 99–124. <https://doi.org/10.1146/annurev-neuro-072116-031526> (2017).
48. Kurzban, R. Does the brain consume additional glucose during self-control tasks? *Evol. Psychol.* **8**, 244–259. <https://doi.org/10.1177/147470491000800208> (2010).
49. Kurzban, R., Duckworth, A., Kable, J. W. & Myers, J. An opportunity cost model of subjective effort and task performance. *Behav. Brain Sci.* **36**, 661–679 (2013).
50. Sagiv, Y., Musslick, S., Niv, Y. & Cohen, J. D. Efficiency of learning vs. processing: Towards a normative theory of multitasking. In *Proceedings of the 40th Annual Meeting of the Cognitive Science Society* (2020). [arXiv:2007.03124v1](https://arxiv.org/abs/2007.03124v1).
51. Musslick, S. & Cohen, J. D. Rationalizing constraints on the capacity for cognitive control. *Trends Cogn. Sci.* **25**, 757–775 (2021).
52. Lewis, J. W. & Van Essen, D. C. Corticocortical connections of visual, sensorimotor, and multimodal processing areas in the parietal lobe of the macaque monkey. *J. Comp. Neurol.* **428**, 112–137 (2000).
53. Jackson, R. L., Rogers, T. T. & Lambon Ralph, M. A. Reverse-engineering the cortical architecture for controlled semantic cognition. *Nat. Human Behav.* **5**, 774 (2021).
54. Schrimpf, M. et al. Brain-Score: Which artificial neural network for object recognition is most brain-like? *bioRxiv* (2018). <https://doi.org/10.1101/407007>.
55. Roads, B. D. & Love, B. C. Enriching ImageNet with human similarity judgments and psychological embeddings. *arXiv* (2020). [arXiv:2011.11015](https://arxiv.org/abs/2011.11015).
56. Xu, Y. & Vaziri-Pashkam, M. Limits to visual representational correspondence between convolutional neural networks and the human brain. *Nat. Commun.* **12**, 2065 (2021).
57. Folstein, J. R., Palmeri, T. J. & Gauthier, I. Category learning increases discriminability of relevant object dimensions in visual cortex. *Cereb. Cortex* **23**, 814–823 (2013).
58. Ahlheim, C. & Love, B. C. Estimating the functional dimensionality of neural representations. *Neuroimage* **179**, 51–62 (2018).
59. Grill-Spector, K., Kourtzi, Z. & Kanwisher, N. The lateral occipital complex and its role in object recognition. *Vision. Res.* **41**, 1409–1422 (2001).
60. Dijkstra, N., Zeidman, P., Ondobaka, S., Van Gerven, M. A. & Friston, K. Distinct top-down and bottom-up brain connectivity during visual perception and imagery. *Sci. Rep.* **7**, 1–9. <https://doi.org/10.1038/s41598-017-05888-8> (2017).
61. Stokes, M., Thompson, R., Cusack, R. & Duncan, J. Top-down activation of shape-specific population codes in visual cortex during mental imagery. *J. Neurosci.* **29**, 1565–1572 (2009).
62. Luo, X., Roads, B. D. & Love, B. C. The costs and benefits of goal-directed attention in deep convolutional neural networks. *Comput. Brain Behav.* **4**(2), 213–230 (2021).
63. Olshausen, B. A. & Field, D. J. Sparse coding of sensory inputs. *Curr. Opin. Neurobiol.* **14**, 481–487 (2004).
64. Vinje, W. E. & Gallant, J. L. Sparse coding and decorrelation in primary visual cortex during natural vision. *Science* **287**, 1273–1276 (2000).
65. Yoshida, T. & Ohki, K. Natural images are reliably represented by sparse and variable populations of neurons in visual cortex. *Nat. Commun.* **11**(1), 872 (2020).
66. Shepard, R. N., Hovland, C. I. & Jenkins, H. M. Learning and memorization of classifications. *Psychol. Monogr. Gen. Appl.* **75**, 1–42 (1961).

67. Nosofsky, R. M., Gluck, M. A., Palmeri, T. J., Mckinley, S. C. & Glauthier, P. Comparing modes of rule-based classification learning: A replication and extension of Shepard, Hovland, and Jenkins (1961). *Memory Cognit.* **22**, 352–369 (1994).
68. Dosovitskiy, A. et al. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. *CoRRabs/2010.1* (2020). [arXiv:2010.11929v2](https://arxiv.org/abs/2010.11929v2).
69. Wilson, R. C., Takahashi, Y. K., Schoenbaum, G. & Niv, Y. Orbitofrontal cortex as a cognitive map of task space. *Neuron* **81**, 267–279 (2014).
70. Mante, V., Sussillo, D., Shenoy, K. V. & Newsome, W. T. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* **503**, 78 (2013).
71. Zeithamova, D., Dominick, A. L. & Preston, A. R. Hippocampal and ventral medial prefrontal activation during retrieval-mediated learning supports novel inference. *Neuron* **75**, 168–179 (2012).
72. Schlichting, M. L., Mumford, J. A. & Preston, A. R. Learning-related representational changes reveal dissociable integration and separation signatures in the hippocampus and prefrontal cortex. *Nat. Commun.* **2015**(6), 1–10 (2015).
73. Chan, S. C., Niv, Y. & Norman, K. A. A probability distribution over latent causes, in the orbitofrontal cortex. *J. Neurosci.* **36**, 7817–7828 (2016).
74. Schuck, N. W., Cai, M. B., Wilson, R. C. & Niv, Y. Human orbitofrontal cortex represents a cognitive map of state space. *Neuron* **91**, 1402–1412 (2016).
75. Anderson, J. R. The adaptive nature of human categorization. (1991).
76. Sanborn, A. N., Griffiths, T. L. & Navarro, D. J. Rational approximations to rational models: Alternative algorithms for category learning. *Psychol. Rev.* **117**, 1144–1167 (2010).
77. Braunlich, K. & Love, B. C. Bidirectional influences of information sampling and concept learning. *Psychol. Rev.* **129**, 213–234 (2022).
78. Vargha-Khadem, F. et al. Differential effects of early hippocampal pathology on episodic and semantic memory. *Science* **277**, 376–380 (1997).
79. Gardiner, J. M., Brandt, K. R., Baddeley, A. D., Vargha-Khadem, F. & Mishkin, M. Charting the acquisition of semantic knowledge in a case of developmental amnesia. *Neuropsychologia* **46**, 2865–2868 (2008).
80. Blumenthal, A. et al. Abnormal semantic knowledge in a case of developmental amnesia. *Neuropsychologia* **102**, 237–247. <https://doi.org/10.1016/j.neuropsychologia.2017.06.018> (2017).
81. Lambon-Ralph, M. A., Jefferies, E., Patterson, K. & Rogers, T. T. The neural and computational bases of semantic cognition. *Nat. Rev. Neurosci.* **18**, 42–55. <https://doi.org/10.1038/nrn.2016.150> (2017).
82. Rougier, N. P., Noelle, D. C., Braver, T. S., Cohen, J. D. & O'Reilly, R. C. Prefrontal cortex and flexible cognitive control: Rules without symbols. *Proc. Natl. Acad. Sci. USA.* **102**, 7338–7343 (2005).
83. Craver, C. F. *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience* (Oxford University Press, Clarendon Press, 2007). <https://philpapers.org/rec/CRAETB>.
84. Love, B. C. Levels of biological plausibility. *Philos. Trans. R. Soc. B: Biol. Sci.* **376**, 20190632. <https://doi.org/10.1098/rstb.2019.0632> (2021).
85. Perich, M. G. & Rajan, K. Rethinking brain-wide interactions through multi-region ‘network of networks’ models. *Curr. Opin. Neurobiol.* **65**, 146–151 (2020).
86. Sexton, N. J. & Love, B. C. Reassessing hierarchical correspondences between brain and deep networks through direct interface. *Sci. Adv.* **8**, 2219. <https://doi.org/10.1126/sciadv.abm2219> (2022).
87. Hutchins, J. & Barger, S. Why neurons die: Cell death in the nervous system. *Anat. Rec.* **253**, 79–90 (1998).
88. Frankle, J. & Carbin, M. The lottery ticket hypothesis: Finding sparse, trainable neural networks. *7th International Conference on Learning Representations, ICLR 2019* (2018). [arXiv:1803.03635v5](https://arxiv.org/abs/1803.03635v5).
89. Ali, A., Ahmad, N., de Groot, E., Johannes van Gerven, M. A. & Kietzmann, T. C. Predictive coding is a consequence of energy efficiency in recurrent neural networks. *Patterns* **3**, 100639 (2022). <https://linkinghub.elsevier.com/retrieve/pii/S2666389922002719>.
90. Friston, K. The free-energy principle: A unified brain theory?. *Nat. Rev. Neurosci.* **11**, 127–138 (2010).
91. Cao, R. New labels for old ideas: Predictive processing and the interpretation of neural signals. *Rev. Philos. Psychol.* **11**(3), 517–546. <https://doi.org/10.1007/s13164-020-00481-x> (2020).
92. Lecun, Y. A path towards autonomous machine intelligence. *OpenReview* (2022).
93. Love, B. C. & Mok, R. M. You can't play 20 questions with nature and win redux. *Behav. Brain Sci.* **46**, e402 (2023).
94. Deng, J. et al. ImageNet: A large-scale hierarchical image database. In *CVPR (IEEE)*, 2009).
95. Zeiler, M. D. & Fergus, R. Visualizing and understanding convolutional networks. In *ECCV, Part I, LNCS 8689*, 818–833 (Springer International Publishing Switzerland, 2014).
96. Kruschke, J. K. ALCOVE: A connectionist model of human category learning. *Psychol. Rev.* **99**, 22–44 (1992).
97. Rehder, B. & Hoffman, A. B. Thirty-something categorization results explained: Selective attention, eyetracking, and models of category learning. *J. Exp. Psychol. Learn. Mem. Cogn.* **31**, 811–829 (2005).
98. Reddy, L. & Kanwisher, N. Coding of visual objects in the ventral stream. *Curr. Opin. Neurobiol.* **16**, 408–414 (2006).
99. Ostwald, D., Lam, J. M., Li, S. & Kourtzi, Z. Neural coding of global form in the human visual cortex. *J. Neurophysiol.* **99**, 2456–2469. <https://doi.org/10.1152/jn.01307.2007> (2008).
100. Pedregosa, F. F. et al. Scikit-learn: Machine Learning in Python Gaël Varoquaux Bertrand Thirion Vincent Dubourg Alexandre Passos PEDREGOSA, VAROQUAUX, GRAMFORT ET AL. Matthieu Perrot. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011). <http://scikit-learn.sourceforge.net>
101. Esteban, O. et al. fMRIPrep: A robust preprocessing pipeline for functional MRI. *Nat. Methods* **16**, 111–116 (2019).
102. Esteban, O. et al. fMRIPrep (2018). <https://zenodo.org/record/6928849>.
103. Gorgolewski, K. J. et al. Nipype (2018). <https://zenodo.org/record/6834519>.
104. Gorgolewski, K. et al. Nipype: A flexible, lightweight and extensible neuroimaging data processing framework in Python. *Front. Neuroinform.* **5**, 13 (2011).
105. Tustison, N. J. et al. N4ITK: Improved N3 bias correction. *IEEE Trans. Med. Imaging* **29**, 1310–1320 (2010).
106. Avants, B. B., Epstein, C. L., Grossman, M. & Gee, J. C. Symmetric diffeomorphic image registration with cross-correlation: Evaluating automated labeling of elderly and neurodegenerative brain. *Med. Image Anal.* **12**, 26–41 (2008).
107. Zhang, Y., Brady, M. & Smith, S. Segmentation of brain MR images through a hidden Markov random field model and the expectation-maximization algorithm. *IEEE Trans. Med. Imaging* **20**, 45–57 (2001).
108. Fonov, V. et al. Unbiased average age-appropriate atlases for pediatric studies. *Neuroimage* **54**, 313 (2011).
109. Jenkinson, M., Bannister, P., Brady, M. & Smith, S. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage* **17**, 825–841 (2002).
110. Jenkinson, M. & Smith, S. A global optimisation method for robust affine registration of brain images. *Med. Image Anal.* **5**, 143–156 (2001).
111. Power, J. D. et al. Methods to detect, characterize, and remove motion artifact in resting state fMRI. *Neuroimage* **84**, 320–341 (2014).
112. Behzadi, Y., Restom, K., Liu, J. & Liu, T. T. A component based noise correction method (CompCor) for BOLD and perfusion based fMRI. *Neuroimage* **37**, 90–101 (2007).
113. Satterthwaite, T. D. et al. Functional maturation of the executive system during adolescence. *J. Neurosci.* **33**, 16249–16261 (2013).

114. Lanczos & C. A precision approximation of the gamma function. *SJNAI*, 86–96 (1964). <https://ui.adsabs.harvard.edu/abs/1964SJNA.....1...86L/abstract>.
115. Abraham, E. et al. Father's brain is sensitive to childcare experiences. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 9792–9797 (2014).
116. Penny, W., Friston, K., Ashburner, J., Kiebel, S. & Nichols, T. *Statistical Parametric Mapping: The Analysis of Functional Brain Images* (The Analysis of Functional Brain Images, Statistical Parametric Mapping, 2007).
117. Wang, L., Mruczek, R. E., Arcaro, M. J. & Kastner, S. Probabilistic maps of visual topography in human cortex. *Cerebral cortex* (New York, N.Y.: 1991) **25**, 3911–3931 (2015). <https://pubmed.ncbi.nlm.nih.gov/25452571/>.

Acknowledgements

This work was supported by ESRC (ES/W007347/1), Wellcome Trust (WT106931MA), and a Royal Society Wolfson Fellowship (18302) to B.C.L., and the Medical Research Council UK (MC_UU_00030/7) and a Leverhulme Trust Early Career Fellowship (Leverhulme Trust, Isaac Newton Trust: SUAI/053 G100773, SUAI/056 G105620, ECF-2019-110) to R.M.M.

Additional information

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1038/s41598-025-89732-4>.

Correspondence and requests for materials should be addressed to X.L.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2025