

A sensory–motor theory of the neocortex

Received: 21 June 2023

Accepted: 26 April 2024

Published online: 27 June 2024

 Check for updates

Rajesh P. N. Rao ^{1,2} ✉

Recent neurophysiological and neuroanatomical studies suggest a close interaction between sensory and motor processes across the neocortex. Here, I propose that the neocortex implements active predictive coding (APC): each cortical area estimates both latent sensory states and actions (including potentially abstract actions internal to the cortex), and the cortex as a whole predicts the consequences of actions at multiple hierarchical levels. Feedback from higher areas modulates the dynamics of state and action networks in lower areas. I show how the same APC architecture can explain (1) how we recognize an object and its parts using eye movements, (2) why perception seems stable despite eye movements, (3) how we learn compositional representations, for example, part–whole hierarchies, (4) how complex actions can be planned using simpler actions, and (5) how we form episodic memories of sensory–motor experiences and learn abstract concepts such as a family tree. I postulate a mapping of the APC model to the laminar architecture of the cortex and suggest possible roles for cortico–cortical and cortico–subcortical pathways.

The predictive coding theory of cortical function, proposed in this journal in 1999 by the author and Ballard¹, has been the subject of increasing attention in recent years^{2–4}. However, as originally proposed, the theory ignored a fundamental aspect of perception, namely, that perception is action based: we move our eyes about three times a second to recognize objects in a scene, orient our heads to localize sounds, use our fingers to identify objects by touch and navigate ourselves in our environment to solve tasks that satisfy our needs. Away from experimentally imposed constraints in the laboratory, perception, in its natural state, can best be viewed as an action-based hypothesis-testing process (also called active sensing, active perception and active inference)^{5–8}.

Recent studies have highlighted the important influence of impending actions across almost all cortical areas (Fig. 1). For example, in mice solving a visual discrimination task using forepaws to rotate a wheel, Zátka-Haas et al.⁹ observed, using widefield calcium imaging, extensive bilateral activity across cortical areas preceding movements on choice trials (left or right action selected) but not on ‘NoGo’ trials (no action is selected) (Fig. 1a, left). They further showed that impending movement could be decoded from cortical activity in most imaged regions by 25 ms before movement (Fig. 1a, right). In the same task, Steinmetz et al.¹⁰ used Neuropixels probes to record spiking activity from thousands of neurons and showed that, not

only does activity in the visual cortex get updated after movement (Fig. 1b, left), but almost all recorded cortical areas had neurons with activities that were predictive of upcoming movements (Fig. 1b, right). Similarly, Stringer et al.¹¹ found that about a third of the population activity of ~10,000 neurons in the visual cortex of awake mice could be predicted from motor actions derived from a video of the mouse’s facial movements (Fig. 1c), suggesting that sensory–motor integration occurs even in the primary sensory cortex (the lack of a similar result in monkeys¹² could be due to the increased functional specialization of primate visual cortical areas compared to that of the rodent cortex).

Actions may be integrated differently across the different layers of a cortical area. For example, Jordan and Keller¹³ showed that both layer 2/3 and layer 5/6 neurons in the mouse primary visual cortex (V1) undergo depolarization before locomotion onset (Fig. 1d, left and middle). While layer 2/3 neurons appear to be computing a difference between motor-related input and bottom–up visual flow input, layer 5/6 responses were consistent with positive integration of visuomotor inputs (Fig. 1d, right). These results complement well-known earlier results on predictive activity in a range of cortical areas such as the visual cortex¹⁴, the parietal cortex¹⁵ and frontal eye fields¹⁶ that anticipate the visual consequences of impending eye movements.

¹Center for Neurotechnology, University of Washington, Seattle, WA, USA. ²Paul G. Allen School of Computer Science and Engineering, University of Washington, Seattle, WA, USA. ✉e-mail: rao@cs.washington.edu

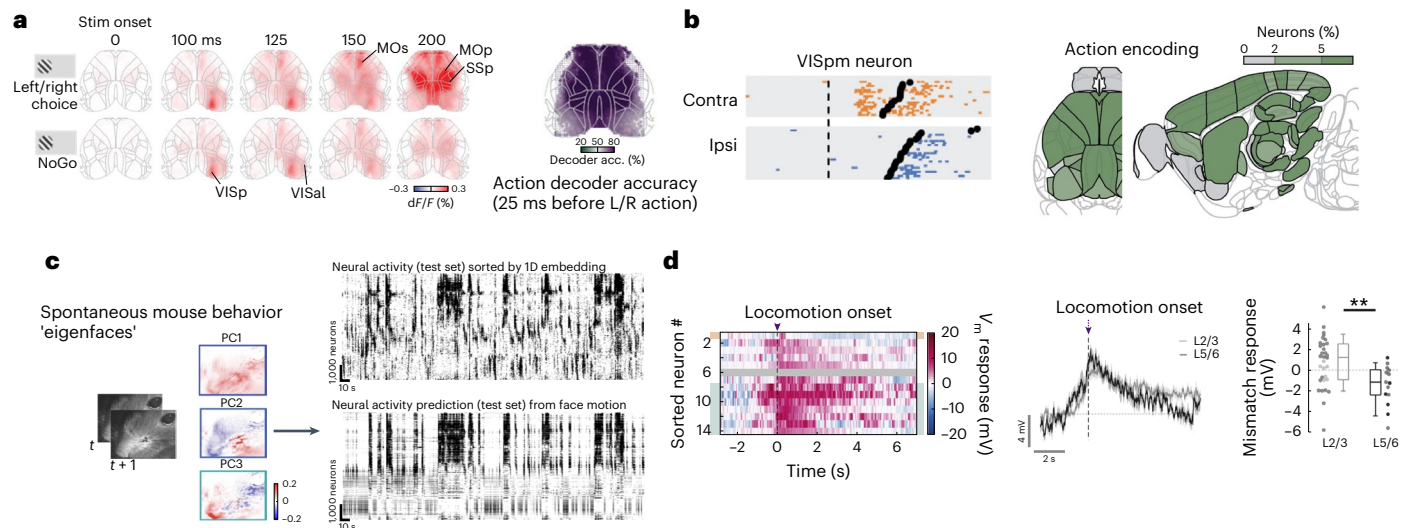


Fig. 1 Widespread influence of actions across the cortex. **a**, Left, widefield calcium imaging reveals bilateral activity (cortical fluorescence dF/F) from 0 to 200 ms after stimulus (stim) onset across cortical areas preceding movements on left or right (L/R) action-selection trials. ‘NoGo’ trials do not show such activity. Right, average action execution decoder accuracy (acc.) 25 ms before movement onset, showing that an impending movement can be decoded from cortical activity in most imaged regions (adapted from ref. 9, CC BY 4.0). MOs, secondary motor cortical area; MOp, primary motor cortex; SSs, primary somatosensory cortex; VISp, primary visual cortex; VISal, secondary visual cortical area. **b**, Left, increased spiking after a correct choice movement in a visual cortex neuron (posteromedial visual area, VISpm) for both contralateral (contra) and ipsilateral (ipsi) stimulus presentations (orange and blue dots, spikes; black dots, movement onset). Right, fraction of neurons in each brain region with pre-movement activity that could be accurately predicted from the animal’s movement (in either the left or right direction) (adapted from ref. 10, Springer Nature Limited). **c**, Motor action information extracted using

principal-component (PC) analysis of a video of a mouse’s facial movements (left, example frames t , $t + 1$; middle, top three principal components) accurately predicted (using reduced-rank regression) about a third of the population activity of ~10,000 neurons (raster representations on the right) measured using two-photon calcium imaging of the visual cortex of awake mice (adapted with permission from ref. 11, AAAS). 1D, one dimensional. **d**, Intracellular recordings in V1 of mice on a spherical treadmill with locomotion coupled to visual flow feedback. Visual flow was halted at random times to generate visuomotor mismatch events. Left, heatmap of average responses before and after locomotion onset across all layer 5/6 neurons. Baseline activity was subtracted from responses by using the average membrane potential in the 2.5 s before locomotion onset before averaging. Middle, average response before and after locomotion onset across all layer 5/6 (LS/6) neurons (black, 14 neurons) compared with all layer 2/3 (L2/3) neurons (gray, 32 neurons). Right, average mismatch responses of layer 2/3 and layer 5/6 neurons (adapted with permission from ref. 13, Elsevier).

The emerging view, as suggested by the studies above, is that almost all cortical areas update their representations on the basis of ‘efference copies’ of upcoming actions (‘corollary discharges’ (ref. 17)) as well as the results of executed actions. Such a view harmonizes well with the anatomical observation that all areas of the neocortex (henceforth, the ‘cortex’), including areas traditionally labeled as sensory cortices, send outputs to subcortical motor regions (refs. 18,19 and references therein) and receive input from these regions. Indeed, Vernon Mountcastle, in his prescient article in 1978 (ref. 20), put forth the hypothesis that a single unifying computational principle might be operating across the entire cortex, proposing the ‘cortical column’ as a modular information-processing unit of the cortex (see also refs. 21–23 for related ideas). This hypothesis is supported by the remarkable anatomical similarities in laminar connectivity patterns across cortical areas^{24,25}, even though the density of cells within laminae may vary across areas. Additional evidence for this hypothesis comes from experiments in which inputs from the optic nerve were diverted via the auditory thalamus to the auditory cortex, causing the auditory cortex to develop visual receptive field properties²⁶.

If there is indeed a common computational principle operating across the cortex, it must be versatile enough to explain capabilities as diverse as (1) learning to recognize an object from multiple visual glimpses through eye and head movements or from multiple tactile sensations through finger movements, (2) solving a complex spatial navigation task using simpler movement sequences and (3) understanding abstract concepts (such as a family tree).

In this Perspective, I suggest APC as a unifying sensory–motor theory of the cortex. APC hypothesizes a canonical cortical module as

consisting of a state-prediction network and an action-prediction network, both implemented within each cortical area. ‘State’ here denotes hidden (or ‘latent’) aspects of the world inferred from sensory inputs, for example, parts of an object to be recognized or one’s location in a building. ‘Action’ refers to not just motor commands but also abstract actions, for example, ‘go to the maternal grandmother node’ in a family tree or ‘perform multiplication’ on two given numbers. APC postulates that feedback from higher cortical areas modulates the dynamics of both state and action networks in lower areas, changing the functions they compute on the fly to suit the needs of the current task. This leads to representations that operate at multiple levels of sensory and motor abstraction, as observed in cortical hierarchies implicated in perception and action^{27–30}. In the following sections, I present computational and neurobiological aspects of APC, comparing emerging studies and experimental results with the model’s predictions. I present simulations chosen to showcase the explanatory breadth of APC. While we have previously investigated APC in the context of machine learning (ML) and artificial intelligence (AI)^{31–33}, here I explore APC as a model of cortical function.

APC Neuroanatomical and physiological motivation

The axonal outputs of layer 5 neurons in almost all cortical areas target subcortical motor centers¹⁹. Even in V1, outputs from layer 5 neurons target the superior colliculus³⁴, which is involved in eye movements (aside from other motor behaviors). Similarly, layer 5 neurons in the primary auditory cortex (A1) send outputs to the inferior colliculus³⁵, which is involved in orienting and defensive motor behaviors³⁶, while

BOX 1

Canonical APC module

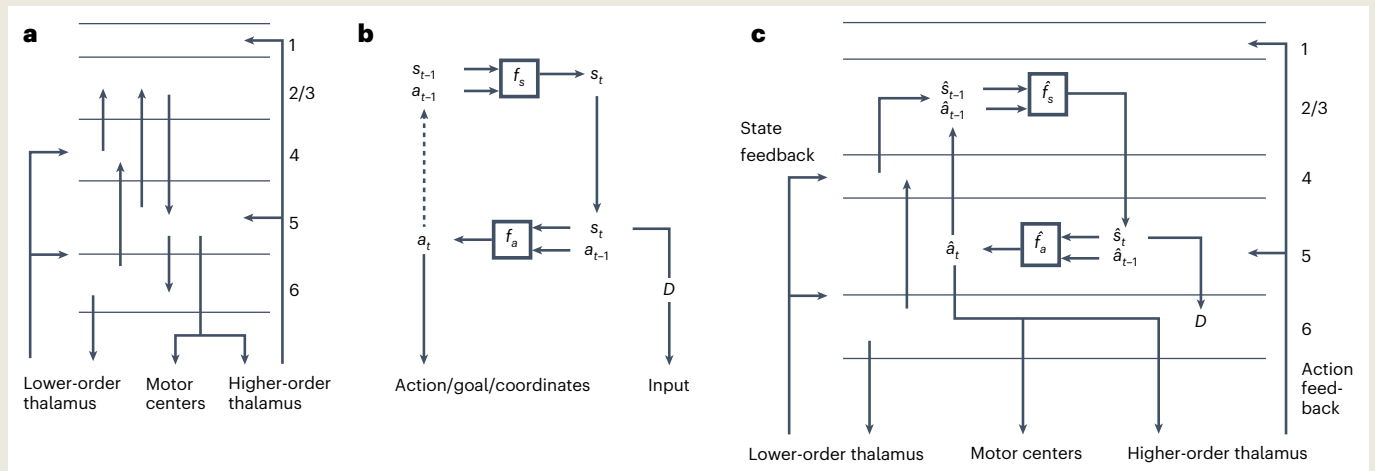
The canonical APC module is motivated by the laminar structure of the cortex (figure in Box 1, panel a) and consists of the following components:

State-transition function. The state-transition function f_s models the physics of the environment and the agent, and predicts the next state s_t , given the previous state s_{t-1} and action a_{t-1} (figure in Box 1, panel b). In general, states and actions are vectors. By learning an approximation \hat{f}_s of f_s from interactions with the world, the agent can learn an ‘internal model’ of the world^{81,87,88,94} (also called a ‘world model’, forward model or generative model) and use it to run simulations of the world, imagine new scenarios, explore what happens when particular actions are executed and plan actions that lead to desirable states. Biologically, the function \hat{f}_s can be implemented by a recurrently connected network of neurons, with the network activity at time t denoting an estimate \hat{s}_t of the state s_t (figure in Box 1, panel c). Neural population activity in visual, auditory and somatosensory cortices, for instance, after processing a sensory stimulus (or more generally, after sequential sampling of the stimulus), can be regarded as the estimated \hat{s}_t for that stimulus computed by the cortical region.

Policy function. An internal model (such as \hat{f}_s above) can be used to plan actions by unrolling the model into the future to explore the

consequences of various action choices, but this mode of selecting actions requires considerable effort and deliberation (‘system 2’ thinking⁶⁸). A more efficient way to select actions (‘system 1’ thinking⁶⁸) is to have a state-to-action ‘policy’ (ref. 44) \hat{f}_a , which maps the current estimated state \hat{s}_t directly to an action \hat{a}_t to achieve the current goal (figure in Box 1, panel c). Biologically, the policy \hat{f}_a can be implemented by a recurrent network of neurons with activity at time t denoting \hat{a}_t .

Coupling perception and action. Given a policy \hat{f}_a for a particular task or goal, the agent can execute an action \hat{a}_t from the policy while simultaneously sending \hat{a}_t as an ‘efference copy’ (or ‘corollary discharge’ (ref. 17)) to the learned model \hat{f}_s to predict the sensory consequences of each action. The agent can then correct its prediction of the new state of the world based on the new sensory observation that resulted from taking the action, using prediction errors as prescribed by predictive coding theory¹. The corrected state estimate \hat{s}_t can in turn be fed as input to the policy \hat{f}_a to generate the next action for the task, continuing until the goal is achieved or the task times out. A biological implementation of this idea involves coupling the state and policy recurrent networks within the laminar structure of a cortical area, as depicted in the figure in Box 1 (panel c).



Canonical APC module. **a**, Depiction of the six-layered laminar structure of a cortical column, showing some of the major connections between layers and with the thalamus (not all connections are shown) (based on refs. 18,24,43). **b**, Canonical APC generative model. The dashed arrow denotes a delay of one time step. D denotes a decoder converting state to input. **c**, One possible implementation of inference in a canonical APC module for the generative model in **b**.

layer 5 neurons in the primary somatosensory cortex send outputs to the spinal cord³⁷, which controls body movements. On the other hand, layer 5 neurons in the primary motor cortex (or M1) (such as Betz cells) have long been implicated in motor function via their axonal projections to the spinal cord³⁸, but the middle and upper layers of traditional ‘motor’ cortical areas such as M1 are involved in sensory processing of feedback from subcortical and cortical sources^{39,40}: for example, layer 2/3 neurons in mouse M1 respond to unexpected sensory perturbations in a visually guided motor task⁴¹, while neurons in mouse M2 encode auditory sensation and expectation⁴². Thus, despite

well-known differences in the laminar densities of neurons in different cortical areas (for example, V1 versus M1), the sensory–motor nature of the cortex is retained. This motivates the idea of a ‘sensory–motor cortical module’ as a canonical feature of the cortex (described further in the section APC module and neuroanatomical implementation).

The figure in Box 1 (panel a) depicts the laminar structure of a typical cortical column and its connectivity (based on refs. 18,24,43). Inputs from a sensory region or a lower cortical area target layer 4 neurons, the outputs of which are then conveyed to the superficial layer 2/3 neurons. These neurons in turn send their axons to the deeper layers,

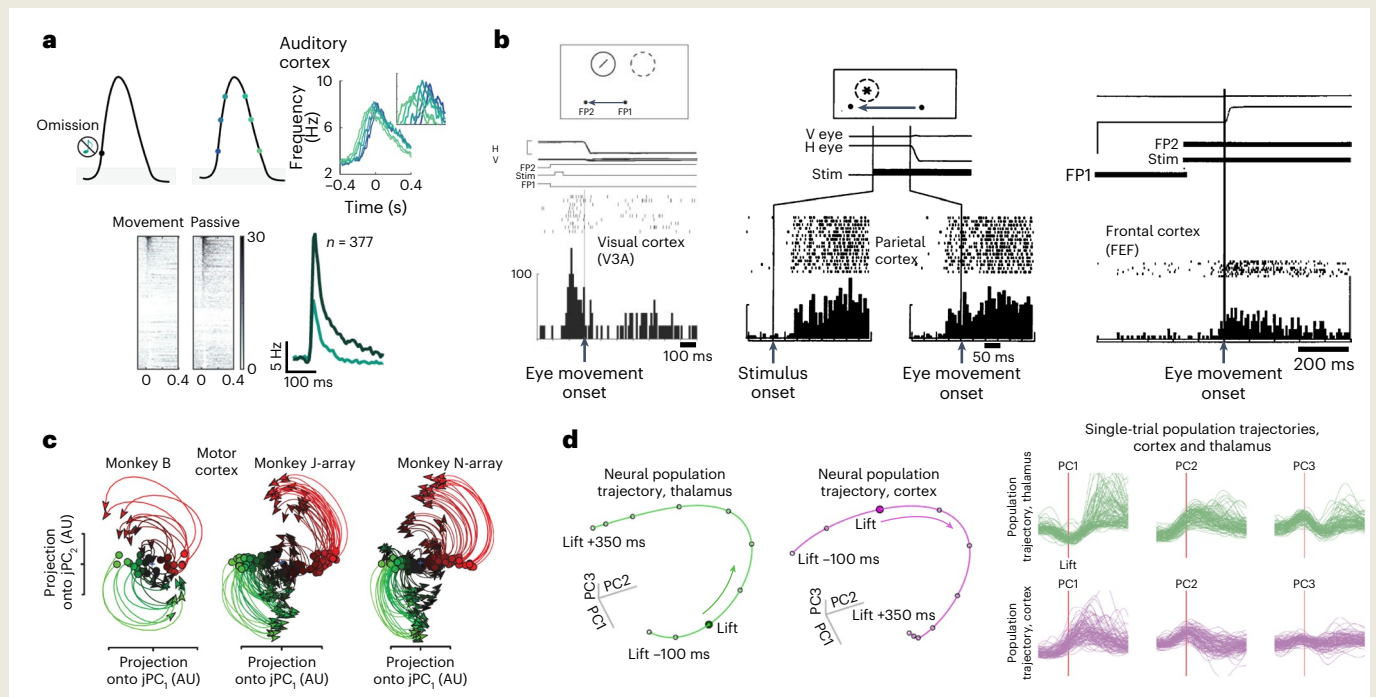
BOX 2

State and action networks in the cortex

State network \hat{f}_s and sensory-motor prediction in the cortex. The APC model predicts that cortical networks implementing the state-prediction function \hat{f}_s should learn to anticipate the sensory consequences of actions (by predicting latent sensory states) and exhibit anticipatory activity before movement. A study by Audette et al.⁹³ found such anticipatory activity in the primary auditory cortex (figure in Box 2, panel a): mice were trained to push a lever with their forelimb, which produced a pure tone at a fixed position early in each movement; after training, omission of this learned movement-associated sound (figure in Box 2, panel a, top left) revealed a large population of auditory cortex neurons firing roughly 200 ms before movement onset, with this activity peaking around the time of the expected tone (figure in Box 2, panel a, top middle and right). The study also found prediction error-like suppression of neural activity for anticipated sounds, consistent with the use of errors in predictive coding for state inference (figure in Box 2, panel a, bottom). Predictive activity anticipating the visual consequences

of an upcoming eye movement has been observed across the cortex including the visual cortex¹⁴, the parietal cortex¹⁵ and the frontal cortex¹⁶ (figure in Box 2, panel b, left, middle and right, respectively). Other types of movement, such as locomotion, can also predictively activate cortical neurons, for example, in V1 (ref. 13 and Fig. 1d). These results are consistent with the APC model's use of a state network \hat{f}_s to learn the sensory consequences of actions.

Action network \hat{f}_a and cortical motor dynamics. The APC model assumes that the network within a cortical area implementing the action-prediction function \hat{f}_a is a recurrent network (figure in Box 1, panel c) for which the outputs, in the case of the motor cortex, encode the dynamics of movement. There is now considerable neurophysiological evidence that, when an animal is performing a movement (for example, reaching toward a goal), the activities of neurons in the motor cortex are well described by a dynamical system of the form $\dot{a} = f(a, u)$ where u is an external input⁹⁵.



State and action networks in the cortex. **a**, Top left, schematic of the omission of an expected sound during a lever-press movement (black curve) in a self-generated sound task in mice⁹³. **a**, Top middle and right, neural activity in the auditory cortex aligned to six positions (colored dots) along the lever-press movement: peak firing always corresponded to the time of the expected sound (occurring on average 36 ms after movement onset), with the activity starting roughly 200 ms before movement onset (as reported in ref. 93). **a**, Bottom, heatmaps (left) of responses for all recorded neurons and the average response (right) for 377 of those neurons responding above a threshold to the sound when it was self-generated by a lever press ('movement') and when it was not self-generated ('passive'). Note the greater suppression (green) for self-generated sounds (adapted from ref. 93, Elsevier). **b**, Predictive activity anticipating the visual consequences of an upcoming eye movement in the visual cortex (left) (adapted with permission from ref. 14, National Academy of Sciences, USA), the parietal cortex (middle) (adapted with permission from ref. 15, AAAS) and the frontal cortex (right) (adapted with permission from ref. 16, American Physiological Society). FP, fixation point; H, horizontal eye movement; V, vertical eye movement; FEF, frontal eye fields. **c**, Two-dimensional projection of population activity of motor cortex neurons in three monkeys performing straight and curved reaching movements, showing rotation of the neural state. Traces are colored from green to red based on the preparatory neural state (circles) for each reach condition (adapted from ref. 95, Springer Nature Limited). jPC₁ and jPC₂, first and second dimensions, respectively, of jPCA space; AU, arbitrary units. **d**, Results from simultaneous population recordings of the motor thalamus and the motor cortex in mice performing reach-to-grasp movements⁹⁷. Left, neural population trajectories for the thalamus (left, green) and the cortex (right, magenta) obtained using trial-averaged PCA. Right, single-trial population activity in the thalamus (top) and the cortex (bottom) (adapted from ref. 97, Springer Nature Limited).

(continued from previous page)

The figure in Box 2 (panel c) illustrates these motor cortical dynamics in monkeys performing reach movements. Such dynamics can be learned and implemented by a recurrent neural network⁹⁶, which in the APC model is the action network \hat{f}_a . The APC action network receives as input not only local recurrent activity but also the current state estimate \hat{s} (figure in Box 1, panel c), which is updated using sensory feedback from other brain regions such as the thalamus. The APC model therefore predicts a close interaction between thalamic inputs and motor cortical outputs of the action network \hat{f}_a . A recent study by Sauerbrei et al.⁹⁷ confirmed the importance of this tight

sensory–motor loop in cortical pattern generation in mice performing dextrous movements. They showed that time-varying thalamic inputs are required for cortical pattern generation. The neural population activity in both the thalamus and the cortex exhibited strong co-modulation in trial-averaged (figure in Box 2, panel d, left) and single-trial (figure in Box 2, panel d, right) activity. Inactivating the thalamus perturbed cortical activity and disrupted limb kinematics⁹⁷, implying that both local dynamics and sensory-derived state contribute to generating cortical motor patterns (see inputs to \hat{f}_a in the figure in Box 1, panel c).

predominantly targeting layer 5 neurons. One class of layer 5 neurons (with thick tufted apical dendrites and firing in bursts) send their axons to subcortical motor centers such as the superior colliculus and other parts of the brainstem^{19,34,43}. Other layer 5 neurons, which do not fire in bursts and have slender apical dendrites, project to the striatum and other cortical regions^{34,43}. There is also a substantial axonal projection from layer 5 back to layer 2/3, signifying recurrent feedback within a cortical column. There are additional projections from layer 5 to layer 6, and layer 6 in turn sends outputs to the parts of the thalamus that send inputs to layer 4.

Computational motivation

Computational considerations point to maintaining a close link between actions and their sensory consequences. In model-based reinforcement learning⁴⁴ and, more generally, in the framework of partially observable Markov decision processes^{45,46}, an intelligent ‘agent’ interacts with the world by executing an action a_{t-1} at time $t-1$, and this causes the agent’s ‘state’ to change from s_{t-1} to s_t ; this change is governed by the state-transition function $f_s(s_{t-1}, a_{t-1})$, which generates a new state s_t , according to a probability distribution $P(s_t | s_{t-1}, a_{t-1})$. When the agent executes an action based on a policy f_a , for example, moving its body by walking or making an eye movement, the hidden state changes to the next state s_t (a new location in a building being navigated or a new part of a scene being recognized) according to $f_s(s_{t-1}, a_{t-1})$. Inferring these hidden states as the agent makes movements is the essence of perception^{47,48}.

APC module and neuroanatomical implementation

The above computational considerations motivate the canonical APC generative model shown in the figure in Box 1 (panel b). The corresponding model for inference and learning, referred to as the canonical APC module, is described in Box 1 and shown in the figure in Box 1 (panel c). This figure also suggests one possible functional mapping of APC’s computational elements onto the cortical laminar structure in the figure in Box 1 (panel a), which builds on previous proposals mapping predictive coding to cortical laminae^{49,50}.

As shown in the figure in Box 1 (panel c), the superficial layer cortical neurons, which receive the filtered sensory inputs from layer 4 and are recurrently connected to each other, are well suited to implementing the state-transition function f_s . The motor output layer 5 neurons, which are also recurrently connected to each other, fit the role of neurons computing the action–policy function f_a . The other class of layer 5 neurons, which convey information to other cortical areas and the striatum, could maintain the current state estimate \hat{s}_t by integrating the state prediction from layer 2/3 and correcting it with prediction errors from the feedforward thalamic inputs to layers 4 and 5/6 (ref. 1). Layer 6 neurons receiving inputs from these state-estimating layer 5 neurons are well placed to compute the prediction for a lower-level area: at the lowest level, layer 6 neurons predict sensory input \hat{l}_t for the input l_t (layer 6 neurons at a higher level would predict the cortical state at a lower level;

for more details, see the section Hierarchical APC and cortical feedback).

Layer 5 motor output neurons, for example, those sending outputs to the superior colliculus, send axon collaterals to higher-order thalamic nuclei^{18,19} and receive motor information from subcortical motor centers such as the superior colliculus regarding actions executed. These thalamic nuclei are therefore in an ideal position to compare the actual executed action a_t (from the superior colliculus or other motor center) and the cortical prediction \hat{a}_t . The resulting action feedback (for example, in the form of action-prediction errors), in addition to sensory feedback (in the form of sensory-prediction errors), can be conveyed by the thalamus back to the cortex to enable the state-transition network \hat{f}_s to correct its state prediction and the action network \hat{f}_a to correct its action prediction. Indeed, it is known that higher-order nuclei such as the pulvinar receive cortical layer 5 inputs and information from the superior colliculus and send axons to superficial layers of area V1, explaining the response of V1 neurons to saccadic eye movements⁵¹.

The implementation suggested above is consistent with growing experimental data on predictive coding in the cortex^{2,4}, with prediction error-like activity reported in superficial layers^{13,52} and predictive activity observed in deeper layers⁵³ and in cortico–cortical interactions^{54,55}. The implementation above also shares similarities with previous canonical circuits for predictive coding⁴⁹ in specifying laminar roles for state estimates and prediction errors but differs in the use of both actions and states to generate predictions, rather than being limited to hidden causes⁴⁹. Evidence for state and action networks in the cortex is summarized in Box 2.

Hierarchical APC and cortical feedback

A characteristic feature of the cortex is the reciprocal nature of connections between cortical areas²⁷: ‘feedforward’ connections from a cortical area A (originating in the superficial layers) to a cortical area B (terminating in layer 4) are invariably reciprocated by anatomically defined ‘feedback’ (or descending) connections from area B to area A (originating in the deeper (and sometimes superficial) layers of B to superficial and deep layers of A). Why are cortical areas reciprocally connected and organized in an approximate hierarchy^{27,56}?

Computational motivation

Consider the problem of going to the grocery store from one’s house. As shown in Fig. 2a, the complexity of the problem can be substantially reduced by dividing the task into subtasks (or ‘subgoals’), dividing each subtask into ‘sub-subtasks’, and so on. Reducing a complex problem to a sequence of easier-to-solve components and reusing these components to solve new problems gets to the heart of compositionality, which is thought to form the basis for cognitive flexibility and fast generalization in humans^{57,58}.

A complex problem can be characterized by its (typically high-dimensional) state-transition function, which governs how the




Fig. 2 | Using hierarchies and compositionality to simplify complex tasks. **a**, Decomposition of the ‘go to the grocery store’ problem into subgoals or subtasks, each of which can be further divided into sub-subgoals or sub-subtasks. Note that the rate of change is faster at the lower levels than at the higher levels, leading naturally to a temporal hierarchy. **b**, A navigation problem in a maze-like building environment with corridors (black) and walls (gray). Blue dot, current location; green square, desired goal location. The structure of the environment (pathways and walls) can be understood in terms of the building’s state-transition dynamics, which in turn can be divided into the simpler transition dynamics

of its compositional elements, namely, the two rooms outlined in yellow and red that appear at several different locations within the reference frame of the environment. These simpler elements can be further decomposed into horizontal and vertical corridors shown on the right that appear at different locations within the local reference frame of each room. **c**, An object (such as a handwritten digit ‘8’) can be divided into parts (loops and curves at the middle level), each of which can be divided into subparts (strokes, lines, smaller curves at the lower level). Each part and subpart is associated with its coordinates (location and transformation) within a local reference frame.

state of the environment changes when one applies an action. Fortunately, in the natural world, the consequences of most actions are local, and these local dynamics tend to be shared across many environments and objects, allowing complex problems to be modeled in terms of simpler lower-dimensional state-transition functions. This is illustrated in the example in Fig. 2b, a simplification of the ‘go to the grocery store’ problem: here, a maze-like environment is modeled using simpler components (yellow- and red-outlined rooms), each composed of even simpler components (corridors).

Interestingly, the same concept can also be applied to visual perception. As illustrated in Fig. 2c, a visual object can be compositionally defined in terms of parts and their locations within the object’s reference frame⁵⁹; the parts can in turn be decomposed into simpler parts within their respective reference frames. Nested compositional representations of objects and environments offer substantial combinatorial flexibility for solving complex problems in terms of simpler, reusable components. The fact that the world we live in and the problems we seek to solve are amenable to compositional solutions makes such an approach attractive, from both a computational and an evolutionary perspective.

The hierarchical APC model

Box 3 describes the APC model’s hierarchical architecture and its neural implementation. The figure in Box 3 (panel e) shows two levels of the model, implemented using top-down ‘contextual inputs’ to connect the higher level to the lower level (see Box 3 for an alternate implementation based on gain modulation).

State inference using predictive coding and compositional learning.

As shown in the figure in Box 3 (panel e, left), the higher-level state neurons maintain an estimate for the state $s_t^{(i+1)}$ at time t and modulate the lower-level state network via top-down feedback given by $H_s^{(i)}(s_t^{(i+1)})$. The lower-level state neurons maintain an estimate for $s_{t,\tau}^i$, where τ denotes a time step at the lower level within the higher-level time interval given by t . The lowest-level state makes a prediction of the input via a ‘decoder’ network D (figure in Box 1, panel b). If D is a linear matrix U , this lowest level of APC is equivalent to the generative model in sparse coding ($I = Us$, where s is sparse⁶⁰). At each time step, the network predicts the next input as a function of previous state and action. Feedforward pathways convey prediction errors to update state estimates¹, while descending pathways convey top-down modulation as described above (see ref. 61 for an example). Prediction errors are also used to learn the weights of the state networks at all levels using predictive coding-based self-supervised learning^{1,32,61}. Such learning approximates error backpropagation, the workhorse of contemporary deep learning, in a biologically plausible manner⁶².

Action inference through planning and reinforcement learning.

As shown in the figure in Box 3 (panel e, right), the higher-level action neurons represent an abstract action (such as ‘open the door’) via an action vector $a_t^{(i+1)}$ at time t . Given the abstract action $a_t^{(i+1)}$, top-down feedback given by the embedding input $H_a^{(i)}(a_t^{(i+1)})$ modulates the lower-level action network and instantiates the goal-specific policy $f_a^{(i)}$, which produces lower-level actions $a_{t,r}^i$. The hierarchical action networks in the APC model can be trained in multiple ways: (1) planning: the state-transition networks can be used to search for sequences of actions, starting from the highest abstraction level, that are likely to result in states with the highest cumulative reward or closest to the goal (see also active inference⁸, planning by inference^{63–65} and model predictive control⁶⁶; see the section Illustrative examples of diverse computational capabilities). Successful actions can be used as ‘labels’ in supervised learning to train the policy networks $f_a^{(i)}$; (2) reinforcement learning: hierarchical reinforcement learning⁶⁷ can be used to train action networks at each level to maximize the total expected reward according to a reward function that may be specific to that level (details in the section Illustrative examples of diverse computational capabilities); (3) policies providing priors for planning: action networks $f_a^{(i)}$ at each level predict a distribution over actions, which can serve as a prior, in a Bayesian sense, for guiding the search for actions in planning. Thus, predicted actions for new tasks will have high uncertainty, requiring effort and deliberation in planning (‘system 2 thinking’ (ref. 68)), while, for frequently encountered tasks, action networks are well trained and will predict actions with high confidence (‘system 1 thinking’ (ref. 68)).

An advantage of continuous-valued state $s^{(i+1)}$ and action $a^{(i+1)}$ vectors is that interpolating or sampling in the neighborhood of learned $s^{(i+1)}$ and $a^{(i+1)}$ generates, on the fly, new state-transition functions $f_s^{(i)}$ and new policy functions $f_a^{(i)}$, opening the door to fast generalization and transfer of knowledge to new tasks. The alternative to continuous states is to use discrete states, as in previous models of predictive coding based on belief propagation or variational message passing (for example, Fig. 10 in ref. 69). Aside from the possibility of fast generalization and transfer, APC’s use of continuous states also allows explicit representation of prediction errors, which in turn allow local optimization of dynamics and learning (under Gaussian assumptions)¹.

Proposed neuroanatomical implementation

I propose that neural populations in a higher cortical area A representing current state and action vectors $s^{(i+1)}$ and $a^{(i+1)}$ modulate the state and action networks $f_s^{(i)}$ and $f_a^{(i)}$ in a lower cortical area B via descend-

BOX 3

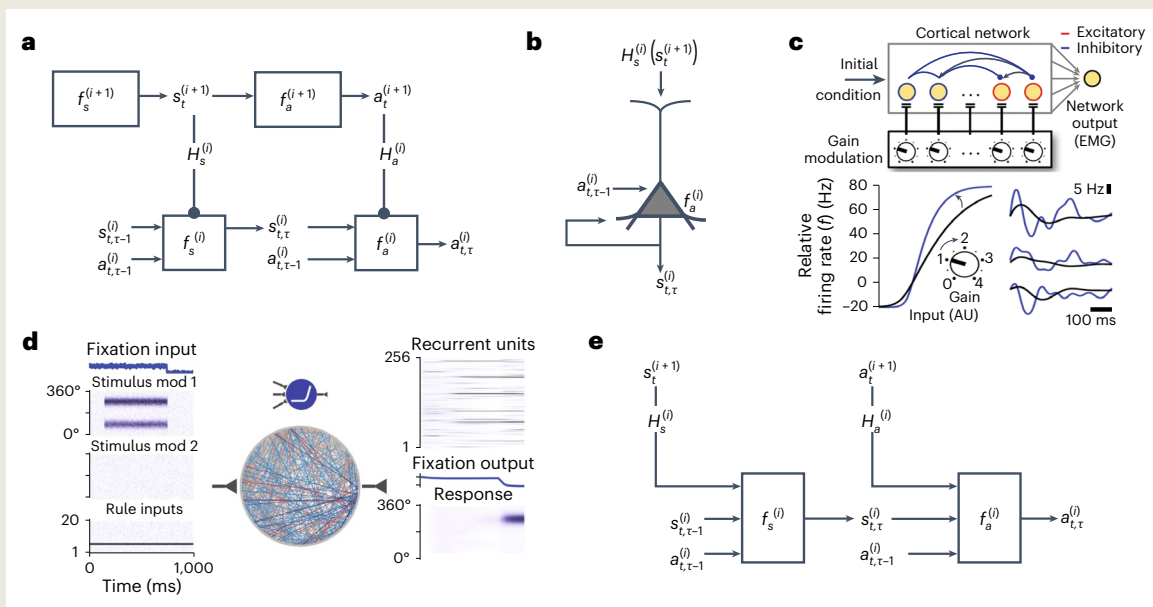
Hierarchical APC model and its neural implementation

In the hierarchical APC model (figure in Box 3, panel a), a state $s^{(i+1)}$ and an action $a^{(i+1)}$ at abstraction level $i + 1$ generate, respectively, a state-transition function $f_s^{(i)}$ and a policy function $f_a^{(i)}$ ('option' (ref. 44)) at the lower level i . These functions interact with each other to generate lower-level states and actions (figure in Box 3, panel a). Each such state and action in turn generates transition and policy functions at an even lower level of abstraction. A lower-level sequence executes for a period of time until a condition is met (for example, a subgoal is reached, a task is completed or times out or there is an irreconcilable error at that level). Control then returns to the higher level, which transitions to a new higher-level state (via $f_s^{(i+1)}$) and action (via $f_a^{(i+1)}$). Such a model captures both the dynamics of states (the 'physics' of the world) and actions ('policies') at different time scales, allowing hierarchical problem solving (Fig. 2). Biologically, the recurrent networks implementing f_s and f_a are governed by specific decay time constants, but differences in recurrent excitation can allow a hierarchy of time scales, as observed across the cortex⁹⁸. Hierarchical dynamics in predictive coding was previously explored in ref. 3 for perception and production of birdsong, with lower-level

dynamics contextualized by slowly varying control parameters supplied by a higher level.

Can the hierarchical model in the figure in Box 3 (panel a) be implemented in networks of neurons? More specifically, how can a population of neurons, representing, for example, the higher-level state vector $s^{(i+1)}$ (or action vector $a^{(i+1)}$), generate a whole function $f_s^{(i)}$ (or $f_a^{(i)}$) at the lower level?

Gain modulation in the cortex. There is considerable evidence for 'gain modulation' in cortical networks^{99–102}, implemented computationally by multiplying the synaptic weights or outputs of neurons by a gain factor. Evidence for gain modulation ranges from multiplicative modulation of tuning curves of visual cortical neurons during attention¹⁰³ to changes in the input–output function of neurons in deep layers of the cortex due to 'top–down' modulatory inputs to their apical dendrites in layers 1 and 2/3 (figure in Box 3, panel b)¹⁰⁰. We can view gain modulation as a biologically plausible way of implementing a 'hypernetwork' (ref. 104), which in ML and AI is a neural network that produces the synaptic weights (or more



Hierarchical APC model. a, Two levels (levels $i + 1$ and i) of a hierarchical APC generative model. Each level has a state-transition function f_s capturing the dynamics of the world at a particular level of abstraction and a policy function f_a specifying that level's actions, goals and coordinates (conditioned on the current highest-level goal or task). Higher-level state and action vectors at time t generate, via top–down networks H_s and H_a , lower-level state-transition and policy functions, allowing the higher level to compose a sequence of states and actions at the lower level to accomplish a goal. **b**, As depicted here for a single pyramidal neuron, I hypothesize that top–down inputs $H_s^{(i)}(s_t^{(i+1)})$ from a higher cortical area to the apical dendrites of lower-area neurons modulate the dynamics of a network of such neurons (for example, via gain modulation^{100,101}), allowing the higher area to change the functions f_s and f_a at the lower level. **c**, Top, multiplicative gain modulation (for example, due to top–down inputs) in the input–output function of neurons in a recurrent network allows the network to generate a rich set of motor cortical dynamics matching experimental data¹⁰². EMG, electromyogram of muscle activity. Bottom, changing the gain from 1 (black) to 2 (blue) (bottom left plot) dramatically alters neuronal firing rates (three example neurons are shown on the right), mimicking quasi-oscillatory motor cortical activity (see figure in Box 2, panels c,d) (adapted from ref. 102, Springer Nature Limited). **d**, The function computed by a recurrent network (center) can be modulated using a nonchanging top–down contextual input or 'rule input' (one-hot vector, bottom left) in addition to recurrent and stimulus inputs (top left), allowing the same network to solve different tasks (output for a specific task is shown on the right) (adapted from ref. 72, Springer Nature Limited). Mod, modality. **e**, Implementation of the APC model in **a** using contextual inputs: higher-level state and action neurons maintaining estimates of $s_t^{(i+1)}$ and $a_t^{(i+1)}$ modulate lower-level state and action networks via top–down contextual inputs $H_s^{(i)}(s_t^{(i+1)})$ and $H_a^{(i)}(a_t^{(i+1)})$, respectively.

(continued from previous page)

plausibly, the gain parameters) for another neural network (called the ‘primary network’). In the APC model, I propose that the higher-level state vector $s^{(i+1)}$ is fed as input to a top-down feedback network $H_s^{(i)}$, which produces the gain values to modulate the lower-level state network $f_s^{(i)}$ (and similarly for the action network). The ability of such a neural mechanism to modulate the function being computed by a cortical network was demonstrated by Stroud et al. (figure in Box 3, panel c), who showed that multiplicative gain modulation of a recurrent network can generate a rich set of motor cortical dynamics matching experimental data¹⁰².

Contextual inputs in the cortex. Aside from gain modulation, higher cortical areas can also change the function being computed by lower cortical networks using top-down contextual inputs. For example, Yang et al.⁷² showed that, by feeding a top-down contextual input (‘rule input’) as a nonchanging input to a recurrent network (in addition to its usual recurrent and external inputs), one can change the input–output function that the network computes, allowing the same network to solve different tasks (figure in Box 3, panel d; see also the

model of Eliasmith and colleagues¹⁰⁵). This is known in ML as the ‘embedding approach’ and can be shown to be equivalent in computational function to hypernetworks¹⁰⁶. In the case of APC, the higher-level state vector $s^{(i+1)}$ (and action vector $a^{(i+1)}$) can be fed as input to a top-down feedback network $H_s^{(i)}$ ($H_a^{(i)}$) that produces an embedding vector, which acts as a contextual input to a lower-level cortical network that computes $f_s^{(i)}$ ($f_a^{(i)}$) (figure in Box 3, panel e). The higher level can therefore control the function being computed at the lower level by changing the embedding vector (contextual input).

The APC model acknowledges the existence of both gain modulation and contextual inputs in the cortex and postulates that either or both of these mechanisms are used for changing the functions $f_s^{(i)}$ and $f_a^{(i)}$ at the lower level according to the current higher-level state vector $s^{(i+1)}$ and the action vector $a^{(i+1)}$. The examples described in the section Illustrative examples of diverse computational capabilities were implemented using the method of contextual inputs. The reader is referred to ref. 61 for examples based on gain modulation and to refs. 32,33 for hypernetwork-based examples.

input connections that target the superficial and deep layers of area B²⁷. Feedforward connections that target layer 4 of area A arise from the lower area B and from the higher-order thalamic region receiving ‘driver input’ from area B^{18,27}. I propose that these feedforward connections carry the state and action feedback (for example, prediction errors) that enable the higher area A to correct its abstract state and action estimates. Such a neural implementation is consistent with studies reporting prediction error-like responses in superficial layers and state-estimation-like responses in deeper layers of the cortex^{2,13}.

A key difference from previous formulations of hierarchical predictive coding^{1,2,49,50} is that APC places subcortical (thalamic) populations center stage in the evaluation of state- and action-prediction errors and their broadcasting to superficial pyramidal cells in the cortex⁷⁰. Additionally, in APC, descending connections from higher to lower cortical areas change the function being computed in the lower area (through top-down modulation), rather than only conveying lower-level state predictions as in traditional predictive coding¹. Higher-area neurons representing action $a_t^{(i+1)}$ modulate the action network $f_a^{(i)}$ in a lower area, changing the policy function that this lower-level network is computing. Neurophysiological evidence for such compositional representations has recently emerged in the premotor cortex⁷¹. Computational models have demonstrated compositionality for task transfer using the embedding space of $a_t^{(i+1)}$ (ref. 72). Hierarchical representations found in the visual^{27,29} and motor systems³⁰ are also consistent with the hierarchical compositional approach espoused by the APC model. Finally, the compositional representations in the cortex postulated by the APC model align well with recent hypotheses regarding compositional hippocampal replay⁷³.

After the state and action networks have been learned for a set of tasks (as described in the section The hierarchical APC model), given a particular task, the topmost state vector in the hierarchy is first inferred from sensory inputs. This vector produces (via that level’s action network $f_a^{(i)}$) the topmost action vector specifying a ‘goal’ or option for the task; I hypothesize that this vector is maintained in the prefrontal cortex. Some of the pre-movement anticipatory activity in Fig. 1a may well reflect such a ‘cognitive’ decision. This abstract action vector is decomposed hierarchically all the way down to elemental actions (for example, muscle control signals in M1). When a sequence of elemental actions is executed and a subgoal is reached, control is returned back to the level above to generate a new subgoal (see the sections Active visual percep-

tion and part-whole learning and Planning and navigation using hierarchical world models for examples) and similarly for all levels. I postulate that such coordination between hierarchical levels for action selection occurs via cortex–basal ganglia–thalamus–cortex loops; I leave the important problem of working out the implementation details in these loops to future research.

Illustrative examples of diverse computational capabilities

The architecture of the APC model was inspired by the hypothesis that evolution may have replicated a common computational principle across the cortex^{20–23}. If that is the case, one would expect the same architecture to be able to solve a diverse set of problems. Inspired by this observation, I provide here examples illustrating the APC model’s diverse capabilities.

Active visual perception and part-whole learning

Human vision can be viewed as an active sensory–motor process that employs eye movements to move the high-resolution fovea to appropriate locations in a scene, gathering evidence for or against competing visual hypotheses^{74,8}. The APC architecture is well suited to modeling such a sensory–motor process, given its integrated state and action networks. To illustrate this capability, we simulated^{31,32} a two-level APC model (figure in Box 3, panel e) in which the lower-level actions emulated eye movements by moving a fovea (‘glimpse sensor’ (ref. 74)) to extract high-resolution information about a small part of the input image within a larger reference frame selected by the higher-level action.

The lower-level action also predicts a new state vector $s_{t,t+1}$, which generates, via a trained decoder, a prediction for the glimpse image expected after the ‘eye movement’. The resulting prediction error was used for state inference and learning. The state networks at both levels were trained to minimize image-prediction errors, while the action networks were trained using reinforcement learning for the task of image reconstruction (for image classification as the task, see ref. 31).

Fig. 3a shows an example of a learned parsing strategy by the two-level APC model. The higher level learned to select actions that cover the input image sufficiently, avoiding blank regions, while the lower level learned to parse subparts inside the reference frame computed by the higher level. Fig. 3a also suggests a potential explanation for why human perception can appear stable despite dramatic changes

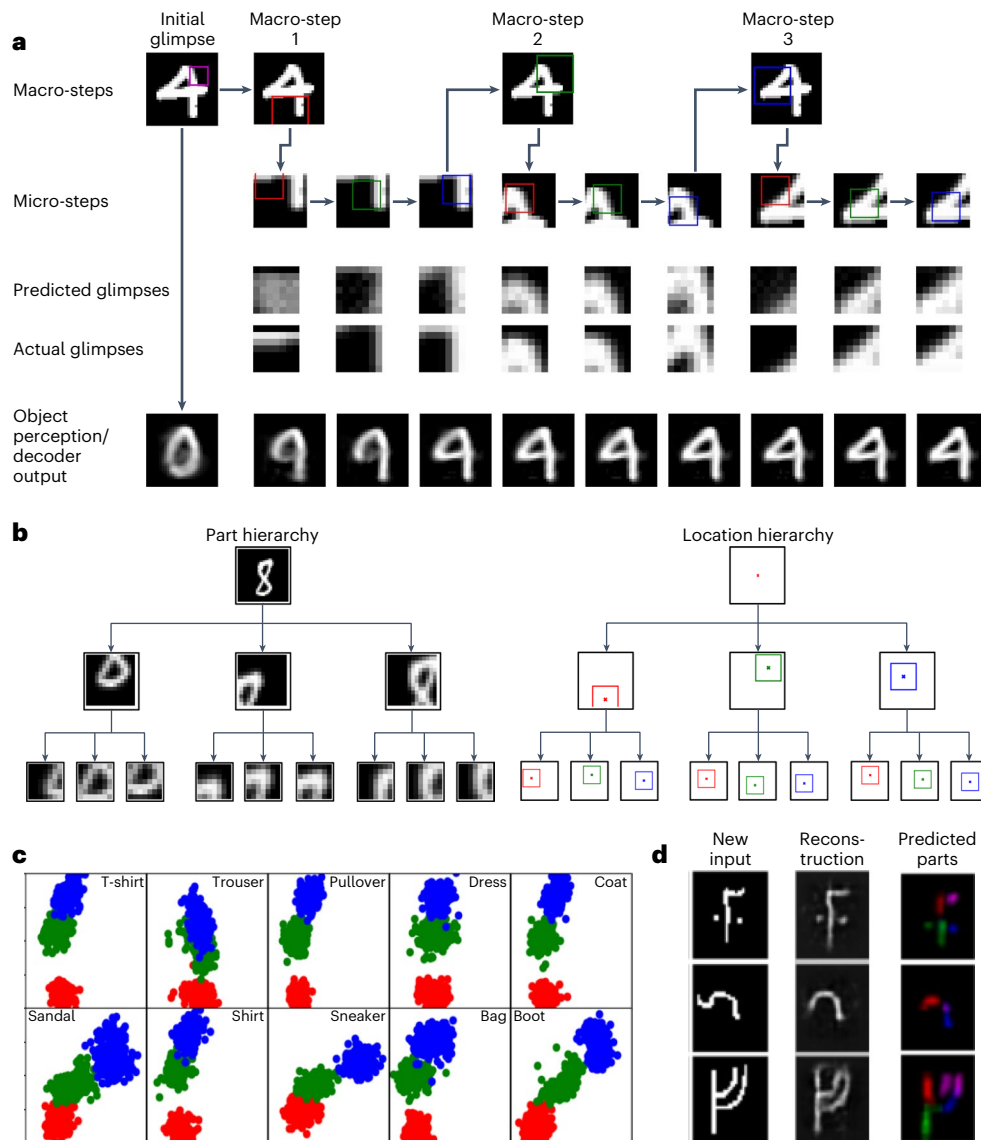


Fig. 3 | Active vision, part-whole learning and transfer of knowledge. **a**, First row, initial glimpse (purple box) and higher-level reference frames selected (red, green and blue boxes) at higher-level time steps ('macro-steps'); second row, regions fixated at lower-level time steps ('micro-steps') within each higher-level reference frame; third and fourth rows, predicted versus actual glimpses; fifth row, the model's 'perception' over time (object reconstructed by a decoder network from the current network state). Note the model's 'perceptual' stability despite jumps in actual glimpses, enabled by predictions of the glimpses similar to visual cortical predictions before eye movements (figure in Box 2, panel b). **b**, The digit '8' is parsed by a trained APC model as a parse tree of parts and subparts (left) and their corresponding coordinates (locations) within their respective reference frames (right). The representation is compositional: the same set

parts and subparts can potentially be reused at other locations and with other transformations to compose new digits. **c**, Higher-level part locations selected by a trained APC model for a particular class of clothing items in the Fashion-MNIST dataset (red, green and blue dots show the average sampled locations fixated in the following order: first, red; second, green; third, blue). Note the differences in the model's fixation strategies between vertically symmetric items (shirts, trousers, bags) and footwear (sandals, sneakers, boots). **d**, An APC model trained on the Omniglot handwritten characters dataset (from 50 different alphabets) can transfer its learned knowledge to predict parts of previously unseen character classes. First column, input image from a new character class. Middle column, APC model's reconstruction of the input. Last column, parts predicted by the model (**d**, adapted with permission from ref. 32, MIT Press).

in our retinal images due to eye movements: the model maintains a stable visual hypothesis that is gradually refined without exhibiting the rapid changes seen in the sampled image regions (Fig. 3a, actual glimpses). This 'perceptual' stability is enabled by the model's ability to predict the expected glimpses for each planned 'eye movement' (Fig. 3a, predicted glimpses), similar to predictive activity observed in the visual cortex before eye movements (figure in Box 2, panel b)¹⁴⁻¹⁶.

Fig. 3b shows a learned part-whole hierarchy for a digit in terms of strokes and mini-strokes along with their locations within nested reference frames. The model learns different parsing strategies for different classes of objects (Fig. 3c). Setting the image-prediction

error input to the network to zero forces the model to predict the next sequence of parts and 'complete' an object³². Finally, compositional learning in the APC model facilitates transfer of learned knowledge to new objects (Fig. 3d).

Planning and navigation using hierarchical world models

Interestingly, the same APC framework used above for active vision can also be used for planning hierarchical actions for tasks such as navigation. Consider the problem of navigating from any starting location to any goal location in a large 'multi-room' building environment such as the one in Fig. 4a (gray, walls; blue circle, current

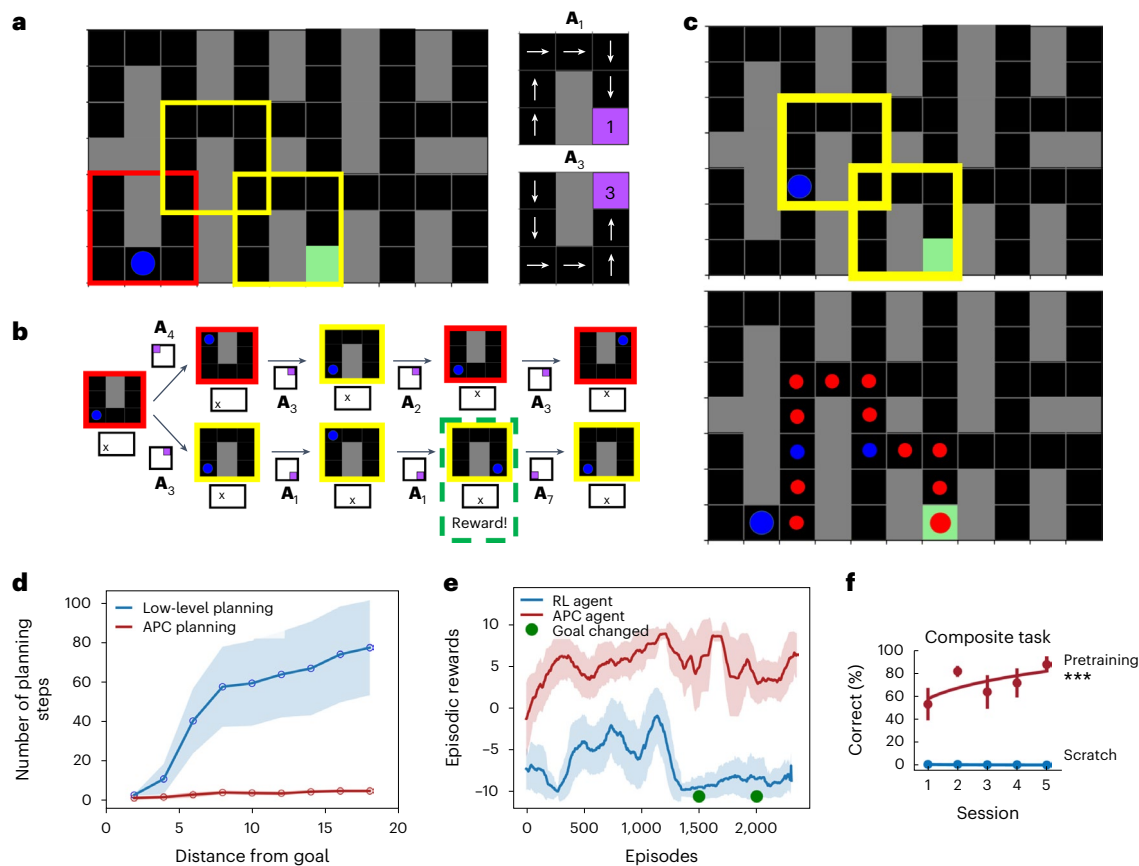


Fig. 4 | Hierarchical planning. **a**, The problem of navigating in a large environment (left) can be reduced to planning using high-level states (red- and yellow-outlined ‘rooms’) and high-level abstract actions (panels on the right show two abstract actions, A_1 and A_3). Blue, current location; gray, walls; green, current goal location. **b**, To navigate to the goal, the APC model uses its learned high-level state network to sample K high-level state–action sequences ($K = 2$ here, shown bifurcating from the initial state). In each sequence, the high-level state is depicted by a predicted room image (red- or yellow-outlined image) and its location (marked by an ‘X’ in the rectangular global frame below the image). High-level actions are depicted as square local frames (next to arrows) with goal locations (purple). **c**, Given the sampled sequences, the model picks the sequence with the highest total reward, executes this sequence’s first (high-level) action to reach the blue location (top) and repeats to reach the goal location with

only three high-level actions (bottom). Small red dot, intermediate location; small blue dot, intermediate goal. **d**, High-level planning by the APC model versus low-level heuristic planning using primitive actions (see text for details). **e**, The APC model can reuse learned high-level actions in new combinations to quickly solve new tasks (green circles, times at which the navigation goal changed); a reinforcement learning (RL) agent needs to relearn a new policy from scratch. Blue- or red-shaded regions in **d**, **e** are 1 s.d. from the mean. **f**, Mice pretrained on two subtasks quickly learned to combine them to solve a new composite task⁷⁵ (compare with the APC model in **e** after a goal change). Blue, performance of mice learning the task from scratch (compare with the reinforcement learning agent in **e** after a goal change) (**a–e**, adapted with permission from ref. 32, MIT Press; **f**, adapted from ref. 75, CC BY 4.0).

location; green square, current goal location). Here, the lower-level states of the APC model are locations in the grid, and lower-level actions are going north, east, south or west, with a large reward at the goal location and smaller negative rewards for each action to encourage shorter paths.

Just as an object consists of parts at different locations, the building environment in Fig. 4a is composed of smaller elements (two 3×3 ‘room types’, S1 (red) and S2 (yellow)) at different locations in the global reference frame of the building. The higher-level states of the APC model are defined by state-embedding vectors S1 and S2, trained to generate, via the top–down network H_s (figure in Box 3, panel a), the lower-level transition functions \hat{f}_s for rooms S1 and S2, respectively.

Similar to how the APC vision model reconstructed an image in the section Active visual perception and part–whole learning by composing parts from subparts, the APC model for planning computes higher-level action embedding vectors A_i (option vectors) that generate, via the top–down network H_a (figure in Box 3, panel a), lower-level policies \hat{f}_a that produce primitive actions (north, east, south or west) from any location in the local reference frame (S1 or S2) to reach a local goal i within that frame. Fig. 4a (right) illustrates two of the eight A_i ,

each trained using reinforcement learning to reach one of the four corners of S1 or S2 (see ref. 32 for details). Defining these policies to operate within the local reference frame of the higher-level state S1 or S2 (regardless of global location in the building) allows the same policy to be reused at multiple locations.

The higher-level state network was trained to predict the next higher-level state. This trained higher-level network was used for planning (using model predictive control⁶⁶): random state–action trajectories of length 4 were generated using the higher-level state network by starting from the current higher-level state and picking at random one of the four higher-level actions A_i for each next higher-level state. The action sequence with the highest total reward was selected, and its first action was executed. This process was repeated. Fig. 4b,c illustrates this high-level planning process using the trained APC model.

Fig. 4d illustrates the efficacy of the APC model’s high-level planning compared to lower-level planning using primitive actions (see ref. 32 for details): the APC model takes significantly fewer planning steps and can reuse its learned higher-level actions in new combinations to quickly solve new tasks (for example, when the goal is changed; Fig. 4e), similar to a recent study in mice⁷⁵ (Fig. 4f).

BOX 4

Model flexibility and predictions

The APC model appears to be flexible enough to perform a diverse set of functions such as: (1) parsing images and learning part–whole hierarchies: the model uses eye movements to parse images and learn hierarchical representations of parts and subparts of objects (details in the section Active visual perception and part–whole learning); (2) invariant perception: learned representations of objects and sequences are transformed by the APC generative model to match current inputs and remain invariant to different types of transformations (translations were considered in the section Active visual perception and part–whole learning, other transformations such as rotations and scaling are included in ref. 33); (3) perceptual stability: inference in the APC model naturally leads to integration of information across actions such as eye movements, leading to perceptual stability (examples in the section Active visual perception and part–whole learning; see also refs. 23,59,74); (4) compositionality and fast transfer of knowledge: by learning compositional representations, the model can compose and generate new objects and action sequences, leading to fast generalization to new inputs and goals (examples in the sections Active visual perception and part–whole learning and Planning and navigation using hierarchical world models); (5) efficient planning: hierarchical state networks in the APC model can be used to solve tasks efficiently (for example, navigating in a large environment) by planning using hierarchical actions (details in the section Planning and navigation using hierarchical world models; see also ref. 107); (6) habit formation: successful plans can be used to learn new policies ('habits'; see the section Planning and navigation using hierarchical world models); alternately, the APC model also allows policies to be learned using hierarchical reinforcement learning (see the section Active visual perception and part–whole learning); (7) reference frames and temporal hierarchies: the APC model provides a neural implementation of nested reference frames²³ and offers an explanation for object-centered parts-based representations in the cortex¹⁰⁸ as well as cortical temporal hierarchies^{28,29}; (8) prediction and postdiction: because the model maintains a temporally stable higher-level state (a 'timeline' (ref. 76)) encoding an entire sequence (past, present and future), the update of this representation during prediction error minimization explains both predictive and postdictive phenomena in perception (for example, flash lag and color phi effects; see ref. 61 for details); (9) generating 'schemas' or 'programs' for solving new tasks: the APC model suggests a neural mechanism (via top–down inputs and/or gain modulation) for generating new sensory–motor 'programs' or 'schemas' on the fly to solve new tasks (Box 3 and ref. 109); (10) binding and episodic memories of perception–action sequences: when coupled with a hippocampus-like associative memory, the model binds multimodal cortical activations at the highest level into an episodic memory, allowing activity recall, prediction based on episodic context and cortical consolidation for fast generalization and learning, as discussed in the section Episodic memories and cortical–hippocampal binding; (11) language and symbolic representations: making state and action representations categorical (for example, ref. 78) and using cortical–hippocampal binding may allow the APC

model to bind sensory and symbolic representations for language processing and cognitive tasks such as arithmetic (see the section Learning abstract concepts and ref. 58); (12) learning abstract concepts: the same sensory–motor architecture used for perception and planning can also be used to model abstract concepts such as family trees (details in the section Learning abstract concepts and ref. 77).

The APC model makes the following predictions:

- The laminar implementation proposed in the figure in Box 1 (panel c) predicts that, in each cortical area, neurons representing the sensory-derived latent state will exhibit predictive activity, representing the output of the state-transition function \hat{f}_s , as a function of both sensory inputs and deeper-layer 'action' inputs; experimental manipulation of deeper-layer activity should allow control of this predictive activity and change this activity similar to how the predicted glimpses change in Fig. 3a as a function of action inputs to \hat{f}_s .
- In each cortical area (including traditional 'sensory' areas), the model predicts a population of neurons in layer 5 representing 'actions', either motor outputs (for example, in M1 or the primary somatosensory cortex) or abstract actions or goals (for example, in the parietal or prefrontal cortex); this action-related population activity in layer 5 should change in a coordinated manner across cortical areas whenever the goal or task is changed, similar to how the lower- and higher-level actions change in the APC model whenever the goal location is changed in Fig. 4a.
- As depicted in the figure in Box 3 (panels a,e), feedback from a higher cortical area should originate from two separate populations (higher-level state- and action-representing neurons) and specifically target two separate populations in the lower area (lower-level state- and action-representing neurons, respectively); this feedback should be modulatory and capable of changing the functional connectivity of their target populations, emulating the effects of the 'hypernetworks' $H_s^{(i)}$ and $H_a^{(i)}$ in the figure in Box 3 (panels a,e). Experimental manipulation of this feedback should selectively change the lower-level functions $\hat{f}_s^{(i)}$ and $\hat{f}_a^{(i)}$ being computed, with the effects becoming apparent in the state predictions being generated by \hat{f}_s (see predicted glimpses in Fig. 3a) and the actions being output by \hat{f}_a (see primitive actions in Fig. 4c (bottom)), respectively.
- APC's hierarchical arrangement predicts the existence of both state and action representations in higher-order cortical areas that encode information over longer time scales than lower areas^{28,29,61,80,98,110}; furthermore, the model predicts that a substantial population-level activity change in the superficial layers of a higher area, denoting a higher-level state change triggered by a subgoal being achieved by the lower area, would cause a large population-level activity change in layer 5, signifying a new high-level action or subgoal being generated (for example, the high-level action leading to macro-step 2 or macro-step 3 in Fig. 3a or a subgoal that reaches the lower yellow square in Fig. 4c (top)).

Episodic memories and cortical–hippocampal binding

Each level of the APC hierarchy learns generic 'basis functions' for representing states and actions based on interactions with the environment. For example, the basis functions learned by the lowest-level state

network when the inputs are natural videos comprise oriented spatiotemporal Gabor filters coding for edges and bars moving at different orientations⁶¹. The neural activity vector is a specific activation pattern coding for the current video segment in terms of the learned

spatiotemporal filters. At the highest level N of the APC model, the specific activation patterns $s^{(N)}$ and $a^{(N)}$ together represent (figure in Box 3, panel e, with $i + 1 = N$) an entire sequence (or timeline⁷⁶) corresponding to the current episode of interaction with the environment, for example, the sequence of glimpses of an object as in Fig. 3a or the sequence of locations visited during navigation as in Fig. 4c (bottom). By ‘binding’ these highest-level neural activation patterns $s^{(N)}$ and $a^{(N)}$ (assumed to correspond to entorhinal cortex activity in the APC model) in a hippocampus-like associative memory, one can store the current sequence of experienced sensations (vision, touch, sound, smell, rewards, etc.), locations (or coordinates within a reference frame) and actions as an episodic memory vector \mathbf{m} (ref. 61).

The projection from the hippocampus back to the entorhinal cortex implies that the fused multimodal information in the current episodic memory vector \mathbf{m} is fed back to enable better prediction (\mathbf{m} plays a role similar to context windows in transformers in AI) and to influence state and action estimation in cortical areas down the hierarchy to the lowest levels. Particularly salient episodic memory vectors may be stored and later compositionally recombined⁷³ or recalled for replay in the cortex when given an internal or external cue, for example, a location where the episode occurred, a sound or smell associated with the episode or a partial visual input marking the beginning of the episode (see ref. 61 for an example).

In summary, the APC model suggests that the cortex encodes generic semantic knowledge about the world within state and action networks that implement nested reference frames. Any particular instantiation of this knowledge invoked by, for example, an interaction with a person or an object, is stored temporarily as an episodic memory vector \mathbf{m} in the hippocampus. This instantiation could be used for reasoning about the current situation or for planning and, if deemed important, could be consolidated within the cortex by updating cortical networks via replay during inactivity or sleep. The idea of fast binding of specific instances (‘fillers’) with generic semantic ‘roles’ is gaining currency in both AI⁵⁸ and hippocampal modeling (for example, the Tolman–Eichenbaum machine⁷⁷; see also ref. 73). The benefits of such a representation, including fast transfer of knowledge and zero-shot learning, can be expected to also accrue to the memory-augmented APC model.

Learning abstract concepts

I briefly sketch here how the same sensory–motor architecture used for perception and planning above could also be potentially used to model abstract concepts. Take, for example, modeling the concept of a family tree. The state–action representations in the APC model can be made categorical (for example, as in ref. 78), allowing states and actions to represent symbols. The states can then represent abstract categories such as father, mother, daughter, uncle, etc., while abstract actions (up, down, etc.) can be used to traverse and define a family tree sequentially. The notion of ‘fast binding’ of cortical representations in hippocampal memory discussed above could be used to bind specific persons to their roles (father, mother, etc.).

Results along these lines were obtained using the Tolman–Eichenbaum machine model⁷⁷, in which a recurrent neural network (similar to the state-transition network in the APC model but for a single level) was used in conjunction with an associative memory to learn the structure of family trees from examples. Extending these ideas to abstract state–action networks for symbolic reasoning in a hierarchical APC model may offer new insights into understanding how cortical–hippocampal networks represent language and solve abstract cognitive tasks such as arithmetic.

Discussion

Inspired by recent results highlighting the influence of actions across most areas of the cortex, I suggested APC as a sensory–motor theory of cortical function. APC proposes that (1) each cortical area implements

both a state-transition network for state prediction and an action network for action (or goal) prediction, and (2) higher-area neurons representing more abstract states and actions modulate lower-area state and action networks via top–down modulatory control to change the functions they are computing, leading to nested reference frames and hierarchical representations of objects, states and actions. A possible neuroanatomical mapping of the APC model to cortical laminar structure was suggested in the section APC module and neuroanatomical implementation.

The APC model lends support to the hypothesis^{20–23} that there may be a unifying computational principle operating across the cortex by showing how the same basic APC architecture can perform a diverse range of computations (see Box 4 for a summary). The APC model shares broad similarities with a number of other models advocating prediction and hierarchy as core aspects of brain function^{1,3,22,23,79–83}, going back to the seminal early work of MacKay⁸⁴ and Albus⁸⁵. The goal of putting action on an equal footing with perception in terms of Bayesian inference and prediction error minimization is in keeping with the theories of free energy minimization proposed by Friston and others^{3,8,69}. In its current formulation, APC addresses action selection via reinforcement learning (see the section Active visual perception and part–whole learning) and planning via model predictive control (as described in the section Planning and navigation using hierarchical world models). The latter is related to planning as inference methods^{63–65} and active inference schemes that optimize expected information gain plus expected value^{8,69}.

Compositionality and the representation of sensory–motor information in cortical columns are also central tenets of the ‘thousand brains’ theory^{23,59}. The close interaction between state-estimation networks and action-computing networks in the APC model is consistent with theories of optimal motor control⁸⁶, especially theories highlighting the importance of internal models in solving the inverse problem of computing optimal motor commands to solve a task^{81,87,88}. However, based on recent evidence pointing to outputs from layer 5 in essentially all cortical areas to subcortical motor centers^{19,34,35,37}, the APC model proposes that all cortical areas include both state-estimation and policy components. M1 is often cited as a uniquely ‘motor’ cortical area missing the sensory input layer 4, with damage to M1 in primates causing permanent loss of distal (although not proximal) movements⁸⁹. However, even M1 receives sensory information from other cortical and subcortical areas⁹⁰, especially in its superficial layers^{39,40}, and could therefore, as suggested by the APC model, predict and estimate state (for example, proprioceptive state) and compute actions based on these state estimates⁹¹.

The APC generative model in the figure in Box 3 (panel a) focuses on hierarchical structure and does not account for cross-modal (sensory to sensory) or hierarchically ‘horizontal’ connections in the neocortex (for example, ref. 92). However, it is possible to extend APC’s generative model to allow cross-modal influences and horizontal interactions to enable more accurate state prediction and estimation. For example, consider a generative model evolved for use by an animal foraging in the forest: the hidden state denoting, for example, a tiger, can generate both a visual cue (stripes) and an auditory cue (rustling sound). In the extended APC model employing such a generative model, the state network in a sensory area (for example, V1) would leverage information from other sensory modalities (for example, from the auditory cortex) via horizontal cortical connections to derive an accurate estimate of the current state of the world. Extending the APC model to account for such cross-modal and horizontal cortico–cortical connections is an important direction for future work.

A large number of unknowns remain such as the exact physiological mechanisms underlying the modulatory interactions between higher-order and lower-order cortical areas across multiple time scales, the role of alpha, beta, theta and gamma oscillations in such interactions and the representation of uncertainty in the cortex.

While there is emerging neurophysiological and neuroanatomical evidence^{2,9–11,13–16,18,41,42,51,55,93} that lends some support to the APC model's predictions (Box 4), there is much that remains to be tested. I hope that the theoretical framework offered by the APC model is helpful in the design of new experiments aimed at uncovering the cortical and subcortical basis of sensory–motor processing and cognition.

References

- Rao, R. P. N. & Ballard, D. H. Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* **2**, 79–87 (1999).
- Keller, G. B. & Mrsic-Flogel, T. D. Predictive processing: a canonical cortical computation. *Neuron* **100**, 424–435 (2018).
- Friston, K. & Kiebel, S. Predictive coding under the free-energy principle. *Philos. Trans. R. Soc. London B Biol. Sci.* **364**, 1211–1221 (2009).
- Jiang, L. P. & Rao, R. P. N. Predictive coding theories of cortical function. *Oxford Research Encyclopedia of Neuroscience* <https://doi.org/10.1093/acrefore/9780190264086.013.328> (Oxford Univ. Press, 2022).
- Halpern, B. P. Tasting and smelling as active, exploratory sensory processes. *Am. J. Otolaryngol.* **4**, 246–249 (1983).
- Lederman, S. J. & Klatzky, R. L. Hand movements: a window into haptic object recognition. *Cogn. Psychol.* **19**, 342–368 (1987).
- Ahissar, E. & Assa, E. Perception as a closed-loop convergence process. *eLife* **5**, e12830 (2016).
- Friston, K. The free-energy principle: a unified brain theory? *Nat. Rev. Neurosci.* **11**, 127–138 (2010).
- Zatka-Haas, P., Steinmetz, N. A., Carandini, M. & Harris, K. D. Sensory coding and the causal impact of mouse cortex in a visual decision. *eLife* **10**, e63163 (2021).
- Steinmetz, N. A., Zatka-Haas, P., Carandini, M. & Harris, K. D. Distributed coding of choice, action and engagement across the mouse brain. *Nature* **576**, 266–273 (2019).
- Stringer, C. et al. Spontaneous behaviors drive multidimensional, brainwide activity. *Science* **364**, eaav7893 (2019).
- Talluri, B. C. et al. Activity in primate visual cortex is minimally driven by spontaneous movements. *Nat. Neurosci.* **26**, 1953–1959 (2023).
- Jordan, R. & Keller, G. B. Opposing influence of top–down and bottom–up input on excitatory layer 2/3 neurons in mouse primary visual cortex. *Neuron* **108**, 1194–1206 (2020).
- Nakamura, K. & Colby, C. L. Updating of the visual representation in monkey striate and extrastriate cortex during saccades. *Proc. Natl Acad. Sci. USA* **99**, 4026–4031 (2002).
- Duhamel, J. R., Colby, C. L. & Goldberg, M. E. The updating of the representation of visual space in parietal cortex by intended eye movements. *Science* **255**, 90–92 (1992).
- Umeno, M. M. & Goldberg, M. E. Spatial processing in the monkey frontal eye field. I. Predictive visual responses. *J. Neurophysiol.* **78**, 1373–1383 (1997).
- Wurtz, R. H., McAlonan, K., Cavanaugh, J. & Berman, R. A. Thalamic pathways for active vision. *Trends Cogn. Sci.* **15**, 177–184 (2011).
- Sherman, S. M. & Guillery, R. W. *Functional Connections of Cortical Areas: A New View from the Thalamus* (MIT, 2013).
- Prasad, J., Carroll, B. & Sherman, S. Layer 5 corticofugal projections from diverse cortical areas: variations on a pattern of thalamic and extrathalamic targets. *J. Neurosci.* **40**, 5785–5796 (2020).
- Mountcastle, V. in *The Mindful Brain* (eds Edelman, G. & Mountcastle, V.) 7–50 (MIT, 1978).
- Creutzfeldt, O. D. Generality of the functional structure of the neocortex. *Naturwissenschaften* **64**, 507–517 (1977).
- Mumford, D. On the computational architecture of the neocortex. II. The role of cortico–cortical loops. *Biol. Cybern.* **66**, 241–251 (1992).
- Hawkins, J. A. *A Thousand Brains: A New Theory of Intelligence* (Basic Books, 2021).
- Douglas, R. J. & Martin, K. A. Neuronal circuits of the neocortex. *Annu. Rev. Neurosci.* **27**, 419–451 (2004).
- Harris, K. D. & Shepherd, G. M. The neocortical circuit: themes and variations. *Nat. Neurosci.* **18**, 170–181 (2015).
- Roe, A. W., Pallas, S. L., Kwon, Y. H. & Sur, M. Visual projections routed to the auditory pathway in ferrets: receptive fields of visual neurons in primary auditory cortex. *J. Neurosci.* **12**, 3651–3664 (1992).
- Felleman, D. & Essen, D. V. Distributed hierarchical processing in the primate cerebral cortex. *Cereb. Cortex* **1**, 1–47 (1991).
- Murray, J. D. et al. A hierarchy of intrinsic timescales across primate cortex. *Nat. Neurosci.* **17**, 1661–1663 (2014).
- Siegle, J. H. et al. Survey of spiking in the mouse visual system reveals functional hierarchy. *Nature* **592**, 86–92 (2021).
- Grafton, S. T. & de C. Hamilton, A. F. Evidence for a distributed hierarchy of action representation in the brain. *Hum. Mov. Sci.* **26**, 590–616 (2007).
- Gklezakos, D. C. & Rao, R. P. N. Active predictive coding networks: a neural solution to the problem of learning reference frames and part–whole hierarchies. Preprint at arxiv.org/abs/2201.08813 (2022).
- Rao, R. P. N., Gklezakos, D. C. & Sathish, V. Active predictive coding: a unifying neural model for active perception, compositional learning, and hierarchical planning. *Neural Comput.* **36**, 1–32 (2024).
- Fisher, A. & Rao, R. P. N. Recursive neural programs: a differentiable framework for learning compositional part–whole hierarchies and image grammars. *PNAS Nexus* **2**, pgad337 (2023).
- Kasper, E., Larkman, A., Lübke, J. & Blakemore, C. Pyramidal neurons in layer 5 of the rat visual cortex. I. Correlation among cell morphology, intrinsic electrophysiological properties, and axon targets. *J. Comp. Neurol.* **339**, 459–474 (1994).
- Stebbins, K., Lesicko, A. & Llano, D. The auditory corticocollicular system: molecular and circuit-level considerations. *Hear. Res.* **314**, 51–59 (2014).
- Xiong, X. et al. Auditory cortex controls sound-driven innate defense behaviour through corticofugal projections to inferior colliculus. *Nat. Commun.* **6**, 7224 (2015).
- Frezel, N. et al. In-depth characterization of layer 5 output neurons of the primary somatosensory cortex innervating the mouse dorsal spinal cord. *Cereb. Cortex Commun.* **1**, tga0052 (2020).
- Rathelot, J. A. & Strick, P. L. Subdivisions of primary motor cortex based on cortico–motoneuronal cells. *Proc. Natl Acad. Sci. USA* **106**, 918–923 (2009).
- Mao, T. et al. Long-range neuronal circuits underlying the interaction between sensory and motor cortex. *Neuron* **72**, 111–123 (2011).
- Hooks, B. M. et al. Organization of cortical and thalamic input to pyramidal neurons in mouse motor cortex. *J. Neurosci.* **33**, 748–760 (2013).
- Heindorf, M., Arber, S. & Keller, G. B. Mouse motor cortex coordinates the behavioral response to unpredicted sensory feedback. *Neuron* **99**, 1040–1054 (2018).
- Holey, B. E. & Schneider, D. M. Sensation and expectation are embedded in mouse motor cortical activity. Preprint at [bioRxiv https://doi.org/10.1101/2023.09.13.557633](https://doi.org/10.1101/2023.09.13.557633) (2023).
- Kim, E., Juavinett, A., Kyubwa, E., Jacobs, M. & Callaway, E. Three types of cortical layer 5 neurons that differ in brain-wide connectivity and function. *Neuron* **88**, 1253–1267 (2015).
- Sutton, R. S. & Barto, A. G. *Reinforcement Learning: An Introduction* 2nd edn <http://incompleteideas.net/book/the-book-2nd.html> (MIT Press, 2018).

45. Kaelbling, L. P., Littman, M. L. & Cassandra, A. R. Planning and acting in partially observable stochastic domains. *Artif. Intell.* **101**, 99–134 (1998).
46. Rao, R. P. N. Decision making under uncertainty: a neural model based on partially observable Markov decision processes. *Front. Comput. Neurosci.* **4**, 146 (2010).
47. von Helmholtz, H. *Handbuch der Physiologischen Optik* Vol. 3 (Voss, 1867).
48. Friston, K., Adams, R. A., Perrinet, L. & Breakspear, M. Perceptions as hypotheses: saccades as experiments. *Front. Psychol.* **3**, 151 (2012).
49. Bastos, A. M. et al. Canonical microcircuits for predictive coding. *Neuron* **76**, 695–711 (2012).
50. Shipp, S. Neural elements for predictive coding. *Front. Psychol.* **7**, 1792 (2016).
51. Miura, S. & Scanziani, M. Distinguishing externally from saccade-induced motion in visual cortex. *Nature* **610**, 135–142 (2022).
52. Keller, G. B., Bonhoeffer, T. & Hübener, M. Sensorimotor mismatch signals in primary visual cortex of the behaving mouse. *Neuron* **74**, 809–815 (2012).
53. Bastos, A. M., Lundqvist, M., Waite, A. S., Kopell, N. & Miller, E. K. Layer and rhythm specificity for predictive routing. *Proc. Natl Acad. Sci. USA* **117**, 31459–31469 (2020).
54. Leinweber, M., Ward, D. R., Sobczak, J. M., Attinger, A. & Keller, G. B. Sensorimotor circuit in mouse cortex for visual flow predictions. *Neuron* **95**, 1420–1432 (2017).
55. Schneider, D. M., Sundararajan, J. & Mooney, R. A cortical filter that learns to suppress the acoustic consequences of movement. *Nature* **561**, 391–395 (2018).
56. Markov, N. T. & Kennedy, H. The importance of being hierarchical. *Curr. Opin. Neurobiol.* **23**, 187–194 (2013).
57. Lake, B. M., Ullman, T. D., Tenenbaum, J. B. & Gershman, S. J. Building machines that learn and think like people. *Behav. Brain Sci.* **40**, e253 (2017).
58. Smolensky, P., McCoy, R. T., Fernandez, R., Goldrick, M. & Gao, J. Neurocompositional computing: from the central paradox of cognition to a new generation of AI systems. *AI Mag.* **43**, 308–322 (2022).
59. Lewis, M., Purdy, S., Ahmad, S. & Hawkins, J. Locations in the neocortex: a theory of sensorimotor object recognition using cortical grid cells. *Front. Neural Circuits* **13**, 22 (2019).
60. Olshausen, B. A. & Field, D. J. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* **381**, 607–609 (1996).
61. Jiang, L. P. & Rao, R. P. N. Dynamic predictive coding: a model of hierarchical sequence learning and prediction in the neocortex. *PLoS Comput. Biol.* **20**, e1011801 (2024).
62. Whittington, J. C. R. & Bogacz, R. An approximation of the error backpropagation algorithm in a predictive coding network with local Hebbian synaptic plasticity. *Neural Comput.* **29**, 1229–1262 (2017).
63. Attias, H. Planning by probabilistic inference. In *Proceedings of the Ninth International Workshop on Artificial Intelligence and Statistics (AISTATS 2003)* (eds Bishop, C. M. & Frey, B. J.) 9–16 (PMLR, 2003).
64. Verma, D. & Rao, R. P. N. Planning and acting in uncertain environments using probabilistic inference. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems* 2382–2387 (IEEE, 2006).
65. Botvinick, M. & Toussaint, M. Planning as inference. *Trends Cogn. Sci.* **16**, 485–488 (2012).
66. Richards, A. *Robust Constrained Model Predictive Control*. PhD thesis, MIT (2004).
67. Botvinick, M. M., Niv, Y. & Barto, A. G. Hierarchically organized behavior and its neural foundations: a reinforcement learning perspective. *Cognition* **113**, 262–280 (2009).
68. Kahneman, D. *Thinking, Fast and Slow* (Farrar, Straus and Giroux, 2011).
69. Friston, K., Parr, T. & de Vries, B. The graphical brain: belief propagation and active inference. *Netw. Neurosci.* **1**, 381–414 (2017).
70. O'Reilly, R. C., Russin, J. L., Zolfaghar, M. & Rohrlich, J. Deep predictive learning in neocortex and pulvinar. *J. Cogn. Neurosci.* **33**, 1158–1196 (2021).
71. Willett, F. R. et al. Hand knob area of premotor cortex represents the whole body in a compositional way. *Cell* **181**, 396–409 (2020).
72. Yang, G. R., Joglekar, M. R., Song, H. F., Newsome, W. T. & Wang, X.-J. Task representations in neural networks trained to perform many cognitive tasks. *Nat. Neurosci.* **22**, 297–306 (2019).
73. Kurth-Nelson, Z. et al. Replay and compositional computation. *Neuron* **111**, 454–469 (2023).
74. Mnih, V., Heess, N., Graves, A. & Kavukcuoglu, K. Recurrent models of visual attention. In *Advances in Neural Information Processing Systems 27* (eds Ghahramani, Z. et al.) 2204–2212 (Curran Associates, 2014).
75. Makino, H. Arithmetic value representation for hierarchical behavior composition. *Nat. Neurosci.* **26**, 140–149 (2023).
76. Hogendoorn, H. Perception in real-time: predicting the present, reconstructing the past. *Trends Cogn. Sci.* **26**, 128–141 (2022).
77. Whittington, J. C. R. et al. The Tolman–Eichenbaum machine: unifying space and relational memory through generalization in the hippocampal formation. *Cell* **183**, 1249–1263 (2020).
78. Hafner, D., Lee, K.-H., Fischer, I. & Abbeel, P. Deep hierarchical planning from pixels. In *Advances in Neural Information Processing Systems 35* (eds Koyejo, S. et al.) 26091–26104 (Curran Associates, 2022).
79. Lee, T. S. & Mumford, D. Hierarchical Bayesian inference in the visual cortex. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* **20**, 1434–1448 (2003).
80. George, D. & Hawkins, J. Towards a mathematical theory of cortical micro-circuits. *PLoS Comput. Biol.* **5**, e1000532 (2009).
81. Wolpert, D. M. & Miall, R. C. Forward models for physiological motor control. *Neural Netw.* **9**, 1265–1279 (1996).
82. Mehta, M. R. Neuronal dynamics of predictive coding. *Neuroscientist* **7**, 490–495 (2001).
83. Heeger, D. J. Theory of cortical function. *Proc. Natl Acad. Sci. USA* **114**, 1773–1782 (2017).
84. Mackay, D. in *Automata Studies* (eds Shannon, C. E. & McCarthy, J.) 235–251 (Princeton Univ., 1956).
85. Albus, J. S. *Brains, Behavior and Robotics* (BYTE, 1981).
86. Scott, S. Optimal feedback control and the neural basis of volitional motor control. *Nat. Rev. Neurosci.* **5**, 532–546 (2004).
87. Jordan, M. I. & Rumelhart, D. E. Forward models: supervised learning with a distal teacher. *Cogn. Sci.* **16**, 307–354 (1992).
88. Kawato, M. Internal models for motor control and trajectory planning. *Curr. Opin. Neurobiol.* **9**, 718–727 (1999).
89. Fetz, E. E. in *Textbook of Physiology* (eds Patton, H. D. et al.) 608–631 (Saunders, 1989).
90. Jones, E. G., Coulter, J. D. & Hendry, S. H. C. Intracortical connectivity of architectonic fields in the somatic sensory, motor and parietal cortex of monkeys. *J. Comp. Neurol.* **181**, 291–347 (1978).
91. Adams, R., Shipp, S. & Friston, K. Predictions not commands: active inference in the motor system. *Brain Struct. Funct.* **218**, 611–643 (2013).
92. Falchier, A., Clavagnier, S., Barone, P. & Kennedy, H. Anatomical evidence of multimodal integration in primate striate cortex. *J. Neurosci.* **22**, 5749–5759 (2002).
93. Audette, N. J., Zhou, W., La Chioma, A. & Schneider, D. M. Precise movement-based predictions in the mouse auditory cortex. *Curr. Biol.* **32**, 4925–4940 (2022).

94. Craik, K. J. W. *The Nature of Explanation* (Macmillan, 1943).
95. Churchland, M. M. et al. Neural population dynamics during reaching. *Nature* **487**, 51–56 (2012).
96. Sussillo, D., Churchland, M. M., Kaufman, M. T. & Shenoy, K. V. A neural network that finds a naturalistic solution for the production of muscle activity. *Nat. Neurosci.* **18**, 1025–1033 (2015).
97. Sauerbrei, B. A. et al. Cortical pattern generation during dexterous movement is input-driven. *Nature* **577**, 386–391 (2020).
98. Chaudhuri, R., Knoblauch, K., Gariel, M.-A., Kennedy, H. & Wang, X.-J. A large-scale circuit mechanism for hierarchical dynamical processing in the primate cortex. *Neuron* **88**, 419–431 (2015).
99. Salinas, E. & Sejnowski, T. J. Gain modulation in the central nervous system: where behavior, neurophysiology, and computation meet. *Neuroscientist* **7**, 430–440 (2001).
100. Larkum, M. E., Senn, W. & Lüscher, H.-R. Top-down dendritic input increases the gain of layer 5 pyramidal neurons. *Cereb. Cortex* **14**, 1059–1070 (2004).
101. Ferguson, K. A. & Cardin, J. A. Mechanisms underlying gain modulation in the cortex. *Nat. Rev. Neurosci.* **21**, 80–92 (2020).
102. Stroud, J. P., Porter, M. A., Hennequin, G. & Vogels, T. P. Motor primitives in space and time via targeted gain modulation in cortical networks. *Nat. Neurosci.* **21**, 1774–1783 (2018).
103. McAdams, C. J. & Maunsell, J. H. R. Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *J. Neurosci.* **19**, 431–441 (1999).
104. Ha, D., Dai, A. M. & Le, Q. V. Hypernetworks. In *5th International Conference on Learning Representations (ICLR 2017)* openreview.net/forum?id=rkpACe1lx (OpenReview.net, 2017).
105. Eliasmith, C. et al. A large-scale model of the functioning brain. *Science* **338**, 1202–1205 (2012).
106. Galanti, T. & Wolf, L. On the modularity of hypernetworks. In *Advances in Neural Information Processing Systems 33* (eds Larochelle, H. et al.) 10409–10419 (Curran Associates, 2020).
107. Tomov, M. S., Yagati, S., Kumar, A., Yang, W. & Gershman, S. J. Discovery of hierarchical representations for efficient planning. *PLoS Comput. Biol.* **16**, e1007594 (2020).
108. Olson, C. R. Brain representation of object-centered space in monkeys and humans. *Annu. Rev. Neurosci.* **26**, 331–354 (2003).
109. George, D. et al. Clone-structured graph representations enable flexible learning and vicarious evaluation of cognitive maps. *Nat. Commun.* **12**, 2392 (2021).
110. Friston, K. J., Rosch, R., Parr, T., Price, C. & Bowman, H. Deep temporal models and active inference. *Neurosci. Biobehav. Rev.* **77**, 388–402 (2017).

Acknowledgements

I thank A. Fisher, D. Gklezakos, P. Jiang, P. Rangarajan and V. Sathish for many discussions and the collaborative work cited in the text. I also thank K. Friston, C. Eliasmith and members of his laboratory, researchers at Numenta, S. Mirbagheri, N. Steinmetz and G. Burachas for discussions and feedback. This work was supported by National Science Foundation EFRI grant 2223495, National Institutes of Health grant 1UF1NS126485-01, the Defense Advanced Research Projects Agency under contract HRO01120C0021, a UW+ Amazon Science Hub grant, a Weill Neurohub Investigator grant, a Frameworks grant from the Templeton World Charity Foundation and a Cherg Jia and Elizabeth Yun Hwang Professorship. The opinions expressed in this publication are those of the author and do not necessarily reflect the views of the funders.

Competing interests

The author declares no competing interests.

Additional information

Correspondence should be addressed to Rajesh P. N. Rao.

Peer review information *Nature Neuroscience* thanks Karl Friston, Aleena Garner, and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

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.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.

© Springer Nature America, Inc. 2024