

# Cognitive computational neuroscience

Nikolaus Kriegeskorte<sup>1\*</sup> and Pamela K. Douglas<sup>2</sup>

**To learn how cognition is implemented in the brain, we must build computational models that can perform cognitive tasks, and test such models with brain and behavioral experiments. Cognitive science has developed computational models that decompose cognition into functional components. Computational neuroscience has modeled how interacting neurons can implement elementary components of cognition. It is time to assemble the pieces of the puzzle of brain computation and to better integrate these separate disciplines. Modern technologies enable us to measure and manipulate brain activity in unprecedentedly rich ways in animals and humans. However, experiments will yield theoretical insight only when employed to test brain-computational models. Here we review recent work in the intersection of cognitive science, computational neuroscience and artificial intelligence. Computational models that mimic brain information processing during perceptual, cognitive and control tasks are beginning to be developed and tested with brain and behavioral data.**

Understanding brain information processing requires that we build computational models that are capable of performing cognitive tasks. The argument in favor of task-performing computational models was well articulated by Allen Newell in 1973 in his commentary “You can’t play 20 questions with nature and win”<sup>1</sup>. Newell was criticizing the state of cognitive psychology. The field was in the habit of testing one hypothesis about cognition at a time, in the hope that forcing nature to answer a series of binary questions would eventually reveal the brain’s algorithms. Newell argued that testing verbally defined hypotheses about cognition might never lead to a computational understanding. Hypothesis testing, in his view, needed to be complemented by the construction of comprehensive task-performing computational models. Only synthesis in a computer simulation can reveal what the interaction of the proposed component mechanisms actually entails and whether it can account for the cognitive function in question. If we did have a full understanding of an information-processing mechanism, then we should be able to engineer it. “What I cannot create, I do not understand,” in the words of physicist Richard Feynman, who left this sentence on his blackboard when he died in 1988.

Here we argue that task-performing computational models that explain how cognition arises from neurobiologically plausible dynamic components will be central to a new cognitive computational neuroscience. We first briefly trace the steps of the cognitive and brain sciences and then review several exciting recent developments that suggest that it might be possible to meet the combined ambitions of cognitive science (to explain how humans learn and think)<sup>2</sup> and computational neuroscience (to explain how brains adapt and compute)<sup>3</sup> using neurobiologically plausible artificial intelligence (AI) models.

In the spirit of Newell’s critique, the transition from cognitive psychology to cognitive science was defined by the introduction of task-performing computational models. Cognitive scientists knew that understanding cognition required AI and brought engineering to cognitive studies. In the 1980s, cognitive science made important advances with symbolic cognitive architectures<sup>4,5</sup> and neural networks<sup>6</sup>, using human behavioral data to adjudicate between candidate computational models. However, computer hardware and machine learning were not sufficiently advanced to simulate

cognitive processes in their full complexity. Moreover, these early developments relied on behavioral data alone and did not leverage constraints provided by the anatomy and activity of the brain.

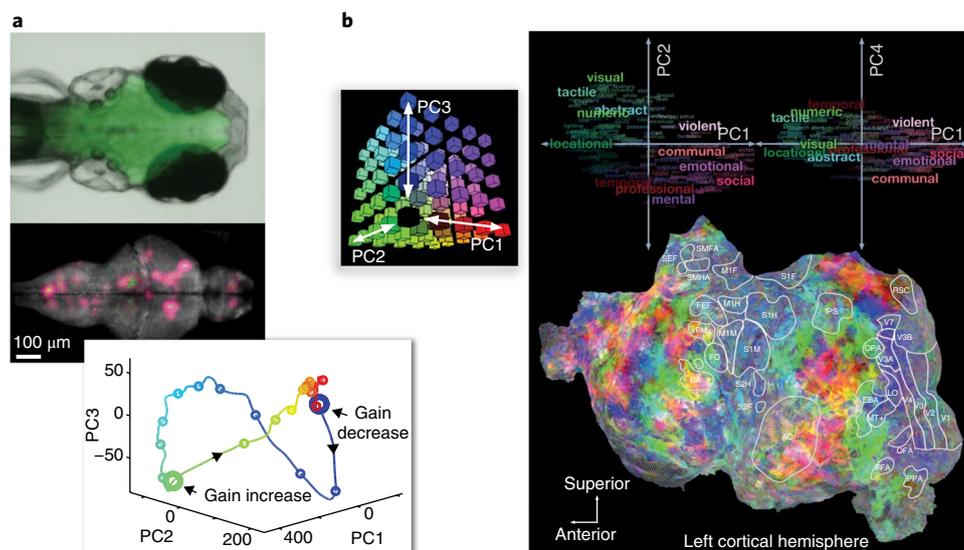
With the advent of human functional brain imaging, scientists began to relate cognitive theories to the human brain. This endeavor came to be called cognitive neuroscience<sup>7</sup>. Cognitive neuroscientists began by mapping cognitive psychology’s boxes (information-processing modules) and arrows (interactions between modules) onto the brain. This was a step forward in terms of engaging brain activity, but a step back in terms of computational rigor. Methods for testing the task-performing computational models of cognitive science with brain-activity data had not been conceived. As a result, cognitive science and cognitive neuroscience parted ways in the 1990s.

Cognitive psychology’s tasks and theories of high-level functional modules provided a reasonable starting point for mapping the coarse-scale organization of the human brain with functional imaging techniques, including electroencephalography, positron emission tomography and early functional magnetic resonance imaging (fMRI), which had low spatial resolution. Inspired by cognitive psychology’s notion of the module<sup>8</sup>, cognitive neuroscience developed its own game of 20 questions with nature. A given study would ask whether a particular cognitive module could be found in the brain. The field mapped an ever increasing array of cognitive functions to brain regions, providing a useful rough draft of the global functional layout of the human brain.

A brain map, at whatever scale, does not reveal the computational mechanism (Fig. 1). However, mapping does provide constraints for theory. After all, information exchange incurs costs that scale with the distance between the communicating regions—costs in terms of physical connections, energy and signal latency. Component placement is likely to reflect these costs. We expect regions that need to interact at high bandwidth and short latency to be placed close together<sup>9</sup>. More generally, the topology and geometry of a biological neural network constrain its dynamics, and thus its functional mechanism. Functional localization results, especially in combination with anatomical connectivity, may therefore ultimately prove useful for modeling brain information processing.

<sup>1</sup>Department of Psychology, Department of Neuroscience, Department of Electrical Engineering, Zuckerman Mind Brain Behavior Institute, Columbia University, New York, NY, USA. <sup>2</sup>Center for Cognitive Neuroscience, University of California, Los Angeles, Los Angeles, CA, USA.

\*e-mail: [n.kriegeskorte@columbia.edu](mailto:n.kriegeskorte@columbia.edu)



**Fig. 1 | Modern imaging techniques provide unprecedentedly detailed information about brain activity, but data-driven analyses support only limited insights.** **a**, Two-photon calcium imaging results<sup>21</sup> show single-neuron activity for a large population of cells measured simultaneously in larval zebrafish while the animals interact with a virtual environment. **b**, Human fMRI results<sup>70</sup> reveal a detailed map of semantically selective responses while a subject listened to a story. These studies illustrate, on the one hand, the power of modern brain-activity measurement techniques at different scales (**a**, **b**) and, on the other, the challenge of drawing insights about brain computation from such datasets. Both studies measured brain activity during complex, time-continuous, naturalistic experience and used principal component analysis (**a**, bottom; **b**, top) to provide an overall view of the activity patterns and their representational significance. PC, principal component.

Despite methodological challenges<sup>10,11</sup>, many of the findings of cognitive neuroscience provide a solid basis on which to build. For example, the findings of face-selective regions in the human ventral stream<sup>12</sup> have been thoroughly replicated and generalized. Nonhuman primates probed with fMRI exhibit similar face-selective regions, which had evaded explorations with invasive electrodes because the latter do not provide continuous images over large fields of view. Localized with fMRI and probed with invasive electrode recordings, the primate face patches revealed high densities of face-selective neurons<sup>13</sup>, with invariances emerging at higher stages of hierarchical processing, including mirror-symmetric tuning and view-tolerant representations of individual faces in the anterior-most patch<sup>14</sup>. The example of face perception illustrates, on one hand, the solid progress in mapping the anatomical substrate and characterizing neuronal responses<sup>15</sup> and, on the other, the lack of definitive computational models. The literature does provide clues to the computational mechanism. A brain-computational model of face recognition<sup>16</sup> will have to explain the spatial clusters of face-selective units and the selectivities and invariances observed with fMRI<sup>17,18</sup> and invasive recordings<sup>14,19</sup>.

Cognitive neuroscience has mapped the global functional layout of the human and nonhuman primate brain<sup>20</sup>. However, it has not achieved a full computational account of brain information processing. The challenge ahead is to build computational models of brain information processing that are consistent with brain structure and function and perform complex cognitive tasks. The following recent developments in cognitive science, computational neuroscience and artificial intelligence suggest that this may be achievable.

1. Cognitive science has proceeded from the top down, decomposing complex cognitive processes into their computational components. Unencumbered by the need to make sense of brain data, it has developed task-performing computational models at the cognitive level. One success story is that of Bayesian cognitive models, which optimally combine prior knowledge about the world with sensory evidence<sup>21–23</sup>. Initially applied to basic sensory and motor processes<sup>23,24</sup>, Bayesian models have begun to engage complex

cognition, including the way our minds model the physical and social world<sup>2</sup>. These developments occurred in interaction with statistics and machine learning, where a unified perspective on probabilistic empirical inference has emerged. This literature provides essential computational theory for understanding the brain. In addition, it provides algorithms for approximate inference on generative models that can grow in complexity with the available data—as might be required for real-world intelligence<sup>25,26</sup>.

2. Computational neuroscience has taken a bottom-up approach, demonstrating how dynamic interactions between biological neurons can implement computational component functions. In the past two decades, the field developed mathematical models of elementary computational components and their implementation with biological neurons<sup>27,28</sup>. These include components for sensory coding<sup>29,30</sup>, normalization<sup>31</sup>, working memory<sup>32</sup>, evidence accumulation and decision mechanisms<sup>33–35</sup>, and motor control<sup>36</sup>. Most of these component functions are computationally simple, but they provide building blocks for cognition. Computational neuroscience has also begun to test complex computational models that can explain high-level sensory and cognitive brain representations<sup>37,38</sup>.

3. Artificial intelligence has shown how component functions can be combined to create intelligent behavior. Early AI failed to live up to its promise because the rich world knowledge required for feats of intelligence could not be either engineered or automatically learned. Recent advances in machine learning, boosted by growing computational power and larger datasets from which to learn, have brought progress at perceptual<sup>39</sup>, cognitive<sup>40</sup> and control challenges<sup>41</sup>. Many advances were driven by cognitive-level symbolic models. Some of the most important recent advances are driven by deep neural network models, composed of units that compute linear combinations of their inputs, followed by static nonlinearities<sup>42</sup>. These models employ only a small subset of the dynamic capabilities of biological neurons, abstracting from fundamental features such as action potentials. However, their functionality is inspired by brains and could be implemented with biological neurons.

The three disciplines contribute complementary elements to biologically plausible computational models that perform cognitive tasks and explain brain information processing and behavior (Fig. 2). Here we review the first steps in the literature toward a cognitive computational neuroscience that meets the combined criteria for success of cognitive science (computational models that perform cognitive tasks and explain behavior) and computational neuroscience (neurobiologically plausible mechanistic models that explain brain activity). If computational models are to explain animal and human cognition, they will have to perform feats of intelligence. AI, and in particular machine learning, is therefore a key discipline that provides the theoretical and technological foundation for cognitive computational neuroscience.

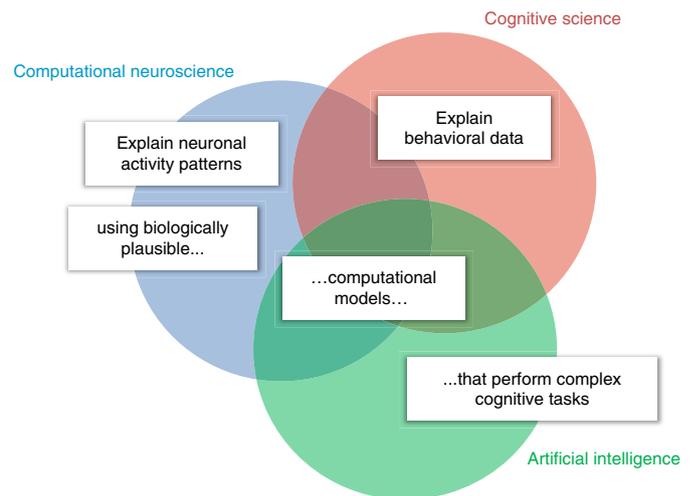
The overarching challenge is to build solid bridges between theory (instantiated in task-performing computational models) and experiment (providing brain and behavioral data). The first part of this review describes bottom-up developments that begin with experimental data and attempt to build bridges from the data in the direction of theory<sup>13</sup>. Given brain-activity data, connectivity models aim to reveal the large-scale dynamics of brain activation; decoding and encoding models aim to reveal the content and format of brain representations. The models employed in this literature provide constraints for computational theory, but they do not in general perform the cognitive tasks in question and thus fall short of explaining the computational mechanism underlying task performance.

The second part of this article describes developments that proceed in the opposite direction, building bridges from theory to experiment<sup>37,38,44</sup>. We review emerging work that has begun to test task-performing computational models with brain and behavioral data. The models include cognitive models, specified at an abstract computational level, whose implementation in biological brains has yet to be explained, and neural network models, which abstract from many features of neurobiology, but could plausibly be implemented with biological neurons. This emerging literature suggests the beginnings of an integrative approach to understanding brain computation, where models are required to perform cognitive tasks, biology provides the admissible component functions, and the computational mechanisms are optimized to explain detailed patterns of brain activity and behavior.

### From experiment toward theory

**Models of connectivity and dynamics.** One path from measured brain activity toward a computational understanding is to model the brain's connectivity and dynamics. Connectivity models go beyond the localization of activated regions and characterize the interactions between regions. Neuronal dynamics can be measured and modeled at multiple scales, from local sets of interacting neurons to whole-brain activity<sup>45</sup>. A first approximation of brain dynamics is provided by the correlation matrix among the measured response time series, which characterizes the pairwise 'functional connectivity' between locations. The literature on resting-state networks has explored this approach<sup>46</sup>, and linear decompositions of the space-time matrix, such as spatial independent component analysis, similarly capture simultaneous correlations between locations across time<sup>47</sup>.

By thresholding the correlation matrix, the set of regions can be converted into an undirected graph and studied with graph-theoretic methods. Such analyses can reveal 'communities' (sets of strongly interconnected regions), 'hubs' (regions connected to many others) and 'rich clubs' (communities of hubs)<sup>48</sup>. Connectivity graphs can be derived from either anatomical or functional measurements. The anatomical connectivity matrix typically resembles the functional connectivity matrix because regions interact through anatomical pathways. However, the way anatomical connectivity generates functional connectivity is better modeled by taking local



**Fig. 2 | What does it mean to understand how the brain works?** The goal of cognitive computational neuroscience is to explain rich measurements of neuronal activity and behavior in animals and humans by means of biologically plausible computational models that perform real-world cognitive tasks. Historically, each of the disciplines (circles) has tackled a subset of these challenges (white labels). Cognitive computational neuroscience strives to meet all the challenges simultaneously.

dynamics, delays, indirect interactions and noise into account<sup>49</sup>. From local neuronal interactions to large-scale spatiotemporal patterns spanning cortex and subcortical regions, generative models of spontaneous dynamics can be evaluated with brain-activity data.

Effective connectivity analyses take a more hypothesis-driven approach, characterizing the interactions among a small set of regions on the basis of generative models of the dynamics<sup>50</sup>. Whereas activation mapping maps the boxes of cognitive psychology onto brain regions, effective connectivity analyses map the arrows onto pairs of brain regions. Most work in this area has focused on characterizing interactions at the level of the overall activation of a brain region. Like the classical brain mapping approach, these analyses are based on regional-mean activation, measuring correlated fluctuations of overall regional activation rather than the information exchanged between regions.

Analyses of effective connectivity and large-scale brain dynamics go beyond generic statistical models such as the linear models used in activation and information-based brain mapping in that they are generative models: they can generate data at the level of the measurements and are models of brain dynamics. However, they do not capture the represented information and how it is processed in the brain.

**Decoding models.** Another path toward understanding the brain's computational mechanisms is to reveal what information is present in each brain region. Decoding can help us go beyond the notion of activation, which indicates the involvement of a region in a task, and reveal the information present in a region's population activity. When particular content is decodable from activity in a brain region, this indicates the presence of the information. To refer to the brain region as 'representing' the content adds a functional interpretation<sup>51</sup>: that the information serves the purpose of informing regions receiving these signals about the content. Ultimately, this interpretation needs to be substantiated by further analyses of how the information affects other regions and behavior<sup>52–54</sup>.

Decoding has its roots in the neuronal-recording literature<sup>27</sup> and has become a popular tool for studying the content of representations in neuroimaging<sup>55–59</sup>. In the simplest case, decoding reveals which of two stimuli gave rise to a measured response pattern.

**Box 1 | The many meanings of “model”**

The word “model” has many meanings in the brain and behavioral sciences. *Data-analysis models* are generic statistical models that help establish relationships between measured variables. Examples include linear correlation, univariate multiple linear regression for brain mapping, and linear decoding analysis. Effective connectivity and causal-interaction models are, similarly, data-analysis models. They help us infer causal influences and interactions between brain regions. Data-analysis models can serve the purpose of testing hypotheses about relationships among variables (for example, correlation, information, causal influence). They are not models of brain information processing. A *box-and-arrow model*, by contrast, is an information-processing model in the form of labeled boxes that represent cognitive component functions and arrows that represent information flow. In cognitive psychology, such models provided useful, albeit ill-defined, sketches for theories of brain computation. A *word model*, similarly, is a sketch for a theory about brain information processing that is defined vaguely by a verbal description. While these are models of information processing, they do not perform the information processing thought to occur in the brain. An *oracle model* is a model of brain responses (often instantiated in a data-analysis model) that relies on information not available to the animal whose brain is being modeled. For example, a model of ventral temporal visual responses as a function of an abstract shape description, or as a function of category labels or continuous semantic features, constitutes an oracle model if the model is not capable of computing the shape, category or semantic features from images. An oracle model may provide a useful characterization of the information present in a region and its representational format, without specifying any theory as to how the representation is computed by the brain. A *brain-computational model (BCM)*, by contrast, is a model that mimics the brain information processing underlying the performance of some task at some level of abstraction. In visual neuroscience, for example, an *image-computable*

*model* is a BCM of visual processing that takes image bitmaps as inputs and predicts brain activity and/or behavioral responses. Deep neural nets provide image-computable models of visual processing. However, deep neural nets trained by supervision rely on category-labeled images for training. Because labeled examples are not available (in comparable quantities) during biological development and learning, these models are BCMs of visual processing, but they are not BCMs of development and learning. *Reinforcement learning models* use environmental feedback that is more realistic in quality and can provide BCMs of learning processes. A *sensory encoding model* is a BCM of the computations that transform sensory input to some stage of internal representation. An *internal-transformation model* is a BCM of the transformation of representations between two stages of processing. A *behavioral decoding model* is a BCM of the transformation from some internal representation to a behavioral output. Note that the label BCM indicates merely that the model is intended to capture brain computations at some level of abstraction. A BCM may abstract from biological detail to an arbitrary degree, but must predict some aspect of brain activity and/or behavior. *Psychophysical models* that predict behavioral outputs from sensory input and *cognitive models* that perform cognitive tasks are BCMs formulated at a high level of description. The label BCM does not imply that the model is either plausible or consistent with empirical data. Progress is made by rejecting candidate BCMs on empirical grounds. Like microscale *biophysical models*, which capture biological processes that underlie brain computations, and macroscale *brain-dynamical and causal-interaction models*, BCMs are models of processes occurring in the brain. However, unlike the other types of process model, BCMs perform the information processing that is thought to be the function of brain dynamics. Finally, the term “model” is used to refer to models of the world employed by the brain, as in *model-based reinforcement learning* and *model-based cognition*.

The content of the representation can be the identity of a sensory stimulus (to be recognized among a set of alternative stimuli), a stimulus property (such as the orientation of a grating), an abstract variable needed for a cognitive operation, or an action<sup>60</sup>. When the decoder is linear, as is usually the case, the decodable information is in a format that can plausibly be read out by downstream neurons in a single step. Such information is said to be ‘explicit’ in the activity patterns<sup>61</sup>.

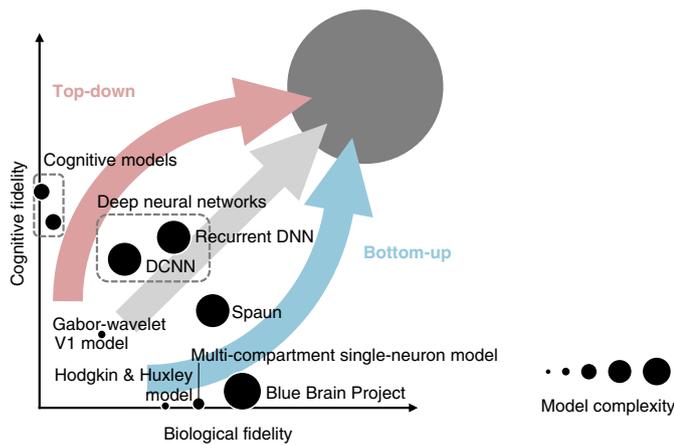
Decoding and other types of multivariate pattern analysis have helped reveal the content of regional representations<sup>55,56,58,59</sup>, providing evidence that brain-computational models must incorporate. However, the ability to decode particular information does not amount to a full account of the neuronal code: it doesn’t specify the representational format (beyond linear decodability) or what other information might additionally be present. Most importantly, decoders do not in general constitute models of brain computation. They reveal aspects of the product, but not the process of brain computation.

**Representational models.** Beyond decoding, we would like to exhaustively characterize a region’s representation, explaining its responses to arbitrary stimuli. A full characterization would also define to what extent any variable can be decoded. Representational models attempt to make comprehensive predictions about the representational space and therefore provide stronger constraints on the computational mechanism than decoding models<sup>52,62</sup>.

Three types of representational model analysis have been introduced in the literature: encoding models<sup>63–65</sup>, pattern component models<sup>66</sup> and representational similarity analysis<sup>37,67,68</sup>. These three methods all test hypotheses about the representational space, which are based on multivariate descriptions of the experimental conditions—for example, a semantic description of a set of stimuli, or the activity patterns across a layer of a neural network model that processes the stimuli<sup>52</sup>.

In encoding models, each voxel’s activity profile across stimuli is predicted as a linear combination of the features of the model. In pattern component models, the distribution of the activity profiles that characterizes the representational space is modeled as a multivariate normal distribution. In representational similarity analysis, the representational space is characterized by the representational dissimilarities of the activity patterns elicited by the stimuli.

Representational models are often defined on the basis of descriptions of the stimuli, such as labels provided by human observers<sup>63,69,70</sup>. In this scenario, a representational model that explains the brain responses in a given region provides, not a brain-computational account, but at least a descriptive account of the representation. Such an account can be a useful stepping-stone toward computational theory when the model generalizes to novel stimuli. Importantly, representational models also enable us to adjudicate among brain-computational models, an approach we will return to in the next section.



**Fig. 3 | The space of process models.** Models of the processes taking place in the brain can be defined at different levels of description and can vary in their parametric complexity (dot size) and in their biological (horizontal axis) and cognitive (vertical axis) fidelity. Theoreticians approach modeling with a range of primary goals. The bottom-up approach to modeling (blue arrow) aims first to capture characteristics of biological neural networks, such as action potentials and interactions among multiple compartments of single neurons. This approach disregards cognitive function so as to focus on understanding the emergent dynamics of small parts of the brain, such as cortical columns and areas, and to reproduce biological network phenomena, such as oscillations. The top-down approach (red arrow) aims first to capture cognitive functions at the algorithmic level. This approach disregards the biological implementation so as to focus on decomposing the information processing underlying task performance into its algorithmic components. The two approaches form the extremes of a continuum of paths toward the common goal of explaining how our brains give rise to our minds. Overall, there is tradeoff (negative correlation) between cognitive and biological fidelity. However, the tradeoff can turn into a synergy (positive correlation) when cognitive constraints illuminate biological function and when biology inspires models that explain cognitive feats. Because intelligence requires rich world knowledge, models of human brain information processing will have high parametric complexity (large dots in the upper right corner). Even if models that abstract from biological details can explain task performance, biologically detailed models will still be needed to explain the neurobiological implementation. This diagram is a conceptual cartoon that can help us understand the relationships between models and appreciate their complementary contributions. However, it is not based on quantitative measures of cognitive fidelity, biological fidelity and model complexity. Definitive ways to measure each of the three variables have yet to be developed. Figure inspired by ref. <sup>122</sup>.

In this section, we considered three types of model that can help us glean computational insight from brain-activity data. Connectivity models capture aspects of the dynamic interactions between regions. Decoding models enable us to look into brain regions and reveal what might be their representational content. Representational models enable us to test explicit hypotheses that fully characterize a region's representational space. All three types of model can be used to address theoretically motivated questions—taking a hypothesis-driven approach. However, in the absence of task-performing computational models, they are subject to Newell's argument that asking a series of questions might never reveal the computational mechanism underlying the cognitive feat we are trying to explain. These methods fall short of building the bridge all the way to theory because they do not test mechanistic models that specify precisely how the information processing underlying some cognitive function might work.

### From theory to experiment

To build a better bridge between experiment and theory, we first need to fully specify a theory. This can be achieved by defining the theory mathematically and implementing it in a computational model (Box 1). Computational models can reside at different levels of description, trading off cognitive fidelity against biological fidelity (Fig. 3). Models designed to capture only neuronal components and dynamics<sup>71</sup> tend to be unsuccessful at explaining cognitive function<sup>72</sup> (Fig. 3, horizontal axis). Conversely, models designed to capture only cognitive functions are difficult to relate to the brain (Fig. 3, vertical axis). To link mind and brain, models must attempt to capture aspects of both behavior and neuronal dynamics. Recent advances suggest that constraints from the brain can help explain cognitive function<sup>42,73,74</sup> and vice versa<sup>37,38</sup>, turning the tradeoff into a synergy.

In this section, we focus on recent successes with task-performing models that explain cognitive functions in terms of representations and algorithms. Task-performing models have been central to psychophysics and cognitive science, where they are traditionally tested with behavioral data. An emerging literature is beginning to test task-performing models with brain-activity data as well. We will consider two broad classes of model in turn, neural network models and cognitive models.

**Neural network models.** Neural network models (Box 2) have a long history, with interwoven strands in multiple disciplines. In computational neuroscience, neural network models, at various levels of biological detail, have been essential to understanding dynamics in biological neural networks and elementary computational functions<sup>27,28</sup>. In cognitive science, they defined a new paradigm for understanding cognitive functions, called parallel distributed processing, in the 1980s<sup>6,75</sup>, which brought the field closer to neuroscience. In AI, they have recently brought substantial advances in a number of applications<sup>42,74</sup>, ranging from perceptual tasks (such as vision and speech recognition) to symbolic processing challenges (such as language translation), and on to motor tasks (including speech synthesis and robotic control). Neural network models provide a common language for building task-performing models that meet the combined criteria for success of the three disciplines (Fig. 2).

Like brains, neural network models can perform feedforward as well as recurrent computations<sup>37,76</sup>. The models driving the recent advances are deep in the sense that they comprise multiple stages of linear-nonlinear signal transformation. Models typically have millions of parameters (the connection weights), which are set so as to optimize task performance. One successful paradigm is supervised learning, wherein a desired mapping from inputs to outputs is learned from a training set of inputs (for example, images) and associated outputs (for example, category labels). However, neural network models can also be trained without supervision and can learn complex statistical structure inherent to their experiential data.

The large number of parameters creates unease among researchers who are used to simple models with small numbers of interpretable parameters. However, simple models will never enable us to explain complex feats of intelligence. The history of AI has shown that intelligence requires ample world knowledge and sufficient parametric complexity to store it. We therefore must engage complex models (Fig. 3) and the challenges they pose. One challenge is that the high parameter count renders the models difficult to understand. Because the models are entirely transparent, they can be probed cheaply with millions of input patterns to understand the internal representations, an approach sometimes called 'synthetic neurophysiology'. To address the concern of overfitting, models are evaluated in terms of their generalization performance. A vision model, for example, will be evaluated in terms of its ability to predict neural activity and behavioral responses for images it has not been trained on.

**Box 2 | Neural network models**

The term “neural network model” has come to be associated with a class of model that is inspired by biological neural networks in that each unit combines many inputs and information is processed in parallel through a network. In contrast to biologically detailed models, which may capture action potentials and dynamics in multiple compartments of each neuron, these models abstract from the biological details. However, they can explain certain cognitive functions, such as visual object recognition, and therefore provide an attractive framework for linking cognition to the brain.

A typical unit computes a linear combination of its inputs and passes the result through a static nonlinearity. The output is sometimes interpreted as analogous to the firing rate of a neuron. Even *shallow* networks (those with a single layer of hidden units between inputs and outputs) can approximate arbitrary functions<sup>123</sup>. However, *deep* networks (those with multiple hidden layers) can more efficiently capture many of the complex functions needed in real-world tasks. Many applications—for example, in computer vision—use feedforward architectures. However, recurrent neural networks, which reprocess the outputs of their units and generate complex dynamics, have brought additional engineering advances<sup>76</sup> and better capture the recurrent signaling in brains<sup>35,124–126</sup>. Whereas feedforward networks are universal function approximators, recurrent networks are universal approximators of dynamical systems<sup>127</sup>. Recurrent processing enables a network to recycle its limited computational resources through time so as to perform more complex sequences of computations. Recurrent networks can represent the recent stimulus history in a dynamically compressed format, providing the temporal context information needed for current processing. As a result, recurrent networks can recognize, predict, and generate dynamical patterns.

Both feedforward and recurrent networks are defined by their architecture and the setting of the connection weights. One way to set the weights is through iterative small adjustments that bring the output closer to some desired output (supervised

learning). Each weight is adjusted in proportion to the reduction in the error that a small change to it would yield. This method is called *gradient descent* because it produces steps in the space of weights along which the error declines most steeply. Gradient descent can be implemented using *backpropagation*, an efficient algorithm for computing the derivative of the error function with respect to each weight.

Whether the brain uses an algorithm like backpropagation for learning is controversial. Several biologically plausible implementations of backpropagation or closely related forms of supervised learning have been suggested<sup>128–130</sup>. Supervision signals might be generated internally<sup>131</sup> on the basis of the context provided by multiple sensory modalities; on the basis of the dynamic refinement of representations over time, as more evidence becomes available from the senses and from memory<sup>132</sup>; and on the basis of internal and external reinforcement signals arising in interaction with the environment<sup>133</sup>. Reinforcement learning<sup>41</sup> and unsupervised learning of neural network parameters<sup>119,134</sup> are areas of rapid current progress.

Neural network models have demonstrated that taking inspiration from biology can yield breakthroughs in AI. It seems likely that the quest for models that can match human cognitive abilities will draw us deeper into the biology<sup>135</sup>. The abstract neural network models currently most successful in engineering could be implemented with biological hardware. However, they only use a small subset of the dynamical components of brains. Neuroscience has described a rich repertoire of dynamical components, including action potentials<sup>108</sup>, canonical microcircuits<sup>136</sup>, dendritic dynamics<sup>128,130,137</sup> and network phenomena<sup>27</sup>, such as oscillations<sup>138</sup>, that may have computational functions. Biology also provides constraints on the global architecture, suggesting, for example, complementary subsystems for learning<sup>139</sup>. Modeling these biological components in the context of neural networks designed to perform meaningful tasks may reveal how they contribute to brain computation and may drive further advances in AI.

Several recent studies have begun to test neural network models as models of brain information processing<sup>37,38</sup>. These studies predicted brain representations of novel images in the primate ventral visual stream with deep convolutional neural network models trained to recognize objects in images. Results have shown that the internal representations of deep convolutional neural networks provide the best current models of representations of visual images in inferior temporal cortex in humans and monkeys<sup>77–79</sup>. When comparing large numbers of models, those that were optimized to perform the task of object classification better explained the cortical representation<sup>77,78</sup>.

Early layers of deep neural networks trained to recognize objects contain representations resembling those in early visual cortex<sup>78,80</sup>. As we move along the ventral visual stream, higher layers of the neural networks come to provide a better basis for explaining the representations<sup>80–82</sup>. Higher layers of deep convolutional neural networks also resemble the inferior temporal cortical representation in that both enable the decoding of object position, size and pose, along with the category of the object<sup>83</sup>. In addition to testing these models by predicting brain-activity data, the field has begun to test them by predicting behavioral responses reflecting perceived shape<sup>84</sup> and object similarity<sup>85</sup>.

**Cognitive models.** Models at the cognitive level enable researchers to envision the information processing without simultaneously

having to tackle its implementation with neurobiologically plausible components. This enables progress on domains of higher cognition, where neural network models still fall short. Moreover, a cognitive model may provide a useful abstraction, even when a process can also be captured with a neural network model.

Neuroscientific explanations now dominate for functional components closer to the periphery of the brain, where sensory and motor processes connect the animal to its environment. However, much of higher-level cognition has remained beyond the reach of neuroscientific accounts and neural network models. To illustrate some of the unique contributions of cognitive models, we briefly discuss three classes of cognitive model: production systems, reinforcement learning models and Bayesian cognitive models.

Production systems provide an early example of a class of cognitive models that can explain reasoning and problem solving. These models use rules and logic, and are symbolic in that they operate on symbols rather than sensory data and motor signals. They capture cognition, rather than perception and motor control, which ground cognition in the physical environment. A ‘production’ is a cognitive action triggered according to an if-then rule. A set of such rules specifies the conditions (‘if’) under which each of a range of productions (‘then’) is to be executed. The conditions refer to current goals and knowledge in memory. The actions can modify the internal state of goals and knowledge. For example, a production may create a subgoal or store an inference. If conditions are met for multiple

## Box 3 | Bayesian cognitive models

*Bayesian cognitive models* are motivated by the assumption that the brain approximates the statistically optimal solution to a task. The statistically optimal way to make inferences and decide what to do is to interpret the current sensory evidence in light of all available prior knowledge using the rules of probability. Consider the case of visual perception. The retinal signals reflect the objects in the world, which we would like to recognize. To infer the objects, we should consider what configurations of objects we deem possible and how well each explains the image. Our prior beliefs are represented by a *generative model* that captures the probability of each configuration of objects and the probabilities with which a given configuration would produce different retinal images.

More formally, a Bayesian model of vision might use a generative model of the joint distribution  $p(\mathbf{d}, \mathbf{c})$  of the sensory data  $\mathbf{d}$  (the image) and the causes in the world  $\mathbf{c}$  (the configuration of surfaces, objects and light sources to be inferred)<sup>140</sup>. The joint distribution  $p(\mathbf{d}, \mathbf{c})$  equals the product of the *prior*,  $p(\mathbf{c})$ , over all possible configurations of causes and the *likelihood*,  $p(\mathbf{d}|\mathbf{c})$ , the probability of a particular image given a particular configuration of causes. A prescribed model for  $p(\mathbf{d}|\mathbf{c})$  would enable us to evaluate the likelihood, the probability of a specific image  $\mathbf{d}$  given specific causes  $\mathbf{c}$ . Alternatively, we might have an implicit model for  $p(\mathbf{d}|\mathbf{c})$  in the form of a stochastic mapping from causes  $\mathbf{c}$  to data  $\mathbf{d}$  (images). Such a model would generate natural images. Whether prescribed or implicit, the model of  $p(\mathbf{d}|\mathbf{c})$  captures how the causes in the world create the image, or at least how they relate to the image. Visual recognition amounts to computing the *posterior*  $p(\mathbf{c}|\mathbf{d})$ , the probability distribution over the causes given a particular image. The posterior  $p(\mathbf{c}|\mathbf{d})$  reveals the causes  $\mathbf{c}$  as they would have to exist in the world to explain the sensory data  $\mathbf{d}$ <sup>141</sup>. A model computing  $p(\mathbf{c}|\mathbf{d})$  is called a *discriminative model* because it discriminates among images—here mapping from effects (the image) to the causes. The inversion mathematically requires a prior  $p(\mathbf{c})$  over the latent causes. The prior  $p(\mathbf{c})$  can constrain the interpretation and help

reduce the ambiguity resulting from the multiple configurations of causes that can account for any image.

Basing the inference of the causes  $\mathbf{c}$  on a generative model of  $p(\mathbf{d}, \mathbf{c})$  that captures all available knowledge and uncertainty is statistically optimal (i.e., it provides the best inferences given limited data), but computationally challenging (i.e., it may require more neurons or time than the animal can use). Ideally, the generative model  $p(\mathbf{d}, \mathbf{c})$  implicit to the inference  $p(\mathbf{c}|\mathbf{d})$  should capture our knowledge not just about image formation, but also the things in the world and their interactions, and our uncertainties about these processes. One challenge is to learn a generative model from sensory data. We need to represent the learned knowledge and the remaining uncertainties. If the generative model is mis-specified, then the inference will not be optimal. For real-world tasks, some degree of misspecification of the model is inevitable. For example, the generative model may contain an overly simplified version of the image-generation process. Another challenge is the computation of the posterior  $p(\mathbf{c}|\mathbf{d})$ . For realistically complex generative models, the inference may require computationally intensive iterative algorithms such as *Markov chain Monte Carlo*, *belief propagation* or *variational inference*. The brain's compromise between statistical and computational efficiency<sup>142–144</sup> may involve learning fast feedforward recognition models that speed up frequent component inferences, crystallizing conclusions that are costly to fluidly derive with iterative algorithms. This is known as *amortized inference*<sup>145,146</sup>.

Bayesian cognitive models have recently flourished in interaction with machine learning and statistics. Early work used generative models with a fixed structure that were flexible only with respect to a limited set of parameters. Modern generative models can grow in complexity with the data and discover their inherent structure<sup>98</sup>. They are called *nonparametric* because they are not limited by a predefined finite set of parameters<sup>147</sup>. Their parameters can grow in number without any predefined bound.

rules, a conflict-resolution mechanism chooses one production. A model specified using this formalism will generate a sequence of productions, which may to some extent resemble our conscious stream of thought while working toward some cognitive goal. The formalism of production systems also provides a universal computational architecture<sup>86</sup>. Production systems such as ACT-R<sup>5</sup> were originally developed under the guidance of behavioral data. More recently such models have also begun to be tested in terms of their ability to predict regional-mean fMRI activation time courses<sup>87</sup>.

Reinforcement learning models capture how an agent can learn to maximize its long-term cumulative reward through interaction with its environment<sup>88,89</sup>. As in production systems, reinforcement learning models often assume that the agent has perception and motor modules that enable the use of discrete symbolic representations of states and actions. The agent chooses actions, observes resulting states of the environment, receives rewards along the way and learns to improve its behavior. The agent may learn a 'value function' associating each state with its expected cumulative reward. If the agent can predict which state each action leads to and if it knows the values of those states, then it can choose the most promising action. The agent may also learn a 'policy' that associates each state directly with promising actions. The choice of action must balance exploitation (which brings short-term reward) and exploration (which benefits learning and brings long-term reward).

The field of reinforcement learning explores algorithms that define how to act and learn so as to maximize cumulative reward.

With roots in psychology and neuroscience, reinforcement learning theory is now an important field of machine learning and AI. It provides a very general perspective on control that includes the classical techniques dynamic programming, Monte Carlo and exhaustive search as limiting cases, and can handle challenging scenarios in which the environment is stochastic and only partially observed, and its causal mechanisms are unknown.

An agent might exhaustively explore an environment and learn the most promising action to take in any state by trial and error (model-free control). This would require sufficient time to learn, enough memory, and an environment that does not kill the agent prematurely. Biological organisms, however, have limited time to learn and limited memory, and must avoid interactions that might kill them. Under these conditions, an agent might do better to build a model of its environment. A model can compress and generalize experience to enable intelligent action in novel situations (model-based control). Model-free methods are computationally efficient (mapping from states to values or directly to actions), but statistically inefficient (learning takes long); model-based methods are more statistically efficient, but may require prohibitive amounts of computation (to simulate possible futures)<sup>90</sup>.

Until experience is sufficient to build a reliable model, an agent might do best to simply store episodes and revert to paths of action that have met with success in the past (episodic control)<sup>91,92</sup>. Storing episodes preserves sequential dependency information important for model building. Moreover, episodic

**Box 4 | Why do cognitive science, computational neuroscience and AI need one another?**

**Cognitive science** needs computational neuroscience, not merely to explain the implementation of cognitive models in the brain, but also to discover the algorithms. For example, the dominant models of sensory processing and object recognition are brain-inspired neural networks, whose computations are not easily captured at a cognitive level. Recent successes with Bayesian nonparametric models do not yet in general scale to real-world cognition. Explaining the computational efficiency of human cognition and predicting detailed cognitive dynamics and behavior could benefit from studying brain-activity dynamics. Explaining behavior is essential, but behavioral data alone provide insufficient constraints for complex models. Brain data can provide rich constraints for cognitive algorithms if leveraged appropriately. Cognitive science has always progressed in close interaction with artificial intelligence. The disciplines share the goal of building task-performing models and thus rely on common mathematical theory and programming environments.

**Computational neuroscience** needs cognitive science to challenge it to engage higher-level cognition. At the experimental level, the tasks of cognitive science enable computational neuroscience to bring cognition into the lab. At the level of theory, cognitive science challenges computational neuroscience to explain how the neurobiological dynamical components it studies contribute to cognition and behavior. Computational neuroscience needs AI, and in particular machine learning, to provide the theoretical and technological

basis for modeling cognitive functions with biologically plausible dynamical components.

**Artificial intelligence** needs cognitive science to guide the engineering of intelligence. Cognitive science's tasks can serve as benchmarks for AI systems, building up from elementary cognitive abilities to artificial general intelligence. The literatures on human development and learning provide an essential guide to what is possible for a learner to achieve and what kinds of interaction with the world can support the acquisition of intelligence. AI needs computational neuroscience for algorithmic inspiration. Neural network models are an example of a brain-inspired technology that is unrivalled in several domains of AI. Taking further inspiration from the neurobiological dynamical components (for example, spiking neurons, dendritic dynamics, the canonical cortical microcircuit, oscillations, neuromodulatory processes) and the global functional layout of the human brain (for example, subsystems specialized for distinct functions, including sensory modalities, memory, planning and motor control) might lead to further AI breakthroughs. Machine learning draws from separate traditions in statistics and computer science, which have optimized statistical and computational efficiency, respectively. The integration of computational and statistical efficiency is an essential challenge in the age of big data. The brain appears to combine computational and statistical efficiency, and understanding its algorithm might boost machine learning.

control enables the agent to exploit such dependencies even before understanding the causal mechanism supporting a successful path of action.

The brain is capable of each of these three modes of control (model-free, model-based, episodic)<sup>89</sup> and appears to combine their advantages using an algorithm that has yet to be discovered. AI and computational neuroscience share the goal of discovering this algorithm<sup>41,90,93–95</sup>, although they approach this goal from different angles. This is an example of how a cognitive challenge can motivate the development of formal models and drive progress in AI and neuroscience.

A third, and critically important, class of cognitive model is that of Bayesian models (Box 3)<sup>21,96–98</sup>. Bayesian inference provides an essential normative perspective on cognition. It tells us what a brain should in fact compute for an animal to behave optimally. Perceptual inference, for example, should consider the current sensory data in the context of prior beliefs. Bayesian inference simply refers to combining the data with prior beliefs according to the rules of probability.

Bayesian models have contributed to our understanding of basic sensory and motor processes<sup>22–24</sup>. They have also provided insights into higher cognitive processes of judgment and decision making, explaining classical cognitive biases<sup>99</sup> as the product of prior assumptions, which may be incorrect in the experimental task but correct and helpful in the real world.

With Bayesian nonparametric models, cognitive science has begun to explain more complex cognitive abilities. Consider the human ability to induce a new object category from a single example. Such inductive inference requires prior knowledge of a kind not captured by current feedforward neural network models<sup>100</sup>. To induce a category, we rely on an understanding of the object, of the interactions among its parts, of how they give rise to its function. In the Bayesian cognitive perspective, the human mind, from infancy, builds mental models of the world<sup>2</sup>. These models may not only be generative

models in the probabilistic sense, but may be causal and compositional, supporting mental simulations of processes in the world using elements that can be re-composed to generalize to novel and hypothetical scenarios<sup>2,98,101</sup>. This modeling approach has been applied to our reasoning about the physical<sup>101–103</sup> and even the social<sup>104</sup> world.

Generative models are an essential ingredient of general intelligence. An agent attempting to learn a generative model strives to understand all relationships among its experiences. It does not require external supervision or reinforcement to learn, but can mine all its experiences for insights on its environment and itself. In particular, causal models of processes in the world (how objects cause images, how the present causes the future) can give an agent a deeper understanding and thus a better basis for inferences and actions.

The representation of probability distributions in neuronal populations has been explored theoretically and experimentally<sup>105,106</sup>. However, relating Bayesian inference and learning, especially structure learning in nonparametric models, to its implementation in the brain remains challenging<sup>107</sup>. As theories of brain computation, approximate inference algorithms such as sampling may explain cortical feedback signals and activity correlations<sup>97,108–110</sup>. Moreover, the corners cut by the brain for computational efficiency, the approximations, may explain human deviations from statistical optimality. In particular, cognitive experiments have revealed signatures of sampling<sup>111</sup> and amortized inference<sup>112</sup> in human behavior.

Cognitive models, including the three classes highlighted here, decompose cognition into meaningful functional components. By declaring their models independent of the implementation in the brain, cognitive scientists are able to address high-level cognitive processes<sup>21,97,98</sup> that are beyond the reach of current neural networks. Cognitive models are essential for cognitive computational neuroscience because they enable us to see the whole as we attempt to understand the roles of the parts.

**Box 5 | Shareable tasks, data, models and tests: a new culture of multidisciplinary collaboration**

Neurobiologically plausible models that explain cognition will have substantial parametric complexity. Building and evaluating such models will require machine learning and big brain and behavioral datasets. Traditionally, each lab has developed its own tasks, datasets, models and tests with a focus on the goals of its own discipline. To scale these efforts up to meet the challenge, we will need to develop tasks, data, models and tests that are relevant across the three disciplines and shared among labs (see figure). A new culture of collaboration will assemble big data and big models by combining components from different labs. To meet the conjoined criteria for success of cognitive science, computational neuroscience and artificial intelligence, the best division of labor might cut across the traditional disciplines.

**Tasks.** By designing experimental tasks, we carve up cognition into components that can be quantitatively investigated. A task is a controlled environment for behavior. It defines the dynamics of a task ‘world’ that provides sensory input (for example, visual stimuli) and captures motor output (for example, button press, joystick control or higher-dimensional limb or whole-body control). Tasks drive the acquisition of brain and behavioral data and the development of AI models, providing well-defined challenges and quantitative performance benchmarks for comparing models. The ImageNet tasks<sup>148</sup>, for example, have driven substantial progress in computer vision. Tasks should be designed and implemented such that they can readily be used in all three disciplines to drive data acquisition and model development (related developments include OpenAI’s Gym, <https://gym.openai.com/>; Universe, <https://universe.openai.com/>; and DeepMind’s Lab<sup>149</sup>). The spectrum of useful tasks includes classical psychophysical tasks employing simple stimuli and responses as well as interactions in virtual realities. As we engage all aspects of the human mind, our tasks will need to simulate natural environments and will come to resemble computer games. This may bring the added benefit of mass participation and big behavioral data, especially when tasks are performed via the Internet<sup>150</sup>.

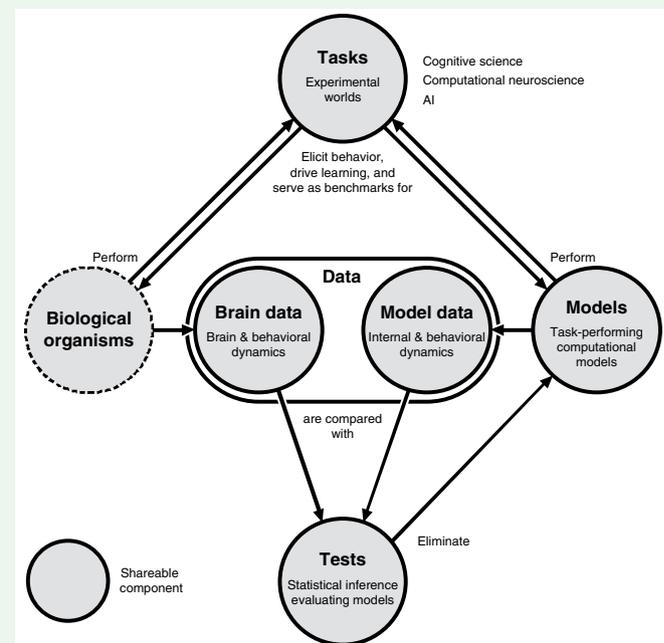
**Data.** Behavioral data acquired during task performance provides overall performance estimates and detailed signatures of success and failure, of reaction times and movement trajectories. Brain-activity measurements characterize the dynamic computations underlying task performance. Anatomical data can characterize the structure and connectivity of the brain at multiple scales. Structural brain data, functional brain data and behavioral data will all be essential for constraining computational models.

**Models.** Task-performing computational models can take sensory inputs and produce motor outputs so as to perform experimental tasks. AI-scale neurobiologically plausible models can be shared openly and tested in terms of their task performance and in terms of their ability to explain a variety of brain and behavioral datasets, including new datasets acquired after definition of the model. Initially, many models will be specific to small subsets of tasks. Ultimately, models must generalize across tasks.

**Tests.** To assess the extent to which a model can explain brain information processing during a particular task, we need tests that compare models and brains on the basis of brain and

behavioral data. Every brain is idiosyncratic in its structure and function. Moreover, for a given brain, every act of perception, cognition and action is unique in time and cannot be repeated precisely because it permanently changes the brain in question. These complications make it challenging to compare brains and models. We must define the summary statistics of interest and the correspondence mapping between model and brain in space and time at some level of abstraction. Developing appropriate tests for adjudicating among models and determining how close we are to understanding the brain is not merely a technical challenge of statistical inference. It is a conceptual challenge fundamental to theoretical neuroscience.

The interaction among labs and disciplines can benefit from adversarial cooperation<sup>134</sup>. Cognitive researchers who feel that current computational models fall short of explaining an important aspect of cognition are challenged to design shareable tasks and tests that quantify these shortcomings and to provide human behavioral data to set the bar for AI models. Neuroscientists who feel that current models do not explain brain information processing are challenged to share brain-activity data acquired during task performance and tests comparing activity patterns between brains and models to quantify the shortcomings of the models. Although we will have a plurality of definitions of success, translating these into quantitative measures of the quality of a model is essential and could drive progress in cognitive computational neuroscience, as well as engineering.



**Interactions among shareable components.** Tasks, data, models and tests are components (gray nodes) that lend themselves to sharing among labs and across disciplines, to enable collaborative construction and testing of big models driven by big brain and behavioral datasets assembled across labs.

**Looking ahead**

**Bottom up and top down.** The brain seamlessly merges bottom-up discriminative and top-down generative computations in per-

ceptual inference, and model-free and model-based control. Brain science likewise needs to integrate its levels of description and to progress both bottom-up and top-down, so as to explain task

performance on the basis of neuronal dynamics and provide a mechanistic account of how the brain gives rise to the mind.

Bottom-up visions, proceeding from detailed measurements toward an understanding of brain computation, have been prominent and have driven the most important recent funding initiatives. The European Human Brain Project and the US BRAIN Initiative are both motivated by bottom-up visions, in which an understanding of brain computation is achieved by measuring and modeling brain dynamics with a focus on the circuit level. The BRAIN Initiative seeks to advance technologies for measuring and manipulating neuronal activity. The Human Brain Project attempts to synthesize neuroscience data in biologically detailed dynamic models. Both initiatives proceed primarily from experiment toward theory and from the cellular level of description to larger-scale phenomena.

Measuring large numbers of neurons simultaneously and modeling their interactions at the circuit level will be essential. The bottom-up vision is grounded in the history of science. Microscopes and telescopes, for example, have brought scientific breakthroughs. However, it is always in the context of prior theory (generative models of the observed processes) that better observations advance our understanding. In astronomy, for example, the theory of Copernicus guided Galileo in interpreting his telescopic observations.

Understanding the brain requires that we develop theory and experiment in tandem and complement the bottom-up, data-driven approach by a top-down, theory-driven approach that starts with behavioral functions to be explained<sup>113,114</sup>. Unprecedentedly rich measurements and manipulations of brain activity will drive theoretical insight when they are used to adjudicate between brain-computational models that pass the first test of being able to perform a function that contributes to the behavioral fitness of the organism. The top-down approach, therefore, is an essential complement to the bottom-up approach toward understanding the brain (Fig. 3).

**Integrating Marr's levels.** Marr (1982) offered a distinction of three levels of analysis: (i) computational theory, (ii) representation and algorithm, and (iii) neurobiological implementation<sup>115</sup>. Cognitive science starts from computational theory, decomposing cognition into components and developing representations and algorithms from the top down. Computational neuroscience proceeds from the bottom up, composing neuronal building blocks into representations and algorithms thought to be useful components in the context of the brain's overall function. AI builds representations and algorithms that combine simple components to implement complex feats of intelligence. All three disciplines thus converge on the algorithms and representations of the brain and mind, contributing complementary constraints<sup>116</sup>.

Marr's levels provide a useful guide to the challenge of understanding the brain. However, they should not be taken to suggest that cognitive science need not consider the brain or that computational neuroscience need not consider cognition (Box 4). Marr was inspired by computers, which are designed by human engineers to precisely conform to high-level algorithmic descriptions. This enables the engineers to abstract from the circuits when designing the algorithms. Even in computer science, however, certain aspects of the algorithms depend on the hardware, such as its parallel processing capabilities. Brains differ from computers in ways that exacerbate this dependence. Brains are the product of evolution and development, processes that are not constrained to generate systems whose behavior can be perfectly captured at some abstract level of description. It may therefore not be possible to understand cognition without considering its implementation in the brain or, conversely, to make sense of neuronal circuits except in the context of the cognitive functions they support.

For an example of a challenge that transcends the disciplines, consider a child seeing an escalator for the first time. She will rapidly recognize people on steps traveling upward obliquely. She might think of it as a moving staircase and imagine riding on it, being lifted one story without exerting any effort. She might infer its function and form a new concept on the basis of a single experience, before ever learning the word "escalator".

Deep neural network models provide a biologically plausible account of the rapid recognition of the elements of the visual experience (people, steps, oblique upward motion, handrail). They can explain the computationally efficient pattern recognition component<sup>42</sup>. However, they cannot explain yet how the child understands the relationships among the elements, the physical interactions of the objects, the people's goal to go up, and the function of the escalator, or how she can imagine the experience and instantly form a new concept.

Bayesian nonparametric models explain how deep inferences and concept formation from single experiences are even possible. They may explain the brain's stunning statistical efficiency, its ability to infer so much from so little data by building generative models that provide abstract prior knowledge<sup>98</sup>. However, current inference algorithms require large amounts of computation and, as a result, do not yet scale to real-world challenges such as forming the new concept "escalator" from a single visual experience.

On a 20-watt power budget, the brain's algorithms combine statistical and computational efficiency in ways that are beyond current AI of either the Bayesian or the neural network variety. However, recent work in AI and machine learning has begun to explore the intersection between Bayesian inference and neural network models, combining the statistical strengths of the former (uncertainty representation, probabilistic inference, statistical efficiency) with the computational strengths of the latter (representational learning, universal function approximation, computational efficiency)<sup>117-119</sup>.

Integrating all three of Marr's levels will require close collaboration among researchers with a wide variety of expertise. It is difficult for any single lab to excel at neuroscience, cognitive science and AI-scale computational modeling. We therefore need collaborations between labs with complementary expertise. In addition to conventional collaborations, an open science culture, in which components are shared between disciplines, can help us integrate Marr's levels. Shareable components include cognitive tasks, brain and behavioral data, computational models, and tests that evaluate models by comparing them to biological systems (Box 5).

The study of the mind and brain is entering a particularly exciting phase. Recent advances in computer hardware and software enable AI-scale modeling of the mind and brain. If cognitive science, computational neuroscience and AI can come together, we might be able to explain human cognition with neurobiologically plausible computational models.

Received: 7 November 2016; Accepted: 11 July 2018;  
Published online: 20 August 2018

## References

1. Newell, A. You can't play 20 questions with nature and win: projective comments on the papers of this symposium. Technical Report, School of Computer Science, Carnegie Mellon University (1973).
2. 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).
3. Kriegeskorte, N. & Mok, R. M. Building machines that adapt and compute like brains. *Behav. Brain Sci.* **40**, e269 (2017).
4. Simon, H. A. & Newell, A. Human problem solving: the state of the theory in 1970. *Am. Psychol.* **26**, 145-159 (1971).
5. Anderson, J. R. *The Architecture of Cognition* (Harvard Univ. Press, Cambridge, MA, USA, 1983).

6. McClelland, J. L. & Rumelhart, D. E. *Parallel Distributed Processing* (MIT Press, Cambridge, MA, USA, 1987).
7. Gazzaniga, M. S. ed. *The Cognitive Neurosciences* (MIT Press, Cambridge, MA, USA, 2004).
8. Fodor, J. A. Précis of The Modularity of Mind. *Behav. Brain Sci.* **8**, 1 (1985).
9. Chklovskii, D. B. & Koulakov, A. A. Maps in the brain: what can we learn from them? *Annu. Rev. Neurosci.* **27**, 369–392 (2004).
10. Szucs, D. & Ioannidis, J. P. A. Empirical assessment of published effect sizes and power in the recent cognitive neuroscience and psychology literature. *PLoS Biol.* **15**, e2000797 (2017).
11. Kriegeskorte, N., Simmons, W. K., Bellgowan, P. S. F. & Baker, C. I. Circular analysis in systems neuroscience: the dangers of double dipping. *Nat. Neurosci.* **12**, 535–540 (2009).
12. Kanwisher, N., McDermott, J. & Chun, M. M. The fusiform face area: a module in human extrastriate cortex specialized for face perception. *J. Neurosci.* **17**, 4302–4311 (1997).
13. Tsao, D. Y., Freiwald, W. A., Tootell, R. B. & Livingstone, M. S. A cortical region consisting entirely of face-selective cells. *Science* **311**, 670–674 (2006).
14. Freiwald, W. A. & Tsao, D. Y. Functional compartmentalization and viewpoint generalization within the macaque face-processing system. *Science* **330**, 845–851 (2010).
15. Grill-Spector, K., Weiner, K. S., Kay, K. & Gomez, J. The functional neuroanatomy of human face perception. *Annu. Rev. Vis. Sci.* **3**, 167–196 (2017).
16. Yildirim, I. et al. Efficient and robust analysis-by-synthesis in vision: a computational framework, behavioral tests, and modeling neuronal representations. in *Annual Conference of the Cognitive Science Society* (eds. Noelle, D. C. et al.) (Cognitive Science Society, Austin, TX, USA, 2015).
17. Kriegeskorte, N., Formisano, E., Sorger, B. & Goebel, R. Individual faces elicit distinct response patterns in human anterior temporal cortex. *Proc. Natl Acad. Sci. USA* **104**, 20600–20605 (2007).
18. Anzellotti, S., Fairhall, S. L. & Caramazza, A. Decoding representations of face identity that are tolerant to rotation. *Cereb. Cortex* **24**, 1988–1995 (2014).
19. Chang, L. & Tsao, D. Y. The code for facial identity in the primate brain. *Cell* **169**, 1013–1028.e14 (2017).
20. Van Essen, D. C. et al. The Brain Analysis Library of Spatial maps and Atlases (BALSA) database. *Neuroimage* **144**(Pt. B), 270–274 (2017).
21. Griffiths, T. L., Chater, N., Kemp, C., Perfors, A. & Tenenbaum, J. B. Probabilistic models of cognition: exploring representations and inductive biases. *Trends Cogn. Sci.* **14**, 357–364 (2010).
22. Ernst, M. O. & Banks, M. S. Humans integrate visual and haptic information in a statistically optimal fashion. *Nature* **415**, 429–433 (2002).
23. Weiss, Y., Simoncelli, E. P. & Adelson, E. H. Motion illusions as optimal percepts. *Nat. Neurosci.* **5**, 598–604 (2002).
24. Körding, K. P. & Wolpert, D. M. Bayesian integration in sensorimotor learning. *Nature* **427**, 244–247 (2004).
25. MacKay, D. J. C. *Information Theory, Inference, and Learning Algorithms*. (Cambridge Univ. Press, Cambridge, 2003)
26. Murphy, K. P. *Machine Learning: A Probabilistic Perspective* (MIT Press, Cambridge, MA, USA, 2012).
27. Dayan, P. & Abbott, L. F. *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems* (MIT Press, Cambridge, MA, USA, 2001).
28. Abbott, L. F. Theoretical neuroscience rising. *Neuron* **60**, 489–495 (2008).
29. Olshausen, B. A. & Field, D. J. Sparse coding of sensory inputs. *Curr. Opin. Neurobiol.* **14**, 481–487 (2004).
30. Simoncelli, E. P. & Olshausen, B. A. Natural image statistics and neural representation. *Annu. Rev. Neurosci.* **24**, 1193–1216 (2001).
31. Carandini, M. & Heeger, D. J. Normalization as a canonical neural computation. *Nat. Rev. Neurosci.* **13**, 51–62 (2011).
32. Chaudhuri, R. & Fiete, I. Computational principles of memory. *Nat. Neurosci.* **19**, 394–403 (2016).
33. Shadlen, M. N. & Kiani, R. Decision making as a window on cognition. *Neuron* **80**, 791–806 (2013).
34. Newsome, W. T., Britten, K. H. & Movshon, J. A. Neuronal correlates of a perceptual decision. *Nature* **341**, 52–54 (1989).
35. Wang, X.-J. Decision making in recurrent neuronal circuits. *Neuron* **60**, 215–234 (2008).
36. Diedrichsen, J., Shadmehr, R. & Ivry, R. B. The coordination of movement: optimal feedback control and beyond. *Trends Cogn. Sci.* **14**, 31–39 (2010).
37. Kriegeskorte, N. Deep neural networks: a new framework for modeling biological vision and brain information processing. *Annu. Rev. Vis. Sci.* **1**, 417–446 (2015).
38. Yamins, D. L. K. & DiCarlo, J. J. Using goal-driven deep learning models to understand sensory cortex. *Nat. Neurosci.* **19**, 356–365 (2016).
39. Krizhevsky, A., Sutskever, I. & Hinton, G. E. ImageNet classification with deep convolutional neural networks. in *Advances in Neural Information Processing Systems 25* 1097–1105 (Curran Associates, Red Hook, NY, USA, 2012).
40. Silver, D. et al. Mastering the game of Go with deep neural networks and tree search. *Nature* **529**, 484–489 (2016).
41. Mnih, V. et al. Human-level control through deep reinforcement learning. *Nature* **518**, 529–533 (2015).
42. LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
43. Cohen, J. D. et al. Computational approaches to fMRI analysis. *Nat. Neurosci.* **20**, 304–313 (2017).
44. Forstmann, B. U., Wagenmakers, E.-J., Eichele, T., Brown, S. & Serences, J. T. Reciprocal relations between cognitive neuroscience and formal cognitive models: opposites attract? *Trends Cogn. Sci.* **15**, 272–279 (2011).
45. Deco, G., Tononi, G., Boly, M. & Kringelbach, M. L. Rethinking segregation and integration: contributions of whole-brain modelling. *Nat. Rev. Neurosci.* **16**, 430–439 (2015).
46. Biswal, B., Yetkin, F. Z., Haughton, V. M. & Hyde, J. S. Functional connectivity in the motor cortex of resting human brain using echo-planar MRI. *Magn. Reson. Med.* **34**, 537–541 (1995).
47. Hyvarinen, A., Karhunen, J. & Oja, E. *Independent Component Analysis* (Wiley, Hoboken, NJ, USA, 2001).
48. Bullmore, E. T. & Bassett, D. S. Brain graphs: graphical models of the human brain connectome. *Annu. Rev. Clin. Psychol.* **7**, 113–140 (2011).
49. Deco, G., Jirsa, V. K. & McIntosh, A. R. Emerging concepts for the dynamical organization of resting-state activity in the brain. *Nat. Rev. Neurosci.* **12**, 43–56 (2011).
50. Friston, K. Dynamic causal modeling and Granger causality. Comments on: the identification of interacting networks in the brain using fMRI: model selection, causality and deconvolution. *Neuroimage* **58**, 303–305 (2011). author reply 310–311.
51. Dennett, D. C. *The Intentional Stance* (MIT Press, Cambridge, MA, USA, 1987).
52. Diedrichsen, J. & Kriegeskorte, N. Representational models: a common framework for understanding encoding, pattern-component, and representational-similarity analysis. *PLoS Comput. Biol.* **13**, e1005508 (2017).
53. Afraz, S.-R., Kiani, R. & Esteky, H. Microstimulation of inferotemporal cortex influences face categorization. *Nature* **442**, 692–695 (2006).
54. Parvizi, J. et al. Electrical stimulation of human fusiform face-selective regions distorts face perception. *J. Neurosci.* **32**, 14915–14920 (2012).
55. Norman, K. A., Polyn, S. M., Detre, G. J. & Haxby, J. V. Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends Cogn. Sci.* **10**, 424–430 (2006).
56. Tong, F. & Pratte, M. S. Decoding patterns of human brain activity. *Annu. Rev. Psychol.* **63**, 483–509 (2012).
57. Kriegeskorte, N. & Kievit, R. A. Representational geometry: integrating cognition, computation, and the brain. *Trends Cogn. Sci.* **17**, 401–412 (2013).
58. Haxby, J. V., Connolly, A. C. & Guntupalli, J. S. Decoding neural representational spaces using multivariate pattern analysis. *Annu. Rev. Neurosci.* **37**, 435–456 (2014).
59. Haynes, J.-D. A primer on pattern-based approaches to fMRI: principles, pitfalls, and perspectives. *Neuron* **87**, 257–270 (2015).
60. Jin, X. & Costa, R. M. Shaping action sequences in basal ganglia circuits. *Curr. Opin. Neurobiol.* **33**, 188–196 (2015).
61. DiCarlo, J. J. & Cox, D. D. Untangling invariant object recognition. *Trends Cogn. Sci.* **11**, 333–341 (2007).
62. Naselaris, T. & Kay, K. N. Resolving ambiguities of MVPA using explicit models of representation. *Trends Cogn. Sci.* **19**, 551–554 (2015).
63. Mitchell, T. M. et al. Predicting human brain activity associated with the meanings of nouns. *Science* **320**, 1191–1195 (2008).
64. Kay, K. N., Naselaris, T., Prenger, R. J. & Gallant, J. L. Identifying natural images from human brain activity. *Nature* **452**, 352–355 (2008).
65. Dumoulin, S. O. & Wandell, B. A. Population receptive field estimates in human visual cortex. *Neuroimage* **39**, 647–660 (2008).
66. Diedrichsen, J., Ridgway, G. R., Friston, K. J. & Wiestler, T. Comparing the similarity and spatial structure of neural representations: a pattern-component model. *Neuroimage* **55**, 1665–1678 (2011).
67. Kriegeskorte, N., Mur, M. & Bandettini, P. Representational similarity analysis - connecting the branches of systems neuroscience. *Front. Syst. Neurosci.* **2**, 4 (2008).
68. Nili, H. et al. A toolbox for representational similarity analysis. *PLoS Comput. Biol.* **10**, e1003553 (2014).
69. Devereux, B. J., Clarke, A., Marouchos, A. & Tyler, L. K. Representational similarity analysis reveals commonalities and differences in the semantic processing of words and objects. *J. Neurosci.* **33**, 18906–18916 (2013).

70. Huth, A. G., de Heer, W. A., Griffiths, T. L., Theunissen, F. E. & Gallant, J. L. Natural speech reveals the semantic maps that tile human cerebral cortex. *Nature* **532**, 453–458 (2016).
71. Markram, H. The Blue Brain Project. *Nat. Rev. Neurosci.* **7**, 153–160 (2006).
72. Eliasmith, C. & Trujillo, O. The use and abuse of large-scale brain models. *Curr. Opin. Neurobiol.* **25**, 1–6 (2014).
73. Eliasmith, C. et al. A large-scale model of the functioning brain. *Science* **338**, 1202–1205 (2012).
74. Hassabis, D., Kumaran, D., Summerfield, C. & Botvinick, M. Neuroscience-inspired artificial intelligence. *Neuron* **95**, 245–258 (2017).
75. Rumelhart, D. E., Hinton, G. E. & Williams, R. J. Learning representations by back-propagating errors. *Nature* **323**, 533–536 (1986).
76. Goodfellow, I., Bengio, Y. & Courville, A. *Deep Learning* (MIT Press, Cambridge, MA, USA, 2016).
77. Yamins, D. L. K. et al. Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proc. Natl Acad. Sci. USA* **111**, 8619–8624 (2014).
78. Khaligh-Razavi, S.-M. & Kriegeskorte, N. Deep supervised, but not unsupervised, models may explain IT cortical representation. *PLoS Comput. Biol.* **10**, e1003915 (2014).
79. Cadieu, C. F. et al. Deep neural networks rival the representation of primate IT cortex for core visual object recognition. *PLoS Comput. Biol.* **10**, e1003963 (2014).
80. Güçlü, U. & van Gerven, M. A. J. Deep neural networks reveal a gradient in the complexity of neural representations across the ventral stream. *J. Neurosci.* **35**, 10005–10014 (2015).
81. Eickenberg, M., Gramfort, A., Varoquaux, G. & Thirion, B. Seeing it all: convolutional network layers map the function of the human visual system. *Neuroimage* **152**, 184–194 (2017).
82. Cichy, R. M., Khosla, A., Pantazis, D., Torralba, A. & Oliva, A. Comparison of deep neural networks to spatio-temporal cortical dynamics of human visual object recognition reveals hierarchical correspondence. *Sci. Rep.* **6**, 27755 (2016).
83. Hong, H., Yamins, D. L. K., Majaj, N. J. & DiCarlo, J. J. Explicit information for category-orthogonal object properties increases along the ventral stream. *Nat. Neurosci.* **19**, 613–622 (2016).
84. Kubilius, J., Bracci, S. & Op de Beeck, H. P. Deep neural networks as a computational model for human shape sensitivity. *PLoS Comput. Biol.* **12**, e1004896 (2016).
85. Jozwik, K. M., Kriegeskorte, N., Storrs, K. R. & Mur, M. Deep convolutional neural networks outperform feature-based but not categorical models in explaining object similarity judgments. *Front. Psychol.* **8**, 1726 (2017).
86. Moore, C. & Mertens, S. *The Nature of Computation*. (Oxford Univ. Press, Oxford, 2011).
87. Borst, J., Taatgen & Anderson, J. Using the ACT-R cognitive architecture in combination with fMRI data. in *An Introduction to Model-Based Cognitive Neuroscience* (eds. Forstmann, B. U. & Wagenmakers, E.-J.) (Springer, New York, 2014).
88. Sutton, R. & Barto, A. *Reinforcement Learning: An Introduction* Vol. 1 (MIT Press, Cambridge, MA, USA, 1998).
89. O’Doherty, J. P., Cockburn, J. & Pauli, W. M. Learning, reward, and decision making. *Annu. Rev. Psychol.* **68**, 73–100 (2017).
90. Daw, N. D. & Dayan, P. The algorithmic anatomy of model-based evaluation. *Phil. Trans. R. Soc. Lond. B* **369**, 20130478 (2014).
91. Lengyel, M. & Dayan, P. Hippocampal contributions to control: the third way in *Advances in Neural Information Processing Systems 20* 889–896 (MIT Press, Cambridge, MA, USA, 2008).
92. Gershman, S. J. & Daw, N. D. Reinforcement learning and episodic memory in humans and animals: an integrative framework. *Annu. Rev. Psychol.* **68**, 101–128 (2017).
93. Schultz, W., Dayan, P. & Montague, P. R. A neural substrate of prediction and reward. *Science* **275**, 1593–1599 (1997).
94. Sutton, R. Integrated architectures for learning, planning, and reacting based on approximating dynamic programming. in *Proceedings of the Seventh International Conference on Machine Learning* 216–224 (Morgan Kaufmann, San Francisco, 1990).
95. Daw, N. D., Niv, Y. & Dayan, P. Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat. Neurosci.* **8**, 1704–1711 (2005).
96. Ma, W. J. Organizing probabilistic models of perception. *Trends Cogn. Sci.* **16**, 511–518 (2012).
97. Fiser, J., Berkes, P., Orbán, G. & Lengyel, M. Statistically optimal perception and learning: from behavior to neural representations. *Trends Cogn. Sci.* **14**, 119–130 (2010).
98. Tenenbaum, J. B., Kemp, C., Griffiths, T. L. & Goodman, N. D. How to grow a mind: statistics, structure, and abstraction. *Science* **331**, 1279–1285 (2011).
99. Tversky, A. & Kahneman, D. Judgment under uncertainty: heuristics and biases. in *Utility, Probability, and Human Decision Making* (eds. Wendt, D. & Vlek, C.) 141–162, [https://doi.org/10.1007/978-94-010-1834-0\\_8](https://doi.org/10.1007/978-94-010-1834-0_8) (Springer Netherlands, Dordrecht, the Netherlands, 1975).
100. Lake, B. M., Salakhutdinov, R. & Tenenbaum, J. B. Human-level concept learning through probabilistic program induction. *Science* **350**, 1332–1338 (2015).
101. Ullman, T. D., Spelke, E., Battaglia, P. & Tenenbaum, J. B. Mind games: game engines as an architecture for intuitive physics. *Trends Cogn. Sci.* **21**, 649–665 (2017).
102. Battaglia, P. W., Hamrick, J. B. & Tenenbaum, J. B. Simulation as an engine of physical scene understanding. *Proc. Natl Acad. Sci. USA* **110**, 18327–18332 (2013).
103. Kubricht, J. R., Holyoak, K. J. & Lu, H. Intuitive physics: current research and controversies. *Trends Cogn. Sci.* **21**, 749–759 (2017).
104. Pantelis, P. C. et al. Inferring the intentional states of autonomous virtual agents. *Cognition* **130**, 360–379 (2014).
105. Pouget, A., Beck, J. M., Ma, W. J. & Latham, P. E. Probabilistic brains: knowns and unknowns. *Nat. Neurosci.* **16**, 1170–1178 (2013).
106. Orhan, A. E. & Ma, W. J. Efficient probabilistic inference in generic neural networks trained with non-probabilistic feedback. *Nat. Commun.* **8**, 138 (2017).
107. Tervo, D. G. R., Tenenbaum, J. B. & Gershman, S. J. Toward the neural implementation of structure learning. *Curr. Opin. Neurobiol.* **37**, 99–105 (2016).
108. Buesing, L., Bill, J., Nessler, B. & Maass, W. Neural dynamics as sampling: a model for stochastic computation in recurrent networks of spiking neurons. *PLoS Comput. Biol.* **7**, e1002211 (2011).
109. Haefner, R. M., Berkes, P. & Fiser, J. Perceptual decision-making as probabilistic inference by neural sampling. *Neuron* **90**, 649–660 (2016).
110. Aitchison, L. & Lengyel, M. The Hamiltonian brain: efficient probabilistic inference with excitatory-inhibitory neural circuit dynamics. *PLoS Comput. Biol.* **12**, e1005186 (2016).
111. Sanborn, A. N. & Chater, N. Bayesian brains without probabilities. *Trends Cogn. Sci.* **20**, 883–893 (2016).
112. Dasgupta, I., Schulz, E., Goodman, N. & Gershman, S. Amortized hypothesis generation. Preprint at *bioRxiv* <https://doi.org/10.1101/137190> (2017).
113. Krakauer, J. W., Ghazanfar, A. A., Gomez-Marín, A., MacIver, M. A. & Poeppel, D. Neuroscience needs behavior: correcting a reductionist bias. *Neuron* **93**, 480–490 (2017).
114. Gomez-Marín, A., Paton, J. J., Kampff, A. R., Costa, R. M. & Mainen, Z. F. Big behavioral data: psychology, ethology and the foundations of neuroscience. *Nat. Neurosci.* **17**, 1455–1462 (2014).
115. Marr, D. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information* (MIT Press, Cambridge, MA, USA, 2010).
116. Love, B. C. The algorithmic level is the bridge between computation and brain. *Top. Cogn. Sci.* **7**, 230–242 (2015).
117. Gal, Y. & Ghahramani, Z. Dropout as a Bayesian approximation: representing model uncertainty in deep learning. Preprint at <https://arxiv.org/abs/1506.02142> (2016).
118. Rezende, D., Mohamed, S., Danihelka, I., Gregor, K. & Wierstra, D. One-shot generalization in deep generative models. *Proc. Int. Conf. Mach. Learn. Appl.* **48**, 1521–1529 (2016).
119. Kingma, D. & Welling, M. Auto-encoding variational Bayes. Preprint at <https://arxiv.org/abs/1312.6114> (2013).
120. Naselaris, T. et al. Cognitive Computational Neuroscience: a new conference for an emerging discipline. *Trends Cogn. Sci.* **22**, 365–367 (2018).
121. Ahrens, M. B. et al. Brain-wide neuronal dynamics during motor adaptation in zebrafish. *Nature* **485**, 471–477 (2012).
122. Kietzmann, T., McClure, P. & Kriegeskorte, N. Deep neural networks in computational neuroscience. Preprint at *bioRxiv* <https://doi.org/10.1101/133504> (2017).
123. Hornik, K. Approximation capabilities of multilayer feedforward networks. *Neural Netw.* **4**, 251–257 (1991).
124. Wyatte, D., Curran, T. & O’Reilly, R. The limits of feedforward vision: recurrent processing promotes robust object recognition when objects are degraded. *J. Cogn. Neurosci.* **24**, 2248–2261 (2012).
125. Spoerer, C. J., McClure, P. & Kriegeskorte, N. Recurrent convolutional neural networks: a better model of biological object recognition. *Front. Psychol.* **8**, 1551 (2017).
126. Hunt, L. T. & Hayden, B. Y. A distributed, hierarchical and recurrent framework for reward-based choice. *Nat. Rev. Neurosci.* **18**, 172–182 (2017).
127. Schäfer, A. M. & Zimmermann, H. G. Recurrent neural networks are universal approximators. *Int. J. Neural Syst.* **17**, 253–263 (2007).
128. O’Reilly, R. C., Hazy, T. E., Mollick, J., Mackie, P. & Herd, S. Goal-driven cognition in the brain: a computational framework. Preprint at <http://arxiv.org/abs/1404.7591> (2014).
129. 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).

130. Schiess, M., Urbanczik, R. & Senn, W. Somato-dendritic synaptic plasticity and error-backpropagation in active dendrites. *PLoS Comput. Biol.* **12**, e1004638 (2016).
131. Marblestone, A. H., Wayne, G. & Kording, K. P. Towards an integration of deep learning and neuroscience. *Front. Comput. Neurosci.* **10**, 94 (2016).
132. Shadlen, M. N. & Shohamy, D. Decision making and sequential sampling from memory. *Neuron* **90**, 927–939 (2016).
133. Roelfsema, P. R. & van Ooyen, A. Attention-gated reinforcement learning of internal representations for classification. *Neural Comput.* **17**, 2176–2214 (2005).
134. Goodfellow, I. et al. Generative adversarial nets. Preprint at <https://arxiv.org/abs/1406.2661> (2014).
135. Kandel, E. R., Schwartz, J. H., Jessell, T. M., Siegelbaum, S. A. & Hudspeth, A. J. *Principles of Neural Science* (McGraw-Hill Professional, New York, 2013).
136. Bastos, A. M. et al. Canonical microcircuits for predictive coding. *Neuron* **76**, 695–711 (2012).
137. Larkum, M. A cellular mechanism for cortical associations: an organizing principle for the cerebral cortex. *Trends Neurosci.* **36**, 141–151 (2013).
138. Fries, P. A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends Cogn. Sci.* **9**, 474–480 (2005).
139. Kumaran, D., Hassabis, D. & McClelland, J. L. What learning systems do intelligent agents need? complementary learning systems theory updated. *Trends Cogn. Sci.* **20**, 512–534 (2016).
140. Yuille, A. & Kersten, D. Vision as Bayesian inference: analysis by synthesis? *Trends Cogn. Sci.* **10**, 301–308 (2006).
141. Helmholtz, H. *Handbuch der physiologischen Optik* (Dover, New York, 1860).
142. Gershman, S. J., Horvitz, E. J. & Tenenbaum, J. B. Computational rationality: a converging paradigm for intelligence in brains, minds, and machines. *Science* **349**, 273–278 (2015).
143. Simon, H. A. Bounded rationality. in *Utility and Probability* (eds. Eatwell, J., Milgate, M. & Newman, P.) 15–18, [https://doi.org/10.1007/978-1-349-20568-4\\_5](https://doi.org/10.1007/978-1-349-20568-4_5) (Palgrave Macmillan, London, 1990).
144. Griffiths, T. L., Lieder, F. & Goodman, N. D. Rational use of cognitive resources: levels of analysis between the computational and the algorithmic. *Top. Cogn. Sci.* **7**, 217–229 (2015).
145. Srikumar, V., Kundu, G. & Roth, D. On amortizing inference cost for structured prediction *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning* 1114–1124 (Association for Computational Linguistics, Stroudsburg, PA, USA, 2012).
146. Bengio, Y., Scellier, B., Bilaniuk, O., Sacramento, J. & Senn, W. Feedforward initialization for fast inference of deep generative networks is biologically plausible. Preprint at <https://arxiv.org/abs/1606.01651> (2016).
147. Ghahramani, Z. Bayesian non-parametrics and the probabilistic approach to modelling. *Philos. Trans. A Math. Phys. Eng. Sci.* **371**, 20110553 (2012).
148. Deng, J. et al. ImageNet: a large-scale hierarchical image database. in *2009 IEEE Conference on Computer Vision and Pattern Recognition* 248–255, <https://doi.org/10.1109/CVPR.2009.5206848> (IEEE, Piscataway, NJ, USA, 2009).
149. Beattie, C. et al. DeepMind Lab. Preprint at <https://arxiv.org/abs/1612.03801> (2016).
150. Griffiths, T. L. Manifesto for a new (computational) cognitive revolution. *Cognition* **135**, 21–23 (2015).

## Acknowledgements

This paper benefited from discussions in the context of the new conference Cognitive Computational Neuroscience, which had its inaugural meeting in New York City in September 2017<sup>20</sup>. We are grateful in particular to T. Naselaris, K. Kay, K. Kording, D. Shohamy, R. Poldrack, J. Diedrichsen, M. Bethge, R. Mok, T. Kietzmann, K. Storrs, M. Mur, T. Golan, M. Lengyel, M. Shadlen, D. Wolpert, A. Oliva, D. Yamins, J. Cohen, J. DiCarlo, T. Konkle, J. McDermott, N. Kanwisher, S. Gershman and J. Tenenbaum for inspiring discussions.

## Competing interests

The authors declare no competing interests.

## Additional information

Reprints and permissions information is available at [www.nature.com/reprints](http://www.nature.com/reprints).

Correspondence should be addressed to N.K.

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