

Expanding perspectives on cognition in humans, animals, and machines

Alex Gomez-Marin¹ and Zachary F Mainen



Over the past decade neuroscience has been attacking the problem of cognition with increasing vigor. Yet, what exactly *is* cognition, beyond a general signifier of anything seemingly complex the brain does? Here, we briefly review attempts to define, describe, explain, build, enhance and experience cognition. We highlight perspectives including psychology, molecular biology, computation, dynamical systems, machine learning, behavior and phenomenology. This survey of the landscape reveals not a clear target for explanation but a pluralistic and evolving scene with diverse opportunities for grounding future research. We argue that rather than getting to the bottom of it, over the next century, by deconstructing and redefining cognition, neuroscience will and should expand rather than merely reduce our concept of the mind.

Address

Champalimaud Neuroscience Programme, Champalimaud Centre for the Unknown, Lisbon, Portugal

Corresponding authors: Gomez-Marin, Alex (agomezmarin@gmail.com) and Mainen, Zachary F (zmainen@neuro.fchampalimaud.org)

¹ Present address: Instituto de Neurociencias, Consejo Superior de Investigaciones Científicas & Universidad Miguel Hernández, Sant Joan d'Alacant, Spain.

Current Opinion in Neurobiology 2016, 37:85–91

This review comes from a themed issue on **Neurobiology of behavior**

Edited by **Alla Karpova** and **Roosbeh Kiani**

<http://dx.doi.org/10.1016/j.conb.2016.01.011>

0959-4388/© 2016 Elsevier Ltd. All rights reserved.

Introduction

Neuroscience aims to explain how the workings of the mind and behavior arise from the biology of interacting cells and molecules. This is arguably the biggest scientific challenge of our era. One of the main reasons that this problem is so hard is obvious: it is the vast complexity of the brain. But a somewhat less obvious part of the problem is agreeing on the question: what exactly are we trying to explain? The term ‘cognition’ is used in neuroscience as shorthand to refer to the most complex functions the brain enables [1]. Cognition has traditionally been defined in terms of a suite of ‘cognitive functions’ — attention, working memory, and so on. But are these really the key building blocks for complex behavior? If not, the definition of cognition should be realigned

on a better-grounded empirical or theoretical approach. But where is that to be found? These questions are critical because without a firm grasp of what exactly is cognition, apparent progress in understanding its neural basis may evaporate as it is realized that the wrong object was under scrutiny.

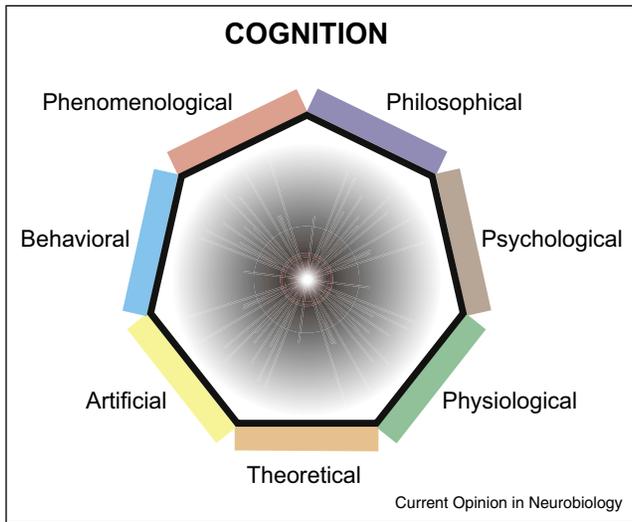
In this review we survey contemporary approaches to brain complexity that might serve as a basis for redefining cognition (see [Figure 1](#)). We begin by very briefly reviewing over 40 years of progress in studying the single neuron correlates of cognition. We emphasize the recent theoretical proposal that the complexity of these cognitive tasks is extremely low, implying that we might indeed be missing something in the definition. From this we launch into a broader discussion of possible theoretical bases for defining cognition. We then turn to consider the more empirical bases found in machine learning, behavioral studies, and neuro-phenomenology.

Neural basis of cognition in monkeys and rodents

The idea of cognition goes back at least to Aristotle (*gignoskein*, to know, the process of acquiring knowledge) but did not really get going in a scientific form until the 19th century. The field of psychology has sought for many years to understand the basis of intelligence — cognition being the name for the higher capacities of the mind. By the time of Ebbinghaus and James in the late 19th century, many forms of cognition were beginning to be subjected to scientific study by psychologists.

Starting in the 1960s neuroscientists began to translate work in cognitive psychology from humans into monkeys. The goal of this work was to reveal direct physical correlates of cognitive functions such as working memory [2]. To do this, they systematically recorded the electrical activity of single neurons in behaving animals. This required them to teach monkeys to perform cognitive tasks. This line of work demonstrated that monkeys not only exhibit behaviors considered to constitute the most basic building blocks of cognition — learning & memory, decision-making — but can perform more complex tasks including selective attention [3], flexible categorization [4], integration of evidence [5], probabilistic reasoning [6], metacognitive judgements [7], numerical reasoning [4], context-dependent decisions [8] and more.

Figure 1



A pluralistic landscape of abstractions to approach the phenomenon of cognition.

Starting in the 1980s, researchers began to show that even rodents can also perform a wide variety of standard cognitive tasks including selective attention [9], flexible sensorimotor mapping [10] and working memory. In the past decade or so, this line of work in freely behaving animals was refined to even more exacting standards, allowing quantitative measurements of single neuron activity in relationship to functions including confidence judgments [11], psychophysical thresholds [12], integration of evidence [13], self-initiated actions [14], cue combination [15], and flexible sensory-motor associations [16]. A burgeoning line of work has taken this approach to even more constrained, head-fixed, mice performing tasks such as memory guided decisions [17], sometimes in a virtual reality [18].

Cognition in just a few dimensions

This work has set the stage for huge progress in ‘opening the black box’, as it allows rodents — which are smaller, cheaper, far easier to experiment with, less ethically demanding — to stand in for monkeys or even humans. Particularly because this allows the use of molecular and genetic manipulations, which have become more and more powerful [19,20], many researchers using macaques have begun to switch to rats or mice or smaller, faster-breeding marmosets, in which genetic manipulations may be performed with greater ease.

Yet the desire to account for behavior on the basis of single spikes in single neurons raises a rather perplexing issue. It seems preposterous to hope to understand the concerted actions of a brain of 100 million neurons by studying a handful of neurons at a time. Yet, somehow this does not seem to be the case. Single neurons or

populations <100 have typically revealed meaningful correlates of target cognitive phenomena. When analyzed using dimensionality reduction techniques, even the most complex tasks mentioned have yielded a dimensionality that is far lower than even the relatively small number of neurons being analyzed [8,21].

What is going on? Ganguli and colleagues have proposed a theory of ‘neural task complexity’ [22**]. They argue that as long as task parameters are limited and neuronal activity patterns vary smoothly across task parameters, that even for what we think of as complex cognitive tasks, the neural task complexity is very small. This work suggests that rather than recording more neurons, in order to make progress we need richer, more complex tasks [22**]. This relates to the important issue of how much of what we measure in behaving animals is due to behavior itself or simply reflects the environment in which the animal is placed in the laboratory. It also raises the more radical possibility that neuroscience may already be outgrowing the traditionally defined set of cognitive functions.

A basis in psychology or in the brain itself?

Nearly 80 years ago, Tolman proclaimed that ‘everything important in psychology (...) can be investigated in essence through the continued experimental and theoretical analysis of the determiners of rat behavior at a choice-point in a maze’ [23]. Our current cognitive taxonomy appears to include a much richer set of behaviors, but much of our basic cognitive toolkit was already old hat to Tolman, having been developed decades before by psychologists such as William James [24]. Granting James’ genius, how much confidence can be taken that the 19th century agenda of cognitive psychology covers the breadth and depth of mental function?

The idea of dissecting cognition into isolatable components, combined with the possibility of studying the neural basis of such a taxonomy of subfunctions (attention, working memory, meta-cognition, etc.) appears to be well-suited to the paradigm set by molecular biology and genetics for the study of biological mechanism. The prominence of neural reductionism, ‘the tendency to seek all causal relations of behavior in brain processes’ [25], has caused some to wonder whether psychological explanation is in danger of extinction [26]. This is a bias that might be exacerbated in the current climate of increased competition for publication and funding. What does it mean to explain cognition [27]? If explaining entails explaining away, then the drive toward reductionistic accounts may ironically turn into a roadblock in our understanding. Indeed, there is reason to think that the reductionist approach to cognition may be premature. Truly novel work does go on at the psychological level, even if it has a harder time finding an audience without a neural substrate. For instance, while Bayesian-inspired

psychological research has made inroads into mechanistic cognitive neuroscience, cognitive psychologists continue to work beyond what is cutting edge in neuroscience on problems such as rapid structural learning and concept learning [28**].

It has been argued that psychology suffers from human-invented categories that do not necessarily conform to nature — ‘we take a man-created word or concept (...) and search for brain mechanisms that may be responsible for the generation of this conceived behavior’ [29]. An alternative is an inside-out or bottom-up approach building from brain mechanisms toward cognitive function. While this strategy offers freedom from potentially awkward hypotheses about the phenomena to be explained, for a system as complex as a brain, a pure bottom up approach appears to be a non-starter. The power of discovery based on sophisticated techniques relying on big data can fail to produce significant results but, in a sense, it can never be wrong. Therefore falsifiable theories and concrete testable hypothesis are needed. Even in an organism for which the complete wiring diagram is available, *Caenorhabditis elegans*, behavioral function cannot be deduced without additional hypotheses. This problem is also illustrated by the heroic but ultimately humbling attempt to model a single cortical column from the bottom-up [30].

Theories and frameworks for cognition

If one is to approach and define cognition at least partly from the top-down, what are the available theoretical frameworks from which this might be done? Here we find several, not necessarily mutually exclusive, alternatives.

Dynamical systems

A long-standing theoretical approach to cognition is based on nonlinear dynamical systems theory, which has been proposed to be a non-computational way to conceive of cognition [31]. Two claimed advantages of conceiving of a cognitive agent as a dynamical system, rather than a digital computer, are the possibility to interact with the world without ‘representing’ it and the possibility to avoid discretizing time [32]. Dynamical systems approaches were the basis for understanding neuronal function at the cellular level [33]. They are also the basis for more ‘biophysical’ network models, which have demonstrated that the dynamics of very simple neuronal circuits can be tuned to perform elementary cognitive functions such as working memory and decision-making [34,35]. Moreover, with appropriate (albeit sometimes implausibly nonlocal) learning rules these networks can be endowed with surprisingly rich dynamics, for example, generating arbitrary dynamical mappings [36,37], providing mechanistic hypothesis for the basis of context-dependent decision-making. Dynamical network models based on engineering rather than learning solutions have begun to be assembled to perform complete tasks [38] but cannot,

say, play a video game. Approaches based on optimization (learning) of relatively simple objective functions have been remarkably successful in capturing the dynamics of complex actions (e.g. standing up) using nonlinear dynamical models of the motor plant and physical environment [39]. Work from this perspective argues that cognition and action are deeply intertwined, an ‘embodied’ perspective reinforced by the loopiness of brain anatomy. Motor-control inspired dynamic systems approaches may provide fertile ground for future neuroscience of cognition.

Reinforcement learning (RL)

Over the past two decades, RL theory has been one of the most impactful theoretical frameworks in neuroscience [40]. With roots in artificial intelligence (AI), to which we return below, the success of RL in neuroscience has been forged by theorists who have taken a very much top-down or normative approach to ask what *should* the brain do. Importantly, RL theorists have also cared deeply about the underlying brain circuitry, particularly at the systems level (area by area), and the behavior, particularly that carried out by the ‘classical’ animal learning theorists. Although the name emphasizes one flavor of learning, RL naturally incorporates decision-making, and strong relationships with microeconomic theory exist [41]. RL deals with optimization based on an externally supplied scalar objective function. Yet it should not be forgotten that learning is a much richer and murkier domain. Beyond RL are forms of learning that work without supervision of any kind, that make use of observation, and that require only a single exposure.

Bayesian decision theory

Decision theory is a branch of mathematical psychology that deals with choosing based on noisy samples. It was inspired by the problem of code-breaking and fueled for many years by the ability to explain human reaction time and performance data across a range of domains [42]. It thus has both a normative basis and empirical support. The widely used diffusion-to-bound model has been linked to single neuron recordings, implemented using the circuit-level dynamic models and derived as a subcase of Bayesian theory. Decision-making is clearly a limited aspect of cognition, but integration is a key process not only in accumulation of evidence, but in working memory, cue combination, and Bayesian inference, etc. Evidence for non-leaky, noiseless integration even in rats [43] and, the ability to account for other aspect of behavioral data, for example, action timing [14], contribute to the attractiveness of these models as a basis for describing cognition.

Quantum cognition is an offshoot of decision theory that borrows the quantum probability formalism based on the Dirac and Von Neumann axioms. Quantum decision models can explain several anomalous effects in sequential

decision-making [44**]. A quantum diffusion-to-bound model was introduced [45], which is a promising start toward other cognitive mechanisms. In the background, as either a beacon or a specter, is whether it is ultimately tied to quantum mechanics or quantum computing, as infamously proposed by Penrose [46].

Information theory

An information processing metaphor, in which the brain is essentially an input/output system has dominated cognitive science since its origins. Information theory quantitatively formalizes the communication of signals using the fundamental concept of entropy. In neuroscience it has been used very effectively to study how neural signals propagate through sensory systems [47] and provides a useful basis for conceptualizing the efficiency of neural representations or codes. Moreover, information theory has been used to define what constitutes a ‘conscious’ process [48,49*] and recently, it has been shown that simple forms of ‘cognitive’ behavior, such as solving a puzzle or cooperating socially, can emerge from principles of entropy maximization, where future entropy acts as a force in the present [50**]. Under the rubric ‘free energy’, and married with Bayesian and RL theory, an information-based approach forms the core of a proposed universal brain model [51].

Theory of computation

Together with information, the notion of computation has long been a dominant abstraction for conceiving the brain’s most complex functions. Perhaps most fundamentally, Kolmogorov complexity provides a definition of randomness based on the shortest computer program that can be run on a Turing machine [52]. Algorithmic complexity may be useful in application to cognition by offering a grounded definition of what is hard or complex. Despite the dominance of the computational metaphor, there has been surprisingly little contact between cognitive neuroscience and most aspects of computer science. Interestingly, it has been suggested that practical advances in the field of computer engineering (e.g. operating systems) are not themselves well-understood by computer scientists [53], which may partly explain why current approaches to thinking about cognition remain relatively uninfluenced by advances in computer engineering.

Machine cognition

The theoretical approaches just outlined are rooted in mathematics and deep concepts — probabilistic inference, learning, entropy, complexity — that are being mined for their potential as an ultimate level of explanation for brain function. These theories have also driven, in parallel, the quest to artificially engineer intelligent machines. In particular, the field of neural networks or machine learning has been closely related to the field of computational neuroscience. What might be the relevance of machine learning,

beyond a basis in some of the same normative theories, to defining cognition?

AI is a practical perspective by which to define what brain functions are truly ‘hard’, and therefore ‘cognitive’, under our current theoretical assumptions. It has long been understood that for ‘symbolic’ AI, problems like chess are easy while problems like vision are impossible. Neural networks based on ‘deep learning’ [66*] appear to have cracked what had been considered difficult problems such as image recognition. As these approaches are pushed, their limits may come into clearer focus before changing again as new algorithms emerge. Already, hybrid approaches combining, for example, neural networks with Turing machines [54], are proliferating.

Neural networks share some core features with brains, including distributed processing over simple units and modifiable connections. But they omit the vast majority of what we know about brains, from spikes to cortical layers. And what makes deep learning successful has not been the incorporation of more exacting neural principles but rather clever algorithms such as the back-propagation learning rule that has little to do with brain architecture [55]. It remains uncertain to what degree machine learning approaches will benefit from incorporating more of the brain’s architecture. Consider how divergent are the architectures of the cephalopod and the mammalian brain. Pessimistically (or not), brains and machines might evolve two different solutions to intelligence built on two different sets of tricks, with neither illuminating much the other.

Cognition informed by behavior

If we get behavior wrong, we will not get cognition right. As we have argued recently [56*] behavior is the foundation of neuroscience, not just a handy tool to investigate the function of the nervous system. When seen as a phenomenon in its own right — not as a mere neural side-effect — the study of behavior becomes invaluable to reveal the nature of cognition.

Bottom up approaches aim to discover, from sophisticated analyses, structure in the copious and high-resolution behavioral data increasingly becoming available. In computational ethology [57], state-of-the-art approaches allow the discovery of stereotyped behaviors directly from the images themselves [58*,59]. Such remarkable data analysis will ultimately need to be incorporated in the context of testable theories. Behavior itself is also a worthy subject of ‘top-down’ theoretical consideration that may inform and perhaps even upend our notion of cognition. Among other potential avenues, one that seems particularly relevant today is to reconsider the idea of negative feedback control, the basis for the ‘cybernetic revolution’ of the 1940s. Maxwell’s seminal 1867 paper on ‘governors’ [67] described autonomous systems that can

maintain certain variables at a reference value despite external disturbances (such as the tray of a waiter as he walks through the restaurant). In fact, negative feedback control is closely related to the concept of homeostasis.

While cybernetic concepts thoroughly transformed engineering practice, they had relatively little impact on the study of behavior. Their potential relevance was, however, eloquently championed by Powers, who proclaimed that ‘behavior is the control of perception’ [60]. Then, to understand behavior is to discover the variables that the organism controls, not how it responds to certain stimuli. The implication of this view is that linear cause-and-effect explanations will fail because they assume that real behavior is open loop, when it is not [61^{*}]. Behavioral approaches that focus on manipulating the input (stimuli, rewards) and measuring the output (responses, actions) tend to confound controlled with controlling variables. The difference is crucial: the former are controlled by the behaving system whereas the latter are controlled by the human studying the behaving system [62]. One might then be studying the properties of the environmental feedback mistaking them for the behavioral function itself. Note, for instance, the distinction between the behavior of a fan and a thermostat: the first controls its output, the second its input. The thermostat might be said to emit heat, but that is not really what it does. What it does (or attempts to do) is to specify the temperature being sensed.

Therefore, rather than concentrating on what an animal is doing, what may be more relevant is what the animal is trying to perceive. This double inversion (from the experimenter’s point of view to the animal’s and from action to perception) has three potentially critical implications for the study of cognition: first, motor output is a side effect of perceptual control and therefore quantitative data collection data is insufficient by itself, second, averaging across individuals may smear out control variables, and third, restrained experimental setups may not let animals control the relevant inputs; freedom in the requisite dimensions is essential. In short, control (circular causality) and subjectivity (animal centrism) may be essential ingredients in behavioral and cognitive neuroscience.

Cognition beyond the brain itself

Finally, we note two important and relatively new approaches that assert that the data relevant to cognition do not lie solely within the realm of neuroscience facts as traditionally construed. First, ‘neurophenomenology’, introduced by Francisco Varela, aims to include introspection as a source of scientific evidence [63]. Neurophenomenology is based on first-person (subjective) moment-to-moment accounts coupled with classical quantification procedures such as neuroimaging. Together, these are proposed to provide mutual constraints to tighten the link

between neurodynamics and conscious experience and, in particular, to offer new insights in the study of variability and spontaneity in cognition [64^{*}]. Second, ‘embodied cognition’, also championed by Varela, centers on the recognition that cognition is embedded in a context beyond the brain itself, including the body as well as broader environmental, biological and cultural systems [65]. Both neurophenomenology and embodied cognition illustrate the nascent fruits of substantive dialogue between philosophy and neuroscience that we hope will continue to instill new perspectives in the years to come.

Conclusions

The neuroscience of cognition is in its infancy. Greatly empowered by new tools, it is still largely hamstrung by the perspectives of the past. In formulating the problems of the future, the field should benefit by paying careful mind not only to neural data but also to the nature of the phenomena to be explained. If the questions we ask are restricted to the same terms they have traditionally been formulated in, potential understanding will be greatly restricted. The various disciplines and frameworks we have summarized here offer rich and widely differing perspectives on how to conceive of cognition. Their mutual incoherence attests to a conceptual vitality that is not merely appropriate but critical to a young field tackling an immense problem. Going forward, it is important for the field to continue to evolve, drawing inspiration from any and all potentially relevant quarters and resisting the tendency to become stuck in a singular, archaic, view of the mind. When we define cognition, we are, after all, helping to define what it is to be human.

Conflict of interests

Nothing declared.

References and recommended reading

Papers of particular interest, published within the period of review, have been highlighted as:

- of special interest
 - of outstanding interest
1. Cromwell HC, Panksepp J: **Rethinking the cognitive revolution from a neural perspective: how overuse/misuse of the term ‘cognition’ and the neglect of affective controls in behavioral neuroscience could be delaying progress in understanding the BrainMind.** *Neurosci Biobehav Rev* 2011, **35**:2026-2035.
 2. Fuster JM, Alexander GE: **Neuron activity related to short-term memory.** *Science* 1971, **173**:652-654.
 3. Desimone R, Duncan J: **Neural mechanisms of selective visual attention.** *Annu Rev Neurosci* 1995, **18**:193-222.
 4. Nieder A, Freedman DJ, Miller EK: **Representation of the quantity of visual items in the primate prefrontal cortex.** *Science* 2002, **297**.
 5. Shadlen MN, Newsome WT: **Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey.** *J Neurophysiol* 2001, **86**:1916-1936.
 6. Kiani R, Shadlen MN: **Representation of confidence associated with a decision by neurons in the parietal cortex.** *Science* 2009, **324**:759-764.

7. Hampton RR: **Multiple demonstrations of metacognition in nonhumans: converging evidence or multiple mechanisms?** *Comp Cogn Behav Rev* 2009, **4**:17.
8. Mante V, Sussillo D, Shenoy KV, Newsome WT: **Context-dependent computation by recurrent dynamics in prefrontal cortex.** *Nature* 2013, **503**:78-84.
9. Carli M, Robbins TW, Evernden JL, Everitt BJ: **Effects of lesions to ascending noradrenergic neurones on performance of a 5-choice serial reaction task in rats; implications for theories of dorsal noradrenergic bundle function based on selective attention and arousal.** *Behav Brain Res* 1983, **9**:361-380.
10. Roberts AC, Robbins TW, Everitt BJ: **The effects of intradimensional and extradimensional shifts on visual discrimination learning in humans and non-human primates.** *Q J Exp Psychol* 1988, **40**:321-341.
11. Kepecs A, Uchida N, Zariwala HA, Mainen ZF: **Neural correlates, computation and behavioural impact of decision confidence.** *Nature* 2008, **455**:227-231.
12. Miura K, Mainen ZF, Uchida N: **Odor representations in olfactory cortex: distributed rate coding and decorrelated population activity.** *Neuron* 2012, **74**:1087-1098.
13. Hanks TD *et al.*: **Distinct relationships of parietal and prefrontal cortices to evidence accumulation.** *Nature* 2015, **520**:220-223.
14. Murakami M, Vicente MI, Costa GM, Mainen ZF: **Neural antecedents of self-initiated actions in secondary motor cortex.** *Nat Neurosci* 2014.
15. Raposo D, Kaufman MT, Churchland AK: **A category-free neural population supports evolving demands during decision-making.** *Nat Neurosci* 2014.
16. Duan CA, Erlich JC, Brody CD: **Requirement of prefrontal and midbrain regions for rapid executive control of behavior in the rat.** *Neuron* 2015, **86**:1491-1503.
17. Guo ZV *et al.*: **Flow of cortical activity underlying a tactile decision in mice.** *Neuron* 2014, **81**:179-194.
18. Harvey CD, Coen P, Tank DW: **Choice-specific sequences in parietal cortex during a virtual-navigation decision task.** *Nature* 2012, **484**:62-68.
19. Bernstein JG, Boyden ES: **Optogenetic tools for analyzing the neural circuits of behavior.** *Trends Cogn Sci* 2011, **15**:592-600.
20. Deisseroth K: **Optogenetics: 10 years of microbial opsins in neuroscience.** *Nat Neurosci* 2015, **18**:1213-1225.
21. Kobak D *et al.*: **Demixed principal component analysis of population activity in higher cortical areas reveals independent representation of task parameters.** *ArXiv Prepr* 2014. ArXiv14106031.
22. Gao P, Ganguli S: **On simplicity and complexity in the brave new world of large-scale neuroscience.** *Curr Opin Neurobiol* 2015, **32**:148-155.
- Sketching a theory of 'neural task complexity', the authors discuss to what extent neural dimensionality should be expected to increase as more neurons are recorded, especially since only a tiny fraction of neurons can be measured to learn about circuit dynamics and emergent behavior. Instead, their quantitative theoretical formalism suggests the design of richer tasks will become increasingly important in the context of recent technological breakthroughs in large-scale recordings.
23. Tolman EC: **The determiners of behavior at a choice point.** *Psychol Rev* 1938, **45**:1.
24. James W: (1902/1992) *The Varieties of Religious Experience: A Study in Human Nature.* Birmingham, AL: The Classics of Psychiatry and Behavioral Science Library; 1901.
25. Lashley KS: **Basic neural mechanisms in behavior.** *Psychol Rev* 1930, **37**:1.
26. Wright CD: **Is psychological explanation becoming extinct?** *Matter Mind* 2009, **249**.
27. Craik KJW: *The Nature of Explanation.* CUP Archive; 1967.
28. Lake BM, Salakhutdinov R, Tenenbaum JB: **Human-level concept learning through probabilistic program induction.** *Science* 2015, **350**:1332-1338.
- With near perfect one-shot performance, higher than available deep learning algorithms, the authors propose a model that quickly learns from handwritten examples and recreates visual concepts by drawing new letters that are indistinguishable from those produced by humans.
29. Buzsáki G: **Theta rhythm of navigation: link between path integration and landmark navigation, episodic and semantic memory.** *Hippocampus* 2005, **15**:827-840.
30. Markram H *et al.*: **Reconstruction and simulation of neocortical microcircuitry.** *Cell* 2015, **163**:456-492.
31. Van Gelder T: **What might cognition be, if not computation?** *J Philos* 1995:345-381.
32. Van Gelder T: **The dynamical hypothesis in cognitive science.** *Behav Brain Sci* 1998, **21**:615-628.
33. Hodgkin AL, Huxley AF: **A quantitative description of membrane current and its application to conduction and excitation in nerve.** *J Physiol* 1952, **117**:500-544.
34. Machens CK, Brody CD: **Design of continuous attractor networks with monotonic tuning using a symmetry principle.** *Neural Comput* 2008, **20**:452-485.
35. Schulze A *et al.*: **Dynamical feature extraction at the sensory periphery guides chemotaxis.** *eLife* 2015, **4**:e06694.
36. Sussillo D, Abbott LF: **Generating coherent patterns of activity from chaotic neural networks.** *Neuron* 2009, **63**:544-557.
37. Maass W, Natschläger T, Markram H: **Real-time computing without stable states: a new framework for neural computation based on perturbations.** *Neural Comput* 2002, **14**:2531-2560.
38. Eliasmith C *et al.*: **A large-scale model of the functioning brain.** *Science* 2012, **338**:1202-1205.
39. Todorov E: **Optimality principles in sensorimotor control.** *Nat Neurosci* 2004, **7**:907-915.
40. Dayan P, Niv Y: **Reinforcement learning: the good, the bad and the ugly.** *Curr Opin Neurobiol* 2008, **18**:185-196.
41. Glimcher PW, Fehr E: *Neuroeconomics: Decision Making and the Brain.* Academic Press; 2013.
42. Ratcliff R: **Methods for dealing with reaction time outliers.** *Psychol Bull* 1993, **114**:510.
43. Brunton BW, Botvinick MM, Brody CD: **Rats and humans can optimally accumulate evidence for decision-making.** *Science* 2013, **340**:95-98.
44. Pothos EM, Busemeyer JR: **Can quantum probability provide a new direction for cognitive modeling?** *Behav Brain Sci* 2013, **36**:255-274.
- As an alternative to classical probability theory, quantum probability is argued to offer important advantages for modeling certain cognitive and psychological processes. Findings that are hard to explain in the classical context, such as context and order dependence, find an arguably more natural explanation with quantum principles, whose axiomatic mathematical foundations and predictions for event probabilities are different than Bayesian rules.
45. Kvam PD, Pleskac TJ, Yu S, Busemeyer JR: **Interference effects of choice on confidence: quantum characteristics of evidence accumulation.** *Proc Natl Acad Sci U S A* 2015, **112**:10645-10650.
46. Penrose R, Hameroff S: **Consciousness in the universe: neuroscience, quantum space-time geometry and Orch OR theory.** *J Cosmol* 2011, **14**:1-17.
47. Rieke F: *Spikes: Exploring the Neural Code.* MIT Press; 1999.
48. Tononi G: **An information integration theory of consciousness.** *BMC Neurosci* 2004, **5**:42.
49. Tononi G, Koch C: **Consciousness: here, there and everywhere?** *Philos Trans R Soc Lond B Biol Sci* 2015, **370**:20140167.
- Building on Tononi's 'Integrated Information Theory' of consciousness, an information-centered conceptual framework is further expanded and

clarified. This is one of the few current serious scientific approaches to consciousness that, beyond the standard 'neural correlates' empirical strategy, reckons consciousness as a fundamental property of certain systems. Based on a few axioms that define phenomenological requirements, the conceptual framework is claimed to account for the quality of experience as well as prescribing quantitatively the degree of consciousness any system may (or should) have.

50. Wissner-Gross AD, Freer CE: **Causal entropic forces**. *Phys Rev Lett* 2013, **110**:168702.

The authors establish a formal link between intelligence and entropy maximization by means of a generalization of entropic forces that can causally act in the present based on future entropy, generating from simple physical principles remarkable adaptive complex behavior (such as stabilizing an inverted pendulum, solving puzzles or behaving cooperatively).

51. Friston K: **The free-energy principle: a unified brain theory?** *Nat Rev Neurosci* 2010, **11**:127-138.

52. Turing AM: **On computable numbers, with an application to the Entscheidungsproblem**. *J Math* 1936, **58**:5.

53. Doyle JC, Csete M: **Architecture, constraints, and behavior**. *Proc Natl Acad Sci U S A* 2011, **108**:15624-15630.

54. Graves A, Wayne G, Danihelka I: **Neural turing machines**. *ArXiv Prepr*. 2014. ArXiv14105401.

55. Gomes L: **Machine-learning maestro michael jordan on the delusions of big data and other huge engineering efforts**. *IEEE Spectr* 2014, **20**.

56. Gomez-Marín A, Paton JJ, Kampff AR, Costa RM, Mainen ZF: **Big behavioral data: psychology, ethology and the foundations of neuroscience**. *Nat Neurosci* 2014, **17**:1455-1462.

A review of the promises and challenges of big data in behavioral neuroscience, arguing that technology alone does not suffice to make progress but needs to be supplemented by insightful theoretical frameworks and careful experimental design, and emphasizing that behavior is not simply a tool to aid neuroscientists but the very foundation of neuroscience.

57. Anderson DJ, Perona P: **Toward a science of computational ethology**. *Neuron* 2014, **84**:18-31.

58. Berman GJ, Choi DM, Bialek W, Shaevitz JW: **Mapping the stereotyped behaviour of freely moving fruit flies**. *J R Soc Interface* 2014, **11**:20140672.

The authors introduce a method that maps the high-dimensional space of postural dynamics to stereotyped behavioral states in freely moving fruit

flies. The generation of behavioral spaces from arbitrary high-resolution videos allows the application of this method to other organisms such as worms or mice, potentially obtaining their behavioral repertoire directly from the data.

59. Wiltschko AB *et al.*: **Mapping sub-second structure in mouse behavior**. *Neuron* 2015, **88**:1121-1135.

60. Powers WT: *Behavior: The Control of Perception*. Chicago: Aldine; 1973.

61. Yin HH: *Computational and Robotic Models of the Hierarchical Organization of Behavior*. Springer; 2013. pp. 319-347.

A comprehensive summary and reformulation of the basic insights of William Powers' perceptual control theory in the study of behavior. The author argues that if behavior is essentially a control process governed by negative feedback, often-missed implications can be drawn about its hierarchical organization and new experimental directions outlined to search for the controlled variables, which then become the relevant animal-centric aspects of behavior to pin down.

62. Marken RS: **Controlled variables: psychology as the center fielder views it**. *Am J Psychol* 2001, **114**:259-282.

63. Varela FJ: **Neurophenomenology: a methodological remedy for the hard problem**. *J Conscious Stud* 1996, **3**:330-349.

64. Fazelpour S, Thompson E: **The Kantian brain: brain dynamics from a neurophenomenological perspective**. *Curr Opin Neurobiol* 2015, **31**:223-229.

A brief but insightful review on the role of introspection in studying spontaneous brain activity. The authors suggest that first-person neurophenomenological methods are an indispensable approach to shed light on the moment-to-moment relationship between spontaneous brain dynamics and ongoing cognition [64*].

65. Varela FJ, Thompson E, Rosch E: *The Embodied Mind: Cognitive Science and Human Experience*. Cambridge, Massachusetts: The MIT Press; 1993.

66. LeCun Y, Bengio Y, Hinton G: **Deep learning**. *Nature* 2015, **521**:436-444.

A comprehensive short review on how deep learning techniques can discover intricate structure in the data and learn representations by means of back-propagation algorithms at different levels of abstraction.

67. Maxwell JC: **On governors**. *Proc. R. Soc. Lond.* 1867, **16**:270-283.