frontiers in Computational Neuroscience

REVIEW
published: 07 December 2017
doi: 10.3389/fncom.2017.00112

Computational Foundations of Natural Intelligence

*Marcel van Gerven**

Computational Cognitive Neuroscience Lab, Department of Artificial Intelligence, Donders Institute for Brain, Cognition and Behaviour, Radboud University, Nijmegen, Netherlands

New developments in AI and neuroscience are revitalizing the quest to understanding natural intelligence, offering insight about how to equip machines with human-like capabilities. This paper reviews some of the computational principles relevant for understanding natural intelligence and, ultimately, achieving strong AI. After reviewing basic principles, a variety of computational modeling approaches is discussed. Subsequently, I concentrate on the use of artificial neural networks as a framework for modeling cognitive processes. This paper ends by outlining some of the challenges that remain to fulfill the promise of machines that show human-like intelligence.

Keywords: natural intelligence, strong AI, cognition, artificial neural networks, machine learning

1. INTRODUCTION

Understanding how mind emerges from matter is one of the great remaining questions in science. How is it possible that organized clumps of matter such as our own brains give rise to all of our beliefs, desires and intentions, ultimately allowing us to contemplate ourselves as well as the universe from which we originate? This question has occupied cognitive scientists who study the computational basis of the mind for decades. It also occupies other breeds of scientists. For example, ethologists and psychologists focus on the complex behavior exhibited by animals and humans whereas cognitive, computational and systems neuroscientists wish to understand the mechanistic basis of processes that give rise to such behavior.

The ambition to understand natural intelligence as encountered in biological organisms can be contrasted with the motivation to build intelligent machines, which is the subject matter of artificial intelligence (AI). Wouldn't it be amazing if we could build synthetic brains that are endowed with the same qualities as their biological cousins? This desire to mimic human-level intelligence by creating artificially intelligent machines has occupied mankind for many centuries. For instance, mechanical men and artificial beings appear in Greek mythology and realistic human automatons had already been developed in Hellenic Egypt (McCorduck, 2004). The engineering of machines that display human-level intelligence is also referred to as strong AI (Searle, 1980) or artificial general intelligence (AGI) (Adams et al., 2012), and was the original motivation that gave rise to the field of AI (Newell, 1991; Nilsson, 2005).

Excitingly, major advances in various fields of research now make it possible to attack the problem of understanding natural intelligence from multiple angles. From a theoretical point of view we have a solid understanding of the computational problems that are solved by our own brains (Dayan and Abbott, 2005). From an empirical point of view, technological breakthroughs allow us to probe and manipulate brain activity in unprecedented ways, generating new neuroscientific insights into brain structure and function (Chang, 2015). From an engineering perspective, we are finally able to build machines that learn to solve complex tasks, approximating and sometimes surpassing human-level performance (Jordan and Mitchell, 2015). Still, these efforts

OPEN ACCESS

Edited by:
Florentin Wörgötter,
University of Göttingen, Germany

Reviewed by:
Sebastian Herzog,
Max Planck Institute for Dynamics and Self Organization (MPG), Germany
Carme Torras,
Consejo Superior de Investigaciones Científicas (CSIC), Spain

***Correspondence:**
Marcel van Gerven
m.vangerven@donders.ru.nl

Received: *01 August 2017*
Accepted: *22 November 2017*
Published: *07 December 2017*

Citation:
van Gerven M (2017) Computational Foundations of Natural Intelligence. Front. Comput. Neurosci. 11:112. doi: 10.3389/fncom.2017.00112

have not yet provided a full understanding of natural intelligence, nor did they give rise to machines whose reasoning capacity parallels the generality and flexibility of cognitive processing in biological organisms.

The core thesis of this paper is that natural intelligence can be better understood by the coming together of multiple complementary scientific disciplines (Gershman et al., 2015). This thesis is referred to as *the great convergence*. The advocated approach is to endow artificial agents with synthetic brains (i.e., cognitive architectures, Sun, 2004) that mimic the thought processes that give rise to ethologically relevant behavior in their biological counterparts. A motivation for this approach is given by Braitenberg's law of uphill analysis and downhill invention, which states that it is much easier to understand a complex system by assembling it from the ground up, rather than by reverse engineering it from observational data (Braitenberg, 1986). These synthetic brains, which can be put to use in virtual or real-world environments, can then be validated against neuro-behavioral data and analyzed using a multitude of theoretical tools. This approach not only elucidates our understanding of human brain function but also paves the way for the development of artificial agents that show truly intelligent behavior (Hassabis et al., 2017).

The aim of this paper is to sketch the outline of a research program which marries the ambitions of neuroscientists to understand natural intelligence and AI researchers to achieve strong AI (**Figure 1**). Before embarking on our quest to build synthetic brains as models of natural intelligence, we need to formalize what problems are solved by biological brains. That is, we first need to understand how adaptive behavior ensues in animals and humans.

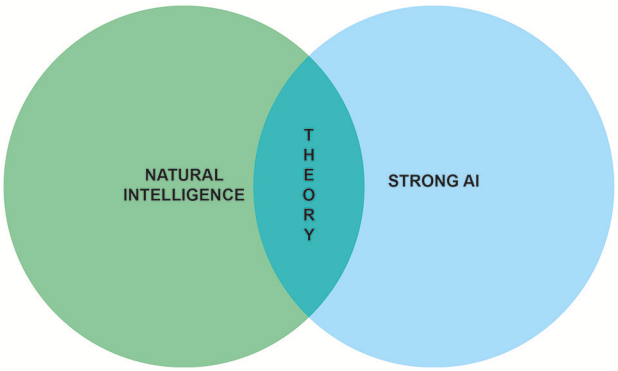

FIGURE 1 | Understanding natural intelligence and achieving strong AI are seen as relying on the same theoretical foundations and require the convergence of multiple scientific and engineering disciplines.

2. ADAPTIVE BEHAVIOR IN BIOLOGICAL AGENTS

Ultimately, organisms owe their existence to the fact that they promote survival of their constituent genes; the basic physical and functional units of heredity that code for an organism (Dawkins, 2016). At evolutionary time scales, organisms developed a range of mechanisms which ensure that they live long enough such as to produce offspring. For example, single-celled protozoans already show rather complex ingestive, defensive and reproductive behavior, which is regulated by molecular signaling (Swanson, 2012; Sterling and Laughlin, 2016).

2.1. Why Do We Need a Brain?

About 3.5 billion years ago, multicellular organisms started to appear. Multicellularity offers several competitive advantages over unicellularity. It allows organisms to increase in size without the limitations set by unicellularity and permits increased complexity by allowing cellular differentiation. It also increases life span since an organism can live beyond the demise of a single cell. At the same time, due to their increased size and complexity, multicellular organisms require more intricate mechanisms for signaling and regulation.

In multicellular organisms, behavior is regulated at multiple scales, ranging from intracellular molecular signaling all the way up to global regulation via the interactions between different organ systems. Hence, the nervous system allows for fast responses via electrochemical signaling and for slow responses by acting on the endocrine system. Nervous systems are found in almost all multicellular animals, but vary greatly in complexity. For example, the nervous system of the nematode roundworm *Caenorhabditis elegans* (*C. elegans*) is made up of 302 neurons and 7,000 synaptic connections (White et al., 1986; Varshney et al., 2011). In contrast, the human brain contains about 20 billion neocortical neurons that are wired together via as many as 0.15 quadrillion synapses (Pakkenberg and Gundersen, 1997; Pakkenberg et al., 2003).

In vertebrates, the nervous system can be partitioned into the central nervous system (CNS), consisting of the brain and the spinal cord, and the peripheral nervous system (PNS), which connects the CNS to every other part of the body. The brain allows for centralized control and efficient information transmission. It can be partitioned into the forebrain, midbrain and hindbrain, each of which contain dedicated neural circuits that allow for integration of information and generation of coordinated activity. The spinal cord connects the brain to the body by allowing sensory and motor information to travel back and forth between the brain and the body. It also coordinates certain reflexes that bypass the brain altogether.

The interplay between the nervous system, the body and the environment is nicely captured by Swanson's four system model of nervous system organization (Swanson, 2000), as shown in **Figure 2**. Briefly, the brain exerts centralized control on the body by sending commands to the motor system based on information received via the sensory system. It exerts this control by way of the cognitive system, which drives voluntary initiation of behavior, as well as the state system, which refers to the intrinsic activity that controls global behavioral state. The motor system can also be influenced directly by the sensory system via spinal cord reflexes. Output of the motor system induces visceral responses that affect bodily state as well as somatic responses that act on the environment. It is also able to drive the secretion of

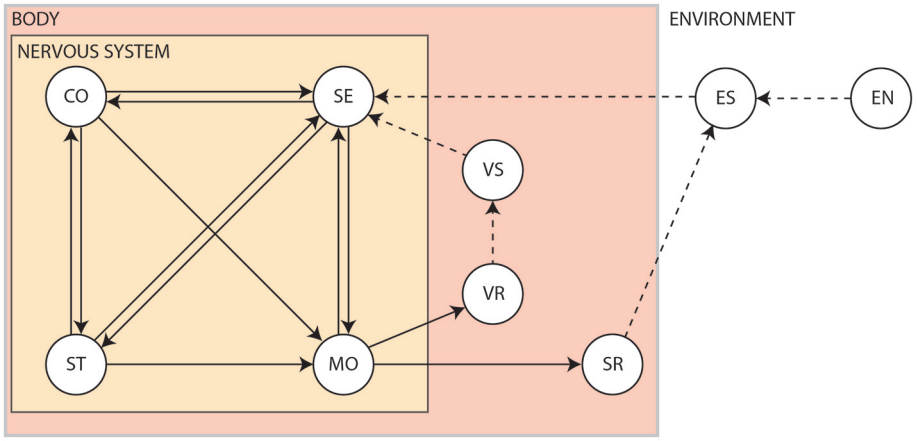

FIGURE 2 | The four system model of nervous system organization. CO, Cognitive system; EN, Environment; ES, Environmental stimuli; MO, Motor system; SE, Sensory system; SR, Somatic responses; ST, Behavioral state system; VR, Visceral responses; VS, Visceral stimuli. Solid arrows show influences pertaining to the nervous system. Dashed arrows show interactions produced by the body or the environment[1].

hormones that act more globally on the body. Both the body and the environment generate sensations that are processed by the sensory system. This closed-loop system, tightly coupling sensation, thought and action, is known as the *perception-action cycle* (Dewey, 1896; Sperry, 1952; Fuster, 2004).

Summarizing, the brain, together with the spinal cord and the peripheral nervous system, can be seen as an organ that exploits sensory input such as to generate adaptive behavior through motor outputs. This ensures an organism's long-term survival in a world that is dominated by uncertainty, as a result of partial observability, noise and stochasticity. The upshot of this interpretation is that action, which drives the generation of adaptive behavior, is the ultimate reason why we have a brain in the first place. Citing Sperry (1952): "the entire output of our thinking machine consists of nothing but patterns of motor coordination." To understand how adaptive behavior ensues, we therefore need to identify the ultimate causes that determine an agent's actions (Tolman, 1932).

2.2. What Makes us Tick?

In biology, ultimately, all evolved traits must be connected to an organism's survival. This implies that, from the standpoint of evolutionary psychology, natural selection favors those behaviors and thought processes that provide the organism with a selective advantage under ecological pressure (Barkow et al., 1992). Since causal links between behavior and long-term survival cannot be sensed or controlled directly, an agent needs to rely on other, directly accessible, ways to promote its survival. This can take the form of (1) evolving optimal sensors and effectors that allow it to maximize its control given finite resources and (2) evolving a behavioral repertoire that maximizes the information gained from the environment and generates optimal actions based on available sensory information.

[1]Figure modified from http://larrywswanson.com/?page_id=1523 with permission.

In practice, behavior is the result of multiple competing needs that together provide an evolutionary advantage. These needs arise because they provide particular rewards to the organism. We distinguish *primary rewards*, *intrinsic rewards* and *extrinsic rewards*.

Primary Rewards

Primary rewards are those necessary for the survival of one's self and offspring, which includes homeostatic and reproductive rewards. Here, homeostasis refers to the maintenance of optimal settings of various biological parameters (e.g., temperature regulation) (Cannon, 1929). A slightly more sophisticated concept is *allostasis*, which refers to the predictive regulation of biological parameters in order to prevent deviations rather than correcting them *post hoc* (Sterling, 2012). An organism can use its nervous system (muscle signaling) or endocrine system (endocrine signaling) to globally control or adjust the activities of many systems simultaneously. This allows for visceral responses that ensure proper functioning of an agent's internal organs as well as basic drives such as ingestion, defense and reproduction that help ensure an agent's survival (Tinbergen, 1951).

Intrinsic Rewards

Intrinsic rewards are unconditioned rewards that are attractive and motivate behavior because they are inherently pleasurable (e.g., the experience of joy). The phenomenon of intrinsic motivation was first identified in studies of animals engaging in exploratory, playful and curiosity-driven behavior in the absence of external rewards or punishments (White, 1959).

Extrinsic Rewards

Extrinsic rewards are conditioned rewards that motivate behavior but are not inherently pleasurable (e.g., praise or monetary reward). They acquire their value through learned association with intrinsic rewards. Hence, extrinsic motivation refers to our tendency to perform activities for known external rewards,

whether they be tangible or psychological in nature (Brown, 2007).

Summarizing, the continual competition between multiple drives and incentives that have adaptive value to the organism and are realized by dedicated neural circuits is what ultimately generates behavior (Davies et al., 2012). In humans, the evolutionary and cultural pressures that shaped our own intrinsic and extrinsic motivations have allowed us to reach great achievements, ranging from our mastery of the laws of nature to expressions of great beauty as encountered in the liberal arts. The question remains how we can gain an understanding of how our brains generate the rich behavioral repertoire that can be observed in nature.

3. UNDERSTANDING NATURAL INTELLIGENCE

In a way, the recipe for understanding natural intelligence and achieving strong AI is simple. If we can construct synthetic brains that mimic the adaptive behavior displayed by biological brains in all its splendor then our mission has succeeded. This entails equipping synthetic brains with the same special purpose computing machinery encountered in real brains, solving those problems an agent may be faced with. In practice, of course, this is easier said than done given the incomplete state of our knowledge and the daunting complexity of biological systems.

3.1. Levels of Analysis

The neural circuits that make up the human brain can be seen as special-purpose devices that together guarantee the selection of (near-)optimal actions. David Marr in particular advocated the view that the nervous system should be understood as a collection of information processing systems that solve particular problems an organism is faced with (Marr, 1982). His work gave rise to the field of computational neuroscience and has been highly influential in shaping ideas about neural information processing (Willshaw et al., 2015). Marr and Poggio (1976) proposed that an understanding of information processing systems should take place at distinct levels of analysis, namely the *computational level*, which specifies what problem the system solves, the *algorithmic level*, which specifies how the system solves the problem, and the *implementational level*, which specifies how the system is physically realized.

A canonical example of a three-level analysis is prey localization in the barn owl (Grothe, 2003). At the computational level, the owl needs to use auditory information to localize its prey. At the algorithmic level, this can be implemented by circuits composed of delay lines and coincidence detectors that detect inter-aural time differences (Jeffress, 1948). At the implementational level, neurons in the nucleus laminaris have been shown to act as coincidence detectors (Carr and Konishi, 1990).

Marr's levels of analysis sidestep one important point, namely how a system gains the ability to solve a computational problem in the first place. That is, it is also crucial to understand how an organism (or species as a whole) is able to learn and evolve the computations and representations that allow it to survive in the natural world (Poggio, 2012). Learning itself takes place at the level of the individual organism as well as of the species. In the individual, one can observe lasting changes in the brain throughout its lifetime, which is referred to as neural plasticity. At the species level, natural selection is responsible for evolving the mechanisms that are involved in neural plasticity (Poggio, 2012). As argued by Poggio, an understanding at the level of learning in the individual and the species is sufficiently powerful to solve a problem and can thereby act as an explanation of natural intelligence. To illustrate the relevance of this revised model, in the prey localization example it would be imperative to understand how owls are able to adapt to changes in their environment (Huo and Murray, 2009), as well as how owls were equipped with such machinery during evolution.

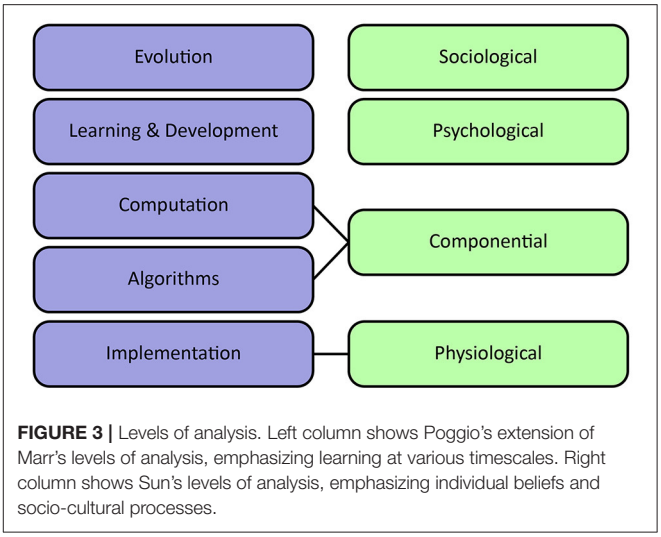

FIGURE 3 | Levels of analysis. Left column shows Poggio's extension of Marr's levels of analysis, emphasizing learning at various timescales. Right column shows Sun's levels of analysis, emphasizing individual beliefs and socio-cultural processes.

Sun et al. (2005) propose an alternative organization of levels of cognitive modeling. They distinguish sociological, psychological, componential and physiological levels. The sociological level refers to the collective behavior of agents, including interactions between agents as well as their environment. It stresses the importance of socio-cultural processes in shaping cognition. The psychological level covers individual behaviors, beliefs, concepts, and skills. The componential level describes inter-agent processes specified in terms of Marr's computational and algorithmic levels. Finally, the physiological level describes the biological substrate which underlies the generation of adaptive behavior, corresponding to Marr's implementational level. It can provide valuable input about important computations and plausible architectures at a higher level of abstraction.

Figure 3 visualizes the different interpretations of levels of analysis. Without committing to a definitive stance on levels of analysis, all described levels provide important complementary perspectives concerning the modeling and understanding of natural intelligence.

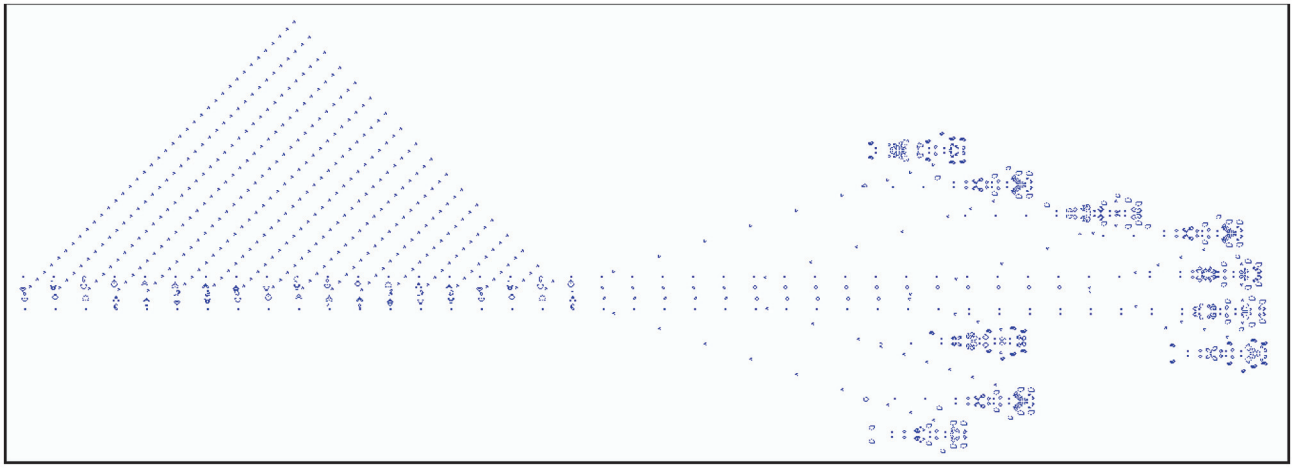

FIGURE 4 | Example of the Game of Life, where each cell state evolves according to a set of deterministic rules that depend on the states of neighboring cells. Depicted is a *breeder* pattern that moves across the universe (here from left to right), leaving behind debris. The breeder produces *Gosper guns* which periodically emit *gliders*; the small patterns that together form the triangular shape on the left-hand side.

3.2. Modeling approaches

The previous section suggests that different approaches to understanding natural intelligence and developing cognitive architectures can be taken depending on the levels of analysis one considers. We briefly review a number of core approaches.

Artificial Life

Artificial life is a broad area of research encompassing various different modeling strategies which all have in common that they aim to explain the emergence of life and, ultimately, cognition in a bottom-up manner (Steels, 1993; Bedau, 2003).

A canonical example of an artificial life system is the cellular automaton, first introduced by von Neumann (1966) as an approach to understand the fundamental properties of living systems. Cellular automata operate within a universe consisting of cells, whose states change over multiple generations based on simple local rules. They have been shown to be capable of acting as universal Turing machines, thereby giving them the capacity to compute any fixed partial computable function (Wolfram, 2002).

A famous example of a cellular automaton is Conway's Game of Life. Here, every cell can assume an "alive" or a "dead" state. State changes are determined by its interactions with its eight direct neighbors. At each time step, a live cell with fewer than two or more than three live neighbors dies and a dead cell with exactly three live neighbors will become alive. **Figure 4** shows an example of a breeder pattern which produces Gosper guns in the Game of Life. Gosper guns have been used to prove that the game of life is Turing complete (Gardner, 2001). SmoothLife (Rafler, 2011), as a continuous-space extension of the Game of Life, shows emerging structures that bear some superficial resemblance to biological structures.

In principle, by virtue of their universality, cellular automata offer the capacity to explain how self-replicating adaptive (autopoeietic, Maturana and Varela, 1980) systems emerge from basic rules. This bottom-up approach is also taken by physicists who aim to explain life and, ultimately, cognition purely from thermodynamic principles (Dewar, 2003, 2005; Grinstein and Linsker, 2007; Wissner-Gross and Freer, 2013; Perunov et al., 2014; Fry, 2017).

Biophysical Modeling

A more direct way to model natural intelligence is to presuppose the existence of the building blocks of life which can be used to create realistic simulations of organisms *in silico*. The reasoning is that biophysically realistic models can eventually mimic the information processing capabilities of biological systems. An example thereof is the OpenWorm project which has as its ambition to understand how the behavior of *C. elegans* emerges from its underlying physiology purely via bottom-up biophysical modeling (Szigeti et al., 2014) (**Figure 5A**). It also acknowledges the importance of including not only a model of the worm's nervous system but also of its body and environment in the simulation. That is, adaptive behavior depends on the organism being both embodied and embedded in the world (Anderson, 2003). If successful, then this project would constitute the first example of a digital organism.

It is a long stretch from the worm's 302 neurons to the 86 billion neurons that comprise the human brain (Herculano-Houzel and Lent, 2005). Still, researchers have set out to develop large-scale models of the human brain. Biophysical modeling can be used to create detailed models of neurons and their processes using coupled systems of differential equations. For example, action potential generation can be described in terms of the Hodgkin-Huxley equations (**Figure 5B**) and the flow of electric current along neuronal fibers can be modeled using cable theory (Dayan and Abbott, 2005). This approach is used in the Blue Brain project (Markram, 2006) and its successor, the Human Brain Project (HBP) (Amunts et al., 2016). See de Garis et al. (2010) for a review of various artificial brain projects.

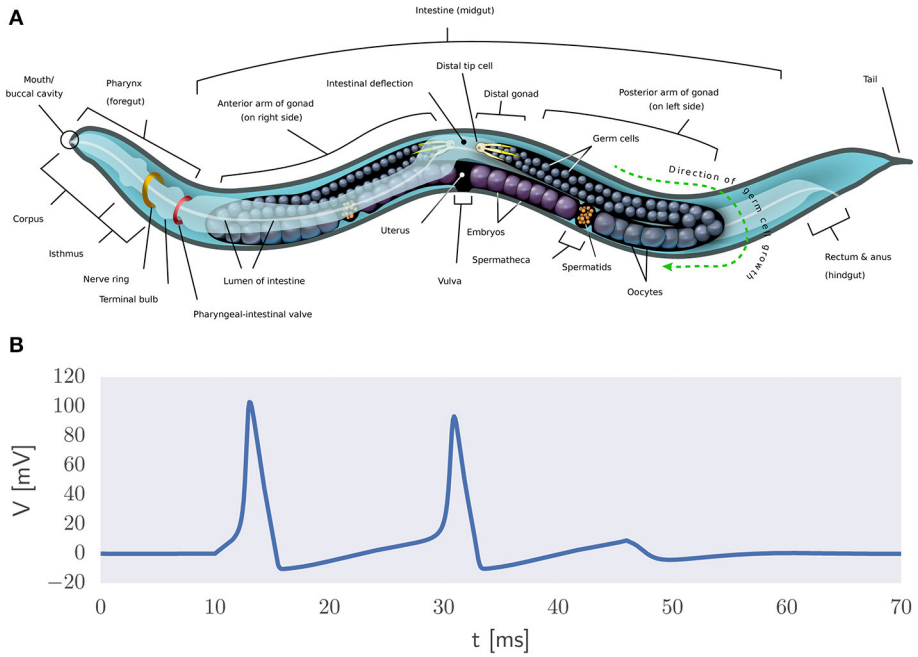

FIGURE 5 | Biophysical modeling. **(A)** Body plan of *C. elegans*[2]. The OpenWorm project aims to provide an accurate bottom-up simulation of the worm acting in its environment. **(B)** Example of action potential generation via the Hodgkin-Huxley equations in the presence of a constant input current.

Connectionism

Connectionism refers to the explanation of cognition as arising from the interplay between basic (sub-symbolic) processing elements (Smolensky, 1987; Bechtel, 1993). It has close links to cybernetics, which focuses on the development of control structures from which intelligent behavior emerges (Rid, 2016).

Connectionism came to be equated with the use of artificial neural networks that abstract away from the details of biological neural networks. An artificial neural network (ANN) is a computational model which is loosely inspired by the human brain as it consists of an interconnected network of simple processing units (artificial neurons) that learns from experience by modifying its connections. Alan Turing was one of the first to propose the construction of computing machinery out of trainable networks consisting of neuron-like elements (Copeland and Proudfoot, 1996). Marvin Minsky, one of the founding fathers of AI, is credited for building the first trainable ANN, called SNARC, out of tubes, motors, and clutches (Seising, 2017).

Artificial neurons can be considered abstractions of (populations of) neurons while the connections are taken to be abstractions of modifiable synaptic connections (**Figure 6**). The behavior of an artificial neuron is fully determined by the connection strengths as well as how input is transformed into output. Contrary to detailed biophysical models, ANNs make use of basic matrix operations and nonlinear transformations as their fundamental operations. In its most basic incarnation, an artificial neuron simply transforms its input $\mathbf{x}$ into a response y through an activation function f, as shown in **Figure 6**. The activation function operates on an input activation which is typically taken to be the inner product between the input $\mathbf{x}$ and the parameters (weight vector) $\mathbf{w}$ of the artificial neuron. The weights are interpreted as synaptic strengths that determine how presynaptic input is translated into postsynaptic firing rate. This yields a simple linear-nonlinear mapping of the form

$$y = f(\mathbf{w}^{\mathsf{T}}\mathbf{x}) . \tag{1}$$

By connecting together multiple neurons, one obtains a neural network that implements some non-linear function $\mathbf{y} = \mathbf{f}(\mathbf{x}; \boldsymbol{\theta})$, where the f_i are nonlinear transformations and $\boldsymbol{\theta}$ stands for the network parameters (i.e., weight vectors). After training a neural network, representations become encoded in a distributed manner as a pattern which manifests itself across all its neurons (Hinton et al., 1986).

Throughout the course of their history ANNs have fallen in and out of favor multiple times. At the same time, each next generation of neural networks has yielded new insights about how complex behavior may emerge through the collective action of simple processing elements. Modern neural networks perform so well on several benchmark problems that they obliterate all competition in, e.g., object recognition (Krizhevsky et al., 2012), natural language processing (Sutskever et al., 2014), game playing (Mnih et al., 2015; Silver et al., 2017) and robotics (Levine et al., 2015), often matching and sometimes surpassing human-level performance (LeCun et al., 2015). Their success relies on combining classical ideas (Widrow and Lehr, 1990; Hochreiter

[2]Figure by K. D. Schroeder, CC BY-SA 3.0, https://commons.wikimedia.org/w/index.php?curid=26958836. Used with permission.

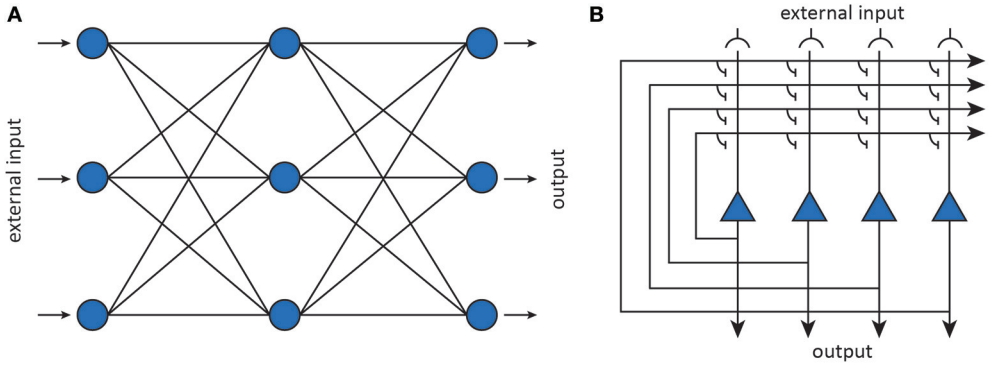

FIGURE 6 | Artificial neural networks (Yuste, 2015). **(A)** Feedforward neural networks map inputs to outputs using nonlinear transformations. **(B)** Recurrent neural networks implement dynamical systems by feeding back output activity to the input layer, where it is combined with external input.

and Schmidhuber, 1997; LeCun et al., 1998) with new algorithmic developments (Hinton et al., 2006; Srivastava et al., 2014; He et al., 2015; Ioffe and Szegedy, 2015; Zagoruyko and Komodakis, 2017), while using high-performance graphical processing units (GPUs) to massively speed up training of ANNs on big datasets (Raina et al., 2009).

Cognitivism

A conceptually different approach to the explanation of cognition as emerging from bottom-up principles is the view that cognition should be understood in terms of formal symbol manipulation. This computationalist view is associated with the cognitivist program which arose in response to earlier behaviorist theories. It embraces the notion that, in order to understand natural intelligence, one should study internal mental processes rather than just externally observable events. That is, cognitivism asserts that cognition should be defined in terms of formal symbol manipulation, where reasoning involves the manipulation of symbolic representations that refer to information about the world as acquired by perception.

This view is formalized by the physical symbol system hypothesis (Newell and Simon, 1976), which states that "a physical symbol system has the necessary and sufficient means for intelligent action." This hypothesis implies that artificial agents, when equipped with the appropriate symbol manipulation algorithms, will be capable of displaying intelligent behavior. As Newell and Simon (1976) wrote, the physical symbol system hypothesis also implies that "the symbolic behavior of man arises because he has the characteristics of a physical symbol system." This also suggests that the specifics of our nervous system are not relevant for explaining adaptive behavior (Simon, 1996).

Cognitivism gave rise to cognitive science as well as artificial intelligence, and spawned various cognitive architectures such as ACT-R (Anderson et al., 2004) (see **Figure 7**) and SOAR (Laird, 2012) that employ rule-based approaches in the search for a unified theory of cognition (Newell, 1991).[3]

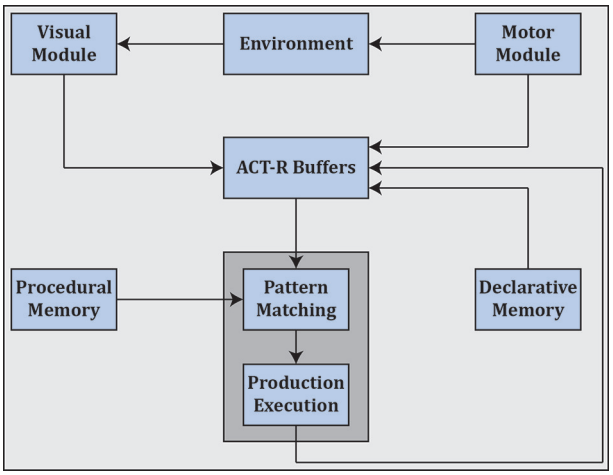

FIGURE 7 | ACT-R as an example cognitive architecture which employs symbolic reasoning. ACT-R interfaces with different modules through buffers. Cognition unfolds as a succession of activations of production rules as mediated by pattern matching and execution[4].

Probabilistic Modeling

Modern cognitive science still embraces the cognitivist program but has since taken a probabilistic approach to the modeling of cognition. As stated by Griffiths et al. (2010), this probabilistic approach starts from the notion that the challenges faced by the mind are often of an inductive nature, where the observed data are not sufficient to unambiguously identify the process that generated them. This precludes the use of approaches that are founded on mathematical logic and requires a quantification of the state of the world in terms of degrees of belief as afforded by probability theory (Jaynes, 1988). The probabilistic approach operates by identifying a hypothesis space representing solutions to the inductive problem. It then prescribes how an agent should revise her belief in the hypotheses given the information provided

[3]In fact, ACT-R also uses some subsymbolic elements and can therefore be considered a *hybrid* architecture.

[4]Figure modified from http://act-r.psy.cmu.edu/about with permission.

by observed data. Hypotheses are typically formulated in terms of probabilistic graphical models that capture the independence structure between random variables of interest (Koller and Friedman, 2009). An example of such a graphical model is shown in **Figure 8**.

Belief updating in the probabilistic sense is realized by solving a statistical inference problem. Consider a set of of hypotheses $\mathcal{H}$ that might explain the observed data. Let $p(h)$ denote our belief in a hypothesis $h \in \mathcal{H}$, reflecting the state of the world, before observing any data (known as the *prior*). Let $p(\mathbf{x} \mid h)$ indicate the probability of observing data $\mathbf{x}$ if h were true (known as the *likelihood*). Bayes' rule tells us how to update our belief in a hypothesis after observing data. It states that the *posterior probability* $p(h \mid \mathbf{x})$ assigned to h after observing $\mathbf{x}$ should be

$$p(h \mid \mathbf{x}) = \frac{p(\mathbf{x} \mid h)p(h)}{\sum_{h \in \mathcal{H}} p(\mathbf{x} \mid h)p(h)} \tag{2}$$

where the denominator is a normalizing constant known as the *evidence* or *marginal likelihood*[5]. Importantly, it can be shown that degrees of belief are coherent only if they satisfy the axioms of probability theory (Ramsey, 1926).

The beauty of the probabilistic approach lies in its generality. It not only explains how our moment-to-moment percepts change as a function of our prior beliefs and incoming sensory data (Yuille and Kersten, 2006) but also places learning, as the construction of internal models, under the same umbrella by viewing it as an inference problem (MacKay, 2003). In the probabilistic framework, mental processes are modeled using algorithms for approximating the posterior (Koller and Friedman, 2009) and neural processes are seen as mechanisms for implementing these algorithms (Gershman and Beck, 2016).

The probabilistic approach also provides a basis for making optimal decisions under uncertainty. This is realized by extending probability theory with decision theory. According to decision theory, a rational agent ought to select that action which maximizes the expected utility (von Neumann and Morgenstern, 1953). This is known as the maximum expected utility (MEU) principle. In real-life situations, biological (and artificial) agents need to operate under bounded resources, trading off precision for speed and effort when trying to attain their objectives (Gigerenzer and Goldstein, 1996). This implies that MEU calculations may be intractable. Intractability issues have led to the development of algorithms that maximize a more general form of expected utility which incorporates the costs of computation. These algorithms can in turn be adapted so as to select the best approximation strategy in a given situation (Gershman et al., 2015). Hence, at the algorithmic level, it has been postulated that brains use approximate inference algorithms (Andrieu et al., 2003; Blei et al., 2016) such as to produce good enough solutions for fast and frugal decision making.

Summarizing, by appealing to Bayesian statistics and decision theory, while acknowledging the constraints biological agents are faced with, cognitive science arrives at a theory of bounded rationality that agents should adhere to. Importantly, this normative view dictates that organisms must operate as Bayesian inference machines that aim to maximize expected utility. If they do not, then, under weak assumptions, they will perform suboptimally. This would be detrimental from an evolutionary point of view.

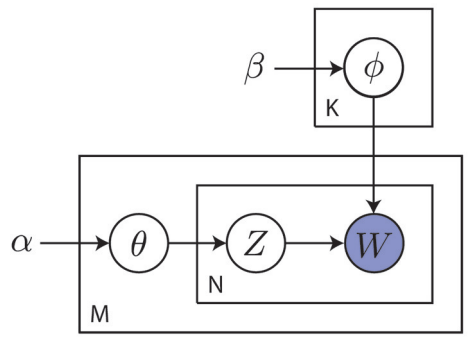

FIGURE 8 | Example of a probabilistic graphical model capturing the statistical relations between random variables of interest. This particular plate model describes a smoothed version of latent Dirichlet allocation as used in topic modeling (Blei et al., 2003). Here, α and β are hyper-parameters, θ_m is the topic distribution for document m, ϕ_k is the word distribution for topic k, z_{nm} is the topic for the n-th word in document m and w_{mn} is a specific word. Capital letters K, M and N denote the number of topics, documents and words, respectively. The goal is to discover abstract topics from observed words. This general approach of inferring posteriors over latent variables from observed data is common to the probabilistic approach.

3.3. Bottom-up Emergence vs. Top-down Abstraction

The aforementioned modeling strategies each provide an alternative approach toward understanding natural intelligence and achieving strong AI. The question arises which of these strategies will be most effective in the long run.

While the strictly bottom-up approach used in artificial life research may lead to fundamental insights about the nature of self-replication and adaptability, in practice it remains an open question how emergent properties that derive from a basic set of rules can reach the same level of organization and complexity as can be found in biological organisms. Furthermore, running such simulations would be extremely costly from a computational point of view.

The same problem presents itself when using detailed biophysical models. That is, bottom-up approaches must either restrict model complexity or run simulations for limited periods of time in order to remain tractable (O'Reilly et al., 2012). Biophysical models additionally suffer from a lack of data. For example, the original aim of the Human Brain Project was to model the human brain within a decade (Markram et al., 2011). This ambition may be hard to realize given the plethora of data required for model estimation. Furthermore, the resulting models may be difficult to link to cognitive function. Izhikevich, reflecting on his simulation of another large biophysically realistic brain model (Izhikevich and Edelman, 2008), states:

[5]Beliefs over continuous quantities can be expressed by replacing summation with integration.

"Indeed, no significant contribution to neuroscience could be made by simulating one second of a model, even if it has the size of the human brain. However, I learned what it takes to simulate such a large-scale system[6]."

Connectionist models, in contrast, abstract away from biophysical details, thereby making it possible to train large-scale models on large amounts of sensory data, allowing cognitively challenging tasks to be solved. Due to their computational simplicity, they are also more amenable to theoretical analysis (Hertz et al., 1991; Bishop, 1995). At the same time, connectionist models have been criticized for their inability to capture symbolic reasoning, their limitations when modeling particular cognitive phenomena, and their abstract nature, restricting their biological plausibility (Dawson and Shamanski, 1994).

Cognitivism has been pivotal in the development of intelligent systems. However, it has also been criticized using the argument that systems which operate via formal symbol manipulation lack intentionality (Searle, 1980)[7]. Moreover, the representational framework that is used is typically constructed by a human designer. While this facilitates model interpretation, at the same time, this programmer-dependence may bias the system, leading to suboptimal solutions. That is, idealized descriptions may induce a semantic gap between perception and possible interpretation (Vernon et al., 2007).

The probabilistic approach to cognition is important given its ability to define normative theories at the computational level. At the same time, it has also been criticized for its treatment of cognition as if it is in the business of selecting some statistical model. Proponents of connectionism argue that computation-level explanations of behavior that ignore mechanisms associated with bottom-up emergence are likely to fall short (McClelland et al., 2010).

The different approaches provide complementary insights into the nature of natural intelligence. Artificial life informs about fundamental bottom-up principles, biophysical models make explicit how cognition is realized via specific mechanisms at the molecular and systems level, connectionist models show how problem solving capacities emerge from the interactions between basic processing elements, cognitivism emphasizes the importance of symbolic reasoning and probabilistic models inform how particular problems could be solved in an optimal manner.

Notwithstanding potential limitations, given their ability to solve complex cognitively challenging problems, connectionist models are taken to provide a promising starting point for understanding natural intelligence and achieving strong AI. They also naturally connect to the different modeling strategies. That is, they connect to artificial life principles by having network architectures emerge through evolutionary strategies (Real et al., 2016; Salimans et al., 2017) and connect to the biophysical level by viewing them as (rate-based) abstractions of biological neural networks (Dayan and Abbott, 2005). They also connect to the computational level by grounding symbolic representations in real-world sensory states (Harnad, 1990) and connect to the probabilistic approach through the observation that emergent computations effectively approximate Bayesian inference (Gal, 2016; Orhan and Ma, 2016; Ambrogioni et al., 2017; Mandt et al., 2017). It is for these reasons that, in the following, we will explore how ANNs, as canonical connectionist models, can be used to promote our understanding of natural intelligence.

[6]From: https://www.izhikevich.org/human_brain_simulation/why.htm

[7]Intentionality or "aboutness" refers to the quality of mental states as being directed toward an object or state of affairs.

4. ANN-BASED MODELING OF COGNITIVE PROCESSES

We will now explore in more detail the ways in which ANNs can be used to understand and model aspects of natural intelligence. We start by addressing how neural networks can learn from data.

4.1. Learning

The capacity of brains to behave adaptively relies on their ability to modify their own behavior based on changing circumstances. The appeal of neural networks stems from their ability to mimic this learning behavior in an efficient manner by updating network parameters $\boldsymbol{\theta}$ based on available data $\mathcal{D} = \{\mathbf{z}^{(1)}, \ldots, \mathbf{z}^{(N)}\}$, allowing the construction of large models that are able to solve complex cognitive tasks.

Learning proceeds by making changes to the network parameters $\boldsymbol{\theta}$ such that its output starts to agree more and more with the objectives of the agent at hand. This is formalized by assuming the existence of a cost function $\mathcal{J}(\boldsymbol{\theta})$ which measures the degree to which an agent deviates from its objectives. $\mathcal{J}$ is computed by running a neural network in forward mode (from input to output) and comparing the predicted output with the desired output. During its lifetime, the agent obtains data from its environment (sensations) by sampling from a data-generating distribution p_{data}. The goal of an agent is to reduce the expected risk

$$\mathcal{J}^*(\boldsymbol{\theta}) = \mathbb{E}_{\mathbf{z}\sim p_{\text{data}}}\left[\ell(\mathbf{z}, \boldsymbol{\theta})\right] \tag{3}$$

where ℓ is the incurred loss per datapoint $\mathbf{z}$. In practice, an agent only has access to a finite number of datapoints which the agent experiences during its lifetime, yielding a training set $\mathcal{D}$. This training set can be represented in the form of an empirical distribution $\hat{p}(\mathbf{z})$ which equals $1/N$ if $\mathbf{z}$ is equal to one of the N examples and zero otherwise. In practice, the aim therefore is to minimize the empirical risk

$$\mathcal{J}(\boldsymbol{\theta}) = \mathbb{E}_{\mathbf{z}\sim \hat{p}}\left[\ell(\mathbf{z}, \boldsymbol{\theta})\right] \tag{4}$$

as an approximation of $\mathcal{J}^*$. In reality, the brain is thought to optimize a multitude of cost functions pertaining to the many objectives it aims to achieve in concert (Marblestone et al., 2016).

Risk minimization can be accomplished by making use of a gradient descent procedure. Let $\boldsymbol{\theta}$ be the parameters of a neural network (i.e., the synaptic weights). We can define learning as a search for the optimal parameters $\boldsymbol{\theta}^*$ based on available training data $\mathcal{D}$ such that

$$\boldsymbol{\theta}^* = \arg\min_{\boldsymbol{\theta}} \mathcal{J}(\boldsymbol{\theta}) . \tag{5}$$

A convenient way to approximate $\boldsymbol{\theta}^*$ is by measuring locally the change in slope of $\mathcal{J}(\boldsymbol{\theta})$ as a function of $\boldsymbol{\theta}$ and taking a step in the direction of steepest descent. This procedure, known as *gradient descent*, is based on the observation that if $\mathcal{J}$ is defined and differentiable in the neighborhood of a point $\boldsymbol{\theta}$, then $\mathcal{J}$ decreases fastest if one goes from $\boldsymbol{\theta}$ in the direction of the negative gradient $-\nabla_{\boldsymbol{\theta}}\mathcal{J}(\boldsymbol{\theta})$. In other words, if we use the update rule

$$\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} - \epsilon \nabla_{\boldsymbol{\theta}} \mathcal{J}(\boldsymbol{\theta}) \tag{6}$$

with small enough learning rate ϵ then $\boldsymbol{\theta}$ is guaranteed to converge to a (local) minimum of $\mathcal{J}(\boldsymbol{\theta})$[8]. Importantly, the gradient can be computed for arbitrary ANN architectures by running the network in backward mode (from output to input) and computing the gradient using automatic differentiation procedures. This forms the basis of the widely used backpropagation algorithm (Widrow and Lehr, 1990).

One might argue that the backpropagation algorithm fails to connect to learning in biology due to implausible assumptions such as the fact that forward and backward passes use the same set of synaptic weights. There are a number of responses here. First, one might hold the view that backpropagation is just an efficient way to obtain effective network architectures, without committing to the biological plausibility of the learning algorithm *per se*. Second, if biologically plausible learning is the research objective then one is free to exploit other (Hebbian) learning schemes that may reflect biological learning more closely (Miconi, 2017). Finally, researchers have started to put forward arguments that backpropagation may not be that biologically implausible after all (Roelfsema and van Ooyen, 2005; Lillicrap et al., 2016; Scellier and Bengio, 2017).

4.2. Perceiving

One of the core skills any intelligent agent should possess is the ability to recognize patterns in its environment. The world around us consists of various objects that may carry significance. Being able to recognize edible food, places that provide shelter, and other agents will all aid survival.

Biological agents are faced with the problem that they need to be able to recognize objects from raw sensory input (vectors in $\mathbb{R}^n$). How can a brain use the incident sensory input to learn to recognize those things that are of relevance to the organism? Recall the artificial neuron formulation $y = f(\mathbf{w}^\mathsf{T}\mathbf{x})$. By learning proper weights $\mathbf{w}$, this neuron can learn to distinguish different object categories. This is essentially equivalent to a classical model known as the perceptron (Rosenblatt, 1958), which was used to solve simple pattern recognition problems via a simple error-correction mechanism. It also corresponds to a basic linear-nonlinear (LN) model which has been used extensively to model and estimate the receptive field of a neuron or a population of neurons (van Gerven, 2017).

[8]In practice, it is more efficient to iterate over subsets of datapoints, known as mini-batches, in sequence. That is, training is organized in terms of epochs in which all datapoints are processed by iterating over mini-batches. Note that, whenever we are not processing all data points in parallel, we are not exactly following the gradient. Therefore, any such procedure is known as *stochastic gradient descent*.

Single-layer ANNs such as the perceptron are capable of solving interesting learning problems. At the same time, they are limited in scope since they can only solve linearly separable classification problems (Minsky and Papert, 1969). To overcome the limitations of the perceptron we can extend its capabilities by relaxing the constraint that the inputs are directly coupled to the outputs. A multilayer perceptron (MLP) is a feedforward network which generalizes the standard perceptron by having a hidden layer that resides between the input and the output layers. We can write an MLP with multiple output units as

$$\mathbf{y} = \mathbf{g}\left(\mathbf{W}\mathbf{f}\left(\mathbf{V}\mathbf{x}\right)\right) \tag{7}$$

where $\mathbf{V}$ denotes the hidden layer weights and $\mathbf{W}$ denotes the output layer weights. By introducing a hidden layer, MLPs gain the ability to learn internal representations (Rumelhart et al., 1986). Importantly, an MLP can approximate any continuous function to an arbitrary degree of accuracy, given a sufficiently large but finite number of hidden neurons (Cybenko, 1989; Hornik, 1991).

Complex systems tend to be hierarchical and modular in nature (Simon, 1962). The nervous system itself can be thought of as a hierarchically organized system. This is exemplified by Felleman & van Essen's hierarchical diagram of visual cortex (Felleman and Van Essen, 1991), the proposed hierarchical organization of prefrontal cortex (Badre, 2008), the view of the motor system as a behavioral control column (Swanson, 2000) and the proposition that anterior and posterior cortex reflect hierarchically organized executive and perceptual systems (Fuster, 2001). Representations at the top of these hierarchies correspond to highly abstract statistical invariances that occupy our ecological niche (Quian Quiroga et al., 2005; Barlow, 2009). A hierarchy can be modeled by a deep neural network (DNN) composed of multiple hidden layers (LeCun et al., 2015), written as

$$\begin{aligned}\mathbf{y} &= \mathbf{f}_{L+1}\big(\mathbf{W}_{L+1}\mathbf{f}_L\big(\mathbf{W}_L \cdots \mathbf{f}_1\big(\mathbf{W}_1\mathbf{x}\big)\cdots\big)\\ &= \mathbf{f}_{\boldsymbol{\theta}}(\mathbf{x})\end{aligned} \tag{8}$$

where $\mathbf{W}_l$ is the weight matrix associated with layer l. Even though an MLP can already approximate any function to an arbitrary degree of precision, it has been shown that many classes of functions can be represented much more compactly using thin and deep neural networks compared to shallow and wide neural networks (Bengio and LeCun, 2007; Bengio, 2009; Le Roux and Bengio, 2010; Delalleau and Bengio, 2011; Mhaskar et al., 2016).

A DNN corresponds to a stack of LN models, generalizing the concept of basic receptive field models. They have been shown to yield human-level performance on object categorization tasks (Krizhevsky et al., 2012). The latest DNN incarnations are even capable of predicting the cognitive states of other agents. One example is the prediction of apparent personality traits from multimodal sensory input (Güçlütürk et al., 2016). Deep architectures have been used extensively in neuroscience to model hierarchical processing (Selfridge, 1959; Fukushima, 1980, 2013; Riesenhuber and Poggio, 1999; Lehky and Tanaka, 2016). Interestingly, it has been shown that the representations

encoded in DNN layers correspond to the representations that are learned by areas that make up the sensory hierarchies of biological agents (Güçlü and van Gerven, 2015, 2017a; Güçlü et al., 2016). Multiple reviews discuss this use of DNNs in sensory neuroscience (Cox and Dean, 2014; Kriegeskorte, 2015; Robinson and Rolls, 2015; Marblestone et al., 2016; Yamins and DiCarlo, 2016; Kietzmann et al., 2017; Peelen and Downing, 2017; van Gerven, 2017; Vanrullen, 2017).

4.3. Remembering

Being able to perceive the environment also implies that agents can store and retrieve past knowledge about objects and events in their surroundings. In the feedforward networks considered in the previous section, this knowledge is encoded in the synaptic weights as a result of learning. Memories of the past can also be stored, however, in moment-to-moment neural activity patterns. This does require the availability of lateral or feedback connections in order to enable recurrent processing (Singer, 2013; Maass, 2016). Recurrent processing can be implemented by a recurrent neural network (RNN) (Jordan, 1987; Elman, 1990), defined by

$$\mathbf{y}_n = \mathbf{f}\left(\mathbf{W}\mathbf{y}_{n-1} + \mathbf{U}\mathbf{x}_n\right) \tag{9}$$

such that the neuronal activity at time n depends on the activity at time $n-1$ as well as instantaneous bottom-up input. RNNs can be interpreted as numerical approximations of differential equations that describe rate-based neural models (Dayan and Abbott, 2005) and have been shown to be universal approximators of dynamical systems (Funahashi and Nakamura, 1993)[9]. Their parameters can be estimated using a variant of backpropagation, referred to as backpropagation through time (Mozer, 1989).

When considering perception, feedforward architectures may seem sufficient. For example, the onset latencies of neurons in monkey inferior-temporal cortex during visual processing are about 100 ms (Thorpe and Fabre-Thorpe, 2001), which means that there is ample time for the transmission of just a few spikes. This suggests that object recognition is largely an automatic feedforward process (Vanrullen, 2007). However, recurrent processing is important in perception as well since it provides the ability to maintain state. This is important in detecting salient features in space and time (Joukes et al., 2014), as well as for integrating evidence in noisy or ambiguous settings (O'Reilly et al., 2013). Moreover, perception is strongly influenced by top-down processes, as mediated by feedback connections (Gilbert and Li, 2013). RNNs have also been used to model working memory (Miconi, 2017) as well as hippocampal function, which is involved in a variety of memory-related processes (Willshaw et al., 2015; Kumaran et al., 2016).

A special kind of RNN is the Hopfield network (Hopfield, 1982), where $\mathbf{W}$ is symmetric and $\mathbf{U} = \mathbf{0}$. Learning in a Hopfield net is based on a Hebbian learning scheme. Hopfield nets are attractor networks that converge to a state that is a local minimum of an energy function. They have been used extensively as models of associative memory (Wills et al., 2005). It has even been postulated that dreaming can be seen as an unlearning process which gets rid of spurious minima in attractor networks, thereby improving their storage capacity (Crick and Mitchison, 1983).

[9] The ability of simple RNNs to integrate information over time remains limited, which led to the introduction of various extensions that perform more favorably in this regard (Hochreiter and Schmidhuber, 1997; Cho et al., 2014; Neil et al., 2016; Wu et al., 2016).

4.4. Acting

As already described, the ability to generate appropriate actions is what ultimately drives behavior. In real-world settings, such actions typically need to be inferred from reward signals r_t provided by the environment. This is the subject matter of reinforcement learning (RL) (Sutton and Barto, 1998). Define a policy $\pi(s,a)$ as the probability of selecting an action a given a state s. Let the return $R = \sum_{t=0}^{\infty} \gamma^t r_{t+1}$ be the total reward accumulated in an episode, with γ a discount factor that downweighs future rewards. The goal in RL is to identify an optimal policy π^* that maximizes the expected return

$$\pi^* = \arg\max_{\pi} \mathbb{E}[R \mid \pi]\,. \tag{10}$$

Reinforcement learning algorithms have been crucial in training neural networks that have the capacity to act. Such networks learn to generate suitable actions purely by observing the rewards entailed by previously generated actions. RL algorithms come in model-free and model-based variants. In the model-free setting, optimal actions are learned purely based on the reward that is gained by performing actions in the past. In the model-based setting, in contrast, an explicit model of the environment is used to predict the consequences of actions that are being executed. Importantly, model-free and model-based reinforcement learning approaches have clear correspondences with habitual and goal-directed learning in neuroscience (Daw, 2012; Buschman et al., 2014).

Various model-free reinforcement learning approaches have been used to develop a variety of neural networks for action generation. For example, Q-learning was used to train networks that play Atari games (Mnih et al., 2015) and policy gradient methods have been used to play board games (Silver et al., 2017) and solve problems in (simulated) robotics (Silver et al., 2014; Schulman et al., 2015), effectively closing the perception-action cycle. Evolutionary strategies are also proving to become an useful approach for solving challenging control problems (Salimans et al., 2017). Similar successes have been achieved using model-based reinforcement learning approaches (Schmidhuber, 2015; Mujika, 2016; Santana and Hotz, 2016).

Another important ingredient required for generating optimal actions is recurrent processing, as described in the previous section. Action generation must depend on the ability to integrate evidence over time since, otherwise, we are guaranteed to act suboptimally. That is, states that are qualitatively different can appear the same to the decision maker, leading to suboptimal policies. Consider for example the sensation of a looming object. The optimal decision depends crucially on whether this object is approaching or receding, which can only be determined by

taking past sensations into account. This phenomenon is known as perceptual aliasing (Whitehead and Ballard, 1991).

A key ability of biological organisms which requires recurrent processing is their ability to navigate in their environment, as mediated by the hippocampal formation (Moser et al., 2015). Recent work shows that particular characteristics of hippocampal place cells, such as stable tuning curves that remap between environments, are recovered by training neural networks on navigation tasks (Kanitscheider and Fiete, 2016). The ability to integrate evidence also allows agents to selectively sample the environment, such as to maximize the amount of information gained. This process, known as active sensing, is crucial for understanding perceptual processing in biology (Yarbus, 1967; Regan and Noë, 2001; Friston et al., 2010; Schroeder et al., 2010; Gordon and Ahissar, 2012). Active sensing, in the form of saccade planning, has been implemented using a variety of recurrent neural network architectures (Larochelle and Hinton, 2010; Gregor et al., 2014; Mnih et al., 2014). RNNs that implement recurrent processing have also been used to model various other action-related processes such as timing (Laje and Buonomano, 2013), sequence generation (Rajan et al., 2015) and motor control (Sussillo et al., 2015).

Recurrent processing and reinforcement learning are also essential in modeling higher-level processes, such as cognitive control as mediated by frontal brain regions (Fuster, 2001; Miller and Cohen, 2001). Examples are models of context-dependent processing (Mante et al., 2013) and perceptual decision-making (Carnevale et al., 2015). In general, RNNs that have been trained using RL on a variety of cognitive tasks have been shown to yield properties that are consistent with phenomena observed in biological neural networks (Song et al., 2016; Miconi, 2017).

4.5. Predicting

Modern theories of human brain function appeal to the idea that the brain can be viewed as a prediction machine, which is in the business of continuously generating top-down predictions that are integrated with bottom-up sensory input (Lee and Mumford, 2003; Yuille and Kersten, 2006; Clark, 2013; Summerfield and de Lange, 2014). This view of the brain as a prediction machine that performs unconscious inference has a long history, going back to the seminal work of Alhazen and Helmholtz (Hatfield, 2002). Modern views cast this process in terms of Bayesian inference, where the brain is updating its internal model of the environment in order to explain away the data that impinge upon its senses, also referred to as the Bayesian brain hypothesis (Jaynes, 1988; Doya et al., 2006). The same reasoning underlies the free-energy principle, which assumes that biological systems minimize a free energy functional of their internal states that entail beliefs about hidden states in their environment (Friston, 2010). Predictions can be seen as central to the generation of adaptive behavior, since anticipating the future will allow an agent to select appropriate actions in the present (Schacter et al., 2007; Moulton and Kosslyn, 2009).

Prediction is central in model-based RL approaches since it requires agents to plan their actions by predicting the outcomes of future actions (Daw, 2012). This is strongly related to the notion of preplay of future events subserving path planning (Corneil and Gerstner, 2015). Such preplay has been observed in hippocampal place cell sequences (Dragoi and Tonegawa, 2011), giving further support to the idea that the hippocampal formation is involved in goal-directed navigation (Corneil and Gerstner, 2015). Prediction also allows an agent to prospectively act on expected deviations from optimal conditions. This focus on error-correction and stability is also prevalent in the work of the cybernetic movement (Ashby, 1952). Note further that predictive processing connects to the concept of allostasis, where the agent is actively trying to predict future states such as to minimize deviations from optimal homeostatic conditions. It is also central to optimal feedback control theory, which assumes that the motor system corrects only those deviations that interfere with task goals (Todorov and Jordan, 2002).

The notion of predictive processing has been very influential in neural network research. For example, it provides the basis for predictive coding models that introduce specific neural network architectures in which feedforward connections are used to transmit the prediction errors that result from discrepancies between top-down predictions and bottom-up sensations (Rao and Ballard, 1999; Huang and Rao, 2011). It also led to the development of a wide variety of generative models that are able to predict their sensory states, also referred to as fantasies (Hinton, 2013). Such fantasies may play a role in understanding cognitive processing involved in imagery, working memory and dreaming. In effect, these models aim to estimate a distribution over latent causes $\mathbf{z}$ in the environment that explain observed sensory data $\mathbf{x}$. In this setting, the most probable explanation is given by

$$\begin{aligned} \mathbf{z}^* &= \arg\max_{\mathbf{z}} p(\mathbf{z} \mid \mathbf{x}) \\ &= \arg\max_{\mathbf{z}} \big[p(\mathbf{x} \mid \mathbf{z}) p(\mathbf{z}) \big] . \end{aligned} \tag{11}$$

Generative models also offer a way to perform unsupervised learning, since if a neural network is able to generate predictions then the discrepancy between predicted and observed stimuli can serve as a teaching signal. A canonical example is the Boltzmann machine, which is a stochastic variant of a Hopfield network that is able to discover regularities in the training data using a simple unsupervised learning algorithm (Hinton and Sejnowski, 1983; Ackley et al., 1985). Another classical example is the Helmholtz machine, which incorporates both bottom-up and top-down processing (Dayan et al., 1995). Other, more recent examples of ANN-based generative models are deep belief networks (Hinton et al., 2006), variational autoencoders (Kingma and Welling, 2014) and generative adversarial networks (Goodfellow et al., 2014). Recent work has started to use these models to predict future sensory states from current observations (Lotter et al., 2016; Mathieu et al., 2016; Xue et al., 2016).

4.6. Reasoning

While ANNs are now able to solve complex tasks such as acting in natural environments or playing difficult board games, one could still argue that they are "just" performing sophisticated pattern

recognition rather than showing the symbolic reasoning abilities that characterize our own brains. The question of whether connectionist systems are capable of symbolic reasoning has a long history, and has been debated by various researchers in the cognitivist (symbolic) program (Pinker and Mehler, 1988). We will not settle this debate here but point out that efforts are underway to endow neural networks with sophisticated reasoning capabilities.

One example is the development of "differentiable computers" that learn to implement algorithms based on a finite amount of training data (Graves et al., 2014; Weston et al., 2015; Vinyals et al., 2017). The resulting neural networks perform variable binding and are able to deal with variable length structures (Graves et al., 2014), which are two objections that were originally raised against using ANNs to explain cognitive processing (Fodor and Pylyshyn, 1988).

Another example is the development of neural networks that can answer arbitrary questions about text (Bordes et al., 2015), images (Agrawal et al., 2016) and movies (Tapaswi et al., 2015), thereby requiring deep semantic knowledge about the experienced stimuli. Recent models have also been shown to be capable of compositional reasoning (Johnson et al., 2017; Lake et al., 2017; Yang et al., 2017), which is an important ingredient for explaining the systematic nature of human thought (Fodor and Pylyshyn, 1988). These architectures often make use of distributional semantics, where words are encoded as real vectors that capture word meaning (Mikolov et al., 2013; Ferrone and Zanzotto, 2017).

Several other properties characterize human thought processes, such as intuitive physics, intuitive psychology, relational reasoning and causal reasoning (Kemp and Tenenbaum, 2008; Lake et al., 2017). Another crucial hallmark of intelligent systems is that they are able to explain what they are doing (Brachman, 2002). This requires agents to have a deep understanding of their world. These properties should be replicated in neural networks if they are to serve as accurate models of natural intelligence. New neural network architectures are slowly starting to take steps in this direction (e.g., Louizos et al., 2017; Santoro et al., 2017; Zhu et al., 2017).

5. TOWARD STRONG AI

We have reviewed the computational foundations of natural intelligence and outlined how ANNs can be used to model a variety of cognitive processes. However, our current understanding of natural intelligence remains limited and strong AI has not yet been attained. In the following, we will touch upon a number of important topics that will be of importance for eventually reaching these goals.

5.1. Surviving in Complex Environments

Contemporary neural network architectures tend to excel at solving one particular problem well. However, in practice, we want to arrive at intelligent machines that are able to survive in complex environments. This requires the agent to deal with high-dimensional naturalistic input, be able to solve multiple tasks depending on context, and devise optimal strategies to ensure long-term survival.

The research community has embraced these desiderata by creating virtual worlds that allow development and testing of neural network architectures (e.g., Todorov et al., 2012; Beattie et al., 2016; Brockman et al., 2016; Kempka et al., 2016; Synnaeve et al., 2016)[10]. While most work in this area has focused on environments with fully observable states, reward functions with low delay, and small action sets, research is shifting toward environments that are partially observable, require long-term planning, show complex dynamics and have noisy and high-dimensional control interfaces (Synnaeve et al., 2016).

A particular challenge in these naturalistic environments is that networks need to be able to exhibit continual (life-long) learning (Thrun and Mitchell, 1995), adapting continuously to the current state of affairs. This is difficult due to the phenomenon of catastrophic forgetting (McCloskey and Cohen, 1989; French, 1999), where previously acquired skills are overwritten by ongoing modification of synaptic weights. Recent algorithmic developments attenuate the detrimental effects of catastrophic forgetting (Kirkpatrick et al., 2015; Zenke et al., 2015), offering a (partial) solution to the stability vs. plasticity dilemma (Abraham and Robins, 2005). Life-long learning is further complicated by the exploration-exploitation dilemma, where agents need to decide on whether to accrue either information or reward (Cohen et al., 2007). Another challenge is the fact that reinforcement learning of complex actions is notoriously slow. Here, progress is being made using networks that make use of differentiable memories (Santoro et al., 2016; Pritzel et al., 2017). Survival in complex environments also requires that agents learn to perform multiple tasks well. This learning process can be facilitated through multitask learning (Caruana, 1997) (also referred to as learning to learn Baxter, 1998 or transfer learning Pan and Fellow, 2009), where learning of one task is facilitated by knowledge gained through learning to solve another task. Multitask learning has been shown to improve convergence speed and generalization to unseen data (Scholte et al., 2017). Finally, effective learning also calls for agents that can generalize to cases that were not encountered before, which is known as zero-shot learning (Palatucci et al., 2009), and can learn from rare events, which is known as one-shot learning (Fei-Fei et al., 2006; Vinyals et al., 2016; Kaiser and Roy, 2017).

While the use of virtual worlds allows for testing the capabilities of artificial agents, it does not guarantee that the same agents are able to survive in the real world (Brooks, 1992). That is, there may exist a reality gap, where skills acquired in virtual worlds do not carry over to the real world. In contrast to virtual worlds, acting in the real world requires the agent to deal with unforeseen circumstances resulting from the complex nature of reality, the agent's need for a physical body, as well as its engagement with a myriad of other agents (Anderson, 2003). Moreover, the continuing interplay between an organism and its environment may itself shape and, ultimately, determine cognition (Gibson, 1979; Maturana

[10] See SHRDLU for an early example of such a virtual world (Winograd, 1972).

and Varela, 1987; Brooks, 1996; Edelman, 2015). Effectively dealing with these complexities may not only require plasticity in individual agents but also the incorporation of developmental change, as well as learning at evolutionary time scales (Marcus, 2009). From a developmental perspective, networks can be more effectively trained by presenting them with a sequence of increasingly complex tasks, instead of immediately requiring the network to solve the most complex task (Elman, 1993). This process is known as curriculum learning (Bengio et al., 2009) and is analogous to how a child learns by decomposing problems into simpler subproblems (Turing, 1950). Evolutionary strategies have also been shown to be effective in learning to solve challenging control problems (Salimans et al., 2017). Finally, to learn about the world, we may also turn toward cultural learning, where agents can offload task complexity by learning from each other (Bengio, 2014).

As mentioned in section 2.2, adaptive behavior is the result of multiple competing drives and motivations that provide primary, intrinsic and extrinsic rewards. Hence, one strategy for endowing machines with the capacity to survive in the real world is to equip neural networks with drives and motivations that ensure their long-term survival[11]. In terms of primary rewards, one could conceivably provide artificial agents with the incentive to minimize computational resources or maximize offspring via evolutionary processes (Stanley and Miikkulainen, 2002; Floreano et al., 2008; Gauci and Stanley, 2010). In terms of intrinsic rewards, one can think of various ways to equip agents with the drive to explore the environment (Oudeyer, 2007). We briefly describe a number of principles that have been proposed in the literature. Artificial curiosity assumes that internal reward depends on how boring an environment is, with agents avoiding fully predictable and unpredictably random states (Schmidhuber, 1991, 2003; Pathak et al., 2017). A related notion is that of information-seeking agents (Bachman et al., 2016). The autotelic principle formalizes the concept of flow where an agent tries to maintain a state where learning is challenging, but not overwhelming (Csikszentmihalyi, 1975; Steels, 2004). The free-energy principle states that an agent seeks to minimize uncertainty by updating its internal model of the environment and selecting uncertainty-reducing actions (Friston, 2009, 2010). Empowerment is founded on information-theoretic principles and quantifies how much control an agent has over its environment, as well as its ability to sense this control (Klyubin et al., 2005a,b; Salge et al., 2013). In this setting, intrinsically motivated behavior is induced by the maximization of empowerment. Finally, various theories embrace the notion that optimal prediction of future states drives learning and behavior (Der et al., 1999; Kaplan and Oudeyer, 2004; Ay et al., 2008). In terms of extrinsic rewards, one can think of imitation learning, where a teacher signal is used to inform the agent about its desired outputs (Schaal, 1999; Duan et al., 2017).

[11]The notion of *wanting* agents was already present in the writings of Thurstone (1923), who wrote: "My main thesis is that conduct originates in the organism itself and not in the environment in the form of a stimulus. [...] All mental life may be looked upon as incomplete behavior which is in the process of being formed. [...] Perception is the discovery of the suitable stimulus which is often anticipated imaginally. The appearance of the stimulus is one of the last events in the expression of impulses in conduct. The stimulus is not the starting point for behavior."

5.2. Bridging the Gap between Artificial and Biological Neural Networks

To reduce the gap between artificial and biological neural networks, it makes sense to assess their operation on similar tasks. This can be done either by comparing the models at a neurobiological level or at a behavioral level. The former refers to comparing the internal structure or activation patterns of artificial and biological neural networks. The latter refers to comparing their behavioral outputs (e.g., eye movements, reaction times, high-level decisions). Moreover, comparisons can be made under changing conditions, i.e., during learning and development (Elman et al., 1996). As such, ANNs can serve as explanatory mechanisms in cognitive neuroscience and behavioral psychology, embracing recent model-based approaches (Forstmann and Wagenmakers, 2015).

From a psychological perspective, ANNs have been compared explicitly with their biological counterparts. Connectionist models were widely used in the 1980's to explain various psychological phenomena, particularly by the parallel distributed processing (PDP) movement, which stressed the parallel nature of neural processing and the distributed nature of neural representations (McClelland, 2003). For example, neural networks have been used to explain grammar acquisition (Elman, 1991), category learning (Kruschke, 1992) and the organization of the semantic system (Ritter and Kohonen, 1989). More recently, deep neural networks have been used to explain human similarity judgments (Peterson et al., 2016). With new developments in cognitive and affective computing, where neural networks become more adept at solving high-level cognitive tasks, such as predicting people's (apparent) personality traits (Güçlütürk et al., 2016), their use as a tool to explain psychological phenomena is likely to increase. This will also require embracing insights about how humans solve problems at a cognitive level (Tenenbaum et al., 2011).

ANNs have also been related explicitly to brain function. For example, the perceptron has been used in the modeling of various neuronal systems, including sensorimotor learning in the cerebellum (Marr, 1969) and associative memory in cortex (Gardner, 1988), sparse coding has been used to explain receptive field properties (Olshausen and Field, 1996), topographic maps have been used to explain the formation of cortical maps (Obermayer, 1990; Aflalo, 2006), Hebbian learning has been used to explain neural tuning to face orientation (Leibo et al., 2017), and networks trained by backpropagation have been used to model the response properties of posterior parietal neurons (Zipser and Andersen, 1988). Neural networks have also been used to model central pattern generators that drive behavior (Duysens and Van de Crommert, 1998; Ijspeert, 2008) as well as the perception of rhythmic stimuli (Torras i Genís, 1986; Gasser, Eck and Port, 1999). Furthermore, reinforcement learning algorithms used to train neural networks for action selection have strong ties with the brain's reward system (Schultz

et al., 1997; Sutton and Barto, 1998). It has been shown that RNNs trained to solve a variety of cognitive tasks using reinforcement learning replicate various phenomena observed in biological systems (Song et al., 2016; Miconi, 2017). Crucially, these efforts go beyond descriptive approaches in that they may explain *why* the human brain is organized in a certain manner (Barak, 2017).

Rather than using neural networks to explain certain observed neural or behavioral phenomena, one can also directly fit neural networks to neurobehavioral data. This can be achieved via an indirect approach or via a direct approach. In the *indirect* approach, neural networks are first trained to solve a task of interest. Subsequently, the trained network's responses are fitted to neurobehavioral data obtained as participants engage in the same task. Using this approach, deep convolutional neural networks trained on object recognition, action recognition and music tagging have been used to explain the functional organization of visual as well as auditory cortex (Güçlü and van Gerven, 2015, 2017a; Güçlü et al., 2016). The indirect approach has also been used to train RNNs via reinforcement learning on a probabilistic categorization task. These networks have been used to fit the learning trajectories and behavioral responses of humans engaged in the same task (Bosch et al., 2016). Mante et al. (2013) used RNNs to model the population dynamics of single neurons in prefrontal cortex during a context-dependent choice task. In the *direct* approach, neural networks are trained to directly predict neural responses. For example, Mcintosh et al. (2016) trained convolutional neural networks to predict retinal responses to natural scenes, Joukes et al. (2014) trained RNNs to predict neural responses to motion stimuli, and Güçlü and van Gerven (2017b) used RNNs to predict cortical responses to naturalistic video clips. This ability of neural networks to explain neural recordings is expected to become increasingly important (Sompolinsky, 2014; Marder, 2015), given the emergence of new imaging technology where the activity of thousands of neurons can be measured in parallel (Ahrens et al., 2013; Churchland and Sejnowski, 2016; Lopez et al., 2016; Pachitariu et al., 2016; Yang and Yuste, 2017). Better understanding will also be facilitated by the development of new data analysis techniques to elucidate human brain function (Kass et al., 2014)[12], the use of ANNs to decode neural representations (Schoenmakers et al., 2013; Güçlütürk et al., 2017), as well as the development of approaches that elucidate the functioning of ANNs (e.g., Nguyen et al., 2016; Kindermans et al., 2017; Miller, 2017)[13].

5.3. Next-Generation Artificial Neural Networks

The previous sections outlined how neural networks can be made to solve challenging tasks and provide explanations of neural and behavioral responses in biological agents. In this final section, we consider some developments that are expected to fuel the next generation of ANNs.

First, a major driving force in neural network research will be theoretical and algorithmic developments that inform why ANNs work so well in practice, what their fundamental limitations are, as well as how to overcome these. From a theoretical point of view, substantial advances have already been made pertaining to, for example, understanding the nature of representations (Anselmi and Poggio, 2014; Lin and Tegmark, 2016; Shwartz-Ziv and Tishby, 2017), the statistical mechanics of neural networks (Sompolinsky, 1988; Advani et al., 2013), as well as the expressiveness (Pascanu et al., 2013; Bianchini and Scarselli, 2014; Kadmon and Sompolinsky, 2016; Mhaskar et al., 2016; Poole et al., 2016; Raghu et al., 2016; Weichwald et al., 2016), generalizability (Kawaguchi et al., 2017) and learnability (Dauphin et al., 2014; Saxe et al., 2014; Schoenholz et al., 2017) of DNNs.

From an algorithmic point of view, great strides have been made in improving training of deep (Srivastava et al., 2014; He et al., 2015; Ioffe and Szegedy, 2015) and recurrent neural networks (Hochreiter and Schmidhuber, 1997; Pascanu et al., 2012), overcoming the reality gap (Tobin et al., 2017), adding modularity to neural networks (Fernando et al., 2017), as well as on improving the efficacy of reinforcement learning algorithms (Schulman et al., 2015; Mnih et al., 2016; Pritzel et al., 2017).

Second, it is expected that as neural network models become more plausible from a biological point of view, model fit and task performance will further improve (Cox and Dean, 2014). This is important in driving new developments in model-based cognitive neuroscience but also in developing intelligent machines that show human-like behavior. One example is to match the object recognition capabilities of extremely deep neural networks with more biologically plausible RNNs of limited depth (O'Reilly et al., 2013; Liao and Poggio, 2016) and achieving category selectivity in a more realistic manner (Peelen and Downing, 2017; Scholte et al., 2017). Another example is to incorporate predictive coding principles in neural network architectures (Lotter et al., 2016). Furthermore, more human-like perceptual systems can be arrived at by including attentional mechanisms (Mnih et al., 2014) as well as mechanisms for saccade planning (Najemnik and Geisler, 2005; Larochelle and Hinton, 2010; Gregor et al., 2014).

In general, ANN research can benefit from a close interaction between the AI and neuroscience communities (Yuste, 2015; Hassabis et al., 2017). For example, neural network research may be shaped by general guiding principles of brain function at different levels of analysis (O'Reilly, 1998; Maass, 2016; Sterling and Laughlin, 2016). We may also strive to incorporate more biological detail. For example, to obtain accurate models of neural information processing we may need to embrace spike-based rather than rate-based neural networks (Brette, 2015)[14]. Efforts are underway to effectively train spiking neural

[12] But see Jonas and Kording (2017) for a critical appraisal of the informativeness of such techniques.

[13] These techniques aim to overcome the interpretability problem raised by Mozer and Smolensky (1989), who state: "One thing that connectionist networks have in common with brains is that if you open them up and peer inside, all you can see is a big pile of goo."

[14] While there surely exists neurobiological evidence for temporal coding with spikes (Segundo et al., 1966; Barrio and Buno, 1990; Bohte, 2004), it remains an

networks (Maass, 1997; Gerstner and Kistler, 2002; Gerstner et al., 2014; O'Connor and Welling, 2016; Huh and Sejnowski, 2017) and endow them with the same cognitive capabilities as their rate-based cousins (Thalmeier et al., 2015; Abbott et al., 2016; Kheradpisheh et al., 2016; Lee et al., 2016; Zambrano and Bohte, 2016).

In the same vein, researchers are exploring how probabilistic computations can be performed in neural networks (Nessler et al., 2013; Pouget et al., 2013; Gal, 2016; Orhan and Ma, 2016; Ambrogioni et al., 2017; Heeger, 2017; Mandt et al., 2017) and deriving new biologically plausible synaptic plasticity rules (Brea and Gerstner, 2016; Brea et al., 2016; Schiess et al., 2016). Biologically-inspired principles may also be incorporated at a more conceptual level. For instance, researchers have shown that neural networks can be protected from adversarial attacks (i.e., the construction of stimuli that cause networks to make mistakes) by integrating the notion of nonlinear computations encountered in the branched dendritic structures of real neurons (Nayebi and Ganguli, 2016).

Finally, research is invested in implementing ANNs in hardware, also referred to as neuromorphic computing (Mead, 1990). These brain-based parallel chip architectures hold the promise of devices that operate in real time and with very low power consumption (Schuman et al., 2017), driving new advances in cognitive computing (Modha et al., 2011; Neftci et al., 2013; Van de Burgt et al., 2017). On a related note, nanotechnology may 1 day drive the development of new neural network architectures whose operation is closer to the molecular machines that mediate the operation of biological neural networks (Drexler, 1992; Strukov, 2011). In the words of Feynman (1992): "There's plenty of room at the bottom."

6. CONCLUSION

As cognitive scientists, we live in exciting times. Cognitivism offers an interpretation of agents as information processing systems that are engaged in formal symbol manipulation. The probabilistic approach to cognition extends this interpretation by viewing organisms as rational agents that need to act in the face of uncertainty under limited resources. Finally, emergentist approaches such as artificial life and connectionism indicate that concerted interactions between simple processing elements can achieve human-level performance at certain cognitive tasks. While these different views have stirred substantial debate in the past, they need not be irreconcilable. Surely we are capable of formal symbol manipulation and decision making under uncertainty in real-life settings. At the same time, these capabilities must be implemented by the neural circuits that make up our own brains, which themselves rely on noisy long-range communication between neuronal populations.

The thesis of this paper is that natural intelligence can be modeled and understood by constructing artificial agents whose synthetic brains are composed of (rate-based) neural networks. To act as explanations of natural intelligence, these synthetic brains should show a functional correspondence with their biological counterparts. To identify such correspondence we can embrace the rich sources of data provided by biology, neuroscience and psychology, providing a link to Marr's implementational level. At the same time, we can use sophisticated machinery developed in mathematics, computer science and physics to gain a better understanding of these systems. Ultimately, these synthetic brains should be able to show the capabilities that are prescribed by normative theories of intelligent behavior, providing a link to Marr's computational level.

The supposition that artificial neural networks are sufficient for modeling all of cognition may seem premature. For example, state-of-the-art question-answering systems such as IBM's Watson (Ferrucci et al., 2010) use ANN technology as a minor component within a larger (symbolic) framework and the AlphaGo system (Silver et al., 2017), which learns to play the game of Go beyond grandmaster level without any human intervention, combines neural networks with Monte Carlo tree search. While it is true that ANNs remain wanting when it comes to logical reasoning, inferring causal relationships or planning, the pace of current research may very well bring these capabilities within reach in the foreseeable future. Such neural networks may turn out to be quite different from current neural network architectures and their operation may be guided by complementary yet-to-be-discovered learning rules.

The quest for natural intelligence can be contrasted with a pure engineering approach. From an engineering perspective, understanding natural intelligence may be considered irrelevant since the main interest is in building devices that do the job. To quote Edsger Dijkstra, "the question whether machines can think [is] as relevant as the question whether submarines can swim." At the same time, our quest for natural intelligence may facilitate the development of strong AI given the proven ability of our own brains to generate intelligent behavior. Hence, biologically inspired architectures may not only provide new insights into human brain function but could also in the long run yield superior curious and perhaps even conscious machines that surpass humans in terms of intelligence, creativity, playfulness, and empathy (Boden, 1998; Moravec, 2000; Der and Martius, 2011; Modha et al., 2011; Harari, 2017).

AUTHOR CONTRIBUTIONS

The author confirms being the sole contributor of this work and approved it for publication.

ACKNOWLEDGMENTS

This work is supported by a VIDI grant (639.072.513) from the Netherlands Organization for Scientific Research. I would like to thank Nadine Dijkstra, Gabriëlle Ras, Andrew Reid, Katja Seeliger and the reviewers for their useful comments.

open question if temporal coding is absolutely necessary for the generation of adaptive behavior. In the end, computing with spikes may have emerged chiefly to promote efficiency and allow long-distance neuronal communication (Laughlin and Sejnowski, 2003).

REFERENCES

Abbott, L. F., Depasquale, B., and Memmesheimer, R.-M. (2016). Building functional networks of spiking model neurons. *Nat. Neurosci.* 19, 350–355. doi: 10.1038/nn.4241

Abraham, W. C., and Robins, A. (2005). Memory retention - the synaptic stability versus plasticity dilemma. *Trends Neurosci.* 28, 73–78. doi: 10.1016/j.tins.2004.12.003

Ackley, D., Hinton, G. E., and Sejnowski, T. (1985). A learning algorithm for Boltzmann machines. *Cogn. Sci.* 9, 147–169. doi: 10.1016/S0364-0213(85)80012-4

Adams, S. S., Arel, I., Bach, J., Coop, R., Furlan, R., Goertzel, B., et al. (2012). Mapping the landscape of human-level artificial general intelligence. *AI Mag.* 33, 25–42. doi: 10.1609/aimag.v33i1.2322

Advani, M., Lahiri, S., and Ganguli, S. (2013). Statistical mechanics of complex neural systems and high dimensional data. *J. Stat. Mech. Theory Exp.* 2013:P03014. doi: 10.1088/1742-5468/2013/03/P03014

Aflalo, T. N. (2006). Possible origins of the complex topographic organization of motor cortex: reduction of a multidimensional space onto a two-dimensional array *J. Neurosci.* 26, 6288–6297. doi: 10.1523/JNEUROSCI.0768-06.2006

Agrawal, A., Lu, J., Antol, S., Mitchell, M., Zitnick, C. L., Batra, D., et al. (2016). VQA: visual question answering. ArXiv:1505.00468, 1–25.

Ahrens, M. B., Orger, M. B., Robson, D. N., Li, J. M., and Keller, P. J. (2013). Whole-brain functional imaging at cellular resolution using light-sheet microscopy. *Nat. Methods* 10, 413–420. doi: 10.1038/nmeth.2434

Ambrogioni, L., Güçlü, U., Maris, E., and van Gerven, M. A. J. (2017). Estimating nonlinear dynamics with the ConvNet smoother. ArXiv:1702.05243, 1–8.

Amunts, K., Ebell, C., Muller, J., Telefont, M., Knoll, A., and Lippert, T. (2016). The Human Brain Project: creating a European research infrastructure to decode the human brain. *Neuron* 92, 574–581. doi: 10.1016/j.neuron.2016.10.046

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., and Qin, Y. (2004). An integrated theory of the mind. *Psychol. Rev.* 111, 1036–1060. doi: 10.1037/0033-295X.111.4.1036

Anderson, M. L. (2003). Embodied cognition: a field guide. *Artif. Intell.* 149, 91–130. doi: 10.1016/S0004-3702(03)00054-7

Andrieu, C., De Freitas, N., Doucet, A., and Jordan, M. I. (2003). An introduction to MCMC for machine learning. *Mach. Learn.* 50, 5–43. doi: 10.1023/A:1020281327116

Anselmi, F., and Poggio, T. A. (2014). *Representation Learning in Sensory Cortex: A Theory*. Tech. Rep. CBMM Memo 026, MIT.

Ashby, W. (1952). *Design for a Brain*. London: Chapman & Hall. doi: 10.1007/978-94-015-1320-3

Ay, N., Bertschinger, N., Der, R., Güttler, F., and Olbrich, E. (2008). Predictive information and explorative behavior of autonomous robots. *Eur. Phys. J. B* 63, 329–339. doi: 10.1140/epjb/e2008-00175-0

Bachman, P., Sordoni, A., and Trischler, A. (2016). Towards information-seeking agents. ArXiv:1612.02605v1, 1–11.

Badre, D. (2008). Cognitive control, hierarchy, and the rostro-caudal organization of the frontal lobes. *Trends Cogn. Sci.* 12, 193–200. doi: 10.1016/j.tics.2008.02.004

Barak, O. (2017). Recurrent neural networks as versatile tools of neuroscience research. *Curr. Opin. Neurobiol.* 46, 1–6. doi: 10.1016/j.conb.2017.06.003

Barkow, J. H., Cosmides, L., and Tooby, J. (eds.). (1992). *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*. New York, NY: Oxford University Press. doi: 10.1017/s0730938400018700

Barlow, H. (2009). "Grandmother cells, symmetry, and invariance: how the term arose and what the facts suggest," in *Cognitive Neurosciences*, ed M. S. Gazzaniga (Cambridge, MA: MIT Press), 309–320.

Barrio, L. C., and Buno, W. (1990). Temporal correlations in sensory-synaptic interactions: example in crayfish stretch receptors. *J. Neurophys.* 63, 1520–1528.

Baxter, J. (1998). "Theoretical models of learning to learn," in *Learning to Learn*, eds S. Thrun and L. Pratt (Norwell, MA: Kluwer Academic Publishers), 71–94. doi: 10.1007/978-1-4615-5529-2_4

Beattie, C., Leibo, J. Z., Teplyashin, D., Ward, T., Wainwright, M., Lefrancq, A., et al. (2016). DeepMind lab. ArXiv:1612.03801v2, 1–11.

Bechtel, W. (1993). The case for connectionism. *Philos. Stud.* 71, 119–154. doi: 10.1007/bf00989853

Bedau, M. A. (2003). Artificial life: organization, adaptation and complexity from the bottom up. *Trends Cogn. Sci.* 7, 505–512. doi: 10.1016/j.tics.2003.09.012

Bengio, Y. (2009). Learning deep architectures for AI. *Found. Trends Mach. Learn.* 2, 1–87. doi: 10.1561/2200000006

Bengio, Y. (2014). "Evolving culture vs local minima," in *Growing Adaptive Machine*, eds T. Kowaliw, N. Bredeche, and R. Doursat (Berlin: Springer-Verlag), 109–138. doi: 10.1007/978-3-642-55337-0_3

Bengio, Y., and LeCun, Y. (2007). "Scaling learning algorithms towards AI," in *Large Scale Kernel Machines*, eds L. Bottou, O. Chapelle, D. DeCoste, and J. Weston (Cambridge, MA: The MIT Press), 321–360.

Bengio, Y., Louradour, J., Collobert, R., and Weston, J. (2009). "Curriculum learning," in *Proceedings of the 26th Annual International Conference on Machine Learning* (Montreal), 1–8. doi: 10.1145/1553374.1553380

Bianchini, M., and Scarselli, F. (2014). On the complexity of neural network classifiers: a comparison between shallow and deep architectures. *IEEE Trans. Neural Netw. Learn. Syst.* 25, 1553–1565. doi: 10.1109/TNNLS.2013.2293637

Bishop, C. M. (1995). *Neural Networks for Pattern Recognition*. Oxford: Oxford University Press.

Blei, D. M., Kucukelbir, A., and McAuliffe, J. D. (2016). Variational inference: a review for statisticians. ArXiv:1601.00670v5, 1–33.

Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent dirichlet allocation. *J. Mach. Learn. Res.* 3, 993–1022. doi: 10.1162/jmlr.2003.3.4-5.993

Boden, M. A. (1998). Creativity and artificial intelligence. *Artif. Intell.* 103, 347–356. doi: 10.1016/S0004-3702(98)00055-1

Bohte, S. M. (2004). The evidence for neural information processing with precise spike-times: a survey. *Nat. Comput.* 3, 195–206. doi: 10.1023/b:naco.0000027755.02868.60

Bordes, A., Chopra, S., and Weston, J. (2015). Large-scale simple question answering with memory networks. ArXiv:1506.02075v1, 1–10.

Bosch, S. E., Seeliger, K., and van Gerven, M. A. J. (2016). Modeling cognitive processes with neural reinforcement learning. BioArxiv, 1–19.

Brachman, R. J. (2002). Systems that know what they're doing. *IEEE Intell. Syst.* 17, 67–71. doi: 10.1109/mis.2002.1134363

Braitenberg, V. (1986). *Vehicles: Experiments in Synthetic Psychology*. Cambridge, MA: The MIT Press. doi: 10.1016/0004-3702(85)90057-8

Brea, J., Gaál, A. T., Urbanczik, R., and Senn, W. (2016). Prospective coding by spiking neurons. *PLoS Comput. Biol.* 12:e1005003. doi: 10.1371/journal.pcbi.1005003

Brea, J., and Gerstner, W. (2016). Does computational neuroscience need new synaptic learning paradigms? *Curr. Opin. Behav. Sci.* 11, 61–66. doi: 10.1016/j.cobeha.2016.05.012

Brette, R. (2015). Philosophy of the spike: rate-based vs spike-based theories of the brain. *Front. Syst. Neurosci.* 9:151. doi: 10.3389/fnsys.2015.00151

Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., et al. (2016). OpenAI gym. ArXiv:1606.01540v1, 1–4.

Brooks, R. A. (1992). "Artificial life and real robots," in *Toward a Practice of Autonomous Systems, Proceedings of First European Conference on Artificial Life*, eds F. J. Varela and P. Bourgine (Cambridge, MA: The MIT Press, Bradford Books).

Brooks, R. A. (1996). "Prospects for human level intelligence for humanoid robots," in *Proceedings of the First International Symposium on Humanoid Robots* (Tokyo), 17–24.

Brown, L. V. (2007). *Psychology of Motivation*. New York, NY: Nova Publishers.

Buschman, T. J., Miller, E. K., and Miller, E. K. (2014). Goal-direction and top-down control. *Philos. Trans. R. Soc. B* 369, 1–9. doi: 10.1098/rstb.2013.0471

Cannon, W. B. (1929). Organization for physiological homeostasis. *Physiol. Rev.* 9, 399–431.

Carnevale, F., de Lafuente, V., Romo, R., Barak, O., and Parga, N. (2015). Dynamic control of response criterion in premotor cortex during perceptual detection under temporal uncertainty. *Neuron* 86, 1067–1077. doi: 10.1016/j.neuron.2015.04.014

Carr, C. E., and Konishi, M. (1990). A circuit for detection of interaural time differences in the brain stem of the barn owl. *J. Neurosci.* 10, 3227–3246.

Caruana, R. (1997). Multitask learning. *Mach. Learn.* 28, 41–75. doi: 10.1109/TCBB.2010.22

Chang, E. F. (2015). Towards large-scale, human-based, mesoscopic neurotechnologies. *Neuron* 86, 68–78. doi: 10.1016/j.neuron.2015.03.037

Cho, K., van Merrienboer, B., Bahdanau, D., and Bengio, Y. (2014). "On the properties of neural machine translation: encoder-decoder approaches," in *Proceedings of the SSST-8, Eighth Work Syntax Semantics and Structure in Statistical Translation* (Doha), 103–111. doi: 10.3115/v1/w14-4012

Churchland, P. S., and Sejnowski, T. J. (2016). Blending computational and experimental neuroscience. *Nat. Rev. Neurosci.* 17, 667–668. doi: 10.1038/nrn.2016.114

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behav. Brain Sci.* 36, 181–204. doi: 10.1017/s0140525x12000477

Cohen, J. D., McClure, S. M., and Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philos. Trans. R. Soc. B* 362, 933–942. doi: 10.1098/rstb.2007.2098

Copeland, B. J., and Proudfoot, D. (1996). On Alan Turing's anticipation of connectionism. *Synthese* 108, 361–377. doi: 10.1007/bf00413694

Corneil, D., and Gerstner, W. (2015). "Attractor network dynamics enable preplay and rapid path planning in maze-like environments," in *Advances in Neural Information Processing Systems 28* (Montreal), 1–9.

Cox, D. D., and Dean, T. (2014). Neural networks and neuroscience-inspired computer vision. *Curr. Biol.* 24, R921–R929. doi: 10.1016/j.cub.2014.08.026

Crick, F., and Mitchison, G. (1983). The function of dream sleep. *Nature* 304, 111–114. doi: 10.1038/304111a0

Csikszentmihalyi, M. (1975). *Beyond Boredom and Anxiety: Experiencing Flow in Work and Play*. Hoboken, NJ: John Wiley & Sons Inc. doi: 10.2307/2065805

Cybenko, G. (1989). Approximation by superpositions of a sigmoidal function. *Math. Control Signals Syst.* 2, 303–314. doi: 10.1007/BF02134016

Dauphin, Y., Pascanu, R., Gulcehre, C., Cho, K., Ganguli, S., and Bengio, Y. (2014). Identifying and attacking the saddle point problem in high-dimensional non-convex optimization. ArXiv:1406.2572, 1–14.

Davies, N. B., Krebs, J. R., and West, S. A. (2012). *An Introduction to Behavioral Ecology, 4th Edn*. Hoboken, NJ: John Wiley & Sons. doi: 10.1037/026600

Daw, N. D. (2012). "Model-based reinforcement learning as cognitive search: neurocomputational theories," in *Cognitive Search: Evolution, Algorithms, and the Brain*, eds P. M. Todd, T. T. Hills, and T. W. Robbins (Cambridge, MA: The MIT Press), 195–208. doi: 10.7551/mitpress/9780262018098.001.0001

Dawkins, R. (2016). *The Selfish Gene, 4th Edn*. Oxford: Oxford University Press. doi: 10.4324/9781912281251

Dawson, M. R. W., and Shamanski, K. S. (1994). Connectionism, confusion, and cognitive science. *J. Intell. Syst.* 4, 215–262. doi: 10.1515/jisys.1994.4.3-4.215

Dayan, P., and Abbott, L. F. (2005). *Theoretical Neuroscience*. Cambridge, MA: MIT Press.

Dayan, P., Hinton, G. E., Neal, R., and Zemel, R. (1995). The Helmholtz machine. *Neural Comput.* 7, 1–16. doi: 10.1162/neco.1995.7.5.889

de Garis, H., Shuo, C., Goertzel, B., and Ruiting, L. (2010). A world survey of artificial brain projects, Part I Large-scale brain simulations. *Neurocomputing* 74, 3–29. doi: 10.1016/j.neucom.2010.08.004

Delalleau, O., and Bengio, Y. (2011). "Shallow vs. deep sum-product networks," in *Advances in Neural Information Processing Systems 24* (Granada), 666–674.

Der, R., and Martius, G. (2011). *The Playful Machine: Theoretical Foundation and Practical Realization of Self-Organizing Robots*. Berlin: Springer Verlag. doi: 10.1007/978-3-642-20253-7

Der, R., Steinmetz, U., and Pasemann, F. (1999). Homeokinesis - a new principle to back up evolution with learning. *Comput. Intell. Model. Control. Autom.* 55, 43–47.

Dewar, R. C. (2003). Information theory explanation of the fluctuation theorem, maximum entropy production and self-organized criticality in non-equilibrium stationary states. *J. Phys. A Math. Gen.* 36, 631–641. doi: 10.1088/0305-4470/36/3/303

Dewar, R. C. (2005). Maximum entropy production and the fluctuation theorem. *J. Phys. A Math. Gen.* 38, L371–L381. doi: 10.1088/0305-4470/38/21/L01

Dewey, J. (1896). The reflex arc concept in psychology. *Psychol. Rev.* 3, 357–370. doi: 10.1037/11304-041

Doya, K., Ishii, S., Pouget, A., and Rao, R. P. N. (eds.). (2006). *Bayesian Brain: Probabilistic Approaches to Neural Coding*. Cambridge, MA: The MIT Press.

Dragoi, G., and Tonegawa, S. (2011). Hippocampal cellular assemblies. *Nature* 469, 397–401. doi: 10.1038/nature09633

Drexler, K. E. (1992). *Nanosystems: Molecular Machinery, Manufacturing, and Computation*. New York, NY: Wiley Interscience. doi: 10.1016/S0010-8545(96)90165-4

Duan, Y., Andrychowicz, M., Stadie, B. C., Ho, J., Schneider, J., Sutskever, I., et al. (2017). One-shot imitation learning. ArXiv:1703.07326v2, 1–23.

Duysens, J., and Van de Crommert, H. W. A. A. (1998). Neural control of locomotion; The central pattern generator from cats to humans. *Gait Posture* 7, 131–141. doi: 10.1016/S0966-6362(97)00042-8

Edelman, S. (2015). The minority report: some common assumptions to reconsider in the modelling of the brain and behavior. *J. Exp. Theor. Artif. Intell.* 3079, 1–26. doi: 10.1080/0952813X.2015.1042534

Elman, J. L. (1990). Finding structure in time. *Cogn. Sci.* 14, 179–211. doi: 10.1016/0364-0213(90)90002-E

Elman, J. L. (1991). Distributed representations, simple recurrent networks, and grammatical structure. *Mach. Learn.* 7, 195–225. doi: 10.1023/A:1022699029236

Elman, J. L. (1993). Learning and development in neural networks - The importance of starting small. *Cognition* 48, 71–99. doi: 10.1016/S0010-0277(02)00106-3

Elman, J. L., Bates, E. A., Johnson, M. H., Karmiloff-Smith, A., Parisi, D., and Plunkett, K. (1996). *Rethinking Innateness: A Connectionist Perspective on Development*. Cambridge, MA: The MIT Press. doi: 10.1017/s0272263198333070

Fei-Fei, L., Fergus, R., Member, S., and Perona, P. (2006). One-shot learning of object categories. *IEEE Trans. Patt. Anal. Mach. Intell.* 28, 594–611. doi: 10.1109/tpami.2006.79

Felleman, D. J., and Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cereb. Cortex* 1, 1–47. doi: 10.1093/cercor/1.1.1

Fernando, C., Banarse, D., Blundell, C., Zwols, Y., Ha, D., Rusu, A., et al. (2017). PathNet: evolution channels gradient descent in super neural networks. ArXiv:1701.08734v1.

Ferrone, L., and Zanzotto, F. M. (2017). Symbolic, distributed and distributional representations for natural language processing in the era of deep learning: a survey. ArXiv:1702.00764, 1–25.

Ferrucci, D., Brown, E., Chu-carroll, J., Fan, J., Gondek, D., Kalyanpur, A. A., et al. (2010). Building Watson: an overview of the DeepQA project. *AI Mag.* 31, 59–79. doi: 10.1609/aimag.v31i3.2303

Feynman, R. (1992). There's plenty of room at the bottom. *J. Microelectromech. Syst.* 1, 60–66. doi: 10.1109/84.128057

Floreano, D., Dürr, P., and Mattiussi, C. (2008). Neuroevolution: from architectures to learning. *Evol. Intell.* 1, 47–62. doi: 10.1007/s12065-007-0002-4

Fodor, J. A., and Pylyshyn, Z. W. (1988). Connectionism and cognitive architecture: a critical analysis. *Cognition* 28, 3–71. doi: 10.1016/0010-0277(88)90031-5

Forstmann, B. U., and Wagenmakers, E.-J. (2015). *Model-Based Cognitive Neuroscience: A Conceptual Introduction*. New York, NY: Springer. doi: 10.1007/978-1-4939-2236-9_7

French, R. M. (1999). Catastrophic forgetting in connectionist networks. *Trends Cogn. Sci.* 3, 128–135. doi: 10.1016/s1364-6613(99)01294-2

Friston, K. J. (2009). The free-energy principle: a rough guide to the brain? *Trends Cogn. Sci.* 13, 293–301. doi: 10.1016/j.tics.2009.04.005

Friston, K. J. (2010). The free-energy principle: a unified brain theory? *Nat. Rev. Neurosci.* 11, 127–138. doi: 10.1038/nrn2787

Friston, K. J., Daunizeau, J., Kilner, J., and Kiebel, S. J. (2010). Action and behavior: a free-energy formulation. *Biol. Cybern.* 102, 227–260. doi: 10.1007/s00422-010-0364-z

Fry, R. L. (2017). Physical intelligence and thermodynamic computing. *Entropy* 19, 1–27. doi: 10.20944/PREPRINTS201701.0097.V1

Fukushima, K. (1980). Neocognitron: a self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biol. Cybern.* 36, 193–202. doi: 10.1007/bf00344251

Fukushima, K. (2013). Artificial vision by multi-layered neural networks: neocognitron and its advances. *Neural Netw.* 37, 103–119. doi: 10.1016/j.neunet.2012.09.016

Funahashi, K.-I., and Nakamura, Y. (1993). Approximation of dynamical systems by continuous time recurrent neural networks. *Neural Netw.* 6, 801–806. doi: 10.1016/s0893-6080(05)80125-x

Fuster, J. M. (2001). The prefrontal cortex - An update: time is of the essence. *Neuron* 30, 319–333. doi: 10.1016/S0896-6273(01)00285-9

Fuster, J. M. (2004). Upper processing stages of the perception-action cycle. *Trends Cogn. Sci.* 8, 143–145. doi: 10.1016/j.tics.2004.02.004

Gal, Y. (2016). Dropout as a Bayesian approximation: representing model uncertainty in deep learning. ArXiv:1506.02142v6, 1–12.

Gardner, E. (1988). The space of interactions in neural network models. *J. Phys. A. Math. Gen.* 21, 257–270. doi: 10.1088/0305-4470/21/1/030

Gardner, M. (2001). *The Colossal Book of Mathematics: Classic Puzzles, Paradoxes, and Problems*. New York, NY: W. W. Norton & Company.

Gasser, M., Eck, D., and Port, R. (1999). Meter as mechanism: a neural network model that learns metrical patterns. *Conn. Sci.* 11, 187–216. doi: 10.1080/095400999116331

Gauci, J., and Stanley, K. O. (2010). Autonomous evolution of topographic regularities in artificial neural networks. *Neural Comput.* 22, 1860–1898. doi: 10.1162/neco.2010.06-09-1042

Gershman, S. J., and Beck, J. M. (2016). "Complex probabilistic inference: from cognition to neural computation," in *Computational Models of Brain and Behavior*, ed A. Moustafa (Chichester, UK: Wiley-Blackwell), 1–17. doi: 10.1002/9781119159193.ch33

Gershman, S. J., Horvitz, E. J., and Tenenbaum, J. B. (2015). Computational rationality: a converging paradigm for intelligence in brains, minds, and machines. *Science* 349, 273–278. doi: 10.1126/science.aac6076

Gerstner, W., and Kistler, W. M. (2002). *Spiking Neuron Models*. Cambridge: Cambridge University Press.

Gerstner, W., Kistler, W. M., Naud, R., and Paninski, L. (2014). *Neuronal Dynamics: From Single Neurons to Networks and Models of Cognition*. Cambridge: Cambridge University Press. doi: 10.1017/CBO9781107447615

Gibson, J. (1979). *The Ecological Approach to Visual Perception*. Boston, MA: Houghton Mifflin. doi: 10.1002/bs.3830260313

Gigerenzer, G., and Goldstein, D. G. (1996). Reasoning the fast and frugal way: models of bounded rationality. *Psychol. Rev.* 103, 650–669. doi: 10.1037//0033-295x.103.4.650

Gilbert, C. D., and Li, W. (2013). Top-down influences on visual processing. *Nat. Rev. Neurosci.* 14, 350–363. doi: 10.1038/nrn3476

Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., et al. (2014). Generative adversarial nets. ArXiv:1406.2661v1, 1–9.

Gordon, G., and Ahissar, E. (2012). Hierarchical curiosity loops and active sensing. *Neural Netw.* 32, 119–129. doi: 10.1016/j.neunet.2012.02.024

Graves, A., Wayne, G., and Danihelka, I. (2014). Neural turing machines. ArXiv:1410.5401, 1–26.

Gregor, K., Danihelka, I., Graves, A., and Wierstra, D. (2014). DRAW: a recurrent neural network for image generation. ArXiv:1502.04623v1, 1–16.

Griffiths, T., Chater, N., and Kemp, C. (2010). Probabilistic models of cognition: exploring representations and inductive biases. *Trends Cogn. Sci.* 14, 357–364. doi: 10.1016/j.tics.2010.05.004

Grinstein, G., and Linsker, R. (2007). Comments on a derivation and application of the 'maximum entropy production' principle. *J. Phys. A Math. Theor.* 40, 9717–9720. doi: 10.1088/1751-8113/40/31/n01

Grothe, B. (2003). New roles for synaptic inhibition in sound localization. *Nat. Rev. Neurosci.* 4, 540–550. doi: 10.1038/nrn1136

Güçlü, U., Thielen, J., Hanke, M., and van Gerven, M. A. J. (2016). "Brains on beats," in *Advances in Neural Information Processing Systems 29* (Barcelona), 1–12.

Güçlü, U., and van Gerven, M. (2017a). Increasingly complex representations of natural movies across the dorsal stream are shared between subjects. *Neuroimage* 145, 329–336. doi: 10.1016/j.neuroimage.2015.12.036

Güçlü, U., and van Gerven, M. A. J. (2015). Deep neural networks reveal a gradient in the complexity of neural representations across the ventral stream. *J. Neurosci.* 35, 10005–10014. doi: 10.1523/JNEUROSCI.5023-14.2015

Güçlü, U., and van Gerven, M. A. J. (2017b). Modeling the dynamics of human brain activity with recurrent neural networks. *Front. Comput. Neurosci.* 11:7. doi: 10.3389/fncom.2017.00007

Güçlütürk, Y., Güçlü, U., Seeliger, K., Bosch, S., van Lier, R., and van Gerven, M. A. J. (2017). "Deep adversarial neural decoding," in *Advances in Neural Information Processing Systems 30* (Long Beach), 1–12.

Güçlütürk, Y., Güçlü, U., van Gerven, M. A. J., and van Lier, R. (2016). "Deep impression: audiovisual deep residual networks for multimodal apparent personality trait recognition," in *Proceedings of the 14th European Conference on Computer Vision* (Amsterdam).

Harari, Y. N. (2017). *Homo Deus: A Brief History of Tomorrow, 1st Edn*. New York, NY: Vintage Books.

Harnad, S. (1990). The symbol grounding problem. *Phys. D Nonlin. Phenom.* 42, 335–346. doi: 10.1016/0167-2789(90)90087-6

Hassabis, D., Kumaran, D., Summerfield, C., and Botvinick, M. (2017). Neuroscience-inspired artificial intelligence. *Neuron* 95, 245–258. doi: 10.1016/j.neuron.2017.06.011

Hatfield, G. (2002). "Perception and the physical world: psychological and philosophical issues in perception," in *Perception and the Physical World: Psychological and Philosophical Issues in Perception*, eds D. Heyer and R. Mausfeld (Hoboken, NJ: John Wiley and Sons), 113–143. doi: 10.1002/0470013427.ch5

He, K., Zhang, X., Ren, S., and Sun, J. (2015). Deep residual learning for image recognition. ArXiv:1512.03385, 1–12.

Heeger, D. J. (2017). Theory of cortical function. *Proc. Natl. Acad. Sci. U.S.A.* 114, 1773–1782. doi: 10.1073/pnas.1619788114

Herculano-Houzel, S., and Lent, R. (2005). Isotropic fractionator: a simple, rapid method for the quantification of total cell and neuron numbers in the brain. *J. Neurosci.* 25, 2518–2521. doi: 10.1523/JNEUROSCI.4526-04.2005

Hertz, J. A., Krogh, A. S., and Palmer, R. G. (1991). *Introduction to the Theory of Neural Computation*. Boulder, CO: Westview Press. doi: 10.1063/1.2810360

Hinton, G. (2013). Where do features come from? *Cogn. Sci.* 38, 1078–1101. doi: 10.1111/cogs.12049

Hinton, G. E., McLelland, J. L., and Rumelhart, D. E. (1986). "Distributed representations," in *Parallel Distributed Processing Explorations in the Microstructure of Cognition*, vol. 1, eds D. E. Rumelhart and J. L. McClelland (Cambridge, MA: MIT Press), 77–109.

Hinton, G. E., Osindero, S., and Teh, Y. (2006). A fast learning algorithm for deep belief nets. *Neural Comput.* 18, 1527–1554. doi: 10.1162/neco.2006.18.7.1527

Hinton, G. E., and Sejnowski, T. J. (1983). "Optimal perceptual inference," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Washington, DC).

Hochreiter, S., and Schmidhuber, J. (1997). Long short-term memory. *Neural Comput.* 9, 1735–1780. doi: 10.1162/neco.1997.9.8.1735

Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. U.S.A.* 79, 2554–2558. doi: 10.1073/pnas.79.8.2554

Hornik, K. (1991). Approximation capabilities of multilayer feedforward networks. *Neural Netw.* 4, 251–257. doi: 10.1016/0893-6080(91)90009-T

Huang, Y., and Rao, R. P. N. (2011). Predictive coding. *WIREs Cogn. Sci.* 2, 580–593. doi: 10.1002/wcs.142

Huh, D., and Sejnowski, T. J. (2017). Gradient descent for spiking neural networks. ArXiv:1706.04698, 1–10.

Huo, J., and Murray, A. (2009). The adaptation of visual and auditory integration in the barn owl superior colliculus with spike timing dependent plasticity. *Neural Netw.* 22, 913–921. doi: 10.1016/j.neunet.2008.10.007

Ijspeert, A. J. (2008). Central pattern generators for locomotion control in animals and robots: a review. *Neural Netw.* 21, 642–653. doi: 10.1016/j.neunet.2008.03.014

Ioffe, S., and Szegedy, C. (2015). Batch normalization: accelerating deep network training by reducing internal covariate shift. ArXiv:1502.03167, 1–11.

Izhikevich, E. M., and Edelman, G. M. (2008). Large-scale model of mammalian thalamocortical systems. *Proc. Natl. Acad. Sci. U.S.A.* 105, 3593–3598. doi: 10.1073/pnas.0712231105

Jaynes, E. (1988). How does the brain do plausible reasoning? *Maximum Entropy Bayesian Methods Sci. Eng.* 1, 1–24. doi: 10.1007/978-94-009-3049-0_1

Jeffress, L. A. (1948). A place theory of sound localization. *J. Comp. Physiol. Psychol.* 41, 35–39. doi: 10.1037/h0061495

Johnson, J., Hariharan, B., van der Maaten, L., Hoffman, J., Fei-Fei, L., Zitnick, C. L., et al. (2017). Inferring and executing programs for visual reasoning. ArXiv:1705.03633.

Jonas, E., and Kording, K. P. (2017). Could a neuroscientist understand a microprocessor? *PloS Comput. Biol.* 13:e1005268. doi: 10.1371/journal.pcbi.1005268

Jordan, M. I. (1987). "Attractor dynamics and parallelism in a connectionist sequential machine," in *Proceedings of the Eighth Annual Conference of the Cognitive Science Society*, 531–546.

Jordan, M. I., and Mitchell, T. M. (2015). Machine learning: trends, perspectives, and prospects. *Science* 349, 255–260. doi: 10.1126/science.aaa8415

Joukes, J., Hartmann, T. S., and Krekelberg, B. (2014). Motion detection based on recurrent network dynamics. *Front. Syst. Neurosci.* 8:239. doi: 10.3389/fnsys.2014.00239

Kadmon, J., and Sompolinsky, H. (2016). "Optimal architectures in a solvable model of deep networks," in *Advances in Neural Information Processing Systems 29* (Barcelona), 1–9.

Kaiser, Ł., and Roy, A. (2017). "Learning to remember rare events," in *5th International Conference on Learning Representations* (Toulon), 1–10.

Kanitscheider, I., and Fiete, I. (2016). Training recurrent networks to generate hypotheses about how the brain solves hard navigation problems. ArXiv:1609.09059, 1–10.

Kaplan, F., and Oudeyer, P.-Y. (2004). Maximizing learning progress: an internal reward system for development. *Embodied Artif. Intell.* 3139, 259–270. doi: 10.1007/b99075

Kass, R., Eden, U., and Brown, E. (2014). *Analysis of Neural Data*. New York, NY: Springer. doi: 10.1007/978-1-4614-9602-1

Kawaguchi, K., Kaelbling, L. P., and Bengio, Y. (2017). Generalization in deep learning. ArXiv: 1710.045468v1, 1–15.

Kemp, C., and Tenenbaum, J. B. (2008). The discovery of structural form. *Proc. Natl. Acad. Sci. U.S.A.* 105:10687. doi: 10.1073/pnas.0802631105

Kempka, M., Wydmuch, M., Runc, G., Toczek, J., and Ja, W. (2016). ViZDoom: a Doom-based AI research platform for visual reinforcement learning. ArXiv:1605.02097v2, 1–8.

Kheradpisheh, S. R., Ganjtabesh, M., and Thorpe, S. J. (2016). STDP-based spiking deep neural networks for object recognition. ArXiv:1611.01421v1, 1–16.

Kietzmann, T. C., McClure, P., and Kriegeskorte, N. (2017). Deep neural networks in computational neuroscience. BioRxiv, 1–23. doi: 10.1101/133504

Kindermans, P.-J., Schütt, K. T., Alber, M., Müller, K.-R., and Dähne, S. (2017). PatternNet and PatternLRP – improving the interpretability of neural networks. ArXiv:1705.05598, 1–11.

Kingma, D. P., and Welling, M. (2014). Auto-encoding variational Bayes. ArXiv:1312.6114, 1–14.

Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., and Rusu, A. A. (2015). Overcoming catastrophic forgetting in neural networks. ArXiv:1612.00796v1, 1–13.

Klyubin, A., Polani, D., and Nehaniv, C. (2005a). Empowerment: a universal agent-centric measure of control. *IEEE Congr. Evol. Comput.* 1, 128–135. doi: 10.1109/CEC.2005.1554676

Klyubin, A. S., Polani, D., and Nehaniv, C. L. (2005b). "All else being equal be empowered," in *Lecture Notes in Computer Science*, Vol. 3630 (Canterbury), 744–753. doi: 10.1007/11553090_75

Koller, D., and Friedman, N. (2009). *Probabilistic Graphical Models: Principles and Techniques*. Cambridge, MA: The MIT Press.

Kriegeskorte, N. (2015). Deep neural networks: a new framework for modeling biological vision and brain information processing. *Annu. Rev. Vis. Sci.* 1, 417–446. doi: 10.1146/annurev-vision-082114-035447

Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). "ImageNet classification with deep convolutional neural networks," in *Advances in Neural Information Processing Systems 25* (Lake Tahoe), 1106–1114.

Kruschke, J. K. (1992). ALCOVE: an exemplar-based connectionist model of category learning. *Psychol. Rev.* 99, 22–44. doi: 10.1037/0033-295X.99.1.22

Kumaran, D., Hassabis, D., and McClelland, J. L. (2016). What learning systems do intelligent agents need? Complementary learning systems theory updated. *Trends Cogn. Sci.* 20, 512–534. doi: 10.1016/j.tics.2016.05.004

Laird, J. E. (2012). *The SOAR Cognitive Architecture*. Cambridge, MA: The MIT Press.

Laje, R., and Buonomano, D. V. (2013). Robust timing and motor patterns by taming chaos in recurrent neural networks. *Nat. Neurosci.* 16, 925–933. doi: 10.1038/nn.3405

Lake, B. M., Ullman, T. D., Tenenbaum, J. B., and Gershman, S. J. (2017). Building machines that learn and think like people. *Behav. Brain Sci.* doi: 10.1017/s0140525x16001837

Larochelle, H., and Hinton, G. E. (2010). "Learning to combine foveal glimpses with a third-order Boltzmann machine," in *Advances in Neural Information Processing Systems 23*, Vol. 40 (Vancouver), 1243–1251.

Laughlin, S. B., and Sejnowski, T. J. (2003). Communication in neuronal networks. *Science* 301, 1870–1874. doi: 10.1126/science.1089662

Le Roux, N., and Bengio, Y. (2010). Deep belief networks are compact universal approximators. *Neural Comput.* 22, 2192–2207. doi: 10.1162/neco.2010.08-09-1081

LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature* 521, 436–444. doi: 10.1038/nature14539

LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 2278–2324. doi: 10.1109/5.726791

Lee, J. H., Delbruck, T., and Pfeiffer, M. (2016). Training deep spiking neural networks using backpropagation. ArXiv:1608.08782, 1–10.

Lee, T., and Mumford, D. (2003). Hierarchical Bayesian inference in the visual cortex. *J. Opt. Soc. Am. A* 20, 1434–1448. doi: 10.1364/josaa.20.001434

Lehky, S. R., and Tanaka, K. (2016). Neural representation for object recognition in inferotemporal cortex. *Curr. Opin. Neurobiol.* 37, 23–35. doi: 10.1016/j.conb.2015.12.001

Leibo, J. Z., Liao, Q., Anselmi, F., Freiwald, W. A., and Poggio, T. (2017). View-tolerant face recognition and Hebbian learning imply mirror-symmetric neural tuning to head orientation. *Curr. Biol.* 27, 62–67. doi: 10.1016/j.cub.2016.10.015

Levine, S., Finn, C., Darrell, T., and Abbeel, P. (2015). End-to-end training of deep visuomotor policies. ArXiv:1504.00702v1, 1–12.

Liao, Q., and Poggio, T. (2016). Bridging the gaps between residual learning, recurrent neural networks and visual cortex. ArXiv:1604.03640, 1–16.

Lillicrap, T. P., Cownden, D., Tweed, D. B., and Akerman, C. J. (2016). Random feedback weights support learning in deep neural networks. *Nat. Commun.* 7, 1–10. doi: 10.1038/ncomms13276

Lin, H. W., and Tegmark, M. (2016). Why does deep and cheap learning work so well? ArXiv:1608.08225, 1–14.

Lopez, C. M., Mitra, S., Putzeys, J., Raducanu, B., Ballini, M., Andrei, A., et al. (2016). "A 966-electrode neural probe with 384 configurable channels in 0.13μm SOI CMOS," in *Solid State Circuits Conference Dig Technical Papers* (San Francisco, CA), 21–23. doi: 10.1109/ISSCC.2016.7418072

Lotter, W., Kreiman, G., and Cox, D. (2016). Deep predictive coding networks for video prediction and unsupervised learning. ArXiv:1605.08104, 1–12.

Louizos, C., Shalit, U., Mooij, J., Sontag, D., Zemel, R., and Welling, M. (2017). Causal effect inference with deep latent-variable models. ArXiv:1705.08821, 1–12.

Maass, W. (1997). Networks of spiking neurons: the third generation of neural network models. *Neural Netw.* 10, 1659–1671.

Maass, W. (2016). Searching for principles of brain computation. BioRxiv, 1–16.

MacKay, D. J. C. (2003). *Information Theory, Inference and Learning Algorithms*. Cambridge: Cambridge University Press. doi: 10.1108/03684920410534506

Mandt, S., Hoffman, M. D., and Blei, D. M. (2017). Stochastic gradient descent as approximate Bayesian inference. ArXiv:1704.04289v1, 1–30.

Mante, V., Sussillo, D., Shenoy, K. V., and Newsome, W. T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* 503, 78–84. doi: 10.1038/nature12742

Marblestone, A. H., Wayne, G., and Kording, K. P. (2016). Towards an integration of deep learning and neuroscience. *Front. Comput. Neurosci.* 10:94. doi: 10.3389/fncom.2016.00094

Marcus, G. (2009). How does the mind work? Insights from biology. *Top. Cogn. Sci.* 1, 145–172. doi: 10.1111/j.1756-8765.2008.01007.x

Marder, E. (2015). Understanding brains: details, intuition, and big data. *PLoS Biol.* 13:e1002147. doi: 10.1371/journal.pbio.1002147

Markram, H. (2006). The blue brain project. *Nat. Rev. Neurosci.* 7, 153–160. doi: 10.1038/nrn1848

Markram, H., Meier, K., Lippert, T., Grillner, S., Frackowiak, R., Dehaene, S., et al. (2011). Introducing the human brain project. *Proc. Comput. Sci.* 7, 39–42. doi: 10.1016/j.procs.2011.12.015

Marr, D. (1969). A theory of cerebellar cortex. *J. Physiol.* 202, 437–470. doi: 10.2307/1776957

Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*.

Cambridge, MA: MIT Press. doi: 10.7551/mitpress/9780262514620.001.0001

Marr, D., and Poggio, T. (1976). *From Understanding Computation to Understanding Neural Circuitry*. Tech. Rep. MIT.

Mathieu, M., Couprie, C., and LeCun, Y. (2016). "Deep multi-scale video prediction beyond mean square error," in *4th International Conference on Learning Representations* (San Juan), 1–14.

Maturana, H., and Varela, F. (1980). *Autopoiesis and Cognition: The Realization of the Living, 1st Edn*. Dordrecht: D. Reidel Publishing Company. doi: 10.1007/978-94-009-8947-4

Maturana, H., and Varela, F. (1987). *The Tree of Knowledge - The Biological Roots of Human Understanding*. London: New Science Library.

McClelland, J. L. (2003). The parallel distributed processing approach to semantic cognition. *Nat. Rev. Neurosci.* 4, 310–322. doi: 10.1038/nrn1076

McClelland, J. L., Botvinick, M. M., Noelle, D. C., Plaut, D. C., Rogers, T. T., Seidenberg, M. S., et al. (2010). Letting structure emerge: connectionist and dynamical systems approaches to cognition. *Trends Cogn. Sci.* 14, 348–356. doi: 10.1016/j.tics.2010.06.002

McCloskey, M., and Cohen, N. J. (1989). Catastrophic inference in connectionist networks: the sequential learning problem. *Psychol. Learn. Motiv.* 24, 109–165. doi: 10.1016/s0079-7421(08)60536-8

McCorduck, P. (2004). *Machines Who Think, 2nd Edn*. Natick, MA: A. K. Peters, Ltd.

Mcintosh, L. T., Maheswaranathan, N., Nayebi, A., Ganguli, S., and Baccus, S. A. (2016). "Deep learning models of the retinal response to natural scenes," in *Advances in Neural Information Processing Systems 29* (Barcelona), 1–9.

Mead, C. (1990). Neuromorphic electronic systems. *Proc. IEEE* 78, 1629–1636 doi: 10.1109/5.58356

Mhaskar, H., Liao, Q., and Poggio, T. (2016). Learning functions: when is deep better than shallow. ArXiv:1603.00988v4, 1–12.

Miconi, T. (2017). Biologically plausible learning in recurrent neural networks for flexible decision tasks. *Elife* 6:e20899. doi: 10.16373/j.cnki.ahr.150049

Mikolov, T., Chen, K., Corrado, G., and Dean, J. (2013). "Efficient estimation of word representations in vector space," in *1st International Conference on Learning Representations* (Scottsdale).

Miller, E. K., and Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annu. Rev. Neurosci.* 24, 167–202. doi: 10.1146/annurev.neuro.24.1.167

Miller, T. (2017). Explanation in artificial intelligence: insights from the social sciences. ArXiv:1706.07269v1, 1–57.

Minsky, M., and Papert, S. (1969). *Perceptrons. An Introduction to Computational Geometry*. Cambridge, MA: MIT Press.

Mnih, V., Badia, A. P., Mirza, M., Graves, A., Lillicrap, T. P., Harley, T., et al. (2016). Asynchronous methods for deep reinforcement learning. ArXiv:1602.01783, 1–28.

Mnih, V., Heess, N., Graves, A., and Kavukcuoglu, K. (2014). "Recurrent models of visual attention," *Advances in Neural Information Processing Systems 27* (Montreal), 1–9.

Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., et al. (2015). Human-level control through deep reinforcement learning. *Nature* 518, 529–533. doi: 10.1038/nature14236

Modha, D. S., Ananthanarayanan, R., Esser, S. K., Ndirango, A., Sherbondy, A. J., and Singh, R. (2011). Cognitive computing. *Commun. ACM* 54, 62–71. doi: 10.1145/1978542.1978559

Moravec, H. P. (2000). *Robot: Mere Machine to Transcendent Mind*. New York, NY: Oxford University Press.

Moser, M.-B., Rowland, D. C., and Moser, E. I. (2015). Place cells, grid cells, and memory. *Cold Spring Harb. Perspect. Biol.* 7:a021808 doi: 10.1101/cshperspect.a021808

Moulton, S. T., and Kosslyn, S. M. (2009). Imagining predictions: mental imagery as mental emulation. *Philos. Trans. R. Soc. B* 364, 1273–1280. doi: 10.1098/rstb.2008.0314

Mozer, M. C. (1989). A focused back-propagation algorithm for temporal pattern recognition. *Complex Syst.* 3, 349–381.

Mozer, M. C., and Smolensky, P. (1989). Using relevance to reduce network size automatically. *Conn. Sci.* 1, 3–16. doi: 10.1080/09540098908915626

Mujika, A. (2016). Multi-task learning with deep model based reinforcement learning. ArXiv:1611.01457, 1–11.

Najemnik, J., and Geisler, W. S. (2005). Optimal eye movement strategies in visual search. *Nature* 434, 387–391. doi: 10.1038/nature03390

Nayebi, A., and Ganguli, S. (2016). Biologically inspired protection of deep networks from adversarial attacks. ArXiv:1703.09202v1, 1–11.

Neftci, E., Binas, J., Rutishauser, U., Chicca, E., Indiveri, G., and Douglas, R. J. (2013). Synthesizing cognition in neuromorphic electronic systems. *Proc. Natl. Acad. Sci. U.S.A.* 110, E3468–E3476. doi: 10.1073/pnas.1212083110

Neil, D., Pfeiffer, M., and Liu, S.-C. (2016). Phased LSTM: accelerating recurrent network training for long or event-based sequences. ArXiv:1610.09513v1, 1–9.

Nessler, B., Pfeiffer, M., Buesing, L., and Maass, W. (2013). Bayesian computation emerges in generic cortical microcircuits through spike-timing-dependent plasticity. *PLoS Comput. Biol.* 9:e1003037. doi: 10.1371/journal.pcbi.1003037

Newell, A. (1991). *Unified Theories of Cognition*. Cambridge, MA: Harvard University Press.

Newell, A., and Simon, H. A. (1976). Computer science as empirical inquiry: symbols and search. *Commun. ACM* 19, 113–126. doi: 10.1145/360018.360022

Nguyen, A., Dosovitskiy, A., Yosinski, J., Brox, T., and Clune, J. (2016). Synthesizing the preferred inputs for neurons in neural networks via deep generator networks. ArXiv:1605.09304, 1–29.

Nilsson, N. (2005). Human-level artificial intelligence? Be serious! *AI Mag.* 26, 68–75. doi: 10.1609/aimag.v26i4.1850

Obermayer, K. (1990). A principle for the formation of the spatial structure of cortical feature maps. *Proc. Natl. Acad. Sci. U.S.A.* 87, 8345–8349. doi: 10.1073/pnas.87.21.8345

O'Connor, P., and Welling, M. (2016). Deep spiking networks. ArXiv:1602.08323, 1–10.

Olshausen, B. A., and Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature* 381, 607–609. doi: 10.1038/381607a0

O'Reilly, R. (1998). Six principles for biologically based computational models of cortical cognition. *Trends Cogn. Sci.* 2, 1–8. doi: 10.1016/s1364-6613(98)01241-8

O'Reilly, R., Hazy, T., and Herd, S. (2012). "The Leabra cognitive architecture: how to play 20 principles with nature and win!," in *The Oxford Handbook of Cognitive Science*, ed S. E. F. Chipman (Oxford: Oxford University Press), 1–31. doi: 10.1093/oxfordhb/9780199842193.013.8

O'Reilly, R. C., Wyatte, D., Herd, S., Mingus, B., and Jilk, D. J. (2013). Recurrent processing during object recognition. *Front. Psychol.* 4:124. doi: 10.3389/fpsyg.2013.00124

Orhan, A. E., and Ma, W. J. (2016). Probabilistic inference in generic neural networks trained with non-probabilistic feedback. ArXiv:1601.03060v4, 1–30.

Oudeyer, P.-Y. (2007). "Intrinsically motivated machines," in *Lecture Notes Computer Science*, Vol. 4850, eds M. Lungarella, F. Iida, J. Bongard, and R. Pfeifer (Berlin: Springer), 304–314. doi: 10.1007/978-3-540-77296-5_27

Pachitariu, M., Stringer, C., Schröder, S., Dipoppa, M., Rossi, L. F., Carandini, M., et al. (2016). Suite2p: beyond 10,000 neurons with standard two-photon microscopy. BioRxiv, 1–14.

Pakkenberg, B., and Gundersen, H. (1997). Neocortical neuron number in humans: effect of sex and age. *J. Comp. Neurol.* 384, 312–320.

Pakkenberg, B., Pelvig, D., Marner, L., Bundgaard, M., Gundersen, H., Nyengaard, J., et al. (2003). Aging and the human neocortex. *Exp. Gerontol.* 38, 95–99. doi: 10.1016/s0531-5565(02)00151-1

Palatucci, M., Pomerleau, D., Hinton, G. E., and Mitchell, T. (2009). "Zero-shot learning with semantic output codes," in *Advances in Neural Information Processing Systems 22*, eds Y. Bengio, D. Schuurmans, J. Lafferty, C. K. I. Williams, and A. Culotta (Vancouver), 1410–1418.

Pan, S. J., and Fellow, Q. Y. (2009). A survey on transfer learning. *IEEE Trans. Knowl. Data Eng.* 22, 1–15. doi: 10.1109/TKDE.2009.191

Pascanu, R., Mikolov, T., and Bengio, Y. (2012). "On the difficulty of training recurrent neural networks," in *Proceedings of the 30th International Conference on Machine Learning* (Atlanta), 1310–1318.

Pascanu, R., Montufar, G., and Bengio, Y. (2013). On the number of response regions of deep feed forward networks with piece-wise linear activations. ArXiv:1312.6098, 1–17.

Pathak, D., Agrawal, P., Efros, A. A., and Darrell, T. (2017). Curiosity-driven exploration by self-supervised prediction. ArXiv:1705.05363, 1–12.

Peelen, M. V., and Downing, P. E. (2017). Category selectivity in human visual cortex: beyond visual object recognition. *Neuropsychologia* 105, 1–7. doi: 10.1016/j.neuropsychologia.2017.03.033

Perunov, N., Marsland, R., and England, J. (2014). Statistical physics of adaptation. ArXiv:1412.1875, 1–24.

Peterson, J. C., Abbott, J. T., and Griffiths, T. L. (2016). Adapting deep network features to capture psychological representations. ArXiv:1608.02164, 1–6.

Pinker, S., and Mehler, J. (eds.). (1988). *Connections and Symbols*. Cambridge, MA: The MIT Press.

Poggio, T. (2012). The levels of understanding framework, revised *Perception* 41, 1017–1023. doi: 10.1068/p7299

Poole, B., Lahiri, S., Raghu, M., Sohl-Dickstein, J., and Ganguli, S. (2016). Exponential expressivity in deep neural networks through transient chaos. ArXiv:1606.05340, 1–16.

Pouget, A., Beck, J. M., Ma, W. J., and Latham, P. E. (2013). Probabilistic brains: knowns and unknowns. *Nat. Neurosci.* 16, 1170–1178. doi: 10.1038/nn.3495

Pritzel, A., Uria, B., Srinivasan, S., Puigdomènech, A., Vinyals, O., Hassabis, D., et al. (2017). Neural episodic control. ArXiv:1703.01988, 1–12.

Quian Quiroga, R., Reddy, L., Kreiman, G., Koch, C., and Fried, I. (2005). Invariant visual representation by single neurons in the human brain. *Nature* 435, 1102–1107. doi: 10.1038/nature03687

Rafler, S. (2011). Generalization of Conway's "Game of Life" to a continuous domain - SmoothLife. ArXiv:1111.1567v2, 1–4.

Raghu, M., Kleinberg, J., Poole, B., Ganguli, S., and Sohl-Dickstein, J. (2016). Survey of expressivity in deep neural networks. ArXiv:1611.08083v1, 1–5.

Raina, R., Madhavan, A., and Ng, A. (2009). "Large-scale deep unsupervised learning using graphics processors," in *Proceedings of the 26th Annual International Conference on Machine Learning* (Montreal), 1–8. doi: 10.1145/1553374.1553486

Rajan, K., Harvey, C. D., and Tank, D. W. (2015). Recurrent network models of sequence generation and memory. *Neuron* 90, 1–15. doi: 10.1016/j.neuron.2016.02.009

Ramsey, F. P. (1926). "Truth and probability," in *The Foundations of Mathematics and other Logical Essays*, ed R. B. Braithwaite (Abingdon: Routledge), 156–198. doi: 10.1007/978-3-319-20451-2_3

Rao, R. P. and Ballard, D. H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat. Neurosci.* 2, 79–87. doi: 10.1038/4580

Real, E., Moore, S., Selle, A., Saxena, S., Suematsu, Y. L., Le, Q., et al. (2016). Large-scale evolution of image classifiers. ArXiv:1703.01041v1, 1–10.

Regan, J. K. O., and Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behav. Brain Sci.* 24, 939–1031. doi: 10.1017/s0140525x01000115

Rid, T. (2016). *Rise of the Machines: A Cybernetic History*. New York, NY: W. W. Norton & Company.

Riesenhuber, M., and Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nat. Neurosci.* 2, 1019–1025.

Ritter, H., and Kohonen, T. (1989). Self-organizing semantic maps. *Biol. Cybern.* 61, 241–254. doi: 10.1007/bf00203171

Robinson, L., and Rolls, E. T. (2015). Invariant visual object recognition: biologically plausible approaches. *Biol. Cybern.* 209, 505–535. doi: 10.1007/s00422-015-0658-2

Roelfsema, P. R., and van Ooyen, A. (2005). Attention-gated reinforcement learning of internal representations for classification. *Neural Comput.* 17, 2176–2214. doi: 10.1162/0899766054615699

Rosenblatt, F. (1958). The perceptron: a probabilistic model for information storage and organization in the brain. *Psychol. Rev.* 65, 386–408. doi: 10.1037/h0042519

Rumelhart, D., Hinton, G., and Williams, R. (1986). "Learning internal representations by error propagation," in *Parallel Distributed Processing, Explorations in the Microstructure of Cognition*, eds D. E. Rumelhart and J. L. McClelland (Cambridge, MA: MIT Press), 318–362.

Salge, C., Glackin, C., and Polani, D. (2013). Empowerment - An introduction. ArXiv:1310.1863, 1–46.

Salimans, T., Ho, J., Chen, X., and Sutskever, I. (2017). Evolution strategies as a scalable alternative to reinforcement learning. ArXiv:1703.03864v2, 1–12.

Santana, E., and Hotz, G. (2016). Learning a driving simulator. ArXiv:1608.01230, 1–8.

Santoro, A., Bartunov, S., Botvinick, M., Wierstra, D., and Lillicrap, T. (2016). One-shot learning with memory-augmented neural networks. ArXiv:1605.06065v1, 1–13.

Santoro, A., Raposo, D., Barrett, D. G. T., Malinowski, M., Pascanu, R., Battaglia, P., et al. (2017). A simple neural network module for relational reasoning. ArXiv:1706.01427v1, 1–16.

Saxe, A., McClelland, J., and Ganguli, S. (2014). "Exact solutions to the nonlinear dynamics of learning in deep linear neural networks Andrew," in *2nd International Conference on Learning Representations* (Banff), 1–22.

Scellier, B., and Bengio, Y. (2017). Equilibrium propagation: bridging the gap between energy-based models and backpropagation. *Front. Comput. Neurosci.* 11:24. doi: 10.3389/fncom.2017.00024

Schaal, S. (1999). Is imitation learning the route to humanoid robots? *Trends Cogn. Sci.* 3, 233–242. doi: 10.1016/s1364-6613(99)01327-3

Schacter, D. L., Addis, D. R., and Buckner, R. L. (2007). Remembering the past to imagine the future: the prospective brain. *Nat. Rev. Neurosci.* 8, 657–661. doi: 10.1038/nrn2213

Schiess, M., Urbanczik, R., and Senn, W. (2016). Somato-dendritic synaptic plasticity and error-backpropagation in active dendrites. *PLoS Comput. Biol.* 12:e1004638. doi: 10.1371/journal.pcbi.1004638

Schmidhuber, J. (1991). "Curious model-building control systems," in *Proceedings of International Joint Conference on Neural Networks*, Vol. 2 (Singapore), 1458–1463. doi: 10.1109/IJCNN.1991.170605

Schmidhuber, J. (2003). "Exploring the predictable," in *Advances in Evolutionary Computing*, eds A. Ghosh and S. Tsutsui (Berlin: Springer), 579–612. doi: 10.1017/CBO9781107415324.004

Schmidhuber, J. (2015). On learning to think: algorithmic information theory for novel combinations of reinforcement learning controllers and recurrent neural world models. ArXiv:1511.09249, 1–36.

Schoenholz, S. S., Gilmer, J., Ganguli, S., and Sohl-Dickstein, J. (2017). "Deep information propagation," in *5th International Conference on Learning Representations* (Toulon), 1–18.

Schoenmakers, S., Barth, M., Heskes, T., and van Gerven, M. A. J. (2013). Linear reconstruction of perceived images from human brain activity. *Neuroimage* 83, 951–961. doi: 10.1016/j.neuroimage.2013.07.043

Scholte, H. S., Losch, M. M., Ramakrishnan, K., de Haan, E. H. F., and Bohte, S. M. (2017). Visual pathways from the perspective of cost functions and deep learning. BioRxiv,1–16.

Schroeder, C. E., Wilson, D. A., Radman, T., Scharfman, H., and Lakatos, P. (2010). Dynamics of active sensing and perceptual selection. *Curr. Opin. Neurobiol.* 20, 172–176. doi: 10.1016/j.conb.2010.02.010

Schulman, J., Levine, S., Moritz, P., Jordan, M., and Abbeel, P. (2015). Trust region policy optimization. ArXiv:1502.05477v4, 1–16.

Schultz, W., Dayan, P., and Montague, P. R. (1997). A neural substrate of prediction and reward. *Science* 275, 1593–1599. doi: 10.1126/science.275.5306.1593

Schuman, C. D., Potok, T. E., Patton, R. M., Birdwell, J. D., Dean, M. E., Rose, G. S., et al. (2017). A survey of neuromorphic computing and neural networks in hardware. ArXiv:1705.06963, 1–88.

Searle, J. R. (1980). Minds, brains and Programs. *Behav. Brain Sci.* 3, 417–424. doi: 10.1017/s0140525x00005756

Segundo, J. P., Perkel, D. H., and Moore, G. P. (1966). Spike probability in neurones: influence of temporal structure in the train of synaptic events. *Kybernetik* 3, 67–82. doi: 10.1007/BF00299899

Seising, R. (2017). Marvin Lee Minsky (1927-2016). *Artif. Intell. Med.* 75, 24–31. doi: 10.1016/j.artmed.2016.12.001

Selfridge, O. (1959). "Pandemonium: a paradigm for learning," in *Symposium on the Mechanization of Thought Processes* (Teddington), 513–526.

Shwartz-Ziv, R., and Tishby, N. (2017). Opening the black box of deep neural networks via information. ArXiv:1703.00810, 1–19.

Silver, D., Lever, G., Heess, N., Degris, T., Wierstra, D., and Riedmiller, M. (2014). "Deterministic policy gradient algorithms," in *2nd International Conference on Learning Representations* (Banff), 387–395.

Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., et al. (2017). Mastering the game of Go without human knowledge. *Nature* 550, 354–359. doi: 10.1038/nature24270

Simon, H. A. (1962). The architecture of complexity. *Proc. Am. Philos. Soc.* 106, 467–482. doi: 10.1007/978-1-4899-0718-9_31

Simon, H. A. (1996). *The Sciences of the Artificial, 3rd Edn*. Cambridge, MA: The MIT Press.

Singer, W. (2013). Cortical dynamics revisited. *Trends Cogn. Sci.* 17, 616–626. doi: 10.1016/j.tics.2013.09.006

Smolensky, P. (1987). Connectionist AI, symbolic AI, and the brain. *Artif. Intell. Rev.* 1, 95–109. doi: 10.1007/BF00130011

Sompolinsky, H. (1988). Statistical mechanics of neural networks. *Phys. Today* 40, 70–80. doi: 10.1063/1.881142

Sompolinsky, H. (2014). Computational neuroscience: beyond the local circuit. *Curr. Opin. Neurobiol.* 25, 1–6. doi: 10.1016/j.conb.2014.02.002

Song, H. F., Yang, G. R., and Wang, X.-J. (2016). Reward-based training of recurrent neural networks for diverse cognitive and value-based tasks. *Elife* 6, 1–51. doi: 10.1101/070375

Sperry, R. W. (1952). Neurology and the mind-brain problem. *Am. Sci.* 40, 291–312.

Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. (2014). Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* 15, 1929–1958. doi: 10.1214/12-AOS1000

Stanley, K., and Miikkulainen, R. (2002). Evolving neural networks through augmenting topologies. *Evol. Comput.* 10, 1–30. doi: 10.1162/106365602320169811

Steels, L. (1993). The artificial life roots of artificial intelligence. *Artif. Life* 1, 75–110. doi: 10.1162/artl.1993.1.1_2.75

Steels, L. (2004). "The autotelic principle," in *Embodied Artificial Intelligence. Lecture Notes in Computer Science*, eds F. Iida, R. Pfeifer, L. Steels, and Y. Kuniyoshi (Berlin; Heidelberg: Springer), 231–242. doi: 10.1007/978-3-540-27833-7_17

Sterling, P. (2012). Allostasis: a model of predictive regulation. *Physiol. Behav.* 106, 5–15. doi: 10.1016/j.physbeh.2011.06.004

Sterling, P., and Laughlin, S. (2016). *Principles of Neural Design*. Cambridge, MA: MIT Press. doi: 10.7551/mitpress/9780262028707.001.0001

Strukov, D. B. (2011). Smart connections. *Nature* 476, 403–405. doi: 10.1038/476403a

Summerfield, C., and de Lange, F. P. (2014). Expectation in perceptual decision making: neural and computational mechanisms. *Nat. Rev. Neurosci.* 15, 745–756. doi: 10.1038/nrn3838

Sun, R. (2004). Desiderata for cognitive architectures. *Philos. Psychol.* 17, 341–373. doi: 10.1080/0951508042000286721

Sun, R., Coward, L. A., and Zenzen, M. J. (2005). On levels of cognitive modeling. *Philos. Psychol.* 18, 613–637. doi: 10.1080/09515080500264248

Sussillo, D., Churchland, M. M., Kaufman, M. T., and Shenoy, K. V. (2015). A neural network that finds a naturalistic solution for the production of muscle activity. *Nat. Neurosci.* 18, 1025–1033. doi: 10.1038/nn.4042

Sutskever, I., Vinyals, O., and Le, Q. V. (2014). "Sequence to sequence learning with neural networks," in *Advances in Neural Information Processing Systems 27* (Montreal), 3104–3112.

Sutton, R. S., and Barto, A. G. (1998). *Reinforcement Learning: An Introduction*. Cambridge, MA: MIT Press. doi: 10.1016/j.brainres.2010.09.091

Swanson, L. W. (2000). Cerebral hemisphere regulation of motivated behavior. *Brain Res.* 886, 113–164. doi: 10.1016/s0006-8993(00)02905-x

Swanson, L. W. (2012). *Brain Architecture: Understanding the Basic Plan, 2nd Edn*. Oxford: Oxford University Press.

Synnaeve, G., Nardelli, N., Auvolat, A., Chintala, S., Lacroix, T., Lin, Z., et al. (2016). TorchCraft: a library for machine learning research on real-time strategy games. ArXiv:1611.00625v2, 1–6.

Szigeti, B., Gleeson, P., Vella, M., Khayrulin, S., Palyanov, A., Hokanson, J., et al. (2014). OpenWorm: an open-science approach to modeling Caenorhabditis elegans. *Front. Comput. Neurosci.* 8:137. doi: 10.3389/fncom.2014.00137

Tapaswi, M., Zhu, Y., Stiefelhagen, R., Torralba, A., Urtasun, R., and Fidler, S. (2015). MovieQA: understanding stories in movies through question-answering. ArXiv:1512.02902, 1–10.

Tenenbaum, J. B., Kemp, C., Griffiths, T. L., and Goodman, N. D. (2011). How to grow a mind: statistics, structure, and abstraction. *Science* 331, 1279–1285. doi: 10.1126/science.1192788

Thalmeier, D., Uhlmann, M., Kappen, H. J., Memmesheimer, R.-M., and May, N. C. (2015). Learning universal computations with spikes. ArXiv:1505.07866v1, 1–35.

Thorpe, S. J., and Fabre-Thorpe, M. (2001). Seeking categories in the brain. *Science* 291, 260–262. doi: 10.1126/science.1058249

Thrun, S., and Mitchell, T. M. (1995). Lifelong robot learning. *Robot. Auton. Syst.* 15, 25–46. doi: 10.1016/0921-8890(95)00004-y

Thurstone, L. (1923). The stimulus-response fallacy in psychology. *Psychol. Rev.* 30:354369. doi: 10.1037/h0074251

Tinbergen, N. (1951). *The Study of Instinct*. Oxford: Oxford University Press.

Tobin, J., Fong, R., Ray, A., Schneider, J., Zaremba, W., and Abbeel, P. (2017). Domain randomization for transferring deep neural networks from simulation to the real world. ArXiv:1703.06907, 1–8.

Todorov, E., Erez, T., and Tassa, Y. (2012). "MuJoCo: a physics engine for model-based control," in *International Conference on Intelligent Robots and Systems* (Vilamoura), 5026–5033. doi: 10.1109/iros.2012.6386109

Todorov, E., and Jordan, M. I. (2002). Optimal feedback control as a theory of motor coordination. *Nat. Neurosci.* 5, 1226–1235. doi: 10.1038/nn963

Tolman, E. (1932). *Purposive Behavior in Animals and Men*. New York, NY: Century.

Torras i Genís, C. (1986). Neural network model with rhythm-assimilation capacity. *IEEE Trans. Syst. Man Cybern.* 16, 680–693. doi: 10.1109/TSMC.1986.289312

Turing, A. M. (1950). Computing machinery and intelligence. *Mind* 49, 433–460. doi: 10.1093/mind/LIX.236.433

Van de Burgt, Y., Lubberman, E., Fuller, E. J., Keene, S. T., Faria, G. C., Agarwal, S., et al. (2017). A non-volatile organic electrochemical device as a low-voltage artificial synapse for neuromorphic computing. *Nat. Mater.* 16, 414–419. doi: 10.1038/NMAT4856

van Gerven, M. A. J. (2017). A primer on encoding models in sensory neuroscience. *J. Math. Psychol.* 76, 172–183. doi: 10.1016/j.jmp.2016.06.009

Vanrullen, R. (2007). The power of the feed-forward sweep. *Adv. Cogn. Psychol.* 3, 167–176. doi: 10.2478/v10053-008-0022-3

Vanrullen, R. (2017). Perception science in the age of deep neural networks. *Front. Psychol.* 8:142. doi: 10.3389/fpsyg.2017.00142

Varshney, L. R., Chen, B. L., Paniagua, E., Hall, D. H., and Chklovskii, D. B. (2011). Structural properties of the Caenorhabditis elegans neuronal network. *PLoS Comput. Biol.* 7:e1001066. doi: 10.1371/journal.pcbi.1001066

Vernon, D., Metta, G., and Sandini, G. (2007). A survey of artificial cognitive systems: implications for the autonomous development of mental capbilities in computational agents. *IEEE Trans. Evol. Comput.* 11, 1–30. doi: 10.1109/TEVC.2006.890274

Vinyals, O., Blundell, C., Lillicrap, T., and Kavukcuoglu, K. (2016). Matching networks for one shot learning. ArXiv:1606.04080v1, 1–12.

Vinyals, O., Brain, G., Fortunato, M., Jaitly, N., and Brain, G. (2017). Pointer networks. ArXiv:1506.03134v2, 1–9.

von Neumann, J. (1966). *Theory of Self-Reproducing Automata*. Champaign, IL: University of Illinois Press.

von Neumann, J., and Morgenstern, O. (1953). *Theory of Games and Economic Behavior, 3rd Edn*. Princeton, NJ: Princeton University Press.

Weichwald, S., Fomina, T., Schölkopf, B., and Grosse-Wentrup, M. (2016). Optimal coding in biological and artificial neural networks. ArXiv:1605.07094v2, 1–10.

Weston, J., Chopra, S., and Bordes, A. (2015). "Memory networks," in *3rd International Conference on Learning Representations* (San Diego), 1–14.

White, R. W. (1959). Motivation reconsidered: the concept of competence. *Psychol. Rev.* 66, 297–333. doi: 10.1037/h0040934

White, S. G., Southgate, E., Thomson, J., and Brenner, S. (1986). The structure of the nervous system of the nematode *C. elegans*. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 314, 1–340. doi: 10.1098/rstb.1986.0056

Whitehead, S. D., and Ballard, D. H. (1991). Learning to perceive and act by trial and error. *Mach. Learn.* 7, 45–83. doi: 10.1007/bf00058926

Widrow, B., and Lehr, M. A. (1990). 30 Years of adaptive neural networks: perceptron, madaline, and backpropagation. *Proc. IEEE* 78, 1415–1442. doi: 10.1109/5.58323

Wills, T. J., Lever, C., Cacucci, F., Burgess, N., and Keefe, J. O. (2005). Attractor ddynamics in the hippocampal representation of the local environment. *Science* 308, 873–876. doi: 10.1126/science.1108905.Attractor

Willshaw, D. J., Dayan, P., and Morris, R. G. M. (2015). Memory, modelling and Marr: a commentary on Marr (1971) 'Simple memory: A theory of archicortex'. *Philos. Trans. R. Soc. B* 370:20140383. doi: 10.1098/rstb.2014.0383

Winograd, T. (1972). Understanding natural language. *Cogn. Psychol.* 3, 1–191. doi: 10.1016/0010-0285(72)90002-3

Wissner-Gross, A. D., and Freer, C. E. (2013). Causal entropic forces. *Phys. Rev. Lett.* 110:168702. doi: 10.1103/physrevlett.110.168702

Wolfram, S. (2002). *A New Kind of Science*. Champaign, IL: Wolfram Media.

Wu, Y., Zhang, S., Zhang, Y., Bengio, Y., and Salakhutdinov, R. (2016). On multiplicative integration with recurrent neural networks. ArXiv:1606.06630v2, 1–11.

Xue, T., Wu, J., Bouman, K. L., and Freeman, W. T. (2016). Visual dynamics: probabilistic future frame synthesis via cross convolutional networks. ArXiv:1607.02586, 1–11.

Yamins, D. L. K., and DiCarlo, J. J. (2016). Using goal-driven deep learning models to understand sensory cortex. *Nat. Neurosci.* 19, 356–365. doi: 10.1038/nn.4244

Yang, G. R., Song, H. F., Newsome, W. T., and Wang, X. (2017). Clustering and compositionality of task representations in a neural network trained to perform many cognitive tasks. BioRXiv, 1–44. doi: 10.1101/183632

Yang, W., and Yuste, R. (2017). *In vivo* imaging of neural activity. *Nat. Methods* 14, 349–359. doi: 10.1038/nmeth.4230

Yarbus, A. L. (1967). *Eye Movements and Vision*. New York, NY: Plenum.

Yuille, A., and Kersten, D. (2006). Vision as Bayesian inference: analysis by synthesis? *Trends Cogn. Sci.* 10, 301–308. doi: 10.1016/j.tics.2006.05.002

Yuste, R. (2015). From the neuron doctrine to neural networks. *Nat. Rev. Neurosci.* 16, 487–497. doi: 10.1038/nrn3962

Zagoruyko, S., and Komodakis, N. (2017). DiracNets: training very deep neural networks without skip-connections. ArXiv:1706.00388, 1–11.

Zambrano, D., and Bohte, S. M. (2016). Fast and efficient asynchronous neural computation with adapting spiking neural networks. ArXiv:1609.02053, 1–14.

Zenke, F., Poole, B., and Ganguli, S. (2015). Improved multitask learning through synaptic intelligence. ArXiv:1703.04200v2, 1–9.

Zhu, Y., Gordon, D., Kolve, E., and Fox, D. (2017). Visual semantic planning using deep successor representations. ArXiv:1705.08080v1, 1–13.

Zipser, D., and Andersen, R. (1988). A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature* 331, 679–684. doi: 10.1038/331679a0

Conflict of Interest Statement: The author declares that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.