





Viewpoints

Deep(er) Learning

 Shyam Srinivasan,^{1,2}  Ralph J. Greenspan,^{1,3,4}  Charles F. Stevens,^{1,2} and  Dhruv Grover¹

¹Kavli Institute for Brain and Mind, University of California–San Diego, La Jolla, California 92093, ²Molecular Neurobiology Laboratory, Salk Institute for Biological Studies, La Jolla, California 92037, ³Division of Biological Sciences, University of California–San Diego, La Jolla, California 92093, and ⁴Department of Cognitive Science, University of California–San Diego, La Jolla, California 92093

Animals successfully thrive in noisy environments with finite resources. The necessity to function with resource constraints has led evolution to design animal brains (and bodies) to be optimal in their use of computational power while being adaptable to their environmental niche. A key process undergirding this ability to adapt is the process of learning. Although a complete characterization of the neural basis of learning remains ongoing, scientists for nearly a century have used the brain as inspiration to design artificial neural networks capable of learning, a case in point being deep learning. In this viewpoint, we advocate that deep learning can be further enhanced by incorporating and tightly integrating five fundamental principles of neural circuit design and function: optimizing the system to environmental need and making it robust to environmental noise, customizing learning to context, modularizing the system, learning without supervision, and learning using reinforcement strategies. We illustrate how animals integrate these learning principles using the fruit fly olfactory learning circuit, one of nature’s best-characterized and highly optimized schemes for learning. Incorporating these principles may not just improve deep learning but also expose common computational constraints. With judicious use, deep learning can become yet another effective tool to understand how and why brains are designed the way they are.

Introduction

At the cusp of a new technological era ushered in by the resurgence of artificial intelligence, machines are now, more than ever, close to achieving human levels of performance (LeCun et al., 2015; Goodfellow et al., 2016). This performance breakthrough is due to a new class of machine learning algorithms called deep learning (LeCun et al., 2015), whose architecture is closer (though still not the same) to that of animals. It has allowed machine learning to become part of our everyday lives, to the extent that we are unaware of even using it, from voice assistants on our smartphones to online shopping suggestions.

Deep learning (LeCun et al., 2015; Goodfellow et al., 2016), a descendent of classical artificial neural networks (Rosenblatt, 1958), comprises many simple computing nodes organized in a series of layers (Fig. 1). The leftmost layer forms the input, and the rightmost layer or output spits out the decision of the neural network (e.g., as illustrated in Fig. 1*a*, whether an image is that of Albert Einstein). The process of learning involves optimizing connection weights between nodes in successive layers to make the neural network exhibit a desired behavior (Fig. 1*b*). This optimization procedure moves backwards through the network in an iterative manner to minimize the difference between desired

and actual outputs (backpropagation). For each layer, errors are minimized at every node one weight at a time (gradient descent). In deep learning, the number of intermediate layers between input and output is greatly increased, allowing the recognition of more nuanced features and decision-making (Fig. 1*a*).

Artificial neural networks originated with the need for algorithms to classify images. Frank Rosenblatt, in the late 1950s, had the ingenious idea of using nervous system-like networks that he called “perceptrons” (Rosenblatt, 1958) to solve this problem. He built on earlier ideas by McCulloch and Pitts (1990), who used a linear threshold function (the neuron is activated when the input exceeds a threshold) to artificially simulate neurons. Rosenblatt felt that the brain is an existence proof model that we can build a system that can sense the environment and take appropriate action. Even at this early stage, he envisioned that perceptrons would one day, “walk, talk, see, write, reproduce itself and be conscious of its existence” (Olazaran, 1996). Although original and transformative, perceptrons had limitations as they could only classify data that were linearly separable (e.g., in two dimensions, data would be separated by a line, in three dimensions by a plane, and so on). This problem was, however, solved with the invention of multilayer perceptrons (Grossberg, 1973) and nonlinear activation functions.

The innovations and discoveries that paved the way for current-day neural networks began in the 1980s. The classification problem was essentially one of mathematical optimization and involved finding connection weights that gave the highest fraction of correct results during training. An issue with optimization problems is that they often get stuck in local minima (the best solution for a limited set of conditions), whereas the solution we desire is a global minimum (i.e., the best solution under all possible conditions). The advent of Hopfield networks

Received Feb. 27, 2018; revised June 30, 2018; accepted July 5, 2018.

This work was supported by National Science Foundation NSF PHY-1444273 to S.S. and C.F.S.; and Air Force Office of Scientific Research FA9550-14-1-0211 to D.G. and R.J.G. We thank Terrence Sejnowski, Jean-Pierre Changeux, David Peterson, Saket Navlakha, and Jen-Yung Chen for helpful discussions and comments.

The authors declare no competing financial interests.

Correspondence should be addressed to either of the following: Dr. Dhruv Grover, Kavli Institute for Brain and Mind, University of California–San Diego, 9500 Gilman Drive, MC 0126, La Jolla, CA 92093-0126, E-mail: dhgrover@ucsd.edu; or Dr. Charles F. Stevens, Molecular Neurobiology Laboratory, Salk Institute for Biological Studies, 10010 North Torrey Pines Road, La Jolla, CA 92037, E-mail: stevens@salk.edu.

DOI:10.1523/JNEUROSCI.0153-18.2018

Copyright © 2018 the authors 0270-6474/18/387365-10\$15.00/0

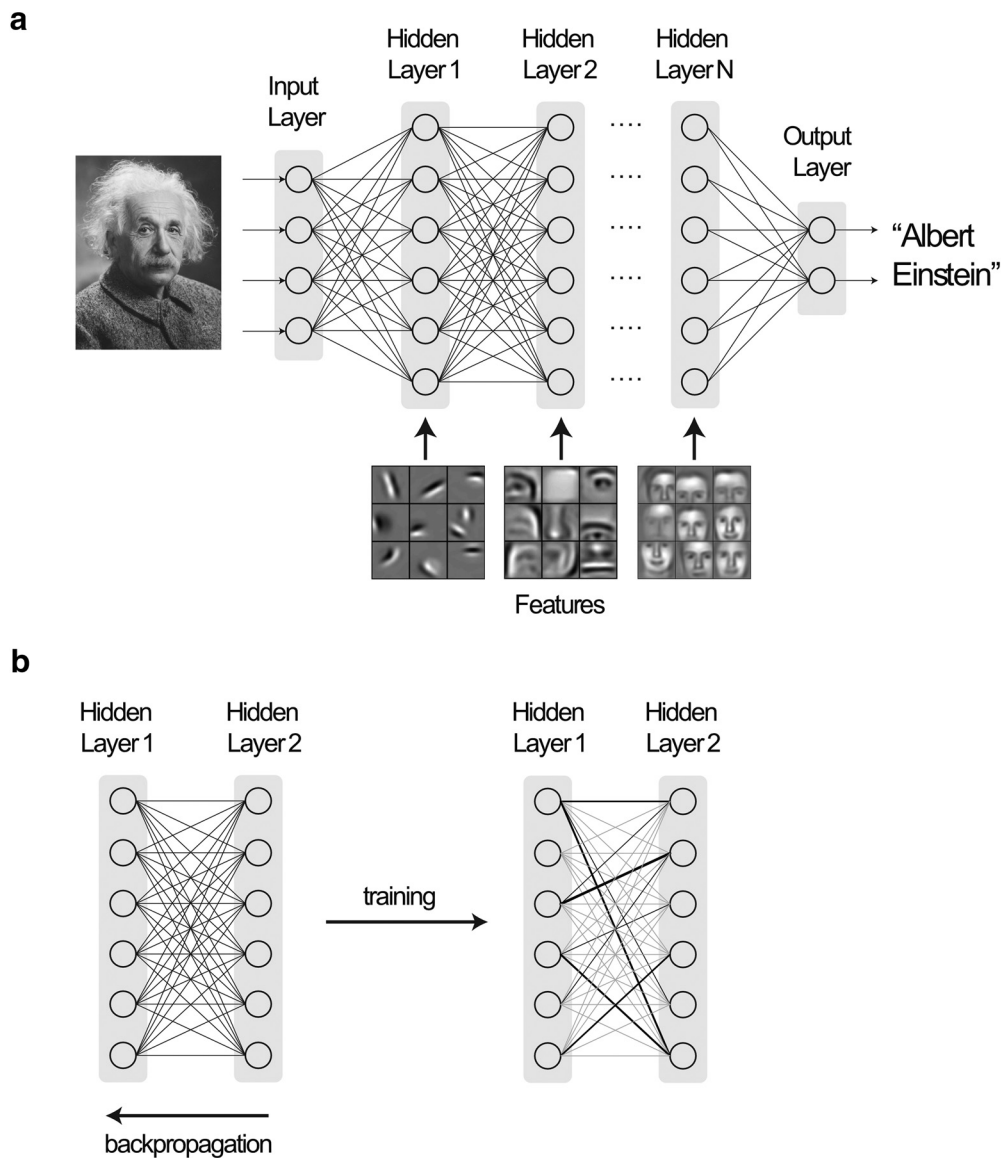


Figure 1. A schematic of a deep learning neural network for classifying images. *a*, The network consists of many simple computing nodes, each simulating a neuron, and organized in a series of layers. Neurons in each layer receive inputs from neurons in the immediately preceding layer, with inputs being weighted by the connection strengths between neurons. Each neuron is activated when the sum of input activity exceeds a threshold, and in turn contributes to the activity of neurons in successive layers. In the figure, the leftmost layer encodes the input, in this specific case, faces. The rightmost layer spits the output, in this case, whether the photo is that of Albert Einstein. The weights between neurons are pruned and perfected by training with millions of labeled trial faces. During each trial, the connection weights are adjusted by backpropagation to produce the right output. After sufficient training, the network evolves to a point where each successive layer in the neural network learns to recognize more complex features (e.g., from lips, nose, eyes, etc., to faces) and classifies correctly. The layer features are adapted from Lee et al. (2009). *b*, Schematic indicating how connection weights between successive layers are altered during training by the backpropagation algorithm to minimize error and produce the right output. This alteration proceeds backwards through the network using a gradient descent approach.

(Hopfield, 1982) showed how content-addressable memory networks could be used to escape a local minimum, thus increasing the range of optimization problems that could be solved. Subsequently, another study showed that, by using a specialized optimization cost function, it was possible to nearly achieve a global minimum (Ackley et al., 1985). What then emerged, in one of the most important discoveries, was a simple process known as backpropagation that reduced errors and converged quickly onto an optimized state (Werbos, 1974; LeCun, 1985; Parker, 1985) and forms the heart of today's deep learning algorithms. An influential 1986 study then experimentally demonstrated the effectiveness of backpropagation in multilayered networks, contributing significantly to its popularization (Rumelhart et al., 1986). Not coincidentally, algorithms that could passably pro-

duce human-like outputs, such as Nettealk (Sejnowski and Rosenberg, 1988), started making their appearance following these innovations. Although current deep learning systems are more sophisticated, the basic engines are not drastically different from earlier ones. Despite the impressive progress, there remained constraints on the performance of deep learning, which were ameliorated only recently because of the exponential increase in computing power.

Computer scientists distinguish among three types of machine learning algorithms, and deep learning algorithms discussed above fall under a category known as supervised learning. It contrasts with two other types, which are called reinforcement and unsupervised learning. These algorithms can be easily illustrated with a real-world example (Sutton and Barto, 1998; Barto

et al., 2004). Consider a scenario where a tourist visits a city for the first time and would like to find the best restaurants for a certain cuisine. One option to generate a list of those that were good would be to randomly pick places. Another option could be to ask a local for their list of recommended places. In the first case, the tourist would arrive at a list by trial-and-error exploration. Although we do not know going in if a restaurant might meet our expectations, we can judge its quality after sampling the food. We observe features that might be representative of the quality of the food (e.g., was there a line out the door, were there locals eating there, etc.), and we might use these indicators when choosing to eat at another place. In essence, the quality of the food reinforces a feature's efficacy, and this type of learning is termed reinforcement learning. In the alternative scenario, we would be given a list of restaurants that have stellar reviews, and by visiting them we can determine for ourselves what makes them great. This form of learning where we develop an intuition for selecting a quality establishment based on a set of examples is supervised learning. With the third category of unsupervised learning, we are given the names of restaurants and asked to classify them. It is up to us to recognize and explain the defining features of a class (e.g., whether a restaurant falls in the category of American diners or taquerias). These three types of learning in machines are subtly different from the two forms of animal learning, classical and operant conditioning, that we describe below.

In the following sections, we first explore the constraints on the performance of deep learning, the reasons for its success and drawbacks. Next, we describe animal learning strategies, with a specific focus on fly olfactory learning. We then suggest ways in which animal brains could inspire further improvements to deep learning. And finally, we propose a fly-inspired algorithm suitable for diverse classification and recognition problems.

The future of deep learning

The first constraint on deep learning is the amount of training data. The bigger the universe of possible scenarios that are fed to the computer, the more it can learn, and the better its predictive performance in associating an unseen example to something already encountered. The second factor is the size of the network. The neural network collectively encodes the rules of association. The bigger the network, the more rules it can store. The third constraining factor is that a massive amount of computing power is required to accommodate the first two factors. The principal reason behind deep learning's success is that it overcame these constraints because of three advances: the availability of huge databases of training examples on the internet, improvements in activation functions to optimize performance, and the emergence of powerful graphical processing units, courtesy of Moore's law, for crunching through huge quantities of data (Koch, 2016).

Because of these advances, when a deep learning algorithm runs on powerful graphical processing unit hardware architecture, it learns associations between features in the data, which it then uses to make decisions. It adapts without the need for explicit programming. It is this feature that allows deep learning to shine when it comes to complex real-world problems. But deep learning's dependence on computational power also exposes one of its weaknesses. It may need to be deployed on machines that have constraints on processing power and network connectivity to the Cloud (edge computing), which nevertheless have to solve real-world tasks, such as image classification, natural language processing, and speech recognition. These three specific tasks have already been perfected by another architecture with parallels to deep learning's layered architecture: animal brains. Animal

brains or natural deep learning systems have been designed by evolution over millions of years to be adapted to the specific task at hand in the face of often constrained resources. Given this level of exquisite optimization, it is likely that we can exploit our knowledge of animal learning architectures to optimize deep learning algorithms.

The best-known example of the nervous system's ability to learn and incorporate experience in decision-making is Pavlovian conditioning. In 1889, Ivan Pavlov showed that a dog repeatedly exposed to the sound of a bell and the aroma of food would salivate to the bell, even in the absence of food (Striedter, 2015). This behavior showed that the dog's brain had created an association between the bell sound and food. The experiment highlights a property central to animal learning: the ability to associate a conditioned stimulus (bell sound) and unconditioned stimulus (food odor) based on experience (repeated trials). The way the brain does Pavlovian conditioning is at the root of much of our learning and memory.

In the century since Pavlov, such conditioning has been observed in myriad animals, ranging from humans to mice and even fruit flies (Tully and Quinn, 1985). The quest has been to understand the molecular, cellular, and circuit-level mechanisms that enable Pavlovian conditioning. One of the best systems where studies have mapped and linked mechanisms at these three levels is the fruit fly. The fly's brain is numerically simpler than mammals, has excellent genetic malleability, and its architecture is known. Researchers have shown that a fly simultaneously exposed to a previously neutral odor and punishment (electric shock) or reward (sugar) remembers the association for long periods of time and avoids or approaches the odor source on subsequent exposures (Tully and Quinn, 1985). In the next section, we build a picture of how the fly learning system works and how this knowledge can inspire improvements to machine learning.

Fly associative learning

Neuroscientists have identified the two circuits responsible for implementing olfactory learning behavior in the fly: the olfactory circuit that senses odors and the higher processing region that associates odors with punishment or reward (i.e., valence). Although a full functional characterization of these circuits is ongoing, much of their architecture and function is known (Aso et al., 2014a; Oswald and Waddell, 2015; Shih et al., 2015). Most odors are mixtures of many chemicals. Mammals and flies use a set of olfactory sensory neuron (OSN) types to detect individual chemical components of these odors. The odors are encoded as 50 dimensional objects by the fly's olfactory system as it has 50 OSN types in its nose, the antennae, and the maxillary palp (Hallem et al., 2004). OSNs pass odor information onto glomeruli (neuropil) at the next stage of the olfactory circuit in the antennal lobe (AL) (Vosshall et al., 2000; Wilson, 2013). Because each glomerulus is OSN-type specific, odors can be distinguished by the unique set of OSN types or glomeruli they activate. Glomeruli, in turn, send odor information, via projection neurons, to the next stage of the circuit, the mushroom body (MB), where information is assembled in a form suitable for learning and memory formation (Eichler et al., 2017).

Unlike the first two stages, AL-MB connections are not topographic. Glomeruli, via projection neurons, seem to contact Kenyon cells (KCs) distributed over the MB without any spatial preference (Caron et al., 2013). While this distributed connectivity is stable over the long term, it differs across animals and even across the two halves of the brain. Although there might be hid-

den structure in connectivity patterns between glomeruli and MB neurons, an analysis of connectivity has shown that it is not significantly different from instances where connections are assigned randomly (Caron et al., 2013). For practical purposes, these connections can be mimicked by a probability distribution. This distributed connectivity is reflected in MB responses to odors, wherein activated KCs are spread over the MB without any preference (Fig. 2) (Turner et al., 2008; Aso et al., 2014a).

Such representation carries three advantages: (1) a potentially large number of odors can be encoded with just 2000 neurons; (2) creating odor tags by activating a sparse set of neurons that seem randomly distributed across the MB ensures that two distinct odors are highly likely to have nonoverlapping odor tags (i.e., they are discriminable) (Lin et al., 2014; Hige et al., 2015; Stevens, 2015); and (3) despite random assignment, theoretical analyses assuming linear rate firing models suggest that no odor information is lost (Stevens, 2015). If neuronal activation is nonlinear, it is possible that some odor information might be lost. Experimental evidence so far suggests that most of the odor information is maintained (Campbell et al., 2013). Thus, while distinct odors have little overlap in activity patterns in the antennal lobe and MB, similar odors will have significant overlap (Schaffer et al., 2018; Srinivasan and Stevens, 2018). This scheme of using sparse and distributed tags for encoding information also affects how more complex objects, such as the association between an odor and valence (e.g., punishment) inputs (Han et al., 1996).

The encoding of odor-valence associations is handled by the circuit downstream from the MB and is responsible for driving the fly's approach or avoidance behavior (Fig. 2) (Aso et al., 2014a; Hige et al., 2015; Oswald and Waddell, 2015; Waddell, 2016). The circuit comprises three components. The first component is the collection of MB output neurons (MBONs), which receive the outputs of KCs within 16 segregated compartments. There are 21 types of MBONs, and each one receives input (mostly) from 1 of the 16 compartments (Fig. 2, line of synapses between a dopaminergic neuron [DAN] and an MBON). These compartments form the second component. Within each compartment, a KC's axonal output synapses (communicates) with the dendrites or inputs of MBONs: a KC usually innervates multiple compartments contacting MBONs in each compartment. MBONs collectively (by increased and/or decreased activation of certain subsets) supply information contributing to an approach or avoidance behavior.

Detailed measurements of two such compartments have revealed that MBON dendrites arborize the compartment, making multiple synapses with each innervating KC (Takemura et al., 2017). A careful analysis has shown that the number of synapses (20 and 28 on average) that each KC makes with an MBON follows a Poisson distribution. Thus, it is highly likely that an odor will activate ~5% of synapses, which will be distributed throughout the compartment (Takemura et al., 2017), recapitu-




Figure 2. Fly olfactory associative learning schematic. Odor information from the fly's nose filters down to the MB, activating a sparse set of KCs that code for a particular odor. These KCs, in turn, synapse with MBONs that influence approach or avoidance behavior. During olfactory conditioning, simultaneous presentation of an aversive stimulus, such as an electric shock, and a neutral odor activates the DAN for shock and KCs encoding the odor. The DANs release dopamine, which modulates the KC/MBON synapse (inset). Repeated training alters the strength of the synapse over the long term. Thus, following training, odor presentation alone is sufficient to activate the MBON that influences an aversive response. Black circles represent active neurons and synapses. Gray circles represent inactive neurons and synapses. Each line of synapses between a DAN and MBON indicates a separate compartment. Increasing the gain of the hypothesized recurrent circuit from MBONs to DANs (as indicated by a dotted line) reduces the amount of training required for learning.

lating the distributed organization and activation of odor tags. Consequently, the advantages of distributed coding in the MB that we outlined above for KCs apply to odor-valence tags too (e.g., the odor-valence tags for two distinct odors are unlikely to overlap, like their KC odor tags). An important caveat is that Takemura et al. (2017) also showed that, in one compartment, which receives non-odor-coding KC axons, MBON dendrites were segregated into subcompartments. This suggests a possible mechanism for segregating learning and memory across sensory modalities.

The third component, the one that creates an association, is a set of 20 types of DANs, which also innervate all compartments. Most DANs convey a specific type of valence; for instance, there is a DAN activated by electric shock, and another that is activated by sugar. When the fly is shocked, the DAN that responds to electric shock gets activated and releases dopamine in the compartment that it innervates (Cohn et al., 2015; Hige et al., 2015). Dopamine has the quality of modifying synapse efficacy, and thus all KC-MBON synapses activated by a particular odor get modified. When done repeatedly, the synapses stay altered over a long term, leading to an association between odor and valence. Ultimately, presentation of the odor alone leads to an avoidance behavior (Fig. 2). Notably, the MBON-DAN-KC network includes feed-forward loops (Aso et al., 2014a) that have yet to be functionally characterized. Here, too, detailed measurements have shown that DAN axons, at least onto KCs, mimic a Poisson distribution. But as dopaminergic modification of KC synapses occurs by volume transmission, DANs would influence nearly all KC synapses.

A notable feature of this circuit and the DANs in particular is that they act as negative or positive reinforcers. For example, as shown in the experiments above, repeated shocking in the pres-

ence of a particular odor progressively weakens the KC-MBON synapse, thereby reinforcing the negative association. It is likely that a similar scheme would reinforce positive associations. This reinforcement ability is fundamental to another form of learning in addition to Pavlovian conditioning, called operant conditioning, which involves learning from the consequences of our actions rather than associations between events. Operant conditioning was first postulated by the psychologist Edward Thorndike in 1908 when he shut a cat in a box that could only be opened by the cat pressing a lever. Thorndike found that the cat's ability to open the box improved over trials as the efficacy of pressing the lever was reinforced, leading to his theory on the law of effect. This theory was taken up by behavioral psychologists, such as Skinner (Skinner, 1963), and has also been shown to operate in fruit flies (Wolf and Heisenberg, 1991), implicating the MB circuitry we described above (Heisenberg, 1998; Liu et al., 1999). Thus, the MB-MBON-DAN circuitry presents a highly flexible machinery for both associative and operant-based conditioning.

Coincidentally, operant conditioning, like other forms of learning, has also made its way to computer science, where it is known as reinforcement learning. In reinforcement learning, the function that is learned is value-based, unlike supervised learning, which is classifier- or predictor-based. With this definition in mind, both operant and classical conditioning are instances of reinforcement learning as both use value-based labels (e.g., heat or sugar rewards) (Barto et al., 2004). However, in classical conditioning, the experimenter controls the association to be formed akin to the training phase of a supervised learning algorithm in computer science (Barto et al., 2004). Thus, the animal conditioning paradigms that we have described do not neatly fit into the three categories of learning envisioned by computer scientists. These nuanced differences might be the reason why animals and machines have so far excelled at different computations.

What we have outlined is nature's scheme to link a wide variety of states to an equally wide variety of valences depending on context. There are three signatures of this circuit. One is that the stimulus (odor) is encoded by a sparse and distributed neuronal ensemble. Second, the connectivity into these neurons and between them and the downstream MBONs seems to be distributed without any spatial preference. Third, downstream synapses onto MBONs are plastic, providing a flexible way to associate stimuli and valence. This network motif has been elaborated by evolution in other circuits and brain regions across myriad species.

The circuit closest to the MB in terms of structure and function is the primary olfactory cortex in mammals (Bekkers and Suzuki, 2013) and its fish homolog, the dorsal posterior telencephalon (Yaksi et al., 2009). Like the MB, connectivity into and from olfactory cortex neurons is distributed without any spatial preference, and this connectivity is reflected in the activity patterns that are similarly distributed and sparse. Moreover, the association layer formed by the outputs of piriform neurons has shown evidence of synaptic plasticity that is driven by descending inputs from the orbitofrontal cortex (Strauch and Manahan-Vaughan, 2018).

The next best example of a brain region with the MB motif is the cerebellum and cerebellar-like circuits across vertebrate species (Marr, 1969; Oertel and Young, 2004; Bell et al., 2008; Farris, 2011). Here, Purkinje cells function as the output neurons of the circuit and receive input from deep cerebellar nuclei via granule cells. Each granule cell extends parallel fibers into the molecular layer where they contact Purkinje cells along the foliol axis. These synapses seem randomly distributed and intriguingly show evi-

dence of synaptic plasticity of the form encountered in MB compartments: associations are formed by the weakening of synapses between granule and Purkinje cells (Bell et al., 2008). This intricate architecture, also termed the Marr motif (Marr, 1969; Stevens, 2015), has been shown to function as an adaptive filter; for example, in weakly electric fish, it allows the fish to filter out the electrical field imbalances caused by itself and isolate external electrical signals from the environment (Bell et al., 2008).

Hallmarks of the MB are also found in three other regions involved in learning: the hippocampus and prefrontal cortex in mammals, and the vertical lobes in molluscs. In the hippocampus, input from the entorhinal cortex to CA3 mirrors the connectivity from the AL to MB, and activity in the CA3 region is sparse and distributed (Haberly, 2001). Additionally, CA3 contains a winner-take-all circuit just like the MB and olfactory cortex. Similar activity patterns are also observed in the prefrontal cortex, and studies have suggested that these circuits perform similar computations as the olfactory cortex and MB (Barak et al., 2013). Last, the connectivity of the vertical lobes in molluscs with their fan-in fan-out networks that experience synaptic plasticity resembles the MB (Shomrat et al., 2011).

Finally, the way the fly association network encodes reward and punishment has parallels with mammalian systems. In the fly, the MB, which encodes odors, coordinates with DANs, the superior medial protocerebrum, crepine neuropil, and lateral horn regions to drive reward and motivated behavior. Similarly, in mammals, the olfactory cortex, which encodes odors, communicates with the basal ganglia (substantia nigra, nucleus accumbens, dorsal raphe nucleus) and prefrontal (orbitofrontal) cortex to drive similar behavior (Scaplen and Kaun, 2016; Waddell, 2016).

The fly learning system contains all of ~2000 neurons (Aso et al., 2014a), housed within the brain that contains on the order of 100,000 neurons: a trivial number compared with the computational power of current day deep learning machines, such as Google translate or AlphaGO (Silver et al., 2016). Even though flies live for only a few weeks, their sophisticated behavioral repertoire would be the envy of any artificial intelligence system. They are capable of remarkable aerobic feats, exhibit both sensory-evoked and socially evoked behaviors in a complex world, and are capable of higher cognitive processes, such as sleep, attention, various forms of associative conditioning, and in some cases even abstract generalization. Therefore, it is quite reasonable to expect that animal systems, such as the humble fruit fly, might yet have a few tricks up their sleeve that might benefit machines.

Advantages of natural deep layered architectures

We now highlight five features of animal systems that might serve as inspiration for improvements to deep learning and artificial neural networks in general.

Performance

Compressing memory

Machines face constant increases in size and complexity of data to be processed. A possible solution, for efficient processing, could come from the way flies and other animals store a huge amount of information with a small set of neurons, by using a randomly distributed connection network (Ganguli and Sompolinsky, 2012; Stevens, 2015).

The potential number of odors that a fly can encode is $\sim 2^{50}$, as there are 50 types of OSNs. The number of neurons dedicated to storing this information in the MB is, however, only 2000. Fly MBs serve as a compressed storage space because evolution has

designed an algorithm resembling a technique pioneered by mathematicians in the field of random projections called compressed sensing (Candès et al., 2006a; Donoho, 2006; Ganguli and Sompolinsky, 2012). Whenever a signal is transmitted between two vector spaces (e.g., AL space to KC space), if the input signal is sampled at random (i.e., multiplied by a random matrix), the resulting signal is called a random projection. When the input signal is sparse and the random matrix has certain specific properties (e.g., a Gaussian matrix), the random projection can be much smaller while still retaining all the information present in the input signal: hence the term compressed sensing. Because the connection matrix from the AL to MB seems randomly distributed, the MB odor tag is akin to a compressed sensing measurement, providing a compact representation of the odor, and in turn, an efficient storage scheme.

Robustness of classification

The fly MB distributed connection network (Caron et al., 2013; Takemura et al., 2017) has another advantage. As theoretical studies have shown, such a network robustly maintains information transmitted between layers (Candès et al., 2006b; Barak et al., 2013; Krishnamurthy et al., 2017; Litwin-Kumar et al., 2017). In the fly and the mammalian olfactory circuit, the connection network mimics a Poisson-Gamma distribution (Stevens, 2015; Srinivasan and Stevens, 2018). This distribution ensures that, while most neurons have low-level activity, a small fraction of the population is highly active for any particular odor. Further, evolution has used another trick to ensure robust classification by suppressing the activity of all but the top 5%–10% of neurons in the population. In the fly, this is implemented as a negative feedback loop between the MB KCs and the inhibitory anterior-paired-lateral neuron (Papadopoulou et al., 2011; Lin et al., 2014). As more KCs are activated, they receive more inhibition from the anterior-paired-lateral neuron, thereby maintaining a sparse odor code. As this sparse neuronal population is randomly distributed (because of the connection network), the chance of overlap between distinct odors is extremely low.

These two mechanisms ensure that the fly can robustly classify odors, even in the presence of transmission noise. Such efficient and robust classification schemes can be easily incorporated into existing algorithms. A recent paper took this precise network and showed theoretically and practically (Dasgupta et al., 2017; S. Navlakha, personal communication) that it could be used to improve search algorithms.

Pruning

In deep learning, a naive system starts with an all-to-all network and lets training prune the connections to improve decision-making. Pruning is, however, a computationally expensive task. With more layers, the computational requirement increases (nonlinearly), and so does the time for propagation of information between distant nodes (Hochreiter et al., 2001). If deep learning is to be deployed on machines that have constraints on processing power and network connectivity to the Cloud, computational requirements must be optimal.

The way the synapses between odor encoding KC neurons and downstream decision-making circuit neurons are pruned in the fly circuit suggests a solution to this issue. These synapses receive DAN input associated with a punishment or reward (Aso et al., 2014b). If the concept of DANs was to be implemented in deep learning, these DANs would feed directly into intermediate layers bypassing the input layer. They would adjust the connections between the sets of computing nodes in different layers, acting as

a knob that can be tuned to adjust the level of learning. A variation of this scheme might reduce some of the computational costs involved with pruning and backpropagation.

Scaling

Another lesson that animal systems can provide toward improving performance is in determining the optimal size of the system required for the task at hand. Evolution has provided an existence-proof example of how this can be done with the visual system, which comes in a range of sizes depending on the animal's environmental niche (Stevens, 2001; Srinivasan et al., 2015). Some (e.g., mice or rats) have small visual circuits and limited visual acuity, whereas others (e.g., primates) have enormous visual circuits and correspondingly high visual acuities (Srinivasan et al., 2015). Remarkably, the number of neurons and visual acuity can be accurately predicted for any system. This phenomenon, known as scaling, has two properties. First, the architecture of the visual circuit (i.e., the hardware) is similar across mammals. Second, the number of processing elements in the circuit and visual ability have the same (power-law) relationship across different brain sizes. Therefore, we can determine the optimal size for any task at hand because we know the design constraint (i.e., the size-performance relationship). Such design constraints would help build optimally sized machine learning systems.

Context-dependent learning

Evolution has designed three forms of memory needed for an animal to thrive in the natural world. The first operates over evolutionary timescales, containing mechanisms hardwired in the brain as they are integral to survival and propagation of the species. A fitting example is the response to pheromones, wherein characteristic scents are programmed to elicit a stereotypic behavior (e.g., approaching potential mates or avoiding predators). These responses are encoded by circuits whose connections are similar across individuals (e.g., in mammals); pheromone-sensing neurons target stereotyped locations in the amygdala. In flies, pheromone-sensing neurons contact the MB along with stereotyped locations to the lateral horn. The parallels in terms of stereotypic organization suggest that the lateral horn might drive similar programmed behaviors in flies (Keene and Waddell, 2007; Nagayama et al., 2014; Root et al., 2014). Intriguingly, a recent study suggests that the lateral horn also plays a role in mediating some learned behaviors (Dolan et al., 2017).

The next two forms of memory are dynamic and implemented by mechanisms that enable the animal to learn and respond appropriately to situations in an ever-changing environment: (1) mechanisms that enable learning from very few experiences (e.g., we quickly learn to avoid fire); and (2) mechanisms that require repeated exposure (e.g., musicians practice for years to become exceptional). This phenomenon, first identified by Herbert Simon (Chase and Simon, 1973; Charness et al., 2005) has since become popular as the 10,000 hour rule. The fly olfactory association network suggests how a system can be adapted for both forms of learning.

In the fly, some associations are learned with a single trial (Ichinose et al., 2015; Lewis et al., 2015). For one such association involving sugar reward learning (Ichinose et al., 2015), researchers showed that the axons of the MBON and the dendrites of the DAN that implement learning, overlap. They suggested that a recurrent circuit between the DAN and MBON (Fig. 2) leads to increased activation of the DAN, which further alters the strength of KC-MBON synapses, mimicking effect of multiple trials with

just one trial. Although further experiments showing synaptic connections from MBON axons to DAN dendrites are required, the study presents an intriguing mechanism for implementing one trial learning in a deep learning network. Recent research in machine learning has explored the use of feedback to dynamically change learning capability. This includes variations of recurrent neural networks that include a memory component (Weston et al., 2014). Interestingly, a recurrent circuit was not found in the second association network (Lewis et al., 2015), suggesting the presence of alternative or additional mechanisms for implementing one-shot learning.

Such one-shot learning algorithms are apt for tasks where event sequences have structure (e.g., handwriting and speech recognition). But even these algorithms do not tailor the amount of training to the event's relative importance (e.g., life-threatening vs rote tasks). Most of our learned experiences fall somewhere in between avoiding fire and becoming a musician. With many of these activities, it is possible to shorten or lengthen the learning time, and a key facilitator of this process is emotion.

The relation between emotion and learning is a key differentiator between animals and machines. A large cohort of studies have shown that emotion can effectively tune our ability to learn. Experiences that emotionally engage us elicit greater levels of attention, which lead to quicker and longer-lasting memories. Furthermore, as emotions go, painful experiences seem to have more salience for memories than joyful ones (Kahneman and Tversky, 1982). This might be a byproduct of evolution as the value of learning an experience that could be lethal might be the difference between life and death. While deep learning variations, such as one shot (Fei-Fei et al., 2006) and convolutional networks (Krizhevsky et al., 2012), have incorporated concepts of faster learning, the next step would be to mimic animal learning systems, wherein a system that flexibly learns is paired with one that assigns priorities to the learning tasks (emotional centers). The fly associational network provides a blueprint for building such a neural network.

Priming

Computers find it challenging to learn and generalize from examples without the need for supervision. Here again, animal systems provide a solution, which might come from understanding the neural basis of a phenomenon known as priming.

Priming, discovered by cognitive neuroscientists in the 1970s (Tulving et al., 1990; Schacter and Buckner, 1998), was observed when subjects were asked to perform word associations. They were first exposed to a series of words without any instructions being given about what the words meant or what they had to do. When later asked to make associations between words, researchers found that word associations were much easier for those words that subjects had previously encountered, even though there was no contextual (instructional) underpinning to the exposure.

A recent study in the fly sheds light on the kind of molecular circuitry that might be used in priming. Not surprisingly, it involves the fly association network outlined previously (Hattori et al., 2017). In this case, however, the odor serves both as a conditioned and unconditioned stimulus. The odor activates the specific set of KCs in the MB for that odor, which in turn drive the activity of a specific MBON. The odor also activates the DAN that innervates this compartment, and repeated exposure to the odor weakens the synapse between the KCs and the MBON. This results in the MBON not being activated after repeated exposure. In this way, without any explicit instruction, the fly can distinguish

previously experienced odors from new ones, providing a mechanistic basis for priming and a possible algorithm for machines to implement it.

Specialization of function in animal systems

Fly and other complex brains are highly modular. They have specialized circuitry for capturing and extracting relevant features of the environment. This modularity extends to centers of the brain that enable higher cognitive function. For instance, in the fly, the MB integrates information from olfactory and visual circuits and then collaborates with valence and decision-making centers in the brain. Correspondingly, in mammals, peripheral sensory information forms one part of the input, along with information from memory and valence and motivation signals from the prefrontal cortex and the basal ganglia. In general, modularization has been proposed (Striedter, 2005) to be an efficient way for brains to solve complex problems by breaking them up into simpler components.

Machines, too, have used this feature of modularity to improve performance. An excellent example comes from comparing AlphaGO (the deep learning system that beat Lee Sedol, the GO champion) to its predecessors. Unlike its predecessors, AlphaGO used two networks to get better at playing GO (Koch, 2016; Silver et al., 2016): a value network and a policy network. The value network gives the reward value of a particular move, and the policy network chooses, from among all possible moves, the one that gives the highest reward. Both ran deep learning algorithms on similar hardware (but geared to different objectives) that in concert produced a champion GO player.

In the future, there will be a need for deep learning systems to handle more complex tasks than playing GO (e.g., a driverless car). It will comprise many subsystems that need to collectively collaborate to achieve a goal. The brain's architecture, and the way its individual substructures encode and exchange information, is a template for more holistic deep learning systems.

Reinforcement learning

Reinforcement learning, in both animals and machines, is a learning process that uses previous experiences to improve future outcomes. The implementation of reinforcement learning in machines can be broken down into three major components (Sutton and Barto, 1998). The first, called a policy module, oversees what the machine will do next based on its current state and input (i.e., how it should handle every possible situation). The second, the reward module, gives a numerical value for every possible state that the machine can take. The machine then attempts to choose transitions that land it in a state with the highest reward, although there is a caveat. A high reward in the short-term may not necessarily lead to a high reward in the long-term. To guard against this, there is a third and probably the most important module called the value module. It gives the long-term value of being in any particular state. Thus, the decision of transitioning to a different state is made not by the state that gives the highest reward but by the state that has the highest value. For instance, while playing chess it might be okay to kill a pawn for a win in the longer term. Furthermore, these states are continually reestimated as the machine experiences more of the world.

Although machines use the policy module in making decisions on which state to transition to next, they occasionally make random choices. Such choices enable them to fully explore the environment and sample hitherto untested states that might yield higher reward values in the short- or long-term.

The efficacy of reinforcement learning becomes apparent if we again consider the AlphaGO system. Unlike previous iterations of deep learning systems, AlphaGO used reinforcement learning to learn and improve its game strategy (Silver et al., 2016). Before each move, it played the remainder of the game many times over using a strategy known as Monte-Carlo tree search. Using a deep neural network to guide its search, AlphaGO chose the move that was most successful in simulation. To do this, it used two networks: a policy network to suggest the moves it should make and a value network to evaluate the reward (winning the game) for the move. Both the policy and value networks were improved during initial training by using trial-and-error reinforcement learning. It is this incorporation of the consequences (as probabilities) of every move with almost limitless memory and computing resources that resulted in Lee Sedol's defeat. Thus, the use of reinforcement learning was instrumental in improving the performance of a GO playing machine, and its computational effectiveness might be one reason for evolution to have used this strategy in animals.

The reinforcement learning of AlphaGO differs from reinforcement learning in flies and other animals in one important aspect. AlphaGO's implementation closely resembles temporal difference learning (Sutton and Barto, 1998), wherein learning is a multistep process and learning at each step depends on the difference between expected and actual rewards. This multistep learning process bears some resemblance to how we learn: we can remember a sequence of events that leads to a reward as opposed to a single event. The MB structure in flies, which exhibits operant conditioning, has parallels with one-step temporal difference learning, but whether a fly is capable of multistep learning, such as AlphaGO, remains to be seen. Interestingly, other insect species, such as bees and ants (e.g., mammals), do possess multistep learning mechanisms and use it for food search and navigation (Collett, 1996; Menzel et al., 2005; Neuser et al., 2008; Ardin et al., 2016; Schwarz et al., 2017).

Recent work in machine learning has outlined how one can take advantage of a network like the fly that does classical conditioning (or correlational learning in computer science parlance). As shown by others (Sutton and Barto, 1998; Wörgötter and Porr, 2005; Florian, 2007), one way to do this is by using a reward function that is delayed by several steps, so that the one-step reinforcement learning scheme can be extended to multiple steps. Intriguingly, cerebellar circuitry, which has parallels with MB, implements an adaptive filter mechanism that requires comparisons of the current stimulus with future stimuli. Such a mechanism could also be effective as a multistep temporal difference learning scheme. A notable difference between animal and machine learning versions of reinforcement learning algorithms is that machine-based reinforcement learning can compute orders of magnitude more future moves than any animal can, possibly even the best of us. While games, such as chess and GO, serve as an excellent test bed for development of artificial intelligence algorithms, ultimately the goal is to solve complex real-world problems, which are not constrained by a defined set of rules. Here, the search space to be traversed is infinite, and whether machine learning's advantage over animal systems in terms of computing future moves still holds, remains to be seen. Therefore, it might be tantalizing to imagine that perhaps animals have evolved to be limited in their predictive capabilities for a reason.

A fly-based artificial neural network

In this essay, we have highlighted the salient features of animal learning that enable them to function effectively in an ever-

changing world. We used the fly MB circuitry to illustrate how these features are implemented and have outlined various benefits of incorporating animal learning algorithms into deep learning. A reasonable next step might be to ask how such a system could be structured. Next, we provide an extremely simple toy illustration to name and classify objects and point out the advantages of such a design.

In the fly, DANs and KCs drive MBON activity and the fly's decision to approach or avoid. This circuitry can be exploited to build better algorithms if we consider that such associations are likely to be at the heart of how we name and classify objects (Quiroga et al., 2005). Algorithmically, the fly circuit could be appropriated in the following way to recognize objects. The KCs would encode the attributes of the object (e.g., a person's image, features of a landmark, etc.) and the DANs would encode the name, with each DAN also having a corresponding MBON that they contact exclusively within a compartment. The first time the machine encounters that object, its DAN would alter the connection between the object's KC tag and MBON. Once the name has been learned, subsequent exposure to the object would activate its MBON.

The presence of DANs in this toy artificial neural network differentiates it from other machine learning implementations by providing it the capability of context-dependent learning and reinforcement learning without the need for additional components. This network has other benefits: (1) The representation would be robust and in a highly compressed form. (2) It would be scalable. Depending on the number of associations we would like to store, we can precisely predict the number of computing elements (neurons and compartments) needed. (3) It will be possible to quickly form schemes for robust classifications into families and superfamilies based on the distance between individual objects. (4) The system's design easily allows it to be modular and hierarchical, with each module having the same architecture.

In conclusion, the confluence of ideas between neuroscience and machine learning has a rich history (McCulloch and Pitts, 1943; Rosenblatt, 1958; Hassabis et al., 2017). Convolutional neural networks show striking similarities with the visual system. When neuroscientists in the 1960s discovered properties of simple and complex cells in the visual system (Hubel and Wiesel, 1962), they were, in essence, describing convolutional neural networks (Krizhevsky et al., 2012) in the brain. Each convolutional neural network layer contains representations of component features of the next layer (e.g., nose, lips, eyes before faces in Figure 1a). It resembles the formation of hierarchical or step-by-step representations of increasingly complex image features by successive regions within the visual cortex (Hubel and Wiesel, 1962; Krizhevsky et al., 2012). Similarly, long short-term memory networks (Hochreiter et al., 2001), a variant of recurrent neural networks, can remember events over a longer term and bear remarkable resemblance to earlier biological discoveries in the 1980s (Hopfield and Tank, 1986; Elman, 1990). And, currently, researchers are exploring how higher cognitive tasks, such as attention, can be implemented in deep learning systems (Cho et al., 2015), which might also yield insights about attentional mechanisms in brains.

Both animal and deep learning systems are computational engines, albeit implemented on differing hardware architectures. Both could be considered as variations of the universal Turing machine and subject to the same universal computational constraints (Marblestone et al., 2016). The advantages that we prescribe can be mutually beneficial to both fields. On the one hand, they will improve deep learning systems. On the other, their in-

corporation will give us fresh insight as to why they exist in brains. For instance, the effectiveness of modularity and reinforcement as strategies is apparent from AlphaGO's performance and provides a possible explanation for why they were used by evolution in brain design. Similarly, if other prescriptions, such as context-dependent learning, priming, or randomization, yield improvements in deep learning systems, they will equally serve as a window into evolution's design objectives in building a better brain. Exploring the bridge between animal and deep learning is therefore an important step in our quest for understanding and recreating true intelligence.

References

- Ackley DH, Hinton GE, Sejnowski TJ (1985) A learning algorithm for Boltzmann machines. *Cogn Sci* 9:147–169. [CrossRef](#)
- Ardin P, Peng F, Mangan M, Lagogiannis K, Webb B (2016) Using an insect mushroom body circuit to encode route memory in complex natural environments. *PLoS Comput Biol* 12:e1004683. [CrossRef Medline](#)
- Aso Y, Hattori D, Yu Y, Johnston RM, Iyer NA, Ngo TT, Dionne H, Abbott LF, Axel R, Tanimoto H, Rubin GM (2014a) The neuronal architecture of the mushroom body provides a logic for associative learning. *Elife* 3:e04577. [CrossRef Medline](#)
- Aso Y, Sitaraman D, Ichinose T, Kaun KR, Vogt K, Belliard-Guérin G, Plaçais PY, Robie AA, Yamagata N, Schnaitmann C, Rowell WJ, Johnston RM, Ngo TT, Chen N, Korff W, Nitabach MN, Heberlein U, Preat T, Branson KM, Tanimoto H, et al. (2014b) Mushroom body output neurons encode valence and guide memory-based action selection in *Drosophila*. *Elife* 3:e04580. [CrossRef Medline](#)
- Barak O, Rigotti M, Fusi S (2013) The sparseness of mixed selectivity neurons controls the generalization-discrimination trade-off. *J Neurosci* 33:3844–3856. [CrossRef Medline](#)
- Barto AG, Dietterich TG, Si J, Barto A, Powell W, Wunsch D (eds.) (2004) Reinforcement learning and its relationship to supervised learning. In: *Handbook of Learning and Approximate Dynamic Programming*. [CrossRef](#)
- Bekkers JM, Suzuki N (2013) Neurons and circuits for odor processing in the piriform cortex. *Trends Neurosci* 36:429–438. [CrossRef Medline](#)
- Bell CC, Han V, Sawtell NB (2008) Cerebellum-like structures and their implications for cerebellar function. *Annu Rev Neurosci* 31:1–24. [CrossRef Medline](#)
- Campbell RA, Honegger KS, Qin H, Li W, Demir E, Turner GC (2013) Imaging a population code for odor identity in the *Drosophila* mushroom body. *J Neurosci* 33:10568–10581. [CrossRef Medline](#)
- Candès EJ, Romberg J, Tao T (2006a) Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *Inf Theory IEEE Trans* 52:489–509. [CrossRef](#)
- Candès EJ, Romberg JK, Tao T (2006b) Stable signal recovery from incomplete and inaccurate measurements. *Commun Pure Appl Math* 59:1207–1223. [CrossRef](#)
- Caron SJ, Ruta V, Abbott LF, Axel R (2013) Random convergence of olfactory inputs in the *Drosophila* mushroom body. *Nature* 497:113–117. [CrossRef Medline](#)
- Charness N, Tuffiash M, Krampe R, Reingold E, Vasyukova E (2005) The role of deliberate practice in chess expertise. *Appl Cogn Psychol* 19:151–165. [CrossRef](#)
- Chase WG, Simon HA (1973) *The mind's eye in chess*. New York, NY: Elsevier Inc.
- Cho K, Courville A, Bengio Y (2015) Describing multimedia content using attention-based encoder-decoder networks. *IEEE Trans Multimed* 17:1875–1886. [CrossRef](#)
- Cohn R, Morante I, Ruta V (2015) Coordinated and compartmentalized neuromodulation shapes sensory processing in *Drosophila*. *Cell* 163:1742–1755. [CrossRef Medline](#)
- Collett T (1996) Insect navigation en route to the goal: multiple strategies for the use of landmarks. *J Exp Biol* 199:227–235. [CrossRef Medline](#)
- Dasgupta S, Stevens CF, Navlakha S (2017) A neural algorithm for a fundamental computing problem. *Science* 358:793–796. [CrossRef Medline](#)
- Dolan MJ, Belliard-Guerin G, Bates AS, Aso Y, Frechter S, Roberts RJ V, Schlegel P, Wong A, Hammad A, Bock D, Rubin GM, Preat T, Plaçais P-Y, Jefferis GSXE (2017) Communication from learned to innate olfactory processing centers is required for memory retrieval in *Drosophila*. [bioRxiv 167312](#).
- Donoho DL (2006) Compressed sensing. *Inf Theory IEEE Trans* 52:1289–1306. [CrossRef](#)
- Eichler K, Li F, Litwin-Kumar A, Park Y, Andrade I, Schneider-Mizell CM, Saumweber T, Huser A, Eschbach C, Gerber B, Fetter RD, Truman JW, Priebe CE, Abbott LF, Thum AS, Zlatić M, Cardona A (2017) The complete connectome of a learning and memory centre in an insect brain. *Nature* 548:175–182. [CrossRef Medline](#)
- Elman JL (1990) Finding structure in time. *Cogn Sci* 14:179–211. [CrossRef](#)
- Farris SM (2011) Are mushroom bodies cerebellum-like structures? *Arthropod Struct Dev* 40:368–379. [CrossRef Medline](#)
- Fei-Fei L, Fergus R, Perona P (2006) One-shot learning of object categories. *IEEE Trans Pattern Anal Mach Intell* 28:594–611. [CrossRef Medline](#)
- Florian RV (2007) Reinforcement learning through modulation of spike-timing-dependent synaptic plasticity. *Neural Comput* 19:1468–1502. [CrossRef Medline](#)
- Ganguli S, Sompolinsky H (2012) Compressed sensing, sparsity, and dimensionality in neuronal information processing and data analysis. *Annu Rev Neurosci* 35:485–508. [CrossRef Medline](#)
- Goodfellow I, Bengio Y, Courville A (2016) *Deep learning*. Cambridge, MA: Massachusetts Institute of Technology.
- Grossberg S (1973) Contour enhancement, short term memory, and constancies in reverberating neural networks. *Stud Appl Math* 52:213–257. [CrossRef](#)
- Haberly LB (2001) Parallel-distributed processing in olfactory cortex: new insights from morphological and physiological analysis of neuronal circuitry. *Chem Senses* 26:551–576. [CrossRef Medline](#)
- Hallem EA, Ho MG, Carlson JR (2004) The molecular basis of odor coding in the *Drosophila* antenna. *Cell* 117:965–979. [CrossRef Medline](#)
- Han KA, Millar NS, Grotewiel MS, Davis RL (1996) DAMB, a novel dopamine receptor expressed specifically in *Drosophila* mushroom bodies. *Neuron* 16:1127–1135. [CrossRef Medline](#)
- Hassabis D, Kumaran D, Summerfield C, Botvinick M (2017) Neuroscience-inspired artificial intelligence. *Neuron* 95:245–258. [CrossRef Medline](#)
- Hattori D, Aso Y, Swartz KJ, Rubin GM, Abbott LF, Axel R (2017) Representations of novelty and familiarity in a mushroom body compartment. *Cell* 169:956–969. [e17. CrossRef Medline](#)
- Heisenberg M (1998) What do the mushroom bodies do for the insect brain? An introduction. *Learn Mem* 5:1–10. [Medline](#)
- Hige T, Aso Y, Modi MN, Rubin GM, Turner GC (2015) Heterosynaptic plasticity underlies aversive olfactory learning in *Drosophila*. *Neuron* 88:985–998. [CrossRef Medline](#)
- Hochreiter S, Bengio Y, Frasconi P, Schmidhuber J (2001) Gradient flow in recurrent nets: the difficulty of learning long-term dependencies. New York, NY: IEEE Press.
- Hopfield JJ (1982) Neural networks and physical systems with emergent collective computational abilities. *Proc Natl Acad Sci U S A* 79:2554–2558. [CrossRef Medline](#)
- Hopfield JJ, Tank DW (1986) Computing with neural circuits: a model. *Science* 233:625–633. [CrossRef Medline](#)
- Hubel DH, Wiesel TN (1962) Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J Physiol* 160:106–154. [CrossRef Medline](#)
- Ichinose T, Aso Y, Yamagata N, Abe A, Rubin GM, Tanimoto H (2015) Reward signal in a recurrent circuit drives appetitive long-term memory formation. *Elife* 4:e10719. [CrossRef Medline](#)
- Kahneman D, Tversky A (1982) The psychology of preferences. *Sci Am* 246:160–173. [CrossRef](#)
- Keene AC, Waddell S (2007) *Drosophila* olfactory memory: single genes to complex neural circuits. *Nat Rev Neurosci* 8:341–354. [CrossRef Medline](#)
- Koch C (2016) How the computer beat the go player. *Sci Am Mind* 27:20–23. [CrossRef](#)
- Krishnamurthy K, Hermundstad AM, Mora T, Walczak AM, Balasubramanian V (2017) Disorder and the neural representation of complex odors: smelling in the real world. Available at <https://www.biorxiv.org/content/early/2017/07/06/160382>.
- Krizhevsky A, Sutskever I, Hinton GE (2012) Imagenet classification with deep convolutional neural networks. In: *Advances in neural information processing systems*, pp 1097–1105. Stateline, NV: neural information processing systems.
- LeCun Y (1985) Une procédure d'apprentissage pour réseau a seuil asym-

- metrique (a learning scheme for asymmetric threshold networks). *Proc Cognitiva Paris* 85:599–604.
- LeCun Y, Bengio Y, Hinton G (2015) Deep learning. *Nature* 521:436–444. [CrossRef Medline](#)
- Lee H, Grosse R, Ranganath R, Ng AY (2009) Convolutional deep belief networks for scalable unsupervised learning of hierarchical representations. In: *Proceedings of the 26th Annual International Conference on Machine Learning*, pp 609–616. Montreal, QC, Canada: Association for Computing Machinery.
- Lewis LP, Siju KP, Aso Y, Friedrich AB, Bulteel AJ, Rubin GM, Grunwald Kadow IC (2015) A higher brain circuit for immediate integration of conflicting sensory information in *Drosophila*. *Curr Biol* 25:2203–2214. [CrossRef Medline](#)
- Lin AC, Bygrave AM, de Calignon A, Lee T, Miesenböck G (2014) Sparse, decorrelated odor coding in the mushroom body enhances learned odor discrimination. *Nat Neurosci* 17:559–568. [CrossRef Medline](#)
- Litwin-Kumar A, Harris KD, Axel R, Sompolinsky H, Abbott LF (2017) Optimal degrees of synaptic connectivity. *Neuron* 93:1153–1164.e7. [CrossRef Medline](#)
- Liu L, Wolf R, Ernst R, Heisenberg M (1999) Context generalization in *Drosophila* visual learning requires the mushroom bodies. *Nature* 400:753–756. [CrossRef Medline](#)
- Marblestone AH, Wayne G, Kording KP (2016) Toward an integration of deep learning and neuroscience. *Front Comput Neurosci* 10:94. [CrossRef Medline](#)
- Marr D (1969) A theory of cerebellar cortex. *J Physiol* 202:437–470. [CrossRef Medline](#)
- McCulloch WS, Pitts W (1990) A logical calculus of the ideas immanent in nervous activity. *Bull Math Biol* 52:99–115. [Medline](#)
- Menzel R, Greggers U, Smith A, Berger S, Brandt R, Brunke S, Bundrock G, Hülse S, Plümpe T, Schaupp F, Schüttler E, Stach S, Stindt J, Stollhoff N, Watzl S (2005) Honey bees navigate according to a map-like spatial memory. *Proc Natl Acad Sci U S A* 102:3040–3045. [CrossRef Medline](#)
- Nagayama S, Homma R, Imamura F (2014) Neuronal organization of olfactory bulb circuits. *Front Neural Circuits* 8:98. [CrossRef Medline](#)
- Neuser K, Triphan T, Mronz M, Poeck B, Strauss R (2008) Analysis of a spatial orientation memory in *Drosophila*. *Nature* 453:1244–1247. [CrossRef Medline](#)
- Oertel D, Young ED (2004) What's a cerebellar circuit doing in the auditory system? *Trends Neurosci* 27:104–110. [CrossRef Medline](#)
- Olazaran M (1996) A sociological study of the official history of the perceptrons controversy. *Soc Stud Sci* 26:611–659. [CrossRef](#)
- Owald D, Waddell S (2015) Olfactory learning skews mushroom body output pathways to steer behavioral choice in *Drosophila*. *Curr Opin Neurobiol* 35:178–184. [CrossRef Medline](#)
- Papadopoulou M, Cassenaer S, Nowotny T, Laurent G (2011) Normalization for sparse encoding of odors by a wide-field interneuron. *Science* 332:721–725. [CrossRef Medline](#)
- Parker DB (1985) Learning-logic: casting the cortex of the human brain in silicon. Cambridge, MA: Massachusetts Institute of Technology, Center for Computational Research in Economics and Management Science.
- Quiroga RQ, Reddy L, Kreiman G, Koch C, Fried I (2005) Invariant visual representation by single neurons in the human brain. *Nature* 435:1102–1107. [CrossRef Medline](#)
- Root CM, Denny CA, Hen R, Axel R (2014) The participation of cortical amygdala in innate, odor-driven behavior. *Nature* 515:269–273. [CrossRef Medline](#)
- Rosenblatt F (1958) The perceptron: a probabilistic model for information storage and organization in the brain. *Psychol Rev* 65:386–408. [CrossRef Medline](#)
- Rumelhart DE, Hinton GE, Williams RJ (1986) Learning representations by back-propagating errors. *Nature* 323:533–536. [CrossRef](#)
- Scaplen KM, Kaun KR (2016) Reward from bugs to bipeds: a comparative approach to understanding how reward circuits function. *J Neurogenet* 30:133–148. [CrossRef Medline](#)
- Schacter DL, Buckner RL (1998) Priming and the brain. *Neuron* 20:185–195. [CrossRef Medline](#)
- Schaffer ES, Stettler DD, Kato D, Choi GB, Axel R, Abbott LF (2018) Odor perception on the two sides of the brain: consistency despite randomness. *Neuron* 98:736–742.e3. [CrossRef Medline](#)
- Schwarz S, Wystrach A, Cheng K (2017) Ants' navigation in an unfamiliar environment is influenced by their experience of a familiar route. *Sci Rep* 7:14161. [CrossRef Medline](#)
- Sejnowski TJ, Rosenberg CR (1988) NETtalk: a parallel network that learns to read aloud. Cambridge, MA: Massachusetts Institute of Technology.
- Shih CT, Sporns O, Yuan SL, Su TS, Lin YJ, Chuang CC, Wang TY, Lo CC, Greenspan RJ, Chiang AS (2015) Connectomics-based analysis of information flow in the *Drosophila* brain. *Curr Biol* 25:1249–1258. [CrossRef Medline](#)
- Shomrat T, Graindorge N, Bellanger C, Fiorito G, Loewenstein Y, Hochner B (2011) Alternative sites of synaptic plasticity in two homologous “fan-out fan-in” learning and memory networks. *Curr Biol* 21:1773–1782. [CrossRef Medline](#)
- Silver D, Huang A, Maddison CJ, Guez A, Sifre L, van den Driessche G, Schrittwieser J, Antonoglou I, Panneershelvam V, Lanctot M, Dieleman S, Grewe D, Nham J, Kalchbrenner N, Sutskever I, Lillicrap T, Leach M, Kavukcuoglu K, Graepel T, Hassabis D (2016) Mastering the game of go with deep neural networks and tree search. *Nature* 529:484–489. [CrossRef Medline](#)
- Skinner BF (1963) Operant behavior. *Am Psychol* 18:503. [CrossRef](#)
- Srinivasan S, Stevens CF (2018) The distributed circuit in the piriform cortex makes odor discrimination robust. *J Comp Neurol*, in press.
- Srinivasan S, Carlo CN, Stevens CF (2015) Predicting visual acuity from the structure of visual cortex. *Proc Natl Acad Sci U S A* 112:7815–7820. [CrossRef Medline](#)
- Stevens CF (2001) An evolutionary scaling law for the primate visual system and its basis in cortical function. *Nature* 411:193–195. [CrossRef Medline](#)
- Stevens CF (2015) What the fly's nose tells the fly's brain. *Proc Natl Acad Sci U S A* 112:9460–9465. [CrossRef Medline](#)
- Strauch C, Manahan-Vaughan D (2018) In the piriform cortex, the primary impetus for information encoding through synaptic plasticity is provided by descending rather than ascending olfactory inputs. *Cereb Cortex* 28:764–776. [CrossRef Medline](#)
- Striedter GF (2005) Principles of brain evolution. Sunderland, MA: Sinauer.
- Striedter GF (2015) Neurobiology: a functional approach. Oxford: Oxford UP.
- Sutton RS, Barto AG (1998) Reinforcement learning: an introduction. Cambridge, MA: Massachusetts Institute of Technology.
- Takemura SY, Aso Y, Hige T, Wong A, Lu Z, Xu CS, Rivlin PK, Hess H, Zhao T, Parag T, Berg S, Huang G, Katz W, Olbris DJ, Plaza S, Umayam L, Aniceto R, Chang LA, Lauchie S, Ogundeyi O, et al. (2017) A connectome of a learning and memory center in the adult *Drosophila* brain. *Elife* 6:e26975. [Medline](#)
- Tully T, Quinn WG (1985) Classical conditioning and retention in normal and mutant *Drosophila melanogaster*. *J Comp Physiol A* 157:263–277. [CrossRef Medline](#)
- Tulving E, Schacter DL (1990) Priming and human memory systems. *Science* 247:301–306. [CrossRef Medline](#)
- Turner GC, Bazhenov M, Laurent G (2008) Olfactory representations by *Drosophila* mushroom body neurons. *J Neurophysiol* 99:734–746. [CrossRef Medline](#)
- Vosshall LB, Wong AM, Axel R (2000) An olfactory sensory map in the fly brain. *Cell* 102:147–159. [CrossRef Medline](#)
- Waddell S (2016) Neural plasticity: dopamine tunes the mushroom body output network. *Curr Biol* 26:R109–R112. [CrossRef Medline](#)
- Werbos PJ (1974) Beyond regression: new tools for prediction and analysis in the behavioral sciences. Cambridge, MA: Harvard University.
- Weston J, Chopra S, Bordes A (2014) Memory networks. *arXiv Prepr arXiv:14103916*.
- Wilson RI (2013) Early olfactory processing in *Drosophila*: mechanisms and principles. *Annu Rev Neurosci* 36:217–241. [CrossRef Medline](#)
- Wolf R, Heisenberg M (1991) Basic organization of operant behavior as revealed in *Drosophila* flight orientation. *J Comp Physiol A* 169:699–705. [Medline](#)
- Wörgötter F, Porr B (2005) Temporal sequence learning, prediction, and control: a review of different models and their relation to biological mechanisms. *Neural Comput* 17:245–319. [CrossRef Medline](#)
- Yaksi E, von Saint Paul F, Niessing J, Bundschuh ST, Friedrich RW (2009) Transformation of odor representations in target areas of the olfactory bulb. *Nat Neurosci* 12:474–482. [CrossRef Medline](#)