Multilevel development of cognitive abilities in an artificial neural network

Konstantin Volzhenin[a,b], Jean-Pierre Changeux[a,1] 📷, and Guillaume Dumas[a,c,1]

Several neuronal mechanisms have been proposed to account for the formation of cognitive abilities through postnatal interactions with the physical and sociocultural environment. Here, we introduce a three-level computational model of information processing and acquisition of cognitive abilities. We propose minimal architectural requirements to build these levels, and how the parameters affect their performance and relationships. The first sensorimotor level handles local nonconscious processing, here during a visual classification task. The second level or cognitive level globally integrates the information from multiple local processors via long-ranged connections and synthesizes it in a global, but still nonconscious, manner. The third and cognitively highest level handles the information globally and consciously. It is based on the global neuronal workspace (GNW) theory and is referred to as the conscious level. We use the trace and delay conditioning tasks to, respectively, challenge the second and third levels. Results first highlight the necessity of epigenesis through the selection and stabilization of synapses at both local and global scales to allow the network to solve the first two tasks. At the global scale, dopamine appears necessary to properly provide credit assignment despite the temporal delay between perception and reward. At the third level, the presence of interneurons becomes necessary to maintain a self-sustained representation within the GNW in the absence of sensory input. Finally, while balanced spontaneous intrinsic activity facilitates epigenesis at both local and global scales, the balanced excitatory/inhibitory ratio increases performance. We discuss the plausibility of the model in both neurodevelopmental and artificial intelligence terms.

cognitive architecture | synaptic epigenesis | global neuronal workspace | artificial consciousness

Understanding the human brain remains the major challenge of biological sciences and has become the focus of considerable attention from both neurobiology and computational sciences (1–3). As a consequence of the recent success of neuro-inspired algorithms, including artificial neural networks and reinforcement learning, the fast-developing field of machine learning continues to look to neuroscience for inspiration. For instance, recent reflections on the connectomic implications of brain hominization in the course of evolution and brain development have underlined the importance of a still undereval-uated notion of multilevel processing in the brain from sensory processing—and its elementary "local" neuronal circuits—to higher brain functions (4). Several recent theories emphasize the hierarchical relationship between local and global processes in the making of higher-level cognition. Among these, the global neuronal workspace (GNW) (5) is exemplary. A hierarchical relationship between local and global processes is also outlined in Kahneman's (6) functional distinction of "system 1" handling fast and nonconscious cognitive processes and "system 2" handling cognitive tasks requiring slower and more concerted conscious effort. While current algorithms in machine learning are now tackling system 1 tasks at the human performance level, a major milestone for modern artificial intelligence would be to provide models capable of approximating system 2 cognition (7).

To a large degree, the strides made in artificial intelligence research during recent decades can be explained by the overwhelming success of error backpropagation. Nevertheless, its biological plausibility remains a matter of considerable debate (8–10). Unlike artificial deep networks, and due to their discontinuous nature, spiking neurons are thought to learn by other means. Concretely, electrophysiological methods have indicated two principal mechanisms of learning in the human brain: Hebbian learning (11) and reinforcement learning (12). The mechanism of epigenesis by synapse selection (or synaptic pruning) is also known to play a significant role in both learning and biological development (13). Moreover, it has been demonstrated that neurogenesis can occur in adults (14) and that astrocytes may be implicated in synaptic modulation during learning (15). Despite solid evidence from joint anatomical, physiological, and

Significance

Multiple biological mechanisms support the unique ability of the brain to develop complex cognitive abilities. Nevertheless, it remains unclear which mechanisms are necessary and sufficient. We propose a neurocomputational model of the developing brain spanning sensorimotor, cognitive, and conscious levels. The model solves three tasks of increasing complexity: from visual recognition to cognitive manipulation and maintenance of conscious percepts. Results highlight two fundamental mechanisms for the multilevel development of cognitive abilities in biological neural networks: 1) synaptic epigenesis, with Hebbian learning at the local scale and reinforcement learning at the global scale; and 2) self-organized dynamics, through spontaneous activity and balanced excitatory/inhibitory ratio of neurons. We emphasize how these core features of human intelligence could guide future development in artificial intelligence.

[1]To whom correspondence may be addressed. Email: changeux@noos.fr or guillaume.dumas@ppsp.team.

molecular investigations in the course of nervous system development, such mechanisms have been underexploited in brain modeling and computer sciences (13, 16). Recent work in neuroscience continues to uncover novel mechanisms at play during learning, which are far from being integrated into computer science. These and other mechanisms may provide important links, both within and between levels of organization, in diverse contexts—ranging from gene networks to long-range neuronal connectivity of the brain (17). Recent neuromorphic approaches and biologically plausible architectures, even including various brain regions such as the prefrontal cortex, basal ganglia, or thalamus, can perform a wide set of different computational tasks with energy efficiency (18, 19).

If detailed key mechanisms of learning in the brain have already been investigated, there is, as yet, no theoretical consensus on how these varied learning mechanisms interact in the brain (20–23). Yet it is known that the brain operates constantly through variation–selection mechanisms at multiple timescales (17) and thus supports the development of a multiscale and dynamical perspective on cognition and learning (24, 25). In artificial intelligence, new approaches have adopted such dynamical and multiscale aspects using attention (26) or social interaction between multiple agents (27). Basic developments inspired by neuroscience have also highlighted the importance of better capturing hierarchical relationships (7, 28, 29). Finally, the idea of biological learning without any inductive bias, typical of artificial neural networks (ANNs), has been criticized. On the contrary, biological investigations point dramatically to the embodied quality of cognition, perception, and action, as well as the important role of embeddedness in natural and social milieus (30, 31).

In this paper, we present a framework that pursues this biologically grounded inquiry of the mechanisms underlying cognitive ability in humans. We introduce a model of synaptic epigenesis that combines multiscale architecture with spike timing dependent plasticity (STDP) and dopamine signaling. We outline the most elementary structural biological requirements for conscious processing (32) that are possibly accessible in simple organisms such as *Drosophila* (33). We hypothesize that synaptic epigenesis unfolds differently at local and global scales and that the requisite conditions for solving complex tasks associated with the GNW include not only local but global epigenesis. We first identify necessary and sufficient conditions for proper learning of perceptual, cognitive, and conscious tasks. Next, we analyze how the identified factors influence performance in these tasks. After introducing the neurocomputational model, results are presented, and the key role of dopamine and inhibitory neurons in the learning of higher cognitive tasks is discussed. Finally, we delineate future perspectives for the field of computational cognitive neuroscience and neuroAI.

Model

Overview of the Model and Its Neuronal Components. We propose a computational model of an elaborate multilevel neural network able to pass cognitive tasks of increasing difficulty (Fig. 1). The task solving develops at three nested levels of hierarchical organization, where synaptic epigenesis may proceed at both local and global scales. The three different levels of structural organization show the increasing complexity of connectomic architecture with each nested level through a continuous progression. The first level, which we refer to as the "sensorimotor level," deals with local sensory processing and classification of visual information: It requires local synaptic epigenesis. At the second
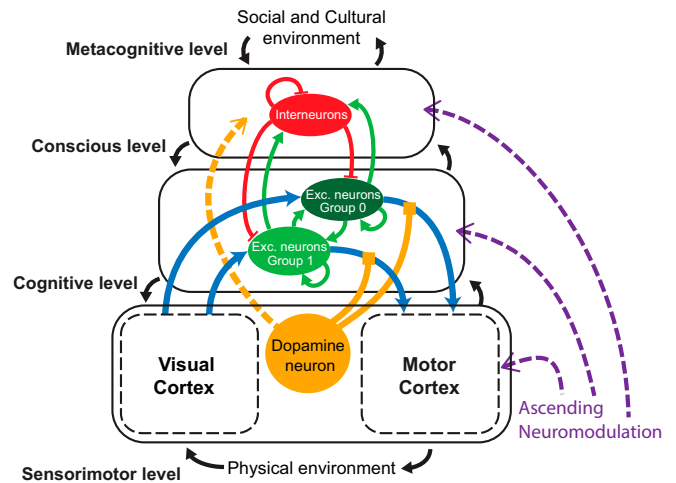

Fig. 1. The full network architecture comprises sensorimotor, cognitive, conscious, and metacognitive levels. This is a schematic representation of the full network and its components. Colors of arrows depict different types of connections: light blue, excitatory long ranged (fixed: between VC and GNW; STDP: between GNW and motor cortex - MC); red, inhibitory (STDP); green, excitatory (STDP); yellow, extracellular dopamine modulation (affects dopamine modulated STDP). This version of the model is used to perform the trace conditioning task on the conscious level. To use the model with the delay conditioning task we do not introduce the interneurons population on the cognitive level.

level, or "cognitive level" (34), the network successfully passes a delay conditioning task: It mobilizes multiple cortical areas, and their integration requires long-range axonal connections. The third level, referred to as the "conscious level," is able to carry a trace conditioning task using an architecture similar to the cognitive level, yet with the addition of the necessary contribution of inhibitory interneurons. And, even though inhibitory neurons are present not only at the highest level of brain organization, we are showing that they are a necessary requirement at this particular step. The conscious level can be further linked to constant interactions with the social and cultural environment and form a so-called metacognitive level; however, we do not study these connections in this work, and this hypothesis is a subject of future studies.

Fig. 1 illustrates the overall neuronal organization of the model network. The three structural levels are all made up of neuronal layers organized hierarchically and nested within each other in a bottom-up and top-down manner (11). Each level may include several differentiated territories, or subnetworks, such as, for instance, the visual cortex or the motor cortex in the lower sensorimotor level. Every part of the network is modeled using neuronal assemblies, except for the first two layers of the local network of the visual cortex. The state of each neuron is determined by the leaky integrate-and-fire model (35), and each assembly includes not only internal connections but external connections to other assemblies of the whole network. An important feature of the model is its evolution with time. In the course of learning, the network connections are subject to epigenesis by selective stabilization of synapses (13). At critical stages of development, stable connections emerge from an initial growth process with overproduction of synaptic contacts with maximal diversity and variability, and then a selection–stabilization regulated by the state of activity of the network takes place, together with the elimination—or pruning—of the unselected ones (13). We examine, at all three levels of organization, how this development happens, both at local and global scales, and how internal and environmental factors affect its evolution.

The intrinsic spontaneous activity of the component neurons is a singular functional component introduced in the model proposed to occur at all its hierarchical levels. Also, the network is not exclusively composed of excitatory neurons as in most standard networks. At the highest level, we introduce inhibitory interneurons and test how the interneuron vs. excitatory neuron ratio affects the performance. Long-range connections (5) which interconnect widely distinct territories of the global network are introduced at the higher cognitive and conscious levels.

Last, throughout this work, we use concomitantly different types of learning. At the local level, the network evolves under Hebbian learning with STDP (36). However, all the global connections are controlled by reinforcement learning (12). To do so, we introduce the dopamine reward system (37), which is modulated by interactions with the external environment (Fig. 1).

Description of the Tasks. The three selected tasks challenge differentially the sensorimotor, cognitive, and conscious levels.

The sensorimotor level is common to all tasks. At this level, we use the MNIST (Mixed National Institute of Standards and Technology) dataset of handwritten digits as visual input to the visual cortex (38). The network performs a number recognition task (39). Following image presentation, the network processes spikes arriving within the first 250 ms, followed by a relaxation period of the same length. Thus, one image is shown every 500 ms during the learning period. At the higher levels, when the full network receives an input from the visual cortex, the network is expected to respond to the question posed by the experimenter:" Is this number larger than x? If yes, push a button." For the case with only "0" and "1" in the input data, we took x equal to "0."

Two other tasks, delay and trace conditioning (40), have been used to challenge the cognitive and conscious levels. In delay conditioning, the unconditioned stimulus (US) immediately follows or coterminates with the conditioned stimulus (CS), whereas, in trace conditioning, the CS and US are separated in time by a "trace" interval. Here, the numerical input operates as the CS. We study two different variations of the task with different temporal locations of the CS and US. During the delay conditioning, we present a number for 150 ms, and the trigger stimulus coincides with the last 50 ms of the input. To create the trace conditioning, we first introduce an input for 150 ms, followed by a pause of 200 ms where the input is removed, and, only after this pause, the GNW receives the trigger stimulus of 50 ms. Whenever the motor cortex produces output, that is, the firing frequency of the motor neuron is above a certain threshold, the environment provides a negative or positive reward, depending on the answer.

Local Network of the Visual Cortex. The architecture of this network (Fig. 2A), inspired by the work of Masquelier and Thorpe (41), consists of three layers representing the simplified model of the visual pathway. The first layer imitates the lateral geniculate nucleus and performs a convolution operation with filters of various angles. Then, in the second layer, only the strongest orientation reaches the V1 area. And, in the last layer (V4 area), the patterns of digits are learned by means of STDP. In the text, we use the notation of "complex" and "simple: cells; thus the overall architecture can be written as S1–C1–S2 (see ref. 37). A detailed description of each layer can be found in *Materials and Methods.*

To train the model and observe the local epigenesis of STDP regulated synapses, we used the MNIST dataset of handwritten digits (Fig. 2B). These inputs consist of images of handwritten
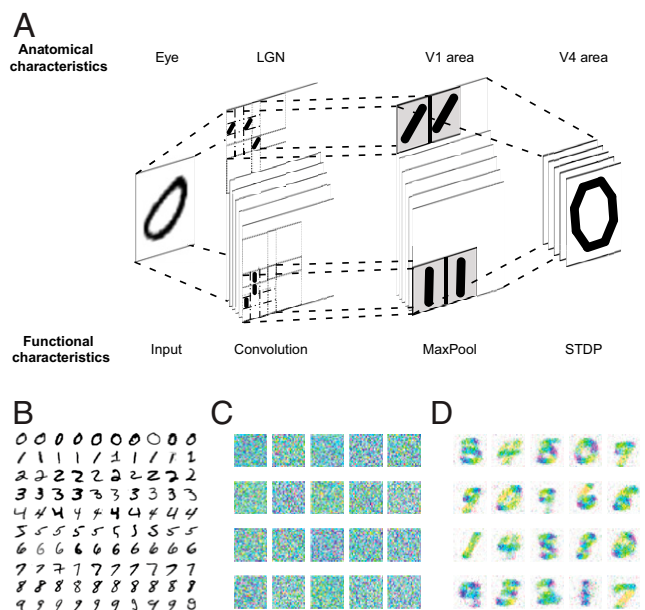

Fig. 2. The local network of the visual cortex performs a classification task on the MNIST dataset. (A) The schematic representation of the network structure with functional descriptions of each layer and corresponding anatomical analogies. The model is based on three processing layers, where the first one performs convolutions on the original image with six different filters of the same scale responsible for six different orientations. The second layer propagates only the strongest orientation in a given area. The connections of the last layer are regulated by STDP. The model is inspired by the paper of Masquelier and Thorpe (41); for more information see the original article. (B) The sample of the MNIST dataset of handwritten digits, 100 examples. (C and D) Pattern reconstructions of the 20 last layer cells before (C) and after (D) learning. The reconstructions were done using the HSV color scheme, where the color of each pixel represents the mean of the intensities throughout the six orientations at that location.

digits from "0" to "9." MNIST is one of the fundamental datasets in machine learning, which makes it reliable and easy to use, as well as an informative benchmark for nontraditional learning methods. The images are introduced to the network continuously, with equal relaxation periods in between. These relaxation phases are needed for the system to eliminate the temporal effect on STDP coming from the previous stimulus, and to bring all membrane potentials back to rest.

After a number of images are presented, the network's plasticity allows it to evolve from a homogeneous state to one that has encoded various patterns by way of Hebbian learning. Fig. 2 C and D shows what the patterns of 20 S2 cells look like before and after the network completes its learning process, respectively. These images are plotted using the HSV (hue, saturation, value) color scheme. Values for each pixel are set to 100%. Hue and saturation are determined in the following way: Each pixel combines information from six synapses (by the number of used orientations). These six synaptic weights are transformed into the complex plane: from w to $w^* = w * exp(i\varphi)$, where φ is the particular orientation. Next, all six weights are added together. The absolute value of the resulting complex number is saturation, and the angle is hue. Using this color scheme, we are able to see not only the strength of the C1–S2 connection but also the dominant orientation (if present) coming from any of the C1 cells (Fig. 2 D).

Full Network. After the local network has undergone the process of Hebbian learning, the weights of C1–S2 connections are frozen, and this network is considered mature. Then, a new architecture is constructed on top of it, to explore global

connections. This intermediate network consists of four fundamental elements: the GNW, the visual and motor cortices, and the Striatum, represented by a single "dopamine" neuron (Fig. 1). The visual cortex is essentially represented by the aforementioned local network with deactivated STDP and fixed synaptic weights. Input to the global network consists of the output of S2, once its trained model has been extracted from the stream of images, features to be read by the full network. In other words, raw data are inaccessible to the full network, and only intermediate features can be accessed by subsequent brain regions. S2 output is passed to the excitatory population in the GNW.

The motor cortex is minimally represented by only a single neuron. The firing of this neuron is interpreted as revealing the interaction with the external environment, such as pushing a button and giving a positive answer to a question. It processes outputs only from the excitatory populations in the GNW.

The GNW is assumed to contain two populations of neurons: excitatory neurons and interneurons. To test the cognitive level with the delay conditioning, we use the simplest version of GNW without interneurons. The excitatory population is divided into two selective groups (by the number of digits we want to use for the task; in our case, it is a binary classification, and we only use digits "0" and "1"). Each selective group is connected to S2 cells in the visual cortex model which are responsible for one specific digit. All these connections are fixed. The connections inside the GNW between both excitatory and interneurons are all to all.

Learning in the GNW–MC connections is accomplished using reinforcement via interaction with the external environment (dopamine-modulated STDP). The internal GNW connections are subject to Hebbian learning (classical STDP) in the same way as C1–S2 connections in the local network. More detailed information about STDP models can be found in *Materials and Methods.*

Results

The proposed model of an elaborate multilevel neural network (Fig. 1) is able to pass behavioral tasks of increasing difficulty. The three levels of structural organization exhibit increasing connectomic complexity with each nested level.

Sensorimotor Level—Visual Recognition Task: Stabilization and pruning. At this level, the neuronal network—described above as the "visual cortex local network"—can perform a recognition task of numbers, using handwritten digits as inputs. This image processing task mobilizes nonconscious perception, by leveraging local synaptic epigenesis as well as the spontaneous activity of the network. Fig. 3 shows the computed multistep evolution of this performance.

The initial state of the LGN-V1 synaptic population is highly homogeneous. The synaptic weights are normally distributed around the mean of 70% with an SD of 2%. When the network is presented with the constant input stream of images, the synaptic population starts to evolve. The evolution of the synaptic strength of the whole synaptic population is presented in Fig. 3B. This graph shows the evolution of the percentiles for the population with a step size of 10, from 0 to 100. It can be seen that all synapses do not behave in the same way. Around two-thirds of population synapses immediately weaken and are, as a consequence, eliminated. At the same time, other synapses are strengthened and selected (see ref. 13).

Fig. 3C describes the evolution of the population of synapses in terms of selective stabilization. The selected synapses are those responsible for the transfer of the signals related to the input information. Others are gradually eliminated, such that they exert a lesser impact on the final layer. However, this process is not linear, and learning patterns change over time. Fig. 3D represents the evolution of the states within the synaptic population. In this schematic representation, a synapse belongs to the green ("selected") area if the synaptic strength (or synaptic weight) is greater than 75%. Similarly, the red ("eliminated") area accounts for synapses with a strength of less than 25%. All others between 25% and 75% fall into the gray ("undetermined") area. This zone contains the entire population at the beginning of the simulation.

Due to nonlinear behavior, we have outlined two distinct stages of learning depending on the growth rate: 1) from the start to ~200 shown images and 2) from ~200 to 2,000. Statistical tests have been performed to establish that the growth rate of selected and eliminated synapses is reliably different during the first and second periods. The t test has shown the P value $< 10^{-8}$ and Cohen's D > 15 for both groups. With
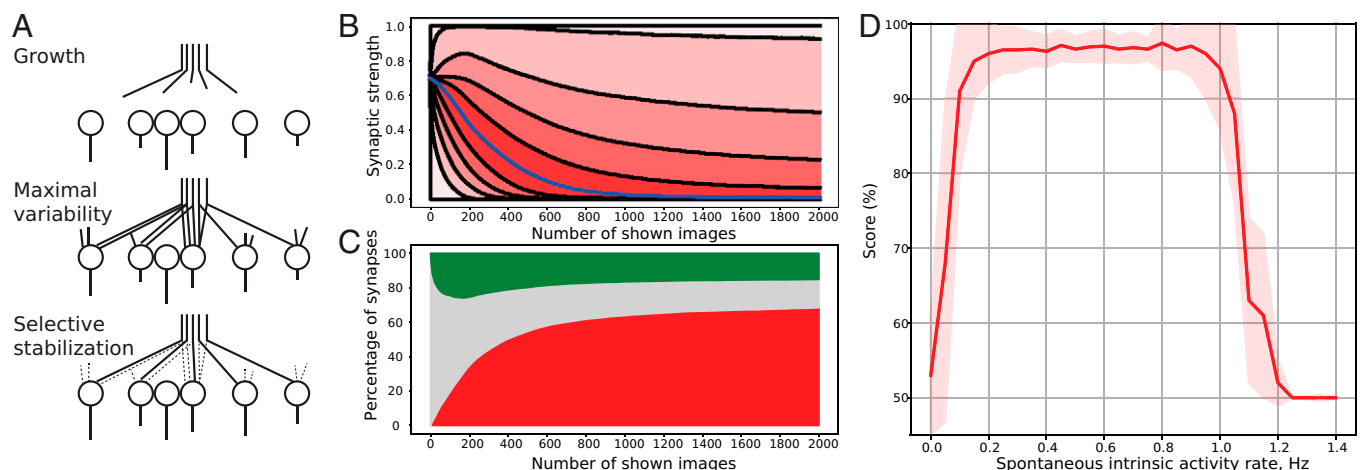

Fig. 3. Local epigenesis and performance of the local network. (*A*) Schematic representation of the hypothesis of epigenesis by selective stabilization of synapses [from Changeux (42)]. (*B*) Evolution of the synaptic strength of the whole synaptic population in the last layer of the local network. Blue line, the median; black lines: zero to hundredth percentiles (with a step of 10). (*C*) The epigenesis process in terms of selective stabilization of synapses, that is, synaptic numbers instead of strength like in *A*. Red group represents eliminated synapses with weights of less than 25%. Green group represents selected synapses with weights of more than 75%. Gray group includes weights between 25% and 75%. (*D*) The performance of the matured network depends on the rate of spontaneous intrinsic activity. There is only a specific window where the network can produce a high performance of more than 90%, with exact rates of the spontaneous intrinsic activity depending on the number of neurons.

the chosen set of parameters, the first period occurs while the model is processing the first ~200 images. During this time, we notice the rapid selective stabilization of half the synaptic population. Simultaneously, the percentage of selected synapses rises either. During this period, V1 cells gain their selectivity toward specific sets of stimuli.

The second stage of learning can be characterized as the fixation and reinforcement of the changes that happened before. The fraction of eliminated synapses monotonously reaches a plateau level with a value of around two-thirds of all synapses. The group of selected synapses, however, experiences "overshooting" and, after 200 images, starts to slowly decrease before reaching a plateau. This behavior can be described as a period when the S2 features become clearer, eliminating redundancy between them and within the spontaneous intrinsic activity.

In addition to synaptic epigenesis, the intensity of spontaneous intrinsic activity is among the factors affecting the performance of the network. In our model, each of the C1 neurons is able not only to produce spikes caused by the S1 layer but also to generate spontaneous intrinsic activity in a form of random Poisson distribution with a fixed rate. We can see that high performance (with the accuracy of classification > 90%) can only be achieved in a specific range of values (Fig. 3D). We observe no or very weak classification with values of spontaneous intrinsic activity close to zero; the same behavior is spotted for large values as well (> 1.0 Hz). The classification reaches high scores of >90% only between 0.1 Hz and 1.0 Hz. Thus, a sufficient but not exceedingly high amount of spontaneous intrinsic activity is a necessary factor in achieving the accurate recognition of numbers. The specific values of the spontaneous activity vary depending on how many neurons are used and how strongly it affects the conductivity of each neuron.

Cognitive Level—Nonconscious Delay Conditioning Task: Long-Ranged Connections and Reward. In this section, we examine which network global architecture is needed for the performance of a delay conditioning task (Fig. 2), again using digits as visual stimuli for the unconditioned and conditioned stimuli. We discover that both the presence of long-range connections and their epigenesis are required. The initial state of the network is homogeneous: The local connections are set to 1% of the maximum weight, and the long-range connections between the GNW and the motor cortex are 50% for all synapses (Fig. 4).

The network uses long-range connections to transmit stimuli between visual and motor cortices via a hub of excitatory neurons, as hypothesized in the GNW model (5). We show that the network can perform a delay conditioned task (Fig. 4A represents the temporal relation between the stimulus and the trigger) without conscious processing. Fig. 4B depicts the activity of the GNW at the beginning and the final stages of the learning process. The stimulus is erased from memory as soon as the input is gone, meaning that it is not processed in any way between visual and motor cortices. However, depending on the initial parameters, our network can start firing constantly in a positive feedback loop, since it contains only excitatory neurons. In this case, the network is not able to hold representations either. The epigenesis process transpires rather quickly, on the scale of ~15 shown images, and all the synapses responsible for the representations of the digit "1" get selected while the representations of the digit "0" get eliminated (Fig. 4C).

The model turns out to be sensitive to the rate of spontaneous intrinsic activity. We have observed that, regardless of configuration, the model can learn how to perform the task only
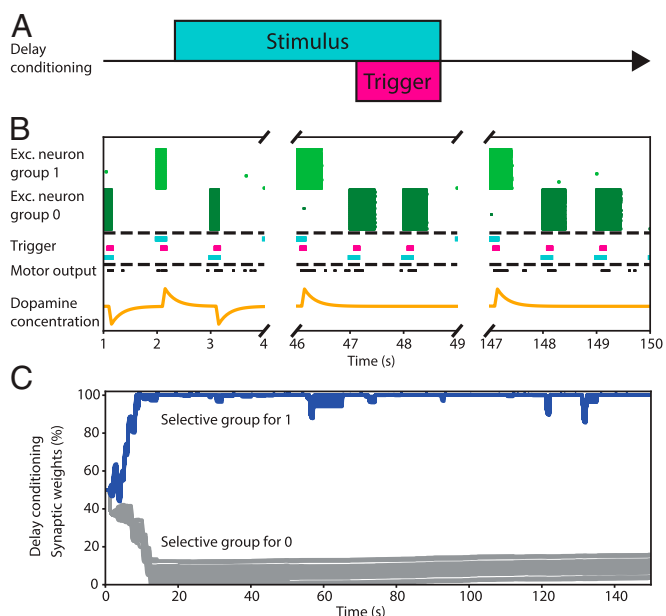

Fig. 4. Epigenesis of the long-range connections network between GNW and motor cortex while performing binary classification under the delay conditioning. (A) Description of the delay conditioning task. (B) Evolution of the firing rate of the two populations of neurons within the GNW (excitatory selective for the digit "1," excitatory selective for the digit "0"). Fifteen neurons per population are shown; each dot shows a neuron spiking. Depending on the network parameters, if the internal excitation in the GNW is not strong enough to cause constant feedback firing, the stimulus representation exists only while the input signal is present. (C) Epigenesis of the global connections between GNW and motor cortex. The network needs less than 20 images to learn the task and makes mistakes only when the visual classification in the local network is wrong. Blue and gray lines correspond, respectively, to one and zero output of the GNW.

when the rate is set within a certain interval (Fig. 5). Moreover, while there is a specific range of values within which the network provides decent classification (accuracy > 95%), excessive activity can disturb the learning process, and the accuracy will remain around the chance level. The performance does not change significantly as long as the spontaneous intrinsic activity is set in the right interval.

Conscious Processing Level—Trace Conditioning Task: The critical role of inhibitory neurons. To date, the models associated with the GNW theory have focused on purely excitatory networks. Here we show the critical role of inhibitory neurons for conscious tasks. Additionally, we demonstrate how the performance on those tasks depends on the interplay between the ratio of excitatory and inhibitory neurons and spontaneous intrinsic activity.

Although the long-range connectivity itself is enough to perform the delay conditioning, it cannot achieve the trace conditioning task by itself (which is represented in Fig. 6A). On this level, the role of interneurons is seen to be critical in the function of the GNW for the maintenance of conscious representation. This is needed to create and hold stable representations of stimuli and erase them when needed. Our model comprises two selective excitatory populations and one interneuron population. Once mature, it can hold a stable representation and erase it upon the presentation of a different stimulus (Fig. 6B). In this case, global epigenesis takes more time, on the order of ~150 shown images, because it requires the local epigenesis in the GNW to occur first (Fig. 6C). Therefore, the mechanism of learning trace conditioning is highly dependent on the ability of working memory to maintain a needed stimulus.
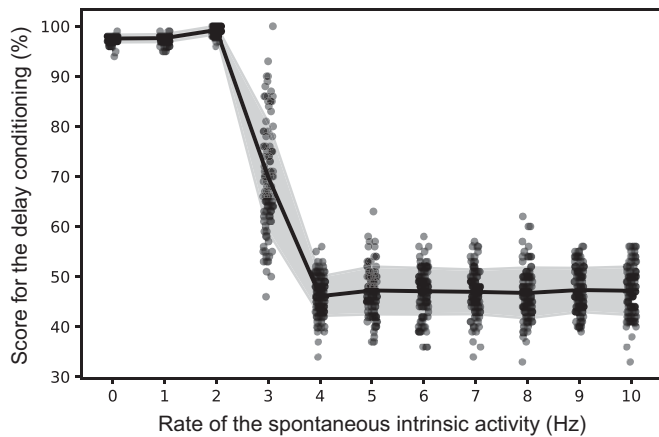
Fig. 5. Effect of spontaneous intrinsic activity on the performance during the delay conditioning on the cognitive level. The accuracy of the binary classification as a function of the rate of spontaneous intrinsic activity is shown. High values of spontaneous intrinsic activity force neurons to fire continuously, in an "epileptic" manner, and the network loses the ability to perform classification tasks. In such a state, the average accuracy is slightly lower than a chance level (50%) because the visual cortex does not have 100% accuracy in digit recognition (Fig. 3D).

Just as in the intermediate level, with the delay conditioning, we have studied how the model behaves with various values of the spontaneous intrinsic activity and the excitatory/inhibitory (E/I) ratio. The performance of the model strongly depends on both parameters, and high accuracy can be achieved only in the specific range of both values (Fig. 7A). In this case, we have identified a very specific range of spontaneous intrinsic activity (8 Hz to 13 Hz), as well as a narrow interval of the E/I ratio which is located between 70% and 95% (maximum accuracy achieved with ratio = 80% and rate = 12 Hz).

The same study has been conducted for the delay conditioning task, to assure that the model can achieve good performance on the same parameter values (*SI Appendix,* Fig. S1). It is a crucial step toward biological plausibility: The network cannot perform the trace conditioning task without being able to work with delay conditioning beforehand. In our case, we have identified, with high accuracy, the region where both tasks are performed.

Discussion

In this paper, we address the issue of how the progressive complexification and integration of an artificial neural network gives rise to cognitive abilities. The current model features the GNW (5) as a primary mechanism for conscious processing; however, it is not strictly limited to a particular brain area. The model is based on three hierarchical levels of information processing: the sensorimotor level; the cognitive level corresponding to the global nonconscious processing of information across different brain regions; and, last, the conscious level corresponding to the autonomous and lasting processing of information, even in the absence of externally applied sensory stimulation. This partition is consistent with the former distinction by Descartes (43) and Kant (44) of three levels of functional processing, "sensibility, understanding and reason" (see refs. 31 and 32). More recently, Daniel Kahneman proposed the dichotomy of a fast "system 1," instinctive and emotional, and a slow "system 2," which is more deliberative and logical (6). The development of computer science and artificial intelligence has led to a strikingly different distinction, such as David Marr's (45) partition of computational mechanisms into what he refers to as "formal

levels": implementation, algorithms, and computations. Marr's claim relied on the underlying assumption that the computations required for a given cognitive task are distinct from the physical implementation, natural or artificial, of the network. They are independent of the cognitive process under study and, most of all, of the level of organization of the network. The word "level" has, in the two cases, entirely different meanings.

In light of the need for grounding cognitive models in biology and validating their genuineness across levels, we followed successive mechanisms in known structural and functional organizations, identified through both the development and evolution of the human brain (46). The first step proceeds from local to global integration of information; a second one moves from nonconscious to conscious processing. This leads to the three structurally and physiologically grounded levels: "perception," with the learning of invariants through local synaptic epigenesis; "cognition," with the learning of nonconscious integration through the long-range connectivity and reward; and, finally, conscious processing, with the learning of maintenance of representation online through interneurons and spontaneous activity. In our particular model, the conscious level does not comprise all the features of consciousness but rather emphasizes the necessary functional elements, such as the working memory. Functional characteristics are, indeed, essential for the study of consciousness, and may even fundamentally help to solve the hard problem (47). Here, the learning process is based on the idea of sequential network maturation, when the local areas are trained before the global connections (48, 49). Those critical periods of learning should display better performance than training the
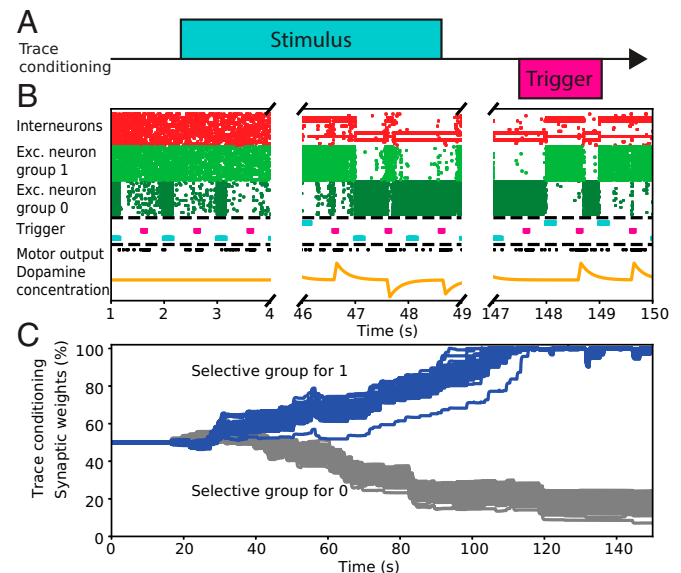

Fig. 6. Epigenesis of the long-range connections network between GNW and motor cortex while performing binary classification under trace conditioning. (A) Description of the trace conditioning task. (B) Evolution of the firing rate of the three populations of neurons within the GNW (inhibitory, excitatory selective for the digit "1," and excitatory selective for the digit "0"). Fifteen neurons per population are shown; each dot shows a neuron spiking. Depending on the network parameters, if the internal excitation in the GNW is not strong enough to cause constant feedback firing, the network gradually learns to sustain the representation of a stimulus in the GNW. When the next stimulus is present, the previous one is erased. The interneuron population also forms two distinct subpopulations selective to two different excitatory groups. (C) Epigenesis of the global connections between GNW and motor cortex. The network needs around 150 images to learn the task, which is significantly more than with the delay conditioning because it also has to spend time learning how to hold representations in the GNW. Blue and gray lines correspond, respectively, to one and zero output of the GNW.
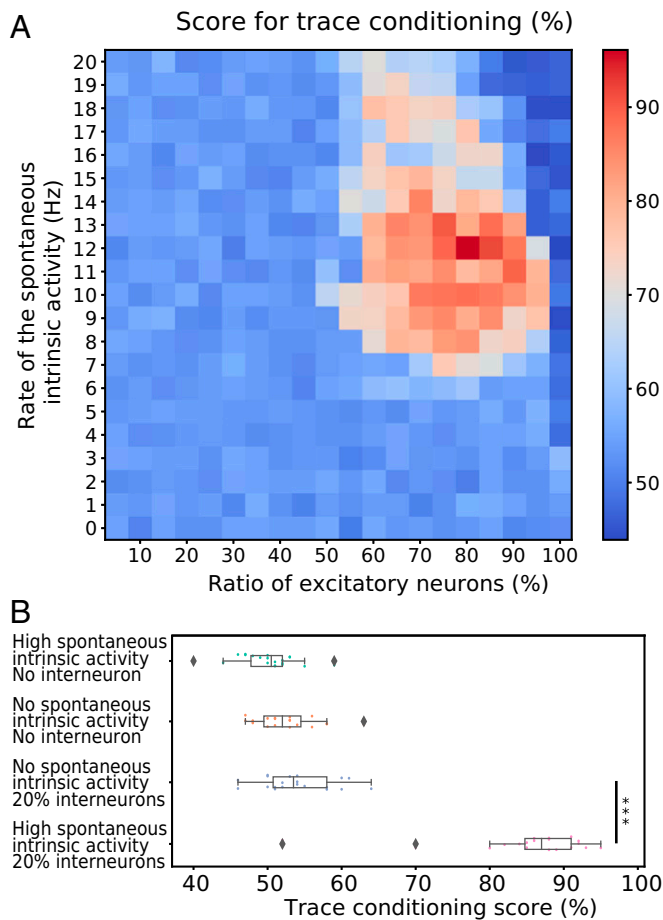
A

Score for trace conditioning (%)

B

Fig. 7. Factors affecting the performance during the trace conditioning on the consciousness level. (*A*) The accuracy of the binary classification as a function of the rate of spontaneous intrinsic activity and the proportion of the interneurons. Unlike the delay conditioning on the cognitive level, the rate of the spontaneous activity plays a great role here, and, in order for the model to show good performance (>75%), both parameters must be within a specific range. (*B*) The accuracy of the model in four distinct states: 1) high spontaneous intrinsic activity and no interneurons, 2) no spontaneous intrinsic activity and no interneurons, 3) no spontaneous intrinsic activity and 20% of neurons are inhibitory, and 4) high spontaneous intrinsic activity and 20% of neurons are inhibitory. The first three states provide poor performance, while the last one, which has the optimal set of parameters, yields high accuracy. Diamonds and triple asterisk are respectively indicating outliers and statistically significant differences (Mann Whitney U test, *P* < 0.001).

whole network simultaneously; however, this hypothesis remains to be tested in future work. The model expresses key insights across the three levels of analyses: First, synaptic epigenesis, modeled here by selection and stabilization of synapses, is a critical mechanism at all levels, from perception to consciousness; second, dopamine is necessary for cognitive tasks to achieve proper credit assignment despite temporal delay between perception and reward; third, interneurons allow the maintenance of self-sustained representation within the GNW in the absence of sensory input, thus enabling the system to solve conscious tasks. Even though it has been debated that trace conditioning is correlated with consciousness (50), this type of test is still considered to be one of the best measures to test conscious processing in the absence of explicit reports (51). Finally, our results show how balanced spontaneous intrinsic activity facilitated epigenesis at both local and global scales; the balanced E/I ratio increased performance. Those observations, in addition to emphasizing general principles at play in the human brain, are in surprising accordance with empirical observations. For instance, without including synaptic pruning, the temporal evolution of synaptic

weights leads to the three classical phases of growth, maximal variability, and stabilization (42). More unexpectedly, the optimal neurobiological parameters predicted for the conscious level are 20% of inhibitory neurons, which is predicted for balanced networks (52). Lastly, the "optimal" 12 Hz of spontaneous activity discovered in our artificial network coincides with the actual recordings of spontaneous spiking in the human cortex. More generally, according to Fig. 7*A*, the optimal range of the spontaneous intrinsic activity lies within 8 and 14 Hz, which is also strongly reminiscent of the highly conserved and omnipresent alpha rhythm in the cerebral cortex (53). This correlation (along with its importance for the system) may be studied in the future as one of the mechanisms supported by the alpha rhythm and how it may relate to stochastic resonance (54).

Computational neuroscience has increasingly relied on multiscale brain models (55, 56) to explain basic cognitive functions (18); active, top-down, and prospective memory retrieval (57, 58); and even syntactic processing for the production of language (59). Some of these models, like ours, have employed spiking neural networks with STDP to study synaptic epigenesis during learning of sensory representation for tactile (60) and visual (41) perception, or to demonstrate how biases in natural statistics can influence population encoding and downstream correlates of behavior (61). These models, however, have tended to focus on a single learning mechanism, applied to a specific task. Here we combine both Hebbian and reinforcement learning across different canonical tasks. While the tasks considered here may appear minimalistic compared to recent sensational breakthroughs coming from artificial intelligence (62, 63), our goal was not to achieve behavioral complexity but to uncover general biological principles through natural plausibility. Moreover, while it is certainly possible for artificial systems to mimic the behavior of a conscious agent, it is far less trivial to demonstrate the genuineness of its semantic understanding and subjective experience (64). From a purely pragmatic viewpoint, the Turing test (65) relies only on the judges being unable to tell whether the agent is a human or a machine (66). We chose to view the problem in another way, by grounding the cognitive architecture in neurobiology and, from there, building the most parsimonious, realistic model capable of solving both perceptual and conscious tasks. This allowed delineating necessary and sufficient biological mechanisms for cognitive abilities in an artificial neural network (67), especially to achieve artificial consciousness (68, 69). Future works should probe the role of other important biological mechanisms such as oscillations (70) or distributed coding for task value (71), and may likewise wish to take into account additional subcortical structures (beyond the striatum), such as the thalamus (72) and the basal ganglia (73). A promising avenue is also to explore the combination of the GNW architecture with predictive processing (74); the top-down feedback required for ignition and conscious access then becomes entangled with the prediction signal of the internal generative model of the world (75). Finally, the GNW architecture is also opening new venues in deep learning: on the one hand, by creating high-level inductive biases that improve out-of-distribution generalization (76, 77), and, on the other, by providing an amodal latent space where alignment of representations across modalities becomes automatic (78, 79).

Recently, a separate distinction was made between basic "conscious" processing and "the self-monitoring of those computations, leading to a subjective sense of certainty or error" (80). This metacognitive dimension of consciousness is also linked to the hypothesis of the social origin of consciousness, which proposes the sense of self as a vestige of the evolutionary skills initially developed for understanding others (81). In this

sense, the hominization of primates seems connected to molecular changes associated with social cognition (82) and, in line with our results, with the increase of long-range connectivity and regulation of the balance between excitatory and inhibitory neurons(4). Beyond these features, the prolonged postnatal brain development with a proper cascade of critical periods (83, 84) leverages the multiple nongenetic interactions with the physical, social, and cultural environment, ultimately giving rise to categorically human-specific cognitive abilities, including the recursivity of language (4, 85). A key perspective is thus to further investigate the role of cultural embedding and intersubjective interaction (86), adding, explicitly, a social dimension to the GNW (17, 87). Some dyadic computational models (88, 89) have already been proposed using either continuous dynamical systems or discrete symbolic representation. Anchoring future computational models in both biological and social realities will not only continue to shed light on the core mechanisms underlying cognition, but it will also help to provide a unique bridge to artificial intelligence toward the only known systems with advanced social consciousness: the human brain (90).

Materials and Methods

This section presents a technical description of the network with all the learning mechanisms used. All the code is available on the GitHub repository: https://github.com/ppsp-team/EpigeneSys/tree/master/GNW (91).

Local Network.

S1 layer. The convolution of the images is done with six 7×7 Gabor filters. The frequency is 0.1 and the SD is 1 for all of them. The angles of the filters start at $\pi/8$ and go with a step of $\pi/6$. The offset of $\pi/8$ was introduced in order not to focus on vertical and horizontal lines which are not as informative in classification. Every time the Gabor kernel overlaps a border of a picture, this picture is extended by replicating the edge pixels. The resulting convolution is again resized to 128×128 pixels, thus creating an S1 layer with the size of $128 \times 128 \times 6$.

C1 layer. Each cell of this layer processes spikes from S1 cells in a 7×7 square (receptive field size of a given C1 cell). A C1 cell propagates only the single strongest orientation from its receptive field at a time. The receptive fields of two adjacent C1 cells are shifted by six pixels, meaning that they have an area of intersection of 7×1 pixels. This leads to a total C1 layer size of $25 \times 25 \times 2$ (25×25 for the propagated values and 25×25 to store the orientations). However, the C1–S2 connections need to be set in a way that helps to identify not only the intensity of the propagated signal but the orientation. To implement this in a more biologically plausible way, the C1 layer has to be transformed. Each element on the 25×25 grid would now contain six C1 cells (one for each orientation) with the winner-takes-all condition: Only one out of six cells can fire at a given time. Now, if we want to trace the signal back, we can determine not only its spatial location but also its orientation. This modification leaves the C1 layer with a size of $25 \times 25 \times 6$ with only 25×25 cells firing each time. We have also implemented the lateral inhibition procedure in a way it is done in the Masquelier and Thorpe (41) model. The C1–S2 connections were governed by STDP.

S2 layer. The S2 layer is represented by a fixed number of integrate-and-fire neurons (we used either 10 or 20 neurons). We have also added a sinusoidal function into the equations of S2 cells to model the alpha waves. By doing this, we achieve a clearing effect of periodic modulations. The wave and the images are not perfectly in phase with each other, and the modulation wave is shifted by an angle of $\pi/8$. This provides better elimination of the spontaneous intrinsic activity and signal amplification. To determine which S2 cells encode which numbers, the mature network was presented with a test set of images. To perform such classification, the winner-takes-all strategy is used once more: A stimulus gets propagated to only one cell which fires first. Then each S2 cell is assigned to a number to which it fires the greatest number of times.

Full Network.

GNW. The population is represented by 100 integrate-and-fire neurons. Eighty neurons are excitatory, and they are connected to the S2 layer of the visual cortex network. The excitatory population is divided into selective groups in a way that the number of groups is equal to the number of digits presented in the input (e.g., two groups if we use only images of "0" and "1"). Each selective group is connected to S2 patterns responsible for one particular digit. The other 20 are interneurons, and they play an inhibitory role instead. The connection pattern in the GNW is all to all.

Such a state of the GNW model was used to perform trace conditioning. To create the intermediate level for performing nonconscious delay conditioning, all the interneurons were deleted from the model. The use of inhibitory neurons for trace conditioning is necessary to create a working memory for the highest, conscious level. The functionality of this level can be more complex for different tasks, but the current model shows that the working memory is essential for trace conditioning as one of the basic markers of conscious processing (51). The GNW does not represent a single brain area in particular but a network of long-range connections that is able to interact with multiple different regions.

Motor cortex. The model of the motor cortex uses only one integrate-and-fire neuron. It gets input from the excitatory population of the GNW. Its output is considered to be a sign of transmission of the information to the external environment, and it induces the activation of the reward pathway.

Striatum. To perform the reinforcement learning, we introduce one integrate-and-fire neuron responsible for changing the extracellular dopamine concentration in the GNW. It can be seen as a model of the striatum. The dopamine concentration is set on a baseline and increases or decreases depending on whether the answer was right or wrong. Before the next image is presented, the concentration exponentially decays back to the baseline.

STDP Mechanisms. In our model, there are three areas where synaptic populations are subject to epigenesis: C1–S2 connections in the local network, and GNW–motor cortex and GNW–GNW connections (inhibitory and excitatory) in the full network. We use different models of STDP to set up different connections.

Classical STDP. This is the main plasticity model in our networks that is used in all excitatory connections except GNW–motor cortex. We used the following equations to model the effect of STDP:

$$\tau_{\text{pre}} \frac{d}{dt} A_{\text{pre}} = -A_{\text{pre}} \qquad [1]$$

$$\tau_{\text{post}} \frac{d}{dt} A_{\text{post}} = -A_{\text{post}}. \qquad [2]$$

This system of equations models two possible scenarios: the firing of a presynaptic neuron and the firing of a postsynaptic neuron. In the first case, the value of A_{pre} is updated by adding an increment value $\Delta A_{\text{pre}} > 0$ (models the effect of further potentiation), and the synaptic weight is updated by A_{post} in a given moment of time (depression). In the second case, everything is the opposite: A_{post} is updated by adding an increment value $\Delta A_{\text{post}} < 0$ for future depression, and the weight is potentiated by the value of A_{pre}.

Symmetrical STDP. The symmetrical model of STDP is based on the STDP measured experimentally in the GABAergic neurons population (92). It is used to govern the plasticity of interneurons in the GNW. The mechanisms for the symmetrical STDP are similar to the classical STDP; however, the equations are slightly different,

$$\tau_{\text{pre}} \frac{d}{dt} A_{\text{pre}} = -A_{\text{pre}} - \frac{\Delta A_{\text{pre}}}{15} \qquad [3]$$

$$\tau_{\text{post}} \frac{d}{dt} A_{\text{post}} = -A_{\text{post}} - \frac{\Delta A_{\text{post}}}{15}, \qquad [4]$$

where $\Delta A_{\text{pre}} > 0$ and $\Delta A_{\text{post}} > 0$. The additional terms are added to keep the values of A_{pre} and A_{post} slightly below zero to model the depression when two neurons fire with a large temporal difference. In other cases, when the temporal distance is short enough, the model produces an effect of the long-term potentiation.

Dopamine modulated STDP. This type of synaptic plasticity is used to introduce reinforcement learning which involves multiple brain regions and feedback from the external environment. The model is based on the work of Izhikevich (37)

on linking STDP and dopamine signaling. It is represented by the following set of equations:

$$\tau_d \frac{d}{dt} d = -d \qquad [5]$$

$$\tau_c \frac{d}{dt} c = -c \qquad [6]$$

$$\tau_s \frac{d}{dt} s = cd \qquad [7]$$

$$\tau_{\text{pre}} \frac{d}{dt} A_{\text{pre}} = -A_{\text{pre}} \qquad [8]$$

$$\tau_{\text{post}} \frac{d}{dt} A_{\text{post}} = -A_{\text{post}}. \qquad [9]$$

The variable c is updated in the same way as the synaptic weight is updated during classical STDP (using Eqs. **3** and **4**). The variable d represents the extracellular dopamine concentration: If it is lower than a baseline, then it launches the negative reward mechanism; if it is higher, it launches a positive reward mechanism. The variable s serves as an increment to the synaptic weight and depends not only on the sign and magnitude of the classical STDP but on the sign and magnitude of

the variable d. The release of dopamine is controlled by a single neuron that fires once when a new image is presented with a random delay from 0 s to 1/32 s.

Author affiliations: ᵃNeuroscience Department, Institut Pasteur, 75015 Paris, France; ᵇLaboratory of Computational and Quantitative Biology, Sorbonne Université, 75005 Paris, France; and ᶜMila - Quebec Artificial Intelligence Institute, Centre Hospitalier Universitaire Sainte-Justine Research Center, Department of Psychiatry, Université de Montréal, Montréal, QC H3T 1C5, Canada

1. Y. LeCun, Y. Bengio, G. Hinton, Deep learning. *Nature* **521**, 436–444 (2015).
2. D. Hassabis, D. Kumaran, C. Summerfield, M. Botvinick, Neuroscience-inspired artificial intelligence. *Neuron* **95**, 245–258 (2017).
3. B. A. Richards *et al.*, A deep learning framework for neuroscience. *Nat. Neurosci.* **22**, 1761–1770 (2019).
4. J.-P. Changeux, A. Goulas, C. C. Hilgetag, A connectomic hypothesis for the hominization of the brain. *Cereb. Cortex* **31**, 2425–2449 (2021).
5. S. Dehaene, M. Kerszberg, J.-P. Changeux, A neuronal model of a global workspace in effortful cognitive tasks. *Proc. Natl. Acad. Sci. U.S.A.* **95**, 14529–14534 (1998).
6. D. Kahneman, *Thinking, Fast and Slow* (Macmillan, 2011).
7. Y. Bengio, The consciousness prior. *arXiv* [Preprint] (2019). https://doi.org/10.48550/arXiv.1709.08568. Accessed 13 October 2020.
8. J. C. R. Whittington, R. Bogacz, Theories of error back-propagation in the brain. *Trends Cogn. Sci.* **23**, 235–250 (2019).
9. T. P. Lillicrap, A. Santoro, L. Marris, C. J. Akerman, G. Hinton, Backpropagation and the brain. *Nat. Rev. Neurosci.* **21**, 335–346 (2020).
10. J. Aru, M. Suzuki, M. E. Larkum, Cellular mechanisms of conscious processing. *Trends Cogn. Sci.* **25**, 1096 (2020).
11. D. O. Hebb, *The Organization of Behavior: A Neuropsychological Theory* (Wiley, 1949).
12. M. Botvinick *et al.*, Reinforcement learning, fast and slow. *Trends Cogn. Sci.* **23**, 408–422 (2019).
13. J. P. Changeux, P. Courrège, A. Danchin, A theory of the epigenesis of neuronal networks by selective stabilization of synapses. *Proc. Natl. Acad. Sci. U.S.A.* **70**, 2974–2978 (1973).
14. M. D. Döbrössy *et al.*, Differential effects of learning on neurogenesis: Learning increases or decreases the number of newly born cells depending on their birth date. *Mol. Psychiatry* **8**, 974–982 (2003).
15. M. M. Halassa, P. G. Haydon, Integrated brain circuits: Astrocytic networks modulate neuronal activity and behavior. *Annu. Rev. Physiol.* **72**, 335–355 (2010).
16. S. Chavlis, P. Poirazi, Drawing inspiration from biological dendrites to empower artificial neural networks. *Curr. Opin. Neurobiol.* **70**, 1–10 (2021).
17. J.-P. Changeux, Climbing brain levels of organisation from genes to consciousness. *Trends Cogn. Sci.* **21**, 168–181 (2017).
18. C. Eliasmith *et al.*, A large-scale model of the functioning brain. *Science* **338**, 1202–1205 (2012).
19. S. K. Esser *et al.*, Convolutional networks for fast, energy-efficient neuromorphic computing. *Proc. Natl. Acad. Sci. U.S.A.* **113**, 11441–11446 (2016).
20. R. Chéreau, L. E. Williams, T. Bawa, A. Holtmaat, Circuit mechanisms for cortical plasticity and learning. *Semin. Cell Dev. Biol.* **125**, 68–75 (2021).
21. P. R. Roelfsema, A. Holtmaat, Control of synaptic plasticity in deep cortical networks. *Nat. Rev. Neurosci.* **19**, 166–180 (2018).
22. F. Zenke, W. Gerstner, S. Ganguli, The temporal paradox of Hebbian learning and homeostatic plasticity. *Curr. Opin. Neurobiol.* **43**, 166–176.
23. G. I. Parisi, R. Kemker, J. L. Part, C. Kanan, S. Wermter, Continual lifelong learning with neural networks: A review. *Neural Netw.* **113**, 54–71 (2019).
24. D. S. Bassett, M. S. Gazzaniga, Understanding complexity in the human brain. *Trends Cogn. Sci.* **15**, 200–209 (2011).
25. J. A. S. Kelso, G. Dumas, E. Tognoli, Outline of a general theory of behavior and brain coordination. *Neural Netw.* **37**, 120–131 (2013).
26. A. Vaswani *et al.*, "Attention is all you need" in *Advances in Neural Information Processing Systems*, I. Guyon *et al.*, Eds. (Curran Associates, 2017), **vol. 30**, pp. 5998–6008.
27. O. Vinyals *et al.*, Grandmaster level in StarCraft II using multi-agent reinforcement learning. *Nature* **575**, 350–354 (2019).
28. S. Sabour, N. Frosst, G. E. Hinton, Dynamic routing between capsules. *arXiv* [Preprint] (2017). https://doi.org/10.48550/arXiv.1710.09829. Accessed 4 December 2017.
29. A. Goyal *et al.*, Object files and schemata: Factorizing declarative and procedural knowledge in dynamical systems. *arXiv* [Preprint] (2020). https://doi.org/10.48550/arXiv.2006.16225. Accessed 16 November 2020.
30. A. M. Zador, A critique of pure learning and what artificial neural networks can learn from animal brains. *Nat. Commun.* **10**, 3770 (2019).
31. R. Pfeifer, M. Lungarella, F. Iida, Self-organization, embodiment, and biologically inspired robotics. *Science* **318**, 1088–1093 (2007).

32. G. A. Mashour, P. Roelfsema, J.-P. Changeux, S. Dehaene, Conscious processing and the global neuronal workspace hypothesis. *Neuron* **105**, 776–798 (2020).
33. D. Grover *et al.*, Differential mechanisms underlie trace and delay conditioning in *Drosophila*. *Nature* **603**, 302–308 (2022).
34. A. Newell, J. C. Shaw, H. A. Simon, Elements of a theory of human problem solving. *Psychol. Rev.* **65**, 151–166 (1958).
35. E. M. Izhikevich, Which model to use for cortical spiking neurons? *IEEE Trans. Neural Netw.* **15**, 1063–1070 (2004).
36. H. Markram, W. Gerstner, P. J. Sjöström, Spike-timing-dependent plasticity: A comprehensive overview. *Front. Synaptic Neurosci.* **4**, 2 (2012).
37. E. M. Izhikevich, Solving the distal reward problem through linkage of STDP and dopamine signaling. *Cereb. Cortex* **17**, 2443–2452 (2007).
38. Y. LeCun, C. Cortes, C. Burges, MNIST handwritten digit database. http://yann.lecun.com/exdb/mnist/. Accessed 30 October 2014.
39. Y. Lecun, L. Bottou, Y. Bengio, P. Haffner, Gradient-based learning applied to document recognition. *Proc. IEEE* **86**, 2278–2324 (1998).
40. R. E. Clark, J. R. Manns, L. R. Squire, Classical conditioning, awareness, and brain systems. *Trends Cogn. Sci.* **6**, 524–531 (2002).
41. T. Masquelier, S. J. Thorpe, Unsupervised learning of visual features through spike timing dependent plasticity. *PLOS Comput. Biol.* **3**, e31 (2007).
42. J.-P. Changeux, *Neuronal Man: The Biology of Mind* (Pantheon, 1985).
43. R. Descartes, *Meditationes de Prima Philosophia* (J. Vrin, Paris, 1641).
44. I. Kant, *Kritik der reinen Vernunft* (Preussische Akademie, Berlin, 1781).
45. D. Marr, *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information* (W. H. Freeman, 1982).
46. H. Lagercrantz, J.-P. Changeux, The emergence of human consciousness: From fetal to neonatal life. *Pediatr. Res.* **65**, 255–260 (2009).
47. A. Cleeremans, C. Tallon-Baudry, Consciousness matters: Phenomenal experience has functional value. *Neurosci. Conscious.* **2022**, niac007 (2022).
48. M. J. Arcaro, M. S. Livingstone, A hierarchical, retinotopic proto-organization of the primate visual system at birth. *eLife* **6**, e26196 (2017).
49. M. J. Arcaro, P. F. Schade, J. L. Vincent, C. R. Ponce, M. S. Livingstone, Seeing faces is necessary for face-domain formation. *Nat. Neurosci.* **20**, 1404–1412 (2017).
50. K. S. LaBar, J. F. Disterhoft, Conditioning, awareness, and the hippocampus. *Hippocampus* **8**, 620–626 (1998).
51. T. A. Bekinschtein, M. Peeters, D. Shalom, M. Sigman, Sea slugs, subliminal pictures, and vegetative state patients: Boundaries of consciousness in classical conditioning. *Front. Psychol.* **2**, 337 (2011).
52. T. P. Vogels, H. Sprekeler, F. Zenke, C. Clopath, W. Gerstner, Inhibitory plasticity balances excitation and inhibition in sensory pathways and memory networks. *Science* **334**, 1569–1573 (2011).
53. F. H. Lopes da Silva, J. P. Pijn, D. Velis, P. C. Nijssen, Alpha rhythms: Noise, dynamics and models. *Int. J. Psychophysiol.* **26**, 237–249 (1997).
54. J. Lefebvre, A. Hutt, F. Frohlich, Stochastic resonance mediates the state-dependent effect of periodic stimulation on cortical alpha oscillations. *eLife* **6**, e32054 (2017).
55. E. M. Izhikevich, G. M. Edelman, Large-scale model of mammalian thalamocortical systems. *Proc. Natl. Acad. Sci. U.S.A.* **105**, 3593–3598 (2008).
56. H. de Garis, C. Shuo, B. Goertzel, L. Ruiting, A world survey of artificial brain projects, Part I: Large-scale brain simulations. *Neurocomputing* **74**, 3–29 (2010).
57. T. Gisiger, M. Kerszberg, J.-P. Changeux, Acquisition and performance of delayed-response tasks: A neural network model. *Cereb. Cortex* **15**, 489–506 (2005).
58. T. Gisiger, M. Boukadoum, A loop-based neural architecture for structured behavior encoding and decoding. *Neural Netw.* **98**, 318–336 (2018).
59. C. H. Papadimitriou, S. S. Vempala, D. Mitropolsky, M. Collins, W. Maass, Brain computation by assemblies of neurons. *Proc. Natl. Acad. Sci. U.S.A.* **117**, 14464–14472 (2020).
60. Y. Kremer, J.-F. Léger, D. Goodman, R. Brette, L. Bourdieu, Late emergence of the vibrissa direction selectivity map in the rat barrel cortex. *J. Neurosci.* **31**, 10689–10700 (2011).
61. T. Chauhan, T. Masquelier, A. Montlibert, B. R. Cottereau, Emergence of binocular disparity selectivity through Hebbian learning. *J. Neurosci.* **38**, 9563–9578 (2018).

62. T. B. Brown *et al.*, Language models are few-shot learners. *arXiv* [Preprint] (2020). 10.48550/arXiv.2005.14165. Accessed 16 July 2021.

63. D. Silver *et al.*, A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science* **362**, 1140–1144 (2018).

64. J. R. Searle, Minds, brains, and programs. *Behav. Brain Sci.* **3**, 417–424 (1980).

65. A. M. Turing, Computing machinery and intelligence. *Mind* **LIX**, 433–460 (1950).

66. G. Dumas, G. C. de Guzman, E. Tognoli, J. A. S. Kelso, The human dynamic clamp as a paradigm for social interaction. *Proc. Natl. Acad. Sci. U.S.A.* **111**, E3726–E3734 (2014).

67. C. Moutard, S. Dehaene, R. Malach, Spontaneous fluctuations and non-linear ignitions: Two dynamic faces of cortical recurrent loops. *Neuron* **88**, 194–206 (2015).

68. L. Blum, M. Blum, A theory of consciousness from a theoretical computer science perspective: Insights from the Conscious Turing Machine. *Proc. Natl. Acad. Sci. U.S.A.* **119**, e2115934119 (2022).

69. A. L. Oliveira, A blueprint for conscious machines. *Proc. Natl. Acad. Sci. U.S.A.* **119**, e2205971119 (2022).

70. W. Singer, Recurrent dynamics in the cerebral cortex: Integration of sensory evidence with stored knowledge. *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2101043118 (2021).

71. W. Dabney *et al.*, A distributional code for value in dopamine-based reinforcement learning. *Nature* **577**, 671–675 (2020).

72. N. Dehghani, R. D. Wimmer, A computational perspective of the role of the thalamus in cognition. *Neural Comput.* **31**, 1380–1418 (2019).

73. F. Villagrasa *et al.*, On the role of cortex-basal ganglia interactions for category learning: A neurocomputational approach. *J. Neurosci.* **38**, 9551–9562 (2018).

74. C. J. Whyte, Integrating the global neuronal workspace into the framework of predictive processing: Towards a working hypothesis. *Conscious. Cogn.* **73**, 102763 (2019).

75. L. Chanes, L. F. Barrett, Redefining the role of limbic areas in cortical processing. *Trends Cogn. Sci.* **20**, 96–106 (2015).

76. A. Goyal, Y. Bengio, Inductive biases for deep learning of higher-level cognition. *ArXiv* [Preprint] (2021). https://doil.org/10.48550/arXiv.2011.15091v3. Accessed 9 November 2021.

77. A. Goyal *et al.*, Coordination among neural modules through a shared global workspace. *ArXiv* [Preprint] (2021). https://doil.org/10.48550/arXiv.2103.01197. Accessed 8 November 2021.

78. R. VanRullen, R. Kanai, Deep learning and the global workspace theory. *Trends Cogn. Sci.* **44**(9), 692-704. (2021).

79. A. Safron, An integrated world modeling theory (IWMT) of consciousness: Combining integrated information and global neuronal workspace theories with the free energy principle and active inference framework; Toward solving the hard problem and characterizing agentic causation. *Front Artif Intell* **3**, 30 (2020).

80. S. Dehaene, H. Lau, S. Kouider, What is consciousness, and could machines have it? *Science* **358**, 486–492 (2017).

81. M. S. A. Graziano, S. Kastner, Human consciousness and its relationship to social neuroscience: A novel hypothesis. *Cogn. Neurosci.* **2**, 98–113 (2011).

82. G. Dumas, S. Malesys, T. Bourgeron, Systematic detection of brain protein-coding genes under positive selection during primate evolution and their roles in cognition. *Genome Res.* **31**, 484–496 (2021).

83. T. K. Hensch, Critical period plasticity in local cortical circuits. *Nat. Rev. Neurosci.* **6**, 877–888 (2005).

84. T. K. Hensch, M. Fagiolini, "Excitatory–inhibitory balance and critical period plasticity in developing visual cortex" in *Development, Dynamics and Pathiology of Neuronal Networks: From Molecules to Functional Circuits*, J. van Pelt et al., Eds. (*Progress in Brain Research*, Elsevier, 2005), vol. 147, pp. 115–124.

85. C. Heyes, Précis of cognitive gadgets: The cultural evolution of thinking. *Behav. Brain Sci.* **42**, e169 (2019).

86. E. Thompson, F. J. Varela, Radical embodiment: Neural dynamics and consciousness. *Trends Cogn. Sci.* **5**, 418–425 (2001).

87. C. Fields, J. F. Glazebrook, Disrupted development and imbalanced function in the global neuronal workspace: A positive-feedback mechanism for the emergence of ASD in early infancy. *Cogn. Neurodyn.* **11**, 1–21 (2016).

88. G. Dumas, M. Chavez, J. Nadel, J. Martinerie, Anatomical connectivity influences both intra- and inter-brain synchronizations. *PLoS One* **7**, e36414 (2012).

89. I. Konvalinka, P. Vuust, A. Roepstorff, A coupled oscillator model of interactive tapping. *Proceedings of 7th Triennial Conference of European Society for the Cognitive Sciences of Music (ESCOM 2009)* (European Society for the Cognitive Sciences of Music, 2009), pp. 242–245.

90. S. Bolotta, G. Dumas, Social neuro AI: Social interaction as the "dark matter" of AI. *Front. Comput. Sci.* 4:846440. (2022).

91. K. Volzhenin, Numerical simulations and data for "Multilevel development of cognitive abilities in an artificial neural network." GitHub. https://github.com/ppsp-team/EpigeneSys/tree/master/GNW. Deposited 20 January 2022.

92. M. V. Tsodyks, H. Markram, The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proc. Natl. Acad. Sci. U.S.A.* **94**, 719–723 (1997).