

Theory and Simulation in Neuroscience

Wulfram Gerstner,^{1*} Henning Sprekeler,^{2,3,4} Gustavo Deco^{5,6}

Modeling work in neuroscience can be classified using two different criteria. The first one is the complexity of the model, ranging from simplified conceptual models that are amenable to mathematical analysis to detailed models that require simulations in order to understand their properties. The second criterion is that of direction of workflow, which can be from microscopic to macroscopic scales (bottom-up) or from behavioral target functions to properties of components (top-down). We review the interaction of theory and simulation using examples of top-down and bottom-up studies and point to some current developments in the fields of computational and theoretical neuroscience.

Mathematical and computational approaches in neuroscience have a long tradition that can be followed back to early mathematical theories of perception (1, 2) and of current integration by a neuronal cell membrane (3). Hodgkin and Huxley combined their experiments with a mathematical description, which they used for simulations on one of the early computers (4). Hebb's ideas on assembly formation (5) inspired simulations on the largest computers available at that time (1956) (6). Since the 1980s, the field of theoretical and computational neuroscience has grown enormously (7).

Modern neuroscience methods requiring extensive training have led to a specialization of researchers, such that neuroscience today is fragmented into labs working on genes and molecules, on single-cell electrophysiology, on multineuron recordings, and on cognitive neuroscience and psychophysics, to name just a few. One of the central tasks of computational neuroscience is to bridge these different levels of description by simulation and mathematical theory. The bridge can be built in two different directions. Bottom-up models integrate what is known on a lower level (e.g., properties of ion channels) to explain phenomena observed on a higher level [e.g., generation of action potentials (4, 8–10)]. Top-down models, on the other hand, start with known cognitive functions of the brain (e.g., working memory), and deduce from these how components (e.g., neurons or groups of neurons)

should behave to achieve these functions. Influential examples of the top-down approach are theories of associative memory (11, 12), reinforcement learning (13, 14), and sparse coding (15, 16).

Bottom-up and top-down models can be studied either by mathematical theory (theoretical neuroscience) or by computer simulation (computational neuroscience). Theory has the advantage of providing a complete picture of the model behavior for all possible parameter settings, but analytical solutions are restricted to relatively simple models. The aim of theory is therefore to purify biological ideas to the bare minimum, so as to arrive at a “toy model” that crystallizes a concept in a set of mathematical equations that can be fully understood. Simulations, in contrast, can be applied to all models, simplified as well as complex ones, but they can only sample the model behavior for a limited set of parameters. If the relevant parameters were known, the whole or a major fraction of the brain could be simulated using known biophysical components—a prospect that has inspired recent large-scale simulation projects (17–21).

How can theory and simulation interact to contribute to our understanding of brain function? In this article, we illustrate the interaction with four case studies and outline future avenues for synergies between theory and simulation. Even though it forms an important subfield of computational neuroscience, the analysis of neuronal data (22–24) is not included in this review.

From Detailed to Abstract Models of Neural Activity: Bottom-Up Theory

The brain contains billions of neurons that communicate by short electrical pulses, called action potentials or spikes. Hodgkin and Huxley's description of neuronal action potentials (4) today provides a basis for standard simulator software (25, 26) and, more generally, a widely used framework for biophysical neuron models. In these models, each patch of cell membrane is described by its composition of ion channels, with specific

time constants and gating dynamics that control the momentary state (open or closed) of a channel (Fig. 1C).

By a series of mathematical steps and approximations, theory has sketched a systematic bottom-up path from such biophysical models of single neurons to macroscopic models of neural activity. In a first step, biophysical models of spike generation are reduced to integrate-and-fire models (27–32) where spikes occur whenever the membrane potential reaches the threshold (Fig. 1B). In the next abstraction step, the population activity $A(t)$ —defined as the total number of spikes emitted by a population of interconnected neurons in a short time window—is predicted from the properties of individual neurons (33–35) using mean-field methods known from physics: Because each neuron receives input from many others, it is sensitive only to their average activity (“mean field”) but not to the activity patterns of individual neurons. Instead of the spike-based interaction among thousands of neurons, network activity can therefore be described macroscopically as an interaction between different populations (33, 35, 36). Such macroscopic descriptions—known as population models, neural mass models, or, in the continuum limit, neural field models (Fig. 1A)—help researchers to gain an intuitive and more analytical understanding of the principal activity patterns in large networks.

Although the transition from microscopic to macroscopic scales relies on purely mathematical arguments, simulations are important to add aspects of biological realism (such as heterogeneity of neurons and connectivity, adaptation on slower time scales, and variability of input and receptive fields) that are difficult to treat mathematically. However, the theoretical concepts and the essence of the phenomena are often robust with respect to these aspects.

Decision-Making: Theory Combines Top-Down and Bottom-Up

Many times a day we are confronted with a decision between two alternatives (A or B), such as “Should I turn left or right?” Psychometric measures of performance and reaction times for two-alternative forced-choice decision-making paradigms can be fitted by a phenomenological drift-diffusion model (37). This model consists of a diffusion equation describing a random variable that accumulates noisy sensory evidence until it reaches one of two boundaries corresponding to a specific choice (Fig. 2A). Despite its success in explaining reaction time distributions, the drift-diffusion model suffers from a crucial disadvantage, namely the difficulty in assigning a biological meaning to the model parameters.

Recently, neurophysiological experiments have begun to reveal neuronal correlates of decision-making, in tasks involving visual patterns of moving random dots (41, 42) or vibrotactile (43) or auditory frequency comparison (44). Compu-

¹School of Computer and Communication Sciences and Brain Mind Institute, School of Life Sciences, Ecole Polytechnique Fédérale de Lausanne, 1015 Lausanne, Switzerland. ²Technische Universität Berlin, 10587 Berlin, Germany. ³Humboldt-Universität zu Berlin, 10115 Berlin, Germany. ⁴Bernstein Center for Computational Neuroscience Berlin, 10115 Berlin, Germany. ⁵Center for Brain and Cognition, Department of Information and Communications Technologies, Universitat Pompeu Fabra, Roc Boronat 138, 08018 Barcelona, Spain. ⁶Institució Catalana de Recerca i Estudis Avançats, Passeig Lluís Companys, 2308010 Barcelona, Spain.

*To whom correspondence should be addressed. E-mail: wulfram.gerstner@epfl.ch

tational neuroscience offers a framework to bridge the conceptual gap between the cellular and the behavioral level. Explicit simulations of microscopic models based on local networks with large numbers of spiking neurons can reproduce and explain both the neurophysiological and behavioral data (39, 45). These models describe the interactions between two groups of neurons coupled through mutually inhibitory connections (Fig. 2C). Suitable parameters are inferred by studying the dynamical regimes of the system and choosing parameters consistent with the experimental observations of decision behavior. Thus, the pure bottom-up path is complemented by the top-down insights of target functions that the network needs to achieve.

Theory has provided a way to identify the connection between the simulation-based biophysical level and the phenomenological drift-diffusion models that were used to quantify behavioral data in the past. This analysis first involves reducing the system of millions of differential equations to only two equations, by the mean-field techniques (33) discussed in the previous paragraph (see Fig. 1A). In this reduced system, different dynamical regimes can be studied by means of a bifurcation diagram (Fig. 2B), leading to a picture in which the dynamics of decision-making can be visualized as a ball rolling down into one of the minima of a multi-well energy landscape. Close to the bifurcation point (DD), finally, the dynamics can be further reduced to a nonlinear one-dimensional drift-diffusion model (46) (Fig. 2A).

Hebbian Assemblies and Associative Memories: Top-Down Concepts

The name of Hebb is attached to two different ideas that are intimately linked (5): the Hebbian cell assembly and the Hebbian learning rule. The former states that mental concepts are represented by the joint activation of groups of cells (assemblies). The latter postulates that these assemblies are formed by strengthening the connections between neurons that are “repeatedly and persistently active together.” These strong connections enable the network to perform an associative retrieval of memories (Fig. 3A): If a neuron that is part of the assembly does not receive an external stimulus sufficient to trigger

its firing, it can nonetheless be activated by lateral excitatory input from active partners in the assembly, such that eventually the stored activity pattern is completed and the full assembly—and thus the memory—is retrieved. Although Hebb wrote down his ideas in words, these soon inspired mathematical models of learning rules and large-scale simulation studies (6). Whereas

the simulations pointed to limitations of Hebb’s original ideas in practical applications, mathematical studies (11, 48, 49) cleared the first paths through the jungle of possible model configurations toward a working model of memory retrieval. In analogy to models of magnetic systems and spin glasses in statistical physics, the picture of memory retrieval as movement toward

a minimum in an energy landscape emerged (Fig. 3B), first in networks of binary units (12, 50) and later for rate models (51, 52). In these frameworks, important questions could be answered, such as that of memory capacity. Because each neuron can participate in many different assemblies, memory capacity is difficult to estimate by pure verbal reasoning. Theory shows that the number of different random patterns that can be stored scales linearly with the number of neurons (53): A network with 10,000 neurons can store about 1000 patterns, but if we double the number of neurons, we can store twice as many.

It took many steps from these abstract concepts to arrive at biologically plausible models of memory. Whereas the energy landscape analogy is restricted to a limited class of abstract models, the concept of memory retrieval as movement toward an attractor state (Fig. 3C) turned out to be robust with respect to model details. The top-down path (Fig. 3, C to E) of model development includes the transition from high-activity states to low-activity states (54); from networks with complete and reciprocal connectivity to ones with sparse and asymmetric connections (55); tuning of inhibition in sparsely active networks (56); transition from rate neurons to spiking neurons (34, 57, 58); addition of a spontaneous state at low firing rates (59, 60); homeostatic control of threshold or plasticity (47, 61); and further steps still to be taken, so as to solve the issue of online learning in associative memory networks (62). The above historical sketch shows the flow of ideas from top to bottom, by adding biological realism and details to well-understood abstract concepts. As research moves along this path, numerical simulations gain in importance, but theory remains the guiding principle. Given the technical challenges of detecting distributed but jointly active neuronal assemblies in the brain, it is not surprising that experimental evidence for associative memories as

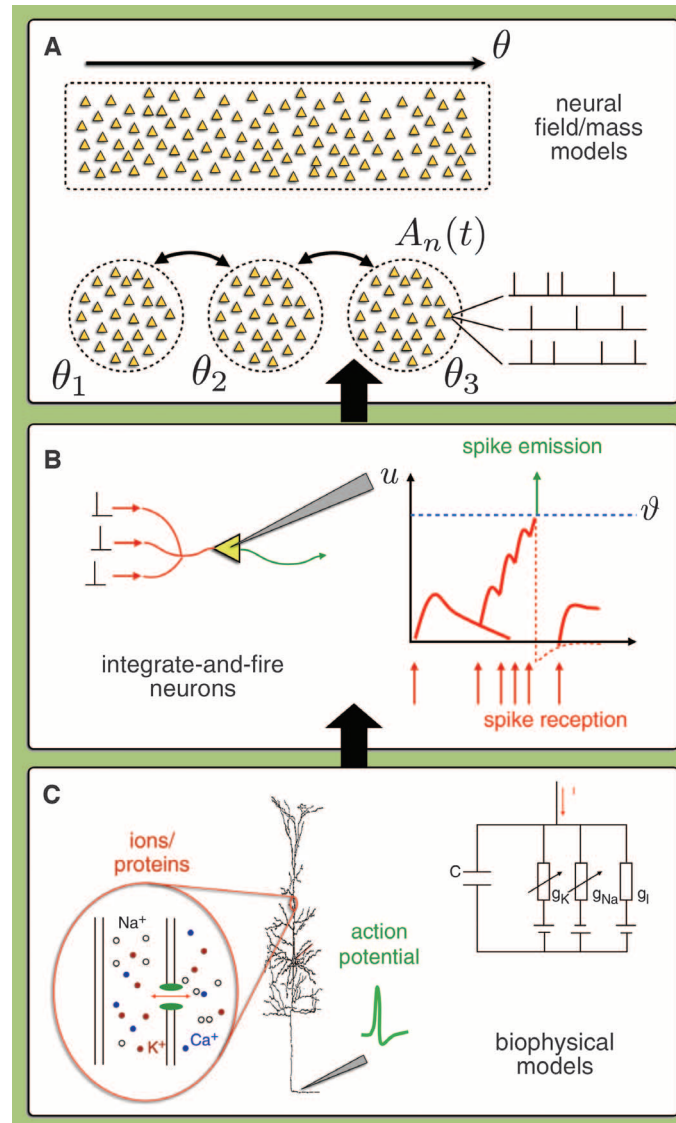


Fig. 1. Bottom-up abstraction scheme for neuronal modeling. **(A)** Neural mass model at the macroscopic scale. **(B)** Integrate-and-fire point neuron model. **(C)** Biophysical neuron model. The ion currents flowing through channels in the cell membrane (left) are characterized by an equivalent electrical circuit comprising a capacity C and a set of time-dependent conductances g , one for each channel type (right). These currents can generate action potentials (green). At the next level of abstraction (B), action potentials are treated as formal events (“spikes”) generated whenever the membrane potential u (solid red line) crosses a threshold ϑ (dashed blue line). After each spike the voltage is reset (dashed red line). Arrivals of spikes from other neurons (red arrows) generate postsynaptic potentials. At the highest level of abstraction (A), groups of neurons interact by their population activity $A_n(t)$ derived mathematically from neuronal parameters. In a continuum description (neural field model), different neuronal populations are characterized by their input characteristics such as the preferred orientation θ of a visual stimulus.

attractor states remains scarce and indirect: Neither invariant sensory representations as observed in individual high-level neurons (63) nor persistent activity during working memory tasks (64, 65) are sufficient to prove attractor dynamics, and the interpretation of multiunit recordings during remapping in the hippocampus as a signature of attractors (66) has been questioned (67), so that novel experimental approaches that enable large-scale recordings of complete neural ensembles (68–70) will be necessary.

Reward-Based Learning: From Behavior to Synapse

Hebbian assemblies play a role in models of decision-making (e.g., two populations representing choices A or B), but the assemblies are fixed and do not change. As such, these models describe static stimulus-response mappings that do not take into account that we can alter our decision strategies, that we can learn from experience. Theories of behavioral learning are an example for a top-down approach that has led from psy-

chology to models of synaptic plasticity, which go beyond a simple Hebbian learning rule.

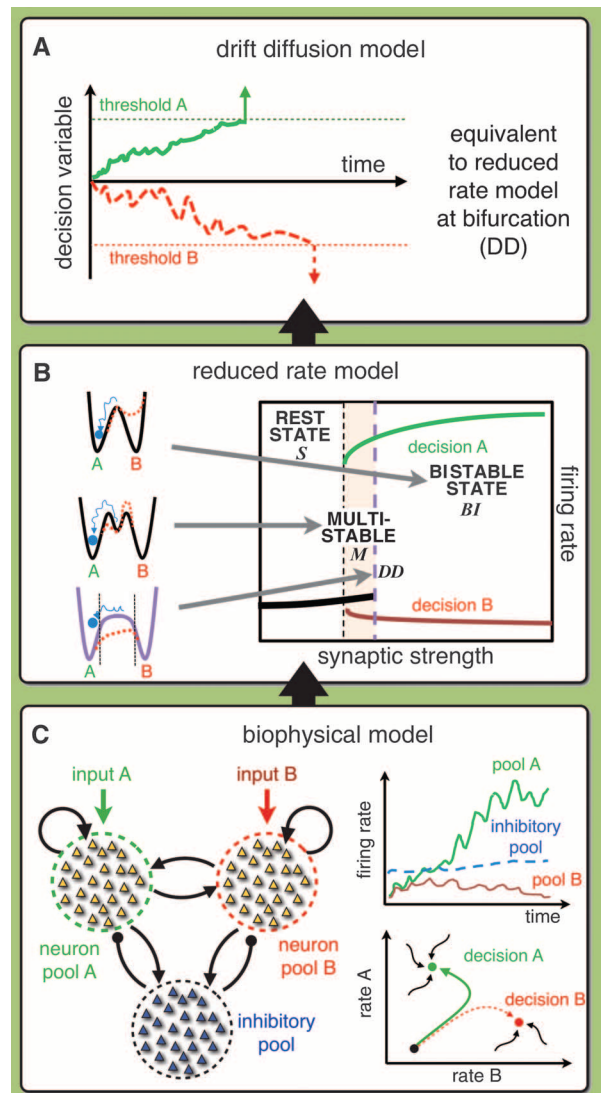
Models of behavioral learning date back to early experimental psychology. Thorndike's "law of effect" (71) stated that animals use behavioral alternatives more often when they were rewarded in the past. The mathematical theory of Rescorla and Wagner (13) extended this idea and suggested that it is not reward per se that drives learning, but the discrepancy between actual and expected outcome—an insight that is also essential in models of temporal difference learning (72). In the 1980s, this line of research on trial-and-error learning in animal psychology and artificial intelligence joined the parallel, mainly engineering-driven line of control theory (73) to form the large research field of reinforcement learning (74).

Although reinforcement learning often makes use of neural architectures (75, 76), its main focus is algorithmic. The renewed interest of computational and theoretical neuroscience in reward-based learning was triggered in the 1990s by two physiological findings. First, Schultz and col-

leagues discovered that dopamine neurons in the midbrain respond to unexpected rewards (77) with activity patterns that resemble the reward prediction error of temporal difference learning (14). Moreover, it was shown that synaptic plasticity, long hypothesized to form the neural basis of learning (78), is under dopaminergic control (79–81) and is thus modulated by reward prediction errors. These findings have led to the idea that the traditional Hebbian view of synaptic plasticity driven by pre- and postsynaptic activity must be augmented by reward prediction error as a third factor (80). The current tasks of computational neuroscience are to compare the learning rules suggested by the top-down approach with experimental data, and to generalize existing concepts in order to evaluate whether, and under what conditions, three-factor learning rules (82–84) of reward-modulated Hebbian synaptic plasticity can be useful at the macroscopic level of networks (85, 86) and behavior (87, 88).

Reinforcement learning is a textbook example for a synergistic interaction of theory

Fig. 2. Binary decision-making. **(A)** Behavioral level: Drift-diffusion model for reaction time experiments (37). A decision is taken (arrow) when the decision variable hits a threshold (first trial, green, choice A; second trial, red, choice B). **(B)** Mesoscopic level: Mathematical model of interacting populations. A decision can be visualized (left) as a ball moving down the energy landscape into one of the minima corresponding to A or B. The black lines in the energy landscape correspond to an unbalanced (50-50) free choice (i.e., no evidence in favor of any decision); the dotted red lines represent the cases with evidence for decision A. The three different landscapes correspond to different parameter regimes, indicated in the bifurcation diagram (right). In the multistable region (M), the spontaneous state is stable, but noise can cause a transition to a decision. In the bistable region (BI), a binary choice is enforced. At the bifurcation point (DD), the landscape around the spontaneous state is flat and the decision dynamics can be further reduced to the drift-diffusion process in (A). The vertical dashed line in the energy landscape indicates the "point of no return" and correspond to the decision thresholds in (A). **(C)** Multineuron spiking model (38, 39). Two pools of neurons (left) receive input representing evidence for choice option A or B, respectively, while shared inhibition leads to a competition between the pools. In the absence of stimuli, the system is spontaneously active, but if a stimulus is presented, the spontaneous state destabilizes and the dynamics evolve toward one of the two decision states (attractor states, bottom right). During a decision for option A, the firing rate of neurons in pool A increases (top right, green line). The mathematical model in (B) can be derived from the spiking model in (C) by mean-field methods (40). The parameter of the bifurcation diagram in (B) is the strength of excitatory self-interaction of neurons in (C).



and simulation. Most algorithms are based on mathematical theory, but without an actual implementation and simulation, it is difficult to predict how they perform under realistic conditions. The undesired increase in learning time with problem complexity observed in simulations of most algorithms poses challenges for the future. The solution could lie in efficient representations of actions (89) as well as environmental and bodily states (90), possibly adapted through Hebbian learning; place cells could serve here as an example (91, 92). Large-scale simulations will be key to evaluating whether neuronal representations across different brain areas are apt to turn complex reward-based learning tasks into simple ones, or whether we need a paradigm shift in reward-based learning.

Large-Scale Simulation Needs Theory

With the growth of computer power, the notion of large-scale simulation continues to change. In the 1950s, networks of 512 binary neurons were explored on the supercomputers of that time (6); both network size and biological realism increased in the 1980s to 9900 detailed model neurons (93). Current simulations of networks go up to 10^9 (94) or 10^{11} (95) integrate-and-fire neurons or 10^6 multicompartment integrate-and-fire models with synaptic plasticity (96).

The structure of neural networks, which is intrinsically distributed and parallel, lends itself to implementations on highly parallel computing devices [e.g., (97, 98)] and has triggered specialized “neuromorphic” design of silicon circuits (99, 100). Future implementations of large networks of integrate-and-fire neurons with synaptic plasticity on specialized parallel hardware (101–104) should run much faster than biological real time, opening the path toward rapid exploration of learning paradigms. Current large-scale implementations on general-purpose or specialized computing devices are feasibility studies suggesting that the simulation of neural networks of the size of real brains is possible. But before such simulations become a tool of brain science, major challenges need to be addressed—and this is where theory comes back into play.

First, theory supports simulations across multiple spatial scales. Even on the biggest computers it is impossible to simulate the whole brain at a molecular resolution. However,

in (for example) a study of the interaction of a drug with a synapse, the synapse could be simulated at the molecular level; the neuron on which the synapse sits could be simulated at the level of biophysical model neurons; the brain area in which the circuit is embedded could be simulated at the level of integrate-and-fire models;

and the remaining brain areas can be summarized in population equations that describe their mean activity. Theory provides the methods to systematically bridge the scales while ensuring the self-consistency of the model.

Second, theories guide simulations of learning. Because learning (at the behavioral level)

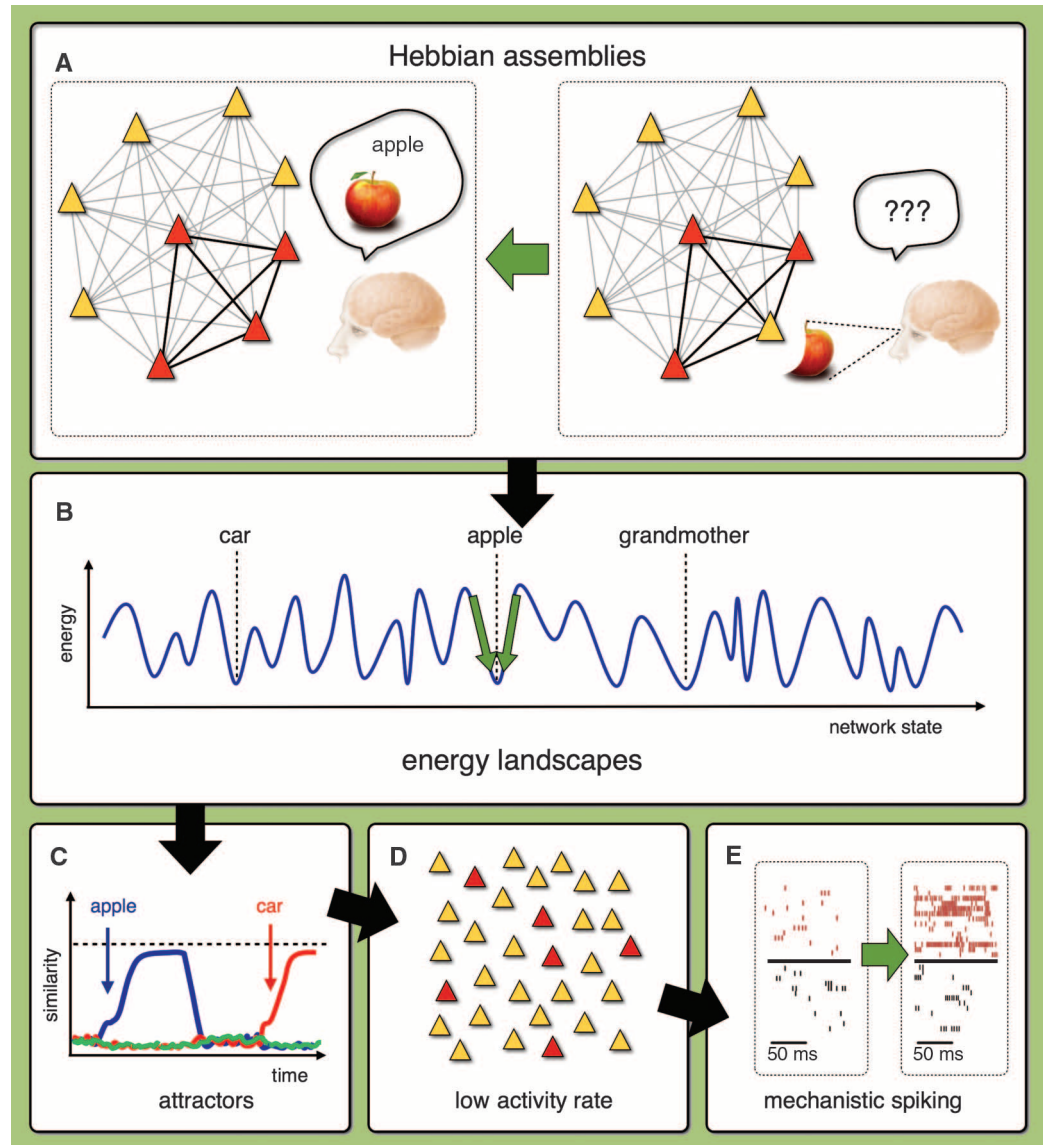


Fig. 3. Top-down evolution of ideas and models from Hebbian assemblies to associative memories. (A) Left: According to Hebb, neurons that are active together (red triangles) during a concept such as “apple” form a cell assembly with stronger excitatory connections with each other (thick lines) relative to connections to and from inactive neurons (yellow). Right: Partial information is sufficient to excite the inactive neuron of the assembly and retrieve the full concept (green arrow indicates flow of network dynamics). (B) The dynamics can be visualized in an energy landscape (energy as a function of network state) as a downward flow into the well corresponding to “apple”; other concepts (car, grandmother) correspond to other minima of the energy. (C) Similarity of the present network state with one of the assemblies. After a short presentation of partial information resembling the concept “apple,” the network moves toward an attractor state with high similarity (blue) to the corresponding assembly. After a reset, the network is ready to respond to a second stimulus (“car,” red). Similarity with other learned patterns (green) remains nearly unchanged during retrieval of one of the memories. The dashed horizontal line indicates maximal similarity, and therefore identity, with a stored assembly. (D) In a brainlike neural network, only few units have elevated firing rates (red). (E) Memory retrieval from partial information also works in networks of excitatory (red) and inhibitory (black) spiking neurons. In the spontaneous state (left), each neuron emits only a few spikes (vertical bars along a horizontal line); during memory retrieval, a few excitatory neurons fire many spikes (right). [Adapted from (47)]

and long-term plasticity of synapses (at the cellular level) evolve on the time scale of hours, experimental progress is slow and data to constrain synaptic plasticity rules are scarce. Even before a candidate plasticity rule is implemented in a simulation, theories allow us to predict, in some cases, that a given rule cannot work for the task at hand; that another rule should work, but might be slow; and that yet another rule should work perfectly if combined with some finely tuned homeostatic process, etc.

Third, regularization theories provide a framework for optimization of parameters. Even at the level of a single neuron, the number of parameters is huge because there are about 200 different types of ion channels (105), the density of which varies along the dendrite. Moreover, in a network of billions of neurons, each with thousands of connections, the number of parameters to specify the connectivity is daunting. How then, can experimental findings ever sufficiently constrain such detailed models?

A solution to this problem could be provided by regularization theory, which penalizes parameter settings that deviate from what is considered plausible. The challenge in the application of regularization to biologically detailed neural networks consists in designing appropriate regularization terms that summarize plausible prior assumptions in a transparent fashion. For example, ion channel distributions along a dendrite can be regularized by imposing plausible density profiles (10, 106) and penalizing the use of more than a few channels for any neuron (107). Connectivity can be regularized by penalizing deviations from simple “connectivity rules” (108, 109).

Fourth, theory drives understanding. It is extremely difficult to control a complex simulation if we do not have an intuition of the basic mechanisms at work. Mathematical abstraction forces us to express intuitions rigorously—to turn ideas into toy models. These not only make qualitative predictions of how a complex simulation should behave under changes of parameters, but also provide the understanding necessary to communicate ideas to others.

Theory Needs Simulations

Although theoretical concepts exist for some brain functions, many of our cognitive abilities still await neurally plausible models—for example, verbal and mathematical reasoning or representations of language and music. Will existing and future concepts developed in mathematical toy models be transferable to systems of the size of a brain with realistic input and output? We believe so, but extensions to large systems cannot currently be tested. In our opinion, mathematical toy models will continue to play a major role in guiding the way we think about neuroscience. However, initiatives for shared modular and reusable code and standardized simulator interfaces (110–113) are important; today the biggest problems addressed in computational neuroscience are often limited to the size of one Ph.D. project.

More generally, the community of theoretical and computational neuroscience would profit from a simulation environment where the ideas developed in the toy models could be tested on a larger scale, in a biologically plausible setting, and where the ideas arising in different communities and labs are finally connected to the bigger whole.

References and Notes

- H. Helmholtz, *Die Lehre von den Tonempfindungen als physiologische Grundlage für die Theorie der Musik* (Vieweg, Braunschweig, ed. 6, 1913). (First published in 1862)
- E. Mach, *Die Analyse der Empfindungen* (Gustav Fischer, Jena, ed. 5, 1906), chapter X; www.uni-leipzig.de/~psycho/wundt/operamach/empfindng/AlysEmln.htm.
- L. Lapicque, *J. Physiol. Pathol. Gen.* **9**, 620 (1907).
- A. L. Hodgkin, A. F. Huxley, *J. Physiol.* **117**, 500 (1952).
- D. O. Hebb, *The Organization of Behavior* (Wiley, New York, 1949).
- N. Rochester, J. Holland, L. Haibt, W. Duda, *IRE Trans. Inf. Theory* **2**, 80 (1956).
- L. F. Abbott, *Neuron* **60**, 489 (2008).
- W. Rall, in *Methods in Neuronal Modeling*, C. Koch, I. Segev, Eds. (MIT Press, Cambridge, MA, 1989), pp. 9–62.
- Z. F. Mainen, J. Joerges, J. R. Huguenard, T. J. Sejnowski, *Neuron* **15**, 1427 (1995).
- E. Hay, S. Hill, F. Schürmann, H. Markram, I. Segev, *PLOS Comput. Biol.* **7**, e1002107 (2011).
- D. J. Willshaw, O. P. Buneman, H. C. Longuet-Higgins, *Nature* **222**, 960 (1969).
- J. J. Hopfield, *Proc. Natl. Acad. Sci. U.S.A.* **79**, 2554 (1982).
- R. A. Rescorla, A. R. Wagner, in *Classical Conditioning II: Current Research and Theory*, A. H. Black, W. F. Prokasy, Eds. (Appleton Century Crofts, New York, 1972), pp. 64–99.
- W. Schultz, P. Dayan, P. R. Montague, *Science* **275**, 1593 (1997).
- H. B. Barlow, in *Sensory Communication*, W. A. Rosenbluth, Ed. (MIT Press, Cambridge, MA, 1961), pp. 217–234.
- B. A. Olshausen, D. J. Field, *Curr. Opin. Neurobiol.* **14**, 481 (2004).
- DARPA SyNAPSE Program (www.artificialbrains.com/darpa-synapse-program).
- BrainCorporation (<http://braincorporation.com>).
- S. Lang, V. J. Dercksen, B. Sakmann, M. Oberlaender, *Neural Netw.* **24**, 998 (2011).
- BrainScales (<https://brainscales.kip.uni-heidelberg.de>).
- H. Markram, *Nat. Rev. Neurosci.* **7**, 153 (2006).
- F. Rieke, D. Warland, R. de Ruyter van Steveninck, W. Bialek, *Spikes—Exploring the Neural Code* (MIT Press, Cambridge, MA, 1996).
- J. W. Pillow *et al.*, *Nature* **454**, 995 (2008).
- R. Quiñero, S. Panzeri, *Nat. Rev. Neurosci.* **10**, 173 (2009).
- M. L. Hines, N. T. Carnevale, *Neural Comput.* **9**, 1179 (1997).
- J. M. Bower, D. Beeman, *The Book of Genesis* (Springer, New York, 1995).
- A. V. M. Herz, T. Gollisch, C. K. Machens, D. Jaeger, *Science* **314**, 80 (2006).
- W. M. Kistler, W. Gerstner, J. L. van Hemmen, *Neural Comput.* **9**, 1015 (1997).
- B. Ermentrout, *Neural Comput.* **8**, 979 (1996).
- N. Brunel, *J. Neurosci.* **23**, 11628 (2003).
- R. Jolivet, T. J. Lewis, W. Gerstner, *J. Neurophysiol.* **92**, 959 (2004).
- R. Brette, W. Gerstner, *J. Neurophysiol.* **94**, 3637 (2005).
- N. Brunel, *Comput. Neurosci.* **8**, 183 (2000).
- S. Fusi, M. Mattia, *Neural Comput.* **11**, 633 (1999).
- W. Gerstner, *Neural Comput.* **12**, 43 (2000).
- H. R. Wilson, J. D. Cowan, *Biophys. J.* **12**, 1 (1972).
- R. Ratcliff, J. N. Rouder, *Psychol. Sci.* **9**, 347 (1998).
- X. J. Wang, *Neuron* **60**, 215 (2008).
- G. Deco, E. T. Rolls, R. Romo, *Prog. Neurobiol.* **88**, 1 (2009).
- K. F. Wong, X. J. Wang, *J. Neurosci.* **26**, 1314 (2006).
- J. I. Gold, M. N. Shadlen, *Trends Cogn. Sci.* **5**, 10 (2001).
- J. I. Gold, M. N. Shadlen, *Annu. Rev. Neurosci.* **30**, 535 (2007).
- R. Romo, E. Salinas, *Nat. Rev. Neurosci.* **4**, 203 (2003).
- L. Lemus, A. Hernández, R. Romo, *Proc. Natl. Acad. Sci. U.S.A.* **106**, 14640 (2009).
- X.-J. Wang, *Neuron* **36**, 955 (2002).
- A. Roxin, A. Ledberg, *PLOS Comput. Biol.* **4**, e1000046 (2008).
- T. P. Vogels, H. Sprekeler, F. Zenke, C. Clopath, W. Gerstner, *Science* **334**, 1569 (2011).
- J. A. Anderson, *Math. Biosci.* **14**, 197 (1972).
- T. Kohonen, *IEEE Trans. Comput.* **C-21**, 353 (1972).
- W. A. Little, *Math. Biosci.* **19**, 101 (1974).
- M. A. Cohen, S. Grossberg, *IEEE Trans. Syst. Man Cybern.* **13**, 815 (1983).
- J. J. Hopfield, *Proc. Natl. Acad. Sci. U.S.A.* **81**, 3088 (1984).
- D. J. Amit, H. Gutfreund, H. Sompolinsky, *Phys. Rev. Lett.* **55**, 1530 (1985).
- D. J. Amit, H. Gutfreund, H. Sompolinsky, *Phys. Rev. A* **35**, 2293 (1987).
- B. Derrida, E. Gardner, A. Zippelius, *Europhys. Lett.* **4**, 167 (1987).
- D. Golomb, N. Rubin, H. Sompolinsky, *Phys. Rev. A* **41**, 1843 (1990).
- W. Gerstner, J. L. van Hemmen, *Network* **3**, 139 (1992).
- D. J. Amit, M. V. Tsodyks, *Network* **2**, 259 (1991).
- D. J. Amit, N. Brunel, *Network* **8**, 373 (1997).
- N. Brunel, V. Hakim, *Neural Comput.* **11**, 1621 (1999).
- J. Triesch, *Neural Comput.* **19**, 885 (2007).
- S. Fusi, *Biol. Cybern.* **87**, 459 (2002).
- R. Q. Quiroga, L. Reddy, G. Kreiman, C. Koch, I. Fried, *Nature* **435**, 1102 (2005).
- J. M. Fuster, G. E. Alexander, *Science* **173**, 652 (1971).
- E. K. Miller, J. D. Cohen, *Annu. Rev. Neurosci.* **24**, 167 (2001).
- T. J. Wills, C. Lever, F. Cacucci, N. Burgess, J. O’Keefe, *Science* **308**, 873 (2005).
- K. J. Jeffery, *Neural Plast.* **2011**, 182602 (2011).
- D. S. Greenberg, A. R. Houweling, J. N. D. Kerr, *Nat. Neurosci.* **11**, 749 (2008).
- B. F. Grewe, F. Helmchen, *Curr. Opin. Neurobiol.* **19**, 520 (2009).
- C. Koch, R. C. Reid, *Nature* **483**, 397 (2012).
- E. L. Thorndike, *Animal Intelligence* (Hafner, Darien, CT, 1911).
- R. Sutton, *Mach. Learn.* **3**, 9 (1988).
- D. P. Bertsekas, *Dynamic Programming and Stochastic Control* (Academic Press, New York, 1976).
- R. Sutton, A. Barto, *Reinforcement Learning* (MIT Press, Cambridge, MA, 1998).
- A. G. Barto, R. S. Sutton, C. W. Anderson, *IEEE Trans. Syst. Man Cybern.* **13**, 835 (1983).
- R. Sutton, A. Barto, in *Learning and Computational Neuroscience: Foundations of Adaptive Networks*, M. Gabriel, J. Moore, Eds. (MIT Press, Cambridge, MA, 1990), pp. 497–537.
- P. Apicella, T. Ljungberg, E. Scarnati, W. Schultz, *Exp. Brain Res.* **85**, 491 (1991).
- S. J. Martin, P. D. Grimwood, R. G. M. Morris, *Annu. Rev. Neurosci.* **23**, 649 (2000).
- P. Calabresi, R. Maj, N. B. Mercuri, G. Bernardi, *Neurosci. Lett.* **142**, 95 (1992).
- J. N. J. Reynolds, J. R. Wickens, *Neural Netw.* **15**, 507 (2002).
- V. Pawlak, J. R. Wickens, A. Kirkwood, J. N. D. Kerr, *Front. Synaptic Neurosci.* **2**, 146 (2010).
- X. Xie, H. S. Seung, *Phys. Rev. E* **69**, 041909 (2004).
- E. M. Izhikevich, *Cereb. Cortex* **17**, 2443 (2007).
- R. V. Florian, *Neural Comput.* **19**, 1468 (2007).
- N. Frémaux, H. Sprekeler, W. Gerstner, *J. Neurosci.* **30**, 13326 (2010).

86. R. Legenstein, S. M. Chase, A. B. Schwartz, W. Maass, *J. Neurosci.* **30**, 8400 (2010).
87. Y. Loewenstein, H. S. Seung, *Proc. Natl. Acad. Sci. U.S.A.* **103**, 15224 (2006).
88. M. H. Herzog, K. C. Aberg, N. Frémaux, W. Gerstner, H. Sprekeler, *Vision Res.* **61**, 95 (2012).
89. A. G. Barto, S. Mahadevan, *Discrete Event Dyn. Syst.* **13**, 341 (2003).
90. S. Mahadevan, M. Maggioni, *J. Machine Learn. Res.* **8**, 2169 (2007).
91. D. J. Foster, R. G. Morris, P. Dayan, *Hippocampus* **10**, 1 (2000).
92. D. Sheynikhovich, R. Chavarriaga, T. Strösslin, A. Arleo, W. Gerstner, *Psychol. Rev.* **116**, 540 (2009).
93. R. D. Traub, R. Miles, R. K. S. Wong, *IEEE Eng. Med. Biol.* **7**, 31 (1988).
94. R. Ananthanarayanan, S. K. Esser, H. D. Simon, D. S. Modha, in *Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis* (2009); <http://dl.acm.org/citation.cfm?id=1654059.1654124>.
95. E. Izhikevich, Simulation of large scale brain models (www.izhikevich.org/human_brain_simulation), 2005.
96. E. M. Izhikevich, G. M. Edelman, *Proc. Natl. Acad. Sci. U.S.A.* **105**, 3593 (2008).
97. K. Obermayer, H. Ritter, K. Schulten, *Parallel Comput.* **14**, 381 (1990).
98. E. Niebur, D. Brettle, in *Advances in Neural Information Processing Systems*, J. Cowan, G. Tesauro, J. Alspecter, Eds. (Morgan Kaufmann, San Mateo, CA, 1994), vol. 6, pp. 904–910.
99. C. Mead, Analog VLSI and neural systems (Addison-Wesley, Boston, 1989).
100. G. Indiveri *et al.*, *Front. Neurosci.* **5**, 73 (2011).
101. J. Schemmel, D. Brüderle, K. Meier, B. Ostendorf, in *IEEE International Symposium on Circuits and Systems*, 10.1109/ISCAS.2007.378289 (2007).
102. S. Furber, A. Brown, in *Application of Concurrency to System Design*, 10.1109/ACSD.2009.17 (2009).
103. M. Xin *et al.*, *Comput. Sci. Eng.* **12**, 91 (2010).
104. S. Millner, A. Grübl, K. Meier, J. Schemmel, M.-O. Schwartz, in *Advances in Neural Information Processing Systems*, J. Lafferty, C. K. I. Williams, J. Schawe-Taylor, R. S. Zemel, A. Culotta, Eds., <http://books.nips.cc/nips23.html> (2010).
105. R. Ranjan *et al.*, *Front. Neuroinform.* **5**, 36 (2011).
106. Q. J. M. Huys, L. Paninski, *PLOS Comput. Biol.* **5**, e1000379 (2009).
107. M. Toledo-Rodriguez *et al.*, *Cereb. Cortex* **14**, 1310 (2004).
108. T. Binzegger, R. J. Douglas, K. A. C. Martin, *J. Neurosci.* **24**, 8441 (2004).
109. S. Ramaswamy *et al.*, *J. Physiol.* **590**, 737 (2012).
110. A. P. Davison *et al.*, *Front. Neuroinform.* **2**, 11 (2008).
111. R. Brette *et al.*, *J. Comput. Neurosci.* **23**, 349 (2007).
112. P. Gleeson, V. Steuber, R. A. Silver, *Neuron* **54**, 219 (2007).
113. P. Gleeson *et al.*, *PLOS Comput. Biol.* **6**, e1000815 (2010).

Acknowledgments: Supported by European Research Council grants 268689 (W.G.) and 295129 (G.D.); the European Community's Seventh Framework Program (grant agreement 243914, Brain-I-Nets, W.G. and H.S.; grant agreement 269921, BrainScaleS, W.G. and G.D.); Swiss National Science Foundation grants 200020 and 132871; and Spanish Ministerio de Economía y Competitividad grants CDS-2007-00012 and SAF2010-16085. We thank T. Vogels for helpful discussions.

10.1126/science.1227356



Theory and Simulation in Neuroscience

Wulfram Gerstner *et al.*

Science **338**, 60 (2012);

DOI: 10.1126/science.1227356

This copy is for your personal, non-commercial use only.

If you wish to distribute this article to others, you can order high-quality copies for your colleagues, clients, or customers by [clicking here](#).

Permission to republish or repurpose articles or portions of articles can be obtained by following the guidelines [here](#).

The following resources related to this article are available online at www.sciencemag.org (this information is current as of January 15, 2015):

Updated information and services, including high-resolution figures, can be found in the online version of this article at:

<http://www.sciencemag.org/content/338/6103/60.full.html>

A list of selected additional articles on the Science Web sites **related to this article** can be found at:

<http://www.sciencemag.org/content/338/6103/60.full.html#related>

This article **cites 89 articles**, 22 of which can be accessed free:

<http://www.sciencemag.org/content/338/6103/60.full.html#ref-list-1>

This article appears in the following **subject collections**:

Neuroscience

<http://www.sciencemag.org/cgi/collection/neuroscience>