

Conceptual illustration, *The Mind-Body problem*.

driven—rather than concept-driven—strategies shape the future industrialization of neuroscience through the rapid emergence of very-large-scale data-mining initiatives. Among recent supranational ventures, the EPFL-IBM consortium “Blue Brain” (1), the European consortium “The Human Brain Project” (HBP) (2), the U.S. consortia BRAIN (3, 4) and “The Human Connectome” (5), and the privately owned Allen Institute (6) all flirt with the possibility of achieving, within the next decades, the full simulation of the human brain (Box 1). Although big-data initiatives have started an impressive thrust in brain research, I question here their impact on how the brain sciences are evolving and highlight the necessity of developing alternative scientific strategies. After briefly reviewing the current advances and hopes that new technologies bring within range of modern brain research, I raise the possibility that, at the same time, scientific conduct is undergoing a radical societal change (section 1). I outline the risks generated by the big-data revolution in brain sciences, discussing various conceptual bottlenecks (sections 2 to 5). I illustrate practical and theoretical limitations that brute-force strategies may encounter in simulating the full brain (sections 6 and 7). I suggest safeguards that should be kept in mind in the new societal context dominated by “economics of promises” (section 8), and conclude with a list of positive recommendations.

1. Big-data initiatives: A worldwide change of scientific strategy in brain studies?

The prevailing consensus in neuroscience is that technology has revolutionized our approach in looking at brain structure and function in relation to behavior (7, 8), and in multiple ways:

1) at the technical level: by extending the power of techniques of circuit identification beyond that already reached by genetic or viral approaches, enabling high-throughput optical manipulation of large-neural ensemble activity with single-cell and single-spike resolution *in vivo* (9–12);

2) at the methodological level: by imposing new standards in experimentation and data acquisition in direct relation with behavior (13, 14);

3) at the data production level: by compiling genomic, structural, and functional databases, the size of which (measured in petabytes) is orders of magnitude larger than that of a complete mammalian genome (15);

4) at the level of analysis: by the application of methods of dimensionality reduction (16, 17) and of pattern-searching algorithms specialized for high-dimensional spaces (18), used previously in statistics, machine learning, and physics;

5) at the modeling level: by the overwhelming development of optimization and Bayesian predictive methods (19, 20) and deep learning approaches (21), made possible by the countless dimension of the data reservoir.

The impact of technical advances on brain research has become such that a major change in

REVIEW

Big data and the industrialization of neuroscience: A safe roadmap for understanding the brain?

Yves Frégnac

New technologies in neuroscience generate reams of data at an exponentially increasing rate, spurring the design of very-large-scale data-mining initiatives. Several supranational ventures are contemplating the possibility of achieving, within the next decade(s), full simulation of the human brain.

I question here the scientific and strategic underpinnings of the runaway enthusiasm for industrial-scale projects at the interface between “wet” (biology) and “hard” (physics, microelectronics and computer science) sciences. Rather than presenting the achievements and hopes fueled by big-data-driven strategies—already covered in depth in special issues of leading journals—I focus on three major issues: (i) Is the industrialization of neuroscience the soundest way to achieve substantial progress in knowledge about the brain? (ii) Do we have a safe “roadmap,” based on a scientific consensus? (iii) Do these large-scale approaches guarantee that we will reach a better understanding of the brain?

This “opinion” paper emphasizes the contrast between the accelerating technological development and the relative lack of progress in conceptual and theoretical understanding in brain sciences. It underlines the risks of creating a scientific bubble driven by economic and political promises at the expense of more incremental approaches in fundamental research, based on a diversity of roadmaps and theory-driven hypotheses. I conclude that we need to identify current bottlenecks with appropriate accuracy and develop new interdisciplinary tools and strategies to tackle the complexity of brain and mind processes.

Introduction

This essay explores how the big-data revolution has started to have an impact on brain sciences and assesses the dangers of letting technology-

Unité de Neurosciences, Information et Complexité (UNIC-CNRS), Gif-sur-Yvette, France. Email: fregnac@unic.cnrs-gif.fr

reference animal models used in neuroscience has occurred in less than 10 years: most state-of-the-art techniques favor the use of few experimental species [e.g., zebrafish, mouse, and marmoset among the vertebrates (22, 23)] and have already consigned to relative oblivion those used traditionally for functional electrophysiology and cognitive mapping (e.g., rat, cat, ferret, and macaque). Simultaneously, outstanding progress in non-invasive imaging techniques (24) such as diffusion tensor imaging (DTI), functional magnetic resonance imaging (fMRI), and ultra high-field MRI, paired with sophisticated neuro-cognitive paradigms (25, 26) and multivariate analysis methods (27, 28), now reaches spatial-scale resolution and temporal precision ranges (25, 27) closer to those used in invasive physiology in nonhuman mammals (29), making cross-species comparison, including humans, feasible in the near future.

Because bold scientific claims increase with technological prowess, the field has also raised its level of self-criticism. Despite major advances in optogenetic control of neural activity patterns (9, 11, 12), “interventionist” neuroscience is still required to show its efficiency in unraveling neural mechanisms causal to behavior (30). Methods must be developed to untangle multiple sources of shared or context-dependent correlations. At a more macroscopic level, localizationist interpretations in brain imaging recently came under scrutiny, both at the paradigmatic and preprocessing level, leading to more controlled definitions of reference or “null” statistics (31, 32). Still unsolved is the obvious difficulty of “putting all together” across scales, when comparing, for instance, neural responses and neurovascular coupling dynamics (33–37). These discrepancies need to be resolved, because they highlight the risks of betting on ill-chosen instrumentation-imposed observables.

The major risks go well beyond technological misuses or misinterpretations. The present trend prefigures a radical societal change in scientific conduct, where new directions in science are launched by new tools rather than by new concepts (38). Many leading scientists and funding agencies now share the view that “progress in science depends on new techniques, new discoveries and new ideas, probably in that order” (39). The pressure has become such that, to receive funding and eventually publish high-impact papers, scientists are often required to use mouse-specific state-of-the-art techniques, irrespective of their adequacy. To some degree, wishful thinking has replaced the conceptual drive behind experiments, as if using the fanciest tools and exploiting the power of numbers could bring about some epiphany.

Although industrialization in scientific methods and practice successfully prevailed in the human genome sequencing project [(40); but see (41, 42)], it is unlikely that a similar brute-force approach will guarantee major advances in understanding brain complexity. Conceptual guidance is required to make the best use of technological advances, regardless of their obvious benefits. “Technology is a useful servant but a dangerous master.” As pointed out by Florian Engert, “the essential ingredient that turns a

useless map” or database “into an invaluable resource” remains “the experimental design employed to gather and analyze the underlying data, and ultimately the thought process, creativity, and ingenuity that went into this design” (43). At a more conceptual level, barrier-breaking innovation paradoxically stems more often from unpredictable “rupture” processes than industrialized approaches. In numerous cases, seminal findings in neuroscience were chance discoveries and daring interpretations. These go well beyond the technological limits of observations and, sometimes, provide the missing but consensual experimental evidence of prior conceptualization formulated centuries earlier. Better tools in hand are just not enough.

2. Bottlenecks in large-scale search studies: Big-data is not knowledge

Provided adequate funding, “big” is easy to acquire and accumulate but hard to classify, interpret, and make sense of. The sea of biological data creates the illusion of knowing “more,” whereas we should rather acknowledge our profound underestimation of how “complex” the brain is. Big data in biology is not limited to acquisition of vast numbers of observables. It further requires

“The search for a unified theory...remains at a rudimentary stage for the brain sciences.”

selection criteria to evaluate their strategic value, and sophisticated handling to extract knowledge. Classically, in information science, one distinguishes four levels in the so-called DIKW pyramids (44), ranging from “data” to “information” to “knowledge” and “wisdom” (understanding). We are currently facing an overflow of data without definite strategies to convert it into knowledge and eventually reach a better comprehension of the living brain.

The most common target in large-scale enterprises flourishing around brain sciences is the generation of biochemical or structural catalogs, most often “static,” taking the form of localizationist atlases in brain-imaging studies or structural inventories at the molecular, cellular, or network level. Of course, static “atlases” imply sophisticated visualization and are sold as tangible deliverables that can be easily understood in layman’s terms. Their use often leads to over-interpretation, when the brain is reduced to a charted globe divided into islands and continents (45–48). Many specialists are aware of the need of rescaling the applicability of instrumental methods and redefining the strict validity range of the conclusions derived from these atlases (49, 50).

Only 20 to 30 years ago, neuroanatomical and neurophysiological information was relatively scarce while understanding mind-related processes seemed within reach. Nowadays, we are drowning in a flood of information. Paradox-

ically all sense of global understanding is in acute danger of getting washed away. Each overcoming of technological barriers opens a Pandora’s box by revealing hidden variables, mechanisms, and nonlinearities, adding new levels of complexity. By reaching the microscopic-scale resolution, advanced technologies have unveiled a new world of diversity and randomness, which was not apparent in pioneer functional studies using spike rate readout or mesoscopic imaging of reduced sensitivity (51–53). This contrast between meso- and microscale functional architectures attests to the necessity of putting more effort in understanding the “regularization” impact of emergence laws—operating in a bottom-up way—across successive levels of integration (see sections 3 and 7). Observations made in parallel with different instruments (sensitive to various spatiotemporal scales) should be combined to build realistic biophysical models to reconcile the loosely related observables across integration levels. In particular, one needs to extract better predictive tools to understand the neural basis of activation processes revealed by brain imaging and find ways of comparing quantitatively state-of-the-art morphological tracing with DTI. Only then could one envision a comprehensive and compressed multiscale functional and structural data repository.

Another approach may be to seek advice from equivalent big-data enterprises in other disciplines such as astrophysics and elementary particle research. Both of these routinely generate petabytes of data. Although particle research does not necessarily conjure up the theoretical viewpoint that we are crucially missing, generations of physicists have been exploring the multiscale complexity of physical matter on the basis of ever-increasing big-data collections (see section 7). Presently, the major difference with brain science is that theorists in particle physics field are involved before—and not after—the hypothesis-driven data are collected. They actively participate in the definition of collective infrastructures and the design of one-of-a-kind equipment shared by the entire experimental collectivity. The recommendation made here is that biologists, who are new to this field, should learn from physicists. As such, the roadmap from data to knowledge could be mapped out in a much clearer fashion and the dead ends, where no one has a clear idea of what to do with all the data, would be far less likely.

To summarize, the trend toward increased measurement sensitivity and more microscopic scales carries its own paradox: A digitized ersatz of lower dimensionality will never account for the multiscale complexity of the full brain. We should adapt our strategic planning so that conceptual efforts grow in a way that is commensurate with technological development—and not follow it, as is presently the case.

3. Bottlenecks in multilevel analysis: The Marr-Poggio conundrum

One of the advertised “blue sky” goals of big-data-driven initiatives is to establish the subcellular and cellular mechanisms causal to behavior

through an exhaustive reductionist analysis. The best-known roadmap for dealing with brain complexity was formulated by David Marr some 35 years ago (54). One way to look at the proposed hierarchy of analysis levels (Fig. 1) is to progress from the global “functional and computational” level, through the intermediate “algorithmic” level down to the “substrate” or “implementation” level. The two higher levels, computational and algorithmic, can be considered as the most generic and abstract, independent of the biological trick used to implement them. Marr argued that whereas “algorithms and mechanisms are empirically more accessible, ...the level of computational theory... is critically important from an information-processing point of view...[because]...the nature of the computations that underlie perception [and, by extension, cognition] depends more upon the computational problems that have to be solved than upon the particular hardware in which their solutions are implemented” (54). Marr was convinced that a purely reductionist strategy, decomposing the global process into its elementary subcomponents, was “genuinely dangerous.” Trying to understand the emergence of cognition from neuronal responses “is like trying to understand a bird’s flight by studying only feathers. It just cannot be done.” Marr’s main intuition was that it is much more difficult to infer from the neural implementation level what algorithm the brain is using (bottom-up) than to reach the algorithmic level from the

**...bottom-up
“emergence”...
remains an open
issue today.**

study of the computational problem that it is trying to solve (top-down along the hierarchy). The bottom-up “emergence” process arising from the interaction of local low-level biological processes remains an open issue today. The way in which sensory neurophysiology has conferred to single-neuron firing the embodiment of high-level psychological properties that can only be sensibly ascribed to a whole behaving organism is a striking example of mereological fallacy (30, 55).

Despite the wealth of produced data, constructionist approaches are thus likely to produce mimicry by a brain ersatz, because of the difficulties of reverse inference (in this case, inferring function and behavior from neural-level activation). This prediction was recently computationally explored, by designing arbitrary experiments on an artificial brain-like artifact, a single microprocessor, to see if popular data analysis methods from neuroscience could elucidate the way in which it processes information and controls behavior (in the present case, three classic videogames) (56). Although the processor’s algorithmic flowchart was known a priori, classical interventionist neuroscience methods failed to explain how the processor works, regardless of the amount of data (30).

The critical point remains that causal-mechanistic explanations are qualitatively different from understanding how a combination of component modules performing the computations at a lower level produces emergent behavior at a higher level.

The first difficulty arises because higher-level concepts are needed to understand the neural implementation level. So, even when causality is demonstrated, it makes sense only when all levels are considered together simultaneously: “Ion channels do not beat, heart cells do. Neural circuits do not feel pain, whole organisms do” (30). Some key studies illustrate the necessity of binding different levels in the experimental design itself—for instance, by linking the neural level with the theoretical context derived from preexisting behavioral knowledge. The supervised learning experiments engineered in single neurons recorded in visual cortex in vivo (57), for example, were conceived as the direct neural implementation (substrate level) of a hypothetical plasticity rule (58) (algorithmic level) derived from associative memory (59) and Ising (60) models (computational level).

A second difficulty comes from Marr’s “multiple realizability” argument, which states that the same function can be achieved through any number of different substrates (30, 54, 61). The impossibility of mapping behavior or function in a unequivocal way on the parametric state of the synaptic or conductance ensemble (defining observed dynamics of the neural net under study) was reproduced in simulation models of *Aplysia* (62, 63) and vertebrate cerebellum (64). This conundrum reveals unexpected complexity whichever way the hierarchy is read, from the computation or macro level to the substrate or micro level, or the reverse.

An additional hidden twist is that the biological substrate level may consist of nested sublevels, each operating at different biophysical scales. Tomaso Poggio emphasized how knowledge of the more elementary steps of information processing is required to account for the complexity of more global computations (65). The key issue is to determine the minimal stratification level needed to preserve the nonlinearities and self-organizing properties at higher integrative levels (66).

Refined electrophysiological studies in the early visual system show clear cases where most spiking-net models—by not giving enough descriptive depth to the biophysical substrate—are too simplified to self-generate low-level feature specificity (orientation selectivity, contrast invariance, and so forth): (i) Rather than the simplified +/- algebra of McCulloch-Pitts neurons, synaptic biophysics in vivo suggests a much richer algebra that includes scaling and division of excitatory inputs by inhibitory ones, where a digital “zero” in the target neuron output could mean either absence of incoming signal (what spiking nets generally assume) or the division or “veto” of an excitatory input by a strong concomitant shunting inhibition (66, 67). (ii) Although orientation selectivity is a hallmark of mammalian cortical organization, this feature selectivity is, in most spiking models, forced in an ad hoc way, by pre-specified wiring rules between thalamus and cortex. Only the orientation preference map appears to be treated as an emergent property resulting from horizontal connection plasticity (68). This oversimplification is challenged when viewed

Box 1. “Big data” projects in brain sciences: Websites

Australia:

The Brain Dialogue: www.cibf.edu.au/australian-brain-alliance

Europe:

The “Blue Brain” Project: bluebrain.epfl.ch/

The “BrainScales” Project: brainscales.kip.uni-heidelberg.de/public/

The Human Brain Project: www.humanbrainproject.eu/

INCF (International Neuroinformatics Coordinating Facility): www.incf.org/

United States:

The Human Connectome Project: www.humanconnectomeproject.org/

www.neuroscienceblueprint.nih.gov/connectome/

BRAIN Initiative Alliance: www.braininitiative.org/

The Allen Institute: observatory.brain-map.org/visualcoding/overview

www.brain-map.org/

Israel:

Brain Technologies: israelbrain.org/

Japan:

Brain Mapping by Integrated Neurotechnologies for Disease Studies (Brain/MINDS):

(Riken BSI): www.bmind.s.brain.riken.jp/

China:

Brain Project: Basic neuroscience, brain diseases and brain-inspired computing in progress (147).

from the conductance level: Voltage-clamp measurements in vivo, even in layer 4, reveal an unexpected level of nonlinear interaction and diversity between excitatory and inhibitory conductances (67, 69–71), which, in V1 simple cells, are hardly detectable (72) or absent at the spiking level (73). The consequence is that the same functional receptive field type, “simple” or “complex,” may indeed be produced by multiple dynamic interaction patterns between excitation and inhibition (71, 74). This unexpected wiring diversity in the synaptic genesis of V1 receptive fields concurs with statistical predictions made by multilayered convolutional models (75). By oversimplifying synaptic integration biophysics and limiting simulations to the spike level, most computational models trivialize the emergence of “higher-order” properties through a purely feedforward cascade (76, 77) when the principal wiring feature of sensory neocortex is—by far—synaptic reverberation and amplification (66).

In view of the weight presently given to spike-based feedforward processing and deep learning, the reexamination of conductance-based versus spike-based computing and the role given to synaptic reentry both sound essential. Bottlenecks in multiscale modeling are rarely addressed in depth, and, although it is agreed that nobody has the definitive solution, this remains a serious blow for “constructionist” models of the brain. Alternative viewpoints should be developed.

4. Bottlenecks in reverse engineering: Lessons learned from the invertebrates

One safe way to handle big-data sets in vertebrates is to avoid the pitfalls known from pioneering studies in paucineuronal networks. Comparative neuroscience offers multiple test studies: (i) small, genetically tractable animal models (78), such as *Caenorhabditis elegans*; (ii) functionally identified clusters of giant cells, in sensory-motor ganglions in *Aplysia* and crustaceans; and (iii) transparent zebrafish, making the online imaging of the whole connectome possible (79). This suggests access to “full brain” descriptions with the reconstruction of causal structuro-functional relations matching canonical neuronal states with species-specific behavioral repertoires (14, 80, 81).

Yet, even with such elementary invariant-like systems, interindividual variability cannot be ignored. A counterintuitive finding in *C. elegans* is that there is no such thing as “simplicity” despite the reduced connectome (302 neurons, 6963 synapses, 890 gap junctions), even at the earliest stage of sensory processing. Averaging neuronal responses of a single olfactory cell is deceptive, because the activation of the same neuron, depending on the context, may lead to several possible behavioral outcomes (82). The main predictive signal of the response is the internal state of the functional assembly in which the cell participates, at the exact time when external inputs become processed. Similar state dependencies in neuronal processing have just started to be explored in vertebrates (83, 84).

Partial understanding of the functional extent and multiscale impact of contextual processing

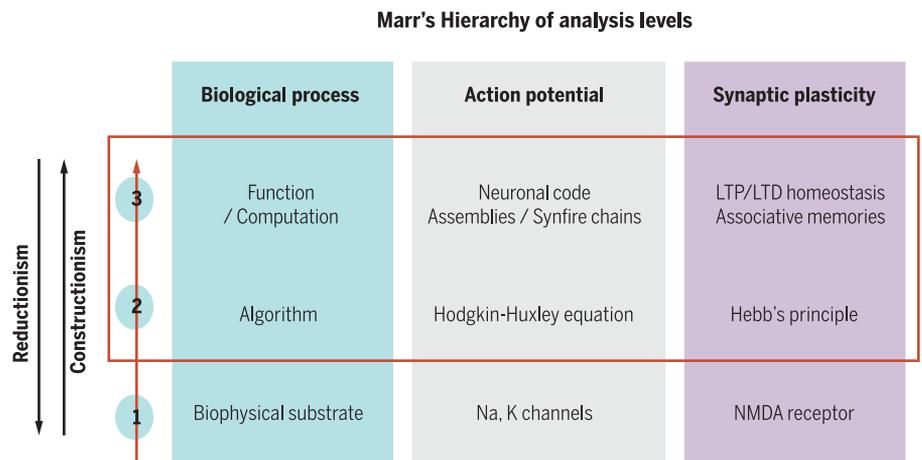


Fig. 1. The hierarchy of analysis levels [inspired by David Marr (54)]. The three levels of Marr's hierarchy illustrated are (from top to bottom) function and computation at the higher level (3), algorithm at the intermediate level (2), and biophysical substrate at the lower level (1). Reductionist approaches progress from levels 3 to 1, whereas constructionism goes the opposite way, from 1 to 3. Two examples of the three-level analysis are given for two different biological processes: action potential (middle column) and synaptic plasticity (right column). The two upper levels of Marr's hierarchy define the field of computational neuroscience (red inset), the scope of which is to identify generic computations and functions and their underlying algorithms, independently of the biophysical substrate of the process under study.

has been obtained in classical studies in the lobster's stomato-gastric ganglion (85). By releasing diffusible neuromodulators, specialized “orchestra conductor” neurons change the conductance repertoire of the other individual neurons and allow them to participate at distinct times in a diversity of functional subnetworks (“assembly reconfigurability”). This feature highlights the impossibility of separating intrinsic (conductance repertoire, genomic expression) from extrinsic (synapses) features. The diffusive nature of the modulatory process and its dependency on the internal mesoscopic state generated by the recurrent synaptic activity open a yet largely unexplored scale of complexity.

A straightforward lesson from invertebrates is that a purely “Lego”-like reconstruction approach—based on the full reconstruction of the brain's connectome and the gene expression, electrical, and morphological determinants profile of the major classes of its neural components (86, 87)—may be doomed from the start. Despite similar evidence in vertebrates, some doubt remains as to whether the versatility of the excitability pattern and the dependency of conductance repertoire expression on past brain states (and modulators) are taken at face value in classifications and nomenclatures of supposedly invariant identity determinants (88). Thus, the dynamic complexity revealed in simpler organisms provides a powerful warning against the use of purely bottom-up constructivist large-scale studies in higher organisms.

5. Bottlenecks in evolutionary leaps: Anthropocentrism from “mouse” to “man”

“Understanding the brain” is often read as understanding the “human” brain. This anthropo-

morphic bias reveals a loss of perspective regarding the essence of living systems: their diversity, their adaptability, and their dependence on evolutionary history. Losing track of this perspective is dangerous, because only broad comparisons offer the potential to distinguish general principles from unimportant implementation details. If paving the way toward “a general theory of the brain” is a worthy goal, as we believe it is, then it is essential to conceive comparative physiology strategies, which allow us to discriminate between species-specific “bags of tricks” and canonical computations shared by living brains (30, 66, 89–92). Certain forms of computation and algorithms seem to be preserved (i.e., gain control, normalization, exponentiation, association, and coincidence detection), but the detailed mechanistic implementations are often species-specific and structure-dependent (30). Industrial-scale efforts are, by their present design, focused on limited behaviors and species, and thus orthogonal to a broad-enough perspective.

A second problem is that the human brain is probably among the most complex of nervous systems. This has led, without much strategic planning other than exploiting the availability of a genetically modifiable mammalian system, to the increasing use of the mouse as a model. Because it is a mammal, it must be similar to the human. Although the mouse model has produced important advances in the study of basic sensory-motor integration principles, it may be less appropriate for studying perceptual processes for modalities (vision) less adapted to its behavioral repertoire and, more obviously still, for higher cognitive functions. This is particularly true in species such as humans and other primates where sensory cortical processing involves

elaborate reciprocal connectivity patterns linking sets of functionally distinct areas (93, 94), which are mostly absent in the mouse cortex.

A wiser alternative could be to refine approaches progressively and recursively according to species-specific behavior, and cognition repertory (95). Search for homologies should be validated on the basis of structural, functional, and cognitive similarities between species. The choice of the right species calls for increased efforts in comparative physiology, which have been downplayed since the start of the mouse dominance era. The choice of the right tasks requires new methods of behavior classification. By applying unsupervised learning methods to the largest possible set of coregistered neural data and behavioral observations, one may hope to achieve substantial dimensionality reduction and obtain an objective mapping of possible behavioral repertoires over a restricted ensemble of reproducible brain states, as has been done successfully in invertebrates (81).

6. Simulating the brain: The cart before the horse—immaturity of paradigms and lack of hypothesis-driven design

A fundamental issue for large database generalization and validation is to provide universal paradigm or task standards that are optimized for the study of specific cognitive functions. For illustration's sake, let us concentrate on an apparently "simple" case study, i.e., how to characterize neural processes involved in low-level visual perception.

In the search for generic sensory integration principles, how can we conceive a "good" stimulus set before we know what the system under study is designed to perceive (96)? The process cannot be formulated without priors, often linked with behavioral observations and hypothesis testing, and should probably be automated only after a progressive, informed, recursive, maybe even "old-fashioned," phase of investigation. Presenting the largest spectrum of input statistics seems the appropriate way to push the sensory system to its information capacity limits (97) and explore the dependency of the neural code on external input statistics (70, 74, 98, 99). However, in practice, the battery of stimuli used to build large data sets faces unacknowledged technical constraints: Stimulus choices are often guided by the efficiency with which strong firing can be evoked—leading to a prevalence of high firing rates, more easily detectable by calcium fluorescence changes—rather than by information theory concepts (rate code/dense spiking versus spike-timing code/sparseness). The cognitive repertoire should also be used more carefully to constrain the choice of species: There is something odd in applying in the mouse, a nearly blind animal (100), a battery of stimulation paradigms based on decades of work on highly visual species (cat, macaque, and human) without paying attention to ethological differences in the reliance on vision [but see (101)]. Indeed, visual cortex may play different roles in different species; for instance, space coding during navigation—in concert with hippocampus—in rodents, versus primate perceptual sketch elaboration

and form or motion extraction—in concert with higher cortical areas—in more visual species. Consequently, testing the responses of mouse primary visual cortex (V1) to a high-contrast classic Hollywood black-and-white movie (102) seems as inappropriate as studying pangolin olfaction with plumes of warm Parisian croissants. Conversely, searching for place or grid cells may be deceiving in nonhuman primate visual cortex when it makes sense in the rodent.

Choosing the right stimulus and species is not the only issue. Since the shift over the past 20 years from the anesthetized-paralyzed preparation to the behaving animal, the standardization of the global context has become a major concern (103). Visual responsiveness in the awake mouse depends heavily on locomotion and full-body action (83), rendering inseparable the sensory and motor components. However, a similar conditional dependency of visual processing has not been confirmed in higher mammals, where primary sensory and motor cortices are much less—or even not at all in the adult—directly interconnected. Consequently, the generalized use of "running-on-a-ball" paradigms in the rodent may have set a new behavioral standard for studying sensory

"Industrial-scale efforts are...orthogonal to a broad-enough perspective."

responses, optimized to increase neural excitability in the rodent only, but reducing the global relevance to vision per se (66).

The overall consequence is that, by imposing such artificial paradigms as the "standard tests" for brain observatories, each resulting data set will yield predictions restricted to specific contexts, but largely unrelated to "natural" behavior. Big-data initiatives in early vision have not yet put enough effort into defining parameters critical of the "naturalness" of the evoked sensory drive. As summarized by Bruno Olshausen, "the problem is not just that we lack the proper data, but that we do not even have the right conceptual framework for thinking about what is happening" (104). Similarly, however impressive they may be, all-optical "interventionist" paradigms do not signal the end of the quest: New conceptual frameworks are needed that "provide the mapping between large-scale neural data and behavior in an algorithmic sense and not just a correlative or even causal way" (30). The practical message here is that both paradigms and context—in which data are acquired—should be rationalized and justified on purely theoretical grounds, before becoming the norm of the industrialization stage.

7. Simulating the brain—The cart without a driver: Missing a strong brain theory

Do we have a clear view of what can be expected from reverse engineering and embodied construc-

tionism? Some of the large-scale initiatives recapitulate earlier constructionist approaches that tried to simulate brain circuits by building models "that are very closely linked to the detailed anatomical and physiological structure" of the brain, in hopes of "generating unanticipated functional insights based on emergent properties of neuronal structure." The first attempts in the 1990s (105–107) were limited by the lack of prediction of rich enough behavioral repertoires and cognitive functions (108). Conversely, more engineering-oriented and simplified blackbox simulations (109) were criticized for their lack of descriptive depth (110). Even so, some success has been obtained by clever built-in top-down constraints. High-performance computing may change the odds (111), and experts agree that large-scale simulation should provide possible breakthroughs in system identification as has been the case for deep learning (112). Nevertheless, given the analytic intractability of the brain, the challenge of "putting all together" remains wide open. The major obstacle remains the lack of unifying theory and the relative paucity of top-down guidance by high-level knowledge derived from psychological studies of the mind.

In this section, I will review three correlated issues: (i) Are there theoretical conjectures indicating that a full spike-based brain simulation is not a realistic target? (ii) How do system and computational neurosciences integrate theory so far? and (iii) Are there alternative roadmaps to re-address what may be considered as an ill-posed problem?

Point 1: Because of their dominant bottom-up drive, the danger of the large-scale neuroscience initiatives is to produce purely descriptive ersatz of brain, sharing some of the internal statistics of the biological counterpart, at best the two first-order moments (mean and variance), but devoid of self-generated cognitive abilities. The numbers will certainly look right, but there is no guarantee that such simulated brains will work. This intuition resonates with theoretical conjectures based on pure logic. As early as the 1980s, a Gedanken experiment was proposed by von der Malsburg which considered two brain-like assemblies, built with the exact same connectivity graph and producing the exact same averaged firing patterns. What would happen if a jitter of a few milliseconds was applied in the arrival time of each occurring spike (while keeping mean rate invariant)? Is there a critical jitter value that should not be exceeded, to keep alive the emergent properties of the graph (113, 114)? The same conjecture could be generalized at the second-order statistics level. Let us imagine that big data makes it possible to build a cortex-like digital machine where the variance of the distributions of synaptic weights afferent to (or efferent from) each neuron could be matched to those directly measured (over time) in the same ensembles of real synapses. Would one predict the mean and variance-equalized artificial network to be as operative as the real brain? Because—in real brains—the efficiency of individual synaptic weights and their spatial distribution are stabilized through associative plasticity and normalization

processes (if our popular learning theories are right), plugging in simulated synapses mean and variance levels devoid of information content would result in an “averaged connectome” without memory of its past interactions with the outside world. Thus, brain simulations elaborated from static and averaged atlases might be likely useless in simulating brain function. Realistic solutions require that the dynamic entity of the simulated brain “grows” and interacts with the same outside world as the real brain, i.e., that both share the same interactive constraints at any point in time to produce the same behavior or implement the same cognitive process.

Point 2: How do system and computational neurosciences integrate theory so far? In a provocative review (103), Carandini assumes the existence of an intermediate level of circuit integration, where canonical operations can be defined as invariant computations repeated and combined in different ways across the brain. To identify them, it becomes necessary to record from a myriad of neurons in multiple brain regions rather than from single neurons. “Understanding computation...provides a language for theories of behavior.” This concept is very close to the algorithmic level of Marr, because it no longer depends on the understanding of the biophysics of the substrate, which may vary from region to region and species to species. However, most consensual canonical principles are not derived from the search of big data but from philosophical or psychological principles arising from past centuries (115). For instance, the current theories of associative synaptic plasticity did not originate with spike-timing-dependent plasticity (STDP) but can be seen as the revival of causality-based rules inherited from psychologists [(116–118), to cite only a few (119)]. Other rules address a more macroscopic level, irrespective of the biological substrate implementation of the underlying mechanisms, such as the psychic laws of the Gestalt school in 1930s (117, 121) or the binding-by-synchrony hypothesis (120). It is only recently that the introduction of top-down constraints satisfying Bayesian optimization (19, 20) seems to provide innovative insights into mesoscopic processing in the brain and the way it adapts to multiple task-driven constraints.

Point 3: Exploiting biological data obtained at different spatial and temporal scales should benefit from earlier concepts developed in statistical physics. Anderson (122) points out that the field of supraconductivity shows the reductionist fallacy (see section 3: Marr-Poggio conundrum). The ability to reduce everything to simple laws does not imply the ability to start from those laws and reconstruct the whole (the brain in biology, the universe in physics). The constructionist hypothesis breaks down when confronted with scale changes and complexity (123). Anderson summarizes the principle of “symmetry breaking” across scales, as follows: (i) The internal structure of a piece of matter or a living brain need not be symmetrical even if the total state of it is (an argument that mean field theories do not always follow); (ii) the macroscopic state of a large system



Fig. 2. Building brain sciences through “economics of promises”? Promises based on data-driven exploration and modeling of the human brain share similarities and even inspiration with the imagery of science fiction. They become the scientific justification for the capture of large-scale funding.

has less symmetry than that obeyed by the microscopic laws which govern it. “In the so-called $N \rightarrow \infty$ limit...matter will undergo mathematically sharp, singular ‘phase transitions’ to states in which the microscopic symmetries...are in a sense violated...Functional structure in a theological sense, as opposed to mere crystalline shape, must also be considered a stage, possibly intermediate between crystallinity and information strings, in the hierarchy of broken symmetries.” A rare echo of this principle can be found in a pioneer multiscale model of emergence of local and global features in the early visual system (75, 124, 125).

Progress should be expected by building novel descriptive frameworks which extract—from zillions of measurements—mesoscopic variables, analogous to the concept of quasiparticles in statistical physics. Solid-state physicists successfully developed “middle way” theories (126) that overcome the limitation that equations for particle interactions become impossible to solve or simulate for more than 10 particles. The introduction of a formalism based on virtual quasiparticles may simplify the analytical treatment of long-distance interactions between numerous elementary bound particles, by an equivalent free quasiparticle with shorter interaction. The search for such macroscopic variables could offer an analytic way of treating neural network dynamics and enrich the present mean-field equation formalism. This would allow the building of new kinds of “stereological” models of gray matter, combining the local-range connectivity of columnar ensembles, the extrasynaptic volume diffusion of second messengers and modulators, and the oscillatory coupling due to physical distance in the three-dimensional (3D) brain [a factor unaccounted for by classical ring (1D) or layered (2D) networks]. Quasiparticles have dual corpuscular and wave

counterparts, which may apply to information diffusion and propagation across cortical networks, for which evidence can be monitored by fast voltage-sensitive dye imaging. Use of such models may reconcile the physics of interacting particles and waves with the functional physiology of long-distance interconnected cortical columns.

The search for a unified theory, as in particle physics, remains at a rudimentary stage for the brain sciences. When changing scales, symmetry breaks introduce major nonlinearities that we cannot account for at present. Thus, the validity of theories and the choice of the relevant explanatory variables remain restricted to certain levels of integration, resulting in simulation attempts that are essentially local and species- and task-dependent. The hope is that understanding mesoscale organization and full network dynamics might reveal a simpler formalism than the microscale level, similar to general laws in statistical thermodynamics (127). The limitation for reverse engineering is that mean-field-like approaches, because of their underlying simplifications, will lose important generative mechanisms of low-level nonlinearities. A more empirical and modest alternative could be to multiply the diversity of proposed multiscale models, selecting those that most efficiently reduce complexity: “A good theoretical model of a complex system should be like a good caricature: It should emphasize those features which are most important and should downplay the inessential details... Since one does not really know which are the inessential details until one has understood the phenomena under study...one should investigate a wide range of models and not stake one’s life (or one’s theoretical insight) on one particular model only” (128). Hence, again, the definition of multiscale data integration and the convergence

to a theoretical understanding must be progressive and recursive.

8. The risks, for basic research, of dominant strategies based on “economics of promises”

Let us leave theory and move to the economics and policy of science. International think-tank meetings for defining a worldwide unified strategy (129, 130) attract public attention and feed the buzz of wide-audience science chronicles. Large-scale brain initiatives are often presented to the public as unselfish but costly science, generating state-of-the-art infrastructures and large data resources open to the community. They are advertised as opening the door for brain-derived information technology (IT) and, in the minds of some high profile IT leaders, paving the way to transhumanism (131, 132).

Part of the original motivation for big data comes from its success in studying simple organisms: for instance, the complete lineage and full reconstruction using electron microscopy of *C. elegans*, initiated in the 1980s, were shared by the entire field, leading to faster progress. However, the justification for the full human brain simulation is more questionable: The metaphor of “mind observatory,” used rhetorically to link it with physics exploratory platforms such as CERN, is misleading. Megascience infrastructures in physics take immediate advantage of shared “unique” instruments, which have been cooperatively designed to collect new experimental data and test explicit hypotheses through an overarching theory. In the brain sciences, however, building massive database architecture without theoretical guidance may turn into a waste of time and money (133, 134).

The “observatory” function itself, i.e., yielding new data that were formerly out of reach because of technical limitations, is not even central to some of the large-scale brain initiatives. For instance, the flagship project (HBP) transformed its original drive (for a better understanding of brain) into a “viewing neuroscope” IT platform built largely on preexisting data. Progress is expected mostly from an alliance of deep learning, neuroinformatics, and neuromorphic computation, and promised to be quantitative enough to sustain virtual medicine applications (135).

This strategic drift illustrates the impact of “megascience,” considered by sociologists of emergent technologies as a new form of societo-scientific culture (131, 132, 136–139). “Economics of promise” are built around a scientific or industrial process (or even a theoretical law) whose justification is primarily based not on scientific or technological arguments but on the promises themselves (as if these were guaranteed to be fulfilled). This trend, which has deep roots linked to what modern society expects from biology in the large sense, has been repeatedly observed in different scientific subfields such as large-scale brain simulation, nanotechnology, stem cells, and synthesis biology (138). It even applies to the myth of Moore’s law that perpetuates itself because of the marketing of chip designers in neuromorphic computing (132, 140).

Plausible reasons have been identified to justify such drastic changes in scientific conduct: rarefaction of funding for basic research in brain science, the necessary requirement of a major translational impact at the societal level, “hype” purposely designed to reach the largest public audience as well as political decision-makers, overselling promises in the public health domain and possible blue-sky industrial outcomes. The attractiveness to politicians, administrators, and funders (whether public or private) of massive and visible one-track programs is obvious (141), but one may consider that high-level “deciders” are not always entirely aware of—or possibly interested in—the downsides of these mammoth programs, or of the obvious weaknesses of their scientific underpinnings. Promises are no longer an extrapolation of the “possible future” (Fig. 2),

“The present trend prefigures a radical societal change in scientific conduct...”

but become the scientific justifications of purely economic and political “bubble” strategies engineered to capture funding on the basis of competitive supranational calls (139, 142).

A side effect is that governmental institutions in Europe and the United States suggest that enough data may be already available on the laboratory shelves, constituting a pile of “siloeed” dormant sources that need to be curated (143, 144). Will this become a cheap pretense used to justify budget reduction in experimental basic neuroscience? It seems indeed easier in terms of budget control to turn scientists into high-tech engineers rather than to fund basic research on a wider spectrum with reduced short-term impact.

There exists a real danger that a few large-scale international projects building the foundations of virtual or in silico neuroscience will massively engage the funds available in basic neurosciences to the detriment of small and medium-size basic research initiatives focusing on integrative, cognitive, or computational neuroscience. One gets the impression that the future of acquisition and exploitation of brain-related data will be shared between a few large-scale continental initiatives or strong industrial-like ventures. The possibility of conflicts of interest (which grows with the size of the consortia), of attempts to self-appropriate knowledge and eventually make a profitable business of it (145, 146), all remind us that it is urgent to define worldwide accepted standards of transparent macro-management and access to data and technologies.

Conclusion

In this Review, I have tried to point out that, although big-data and technological advances undeniably have immense value for future de-

velopments, the expedient industrialization of neuroscience and the potential long-term importance of the personal, political, and commercial incentives driving it are causes for concern. Systematic and streamlined approaches are not appropriate for all facets of brain research, and the interpretation of massive data sets collected without appropriate forethought may turn out to be impossible. Given the exponentially increasing rate at which big data are being collected, exabyte information will be accumulated before the end of the next decade. Out of this magma, it may be difficult to tease out of the hypothetical key principles that might help resolve the main questions that should have been at the root of their design and made explicit all along.

Megascience dominance, if improperly managed, may lead to the drying up of traditional funding channels and the disappearance of smaller-scale and rationally designed research programs, which are still the major source of breaking discoveries. To master megascience development and reduce negative side effects, current strategies could be greatly improved by the following:

1) rationalizing the codesign of the choice of experimental models (choice of species, precise targeting of behavioral specificity) and the justification of appropriate techniques (sensitivity range of the instrumentation, spatial and temporal scale ranges to be explored);

2) clarifying the hidden scientific assumptions associated with each instrumentation type and interrelating explanatory variables (i.e., conductance, spike rate, calcium fluorescence, metabolic or hemodynamic signals) despite their biophysical diversity;

3) clarifying the hidden impact of preprocessing steps and statistical methods to reduce across-study heterogeneity;

4) developing more efficient recursive loops between experiments and theory-driven top-down predictions, to confront a larger diversity of brain models and compare their predictive power;

5) building innovative theoretical frameworks not only inspired by computational neuroscience, mathematics, and psychology, but also enriched by complementary fields used to deal with complex systems of high dimensionality (statistical physics, thermodynamics, astrophysics);

6) vetting the most relevant experimental paradigms, to define in an unbiased way the parametric features and the reproducibility of the stimulation context necessary to the constitution of large-data set repositories;

7) allowing open access—to scientists and modelers—to the entire data reservoir and its data sharing, devoid of selective control by the ownership claims of grant funders.

These changes in scientific planning will undoubtedly require the generalized practice of interdisciplinarity between physics and biology, focusing on the major bottlenecks (129, 130). Only in this way, can we hope to improve our critical skills and collectively optimize our capacity to better anticipate the challenges we face in exploring uncharted levels of complexity.

REFERENCES AND NOTES

- H. Markram, *Nat. Rev. Neurosci.* **7**, 153–160 (2006).
- H. Markram, *Sci. Am.* **306**, 50–55 (2012).
- A. P. Alivisatos et al., *Neuron* **74**, 970–974 (2012).
- L. A. Jorgenson et al., *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **370**, 20140164 (2015).
- T. B. Leergaard, C. C. Hilgetag, O. Sporns, *Front. Neuroinform.* **6**, 14 (2012).
- C. Koch, R. C. Reid, *Nature* **483**, 397–398 (2012).
- E. R. Kandel, H. Markram, P. M. Matthews, R. Yuste, C. Koch, *Nat. Rev. Neurosci.* **14**, 659–664 (2013).
- T. J. Sejnowski, P. S. Churchland, J. A. Movshon, *Nat. Neurosci.* **17**, 1440–1441 (2014).
- E. S. Boyden, F. Zhang, E. Bamberg, G. Nagel, K. Deisseroth, *Nat. Neurosci.* **8**, 1263–1268 (2005).
- K. Deisseroth, M. J. Schnitzer, *Neuron* **80**, 568–577 (2013).
- A. M. Packer, L. E. Russell, H. W. Dalglish, M. Häusser, *Nat. Methods* **12**, 140–146 (2015).
- V. Emiliani, A. E. Cohen, K. Deisseroth, M. Häusser, *J. Neurosci.* **35**, 13917–13926 (2015).
- M. Leinweber et al., *J. Vis. Exp.* **50885**, e50885 (2014).
- R. Portugues, C. E. Feierstein, F. Engert, M. B. Orger, *Neuron* **81**, 1328–1343 (2014).
- E. Landhuis, *Nature* **541**, 559–561 (2017).
- M. Rabinovich, R. Huerta, G. Laurent, *Science* **321**, 48–50 (2008).
- J. P. Cunningham, B. M. Yu, *Nat. Neurosci.* **17**, 1500–1509 (2014).
- M. Helmstaedter, *Neuron* **86**, 25–28 (2015).
- D. C. Knill, A. Pouget, *Trends Neurosci.* **27**, 712–719 (2004).
- K. Friston, *Nat. Rev. Neurosci.* **11**, 127–138 (2010).
- Y. LeCun, Y. Bengio, G. Hinton, *Nature* **521**, 436–444 (2015).
- D. Cyranoski, *Nature* **514**, 151–152 (2014).
- J. F. Mitchell, D. A. Leopold, *Neurosci. Res.* **93**, 20–46 (2015).
- D. Le Bihan, *Looking Inside the Brain: The Power of Neuroimaging* (Princeton Univ. Press, NJ, 2014).
- J. D. Haynes, G. Rees, *Nat. Rev. Neurosci.* **7**, 523–534 (2006).
- S. Dehaene, J. P. Changeux, *Neuron* **70**, 200–227 (2011).
- J. R. King, S. Dehaene, *Trends Cogn. Sci.* **18**, 203–210 (2014).
- J. D. Haynes, *Neuron* **87**, 257–270 (2015).
- J. Goense, Y. Bohraus, N. K. Logothetis, *Front. Comput. Neurosci.* **10**, 66 (2016).
- J. W. Krakauer, A. A. Ghazanfar, A. Gomez-Marín, M. A. Machler, D. Poeppel, *Neuron* **93**, 480–490 (2017).
- D. M. Cole, S. M. Smith, C. F. Beckmann, *Front. Syst. Neurosci.* **4**, 8 (2010).
- A. Eklund, T. E. Nichols, H. Knutsson, *Proc. Natl. Acad. Sci. U.S.A.* **113**, 7900–7905 (2016).
- N. K. Logothetis, J. Pauls, M. Augath, T. Trinath, A. Oeltermann, *Nature* **412**, 150–157 (2001).
- N. K. Logothetis, B. A. Wandell, *Annu. Rev. Physiol.* **66**, 735–769 (2004).
- Y. B. Sirotin, A. Das, *Nature* **457**, 475–479 (2009).
- M. L. Schölvinck, A. Maier, F. Q. Ye, J. H. D. D. D. A. Leopold, *Proc. Natl. Acad. Sci. U.S.A.* **107**, 10238–10243 (2010).
- B. X. Huo, J. B. Smith, P. J. Drew, *J. Neurosci.* **34**, 10975–10981 (2014).
- F. Dyson, *Imagined worlds*. The Jerusalem-Harvard Lectures (Harvard Univ. Press, Cambridge, 1997).
- C. Lange, in *Nobel Lectures. Peace, 1901–1925*, F. Haberman, Ed. (Elsevier, Amsterdam, 1972).
- F. S. Collins, M. Morgan, A. Patrinos, *Science* **300**, 286–290 (2003).
- L. Roberts, *Science* **248**, 804 (1990).
- K. R. Chi, *Nature* **538**, 275–277 (2016).
- F. Engert, *Neuron* **83**, 1246–1248 (2014).
- C. Zins, *J. Am. Soc. Inf. Sci. Technol.* **58**, 479–493 (2007).
- M. F. Glasser et al., *Nature* **536**, 171–178 (2016).
- A. G. Huth, W. A. de Heer, T. L. Griffiths, F. E. Theunissen, J. L. Gallant, *Nature* **532**, 453–458 (2016).
- M. Behaim, *The Erdapfel*, Der Behaim-Globus in Nürnberg; <http://nuernberg.bayern-online.de/die-stadt/wissenswertes/der-behaim-globus/>.
- M. O. Gewaltig, A simulated mouse brain in a virtual mouse body; www.youtube.com/watch?v=ldXEUUVkDwU.
- J. Gonzalez-Castillo et al., *Proc. Natl. Acad. Sci. U.S.A.* **109**, 5487–5492 (2012).
- J. Stelzer, G. Lohmann, K. Mueller, T. Buschmann, R. Turner, *Front. Hum. Neurosci.* **8**, 462 (2014).
- K. Ohki, R. C. Reid, *Curr. Opin. Neurobiol.* **17**, 401–407 (2007).
- S. L. Smith, M. Häusser, *Nat. Neurosci.* **13**, 1144–1149 (2010).
- P. O. Kanold, I. Nelken, D. B. Polley, *Trends Neurosci.* **37**, 502–510 (2014).
- D. Marr, *Vision* (MIT Press, Cambridge, 1982).
- J. Bickle, *Top. Cogn. Sci.* **7**, 299–311 (2015).
- E. Jonas, K. P. Kording, *PLOS Comput. Biol.* **13**, e1005268 (2017).
- Y. Frégnac, D. Shulz, S. Thorpe, E. Bienenstock, *Nature* **333**, 367–370 (1988).
- E. L. Bienenstock, L. N. Cooper, P. W. Munro, *J. Neurosci.* **2**, 32–48 (1982).
- T. J. Sejnowski, *J. Math. Biol.* **4**, 303–321 (1977).
- M. M. Nass, L. N. Cooper, *Biol. Cybern.* **19**, 1–18 (1975).
- R. P. Cooper, D. Peebles, *Top. Cogn. Sci.* **7**, 243–258 (2015).
- A. A. Prinz, D. Bucher, E. Marder, *Nat. Neurosci.* **7**, 1345–1352 (2004).
- E. Marder, J. M. Goaillard, *Nat. Rev. Neurosci.* **7**, 563–574 (2006).
- E. De Schutter, O. Ekeberg, J. H. Kotaleski, P. Achard, A. Lansner, *Trends Neurosci.* **28**, 562–569 (2005).
- T. Poggio, *Visual Algorithms* (MIT, Cambridge, 1982).
- Y. Frégnac, B. Bathellier, *Neuron* **88**, 110–126 (2015).
- L. J. Borg-Graham, C. Monier, Y. Frégnac, *Nature* **393**, 369–373 (1998).
- J. A. Bednar, C. K. I. Williams, in *From Neuron to Cognition via Computational Neuroscience*, M. A. Arbib, J. J. Bonaio, Eds. (MIT Press, Cambridge, 2016), pp. 409–432.
- C. Monier, J. Fournier, Y. Frégnac, *J. Neurosci. Methods* **169**, 323–365 (2008).
- P. Baudot et al., *Front. Neural Circuits* **7**, 206 (2013).
- J. Fournier et al., *J. Neurosci.* **34**, 5515–5528 (2014).
- N. C. Rust, O. Schwartz, J. A. Movshon, E. P. Simoncelli, *Neuron* **46**, 945–956 (2005).
- J. Touryan, G. Felsen, Y. Dan, *Neuron* **45**, 781–791 (2005).
- J. Fournier, C. Monier, M. Pananceau, Y. Frégnac, *Nat. Neurosci.* **14**, 1053–1060 (2011).
- R. Linsker, *Proc. Natl. Acad. Sci. U.S.A.* **83**, 8390–8394 (1986).
- T. Serre, A. Oliva, T. Poggio, *Proc. Natl. Acad. Sci. U.S.A.* **104**, 6424–6429 (2007).
- C. F. Cadieu et al., *PLOS Comput. Biol.* **10**, e1003963 (2014).
- C. I. Bargmann, E. Marder, *Nat. Methods* **10**, 483–490 (2013).
- M. B. Ahrens, F. Engert, *Curr. Opin. Neurobiol.* **32**, 78–86 (2015).
- J. Larsch, D. Ventimiglia, C. I. Bargmann, D. R. Albrecht, *Proc. Natl. Acad. Sci. U.S.A.* **110**, E4266–E4273 (2013).
- J. T. Vogelstein et al., *Science* **344**, 386–392 (2014).
- C. I. Bargmann, *BioEssays* **34**, 458–465 (2012).
- C. M. Niell, M. P. Stryker, *Neuron* **65**, 472–479 (2010).
- E. Zagha, D. A. McCormick, *Curr. Opin. Neurobiol.* **29**, 178–186 (2014).
- P. Meyrand, J. Simmers, M. Moulins, *Nature* **351**, 60–63 (1991).
- G. A. Ascoli et al., *Nat. Rev. Neurosci.* **9**, 557–568 (2008).
- J. DeFelipe et al., *Nat. Rev. Neurosci.* **14**, 202–216 (2013).
- H. Markram et al., *Cell* **163**, 456–492 (2015).
- D. J. Anderson, P. Perona, *Neuron* **84**, 18–31 (2014).
- J. Fournier, C. M. Müller, G. Laurent, *Curr. Opin. Neurobiol.* **31**, 119–126 (2015).
- M. Hemberger, L. Pammer, G. Laurent, *Curr. Opin. Neurobiol.* **41**, 24–30 (2016).
- M. M. Yartsev, *Science* **358**, 466 (2017).
- R. J. Douglas, K. A. Martin, *Curr. Biol.* **22**, R1033–R1038 (2012).
- M. E. Laramée, D. Boire, *Front. Neural Circuits* **8**, 149 (2015).
- A. Gomez-Marín, J. J. Paton, A. R. Kampff, R. M. Costa, Z. F. Mainen, *Nat. Neurosci.* **17**, 1455–1462 (2014).
- H. B. Barlow, *J. Physiol.* **119**, 69–88 (1953).
- P. Dayan, L. Abbott, *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems* (MIT Press, Cambridge, 2002).
- W. E. Vinje, J. L. Gallant, *Science* **287**, 1273–1276 (2000).
- B. R. Cowley, M. A. Smith, A. Kohn, B. M. Yu, *PLOS Comput. Biol.* **12**, e1005185 (2016).
- M. Baker, *Nature* **502**, 156–158 (2013).
- D. J. Wallace et al., *Nature* **498**, 65–69 (2013).
- M. Hawrylycz et al., *Proc. Natl. Acad. Sci. U.S.A.* **113**, 7337–7344 (2016).
- M. Carandini, *Nat. Neurosci.* **15**, 507–509 (2012).
- B. Olshausen, in *20 Years of Computational Neuroscience*, J. M. Bower, D. Beeman, Eds. (Springer, New York, 2013).
- J. M. Bower, D. Beeman, *The Book of GENESIS: Exploring Realistic Neural Models with the General NEural Simulation System* (Telos, New York, 1998).
- G. M. Edelman, *Science* **318**, 1103–1105 (2007).
- J. Macera, P. H. Goodman, F. C. Harris Jr., R. Drewes, J. B. Maciokas, *Robot. Auton. Syst.* **46**, 97–110 (2004).
- C. Eliasmith, O. Trujillo, *Curr. Opin. Neurobiol.* **25**, 1–6 (2014).
- C. Eliasmith et al., *Science* **338**, 1202–1205 (2012).
- L. Sanders, Simulated brain mimics human quirks: Model doesn't just calculate but turns decisions into actions. *Science News* (29 November 2012); www.sciencenews.org/article/simulated-brain-mimics-human-quirks.
- E. M. Izhikevich, G. M. Edelman, *Proc. Natl. Acad. Sci. U.S.A.* **105**, 3593–3598 (2008).
- K. R. Chi, *Nature* **531**, S16–S17 (2016).
- C. von der Malsburg, in *Brain Theory*, G. Palm, A. Aertsens, Eds. (Springer, Berlin, 1986), pp. 161–176.
- E. D. Lumer, G. M. Edelman, G. Tononi, *Cereb. Cortex* **7**, 228–236 (1997).
- Y. Fregnac, in "Voices": BRAIN Initiative and Human Brain Project: Hopes and reservations. *Cell* **155**, 265–266 (2013).
- W. James, *Psychology: Briefer Course* (Harvard Univ. Press, Cambridge, 1890).
- Y. Delage, *Le Réve: Etude Psychologique, Philosophique et Littéraire [The Dream: A Psychological, Philosophical and Literary Study (in French)]* (Presses Universitaires de France, Paris, 1919).
- D. Hebb, *The Organization of Behavior* (Wiley, New York, 1949).
- Y. Frégnac, *Nature Encyclopedia in Cognitive Science* **2**, 320–329 (2002).
- W. Singer, *Annu. Rev. Physiol.* **55**, 349–374 (1993).
- P. W. Milner, *Psychol. Rev.* **64**, 242–252 (1957).
- P. W. Anderson, *Science* **177**, 393–396 (1972).
- D. Noble, *Interface Focus* **2**, 55–64 (2012).
- R. Linsker, *Proc. Natl. Acad. Sci. U.S.A.* **83**, 8779–8783 (1986).
- R. Linsker, *Proc. Natl. Acad. Sci. U.S.A.* **83**, 7508–7512 (1986).
- R. B. Laughlin, D. Pines, J. Schmalian, B. P. Stojkovic, P. Wolynes, *Proc. Natl. Acad. Sci. U.S.A.* **97**, 32–37 (2000).
- G. Deco, V. K. Jirsa, P. A. Robinson, M. Breakspear, K. Friston, *PLOS Comput. Biol.* **4**, e1000092 (2008).
- V. Y. Frenkel, *Yakov Ilich Frenkel: His Work, Life, and Letters* (Birkhäuser Verlag, Basel/Boston, 1996).
- Z. J. Huang, L. Luo, *Science* **350**, 42–44 (2015).
- S. Grillner et al., *Nat. Neurosci.* **19**, 1118–1122 (2016).
- L. Ferry, *La révolution transhumaniste [The Revolution of "Transhumanism"]*. (Plon, Paris, 2016).
- J.-G. Ganascia, *Le mythe de la singularité [The Myth of Singularity (in French)]*. Science Ouverte (Seuil, Paris, 2017).
- Y. Frégnac, G. Laurent, *Nature* **513**, 27–29 (2014).
- Z. F. Mainen, M. Häusser, A. Pouget, *Nature* **539**, 159–161 (2016).
- H. Markram, *Funct. Neurol.* **28**, 145–151 (2013).
- U. Felt, B. Wynne, "Taking European knowledge society seriously," *Report of the Expert Group on Science and Governance to the Science, Economy and Society Directorate, Directorate-General for Research* (European Commission, Brussels, 2007).
- R. Jones, *Nat. Nanotechnol.* **3**, 65–66 (2008).
- Sciences et Technologies émergentes: pourquoi tant de promesses? M. Audétat, Ed., *Emerging Sciences and Technologies* (Hermann, 2015).
- F. Panese, in *Sciences et Technologies émergentes: pourquoi tant de promesses*, M. Audétat, Ed. (Hermann, Paris, 2015), pp. 165–193.
- S. Loeve, in *Sciences et Technologies émergentes: pourquoi tant de promesses?* M. Audétat, Ed. (Hermann, Paris, 2015), pp. 91–113.
- M. Wadman, *Nature* **495**, 19 (2013).
- D. B. Pedersen, V. F. Hendricks, *Philos. Technol.* **27**, 503–518 (2014).
- A. R. Ferguson, J. L. Nielson, M. H. Cragin, A. E. Bandrowski, M. E. Martone, *Nat. Neurosci.* **17**, 1442–1447 (2014).
- S. Choudhury, J. R. Fishman, M. L. McGowan, E. T. Juengst, *Front. Hum. Neurosci.* **8**, 239 (2014).
- S. Theil, *Sci. Am.* **313**, 36–42 (2015).
- L. Schneider, *The laborious delivery of Markram's brainchild*. <https://forbettecience.wordpress.com/2016/07/15/the-laborious-delivery-of-markrams-brainchild/>.
- L. Wang, *Natl. Sci. Rev.* **4**, 258–263 (2017).

ACKNOWLEDGMENTS

I thank G. Laurent and F. Engert for their supportive scientific interaction in an early draft of this text. I thank M. Yartsev, K. Grant, K. Petersen, F. Frégnac-Clave, and the two anonymous reviewers for helpful comments in the final steps of this manuscript.

10.1126/science.aan8866