iPSCs to cell lines from healthy donors that share no or only partial genetic background poses considerable risks for interpreting a supposedly disease-related phenotype. This becomes particularly relevant for most late-onset diseases that typically show slow progression of pathophysiological changes and therefore are expected to display only subtle changes in vitro, which may not fully manifest during a short study period.

Recent advances in gene-editing technologies enable the targeted modification of human cells for gene disruptions, genetic repair, or insertion of reporter genes (*14*). The ability to modify single base pairs, thereby seamlessly correcting or introducing disease-causing mutations in human pluripotent stem cells, allows the creation of genetically controlled experimental model systems in which the disease-causing genetic variation is the sole experimental variable (*19*–*21*). This could substantially simplify the analysis of the interaction between these genomic variants and disease phenotypes, thereby revealing new insights into the pathophysiology of monogenic and complex diseases.

The combination of iPSC, gene editing, and genome-wide technologies gives us the opportunity to systematically and faithfully model human disease in relevant human cell types. Acknowledging the inherent limitations and subjecting iPSC technology to the same rigor that has become standard in other model systems will make it an indispensable tool for biomedical research.

References

1. S. Yamanaka, *Cell Stem Cell* **10**, 678 (2012).
2. T. T. Onder, G. Q. Daley, *Curr. Opin. Genet. Dev.* **22**, 500 (2012).
3. G. L. Boulting *et al.*, *Nat. Biotechnol.* **29**, 279 (2011).
4. C. Bock *et al.*, *Cell* **144**, 439 (2011).
5. H. Kim *et al.*, *Cell Stem Cell* **8**, 695 (2011).
6. F. Soldner *et al.*, *Cell* **136**, 964 (2009).
7. R. Lister *et al.*, *Nature* **471**, 68 (2011).
8. B. W. Carey *et al.*, *Cell Stem Cell* **9**, 588 (2011).
9. M. Stadtfeld *et al.*, *Nat. Genet.* **44**, 398, S1 (2012).
10. S. Mekhoubad *et al.*, *Cell Stem Cell* **10**, 595 (2012).
11. A. Gore *et al.*, *Nature* **471**, 63 (2011).
12. S. M. Hussein *et al.*, *Nature* **471**, 58 (2011).
13. L. Cheng *et al.*, *Cell Stem Cell* **10**, 337 (2012).
14. D. Hockemeyer *et al.*, *Nat. Biotechnol.* **29**, 731 (2011).
15. The 1000 Genomes Project Consortium, *Nature* **490**, 56 (2012).
16. J. A. Tennessen *et al.*, *Science* **337**, 64 (2012).
17. M. R. Nelson *et al.*, *Science* **337**, 100 (2012).
18. J. Majewski, T. Pastinen, *Trends Genet.* **27**, 72 (2011).
19. F. Soldner *et al.*, *Cell* **146**, 318 (2011).
20. K. Yusa *et al.*, *Nature* **478**, 391 (2011).
21. V. Sebastiano *et al.*, *Stem Cells* **29**, 1717 (2011).

10.1126/science.1227682

NEUROSCIENCE

Building the Human Brain

A computational model of the human brain simulates complex behaviors.

Christian K. Machens

The human brain is exceedingly complex and studying it encompasses gathering information across a range of levels, from molecular processes to behavior. The sheer breadth of this undertaking has perhaps led to an increased specialization of brain research and a concomitant fragmentation of our knowledge. A potential solution is to integrate all of this knowledge into a coherent simulation of the brain (*1*). However, simply "building" a brain from the bottom up by replicating its parts, connections, and organization fails to capture its essential function—complex behavior (*2*). Instead, just as engineers can only construct cars and computers because they know how they work, we will only be able to construct a brain if we know how it works—that is, if we understand the computations that are carried out in individual brain areas, and how these computations are implemented on the level of neural networks. On page 1202 of this issue, Eliasmith *et al*. (*3*) make headway toward this benchmark by presenting just such a large-scale computational model of the human brain that can simulate a variety of complex behaviors.

The model of Eliasmith *et al*., called the Semantic Pointer Architecture Unified Network (or "Spaun"), can observe visual images and indicate its responses with a physical model of an arm. The authors show that Spaun can perform eight quite diverse tasks, all of which involve the presentation of various images (mostly numbers) and subsequent motor responses (numbers drawn with the arm). The tasks range from simple perceptual tasks (image recognition) to working memory tasks (sequence recall) and reinforcement learning tasks (a gambling task) to complex cognitive tasks [aspects of an intelligence quotient (IQ) test]. Spaun performs all of these tasks based on the activity of 2.5 million simulated neurons that are organized into subsystems resembling different brain areas, and wired up to provide the necessary functionality.

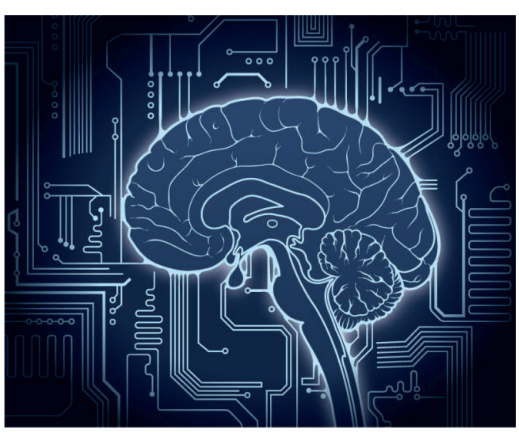

CREDIT: Y. HAMMOND/*SCIENCE*

The incoming visual image is first "compressed" so that any irrelevant or redundant image parts are eliminated. Compression is achieved through a so-called hierarchy of restricted Boltzmann machines (in a feed-forward neural network) in which successive layers extract more and more complex features of the visual input (*4*). The authors associate these layers with areas of the brain that comprise the ventral visual stream (the primary and secondary visual cortex, the extrastriate cortex, and the inferior temporal cortex). On the motor side, Spaun turns a simple, internally generated command signal (for example, how to draw the number six) into a complex arm movement that consists of many elementary motions. The relevant computations are based on optimal control theory (*5*) and are associated with supplementary motor areas and the primary motor cortex. This combined compression and expansion of information solves the "curse of dimensionality" in dealing with the environment—the problem of handling an incredibly vast amount of sensory information and choosing from an incredibly vast amount of possible motor outputs.

The actual cognitive machinery of Spaun consists of two intertwined components: a working memory system [prefrontal cortex

Champalimaud Neuroscience Programme, Champalimaud Centre for the Unknown, Av. Brasilia, 1400-038 Lisbon, Portugal. E-mail: christian.machens@neuro.fchampalimaud.org

(PFC)] and an action selection system (basal ganglia and thalamus). The latter controls the current brain state and is partly inspired by the theory of reinforcement learning and a contemporary model of the basal ganglia (*6*). Spaun's working memory is based on an innovative amalgam of neural integrators taken from (computational) neuroscience (*7*), and so-called convolution memories, taken from (mathematical) psychology (*8*). The neural integrators provide Spaun with a network mechanism to store information, whereas the convolution memories provide a memory-efficient algorithm that allows Spaun to bind newly arriving information with already stored information. Spaun thereby replicates behavioral effects such as primacy and recency: a tendency to remember the first and last item in a list better than any others.

An additional working memory system is used to automatically infer relations between past and present stimuli. This automatic inference amounts to a simple case of syntactic generalization and is made possible by the way Spaun handles the representation of numbers. These representations provide a direct link to the symbolic computations common to the connectionist (and computer science) literature. Spaun uses these computations to pass some basic aspects of an IQ test.

The systems in Spaun allocated to the PFC bridge the gap between abstract, symbolic computations and the activity of single neurons. Particularly with respect to the convolution memories, Eliasmith *et al*. make some interesting predictions of how firing rates of neurons (the average number of electric impulses or "spikes" per unit time) should evolve during sequential working memory tasks. These predictions are worth investigating experimentally, especially because sequential working memory tasks have been employed in monkey electrophysiology, so that the right type of data may already be known.

In each of Spaun's areas or modules, the actual information is processed through populations of spiking (active) neurons. The link between the high-level computations performed by the networks and the low-level computations performed by individual neurons is constructed by means of the "neural engineering framework" (*9*), which specifies how to implement arbitrary mathematical vector operations in spiking neural networks. The framework assumes that information is read out linearly from neural firing rates and is transformed through the nonlinearities of neural activation functions. Consequently, the information processed by each area is distributed across neurons, and thereby roughly matches some of the known electrophysiological features of the areas, such as the diverse tuning of neural firing rates to sensory stimuli or motor outputs.

Given the scope of the model engineered by Eliasmith *et al*., it is not surprising that many aspects of Spaun deviate from real brains. For instance, the spiking activity within many areas differs in several aspects (including basic statistical aspects) from that measured in real brains. To what extent this problem can be remedied in future work, or to what extent these discrepancies point toward fundamentally different computations in the brain, is currently unclear. Spaun's principal shortcoming is that it is essentially hard-wired and cannot learn any new tasks. Its architecture is quite flexible and not bound to particular tasks, and several parts of Spaun are learned (such as the visual hierarchy or the values of actions). However, learning in its broadest sense, such as learning completely new tasks, is one of the issues that the authors have deliberately—and perhaps wisely—side-stepped. Indeed, just as Spaun falls short on this point, so does our understanding of the brain. By assembling a large amount of brain know-how into one model, Eliasmith *et al*. have provided a coherent theory of how the brain works (with the exception of learning). To paraphrase the statistician George Box, their model is likely to be wrong, but it is certainly useful. Moreover, the authors have provided an opening gambit for top-down approaches to large-scale simulations of the brain. Spaun levels the playing field by setting a new goal and a new benchmark for such simulations: to not simply incorporate the largest number of neurons or the greatest amount of detail, but to reproduce the largest amount of functionality and behavior.

References

1. H. Markram, *Nat. Rev. Neurosci.* **7**, 153 (2006).
2. E. M. Izhikevich, G. M. Edelman, *Proc. Natl. Acad. Sci. U.S.A.* **105**, 3593 (2008).
3. C. Eliasmith *et al., Science* **338**, 1202 (2012).
4. G. E. Hinton, R. R. Salakhutdinov, *Science* **313**, 504 (2006).
5. E. Todorov, *Nat. Neurosci.* **7**, 907 (2004).
6. K. Doya, *Nat. Neurosci.* **11**, 410 (2008).
7. C. D. Brody, R. Romo, A. Kepecs, *Curr. Opin. Neurobiol.* **13**, 204 (2003).
8. T. A. Plate, *IEEE Trans. Neural Netw.* **6**, 623 (1995).
9. C. Eliasmith, *Neural Comput.* **17**, 1276 (2005).

10.1126/science.1231865

ECOLOGY

Integrating the Captive and the Wild

Modern conservation management increasingly integrates approaches developed in zoos with those from the wild to actively manage populations.

Kent H. Redford,[1] Deborah B. Jensen,[2] James J. Breheny[3]

Conservation biology was founded with a focus on the plight of species by a group of scientists that included representatives of the zoo and botanical garden communities (*1*). In the decades since then, zoos have sought to minimize their impact on populations of wild animals and to contribute to field conservation while maintaining genetically and demographically sustainable captive populations and using these animals to help educate the public about conservation. In a recent article, Lacy argued that zoos are not achieving genetic sustainability for target animal populations (*2*). However, zoos have contributed a set of approaches to species management that are being integrated with those from field conservation to create hybrid forms of species management better suited to present-day conditions.

Examples abound of the new challenges facing species conservation. Novel and emerging diseases threaten wildlife populations that will require new, active methods of veterinary management (*3*). The effects of climate change are driving us to consider moving populations of species outside their current geographic ranges to new areas where they may survive (*4*). In these circumstances, the view that species can be effectively conserved with minimal management simply by creating large areas of natural habitat no longer holds true.

[1]Archipelago Consulting, Irvington, NY 10533, USA. [2]Woodland Park Zoo, Seattle, WA 98130, USA. [3]Wildlife Conservation Society, Bronx, NY 10460, USA.

Published by AAAS

Building the Human Brain

Christian K. Machens

Science **338** (6111), 1156-1157.
DOI: 10.1126/science.1231865