Leading Edge
Minireview

Cell

A Biological Imitation Game

Christof Koch[1,*] and Michael A. Buice[1]
[1]Allen Institute for Brain Science, Seattle, WA 98109, USA
*Correspondence: christofk@alleninstitute.org
http://dx.doi.org/10.1016/j.cell.2015.09.045

The digital reconstruction of a slice of rat somatosensory cortex from the *Blue Brain Project* provides the most complete simulation of a piece of excitable brain matter to date. To place these efforts in context and highlight their strengths and limitations, we introduce a Biological Imitation Game, based on Alan Turing's Imitation Game, that operationalizes the difference between real and simulated brains.

The much awaited *opus magnum* of the "Blue Brain Project" ("BBP") is appearing in this issue of *Cell* (Markram et al., 2015). The brain child, so to speak, of the electrophysiologist Henry Markram, the BBP is a large-scale, international collaborative effort that seeks to simulate a synthetic brain using biologically realistic models of synapses and nerve cells. The BBP is an exceptional (neuro)-engineering effort based on a bold and long-term vision. To be successful, it requires meticulous attention to detail, a highly interdisciplinary team, and large-scale and stable resources. Financially, it is backed by the Swiss government, with more recent contributions by the pan-European Human Brain Project.

Initiated a decade ago, the ambitious aims of the BBP are to amass all relevant knowledge pertaining to the mammalian brain; to distill the associated information into an integrated, standardized, and open-access database; and to endow this massive yet static entity with life by simulating spontaneous and sensory-driven synaptic activity and the associated cellular electrical responses using partial differential equations. The *Cell* paper describes a first-draft digital reconstruction of a section of rat somatosensory cortex. It is, without doubt, the most complete simulation to date of a piece of excitable brain matter.

The utility of such massive simulations has been much debated—in particular, as they relate to the ill-defined goal of "understanding the brain" that is on the masthead of the Human Brain Project and the other large-scale brain projects initiated over the last 3 years (Kandel et al., 2013). So before we dive into the details, let us step back and consider how best to think about this category-defying model.

Introducing a Biological Imitation Game

In an effort to circumvent the question "can machines think"?, the mathematician Alan Turing (Turing, 1950) introduced "The Imitation Game" in which an interrogator is tasked with determining the identity of two examinees as "male" or "female." This game is imagined to be played by human beings. Turing introduced as a measure of the thinking ability of machines the length of time the Interrogator needed to play before realizing that one of the other players had been replaced by a machine. The longer this takes, the more this machine could be argued to operationally behave like a thinking human.

Let us imagine our own Imitation Game, in which a lead investigator running a lab is recording from neocortical neurons in a rat but has also devised a sophisticated computer model of its electrical behavior. Being quite busy with grant writing and other administrative tasks, she rarely enters the lab, relying instead on reports from her students and postdocs.

After some discussion, they agree on a protocol X—say, stimulate sensory input into somatosensory cortex while recording from layer 5 pyramidal neurons. The experimentalists in her lab judiciously carry out these manipulations and bring her a set of plots—extra-cellular recordings, current source density plots, and so on. Their computational colleagues do likewise and generate the same type of graphs. Let's call these, respectively, A(X) and B(X).

For the purposes of this operational definition, we assume that the investigator does not know whether A or B is the experimental data, taken here as ground truth. Both look superficially similar. To distinguish between the real data, A(X), and the simulated data, B(X), the investigator must devise clever experiments for her group to carry out on both, probing for the simulation's failure modes. Clearly, the longer that it will take the investigator, given her accumulated knowledge and insights about real brains, to confidently identify the provenance of the data, the more the model can be said to mimic the real system. We take as a measure of the quality of the simulation the length of time it takes her to confidently identify the provenance of the data and the level of sophistication of question required to do so.

Thus, the question of how good any one computer simulation of reality is can be replaced by an operational measure of how long an expert can be fooled by the simulation. As the model becomes more and more sophisticated, this will take longer and longer. In the limit of a completely faithful digital simulacrum, reliably judging which is real and which is synthetic will become impossible. Note that, in the spirit of Turing, our operational definition bypasses eristic discussions of what is meant by "understanding" the brain.

Looking Under the Hood

We will now summarize the approach taken by the BBP team. The model simulates a segment of 2-week-old rat somatosensory cortex, 2 mm tall and 210 μm in radius, about the size of a cortical column. The individual nerve cells are derived from the Markram laboratory's heroic efforts over almost two decades to record and label more than 14,000 neurons in slices taken from

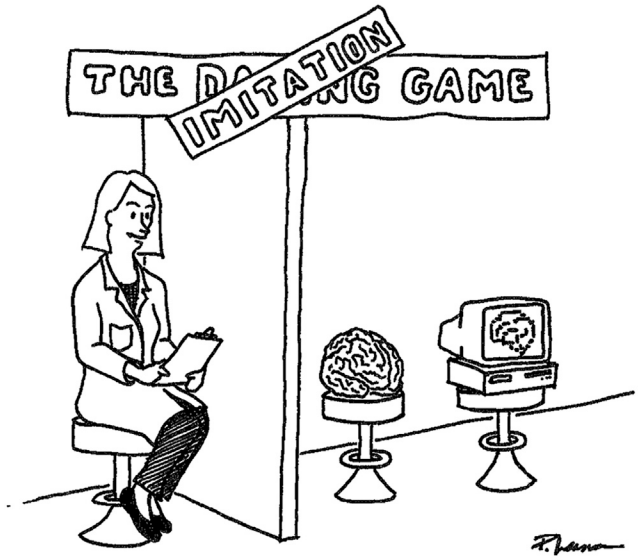

Bachelor #1: What is your favorite neurotransmitter to release during spreading waves? Illustration by Phil Lesnar, courtesy of the Allen Institute.

young male Wistar rats (e.g., Markram et al., 1997 and 2004). The dendritic trees and local axonal morphologies of 1,009 of these cells are completely digitally reconstructed in 3D. The cells fall into 55 morphological classes, as determined by the layer in which the cell bodies were located and their morphology: 42 GABAergic and 13 glutamergic types. This is roughly in line with the literature (Harris and Shepherd 2015). The electrical properties of these neurons are inferred by recording from the cell bodies of 3,900 neurons in young rat slices. Using a standard electrophysiological protocol, 11 types of firing patterns—ten associated with inhibitory and one (continuous adapting) with excitatory cells—can be distinguished. From a combined dataset of 511 morphologically and electro-physiologically characterized neurons, the BBP team infers 207 morpho-electrical types (207 out of a possible 55 × 11 = 605 cell types).

Next, the team populates their column with 31,320 neurons drawn from these 207 types, with multiple clones of each cell type randomly and independently spatially distributed across five layers (L1, L2/3, L4, L5, and L6) in accordance with measured cell densities and immunohistochemical staining.

The electrical behavior of each neuron is modeled using a classical Hodgkin-Huxley formalism extended to include 13 voltage- and calcium-dependent conductances inserted into the membrane at the cell body and the proximal and distal dendritic tree (Hay et al., 2011).

At this stage, these neurons with their thick dendrites and thin non-myelinated axons are scattered throughout the cylindrical volumes, a dense jungle with roots, flowering plants, brushes, trees, and their canopy intertwined. Both cortex and forest share a predominantly vertical organization. They differ as only the former have highly specialized junctions among their members, chemical synapses that convert electrical activity in presynaptic axonal boutons into neurotransmitter release and back into electrical activity at postsynaptic sites on spines, dendrites, or cell bodies.

The Connectomics Algorithm

But where to place the all-important synapses? Here, we come to the beating heart of the BBP, an algorithmic approach to reconstruct synaptic connections (described in detail in a companion paper [Reimann et al., 2015]). Whenever an axonal profile comes into close encounter with a dendrite, a potential synapse is marked. Such a potential synapse is the child of chance, based on the probabilities that an axonal bouton finds itself near a dendrite. This is known as Peter's rule, stipulating that axons seek out dendrites randomly—the higher the bouton density and the larger the dendrite, the bigger the connection probability. By this measure, most neurons connect to most other neurons (there are, after all, close to one billion synapses in each mm^3 of gray matter). Peters' rule places a hard constraint on the connectivity, in the sense that both axons and dendrites need to be simultaneously present for a synapse to form.

Reimann and colleagues (Reimann et al., 2015) know from light- and electron-microscopic data derived by a great many labs, including their own, that any one functional connection among two neurons is implemented by many anatomical synapses. Put differently, if two cortical neurons are connected, they form anywhere between a handful and perhaps 20 anatomical synapses; only rarely does a single synapse connect two neurons. The observation that simple axonal-dendritic proximity—which would predict a majority of singleton connections—does not explain connectivity highlights the limit of Peters' rule. It also points to the existence of learning rules that prune and grow synaptic connections.

The BBP captures such observed anatomical and morphological constraints in a probabilistic connectomics algorithm that prunes the vast number of spatial appositions, leaving a much smaller number of actual anatomical synapses: 638 million potential synapses yield 37 million actual synapses that form 8.1 million connections, with 3.6 excitatory and 13.9 inhibitory synapses per connection (all values given here are averages; the paper gives statistical distributions). These data are compatible with a recent electron-microscopic reconstruction of a sliver of cortical tissue (Kasthuri et al., 2015).

The physiology of these 37 million synapses is based on extrapolated data from paired-cell recordings, including their all-important short-term facilitating or depressing dynamic behavior. 207 types of neurons could, in principle, be connected using 207 × 207 = 42,849 different types of synapses. Unfortunately, for the vast majority of synaptic types, no experimental recordings are available, and reasonable inter- and extrapolations and assumptions must come to the rescue.

Booting up the Virtual Slice

With synaptic input in place, fortified by somatic depolarization of cells to facilitate their firing (mimicking arousing neuromodulator substances active under in vivo conditions) and some spontaneous synaptic release, the transmembrane potential of each cell evolves according to the nonlinear cable equation. This is a partial differential equation that describes how the membrane

potential evolves in thin and elongated neuronal processes whose membrane properties are described by capacitive and conductive processes (Koch 1999). The numerical simulation uses a parallel version of the NEURON software that has been instrumental to the field (Carnevale and Hines 2006), running on an IBM Blue Gene/Q supercomputer at the Swiss National Supercomputing Center (CSCS) in Lugano.

Powering up this digital slice yields reasonable-looking firing behavior and slow oscillatory bursts at 1 Hz. The simulation goes through some finger exercises, demonstrating its neurobiological authenticity. The overall network activity is very sensitive to the extracellular calcium concentration $[Ca^{2+}]_o$ and the level of somatic depolarization of all neurons. $[Ca^{2+}]_o$, varying between 1 mM under in vivo and 2 mM under in vitro conditions affects the probability of synaptic release. In effect, $[Ca^{2+}]_o$ controls the network's susceptibility to synaptic perturbation, while the depolarization level controls the spontaneous firing of the neurons. Manipulating these variables flips the network between two qualitatively different dynamic regimes. Under in vitro conditions ($[Ca^{2+}]_o$ = 2 mM and less somatic depolarization), the network is highly responsive to perturbations, producing stereotypical and synchronized responses, such as spreading waves (Figure 14 in Markram et al., 2015), while under in vivo conditions ($[Ca^{2+}]_o$ closer to 1 mM and greater somatic depolarization), responses are graded and asynchronous. This transition is quite sharp and is mediated by differential changes in the excitatory to inhibitory synaptic balance.

Given the very large number of approximations and extrapolations that underlie the BBP model, the fact that neurons do not erupt into either a paroxysm of epileptic discharge nor descend into a coma-like state of electrical silence but act, to a first order, as neurons in slices do, is a remarkable achievement. This is an important initial success in our Imitation Game, for the very reason that it allows the game to continue. Our investigator can begin asking tougher questions (of course, it also begs the question of whether all of these details were necessary in the first place, to which we will return later on).

Pushing further, the BBP team highlights aspects of in vivo activity that are recapitulated by the model—in particular, the extent to which groups of neurons are active in a correlated manner (for instance, when spikes fire near simultaneously), have no particular temporal relationship, or are anti-correlated. Indeed, the digital simulacrum exhibits the same correlational structure as predicted on theoretical grounds and observed in rodent cortex by Renart et al. (2010). In particular, activities in excitatory and inhibitory neurons track each other, generating negative correlations in synaptic structure that cancel strong shared input. Similarly, the simulation shows repeating triplet structures in the synchronous regime that do not appear in the asynchronous regime of low calcium. Finally, a recent experimental study of Okun et al. (2015) establishes the existence of many highly correlated "chorister" neurons and a few isolated "soloists." The former are neurons whose activity is tightly correlated with the average activity of the network that they are embedded in, while the latter appear to actively fire independently of the rest of the network. The simulation recapitulates this observation as well.

These demonstrations are comforting validations of the model in terms of our Imitation Game. Of course, the transition observed by the BBP team is clearly describable by simple models (there is a long history, with notable contributions from Wilson and Cowan (1972) and Amari (1975) of applied mathematics devoted to the solution space of so-called "neural field theories," demonstrating complex behavior as neural characteristics are varied (see Bressloff, 2014), including spreading waves (Beurle, 1956).

The most interesting observation of the present study is the sensitivity of network behavior to extracellular calcium concentration and that propagating waves may be an artifact of too high $[Ca^{2+}]_o$ levels (which would not, however, explain the waves seen in Xu et al., 2007). Yet the enormous effort of the BBP to carefully reconstruct details of the cortical circuit are probably irrelevant to this prediction, as the dependency on extracellular calcium appears to enter through an empirical model of synaptic neurotransmitter release and not through any network dynamics. In the same vein, the mechanism of Renart et al. (2010) was obtained in a model of conductance-based integrate-and-fire neurons. The additional detail added by the BBP does not affect the mechanism.

Uncovering the Limits of the Model

In the Imitation Game, there is one cunning strategy that is guaranteed to lead to the denouement of the impostor. Simply probe further and further into the micro-structure: while it is not clear whether there is any "ultimate" level of reality—elementary particles such as the Higgs boson, fields, or superstrings (depending on the energy invested in probing the micro-structure) —simulations abruptly bottom out. To wit, the investigator could ask her lab members to resolve fine details of the macroscopic Na^+, K^+, Ca^{2+}, and Cl^- membrane currents. If appropriate electrodes and amplifiers are used, the recordings will uncover discrete, microscopic, and stochastic ionic channels, while the current BBP model stops with the continuous and deterministic Hodgkin-Huxley currents. Bingo!

And therein lies an important lesson. If the real and the synthetic can't be distinguished at the level of firing rate activity (even though it is uncontroversial that spiking is caused by the concerted action of tens of thousands of ionic channel proteins), the molecular level of granularity would appear to be irrelevant to explain electrical activity. Teasing out which mechanisms contribute to any specific phenomena is essential to what is meant by understanding.

Markram et al. claim that their results point to the minimal datasets required to model cortex. However, we are not aware of any rigorous argument in the present triptych of manuscripts (Markram et al., 2015; Ramaswamy et al., 2015; Reimann et al., 2015), specifying the relevant level of granularity. For instance, are active dendrites, such as those of the tall, layer 5 pyramidal cells (Hay et al., 2011), essential? Could they be removed without any noticeable effect? Why not replace the continuous, macroscopic, and deterministic HH equations with stochastic Markov models of thousands of tiny channel conductances? Indeed, why not consider quantum mechanical levels of descriptions? Presumably, the latter two avenues have not been chosen because of their computational burden and the

intuition that they are unlikely to be relevant (Strassberg and DeFelice, 1993; Koch and Hepp, 2006). The Imitation Game offers a principled way of addressing these important questions: only add a mechanism if its impact on a specific set of measurables can be assessed by a trained observer.

Consider the problem of numerical weather prediction and climate modeling, tasks whose physico-chemical and computational complexity is comparable to whole-brain modeling. Planet-wide simulations that cover timescales from hours to decades require a deep understanding of how physical systems interact across multiple scales and careful choices about the scale at which different phenomena are modeled. This has led to an impressive increase in predictive power since 1950, when the first such computer calculations were performed (Bauer et al., 2015). Of course, a key difference between weather prediction and whole-brain simulation is that the former has a very specific and quantifiable scientific question (to wit: "is it going to rain tomorrow?"). The BBP has created an impressive initial scaffold that will facilitate asking these kinds of questions for brains.

What is also noteworthy is that, with the publication of this paper, the BBP has publicly released a treasure trove of morphological, electrical, and connectional data collected over the last years by members of the Markram lab that was used in this model. For more details, see the companion manuscript (Ramaswamy et al., 2015). The experience of the Allen Institute (www.brain-map.org) has demonstrated the impact of releasing all relevant data and models soon after they pass an internal quality control stage (usually several years prior to publication of any associated results). The release of the massive database of the BBP is thus a very positive development.

It is clear that physically detailed, whole-brain simulations enable neuroscientists to answer specific questions that are difficult to address experimentally. The BBP has provided a powerful tool in this regard. Such simulations may be essential to develop therapeutics for brain-based diseases. Whole-brain models can also be said to encapsulate the present state of knowledge, as per the Imitation Game. However, any simulation-gained knowledge must be supported and complemented by theories that, by isolating the relevant variables, enable us to deeply understand the most organized and highly excitable piece of matter in the known universe.

REFERENCES

Amari, S. (1975). Biol. Cybern. *17*, 211–220.

Bauer, P., Thorpe, A., and Brunet, G. (2015). Nature *525*, 47–55.

Beurle, R.L. (1956). Phil. Trans. Roy. Soc. B. *240*, 55–94.

Bressloff, P.C. (2014). Waves in Neural Media: From Single Neurons to Neural Fields (Springer).

Carnevale, N.T., and Hines, M.L. (2006). The NEURON Book (Cambridge University Press).

Harris, K.D., and Shepherd, G.M.G. (2015). Nature Rev. Neurosci. *18*, 170–181.

Hay, E., Hill, S., Schumann, F., Markram, H., and Segev, I. (2011). PLoS Comp. Biol. *7*, e1002107.

Kandel, E.R., Markram, H., Matthews, P.M., Yuste, R., and Koch, C. (2013). Nat. Rev. Neurosci. *14*, 659–664.

Kasthuri, N., Hayworth, K.J., Berger, D.R., Schalek, R.L., Conchello, J.A., Knowles-Barley, S., Lee, D., Vázquez-Reina, A., Kaynig, V., Jones, T.R., et al. (2015). Cell *162*, 648–661.

Koch, C. (1999). Biophysics of Computation (New York, NY: Oxford University Press).

Koch, C., and Hepp, K. (2006). Nature *440*, 611–612.

Markram, H., Lübke, J., Frotscher, M., Roth, A., and Sakmann, B. (1997). J. Physiol. *500*, 409–440.

Markram, H., Toledo-Rodriguez, M., Wang, Y., Gupta, A., Silberberg, G., and Wu, C. (2004). Science *5*, 793–807.

Markram, H., Muller, E., Ramaswamy, S., Reimann, M.W., Abdellah, M., Sanchez, C.A., Ailamaki, A., Alonso-Nanclares, L., Antille, N., Arsever, S., et al. (2015). Cell *163*, this issue, 456–492.

Okun, M., Steinmetz, N.A., Cossell, L., Iacaruso, M.F., Ko, H., Barthó, P., Moore, T., Hofer, S.B., Mrsic-Flogel, T.D., Carandini, M., and Harris, K.D. (2015). Nature *521*, 511–515.

Ramaswamy, S., Courcol, J.-D., Abdellah, M., Adaszewski, S., Antille, N., Arsever, S., Guy Antoine, A.K., Bilgili, A., Brukau, Y., Chalimourda, A., et al. (2015). Front. Neural Circuits *9*, 44.

Reimann, M.W., Muller, E.B., Ramaswamy, S., and Markram, H. (2015). Front. Comput. Neurosci. *9*, 28.

Renart, A., de la Rocha, J., Bartho, P., Hollender, L., Parga, N., Reyes, A., and Harris, K.D. (2010). Science *327*, 587–590.

Strassberg, A.F., and DeFelice, L.J. (1993). Neural Comput. *5*, 843–855.

Turing, A.M. (1950). Mind *59*, 433–460.

Wilson, H.R., and Cowan, J.D. (1972). Biophys. J. *12*, 1–24.

Xu, W., Huang, X., Takagaki, K., and Wu, J.Y. (2007). Neuron *55*, 119–129.