Neuron
NeuroView

Architecting Discovery: A Model for How Engineers Can Help Invent Tools for Neuroscience

Edward S. Boyden[1,2,3,4,5,6,*] and Adam H. Marblestone[1,*]
[1]MIT Media Lab, Massachusetts Institute of Technology, Cambridge, MA 02139, USA
[2]McGovern Institute, Massachusetts Institute of Technology, Cambridge, MA 02139, USA
[3]Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139, USA
[4]Department of Biological Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139, USA
[5]Koch Institute, Massachusetts Institute of Technology, Cambridge, MA 02139, USA
[6]Lead Contact
*Correspondence: esb@media.mit.edu (E.S.B.), amarbles@media.mit.edu (A.H.M.)
https://doi.org/10.1016/j.neuron.2019.03.023

How can engineers help invent tools for neuroscience? We here explore a model for how engineers can choose problems and map out possible technologies that help address them. We also discuss design principles of tools and the role of failure.

Twentieth-century scientific research yielded many foundational discoveries that resulted in inventions such as computers and wireless communication. It was a productive century. Why? One factor may be that these successes were rooted in relatively mature sciences such as physics and chemistry. For example, in physics, there are a small number of building blocks—such as protons and electrons—and a small number of ways they interact—through the laws of quantum mechanics and special relativity. In chemistry, there is a fairly short list of kinds of atoms, enumerated in the periodic table. In contrast, some of the sciences we struggle with in the twenty-first century, like neuroscience, involve many building blocks, which interact in many different ways. For example, consider the many different kinds of brain cell, the number of which is unknown even for the normal human brain, much less in the diversity of neurological and psychiatric conditions. It should be no surprise that we struggle to explain how the brain computes. In addition, treatments for brain diseases remain relatively few; brain drugs cost perhaps $1 billion each for clinical development, taking a decade for clinical approval, and even those that make it to human trials fail 90% of the time to be approved for human use (Miller, 2010). Because of the brain's complexity, it shouldn't be surprising that many of our hypotheses about the mechanisms underlying biological functions, or about intervention targets for diseases, end up proven wrong.

How should engineers think about building tools that help confront neuroscience challenges today? One metaphor that is sometimes used nowadays is the "moonshot." The actual Apollo program was based on solid physics knowledge. It was a tour-de-force engineering and management challenge, no doubt, but the number of foundational scientific unknowns was small by 1962 (for example, rockets had already taken humans into outer space). But imagine trying to land on the moon in the year 1600—before we understood calculus, mechanics, and aerodynamics. All the money on the planet, in the year 1600, might not have gotten you to the moon. For some neuroscience challenges, we may be, metaphorically speaking, closer to the year 1600 than 1960. This doesn't mean that systematic, concerted, and organized efforts can't work in neuroscience, but it means that we have to choose the *targets* for those efforts correctly. Another concept that sometimes arises is "big data." Sheer quantity, of course, is not enough; data ideally would be sufficiently precise, and have the right spatial and/or temporal resolution, to pinpoint fundamental building blocks and how they interact. For example, brain mapping technologies that can reveal synaptic connections may help scientists answer questions that are not amenable to analyses with lower-resolution methods.

Thus, one powerful way of approaching problems of neuroscience is to try to understand complex systems in terms of their underlying building blocks and interactions, as in fields like physics and chemistry before: to understand networks and circuits in terms of cells, and cells in terms of their component biomolecules. From these maps of the building blocks and interactions, we might be able to devise new hypotheses in a more systematic way. In short, getting to the ground truth, which we define for the purposes of this essay as a level of precision of analysis that allows building blocks and their fundamental interactions to be directly understood, can be worth it.

Some brain diseases have been confronted successfully, sometimes through highly unexpected means. Julius Wagner-Jauregg, who won the Nobel Prize in 1927, took patients with dementia paralytica, associated with late-stage syphilis, and inoculated them with malaria. Malaria caused a high fever that would kill the bacterium that causes syphilis. Of course, malaria sometimes killed the patient, too. One year later, antibiotics were discovered. Antibiotics directly confront the bacterium, with fewer side effects. Of course, treating most brain diseases is far more complex than fighting a single pathogen that can be selectively killed by a drug. Brain diseases often involve subtle changes in different brain cell types, over long periods of time. This might make it even more important, for many brain diseases, to acquire maps of the building blocks and their interactions.

In general, there is a lack of technology for seeing and perturbing complex biological systems, including the brain, with sufficient precision. Building tools that can

get down to ground-truth levels of description is hard to do. In neuroscience, however, there isn't universal agreement about what the ground-truth level of description is—do we need to map individual biomolecules, or synapses, or cell types, before we can devise truly powerful explanations? And if we assume a certain level to be ground truth and try to scale up the relevant technologies to map the brain at that level of description, what if we are wrong (Marblestone and Boyden, 2014)? Nevertheless, hope emerges from examples like the crab stomatogastric ganglion, where the ability to record from every cell, to map the connectivity between them, and to derive molecular profiles of the component cells, has yielded many fundamental insights into neural computation (Prinz et al., 2004; Schulz et al., 2006), suggesting that integrative observations of the activity, wiring, and molecular composition of neurons throughout entire neural circuits may be useful.

Another reason building tools in neuroscience is hard is that technology inventors in neuroscience must often think backward from biological problems to design relevant tools. For example, one class of biological problem is to understand what kinds of behaviors or pathological states a set of neurons can, through their electrical activity, initiate or sustain. Optogenetic tools (Boyden, 2011) enable the control of the electrical activity of specific cells with light and thus are useful for causally investigating the contribution of specific cells to a behavior or pathological state. However, the question of what a ground-truth level of precision is sometimes arises. As one example, some studies using optogenetics treat electrical potentials in neurons as an abstraction layer, ignoring the underlying ion chemistries associated with the electrical pulses. But in some cases, this assumption breaks down—for example, light-activated chloride and proton pumps, and light-activated chloride channels, can be used to silence neuronal electrical activity, but in each of these cases, the ions translocated during optogenetic silencing can sometimes affect synaptic communication (Mahn et al., 2016; Raimondo et al., 2012).

Making a powerful, easy-to-use technology sometimes requires serendipity. But in contrast to the expensive, slow luck required to develop a drug, the kind of luck needed here may, in principle, be made far cheaper and faster to obtain. For example, inventing a new technology sometimes requires you to stumble across the right reagent, such as with the aforementioned optogenetic tools or as with early characterizations of CRISPR by scientists interested in producing better yogurt (Barrangou et al., 2007). What if we could invent new tools by searching for such reagents systematically, in various ecological niches, in a high-throughput way? Innovation sometimes requires you to connect the dots across fields, as in the case of next-gen nucleic acid sequencing, where a combination of biological, chemical, and optical ideas were fused into an impactful technology. What if we could invent new tools by connecting such dots in a highly scalable, deliberate way? In our group at MIT, and with our collaborators, we are beginning to explore whether one might develop a discipline around the effective development of new tools, which one might call "architecting discovery." This discipline may be a learnable, teachable skill. Although these are early days, and much is in flux, architecting discovery may involve, at its core, asking at least three kinds of question, listed below.

The first question is: how do we pick the class of problem that the tool will solve or the space of hypotheses it will allow users to probe? One strategy is to look at examples of specific hypotheses that researchers would like to test and then look for common intellectual structure to these hypotheses, so that a single new tool might help many groups solve problems in parallel. As noted above, for example, optogenetics enables scientists to confront questions asking whether a specific cell type or neural pathway is causally involved in a behavior or pathological state. For expansion microscopy (Wassie et al., 2019), the general class of questions we considered was, could we analyze a neural circuit in terms of its component cells and biomolecules, imaging across scales?

A corollary of this principle is, when you are inventing a new tool, don't only take individual requests from other scientists. It can helpful to consider multiple scientific problems, or multiple hypotheses, and to then try to devise a technology that would help many groups confront many members of this entire set of problems or hypotheses. Is there an underlying, fundamental capability that is missing? Then, think backward from this set of problems, ideally surveying multiple fields of science and engineering for potential solutions. Perhaps you will run out of skills or knowledge, which is okay; if you have a concrete visualization of the class of problem, that can help you find and motivate collaborators with appropriate solution-domain knowledge.

The second question is, how well can you envision or roadmap out the possible solutions so that you can pick the best one possible? One trick is to look at what people are doing, and then imagine what it would mean to do the most opposite thing you can think of. This is obviously just a heuristic, but sometimes it can help. For example, for expansion microscopy, if everyone else is zooming in to see a biological system better, what happens if we blow up the biological system (in an even way, of course), instead? Another strategy is to map the landscape of currently practiced and conceived activities in an area, work to understand the fundamental limits for each one, and then ask why current approaches fall short of these fundamental limits; perhaps a stalled approach can be rebooted, if we understand its bottlenecks and can find a way to work around them by bringing in ideas from another domain. Carrying out such a process of identifying ideas that can be rebooted requires good ways to search the space of scientific approaches and papers, and potentially is assisted by being able to work in rapid communication with a large collaborative network with diverse expertise. Many ideas are hiding in plain sight, perhaps in a partly failed or misunderstood state that obscures their true value. Maybe a new scientific journal, that republishes old and forgotten results, or obscure-seeming methods, would therefore be a useful thing to read? Sometimes the knowledge underlying a major scientific tool discovery is present, but the wisdom to know that the knowledge is important is not yet present.

More rigorously, you could try to split the set of possible candidate solutions to a problem into subsets that are

non-overlapping but that collectively tile the space of possibility. The idea is to iteratively split the space of possibilities into two (or more) subsets, over and over—for example, split the set of possible solutions into one set with property A and one set that does not have property A. For example, you could take the space of all possible energy sources and split the set into two subcategories—renewable sources and non-renewable sources. Then you could split the renewables into two subsets—solar and non-solar. And already the act of splitting is forcing us to think of new ideas. What are the non-solar renewable energy sources, for example? Could we take advantage of tides, caused by the moon? The outcome of this exercise is a diagram that looks like a tree, with the forking branches representing different subsets of ideas, and the leaves of the tree representing potential projects that could be tested via calculations or via pilot studies (because of this shape, we sometimes call this method making "tiling trees" in our group). Thinking like this helped Fritz Zwicky, an astrophysicist, to predict—in the early part of the twentieth century—many phenomena that are now being explored in astrophysics today, like dark matter and gravitational lenses.

In the earliest days of optogenetics, we arrived at the core idea by going through the laws of physics systematically and thinking about what forms of energy (mechanical, magnetic, optical, etc.) could be delivered to the brain and used to control neurons. In parallel, we thought about what molecular strategies could be used to make specific sets of neuron sensitive to those forms of energy (nanoparticles, magnetic beads, light-activated proteins, etc.). Then, serendipitously, it turned out that the natural world had evolved a class of protein that, we found, contained members that worked to make genetically targeted brain cells sensitive to light.

The third question is, once you have chosen a problem and have a roadmap of the space of possible solutions, how do you choose which path is the best? In neuroscience, or other complex fields which involve lots of building blocks and interactions, one useful heuristic is a principle that you might call "the principle of applied laziness." In messy systems, like biological ones, the technologies that work really well will sometimes be very simple. Optogenetics, for example, involved essentially single genes encoding for microbial opsins. Some of this simplicity is serendipitous, of course—for example, in optogenetics, the fact that mammalian neurons had sufficient endogenous levels of all-trans-retinal, the chemical co-factor required for microbial opsin function, was not predictable in advance, but had to be discovered by trying it out (Boyden et al., 2005). For a tool to be simple, there is sometimes an element of discovery involved. Had optogenetics required chemical co-factors to be administered into the living mammalian brain before optical illumination, this would have made many neuroscience experiments more complex and potentially less robust. It would also have made the tool harder to disseminate. Note well, the fact that an innovation in its final form may be simple does not mean that the path to inventing it or realizing it is always simple. But if in the early stages of a project we aim for "constructive failures"—failures that teach us something new, perhaps because they show us something that has never been seen before—they may show us a better path to a truly impactful solution, orthogonal to what everyone else is doing. From such wisdom, one can then sometimes move on to a simple, powerful design, avoiding the complex in favor of the simple, precise, and/or robust.

Sometimes there is a tension in biomedicine: do we pursue basic science for curiosity's sake, or do we focus on translational work? Of course, both basic science and translation are important. One way to connect the two is to consider a third path: what if we could accelerate our understanding of the basic science, via new technology, so that we could reduce the risk of translational work and magnify the impact of the latter? To return to the moonshot analogy: suppose it's the year 1600. If someone were to propose to land on the moon back then, would it have been possible for scientists to resist the temptation to just go for it, and instead to take a step back and accelerate the invention of calculus, work out the laws of mechanics, and begin experiments on aerodynamics? In the end, of course, the necessary progress occurred, but over a period of centuries. Is there a way of approaching the problem that would have accelerated this progress, achieving deliberately what the history of science has suggested was achieved by happenstance? Considering the path of science and the events, structures, and ways of thinking that drive scientific revolutions is beyond the scope of this essay. But perhaps the perspectives contained within provide a model for how, in neuroscience, we might think about ways to take a step back to reduce the risk of, and accelerate, our path.

ACKNOWLEDGMENTS

E.S.B. was funded by John Doerr, Open Philanthropy, the HHMI-Simons Faculty Scholars Award, and the NIH Director's Pioneer Award. We thank John Doerr, Jini Kim, Tom Kalil, Patrick Collison, Michael Nielsen, Laura Deming, Joi Ito, Tom Skalak, and Clara Wu Tsai for feedback on the ideas here contained.

REFERENCES

Barrangou, R., Fremaux, C., Deveau, H., Richards, M., Boyaval, P., Moineau, S., Romero, D.A., and Horvath, P. (2007). CRISPR provides acquired resistance against viruses in prokaryotes. Science *315*, 1709–1712.

Boyden, E.S. (2011). A history of optogenetics: the development of tools for controlling brain circuits with light. F1000 Biol. Rep. *3*, 11.

Boyden, E.S., Zhang, F., Bamberg, E., Nagel, G., and Deisseroth, K. (2005). Millisecond-timescale, genetically targeted optical control of neural activity. Nat. Neurosci. *8*, 1263–1268.

Mahn, M., Prigge, M., Ron, S., Levy, R., and Yizhar, O. (2016). Biophysical constraints of optogenetic inhibition at presynaptic terminals. Nat. Neurosci. *19*, 554–556.

Marblestone, A.H., and Boyden, E.S. (2014). Designing tools for assumption-proof brain mapping. Neuron *83*, 1239–1241.

Miller, G. (2010). Is pharma running out of brainy ideas? Science *329*, 502–504.

Prinz, A.A., Bucher, D., and Marder, E. (2004). Similar network activity from disparate circuit parameters. Nat. Neurosci. *7*, 1345–1352.

Raimondo, J.V., Kay, L., Ellender, T.J., and Akerman, C.J. (2012). Optogenetic silencing strategies differ in their effects on inhibitory synaptic transmission. Nat. Neurosci. *15*, 1102–1104.

Schulz, D.J., Goaillard, J.M., and Marder, E. (2006). Variable channel expression in identified single and electrically coupled neurons in different animals. Nat. Neurosci. *9*, 356–362.

Wassie, A.T., Zhao, Y., and Boyden, E.S. (2019). Expansion microscopy: principles and uses in biological research. Nat. Methods *16*, 33–41. Published online December 20, 2018.