

From “Poke and Hope” to “Predict and Treat”

Philip Sabes

CEO, Arbor Neuroscience · September 2026

A young woman is prescribed an antidepressant. Eight weeks later, it hasn't helped, so she tries another. At first she feels better, but after a few months the depression is back. The slow cycle continues for years. New drugs, adjustments, side effects, therapies, and waiting. Each time her clinician learns only that the last thing didn't work, not why.

Eventually she's offered neuromodulation, a treatment that delivers energy into the brain to change the patterns of activity in a target circuit. The same approach can treat OCD and Parkinson's; the difference is where you aim. For her, though, the process is too familiar: intervene, wait, come back.

But neuromodulation doesn't have to work this way. Instruments that deliver energy could also measure what the circuit is doing. We could follow the causal chain: the intervention, the circuit's response, the change in symptoms. Each adjustment could then inform the next – for her, and for the next patient.

We're building Arbor Neuroscience to make that possible.

Arbor is a Focused Research Organization, part of Convergent Research, and the new name for Forest Neurotech. Forest built an instrument to measure activity across the brain. Arbor will use it to find out what neuromodulation actually does – and to make it predictable.

From tools to understanding

By 2017, brain-computer interfaces already worked. People with paralysis were controlling computers and robotic arms, but only when hooked up to lab equipment. The bottleneck was hardware. That's why I left my UCSF faculty position that year to co-found Neuralink, where we scaled work that had started in my lab: a robot that sews flexible, micron-scale electrode threads into cortex.

Most of the excitement around brain interfacing was, and still is, focused on BCI devices that connect people to something else: motor or sensory prostheses for now, and in the longer term to each other, to AI, or to some virtual future.

But I became drawn the other way, to devices that connect the brain to itself: to understand what a circuit is doing, and how to change it. The scientific questions are just as deep, and people need the answers now. Here, the bottleneck is understanding.

We don't know how to fit it to the patient

For millions of people with depression, addiction, or chronic pain, existing treatments haven't worked. But their symptoms arise from patterns of activity in neural circuits, and some of those patterns can be changed by delivering energy to the right circuits. The same tool, aimed at different circuits, can treat different diseases.

Deep brain stimulation has been used in more than 300,000 people worldwide¹. Transcranial magnetic stimulation (TMS) has been approved for depression since 2008, is delivered in outpatient offices, and is routinely reimbursed. Focused ultrasound can reach deep targets with millimeter-scale precision through the skull. And when these therapies work, they work astonishingly well: tremor disappears, or decades-old depression lifts.

Yet most people who might benefit never receive these treatments. Neuromodulation is approved for only a handful of diagnoses, and even where it's approved, it's a last-line treatment. After a standard course of TMS, for example, roughly a third of patients see their depression clear². But we can't say which third in advance. We don't know how to choose the right treatment for a particular patient, or how to refine it while it's being delivered. And when it fails, we can't explain why.

Two problems underlie this.

Patients are heterogeneous. A diagnosis is a category, agreed on by convention. Symptoms are what a particular person actually experiences, and two people with the same diagnosis can look very different. It's the symptoms that likely map onto circuits, not the category.

But brain targets are chosen from the diagnosis. This is true even in trials that personalize targeting from a patient's own anatomy. For example, Mayberg and colleagues have identified a specific intersection of fiber pathways, mapped in each patient, as a DBS target for depression³. That's real anatomical personalization, but the target itself has been linked to particular symptoms of depression, low mood and a slowing of movement and thought⁴. Two people with the same diagnosis might differ on other symptom axes, such as anxiety, rumination, or an inability to experience pleasure, and these likely arise from different circuits⁵. With better symptom-to-circuit maps, we could choose the target from the patient rather than the diagnosis⁶.

And we can't see what treatment is doing. Intervention parameters are set from clinician experience and heuristics, then refined against what can be observed in the clinic: a thumb twitch, a side effect, a patient's report. Those are valid signals, but none is a direct readout of circuit response. This doesn't matter if symptom changes are visible in seconds, as in tremor. But for most symptoms, treatment effects are not immediate. So treatment begins: writing to the brain with little or no way to read the circuit response, and waiting weeks or months until the patient returns for adjustment. The cycle repeats until something works or the treatment is abandoned. Poke and hope.

The knowledge has to come first

I spent a couple of years trying to solve this with a startup company. The premise of Integral Neurotechnologies was that a sufficiently capable clinical device could break this bottleneck: give ordinary patients a high-precision, bidirectional deep brain interface, so that every treatment would become a chance to learn how interventions change circuits. We ran into what looked like a fundraising problem, but fundraising was just the symptom. Neither we nor anyone else had the data to settle key design decisions, such as the number and spatial resolution of the electrodes, or how much current is needed. We were stuck in a fundamental chicken-and-egg problem: no evidence without the device, no device without the evidence.

So what's most needed now isn't a commercial device. It's knowing enough about what interventions do that we can choose a patient's treatment, refine it, and eventually specify the device that delivers it. To get there, we can build and iterate on research instruments. That's what breaks the chicken-and-egg.

What we're building

Building that knowledge means building models, and building the experimental platforms and long-term datasets those models require.

Two kinds of models are needed.

The first is a model that lets us see what a treatment is doing. We need to see enough of the causal chain to find signals that both move when we intervene and tell us where symptoms are heading. Those signals become the feedback variables that drive the model: given where a patient's circuits are now and what we do next, what happens to their symptoms?

The second is a map from symptoms to circuits, which tells us where to aim in the first place. A great deal of the raw material already exists either in public datasets or in individual labs and clinics: functional and anatomical imaging, brain injury and clinical

outcome data, and gene expression maps. We'll start with the available data and improve the map as new data arrives, from partners with clinical and interventional data to share and from our own experiments aimed at gaps the model reveals. The goal is to predict, for a particular person with particular symptoms, which circuits need intervention and how to intervene. Because symptoms cut across diagnostic categories, a map built this way may extend to conditions we never studied directly.

Together, these point toward treatment that fits the patient. First, a patient's assessment identifies the dimensions of their symptoms, and the treatment plan is determined by a symptom-to-circuit mapping rather than a diagnostic label. Then, whether the therapy is an implanted device, a wearable, or a clinic-based system, its parameters are tuned from the circuit readout rather than from what the patient can report weeks later. When something doesn't work, the circuit response tells you where it failed.

Where we start

We already have a powerful asset. Forest Neurotech, founded by Sumner Norman, Tyson Aflalo, and Will Biederman, built a wearable device for functional ultrasound imaging of the brain, the Forest 1⁷. It measures localized changes in blood flow driven by neural activity, the same class of signal as fMRI, but at better spatial and temporal resolution, and it reaches deep brain structures. The next version will also deliver focused ultrasound neuromodulation: a deep, bidirectional window onto the human brain.

These are research instruments, and they work in people who have had a section of skull replaced by an implant that ultrasound can see through – usually after surgery for trauma, stroke, or a tumor. Most implants are made of ultrasound-opaque materials, but we estimate that tens of thousands of people in the US have an implant that permits functional ultrasound imaging. Many of them live with problems of memory, attention, mood, and sleep – among the symptoms that are hardest to treat in psychiatric and neurological disease.

So our first program is to treat those symptoms while watching what the stimulation actually does, in the same person, over months, first in the clinic and eventually at home. What makes this approach work is not a large population, but the depth of data from each participant: seeing what each intervention does, and using that to choose the next. A first study is underway with the Barking, Havering and Redbridge University Hospitals NHS Trust and the University of Plymouth, funded by ARIA, testing whether we can shift mood-related circuit activity.

Building these models will take more than one tool or one study population. In parallel with the human work, we'll run longitudinal animal studies, where ultrasound and implanted electrodes can be combined to stimulate and record across a wider range of spatial and temporal scales.

The science determines what comes next

By the end of this first phase we intend to have an integrated platform that can stimulate a circuit and watch the response across the brain; longitudinal datasets, in people and in animals; a first symptom-to-circuit map, built from those data together with the datasets that already exist; and demonstrations of the loop closing, with a model choosing where and how to intervene, a measurement showing what that did, and the next choice following from the measurement.

We don't expect to do this all ourselves. Our clinical and animal work will be done with hospital and academic collaborators. We also expect to work with groups that have already collected large datasets, and to design new studies with partners – either sharing our tools or using theirs – to generate data neither of us could produce alone. What Arbor does is assemble a multidisciplinary group to work exclusively on this problem: building the instruments, generating the causal, longitudinal data these models need, and building the models themselves.

The FRO model exists for problems like this: too integrated for a single lab, too early for a company. FROs have typically been time-limited. We've sized a first phase to the core funding in place, and it's enough to tell us what's working, what isn't, and what to do next. How long Arbor runs after that depends on what's still missing, and whether Arbor is the right place to build it. Arbor is scoped to the problem rather than to a product, and it will evolve as the science does.

For now, our goals are clinical, but the implications are larger. Most of us already use coarse tools – coffee, alcohol, or pills – to change our mental states. The circuits that underlie these states are identifiable, and we think it will eventually be possible to reach them precisely and intentionally. We start with disease because the problems are urgent and solvable, and because we expect it's the fastest path to the knowledge that everything else will require.

Come build it with us

Arbor has a new name and a refocused mission, and now we're hiring. We need people to design and run the studies, to build symptom-to-circuit maps and turn data into prediction, to build the instruments, and to build the organization.

This is a place to do ambitious science without a product roadmap deciding what gets built. If that's the work you want, let's talk.

2. Hutton et al., 2023 – <https://doi.org/10.1016/j.brs.2023.10.001>
3. Alagapan et al., 2023 – <https://doi.org/10.1038/s41586-023-06541-3>
4. Figuee et al., 2025 – <https://doi.org/10.1016/j.neurot.2025.e00637>
5. Williams, 2016 – [https://doi.org/10.1016/S2215-0366\(15\)00579-9](https://doi.org/10.1016/S2215-0366(15)00579-9)
6. Taylor et al., 2026 – <https://doi.org/10.1038/s41380-026-03535-1>
7. Forest 1 details – arborneuroscience.org/https://www.arborneuroscience.org/technology.html