

SEIZURES, IN THEORY: Computational Neuroscience and Epilepsy

By Katharine Miller

With the headline “Easing Epilepsy With Battery Power,” the *New York Times* on March 24, 2014, described an implantable device for controlling epileptic seizures in patients who do not respond to medication. Developed by NeuroPace and recently approved by the FDA, the RNS[®] System is trained to recognize an individual patient’s seizure pattern and then deliver electrical stimulation to stop seizures before they can take off.

For some patients, the device is a godsend, yet it works for only a subset of patients and even for those, its effectiveness is limited: “Fifty-five percent of patients experienced a 50 percent or greater reduction in seizures two years post implant,” the

company’s press release declared, and most will continue to take medication. While the NeuroPace RNS[®] System could certainly be considered a victory for computation (it uses machine learning and could benefit an estimated 400,000 Americans), there’s no question that better treatments are still needed. In recent years, even as medicines and surgical techniques have reduced seizure frequency for roughly 80 percent of patients with epilepsy, many people remain treatment-resistant.

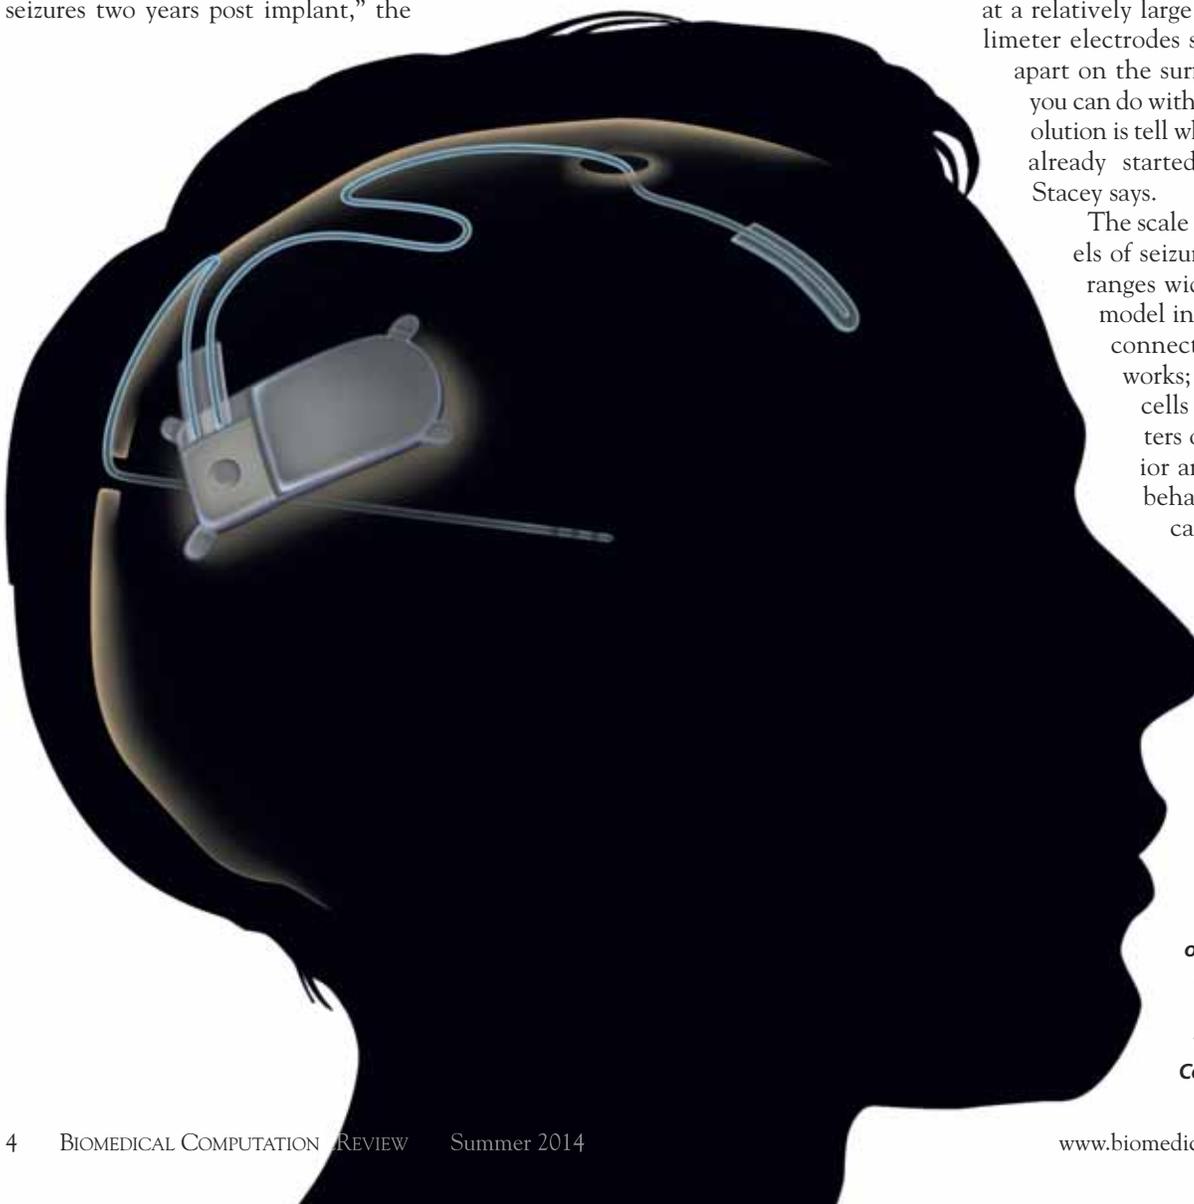
During a seizure, voltage activity in the brain becomes synchronous. Interconnected neurons go from a state of independent pro-

cessing to being connected in a massive cascade, says **William Stacey, MD, PhD**, assistant professor of neurology and biomedical engineering at the University of Michigan. It’s what engineers would call a feed forward loop: Because one neuron fires, another one does until they are all firing together. “What makes a system in its normal behavior suddenly go into this self-sustaining avalanche?” Stacey asks. It’s a question that has long puzzled clinicians and researchers alike.

Whether computational approaches can provide a helpful answer will require a bridging of the gap between the scales of clinical and computational research, Stacey says. Clinicians measure electrical activity at a relatively large scale—using four-millimeter electrodes spaced one centimeter apart on the surface of the brain. “All you can do with that type of spatial resolution is tell when an area of brain has already started to have a seizure,” Stacey says.

The scale of computational models of seizure, on the other hand, ranges widely. Some researchers model individual cells and then connect them into small networks; others describe similar cells using lumped parameters of their average behavior and then simulate their behavior to see if it replicates reality; still others create mathematical models of dynamic networks across

Neuropace recently announced FDA approval of its RNS[®] System for detecting seizures and delivering deep brain stimulation (DBS) to stop them. The device is implanted in the cranium with either one or two leads for detecting the seizure and providing neurostimulation to the targeted brain areas. Courtesy of NeuroPace.



the entire brain.

Many of these models are difficult to validate experimentally. That's because there's currently no way to know if the connections in a physiological model are accurate and it's not possible to measure the network dynamics across the entire human brain, Stacey says. But that is changing. "We stand at the cusp of a very rich time in unraveling the dynamics of seizures," he says. Computational models are getting bigger and brain recordings are getting smaller. "As soon as they are at the same level—and we're close—then everything on the computer can be validated and we'll be able to play with the model to produce predictions."

The Devil in the Details

Modeling individual cells and connecting them into networks to study what makes them go haywire in epilepsy is one appealing approach, Stacey says. "It's a very intriguing problem for people interested in dynamics," he says. "And it allows us to model the brain's actual physiology, though it can be difficult to validate that the neuronal connections in such models are accurate."

It's also very easy to make a network have a seizure using a model of a cell. In a normal brain, negative feedback keeps firing neurons from getting out of control. "It's very easy to break that feedback in a computer model," Stacey notes. "It makes you wonder why everybody doesn't have seizures."

Yet researchers who build physiologically detailed network models and simulations of epilepsy say they are valuable for generating hypotheses that get tested in the lab and then iterated back through the model. Theoden Netoff, PhD, associate professor of biomedical engineering at the University of Minnesota, is one such researcher. He wondered whether computer models might provide a better understanding of how and why deep brain stimulation (DBS), which is sometimes used to treat epilepsy by sending regularly scheduled electrical energy to the brain, stops or shortens some seizures but not others. The team was particularly focused on determining whether changing the frequency of DBS

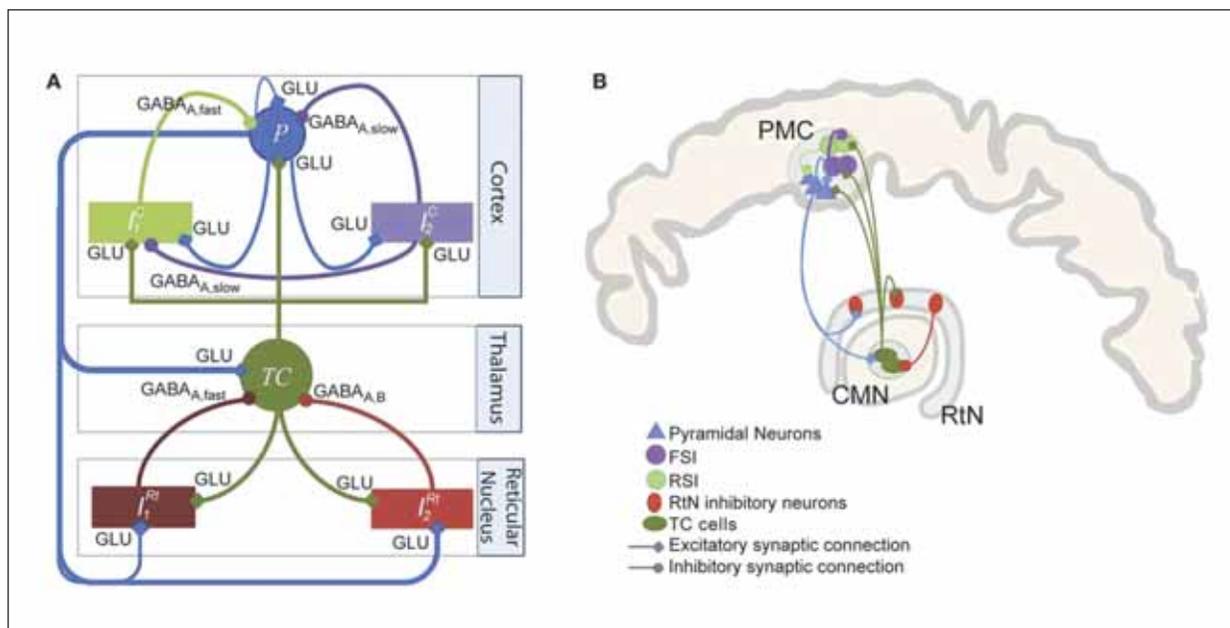
would shorten (or lengthen) the duration of so-called tonic-clonic seizures, in which a person first goes rigid (the tonic phase) and then starts to jerk uncontrollably (the clonic phase).

Netoff and his colleagues used a standardized computer model of an individual brain cell to build a 3,000-cell excitatory neuronal network that exhibits network statistics not unlike those in a rat visual cortex. The network is also capable of epileptic activity (it can both synchronize and desynchronize). They then added various frequency pulses of stimulation to simulate the model network's response to DBS. The result: The model predicts that DBS frequency affects the duration of the different phases of seizure in a way that is directly related to the neuron-firing rate and the level of synchronicity. For example, during the tonic phase, using a DBS frequency that matched the neuronal firing rate brought the tonic phase to a close more rapidly, while a frequency slightly below the neuronal firing rate shortened the clonic phase.

Indeed, in a computer simulation, when an adaptive algorithm controlled the frequency of DBS, it was more effective in truncating seizures. Netoff is currently running experiments to test these predictions.

Lumping It

Because it is difficult to use detailed models to study the extensive brain regions involved in epilepsy, some researchers are using lumped parameter models (also known as macroscopic models or neural mass models), that use average behaviors of particular cell types. Fabrice Wendling, PhD, research scientist at Laboratoire Traitement du Signal et de L'Image, Université de Rennes 1, in Rennes, France, who has used this approach for some time, noticed that these models couldn't recreate one of the signatures of epilepsy: high-frequency oscillations known as fast ripples. Concerned that his macroscopic models might be missing something, Wendling set about decoding the parameters of the

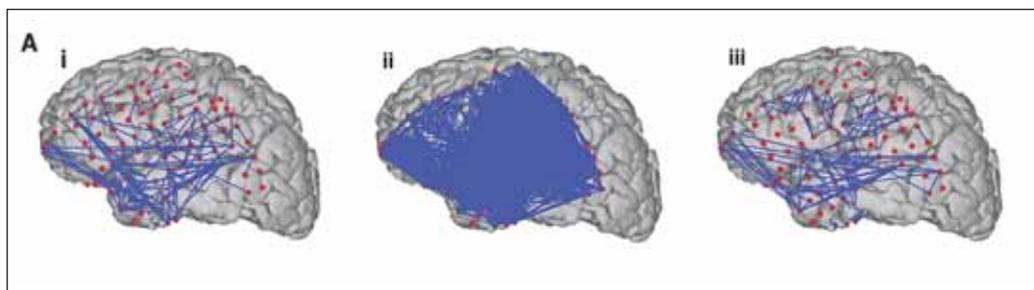


Wendling uses lumped parameter models to simulate seizures in the brain. For example, in this model (A) of the thalamocortical loop, three compartments (cortical, thalamic and reticular) each contain relevant subpopulations of neurons connected in a way that is compatible with brain connectivity patterns (B) inferred from the literature (PMC = pre-motor cortex; RtN = Reticular Nucleus; CMN = centromedian nucleus of the thalamus). The model then simulates the average behavior of those regions rather than the detailed behavior of each neuron. Reprinted from Mina F, et al., *Modulation of epileptic activity by deep brain stimulation: a model-based study of frequency-dependent effects*, *Frontiers in Computational Neuroscience*, 7:94 (2013).

The work, which was published in *Frontiers in Neural Circuits* in February 2013, suggests that a closed-loop feedback system that can adjust DBS frequency in response to changes in the neuron-firing rate would offer greater control over seizure duration.

macroscopic model by relating them to the parameters in more detailed models. By developing a detailed model for the same system that he was modeling macroscopically, he was able to see what lay behind the macroscopic model and understand why it

couldn't exhibit fast ripples. Essentially, such ripples develop in the detailed model when specific sets of pyramidal neurons are weakly synchronized. "It makes sense that the lumped model can't see the fast ripples because it assumes the activity in each subpopulation of cells is highly synchronized," Wendling says. When the researchers increase excitability in both models, however, the same sharp epileptic spikes appear. "Once both models can generate the same type of epileptic activity (for example, epileptic spikes) then it's much easier to see which parameters at the detailed level correspond to the macroscopic



Kramer and his colleagues construct functional networks of brain dynamics during seizure. In part A of this graphic, the red dots represent the locations of electrodes on the brain, with blue lines showing coupled firing between various brain locations (i) just before seizure; (ii) at seizure initiation; and (iii) mid-seizure. The density of lines suggests that coupling is high at initiation but then becomes fractured in the middle of the seizure. Part B displays these same networks among 100 electrodes in one patient's brain every five seconds, with the electrodes arrayed around the edges of circles. The seizure period is shaded pink and shows coupling as the seizure starts, followed by decoupling in the middle of the seizure, and intense coupling again at the end. Reprinted from Kramer MA, et al., *Coalescence and Fragmentation of Cortical Networks during Focal Seizures*, *J. Neuroscience* 30(30:10076-10085 (2010).



parameters," Wendling says. The work was published in the *European Journal of Neuroscience* in 2012.

Since that time, Wendling has used his macroscopic model to help understand the relationship between DBS frequency and treatment success. For example, his team created a model of the thalamocortical net-

work that includes multiple cell types in several compartments of the brain. They then trained the network to reproduce a particular patient's EEG recordings during seizure, and simulated various frequencies of DBS on the network. These simulations reproduced the patient's unusual and interesting response to DBS: His seizures typically stopped in response to low and high but not intermediate frequency stimulation. The work, reported in July 2013 in *Frontiers in Computational Neuroscience*, posits a possible explanation based on what happened in the model—low-frequency stimulation inhibited the feed-forward nature of the patient's seizure

while high frequency stimulation inhibited thalamic output. Intermediate frequency stimulation, on the other hand, just kept the epileptic dynamics going.

Wendling says he's optimistic that DBS will prove valuable as a therapy for epilepsy once there's a better understanding of how to use it optimally. And to gain that understanding, he says both detailed and macroscopic approaches will be

useful. "They are complementary and necessary," Wendling says. "What you can do with one approach you cannot do with the other and vice versa."

The Whole Enchilada

Some researchers take an even broader view of the network dynamics in epilepsy. They look at the entire system rather than one piece of it. **Mark Kramer, PhD**, assistant professor of mathematics and statistics at Boston University, for example, looks at seizure dynamics across the entire brain during the duration of the seizure. He then creates computer models to connect data to mechanisms. The goal: to help surgeons decide which part of the brain to cut out; or define optimal targets for stimulation by a device such as the one made by NeuroPace.

In work published in 2010, Kramer and his colleagues used electrocorticogram data—electrical activity measured directly on the surface of the brain's cortex—to build functional networks of the coupling and decoupling of brain areas during the course of a seizure. These networks reveal more coupling at the beginning of a seizure, less in the middle, and then more again at the end, suggesting that seizures are not simply hypersynchronous events but instead exhibit more subtle dynamics. A greater understanding of the coupling and decoupling of brain areas during seizure might suggest ways to

prevent the seizure from spreading across the brain by surgically firewalling certain connections, Kramer suggests. “Ideally, network tools could help us refine what surgeons cut out,” he says. “That’s one of our goals. We’re not there yet.”

Kramer is also interested in how seizures end. Recent research suggests that synchrony increases just before the seizure ends. “It gets more and more similar and then the brain shuts down,” Kramer says. He hypothesizes that seizures end because they cross some kind of tipping point or critical transition. Moreover, perhaps when seizures keep going and going (a condition called status epilepticus), the brain’s rhythmic activity tries to slow but then speeds up again, repeatedly approaching an ending but not quite making it. “What was nice about the hypothesis was that it led to specific testable measures,” Kramer says. The model simulation of the tipping point theory replicated the expected brain dynamics, with the same features of rhythmic slowing, increased coupling, and flickering between seizure and non-seizure states that had been observed in functional networks during the transition. The work was published in *Proceedings of the National Academy of Sciences (PNAS)* in 2012.

“It’s a different way to think about seizure termination, focusing on the mathematical mechanisms rather than biophysiology,” Kramer says. It’s possible, for example, that the mathematical constraints might help rule out other models that don’t fit the predicted pattern.

Stacey took an even broader approach to the tipping point question in a recent collaboration with **Viktor Jirsa** (physics) and **Christophe Bernard** (neuroscience), both at the Université de Marseille in France. They found that seizure dynamics in any species can be described by a common set of abstract mathematical equations. They validated the equations with data from humans, monkeys, rats, mice, zebrafish, and flies. This work, to be published in the journal *Brain* in 2014 (in press), suggests that seizures are, in fact, among “the normal repertoire of brain activities,” Stacey says. Moreover, they suggest that treatments should be directed toward altering dynamical properties of the brain rather than specific pathways.

A Question of Control

Some researchers are betting that work like Kramer’s and Stacey’s will yield a greater understanding of seizure dynamics that could eventually lead not only to better

treatments for epilepsy, but even to a cure. **Paul Carney, MD**, professor of pediatric neurology at the University of Florida College of Medicine and director of the University’s Center Of Excellence for Epilepsy Research and Comprehensive Pediatric

THE PROBLEM OF PREDICTION

The NeuroPace RNS® device relies on seizure detection—spotting a seizure just as it’s starting, typically only seconds before onset. By that point, Stacey says, the seizure is already underway. Prediction, along with the possibility of true prevention, has to occur sooner. “Is something already burning or is there just heat and smoke?” Stacey says. “It’s a lot easier to put out a fire before the flames start.”

People with epilepsy would welcome a way to know when a seizure is coming, Carney says. “If I told you that you’d have a seizure today at 1 p.m., you could arrange your life around it.”

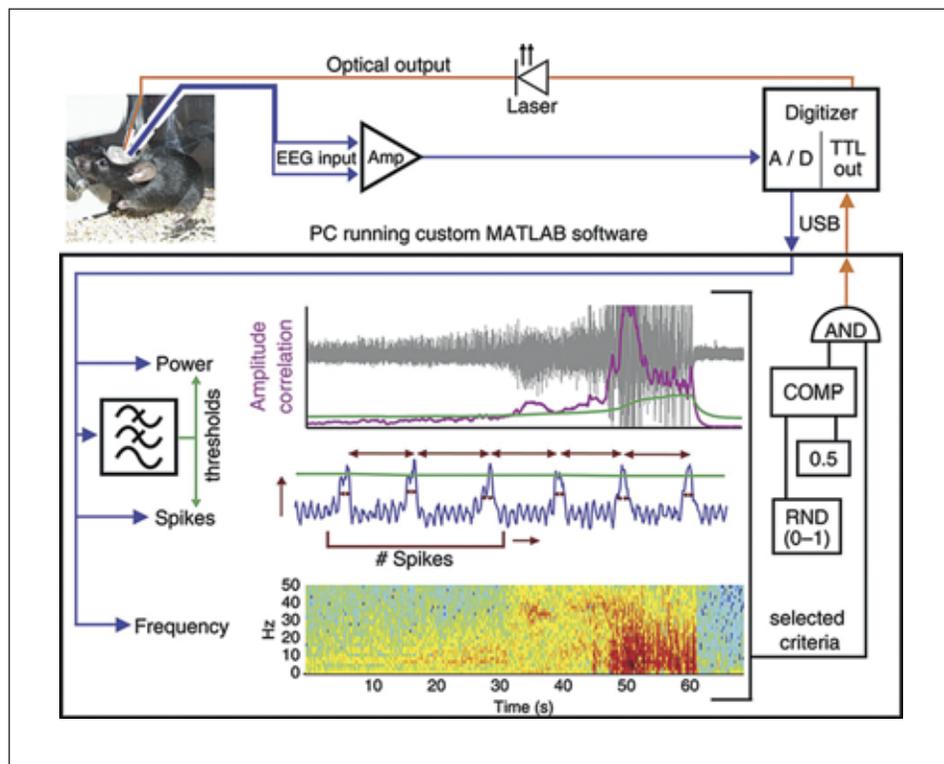
About 20 years ago, researchers got very excited about using computation to predict seizures well before they start. They set about looking for patterns in the electroencephalograms (EEGs) of people having seizures to see if they could predict seizure onset well before it actually starts, and at least a minute ahead of time. An initial flurry of promising algorithms didn’t pan out because they used flawed statistical methods.

Then in 2007, a company called Neurovista published an appropriate statistical framework for prediction. They went on to develop an implantable prediction device that can show patients, on a handheld device, whether likelihood of a seizure is low (blue light), moderate (white light) or high (red light) in the next few hours. A 2013 paper in *Lancet Neurology* reported that for 8 out of 11 patients tested, the device predicted seizures accurately between 56 and 100 percent of the time. For investors and the FDA, that apparently wasn’t enough of a home run, Stacey says. Funding dried up and the future of the prediction device is uncertain.

At this point, Carney says, “The field is back to trying to understand what’s going on rather than trying to predict seizures based on what we know now.”

Epilepsy is especially optimistic about control theory, an approach borrowed from finance, weather, and airplane cruise control or autopilot. “The airplane makes subtle adjustments as you fly,” Carney says. In the brain, he says, there’s also a controller that

applies gentle adjustments to keep things within a certain dynamical range. “Can we take advantage of those intrinsic mechanisms to prevent seizures?” he wonders. Perhaps as a seizure is ramping up, there might be a point when intervention (turning on a stimulator or taking a medication) would keep the brain out of the danger zone. “Rather than responding to the hurricane,



In an optogenetic closed-loop system for stopping seizures, Krook-Magnusson and her colleagues fed EEG signals coming from the mouse brain (blue arrows) into real-time seizure detection software containing several possible algorithms for recognizing changes in features such as signal power, spikes, or frequency. The software was tuned to recognize certain thresholds for seizure in each mouse. Once detected, the experimental protocol called for administration of light (orange arrows) in half of the events in random fashion. The result: Optogenetic control reduced the frequency and duration of seizures in the mice. Reprinted with permission from Krook-Magnusson E, et al., On-demand optogenetic control of spontaneous seizures in temporal lobe epilepsy, Nature Communications 4:1376 (2013).

you break it up in advance.”

Unlike DBS, which Carney describes as a black box, control theorists would start by figuring out what features in the brain can be acted on to provide the necessary control.

One approach that is already showing great potential is optogenetics: Using a pulse of light to activate genes involved in

epilepsy. A device for detecting and then automatically and optogenetically stopping spontaneous temporal lobe seizures recently proved effective in transgenic mice. The research team, led by **Esther Krook-Magnusson, PhD**, postdoctoral fellow in the department of anatomy and neurobiology at the University of California, Irvine, used two breeds of mice, each designed to express light-sensitive proteins that would either inhibit certain excitatory brain cells or activate the power of inhibitory (GABAergic) cells. They then implanted the mice with electrodes for detecting seizures and an optical fiber for delivering light to the target cells. First, the detector had to be trained on the specific mouse’s seizure data, a not insignificant hurdle because temporal lobe seizures are tricky to detect. Detection also had to be fast, because it would occur only seconds before a seizure would otherwise start. “Computations have to be done efficiently and at an appropriate time scale,” Krook-Magnusson says.

For both breeds of mice, the device reduced seizures and seizure duration with no obvious side effects. “Since it is ‘on demand’ rather than continuous treatment, we’re not interrupting good network activity,” Krook-Magnusson says.

The work offers a tool for understanding the roles of specific cell types in causing and stopping seizures, and might lead to new pharmacologic approaches, she says. There’s also the possibility of using optogenetics to treat humans, although currently the idea of transfecting a human brain with a virus carrying the necessary genes is out of favor, Carney notes. “We have not been able to convince reviewers that optogenetics has a clinical future,” he says.

But the epilepsy field’s interest in control theory goes beyond optogenetics. In his 2012 book, *Neural Control Engineering*, **Steven Schiff, MD, PhD**, director of the Penn State Center for Neural Engineering, and a pioneer of using computational neuroscience to study epilepsy, proposes applying non-linear control theory to models of epilepsy at all scales—neuronal, lumped, and whole-brain—and paints a picture of where control theory could take the field.

According to Carney, Schiff’s interest in control theory reflects a shift in computational neuroscience away from a signal processing approach to epilepsy and toward more advanced dynamical modeling. “Ultimately we want prevention and cure,” Carney says. “We have treatments right now. But computational neuroscience lets you take experiments or results to the next level.” □