



$$V = IR$$

$$PV = nRT$$

3

f

4

7

6

7

8^{va} (un...)

...as I pondered weak
and weary
over lost Lenore,
quothe the Raven,
"Nevermore."

60-70% Si
15-20% Flux (alkali)
10-20% Mn, Ba, PbO₂,
Ca

Mustard wife

How Many Neurons Must One Man Have, Before You Call Him a Man?

by Timothy Marzullo and Hiram Parikh

We all decided to tell our stories, but one by one, and Stan was first. "We've got a long way to go," preambled Dean, "and so you must take every indulgence and deal with every single detail you can bring to mind—and still it won't all be told. Easy, easy," he cautioned Stan, who began telling his story, "you've got to relax too." Stan swung into his life story as we shot across the dark. He started with his experiences in France but to round out ever-growing difficulties he came back and started at the beginning with his boyhood in Denver.

—Jack Kerouac, *On the Road*

One of the great unsolved problems in science is deciphering the neural code for long-term memory or how the brain stores and represents information. Simply put: How are we that we are? How is long-term memory stored? How is it encoded? How is it lost? Can we hope to build a machine to do what we do naturally? This is an introduction to the both ancient and contemporary problem of the neural code and we explore a series of thought experiments examining how the brain encodes long-term memory and possible theoretical storage limits.

Contemporary neuroscience research largely rests on the assumption that the activity of neurons encodes memory. Typical memory research involves identifying what features of the electrical discharges of individual neurons correlate best with some parameter of a behavioral memory task. Such techniques first arose in the 1960s with Edward Evart's seminal work recording the

Digital Object Identifier 10.1109/MPOT.2008.930065

electrical discharges of individual neurons in the motor cortex of monkeys as they performed trained movements. Advances in technology over the past decades have allowed researchers to begin recording multiple neurons simultaneously, enabling access to the “black box” of the brain to test theories of memory. Experiments have ranged from psychological (memory recall tasks in humans), to the neurophysiological (testing rats/monkeys on memory tasks while recording brain activity), to the molecular (examining the relationships between proteins in the synapses during electrical activity excitation or depression). Compared to short-term and working memory, modeling and examining declarative long-term memory (facts and songs you know, autobiographical details, lovers you have had, where you were on 11 September 2001) remains elusive.

Obviously, we do not remember everything; even the most celebrated autistic savants do not store information like video recorders. All of us remember only a fraction of the information that impinges on our ears, nose, and eyes. The information in the brain can be represented by the following inequality:

Total information input > Retained Information.

Whatever the final nature of the information, it must be physically stored somewhere in the brain. Since the brain has a finite number of neurons and a limited amount of wiring, the brain’s capacity is therefore finite. In terms of information theory, let X be the message—the face, sound, or episode. The brain acts as the channel that encodes the message as Y (see Fig. 1).

The message to be encoded is X, and is represented by some encoding scheme as Y in the brain. This process is probabilistic, so by observing output Y, we can guess the most probable input sequence X. Dependent on the encoding process, the encoded message Y contains mutual information (I) about the input signal X and vice-versa. In simple terms, the channel capacity theorem imposes a theoretical limit based on the encoding scheme employed to represent the information X as Y in the brain. Thus, by choosing an appropriate encoding scheme, we can increase the channel capacity (C) and reconstruct the input sequence with a low error rate.

Encoding memory

How does the brain (the channel)

actually encode the information? As we see from Shannon’s equation, using different encoding schemes can utilize the capacity of the brain with varying degrees of efficiency.

Before discussing encoding schemes, let us begin with a brief primer on the physiology of neurons in the brain. The discharging of a neuron causes an impulse to travel down its axons. Such an impulse is called an action potential, or in neuroscientific jargon, a “spike,” to aptly describe its appearance on an oscilloscope. This is an “all-or-nothing” event and can be thought of as a digital pulse represented as a 1 when there is a spike and a 0 when there is not a spike. Thus, at a scale of 1 ms, a “spike train” of a single neuron can be represented as a binary string. For example, the string 110010001 denotes that four spikes occurred in a 9 ms period (see Fig. 2).

Cracking the neural code boils down to looking at the string of 0s and 1s that represents the spiking activity of individual neurons and deciphering exactly what these strings or spike trains represent. One problem lies in the fact that the strings are stochastic, with the same stimulus often represented by subtly different strings, making generalization difficult, if not impossible. Another problem is determining what kind of encoding scheme is being employed.

There are many theories regarding the encoding schemes on how spike trains represent memory. Here, for pedagogical reasons, we focus on two codes: ensemble codes and gnostic codes. The codes are relatively simple to understand, representing two extremes of memory coding schemes, and evidence exists in the neuroscience literature that the brain may indeed use both these codes (see sidebars “Spiking


Models of Memory” and “Calculating Memories Yourself”).

Ensemble codes can be thought of as an orchestral concert in which the individual instruments submit to the will of the ensemble, and, *as a combination*, encode the music. Gnostic codes are a more “solo” version of the idea. All neurons are *prima donnas* and each encodes one specific object. Such gnostic codes are popularly called “grandmother codes,” meaning that a specific neuron encodes the memory of your grandmother. See Fig. 3 for a representation of the two codes. Here we will examine the requirements of the gnostic code and the ensemble code to calculate the theoretical bounds for each code.

The lower limit: Using the ensemble code

In ensemble codes, individual neurons act in combination to encode events. The brain may indeed use such a code for reasons of efficiency. Consider three neurons that can encode a ‘1’ or a ‘0’. By themselves, the total encoding limit for the three neurons is three events, but together the three neurons can encode $2^3 = 8$ events. Given the diversity of the external world, and the variety of things we need to remember, such a code is indeed more attractive than the grandmother code. Consider all the photons that hit your retina, all the pressure waves that travel across your inner ear, all the tastes you encounter, and all the things you touch. It makes computational sense to reduce the metabolic cost of encoding memory states by using the more efficient ensemble code. Multiple studies have demonstrated that the brain indeed does so, especially in the context of sensory information.

For example, an ensemble code is



Shannon’s Information Capacity Theorem

$$C = \max_{p(x)} I(X; Y)$$

C = Channel capacity, I = Mutual information

Fig. 1 Shannon’s information capacity and brain storage.

used for the perception of sound location in space. Via the grandmother neuron hypothesis, individual neurons would encode for particular sound locations. The experimental evidence, however, does not support this. Auditory cortex neurons are broadly tuned for sound source location; one neuron alone only weakly encodes multiple different locations in space. Researchers at the University of Michigan in 1999 recorded from the auditory cortex of anesthetized cats as broadband sounds were presented in 3-D space. Rather than finding specific neurons encoding for space, mathematical analysis of the firing patterns (using artificial neural network techniques), revealed that spike counts and spike timing within an ensemble of neurons collectively encoded sound location. With a sample size of 128 recorded neurons, sound localization comparable to an animal's behavioral performance of localization was achieved. Thus, the brain appears to use an ensemble code for some functions.

Now let us consider the theoretical limits of an ensemble code, and whether

it can indeed be used to code memory. Consider the activity of an individual neuron over a given period, say 5 ms. Consider a temporal precision (the lower limit at which timing of a single

The question of even how to begin to measure the capacity of someone's long-term memory remains unanswered.

electrical discharge still matters) of 1 ms. Thus, in any spike train a variety of different positions exist for the spikes, and the total number of different states in a certain window can be calculated using Case 1 in Fig. 4. For example, in 5 ms a neuron could theoretically encode one of a potential 32 different states.

A spike, however, typically lasts 1 ms and is followed by a refractory period (the minimum time before the neuron is ready to fire another spike). If we incorporate the refractory period, then the equation becomes more complex since certain patterns are now impossible; for example, a pattern like "111." Case 1 is then expanded to Case 2 from Fig. 4.

In a 5 ms spike train with a 1 ms refractory period, the number of potential "states" reduces to 13.

To compute the number of states a *population* of neurons encodes, we simply use the number of states S possible for a single neuron with the given parameters and plug it into the equation at the bottom of Fig. 4. P is the number of neurons in the ensemble, we raise P to the number of states S , and we subtract one from this quantity to correct for the all-zero case when none of the neurons fire. Thus, for two neurons with 5 ms spike trains, temporal precisions of 1 ms, and a 1 ms refractory periods, the total number of potential states that can be encoded by the two neurons is 8,191.

Now let us use physiological constraints to calculate the number of states real neurons could potentially encode. Given inherent chemical and electrical delays in the brain, experiments have shown that for someone to be conscious of a stimulus takes about 200–300 ms. Also, the refractory period is typically 2 ms. Given a 300 ms window, a 1 ms precision, and a 2 ms refractory period, the number of independent states *per ensemble of neurons* increases dramatically. With one neuron, it is possible to encode $8,3269 * 10^{49}$ states, and with 2 neurons, the number expands to become more than the number of atoms in the universe (speculated to be $\sim 10^{81}$). If such a code is true, then what are the rest of the neurons in the brain doing? Are the rest of the neurons in our skulls simply serving as highly redundant modules, making most of the brain's neurons mute spectators?

The upper limit: Using grandmother neurons

Let us take another popular coding scheme, known in various forms as the grandmother hypothesis or gnostic neuron hypothesis. This hypothesis states memory is encoded by neurons that respond only to highly specific complex stimuli, such as the face or voice of your grandmother. Dr. Jerzy Konorski first suggested the grandmother code




Fig. 2 A trace of a neural discharge represented as a binary code. The neural trace is a real discharge from a neuron collected during experiments in our lab.

in 1967, and recent findings have demonstrated face selectivity in the inferior temporal cortex of macaque monkeys and humans. In 2005, researchers, quite dramatically and humorously, found neurons in humans that exclusively responded to the faces of celebrities Halle Berry and Jennifer Aniston. It was also reported that some neurons responded to multiple images, a surprisingly high level scheme. Specifically, one neuron responded to pictures of Jennifer Aniston and Brad Pitt standing together, along with other pictures of the Sydney Opera House. Such findings are very compelling in support of the grandmother neuron hypothesis.

Though such neurons may indeed be used to encode extremely important stimuli that need to be processed quickly (such as faces), we use the ploy of *reductio ad absurdum* to determine the upper limits of such a grandmother code that takes into account neurons that encode only combinations of specific objects or events. Take an author of this manuscript, Hirak Parikh. Consider reasonable values for the number of events he experiences a year that he may consciously remember: 30 people, 100 events, 50 songs, five health states, 20 locations, ten cars (or modes of travel), and ten emotions. Assuming that some neurons in the brain encode objects/stimuli only in combination, this yields $30 \times 100 \times 50 \times 5 \times 20 \times 10 \times 10 = 1.5 \times 10^9$ (or 1.5 billion) states. Assuming, conservatively, that he started remembering things at the age of five, this yields (22 years of memory \times 1.5 billion) *33 billion* states. Given that the neocortex contains \sim 22 billion neurons, there are insufficient neurons in the neocortex to store a truly combinatoric grandmother code.

This approximation seems intrinsically unrealistic, but the question is, as always, of resolution. How much of your grandma do you actually remember? And how many neurons encode such features of your grandma? We scoured the literature attempting to find some calculation of "how much" declarative memory an adult has, but research on the topic is thin. The question of how to even begin measuring the capacity of someone's long-term memory remains unsolved. We encourage the readers of this article to attempt to measure memory capacities through similar thought experiments. If you had to design an experiment determining the storage capacity of someone's long-term memory, how would you do it? Have someone tell you 5% of their knowledge, and multiply such data by 20? How could an individual

even determine whether they gave you 5% of his/her knowledge? Such methods are impractical, and, if anything, will probably serve only to underestimate the capacity of someone's memory.

The memory equation

Using our approximations above, if we use the ensemble code alone, we have an excess of neurons in the brain to encode memory; if we use the grandmother code alone, we have a dearth of neurons in the brain to encode memory. Yet, as we have discussed, experimental evidence exists for both kinds of codes in the brain. Between these two extreme kinds of codes, there may exist other schemes that are employed by the brain to represent memory, such as neuron-assembly type codes using oscillations or hybrid codes using aspects of both the ensemble code and the grandmother code. Thus, in order to account for the total number of neurons encoding memory that could fall within reasonable bounds of the amount of neurons in the human neocortex, we present the following equation:

$$\# \text{ of cells encoding memory, } N = E + G + \sum C$$

where N is the total number of neurons used to encode long-term memory, E is the number of ensemble neurons, G is the number of grandmother neurons, and $\sum C$ represents the sum of neurons using additional codes that we did not address here.


Daniel Tammet and exemplar representation of memories

Some people have capabilities of memory far beyond the normal person. Take the recent example of Daniel Tammet, who in 2004 memorized π to 22,514 decimal places. Memorizing π is a typical test of memory prowess. For example, remembering a number like 12233344445555... is easy. Claude E. Shannon demonstrated that if you can detect a pattern in a sequence, then you could exploit the pattern to reduce the number of bits required to represent it. But, since the sequence of digits in π is completely random, you cannot exploit any pattern of digits to achieve compres-



ENCODING GRANDMA MARZULLO

ENSEMBLE CODE



GRANDMOTHER CODE




Fig. 3 Ensemble code versus grandmother code.

sion, and you have to store all the digits in the sequence.

How then could the two encoding schemes we discussed above calculate π this way? With the grandmother code, you can imagine a chain of neurons, each encoding a single number of π . Thus, Daniel Tammet would need 22,514 grandmother neurons to encode his impressive feat. Using an ensemble code, with our values described above, would require only one or two neurons to encode 22,514 events. In either case, the number of neurons to memorize π , given our codes, is small (a cubic mm of neocortex contains ~50,000 neurons). So how is Daniel Tammet able to memorize

π while we cannot? What is his secret? Daniel Tammet has Asperger's syndrome (a form of autism), which causes him to obsess over counting and minutiae. In addition, he is synesthetic (his senses overlap multiple modalities, that is, numbers have colors, shapes, and personalities in his mind). He describes his method of memorizing π by imagining a long squiggly line, much like a mountain range, which represents the entire 22,514-digit approximation of π . Thus, perhaps Daniel is encoding π in multiple, overlapping, redundant ways, and this dramatically improves his memory skills beyond the typical human.

Spiking models of memory

Spike train loops, or engrams, though first introduced by the German biologist Richard Semon in "Die Mneme" in 1904, were brought to popularity as a mechanism of memory encoding by Karl Lashley and Donald Hebb in the early part of the 20th century. How are such spike loops maintained? The large body of LTP (long-term potentiation) research suggests modification of synaptic proteins preserves the maintenance of possible spike train loops. Karim Nader of McGill University (2003) notes succinctly:

"The current view of memory consolidation posits that short-term memory is independent of protein synthesis and that long-term memory depends on the synthesis of new proteins. Furthermore, it is a dialogue between the synapses and the proteins made in the neuron body that is at the heart of the plasticity process. The chemical view lies at one extreme, positing that specific memories are encoded at the level of the genome. At the other extreme are models that are independent of protein synthesis and can be explained only by changes at the co-activated synapses."

The notion that memory may be stored chemically began to arise in popularity with the discovery in the 1950s that hereditary information is encoded by nucleic acid sequences (DNA). This led to some rather notorious research in the 1960s and 1970s in which slurries of brain tissue from a trained rat were injected into a naïve rat with concurrent transfer of the memory. Though initially very promising (the memory peptide was named "scotophobin" for a fear of the dark), the effects turned out to be caused by the general effects of hormones rather than transferable memory traces.

Clearly though, changes in the molecular structure of neurons themselves need to occur for the encoding and retrieval of memory, for the fact that a brain can be temporarily shut down via cold trauma or anesthesia, without loss of long-term memory, proves that the brain's memories are represented as structural changes in the neurons themselves, not as spike train loops. Our pedagogical work here does not ignore the modeling of memories based on synaptic/dendritic modification; *spikes are the manifestation of the underlying structural changes*. Otherwise, the extreme counter view to a spiking model of memory storage would be that if every neuron in the brain spiked in exactly the same pattern, on two different occasions, the state of the brain may still be different because only the chemical modification/content of synapses matter for memory encoding. Thus we believe, in our work here, we can model memory as a spike code while still not rejecting that synaptic modifications are behind this process.

Leaky tank model of memory

As shown, the encoding schemes employed by the brain determine how efficiently information is stored via Shannon's information capacity theorem. The process of memory storage, however, is a dynamic one—we learn, but we also forget. Independent of the encoding scheme, we attempt here to model the dynamic nature of memory storage and loss. We propose to model this dynamic nature of memory by introducing our "leaky tank" model (Fig. 5). We assume the brain, regardless of the encoding scheme, has a finite capacity. At birth, we start *tabula rasa* (the blank slate), and as we learn and grow, our memory tank gradually fills with information. We know that each of us possesses a different capacity for absorbing new information, and this is modeled by the diameter of the input pipe (A) and the input rate ($d\eta/dt$) of information.

The more frustrating and annoying aspect of memory is the loss of the stored information. We model this by a small leak at the bottom of the tank. Some of us are better at retaining information, which is modeled by the diameter "a" of the leak pipe. We also make the assumption that the more we know, the more "pressure" there is to lose stored information. Indeed, it may be possible that the reason children have better memorization abilities than adults is simply because the tank is still filling and the leakage pressure is low.

Studies on learning have demonstrated that frequent recall and review is necessary to retain information. In a sense, we then are pouring the same information into the tank over and over again to keep the knowledge in storage. Indeed, a recent paper by Karpicke and Roediger in *Science* showed that frequent testing of material helps with better long-term memory.

In addition, two amnesias, retrograde (forgetting the past) and anterograde (inability to remember new things), can be explained by our model. In retrograde amnesia, the tank is completely emptied by a traumatic event, but equilibrium can be restored by the gradual addition of new information. In anterograde amnesia, the input pipe has been sealed completely by a traumatic event, and new information cannot enter the tank, as in the movie *Memento*, where the protagonist Leonard cannot remember new events since sustaining a head injury. However, we note that this model is a simplification of the complexities of am-

Calculating memories yourself

If you would like to calculate your own estimates of how many states a single neuron could encode using Case 2 from Figure 3, our brief MATLAB source code is below:

```
%State Calculation
%MATLAB code to calculate the different combinations of
states (S), taking in
%account the refractory period (R), precision (P) and
window size (W)
%Parameters

W=5; %W=window (ms)
P=1; %P=precision (ms)
R=1; %R=refractory period (ms)

N = round(W/P);
S = 0;

for r = 1:ceil(N/(R+1))
    S = S + nchoosek(N-(r-1)*R,r);
end
S = S + 1; % for the all-zero state
disp(S)
```

nesia; for instance, in retrograde amnesia the memories often return with time. In that case, access rather than storage of the memories is impaired.

See Fig. 5 for a mathematical description of the model. If these parameters can be measured by some experiment, we might be able to determine the limits of memory storage and loss.


Then they got down to business. They sat on the bed crosslegged and looked straight at each other. I slouched in a nearby chair and saw all of it. They began with an abstract thought, discussed it, reminded each other of another abstract point forgotten in the rush of events.

—Jack Kerouac,
On the Road

Prior speculations and the feasibility of theoretical limits

Our calculations are thought experiments and are intended to ask the ques-

tion: how much of a memory is actually represented in the brain? Sometimes our brains fool us in believing we know more about an event than we actually



Case 1: (w/o considering the refractory period)

of states = $S = 2^N = 2^5 = 32$ (N=bins)

For 'p' neurons, each with 's' states

Total # possible codes (T) are

$T = p^S - 1$ (∵ all-zero ensemble is the null-state)

∴ For N=2, S=32 states/neuron

If p=2 neurons

Total is: $T = 2^{32} - 1 \approx 4.2 \times 10^9$ combinations or codes

Case 2

If we consider the refractory period, then certain states are invalid.

The equation to calculate the possible states is:

$$S = 1 + \sum_{r=1}^{\lfloor \frac{N}{R+1} \rfloor} \binom{N-(r-1)R}{r} \quad \begin{matrix} N = \text{bins} \\ R = \text{refractory period} \end{matrix}$$

For R=1 and N=5


$$S = 1 + \sum_{r=1}^3 \binom{5-(r-1)}{r}$$

S = 13 states/neuron

If p=2 neurons

$T = 2^{13} - 1 = 8191$ combinations or codes

Fig. 4 Representing the ensemble code: Determining how many states can be encoded with two neurons in 5 ms, with or without the refractory period.



ODE for memory storage

$$M(t) = M_{in} - M_{out}$$

Total memory

$$\frac{dM}{dt} = \underbrace{A \cdot \frac{d\eta}{dt}}_{\text{total inflow}} - \underbrace{a \cdot \frac{df}{dt}}_{\text{total outflow}} \quad \text{--- ①}$$

Since, potential energy (PE) is converted to kinetic energy (KE)

$$mgh = \frac{1}{2} m v^2 \quad \text{--- ②}$$

$$v = \sqrt{2gh} \quad \text{--- ③}$$

(Torricelli's equation)

Here $\frac{df}{dt} = v$

and $\frac{dh \cdot D}{dt} = A \cdot \frac{d\eta}{dt} \quad \text{--- ④}$
 (D = diameter of tank)

Substituting ③, ④ in ①

Dynamic Memory Equation

$$\frac{dM}{dt} = \left(\frac{dh \cdot D}{dt} - a \sqrt{2gh} \right)$$

Fig. 5 The leaky tank model of memory

baboons, in ongoing experiments, can learn up to 3,500–5,000 pictures. The limit of human memory remains unknown (even how to measure such a thing is unknown). Is it even reasonable to think of long-term memory as some cold storage mechanism with an upper limit? Indeed, as Hunter pointed out in 1988, calculating “bits” of memory may be inherently flawed.

“One [meaning of memory] refers to the accumulated experience of a human being; it is the difference between a naïve youth and a wise elder. Human memory is not confined to the ability to recognize previously seen objects or accurately reproduce previously seen texts. It is our memories that allow us to recognize and refine categories of objects or events, to generate reasonable expectations and plans, or to focus at-

ention on salient aspects of the environment, to name but a few of the tasks memory subserves. In short, human memory is the ability to bring previous experiences to bear on new situations.”

Though contemporary scientists and engineers have made preliminary attempts, measuring and estimating memory capacity of humans in general, and the neocortex in particular, is a difficult problem to solve (or even define). For example, Jeff Hawkins and his colleagues at Numenta have designed a pattern recognition algorithm based on the structure of the neocortex. As approximations of human memory come increasingly closer to reality, we can begin to build real models to take advantage of the memory encoding capacity of the brain and perhaps create artificial machines that can replicate this at a much larger scale.

Finally, though we discussed spiking-based sources of memory encoding, we did not incorporate the role of brain oscillations, which have been described in recent literature as a mechanism of memory encoding through “binding” multiple levels of organization of the brain. Future modeling research may eventually shift focus from “spiking models” of memory to one where oscillating membrane potential voltages, including sub-threshold discharges, combined with spiking, may bring us closer to cracking the elusive code for long-term declarative memory.

Acknowledgments

The authors were idling in the lab and came up with these thought experiments to attempt to measure the limits of popular encoding schemes. This work was then submitted as an educational poster at the Society for Neuroscience conference in 2006. We would like to thank the many scientists and engineers at the conference that provided valuable suggestions after being tricked and trapped into listening to our presentation. We also thank Gregory Gage for helpful insight and production values at the same conference, as well as Miss Trigger for her careful copy editing of this manuscript. The illustrations in this work were drawn by HIRAK PARIKH.

Read more about it

- G. Buzsaki, *Rhythms of the Brain*. New York: Oxford Univ. Press, 2006.
- Y. Dudai, “How big is human memory, or on being just useful enough,” *Learning and Memory*, vol. 3, pp. 341–365, Jan. 1997.
- T.M. Cover and J.A. Thomas, *Elements of Information Theory*. New York: Wiley, 1991.
- J. Fagot and R.G. Cook, “Evidence for large long-term memory capacities in baboons and pigeons and its implications for learning and the evolution of cognition,” *Proc. Natl. Acad. Sci. USA*, vol. 103, no. 106, pp. 17564–17567, Nov. 2006.
- S. Finger, “The nature of the memory trace,” *Origins of Neuroscience*. New York: Oxford University Press. 1994.
- U. Frey and R.G. Morris, “Synaptic tagging and long-term potentiation,” *Nature*, vol. 385, no. 6616, pp. 533–536, Feb. 1997.
- S. Furukawa et al., “Coding of sound-source location by ensembles of

cortical neurons," *J. Neurosci.*, vol. 20, no. 3, pp. 1216–1228, Feb. 2000.

• C.G. Gross, "Genealogy of the 'grandmother cell,'" *Neuroscientist*, vol. 8, no. 5, pp. 512–518, Oct. 2002.

• D.O. Hebb, *The Organization of Behavior: A Neuropsychological Theory*. New York: Wiley-Interscience, 1949.

• L. Hunter, "Estimating human cognitive capacities: a response to Landauer," *Cogn. Sci.*, vol. 12, no. 2, pp. 287–291, Apr. 1988.

• J.D. Karpicke and H.L. Roediger III, "The critical importance of retrieval for learning," *Science*, vol. 319, no. 6865, pp. 966–968, Feb. 2008.

• J. Konorski, "Some new ideas concerning the physiological mechanisms of perception," *Acta Biol. Exp. (Warsz)*, vol. 27, no. 2, pp. 147–161, 1967.

• T.K. Landauer, "How much do people remember? Some estimates of the quantity of learned information in long-term memory," *Cogn. Sci.*, vol. 10, no. 4, pp. 477–493, Oct. 1986.

• A.R. Luria, *The Mind of a Mnemonist: A Little Book About a Vast Memory*. Cambridge, MA: Harvard Univ. Press, 2006.

• K. Nader, "Response to Arshavsky: challenging the old views," *Trends Neurosci.*, vol. 26, no. 9, pp. 466–467, Sept. 2003.

• B. Pakkenberg and H.J. Gundersen, "Neocortical neuron number in humans: effect of sex and age," *J. Comp. Neurol.*, vol. 384, no. 2, pp. 312–320, July 1997.

• R.Q. Quiroga et al., "Invariant visual representation by single neurons in the human brain," *Nature*, vol. 435, no. 7045, pp. 1102–1107, June 2005.

• B. Setlow, "Georges Ungar and memory transfer," *J. Hist. Neurosci.*, vol. 6, no. 2, pp. 181–192, Aug. 1997.

• C.E. Shannon, "A mathematical theory of communication," *Bell Syst. Tech. J.*, vol. 27, pp. 379–423, 623–656, July and Oct. 1948.

• D. Tammet, *Born On a Blue Day: Inside the Extraordinary Mind of an Autistic Savant*. New York: Free Press, 2007.

• M.S. Zaragoza and S.M. Lane, "Source misattributions and the suggestibility of eyewitness memory," *J. Exp. Psychol. Learn. Mem. Cogn.*, vol. 20, no. 4, pp. 934–945, July 1994.

About the authors

Timothy Marzullo received his Ph.D. in neuroscience in Daryl Kipke's Neural Engineering Lab at the University of Michigan. His interests include neurotechnology and neocortex electrophysiology.

Hirak Parikh is working toward his Ph.D. in biomedical engineering in Daryl Kipke's Neural Engineering Lab at the University of Michigan. His interests include modeling the neocortex's laminar structure and its contribution to computation.

FINISH IT!

www.finish-it.org

PERFECT THE DEGREE YOU'VE earned. Take the FE Exam.

LAUNCH THE CAREER YOU deserve. Turn that little ee into a big PE.

E Find out what it takes. Ask NCEES. www.finish-it.org

L I C E N S U R E : Y O U ' R E W O R T H I T

IEEE potentials

the magazine
for high-tech innovators



Got a machine head?

You don't know Jack

Go with the (micro) flow

Neurons make the man

Finding your identity management system

Fuzzy networked control