From Synapses to Circuitry: Using Memristive Memory to Explore the Electronic Brain

12 authors, including:

Richard Carter
Hewlett Packard Enterprise
28 PUBLICATIONS   632 CITATIONS

SEE PROFILE

Hisham Abdalla
University of Maryland, College Park
14 PUBLICATIONS   418 CITATIONS

SEE PROFILE

M.s. Qureshi
Boston University
19 PUBLICATIONS   334 CITATIONS

SEE PROFILE

Jasmin Léveillé
Boston University
14 PUBLICATIONS   143 CITATIONS

SEE PROFILE

Computer

FEBRUARY 2011

http://www.computer.org

NANOSCALE ARCHITECTURES

◆IEEE

IEEE
Φcomputer
society

From Synapses to Circuitry: Using Memristive Memory to Explore the Electronic Brain

Greg Snider, Rick Amerson, Dick Carter, Hisham Abdalla, and Muhammad Shakeel Qureshi
Hewlett-Packard Laboratories

Jasmin Léveillé, Massimiliano Versace, Heather Ames, Sean Patrick, Benjamin Chandler, Anatoli Gorchetchnikov, and Ennio Mingolla, *Boston University*

In a synchronous digital platform for building large cognitive models, memristive nanodevices form dense, resistive memories that can be placed close to conventional processing circuitry. Through adaptive transformations, the devices can interact with the world in real time.

Building an electronic brain is daunting, in large part because researchers have only the faintest notion of how to do it. Indeed, in light of the biological-scale constraints on power dissipation and volume (roughly 20 W stuffed into a shoebox), the task seems well beyond current technology.

Consider, for example, one dimension of the problem: how to translate synapses into their electronic equivalent. Biological synapses are dense—the cortex needs roughly 10^{10} synapses per square centimeter. They also consume miniscule power; have complex, nonlinear dynamics; and, in some cases, can maintain their memory for decades. Until recently, these characteristics translated to one more unreachable goal for those aspiring to build electronic brains, particularly large models.

In the past few years, however, work on memristive devices[1] has gained momentum, which could bring designers closer to an electronic brain architecture that can adaptively interact with the world in real time. Although memristive devices alone don't solve power and volume problems, they have several attractive features that make them candidates for implementing synaptic memory in intelligent machines:[2]

- their resistance is generally nonlinear, and circuitry can alter it electrically;
- they can be packed into crossbars to form dense memories; and
- many memristive materials are compatible with CMOS processes.

The last characteristic is significant because it means that designers can integrate dense, memristive memories with conventional circuitry[3] and thus place memory and computational circuits closer together. The decreased distance significantly reduces the power dissipation in sending and receiving information between the two components while increasing the data bandwidth between them.

To bootstrap the process of building intelligent hardware, Hewlett-Packard and Boston University are jointly developing the Cog Ex Machina hardware architecture along with Cog, a software framework that runs on top of it. Together, the two provide a low-cost, flexible all-digital platform for building large brain models that can interact with a simulated or real environment in real time.

Although the platform does not yet achieve the goals of intelligent behavior and biological-scale power and volume, it does offer a way for researchers to build
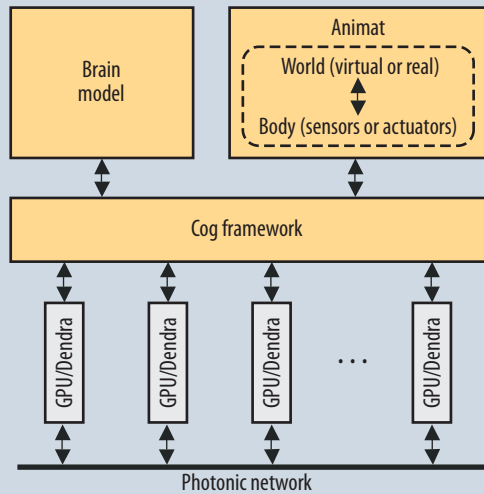
Figure 1. High-level view of the Cog Ex Machina platform. The hardware consists of accelerator nodes, currently GPUs but eventually Dendra chips, which communicate through a photonic network. Researchers build abstract brain models on the Cog software platform. Cog hides the underlying hardware, allowing brain models to run on many hardware implementations. Brain models use Cog to interact with the real world through robotic sensors and actuators or with a virtual world using a software animat.

models quickly and at relatively low cost and to adapt the platform to fit new algorithms. In addition, Cog effectively abstracts away the details of the underlying hardware, so researchers can continue building models on the platform as the underlying hardware technology advances.

HARDWARE ARCHITECTURE

As Figure 1 shows, Cog Ex Machina's hardware architecture comprises multicore hardware accelerators for inference and learning. Because the platform has a digital hardware foundation, fabrication risk is low, and users are free to implement a wide variety of neuromorphic algorithms. The "Analog versus Digital Platforms for Neuromorphic Computing" sidebar describes the limitations of an analog approach relative to our platform. We currently use off-the-shelf graphical processing units (GPUs) to implement the accelerators, but we plan to transition these to Dendra chips, which are massively parallel processor chips that integrate hundreds of cores with memristive memory banks to reduce power and area by several orders of magnitude. Each accelerator, whether GPU or Dendra chip, communicates with other accelerators through a photonic network.

Distributed across the accelerators is the Cog software framework, which supplies researchers with a set of primitives for building massively parallel neuromorphic models. To study how models interact with their environments,

designers can plug in animats, either software creatures embedded in virtual environments or robots built from actuators and cameras, touch sensors, or accelerometers, enabling the model to interact with the real world in real time. Many robotic applications are possible (such as a robot quickly searching for trapped people, hazards, or hotspots before the fire department enters a burning building), but nearly any machine with a nontrivial interface (cell phones, remote controls, cars) could benefit from embedded intelligence.

Power

With the continuing reduction in CMOS feature size, capacitive signaling losses in the wiring are increasingly dominating a chip's power budget. These losses are of particular concern in cognitive architecture design because the brain's wiring is 3D and extremely dense.[4,5] Moreover, brain computation is massively parallel, reading and modifying enormous amounts of memory (synapses) continuously. In light of these complexities, to minimize signaling losses, designers have little choice but to place dense, low-power memory very close to the computational circuits that read and write it. The memristive memory we are developing for our Cog Ex Machina architecture enables us to do exactly that.

Dendra chip

The Dendra chip is a tiled array of *transform engines*, each of which includes a simple processor and a large bank of memristive memory built from dozens of memristive crossbars. Transform engines intercommunicate through a fabric on-chip, and over a network between chips. Each transform engine typically reads and modifies its entire memory bank every 10 ms, so designers must keep wires as short as possible to minimize power consumption.

SOFTWARE ARCHITECTURE

Cog users express a brain model as an arbitrary, directed graph, such as that in Figure 2. The nodes hold computational state and exchange information through edges. All nodes execute one computational step in parallel and then pass messages through edges before executing the next computational step. The graph is clocked at 100 Hz for real-time applications, allowing each node 10 ms to complete its computation and communication at each clock step.

An edge relays information from one node to another using a *tensor field*—a discrete, multidimensional array of tensors. Each tensor, in turn, is a multidimensional numerical array (scalar, vector, dyad, and so on). Computational nodes implement adaptive transformations that use incoming tensor fields to produce output tensor fields on outgoing edges. The transfer function varies over time because an adaptive transformation

ANALOG VERSUS DIGITAL PLATFORMS FOR NEUROMORPHIC COMPUTING

Analog neuromorphic computation using subthreshold CMOS[1] offers a potentially power-efficient path for solving the nonlinear differential equations that pervade many neural models. Several ongoing neuromorphic hardware projects such as Stanford's Brains in Silicon project,[2] Fast Analog Computing with Emergent Transient States (FACETS),[3] and IBM's SyNAPSE project[4] continue this tradition.

Despite the potential power advantages, analog computation carries risk. Because much of the neural dynamics are hardwired in the analog circuitry, an analog platform is not very flexible. In addition, its operation has a fixed time scale,[3] and computation is nondeterministic because of device variation and parasitics. Finally, the fabrication technology is hardly mainstream in a digitally dominated industry.

In addition, learning in analog neuromorphic architectures is often restricted to one or at most a few learning laws, with spike-timing-dependent plasticity (STDP) being the most popular. However, it is not at all clear that STDP is either necessary or sufficient to implement the wide variety of processing in brains. Systems that use dynamical nanodevices, such as the memristive devices we have developed for analog synaptic memories, also face challenges in stabilizing those memories to deal with noise and system dynamics.

Digital computers, on the other hand, are mainstream, algorithmically flexible, and stable, but they are inefficient at numerical integration. In other words, digital computers are very good at alge-bra but not so good at mathematical analysis. We are addressing this dilemma by combining a digital hardware platform with algorithmic transformations that recast neural algorithms from the analysis domain to the algebraic domain. Both the contrast normalization and boundary completion examples illustrate this. In each case, the original neural algorithm was expressed with differential equations, but was implemented with Cog using a transformation to a noniterative algebraic algorithm. It is unknown how well this approach will work in general—it is a risk. However, we feel that the flexibility and reduced fabrication risk of a software digital platform adequately compensates for any downside.

References

1. C. Mead, *Analog VLSI and Neural Systems*, Addison Wesley Longman, 1989.
2. K. Boahen, "Brains in Silicon"; www.stanford.edu/group/brainsinsilicon.
3. J. Schemmel, J. Fieres, and K. Meier, "Wafer-Scale Integration of Analog Neural Networks," *Proc. Int'l Joint Conf. Neural Networks* (IJCNN 08), IEEE Press, 2008, pp. 431-438; http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=4633828.
4. R. Ananthanarayanan et al., "The Cat Is Out of the Bag: Cortical Simulations with 109 Neurons, 1013 Synapses," *Proc. Conf. High-Performance Computing Networking, Storage and Analysis* (SC 09), ACM Press, 2009; http://doi.acm.org/10.1145/1654059.165412.

can change its internal state as a function of its inputs' history, such as feedback from other transformations that it drives.

Adaptive transformations are either linear or nonlinear. Linear transformations, which make up more than 99 percent of the computation, implement a form of tensor convolution that extracts subfields from input tensor fields and then multiplies or contracts those subfields with internal tensor kernels to produce output tensor fields. The kernels can vary over the field and can self-modify their internal state through learning laws such as Hebb rule derivatives. Because linear transformations are computationally regular, parallel hardware can efficiently implement them.

Although nonlinear transformations constitute less than 1 percent of the total computation in most models, they implement essential irregular functions and nonlinear dynamics, transforming one or more input tensor fields to output tensor fields. Through outgoing tensor fields, they can supply feedback to guide learning in the linear transformations connected to their inputs. And, because they hold state, they can also adapt and have time-varying transfer functions.

Each adaptive transformation, whether linear or nonlinear, is itself a massively parallel computation. Cog automatically maps these computations onto the underlying platform. If a single core is all that's available, Cog time-multiplexes all computation on that core. If a single multicore CPU is available, Cog spreads the computation across the cores. When faced with a mix of CPUs, GPUs, and Dendra chips across a network, Cog maps the linear
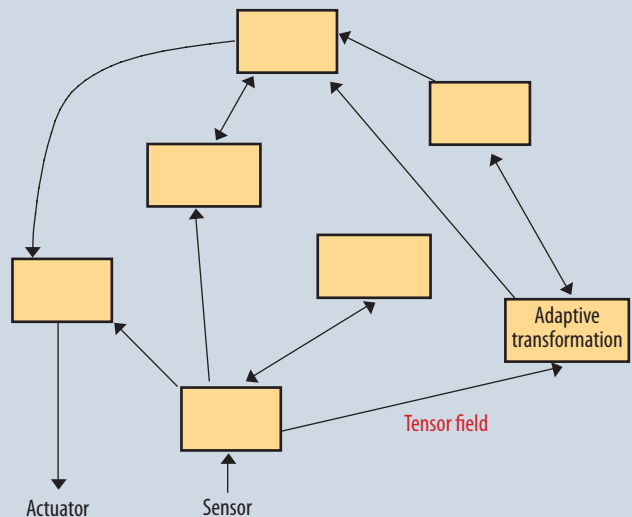

Figure 2. Building a brain model in Cog. Models are directed graphs of adaptive transformations that execute concurrently, exchanging information through tensor fields. Computation is deterministic and race free.
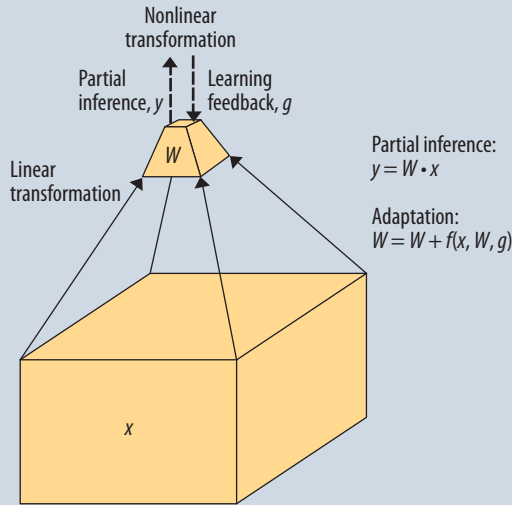
Partial inference:
$$y = W \cdot x$$

Adaptation:
$$W = W + f(x, W, g)$$

Figure 3. How learning occurs. Linear and nonlinear transformations cooperate to implement the learning that corresponds to long-term memory. Cog combines a linear transformation's current state, *W*, with the input, *x*, to produce *y*, a partial inference. It then combines *g*, *x*, and *W* to implement learning. The learning function, *f*, determines the type of learning that occurs.

transformations onto the regular computational engines (GPUs, Dendras), maps nonlinear transformations onto the CPU cores, and implements tensor field communication through local memory or through messages passed across the network.

None of these platform attributes are visible to users, which frees them to focus on their cognitive models. Adding computational resources either speeds up model execution or, in a real-time environment, increases the size of the model that Cog can execute.

These abstractions might seem distant from their biological counterparts, but a rough correspondence exists. Tensor fields moving along the graph edges are similar to the information that axon bundles convey in a human nervous system. Linear transformations are analogous to the computation performed in the dendritic trees of neuron populations, with the learning or adaptation analogous to the modifications of synaptic weights. This is the storage of long-term memory. Nonlinear transformations correspond to the nonlinear dynamics of populations of neuron bodies or somas—the storage of medium- and short-term memories.

Learning

Linear transformation adaptation generally occurs over a much longer time scale relative to nonlinear transformations, and this slower adaptation is what constitutes learning. As Figure 3 shows, feedback from a nonlinear transformation guides learning. Cog holds the actual learned state, *W*, within a linear transformation. It then convolves or correlates *W* with an input, *x* (part of a tensor field), to produce an output tensor field, *y*. We call this a *partial inference*. Cog uses the partial inference to drive a nonlinear transformation, which can respond by feeding back a learning field, *g*, to the linear transformation. The linear transformation uses *g*, *x*, and *W* (its current state) to update its learned state.

Through configuration and appropriate feedback,[6] we can implement a wide variety of classical learning laws, which fall into four categories:

- *Hebb rule derivatives*, including classic Hebbian, Hebb plus passive decay, presynaptically gated decay (outstar), postsynaptically gated decay (instar), Oja, dual OR, and dual AND;
- *threshold-based rules*, including Covariance 1, Covariance 2, BCM (Dayan and Abbott), original BCM (oBCM), IBCM, and Bienenstock, Cooper, and Munro (BCM) theory (Law and Cooper);
- *feedback-based rules*, including back-propagation, Harpur's rule, and contrastive divergence; and
- *temporal-trace-based rules*, including Rescorla Wagner, temporal difference, and Foldiak.

Development environment

Brains interact with environments, receiving information from sensors and conveying motor commands so that the bodies they govern can move about and affect the world around them. To make development easier, Cog offers a set of virtual environments with which brain models can interact. The environments, which vary in complexity, run synchronously with the brain model.

Because Cog is a synchronous, digital architecture, users can halt and restart model execution without perturbing the computation, so that they can peek inside their model and debug it if necessary. If enough computational resources are available, Cog can run faster than real time in a virtual environment.

Figure 4 is a screenshot of the graphical user interface (GUI) in the Cog debugger. The left side of the screen shows a graph of the network being debugged; in this case, a feed-forward network is implementing a simple form of boundary completion. The user can click on various points within that model to view its internal state, which the GUI displays on the right side.

BUILDING APPLICATIONS

Cog can implement a wide range of neuromorphic algorithms. To give an idea of Cog's capabilities, we describe four small applications that we built using our framework: contrast normalization, independent component analysis, learning of orientation maps with ocular dominance, and boundary completion.
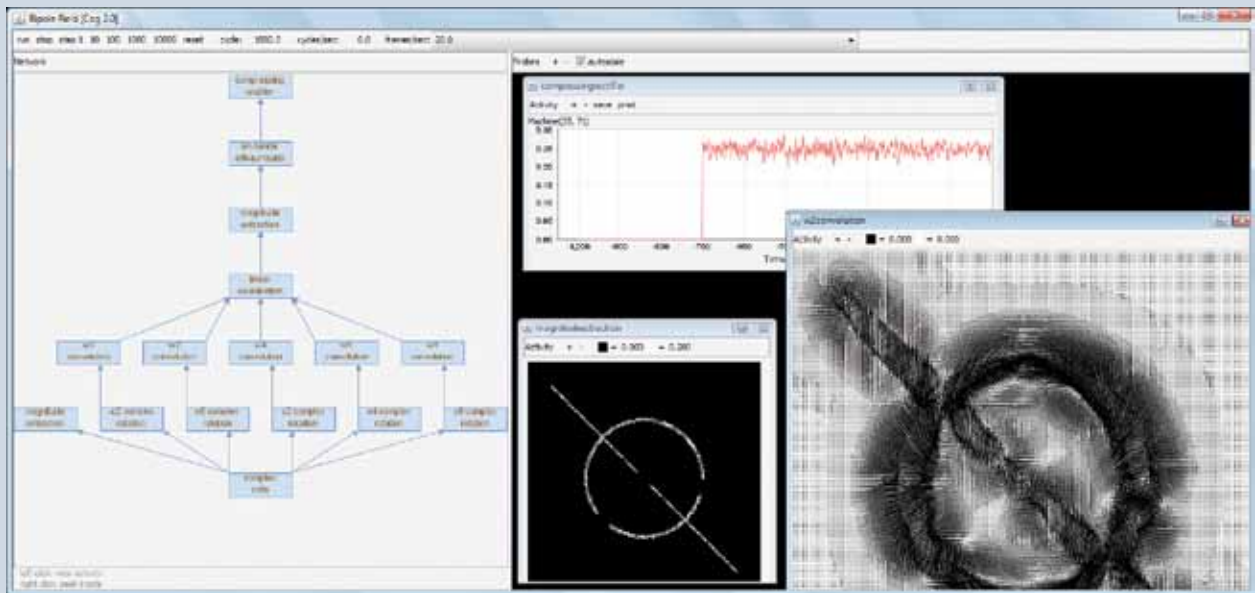
Figure 4. Debugging with Cog. (left) The GUI displays the network structure, which the user is probing to display the internal state of some of the model's adaptive transformations. (right) Each window names the adaptive transformation being displayed. The bottom two windows show snapshots of the tensor fields generated and transmitted by two of the network transformations, while the top window displays a history of a single state variable within a third transformation.

Contrast normalization

Figure 5 illustrates contrast normalization using the Retinex algorithm[7] as implemented with Cog. The image in the figure is of a parking garage, which contains a dynamic range that a camera cannot capture in all its detail, but that a human eye can. In Figure 5a, standard photographic compression loses detail in the brightest and darkest regions. However, one linear and one nonlinear transformation on Cog can easily implement the Retinex algorithm, which approximates retinal processing, to capture detail in both shadow and glare. Figure 5b shows the dramatic difference. Filtering is clearly nonlocal, since the brightest areas in the original image do not always correspond to the brightest areas in the processed image. This early preprocessing is essential for handling the real-world video streams that the brain model receives as input.

Independent component analysis

Scientists believe that most of the information in natural images is contained in the scene's edges[8] and that edge filters can capture and compress that information to simplify later processing. Figure 6 shows an example of a model built with Cog that uses BCM theory to learn the independent components of natural scenes.

In the application, a simple network implemented on Cog uses three transformations in series to approximate early processing in the visual system. The first transformation enhances edges using a difference-of-Gaussians

(a)

(b)

Figure 5. Contrast normalization of a high-contrast scene in a parking garage. (a) Standard photographic compression (jpeg) of the scene loses details in deep shadow and bright regions. (b) The Retinex algorithm as implemented with Cog recaptures the details while maintaining local contrast. (Image data courtesy of R. Brinkworth and D. O'Carroll, "Robust Models for Optic Flow Coding in Natural Scenes Inspired by Insect Biology," *PLoS Computational Biology*, vol. 5, no. 11; www.ploscompbiol.org/article/info%3Adoi%2F10.1371%2Fjournal.pcbi.1000555;jsessionid=4BEF883BF3CAA0B158302B10E2E74AA1.ambra02).
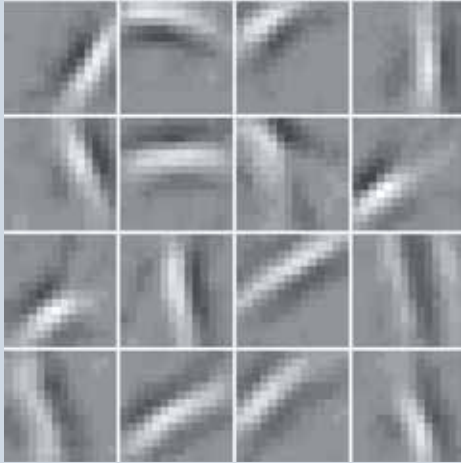
Figure 6. Learning independent components of natural scenes. A network built with Cog for learning independent components preprocessed random patches from natural scenes using a difference-of-Gaussians filter and then processed that output using the BCM neuron model combined with Hebbian learning. The result was the learning of the 16 Gabor-like edge filters shown.
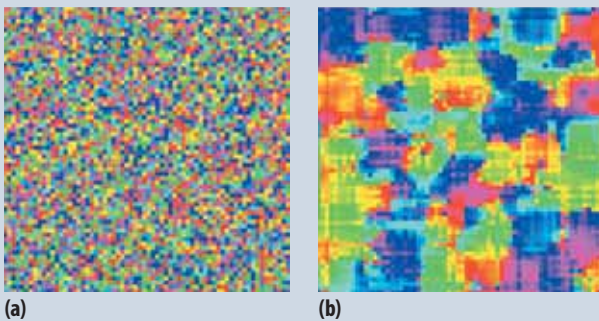

(a)   (b)

Figure 7. Topographic map of orientation selectivity. (a) Before learning, there are no identifiable clusters of orientation selectivity. (b) Clusters emerge after learning, with different colors denoting particular orientations.
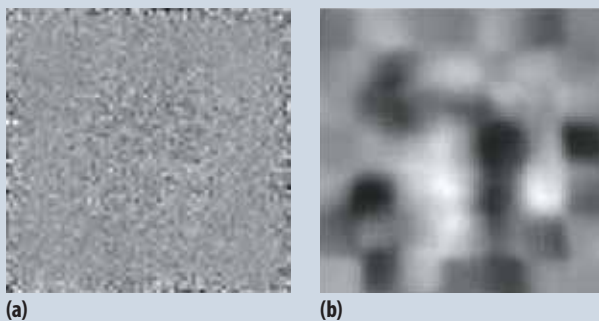

(a)   (b)

Figure 8. Topographic map of ocular dominance. (a) Before learning, there are no identifiable columns of ocular dominance. (b) Columns emerge after learning, with white denoting left eye, black denoting right eye, and gray denoting intermediate.

filter. The second, a linear transformation, applies convolution with adaptive kernels that implement simple Hebbian learning. Finally, the third transformation implements a competitive field of BCM theory neurons.[9] Random patches from natural images (photographs of trees, grass, fields) drive the first layer, and the last two layers respond by adapting to their inputs. The resulting network self-organizes into a set of edge filters, shown in the figure, that react strongly to edges (the "independent components" of natural scenes) in visual inputs (approximating simple cells in the V1 cortical region of the vision system).

Orientation maps and ocular dominance

Experiments show that simple cells (neurons) in the V1 area of the visual cortex behave as edge filters and that the orientation of those edge filters varies smoothly over the surface of V1. Experiments also show that V1 cells receive visual input information from both eyes, but that signals from only one eye dominate any given cell. This ocular dominance organization occurs in clumps or bands, depending on the species.

Figures 7 and 8 illustrate a Cog model of the self-organization of V1 for both orientation and ocular dominance. In this example, we input a series of random images, heavily filtered to form blob-like images.[10] We then topographically connected two views of the image, simulating two eyes, to the simulated V1. Figure 7 shows the resulting self-organization of the orientation filters, and Figure 8 shows the resulting ocular dominance clumping.

Boundary completion

A person can easily recognize visual objects even when noise or foreground objects obstruct or partially occlude them. People somehow infer the missing information and fill in the scene to form completed percepts—presumably to facilitate recognition in later cortical stages. Boundary completion, a simple form of this filling-in process, com-
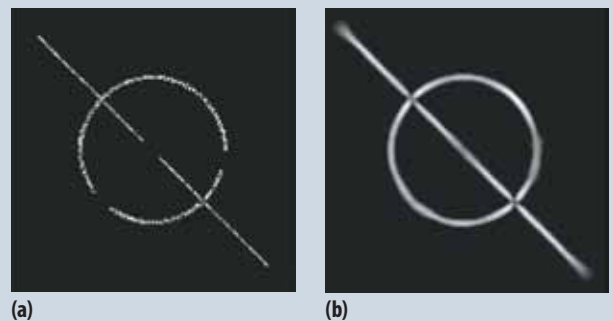

(a)   (b)

Figure 9. Simple boundary completion. (a) A noisy figure with a broken line and broken circle becomes (b) smooth and filled in using a boundary completion algorithm implemented with Cog.

pletes broken and curved edges and lines using the Gestalt principles of proximity and good continuation.

Although a self-organizing boundary-completion model exists, it is computationally expensive to learn.[11] We found it much cheaper simply to prewire the mechanism into a Cog-built model. The model can then exploit a vast number of mathematical tricks, such as steerable filters, fast Fourier transform, tensor convolution, and data compression.[12] Relative to the traditional boundary completion model,[13] a Cog-built model with this mechanism can reduce implementation energy by at least four orders of magnitude.

Figure 9 shows the input and output of a Cog-built model that implements a simple form of boundary completion. In this example, the user has implemented the feed-forward network in Figure 4 to complete the noisy input image of a broken circle and broken line.

Although a general theory of cognition does not yet exist, researchers do recognize that platform flexibility is essential as they plow through the fog and uncertainty of learning to build intelligent machines. Cog has many features that offer this flexibility. Its all-digital hardware foundation reduces technological and fabrication risk. Its placement of memristive synaptic memory banks close to their associated processing elements reduces CV2f power losses by several orders of magnitude—a reduction critical in processing neuromorphic algorithms, which deeply entangle memory and computation.

Cog's tensor framework mechanisms are perhaps nonbiological, but they are well-matched to our underlying CMOS/memristive technology. The framework is also expressive, pulling in linear algebra, geometry, and analysis into a single foundation, and enables exploitation of much mathematical and engineering knowledge—for example, information and coding theory, digital signal processing, non-Euclidean coordinate systems, tensor convolution, normalized convolution, and fast Fourier transforms, among many others. The framework also supports a wide variety of learning laws and network models.

Perhaps the most important architectural attribute is the nearly complete decoupling of the software abstractions for building brains (tensor fields and adaptive transformations) from the underlying hardware platform. Not only does this provide portability among existing and future platforms, it allows us to quickly modify the software architecture to accommodate new or unexpected algorithmic problems as they arise. **C**

Acknowledgments

References

1. L.O. Chua and S.M. Kang, "Memristive Devices and Systems," *Proc. IEEE*, vol. 64, no. 2, 1976, pp. 209-223.
2. D.B. Strukov et al., "The Missing Memristor Found," *Nature*, 1 May 2008, pp. 80-83.
3. Q. Xia et al., "Memristor/CMOS Hybrid Integrated Circuits for Reconfigurable Logic," *Nano Letters*, vol. 9, no. 10, 2009, pp. 3640-3645.
4. B. Madappuram et al., "On Brain-Inspired Connectivity and Hybrid Network Topologies," *IEEE Symp. Nanoscale Architectures* (NANOARCH 08), IEEE Press, June 2008, pp. 54-61.
5. D.S. Basset et al., "Efficient Physical Embedding of Topologically Complex Information Processing Networks in Brains and Computer Networks," *PLoS Computational Biology*, Apr. 2010, e1000748.
6. A. Gorchetchnikov et al., "General Form of Learning Algorithms for Neuromorphic Hardware Implementation," *BMC Neuroscience*, vol. 11 (supp. 1), 2010, p. 91; www.biomedcentral.com/1471-2202/11/S1/P91.
7. D.J. Jobson, Z. Rahman, and G.A. Woodell, "Properties and Performance of a Center/Surround Retinex," *IEEE Trans. Image Processing*, Mar. 1997, pp. 451-462.
8. A. Bell and T. Sejnowski, "The 'Independent Components' of Natural Scenes Are Edge Filters," *Vision Research*, Dec. 1997; pp. 3327-3338.
9. N. Intrator and L. Cooper, "Objective Function Formulation of the BCM Theory of Visual Cortical Plasticity: Statistical Connections, Stability Conditions," *Neural Networks*, vol. 5, no. 1, 1992, pp. 3-17.
10. R. Mikkulainen et al., *Computational Maps in the Visual Cortex*, Springer, 2005.
11. S. Grossberg and J. Williamson, "A Neural Model of How Interlaminar Connections of Visual Cortex Develop into Adult Circuits that Carry Out Perceptual Grouping and Learning," *Cerebral Cortex*, vol. 11, no. 1, 2001, pp. 37-58.
12. E. Franken et al., "An Efficient Method for Tensor Voting Using Steerable Filters," LNCS 3954, Springer, 2006, pp. 228-240.
13. S. Grossberg and E. Mingolla, "Neural Dynamics of Perceptual Grouping—Textures, Boundaries, and Emergent Segmentations," *Perception and Psychophysics*, vol. 38, no. 2, pp. 141-171.

Greg Snider *is a senior researcher at Hewlett-Packard Laboratories and the principal HP investigator in the DARPA SyNAPSE project. His research interests include hardware and software architectures. Snider received an MS in scientific instrumentation from the University of California, Santa Barbara, and an MSEE from Stanford University. Contact him at snider.greg@hp.com.*

Rick Amerson *is a managing consultant for Hewlett-Packard Laboratories. His research interests include computer architecture, instruction sets, memory, computation, and algorithms. Amerson received an SM in management of*

technology from the Massachusetts Institute of Technology. Contact him at frederic.amerson@hp.com.

Dick Carter is a consulting research scientist at Hewlett-Packard Laboratories. His research interests include parallel algorithms and GPU computing. Carter received an MSEE from Stanford University. He is a member of IEEE. Contact him at richard.carter@hp.com.

Hisham Abdalla is a research associate at Hewlett-Packard Laboratories. His research interests include neuromorphic engineering, robotics, computer architecture, and floating-point gate array design. Abdalla received a PhD in electrical and computer engineering from the University of Maryland at College Park. Contact him at hisham.abdalla@hp.com.

Muhammad Shakeel Qureshi is a researcher and circuit designer at Hewlett-Packard Laboratories. His research interests include low-power circuits, MEMS sensor interfaces, optoelectronics, and mixed-signal and radio-frequency circuits for communication blocks. Qureshi received a PhD in electrical and computer engineering from the Georgia Institute of Technology. Contact him at shakeel.qureshi@hp.com.

Jasmin Léveillé is a postdoctoral associate in the Department of Cognitive and Neural Systems at Boston University. His research interests include human and computer vision and neural architectures for high-performance computing. Léveillé received a PhD in cognitive and neural systems from Boston University. Contact him at jasminl@cns.bu.edu.

Massimiliano Versace is a senior research scientist in the Department of Cognitive and Neural Systems at Boston University and director of the university's Neuromorphics Laboratory. He is also codirector of technology outreach at NSF's Science of Learning Center's Center of Excellence for Learning in Education, Science, and Technology (CELEST). He is a principal investigator for the Boston University subcontract with Hewlett-Packard in the DARPA SyNAPSE project. Versace received a PhD in cognitive and neural systems from Boston University. Contact him at maxversace@gmail.com.

Heather Ames is a research scientist in the Department of Cognitive and Neural Systems at Boston University, codirector of technology outreach at NSF's CELEST, and a member of its governing board to facilitate technology transfer from models to private industries and laboratories. Her research interests include speech modeling, analysis of learning and homeostasis in neural network systems, and the technology transfer of brain-based applications. Ames received a PhD in cognitive and neural systems from Boston University. Contact her at heather.m.ames@gmail.com.

Sean Patrick is a doctoral student in the Department of Cognitive and Neural Systems at Boston University. His research interests include robotics and prosthetics; adaptive, neural, and evolutionary algorithms; and computational neuroscience. Patrick received a BS in electrical and computer engineering from the Worcester Polytechnic Institute. Contact him at sean.patrick.619@gmail.com.

Benjamin Chandler is a doctoral student in the Department of Cognitive and Neural Systems at Boston University and ACES associate at the university's Center for Computational Science. His research interests include large-scale simulation, homeostatic plasticity, and neuromorphic computing. Chandler received a BS in cognitive science from Carnegie Mellon University. He is a student member of IEEE. Contact him at bchandle@gmail.com.

Anatoli Gorchetchnikov is a research assistant professor in the Department of Cognitive and Neural Systems at Boston University and leader of the Modular Neural Exploring and Traveling Agent project within DARPA's SyNAPSE program. His research interests include the biologically detailed modeling of the hippocampus and related areas involved in spatial navigation and episodic memory, the application of biological models to robotic navigation, and the development of software tools to increase modeling productivity. Gorchetchnikov received a PhD in cognitive and neural systems from Boston University. Contact him at tangorn@gmail.com.

Ennio Mingolla is a professor and chair, Department of Cognitive and Neural Systems, Boston University. His research interests include the development and empirical testing of neural network models of visual perception and the transition of these models to technological applications, visual psychophysics, and the computational modeling of brain processes. Mingolla received a PhD in experimental psychology from the University of Connecticut. Contact him at ennio@cns.bu.edu.

cn **Selected CS articles and columns are available for free at http://ComputingNow.computer.org.**