Machine Intelligence from Cortical Networks (MICrONS)

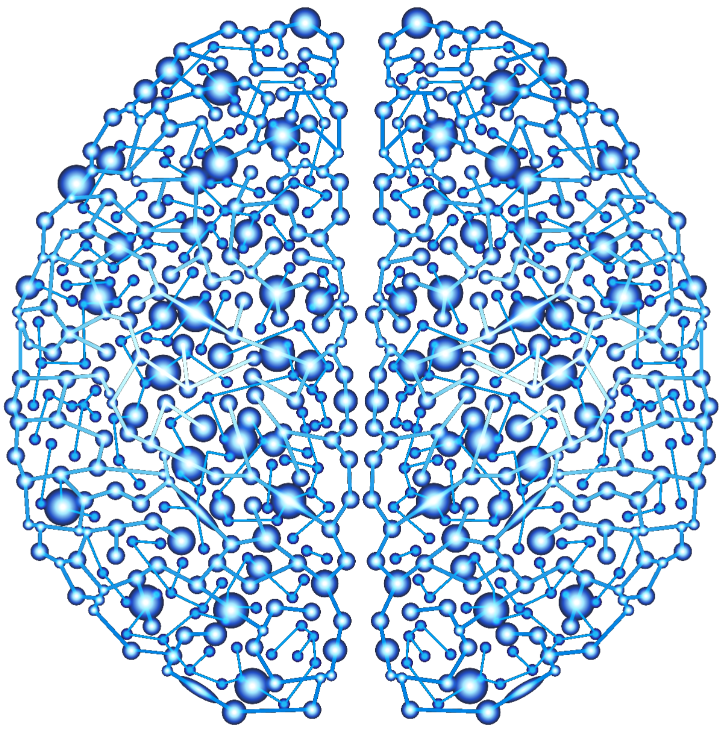

Proposers' Day Abstracts

July 17, 2014

Technical Area 1

Propose an algorithmic framework for information processing that is consistent with existing neuroscience data, but that cannot be fully realized without additional specific knowledge about the data representations, computations, and network architectures employed by the brain

Cortical Theory for Learning Invariances

Lead Investigator: Jeff Hawkins
Team members: Subutai Ahmad, Scott Purdy, Chetan Surpur
Organization: Numenta, Inc., www.numenta.com

Numenta research and capabilities:

Numenta has two broad research agendas:

1) Understand the computational principles of the cortex, and
2) Implement intelligent systems that operate on those principles

Numenta is unique in that we use three complementary approaches. First, we create theoretical models that closely follow biological principles and known neuroscience data. Initially published in the book On Intelligence, we have continued to develop Hierarchical Temporal Memory (HTM), a detailed computational theory of the neocortex. Second, we implement the models as optimized, high quality software modules and test them on numerous datasets. Finally, we release our implementation in the form of GPL open source software as well as commercial products.

Research Proposal: Invariance Learning

We have made progress recently on laying out the canonical circuit within cortical layers and minicolumns. The goal is to demonstrate how laminar circuits implement sequence memory, feedback, sensorimotor control, hierarchy, invariance, learning and attention. For this proposal we would like to focus on one aspect of this, specifically the ability of the cortex to learn arbitrary invariances.

Neurons in V4 and IT have been shown to represent highly invariant but specific visual categories. A holy grail for cortical models is to show how such invariances can be learned (i.e. without hardcoding) for diverse sensory modalities. This capability is acknowledged to be the key to robust pattern recognition and hierarchies. As stated in DiCarlo & Cox (2007), "*Understanding the brain mechanisms that underlie this ability would be a landmark achievement in neuroscience.*"

We have a theoretical model that explains how the cortex learns invariances. We wish to lead a team to expand the scope of this work and demonstrate robust learning of invariances for a sensory modality. The goal is to implement a cortical model that learns invariant representations for a large number of categories. The model must be faithful to biological constraints and be general (i.e. independent of any specific input modality).

We are searching for team members that have: 1) deep expertise and data collection in a sensory area such as vision, audition, 3D sensors, 2) access to significant hardware and software resources for running large experiments, or 3) the ability to design and conduct neuroscience experiments to test the theory.

Interested teams should contact: Subutai Ahmad, VP Research, sahmad@numenta.com

Modeling neural computation using biologically inspired online algorithms

Dmitri "Mitya" Chklovskii

Simons Center for Data Analysis, Simons Foundation, New York, NY

How does the brain compute? Although this question has been asked many times before I believe that now is the right time to answer it. The cause for such optimism is the recent wealth of experimental datasets acquired using cutting-edge experimental techniques in combination with the development of theoretical tools from computer science and electrical engineering.

Our ability, for the first time, to both image activity in neuronal populations using Ca indicators and to reconstruct connectomes of the same population using electron microscopy (EM) yields an unprecedented amount of information about the same neuronal circuit [1,2]. This information will guide us in formulating a theoretical framework of neural computation. Specifically, we propose to use machine learning algorithms developed for online (or streaming data) setting [4] to model the function of neuronal circuits [7].

We have made the first step towards understanding neuronal computation by modeling single-neuron computation as an online factorization of the incoming data matrix [7]. Such approach allows us to derive neuronal activity dynamics, as well as synaptic learning rules from a principled objective function. Thus, we can course-grain the unmanageable complexity of ion-channel physiology and generate experimentally testable predictions. To test these predictions, we will rely on both anatomical and physiological datasets. Interestingly, the enormous size of experimental datasets [3] may require for their analysis the use of online algorithms [4] similar to the ones used to model neuronal function.

My group is uniquely positioned to carry out this research program. It combines expertise in analyzing both anatomical [5,6] and physiological [8] data as well as formulating biologically plausible machine learning algorithms [7].

[1] Bock DD, et al Reid, RC: Network anatomy and in vivo physiology of visual cortical neurons. *Nature* 2011, 471:177-182.

[2] Briggman, K. L., Helmstaedter, M. & Denk, W. Wiring specificity in the direction selectivity circuit of the retina. *Nature* 471, 183–188 (2011).

[3] Hayworth, KJ, et al Lichtman JW. Imaging ATUM ultrathin section libraries with WaferMapper: A multi-scale approach to EM reconstruction of neural circuits. *Frontiers in Neural Circuits* 8 (2014): 68.

[4] Shalev-Shwartz, S., Online Learning and Online Convex Optimization. *Foundations and Trends in Machine Learning*, vol. 4, pp. 107-194, 2011.

[5] Takemura S, et al, Chklovskii DB: A visual motion detection circuit suggested by Drosophila connectomics. *Nature* 2013, 500:175-181.

[6] Plaza, S, Scheffer, L, Chklovskii DB, (2014) Towards large-scale connectome reconstructions. *Current Opinion in Neurobiology 25*: 201-210.

[7] Hu, T, et al, Chklovskii, DB. A neuron as a signal processing device. *Asilomar Conference on Signals, Systems and Computers,* Pacific Grove, CA, Nov. 3-6 2013.

[8] Pnevmatikakis, EA, et al Paninski L,Bayesian spike inference from calcium imaging data. *Asilomar Conference on Signals, Systems & Computers*, Pacific Grove, CA, 2013.

Information process architectural approach to the cortex

L. Andrew Coward and Tom Gedeon
Australian National University

Some modern electronic systems have extreme complexities. For example, the systems that control sections of large telecommunications networks may have hundreds of billions of transistors, many thousands of user features, and require many thousands of man years of effort to design. These systems are nevertheless understood by human beings, in the sense that they can be designed, constructed, modified, and repaired as required.

Although there is minimal direct resemblance between such systems and the brain, the techniques used to organize the knowledge of such systems so that they can be understood can be applied to the brain. It can be demonstrated that the needs to economize on information handling resources and to modify features with no undesirable side effects on other features places some strong constraints on the organization of electronic system architectures. It can also be demonstrated that the needs to economize on physiological resources and learn without interference with prior learning result in natural selection pressures that strongly constrain brain architectures, although into a qualitatively different form from electronic systems designed under external intellectual control.

The constraints on brains result in the emergence of a range of subsystems performing different types of information processes, analogous with the memory and processing structures in electronic systems, but again qualitatively different. One of these subsystems is a structure that defines and detects information conditions on different levels of complexity within the information available to the system

The properties of the theoretical condition define/detect subsystem can be mapped to cortical anatomical structures. The way in which cortical processes are utilized by other brain structures to achieve cognitive behaviours further constrains the information processes performed by the cortex. Cognitive processes can be understood as combinations and sequences of information processes performed the cortex and utilized by other structures[1]. Key information processes include indirect activation of receptive fields on the basis of frequent past simultaneous activity which enables semantic memory, and indirect activation of receptive fields on the basis of simultaneous past change which enables retrieval of memories of specific events.

The theoretical considerations constrain cortical architecture, but do not specify exactly how a given information process will be implemented by anatomy in detail. However, the constraints give a strong starting point for studying what information processes are implemented by different cortical structures, and predicts limits within which anatomy must implement cognition. Anatomical, physiological and chemical studies can then determine exactly how the cortex implements the different types of information process.

---

[1] Coward, L. A. (2013). Towards a Theoretical Neuroscience: from cell chemistry to cognition. Springer: Amsterdam.

Automated serial electron microscopy of brain samples

JW Lichtman, Center for Brain Science Harvard University, Cambridge MA 02138

In order to generate large volumetric datasets of brain tissue, my colleagues and I have developed an automated pipeline to section, image, and analyze neural circuits in brain samples. We section brain using a standard ultramicrotome with a device that automatically picks of the samples and puts them on a tape. We section below 30nm in thickness and can collect up to ~12,000 sections per day. The sections are then cut into strips, placed on a silicon wafer, and imaged using scanning electron microscopy. To speed image acquisition we have developed automated means to locate the positions of sections on the wafer for high resolution imaging (a program called "Wafermapper", Hayworth, Morgan et al., 2014). We have also helped optimize image acquisition by participating in the development of a multibeam scanning electron microscope (Carl Zeiss Inc) particularly for brain imaging. This device, now resident in my laboratory, uses 61 beams to image at speed approaching 1 billion pixels per second at 4nm resolution using secondary electrons. These tools may be useful for the acquisition of large brain samples.

High Temporal, Spatial and Algorithmically Explicit Representation of the Human Brain: *When*, *Where* and *What* of Visual Recognition

PI: Aude Oliva, Computer Science and Artificial Intelligence Laboratory (CSAIL), MIT
Email: oliva@mit.edu, Phone: 617 452 2492, Web: http://cvcl.mit.edu/aude.htm
Collaborators: Antonio Torralba (torralba@csail.mit.edu), CSAIL, and Dimitrios Pantazis (pantazis@mit.edu), McGovern Institute for Brain Research, MIT.

Every time we open our eyes, visual information flows into various parts of our brain, with each region interpreting different aspects of what we are seeing. Integrating the temporal and spatial dynamics of this processing stream has posed a decades-long challenge to neuroscience. Using a novel neuroimaging technique [1], we can now map the flow of visual object information in the whole human brain and identify the stages of visual recognition processes at the millisecond and millimeter scales. The next step is to determine the type of computation that a human cortical region or networks of region perform. Using convolutional neural networks (CNNs) and large datasets (millions of labeled images, [2]) tailored to better represent human visual experience, we can compare the representations of various layers in the human and artificial neural networks, and adapt the artificial networks to emulate the process of categorization of complex visual stimuli, including recognition, familiarity, memory and decision making. Our approach demonstrates the possibility of a large-scale view of the dynamics and algorithms of recognition at the scale of processing steps across the whole human brain.

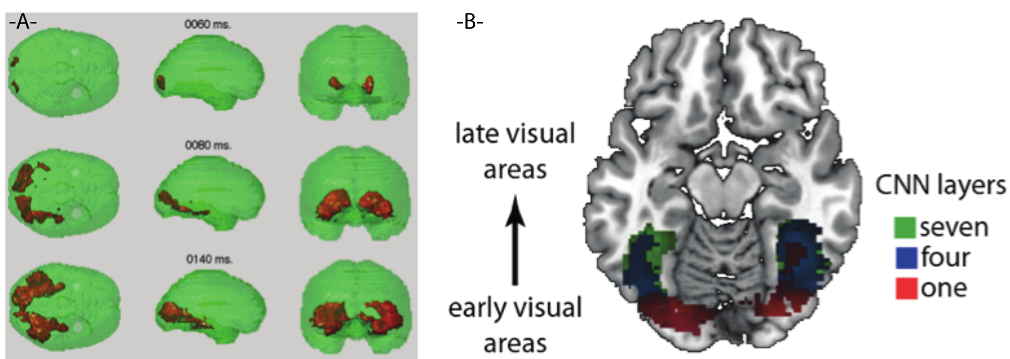

Figure: A) Spatial brain activity over time. The signal starts in the V1 region as early as 60ms and quickly spreads deeper into the brain enabling visual recognition. B) Visualizing the spatial correlation of CNN layers with human ventral visual system.

References:

[1] Cichy, R.M., Pantazis , D., & Oliva, A. (2014). Resolving human object recognition in space and time. *Nature Neuroscience*, 17(3), 455-464.
[2] Zhou, B., Xiao, J., Lapedriza-Garcia, A., Oliva, A., & Torralba, A. (2014). Learning deep features for scene recognition using PLACES database. Submitted Manuscript.

Sparse coding and hierarchical inference in cortical circuits

Bruno A. Olshausen, lead PI (UC Berkeley)
Chris Rozell (GeorgiaTech), Fritz Sommer, Trevor Darrell (UC Berkeley)

Our team aims to elucidate the neural mechanisms of sparse coding and hierarchical inference in cortical circuits, and to leverage these new insights to improve the state of the art in computer vision and image analysis.

There is now much evidence that nervous systems utilize sparse representations at early stages of sensory coding. At the same time, state-of-the-art advances in computer vision and signal processing have begun to employ sparse representations because of their proven effectiveness in tasks such as recognition and compression. However, the exact manner in which nervous systems compute sparse codes is still unknown. Elucidating these mechanisms through a combination of neurophysiological recording and connectomics would not only improve our understanding of the brain, but it could lead to new and more efficient methods for computing sparse representations in practical, technological applications. In particular, we may find new ways to compute sparse codes using dynamical systems that provide robust and informative representations of time-varying input.

Beyond sparse coding, one must understand how representations are transformed at higher levels of representation, and how feedback from higher levels modifies and improves representations at lower levels. It is well known that cortical systems represent sensory information in a hierarchy of processing stages, but how exactly signals are transformed from one stage to the next is largely unknown. The computer vision community has recently had astounding success using deep networks for recognition that are inspired by the hierarchical structure of cortex, and so presumably having more detailed information about this system would enable further improvements in computer vision systems. In particular, understanding the structure of feedback projections and the role they play in information processing and perception could prove extremely valuable. One theory is based on hierarchical Bayesian inference, and proposes that higher levels send 'priors' to lower levels to disambiguate representations. However, this theory has yet to be fleshed out and implemented in a detailed fashion that informs us about the neurobiology. Thus, a second thrust we propose is to develop a detailed model of hierarchical Bayesian inference, building on the already successful deep learning models. At the same time we will characterize the neuroanatomical details of feedback circuitry, first in LGN and then in V1, through connectomics. By guiding this work in a theory driven manner, we aim to gain new insights and constraints about the role of feedback connections and their role in information processing. We also aim to improve the state of the art in computer vision, as measured by performance on benchmark tasks, by using the principles of hierarchical inference ascertained from cortical feedback circuits.

Our team possesses expertise in sparse coding models and unsupervised learning algorithms, especially biophysically realistic mechanistic models of these processes. We also have experience in developing deployable vision systems based on efficient large scale algorithms and infrastructures, specifically the recent yet already widely adopted open source deep learning software framework developed at Berkeley (c.f., caffe.berkeleyvision.org, and it's precursor, "DeCAF").

Computational Advantages of SDR-Based Event Recognition Method

Gerard (Rod) Rinkus, Neurithmic Systems LLC, rod@neurithmicsystems.com

The goal of IARPA's MICrONS Program is to revolutionize machine intelligence by emulating the brain's computing "primitives" and the large-scale architecture in which they are embedded. Neurithmic Systems is developing a canonical cortical circuit model (TEMECOR, Sparsey®) and seeks to build a MICrONS Program based on it. Sparsey is a generic spatiotemporal probabilistic learning / inference algorithm and is currently being applied to video and multi-modal event recognition problems under ONR and DARPA UPSIDE support.

Neurithmic Systems LLC seeks to team with experimentalists and other modelers interested in developing a cortically-inspired generic event recognition system based centrally on Sparsey. Like many other model, e.g., HMAX models, Deep Learning models, Sparsey realizes the benefits of being stackable to arbitrary depth. However, unlike these and other model classes, it has been developed from the outset as a spatiotemporal model and as a model of both episodic and semantic memory. However, the strongest and clearest reason for pursuing a Sparsey-based model is that it possesses world-beating computational time efficiency. That efficiency depends crucially on its use of sparse distributed coding (SDC). For a computational approach to be a serious contender in general intelligent mining/processing of "big data", it must have world-beating computational time complexity for *both* learning and inference, or in the language of databases, both storage and best-match (not merely exact-match) retrieval of data. Sparsey possesses what we call "fixed-time complexity", meaning that *for any particular problem size,* and thus, for the particular model instance large enough to solve it, the algorithm's run time remains fixed for the life of the model instance. Sparsey possesses *fixed time* complexity for *both* learning and best-match retrieval, a.k.a., probabilistic inference, pattern recognition. *No other published information-processing algorithm of any kind has this capability!*

Sparsey's canonical algorithm operates similarly in all macs at all levels of an arbitrarily deep hierarchical model. This algorithm has the property that it maps more similar inputs to more similar (more highly intersecting) SDCs. In order to learn arbitrarily nonlinear categories, i.e., of the type referred to by Bengio (2012) as instances of the "AI Set", Sparsey relies on supervised learning, which is implemented as cross-modal unsupervised learning.

For MICrONS, we intend to put together a project that will include two main efforts. First, we would like to team with experimentalists with keen interest in imaging the large-scale population activity, at single-neuron resolution, such as possible with 2-photon calcium imaging. As part of this effort, we would like to both vet existing aspects/predictions of the model and develop increasing detail, e.g., modeling different cortical lamina, and coverage, e.g., adding a hippocampus analog, thalamic detail, etc. The second effort will focus on applying the Sparsey-based system to problems of interest to IARPA, e.g., recognizing events in video as well as to events in multi-modal data streams. We believe that due to its fixed-time complexity, our approach is highly likely to be able to scale to "Big Data"-sized problems.

MICrONS Proposers' Day Abstract

Lead Investigator: Dr. Alexander D. Wissner-Gross[1,2,3,*]
[1] Gemedy, Inc.; [2] Harvard University Institute for Applied Computational Science;
[3] MIT Media Laboratory; * Email: alexwg@gemedy.com; Web: http://www.alexwg.org

Recent advances in fields ranging from cosmology to computer science have hinted at a possible deep connection between intelligence and entropy maximization, but no formal physical relationship between them has historically been established. Recently, we proposed [1,2,3] a first step toward such a relationship in the form of a causal generalization of entropic forces that we found could cause two defining behaviors of the human "cognitive niche"—tool use and social cooperation—to spontaneously emerge in simple physical systems. Our results suggested a general thermodynamic model of adaptive behavior as a nonequilibrium process in open systems. Encouraged by this progress in reproducing macroscale cognitive adaptive behavior from an exceptionally simple thermodynamic principle, we are currently investigating **network architectures and algorithmic frameworks for neural information processing that are consistent with existing microscale and macroscale neuroscience data, in which the fundamental cortical computing primitive acts as a causal entropy maximizing module, but that cannot be fully realized without additional mesoscale cortical microcircuit knowledge**.

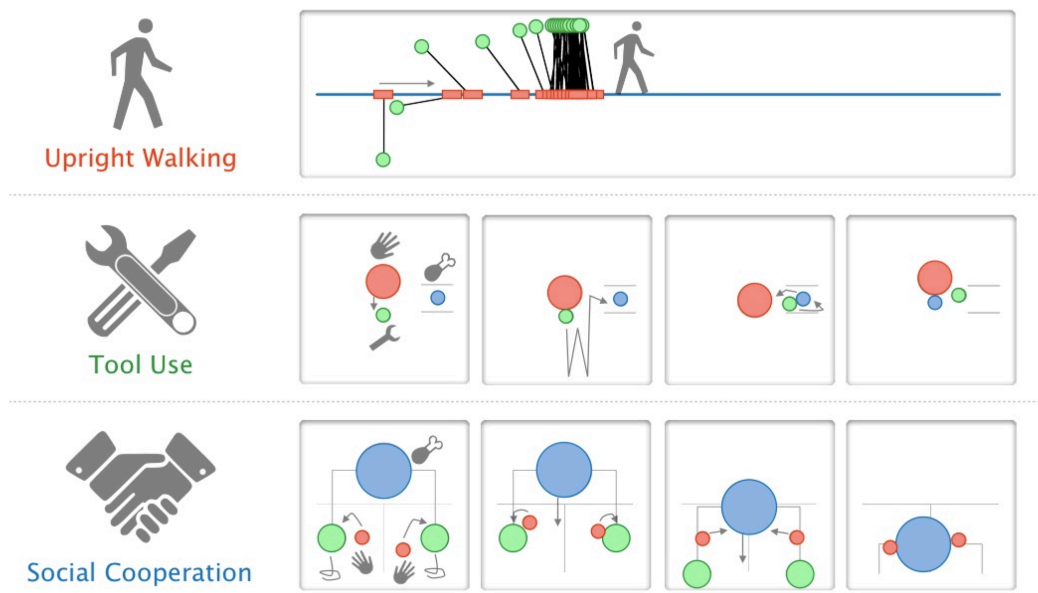

Figure 1. Examples of adaptive behaviors generated by causal entropy maximization [1].

References

[1] A. D. Wissner-Gross, et al., "Causal entropic forces," Phys. Rev. Lett. 110, 168702 (2013). Available at: http://www.alexwg.org/publications/PhysRevLett_110-168702.pdf

[2] D. Monroe, "Model suggests link between intelligence and entropy," Physics 6, 46 (2013). Available at: http://physics.aps.org/articles/v6/46

[3] A. D. Wissner-Gross, "A new equation for intelligence," TED Talk (2013). Available at: http://www.ted.com/talks/alex_wissner_gross_a_new_equation_for_intelligence

Compositional Models of Vision; A.L. Yuille (UCLA)

A fundamental challenge to understanding the visual system is to determine how it can deal with the enormous complexity of natural images and visual tasks. Analogous problems arise in other aspects of intelligence.

We proposes a solution to this problem based on the compositional hypothesis. This allows us to take advantage of a rich class of graphical models which have been developed to address different visual tasks. Recent work (Yuille and Mottaghi 2013) suggests ways that models of this type can be implemented by a visual architecture that has similarities to the visual cortex. We now briefly describe the representation, inference algorithms, and the learning of these models.

Objects and other visual structures are *represented* hierarchically in terms of compositions of more elementary parts which, recursively, are encoded in terms of subparts. Low level parts correspond to features like edges, high-level concepts are objects, and intermediate-level concepts are like gestalt groupings. Objects are represented in a distributed hierarchy in terms of their parts/subparts where the higher levels of the hierarchy only encode coarse summary descriptions of the objects (e.g., horse in a field) with the object details being specified explicitly at lower levels (consistent with Lee and Mumford's high resolution buffer hypothesis). Parts and subparts are *shared* between objects, yielding enormous gains in efficiency of representation and inference (Yuille and Mottaghi 2013). These models have much more explicit representations than alternative hierarchical models, such as HMax and Deep Neural Networks, and can perform a larger variety of visual tasks by accessing different levels of the hierarchy. For example, the top layer nodes identify the object(s) and give coarse localization while the lower-level nodes represent the positions of the object parts and the boundary of the object.

This compositional architecture enables an efficient *inference* algorithm, perhaps implemented by neural circuits, where objects are detected by recursively composing hypotheses for their parts and subparts. This enables rapid hierarchical inference (consistent with Thorpe et al.) with top-down processing required only to resolve low-level ambiguities. The intuition is that the lower-levels of the hierarchy represent low-level hypotheses about images (e.g., the presence of an edge, or a segmentation boundary). These representations are ambiguous during the bottom-up pass and hence multiple hypotheses must be considered. As these hypotheses are propagated up the hierarchy there is greater context and the visual structure are larger, more complex, and hence less ambiguous. In the top-down pass the higher-levels provide context to disambiguate the lower-levels. In addition, alternative algorithms can be specified on this architecture such as analysis by synthesis where a high-level node is activated (i.e. by priming) and an image of the object is generated. Similarly the architecture enables us to implement top-down attention. We note that compositional models are more similar to "programmable neural networks" (Valiant – Circuits of the Mind) where the network can function in different modes (e.g., top-down purely generative, bottom-up and top-down, attention, priming).

There has been preliminary work on *learning* this hierarchical architecture in an unsupervised manner (L. Zhu et al. 2008, L. Zhu et al. 2010). This learning follows the same compositional strategy used in inference. We recursively cluster elementary parts/suparts to form more complex structures.

Compositional models are successful on complex visual tasks evaluated on challenging datasets. They appear to be consistent with known properties of mamallian visual systems and it would be very exciting to test them in more detail and to modify them based on experimental findings.

Technical Area 2

Collect and analyze high-resolution data on the structure and function of cortical microcircuits believed to embody the cortical computing primitives underlying key components of the proposed framework

New tools for high-throughput electron microscopy of brain tissue

Zhihao Zheng [1], Rick Fetter [1], John Price [2], Dan Milkie [3], Omar Torrens [3], Eric Perlman [1], Bill Karsh [1], David Peale [4], Harald Hess [1], Albert Cardona [1], Stephan Saalfeld [1], Davi Bock [1]

Abstract:
We are at the beginning of a project to image the entire brain of the fruit fly *Drosophila melanogaster* using serial section transmission electron microscopy (ssTEM). In the service of this goal and anticipated future projects, we are building custom hardware for high-throughput TEM imaging. These tools include:

(1) An interferometric microtome ("iTome") for automated sectioning and pickup of serial 30 nm sections on conventional slot grids;
(2) A next-generation Transmission EM Camera Array ("TEMCA2") for high throughput data acquisition (net ~50 megavoxels/second, 4x4x30 nm/voxel);
(3) A fast piezo-driven TEM sample stage, capable of stepping and settling to nanometer stability in ~25-35 ms (versus ~3 seconds for a conventional stage); and
(4) A multi-sample Autoloader, enabling automated sample exchange without breaking vacuum and 24/7 unattended image acquisition.

To our knowledge, the TEMCA2 plus fast stage combination is the highest-throughput, highest signal-to-noise electron microscopy imaging system currently deployed for use in neural circuit reconstruction. The sample handling disadvantages of ssTEM during sectioning and imaging are well on their way to being mitigated by the iTome and the Autoloader, respectively.

As proof-of-principle application of the TEMCA2 plus fast stage, we recently acquired a pilot image data set from a manually cut series of 4000 ~35 nm whole-brain sections. 500 serial sections were imaged using TEMCA2 and the fast stage at 4x4 nm/pixel, resulting in a 5 TB dataset suitable for developing downstream software components (stitching, registration, intensity correction, automated segmentation). We are currently assessing traceability of fine neurites in the dataset; early results in the medial lobes of the mushroom body are promising.

[1] Janelia Farm Research Campus, Howard Hughes Medical Institute, Ashburn, VA.

[2] Hudson-Price Designs, LLC, Hingham, MA.

[3] Coleman Technologies, Inc., Newtown Square, PA.

[4] Integral Physics to Engineering, LLC, San Diego, CA.

MICrONS Proposers' Day
July 17, 2014

Polarization Focal Plane Sensing with Plasmonic Nanostructures for Functional Neural Imaging

Viktor Gruev and Barani Raman

Washington University in St. Louis, St. Louis, MO.

Recording neural activity using light has opened up unprecedented possibilities for understanding functionality of the nervous system. Probing neural activity via light offers great advantages over electrophysiology. Current state-of-the-art techniques for recording optical signals from neurons require converting electrical signals into optical signals via molecular reporters. However, use of such reporters has major limitations. To overcome these limitations, we have developed a novel technique for imaging intrinsic optical signals of neurons via polarization-imaging sensors.

The polarization imaging sensor is realized by monolithic integration of plasmonic nanowires filters with CMOS imaging array to detect polarization properties of light at every imaged frame. The nanowire polarization filters are optimized and fabricated in the cleanroom facilities at Washington University and capable of detecting linearly polarized properties of the imaged environment.

The polarization contrast of the scattered light from the neural tissue recorded with our sensor is directly correlated with neural activity due to intrinsic changes in neural cells during action potentials. Hence, a direct approach for recording neural activity is to monitor changes in polarization signatures acquired over time and space from neural tissue. This imaging technique allows for imaging of neural activity with high temporal resolution across a large spatial area of the brain. Furthermore, this technique avoids the use of molecular contrast dyes and records the intrinsic optical signal of neural cell.

The high spatial resolution of the sensor allows for imaging large population of neurons simultaneously over large spatial area in the brain. Preliminary experiments conducted *in vivo* on the locust's olfactory circuits downstream to sensory neurons indicate that the polarization imaging system can discriminate various odors by analyzing spatial-temporal patterns using tailored machine learning algorithms.

In the future, we would like to validate this new optical neural recording tool by correlating it with simultaneous intracellular electrophysiology signals obtained from the insect brain. Furthermore, we will use this neural recordings technique to understand how populations of neurons in the insect antenna lobe and mushroom body (analogous to the vertebrate olfactory bulb and olfactory cortex) process information and perform the following pattern recognition problems: robust odor discrimination and perform adaptive filtering to suppress redundancy while retaining sensitivity to novel cues. We will use this approach to determine how the two primary coding dimensions: time and space, are used in biological signal processing. In order to address these questions, we will develop new class of polarization imaging sensor with plasmonic gold structures leading to high polarization sensitivity necessary to improve the signal to noise ratio of this imaging technique. Signal processing algorithms based on machine learning techniques will be developed to determine how time and space are used for information processing in the brain.

Wei-Chung Allen Lee
Harvard Medical School, Boston, MA

Our group is interested in understanding the organizational principles underlying information processing in neuronal circuits. A neuron's function is fundamentally dependent on how it is connected within its network. Therefore, understanding the relationship between connectivity – circuit structure – and cellular function will help us understand how neurons and networks transform information to bring about perception and behavior.

We approach this problem by combining recent advances in non-linear light and high-throughput electron microscopy (EM) to perform detailed mapping of neuronal sensory physiology and network anatomy. To discover fundamental, conserved organizational principles, we apply our methods to multiple neuronal circuits for which we can characterize neuronal function, extract their wiring, and analyze the interplay. I will describe our work in two parts: our biological questions and our recent technological advances.

To begin addressing the complexity of the cortical network we examined the interplay between circuit structure and neuronal function of the sparsely connected pyramidal cell network in the visual cortex. We used volumetric *in vivo* two-photon calcium imaging of a genetically-encoded calcium indicator to measure the time-resolved responses of a large population of identified neurons to an array of visual stimuli in the awake mouse. We subsequently reconstructed the local excitatory neuronal network using large-scale electron microscopy. In contrast to connectivity of local inhibitory circuitry, we find that connectivity between excitatory neurons exhibits functional specificity. Pyramidal neurons with similar sensory physiology are more highly interconnected with one another both within and between neocortical processing lamina. Moreover, similarly tuned neurons converge onto downstream excitatory targets. Finally, we observe topological organization of synaptic input between neurons connected by multiple synapses. This wiring specificity may act as a substrate for computations underlying cortical sensory processing.

Our recent work demonstrates our ability to perform detailed mapping of neuronal network anatomy and sensory physiology. Although our high-throughput transmission electron microscope camera array (TEMCA) has increased the speed of imaging, we've continued to rely on humans for laborious manual sample collection and handling. Automated tape-collecting ultramicrotomes (ATUMs) have produced massive improvements in the reliability and speed of section pickup that allow collection of over 10,000 serial sections with minimal human intervention. However, in its typical form, the Kapton polyimide collection substrate is incompatible with fast TEM imaging. In an effort to synergistically bridge automated sample collection and high-speed TEM imaging, we have begun developing a novel TEM compatible tape substrate: Grid-Tape which we describe here.

Identification of algorithmic principles of intermediate vision using two-photon microscopy and evolving visual stimuli

Kristina J. Nielsen and Charles E. Connor
Zanvyl Krieger Mind/Brain Institute, Department of Neuroscience, Johns Hopkins University

PROBLEM: OBJECT/SCENE VISION is arguably the most remarkable weakness in current machine intelligence. Extracting real-world information from natural visual images has proven to be an essentially insoluble computational problem. As a result, human observers are still required for determining what is in an image and what is happening in an image, and human controllers are still required to navigate intelligently through dynamic real-world environments. The huge gap between machine vision and biological vision makes this a prime target for mining computational algorithms from neural circuits.

EXPERIMENTAL TARGET: NEURAL ALGORITHMS OF INTERMEDIATE VISION remain almost entirely unknown. Research on biological vision has focused on **(a)** pixel-level processing of orientation, color, and motion in primary visual cortex (V1), and **(b)** endstage signals for object identity and other semantic-level information in inferotemporal cortex (IT) and prefrontal cortex. The algorithms that transform (a) to (b) are implemented in intermediate cortical processing stages such as area V4. Understanding these intermediate transformations is the only way to replicate biological vision in computers.

METHOD: We would use 2-PHOTON IMAGING OF NEURAL POPULATION ACTIVITY to infer local circuit algorithms in area V4 of macaque monkeys, an animal model with extremely close functional and anatomical homology to human vision. V4 has only been studied with microelectrode recording from randomly sampled neurons at disparate locations. 2-photon imaging would provide the first opportunity to observe signaling in 100s of densely packed, closely interacting neurons within a local cortical circuit. The basic processing module of the brain is the cortical column, a 0.5 mm diameter column of interconnected neurons. 2-photon imaging is the first technique for observing the information processing carried out by a cortical column.

EXPERIMENTAL STRATEGY: Inferring algorithmic principles from 2-photon data will require a new strategy for evoking a wide range of activity patterns in a V4 column. We would adapt our previous strategy of evolving visual stimuli guided by responses of individual neurons. Here, we would guide stimulus evolution with high-dimensional metrics for strength and variety of population activity patterns. As algorithmic models develop, they would be used to optimize the informational value of evolving stimuli directly.

COMPUTATIONAL DEVELOPMENT: We will recruit team members to develop and implement intermediate visual processing algorithms based on our neural circuit analyses. We envision an iterative process in which neural measurements inspire initial computational models, which can then be used to guide stimulus evolution and test more specific hypotheses about circuit functions, thus constraining models of increasing specificity and complexity.

Extraction of neuronal network processing algorithms from high speed optical recording of membrane potential in genetically defined classes of neurons in motor cortex.

Pieribone Laboratory - Yale University School of Medicine / The John B. Pierce Laboratory, Inc.

To decipher the transfer functions of cortical neuronal circuitry it will be necessary to record the activity a large collections of identified classes of neurons during the cortical network processing underlying animal behaviors. Electrode-based recording methods remain limited in there ability to identify the specific class of neuron which is being recorded, are highly invasive, cannot generally be maintained for extended periods of time and are restricted in the number of neurons that can be recorded simultaneously. Optical recording of defined class of neurons expressing a genetically encoded probe of membrane potential remains an attractive addition, or possibly, alternative to traditional electrode-based recording. Our laboratory has developed the most advanced genetically encoded fluorescent membrane potential probe, ArcLight (*Neuron*, 2012) which remains the most viable probe for *in vivo* monitoring of neuronal activity. Newer probes offer faster or larger signals *in vitro* but no other currently produce viable signals *in vivo*. In addition our laboratory has developed small mobile, fluorescence imaging systems to allow high speed capture of neuronal activity in freely moving rodents.

We are using the combination of these technologies to image motor cortex in behaving rodents while they execute natural behaviors (i.e. walking, lever pressing, righting, etc.). Our goal is to simultaneously record the electrical activity of a range of different cortical cell types during stereotyped behaviors in rodents over a large area of motor cortical area (2 x 2 mm). The activity patterns recorded in these image sequences is then correlated to high resolution positional data for the animal's body and limbs recorded simultaneously using an array of cameras trained on the animal.

We are seeking collaborators to construct computer algorithms that determine the animal's three dimensional position in time and are able to (if possible) determine body position and movements based on cortical electrical activity. We can use *cre*-based viral vectors to drive expression in a number of available mouse models that target subtypes of cortical neurons. We will compare the electrical activity of a range of different defined cortical cell types during the repeated execution of highly stereotyped motor actions.

Estimating Cortical Graph Structure

Carey E. Priebe, Randal Burns
& a cast of dozens ...
Johns Hopkins University

June 30, 2014

Many contemporary theories of neural information processing suggest that the neocortex employs algorithms composed of repeated instances of a limited set of computing primitives. There is a recognized need for tools for interrogating the structure of the cortical microcircuits believed to embody these primitives. The *cortical column conjecture* suggests that neurons in the neocortex are connected in a graph that exhibits motifs representing repeated processing modules.

We consider theory and methods for inference regarding the structure of the cortical graph. We model the cortical graph as a hierarchical stochastic block model (HSBM) with induced subgraphs which are themselves independent stochastic block models. Our focus is on extracting, and then estimating the structure of, the cortical graph, for the purpose of subsequent modeling and algorithm development.

Let $G \sim \text{HSBM}\big((R_1, n_1, \rho_1, B_1), \cdots, (R_m, n_m, \rho_m, B_m), P_{m\times m}\big)$.

Consider a three-step algorithm:
(1) *graph clustering* to find the induced subgraphs;
(2) *clustering graphs* to find the repeated motifs;
(3) *clustering graph clusters* to estimate the motif structures.

Theorem:
Under suitable eigenvalue assumptions on the HSBM, our three-step algorithm yields consistent estimates for the cortical graph parameters.

Reconstructing neural circuits via social computing

Sebastian Seung
Princeton Neuroscience Institute and EyeWire.

July 2, 2014

Artificial intelligence via machine learning has radically improved the accuracy of neural circuit reconstruction from serial EM images, but human intelligence is still needed. At the present time, AI speeds up humans but does not replace them, because serial EM images contain many locations that are difficult even for human experts to disambiguate. EyeWire is an example of a system that achieves fast and accurate neural circuit reconstruction by combining human and artificial intelligence, an approach known as social computing. EyeWire is scalable to large numbers of humans, with solutions for the problems of crowd wisdom, crowd learning, and incentives.

Ekaterina Taralova,
Columbia University,
etaralova@cs.cmu.edu

Abstract

To create machine learning algorithms based on the function and activity of cortical microcircuits we need to build computational models of visual stimuli and of the spatio-temporal firing patterns they elicit in the central nervous system of animals and humans. To achieve this goal, we will combine the pioneering work in calcium imag developed in Dr. Rafael Yuste's lab with my expertise in computer vision and machine learning. We will harness experimental data o neuron activity and relate it to the visual stimu that bring about this activity.

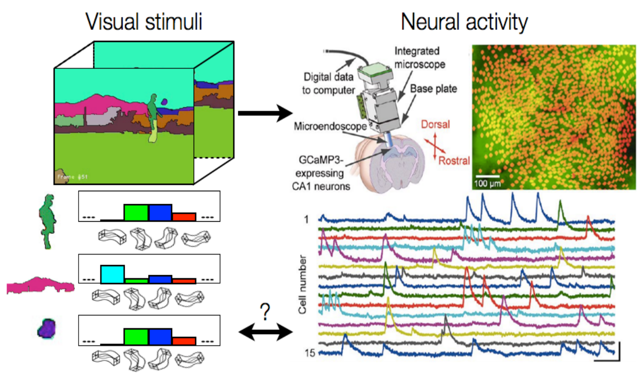

First, we need powerful representations for the visual stimuli. Building image and vid representations typically involves four steps: feature extraction, quantization, encoding, and pooling. While there have been large advances in feature extraction and encoding, the questions of how to quantize video features and what kinds of regions to pool them over have been relatively unexplored. To tackle the challenges present in visual data, it is necessary to develop robust quantization and pooling methods.

In the task of action classification in videos, I have proposed a new method, Source Constrained Clustering, which quantizes features into a codebook that generalizes better across actions. The main insight is to incorporate readily available labels of the sources generating the data, e.g., the people who performed each action.

In the pooling step, it is common to pool feature vectors over local regions. The regions of choice include the entire video, coarse spatio-temporal pyramids, or cuboids of pre-determined fixed size. A consequence of using indiscriminately chosen cuboids is that widely dissimilar features may be pooled together if they are in nearby locations. It is natural to consider pooling video features over supervoxels, for example, obtained from a video segmentation. However, since videos can have a different number of supervoxels, this produces a video representation of variable size. I have proposed a fixed size video representation, Motion Words, where we pool features over video segments.

The ultimate goal of video segmentations is to recover object boundaries, often grouping pixels from regions of very different motion. However, in the context of Motion Words, it is important that regions preserve motion boundaries. I have proposed a supervoxel segmentation, Globally Consistent Supervoxels, which respects motion boundaries and provides better spatio-temporal support for Motion Words.

Finally, it is essential that the representations we build are interpretable, that is, we want to be able to visualize and understand *why* videos are similar. Motion Words enable localization of common regions across videos in both space and time. This gives us the power to understand which regions make videos similar.

Andreas Tolias[1], Mattias Bethge[2], Liam Paninski[3], Michael Roukes[4], Thanos Siapas[5]

[1] Department of Neuroscience, Baylor College of Medicine, Department of Applied and Computational Mathematics, Rice University

[2] Werner Reichardt Centre for Integrative Neuroscience and Institute of Theoretical Physics, University of Tübingen, Germany [3] Bernstein Centre for Computational Neuroscience, Tübingen, Germany

[3] Department of Statistics, Grossman Center for the Statistics of the Mind, University of Columbia

[4] Department of Physics, Mathematics and Astronomy, Caltech

[5] Division of Biology, Division of Engineering and Applied Science, Caltech

Our aim is to integrate experimental and theoretical approaches in order to develop novel machine learning algorithms inspired by cortical microcircuits.

Our experimental aim is to understand the rules by which different types of neurons in the neocortex connect to each other and work together to process information. We want to determine what constitutes the elementary circuit motifs in the neocortex and to characterize their structure and the computations that these modules implement. We combine electrophysiological, imaging, and molecular tools with behavioral and computational approaches to dissect the functional architecture of inhibitory and excitatory microcircuits in the visual system of mice and monkeys. From an anatomical perspective we are mapping out the detailed wiring diagram of the cortical microcircuit using high-throughput multi-cell patch clamp recordings. This enables us to decipher the local circuit diagram including information about synaptic strength and characterize rules of plasticity. Using electrophysiological and imaging methods we characterize the activity structure of large populations of neurons to understand the nature of the neural code. To this end, developed an *in vivo* 3D high-speed, random-access two-photon microscope that is capable of simultaneous 3D motion tracking. This enables us to record the activity of nearly all of the hundreds of cells (up to 500 neurons) in small volumes of the cortex to characterize the structure of microcircuit population activity during visual processing. This method currently provides the largest number of neurons that can be recorded densely *in vivo* in 3D at such high rates in light scattering tissue.

Our team has strong expertise in the development of statistical methods to analyze the organizing principles of high-dimensional neural data. Our ultimate goal is to apply the principles and canonical algorithms we learn from cortical circuits to build the next generation artificial neural networks that will be based on a set of new computational primitives inspired by neuroscience. We are particularly interested to study the role of feedback in neural computation and apply this in deep neural network architectures to solve machine-learning problems that are currently difficult with current state of the art machine learning methods.

Optical technologies for high-speed, single cell specific functional imaging and control of large neuronal population

Alipasha Vaziri *- Research Institute of Molecular Pathology (IMP) & University of Vienna*

Knowledge of structural connectivity in neuronal circuits is necessary to understand how sensory information is represented and processed by neuronal circuits. However, as some examples of well-characterized neuronal architectures illustrate, structural connectivity is not sufficient to predict how stimuli are mapped onto neuronal activity patterns and how the collective dynamics of the network leads to behavior. Addressing this challenge, a prerequisite for discovering the computational and algorithmic principles used by the brain, has been hampered by lack of appropriate tools that allow application of defined high resolution spatiotemporal excitation patterns while simultaneously capturing the dynamics of the entire network. Our recent work has contributed towards closing this gap and moving towards a dynamic map of neuronal circuits.

The spectral width of a femtosecond pulse can be used as an independent degree of freedom to "sculpt" the spatial light distribution in the sample and generate axially highly localized sheet-like light distributions. This has been exemplified by the technique of temporal focusing. We have used this technique to demonstrate scanning-less two-photon optogenetic activation of individual neurons embedded in a pool of genetically identical neurons within 1-2 ms [1]. Further, such sculpted light sources can be combined with galvanometers or spatial light modulators to generate arbitrary spatiotemporal excitation patterns on a neuronal population [2].

Using the same principle of temporal focusing we have also developed a high-speed two-photon technique for brain-wide Ca^{2+}imaging [3]. Thereby highly localized discs of light with diameters up to ~120 μm have been generated in the samples which were rapidly scanned to cover a volume. Using this approach we have demonstrated whole-brain volumetric Ca^{2+} imaging in *C. elegans* at 5Hz volume rate [3].

In order to functionally image even larger volumes at higher speed we have recently established light-field microscopy in combination with 3D deconvolution [4]. Thereby we have demonstrated volumetric Ca^{2+}-imaging of the entire larval zebrafish brain at 20Hz. Using cell identification and segmentation approaches we could follow the activity of ~5000 neurons distributed throughout the brain while animals were exposed to olfactory stimulations [4]. The extension of both techniques to larger volumes, to scattering media and to other systems including rodents is being currently pursued.

Ultimately, the combination of the high speed functional imaging tools with our optogenetic methods will be crucial to move towards a dynamic map of neuronal circuits and the prerequisites for discovering brain algorithms, understanding of cortical computation and to exploit these findings to enhance machine intelligence.

1. Andrasfalvy, B., et al., *Two-photon Single Cell Optogenetic Control of Neuronal Activity by Sculpted Light*. PNAS, 2010. **107** (26):11981-6.
2. Losonczy, A., et al., *Network mechanisms of theta related neuronal activity in hippocampal CA1 pyramidal neurons*. Nat Neurosci, 2010. **13**(8): p. 967-72.
3. Schrodel, T., et al., *Brain-wide 3D imaging of neuronal activity in Caenorhabditis elegans with sculpted light*. Nature Methods, 2013. **10**(10): p. 1013-1020.
4. Prevedel, R., et al., *Simultaneous whole-animal 3D imaging of neuronal activity using light-field microscopy*. Nat Meth, 2014. **11**(7): p. 727-730.

Top Down Modeling and Efficient Automatic Inference for Anatomy and Simulation

Frank Wood — fwood@robots.ox.ac.uk — http://www.robots.ox.ac.uk/~fwood/
Associate Professor, Department of Engineering Science, University of Oxford

Understanding the brain and emulating its function requires both an improved understanding of anatomy and a better understanding of the functional and computational behavior of neural circuits. We propose an effort that addresses both using advanced automatic inference technology.

Our proposed effort is rendered feasible by the existence of advanced tools for general purpose inference currently under development at Oxford (recently matured enough for our purposes). The development of these tools, called probabilistic programming systems (Anglican [Wood et al., 2014] and Probabilistic-C [Paige and Wood, 2014]), was inspired by the need to define and represent complex shape and space priors, particularly those required to improve computational neuroanatomy via top-down, biologically parameterized regularisation. We can now this now.

These tools can also automatically invert arbitrary forward code-based simulators, for instance whole-circuit neural dynamics models, given observed data.

Top-Down Regularization for Neuroanatomy

In close collaboration with partner neuroscientists, we propose to develop expressive, biologically parameterised, top-down model-based regularization for automated anatomy reconstructions from high resolution neural imagery.

Recent work in semi-automated neuroanatomy are accurate and promising [Roberts et al., 2011, Unger et al., 2009, Jarrell et al., 2012, Bock et al., 2011]. In work towards full automation, feature engineering significantly shapes the state of the art [Jain et al., 2007, Andres et al., 2008, Venkataraju and Paiva, 2009]. We propose to build on these methods by developing expressive top-down model-based regularization for neural process tracing that derives from recent work on fully unsupervised tracking [Neiswanger et al., 2014] and that is implemented and executed in our probabilistic programming systems.

Automatic Neural Circuit Simulator Inversion

Another compelling use case of probabilistic programming tools is the automatic inversion of simulators. This quite new capability means that testing model-based hypotheses about, for instance, circuit function can involve no additional effort beyond writing a forward model of the circuit function and using probabilistic programming tools to invert it given features of observed function. In close collaboration with functional anatomists and biologically inspired computing partners we propose to test hypotheses about neural circuit function via automatic inversion of such simulators.

As an example of how this works, consider an ongoing collaboration of the same sort in a different field. A collaborator had already written a forward simulator for the stability performance of a particular kind of ocean-going, re-positionable oil platform when connected to the seafloor. By adding just tens of lines of probabilistic programming code we were able to automatically invert this existing simulator. Running the resulting probabilistic program yielded the posterior distribution of input parameters to the simulator whose values were known only up to unsatisfactory tolerances before the inclusion of observation information. More crucially, knowing the distribution of these values allowed novel explorations of scientific hypotheses. Efficient exploration of circuit function hypotheses via automatic simulator inversion is much the same and is the goal of this effort.

Technical Area 3

Generate computational neural models of cortical microcircuits informed and constrained by this data and by the existing neuroscience literature to elucidate the nature of the cortical computing primitives

Biologically realistic circuitry from hippocampal neuron properties

David J. Hamilton, Alexander O. Komendantov, Christopher L. Rees, Charise White, Diek W. Wheeler, Giorgio A. Ascoli*

Center for Neural Informatics, Structures, and Plasticity; Krasnow Institute for Advanced Study, George Mason University; Fairfax, VA (*Ascoli@gmu.edu)

A list of neuron types is necessary to build biologically realistic simulations of brain circuits and function. Yet even for the rodent hippocampal formation, arguably among the most intensively investigated neural systems, we lack a complete and accessible catalogue of neuron types and properties. Critically, much hard-fought knowledge is not readily accessible, but rather scattered in an oceanic and ever-growing literature. Systematic organization of all existing neuron type knowledge might revolutionize neuroscience akin to the impact the Period Table of the Elements had on chemistry 150 years ago. What organizing principles are suitable for a putative 'periodic table of the neurons'? Neurons are commonly described based on their observed properties, often along three main feature sets: morphological (axonal and dendritic shape, location, and connectivity), physiological (rhythms, spiking activity, synaptic plasticity), and molecular (neurotransmitter, protein and gene expression, transcription factors).

We are systematically organizing the published information on hippocampal neurons into a comprehensive knowledge base (Hippocampome.Org), including the "closed-loop" of dentate gyrus, CA3, CA2, CA1, subiculum, and entorhinal cortex, and focusing on normal, adult or adolescent rodents. Each neuron type is specified in terms of the main neurotransmitter released and known aspects of its anatomy, physiology, and biochemistry. Every property is documented with a pointer to, and excerpt from, relevant published evidence. Initial anatomical properties include the sub-region and layer locations of the axons, somata, and dendrites, as well as any reported target specificity. Physiological properties include passive characteristics (e.g. input resistance and time constant), firing rate and patterns elicited by somatic current injection, and sub-threshold responses. Molecular properties include expression of calcium-binding proteins (PV, CB, CR), neuropeptides (CCK, SOM, VIP, NPY), and several other major biochemical markers.

The simple description of the axonal and dendritic spatial patterns for all neuron types allows the derivation of the full "potential connectome" of the hippocampal circuit. Computational simulations demonstrate that this kind of information powerfully constrains biologically realistic neural network models. Specifically, our recent results suggest that potential connectivity may provide a neural correlate for the gating of memory formation by background knowledge. This mechanism ensures superior learning capability for real associations relative to spurious co-occurrences.

Biomimetic multiscale modeling of cognitive architecture and active sensing

JC Principe, S Dura-Bernal, CE Schroeder, A Keil, M Ding, WW Lytton

We propose a Bayesian approach to sensory processing, using a hierarchical, distributed architecture of dynamic processing elements. The network self-organizes, learning parameters and features that explain the input data. Several key features exploit and extend predictive coding schemes: 1) bidirectional (top–down and bottom–up) processing enable perceptual inference using both sensory data and empirical beliefs about causes from higher layers; 2) dynamic components are at the core of the model, allowing beliefs about temporal context to influence perceptual inference; 3) only salient features of the input data are stored, forming a compressed and sparse representation; 4) re-utilization of the same model within and across hierarchical levels allows for efficient software/hardware implementations, reminiscent of cortical microcircuits.

An instantiation of this architecture was employed to recognize objects in videos using multiple nodes in a layered tree structure. Each node consisted of two modules: a state–space model, which extracted features, and an invariance-learning model, which inferred causes by pooling over states. Inference and parameter learning both minimized the same energy function with respect to either the states or the causes. Preliminary results show that the architecture has features that resemble the functional organization of lower visual cortex. The model showed better classification by leveraging temporal and top-down contextual information during inference, and was able to disambiguate a synthetic video from correlated noise.

We further incorporate a key biological feature utilized by the visual system: eye movement control (active sensing). This deviates from the classical approach in vision research, where stimulus onset is regarded as the critical reference point in time. Instead, we now conceptualize visual parsing and object identification as a continuous loop where a saccade marks the onset of a new sampling episode. Including the motor and sequencing components is an enormous asset, because rather than indirectly inferring the brain's sampling strategy, we can directly observe it through the spatiotemporal pattern of eye movements.

We will use a multi-level, multi-method approach (single cell, CSD, LFP, fMRI), to test the model predictions, by characterizing the spatial and temporal dynamics of these interactive processes, and inform and constrain computational modeling. We aim to build biomimetic computational models that map these Bayesian computations onto the known anatomy and physiology of brain circuitry, linking the different spatial and temporal scales in the brain. We will exploit our previous biomimetic models of several neocortical regions (S1, V1, M1, PFC), where multiscale modeling allowed us to capture the complexity of dendritic processing (micron scale) in the context of a large networks (mm scale) and networks of networks (cm scale). Our prior models have accurately reproduced physiological properties observed in vivo, including firing rates, LFPs, oscillations and traveling waves. Here, we will continue to focus on ensemble function (e.g., co-firing) rather than on individual spike trains. We consider oscillatory activity in neural assemblies as a relevant process that emerges from, but also constrains, interactive neural dynamics. We replace the concept of vision as a unidirectional feedforward sweep through a hierarchy of areas with a model of visual areas connected on a many-to-many basis, with many strong short feedback circuits complementing some sequential connectivity in both directions. As in the past, we will use information theoretic tools, including normalized transfer entropy (nTE), Granger causality and graph theory, to quantify information transfer, and gain insights into the relation between network structure and function. These analyses will help us build novel machine learning algorithms based on the computing principles of cortical microcircuits.

Sensory integration and control for on-the-fly Learning

Marwan Jabri, Neuromorphic LLC, marwan.jabri@neuromorphic.com

Abstract - Neuromorphic LLC is a startup founded by veteran scientists and engineers, with over 60 man-years in academic research and commercialization. Our areas of expertize include:

- Analog and digital neural computing integrated circuits and system [Mhaidat *et al*, Coggins *et al*], with on chip learning by the means of perturbations [Jabri & Flower, Flower & Jabri].
- Sensorimotor control using computational models of the basal ganglia, ventral tegmental area and the cerebellum [Huang *et al*].
- Computational models of area MSTd using unsupervised learning [Jabri *el al*]
- Computational model of the superior colliculus and cerebellum for gaze shift [Wang *et al*]
- Computational model of a V1-V4-PIT-AIT visual pathway with unsupervised learning for online and on-the-fly learning of objects and subjects [Patent pending]

We are also experienced in building and deploying products. In the past ten years we have developed and deployed products that are today in hundreds of millions of handsets and smartphones, and more than 100 mobile operator infrastructure (software and hardware systems).

We are interested in sensorimotor integration and control that enables robots to exhibit *compelling human robot user experience*. Critical to this is the ability of a robot to learn and perceive objects/subjects in its environment. The learning has to be performed on the fly, and the number of objects/subjects to be learnt is in the tens to the hundreds.

References

K. Mhaidat, M. Jabri, D. Hammerstrom, “Representation, Methods, and Circuits for Time Based Conversion and Computation”, Inter. J. of Circuit Theory and Appl., Wiley, 39, 299-311, 2011.

R. Coggins, M. Jabri, B. Flower and S. Pickard, "An Analogue VLSI Neural Network for Intracardiac Morphology Classification", IEEE J. of Solid State Circuits, 30 (5), 542-550, 1995.

M. Jabri and B. Flower, "Training Analog neural networks using weight perturbation", in IEEE Transactions on Neural Networks, Vol. 3, No 1, pp. 154-157, January 1992.

B. Flower and M. Jabri, "Summed Weight Neuron Perturbation: a O (N) improvement over Weight Perturbation", NIPS5, Morgan Kauffmann Publishers, pp 212-219, 1993.

M. Jabri, J. Huang, O. Coenen and T. Sejnowski, "Models of basal ganglia and cerebellum for sensorimotor integration and predictive control." ISOP, Sensor Fusion and Decentralized Control in Robotic Systems III, 2000, pp 125-137.

M. Jabri, K. Park, S. Lee, and T. Sejnowski, “Properties of Independent Components of Self-Motion Optical Flow”, Proceedings of ISMVL, IEEE Computer Press, 2000, pp 355-362.

X. Wang, J. Jin and M.A. Jabri; Neural network models for the gaze shift system in the superior colliculus and cerebellum, Neural Networks, 15, pp: 811-832, 2002.

M. Jabri and X. Li; A Human Companion System, US patent Pending; Dec 25, 2013.

Simulation Environments and Tools for Constructing Neural Models of Cortical Microcircuits

Jeffrey L. Krichmar, Nikil Dutt, Michael Beyeler, Kris Carlson
Department of Cognitive Sciences, Department of Computer Science
University of California, Irvine, Irvine CA, 92697-5100 USA

Our group has developed a suite of tools for simulating brain-scale cortical circuits. These spiking neural networks (SNNs) are constrained by physiological details, such as synaptic conductances, spike-timing dependent plasticity, neuromodulation, and firing patterns, as well as anatomical details, such as thalamocortical projections, cortico-cortical connectivity, and cortical column microcircuitry. The simulation environment can be used to construct, and run SNNs quickly and efficiently by leveraging the parallel processing power of GPUs. For example, we constructed a motion perception model based on the cortical visual stream that exhibited both component and pattern motion selectivity found in cortical area MT, generated speed tuning curves that are in agreement with electrophysiological data, and reproduced behavioral responses in a forced choice task (*1*). We developed a spiking neural network model of visual cortex area V1 and thalamus that showed how cholinergic neuromodulation can affect both top-down and bottom up attention (*2*). In a laminar model of dorsolateral prefrontal cortex (dlPFC), we demonstrated how varying the levels of dopamine (DA) and norepinephrine (NE) in dlPFC could affect working memory (*3*). When DA and NE were outside the optimal levels, there was noise in the circuit, abolishment of sustained activity, and behavioral deficits in working memory tasks. Despite the success our group and others have had in creating these cortical circuits, tuning the enormous number of open parameters in these models becomes a difficult challenge as the desire for more biologically realistic cortical circuits increases. Therefore, we have developed a framework that utilizes evolutionary algorithms to automate the parameter tuning of SNNs (*4*). The objective function is based on known neurobiological constraints and therefore is a general-purpose tool for constructing cortical and other neural circuits. Our group has expertise creating cortical circuits constrained by neurobiology and analyzing how these cortical primitives give rise to cognitive behavior. Our simulation environment, tuning framework, source code, and analysis scripts are publicly available at: http://www.socsci.uci.edu/~jkrichma/CARLsim/.

1. M. Beyeler, M. Richert, N. D. Dutt, J. L. Krichmar, Efficient Spiking Neural Network Model of Pattern Motion Selectivity in Visual Cortex. *Neuroinformatics*, (2014); published online EpubFeb 5 (10.1007/s12021-014-9220-y).
2. M. C. Avery, N. Dutt, J. L. Krichmar, Mechanisms underlying the basal forebrain enhancement of top-down and bottom-up attention. *The European journal of neuroscience* **39**, 852-865 (2014); published online EpubDec 5 (10.1111/ejn.12433).
3. M. Avery, N. Dutt, J. L. Krichmar, A large-scale neural network model of the influence of neuromodulatory levels on working memory and behavior. *Frontiers in Computational Neuroscience* **7**, (2013); published online Epub2013-October-3 (10.3389/fncom.2013.00133).
4. K. D. Carlson, J. M. Nageswaran, N. Dutt, J. L. Krichmar, An efficient automated parameter tuning framework for spiking neural networks. *Frontiers in Neuroscience* **8**, (2014).

Dhireesha Kudithipudi, Ph.D.
NanoComputing Research Lab, Rochester Institute of Technology
Email: dxkeec@rit.edu
Office: (585)-475-5085
Cell: (210)-262-6950

Subcortical brain structures, *e.g.* the thalamus, play a central role in sensory integration, prioritization, attention, and even consciousness. Although most studies in neuromorphic computing have focused on sensory transduction or emulation of circuits within the cerebral cortex, changes at the subcortical levels such as the diencephalon also play a fundamental role in sensory plasticity. Some researchers even refer to the thalamus as the seventh cortical layer. For example, of the ~50 nuclei in the thalamus only 5 are used for sensory relay and the remaining 45 nuclei participate in the complex cortical-subcortical networks and have no primary sensory inputs. These networks of structures are crucial to perception and are modulated through attention via both subcortical and cortical inputs. Studies in the related fields of neuroscience and neuropsychology also are recently converging to show the importance of thalamus in learning. On the other hand, a vast majority of today's computational neuroscience models rely on associative memory models, which are important for intelligence, but it is still unclear how memory feeds back for early prediction and decision-making. Few research groups have focused on modeling specialized signaling properties of thalamic relay neurons or large-scale software models of thalamocortical systems. While these efforts are orthogonal within a subdomain, system designers have failed to capture the essential interplay between thalamic nuclei and cortical brain regions that are necessary for achieving complex behavior and seamless interaction.

We are currently focused on designing the compute primitives for cortical processor (shown in Fig. 1) based on spatio-temporal processing. Initial studies have been performed to capture the rich temporal dynamics of thalamocortical interconnectivity and most particularly the reciprocal nature of the thalamocortical neuronal loop function and the interface modules to the cortical processor.

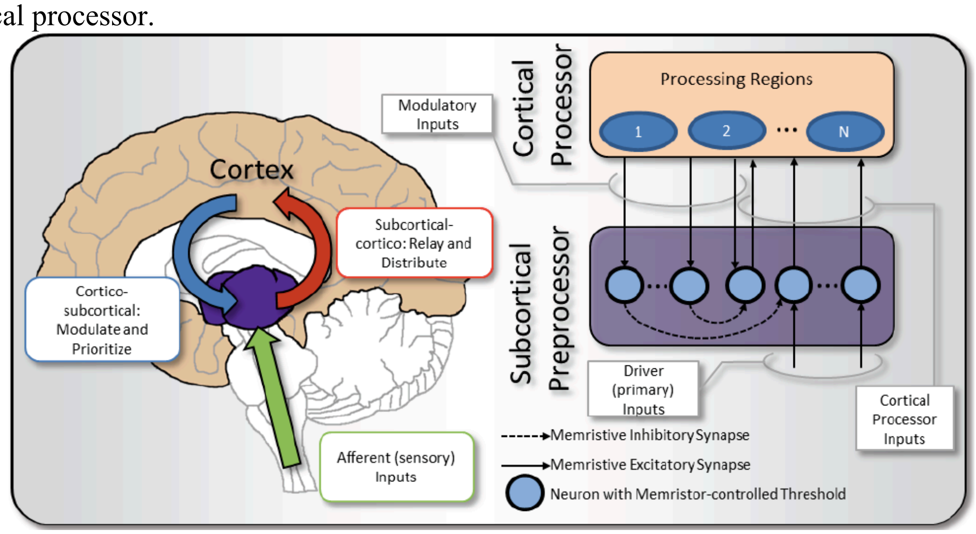

Figure 1: Congruence between a human brain and the proposed high-level processing system. The subcortical core preprocesses and makes predictions based on afferent signals. It also facilitates communication channels between different regions of the cortex.

PIs: Nima Mesgarani & Shihab Shamma

Columbia University, University of Maryland

Title: Biologically informed artificial neural network models

Abstract:

Understanding and modeling the cognitive abilities of the brain requires high-resolution neuroimaging techniques that have only recently become available. Concurrently, advances in computing power have led to the re-emergence of artificial neural network (ANN) models as the dominant paradigm in machine learning. These analytically intractable models have generated both enthusiasm and skepticism in the field, reflecting their widespread albeit cryptic success. While ANNs are loosely modeled after the biological systems, they lack several known neural mechanisms that significantly contribute to the robust cortical representation of speech in the auditory cortex. Our goal in this proposal is to form an integrated research approach where reverse-engineering methodologies are used to determine the computation and organization of artificial neural networks leading to new biologically informed models simulating the functional properties of neural mechanisms. These models will show superior predictive power, reduce the performance gap with biological computing, and advance our knowledge of how the brain represent and processes speech.

This is an interdisciplinary project, which lies at the intersection of neuroscience, speech engineering, and linguistics. The **innovation** of this project lies in its integration of methods and expertise across various disciplines, including system identification, signal processing, neurophysiology, and systems neuroscience. The **long-term goal** of this research is to define and model the neural basis of speech perception in humans. The aim of this proposal is to **analyze** and **transform** the artificial neural network models to accurately reflect the computational and organizational principles of biological systems through **two specific objectives**: I) to determine the acoustic and phonetic feature encoding, connectivity, and the hierarchical emergence of categories and invariance in artificial multilayer (deep) neural networks II) to determine the functional properties of neural mechanisms distinguishing biological and artificial systems, and incorporate them in new models.

Understanding and modeling how the neural mechanisms contribute to human's ability to robustly comprehend speech is a fundamental open scientific problem. An accurate computational model that can adequately explain the neural transformation and human cognition will have overarching impact in many disciplines including speech engineering, cognitive and systems neuroscience, speech and language pathology, and neurolinguistics.

Cortical primitives for statistical inference in a structured world

Xaq Pitkow[1,2], Ankit Patel[1], Rich Baraniuk[1]
[1]Rice University, Department of Electrical and Computer Engineering
[2]Baylor College of Medicine, Department of Neuroscience

Our overarching hypothesis is that cortical computations use deep, nonlinear recurrent networks to create internal models of the natural world and infer hidden states, by reformatting and integrating sense data while respecting uncertainty. Our theories about this involve three themes:

- The brain looks for *change*.
- The brain weighs *uncertainty*.
- The brain is deeply *nonlinear*.

Change: Animals seeks out actionable information, which is distinguished from the general background due to some sort of change over space or time. Thus change detection is a fundamental operation performed by neural circuits. However, there are many variants of change detection, ranging from near-complete differentiation to fractional differentiation, implemented by circuit motifs such as delayed feedforward inhibition (Fig A) or recurrent inhibition (Fig B). Detailed knowledge of cortical circuitry can reveal new variants that evolution has selected as particularly useful.

Uncertainty: Machine learning has advanced enormously by acknowledging the importance of weighing evidence probabilistically. Ample evidence shows that animals also make decisions based on probabilistic reasoning. These two lines of evidence suggest that an understanding of how the brain represents uncertainty could lead to new primitives for machine intelligence.

Bayes' rule trivially ensures that all neural representations $\mathbf{r}$ about the world s are probabilistic, according to $p(s|\mathbf{r}) \propto p(\mathbf{r}|s)p(s)$. A non-trivial probabilistic computation would have to combine information between neural populations in a way that is consistent with probabilistic inference. This creates two restrictions: the combination must depend on at least two aspects of each population activity, such as an estimate and its uncertainty (Fig C, D); and the resultant activity must accumulate in accordance with the reliability of its inputs. Past experiments have reported neural representations that are consistent with these restrictions. Other probabilistic operations, notably including marginalizing over nuisance variables (arguably the most difficult and fundamental aspect of natural inference), can in simple cases be implemented by arbitrary random networks.

Deep nonlinearities: Natural tasks are highly nonlinear functions of sense data. Nonlinear networks can learn to solve these tasks by careful bending of the response space. Deep networks have proven to be more efficient at solving real-world nonlinear tasks than shallow architectures. However, existing deep networks have been designed *ad hoc*. Our team has proposed principles for deep feedforward architectures based on probabilistic semantics for a hierarchical world. We aim to extend these ideas to fully recurrent networks, accounting for statistical structure by appealing to graphical modeling, and accounting for the limited flow of information (Fig E).

Exact inference in graphical models is intractable, but approximate schemes based on message-passing algorithms provide a plausible alternative (Fig F). Neurons are naturally described as message-passing devices, but the particular manner of synthesizing messages depends on the nonlinear response properties of neurons. Combining our statistical inference framework with biological data about cortical connectivity and nonlinear responses will suggest new cortical computing primitives as the approximations that the brain uses for inference.

Title: Cortical Networks Research at IBM

Name: Guillermo Cecchi*, James Kozloski, Janusz Marecki, Mattia Rigotti*, Mark Ritter*, Gerald Tesauro*, Roger Traub and Yuhai Tu (* MICrON Proposers Conference participants)

Organization: IBM T. J. Watson Research Center

Abstract

Neural Modeling: IBM Research has strong effort in neural modeling. Coupled with the right high resolution brain data, these models can serve to elucidate cortical computing primitives, inspire and help develop novel Machine learning (ML) algorithms. There are three areas of modeling that are relevant for MICrON:

(1) Detailed mesoscopic thalamocortical models – A model a single column with thousands of realistic neurons was developed (R. Traub et al, J. Neurophysiol, 2005). The model has been successful in explaining and predicting cortical dynamics such as oscillations and seizures (e.g., Caracedo et al, J. Neurosci., 2013). Realistic multi-column model is being developed.

(2) Neural network models linking neural activity and behavior – We developed a neural network model that extracts the hidden context variables (M. Rigotti et al, NeuroImage, 2010). We are studying the neural basis for context-dependent data representation (M. Rigotti et al, Nature, 2013), and its application for novel context-dependent ML algorithms.

(3) Modeling of global architecture and dynamics for perception learning -- We have implemented linear and non-linear predictive models of large-scale spatio-temporal imaging data including fMRI and calcium imaging (Neuroimage 2011, IEEE 2011, JMLR 2013) in HPC (Blue Gene). We are studying how the predictive dynamical components can be incorporated in a novel ML framework to mimic brain's perceptual learning capability.

Data needed: To constrain and verify these models, **we seek to collaborate with experimental neuroscience labs** to get the relevant brain data on: **1) Anatomy** -- EM reconstructions, DTI, axonal tracing to identify inter- and intra-area connectivity; **2) Function** -- multi-area electrode array recordings, calcium imaging, high-resolution fMRI; **3) Behavior** -- simultaneous measuring context-dependent behavior performance and neural activity in higher animals.

Machine Learning: IBM Resaerch has extensive expertise in all areas of ML, including neuroscience-related algorithmic innovations and applications, such as reinforcement learning (RL), Informax, ML for neuroimaging data analysis, and Deep neural network learning. Our general goal related to MICrON is to develop new machine learning (ML) algorithms to carry out information processing tasks in complex environments based on novel ML frameworks with insights and constraints from neuroscience (cortical and thalamocortical circuits and their dynamics) underlying behavioral and context based perception in human (and higher animals).

Sensory processing, plasticity and pattern recognition in a complex 'mini' brain

Brian H Smith[1], Ramon Huerta[2], Maxim Bazhenov[3]

[1]Arizona State University, [2]University of California San Diego, [3]University of California Riverside

Our goal is to replicate through computational modeling and machine learning algorithms a network equivalent to the cortical structures that integrate and learn about multimodal information: olfactory, tactile, and visual modalities. We use an animal model – the insect (honey bee) brain – that contains approx 10^6 neurons, and it behaviorally gives rise to learning about and generalization among stimuli important for locating food and navigation through a complex visual, tactile, auditory and olfactory world. An insect brain is small enough that we stand a reasonable change of modeling the entire brain, including regions that perform computations similar to areas of the mammalian brain (e.g. the thalmo-cortical-hippocampal system). One specific anatomical region is the Mushroom Body (MB). It contains 30% of the total neural mass in the brain, and it performs computations similar to the mammalian cortical areas. Our group is capable of recording the spatio-temporal activity patterns from large numbers of cells in the MBs simultaneously in the input and output layers. We then fit models that relate the input-output activity recorded *in vivo*. We also have built pattern recognition algorithms that utilize the structural organization, neural coding and learning capabilities (nonassociative, associative, operant) observed in the Mushroom Bodies that can achieve competitive or better performance than state-of-the-art algorithms for a small number of examples. Our strength also lies in the ability to relate machine learning algorithms to realistic conductance based models of neurons. By relating realistic models, anatomical structure, *in vivo* spatio-temporal recordings and behavioral conditioning we have unique capabilities to pinpoint the main principles that achieve efficient learning of spatio-temporal input patterns. We also think that it is critically important to understand how novel pattern recognition algorithms can cope and take advantage of the nonlinear dynamic nature of the natural stimuli. Most of the state-of-the-art algorithms assume static data or at best linear models in the form of primitive linear dynamical systems of hidden Markov models. The reality of neural processing is that these systems are highly dissipative nonlinear dynamical systems. Ignoring this complexity hampers the ability to truly understand how the human brain computes natural data.

MICrONs Poster Abstract

The Georgia Tech Research Institute (GTRI), the applied research arm of Georgia Institute of Technology, in collaboration with a faculty member from the resident instruction side of Georgia Tech, are engaged in research in combining neuroscience results with modeling and machine learning . The team consists of Elizabeth Whitaker (GTRI- Artificial Intelligence (AI), modeling, machine learning), Ethan Trewhitt (GTRI-Software Engineer, modeling, machine learning), Eric Schumacher (School of Psychology–Neuropsychology). We worked together on the IARPA ICARUS project, and on an internal IR&D project developing Brain-based Architectures for Training. We are also working together on the DARPA Narrative Networks projects, integrating neuroscience findings from the fMRI studies at the School of Psychology into narrative composition tools that can be used to experiment with and understand how stories influence behavior.

The objective of the ICArUS Program was to construct brain-based computational models of the process known as sensemaking. Sensemaking, a core human cognitive ability, underlies intelligence analysts' ability to recognize and explain relationships among sparse and ambiguous data. The GTRI/GT team produced a spiking neuron model of a brain component which was integrated into the larger ICArUS team's overall model of sensemaking. The model of the anterior cingulate cortex was configured to interact with other ICArUS components to solve the ICArUS challenge problems. The ACC was configured for uncertainty monitoring, conflict monitoring, recognition of surprise and reward and punishment recognition and learning for performing its role in the challenge problem. GTRI provided the modeling, software design, prototyping and integration with the other model components, while the neuroscience domain expertise, joint model design and vetting were done by the Georgia Tech School of Psychology neuropsychologist.

The objective of the Brain-based Architecture for Training research was the development of an architecture and approach for conducting training activities based on neuroscientce models of student reasoning, learning, and emotion. We will integrate lessons from brain-based models of human learning and reasoning with more traditional student modeling, teaching and learning theories. We are exploring the integration of case-based reasoning and scenario generation techniques, the use of stories, to drive training content in an agent-based architecture for training and mission rehearsal.

In the DARPA Narrative Networks project, Dr. Schumacher is using fMRI experiments and multivoxel pattern analysis to study the effect of suspense in narrative on the ability of the subject to remember the lessons which come after, which combined with other neuroscience findings from the Narrative Networks project is then applied to and integrated into a tool for narrative composition which will incorporate suspense and other findings into the composition to generate a narrative that will influence a target audience.

The GTRI team has researched and applied machine learning of several types. In the DARPA Integrated Learning Project we provided the case-based learning approach which collaborated with learning approaches from other universities to solve problems. We are currently working on an internal research project which is using hybrid cognitive learning and reasoning approaches to solve a variety of problems. These approaches can be applied at the neural data to help build models

Betty Whitaker

Technical Area 4

Implement novel machine learning algorithms that use mathematical abstractions of the identified cortical computing primitives as their basis of operation

Filipp Akopyan
IBM Research - Almaden

From Neuroscience to Machine Learning: An end-to-end story

Under the auspices of DARPA SyNAPSE Program, IBM has developed an end-to-end ecosystem that includes new neuroscience data [1], novel, non-von Neumann hardware [2], new simulator [3], new neuron models [4], new programming language [5], and new applications [6].

Leveraging our investment, to build a vertically integrated team, IBM is looking for collaborations with

- experimental neuroscientists who can provide operationalizable, quantitative data – within the time frame of IARPA MICrONS project – relating to neuroanatomy, neurophysiology, neuroplasticity, neurotransmitters.
- computational neuroscientists who can bridge the gap between neuroscience data and machine learning via concrete models.

[1] Dharmendra S. Modha and Raghavendra Singh, "Network Architecture of the Long Distance Pathways in the Macaque Brain", Proceedings of the National Academy of the Sciences USA, 2010.

[2] Paul Merolla, et al., "A Digital Neurosynaptic Core using Embedded Crossbar Memory with 45pJ per spike in 45nm", IEEE Custom Integrated Circuits Conference, September 2011.

[3] Robert Preissl, et al., "Compass: A Scalable Simulator for an Architecture for Cognitive Computing", IEEE SC 2012.

[4] Andrew S. Cassidy, et al., "Cognitive Computing Building Block: A Versatile and Efficient Digital Neuron Model for Neurosynaptic Cores," Proceedings of the International Joint Conference on Neural Networks in Dallas, TX, August 2013.

[5] Arnon Amir, et al., "Cognitive Computing Programming Paradigm: A Corelet Language for Composing Networks of Neurosynaptic Cores," Proceedings of the International Joint Conference on Neural Networks in Dallas, TX, August 2013.

[6] Steve K. Esser, et al., "Cognitive Computing Systems: Algorithms and Applications for Networks of Neurosynaptic Cores," Proceedings of the International Joint Conference on Neural Networks in Dallas, TX, August 2013.

Lambda Labs: MICrONS Abstract

Stephen Balaban s@lambdal.com
David Nicholaeff dnic@lanl.gov
July 6th, 2014

Lambda Labs seeks an academic, industrial, or national research laboratory to collaborate with in order to improve algorithms for machine intelligence. To construct this new class of models, we will draw upon our experience implementing modern Deep Learning methods, data parallel programs, and large-scale systems.

Current models for brain-like computing include Convolutional Neural Networks (CNNs), Autoencoders, and Restricted Boltzmann Machines (RBMs). These methods break away from the former paradigm of pattern recognition: hand-engineered features with a generic classifier. They currently achieve state-of-the art performance on a number of benchmarks. However, these methods still exhibit a large performance gap, both in accuracy and energy efficiency, to biological neural networks.

Current spiking neuron simulations, like Spaun, simulate millions of neurons. One second of real-time activity takes 2.5 hours of compute time. Data parallel programming techniques on GPUs using frameworks such as OpenCL and CUDA often lead to order of magnitude speedups over non data parallel implementations. We hope to leverage this technology to allow our group to conduct rapid prototyping and testing of different models and hypotheses.

Machine Intelligence from Cortical Networks Proposer's Day

Abstract

NVIDIA Corporation

NVIDIA Corporation has been a pioneer in visual computing for two decades. NVIDIA has deep experience in both the art and science of using computers to analyze and create images. NVIDIA is the inventor of the GPU, a highly parallel processor that excels at accelerating computational tasks in areas such as signal processing, video enhancement and machine learning. NVIDIA is also the inventor of CUDA, an advanced language and ecosystem for high-performance parallel computing. NVIDIA has expertise implementing and optimizing algorithms on parallel architectures. This includes strong experience with machine learning algorithms and systems, in both training and deployment. In recent years, nearly all of the top-place research teams competing in the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) have been using GPU-enabled deep neural networks.

NVIDIA is seeking to partner with teams that are attempting to significantly advance the state-of-art in machine learning. NVIDIA can apply strong expertise in implementing novel machine learning algorithms for computationally intense approaches that apply to very large datasets. NVIDIA prefers to partner with teams that are working with problems and systems that require extremely large amounts of computation and data processing.

Qelzal Corp.

Qelzal Corp. was founded by Olivier JMD Coenen and Ping Wang to bring brain-inspired technologies to the world and demonstrates their advantages in key areas over conventional technology.

The team has over 45 years of combined experience with expertise in physics, brain-inspired processing, spiking neural networks, machine learning, deep networks, control and robotic engineering with degrees from Stanford, UCSD & McGill in Physics, Neuroscience, Electrical Engineering, Computer Science and Biology.

OC was cofounder CEO at Intentiva Inc., graduated from Founder Institute and selected by Plug and Play accelerator. He worked on a DARPA project, which designed an "artificial nervous system for UAVs" (Unmanned Aerial Vehicles) at Brain Corp. PW is founder of Ansir, an innovation center and was founder CTO, business strategist at VEA Tech. and business development officer at Mytek, with experiences at HP and Agilent and in a couple of Silicon Valley startups.

OC established and directed the Neuroscience Group at the Sony Computer Science Laboratory, Paris for 7 years, where his group developed technologies in signal processing, machine learning and robotic control inspired from brain processing. He created and was leader or scientific coordinator on six funding grants, of which two major projects totaling over 12 million dollars in funding and received 16 academic awards and fellowships.

Olivier Coenen, CEO, Pres.; Ping Wang, CTO, Bus.Dev.;
olivier.coenen@qelzal.com; ping.wang@qelzal.com; www.oliviercoenen.com;
www.linkedin.com/in/hpwang; Tél: 650 427-0360; www.Qelzal.com
Some publications at http://goo.gl/hEZKMz on ResearchGate.

Synchrony of Thalamocortical Inputs Maximizes Cortical Reliability, Hsi-Ping Wang, et al., **Science** 328, 106 (2010);DOI: 10.1126/science.1183108

Reverse and Forward Engineering High-Level Vision
Cox Lab at Harvard University
MICrONS Proposer's Day Abstract

Abstract:

Humans recognize visual objects with such ease that it is easy to overlook what an impressive computational feat this represents. Any given object in the world can cast an effectively infinite number of different images onto the retina, depending on its position relative to the viewer, the configuration of light sources, and the presence of other objects in the visual field. In spite of this extreme variation, biological visual systems are able to effortlessly recognize at least hundreds of thousands of distinct object classes, reason about their structure, and guide action — a feat that no current artificial system can come close to achieving.

Our laboratory seeks to understand the computational underpinnings of object recognition and high-level vision, through a concerted effort on two fronts: first, we endeavor to understand the workings of biological visual systems using a variety of experimental techniques, ranging from microelectrode recordings and 2-photon calcium imaging to visual psychophysics; second, we attempt to instantiate what we have learned into artificial object recognition systems, leveraging recent advances in parallel computing to build systems that begin to approach the scale of natural systems. We believe that an integrative approach combining both systems neuroscience and computer science research holds great progress to accelerate our understanding of cortical computation. Our poster presentation seeks to highlight both of these core areas of research.

Our group has done significant work to establish rodents as a more experimentally-accessible and tractable model system for studying high-level visual processing. Our group was the first to demonstrate sophisticated, invariant object recognition abilities in rodents (Zoccolan et al. 2009), and we have built up significant infrastructure for training rodents to perform complex visual tasks. We pair these behaviors with high-throughput electrophysiology and 2-photon calcium imaging, to probe the nature of population responses in cortex, especially as a function of visual experience. We have found that rat extrastriate cortical neurons have many key features in common with primate visual cortical neurons, suggesting that the rat is a promising model for studying fundamental cortical primitives.

On the computational side of the lab, we develop biologically-inspired algorithms and models aimed at solving practical vision problems in a variety of areas, including face and object recognition (Pinto et al. 2009, Pinto et al. 2011, Scheirer et al. 2014), saliency prediction (Vig et al. 2013, Vig et al. 2014), and robotic navigation (Milford et al. 2014). Recently, we have begun to explore the incorporation of biological data as a constraint on machine learning algorithms (Scheirer et al. 2014), and we believe that such approaches hold great promise in driving progress in biologically-inspired machine learning.

Jeremy Holleman
University of Tennessee, Knoxville

In this presentation we will give an overview of unconventional computing efforts at the University of Tennessee and their applicability towards the goals of the MICrONS program. We will describe DeSTIN, an architecture for machine learning that has been demonstrated to be well-suited to analog hardware implementations. An analog CMOS deep learning engine based on DeSTIN has been demonstrated with a computational efficiency of about 1 TOPS/W, a 280x improvement over a synthesized digital equivalent. We will also describe the Neuroscience-Inspired Dynamic Architecture (NIDA), a spike-based computational framework designed through evolutionary optimization. Finally, we will provide an overview of the capabilities within the University of Tennessee in the broader area of unconventional and bio-inspired computation.

Neuromorphic Cortical Architectures for Bio-inspired Learning Machines

Emre Neftci and Gert Cauwenberghs
Institute for Neural Computation, University of California San Diego

The constraints on digital computers shaped current machine learning algorithms and their implementation such that they rely on discrete updates, exact arithmetics, off-line search through deep memory, and minimal interprocess communication.

In contrast, the brain's cognitive power emerges from a collective form of computation extending over very large ensembles of sluggish, imprecise, and unreliable analog components. This realization spurred scientists and engineers to explore the remarkable mechanisms underlying biological cognitive computing by reverse engineering the brain in "neuromorphic" silicon, providing a means to validate hypotheses on neural structure and function through "analysis by synthesis".

Our team exploits the fact that the brain and its corresponding neuromorphic emulations are free from the constraints imposed by traditional digital machine learning to offer a novel standpoint for understanding and implementing highly efficient and effective bio-inspired learning machines. To reach this goal, we design large-scale and real-time adaptive, general-purpose spiking neuromorphic supercomputers; and map learning algorithms and representations onto these spike-based neural substrates.

Our work on neural supercomputing infrastructure has contributed a hierarchical address event routing (HiAER) integrate-and-fire array transceiver (IFAT) system for large-scale hybrid analog-digital spike-based neuromorphic computing. Scaling to millions of multi-compartment integrate-and-fire (I&F) neurons and billions of conductance-based synapses with locally dense and globally sparse dynamically reconfigurable synaptic connectivity and plasticity, the HiAER-IFAT system serves as a computational tool for large-scale systems neuroscience as well as a real-time and low-power silicon platform for neocortical vision and audition.

Our work on bio-inspired learning architecture identified conditions under which such networks of I&F neurons support Monte Carlo neural sampling from Boltzmann probability distributions. We further formulated a form of spike-timing dependent plasticity that augments these I&F networks with on-line learning in complex environments, implementing a variant on contrastive divergence learning in equivalent Boltzmann machines.

We envision that this research will benefit from and contribute to synergies with other efforts focusing on high-throughput connectomics and systems neuroscience data to constrain our learning machines, and in identifying cortical primitives of computation and learning. Such primitives may draw on further advances in thermodynamical foundations of machine learning with biophysically realistic neural sampling.

The IARPA MICrONS program seeks to build a new generation of Machine Learning algorithms in Technical Area 4.

Current ML algorithms have shown rapid progress over the past decade. More recently, learning algorithms using deep architectures have shown object detection, recognition and classification results that rival human performance and in some cases even surpass them. Most of these gains have been made feasible by inspirations produced by a parallel explosion in knowledge and data from advances in neuroscience.

Nonetheless, in several aspects, even these learning algorithms severely lag human capabilities, even that of small children let alone expert adults. For example, such algorithms need a set number of predefined categories; are not flexible enough to generalize beyond them; are poor at explaining and evidence recounting, need a large amount of annotated training data and a huge training time that may run into weeks even using HPC facilities.

This exposes a need to build self-organizing, compositional (hierarchical) learning architectures that are, by design, capable of model refactoring and learning deep semantics. For example, we anticipate that this needs a paradigm shift away from learning and sharing of simple and complex features towards learning invariances. The learning architecture should be composed of computational units (motifs) that allow sharing and construction of nested invariances. This would, in turn, make the architecture capable of sharing (replicating) these motifs and invariances across the network thus enabling faster learning as well as learning of novel concepts in an efficient manner, perhaps, approaching human capabilities of near zero-shot or one-shot learning. We expect that such motifs will be available to us through the work carried out in TA3. Nevertheless, in anticipation, our approach would engineer such architectures inspired from extant knowledge about cortical circuitry, from the level of the neurons to neural microcircuits and cortical columns with the aim to ingest models from TA3 as and when they become available.

Team SRI has vast experience in building large and complex machine learning systems for speech, image and video parsing and understanding for a variety of government and commercial applications. SRI has one of largest computer vision and machine learning groups in the country and several decades of experience in working on related projects with all government agencies. Example domains include human behavior modeling and analysis, robotics, biometrics, security & surveillance, and healthcare. Dr. Graham Taylor (U. Guelph) has extensively worked on deep learning architectures with an emphasis on time series data with applications in computer vision, motion capture, climate data, speech recognition and finance.

We are in the process of putting together a strong team will complementary capabilities to build the next generation of ML architectures from cortical representations. We are seeking collaborations in the area of computational neuroscience to bridge the gap to TA3.

Maneesh Singh
SRI International - Princeton