

ABSTRACTS

Monday, March 2, 2009

Time: 09:00 – 09:30 am

Title: Reverse Engineering The Human Visual System

Presenter: James S. Albus

Krasnow Institute for Advanced Studies

George Mason University, Fairfax

Abstract: A model of computation and representation in the brain is presented that suggests a viable design for a real-time computationally-equivalent model of the human vision system. The model combines a reference model architecture¹ developed for intelligent machine systems with a biologically inspired model of computation and representation in the brain. The proposed model is designed to be both architecturally similar in structure, and functionally equivalent² in computational power, to the human cortex. It models the visual cortex as a hierarchically layered graph of cortical regions, each of which is an array of Cortical Computational Units (CCUs). It models a CCU as .2 mm² chunk of cortical real estate, plus associated subcortical nuclei, that contain three parts:

1. An abstract data structure with slots for attributes, state, and pointers.
2. A set of computational processes that window, segment, and group inputs into patterns (i.e. entities or events); compute and filter attributes and state of the patterns; and set or break links (i.e. pointers) between CCUs.
3. A set of processors (i.e., synapses and neurons) that implement (1) and (2).

The proposed model addresses both: a) receptive field hierarchies that are defined by anatomy, and b) entity/event hierarchies that are defined by pointers that link patterns of pixels and signals to entities, events, situations, and episodes. It models the looping connections between the cortex and underlying subcortical structures, particularly the thalamic nuclei. It models how behavioral goals and priorities are generated at various echelons in the frontal cortex, and how these are decomposed into tasks and plans through a hierarchy of echelons of control. It models how information regarding behavioral goals and priorities is used to select and modify algorithms in the visual processing hierarchy². The proposed model is mathematically tractable in its ability to describe the computational processes and data structures for representation of knowledge that are hypothesized to be embodied in the human visual cortex. It rests on mature mathematical concepts developed in control theory, computer science, and signal processing. The model is amenable to implementation in software on highly parallel supercomputers for demonstration of real-time performance. The measure of success will be a demonstration of its ability to process images from a variety of sensors to a level of perceptual insight that is measurably comparable to human scene understanding.

¹ Albus, et al (2002) 4D/RCS Reference Model Architecture for Unmanned Vehicle Systems, Version 2.0, NISTIR 6910, National Institute of Standards and Technology, Gaithersburg, MD

² Functional equivalence can be defined as producing the same input/output behavior. Given the same input stimulus, the engineered system should produce statistically the same output behavior as the biological system.

Time: 09:30 – 10:00 am

Title: Fusion-Based Robust Signal Processing by Humans and Machines

Presenter: Misha Pavel

Division Head, Biomedical Engineering

Oregon Health & Science University

Abstract: Although many existing automatic pattern recognition systems have achieved considerable success over the past fifty years, the performance of most of these systems

deteriorates drastically in unpredictable and changing environments. These systems are typically trained on labeled training data sets, rely on the use of dimension reduction, and as a result, their performance is governed by the extent that the training set represents the statistics of the application environment. The major shortcoming of these engineering systems is that their performance significantly depends on context and environmental settings. In contrast, biological systems seem to be much more resilient to the environmental changes that are not relevant to their recognition tasks. The differences in the performance between machine and biological learning approaches motivated us to explore new methods for pattern recognition based on high-dimensional representation, learning with partial information, and on rapidly adapting information fusion.

In the first part of my presentation, I will discuss the notion of robustness and some of the reasons for the deterioration in performance in typical pattern-recognition systems that are based on dimension reduction. I will then illustrate several relevant properties of human perceptual systems. In the third part of the talk I will describe ways that a pattern-recognition system can confront the problem of unpredictable and changing environmental conditions. In both of these approaches, the system incorporates high-dimensional representation.

Time: 10:30 – 11:30 am
Title: Nonstandard engineering principles of brain circuits
Presenter: Richard Granger
Professor, Psychological & Brain Sciences
Director, Brain Engineering Laboratory
Dartmouth College
Abstract: Brain components are slow (milliseconds), sparsely connected ($\sim p < 0.01$), and low-precision (1-2 bits per synapse), yet they outperform competing approaches in a range of fundamental applications such as visual and auditory recognition. Brain circuits are circuits, albeit with nonstandard engineering designs, and they are becoming understood computationally. We characterize typical brain circuit architectures and their operating rules; show examples of how such circuits give rise to novel algorithms for learning; and illustrate instances of these methods in multiple fielded applications. The results indicate engineering principles for design of novel parallel processing elements.
Bio: Richard Granger received his Bachelor's and Ph.D. from MIT and Yale, and is Professor of Psychological and Brain Sciences at Dartmouth. He directs Dartmouth's interdisciplinary Brain Engineering Laboratory, with research projects ranging from computation and robotics to cognitive neuroscience. He has authored more than 100 scientific papers and numerous patents, is co-inventor of FDA-approved devices and drugs in clinical trials, and has been the principal architect of a series of advanced computational systems used in military, commercial and medical settings.

Time: 11:00 – 11:30 am
Title: What can you do with your brain-inspired computer now that you've built it?
Presenter: James Anderson
Department of Cognitive and Linguistic Sciences,
Brown University, Providence
Abstract: History shows that software is more difficult and much slower to develop than hardware. If you do not even know the ultimate applications, it is even more difficult. Computers without software are of no value. By necessity, software that runs on a brain-like adaptive computer will be vastly different than the logic-based software of the traditional computer. Because developing software and finding appropriate problems is so hard and yet so important, we summarize below our attempts to develop software techniques and applications for a massively parallel adaptive computer. Over the last

five years, a group of researchers from Aptima, Inc., Brown University, and Alion Science and Technology have constructed a programmable parallel architecture for cognitive computing. Our Project is named the Ersatz Brain Project because we feel that even the most successful intelligent computing system will (for now) only be capable of weakly approximating the flexibility and integrative power of human cognition. The work is biologically inspired in structure and cognitively inspired in function but is often not biologically realistic in its approximations and details. Our ultimate goal is to make a computing system that can be as easily programmed and applied as a traditional von Neumann computer but that will work with massively parallel adaptive computers. We fondly hope what this architecture does well will be in consonance with the operations also performed well by human cognition. The Ersatz Brain approach is based on the use of attractor neural net-based dynamical systems operating at larger and larger scales of organization but using basically the same mechanisms across scales. The discrete aspect of the system -- essential for high level cognition, categorization, and language applications -- arises from the extensive use of discrete attractor states arising from an underlying continuous system. The analog aspect of the dynamical system -- essential for sensory processing, perception, and data integration -- provides robustness, intrinsic time domain behavior, and allows for novel means of programming the resulting system by controlling dynamical system parameters and data topography. Individual neurons are not the basic computing elements of the Ersatz system. The lowest level Ersatz computational module is a dynamical system based on a recurrent, attractor neural network. (In the real brain, cortical columns, a universal component of mammalian cerebral cortex, seem to have about the right size and properties. There are roughly a million columns in human cerebral cortex.) Individual modules are connected laterally to other modules forming two-dimensional arrays of modules we call the Network of Networks [NofN]. Associative learning can be performed using standard learning algorithms: Hebb, LMS, etc. Larger and larger scales of organization can be formed through associative linkage of the modules to form module assemblies and then linkage of the module assemblies to form multiple array structures for data integration and task performance. We note that, historically, machine learning and many proposed "brain like" computer architectures take a severely restricted and uninteresting computational form from the cognitive point of view. Machine learning is basically pattern recognition and is almost identical in philosophy, application and performance limitations to 1920's behaviorism. Going beyond the severe limitations of behaviorism applied to human cognition was the primary impetus for the highly successful cognitive revolution of the 1970's. The Ersatz approach is far more general and flexible than behaviorism. Although simple versions of Ersatz systems can perform traditional pattern recognition effectively, they can, in addition go beyond them, to more powerful and general applications. Specific tasks -- most notably, the controlled interaction of memory with new data -- can be performed through the use of programmable network based filters where one NofN array "programs" other NofN arrays, allowing for system data integration, inference, contextual disambiguation, and rapid task performance often without the need for additional learning

Time:	11:30 am - 12:00 noon
Title:	A mathematical canonical cortical circuit model that can help build future-proof parallel architectures.
Presenter:	Dileep George Founder and CTO, Numenta, Inc.
Abstract:	Building a parallel software/hardware platform for cortical simulations while the algorithms are still evolving is a challenging task. A mathematical abstraction for different circuit models can help to parameterize the architecture to support future version of the algorithms. This talk will explore how a mathematical model for cortical circuits can be built from the assumptions of Bayesian inference in a spatio-temporal hierarchy and how that model can guide hardware implementations.

It is well known that the neocortex is organized as a hierarchy. Hierarchical Temporal Memory (HTM) is a theory of the neocortex that models the neocortex using a spatio-temporal hierarchy. The HTM hierarchy is organized in such a way that the higher levels of the hierarchy incorporate larger amounts of space and longer durations of time. The states at the higher levels of the hierarchy vary at a slower rate compared to the lower levels. It is speculated that this kind of organization leads to efficient learning and generalization because it mirrors the organization of the world.

I will start this talk by demonstrating the recent advances at Numenta in using HTM for object recognition. We are able to recognize objects in clutter with a high degree of accuracy. Top-down attention based feedback is used to recognize multiple objects in a scene. Feedback is used to segment out objects from clutter.

I will then describe how the assumptions of Bayesian inference in a spatio-temporal hierarchy can lead to a mathematical model for cortical circuits. An HTM node is abstracted using a coincidence detector and a mixture of variable memory Markov chains. Bayesian belief propagation equations on this HTM node give a set of functional constraints for the cortical circuits. Anatomical and physiological data provide a set of organizational constraints. The combination of these two constraints can be used to derive a set of cortical circuits that explain many anatomical and physiological features and predict several other. I will discuss how such a mathematical model can be used to parametrize parallel software and hardware architectures so as to build a common platform that can support future algorithm changes.

Time: 12:00 noon – 12:30 pm
Title: The ALPS-EA for Robust, Massively Parallel Optimization
Presenter: Greg Hornby
Senior Scientist, UC Santa Cruz-NASA, NASA Ames Research Center, California
Abstract: Evolutionary Algorithms (EAs), a family of optimization algorithms inspired by natural evolution, have been used in a variety of design domains such as evolving analog circuits, antennas, MEMS devices, sorting networks, recurrent neural networks and others. Unlike most optimization algorithms, which were designed for single processors, EAs naturally parallelize to massively parallel computer systems. One of the main problems with EAs, and all optimization techniques, is that of prematurely converging in a mediocre local optima which cannot be escaped. Since this premature convergence problem prevents the effective use of large numbers of compute cycles, it is a fundamental problem that must be overcome for Massively Parallel Adaptive Computing to reach its full potential. Here we present our idea for reducing this problem called the Age-Layered Population Structure (ALPS) EA. ALPS maintains multiple, semi-independent populations at different "ages" with different degrees of convergence and regularly starts new searches in new parts of the search space. The end result is a robust search algorithm that is designed to fully take advantage of the compute cycles available in a massively parallel environment. Since EAs are a form of biologically inspired, adaptive computer and have been used for evolving both circuits and neural networks I hope this fits within the area of interest of this workshop.

Time: 1:30 pm – 2:00 pm
Title: Stable learning in networks of unreliable, memristive nanodevices
Presenter: Greg Snider
Researcher, Information and Quantum Systems Lab, HP Labs, California
Abstract: Neuromorphic circuits have teased us for fifty years with their potential for creating autonomous, intelligent machines that can adaptively interact with uncertain and changing environments. Although there are many stumbling blocks to achieving that vision, a primary problem has been the lack of a small, cheap circuit that can emulate the

essential properties of a synapse. Brains require synapses, and lots of them (an estimated 10^{14} in the human brain), but only about 1/10000 as many neurons, so synapse circuit design dominates the implementation problem. Memristive nanodevices may fill the role of an electronic analog of biological synapses: they are essentially analog memories that can be switched between extreme states in 20 nanoseconds or less, yet maintain their state for years when power is removed. They can also be manufactured at biological scale densities (more than 10^{10} devices per cm^2) and integrated with conventional CMOS. In this talk I will present some neuromorphic nano/CMOS architectures along with corresponding simulations showing how memristive nanodevices fabrics can be integrated with conventional CMOS circuitry to form networks capable of stable learning in changing environments (thus addressing the “catastrophic interference” problem), even though the nanodevices themselves show large variations in electrical properties. The circuits, like biological brains, are inherently defect-, fault-, and failure-tolerant.

Time: 2:00 pm – 2:30 pm
Title: VLSI Implementations of Very Large Scale Neuromorphic Circuits - Achievements, Challenges and Hopes
Presenter: Karlheinz Meier
Professor, Kirchoff-Institute for Physics,
Ruprecht-Karls-Universität Heidelberg, Universität Heidelberg, Germany
Coordinator, FACETS, Europe
Abstract: The methodologies for design, simulation, construction and operation of very large scale neuromorphic circuits in hardware differ considerably from those known in conventional microelectronics. Special challenges exist in many areas like interconnection technologies, the realisation of distributed long and short term memory, the implementation of adaptation and plasticity mechanisms in electronic circuits, the mapping from biological databases into hardware systems and an efficient exploitation of the inherent fault tolerance when implementing neural circuits in deep-submicron technologies. In the talk we present the work currently carried out in the European FACETS project, where a system based on wafer scale integration with 50 Million plastic synapses and up to 200.000 adaptive-exponential spiking neurons per wafer is now coming into operation. The status of that system as well as possible roadmaps for future developments aiming at the emulation of a substantial fraction of a mammalian brain are discussed. We will also address the issues of technology support, required funding and the challenge of integrating neuroscience in a technology project.

Time: 2:30 pm – 3:00 pm
Title: Neurogrid: Emulating a million neurons in the cortex
Presenter: Kwabena Boahen
Associate Professor, Bioengineering, Stanford University
PI, Brains in Silicon Lab, Stanford University
Abstract: The digital technique used to simulate neural activity has not changed since Hodgkin and Huxley pioneered ion-channel modeling in the 1950s. Since then, progress has come through miniaturization, which doubles computer performance every eighteen months (Moore’s Law). This trend has plateaued in recent years, making real-time cortex-scale simulations unattainable in the foreseeable future, even for the fastest supercomputers. During the same time period, advances in neural recording and imaging techniques have greatly enhanced our ability to characterize the brain’s structure and function, which now truly trumps our ability to simulate its behavior. Fortuitously, with the recently developed ability to emulate (i.e., simulate in real-time) various types of ion-channels as well as arbitrary patterns of synaptic connections, the analog technique pursued by neuromorphic engineers over the past two decades has matured. By employing the analog technique, Neurogrid will emulate a million neurons in the cortex—rivaling the performance of 20–200 IBM Blue Gene racks on this particular task—at under a

thousandth the cost. While neuronal-level mechanisms have been linked to network-level functions through computational modeling (e.g., generation of brain rhythms), scaling these models up to the area- and system-levels (where cognition emerges) has proved difficult. In the visual system alone, there are three dozen cortical areas, each with its own representation of the visual scene. It is not understood how conflicting information in these areas is reconciled. By performing simulations at a scale large enough to include interactions between multiple cortical areas yet detailed enough to account for what is known about brain function and neuronal structure, Neurogrid will help neuroscientists vet various hypotheses about how the brain works.

Time: 3:30 pm – 4:00 pm

Title: Neuromorphic Target Cues

Presenter: Bruce Schachter

Northrop Grumman, ATR & Image Exploitation Technology Center

Abstract: Automatic Target Recognizers are typically based upon a mixed bag of signal processing, image processing and pattern recognition paradigms. We describe an alternative approach for visible band ATR based upon a model of the human visual cortex, as derived from the latest neuroscience literature. The human vision system has two parallel streams of processing: (1) a dorsal stream to answer the question of where and (2) a ventral stream to answer the (target recognition) question of what. We have developed algorithms to mimic the modules of the dorsal process. A human-in-the-loop performs the target recognition (what) task — with tight coupling and feedback from man to machine — from neural processing to neuromorphic processing. The algorithms are being designed for eventual implementation in neuromorphic hardware with very restricted size, weight and power requirements. They are also directed at a new breed of video cameras with resolutions of greater than $10^{**}8$ pixels/frame. The video data is easily split into several data streams making parallelization of the processing quite straightforward. Based upon current real-time simulations, processing requirements are in the 2-4 TFLOP range with a power budget of under 10 watts.

Time: 4:00 pm – 4:30 pm

Title: When the storage device becomes the computer

Presenter: Bob (Robert) Thibadeau

Professor, Computer Science, CMU

Chief Technologist, Seagate Technology

Abstract: ...

Time: 4:30 pm – 5:00 pm

Title: Massive Data Computing

Presenter: Pradeep Dubey, Intel Corporation

Abstract: Computing platforms and the World Wide Web are undergoing significant architectural transitions driven by the unprecedented convergence of the need to process massive amounts of data with the availability of massive amounts of compute resources. This has significant algorithmic implications for traditional approaches to many common computational problems in visual computing and analytics. This further has the potential to enable a new class of Connected Computing applications. The service-oriented focus and real-time nature of these emerging applications make computing more implicit and capable of delivering a significantly more immersive experience to a broader class of end-users. This talk proposal will explore the compute-platform implications of this trend.

Time: 5:00 pm – 5:30 pm

Title: PetaVision: A Software Architecture for Performing Petascale Simulations of Visual Cortex

Presenter: Craig Rasmussen, Los Alamos National Lab

Abstract: The computational power of the new Petascale class of machines being installed at a few locations around the world rivals in many respects the computational power of the mammalian brain. This begs the question as to whether simulations on these machines can achieve mammalian level performance at tasks like visual object recognition. The goal of the PetaVision project is to develop a software framework for performing simulations of visual cortex at a massive scale. Visual computation is realized through the temporal dynamics of an anatomically-derived network of spiking neurons responding to images presented to an artificial retina. The potential of the software architecture of PetaVision was demonstrated in early experiments run on Los Alamos National Laboratories' Roadrunner computer. These experiments reached a peak of 1.14 Petaflops. The software architecture of PetaVision is described and the challenges remaining before synthetic visual cognition can be achieved.

Time: Submitted (family urgency prevented Randy to be available)

Title: Large scale learning models of visual object recognition

Presenter: Randall O'Reilly

Professor, Department of Psychology, Institute of Cognitive Science
Center for Neuroscience, University of Colorado Boulder

Abstract: We have been developing large-scale neural models of visual object recognition, based on the known anatomy and physiology of the relevant visual pathways in the brain. These models learn to robustly recognize novel object instances from trained categories (e.g., a new model of car), with high 90's percent accuracy. They have been trained on over 209 different object categories. These models can have many potential applications, in addition to providing insight into the computational functions of the visual system. Their large size and the amount of training time required make them an ideal target, and challenge, for hardware implementations.