

Summary Report from 2015 Neuro-Inspired Computational Elements (NICE) Workshop

*Information Processing and Computation Systems beyond
von Neumann/Turing Architecture and Moore's Law Limits*



<http://nice.sandia.gov>

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.



U.S. DEPARTMENT OF
ENERGY

Office of
Science



I A R P A
BE THE FUTURE



**2015 Neuro-Inspired Computational Elements Workshop:
Information Processing and Computation Systems beyond von Neumann/Turing
Architecture and Moore's Law Limits**

Hosted by Sandia National Laboratories

February 23-25, 2015

Workshop Advisory Committee

Dan Hammerstrom (DARPA)

Karlheinz Meier (EU HBP, University of Heidelberg)

R. Jacob Vogelstein (IARPA)

Robinson Pino (DOE Office of Science)

Workshop Organizing Committee

Murat Okandan (Sandia National Laboratories)

James Bradley Aimone (Sandia National Laboratories)

Workshop and Presentation Material

All presentation material from the Workshop can be found at nice.sandia.gov.

Special Acknowledgements

Linda Wood, Event Organization

Daniel Grubbs, Workshop Report Editor

Event Sponsors



U.S. DEPARTMENT OF
ENERGY

Office of
Science



I A R P A

Table of Contents

1. Executive Summary	1
2. Summary of Findings.....	2
3. Recommendations.....	3
1. Introduction	4
1.1 Workshop Introduction	4
1.2 DOE Context	5
1.2.1 DOE Strategic Plan	5
1.2.2 The Advanced Scientific Computing Research Program	5
2. Application Areas.....	6
2.1 High Performance Computing/Hybrid Architectures	6
2.2 Robotics/control systems	6
2.3 Big Data/High Velocity Data, Internet-of-Things	7
2.4 Cyber	7
2.5 Neural architecture, theory and algorithm exploration	7
3. Technical Approaches	7
3.1 Software implementation (conventional hardware).....	7
3.2 Combined software/improved architecture (new devices).....	7
3.3 Novel Architectures.....	8
4. Pathways and Resources.....	8
4.1 Large Scale Programs (BRAIN Initiative, EU HBP).....	8
4.2 New Research Projects / Institutional Level.....	8
4.3 Commercial Development/Evaluation	9
Appendix A: List of Presentations	10
Appendix B: List of Registered Attendees	11
Appendix C: Workshop Agenda	14
Appendix D: Presentation Abstracts	15
Appendix E: Website Data for nice.sandia.gov	27

1. Executive Summary

The third Neuro-Inspired Computational Elements (NICE) Workshop was held February 23-25, 2015 in Albuquerque, New Mexico. The goal of the Workshop was to bring together researchers from different scientific disciplines and application areas to provide a common point from which to develop the next generation of information processing/computing architectures that go beyond stored program architecture and Moore's Law limits.

Conventional, stored program architecture systems are designed for algorithmic and exact calculations. Many high impact problems, however, involve large, noisy, incomplete, "natural" data sets that do not lend themselves to convenient solutions from current systems. This Workshop series has focused on addressing these problems by building upon the convergence among neuroscience, microelectronics, and computational systems to develop new architectures designed to handle these natural data sets. An international group of nearly 100 registered attendees with a wide range of expertise and experience (e.g., neuroscience, systems, microelectronics, applications, and potential funding agencies) participated in this year's event by offering knowledgeable insight regarding:

- Applications that look for solutions beyond the capabilities of current computational systems
- Technical approaches that are at the early to middle stages of development for new computational systems
- Pathways and resources to accelerate the development of these new systems

A specific focus for this year's Workshop was the value proposition for neuro-inspired/neuromorphic computing: what these systems offer, or may offer, that exceeds the current and forecasted capabilities of conventional computing approaches. The speakers were requested to address this specific point in relation to their work and offer their views on potential next steps. Information from the speakers, combined with interactions and discussions with the broader set of Workshop participants, has led to the identification of a key finding and a key recommendation.

Key Finding – Neuro-inspired/neuromorphic systems could provide value in two, tightly coupled tracks: 1) as a new approach for analyzing, making sense of data, and predicting and controlling systems, 2) as a platform for understanding neural systems and testing hypotheses generated by neuroscience.

Key Recommendation – For both conventional computing and emerging approaches, a coordinated effort across multiple disciplines and application areas is needed to:

1. Establish appropriate metrics,
2. Develop performance parameters with specific application cases to evaluate current systems,
3. Support future development of neuro-inspired/neuromorphic systems.

This report provides a concise summary of the highlights of the Workshop, key findings, and recommendations for next steps in the context of DOE Office of Science goals and objectives.

2. Summary of Findings

The Workshop concluded with a breakout session at which participants discussed and helped articulate the most pertinent findings from the presentations and discussions. The results of this session shaped the following list of eight key findings:

Finding 1– Neuro-inspired/neuromorphic systems could provide value in two, tightly coupled tracks: 1) as a new approach for analyzing, making sense of data, and predicting and controlling systems, 2) as a platform for understanding neural systems and testing hypotheses generated by neuroscience.

Finding 2– Increasingly, the level of interest in brain-inspired computing approaches is moving beyond academic circles to broader government and industrial communities.

Finding 3– Large-scale projects and programs are underway in neuroscience, scientific computing, neural algorithm discovery, and hardware development and operation; and there is early interest in application metrics/definition for future systems evaluation.

Finding 4– Although current machine learning and other neuro- or bio-inspired systems have demonstrated valuable functions, further developments are necessary to achieve the higher levels of functionality desired for wider-spectrum applications.

Finding 5– High throughput techniques in experimental neuroscience are helping influence more advanced computational theories of neural function, but the community’s capability to translate these computational neuroscience concepts into real mathematical theories and application-useful algorithms is still immature.

Finding 6– Notable unresolved questions still facing the community include the level of neurobiological fidelity necessary for application impact, the necessity of future biological knowledge from experimentation to achieve neural computing’s goals, and the best strategies for achieving learning (both theoretically and in real systems).

Finding 7– The community appears to be approaching a general consensus that spike-based computation provides real, differentiating advantages over classic digital or analog approaches.

Finding 8– There are several theories and frameworks, each of which were presented at the Workshop, that are ready for implementation (HTM, stochastic computing, Leabra, Spaun, and others) on emerging neuromorphic/neuro-inspired computing platforms.

3. Recommendations

Relative to the key findings described above, and based on participant input during the breakout session, this report offers several recommendations for leveraging capabilities and investments to collectively advance the development and application of neuro-inspired/neuromorphic systems.

Recommendation 1 – Establish a coordinated effort.

For both conventional computing and emerging approaches, a coordinated effort across multiple disciplines and application areas is needed to 1) establish appropriate metrics, 2) develop performance parameters with specific application cases to evaluate current systems, and 3) support future development of neuro-inspired/neuromorphic systems.

Recommendation 2 – Maintain the multi-disciplinary nature of NICE Workshops.

While seeing details of experimental neuroscience is not immediately useful to applications and seeing microelectronic device power consumption metrics are not critical for colleagues involved in biological experiments, having the visibility across the full spectrum is necessary and should be maintained in future NICE Workshops.

Recommendation 3 – Use neuro-inspired platforms to develop, test, and refine theories, algorithms and systems.

Current and emerging neuro-inspired platforms appear to have value today and should be used to help construct, test, verify and refine new theories from neuroscience and new approaches such as stochastic computing and machine learning. Likewise, there is value in developing novel hardware platforms, particularly those that incorporate online plasticity and low energy communication strategies.

Recommendation 4 – Quantify and communicate the value proposition.

The perceived value of applying existing neural algorithm frameworks to real-world applications should be quantified and communicated broadly. However, it is important to acknowledge that these are only the “tip of the iceberg” in neural algorithms, and the development of new neural algorithms from the community’s growing knowledge of neuroscience is critical.

Recommendation 5 – Demonstrate specific application cases.

The community needs to develop a stronger application story. In lieu of the ‘killer app,’ which is not clearly visible or readily defined for neural-inspired systems today, the community should be actively developing clear application examples that demonstrate capabilities beyond current conventional computing approaches.

Recommendation 6 – Develop stronger mathematical neural theory.

Stronger mathematical neural theory is required. Improvements in this field will facilitate the transition of conceptual or simulation based computational neuroscience theories to application-useful machine learning tools.

1. Introduction

1.1 Workshop Introduction

The NICE Workshop has been held for the last 3 years (2013-2015) with the goal of bringing together scientists, engineers and stakeholders from a wide spectrum—ranging from experimental neuroscience, computational neuroscience, algorithms, high performance computing, hardware and all the way to applications. In all technological development timelines, including microelectronics and computers, a critical need has driven and funded the early stages of development. In the case of microelectronics, it was the navigation and guidance needs of defense systems, while in the case of computers, it was the neutron diffusion simulations and weather prediction and forecasting needs that provided the “killer app.” Such clear application cases are not currently defined for neuro-inspired/neuromorphic computing approaches, but may become more obvious once current projects gain further traction. The NICE Workshop community has seen and felt the need to articulate clearly the “value proposition” of this new approach. As a result, there are significant efforts to provide the metrics, comparison cases, and application examples to support further development of the science, technology, and ecosystem around this activity. NICE has succeeded in achieving its initial goal of providing a “nucleation point” for these discussions. Continuing the Workshop series will help achieve ancillary goals by accelerating current activities and supporting project and program development for neuro-inspired/neuromorphic computing. The intended result is the development of solutions for critical applications that might not be feasible within the current computing paradigm.

Several *application areas*, *distinct technical approaches* and *technology development paths* were identified during the presentations and discussions at this year’s Workshop. The body of this document is structured to provide additional insight within each of these categories. Figure 1 summarizes some of the interactions across the categories. Further detailed information can be found at nice.sandia.gov.

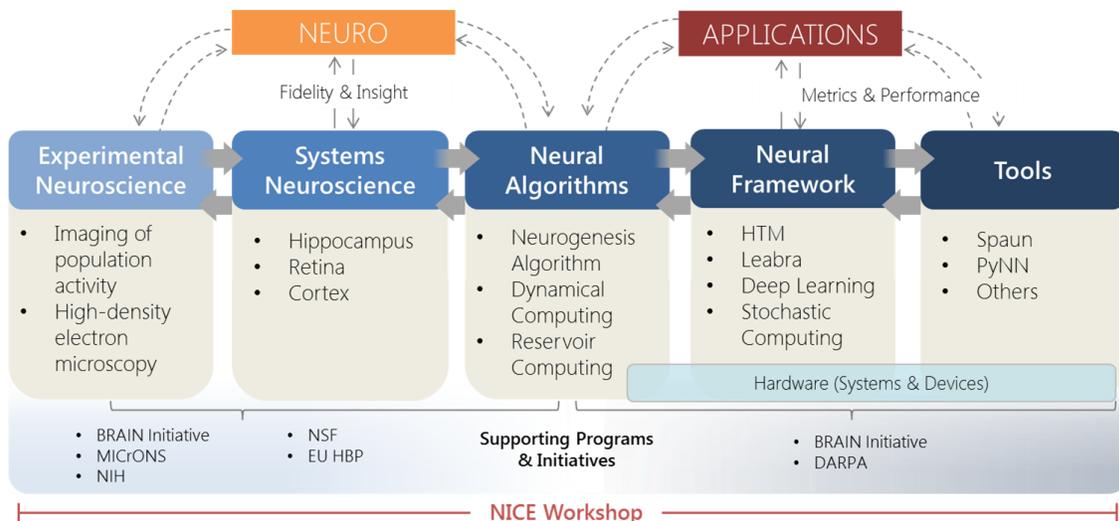


Figure 1 – A wide spectrum of highly interdependent scientific and technical disciplines are at the core of neuro-inspired/neuromorphic computing approaches and applications, all of which were represented at the 2015 NICE Workshop. These disciplines interact at relatively local scales, but a process for transitioning knowledge across disciplines appears to be maturing. The Workshop has a goal of furthering this process to influence and provide critical, longer-range interactions.

1.2 DOE Context

The NICE Workshop is motivated not only by the advancement of understanding of next generation information processing/computing architectures, but also by contributions to DOE strategic goals through the identification of applications, approaches, and resources that would benefit from accelerated development of these new architectures and systems. The Workshop supports specific objectives of DOE's Strategic Plan (2014-2018) and challenges being addressed by the Advanced Scientific Computing Research program.

1.2.1 DOE Strategic Plan

DOE's most recent strategic plan (March 2014) offers twelve strategic objectives related to high-level goals in the areas of: 1) Science and Energy, 2) Nuclear Security, and 3) Management and Performance. Strategic Objective 3 within the Science and Energy goal is intended to *"deliver the scientific discoveries and major scientific*

*tools that transform our understanding of nature and strengthen the connection between advances in fundamental science and technology innovation."*¹ One of DOE's stated strategies to accomplish this objective is to "pursue scientific discoveries that lay the technological foundation to extend our understanding of nature and create new technologies that support DOE's energy, environment, and security missions." Achieving these discoveries will, in many cases, require new computing architectures that overcome challenges in analyzing massive, natural datasets. The critical contributions from these architectures are recognized through the objective's concentration on "advanced scientific computing to analyze, model, simulate, and predict complex phenomena."

*"We must ensure that DOE continues to lead basic research in the physical sciences, **develop the next generation of computation technology,** and develop and maintain world-class scientific user facilities."*

– US Secretary of Energy Ernest Moniz¹

1.2.2 The Advanced Scientific Computing Research Program

Relative to DOE's concentration areas for Strategic Objective 3, the Advanced Scientific Computing Research (ASCR) program has a goal to "discover, develop, and deploy computational and networking capabilities to analyze, model, simulate, and predict complex phenomena important to DOE."² For decades, ASCR and its predecessor programs have enabled scientists to gain new insights into technical challenges through advancements in applied mathematics and computer science research. The program continues to focus research efforts in these areas, as well as in next generation networking and scientific discovery through advanced computing (SciDAC).

Some of today's high impact technical challenges deal with large, natural data sets that make experiments nearly impossible or inordinately costly on current systems. For such challenges, ASCR seeks to develop exascale computing capabilities that "enable the solution of vastly more accurate

¹ http://www.energy.gov/sites/prod/files/2014/04/f14/2014_dept_energy_strategic_plan.pdf

² <http://science.energy.gov/ascr/about/>

predictive models and the analysis of massive quantities of data, producing advances in areas of science and technology that are essential to DOE and Office of Science missions and, in the hands of the private sector, drive U.S. competitiveness.”³ ASCR has begun investing in solutions to these challenges which, when achieved, will “result in not only exascale systems but also in affordable, energy efficient petascale systems and high-end desktops.”⁴ In addition, a recent report created by the Advanced Scientific Computing Advisory Committee (ASCAC) Data Subcommittee highlights “data challenges arising in multiple science domains.”⁵ The following excerpt describes one of the stated exascale computing challenges related to biology and genomics:

“... perhaps most importantly, to learn how nature's design of a biologically-based knowledge machine beats all estimates for low power systems to compute or store information when compared to even the most power efficient hardware systems being conceived for the future.”⁶

2. Application Areas

2.1 High Performance Computing/Hybrid Architectures

One of the critical application areas for new computing architectures is in High Performance Computing (HPC) which envisions extremely large data sets and increasingly harder problems for the next generation of computing systems (i.e., Exascale). An interesting concept, which has been discussed at this Workshop and other new computing technology-related venues is a hybrid approach, where the neuro-inspired systems can perform the functions that are too costly (in terms of energy, time and hardware) in conjunction with conventional computing systems that are specifically designed for numerical computation tasks. Such accelerators could be embedded in an HPC system to support the discovery and analysis tasks at higher energy-efficiency levels than a conventional system would be able to achieve by itself.

2.2 Robotics/control systems

Next-generation microelectronic devices have enabled robotics (UAVs, industrial robots, self driving vehicles, etc.) to steadily increase their functionality and flexibility. Despite these improvements, the higher level functions are still performed by a human supervisor, usually necessitating a bi-directional high bandwidth data link between the platform and the supervisor or on-board personnel. While this can be accommodated in certain situations, availability of a higher functioning, self-contained, semi-autonomous system would greatly change the utilization scenarios in many use cases. Energy efficiency, real-time functionality and physical footprint of the systems are of critical concern on most applications where power, space, and response time are at a premium.

³ <http://science.energy.gov/ascr/about/>

⁴ <http://science.energy.gov/ascr/about/>

⁵ <http://science.energy.gov/ascr/about/>

⁶ <http://science.energy.gov/~media/40749FD92B58438594256267425C4AD1.ashx>

2.3 Big Data/High Velocity Data, Internet-of-Things

In many of the current consumer-facing and enterprise applications, data volume and the rate of data processing have stretched conventional computing resources to the limits. Especially in applications where the “sense-making” of the data analysis is not clearly defined or available, there is an opportunity for new approaches that can detect patterns and anomalies and guide further analysis by other, conventional means. In some scientific experiments, the data volume has been pre-compressed due to sensor characteristics or pre-processing algorithms. These endeavors could also benefit from a new approach to collecting, analyzing and interpreting the physical mechanisms underlying the observations, rather than relying on pre-determined compression, filtering, and analysis approaches.

2.4 Cyber

Cyber security is an example of a critical problem with extremely high data rates and volumes. New computational approaches could have immediate, high impact in this domain. Many current network management and cyber security protocols are reactionary: after an intrusion or problem is detected and analyzed, a countermeasure can be developed and deployed. A higher level of functionality in anomaly detection and the ability to correlate multiple, seemingly unrelated features without being specifically programmed to do so could enable new approaches for the management and maintenance of networks at varying scales of deployment.

2.5 Neural architecture, theory and algorithm exploration

As discussed in great depth at the Workshop, the datasets and experimental findings in neuroscience are vast, and not always well linked. The ability to formulate, test, verify, or disprove theories around neural circuits, computation and other scientific questions with the current and emerging neuro-inspired/neuromorphic platforms at large scales (time, populations, variations in parameters, etc.) would be a critical capability which does not currently exist. Using these results to formulate better brain-machine interfaces, neural prosthesis, and medical treatments is a worthy goal.

3. Technical Approaches

3.1 Software implementation (conventional hardware)

Many of the current approaches to neuro-inspired computing (convolutional networks–machine learning, Hierarchical Temporal Memory or HTM, Leabra, Spaun) are implemented in software, on conventional computing resources. This implementation provides the highest level of flexibility and accessibility for the evaluation and further development of current approaches. The time required to perform the simulations or learning cycles, however, tend to be 10-100x slower than real time. Great results from software implementation are also informing the technical approaches that follow within this section of the report.

3.2 Combined software/improved architecture (new devices)

Systems specifically designed and constructed to simulate and study neural circuits (SpiNNaker, BrainScaleS, Neurogrid, Field-Programmable Analog Arrays, IFAT, etc.) have been under development

for around a decade, or longer. These systems provide significant improvements in certain metrics (energy efficiency, time to execute/model, biological realism) but give up some of the flexibility and familiarity that is available in the software based approaches. Workshop participants acknowledged that these systems are poised for large scale simulations and evaluation of theories, frameworks, and algorithms generated by the community. These systems will also help to accelerate the discovery and development of the next generation of computing, control, and analysis systems, and be the testing ground for new directions in neuroscience.

3.3 Novel Architectures

It might be possible to implement (in hardware) a completely new way of representing, processing, storing and recalling information, based on the spatio-temporal, sparse, hierarchical features of spiking neural networks. These systems are further away from the conventional, symbolic computation model but could be the key to breaking through some conventional computing barriers. Other technical approaches will critically inform how such systems can be built, optimized and utilized in a wide variety of applications. One distinct feature is the absence of a clear “hardware/software” division. While there will still be digital, analog, electronic, optical and/or other novel devices that can be locally programmed and inspected (“peek and poke”), operation and functionality of the complete system will require new tools, metrics, and interfaces.

4. Pathways and Resources

4.1 Large Scale Programs (BRAIN Initiative, EU HBP)

Several existing large, multi-national programs are supporting a cross-cutting assembly of scientific fields including neuroscience and neuromorphic computing. Applications are generally viewed as a later-stage product of this activity, with scientific output being the primary goal. The interactions at the NICE Workshop fostered new connections and strengthened existing collaborations. A sense of new, unexplored areas and great potential was tempered with a warning of previous epochs, “Third Wave of ...” (Neural Networks, Neuromorphic Computing, and New Computing Paradigm). Large scale efforts in the 1970s, ‘80s and following decades, which attempted to achieve higher levels of functionality than contemporary computers created the foundation for current efforts. A potential development path around high performance computing activities was also presented, which would further support activity in neuro-inspired/neuromorphic computing and provide specific application goals and metrics to drive development efforts.

4.2 New Research Projects / Institutional Level

Many universities, research organizations, and foundations have increased their levels of activity in neuroscience and related disciplines including neuro-inspired/neuromorphic computing, with some focusing on more of the early, scientific goals and others positioning themselves in various intersections of the wide spectrum of involved disciplines. There is a high level of interest in neuroscience and neuro/bio-inspired themes in the academic realm, evidenced by numbers of both faculty and students who are getting involved, as well as by expressions of interest by active researchers.

4.3 Commercial Development/Evaluation

In the commercial sector, organizations that routinely handle large volumes of data and derive value from processing, storing and analyzing these data sets, and providers of systems that enable this activity have been actively exploring alternative computing approaches. Machine learning and associated activity has been a great example of new functionality that was enabled by increasing computing power that, in turn, has provided improved methods to analyze data. No clear candidate for next-generation devices or architectures yet exists to supplement or replace conventional computing and CMOS microelectronics; early evaluation is valuable to both the systems manufacturers and service providers. We expect this “pre-competitive” arena will become even more valuable and could see increasing support and activity related to commercial development.

Appendix A: List of Presentations

Speaker	Affiliation	Presentation Title
<i>Brad Aimone</i>	Sandia National Labs.	Adaptive Neural Algorithms: the What, Why, and How
<i>Kris Carlson</i>	U. of California, Irvine	Large-Scale, Biologically Detailed Neuromorphic Networks: Taming the Beast
<i>Sek Chai</i>	SRI International	Computational Noise Resiliency in Deep Learning Architectures
<i>Ralph Etienne-Cummings</i>	Johns Hopkins U.	Seeing with Spikes: From Motion Detection to Object Recognition
<i>Paul Franzon</i>	North Carolina State U.	Hardware Acceleration of Sparse Cognitive Algorithms
<i>Jeremy Freeman</i>	Janelia Farm/HHMI	Measuring and manipulating neural computation
<i>Steve Furber</i>	U. of Manchester	The SpiNNaker Project
<i>Nathan Gouwens</i>	Allen Institute	High-throughput experimental and computational characterization of cortical circuit components
<i>Dan Hammerstrom</i>	DARPA	Neurocomputing
<i>Jennifer Hasler</i>	Georgia Tech	Moving towards large, power efficient neuromorphic systems
<i>Jeff Hawkins</i>	Numenta	Reverse Engineering the Neocortex: Implications for Machine Learning and Machine Intelligence
<i>Marwan Jabri</i>	Neuromorphic LLC	-Biologically-Inspired Unsupervised Learning of Higher Order Visual Features
<i>Garrett Kenyon</i>	Los Alamos National Lab.	Deep, Sparse Representations of Form, Depth and Motion
<i>Konrad Kording</i>	Northwestern U.	At which level do we want to be neuro-inspired?
<i>Dhiresha Kudithipudi</i>	Rochester Inst. of Tech.	Traversing the Application Landscape of Neuromemristive Computing
<i>Tai Sing Lee</i>	Carnegie Mellon U.	Neural circuits for learning internal models of the environment
<i>Gary Marcus</i>	New York University	The Atoms of Neural Computation
<i>Helen Li</i>	U. of Pittsburgh	The neuromorphic computing leveraging emerging devices
<i>Karlheinz Meier</i>	U. of Heidelberg	Mixed-signal accelerated Systems Progress, Results, Plans
<i>Murat Okandan</i>	Sandia National Labs	Neuro-inspired Computational Engines: Beyond von Neumann/Turing Architecture and Moore's Law Limits
<i>Randal O'Reilly</i>	U. of Colorado Boulder	Biologically-inspired error driven learning in thalamocortical circuits
<i>Igor Ovchinnikov</i>	UCLA	Toward Cohomological Neurodynamics
<i>Robinson Pino</i>	DOE Office of Science	-
<i>Xaq Pitkow</i>	Rice University	How can we know if the brain is doing a good job?
<i>Paul Rhodes</i>	Evolved Machines	Quantifying the utility of neural circuitry
<i>Fred Rothganger</i>	Sandia National Labs	Can memristors learn?
<i>Catherine Schuman</i>	U. of Tennessee, Knoxville	A Programmable Array of Neuromorphic Elements
<i>Sebastion Seung</i>	Princeton University	Connectomics: from retina to cortex
<i>Tarek Taha</i>	U. of Dayton	- Neuromorphic Computing on Memristor Cross-Bar Arrays
<i>R. Jacob Vogelstein</i>	IARPA	-
<i>Lloyd Watts</i>	Neocortex	Event-Driven Simulation of Spiking Neural Networks
<i>Ken Whang</i>	NSF	Neuro-inspired computing at NSF
<i>Winfried Wilcke</i>	IBM	The IBM Cortical Learning Center
<i>Alan Yuille</i>	UCLA	Complexity and Compositionality

Appendix B: List of Registered Attendees

Agarwal	Sapan	Sandia National Labs
Aidun	John B.	Sandia National Labs
Aimone	Brad	Sandia National Labs
Antoniades	John "Yiannis"	BAE Systems
Baatar	Chagaan	Office of Naval Research
Birch	Gabriel C.	Sandia National Labs
Broom-Peltz	Brian E.	New Mexico Consortium
Buesing	Lars H.	Columbia University
Candelaria	Jon	Semiconductor Research Corporation
Cao	Yu	Arizona State University
Carlson	Kristofor D.	University of California, Irvine
Caudell	Thomas P.	University of New Mexico
Chai	Sek	SRI International
Chance	Frances	Sandia National Labs
Chavarria	Daniel G.	Pacific Northwest National Laboratory
Chavez	Wesley A.	New Mexico Consortium
Chen	Yiran	University of Pittsburgh
Chiasson	John	Boise State University
Clem	Paul	Sandia National Labs
Cox	Jonathan	Sandia National Labs
Daffron	Christopher	University of Tennessee, Knoxville
Dean	Mark E.	University of Tennessee
DeBenedictis	Erik	Sandia National Labs
Donaldson	Jonathon W.	Sandia National Labs
Draelos	Timothy	Sandia National Labs
Draper	Bruce L.	Sandia National Labs
Edwards	Arthur	Boise State University
Eilert	Sean S.	Micron
Etienne-Cummings	Ralph	Johns Hopkins University
Faust	Aleksandra	Sandia National Labs
Fisher	Dimitry	Brain Corporation
Flynn	Daniel C. "Dan"	Northrop Grumman
Follett	David R.	Lewis Rhodes Labs
Franzon	Paul D	NCSU
Furber	Stephen B.	The University of Manchester
Goudarzi	Alireza	University of NM
Gouwens	Nathan W.	Allen Institute for Brain Science
Graham	Dillon	Sandia National Labs
Greenwald	Hal	The Mitre Corporation
Hammerstrom	Daniel	DARPA

Hasler	Jennifer O.	Georgia Institute of Technology
Hawkins	Jeff	Numenta
Healy	Michael J.	University of New Mexico
Hendrickson	Bruce	Sandia National Labs
Hermina	Wahid	Sandia National Labs
Hill	Aaron	Sandia National Labs
Hsia	Alexander	Sandia National Labs
Imaino	Wayne	IBM
Just	Wayne A.	University of NM
Karnik	Tanay	Intel
Kelly	Sean T.	Northrop Grumman
Kenyon	Garrett T.	Los Alamos National Lab
Kudithipudi	Dhireesha	Rochester Institute of Technology
Kunde	Gerd J.	New Mexico Consortium
LaCasse	Charles F.	Sandia National Labs
Li	Hai	University of Pittsburgh
Loncaric	Josip	Los Alamos National Lab
Lundquist	Sheng	New Mexico Consortium
Mandal	Soumyajit	Case Western Reserve University
Marcus	Gary M.	New York University
Matthew	Marinella	Sandia National Labs
Meier	Karlheinz	Heidelberg University
Merkel	Cory E.	Rochester Institute of Technology
		Space and Naval Warfare Systems
Migliori	Benjamin	Center Pacific
Miller	LeAnn	Sandia National Labs
Mitchell	Alan	Sandia National Labs
Naegle	John H.	Sandia National Labs
Neftci	Emre	UC San Diego
Nugent	Alex	M. Alexander Nugent Consulting
Okandan	Murat	Sandia National Labs
O'Reilly	Randy	University of Colorado, Boulder
Ormandy	Roman	Embody Corp.
Ortiz	Keith	Sandia National Labs
Ovchinnikov	Igor V.	UCLA
Parekh	Ojas	Sandia National Labs
Parekh	Ojas	Sandia National Labs
Piekniewski	Filip	Brain Corporation
Pitkow	Zachary	Rice University
Plimpton	Steve	Sandia National Labs
Rao	Valluri R.	Intel Corporation
Rhodes	Paul	Evolved Machines, Inc.
Robinson	Pino	US Department of Energy

Rothganger	Frederick H.	Sandia National Labs
Ryu	Hyunsurk	Samsung
Salter	Thomas S.	Lab. for Physical Sciences, University of Maryland
Schultz	Peter F.	New Mexico Consortium
Schuman	Catherine	University of Tennessee
Shabani	Alireza	Google
Shepanski	John	Northrop Grumman
Skorheim	Steven W.	UCR
Speed	Ann	Sandia National Labs
Sukumar	Sreenivas R.	ORNL
Taha	Tarek	University of Dayton
Theiler	Max	New Mexico Consortium
Thibadeau	Robert	Bright Plaza, Inc.
Townsend	Duncan	Lewis Rhodes Labs
Tseng	Gilbert H.	Northrop Grumman
Turner	Christian	Sandia National Laboratories
Vahle	Michael O.	Sandia National Labs
Verzi	Stephen J.	Sandia National Labs
Vineyard	Craig M.	Sandia National Labs
Vishal	Saxena	Boise State University
Vogelstein	Jacob	IARPA
Vrudhula	Sarma	Arizona State University
Wagner	John	Sandia National Labs
Wang	Kang	UCLA
Watts	Donald Lloyd	Neocortix, Inc.
Whang	Kenneth	National Science Foundation
Whiting	Mark A.	PNNL
Wilcke	Winfried W.	IBM
Wu	John	Lawrence Berkeley National Lab
Yakopcic	Chris	University of Dayton
Yu	Shimeng	Arizona State University
Zhang	Xinhua	University of NM

Appendix C: Workshop Agenda

Monday, February 23	
8:00-8:10	Welcome/Logistics
8:10-8:30	Bruce Hendrickson
8:30-9:30	Programmatic Panel: Dan, R. Jacob, Karlheinz
9:30-10:00	Brad Aimone
10:00-10:15	Break
10:15-11:00	Sebastian Seung (plenary 1)
11:00-11:30	Tai Sing Lee
11:30-12:15	Xaq Pitkow
12:15-1:30	Lunch
1:30-2:00	Gary Marcus (plenary 2)
2:00-2:30	Lars Buesing
2:30-3:00	Jeff Hawkins
3:00-3:15	Break
3:15-3:45	Alan Yuile
3:45-4:15	Randal O'Reilly
4:15-4:45	Konrad Kording
4:45-6:00	Round-table/open-mic

Tuesday, February 23	
8:00-8:10	Robinson Pino
8:10-8:30	Ken Whang
8:30-8:50	Kris Carlson
8:50-9:10	Dhiresha Kudithipudi
9:10-9:30	Jennifer Hasler
9:30-10:00	Winfried Wilcke
10:00-10:15	Break
10:15-10:30	Lloyd Watts
10:30-10:45	Robert Thibadeau
10:45-11:15	Steve Furber
11:15-11:45	Ralph Etienne-Cummings
11:45-12:15	Murat Okandan
12:15-1:30	Lunch
1:30-2:00	Nathan Gouwens
2:00-2:30	Paul Franzon
2:30-3:00	Jeremy Freeman
3:00-3:15	Break
3:15-3:45	Sek Chai
3:45-4:15	Paul Rhodes
4:15-4:45	Karlheinz Meier (technical)
4:45-6:00	Round-table/open-mic

Wednesday, February 23	
8:00-8:10	Welcome/previous day overview
8:10-8:30	Tarek Taha
8:30-8:50	Fred Rothganger
8:50-9:10	Marwan Jabri
9:10-9:30	Catherine Schuman
9:30-10:00	Garrett Kenyon
10:00-10:15	Break
10:15-10:45	Igor Ovchinnikov
10:45-11:05	Helen Li
11:05-11:25	Matt Marinella
11:25-11:45	Roman Ormandy
11:45-12:00	Wrap-up/Next Steps/ Workshop adjourns
12:00-1:30	Lunch
1:30-4:00	"Neuromorphic Computing Strategy Paper" open session

Appendix D: Presentation Abstracts

Brad Aimone, *Sandia National Laboratories*

Adaptive Neural Algorithms: the What, Why, and How

Arguably, the aspect of true biological neural computation that is least captured by modern machine learning methods is, ironically, the ability to learn continuously. There are a number of reasons that methods in online learning have been slow to materialize, ranging from the application demand to high costs of software or hardware implementation. However, I contend that the biggest challenge to online learning is the lack of a strong theoretical approach for incorporating new information into existing models. Although existing methods struggle with the prospects of incorporating novel information, the brain's ability to learn provides us with a blueprint from which to develop new strategies for continuous adaptation. My talk summarizes some lessons from observing different neural systems and their different learning strategies, with particular emphasis on the variability of neural learning mechanisms and the fact that learning in neural circuits is often specifically suited to the overall region's functions.

Kristofor D. Carlson, *UC Irvine*

Large-Scale, Biologically Detailed Neuromorphic Networks: Taming the Beast

Neuromorphic engineering takes inspiration from biology to design brain-like systems that are extremely low-power, fault-tolerant, and capable of adaptation to complex environments. The design of these artificial nervous systems involves both the development of neuromorphic hardware devices and the development neuromorphic simulation tools. In this presentation, I describe CARLsim, a simulation environment developed by our group that can be used to design, construct, and run spiking neural networks (SNNs) quickly and efficiently using graphics processing units (GPUs). The simulation environment utilizes the parallel processing power of GPUs to simulate large-scale SNNs. I discuss recent improvements in the latest version of CARLsim that include advanced plasticity mechanisms, more complete documentation, and new data analysis and visualization tools. Finally, I discuss an automated parameter tuning framework that utilizes the simulation environment and evolutionary algorithms to tune SNNs. We believe the simulation environment and associated parameter tuning framework presented here can accelerate the development of neuromorphic software and hardware applications by making the design, construction, and tuning of SNNs an easier task.

Sek Chai, *SRI*

Computational Noise Resiliency in Deep Learning Architectures

NICE 2015 Workshop was a great success. There were many positive remarks throughout the event regarding the technical quality and breadth of the discussion. The credit also goes to the Workshop organizers for putting together a strong program. The discussions on brain modeling and connectome were good because they added a sense of perspective on how little we understand human cognition. One of the focus topics this year is applications. I am glad to see several videos and presentations showing neuro-inspired/enabled applications. I have talked to some "new" attendees that these help them understand and connect to our work.

I presented on the topic of noise resiliency of learning algorithms, to attempt to connect neuromorphic research to other computing research topics such as near-threshold, approximate, and fault tolerant computing. I conjecture that noise needs to be addressed in architectures with very dense neuron and synaptic connections. Furthermore, nanoscale VLSI designs will have issues (e.g. device variations, power density, and aging effects), which makes it equally important to study noise. I showed examples how we handle sensor noise, and how we already leverage noise in training our deep learning systems to improve algorithmic performance. Other speakers have related coverage on noise: Karlheinz Meier noted the different noise types considered in hardware design; Xaq Pitkow described a noise correlation as a theoretical basis for learning; Randal O'Reilly described error-driven learning using an autoencoder.

For next year, NICE should continue to cover key topics from theoretical neuroscience to hardware design. It would be good to get updates on US-based initiatives from IARPA MICrONS and Cortical Processor Algorithm Evaluation efforts. I would like to see this community report on efforts in benchmarking and defining key metrics to help evaluate different learning approaches.

Ralph Etienne-Cummings, *Johns Hopkins University*

Seeing with Spikes: From Motion Detection to Object Recognition

Visual motion estimation and object recognition are computationally intensive, but important tasks for sighted animals. As can be expected, the recognition of fast moving objects is more useful, but also much more difficult to compute. Replicating the robustness and efficiency of biological visual information processing in artificial systems would significantly enhance the capabilities of future robotic systems. Twenty-five years ago, Carver Mead outlined his argument for replicating biological processing in silicon circuits. His vision served as the foundation for the field of neuromorphic engineering, which has experienced a rapid growth in interest over recent years as the ideas and technologies mature. Replicating biological visual sensing was one of the first tasks attempted in the neuromorphic field. In this talk we discuss the tasks of visual motion estimation and object recognition *using asynchronous spike trains of change events at the retina. We describe the tasks, present the progression of work from the early first attempts through to the modern day state-of-the-art, and provide an outlook for future directions in the field. In particular, we argue that current frame-based computer vision is anachronistic and must be replaced by an event-based approach that more closely matches the information being sampled with the computation method. It is not surprising that living organisms also uses event based processing to understand their visual world.

Paul Franzone, *NCSU*

Hardware Acceleration of Sparse Cognitive Algorithms

NCSU is currently conducting a study on the performance and cost of implementing accelerators for HTM and related algorithms. The presentation includes the approach and preliminary data.

Steve Furber, University of Manchester

The SpiNNaker Project

Just two years after the world's first stored program ran its first program at Manchester in 1948, Alan Turing published his seminal paper on "Computing Machinery and Intelligence." The paper opens with the words: "I propose to consider the question, 'Can machines think?'" Turing then goes on to explore this question through what he calls "The Imitation Game," but which subsequent generations simply call "The Turing Test".

Despite spectacular progress in the performance and efficiency of machines since Turing's time, we have yet to see any convincing demonstration of a machine that can pass his test. This would have surprised Turing—he believed that all that would be required was more memory. Although cognitive systems are beginning to display impressive environmental awareness, they do not come close to the sort of "thinking" that Turing had in mind.

My take on the problems with true artificial intelligence are that we still haven't worked out what natural intelligence is. Until we do, all discussion of machine intelligence and "the singularity" are specious. Based on this view, we need to return to the source of natural intelligence, the human brain.

The SpiNNaker project has been 15 years in conception and 8 years in construction, and is now ready to contribute to the growing global community (exemplified by the EU Human Brain Project) that is aiming to deploy the vast computing resources now available to us to accelerate our understanding of the brain, with the ultimate goal of understanding the information processing principles at work in natural intelligence. SpiNNaker is a massively-parallel computer system, ultimately to incorporate a million ARM processor cores (the largest machine to date has 100,000 cores) with an innovative lightweight packet-switched communications fabric capable of supporting typical biological connectivity patterns in biological real time.

Jeff Hawkins, Numenta

Reverse Engineering the Neocortex: Implications for Machine Learning and Machine Intelligence

We are making significant progress in reverse engineering the neocortex. I presented a framework for neocortical theory called Hierarchical Temporal Memory (HTM). The core tenets of this theory are the cortex is composed of hierarchy of nearly identical regions. Each region learns and recalls sequences of patterns. The cellular layers that comprise each region implement sequence memory for different aspects of inference and motor generation. All aspects of the theory rely on sparse distributed representations.

As part of HTM theory, I described a detailed model of how a layer of neurons learns sequences. This model requires neurons with thousands of synapses arranged on active dendrites, closely mirroring the structure of biological neurons. I claimed that neuromorphic hardware intended for machine intelligence and cortical modelling must accommodate these more realistic neurons and the higher connectivity they require.

I presented a research roadmap showing the progress made so far and what remains to be done. All of our experiments, software, and documentation are freely available in an open source project called NuPIC. All aspects of NuPIC more than doubled over the last year reflecting increasing interest of HTM theory.

I showed diverse applications of HTM technology for automated analytics of streaming data, including machine anomaly detection, cyber-security, geo-spatial tracking, and natural language processing.

We believe we have passed a threshold in understanding how the cortex works. Machine intelligence based on cortical principles is possible in the near future.

Marwan Jabri, Neuromorphic, LLC

Biologically-Inspired Unsupervised Learning of Higher Order Visual Features

An important aspect of machine learning for visual pattern recognition is the understanding of how higher order visual feature detectors (tuned processing elements) develop. Understanding how cortical areas such as V4, posterior inferotemporal (PIT), and anterior inferotemporal cortices (AIT), could help shed some light.

We present an architecture and unsupervised learning algorithms inspired from the primate visual neural processing pathway. The architecture includes a V1 (simple and complex layers), and layers representing V4, PIT and AIT equipped with lateral excitatory and inhibitory projections. The V4 layer consists of two sublayers, integration and pooling. Hebbian learning occurs in the V4 integration layer, PIT and AIT. We show how complex visual features detectors can form in these higher cortical areas, in particular, face like tunings are observed after learning in the subsystem representing AIT on images of faces.

We apply the architecture and learning algorithms to the task of face recognition from the LFW and proprietary datasets. The output of the AIT layer is used as input features to a two-layer multi-layer perceptron trained for labelling. We obtain very encouraging results on fairly challenging recognition conditions, which include multiple facial poses, illuminations/brightness, and face rotations, with over 89% success rate on LFW datasets and 95% on a proprietary dataset.

Garrett Kenyon, Los Alamos National Laboratory

Deep, Sparse Representations of Form, Depth and Motion

- Subtitle: "The 3rd Age of Neural Computing"
- Sub-subtitle: "A Universal Cortical Processing Module"

Sparse predictive coding modules have emerged as viable candidates for defining a universal cortical processor. Sparse autoencoders are self-organizing and can explain many of the linear and nonlinear response properties of simple cells in the primary visual cortex. Moreover, sparse predictive coding modules can be strung together into essentially arbitrary topologies. Here, we demonstrate how sparse predictive coding modules can be used to represent form, motion and depth features in a manner that

enables subsequent categorization. We use an open source, high-performance neural simulation toolbox called PetaVision. In a typical simulation, we use either single images, sequences of video frames, or stereo image pairs. A layer of cortical neurons then learns an optimal set of features for representing that input as accurately as possible while using as few active elements as possible, a process that can be simulated in a neurally-plausible manner using lateral synaptic inhibition. The same process is then repeated to learn the receptive field properties of subsequent layers arranged in a hierarchical sequence. We then test the resulting sparse representations 3 ways. First, we test for the ability of a multi-layer sparse hierarchy to support good classification performance on object detection tasks, thereby assessing how the sparse predictive coding module represents form in a viewpoint invariant manner. Second, we test for the ability of a multi-layer hierarchy trained on short video sequences to enable good discrimination between different types of human actions, thereby assessing how a sparse representation of motion enables better discrimination of spatiotemporal patterns. Third, we test for the ability of a sparse representations trained on stereo image pairs to reconstruct accurate depth maps. Our results illustrate how sparse predictive coding can be applied to a range of visual processing modalities. Our results thus support the hypothesis that sparse predictive coding can be used to define a universal cortical processing module that can be configured into arbitrary topologies for solving difficult classification tasks.

Konrad Kording, *Northwestern University*

At which level do we want to be neuro-inspired?

When taking inspiration from biology it is important to specify at which level we want to be inspired. I argue that one largely neglected level of abstraction is by the joint representation of information by a whole group of neurons. This could also be helpful because it may sidestep the problem that a lot of findings in neuroscience may not generalize and the problem that computational principles may be in conflict with computational principles. For example, fourier transforms can be implemented very rapidly in current hardware, so if the biological solution is close to a fourier transform we might rather use the idea that can be rapidly implemented. The track record of being neurally inspired beyond rather superficial ideas is not very good. However, we may hope that the explosion of new techniques in neuroscience may give us another shot at being inspired.

To make sense of the emerging large datasets we need ways of automatically analyzing them, ways of converting them into human communicable meaning. Towards this end I have been working with Eric Jonas on automatically classifying cells based on connectomics data. Our idea is to model at the same time the distance dependent connectivity and other aspects of the data. We find that the automatic technique is not much worse than humans at relevant tasks such as synapse prediction and shows some agreement with human anatomists.

Dhiresha Kudithipudi, *Rochester Institute of Technology*

Traversing the Application Landscape of Neuromemristive Computing

In contrast to conventional computer architectures, the human nervous system is inherently mixed-signal, massively parallel, approximate, and plastic, giving rise to its incredible processing ability, low

power, and capacity for adaptation. We realize this new class of architectures using NeuroMemristive systems (NMS), which encompass memristive/CMOS devices for synapses and neurons. NMSs abstract the computational principles found in the nervous system rather than mimicking them. In this work, we present different spatio-temporal processing applications such as epileptic seizure detection and speech emotion recognition. The overall architecture design was based on the echo state network and the classification accuracies of these systems are comparable to state-of-the-art software and hardware systems, with 3x lower power. The circuit-level implementation costs (power, area, etc.) as well as higher-level metrics, such as learning capacity, are compared using both conventional and unconventional activation functions. Device models based on experimental data are used to capture the non-ideal memristor behaviors, such as bistable non-linear devices. We demonstrate for a speech recognition application that in spite of using bi-stable devices, with limited synaptic weights; the performance of the NMS is comparable to the performance of a software-based realization. Robust and novel methods of training NMS on-chip from different data related to the trained off-chip network are proven on early visual system feature detectors. When NMS was tested with Caltech 101 dataset we achieved classification accuracy similar to support vector machines. Finally, an NMS based security primitive for Keccak algorithm is shown which mitigates side-channel attacks.

Tai Sing Lee, Carnegie Mellon University

Neural circuits for learning internal models of the environment

I advocated and presented a top-down approach to understanding cortical neural circuits based on (1) developing computational models that can solve vision problems and (2) training the circuitry of such models with stimuli from natural scenes. A key insight from perception is that vision is constructive in nature. Many early vision problems, such as 3D shape from shading or from binocular images, can be formulated as inferring the internal models and their parameters that can then synthesize images to match the input images. An important model in computer vision for solving this class of problems is called Markov random field. MRF uses local interaction to mediate contextual information to remove local ambiguity. It also explicitly or implicitly generates images synthesized higher order representation (e.g. 3D shapes) or transformation (e.g. transform between left eye and right eye images) as a way to infer abstract visual concepts such as shapes and depth. A neural implementation of Markov random field is Boltzmann machine, which represents the continuous variable at each node of MRF using a population of neurons with tunings that tile the domain of the variable to be inferred. Training a Boltzmann machine with disparity data derived from 3D natural scenes yielded a connection matrix that describes well the functional interaction of disparity-tuned neurons we measured using multi-electrode recording in macaque V1. Simulation of such a network model predicted neurophysiological observations on how the disparity tunings of V1 neurons sharpen over time, or with an increase in stimulus size, and the phenomena of filling-in of surround disparity signals. This work shows that MRF and Boltzmann machine provide a viable model for characterizing and conceptualizing the computational neural circuits in the visual cortex. However, I also presented some new predictive coding evidence and models that illustrate how interneurons can encode sophisticated compatibility constraints beyond what are possible in the pairwise connectivity furnished by the Markov random field model or Boltzmann machine.

Hai (Helen) Li, University of Pittsburgh

The neuromorphic computing leveraging emerging devices

In past years, the development of neuromorphic computing systems has been extensively studied. However, conventional CMOS technology approaching the scaling limit is not sufficient to meet the requirement of neuromorphic computation systems demanding extremely high level parallel operations and large scale storage. The emerging devices that generally feature non-volatility, high cell density, nanosecond access time and low operation, instead, demonstrate great potentials. Even thorough, how to leverage these devices and maximize their advantages in large scale systems remain unsolved. Some key challenges include but not limited to: 1) the feasible device characteristics that can provide sufficient design space; 2) the circuit types and structures that can utilize the advantages of emerging devices; 3) general or specific architectures that enable the use of neuromorphic circuit components; 4) an automatic design flow that is transparent to user; 5) compiler and OS that can seamlessly integrate various applications to the emerging devices based neuromorphic accelerators or systems. In summary, new devices such as memristors have triggered the rebirth of neuromorphic computing and demonstrated its potential and significance in computation intelligence. A holistic scheme integrating the efforts on device, circuit, architecture, design automation, compiler, etc. is very necessary to fully leverage the advantages of these new devices.

Karlheinz Meier, Heidelberg University

Mixed-signal accelerated Systems Progress, Results, Plans

A physical model neuromorphic system has been developed in the EU BrainScaleS project and is scaled-up in the framework of the EU Human Brain Project (HBP). The system features local analogue computing with binary spike communication via configurable point-to-point asynchronous links in continuous time. All time constants are scaled by a factor 10.000, so the system is accelerated and can compress a biological day to 10 s in electronics time. Neurons are implemented as Adaptive-Exponential (AdEx) neurons with up to 16.000 synaptic inputs. Synapses feature short-term plasticity as well as STDP. A maximum of 200.000 neurons and 50.000.0000 synapses are integrated on an 8 inch wafer using wafer-scale-integration technology. A total of 20 wafer modules featuring 4 Million neurons and 1 Billion synapses are assembled in a phase 1 neuromorphic system in HBP. The system offers a very high degree of configurability and a complete software package for non-expert users.

The main objective of the accelerated approach is to evaluate large parameter spaces through rapid prototyping and to study plasticity and learning processes, which are otherwise inaccessible.

Several experiments have already been carried out. An implementation of the insect olfactory system performing multivariate data analysis demonstrates the advantages of neuromorphic computing in a single experiment. The circuits perform as good as classical machine learning implementations but consume a Million times less energy than conventional computers and converge 10.000 times faster than biological implementations. Also, the implementation uses elements with a large variability of typically 20%, which is compensated by neuron populations. A new development has shown that networks of point neurons can represent probability distributions in any space of binary random

variables and perform stochastic inference in this space. Any type of graphical model can be transferred to Boltzmann machines.

Current improvement plans include the implementation of structured neurons with branching dendrites, dendritic spikes and back-propagating action potentials. Another improvement is realising a plasticity processor that has access to internal and external signals to control the network structure on-the-fly. The capability to implement structural plasticity is of particular importance for the HTM model.

A collection of publications describing the system and the experiments performed so far can be found on this website: <http://www.kip.uni-heidelberg.de/user/meierk/research/publications>

Murat Okandan, *Sandia National Laboratories*

Neuro-inspired Computational Engines: Beyond von Neumann/Turing Architecture and Moore's Law Limits

As we reach the performance limits of conventional computing approaches, neuro-inspired computing presents an attractive development path for continuing to improve our computing capabilities, as well as enabling new data analysis, prediction and control capabilities. How information is represented, processed, stored and recalled in these new systems, namely a sparse, spatio-temporal, hierarchical representation scheme, is at the core of this new approach, in contrast to symbolic representation and manipulation in conventional computing systems. By leveraging current computing systems, emerging neuromorphic computing platforms, neuro-inspired algorithms, novel devices and insights from neuroscience, it is expected that these new systems can be built, tested and improved at an accelerating pace. Some of the critical applications that will drive maturation of the first instances of these platforms exist in the academic, commercial and government sector, which will provide opportunities for rapid development of capabilities with the required coordination.

Randal O'Reilly, *University of Colorado, Boulder*

Biologically-based Error Driven Learning in Thalamocortical Circuits

The question of whether the brain uses something like error backpropagation to learn has been fraught with controversy since the algorithm was developed in the 1980's. Given its impressive computational power, which is being newly re-appreciated, and some promising initial progress, I have been working on resolving three major biological problems: (a) How to propagate the errors? (b) How can local synaptic mechanisms implement the proper learning rule? and (c) Where do the supervised target values come from? In 1996, I showed that bidirectional connectivity can propagate errors in a very natural and automatic manner, which closely approximates the backpropagation error gradient. A few years ago, I showed how a STDP-based learning mechanism produces an appropriate synaptic learning dynamic. Currently, we are leveraging certain properties of the thalamocortical circuit to implement hierarchical auto-encoder networks that could bootstrap a full solution to the problem of where the target values come from. This new framework also includes predictive learning over time, and initial results show that this learning in the early visual pathway can improve figure-ground discrimination in cluttered visual scenes. More generally, it has been argued that bidirectional connectivity is essential for

consciousness, and a number of associated functional benefits, and our models are the only high-performance object recognition models that incorporate full bidirectional connectivity and associated constraint satisfaction dynamics. All of this work takes place within a unified biologically-based cognitive architecture called Leabra, which we think provides a compelling balanced model of learning and processing in the brain.

Xaq Pitkow, Rice University

How can we know if the brain is doing a good job?

To understand and emulate the brain, we need to explain both how sensory signals drive neural responses (encoding), and how neural responses drive behavior (decoding). Models of encoding have grown in sophistication, while models of decoding have lagged behind, considering only linear readout of neural activity. Since these models guide our data analyses, we have been stuck with limited options to analyze experimental data. The standard data analysis method boils down to computing linear correlations between the choices of a behaving animal and its neural activity. This is a reasonable measure for simple tasks and the right brain areas, where the mean activities of neurons encode the relevant task variables. Indeed, recent theoretical work shows how to use this measure to infer a linear readout that explains the animal's behavior. However, for natural conditions and natural tasks, this is not possible: mean neural responses are confounded by task-irrelevant variables, so those means do not provide information about the task. We show how this induces a nonlinear code. Here we introduce a new way of thinking about decoding that is appropriate for these more challenging, nonlinear natural tasks. This comes with a practical experimental test of the quality of an animal's nonlinear decoding strategy. When we apply this test to neural responses recorded from primate visual cortex, we find intriguing evidence that animals are using efficient nonlinear decoders. Knowing the nonlinear code is a central component of extracting algorithms for neuromorphic computation, so these results provide a helpful tool in understanding when we have identified the relevant computations in the brain.

Fred Rothganger, Sandia National Laboratories

Can memristors learn?

My presentation consisted of three unrelated ideas for consideration by the community:

1. Cybernetics should be the framework for a "grand unified theory" of the brain -- The field of Cybernetics proper has been developing theory for the last 70 years on many themes that come up in our discussions. It would behoove us to be aware of the framework that already exists. Why are communication and control important concepts for a theory of the brain? At the cellular level, a living system concerns itself with regulating the production of various molecular products. These in turn form a network of interactions that adapt the behavior of the system to the environment, thus sustaining the organism. Molecular communication between cells enables multicellular organisms, and ultimately brains. Our brain processes signals and computes actions using a large number of interacting elements, and its goal is also to sustain the existence of the organism (or, some may argue, the species).
2. The backpropagation algorithm is not viable on a memristor crossbar -- We measured a TaOx

memristor on the bench using random-pulse sampling, and used the resulting data in a table-lookup ("non-parametric") style emulation. The Backprop algorithm required very finely controlled voltage pulses right at the programming threshold of the device. The expected programming noise was orders of magnitude larger than the desired step size. We extrapolate that any algorithm requiring small finely-controlled steps will be unsuitable. We proposed a random walk method, whimsically named LOTTO (Lazy Optimization Through Tesselated Oversampling), which more closely fit the characteristics of the memristor. I noted that Dhireesha Kudithipudi's earlier talk mentions other algorithms which cooperate with the characteristics of memristors and may be much more effective than random walk.

3. Use N2A to write neural algorithms -- N2A (Neurons to Algorithms) is a purely declarative language (comparable to NeuroML and NineML) for describing large dynamical systems. It is object-oriented, in the sense that the user describes classes of components which get instantiated at runtime, as opposed to describing specific instances of components (contrast with Simulink, PowerSim, VinSim, etc.) It represents neural components in an ontological structure, much like classic object-oriented programming languages (C++, Java, Smalltalk, etc.) It has concepts like inheritance and inclusion, with well-defined rules for combining models into larger systems. Our goal is ultimately to write a model of the whole brain, as a community effort.

Catherine Schuman, University of Tennessee, Knoxville

A Programmable Array of Neuromorphic Elements

Neuroscience-inspired dynamic architecture (NIDA) networks are spiking neural networks with three main features: very simple neuron implementations (with respect to biological neurons), synapses that have programmable delay lines and thus have memory, and optimization of structure and parameters using evolutionary optimization. NIDA networks have been successfully designed for applications in control, anomaly detection, and classification. Dynamic Adaptive Neural Network Arrays (DANNAs) are hardware implementations of NIDA networks, currently implemented on FPGAs. DANNAs are arrays of programmable neuromorphic elements, where each element may be programmed as a neuron, synapse, or connector. This programmability allows for various network structures (various numbers of neurons and synapses and various connectivity patterns) to be implemented on DANNAs. We are creating a software development kit that will interact with DANNA networks in hardware and will include user interface, visualization, array control, array interface, and array monitoring capabilities. We are also developing a VLSI implementation of DANNA networks that will allow for larger array sizes, lower power consumption, higher clock speeds, and more advanced monitoring capabilities than the FPGA implementations. We also plan to implement on-line learning mechanisms that allow for parameter and structural evolution in DANNA networks in hardware.

Lloyd Watts, Neocortix

Event-Driven Simulation of Spiking Neural Networks

Abstract: Event-Driven Simulation of Spiking Neural Networks with point-neuron primitives and constant-current synaptic pulses is capable of simulating realistic spiking neural dynamics, provided that these primitive elements are suitably combined to make realistic composite neurons, and provided that

the synapse primitive is capable of representing a summing state variable. We give examples of adapting tonic bursting neurons (modeling Calcium-dependent Potassium Channels), the full set of Izhikevich benchmark neural patterns, nonlinear propagation of Sodium spikes in a Dendritic Tree, and Chaotic Spiking behavior. Compared to differential equation solving methods, event-driven simulation is extremely fast (even faster with GPU support in modern implementations like CARLsim), and has no unwanted synchronization artifacts associated with overly coarse choice of integration time-step in differential equation methods. Event-driven methods are useful both for very fast simulations and may also be suitable for direct implementation on multi-processor hardware such as SpiNNaker.

Ken Whang, *NSF*

Neuro-inspired computing at NSF

NSF has broad interests at the intersection of brain and computing research. This increasingly rich space of interactions includes work focused on brain-like or brain-enabled understanding of machines, and on machine-like or machine-enabled understanding of brains. Productivity of these interactions depends on deep understanding of what is known or knowable, a good sense of what is useful and feasible to abstract out, and ongoing determination of what is helpful in principle and in practice. NSF has been actively seeking community input to help inform its efforts in this area. Funding opportunities include interdisciplinary programs (e.g., Integrative Strategies for Understanding Neural and Cognitive Systems) and disciplinary programs primarily within the CISE and Engineering directorates. More information may be found under the Funding tab at <http://www.nsf.gov/brain>.

Winfried Wilcke, *IBM*

The IBM Cortical Learning Center

The talk will give an overview of the new IBM Cortical Learning Center CLC. It will complement existing project in IBM Research by its strong focus on learning, especially unsupervised and continuous on-line learning. The underlying technology is the Hierarchical Temporal Memory model of the neo-cortex, which has been pioneered by Jeff Hawkins and his company Numenta and which is the subject of collaboration between IBM Research and Numenta. The talk will not give a detailed overview of HTM, because that is covered in an earlier talk in the Workshop, but it will describe the motivation for selecting HTM and outline the research direction for the CLC, which includes continued development of the HTM algorithms, software implementations, applications and HTM specific hardware. Some applications will be demonstrated as videos and the saccadic vision application, which is under development, will be discussed in some detail.

Alan Yuille, *UCLA*

Complexity and Compositionality

This talk introduces compositional models of visual objects. We illustrate these models on simple datasets demonstrating their ability to do unsupervised learning in presence of background clutter, to discover multiple object classes, and to learn from a small numbers of examples. We describe a hierarchical visual architecture which enables multiple object classes to be efficiently stored and rapidly

accessed by a bottom-up and top-down inference algorithm. The objects are represented in a hierarchically distributed manner in terms of parts and subparts, which are constructed recursively by part-subpart compositions. Part sharing between different objects yields efficiency in representation and inference. Parts are represented more coarsely at higher level of the hierarchy, so that the upper levels give coarse summary descriptions (e.g., there is a horse in the image) while the lower levels represents the details (e.g., the positions of the legs of the horse). This hierarchically distributed representation obeys the executive summary principle, meaning that a high level executive only requires a coarse summary description and can, if necessary, get more details by consulting lower level executives. Theoretical analysis shows that this architecture can yield exponential gains in terms of representation and inference.

Appendix E: Website Data for nice.sandia.gov

VISITORS* 9,062 Avg/day: 20 The total number of unique visitors to the site	PAGE VIEWS* 44,625 Avg/day: 100 The total number of pages viewed by all visitors	SESSIONS* 25,001 Avg/day: 56 The total number of visits to the site by new and returning visitors
--	---	--

*Time period: January 1, 2014 – March 20, 2015

