Brain: Analysis Synthesis Computation

Organizers: Hanspeter Pfister, Georg Seelig

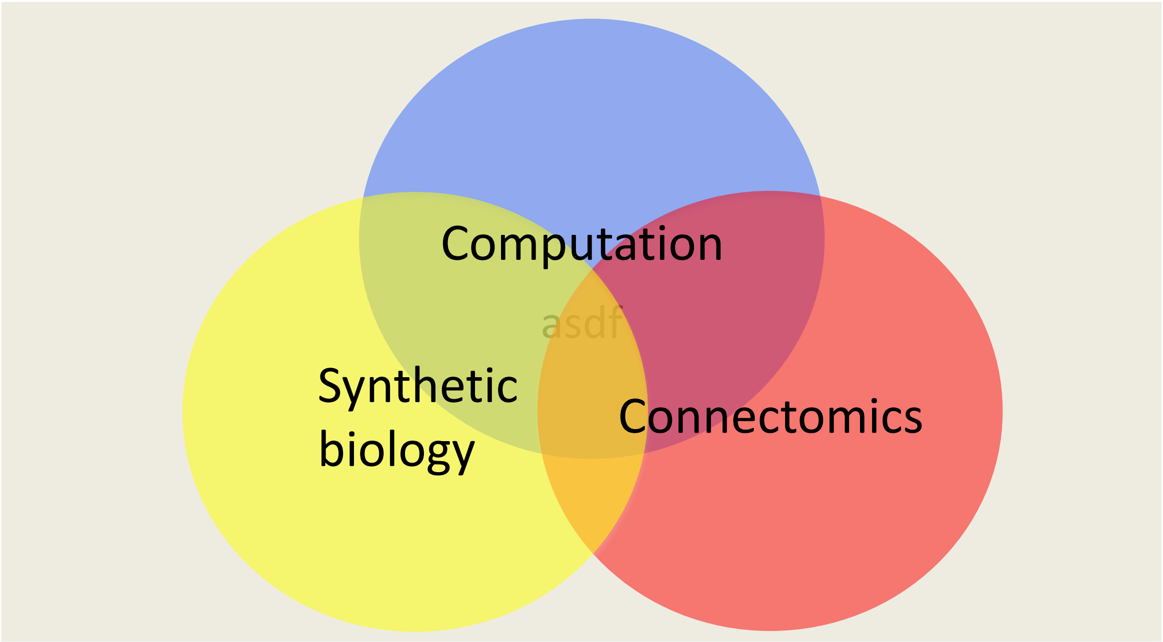

- Top-down: How can brain circuit analysis inform computational methods?
- Bottom-up: How can a complex structure like the brain be "computed" by simple cells?

Participants

- Expected 40 or fewer
- ~75 accepted, ~60 attendees
- ~10 DARPA, ~20 ISAT members, ~30 outside invitees
- Fields represented:
  - Neuroscience, CS, Engineering, Biology
  - Industry: Defense, Private (Qualcomm, Brain Corp.)

Agenda

- Day 1
  - 3 keynotes to set the tone: Adrienne Fairhall (UW), Jeff Lichtman (Harvard), Robert Kirkton (Duke)
  - Participant readout of prepared statements (received 25 statements)
  - Panel: Current and past government efforts; NIH: Susan Volman; Darpa: Geoff Ling (DSO), Gill Pratt (DSO), Dan Hammerstrom (MTO))
  - Breakouts along the three major themes (A/S/C)
  - Dinner with keynote (Jack Gallant, Berkeley)
- Day 2
  - DARPA PM for a Day breakouts
  - Summary plenary session

Jack Gallant: Evening keynote

Mapping cognitive function in the human brain

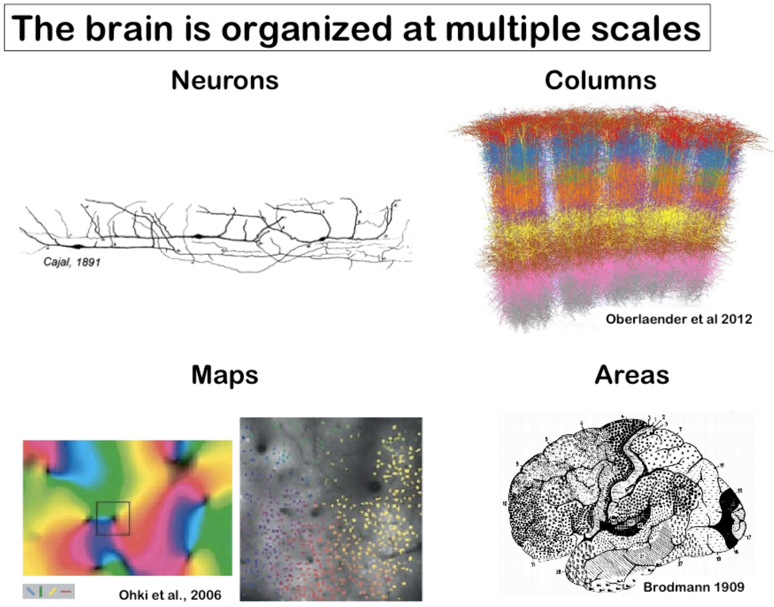

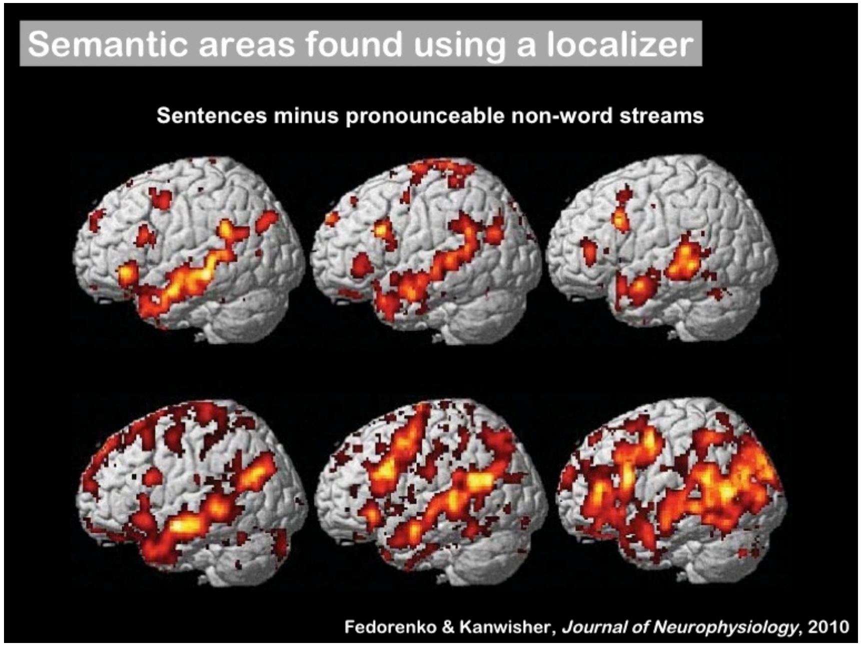

Jeff Lichtman: Analysis

Connectomics

The sparse world view

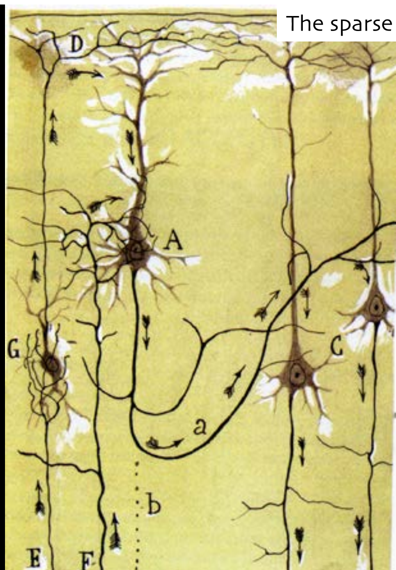

Ramón y Cajal, S. Estudios Sobre la Degeneración y Regeneración del Sistema Nervioso (Moya, Madrid, 1913–1914)

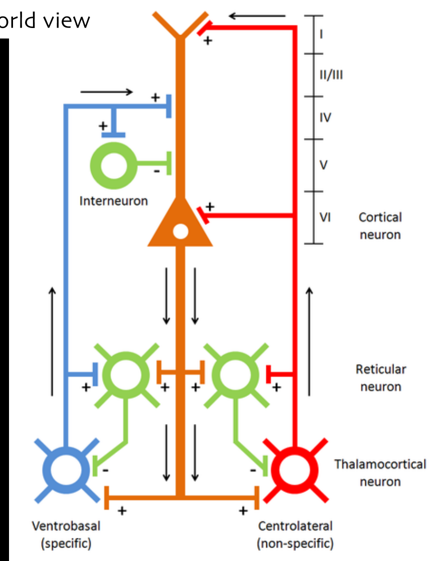

Llinas, Leznik, Urbano 2002

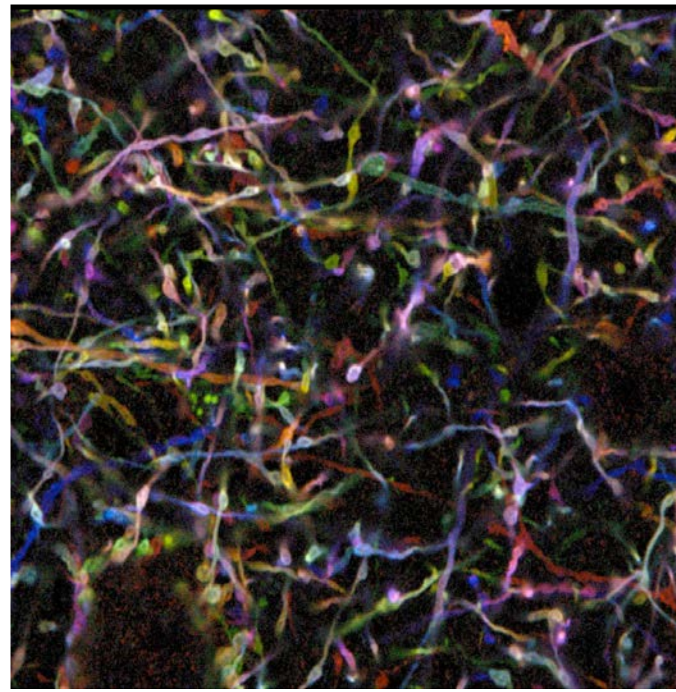

Cerebral cortex inhibitory axons only via infection with a brainbow virus

But what if we want to see *all* the connections?

Dawen Cai et al., 2013 & Luke Bogart, Takao Hensch

Rob Kirkton: Synthesis

Genetic Engineering of Excitable Cells

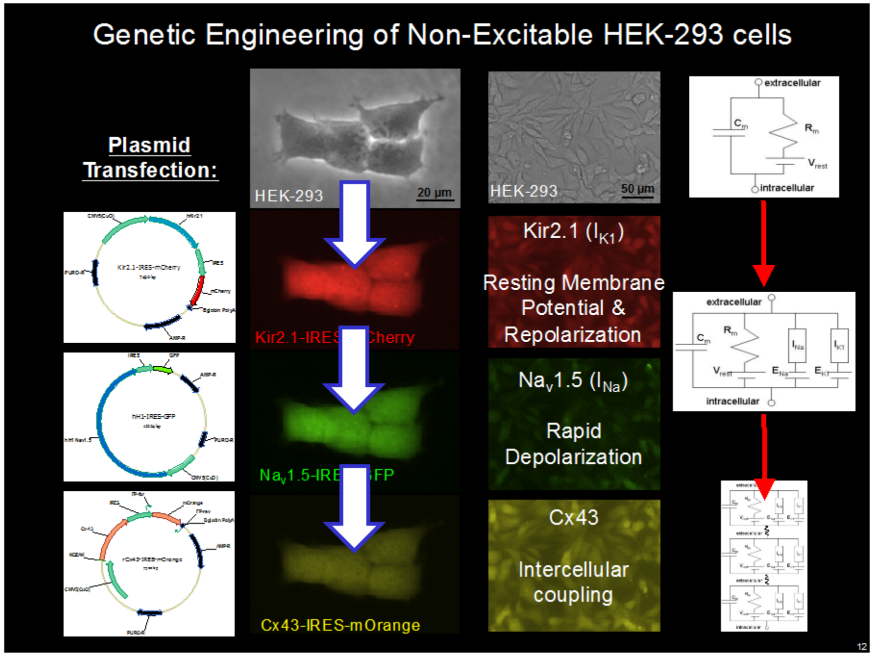

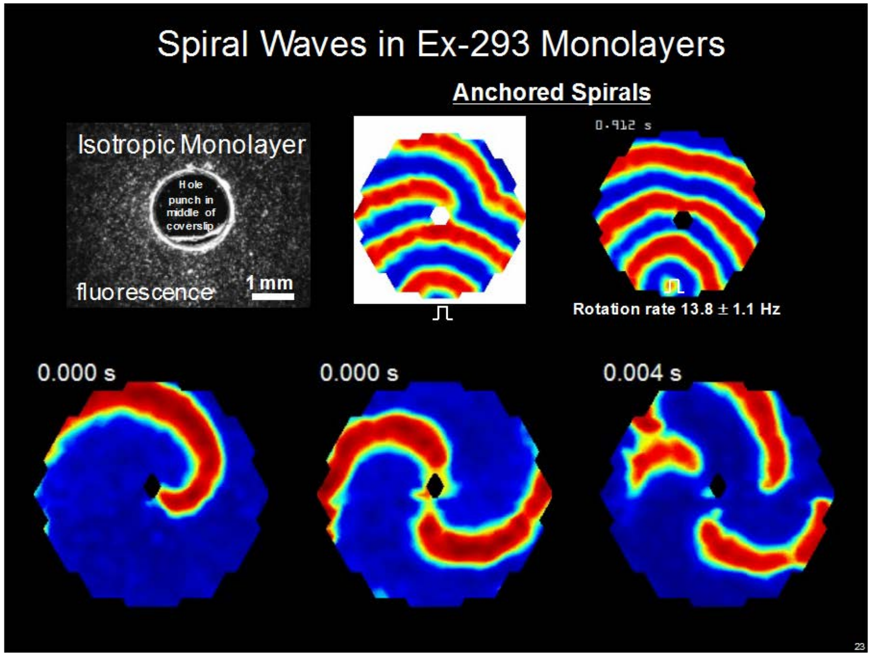

Approved for Public Release,
Distribution Unlimited

Adrienne Fairhall: Computation

Simulations in neuroscience: what, how, why

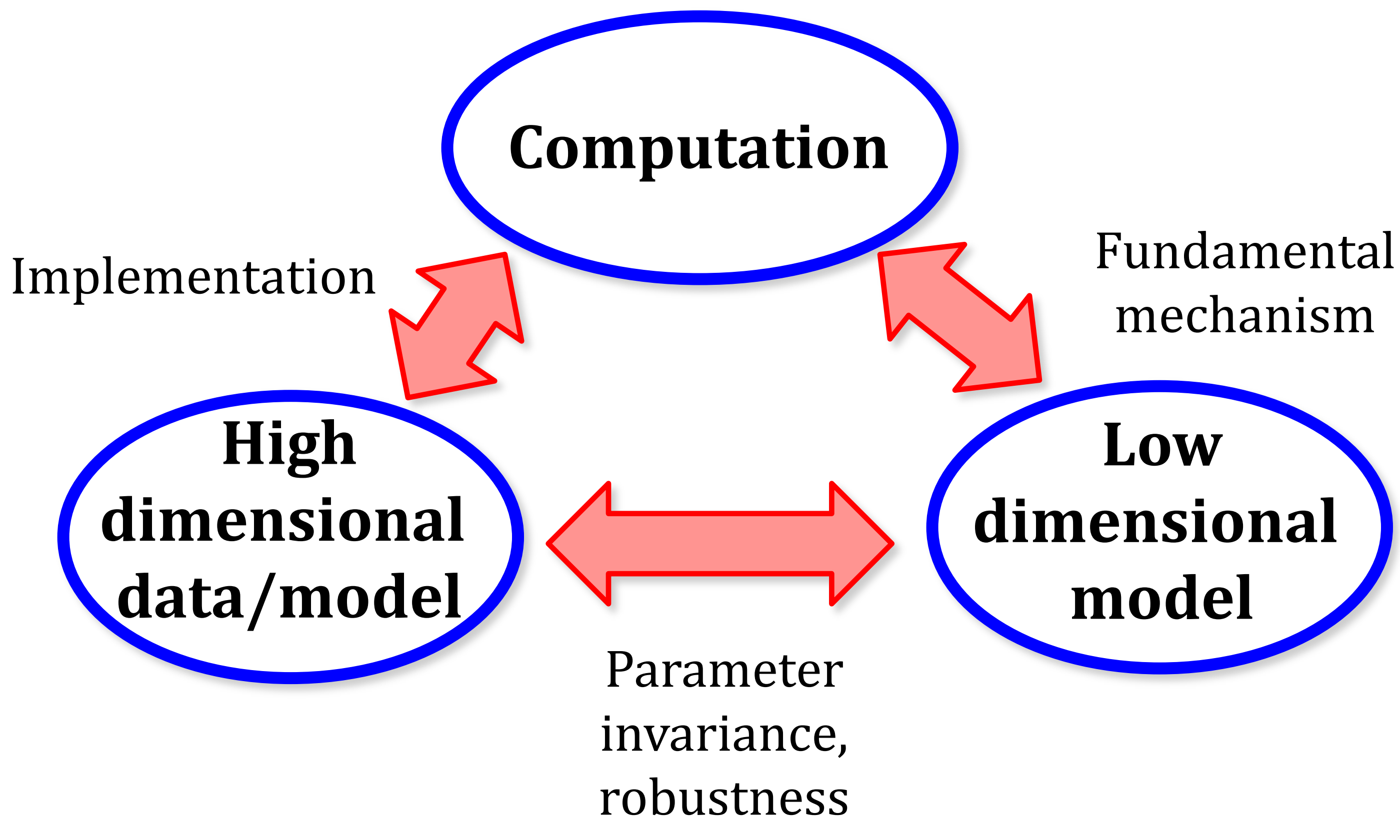

Breakout along 3 major themes

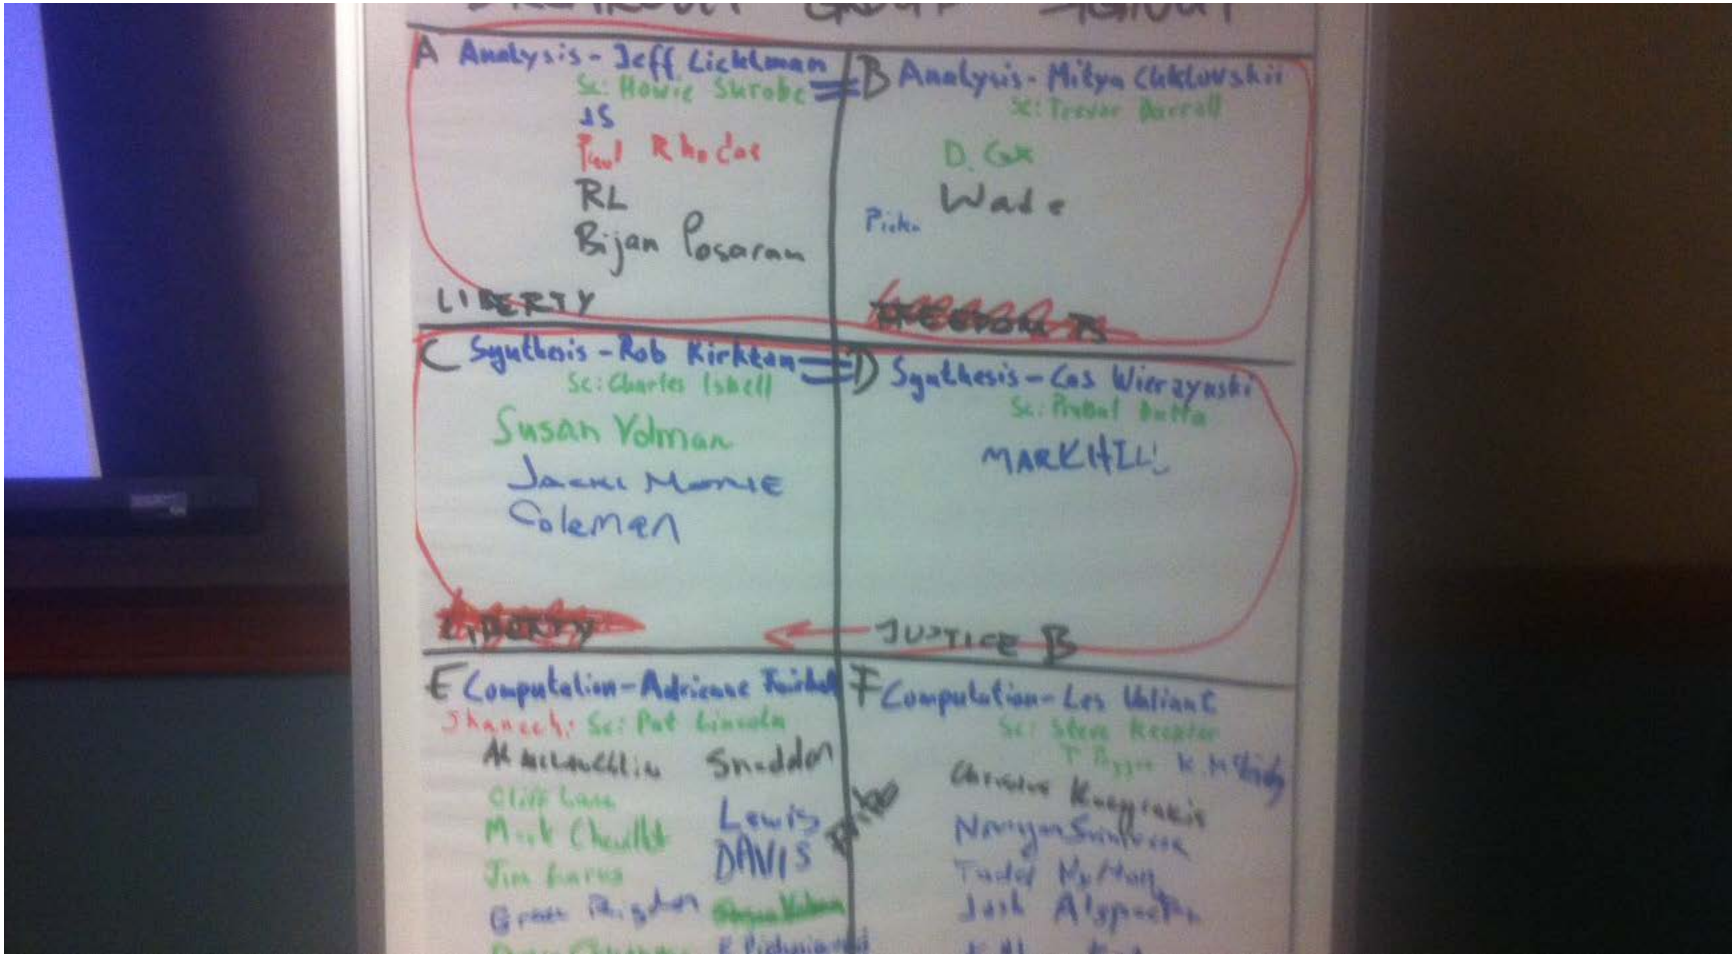

Approved for Public Release,
Distribution Unlimited

DARPA PM for a Day

- Develop pitches for a DARPA program
- Small breakout groups
- Use Heilmeier catechism:
  - What are you trying to do? Articulate your objectives using absolutely no jargon.
  - How is it done today, and what are the limits of current practice?
  - What's new in your approach and why do you think it will be successful?
  - Who cares?
  - If you're successful, what difference will it make?
- Plenary session with 1-2 slides per group
- Vote on the top three ideas

Program suggestions

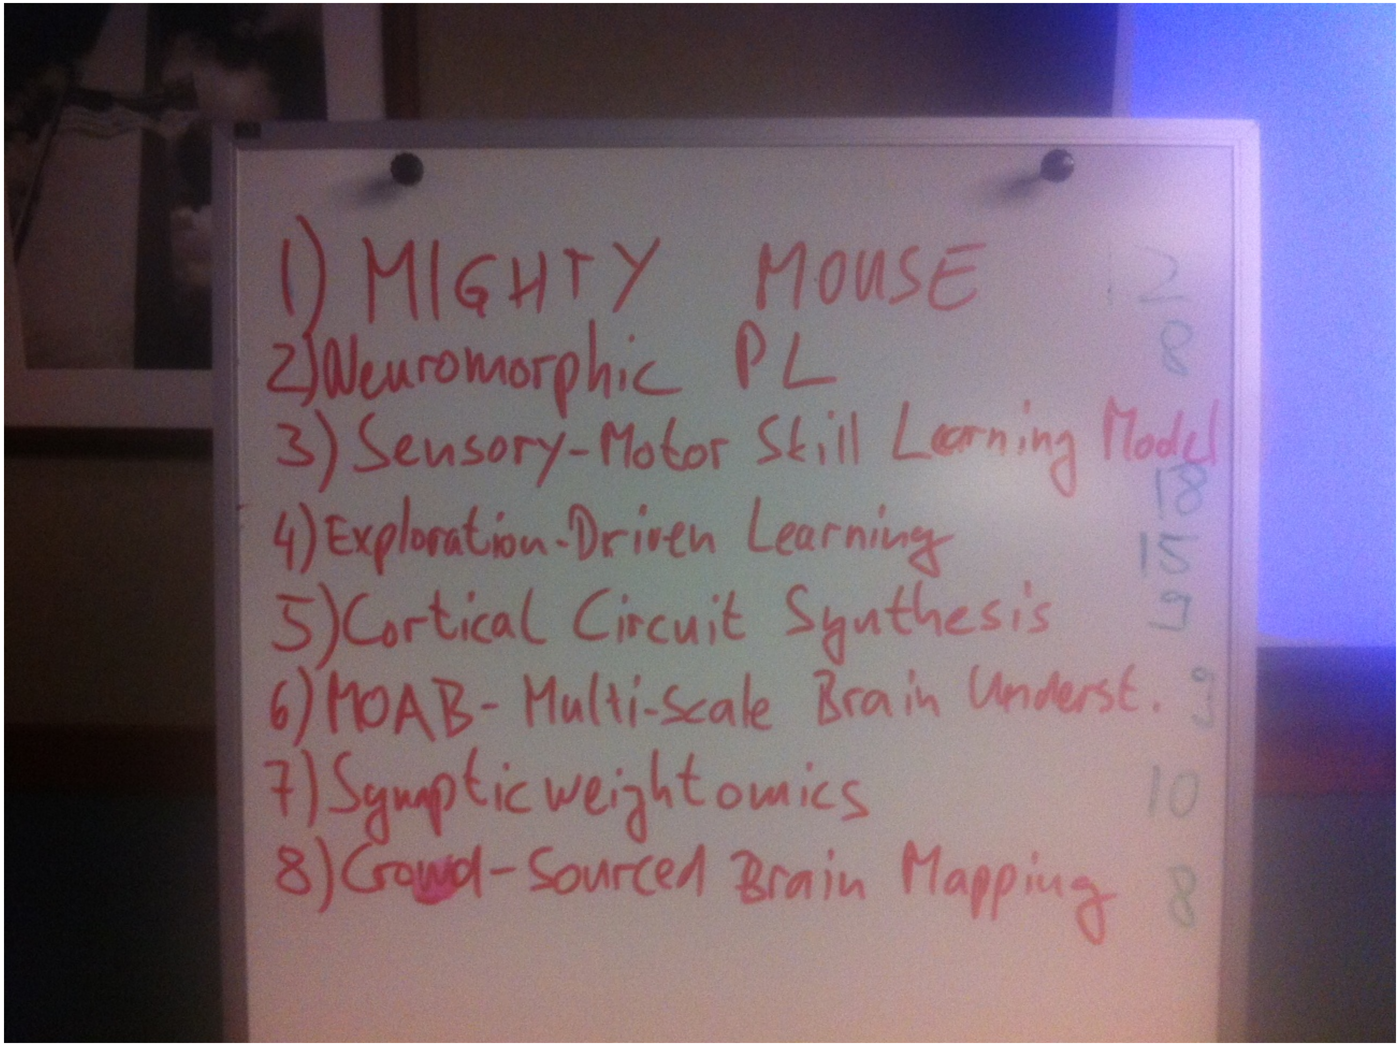

Some common themes

- What are the computations that brains are good at?
  - Sensorimotor integration: combining sensory information (object recognition, vision,..) with motor information (information about the state of the body)
  - Output to motor planning: how sensory inputs are integrated to transform into motor planning and action, how can this be scaled up realistic levels
  - Learning (closed loop learning)
- Can similar brain tasks be simulated? Is there a theory?
  - Theory: go from low-level (neuron) description to higher level description; coarse-grained over complex and diverse building blocks
  - Machine learning (neural network models), closed loop learning, learning of sensorimotor tasks
  - Experimental constraints on theory: measuring synaptic weights (Neural network models (deep believe networks): object recognition, etc all have the same "connectome," substrate of learning is synaptic weights)

Some common themes

- How are these computations related to the circuits (e.g. cortical columns, what kind of computation does a cortical column implement)?
- Experiments
  - Connectivity: functional (synaptic weights…), change during learning, also neuromodulaters can reconfigure circuits by adding and removing neurons, static connectome just (a subset of) potential connection
  - Speed up connectivity measurements (by orders of magnitude, cf. genomics)
  - Novel high throughput methods of functional activity measurements including cell type specificity based on biophysical properties of neurons/brains; novel methods to read out information deeper in the brain than is possible with photostimulation
  - Build (small) brains from scratch to confirm our understanding
- Goals:
  - better understanding of learning->better training of people
  - Integrate neuromorphic engineering with moving robots ()

Conclusions

- Great talks, high-profile participants, good interactions between participants with very different backgrounds
- Many diverse ideas, discussion moved into many directions but, several common themes emerged but each might require a follow-up workshop
- If we did this again, we might try to encourage participants to think further outside the box (anonymous submission of ideas, prize for the craziest idea)
- Great feedback from attending DARPA PM, may have influenced at least one program currently being developed

Appendix A

Breakout Summaries Day 1

Analysis Breakout

Leader: Jeff Lichtman

Scribe: Dmitry Chklovskii

Members:

Howie Shrobe, Paul Rhodes, Basjan Posaran, Trevor Darrell, David Cox, Hanspeter Pfister, Johannes Seelig, R.L., Dave Warner, Filip Piekniewski

What do we need to know to understand how the brain works?

- Pick two target computations:
  - Algorithms that only humans can do, too difficult for computers – cognitive reasoning, high-level vision
  - Algorithms that computers can do but brains do more efficiently and robustly – low-level vision

- Data collection:
  - Map structure: Connectome mapping is polynomial. Vast information is stored in connection weights not in activity
  - Map activity: Activity mapping is exponential. Different scales spatial and temporal. How many neurons ?
  - Develop techniques. Combine different techniques on your favorite system

- Interpreting data:
  - To understand a complex system you break it up into basic components study them with all you've got then integrate them into the full system, elementary building blocks, enumerate them, how to put them together to do something complicated
  - The math for network analysis both static and dynamic does not exist. Better software and theory tools to analyze existing datasets
  - Structure-function relationships. All scales. Connect scales. There is a need to understand functional maps. There is not one connectome
  - Vertebrate neurons easier to simulate than invertebrates?
  - Cognitive?

- Basic principles
  - What does a neuron do? Basic computation of a neuron
  - Canonical cortical circuit

Synthesis

Compare simulations in synthetic systems with experiments

DARPA ISAT :: BASC :: Synthesis++

Susan Volman, Rob Kirkton, Prabal Dutta,
Charles Isabel, Todd Coleman, Jacki Morie,
Mark Hill, Cas Wierzynski

Three questions

- IS for NS: What could info science do in the next decade to help neuroscience? (3D printing)
- NS for CA: What understanding can we extract from neuroscience to build better computers?
- PL to C. Elegans: How can we model cells to better understand their operation, and synthesize them from programming languages (“neurolog -> cells”)?

What’s hard?

- Brain is complex
  - Brain has lots of scales
  - Brain connections regulated by their use!
  - Inhibit/excitatory -> measure strength of those connections?
  - Why do we need these connections? Evolutionary redundancy?
  - How much of this complexity is really needed?
  - How much redundancy is there?
  - Engineers wouldn’t design a brain that way?
  - Too much data in the surface, but want data from the volume!
- Architecture
  - CS: Comp. Arch. Era, post-Dennard scaling → power wall
  - Power is the pain point now, looking for other approaches
- Programming
  - How do we “program” neurons? Abstractions? Compilers? Runtime?

Potential directions?

- Neuroscience Today
  - Can grow cells in 2D
  - Can image 1-2 layers deep with 10k fps camera
  - Can extract information from the surface
- Neuroscience Tomorrow
  - 3D print arbitrary cell patterns
  - 3D print blood supply to avoid rapid interior cell necrosis
  - 3D scanning of deep tissues (beyond 1-2 layers) in vivo. Quantum dots?

- SynBio Tomorrow
  - What would it take to build a c. elegans?
  - Want to manipulate cell/circuits/brains in a way that one understands
  - Programming language for neurons: from verilog to neurolog
    - PL to: Put in chemicals, do some math, get some results
    - What are the right abstractions?
  - Better understanding of neurons, using them to fix neurological problems

What is the state of the art?

- Can grow heart and brain cells in 2D
- Can image 1-2 layers deep with 10k fps camera

Computation Group A

Adrienne Fairhall, Pat Lincoln, Al McLaughlin, Cliff Lau, Mark Chevillet, Jim Larus, Grace Rigdon, Dave Gunning, Robert Sneddon, Tony Lewis, Joel Davis, Filip Piekinewski

Questions for Evolving Operational Research Initiatives

- **What investments that need to be made in the field of Computational Neuroscience that will help shape the next decade of growth and development (theory, methods, models, datasets)?**
  - Where are we with vision? What's left to understand?
  - What other modalities can we start approaching? Audition? Motor control?

- **What are the limiting factors stifling current research initiatives?**
  - Do we have enough computing power? Do we have the right kind of computing power? Do we even have the algorithms/theory to make any sense of the simulation?
  - Do we have any good neural recording modalities? What are the limits of current modalities, and what are the possibilities for new modalities?

- **How do we define quantitative metrics and standards?**
  - Do we care if we predict neural activity? Or do we care if we predict human behavior?

- **What is the application space?**
  - Simulate human sensory capabilities (sensory fusion, audition, somatosensation, olfaction, motor control, decision-making, sense making, navigation)
  - What computations COULD the brain support other than those that it already supports?

Computation Group B

Leader: Les Valiant

Scribe: Steve Keckler

Members:

Josh Alspector, Todd Hylton, Christos Kozyrakis, John Mandferdelli, Bill Mark, Kathryn McKinley, Tommy Poggio, Fritz Sommer, Narayan Srinivasa

8/12/13

- Two definitions of neuromorphic computation
  1) Modeling the brain for the brain's sake
  2) Inspired by brain to solve problems
- Challenges
  - Normative computational model of the brain
    - Need a theory that can be tested
    - Deriving/defining parameters for experimentalists to measure
  - Brain can't be studied in isolation
    - Brain/body/mind environment is necessary
  - Theory of cumulative learning
    - Closed-loop, on-line learning
    - As opposed to conventional machine learning which requires lots of labeled examples for training
    - What are the data representations in the brain?
  - Problem space in which brains are better than what we can do conventionally
    - What are the areas in which brain-inspired algorithms are better than non-brain-inspired algorithms.

Appendix B

Breakout Summaries Day 2

Cortical Circuit Synthesis

Members:

Tommy Poggio, Filip Piekniewski, Bijan Pesaran, Paul Rhodes, Pat Lincoln, John Manferdelli

Project Synthesis

- Synthesize data-constrained cortical circuitry on a large scale
- Use 3-d reconstruction from a 0.5x1x1 mm volume of sensory cortex in which sensory activity has also been observed
- Start with three layer 0.5 mm deep olfactory cortex of mouse, visual cortex in Phase 2
- Embed representation of invariances in the synthetic wiring of a cortical hierarchy during exposure to the sensory world
- Goal: the simulated synthesis of an artificial cortex wired with sensory experience, enabling invariant object recognition

Heilmeier

- Synthesize artificial cortex
- Today, negligible attention to wiring synthesis, neuronal/synaptic elements not very neural
- Has been limited by lack of wiring and activity data at sufficient scale
- Data-constrained self-organizing circuitry, linking connectivity and sensory activity
- Represents a new paradigm for machine learning of invariant object recognition
- This is a path towards neurally inspired sensory, then sensorimotor, then intelligent machines

Mighty Mouse

Members:

David Cox, Adrienne Fairhall, Alex Firpi, Mark Hill, Charles Isbell

Mighty Mouse

- **Understand & build mouse (subset) to move in familiar environment (vision, memory, motor)**
- Enablers
  - – Accessible genetics tools, e.g., labeling
  - – Machine learning, big data
  - – Mouse medium complex, cheap, etc
- Parallel programs: non-dependent, cross-fert.
  - – Top-down: Machine learning, AI, robots
  - – Bottom-up: Neural analysis to inform synthesis

Neuromorphic Programming Language

- At top
  - Expressive, but maybe not Turing-complete
  - Hides (much of) low-level issues
- At bottom
  - Neuromorphic integrated circuits
  - Wetware, e.g., cells or multiple neurons
  - Find computation descriptions of dynamical systems

DARPA Offices

- I20 – computer / decisions
- MTO – Microelectronics – better chips
- DSI Def Sci Off – medicine / prosthetics

–

Exploration Driven Learning

Members:

Bob Laddaga, Tony Lewis, Jeff Lichtman, Fritz Sommer, Christos Kozyrakis, Jim Larus

Methods of exploration-driven learning in unknown environments

- **Objectives**
- Understand drivers of a child's curiosity for learning a model of the world
- Learn from few examples chosen by your actions
- No separate training phase

- **State of the art**
- Neuroscience
  - traditionally separates perception and action
  - adult capabilities and their development
- Robotics
  - programmed actions or teleoperation
- Machine learning usually separates training (with access to all data) and inference phase
- Some work in theoretical psychology, internally-driven reinforcement learning

- **New ideas**
- Internally-driven autonomous learning algorithms
- Method for self-organizing behavioral repertoire
- Algorithms for learning temporal sensor-motor contingencies

- Build qualitatively new class of normative computational models for neuroscience
- Self-organize behavioral repertoire

- **Impact**
- Dawn of new age of robotics, neuroscience, artificial intelligence and data mining

Sensory-Motor Skill Learning Model

Members:

Eric Pohlmeyer, Johannes Seelig, Maryam Shanechi, Bill Mark, Karthryn McKinley

DARPA Pitches

1. Micro-canonical motor-sensory skill learning model
2. Model of cortical column learning
3. Computational paradigm for brain-like learning programmable system

Objectives

- Micro-canonical sensory-motor skill learning model
  - Build a simplified constrained system to study learning – Brain-machine-interface paradigm
  - Build a learning system
  - Understand learning in the brain
  - Help build better artificial learning models
  - Create clinical understanding (restoration, augmentation, rehabilitation?)
  - Turn the brain learning model into a programmable system
  - Use the model to enhance training people in new skills

Current Practice

- Motor cortex only recordings – not integrated
  - No cerebellum, basal ganglia
- Brain Machine Interface is not used as a tool for scientific research
- Limited connectivity and cell type knowledge
- Open loop decoding vs. closed loop control

What's New

- Closed loop control for understanding learning
- Simultaneous recording from multiple areas
- Using the closed loop model to enhance skills training
- Possible leading to a programming model
- Why does it work?
  - Constrained model in which to study learning
  - Well characterized
  - integrates system engineering tools together with brain science

Crowd-Sourced Brain Mapping

Members:

Dave Gunning, Jack Gallant, Steve Keckler, Todd Coleman, Cliff Lau, Rob Kirkton, Prabal Dutta

BASC :: Heilmeier Exercise

Dave Gunning, Jack Gallant, Steve Keckler, Todd Coleman, Cliff Lau, Rob Kirkton, Prabal Dutta

Areas of investment

- Measurement
  - Goal: Scalable brain mapping. Approach: Crowd-source it
  - Portable mapping: fMRI -> EEG Hat -> fNIR Hat -> Inexpensive “Google Hat”
  - What does a cortical column do and how does it interact with adjacent columns, layers, and thalamus? What information is encoded in these action potential signals?
- Simulation/Emulation
  - SimBioTic Computing: architectures that blend mammals and machine (AugCog?)
    - Harvest (and aggregate) Intuition: P300 -> Computer -> Conscious attention
    - Cells+Silicon=Cellicon? => Dog’s olfactory+chips -> Skunk alert
  - Why is the brain more efficient and robust at recognition tasks than computers
  - Cortical computer: simulate cortical columns (10k neurons) and their I/O
  - Can we grow cortical columns in the lab?
- Killer Apps
  - Google Hat: think it, and the answer appears (or is projected through direct stimulation)
  - Image labeling: pictures in; pixel-level taxonomy out
  - OCR: Digitize a century’s worth of DoD handwritten notes
  - Artificial Analyst: automating event detection
  - E2E Sensory-motor control: BCI for Predator targeting

Synapticweightomics

Members:

Leslie Valiant, Susan Volman, Casimir Wierzynski, Ramesh Raskar, David Warner

Ideas

- Physical properties of neuron morphologies - what might they be good for
- Synthetic excitable cells in 3D as computing devices
- Field effects - theta during learning, spindles during sleep, shape sleep architecture.
- "Vasculature" for neuromorphic hardware (how dealing with waste dictates structure of brain circuits; how do deal with heat)
- Cultured neurons with quantum dots as device
- Use neuroscience (fMRI, EEG, p300) to understand how best to display information. Use domain experts as subjects.
- Synaptome: distribution of synaptic weights. High-throughput tools for this measurement, which is key for constraining theory.

Synapticweightomics

- We want a high-throughput way to measure synaptic weights in the brain
- Currently done with paired recordings (10 pairs per day?) [and soon with connectomic techniques?]
- New because we want a high-throughput measurement: combine two-photon excitation & recording, structural EM, something else?
- Critical variable for (1) providing the "ground rules" constraining theories of brain function and (2) neuromorphic implementations. Also a key enabling technology for basic neuroscience.
- If successful, tool will (1) help understand what algorithms the brain uses (2) provide a useful biomarker for disease.