OFFICE OF THE DIRECTOR OF NATIONAL INTELLIGENCE

Machine Intelligence from Cortical Networks (MICrONS)

LEADING INTELLIGENCE INTEGRATION

R. Jacob Vogelstein, PhD
Intelligence Advanced Research Projects Activity (IARPA)
24 February 2014

Goal

Create a new generation of machine learning algorithms with human-like performance characteristics by using *cortical computing primitives* as their basis of operation

Definitions and Assumptions

- Cortical Computing Primitives
  - "Standard" operations performed by a given cortical area(s)
  - Combined into larger networks for higher-level functions
- Assumes some modularity of cortex
- Recurring microcircuit motifs embody primitives
  - Composed of $O(10^2)$–$O(10^4)$ neurons
  - Approximately 50–1,000 microns in diameter*

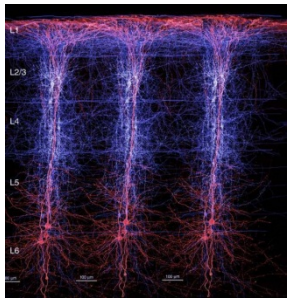

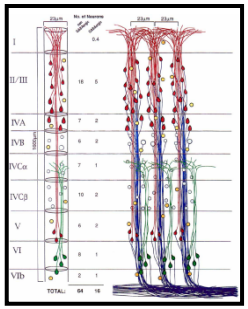

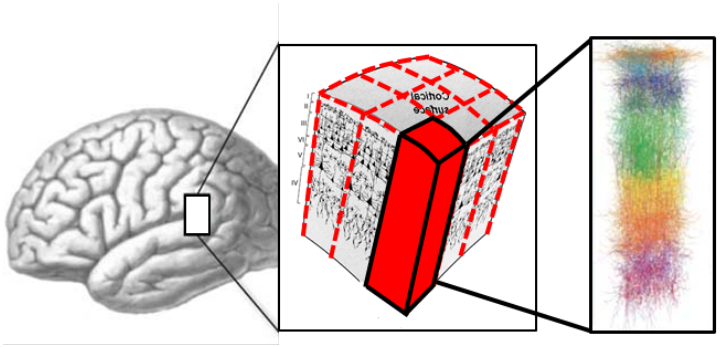

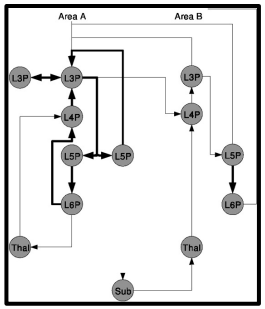

BBP/EPFL; Peters and Sethares, 1996; Douglas and Martin, 2004

Approach

- Reverse-engineer the cortical microcircuits that instantiate the cortical computing primitives
  - Identify their structure, function, and parameterization with high-resolution (single-synapse / single-spike) brain mapping tools
  - Capture results as attributed graphs / annotated schematics
- Identify the functions (primitives) performed by these microcircuits
- Use primitives as building blocks for novel machine learning algorithms

=

$X(t)\,Y(t)$

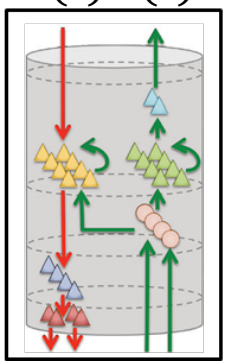

= $P(Y(t)) = \int_{t-T}^{t} f(X(\tau), \theta(\tau))d\tau$

Cortical Computing Primitive

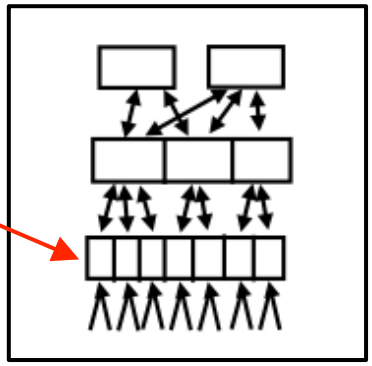

BBP/EPFL; Numenta, Inc.

Roadmap
(an example, not an endorsement)

Pre-existing work

MICrONS R&D

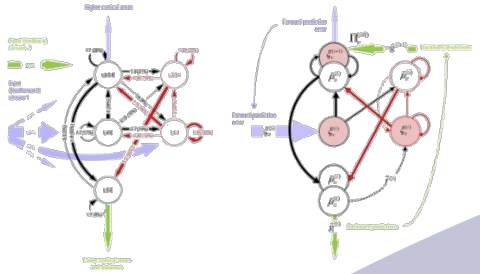

Interpretation of cortical microcircuits

Mathematical framework

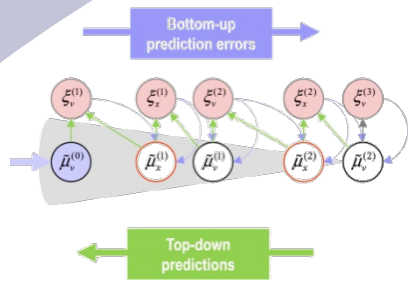

Bastos et al., 2012

Roadmap
(an example, not an endorsement)

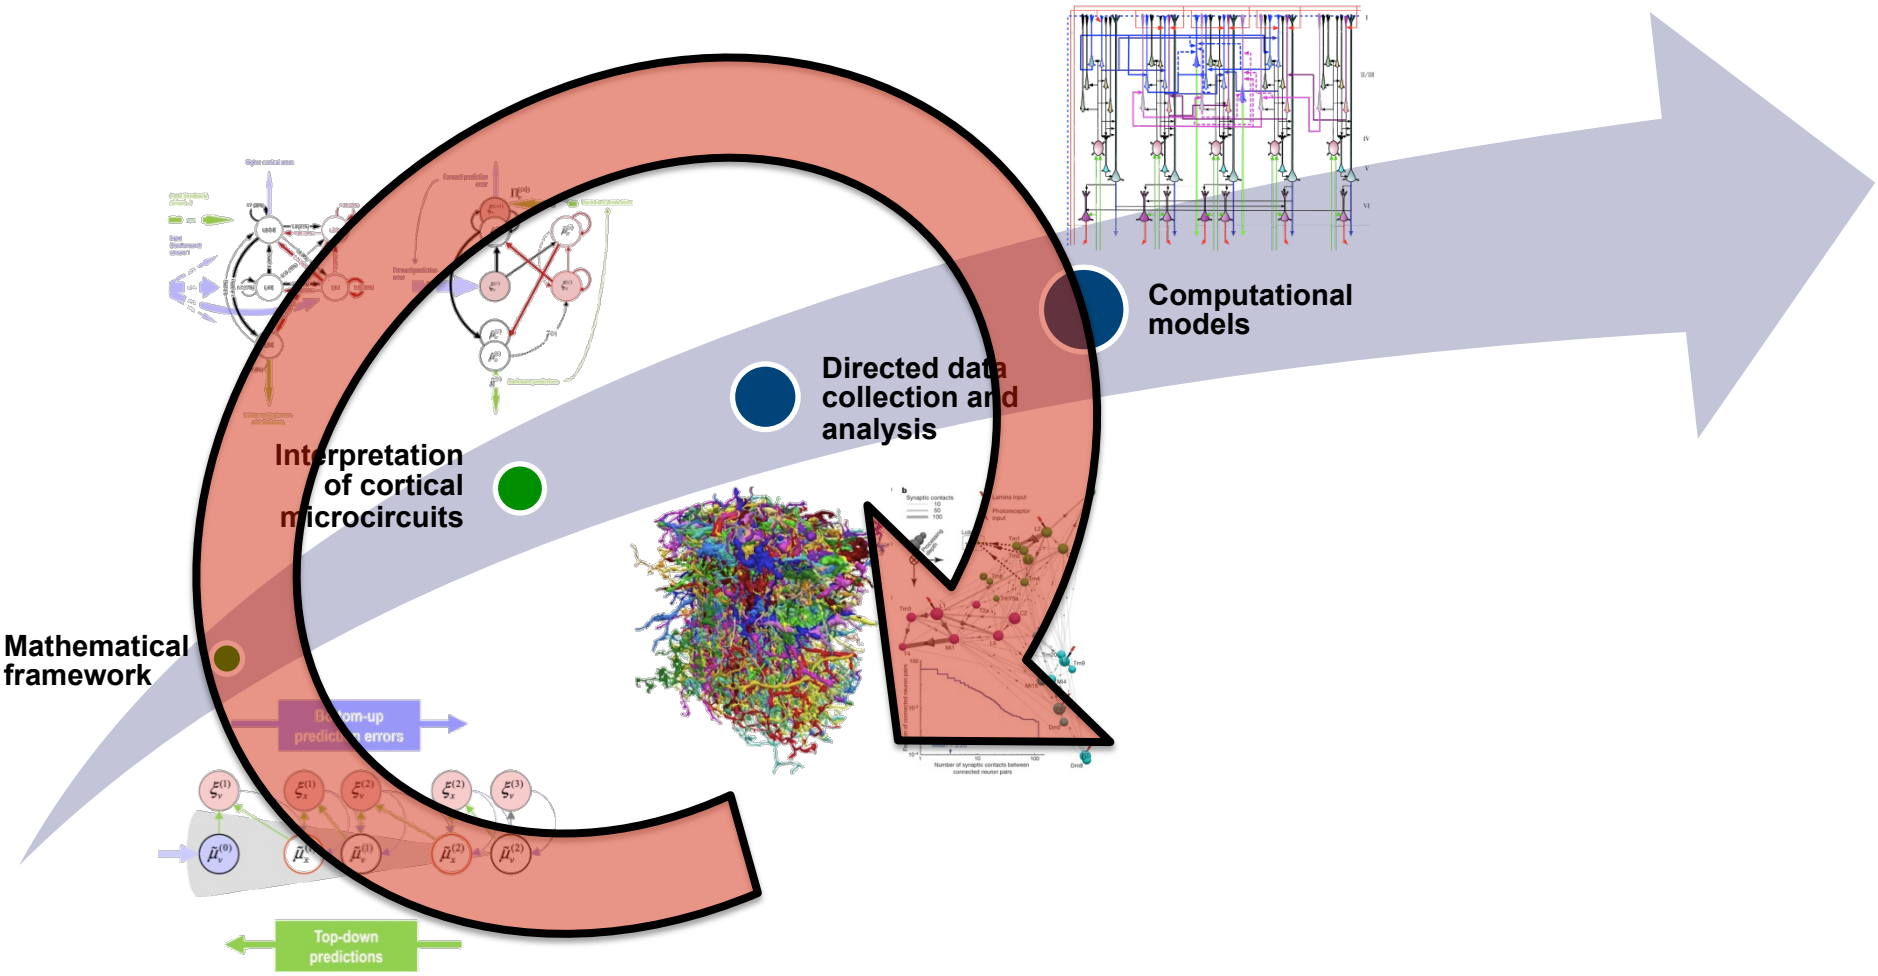

Bastos et al., 2012; Chklovskii; Takemura et al., 2013; George & Hawkins, 2009

Roadmap
(an example, not an endorsement)

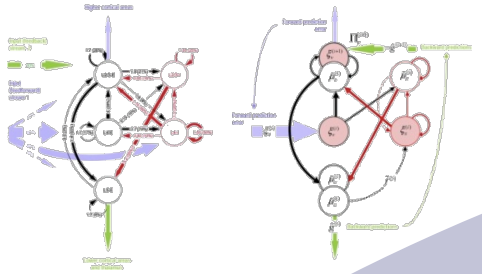

Interpretation of cortical microcircuits

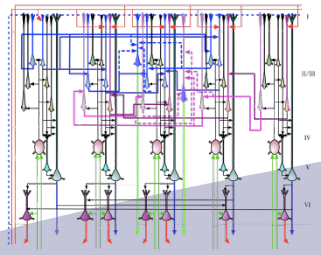

Computational models

Directed data collection and analysis

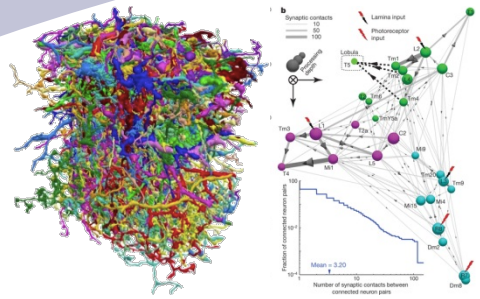

Machine learning algorithm

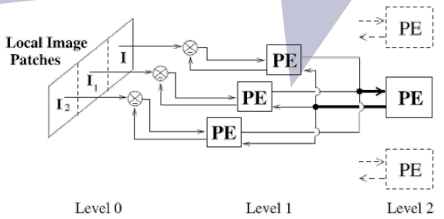

Mathematical framework

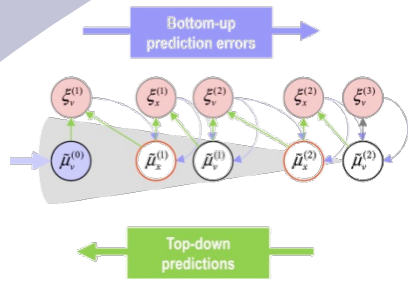

Bastos et al., 2012; Chklovskii; Takemura et al., 2013; George & Hawkins, 2009; Rao & Ballard, 1999

Roadmap
(another example, also not an endorsement)

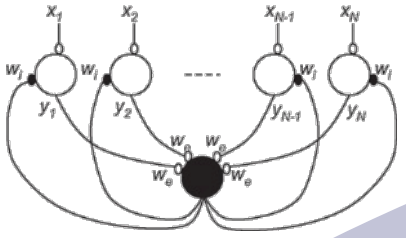

Interpretation of cortical microcircuits

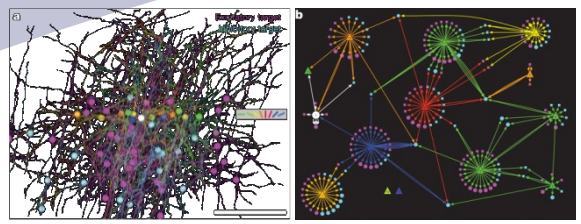

Directed data collection and analysis

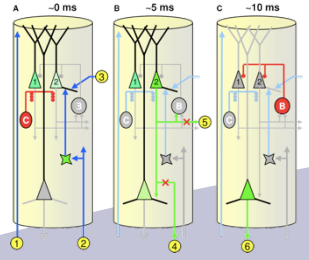

Computational models

Machine learning algorithm

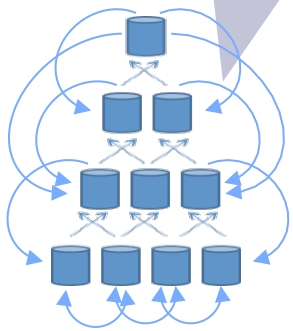

Mathematical framework

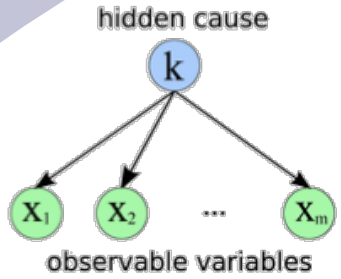

Nessler et al., 2013; Chan et al., 2010; Bock et al., 2011; Rinkus, 2010

Challenges and Reasons for Optimism

- "The largest challenge… is not in collecting the data… [but] to understand what we find." (Bargmann & Marder, Nature Methods, 2013)
- There is hope:
  - "Choosing well among brain regions, and combining connectomes with molecular and functional information about the same cells… will lead to the most informative results." (Bargmann & Marder, Nature Methods, 2013)
  - "… A few testable models emerged directly from analyses of anatomy. One was the concept of a motif… Perhaps these motifs perform a canonical computation, or a few canonical computations, so that solving a few of them effectively solves a larger piece of the diagram." (Bargmann & Marder, Nature Methods, 2013)
  - Much groundwork has been laid by you and others, some of which we will learn more about through today's presentations
- And, we need not understand everything about the brain to be successful in this endeavor

Ingredients for Success

- People
  - Experimental neuroscientists: High-resolution structural and functional imaging
  - Computer scientists: Image processing, data management
  - Applied mathematicians: Analyzing brain “graphs” and identifying motifs
  - Computational neuroscientists: Modeling neurons and neural circuits
  - Machine learners: Constructing algorithms from cortical computing primitives
  - Systems engineers: Systems engineering
- Ideas: A framework for interpreting and exploiting cortical microcircuits
  - Principles for selecting which data to collect and analyze
  - Mathematics for describing the operation of cortical microcircuits
  - Insights for extracting meaningful primitives from observed data about microcircuits
  - Strategies for employing cortical computing primitives for machine learning
- Money
- Direction
  - Integrated teams
  - Application-driven science

Desired Results

- Most interested in novel machine learning algorithms derived from cortical microcircuits that meet the following criteria:
  - Perform ***spatiotemporal*** pattern recognition, classification, anomaly detection, and/or other similar tasks
  - Operate on ***streaming*** data "of interest" to the intelligence community (e.g., video, audio, multimodal, hyperspectral, financial, etc.)
  - Exhibit performance characteristics superior to the conventional state of the art in terms of:
    - Accuracy given limited training data
    - Robustness to novel (untrained) sources of noise
    - Generalization capability
- Less interested in:
  - Algorithms for motor control or robotics (excluding perception-action loops)
  - Algorithms that achieve the same performance as today's state of the art at lower size, power, and/or cost
  - Hardware (for hardware's sake)

Impact

- Short term:
  - Significantly advance our capabilities in automated pattern recognition, classification, anomaly detection, etc. with algorithms that exhibit human-like performance characteristics superior to the conventional state of the art
- Long term:
  - Establish and demonstrate feasibility of a "pipeline" for innovation through integration of neuroscience and machine learning
  - Revolutionize machine intelligence and neural computing technologies

Questions?