Accelerating Returns in Neuroscience in 2^{2^2} Tweets

Thomas Dean
Google Research

Kavli Futures Symposium: Toward Next-gen, Open-source Neurotechnology Dissemination, October 2017

0. *Move beyond linear thinking. Anticipate exponential scaling. Accelerate science by opportunistically exploiting existing technology. Avoid banal tweets and gothic fonts.*

1. *We started in 2013 on structural connectomics, developed new infrastructure and software for petabyte datasets and petaflop computations, highly leveraged machine learning.*

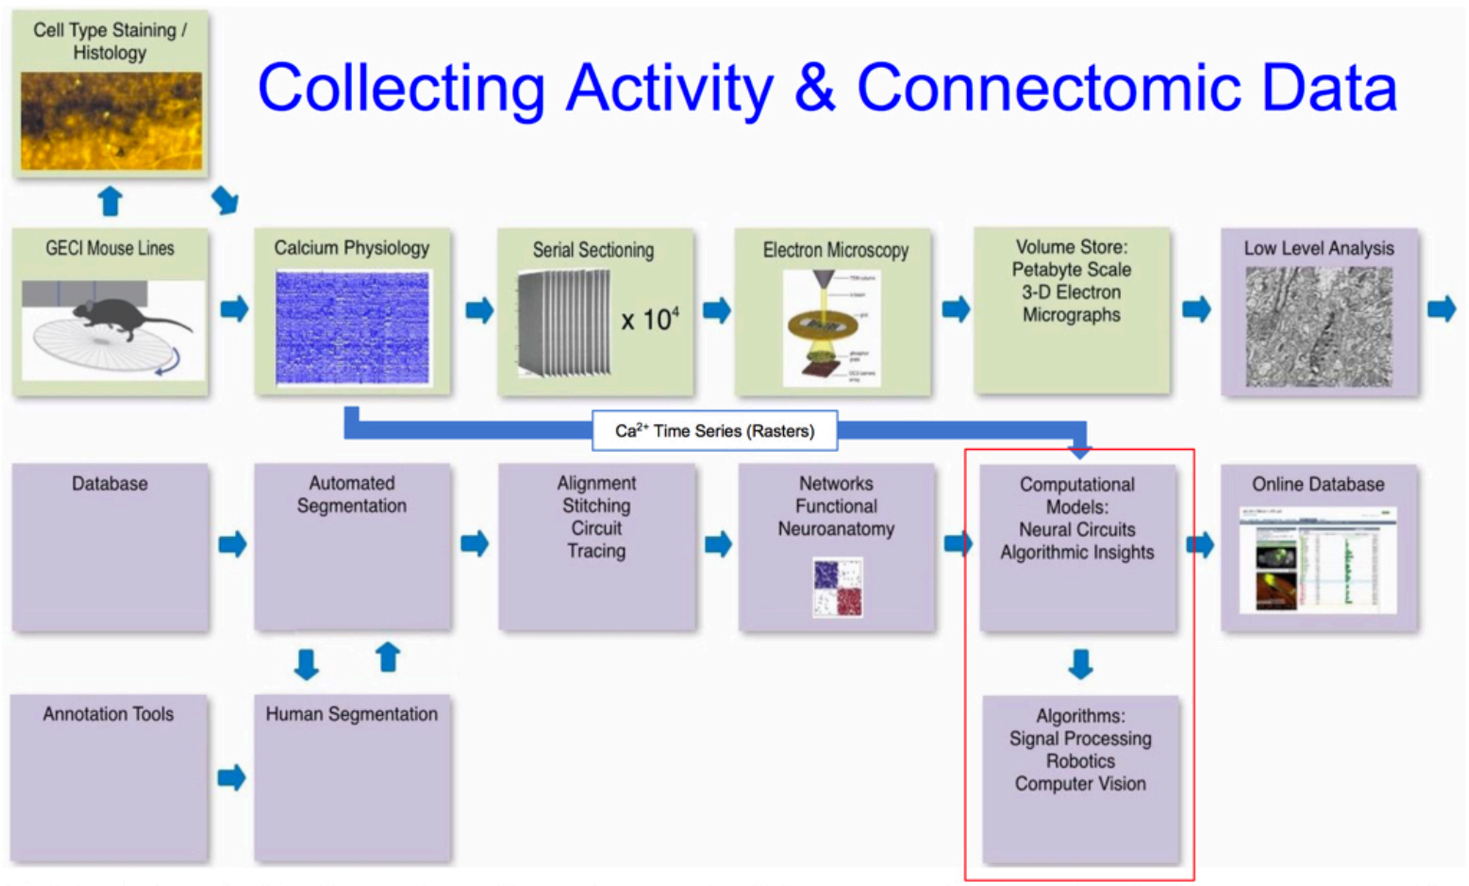

Christof Koch. Project MindScope. In *Frontiers in Computational Neuroscience, 2012 Bernstein Conference*, number 33. 2012.

What we concluded was feasible in 2013:
https://arxiv.org/abs/1307.7302
Technology Prospects for Scalable Neuroscience

Automated reconstruction of dataset obtained by serial block-face electron microscopy from a male zebra finch brain, achieving a mean error-free neurite path length of 1.1 mm, an order of magnitude better than previously published approaches applied to the same dataset. *BioRxiv*, October 2017.

0. Move beyond linear thinking. Anticipate exponential scaling. Accelerate science by opportunistically exploiting technology. Avoid pithy tweets and gothic fonts.

1. We started in 2014 on structural connectomics, developed new infrastructure and software for petabyte datasets and petaflop computations, highly leveraged machine learning.

2. Suppose we didn't have the WWW and Google Search but did have hundreds of data centers with millions of cores, super-fast fiber networking and exabytes of fast storage?

3. Invent the WWW, index all of the world's knowledge and build powerful search engines would be the first order of business. Information sharing rules. Hindsight is 20:20.

4.

Access to data is what is holding us up now. Where can we look for new breakthroughs to maintain the law of accelerating exponential returns for neural recording?

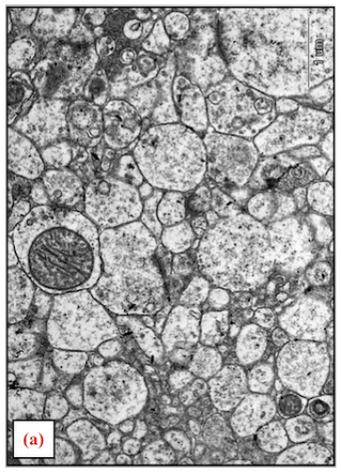

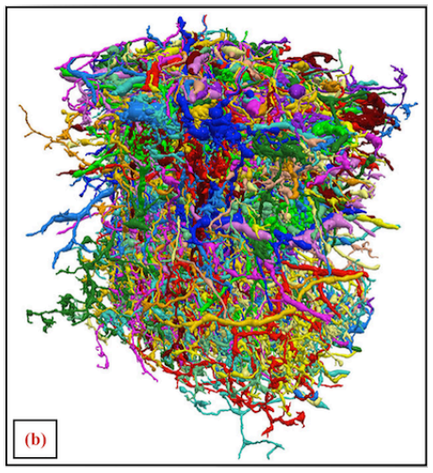

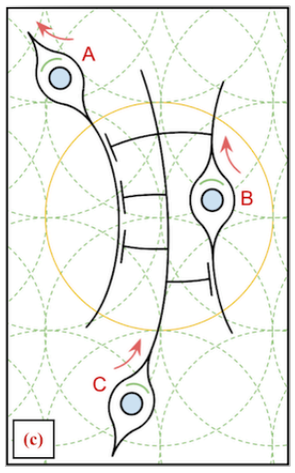

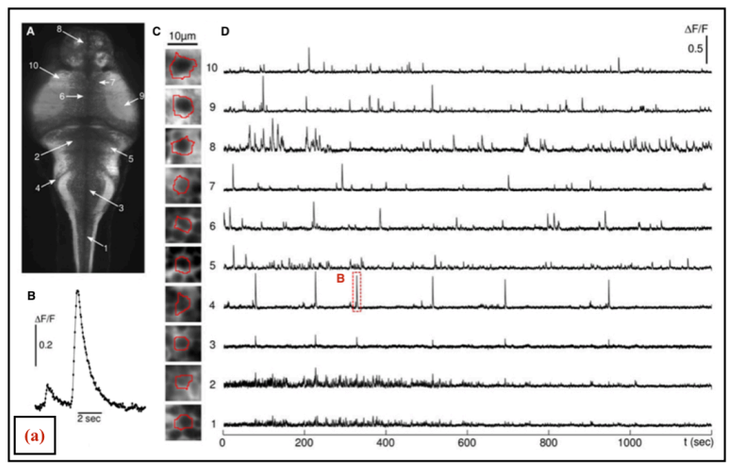

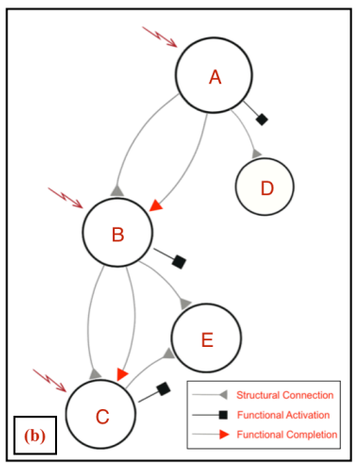

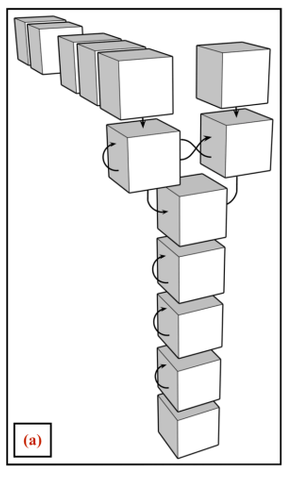

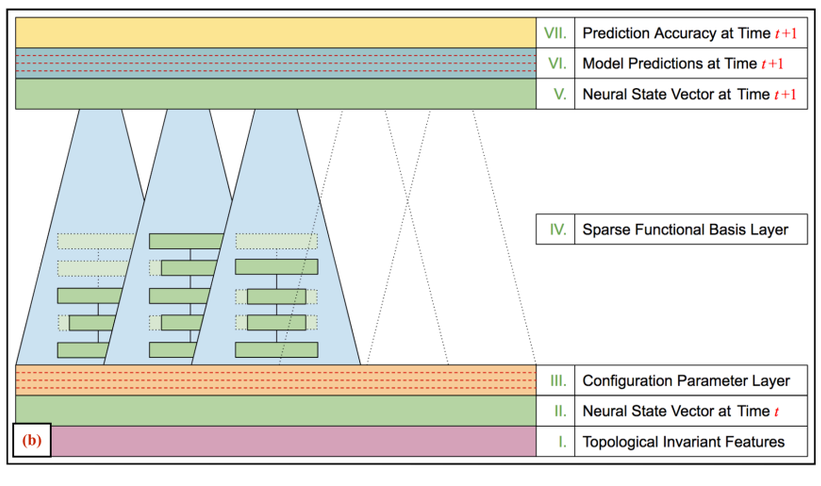

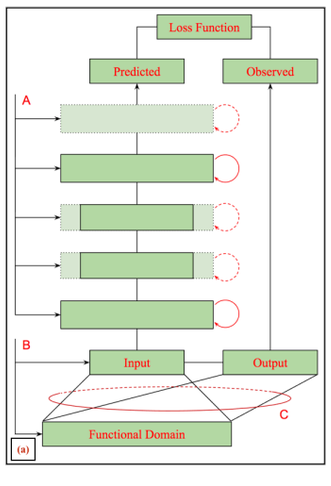

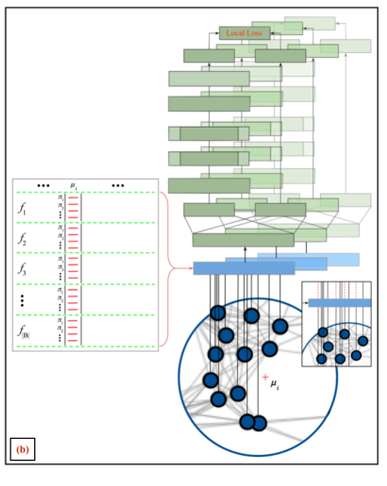

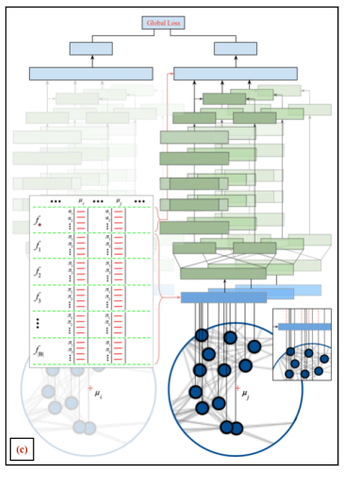

What we concluded will be feasible in 2019:
http://arxiv.org/abs/1710.05183
Inferring Mesoscale Models of Neural Computation

We argue that the geometry of neural circuits is essential in explaining the computational strategies inherent in biological information processing. We propose a blueprint for how to employ tools from machine learning to automatically infer a mesoscale model of neural computation by combining functional and structural data, with an emphasis on learning and exploiting regularities and repeating motifs in neural circuits. *arXiv*, October 2017.

4. *Access to data is what's holding us up now. Where can we look for new breakthroughs to maintain the law of accelerating exponential returns for neural recording?*

5. *In a word, artificial intelligence: exponential scaling that feeds on itself. Hypothetically, what would you do if you paid the bills and were faced with the following tradeoff?*

6. *Work on brains for ten years or work on AI for two years and then work on brains and learn at least as much about brains in half of the time. Without data we slow to a crawl.*

7. *Tens of thousands of engineers are working on AI. They are driven by huge economic incentives, inspired by high expectations and convinced of the benefits of sharing ideas.*

0. If you were on top of deep learning two years ago and haven't been keeping track, you will need to run to catch up with your peers and some of your peers will be AI.

1. The field is moving very fast. Recurrent and convolutional nets are so yesterday's news, they are being supplanted by next-generation attentional networks. So what else is new?

2. What about consciousness, attention, analogical reasoning, one-shot learning, sophisticated continuous dialog with empathy and subtle theory-of-mind reasoning? Assistants.

3. It doesn't take much of a hint from nature to inspire an engineer. Build a conscious agent? Check out Stanislas Dehaene 2014, Michael Graziano 2015, Yoshua Bengio 2017.

4. Automatic programming? Start with code search. Index billions of lines of well-engineered code. Emulate the art of inspired imitation. Pattern memory + creative plagiarism.

5. Create a new neural network architecture? Build systems that employ reinforcement learning to search parameter space and design novel neural network architectures.

6. What was the most important advance in neural networks of the last decade? Scalable infrastructure to expedite end-to-end training. Engineers hate to wait days for results.

7. What will be the most important accelerator for artificial intelligence in the coming decade? Exponentially faster infrastructure for searching NN space. Foresight 20:20.

Supplementary Slides

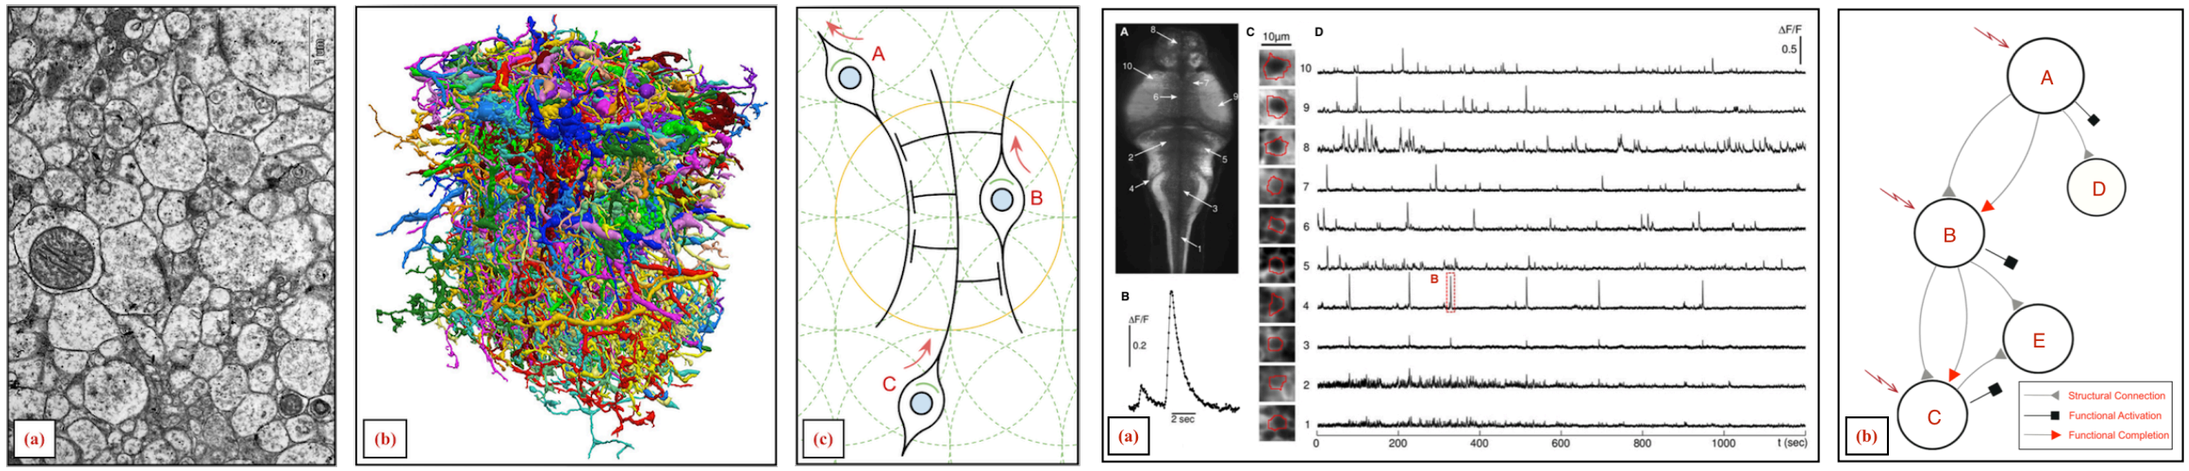

We argue that the geometry of neural circuits is essential in explaining the computational strategies inherent in biological information processing. [http://arxiv.org/abs/1710.05183]

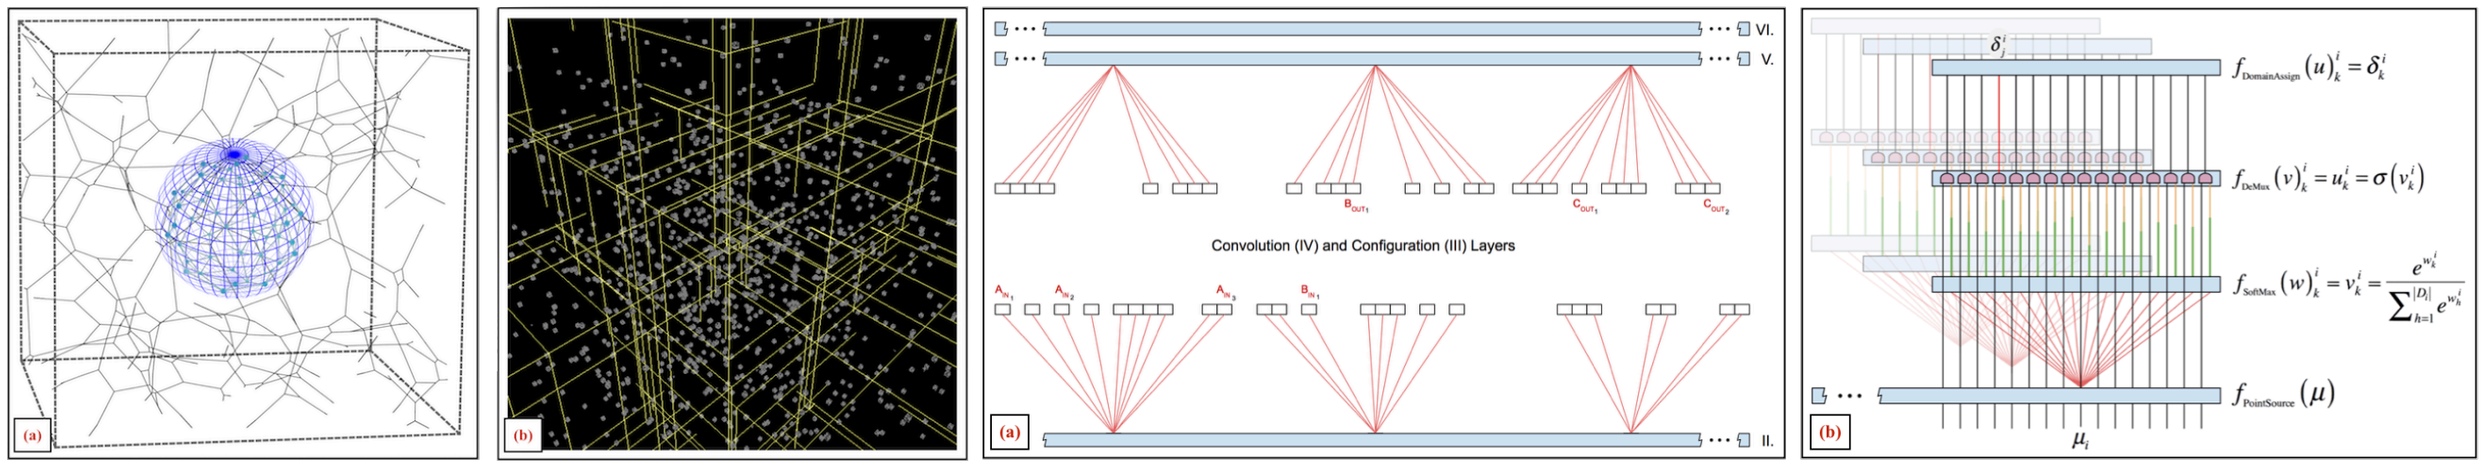

We describe how to employ tools from machine learning to automatically infer a mesoscale model of neural computation by combining functional and structural data, with an emphasis on identifying and exploiting regularities and repeating motifs in large-scale neural circuits.