Expansion Circuits

Brain Tissue as a Computational Substrate

BrainMaker[†]

Thanks to the many collaborators, consultants and supporters who helped:

- Ed Boyden, Jae-Byum Chang, Manos Karagiannis, Adam Marblestone (MIT)
- Winfried Denk, Julia Buhmann (Max Planck Institute, University of Heidelberg)
- Mark Ellisman[‡] (UCSD, National Center for Microscopy and Imaging Research)
- Dan Fletcher, Eric Jonas, Dongjin Seo (University of California Berkeley)
- Viren Jain, Peter Li, Art Pope, Jon Shlens, Rahul Sukthankar (Google Research)
- Dyche Mullins, Robert Burton (University of California San Francisco)
- Clay Reid, Stephen Smith, Costas Anastassiou (Allen Institute for Brain Science)
- Davi Bock, Tamir Gonen, Gerry Rubin, Karel Svoboda (HHMI Janelia Farm)
- Wei-Chung Lee, Xiaowei Zhuang, Guisheng Zhong (Harvard University)

[†] Codename chosen for its allusion to rainmaker and the maker movement. In business speak, a rainmaker is a person or product that generates substantial new business or additional cash flow from sources that are outside the established business channels.

[‡] Mark Ellisman was particularly generous with his time, sharing relevant data, and considerable expertise in electron microscopy.

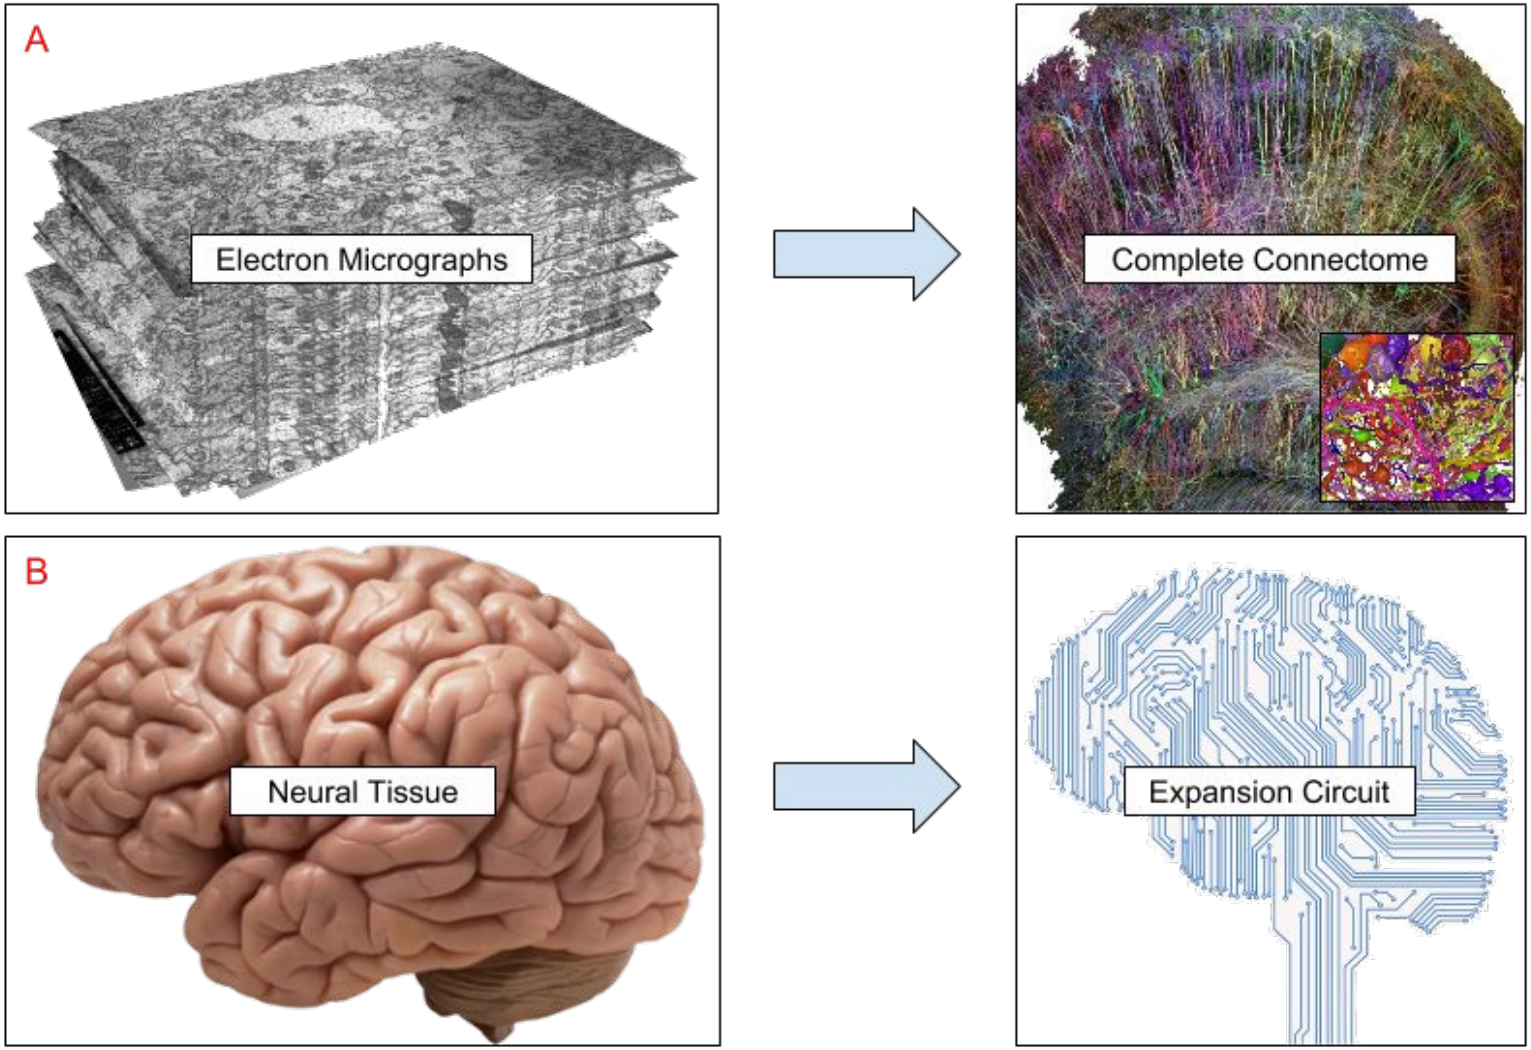
A
Electron Micrographs
Complete Connectome
B
Neural Tissue
Expansion Circuit

Macroscopic

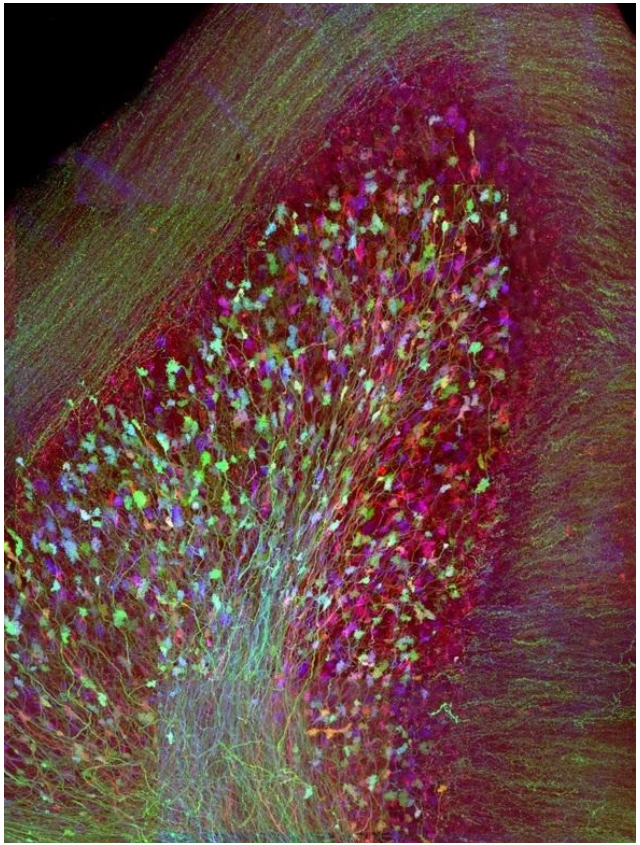

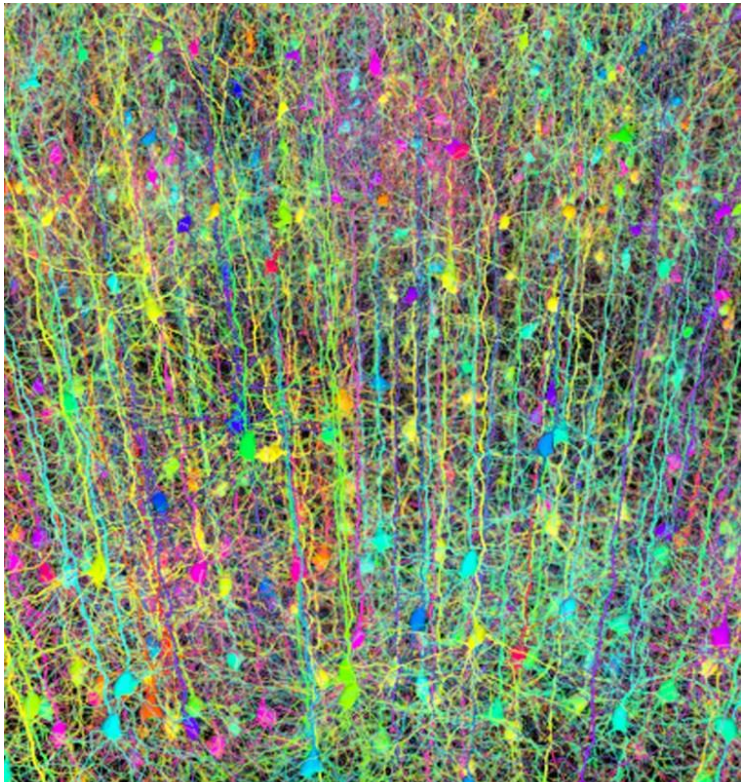

Microscopic

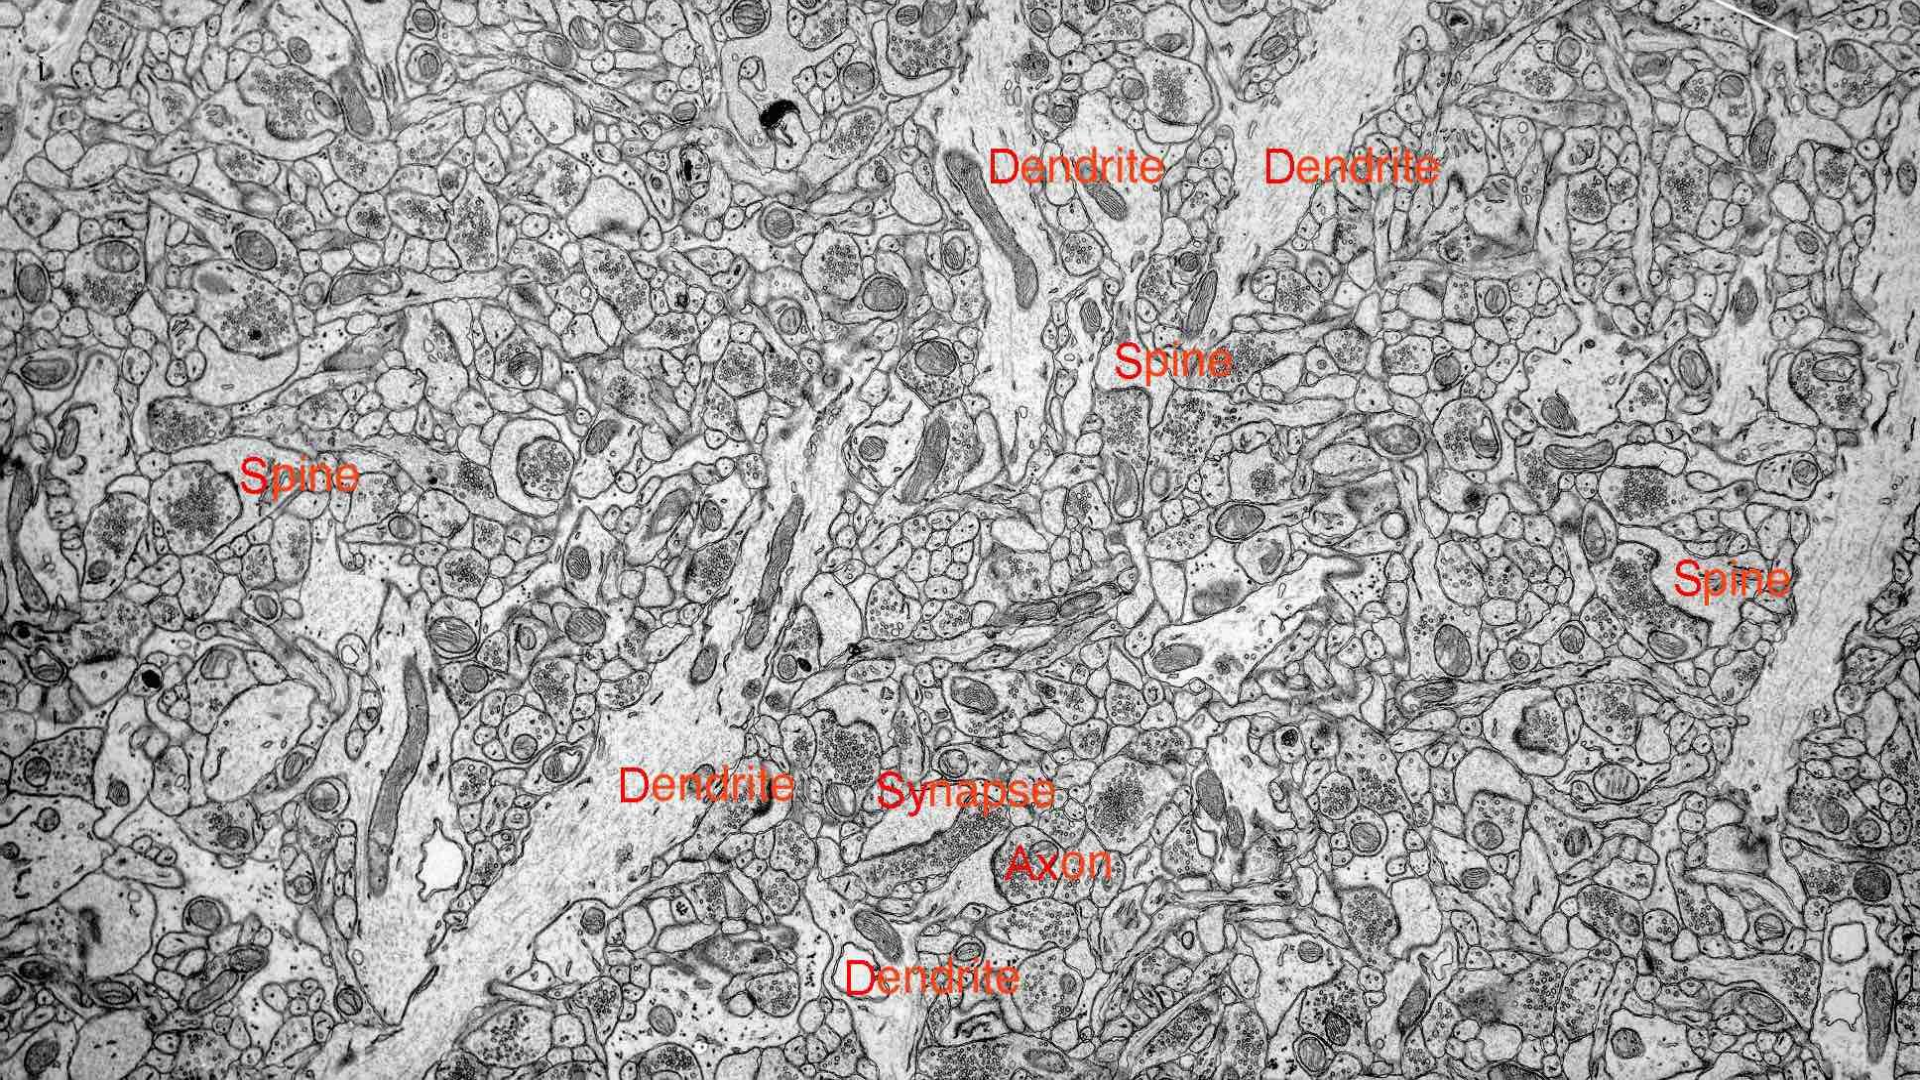
Dendrite
Dendrite
Spine
Spine
Spine
Dendrite
Synapse
Axon
Dendrite

Cytosol [Intracellular]

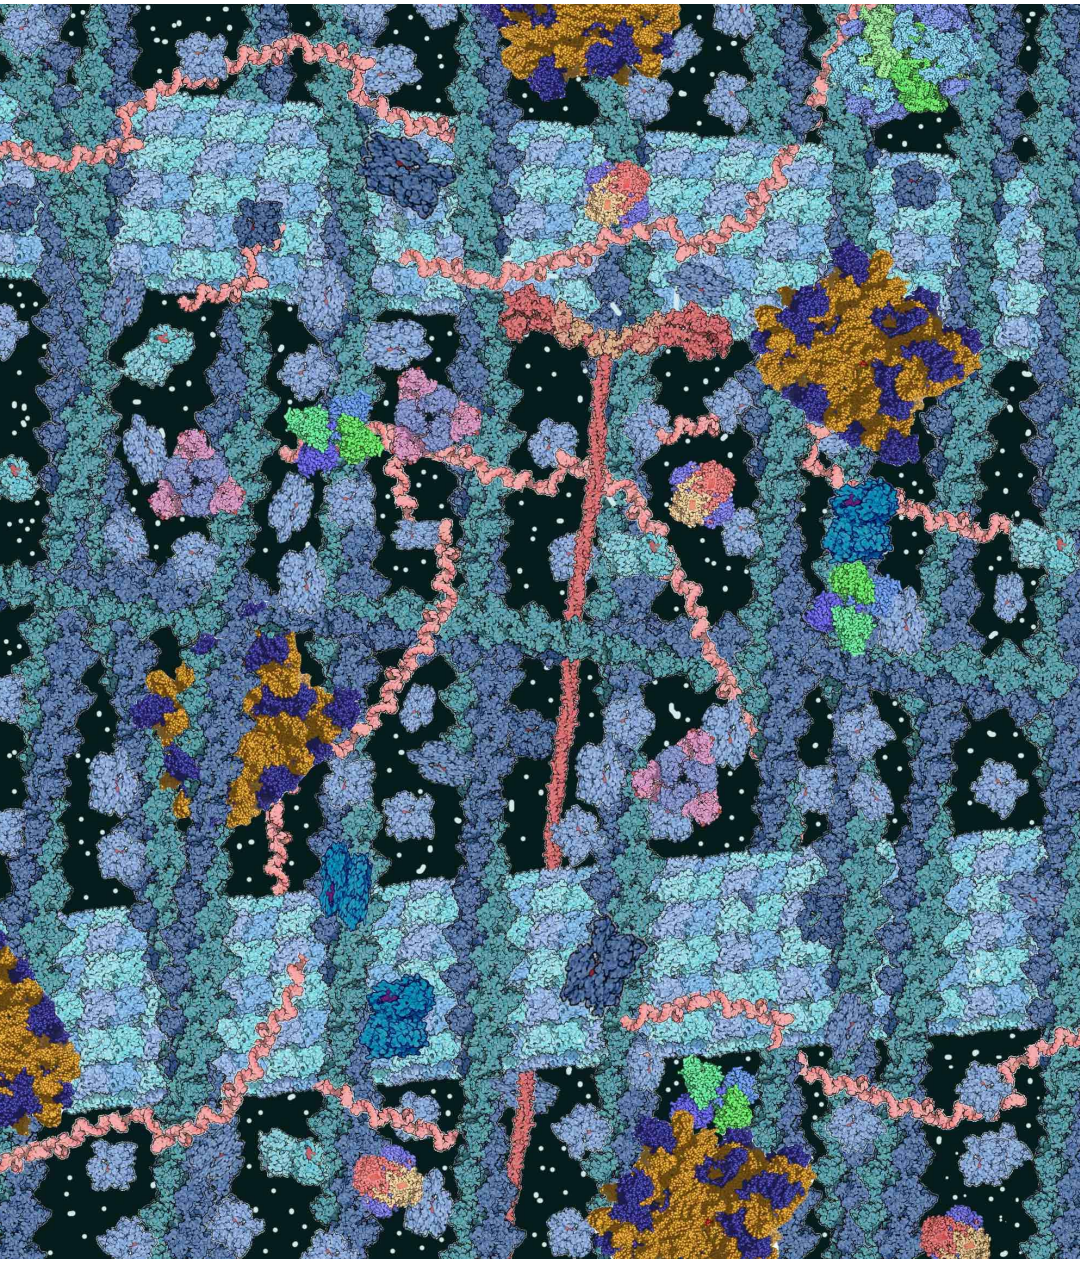

Dense Lipid Tags
[Conjugation]

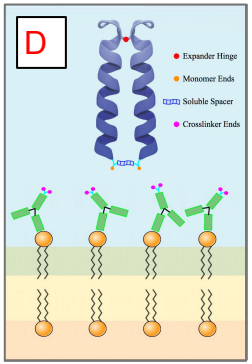

Synthesize Mesh[2]
[Polymerization]

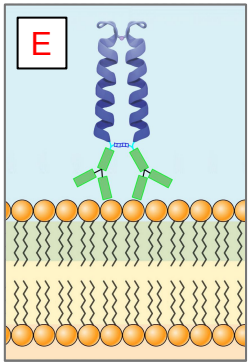

Dissolve Spacers
[Selective Digestion]

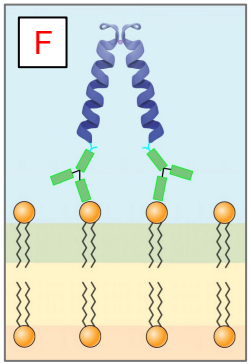

Functionalize Mesh
[Metallization]

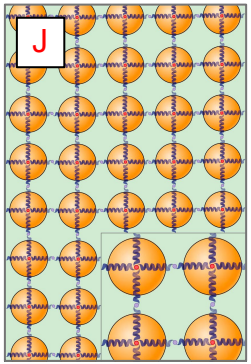

Install Electronics
[Circuit Fabrication]

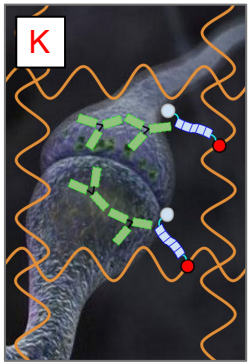

Integrate Power Grid
[Power Distribution]

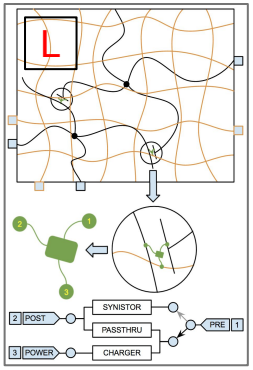

Prepare Tissue
[Permeabilization]

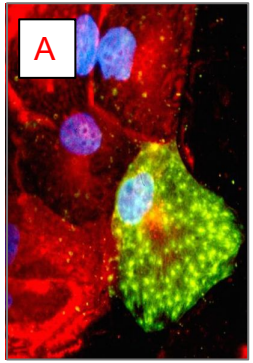

Fluorescent Tags
[Conjugation]

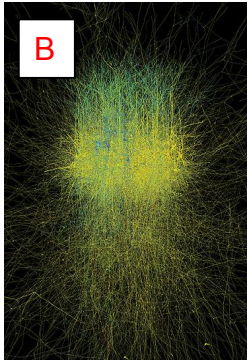

Synthesize Lattice[1]
[Polymerization]

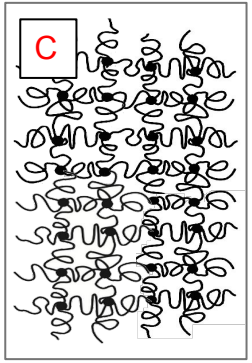

Anchor Tags
[Lattice Stabilization]

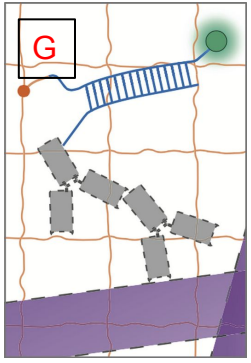

Protein Hydrolysis
[Homogenization]

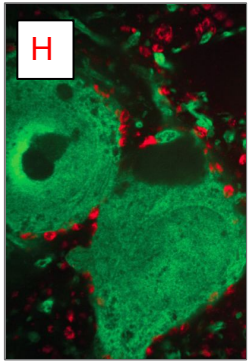

Aqueous Dialysis
[Polymer Expansion]

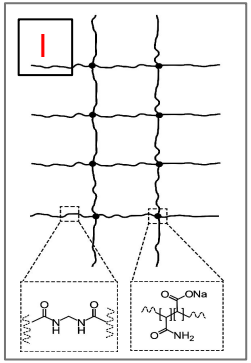

[1] Lattice = expansion network (superabsorbent) polymer.

[2] Mesh = membrane covering (conductive) polymer.

Expansion Microscopy

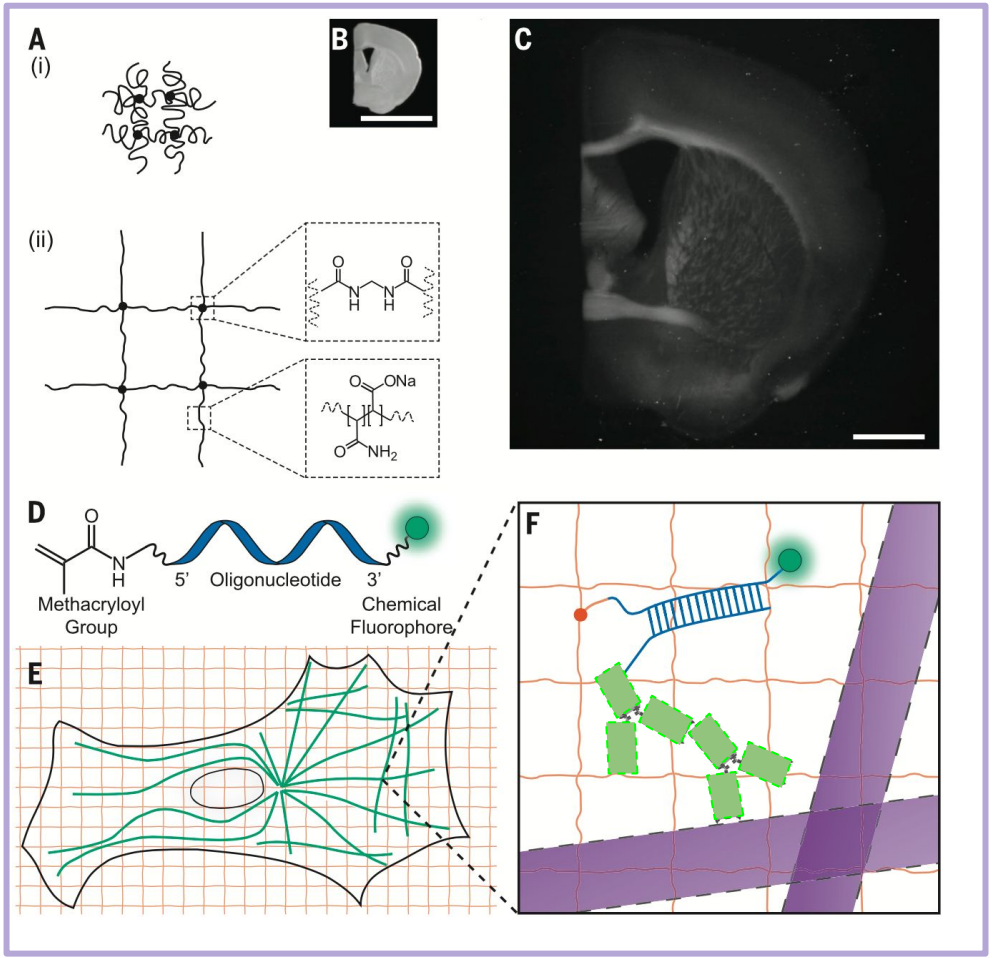

Expansion Microscopy Protocol

A. Halt reactions, enhance structure, permeabilize cell membranes
B. Tag molecules by selectively binding molecular attachment sites
C. Diffuse monomer[1] solute, induce polymerization and crosslinking

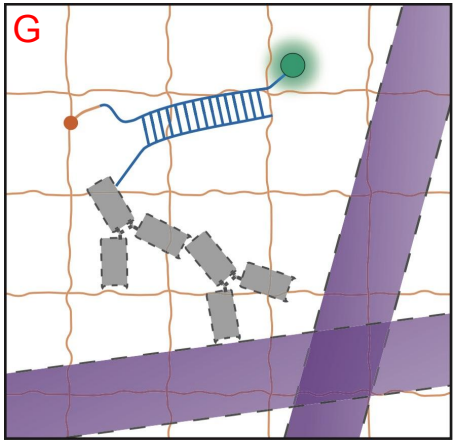

G. Anchor the tag molecules and co-register with the polymer lattice
H. Extract lipids, hydrolyze proteins reducing simple peptide chains
I. Infuse with inert aqueous solution, expand, repeat until plateaus

[1] Here we use a polyacrylamide superabsorbent polymer or *hydrogel* as the expandable basis for a polymer-mesh lattice anchoring all tagged molecules.

Expansion Circuitry Protocol

D. Tag the phospholipid polar head groups with four-way crosslinker
E. Diffuse monomer[2] solute, induce polymerization and crosslinking

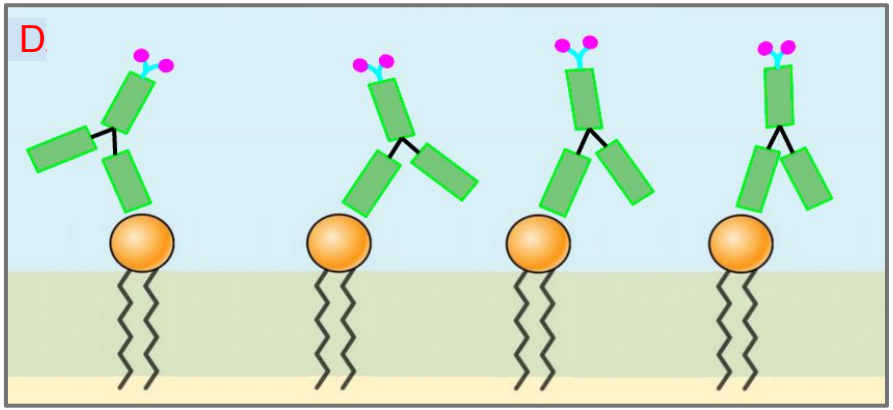

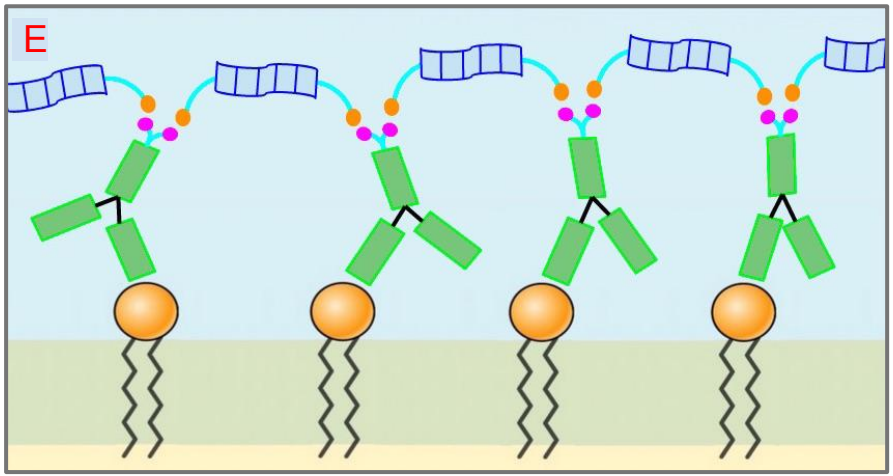

[2] In this case, we could use either a conducting polymer like poly(3,4-ethylenedioxythiophene) (PEDOT) or a suitable non-conducting polymer and later metallize it.

Expansion Circuitry Protocol

D. Tag phospholipid polar head groups with a four-way crosslinker
E. Diffuse monomer[2] solute, induce polymerization and crosslinking

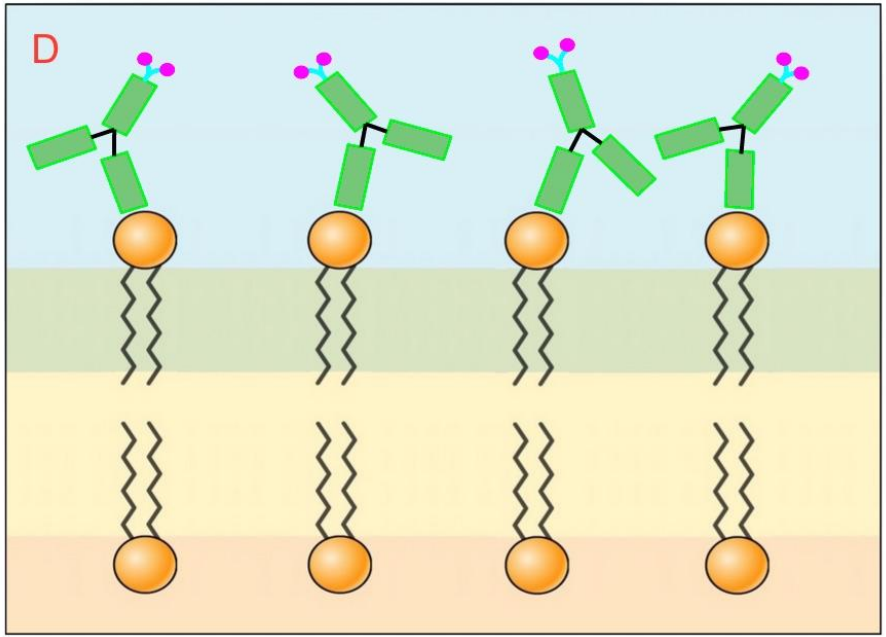

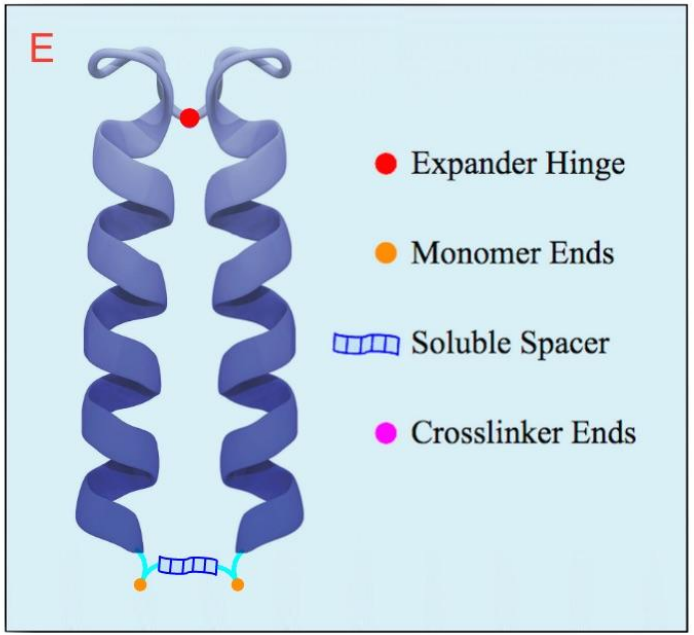

[2] In this case, we could use either a conducting polymer like poly(3,4-ethylenedioxythiophene) (PEDOT) or a suitable non-conducting polymer and later metallize it.

Expansion Circuitry Protocol

D. Tag phospholipid polar head groups with a four-way crosslinker
E. Diffuse monomer[2] solute, induce polymerization and crosslinking

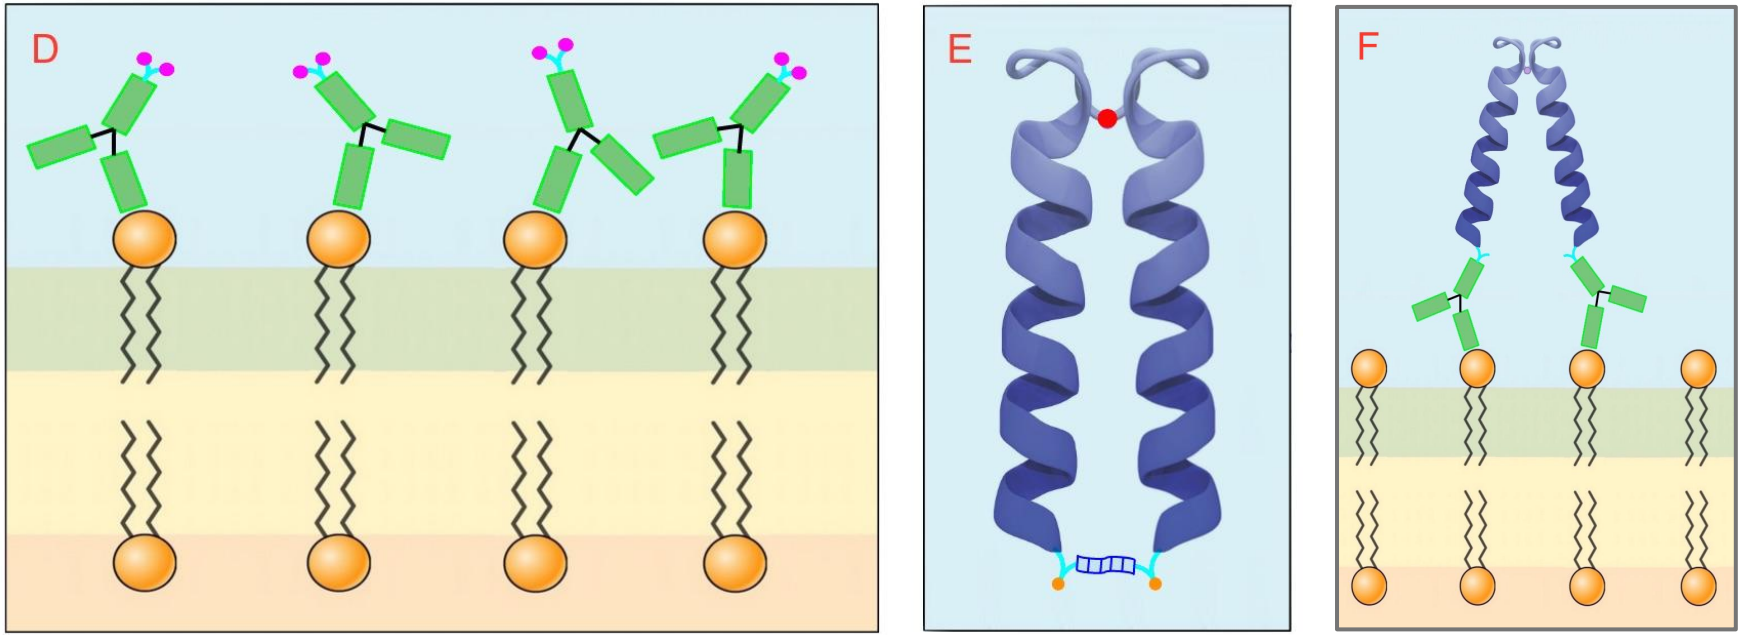

F. Remove spacers with appropriate solvent or cleavage enzyme

[2] In this case, we could use either a conducting polymer like poly(3,4-ethylenedioxythiophene) (PEDOT) or a suitable non-conducting polymer and later metalize it.

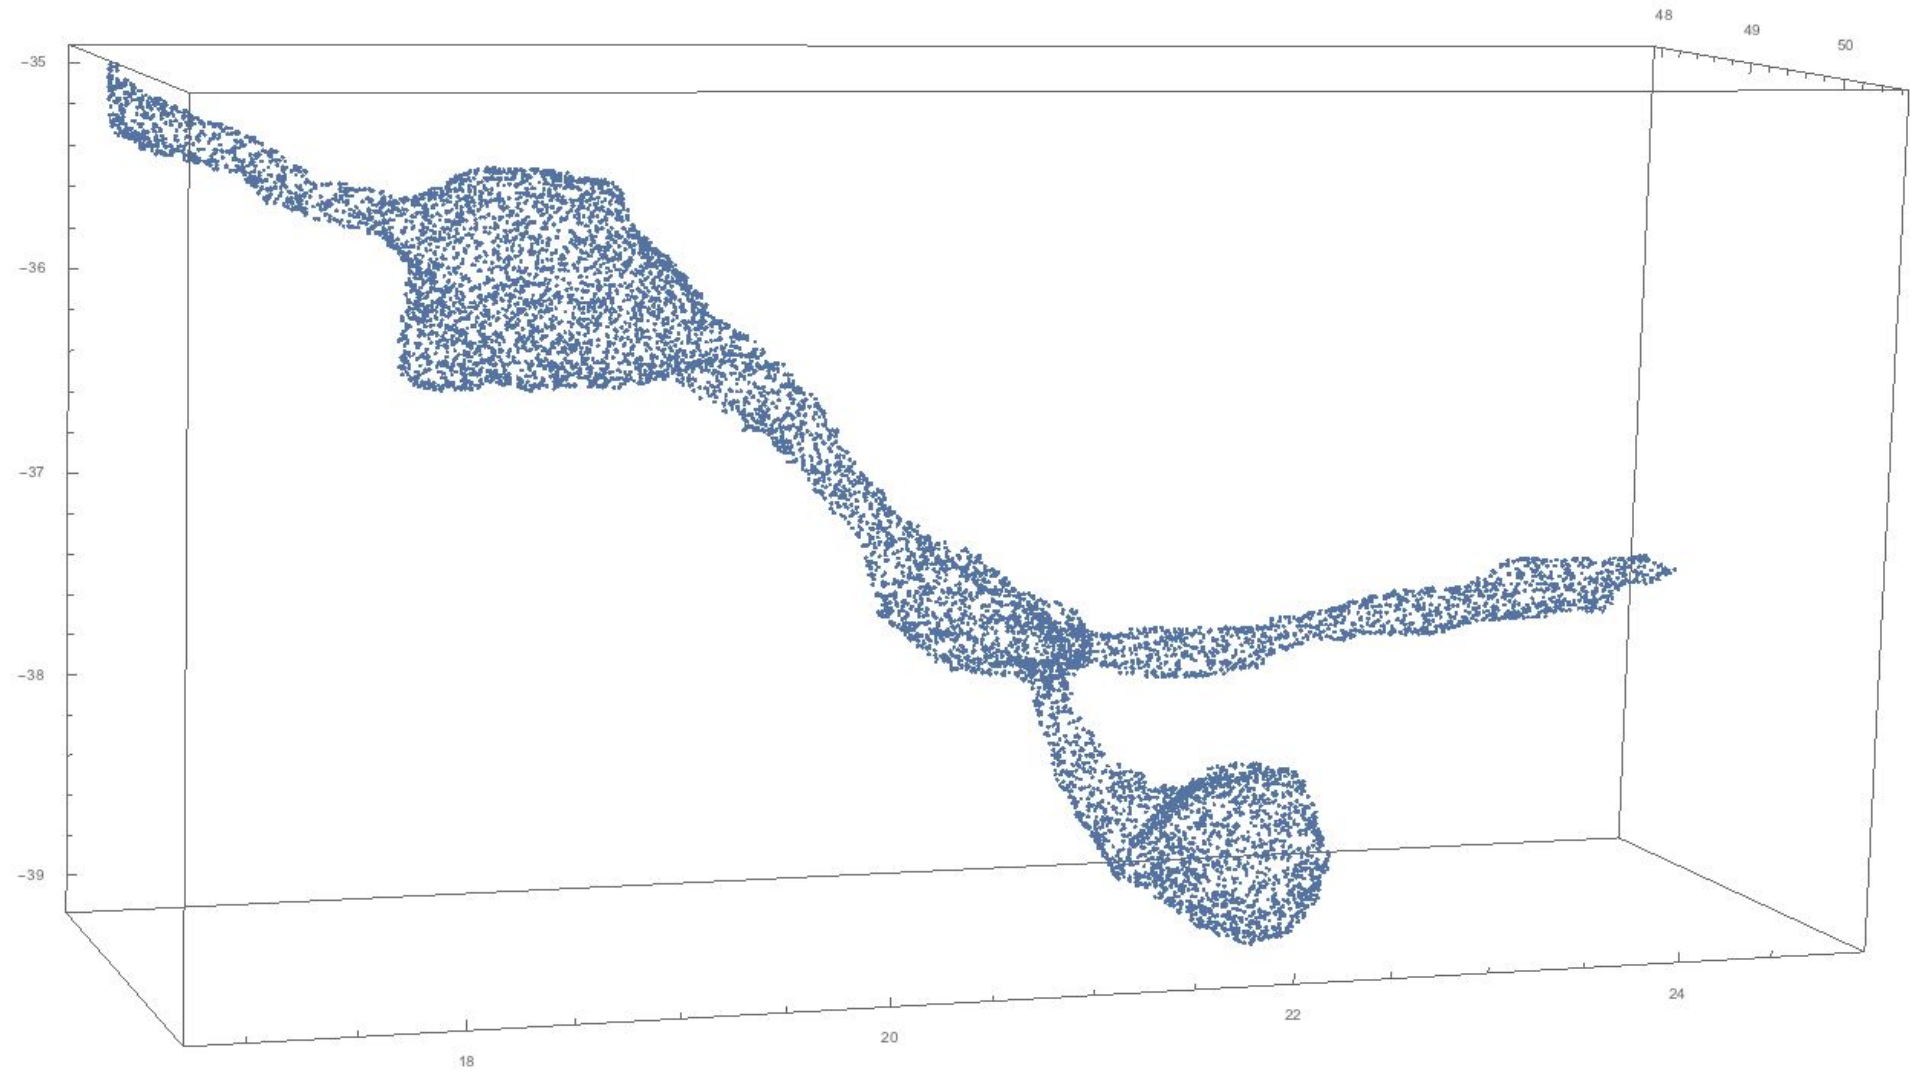
48
49
50
−35
−36
−37
−38
−39
18
20
22
24

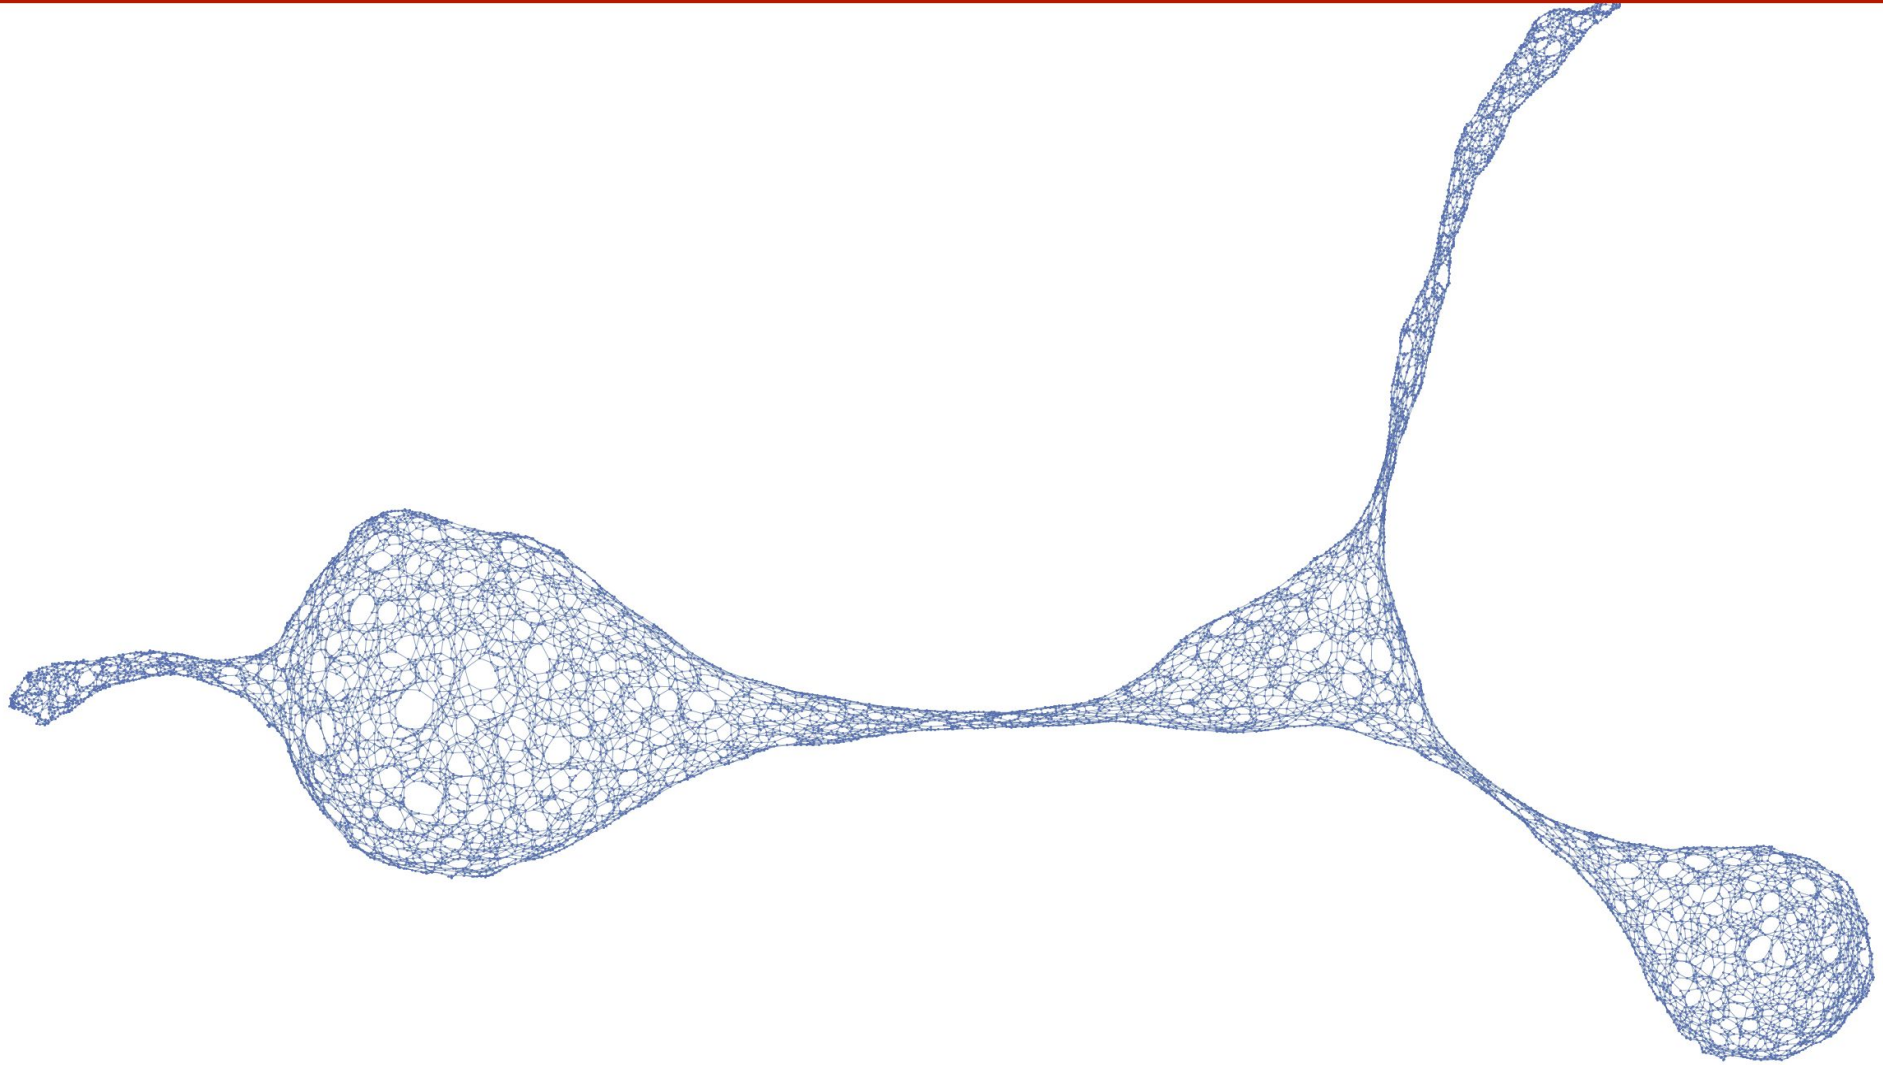

Problem with Lipids

How many [antibodies] can dance on the [polar head group] of a [lipid]?

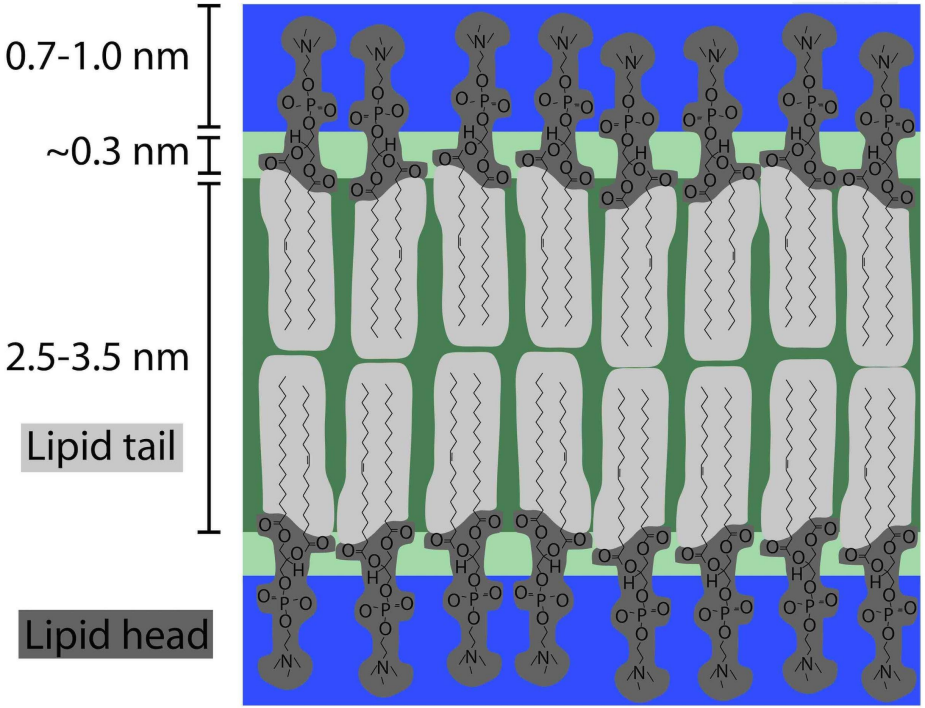

Cytoskeletal Proteins

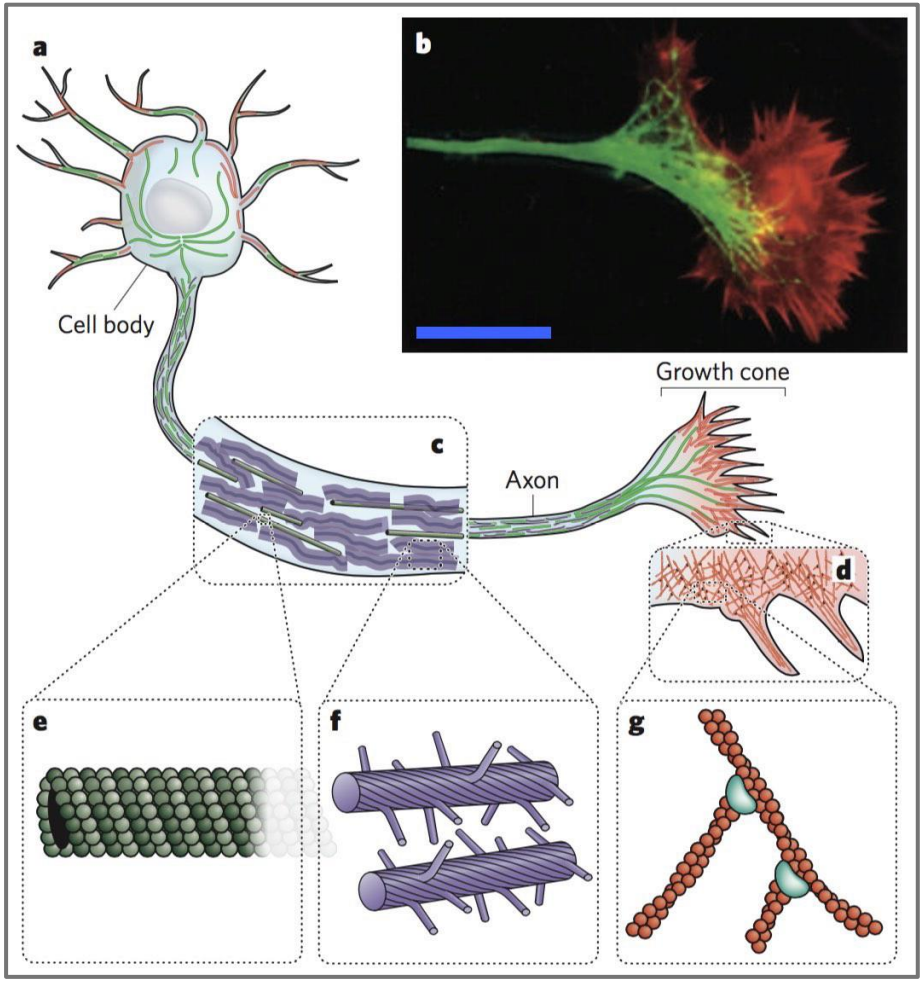

Cytoskeletal Proteins

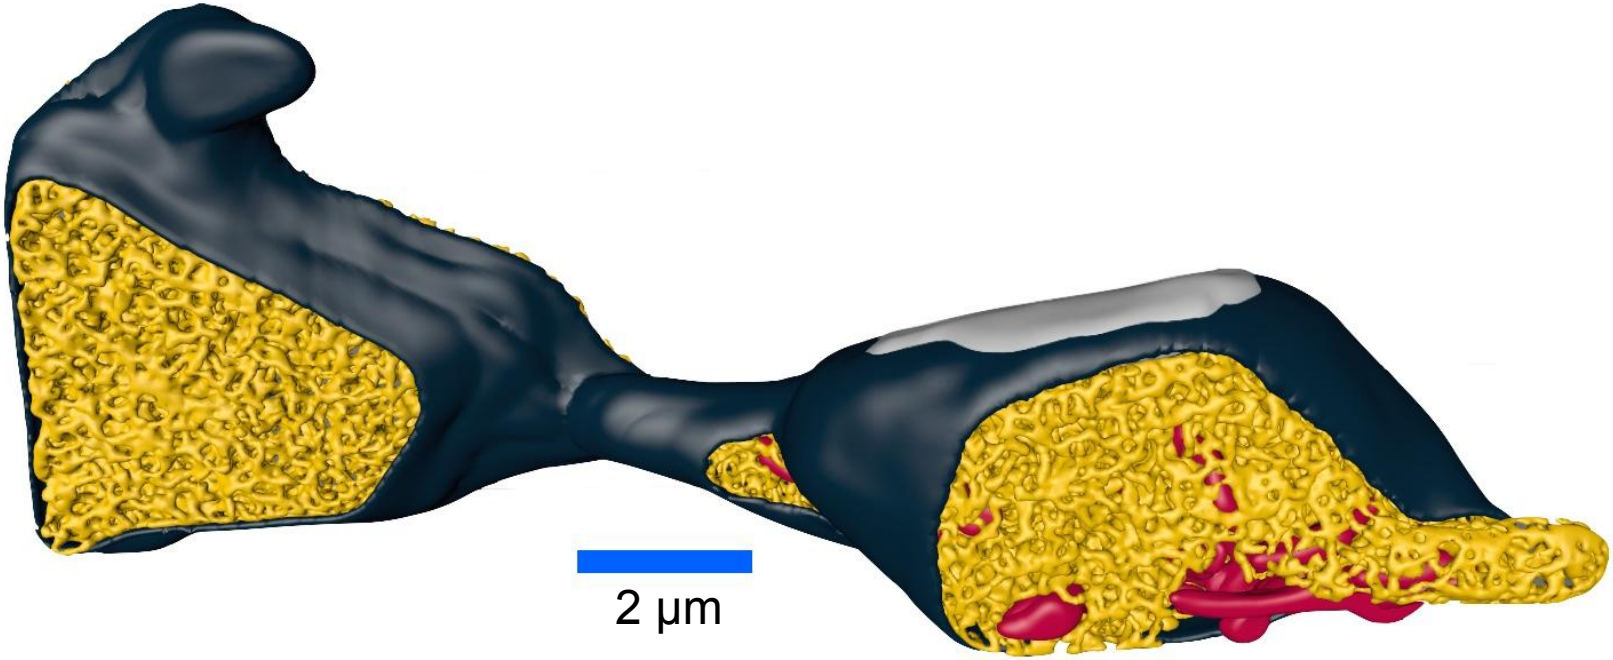

Expansion Circuitry Protocol

J. If polymer mesh coating is not conductive, metallize surface
K. Install active electronic components:
   - ➢ Tag pre- and post-synaptic neurons with connectors
   - ➢ Attach low-power memristor[3] devices to connectors

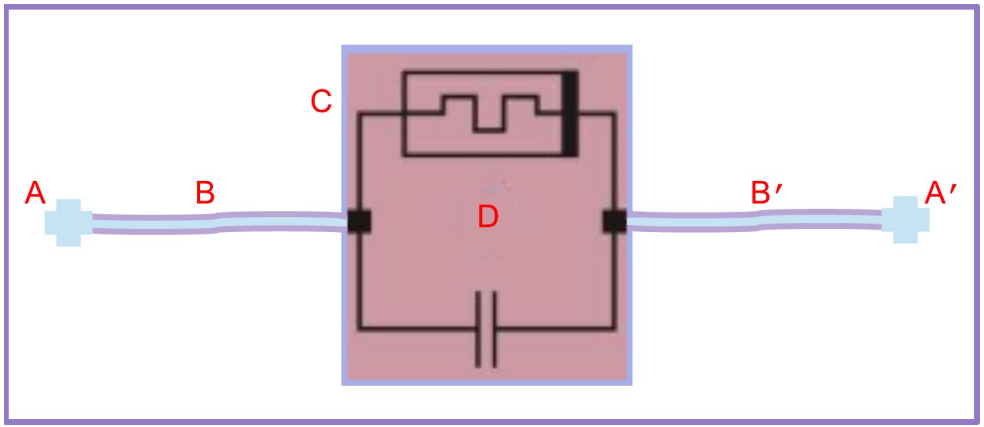

L. Install power distribution network:
   - ➢ Third-rail, single multiplexed or two separate meshes
   - ➢ Ultrasound, piezoelectric harvester, capacitor storage

[3] A memristor regulates the flow of electrical current in a circuit and remembers the amount of charge that has previously flowed through it.

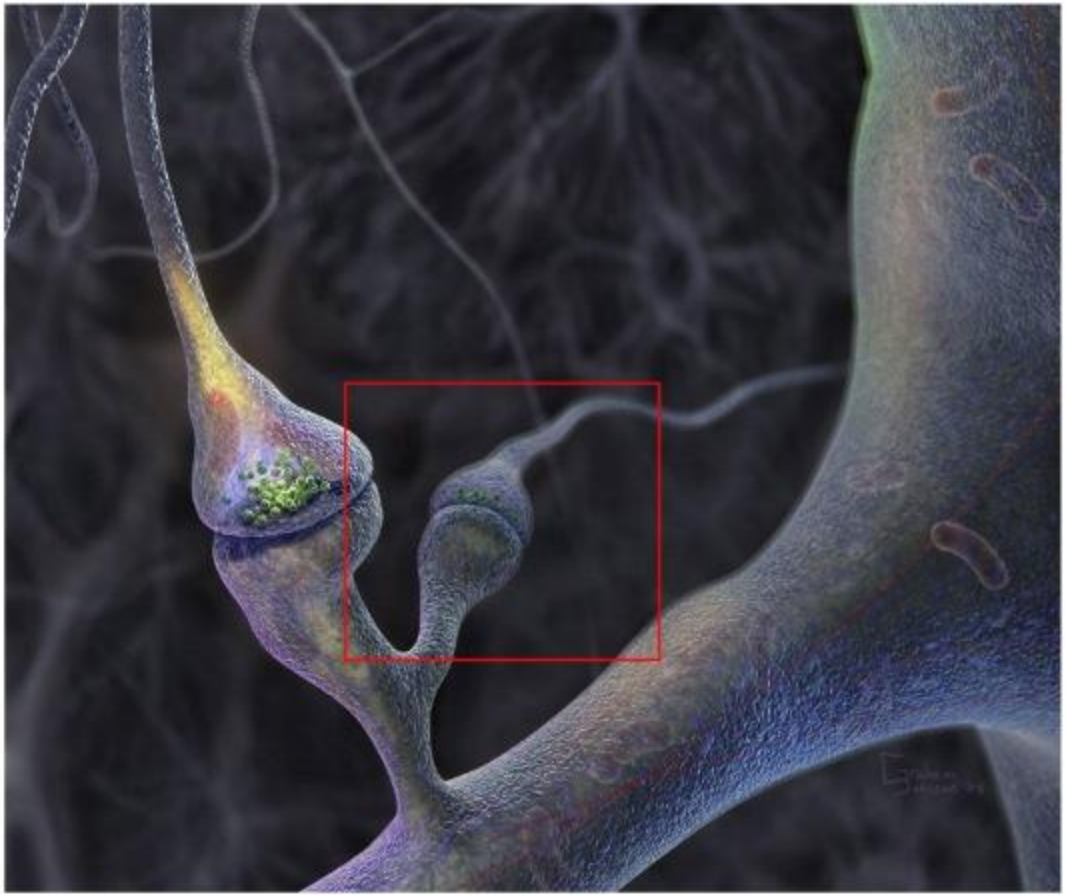

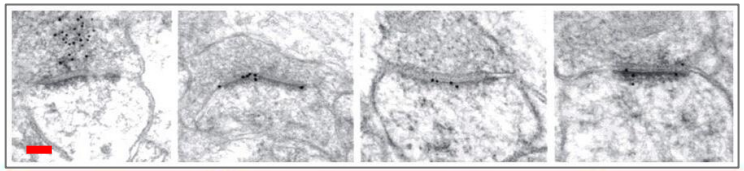

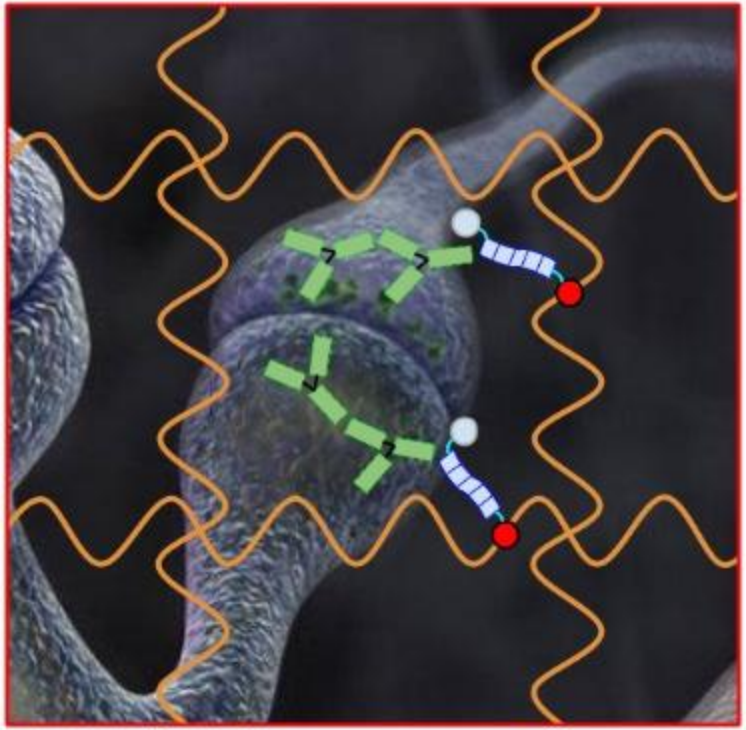

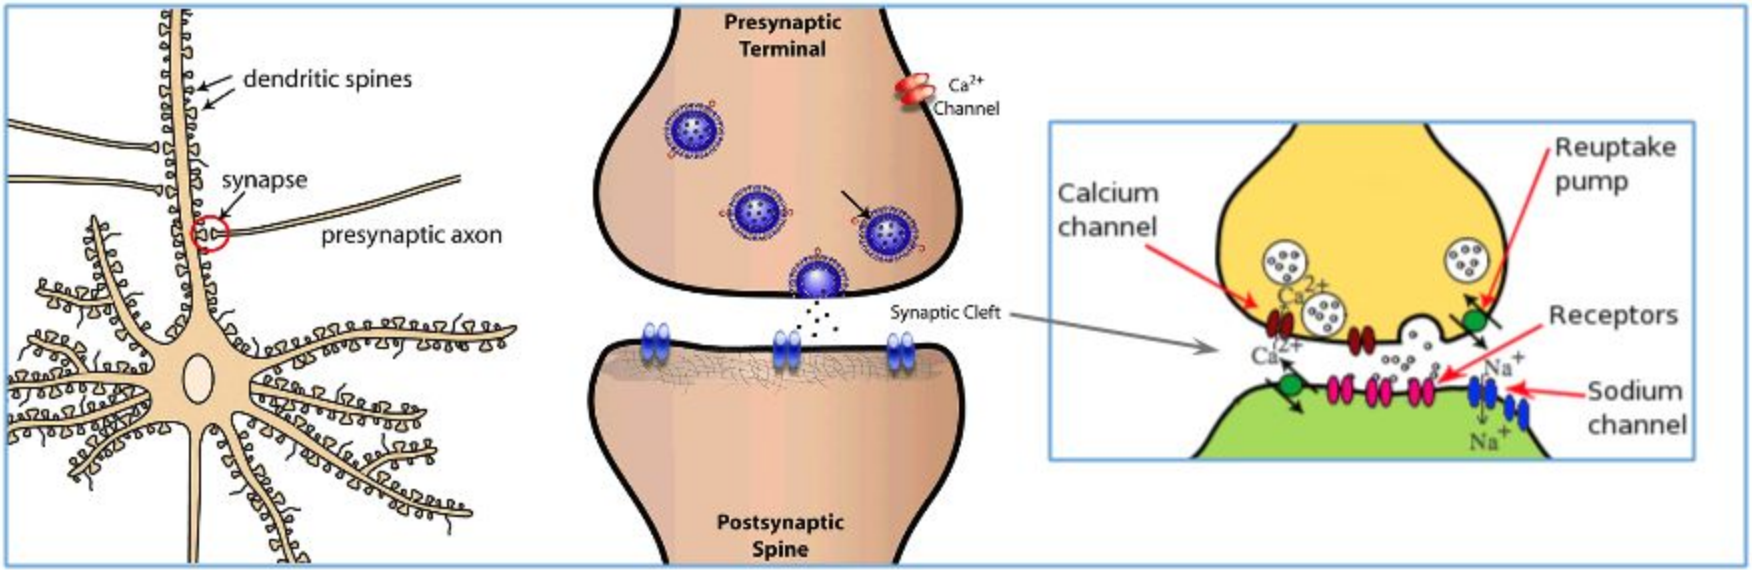
dendritic spines
synapse
presynaptic axon
Presynaptic Terminal
Ca^{2+} Channel
Synaptic Cleft
Postsynaptic Spine
Calcium channel
Reuptake pump
Receptors
Sodium channel
Ca^{2+}
Na^+

Circuit Fabrication

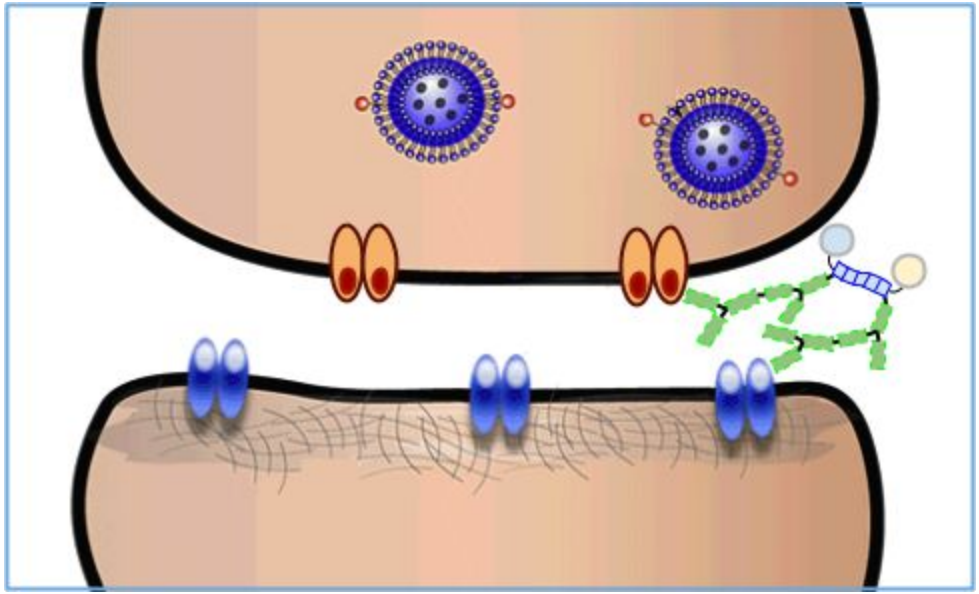

Expansion Circuitry Protocol

J. If the polymer mesh coating is not conductive, metallize surface
K. Install active electronic components:
    1. Identify protein classes to differentiate inputs and outputs
    2. Tag pre- and post-synaptic neurons with I/O connectors
    3. Fabricate synthetic synapses with antibody-keyed leads
    4. Expand tissue to facilitate diffusion of synthetic synapses
    5. Conjugate synthetic-synapse leads with I/O connectors

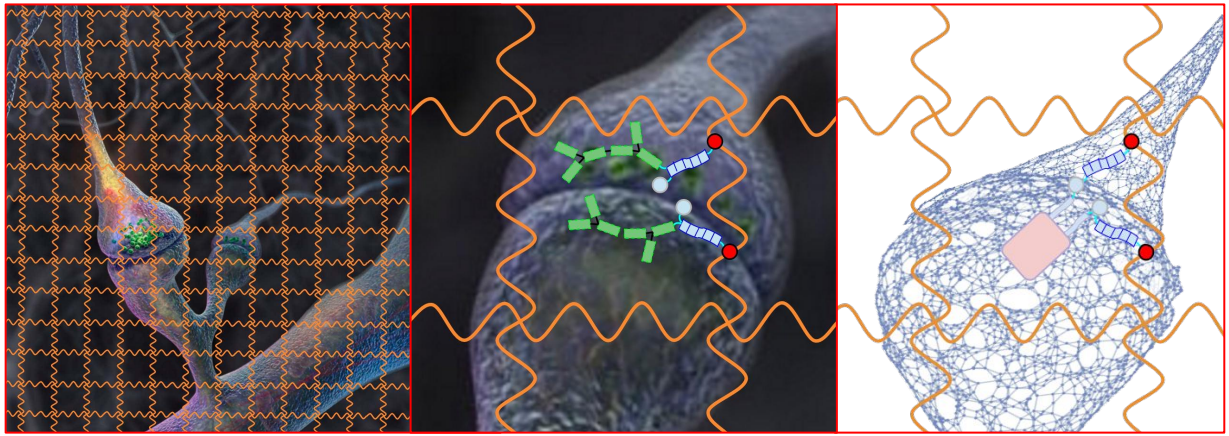

Equivalent Circuit

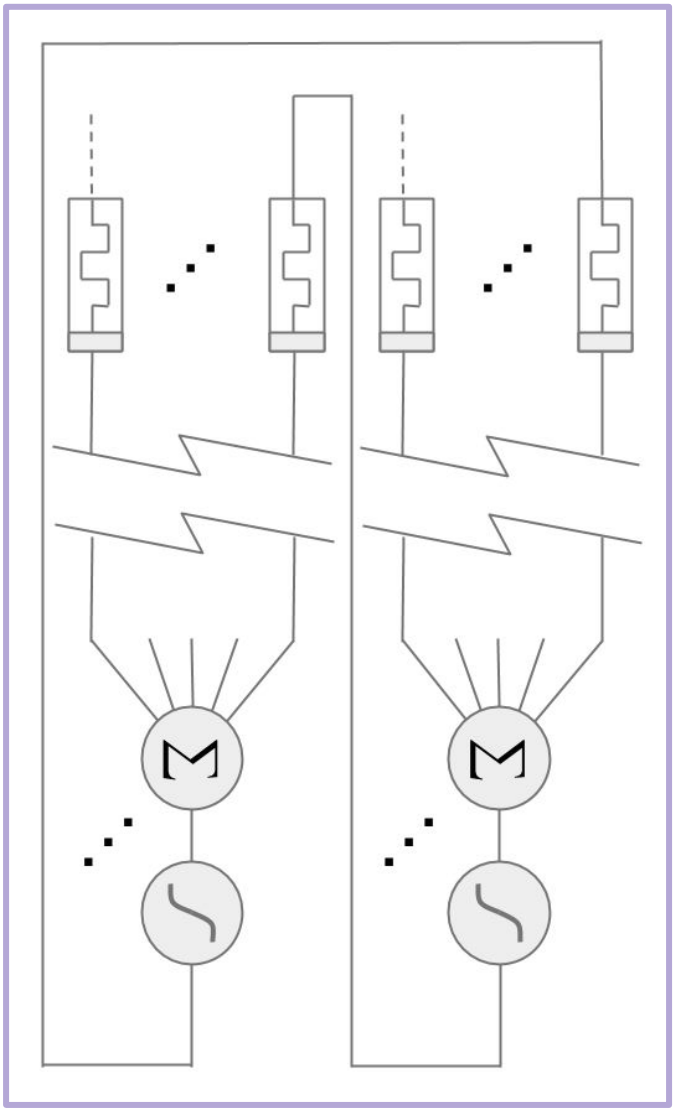

CMOS Spiking Neuron

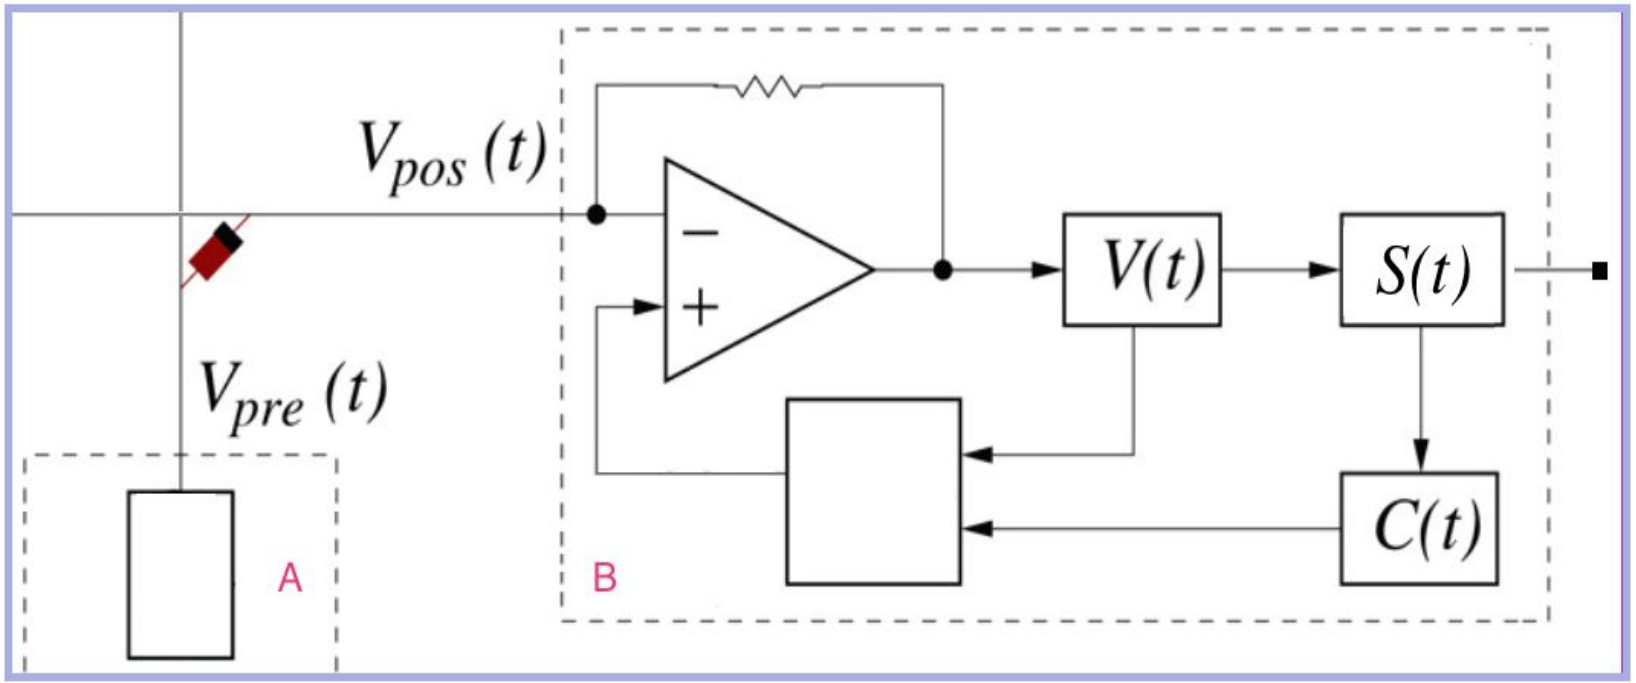

Memristor Circuits

LEAKY INTEGRATE AND FIRE MODEL

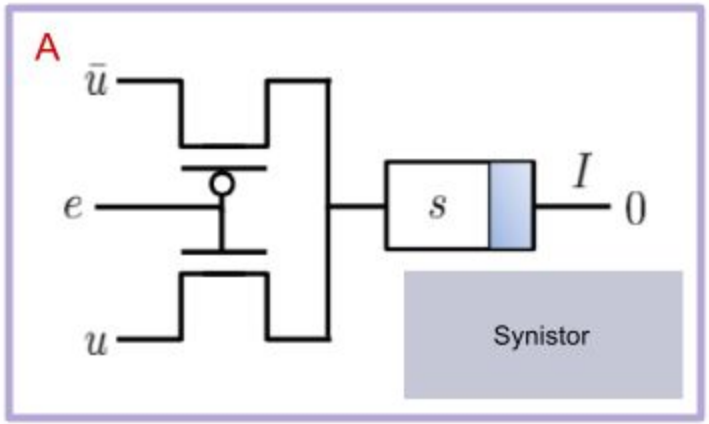

HODGKIN-HUXLEY MODEL [GOLD STANDARD]

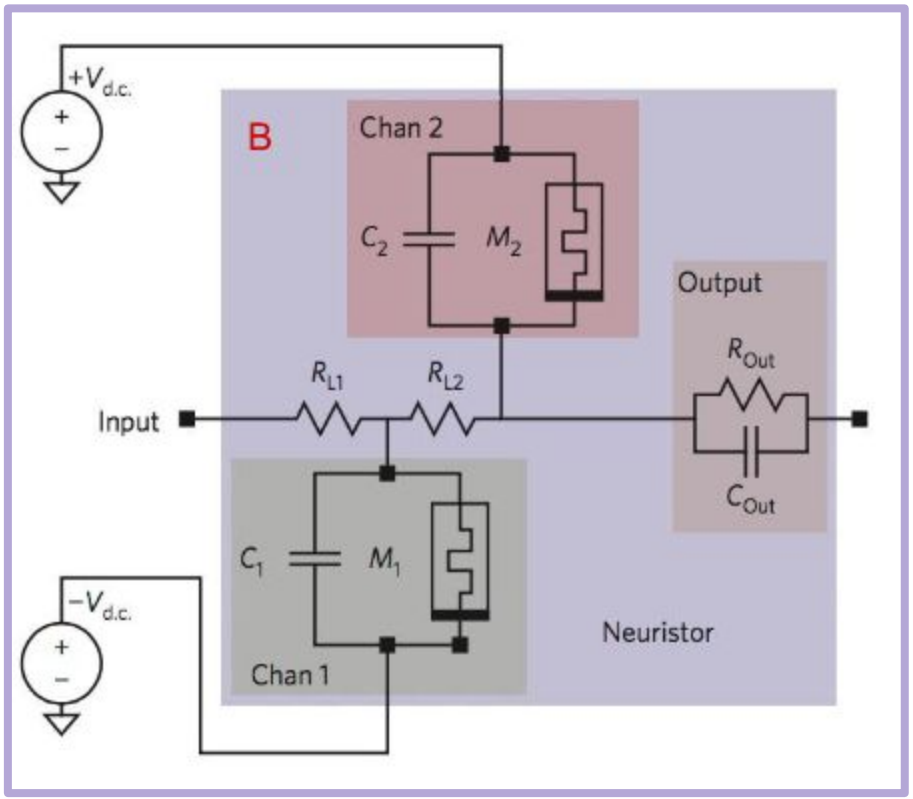

ION CHANNELS: #1 EXCITATORY & 2 INHIBITORY

Memristor Circuits

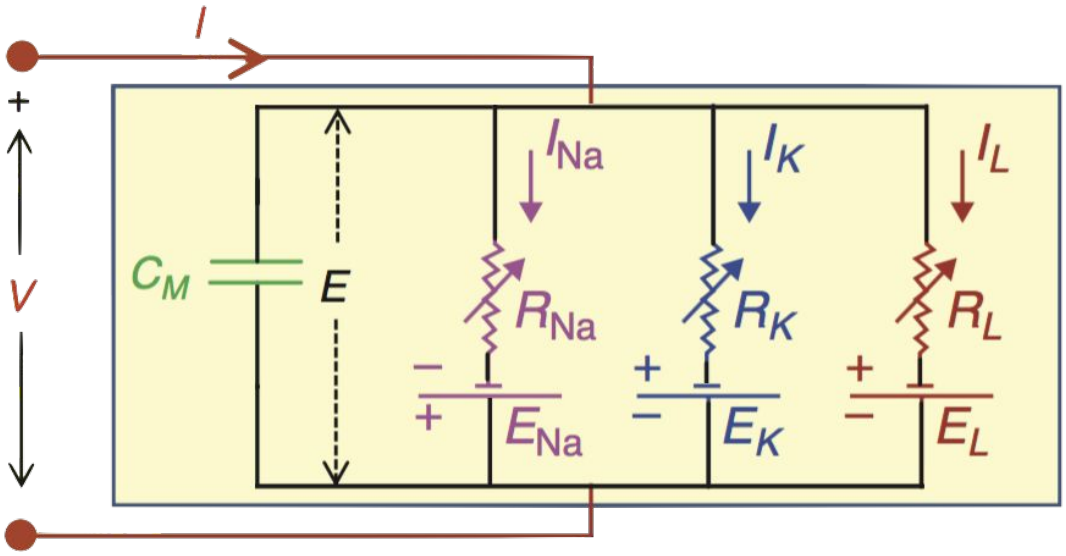

Original Variable-Resistor Hodgkin-Huxley Model

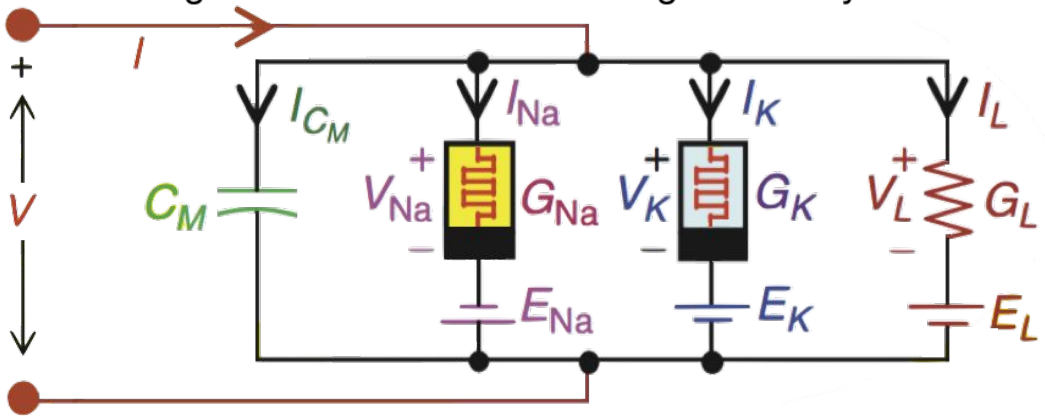

Biologically-Accurate Memristor Hodgkin-Huxley

Power Distribution

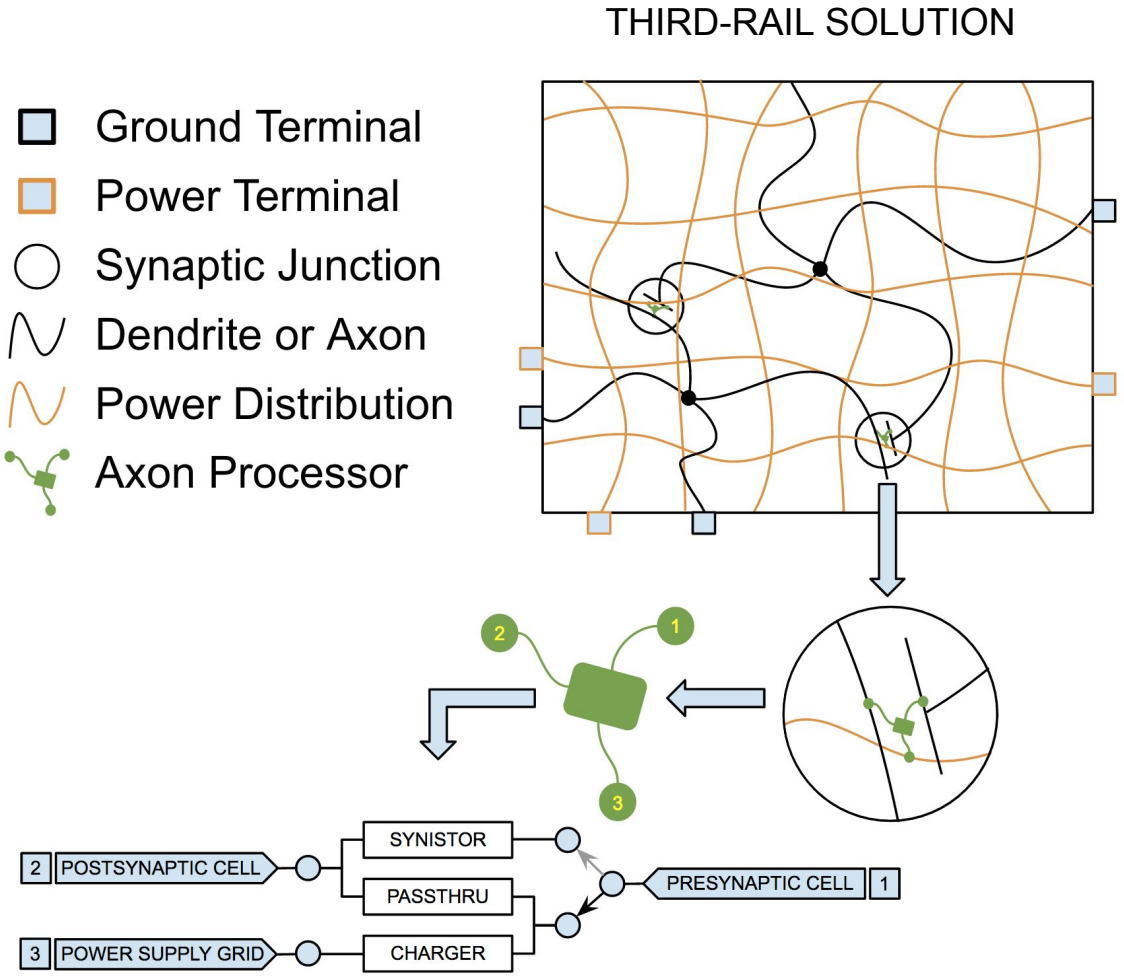

Power Distribution

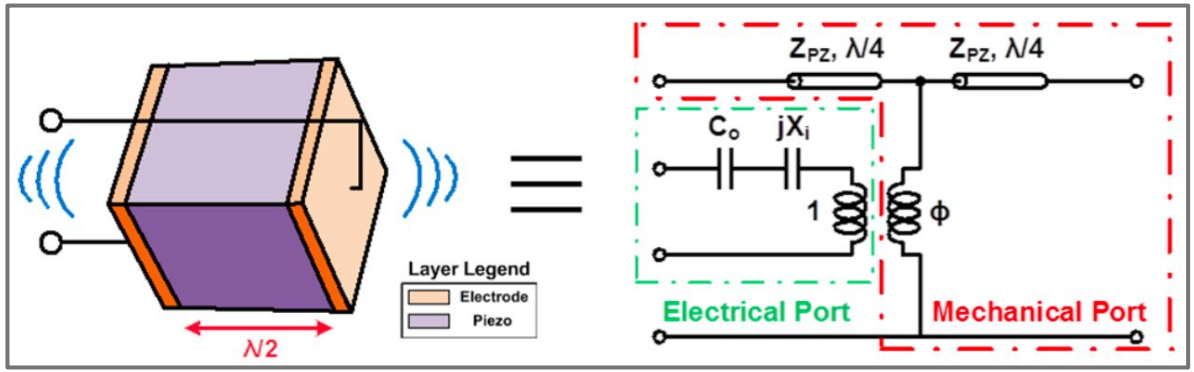

ULTRASONIC POWER HARVESTER SOLUTION

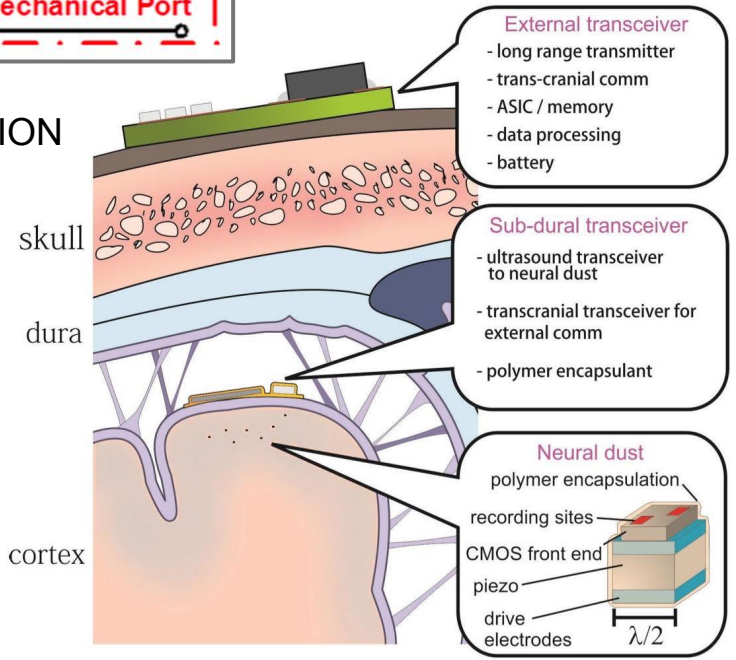

Interconnect Matrix

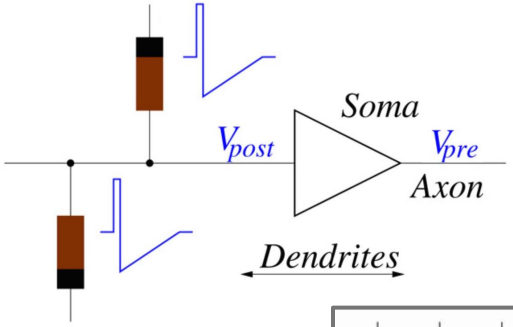

TWO LAYER NEURAL NETWORK

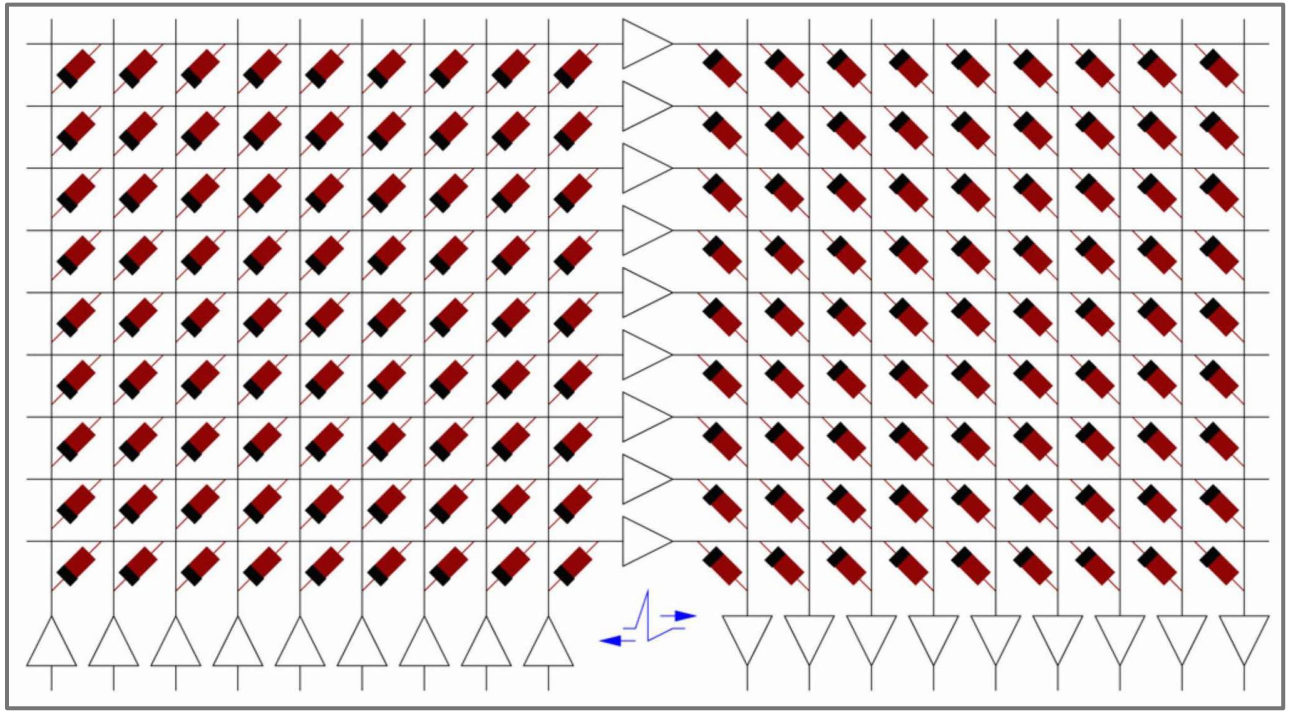

Interconnect Matrix

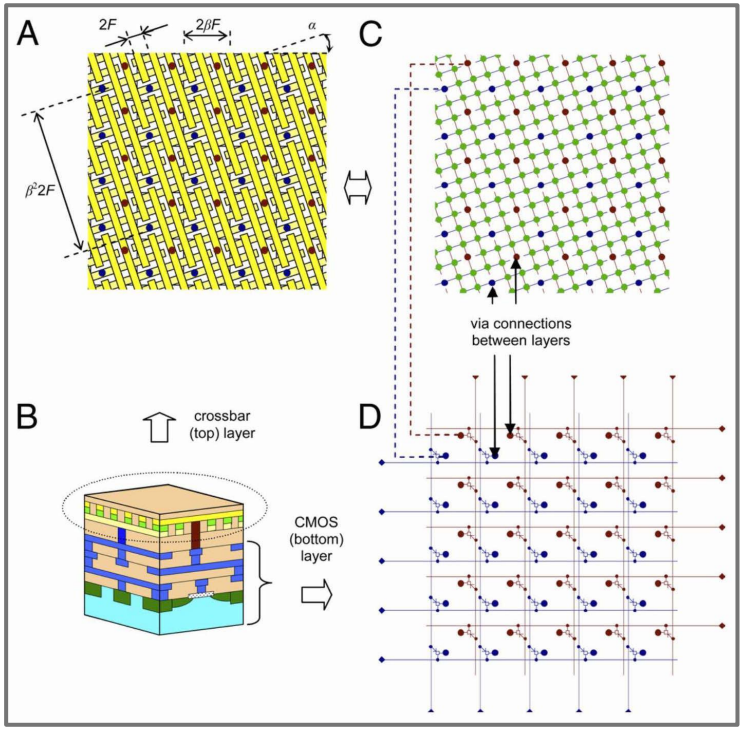

TWO-DIMENSIONAL CROSSBAR

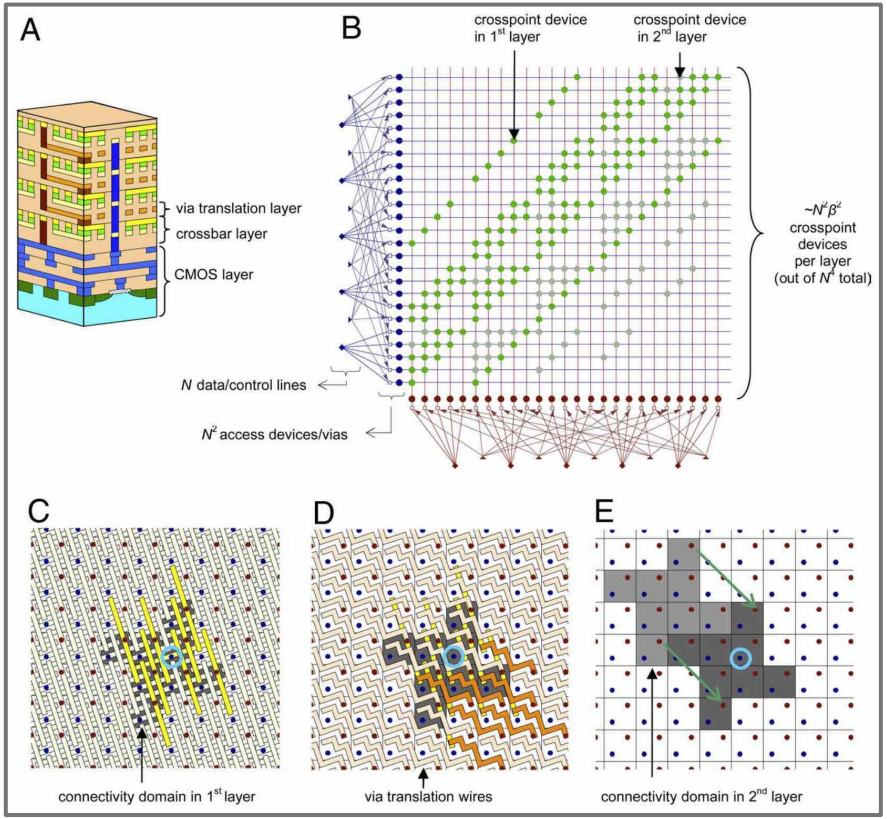

THREE-DIMENSIONAL CROSSBAR

Interconnect Matrix

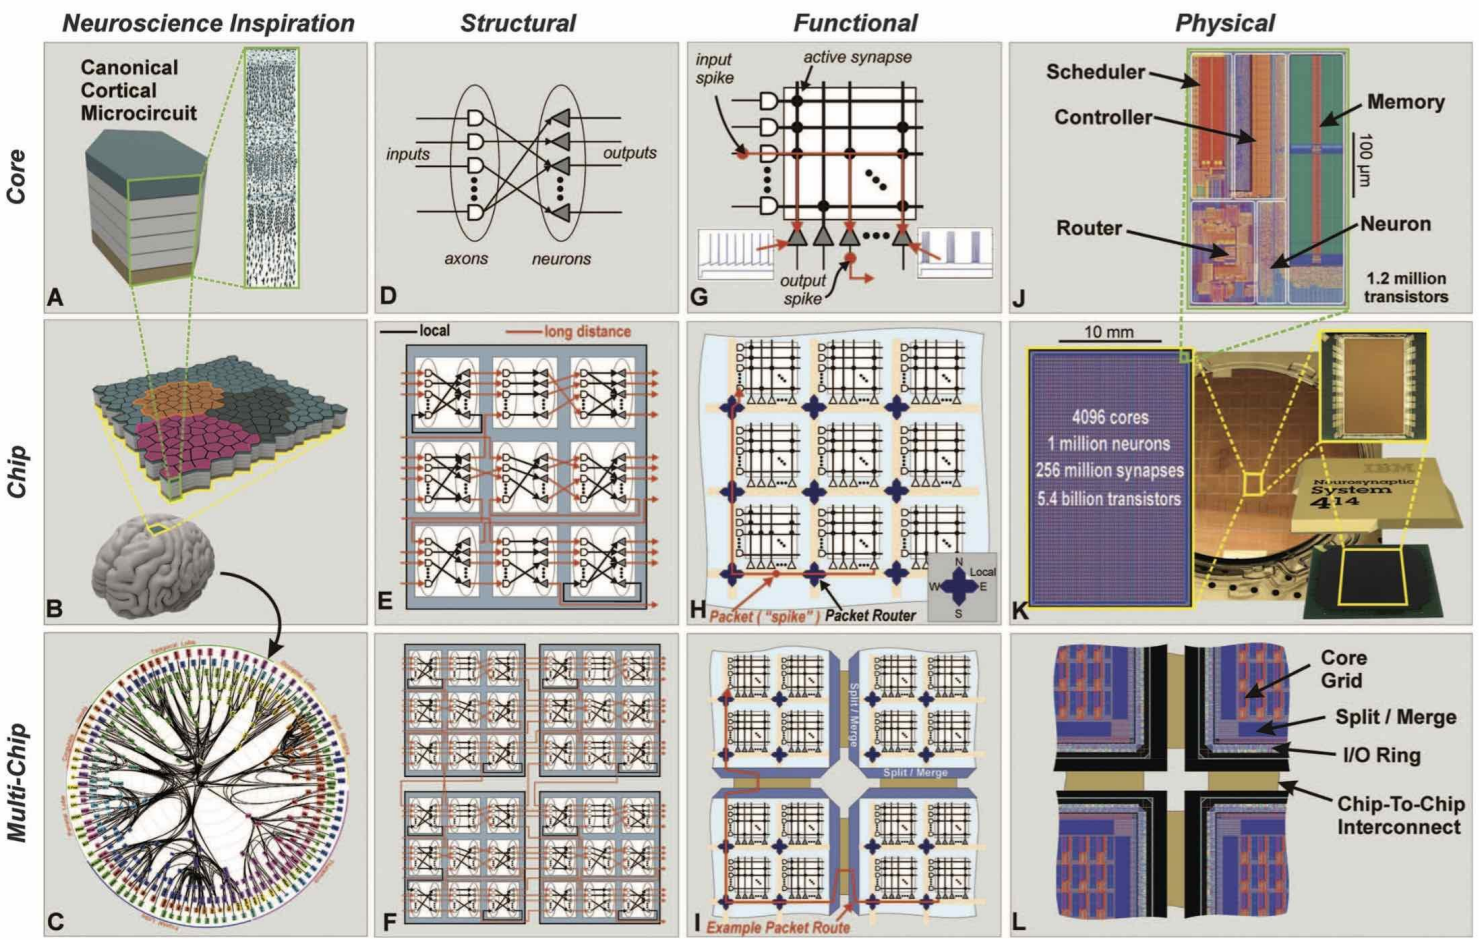

IBM TRUE NORTH NEUROMORPHIC COMPUTING [DARPA SYNAPSE PROGRAM]

Expansion Circuitry: Modeling and Emulation†

I. Scalable Expansion Circuit Emulation — See Appendix A for relevant research papers.
   1. Build Basic Emulator — SPICE equivalent circuits, replicate STDP & BP results.
   2. Add FEA Capabilities — crosstalk, substrate leakage, conductance properties.
   3. Scalable Emulation — parallel SPICE, LUT, train DNN reaction diffusion for FEA.
   4. Power Distribution — one- and two-wire third-rail, ultrasonic wireless solutions.

II. Low-Power Analog Circuit Design — See Appendix B for relevant research papers.
   1. Incumbent Technologies — CMOS, MOSFET, TFET in the subthreshold regime.
   2. Memristor Technologies — TiO_2, NbO_2 thermal, polymeric and spintronic devices.

III. Reaction Diffusion Simulation — See Appendix C for relevant research papers.
   1. Incumbent Technologies — Blue Brain, GENESIS, MCell, NEST, NEURON.
   2. Models — FitzHugh-Nagumo, Hodgkin-Huxley, Leaky Integrate and Fire.
   3. Parallelism — SIMD (GPU), MIMD (HYPERCUBE), Monte-Carlo Sampling.

[†] Tentative expansion-circuit modeling, design and scalable emulation objectives for year-one of the BrainMaker project.

Supplements

Expansion Circuitry: Flights of Fancy

Applications of expansion circuits listed in increasingly fanciful order:

- Walgreen’s $19.95 DIY Countertop Connectomic Continuity Kit
- Ultrasonic Interrogation of Expansion Circuits for Connectomics
- Expansion-Circuit Software Development & Hardware Emulator
- Sub Milliwatt Powered Fly Brain Expansion-Circuit Drone Pilots
- Cordwainer† Project — 3D Deposition Expansion-Circuit Printer

[†] Cordwainer — an anglicized variant of the Anglo-Norman "cordewaner", from Old French cordovan (“(leather) of Cordova”). Cordwainer Smith was the pen-name used by American author Paul Myron Anthony Linebarger for his science fiction works. Linebarger was a noted East Asia scholar and an expert in psychological warfare. The eponymous CordWainer Project refers to a disturbing science fiction story by Smith with the title “Think Blue, Count Two” in which the brain of a mouse is plasticized, rendered essentially immortal and integrated into the automated pilot of an interstellar spaceship scheduled to make a journey of hundreds of lightyears with no company other than a cargo of cryogenically-preserved hibernating colonists.

Biological Brains & Building Artificial Brains

- Suppose our objective is to build an artificial brain
- We are already pretty good at building computers
- What can we learn from studying biological brains
  - ➢ von Neumann architecture serial execution
  - ➢ versus associative, highly-parallel, in-place

- What have we already learned from neuroscience
- Macroscale: +1 — Microscale: +1 — Molecular: ?
- What is it specifically about the wiring of the brain
  - ➢ rich collection of network computing motifs
  - ➢ new parallel algorithms, proof-of-feasibility

Killer Application for Resistive Cross-Point Devices: Training Deep Neural Networks?

```
@article{GokmenandVlasovCoRR-16,
       author = {Tayfun Gokmen and Yurii Vlasov},
        title = {Acceleration of Deep Neural Network Training with Resistive Cross-Point Devices},
      journal = {CoRR},
       volume = {arXiv:1603.07341},
         year = {2016},
     abstract = {In recent years, deep neural networks (DNN) have demonstrated significant business
impact in large scale analysis and classification tasks such as speech recognition, visual object
detection, pattern extraction, etc. Training of large DNNs, however, is universally considered as
time consuming and computationally intensive task that demands datacenter-scale computational
resources recruited for many days. Here we propose a concept of resistive processing unit (RPU)
devices that can potentially accelerate DNN training by orders of magnitude while using much less
power. The proposed RPU device can store and update the weight values locally thus minimizing data
movement during training and allowing to fully exploit the locality and the parallelism of the
training algorithm. We identify the RPU device and system specifications for implementation of an
accelerator chip for DNN training in a realistic CMOS-compatible technology. For large DNNs with
about 1 billion weights this massively parallel RPU architecture can achieve acceleration factors of
30,000X compared to state-of-the-art microprocessors while providing power efficiency of 84,000
GigaOps/s/W. Problems that currently require days of training on a datacenter-size cluster with
thousands of machines can be addressed within hours on a single RPU accelerator. A system consisting
of a cluster of RPU accelerators will be able to tackle Big Data problems with trillions of
parameters that is impossible to address today like, for example, natural speech recognition and
translation between all world languages, real-time analytics on large streams of business and
scientific data, integration and analysis of multimodal flows from massive numbers of sensors.}
}
```

Killer Application for Resistive Cross-Point Devices: Training Deep Neural Networks?

```
@inproceedings{BojnordiandIpekHPCA-16,
    author = {M.N. Bojnordi and E. Ipek},
     title = {Memristive Boltzmann Machine: Hardware Acceleration for Optimization & Deep Learning},
 booktitle = {Proceedings of the International Symposium on High Performance Computer Architecture},
      year = {2016},
  abstract = {The Boltzmann machine is a massively parallel computational model capable of solving a
broad class of combinatorial optimization problems. In recent years, it has been successfully
applied to training deep machine learning models on massive datasets. High performance
implementations of the Boltzmann machine using GPUs, MPI-based HPC clusters, and FPGAs have been
proposed in the literature. Regrettably, the required all-to-all communication among the processing
units limits the performance of these efforts. This paper examines a new class of hardware
accelerators for large-scale combinatorial optimization and deep learning based on memristive
Boltzmann machines. A massively parallel, memory-centric hardware accelerator is proposed based on
recently developed resistive RAM (RRAM) technology. The proposed accelerator exploits the electrical
properties of RRAM to realize in situ, fine-grained parallel computation within memory arrays,
thereby eliminating the need for exchanging data between the memory cells and the computational
units. Two classical optimization problems, graph partitioning and boolean satisfiability, and a
deep belief network application are mapped onto the proposed hardware. As compared to a multicore
system, the proposed accelerator achieves 57× higher performance and 25× lower energy with virtually
no loss in the quality of the solution to the optimization problems. The memristive accelerator is
also compared against an RRAM based processing-in-memory (PIM) system, with respective performance
and energy improvements of 6.89× and 5.2×.},
}
```

Killer Application for Resistive Cross-Point Devices: Training Deep Neural Networks?

```
@article{EsseretalCoRR-16,
       author = {Steven K. Esser and Paul A. Merolla and John V. Arthur and Andrew S. Cassidy and
Rathinakumar Appuswamy and Alexander Andreopoulos and David J. Berg and Jeffrey L. McKinstry and
Timothy Melano and Davis R. Barch and Carmelo di Nolfo and Pallab Datta and Arnon Amir and Brian
Taba and Myron D. Flickner and Dharmendra S. Modha},
        title = {Convolutional Networks for Fast, Energy-Efficient Neuromorphic Computing},
      journal = {CoRR},
       volume = {arXiv:1603.08270},
         year = {2016},
     abstract = {Deep networks are now able to achieve human-level performance on a broad spectrum
of recognition tasks. Independently, neuromorphic computing has now demonstrated unprecedented
energy-efficiency through a new chip architecture based on spiking neurons, low precision synapses,
and a scalable communication network. Here, we demonstrate that neuromorphic computing, despite its
novel architectural primitives, can implement deep convolution networks that (i) approach state-of-
the-art classification accuracy across 8 standard datasets, encompassing vision and speech, (ii)
perform inference while preserving the hardware's underlying energy-efficiency and high throughput,
running on the aforementioned datasets at between 1100 and 2300 frames per second and using between
25 and 325 mW (effectively > 5000 frames / sec / W) and (iii) can be specified and trained using
backpropagation with the same ease-of-use as contemporary deep learning. For the first time, the
algorithmic power of deep learning can be merged with the efficiency of neuromorphic processors,
bringing the promise of embedded, intelligent, brain-inspired computing one step closer.}
}
```