Neuro-Visualization at the Network Edge

Kunal Lillaney

Advisor: Dr. Randal Burns

Johns Hopkins University

HBP CodeJam Workshop #7

13 January 2016

OPEN CONNECTOME PROJECT

COLLECTIVELY REVERSE-ENGINEERING THE BRAIN ONE SYNAPSE AT A TIME.

Overview

- Human visualization drives analysis in this field
- Visualization of petascale neuroscience imaging
  - Stored on the cloud or at data center
  - Internet latencies ruin user experience
- Deploy distributed caching
  - To offload server I/O and rendering
  - To reduce network latency
- Customized to neuroscience data patterns
  - Combination of multi-channel data
  - High selectivity and reduced-dimension projects

OPEN CONNECTOME PROJECT

COLLECTIVELY REVERSE-ENGINEERING THE BRAIN ONE SYNAPSE AT A TIME.

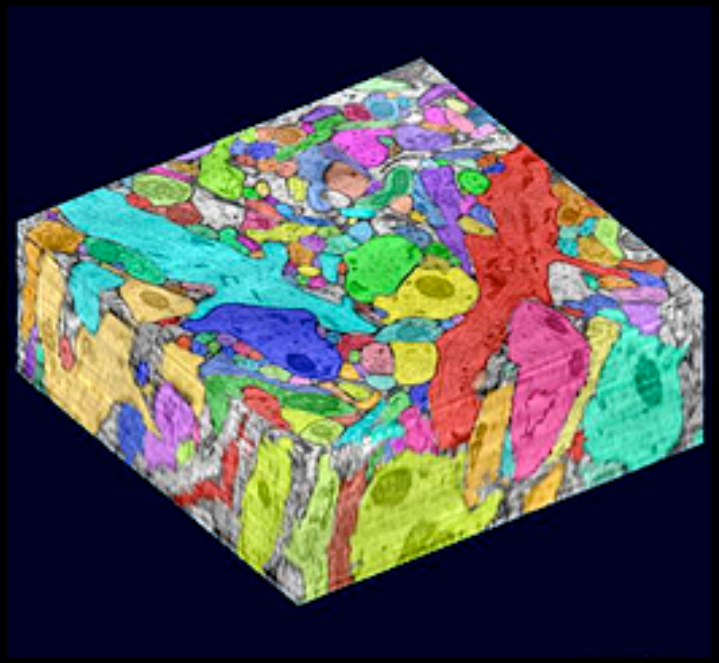

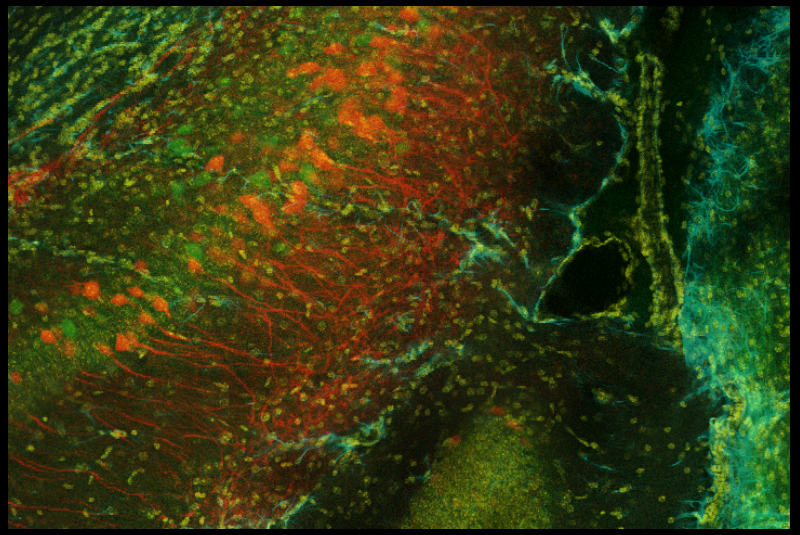

Open-science, data-intensive analysis of the brain

- Peta-scale storage linked with HPC
- Computational vision of brain structure
- Spatial queries (clusters, volumes, distributions)

ARCHITECTURE

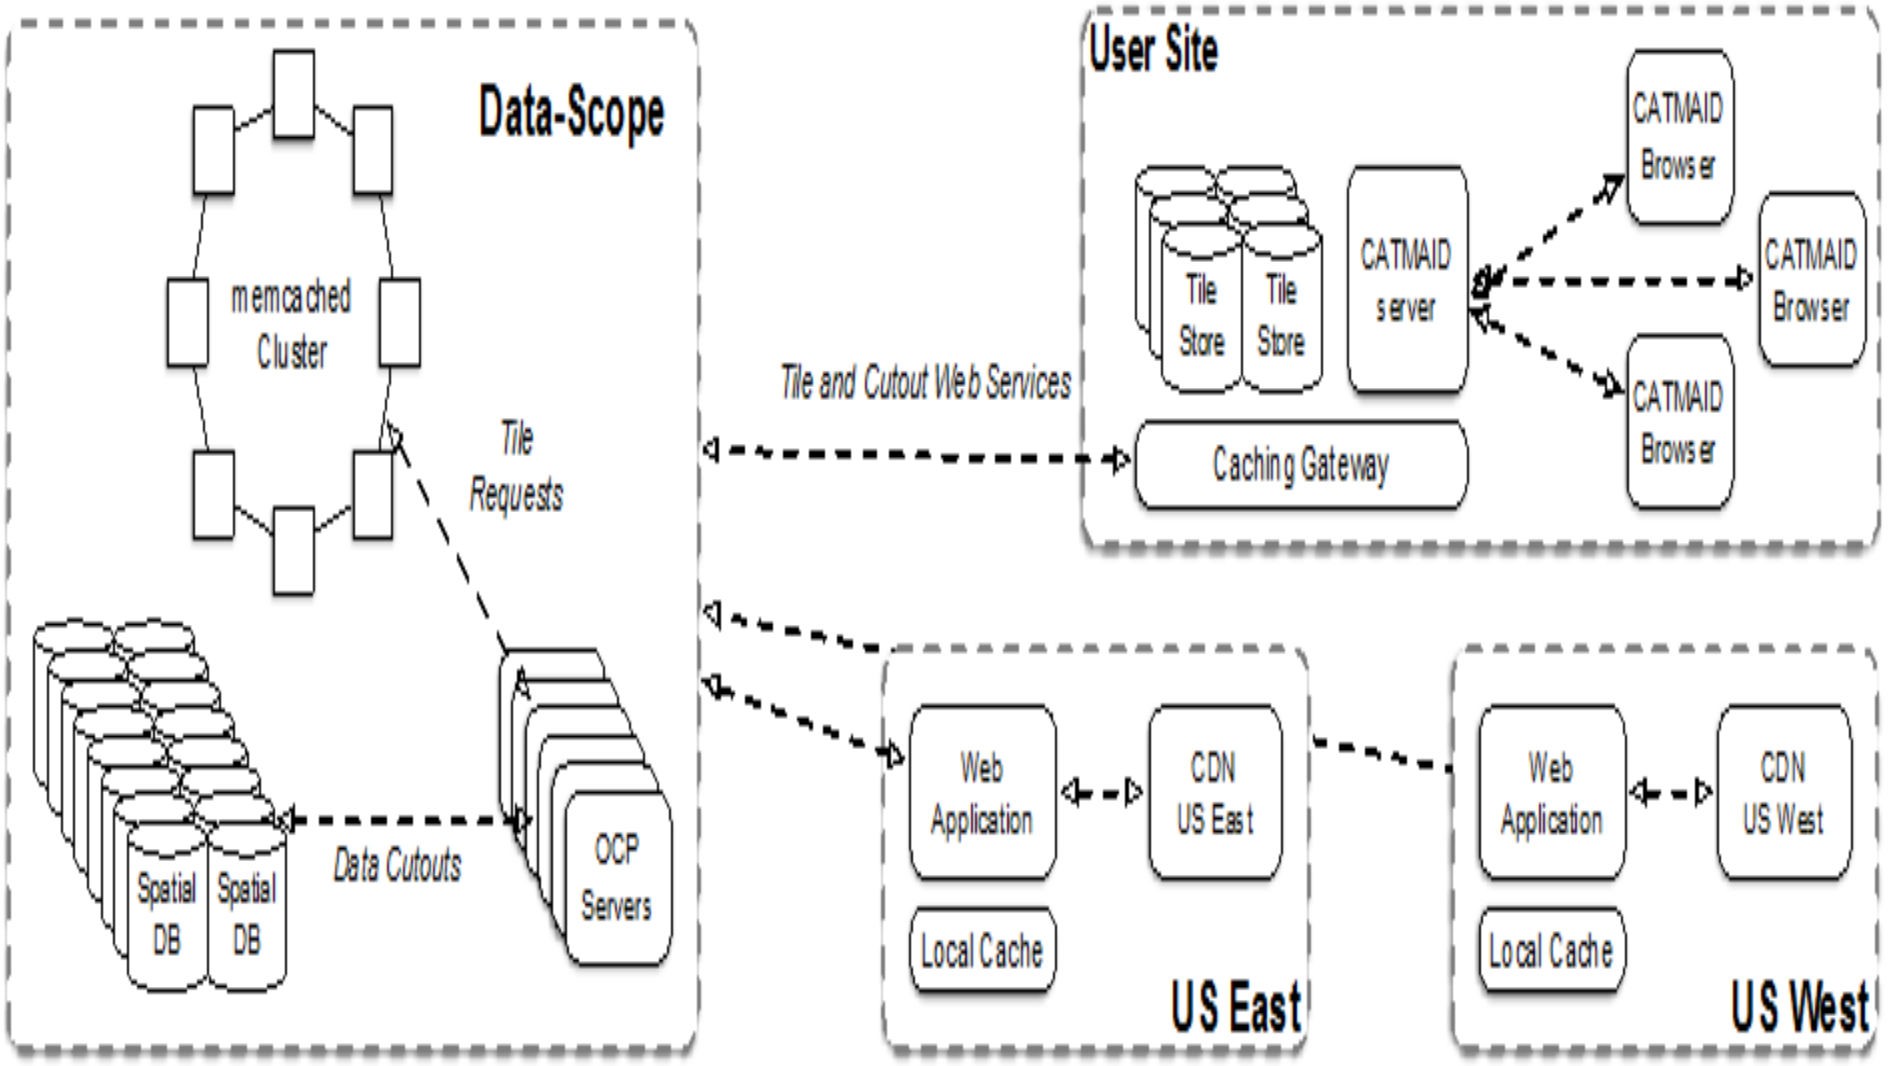

Spatial Database

- Dense 3D or 4D spatial array pationtioned into cuboids
- Space filling curve and Multi-resolution zoom pyramid
- Support for Neuron, Synapse, Segment and more annotation types
- Store ~100TB of imaging data

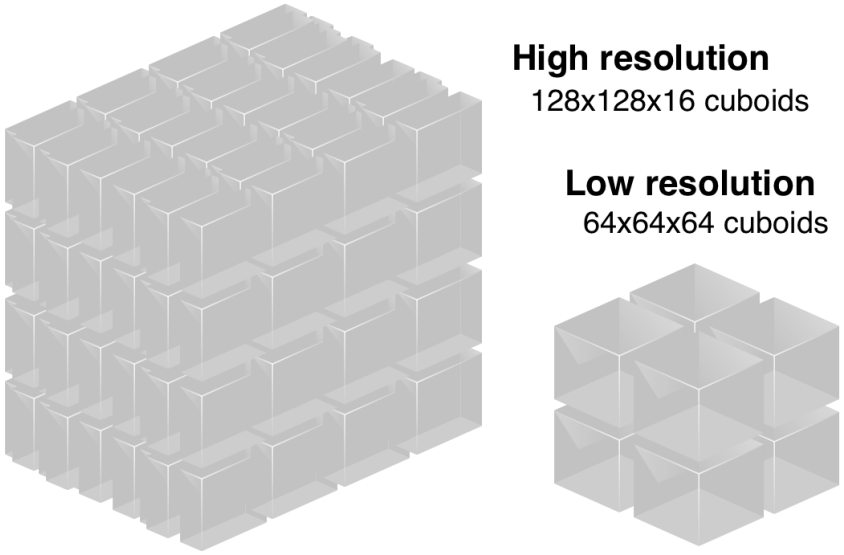

Z order space filling curve

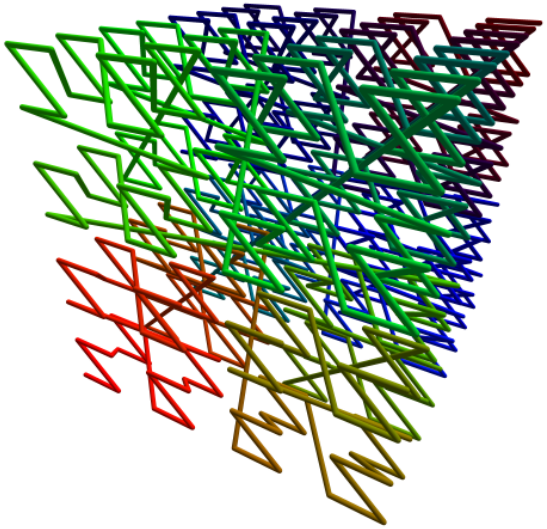

CATMAID

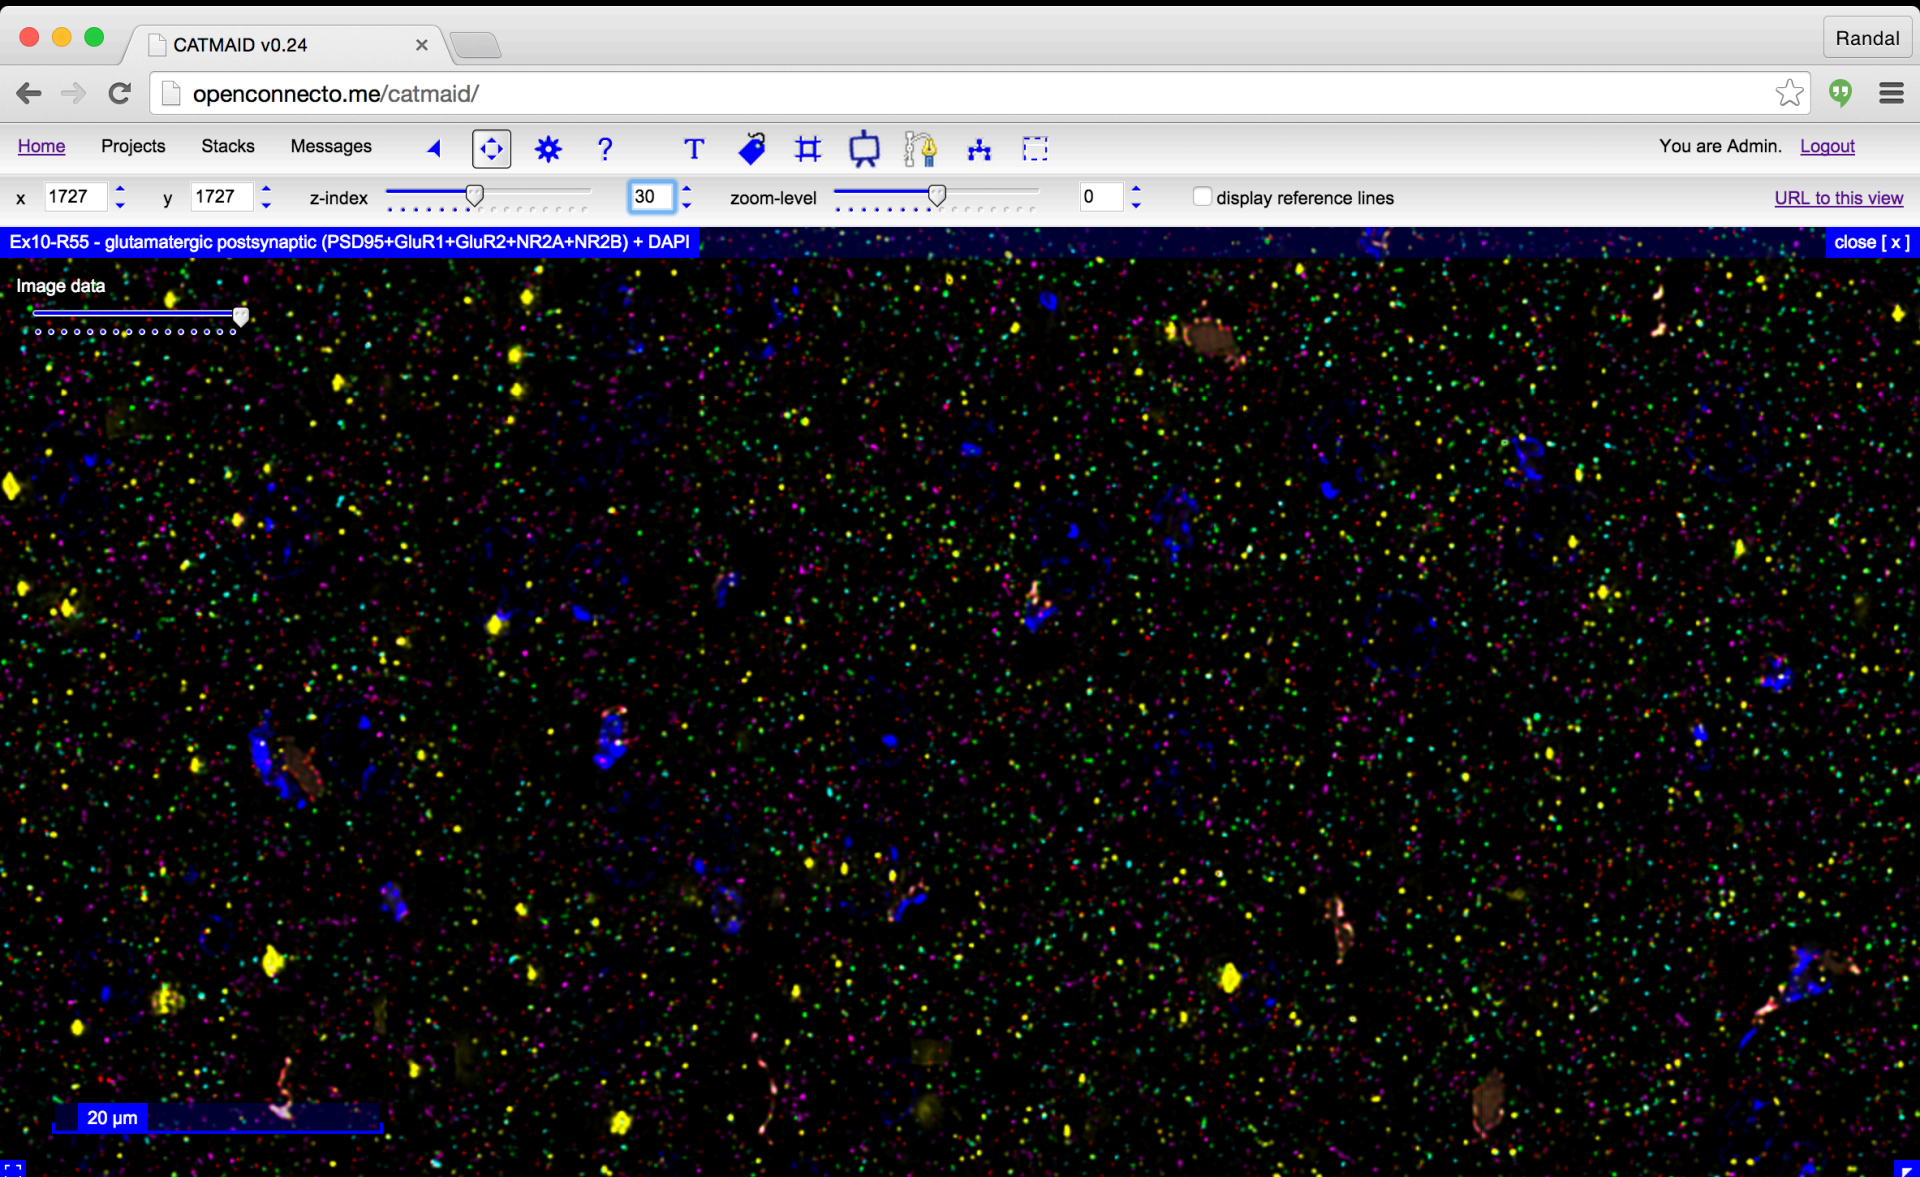
CATMAID v0.24
Randal
openconnecto.me/catmaid/
Home
Projects
Stacks
Messages
You are Admin.
Logout
x 1727
y 1727
z-index
30
zoom-level
0
display reference lines
URL to this view
Ex10-R55 - glutamatergic postsynaptic (PSD95+GluR1+GluR2+NR2A+NR2B) + DAPI
close [x]
Image data
20 µm

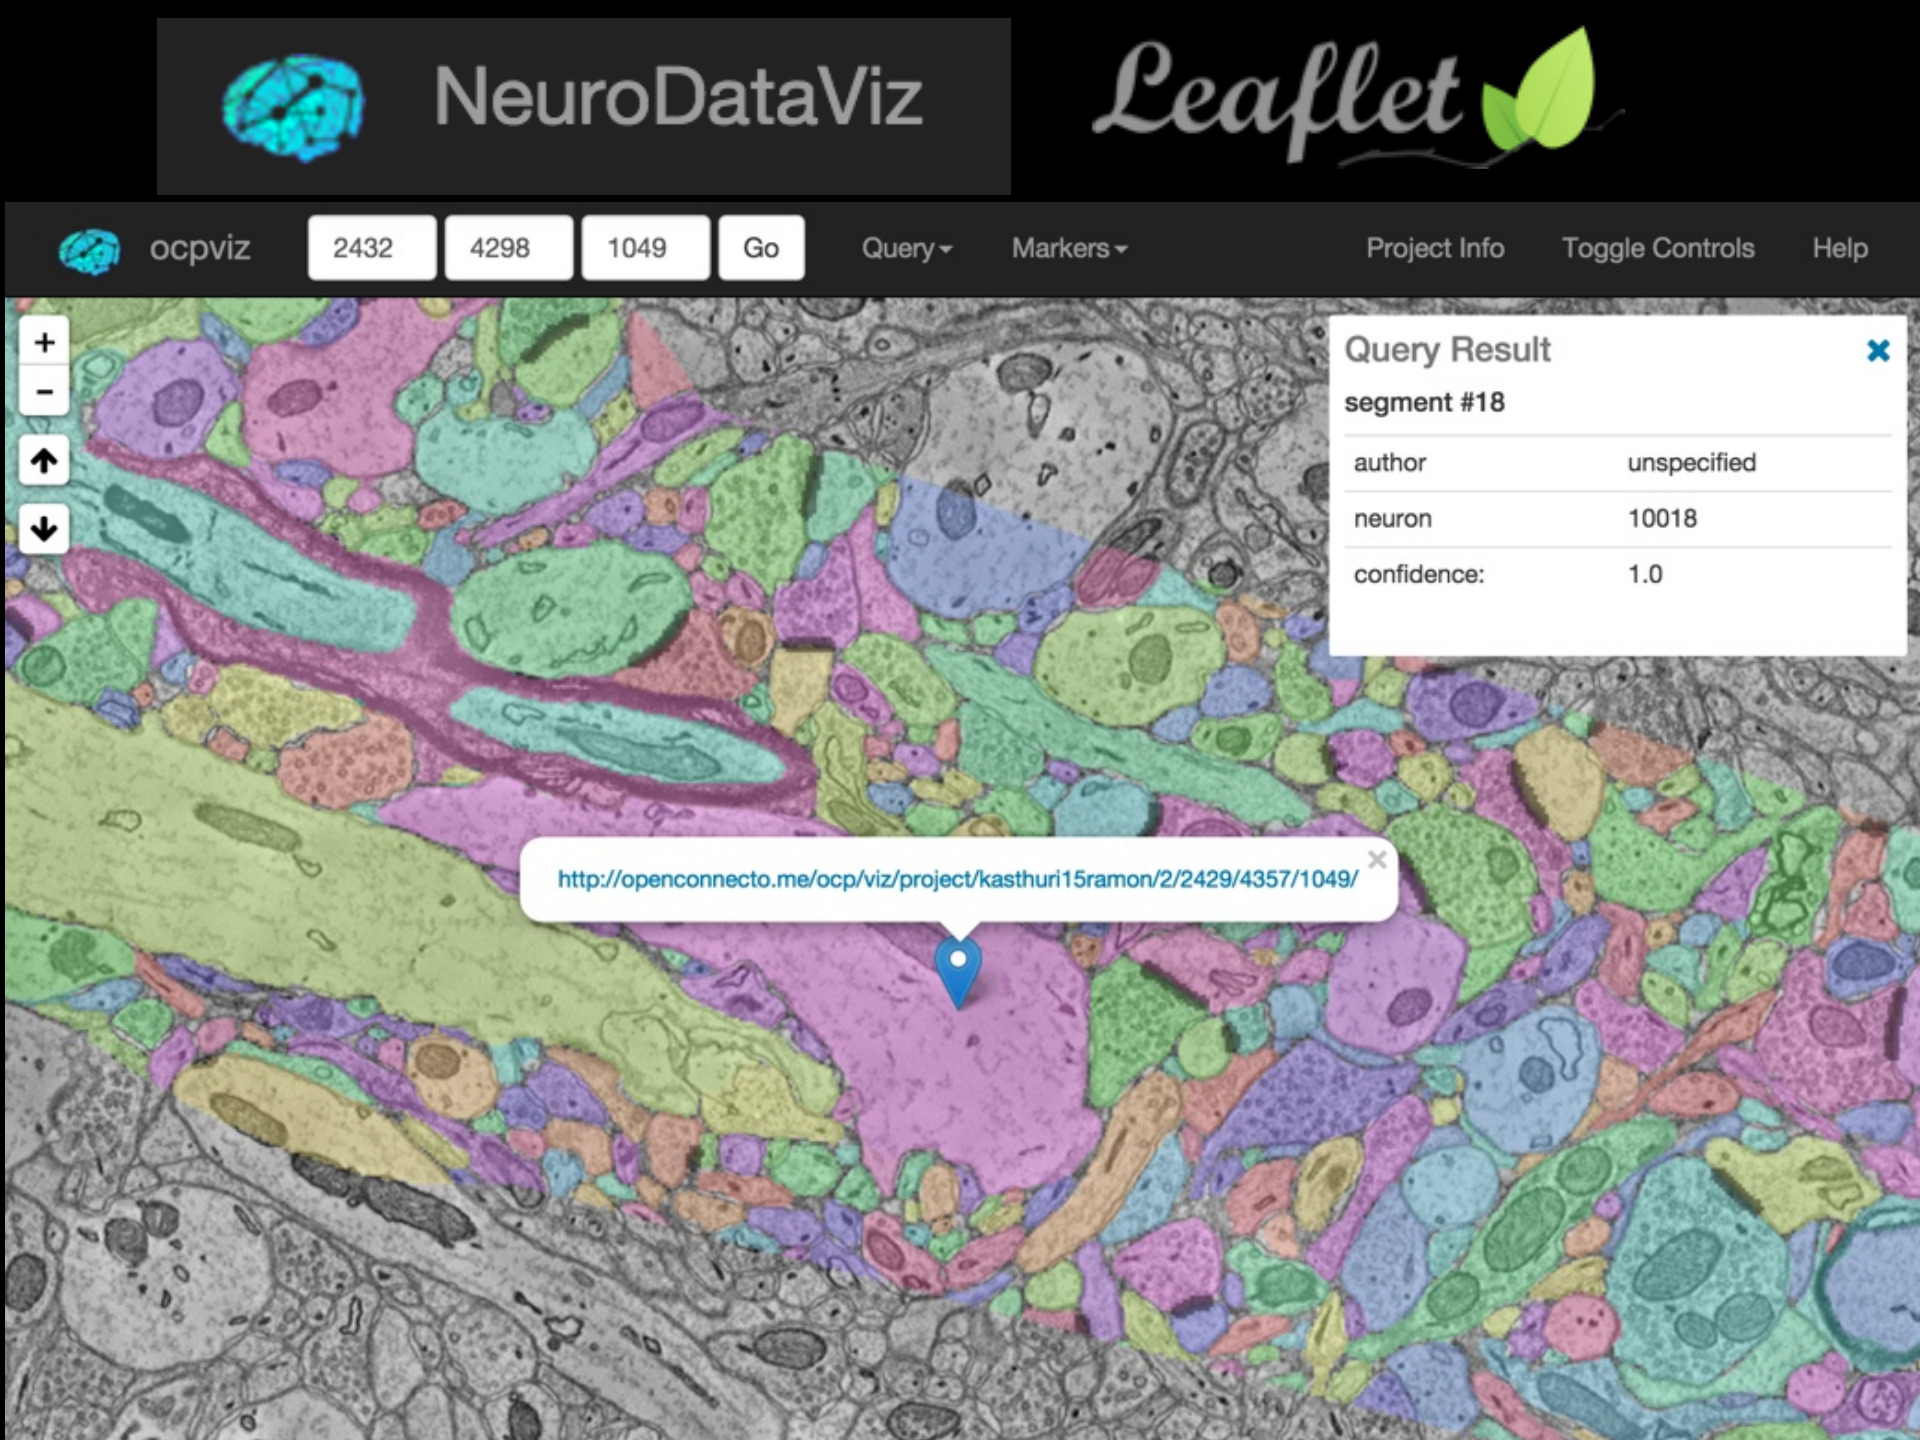
NeuroDataViz
Leaflet
ocpviz
2432
4298
1049
Go
Query
Markers
Project Info
Toggle Controls
Help
Query Result
segment #18
author unspecified
neuron 10018
confidence: 1.0
http://openconnecto.me/ocp/viz/project/kasthuri15ramon/2/2429/4357/1049/

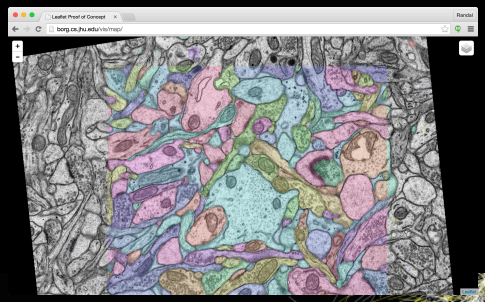

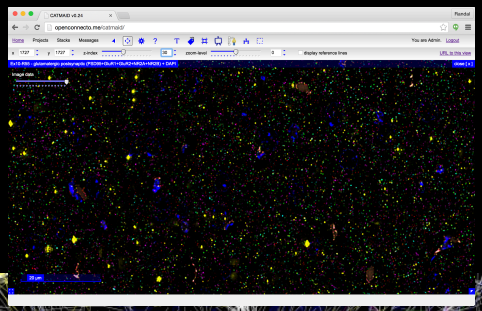

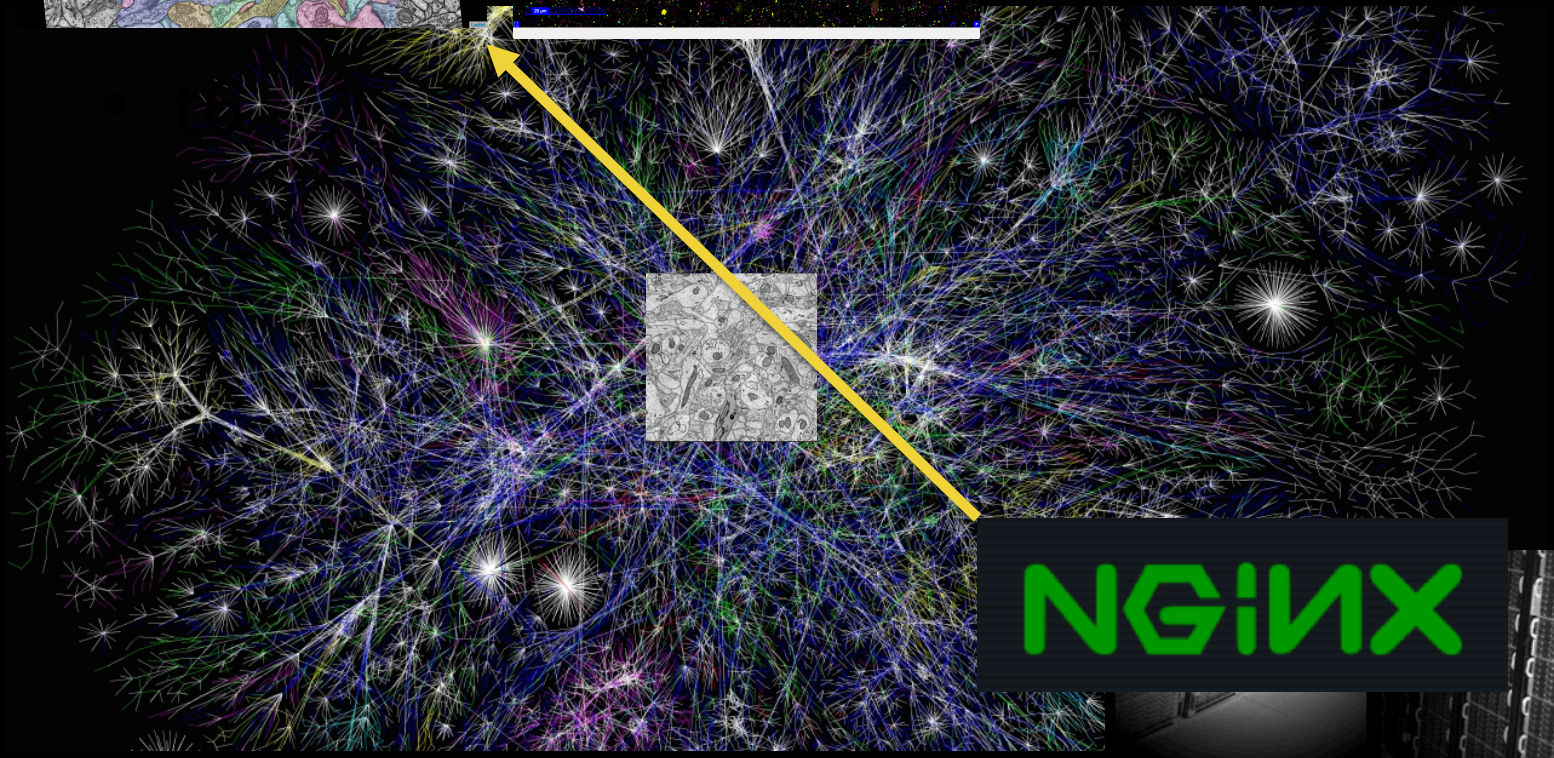

OPEN CONNECTOME PROJECT

COLLECTIVELY REVERSE-ENGINEERING THE BRAIN ONE SYNAPSE AT A TIME.

System Goals

- Visual flow (24+ frames per second)
- Tolerable latency:

  ~100ms initial load (must be < 1 second)
- Need to deliver:
  - Up to 30 512x512 image tiles for each view
  - 6 per layer, up to 5 layers
- Can't do it with Internet latencies

Must push data to network edge, near browser!

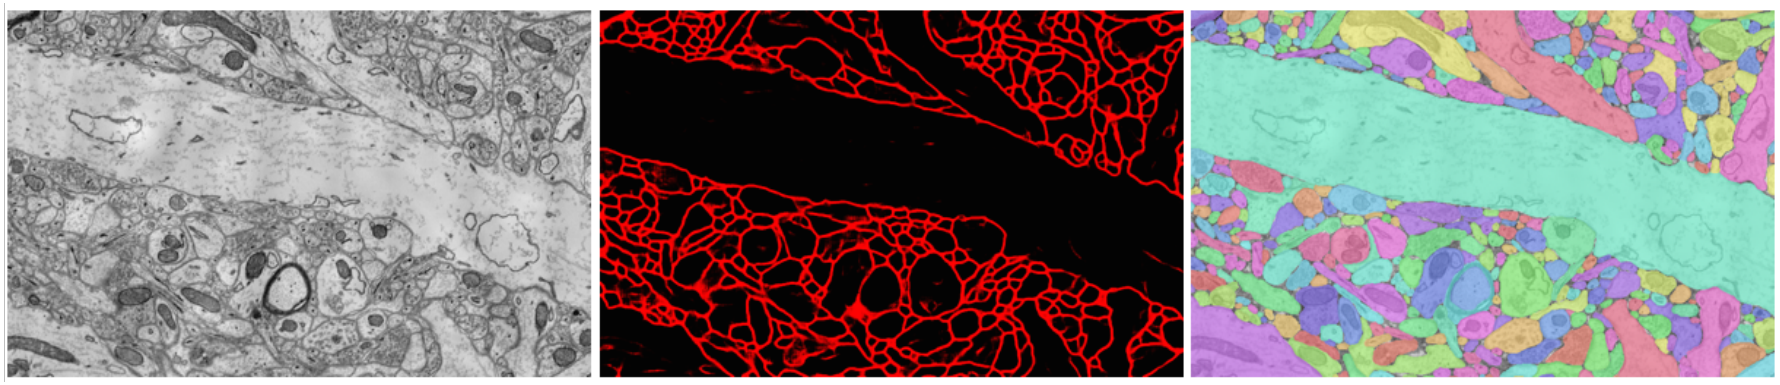

Content-Distribution Network?

- Ingest content, push toward consumer
  - Requires knowledge of content to be consumed
- Does not match our data usage

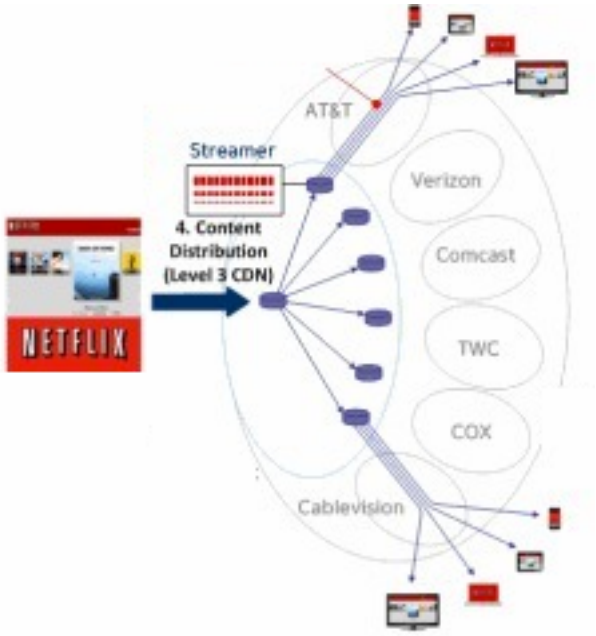

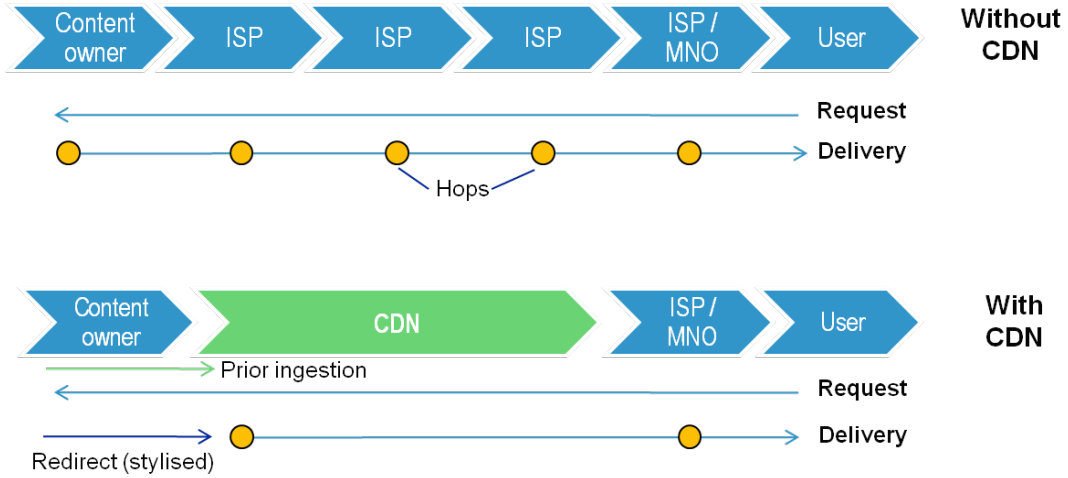

Spatial Data and Usage Patterns

- Small regions of interest in massive data sets
- Dynamic materializations of 2-d tiles
  - From 3-d or 4-d databases
  - Any (axis orthogonal) cutting plane

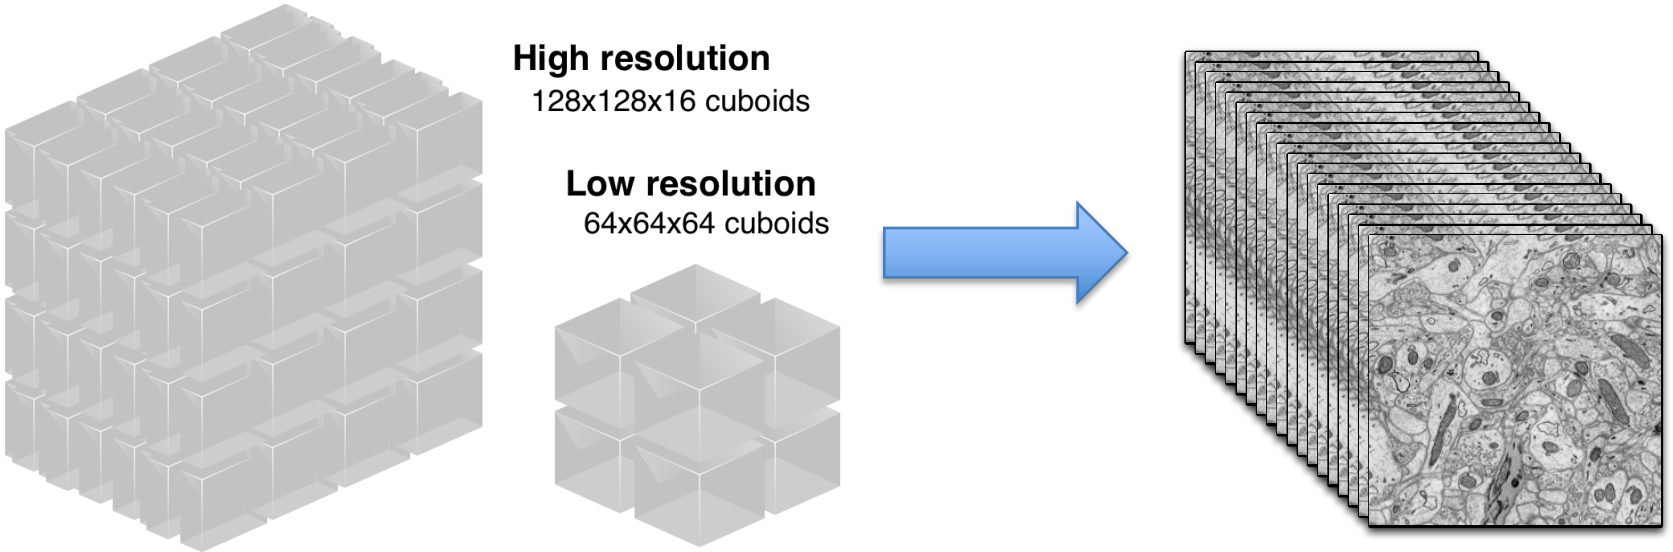

Spatial Data and Usage Patterns

- Exponentially many combinations of channels from the same data set (flattened for performance)

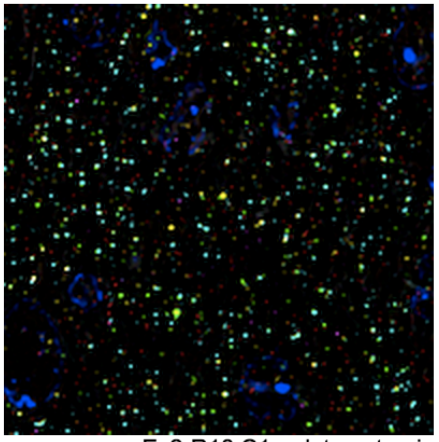

Ex2-R18 C1 - glutamatergic postsynaptic: PSD95 (cyan), GluR1 (yellow), GluR2 (magenta), NR2A (red), NR2B (green), DAPI (blue)

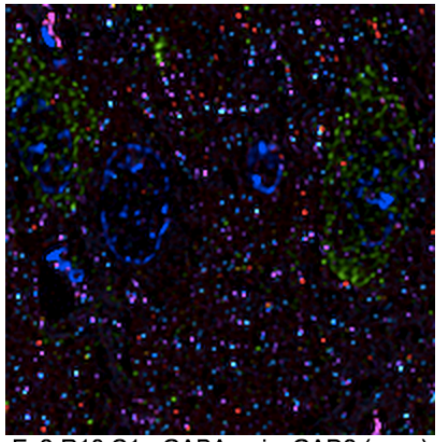

Ex2-R18 C1 - GABAergic: GAD2 (cyan), vGAT (yellow), gephyrin (magenta), GABAARa1 (red), PV25 (green), and DAPI (blue).

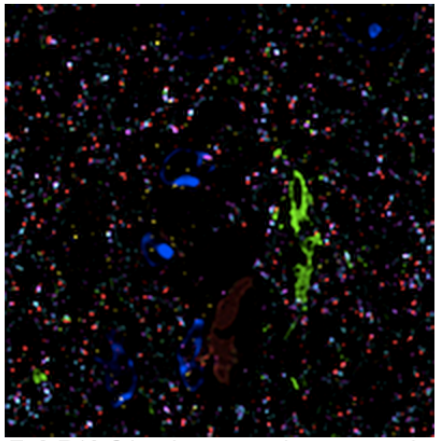

Ex2-R18 C1 - glutamatergic presynaptic: Synapsin1 (cyan), vGluT1 (yellow), vGluT2 (magenta), PSD95 (red), GFP (green), and DAPI (blue).

Must push data to network edge AND must dynamically manage data contents (Caching)!

Caching Architecture

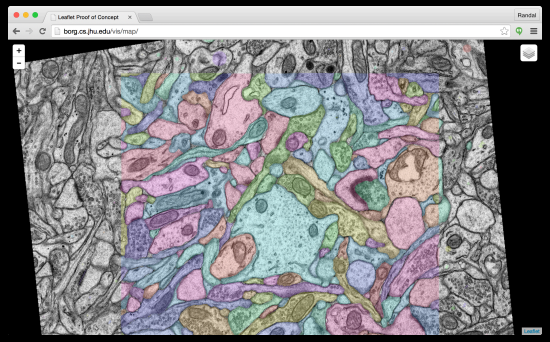

Local Network

Cloud

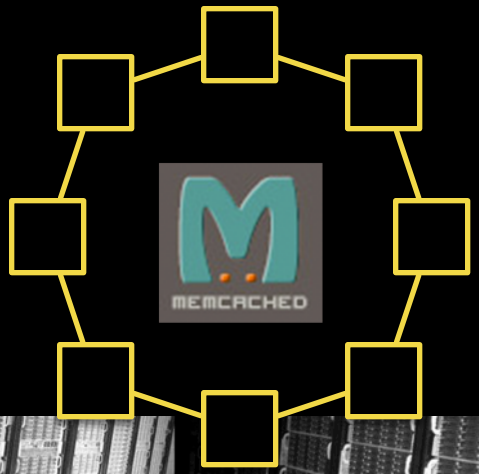

NGINX

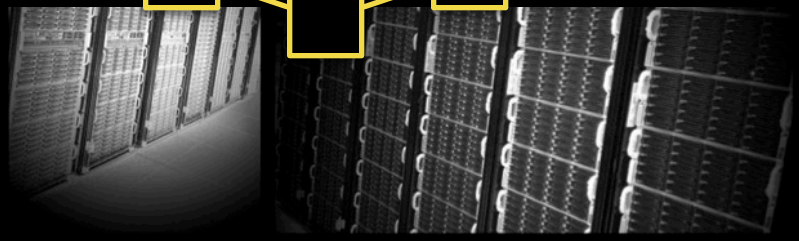

Data Store (PBs)

Disk cache (TBs)

OPEN CONNECTOME PROJECT

COLLECTIVELY REVERSE-ENGINEERING THE BRAIN ONE SYNAPSE AT A TIME.

Tile Request: Initial/Cache Miss

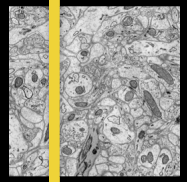

NGINX

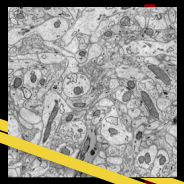

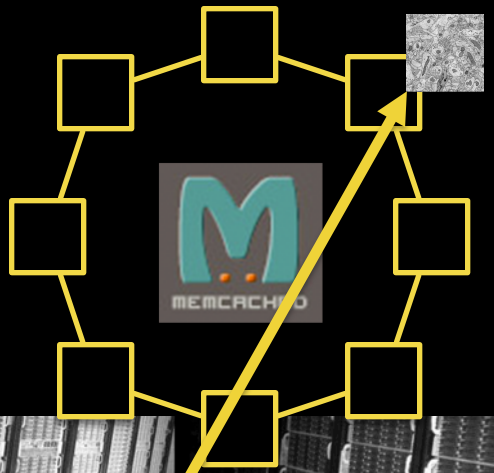

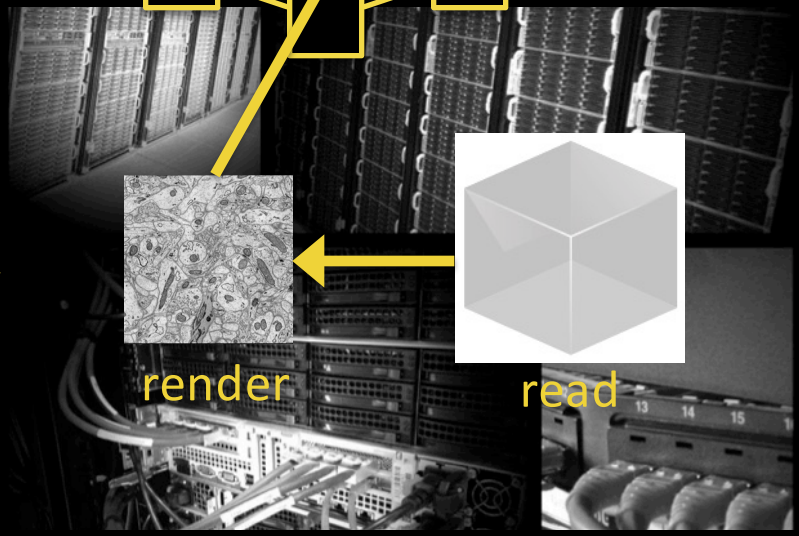

render

read

OPEN CONNECTOME PROJECT

COLLECTIVELY REVERSE-ENGINEERING THE BRAIN ONE SYNAPSE AT A TIME.

Cache Prefetch: Background Load

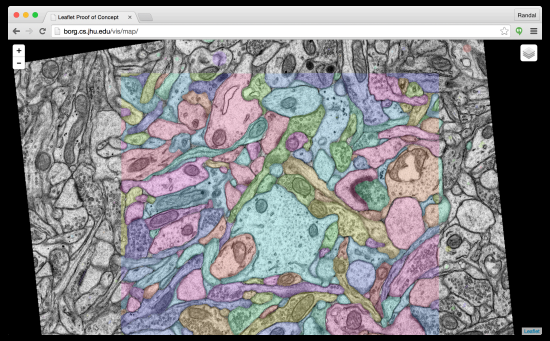

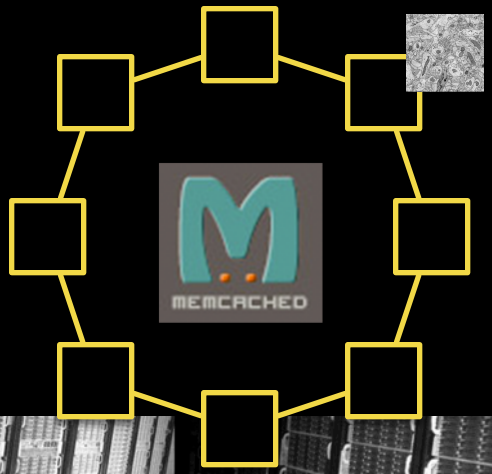

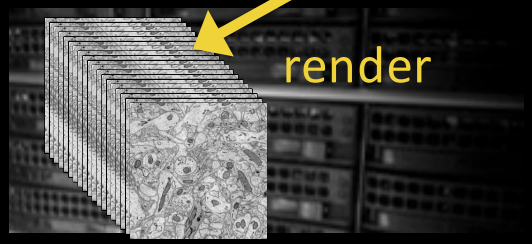

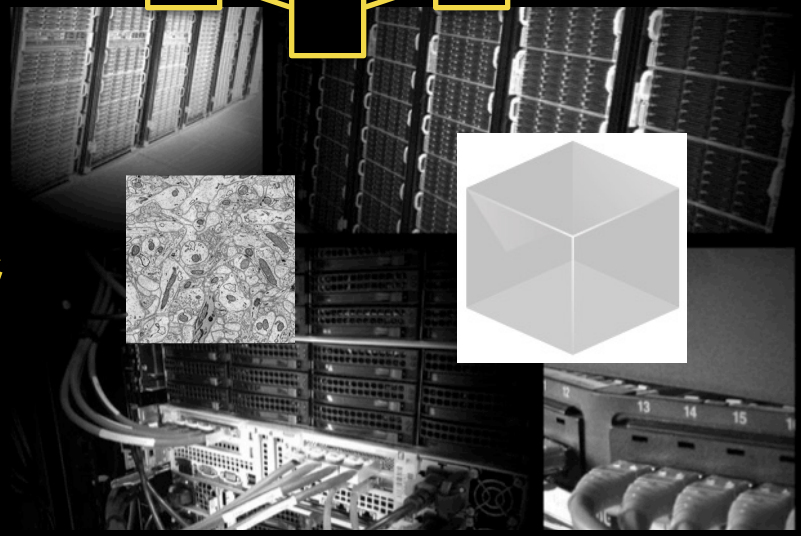

OPEN CONNECTOME PROJECT

COLLECTIVELY REVERSE-ENGINEERING THE BRAIN ONE SYNAPSE AT A TIME.

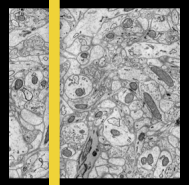

Disk Cache

- Local performance to remote data
- No computation
  - Tiles pre-rendered
- Visual flow
  - When scrolling back and forth through tiles

Deployment Options

- Tile cache collocated with server
  - Reduce I/O load on data servers
  - Offload rendering
- Tile cache in Amazon West, servers in Amazon east
  - All of above and content distribution
  - Reduce Internet latencies
- Tile cache on laptop/workstation with SSD
  - Maximize frame rates
  - Create user experience needed to visualize complex neural structures

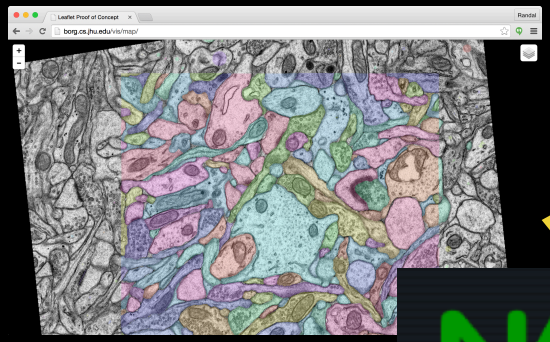

Why memcached?

- Background loading is not instantaneous
- Avoids server load
  - No computation for rendering
  - No I/O or NoSQL queries
- Consistent interfaces for dynamic data don't use tile cache

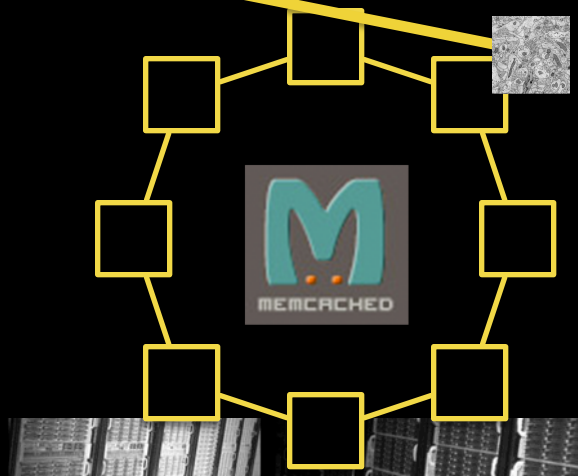

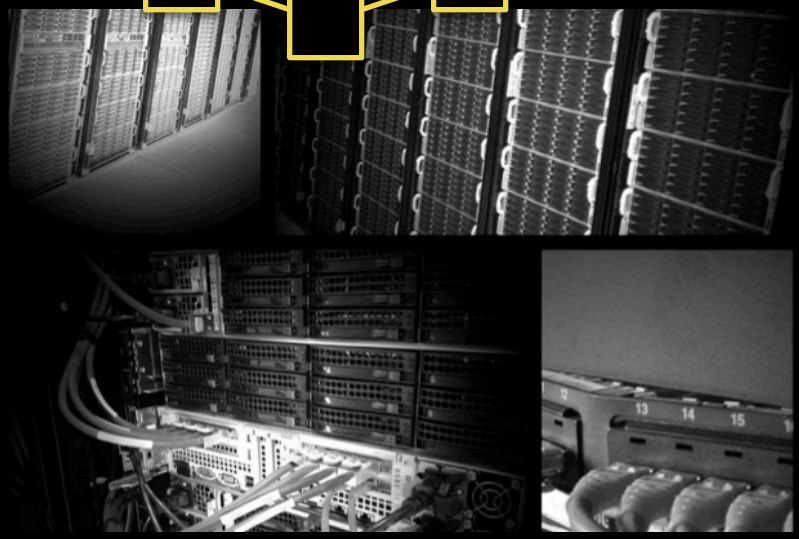

OPEN CONNECTOME PROJECT

COLLECTIVELY REVERSE-ENGINEERING THE BRAIN ONE SYNAPSE AT A TIME.

So What?

- Local performance to remote data
  - Eliminate Internet latency
  - Terabyte cache (on workstation) of petascale data
- User experience
  - Internet latency to first images
  - Local performance for most usage
  - Occasional stall for cache miss
- Open-source, tile caching for spatial data
  - https://github.com/openconnectome/tilecache
  - Not used outside of OCP managed installations today

OPEN CONNECTEAM

Alex Baden
Daniel Berger
Randal Burns
Davi Bock
Albert Cardona
Mark Chevillet
Kwanghun Chung
Ming Chuang
Forrest Collman
Steven Cook
Karl Deisseroth
Scott Emmons
Jeremy Freeman
Will Gray Roncal
Logan Grosenick
Greg Hager
Kristen Harris
Sean Hill
Bobby Kasthuri
Misha Kazhdan
Greg Kiar
Dean Kleissas
Kwame Kutten
Wei-Chung Allen Lee
Jeff Lichtman
Kunal Lillaney
Larry Lindsey
Priya Manavalan
Disa Mhembere
Michael Miller
Dan Naiman
Patrick Parker
Eric Perlman
Carey Priebe
Clay Reid
Stefan Saalfeld
Guillermo Sapira
Anish Simhal
Ayushi Sinha
Stephen Smith
Alexander Szalay
Raju Tomer
R. Jacob Vogelstein
Joshua Vogelstein
Nick Weiler
Li Ye
Da Zheng
thatweare

OPEN CONNECTOME PROJECT

COLLECTIVELY REVERSE-ENGINEERING THE BRAIN ONE SYNAPSE AT A TIME.


Questions?

- Website: neurodata.io
- Documentation : docs.neurodata.io
- Github: openconnectome
- CATMAID : openconnecto.me/catmaid/
- support@neurodata.io

Image Used for Demonstrational and Educational Purposes

- http://upload.wikimedia.org/wikipedia/commons/d/d2/Internet_map_1024.jpg
- http://broabandtrafficmanagement.blogspot.com/2011/08/resource-cdn-explained.html
- http://stopthecap.com/wp-content/uploads/2014/02/netflix-cdn.png