The brain: networks of neurons

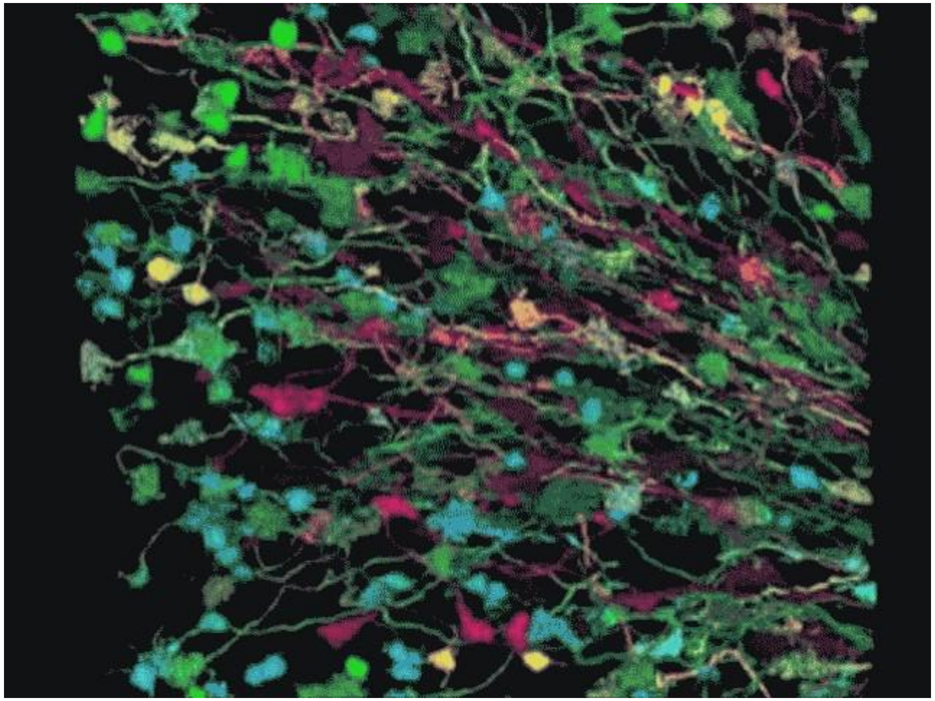

Mouse cerebellum **Lichtman et al. 2008**

The brain: networks of neurons

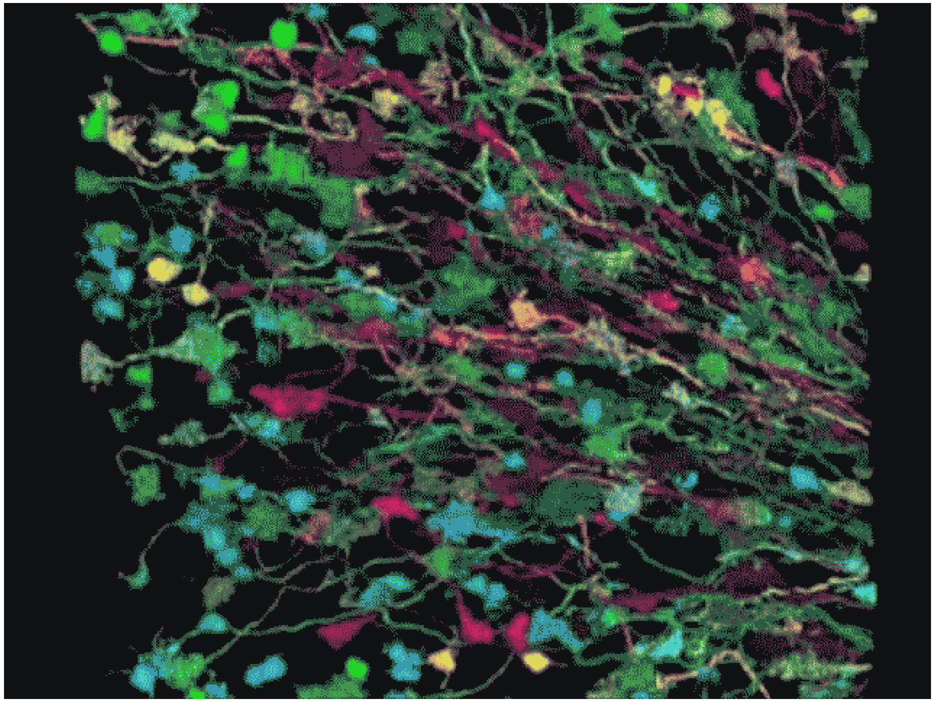

Mouse cerebellum **Lichtman et al. 2008**

1 mm^3
$\sim$1.5 year, 1 $petabyte$ (10^{15})

Mouse brain:
~2000 years, 1 $exabyte$ (10^{18})

Human cortex:
$\sim10^7$ years, 10 $zettabytes$ (10^{21})

World storage: 300 $exabytes$

~ on January 2015

Scalable Inference in Large Neural Systems

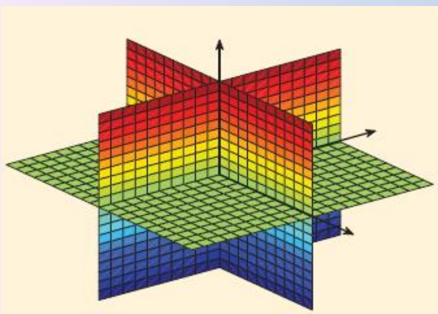

Allie Fletcher

Nov 5, 2015

UC Santa Cruz Electrical Engineering

UC Berkeley Redwood Center, HWNI

UCLA Statistics, Math, & EE– Jan 2016

Outline

- Neural mapping via multineuron excitation
  - Background: Neuroscience
  - Optogenetic stimulation
  - Estimation via graphical models
- Network connectivity via Ca^{2+} imaging
  - Large scale but indirect and smoothed
  - Network model captures dynamics
  - Scalable accurate algorithm
- Receptive Field of Retinal Ganglion Cells
  - Spatio-temporal filtering
  - Improved identification with limited data

Neuroscience: Estimation Challenges

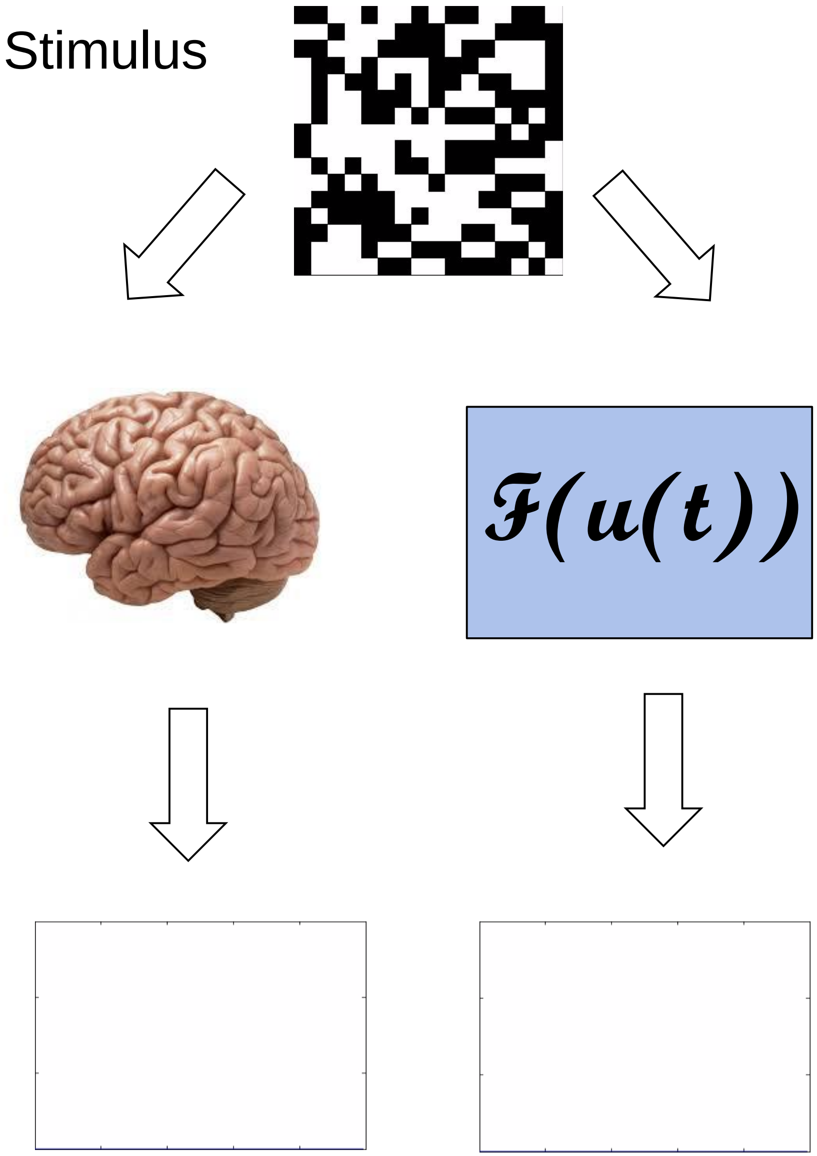

- Large-scale
  - ~86 billion neurons
  - V1 alone 140 million
- Sparse: ~1000 connections
- Complex:
  - Nonlinear dynamics
  - Feedback
- Indirect measurements
- Incomplete data
- Limited *in vivo* collection

I. Connectivity Detection to a Single Neuron

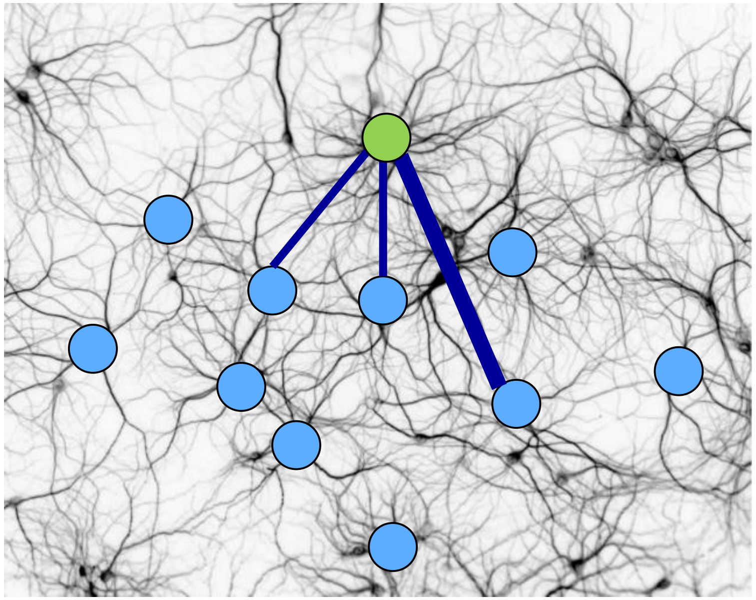

Problem:
Detect connections to one neuron

Goal:
Reduce trials and computation

How:
Optogenetics
Improved “decoding” via subset stimulation

Neuron: Basic Anatomy

Neuron: Basic cell for information processing

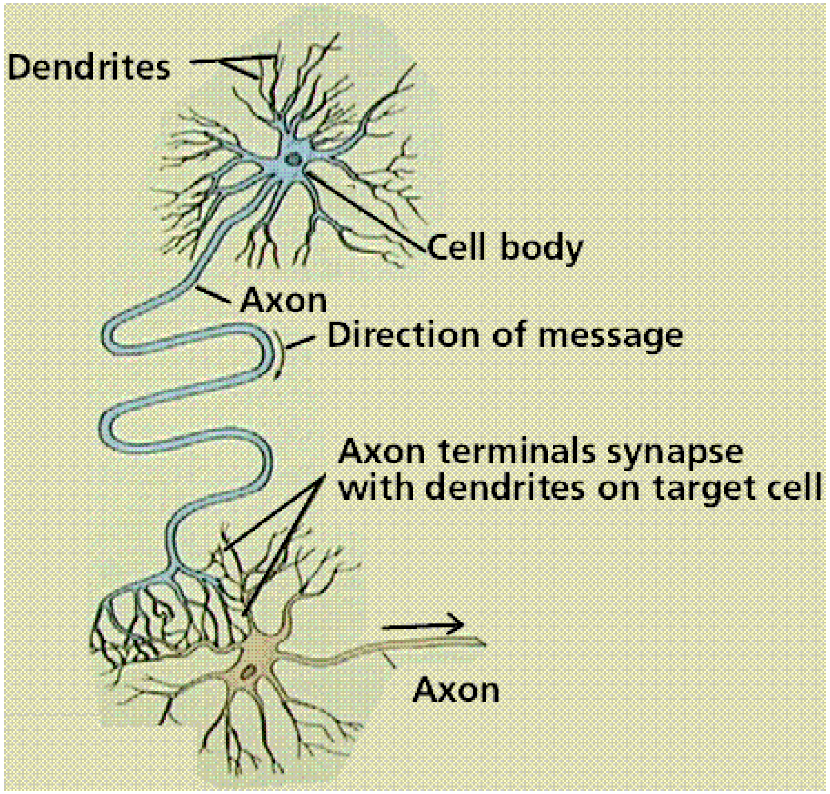

- **Components**
  - Dendrites:
    Filaments receive signals
  - Soma: cell body
  - Axon:
    Outputs electrical signals to neurons or motor functions
  - Synapses:
    Junctions of axons and dendrites

Neuron: Linear-Nonlinear Poisson Model

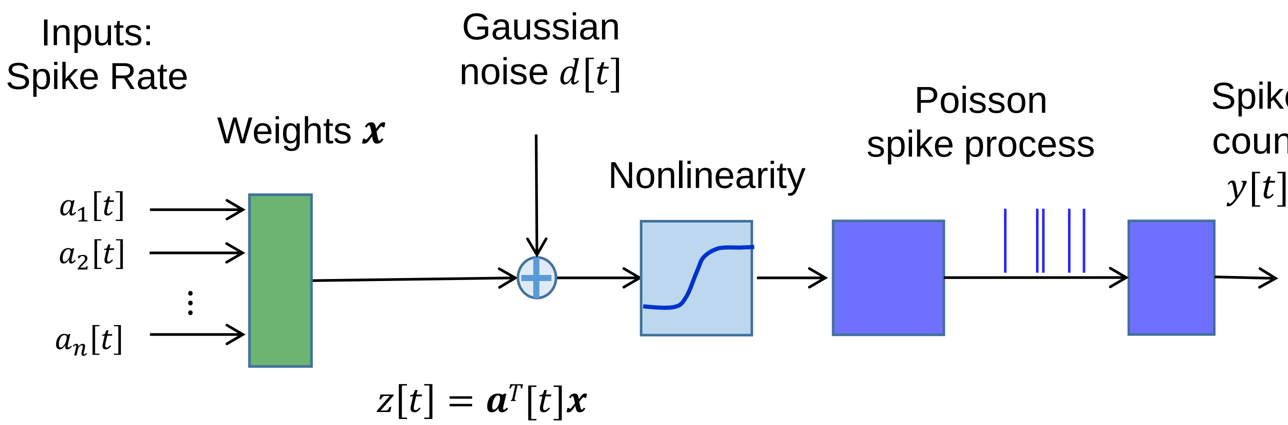

$$z[t] = \sum_j x_j a_j[t]$$

$$\lambda[t] = f(z[t] + d[t])$$

$$y[t] \sim Poisson(\lambda[t])$$

- Functional Model: simple input-output
  - Effectively captures average rates
  - Windows: 10 to 200 spikes/s
- Three-stage LNP: Linear + Nonlinearity + Poisson process
- Connectivity: Identify weights $\boldsymbol{x}$
- Biological models with feedback for precise timing later…

Classic Connectivity Detection

Measure at post-synaptic neuron of interest

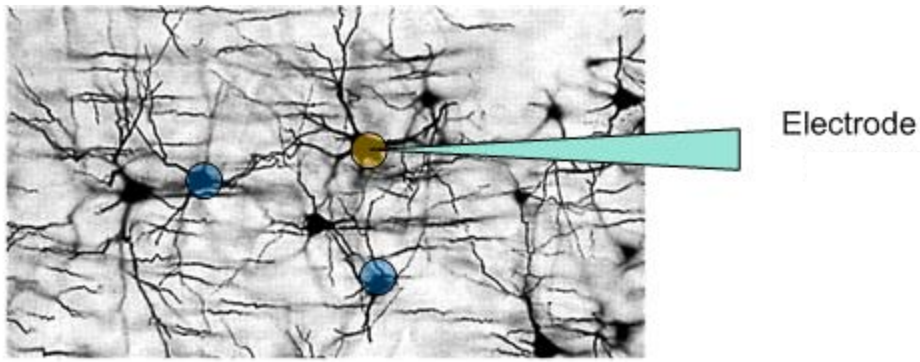

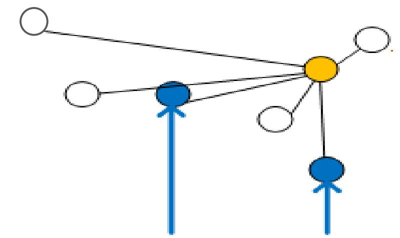

Stimulate potential pre-synaptic neurons: one at a time

- Problems: many measurements per test neuron
  - Most neurons are not connected
  - Noisy system with many exogenous inputs
  - Test neurons die
- Can we do better?

New Technologies: Optical Stimulation

Connectivity Detection:

- Genetically modified neuron
  - Photosensitive protein
- Optically activate test neurons
- Greater spatial precision
  - Pinpoint individual neurons
  - Multiple neurons at once

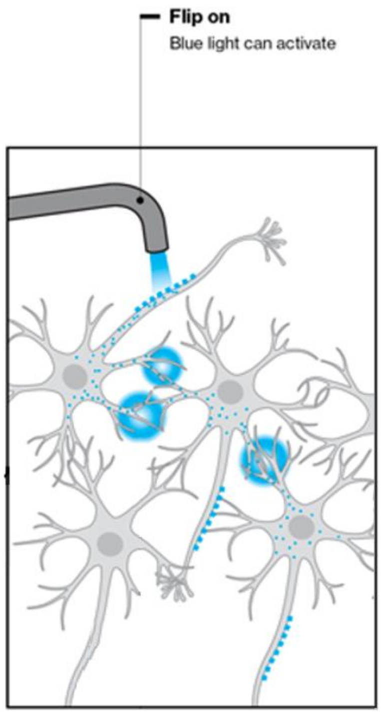

Wang, Hasan & Seung 2009

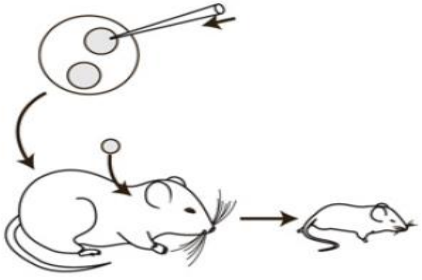

Transgenic mice

Connectivity via Multineuron Excitation

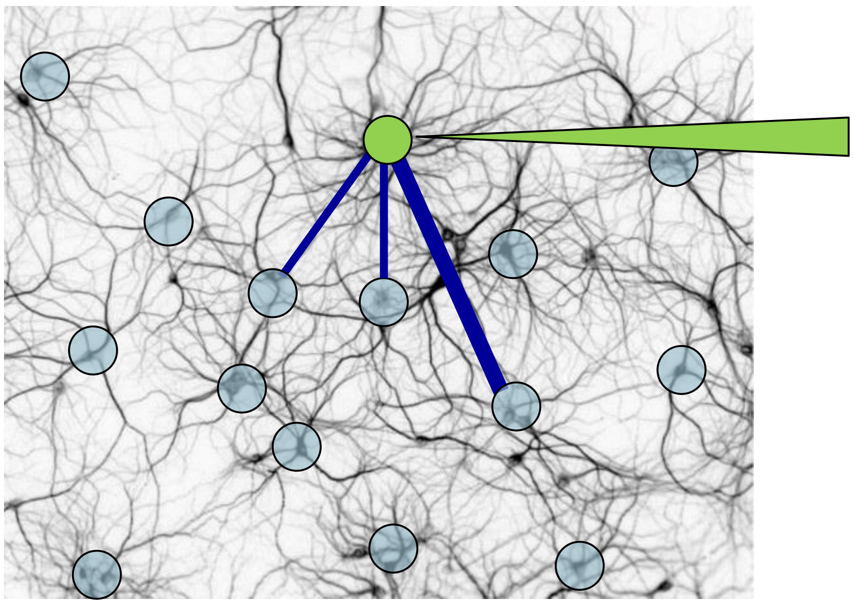

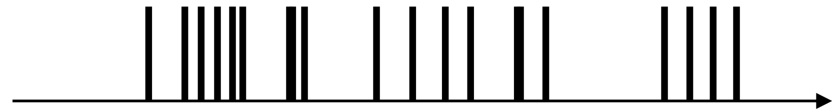

Spike recording
Average over time windows

Potential presynaptic neurons

Synapses or connections

Detection via Multineuron Excitation

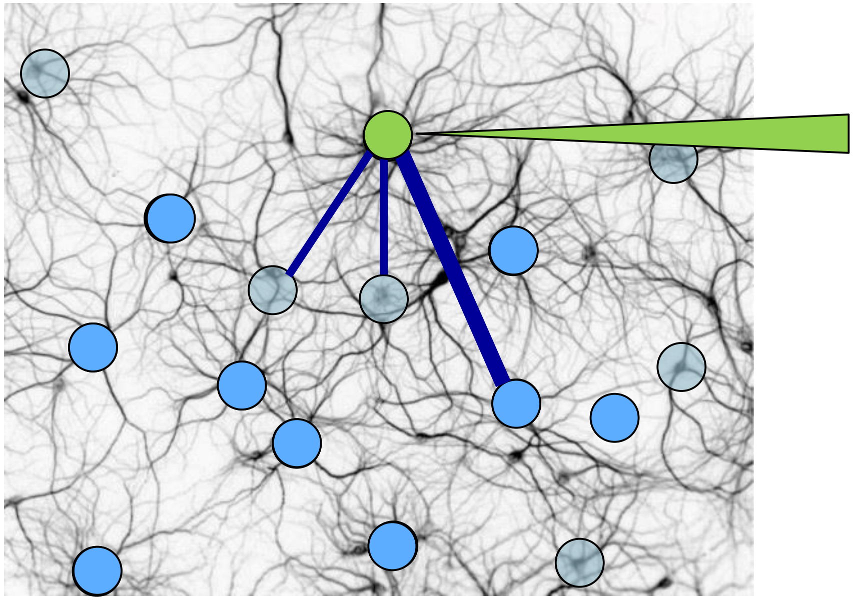

Spike recording

Strong response

No connection

Connections

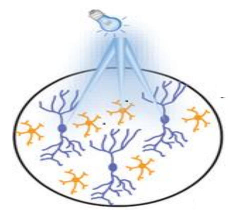

- Stimulate subsets of neurons at a time
  - Increase probability of response
  - Fewer wasted trials

Detection via Multi-Neuron Stimulation

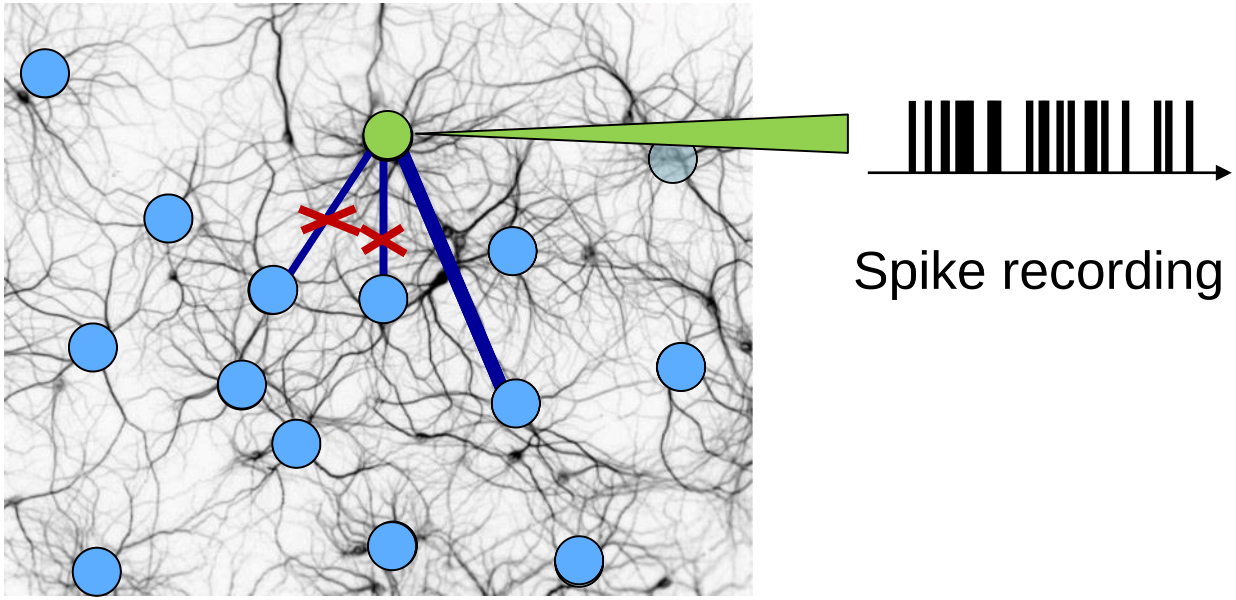

- Stimulate subsets of neurons at a time
  - Increase probability of a response
  - Less measurements wholly wasted
- Benefits
  - Weak sub-threshold connection
  - More reliable with less data

Hu & Chklovskii NIPS 2010, Fletcher et al NIPS 2011

Multi-neuron Excitation: Model

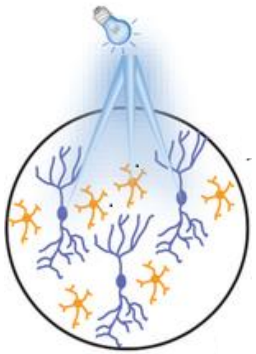

Potential neurons: 1:N

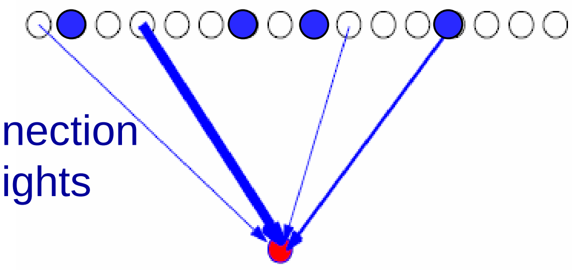

Connection weights

Measurement neuron

j^{th} neuron

Excited

Not excited

Excited: {2, 7, 9, 13}

“1” entries: {2, 7, 9, 13}

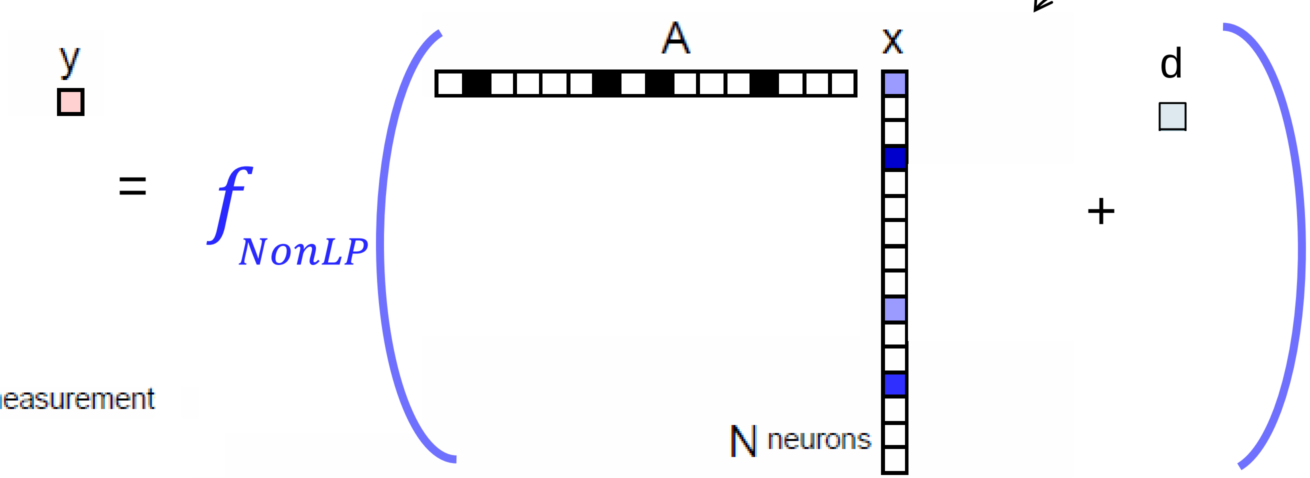

j^{th} element

1

0

Multi-neuron Excitation: Model

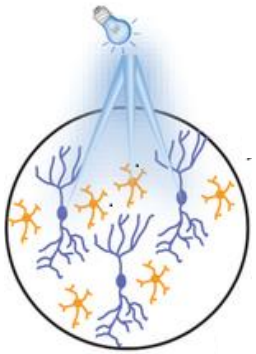

Potential neurons: 1:N

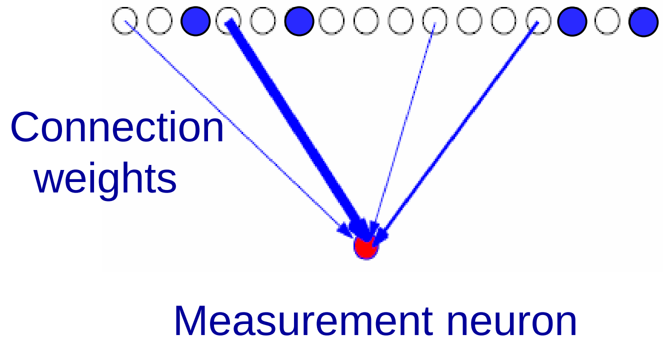

j^{th} neuron

Excited

Not excited

Second trial

Different random subset

Excited: {3, 6, 14, 16}

“1” entries: {3, 6,14,16}

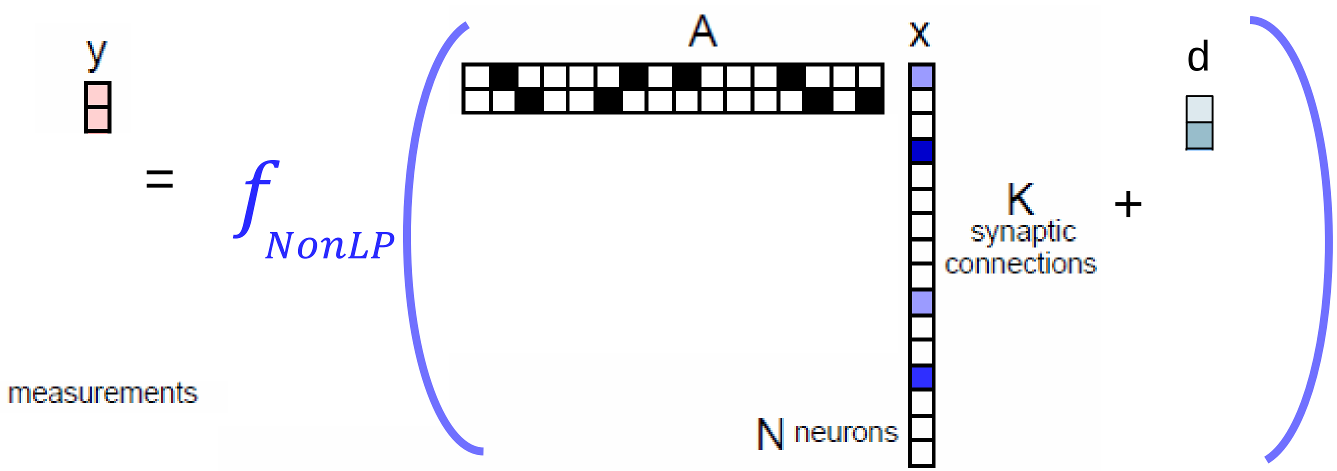

j^{th} element

1

0

Multi-neuron Excitation: Model

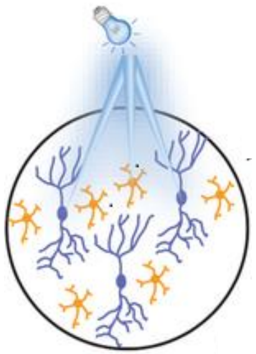

Potential neurons: 1:N

Connection weights

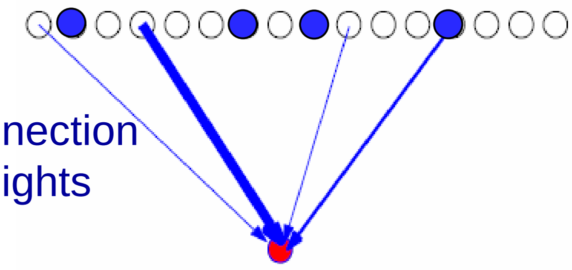

Measurement neuron

j^{th} neuron

Excited

Not excited

m^{th} excitation

m trials or “measurements”

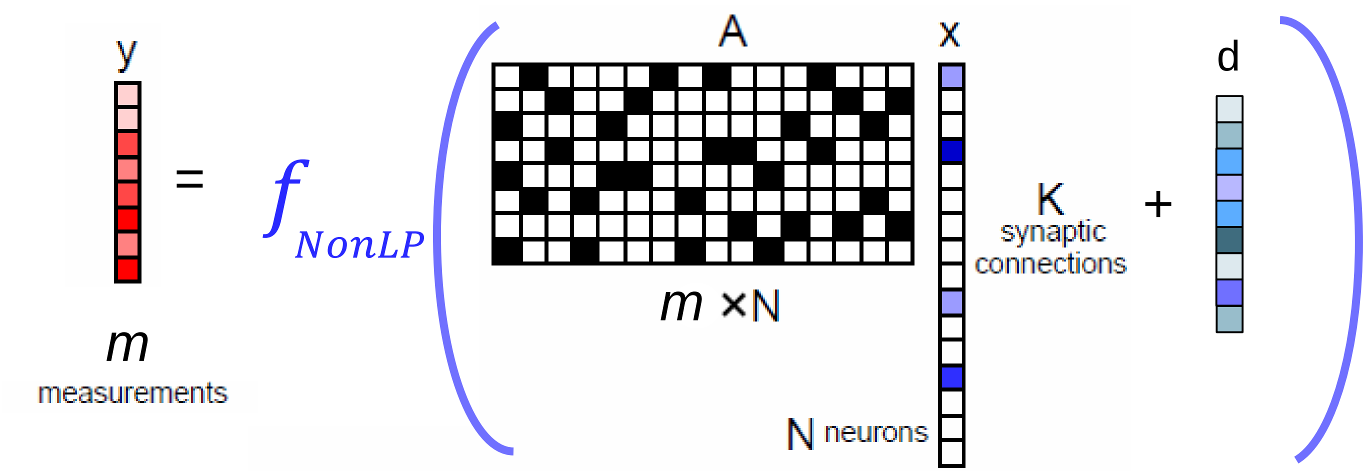

j^{th} element

1

0

Bayesian Nonlinear Generalized Linear Models

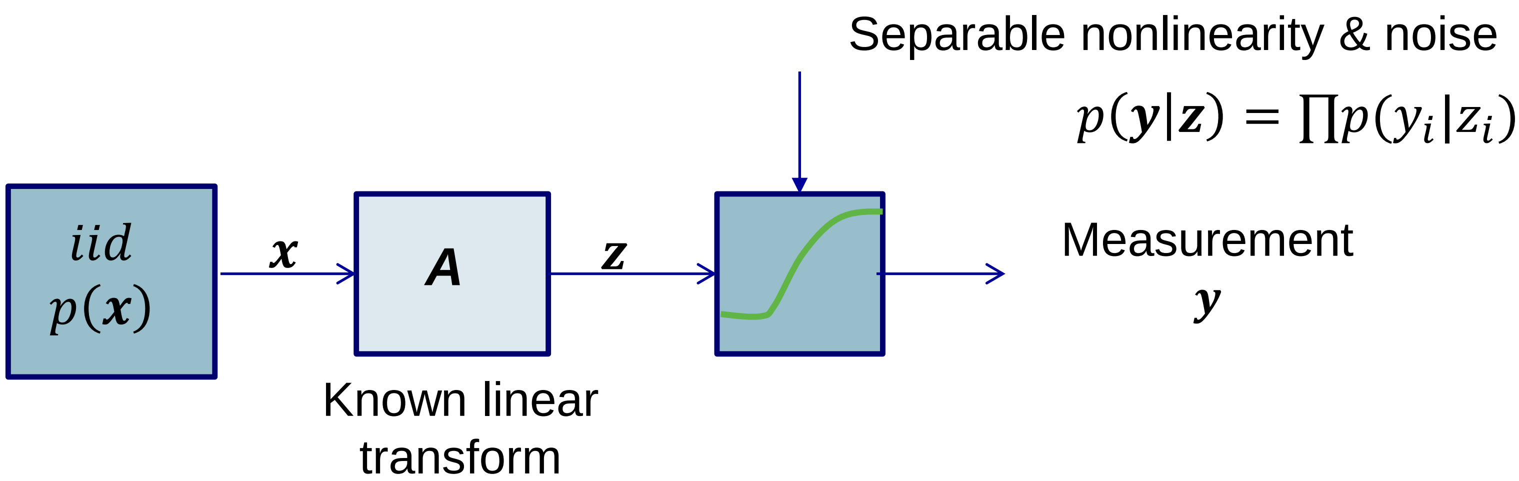

- Problem: Estimate $\boldsymbol{x}$ and $\boldsymbol{z}$ given $\boldsymbol{y}$ and $\boldsymbol{A}$
- Bayesian formulation: general system class
  - Prior $P_X(\boldsymbol{x})$ incorporates constraints, like sparsity
  - $P_{Y|Z}(\boldsymbol{y}|\boldsymbol{z})$ models output: nonlinearities, randomness
- Challenge: optimal estimation is hard
  - Components of vector **x** are coupled in **z**

Example: Sparse recovery

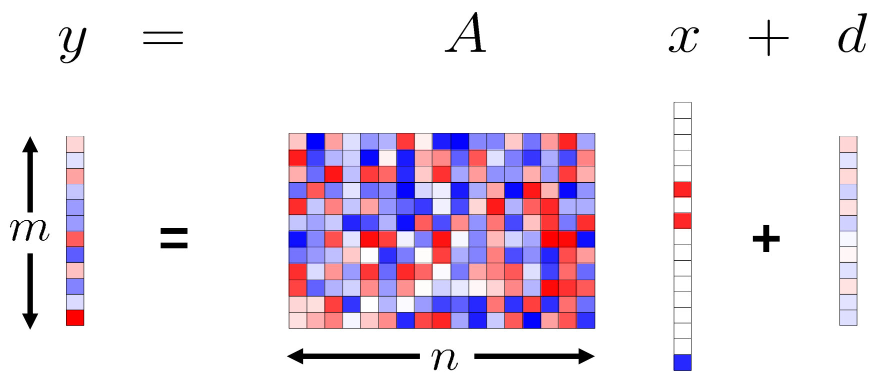

$p(\boldsymbol{x})$ **= Sparse prior**
$p(\boldsymbol{y}|\boldsymbol{z})$ **= Gaussian**

- Problem: Given **A** and **y**, recover sparse **x**
- Many applications
  - Communication channels, linear inverse problems
  - Wavelet image reconstruction
  - Regularized linear regression, classification
  - Compressed sensing…
- Now many algorithms, theoretical analyses, …

Divide & Conquer with Graphical Models

- Subdivide & conquer

$$p_{X|Y}(\boldsymbol{x}|\boldsymbol{y}) = \prod_{j=1}^{n} p(x_j) \prod_{i=1}^{m} p(y_i|\boldsymbol{x})$$

- Few "smaller" components
  - Few variables
  - Limited connections
- Message passing:
  - Iteratively update marginals
  - Global estimation local
- But random A is dense!

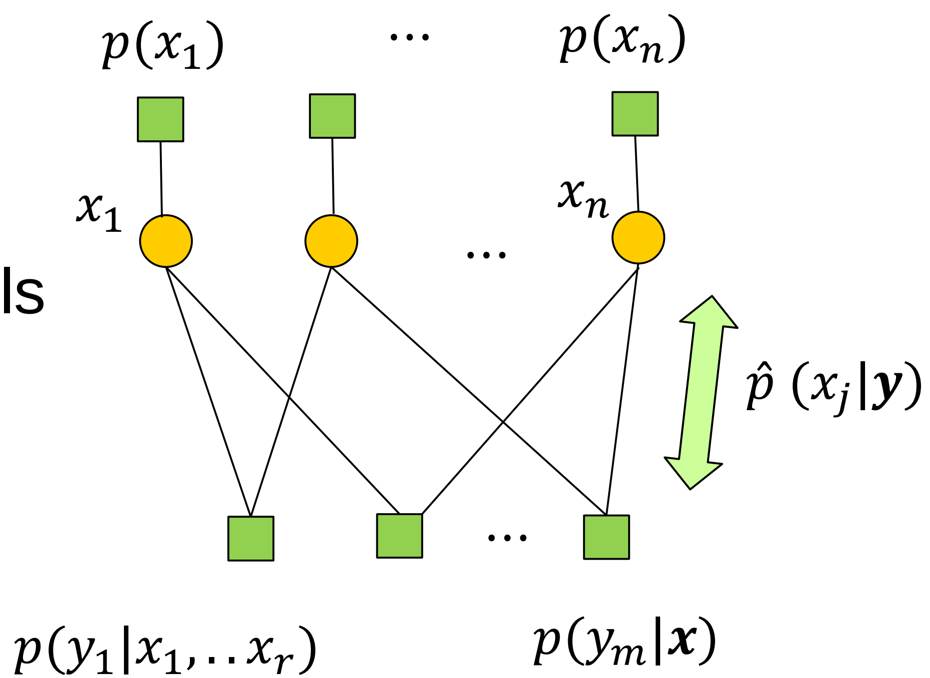

"Graphical Model" for GLM

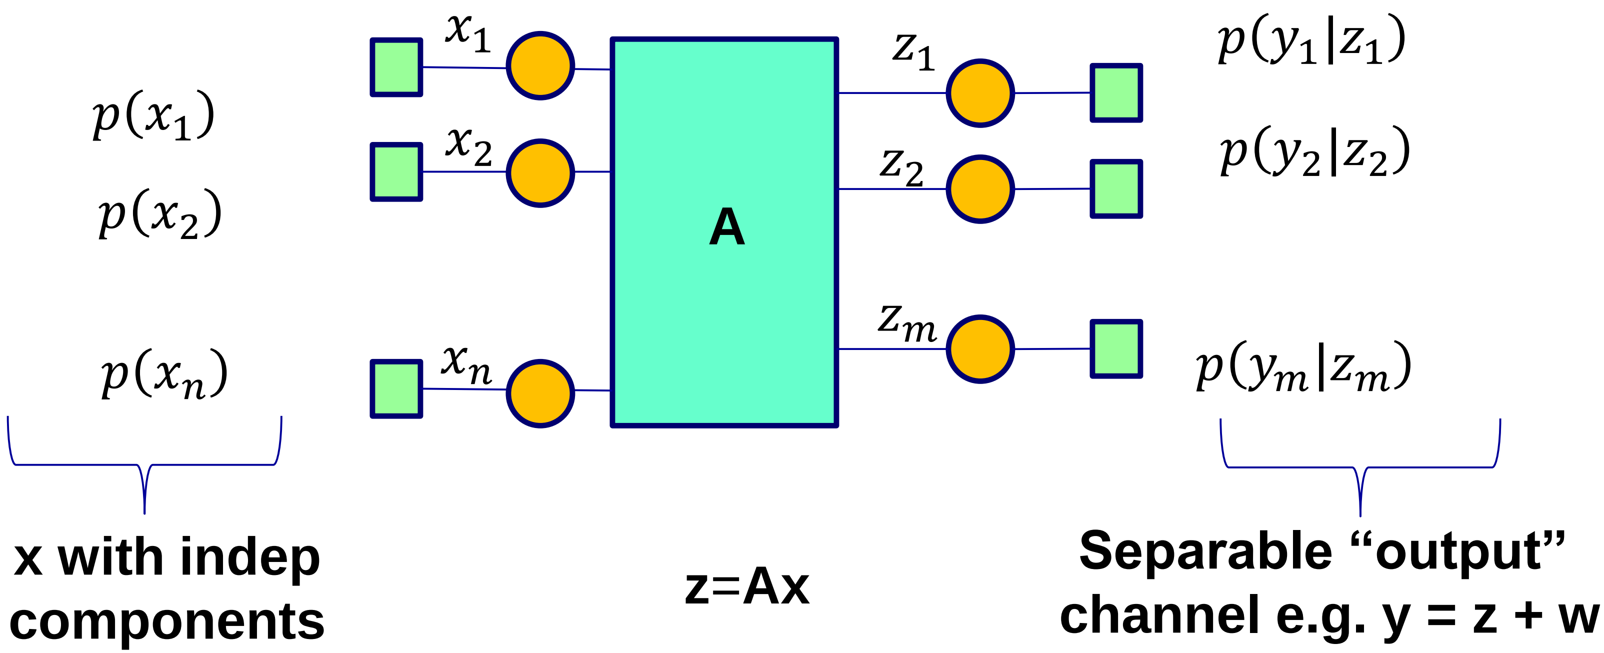

- Assume separable priors & likelihoods
- The posterior density $p(x|y)$ factors into:
  - $m + n$ scalar terms; and
  - Linear constraint $z = Ax$

Generalized Approximate Message Passing (GAMP)

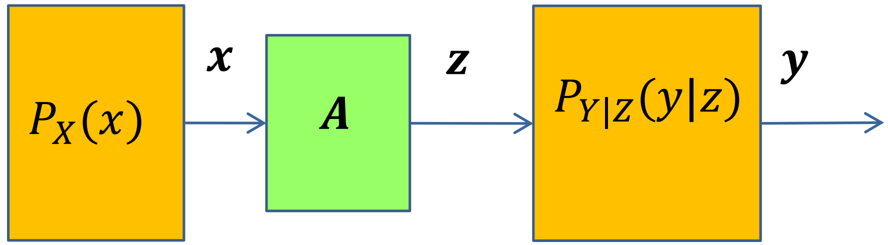

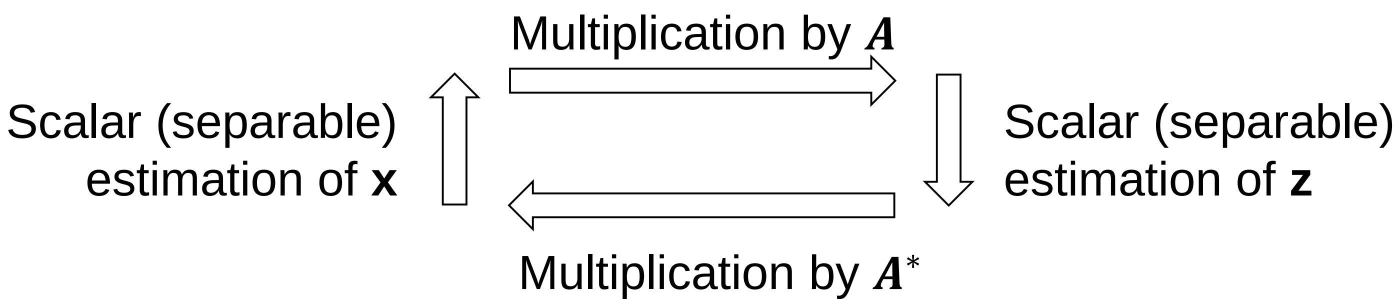

- Gaussian & quadratic approximations
- Asymptotic guarantees
- Low complexity: $O(mn)$ each iteration
- Classic AMP*: separable distributions & AWGN
- GAMP: KNOWN nonlinearities

AMP Donoho, Maleki, Montanari 09, Bayati & Montanari 10, GAMP Rangan et al 10, HyGAMP Fletcher et al 11, …

Theorem: Joint Estimation & Learning with Adaptive GAMP

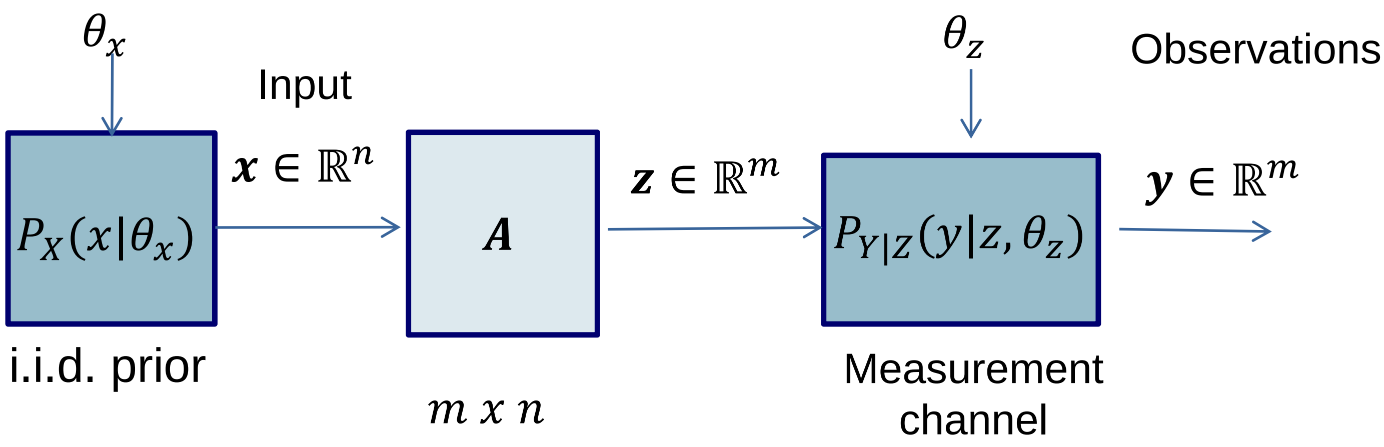

- GLM with unknown parameters θ_x and θ_z
  - Unknown prior, nonlinearities, noise…
- Joint estimation learning problem: Given y and A:
  - Estimate input x and z,
  - Learn parameters θ_x and θ_z in distribution
  - **Consistent estimator**

Kamilov, Fletcher, Rangan, Unser NIPS 2012, Trans IT 2014

Neural Mapping via Adaptive-GAMP: NeuRAMP

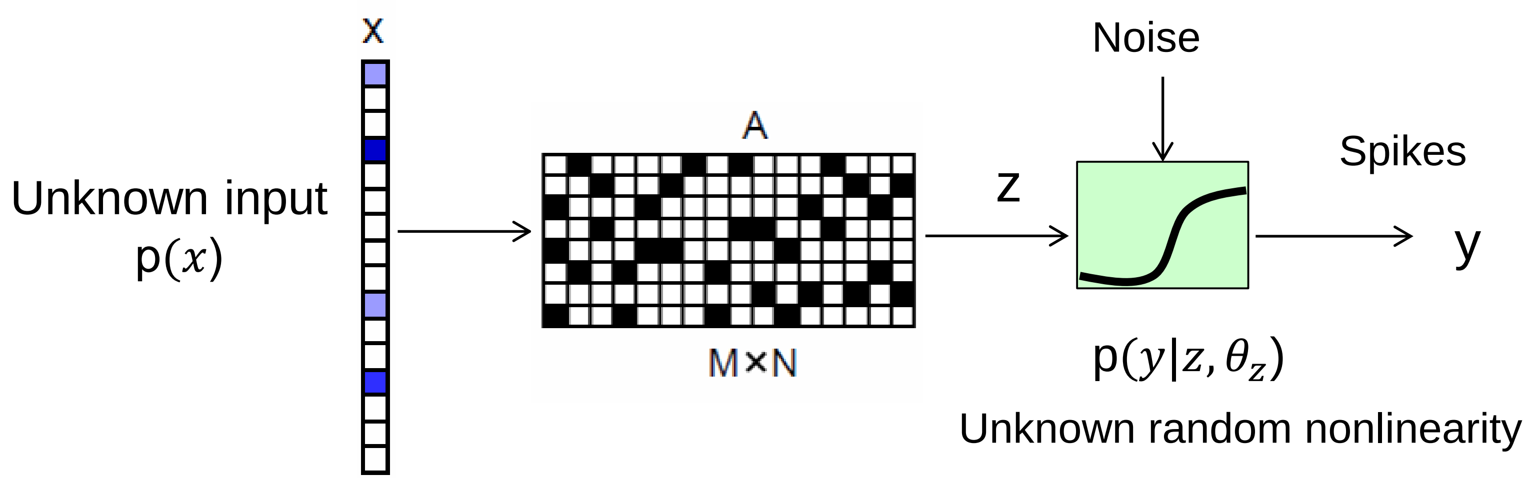

- Problem: For neural LNP model:
  - Incorporates sparsity on prior $P_X(\boldsymbol{x})$
  - $P_{Y|Z}(\boldsymbol{y}|\boldsymbol{z},\theta_z)$ models unknown output nonlinearities
- Jointly: Estimate weight vector **x**

  Learn the nonlinearity
- Computationally fast
- Improved estimates with fewer measurements

Fletcher et al. NIPS (2011)

Simulation: MSE of Synaptic Weight Estimates

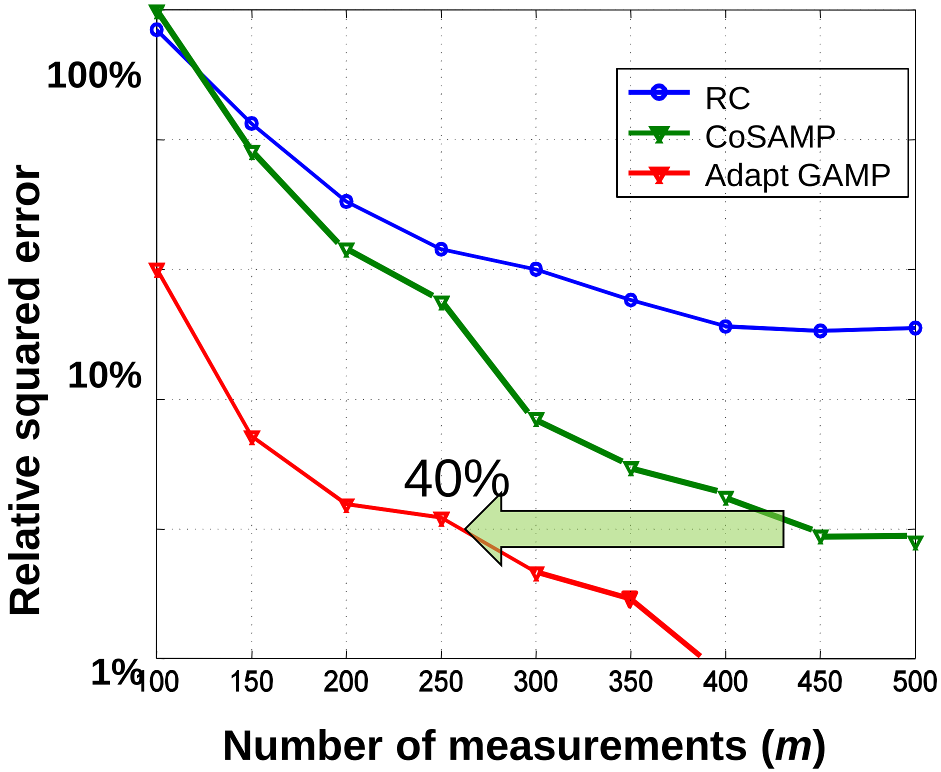

- RC: Reverse Correlation
  - Linear estimation
  - No sparsity
- *CoSaMP: Greedy CS
  - Ignores nonlinearities
- **Adaptive GAMP:
  - Lower MSE
  - Fewer measurements

m = number of trials, random excitation ~40
n = 500,
k =30 or 6% Bernoulli-Gayssuab weights
100 spike windows: 10.4s
Spike rate: 10 spikes/s with.4 second reset
300 trials: ~1 hour

***Hu & Chlovskii NIPS 2010**
****Fletcher et al, NIPS 2011**

Simulation: NeuRAMP Connectivity Detection

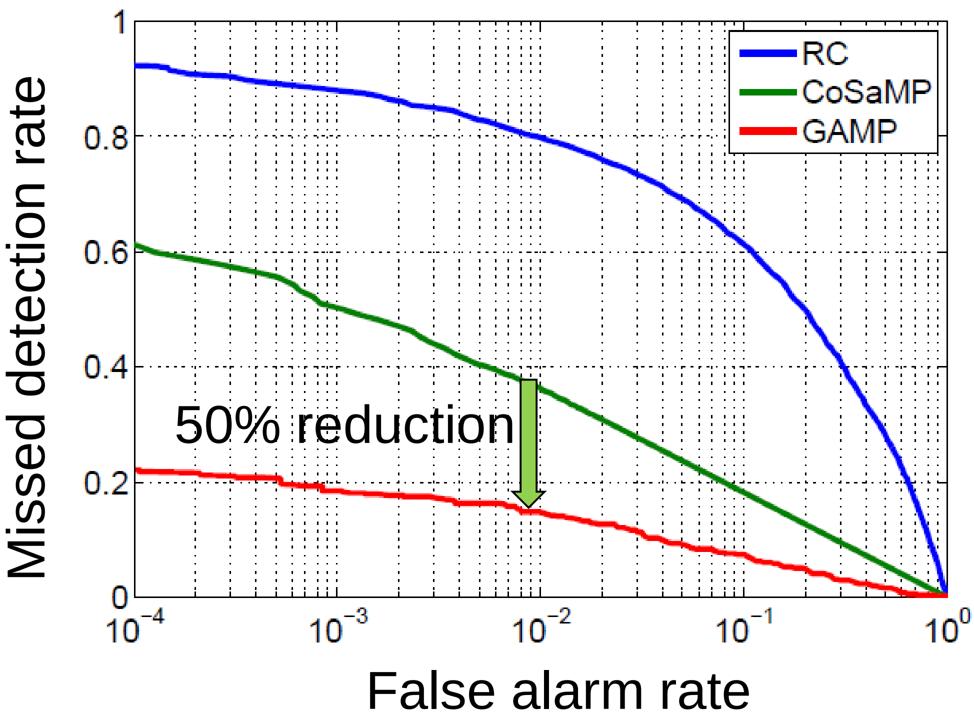

- Factors:
  - Models nonlinearities
  - Optimal learning
  - Incorporates sparsity
  - Similar complexity
- GAMP outperforms RC & CoSaMP
  - 75% and 50% lower missed detects

"Simultaneous all optical manipulation and recording…", AM Packer, Dec 2014 Nature

****"Block sparse" filters: Salamander visual receptive field**

Moving forward: Scalable adaptive block

- Generalized Approximate Message Passing (GAMP)
- Improved neural connectivity detection
  - Data limited
  - Unknown nonlinearities, sparsity levels
- Scalable adaptive GAMP block
  - Linear mixing blocks
  - Low complexity: scalable
  - Extensible
  - Rich input output models
  - Adaptive

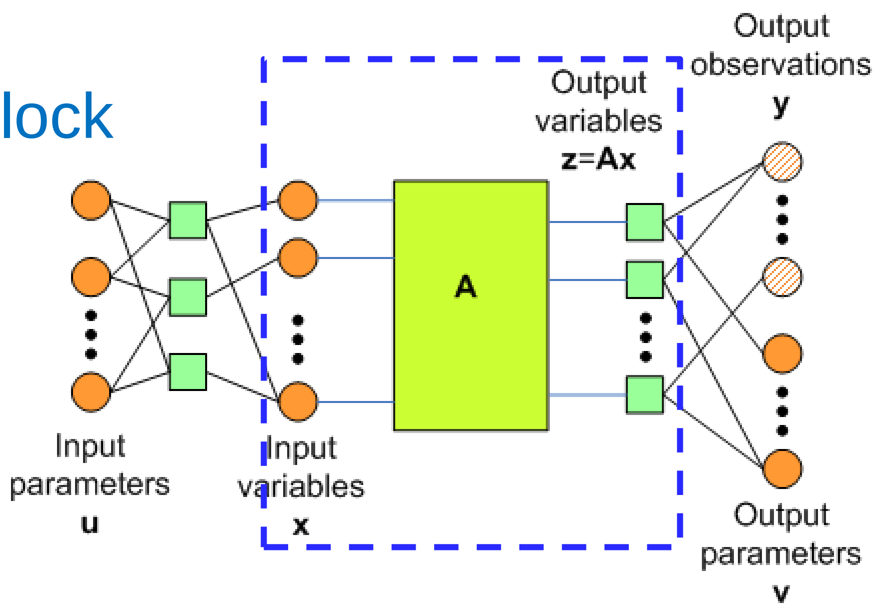

Outline

- Connectivity via multineuronal stimulation
  - Iterative fast adaptive GAMP framework
- Network connectivity via Ca^{2+} imaging
  - Large scale *in vivo* layers
  - Remarkable spatial resolution
  - Indirect, temporally smoothed
  - Network model - captures dynamics
  - Scalable accurate EM algorithm
- Receptive field of retinal ganglion cells
  - Space-time salamander response to stimuli
  - Improved identification with limited data

Inference of Network Connectivity

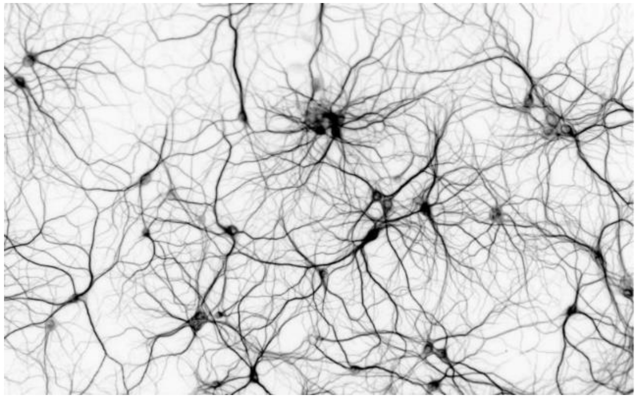

Network of neurons:
unknown synaptic connectivity

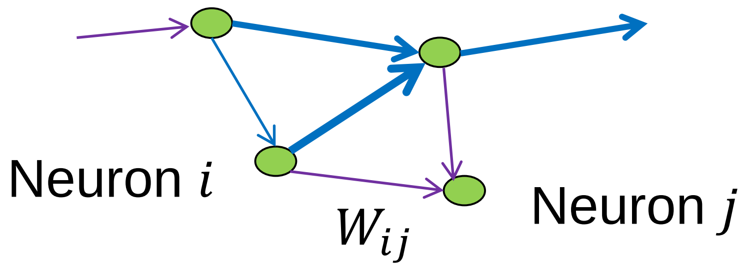

Estimated connectivity map
with synaptic weights W_{ij}

New Technologies: Calcium Imaging

- Fluorescent Ca^{2+} indicators
  - Genetically encoded
  - Chemical dyes
- Spiking: Ca^{2+} influx
- Large populations in parallel

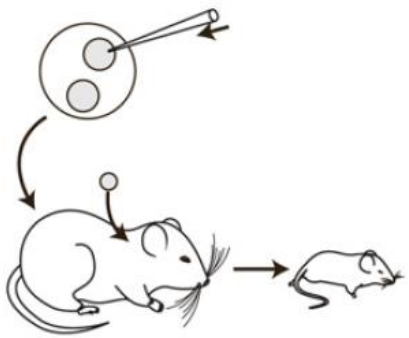

Transgenic mice

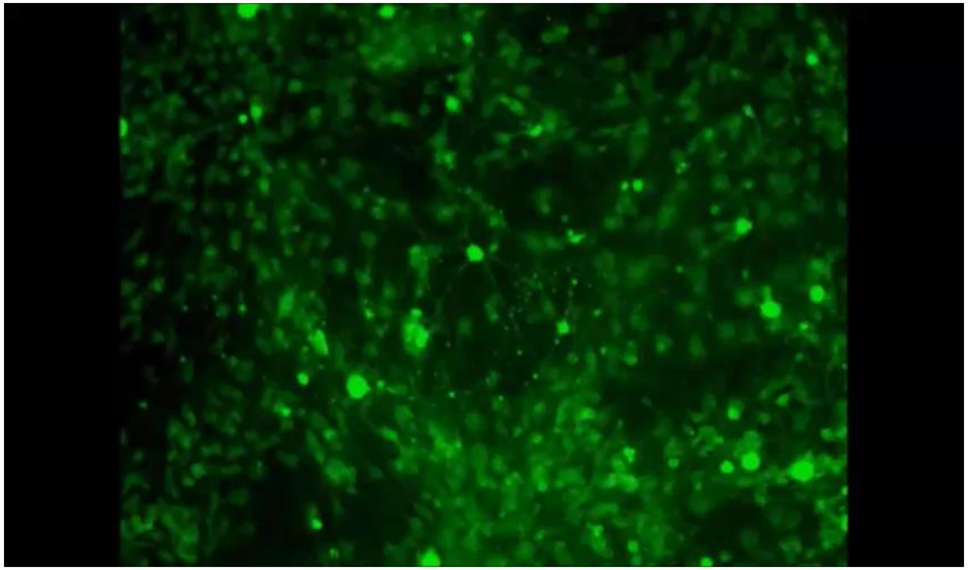

Mei Zhang 2009

Two photon Ca^{2+} imaging: depth acquisition

- Fluorescence & spiking:
  - Ca^{2+} influx
- Two-photon imaging
  - Raster scan
  - Depth acquisition
- Spatial resolution
  - Image into the cortex

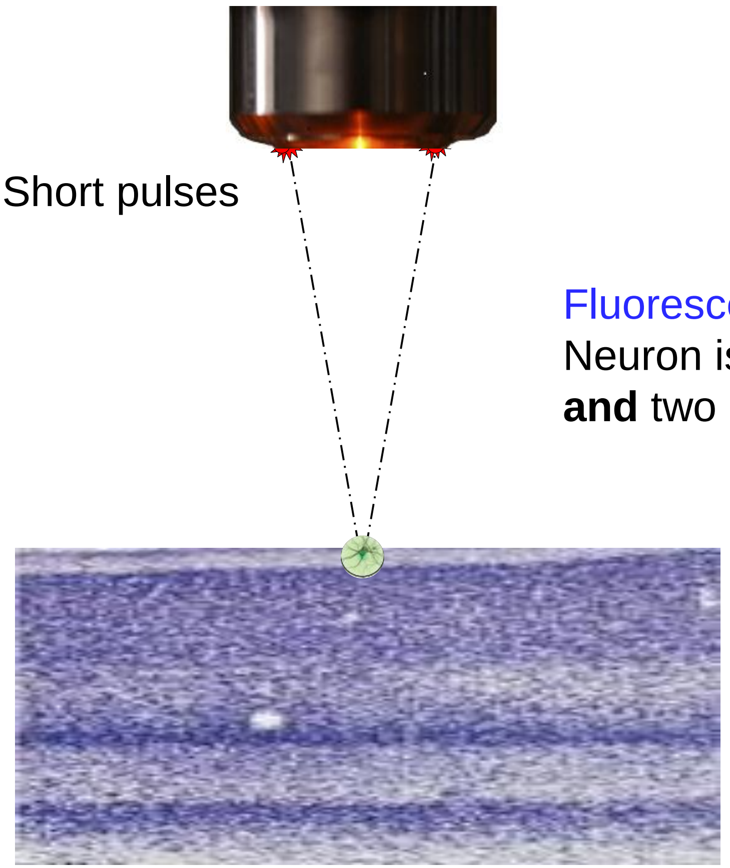

Two photon Ca^{2+} imaging: depth acquisition

- Fluorescence & spiking:
  - Ca^{2+} influx
- Two-photon imaging
  - Raster scan
  - Depth acquisition
- Spatial resolution
  - Image into the cortex

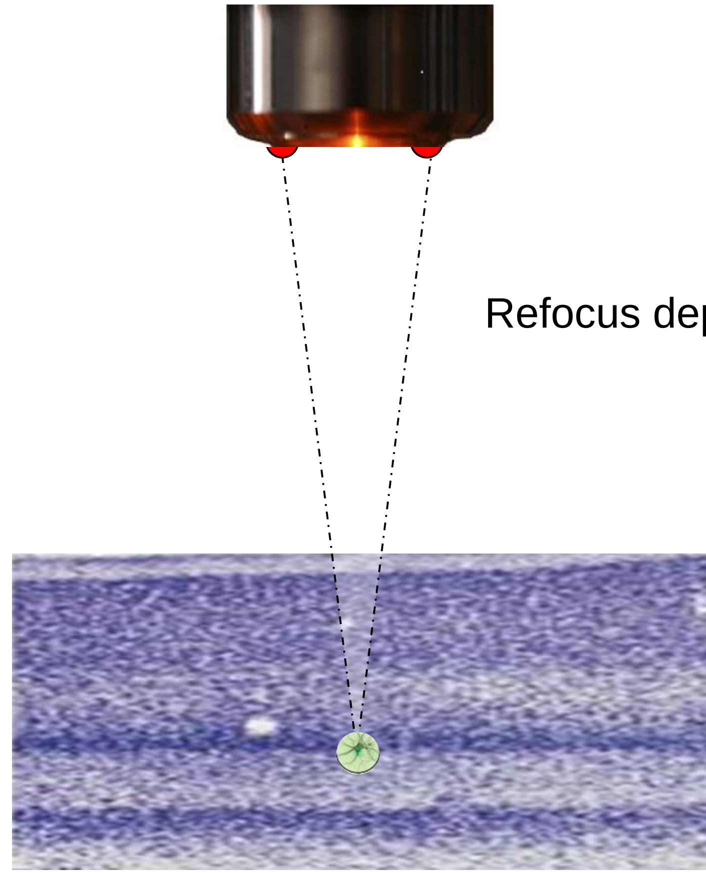

Refocus depth

Cortical layers

Cortex ~ 2-4 mm

Calcium Imaging: Connectivity Detection Problem

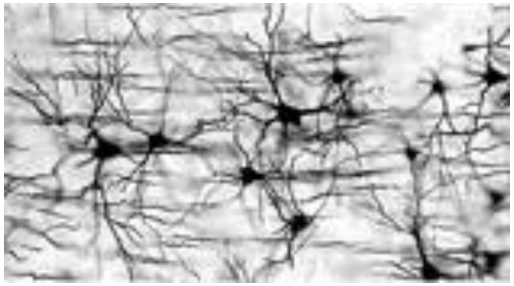

Network of neurons:
unknown synaptic connectivity

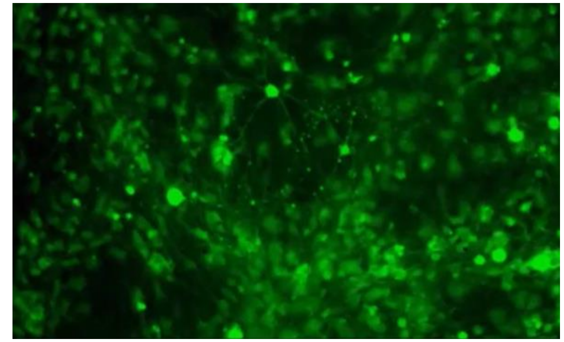

Ca^{2+} fluorescence movie

Ca^{2+} fluorescence image
MPI *2012*

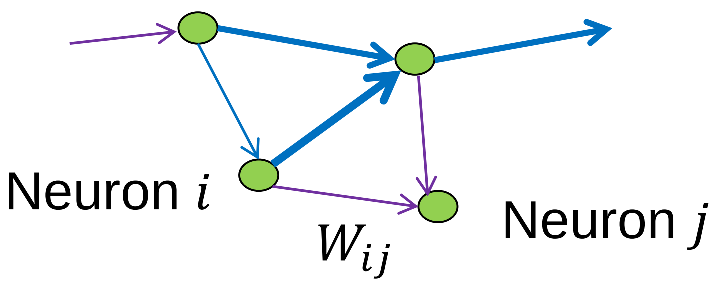

Estimated connectivity map
with synaptic weights W_{ij}

Ca^{2+} Imaging: Strengths

- Parallel measurements ($\sim 10^3$ neurons)
- *In vivo* or *in vitro*
- High spatial resolution (neuronal level, sub-μm)
- Image below surface

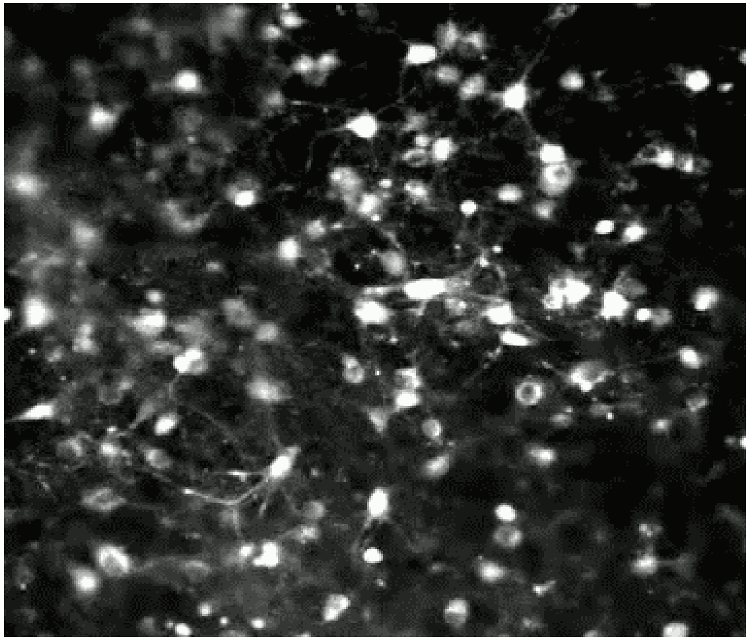

M. Kuykendal and G. Guvanasen, Georgia Tech

Challenges with Ca^{2+} Imaging

Calcium Fluorescence

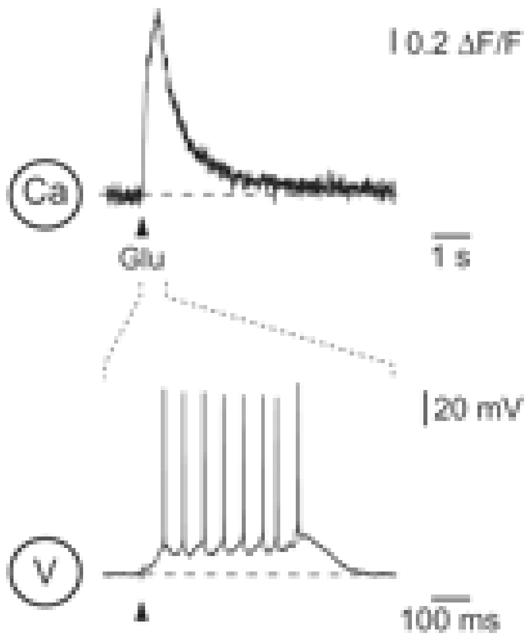

[Stociek et al. (2003)]

Glutamate induced spiking

- Indirect fluorescence traces
- Nonlinear dynamics
- Large data sets
- Exogenous inputs
- Heavy temporal blurring
  - Ca^{2+} long decay: τ_c~0.5 s **
- Low frame rate
  - Frames: 10-100 ms*
  - Interneuron dynamics: 1-3 ms
- Need super-resolution

**GCaMP5, newer indicators faster

NIPS 2013 & NIPS 2015 workshops:
"Statistical Methods For Understanding Neural Systems"

Network Inference: Causality Crucial

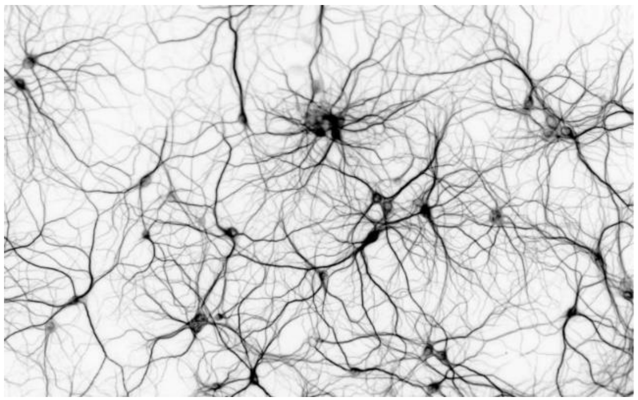

Network of neurons:
unknown synaptic connectivity

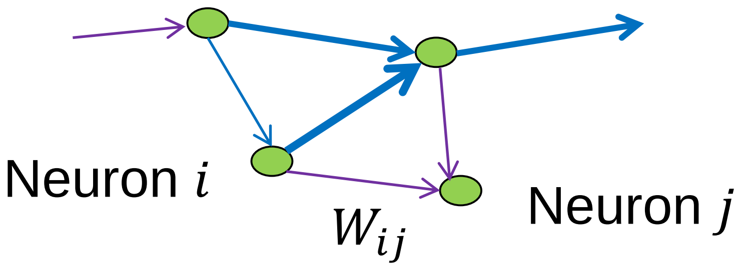

Estimated connectivity map
with synaptic weights W_{ij}

Elements of the network : time-varying electrochemical devices

Neuronal Model: Integrate and Fire

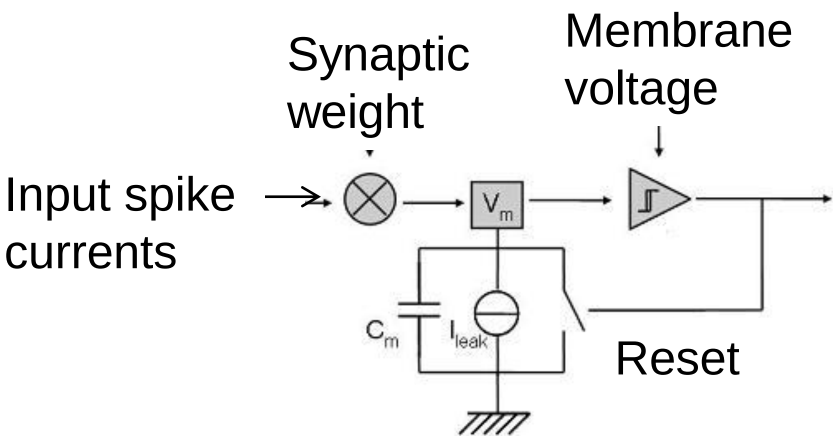

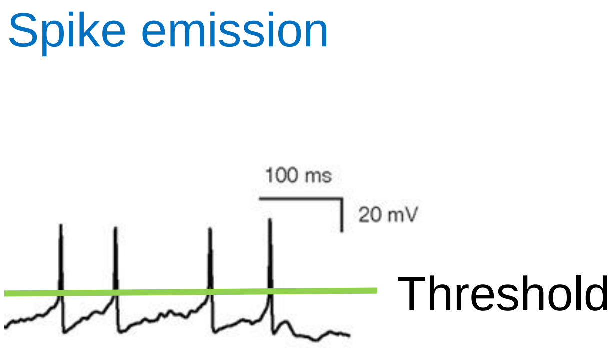

- Electrochemical dynamic model:
- $V_i(t)$ = neuron i potential $I_{ij}(t)$ = current j^{th} to i^{th} neuron
- Integrate phase: Potential builds: $V_i(t) \leq V_{th}$

$$\underbrace{RC_m \frac{dV_i(t)}{dt}}_{\text{Charge increase}} = \underbrace{-V(t)}_{\text{leakage}} + \underbrace{\sum_{j=1}^{N} R_{ij} I_{ij}(t)}_{\text{Incoming current}}$$

- Fire: discharges spike, reset: $V_i(t) = V_{th} \Rightarrow V_i(t^+) = V_{reset}$

[Lapicque (1907)]

Calcium Imaging: Connectivity Detection Problem

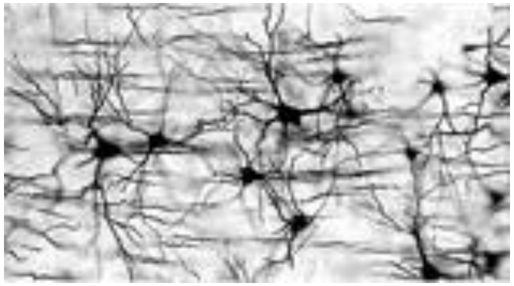

Network of neurons:
unknown synaptic connectivity

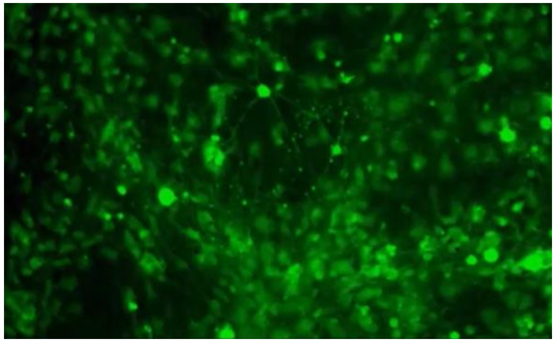

Ca^{2+} fluorescence movie

Ca^{2+} fluorescence image
MPI *2012*

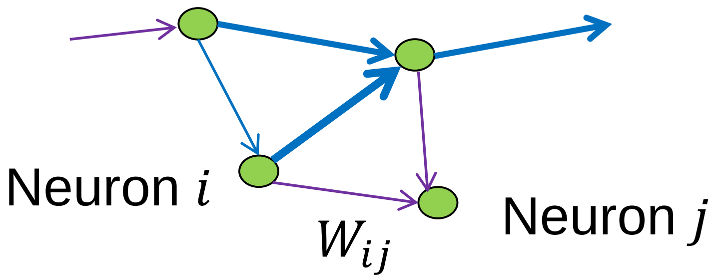

Estimated connectivity map
with synaptic weights W_{ij}

Each Neuron: Discrete-Time Neural Model

- Voltage: integrate and fire:

$$v_i^{k+1} = (1-\alpha)v_i^k + \sum_{j=1}^n W_{ij}\, s_j^k + d_{v,i}^k \qquad \text{[Integrate]}$$

$$\text{if} \quad v_i^{k+1} \geq \mu \Rightarrow \; s_i^k = 1,\, v_i^{k+1} = 0, \qquad \text{[spike \& reset]}$$

$$\text{else } v_i^{k+1} < \mu \Rightarrow \; s_i^k = 0 \qquad \text{[no spike]}$$

Calcium fluorescence

$$z_i^{k+1} = (1-\beta)z_i^k + s_i^k + d_{z,i}^k \qquad \text{[z: Ca}^{2+}\text{]}$$

$$y_i^k = a\, z_i^k + d_{y,i}^k \qquad \text{[y: fluorescence]}$$

W_{ij} = "weight" = integrated voltage change from spike current

- Ca^{2+} fluorescence: dynamical system also
- Nonlinear state space
- Need: connectivity, spike times, voltages, calcium…

Mischenko, Vogelstein, Paninski (2010), Yasuda(2004), Vogelstein et al (2010)
Fletcher et al, COSYNE 2014, NIPS 2014

Summary: System

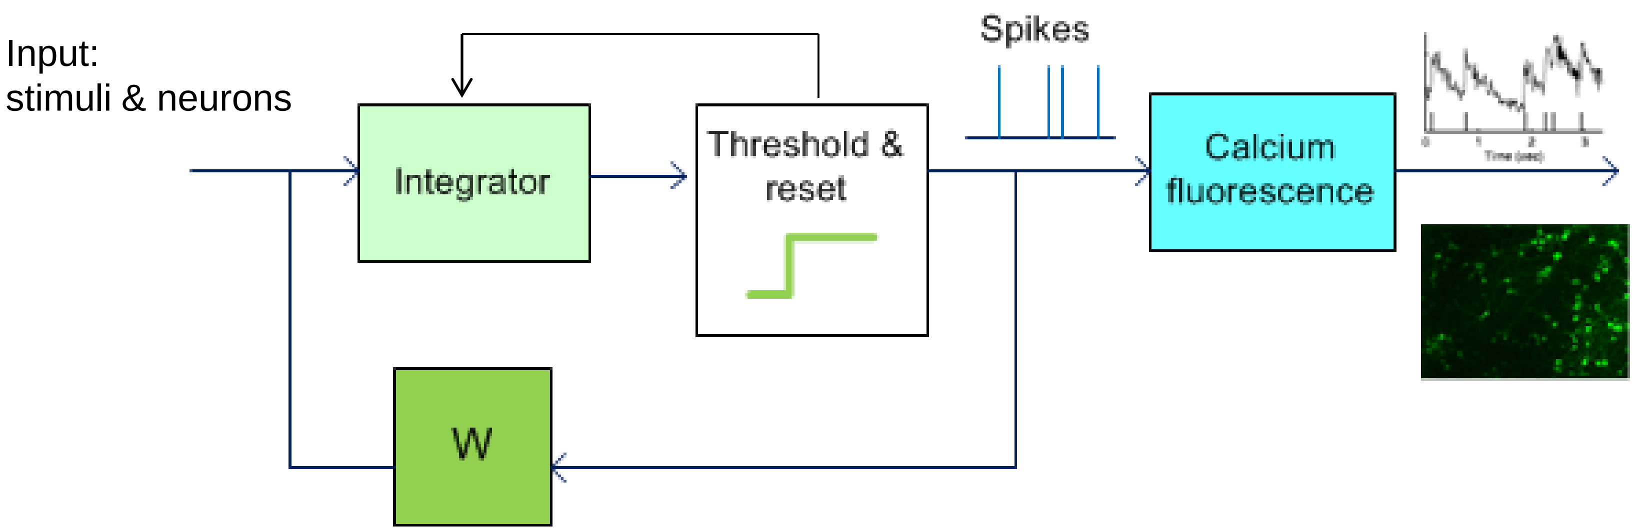

$\boldsymbol{v}^{k+1} = (1-\alpha)\boldsymbol{v}^k + \boldsymbol{W}\boldsymbol{s}^k + \boldsymbol{d}_x^k$ ⟵ Membrane voltage integration

$\boldsymbol{s}^k = 1$, $\boldsymbol{v}^k = 0$ when $\boldsymbol{v}^k > \mu$ ⟵ Spike and reset

$\boldsymbol{z}^{k+1} = (1-\beta)\boldsymbol{z}^k + \boldsymbol{d}_z^k + \boldsymbol{s}^k$ ⟵ Bound Ca^{2+} concentration

$\boldsymbol{y}^k = a\,\boldsymbol{z}^k + \boldsymbol{d}_y^k$ ⟵ Fluorescence

Maximum Likelihood Estimation of Connectivity

- ML estimate:

$$\hat{W} = \arg\max_{W} \log p(\underbrace{y}_{\text{Observations}}|\underbrace{W}_{\text{Desired parameters e.g. connectivity matrix}}) - \lambda \|W\|_1$$

  – Can add regularization term to impose sparsity

Expectation Maximization Algorithm

$$\boldsymbol{v}^{k+1} = (1-\alpha)\boldsymbol{v}^k + \boldsymbol{W}\boldsymbol{s}^k + \boldsymbol{d}_x^k$$
$\boldsymbol{s}^k = 1$, $\boldsymbol{v}^k = 0$ when $\boldsymbol{v}^k > \mu$

$$\boldsymbol{z}^{k+1} = (1-\beta)\boldsymbol{z}^k + \boldsymbol{d}_z^k + \boldsymbol{s}^k$$
$$\boldsymbol{y}^k = a\boldsymbol{z}^k + \boldsymbol{d}_y^k$$

- Want regularized ML:
  $\widehat{W} = \arg\max_{W} \log p(y|W) - \lambda \|W\|_1$
- Problem: Hidden states (v^k, z^k, s^k)
- Use EM iterations

Distribution of hidden states (v^k, z^k, s^k)

E step

M step

Estimate distribution of hidden states given parameters

Easy for Ca^{2+} imaging problem

W is linear in likelihood

LASSO problem

Exact computation hard

Update parameter estimate given density on hidden states:

High-dim integral or Gibbs sampling

[Mischenko et al., 2010]

Parameter estimate W

Decoupling for the E-Step

- Want: State estimates for nonlinear system
  - High dimensional : N Neurons, 3N states
- Key insight: System decouples: scalar iterations $\boldsymbol{q}^k = \boldsymbol{W}\boldsymbol{s}^k$

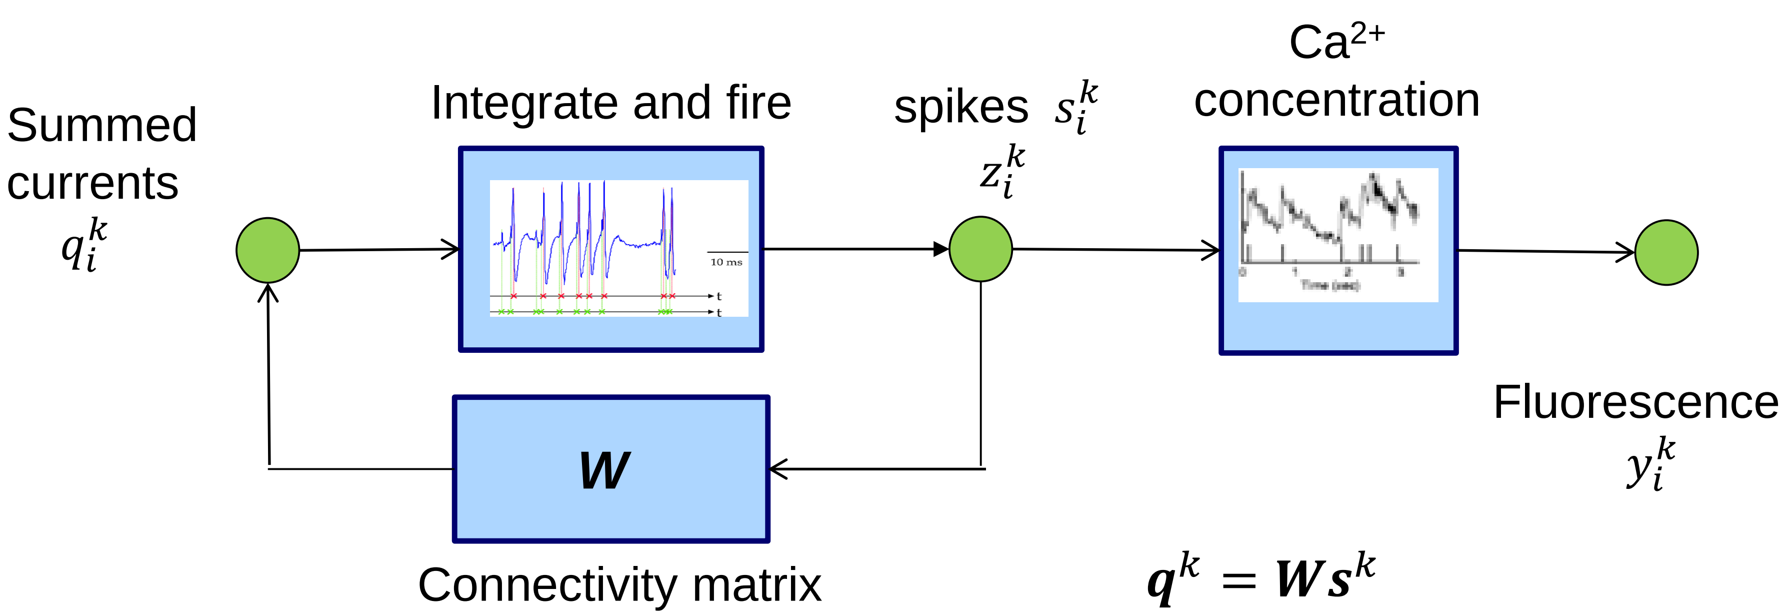

Decoupling for the E-Step

N scalar 1-dim systems

$$v_i^{k+1} = \begin{cases} (1-\alpha)v_i^k + q_i^k & v_i^k \le \mu \\ 0 & v_i^k > \mu \end{cases}$$

$$s_i^k = 1 \quad \text{when} \quad v_i^k \ge \mu$$

N scalar 1-dim systems

$$z_i^{k+1} = (1-\beta)z_i^k + s_i^k$$

$$y_i^k = a z_i^k + d_i^k$$

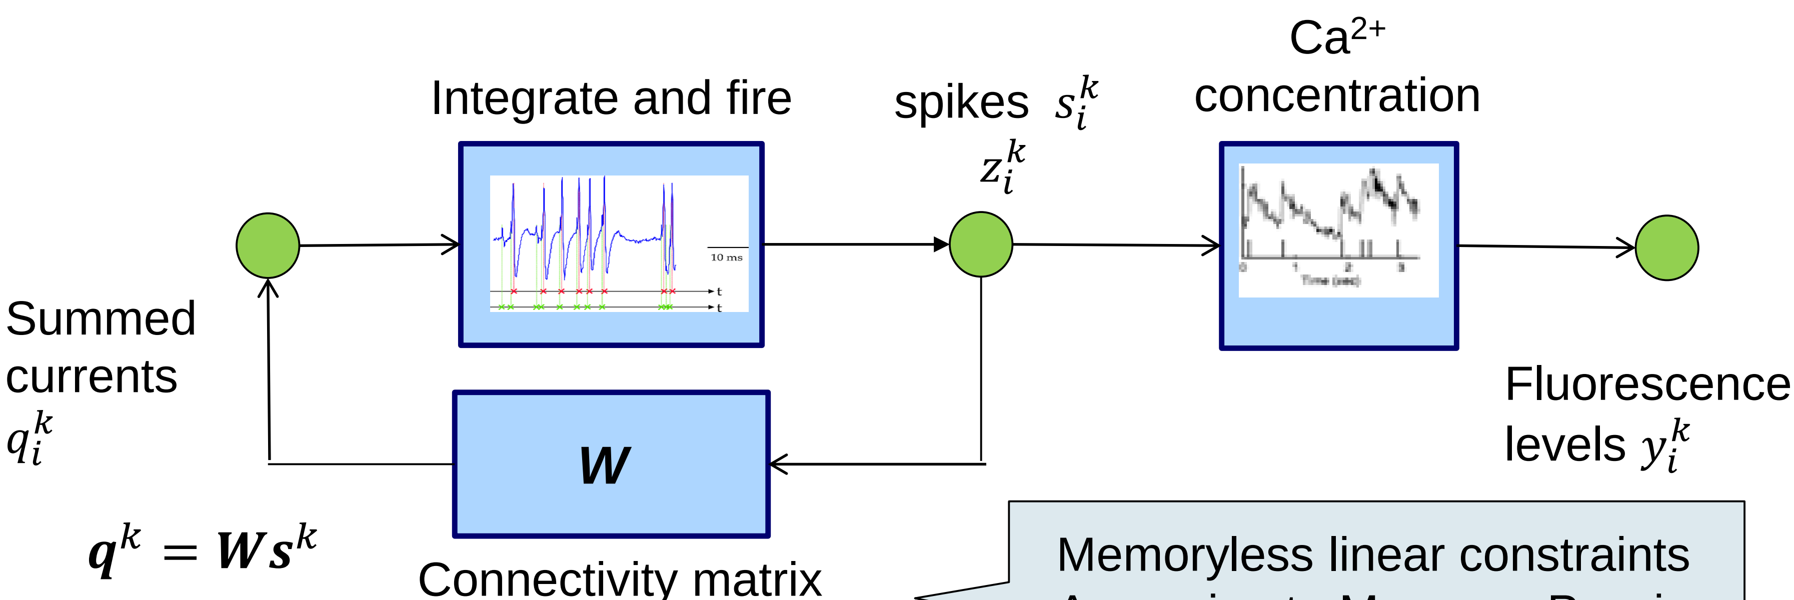

$$\boldsymbol{q}^k = \boldsymbol{W}\boldsymbol{s}^k$$

Memoryless linear constraints
Approximate Message Passing

Decoupling for the E-Step

N scalar Forward Backward

N scalar Forward Backward

$$v_i^{k+1} = \begin{cases} (1-\alpha)v_i^k + q_i^k & v_i^k \le \mu \\ 0 & v_i^k > \mu \end{cases}$$

$$s_i^k = 1 \quad \text{when} \quad v_i^k \ge \mu$$

$$z_i^{k+1} = (1-\beta)z_i^k + s_i^k$$

$$y_i^k = a z_i^k + d_i^k$$

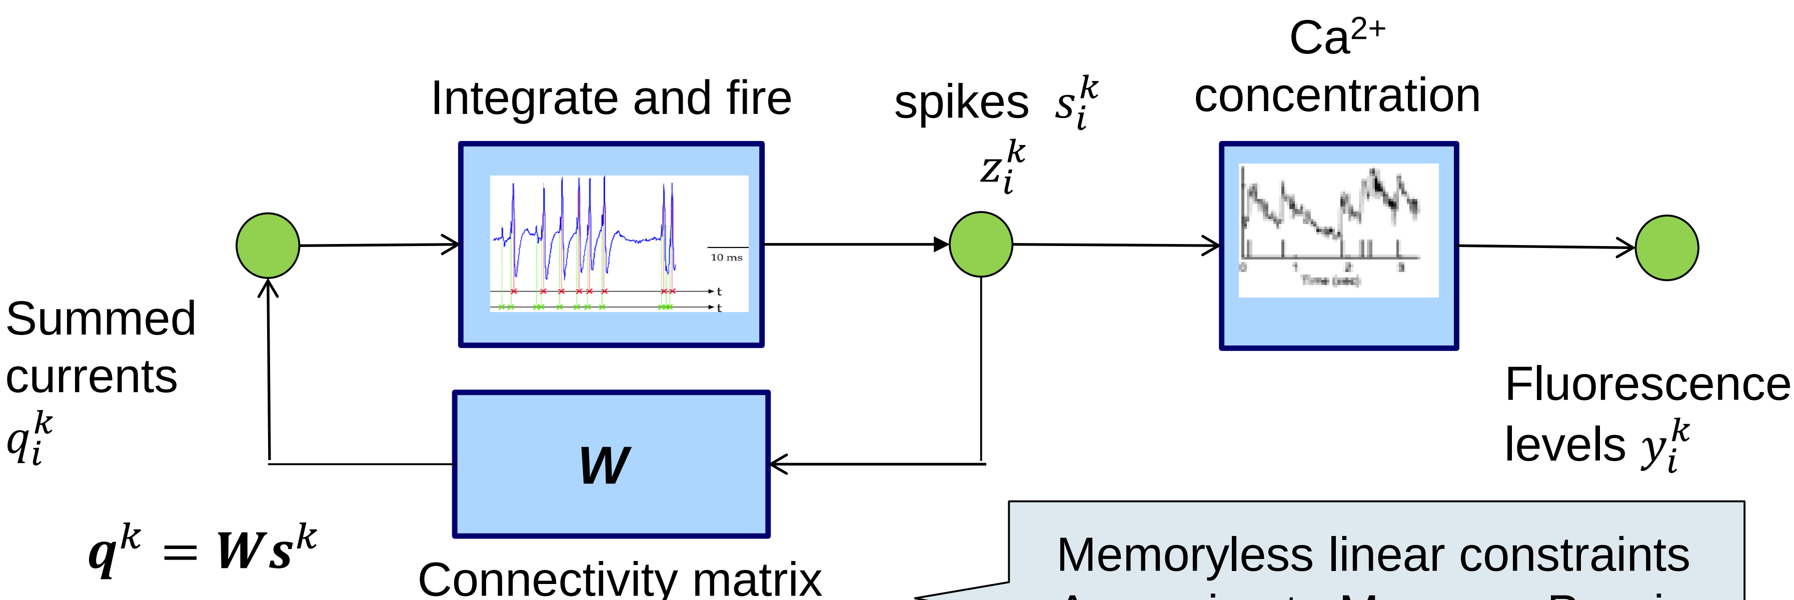

$$\boldsymbol{q}^k = \boldsymbol{W}\boldsymbol{s}^k$$

Memoryless linear constraints
Approximate Message Passing

EM Overview

- Fluorescence movie: estimate
  - Connectivity matrix W, spike times, voltages
- Scalable block
  - High-dimensional, nonlinear dynamical system

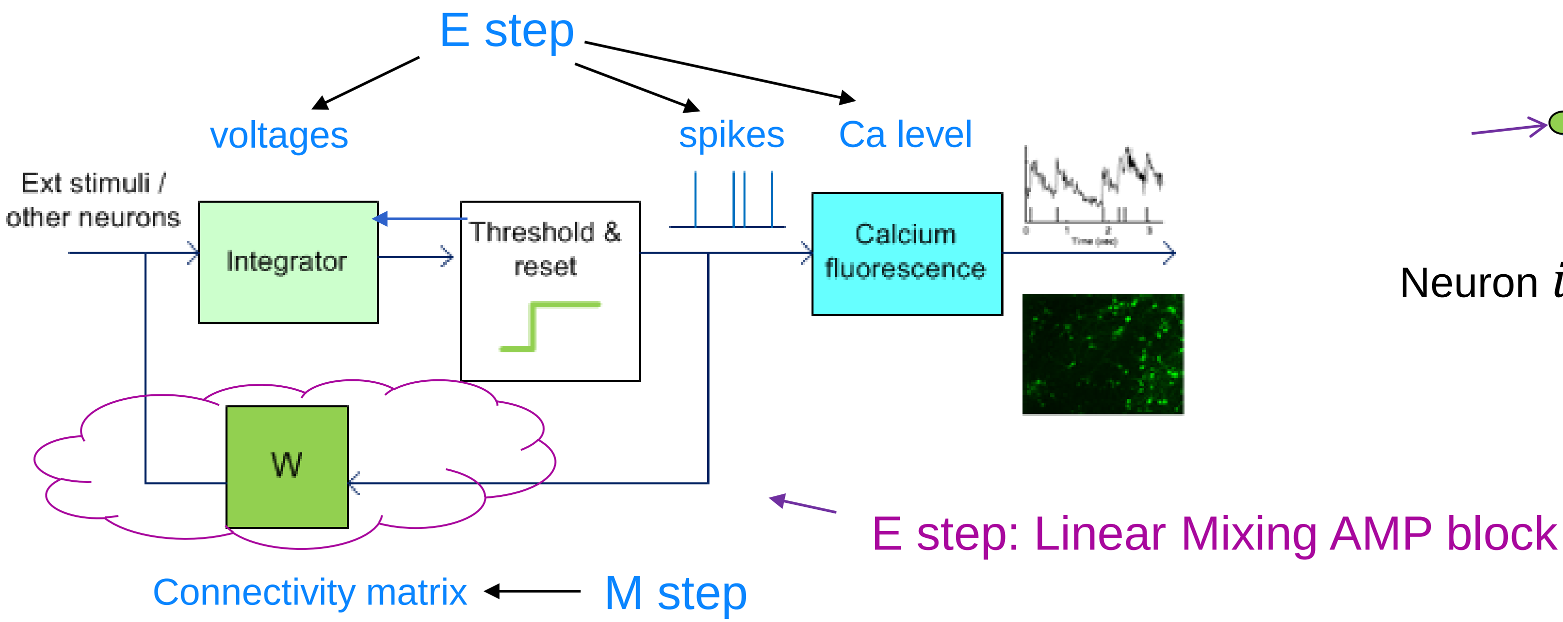

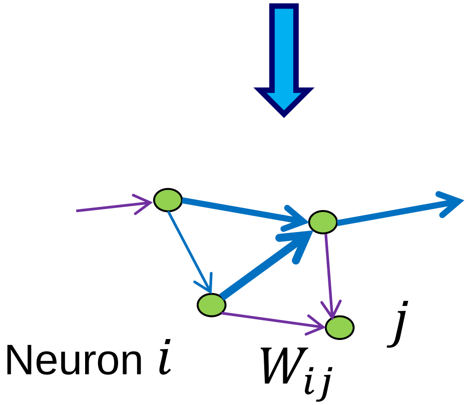

Simulation Results: Accuracy of the Weights

Accurate estimation

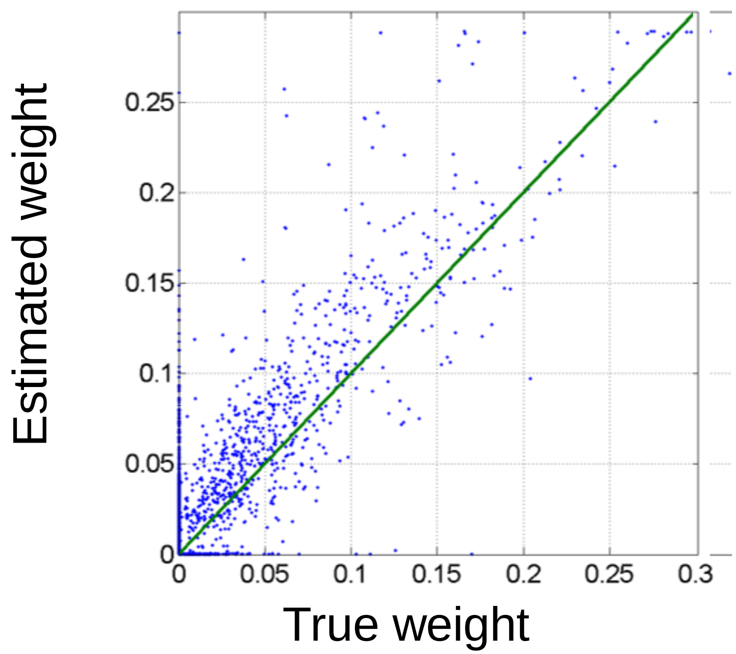

- Neuron model **[Sayer (1990)]**
  - 100 guinea pig cortical neurons
  - Synchronized bursting: 10 spikes/s
  - 10% sparse random connectivity
  - 20 *ms* integrate-fire τ_c
  - 2 *ms* inter-neuron conduction time
  - 1 *ms* time step
- Ca^{2+} imaging model
  - 100 frames/s, 100 s trials
  - 10000 $\mathscr{F}\ell$ values per neuron
  - 500 *ms* Ca^{2+} τ_c
  - Fluorescence SNR = 20 dB

Fletcher et al NIPS '14

Accuracy of the Weights

Fast convergence

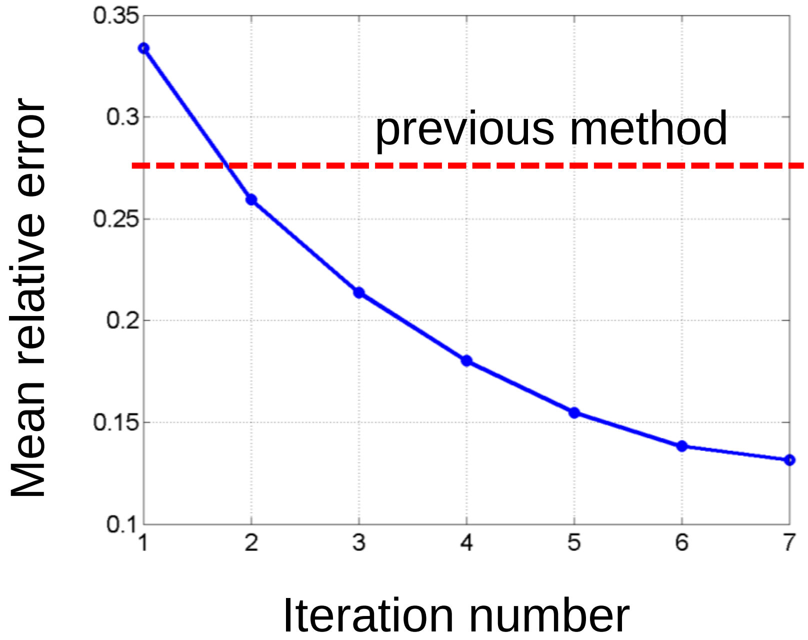

•Relative MSE = $\frac{E\left(W_{ij}-\widehat{W}_{ij}\right)^2}{E\left(W_{ij}\right)^2}$ = 0.12

- Neuron model **[Sayer (1990)]**
  - Guinea pig cortical column
  - 100 neurons
  - 10% sparse random connectivity
  - Synchronized bursting: 10 spikes/s
  - 20 *ms* integrate-fire τ_c
  - 2 *ms* inter-neuron conduction
  - 1 *ms* time step
- Data Collection
  - 100 frames/s, 100 s trials
  - 10000 $\mathscr{F}\ell$ samples per neuron
  - Ca^{2+} τ_c = 500 *ms*, $\mathscr{F}\ell$ SNR = 20 dB

- Previous work MSE=0.28, same parameters **Mischenko et al. (2010)**
- Accurate, low complexity $O(N)$ per iteration, instead of $O(*^N)$

Fletcher et al NIPS '14

UCSC

Network Connectivity Summary

- Network analysis from Ca^{2+} imaging
  - Attack temporal resolution issues for neural dynamics
  - Scalable EM algorithm
- Rich, flexible modeling framework
  - Incorporates nonlinearities, indirect measurements, dynamics
- Computationally scalable solution
  - Linear in number of measurements
- Demonstrated performance
  - Outperforms existing techniques
  - Allows more biologically plausible model with feedback

http://gampmatlab.sourceforge.net/ Phil Schniter!!!

Outline

- Connectivity via multineuronal stimulation
  – Iterative fast adaptive GAMP framework
- Network connectivity via Ca^{2+} imaging
  – Network model - captures dynamics
  – Graphical models
  – Scalable accurate based algorithm
- Receptive field of retinal ganglion cells
  – Space-time salamander response to stimuli
  – Improved identification with limited data

Receptive Field Identification

- Retinal ganglion cell (RGC)
  Sensitive to light in its field of view,
  or *receptive field*

- Tuned to some local features
  in time and space (curve, edge, etc)

- Response estimation of RGCs:
  - Expose retina to image
  - Measure response via electrode
  - Fit model

- Challenge: Model is often nonlinear

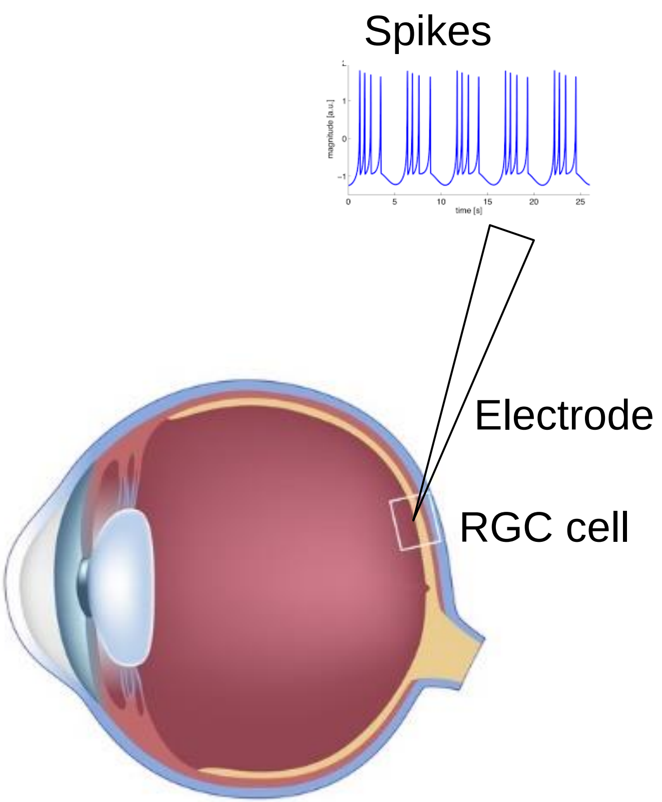

Salamander Receptive Field Identification

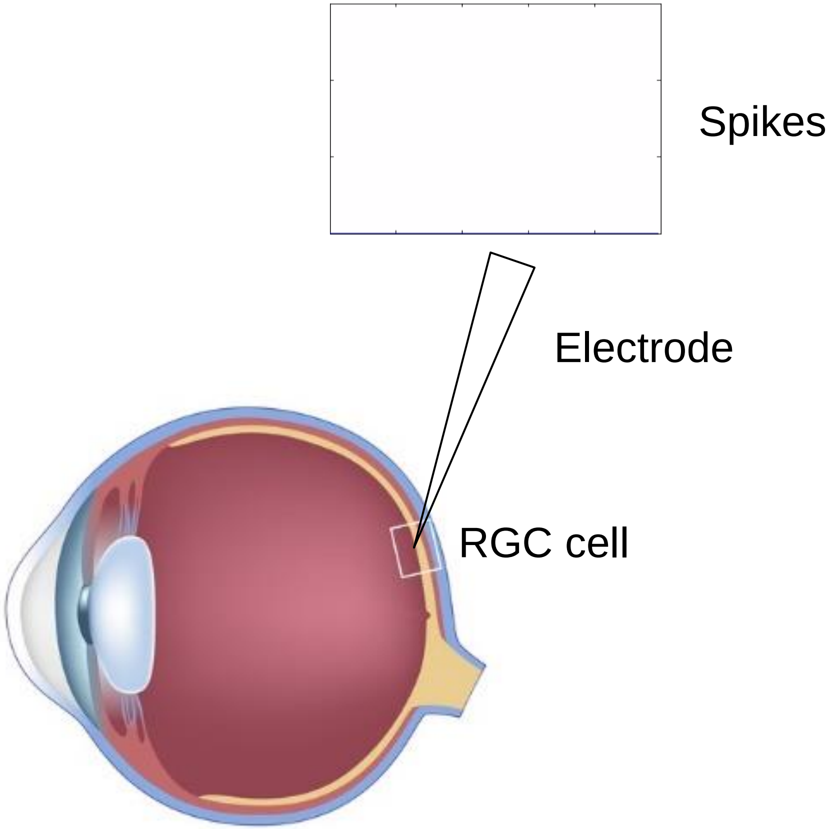

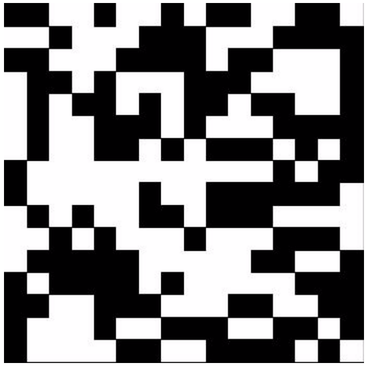

Visual stimuli

Retinal Ganglion Cell LNP Model

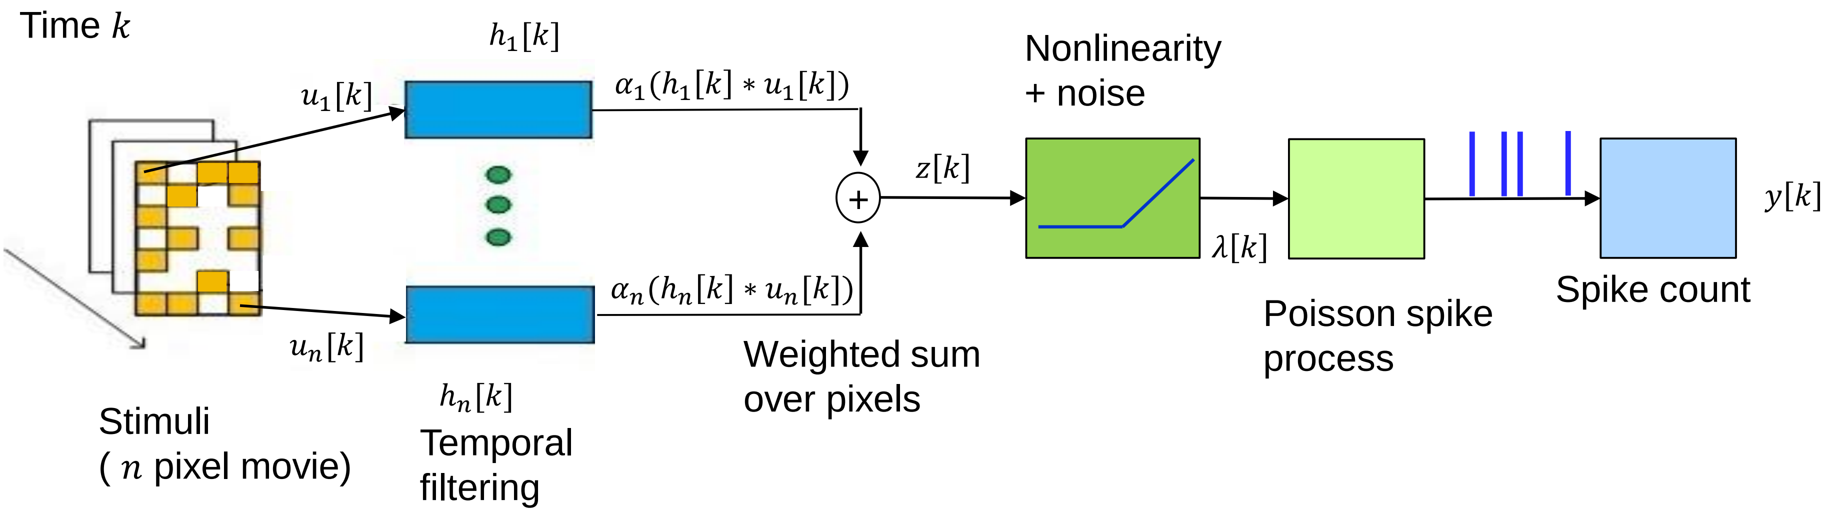

- Linear-Nonlinear Poisson model. Given stimuli $u_i[k]$
  - Filtering over time and space: $z[k] = \sum_{i=1}^{n} \alpha_i(h_i[k] * u_i[k])$
  - Nonlinear phase: $y[k] = Poisson\big(f(z[k] + d[k])\big)$

- Identification problem: Given stimuli $u_i[k]$ and spike counts $y[k]$
  - Estimate weighted filters $w_i[k] = \alpha_i h_i[k]$,
    Describes space-time response of neuron to pixel i
  - Each $h_i[k]$ is an L tap filter
  - Estimate nonlinearity $f(\cdot)$

Fletcher et al. NIPS 2012

Structured Matrix View of Dynamic LNP

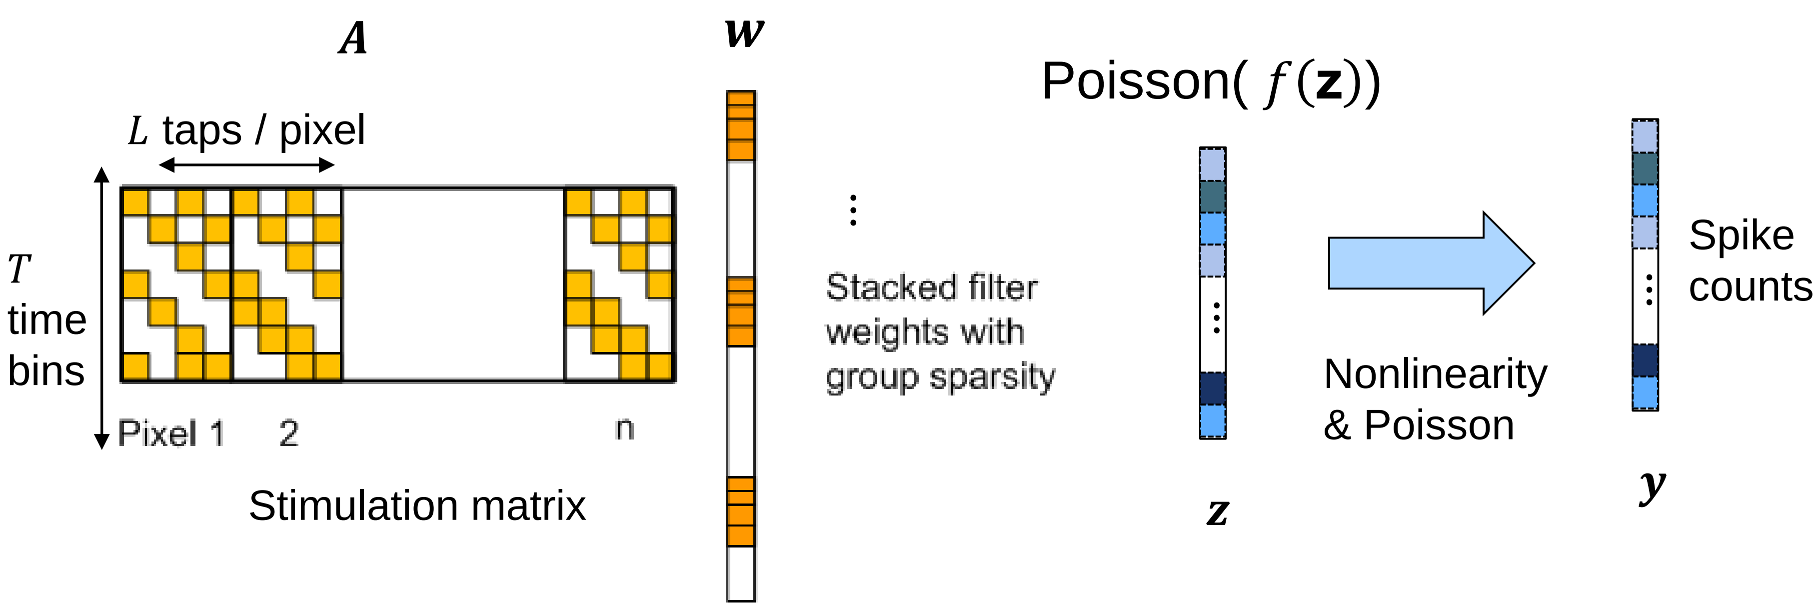

- LNP model: cascade of linear and nonlinear system
  - **A** rows: n pixel values at L delays (L = filter taps)
- Weights have a group sparse constraint:
  - RGC sensitive to small image region (spatial sparsity)
  - Coefficients of filter of one stimuli are on or off together

Classical Methods and their Limitations

- Linear methods:
  - Matched filter: $\hat{\boldsymbol{w}} = \frac{1}{n}\boldsymbol{A}^{\boldsymbol{T}}\boldsymbol{y}$ (also called STA)
  - Linear MMSE: $\hat{\boldsymbol{w}} = \left(\boldsymbol{A}^{\boldsymbol{T}}\boldsymbol{A} + \sigma^2\boldsymbol{I}\right)^{-1}\boldsymbol{A}^{\boldsymbol{T}}\boldsymbol{y}$ (also called RC)
  - Also, linear least squares / zero forcing
  - Simple but does not exploit sparsity

- Compressed sensing methods

$$\hat{\boldsymbol{w}} = \arg\min_{w} ||\boldsymbol{y} - \boldsymbol{A}\boldsymbol{w}||_2^2 + \lambda||\boldsymbol{w}||_1$$

  - Exploit sparsity of **w**
  - Many methods: LASSO, OMP, CoSAMP,…
  - Could also incorporate group sparsity via group Lasso
  - But, does not account for output nonlinearities

Hybrid Algorithm for Structured Input: GAMP

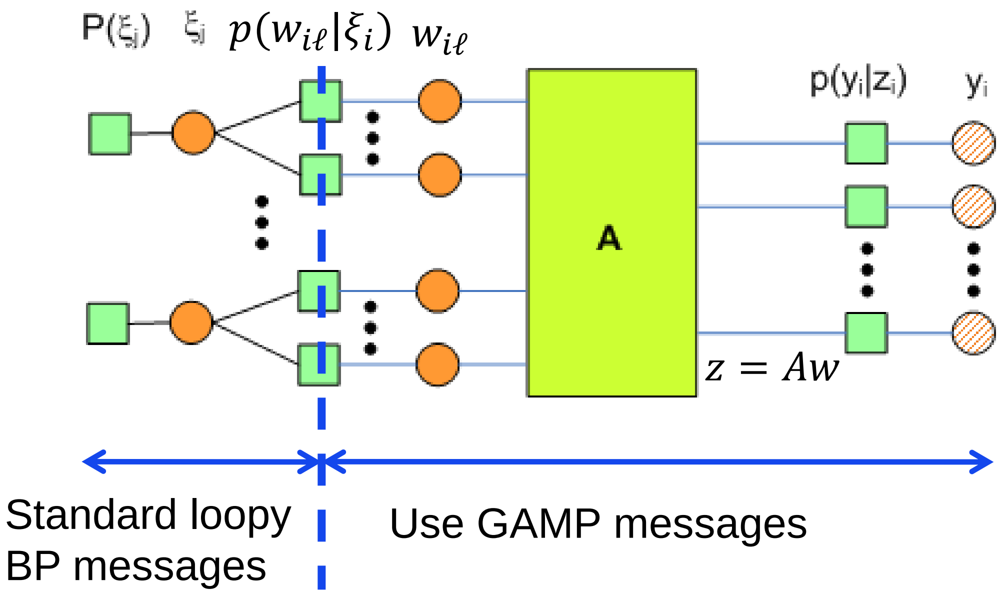

- Introduce binary variables ξ_j to correlate sparsity over time
- Apply GAMP in a "turbo" manner with loopy BP
- Low complexity: each ξ_j is binary
- More general than group OMP and group lasso;
  - Similar complexity,
  - Better performance **[Rangan, Fletcher, Goyal & Schniter '12]**

Receptive Field Computation Considerations

- Problem size: 3630 variables
  - 11 x 11 pixels, 30 taps per pixel
  - 190,000 measurements (~30 min at 10 ms sampling)
  - Structured A matrix is 190,000 by 3630
- AGAMP iteration cost: multiplying by **A** & **A***
  - Exploit block Toeplitz structure and entries are 0-1
- For larger problems, algorithm is parallelizable
  - Graphical methods: inherent decomposable
  - Parallelize multiplications across rows/columns of **A**
- Cannot theoretically guarantee convergence
  - Demonstrate performance experimentally

Experimental Results: Cross-Validation

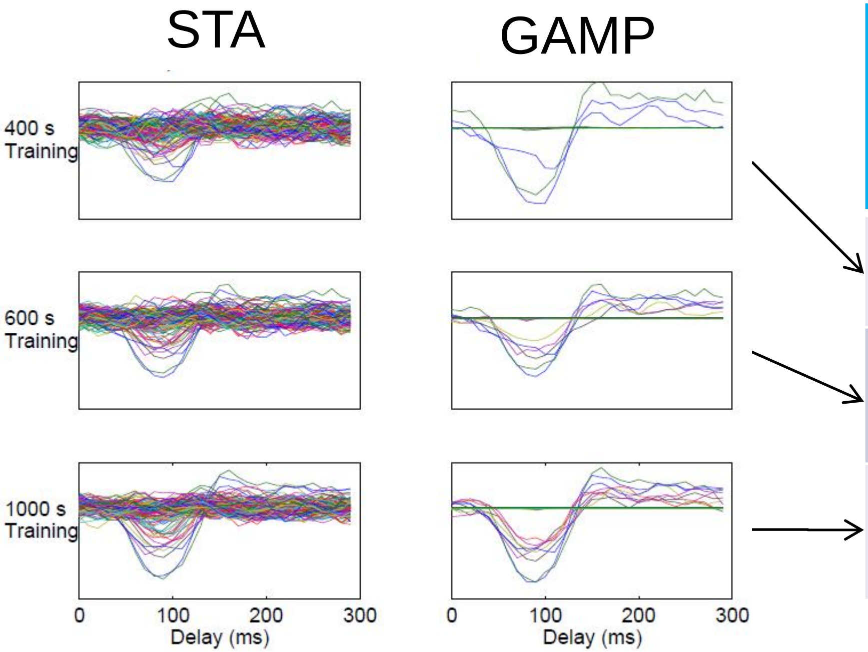

Num training samples	Cross-validation score	
	STA	G-AMP
25000	0.906	0.917
50000	0.914	0.921
100000	0.918	0.923

- Validated on data used in training (190000 total samples)
- Cross-validation score = Geometric mean of likelihood of spike rate
- GAMP: same error, 25000 samples versus 100,000

Data: Anthony Leonardo
Janelia Farm

Salamander Retinal Response (Spatial)

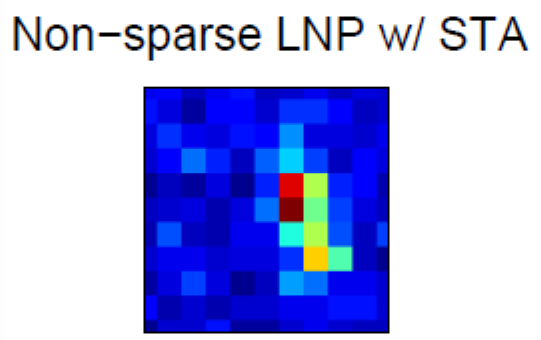

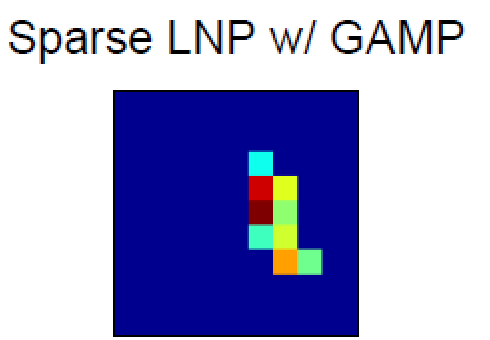

Spatial receptive field
estimates for 11x11 pixel area

0 1

(normalized)

- Spatial receptive field:
  - Plot estimated 11x11 response magnitudes.
  - Color = 30-tap filter magnitude for each pixel
- Standard STA estimate shows noisy (spurious) responses outside
- GAMP method removes noise
  - Shows only a response in a small area

Filter response 11 x 11 pixels for salamander RGCs
Data from Anthony Leonardo, Janelia Farm

AMP++ methods: Applications in Imaging

- Hybrid-AMP can incorporate complex structure
  - Incorporate dependencies between wavelet coefficients
  - Hierarchical models, etc

http://gampmatlab.sourceforge.net/

Algorithm	NMSE (dB)	Comp time (secs)
MHT+IRWL1	-14.37	363
CoSAMP	-16.90	25
SPLG1	-18.06	536
MCMC	-20.10	742
Turbo-GM	-20.74	51

Lowest MSE and almost fastest computation

[Som, Schniter, 2011]

Thank you

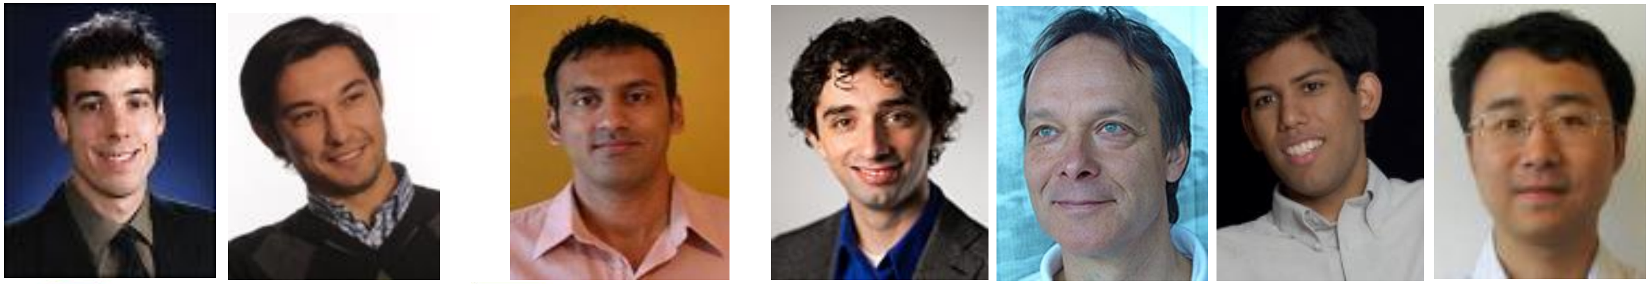

Much thanks! To: Fritz Somer, Surya Ganguli, Eero Simoncelli, Eftychios Pnevmatikakis,Matthias Bethge, Anthony Leonardo, Evan Lyall, students of Deisseroth & Adesnik labs, Liam Paninski, Mayur Mudigonda, Jascha Sohl-Dickstein…

- Funding: NSF CAREER,ONR, AFOSR, Qualcomm

Moving forward

- Better account for exogenous effects
  - Correlations across neurons of interest**
  - Larger areas via “shotgun” techniques **[Pnevmatikakis et al. 2013]**
- Data sets in collaboration
  - Paninski Lab, Columbia; Tolias Lab, Baylor, Allen Institute
  - Validate methods: first cultured without exogenous
- Model new Ca^{2+} indicators:
  - GCaMP6f (2013): faster decay, rise time = 50-75 ms, ….
  - Nonlinear fluorescence model
- Theory
  - Convergence issues of GAMP via new ADMM GAMP
  - Networks of low-dimensional, nonlinear dynamical blocks
  - Provable results for structured non-iid, structured A

Future work : dynamical networks

High dimensional inference for dynamical systems

- *Generalized linear dynamical networks:*
  - Underlying low-dimensional, nonlinear dynamical blocks
  - Linear memoryless constraints, graphical models
- Many phenomena
  - Neural systems, communication networks, particles, media, …
  - Extends GLM to include networked dynamics
- Can we extend methods for:
  - Scalable estimation algorithms?
  - Learning connectivity?
  - Provable guarantees?

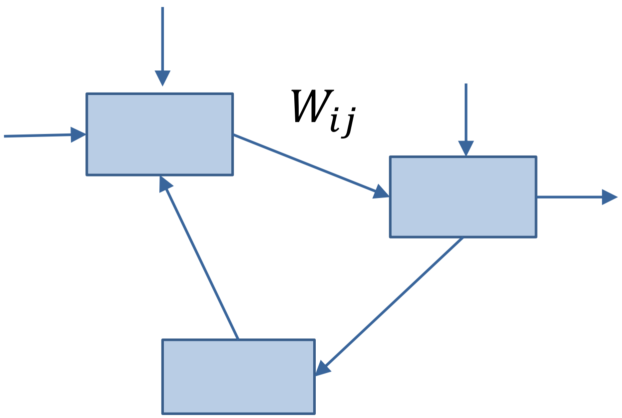

$$x_i^{k+1} = f_i(x_i^k, u_i^k)$$

Joint Estimation and Learning for GLMS

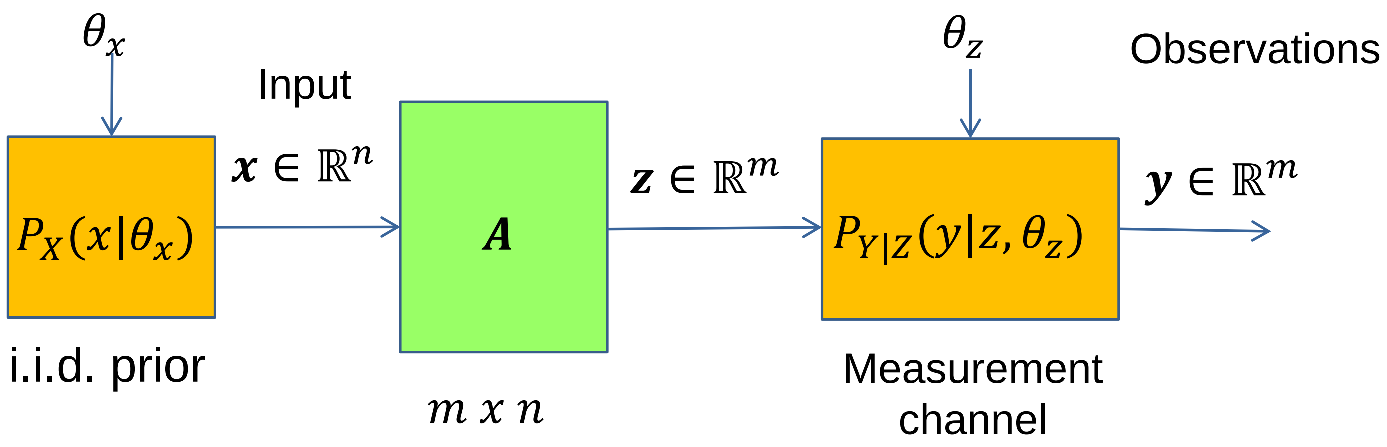

- GLM with unknown parameters θ_x and θ_z
  - Unknown prior, nonlinearities, noise…
- Joint estimation learning problem: Given y and A:
  - Estimate input x and z,
  - Learn parameters θ_x and θ_z in distribution
  - **Consistent**

Fletcher, Rangan NeuRAMP NIPS 2011

Kamilov, Fletcher, et al NIPS 2012, Trans IT 2014

Future work : broader

High dimensional inference for dynamical systems

- *Generalized linear dynamical detworks:*
  - Underlying low-dimensional, nonlinear dynamical blocks
  - Linear memoryless constraints, graphical models
- Many phenomena
  - Neural systems, communication networks, particles, media, …
  - Extends GLM to include networked dynamics
- Can we extend methods for:
  - Scalable estimation algorithms?
  - Learning connectivity?
  - Provable guarantees?

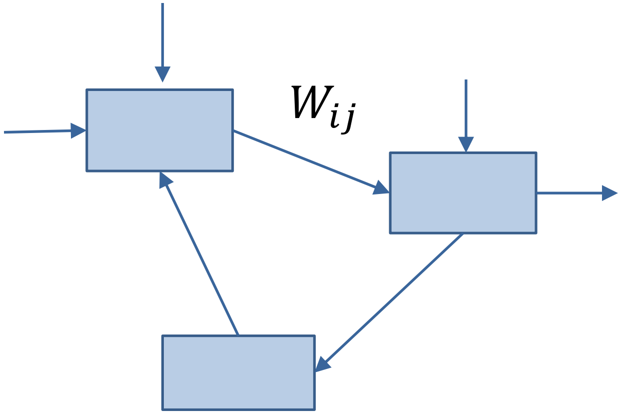

$$x_i^{k+1} = f_i(x_i^k, u_i^k)$$