A theory of neural dimensionality, dynamics and measurement

Surya Ganguli

Dept of Applied Physics

And, by courtesy,

Neurobiology
Electrical Engineering

collaboration with Shenoy Lab

Stanford University

Peiran Gao

Eric
Trautmann

An exponential Moore's Law for the number of recorded neurons

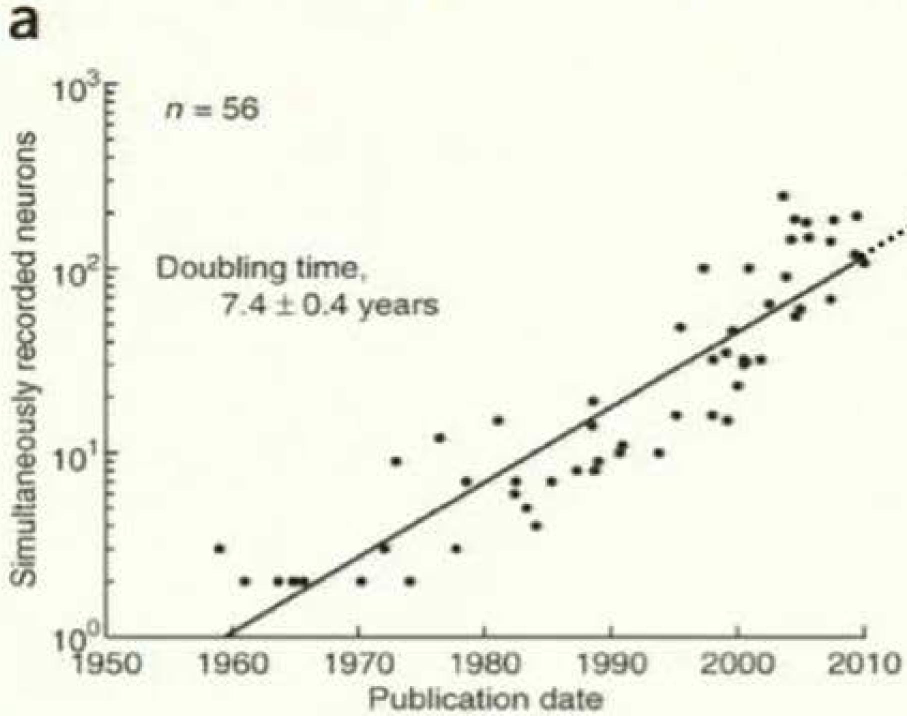

Stevenson & Kording, 2011

Multielectrode recordings allow us to record from 10^2 to 10^3 neurons.

Mammalian circuits controlling complex behaviors contain > 10^6 to 10^9 neurons.

Are we in an anti-Goldilocks moment? (122 years to get 5 orders more)

Too many neurons so that data analysis is not easy.
Not enough neurons to really understand circuit computation?

An example dataset: the single neuron view

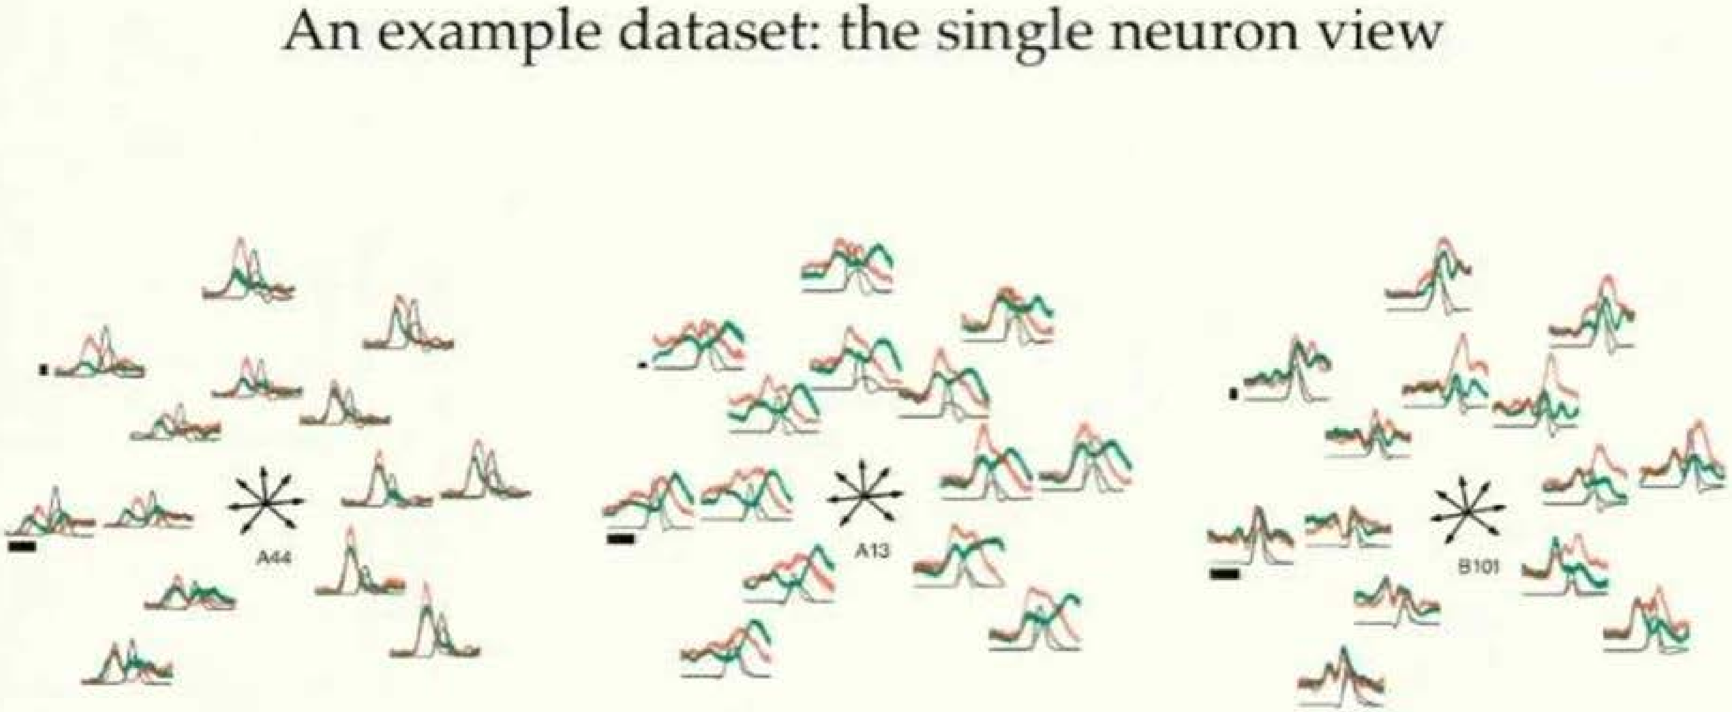

Churchland and Shenoy, J. Neurophys. 2007

Trial averaged firing rates from 3 neurons while a monkey is reaching to targets at 7 directions, two lengths and two speeds (red / green)

There are about 100 more neurons like these.

How are such datasets analyzed?

Dynamical portraits of circuit computation via dim reduction

locust, antenna lobe

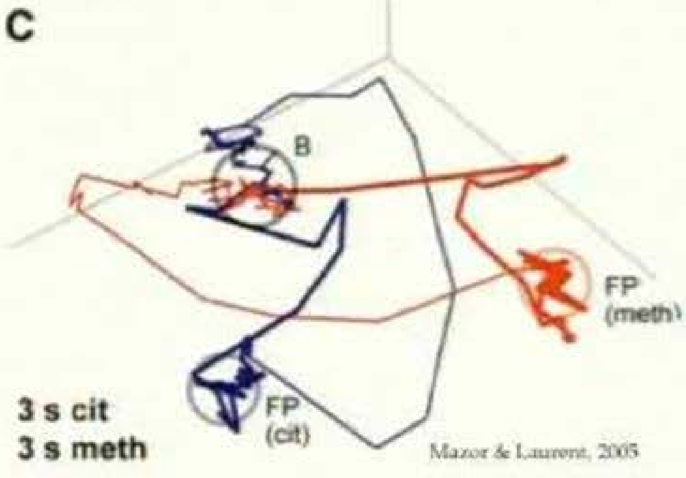

Mazor & Laurent, 2005

monkey, PFC

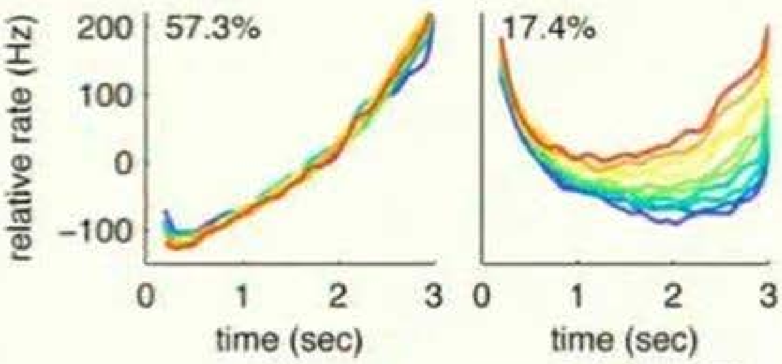

Machens et al., 2010

monkey, motor/pre-motor cortex

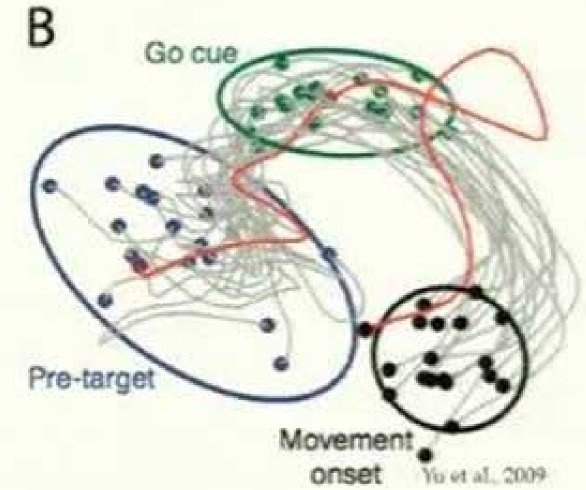

Yu et al., 2009

zebra fish, whole brain

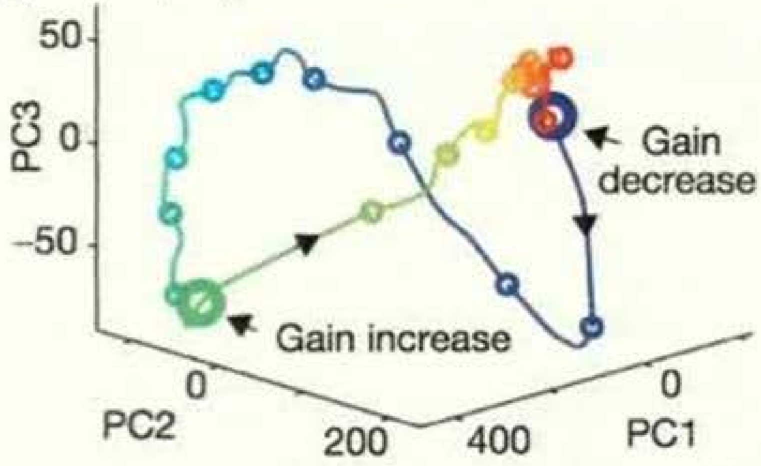

Ahrens et al., 2012

Fundamental conceptual questions

In a wide variety of neuronal recordings, measured neuronal dimensionality is far less than the number of neurons.

What is the interpretation of this empirical observation?

What is the origin of this underlying simplicity?

While we now record from many neurons (O(100)); brain circuits controlling behavior have many more unrecorded neurons (O(1 billion) in primate motor cortex).

How would the dimensionality change if we recorded more neurons?

How would the dynamical portraits change if we recorded more neurons? Can we trust them with such small numbers of neurons?

What (if anything) can we learn about large dynamical networks at such an overwhelming level of under sampling?

The need for a theory of dimensionality and dynamics

In primate motor cortex there are
O(1 billion) neurons controlling O(650) skeletal muscles.

In these experiments, O(100) neurons were recorded.

The PCA dimensionality (~ 70% variance explained) across all 8 reaches is 7.

The PCA dimensionality (~ 70% variance explained) for one reach is 3.3.

Where do these numbers come from – how large could they possibly be?

New mathematical definition of neuronal task complexity:
1) Upper bound dimensionality.
2) Tell us how many neurons we need to record.

New definition of neural task complexity

Neural dimensionality

Theorem:
dimensionality ≤ task complexity

Motor cortical data is as high dimensional as possible given task complexity

Future experiments:
recording more neurons w/o increase in task complexity ≠ richer datasets

Neural measurement

Conditions for accurate recovery of dynamic portraits

Random projection theory:
# of neurons required
~ log(neural task complexity)

Past results:
existing dynamic portraits are likely to be accurate despite recording few neurons

Neural Dimensionality and Task Complexity: Intuition

	fix angle, vary time	fix time, vary angle	vary both angle/time
max dim	$\sim T / \tau$	$\sim 2\pi / \Delta$	$\sim T / \tau \times 2\pi / \Delta$

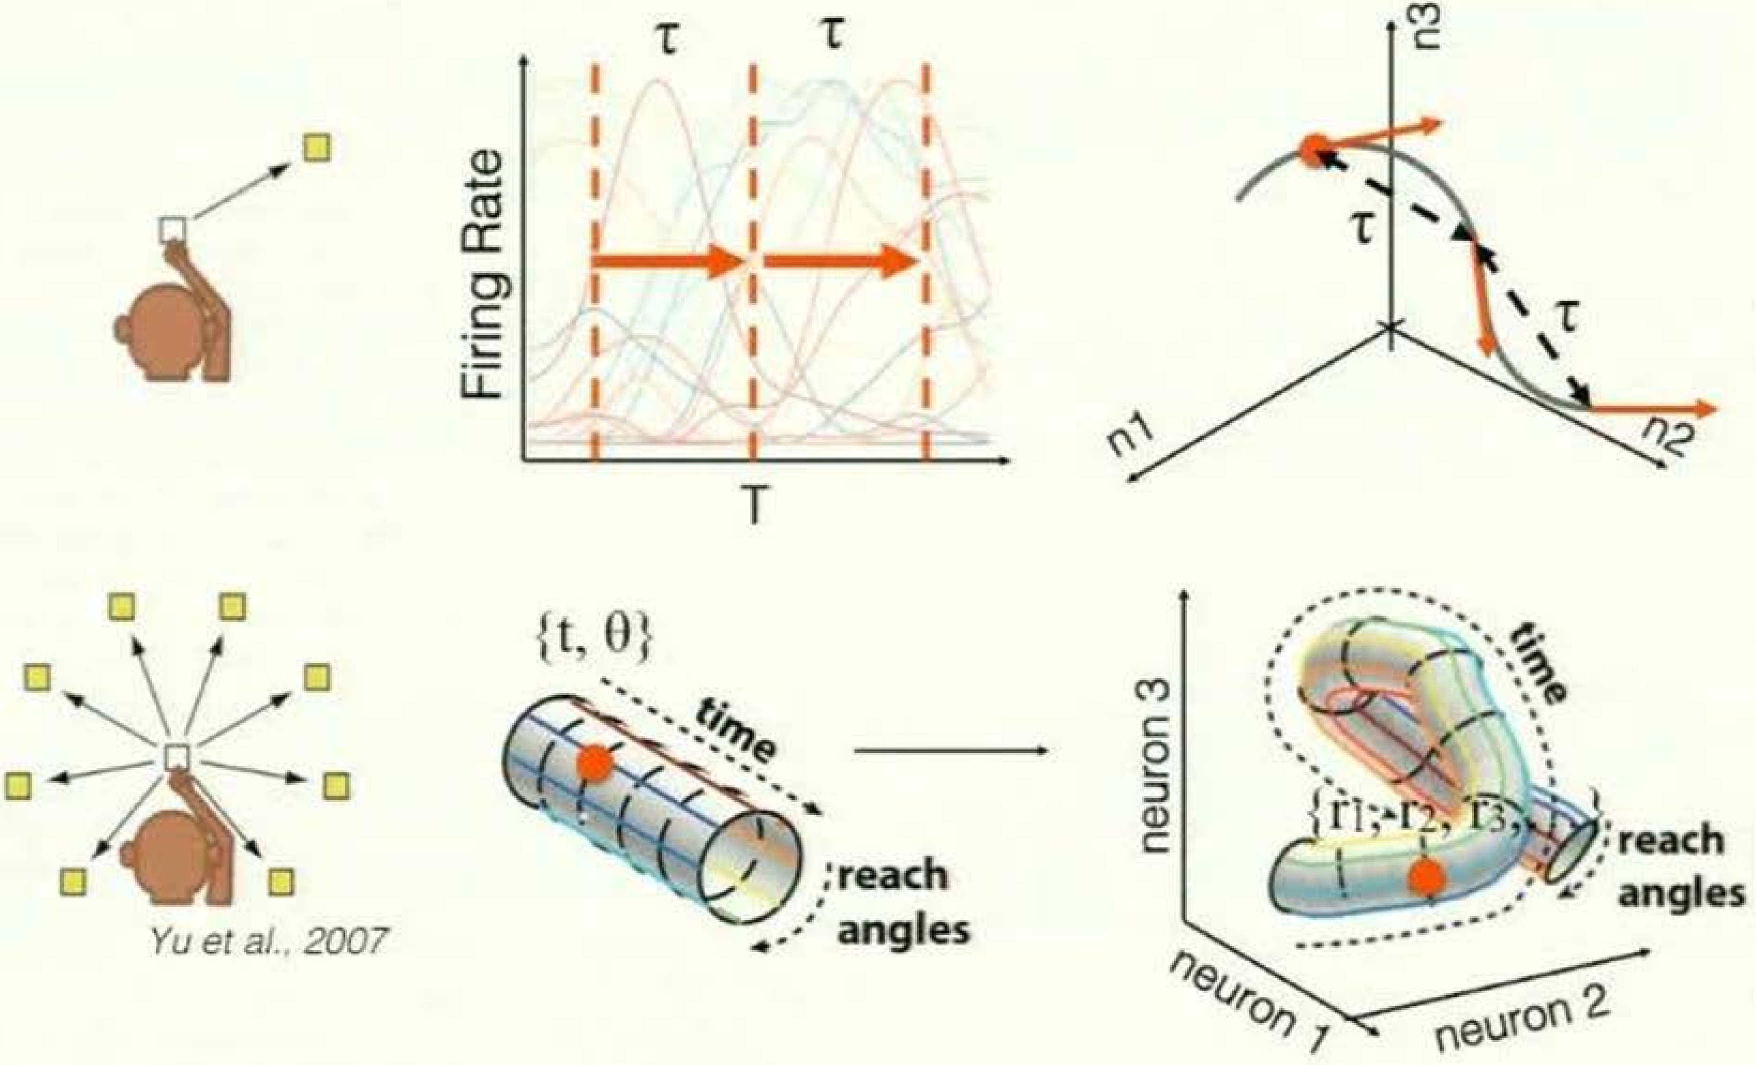

Yu et al., 2007

Neural Dimensionality and Task Complexity: Theory

Task parameters: $p_1, p_2, \ldots, p_K$ (time, speed, angle, distance etc.)

Over ranges: $L_1, L_2, \ldots, L_K$

With neural correlation lengths: $\lambda_1, \lambda_2, \ldots, \lambda_K$

Define task complexity:

$$c \frac{L_1}{\lambda_1} \frac{L_2}{\lambda_2} \cdots \frac{L_K}{\lambda_K}$$

Our theory provides:

1) A way to quantitatively extract neural correlation length parameters $\lambda_1, \lambda_2, \ldots, \lambda_K$
2) and the proportionality constant c, such that we can prove…

*participation ratio of PCA eigen spectrum ~70% var explained

A theorem:

neural dimensionality*
≤
min(task complexity, # of recorded neurons)

$$D = \frac{(\sum_i \lambda_i)^2}{\sum_i \lambda_i^2}$$

Neural Dimensionality in Motor Cortex

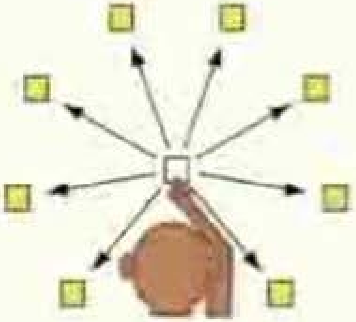

Yu et al., 2007

109 neurons	Dimensionality	Task complexity
Single reach	3.3	4.2
Multiple reaches	7	10

Implication: neural dimensionality not small; but almost as large as possible given task constraints

Prediction #1: vary task complexity by varying T, dimensionality should <u>vary linearly with T</u>

Prediction #2: vary # of neurons in the dataset, dimensionality should be <u>unchanged</u>

Neural Dimensionality in Motor Cortex

Implication: task complexity, not # of neurons, is the main limit on neural dimensionality

Prediction #1: vary task complexity by varying T, dimensionality should vary linearly with T

Prediction #2: vary # of neurons in the dataset, dimensionality should be unchanged

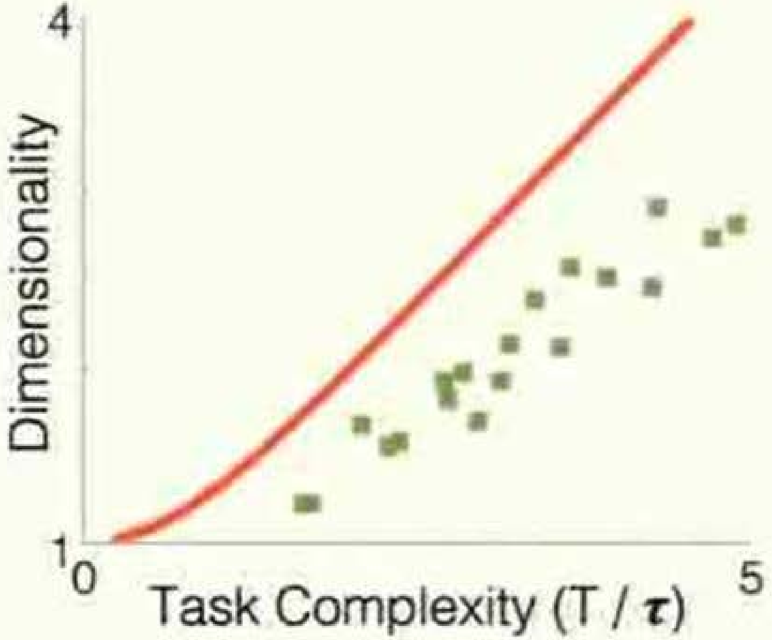

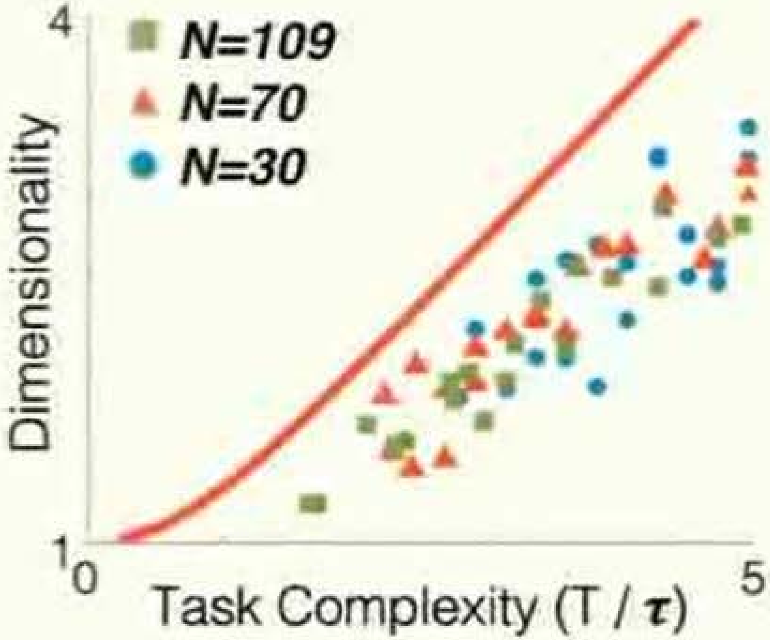

Measuring the Dynamic Portrait under Sub-sampling

When are portraits from relatively few neurons = those from all neurons?

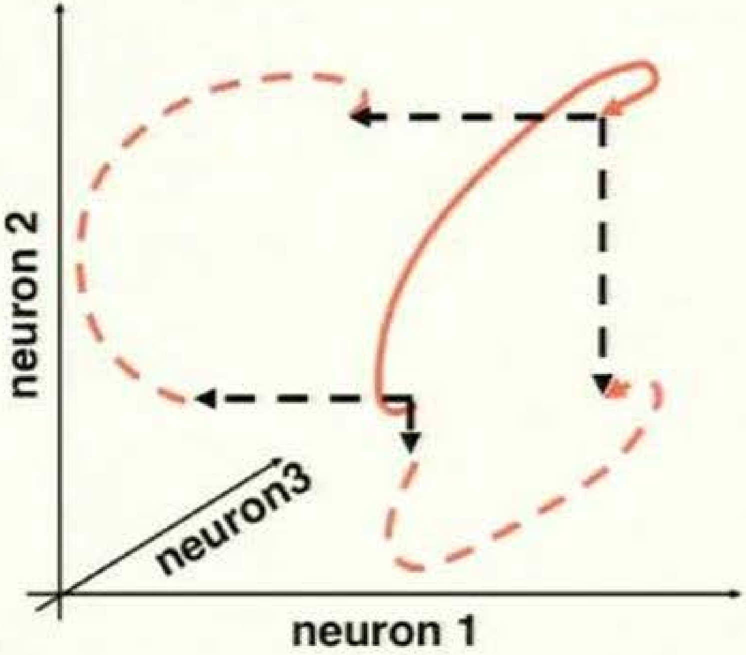

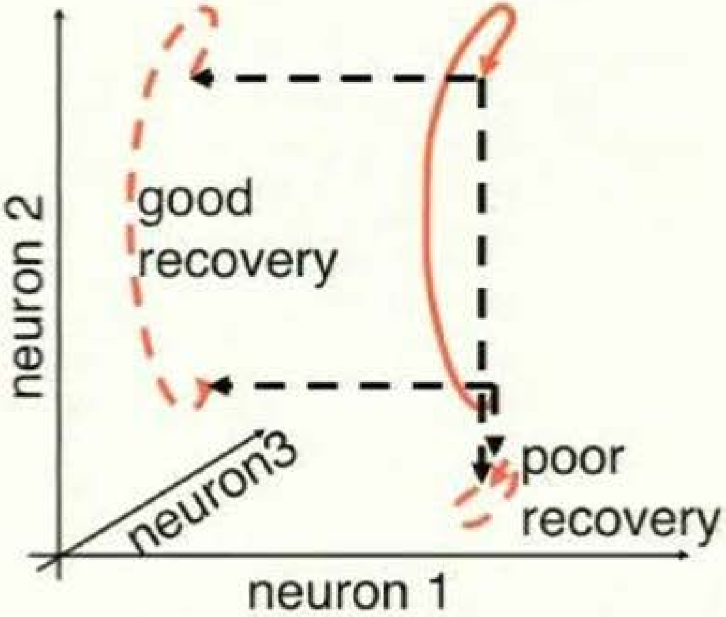

When patterns of neural activity are distributed across neurons, we can accurately recover dynamic portraits despite subsampling

The act of neuronal measurement as a random projection

If neural manifold is randomly oriented:

An experiment we can do: measure a random subset of M neurons

random subset

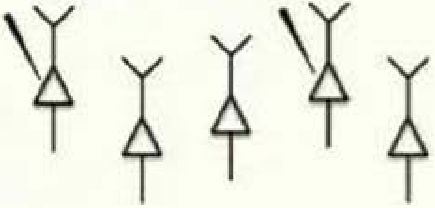

is equivalent to

An experiment we cannot yet do: measure M random linear combinations (i.e. random projections) of *all* neurons

random projection

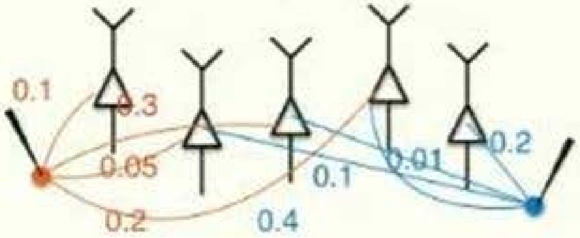

A larger context: random projections

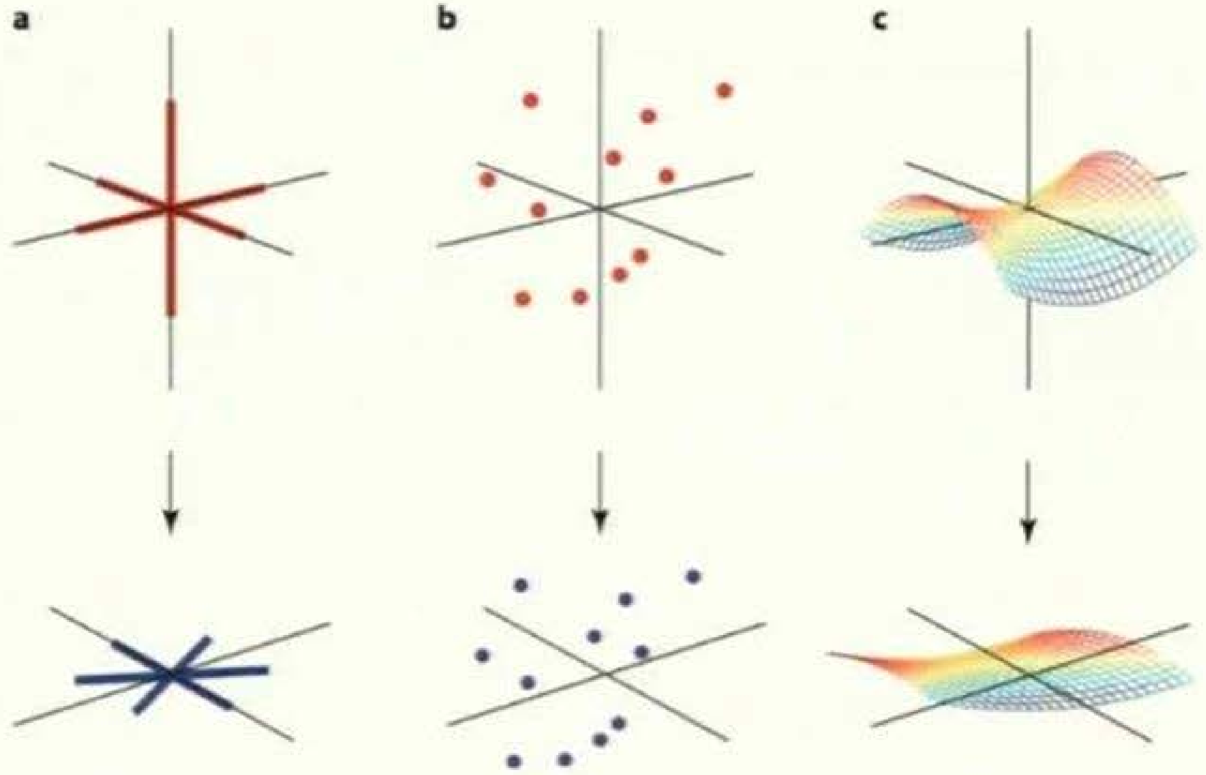

$\mathbf{x} = \mathbf{A}\mathbf{s}$ is a random projection from a N dim space down to an M dim space

Data / interesting signals live on a K-dim submanifold in N-dim space

When will the geometry of this manifold be preserved under a random proj. ?

Distortion: $D_{ab} = (\| \mathbf{A}\mathbf{s}^a - \mathbf{A}\mathbf{s}^b \|^2 - \| \mathbf{s}^a - \mathbf{s}^b \|^2) / \| \mathbf{s}^a - \mathbf{s}^b \|^2$

A consequence of neuronal measurement as a random projection

By adapting random projection theory:

$$\#\ \text{neurons needed} = \frac{1}{\text{distortion}^2}\left(c_1 \log(\text{ task complexity}) + c_2\right)$$

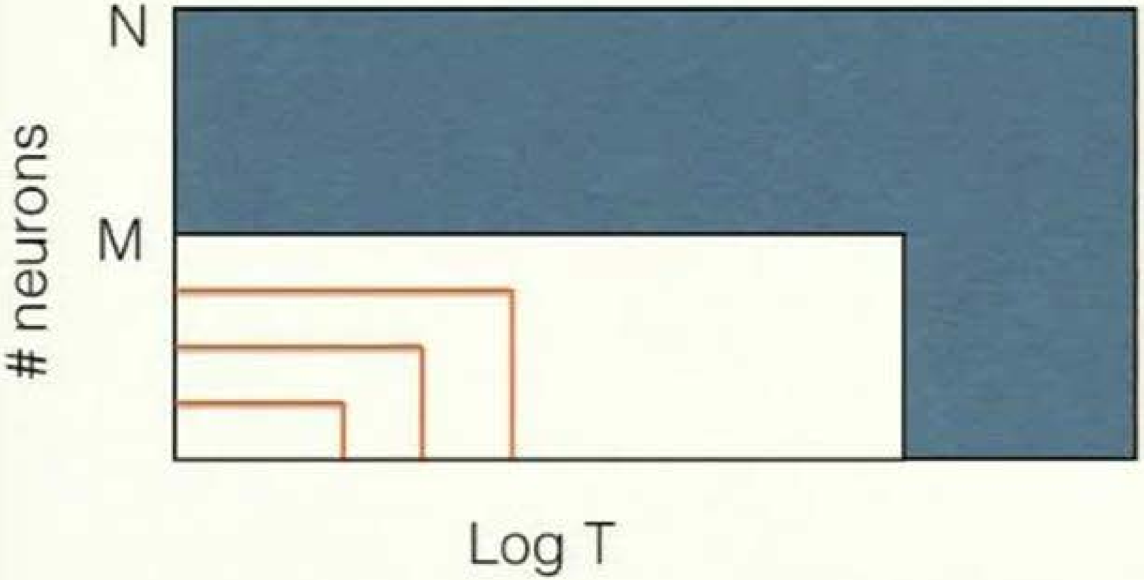

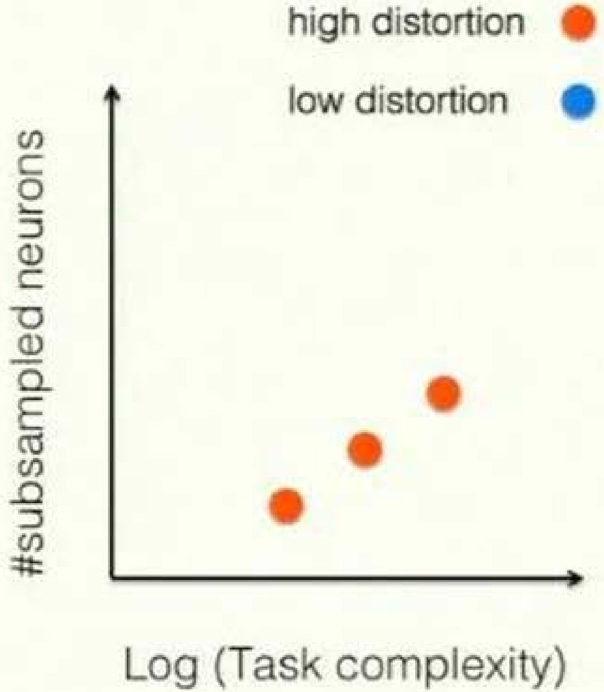

To keep the same level of desired distortion, # of neurons need only scale logarithmically with task complexity (good news!)

To maintain accuracy of the recovered portraits, # of neurons required ~ log(task complexity)

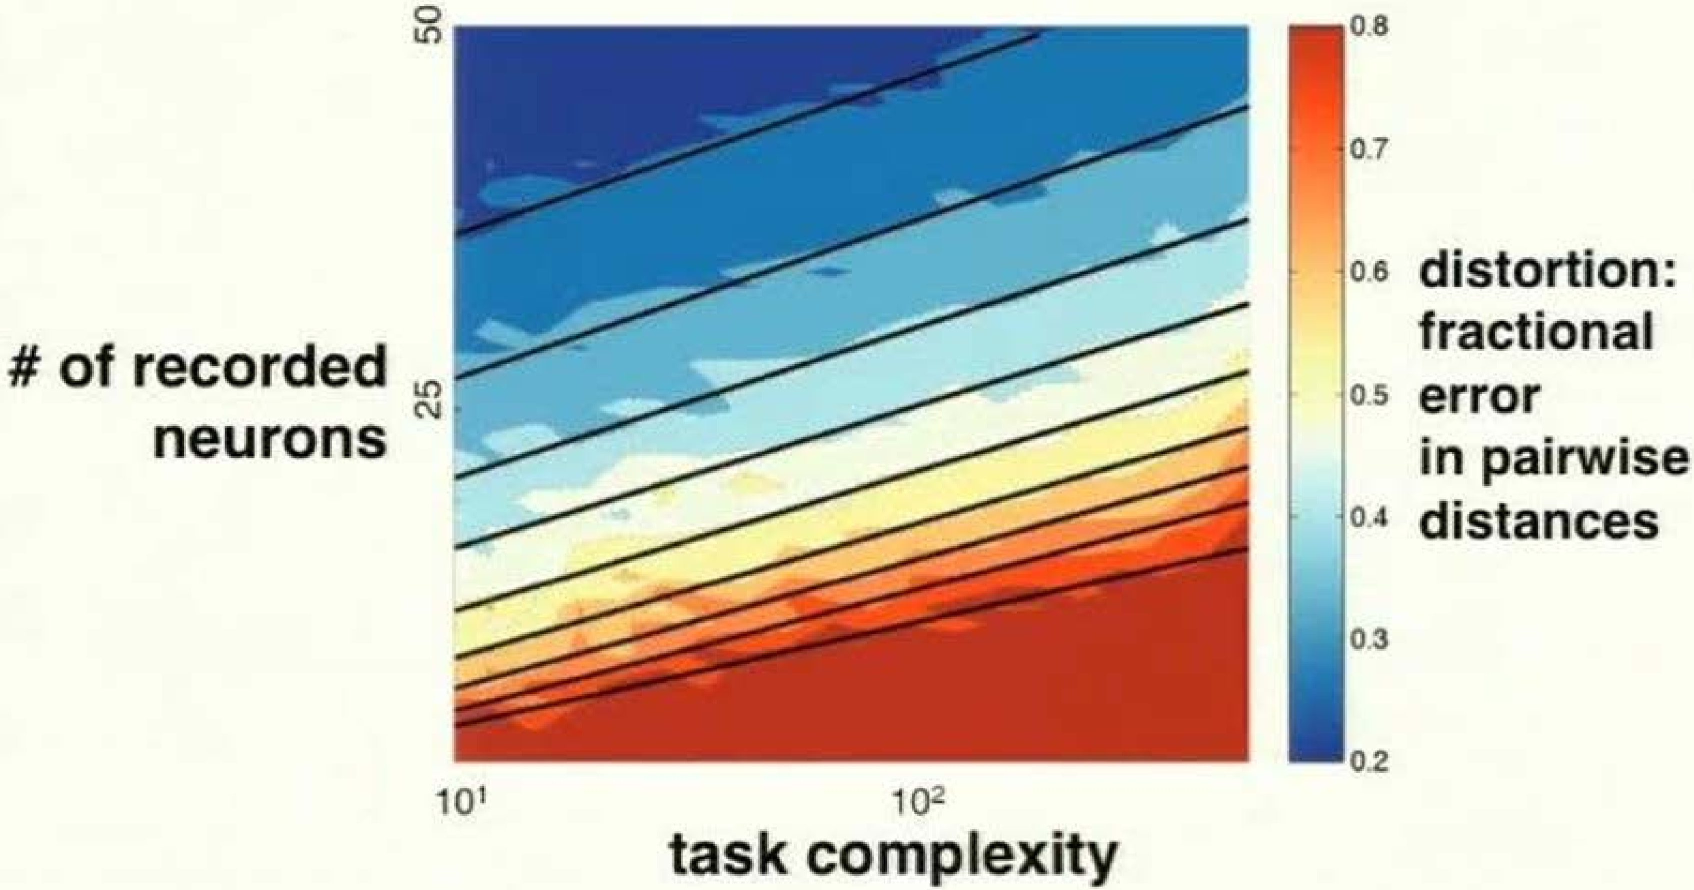

distortion contours of motor cortical data

A consequence of neuronal measurement as a random projection

By adapting random projection theory:

$$\textbf{\# neurons needed} = \frac{1}{\text{distortion}^2}\left(c_1 \, \textbf{log(task complexity)} + c_2\right)$$

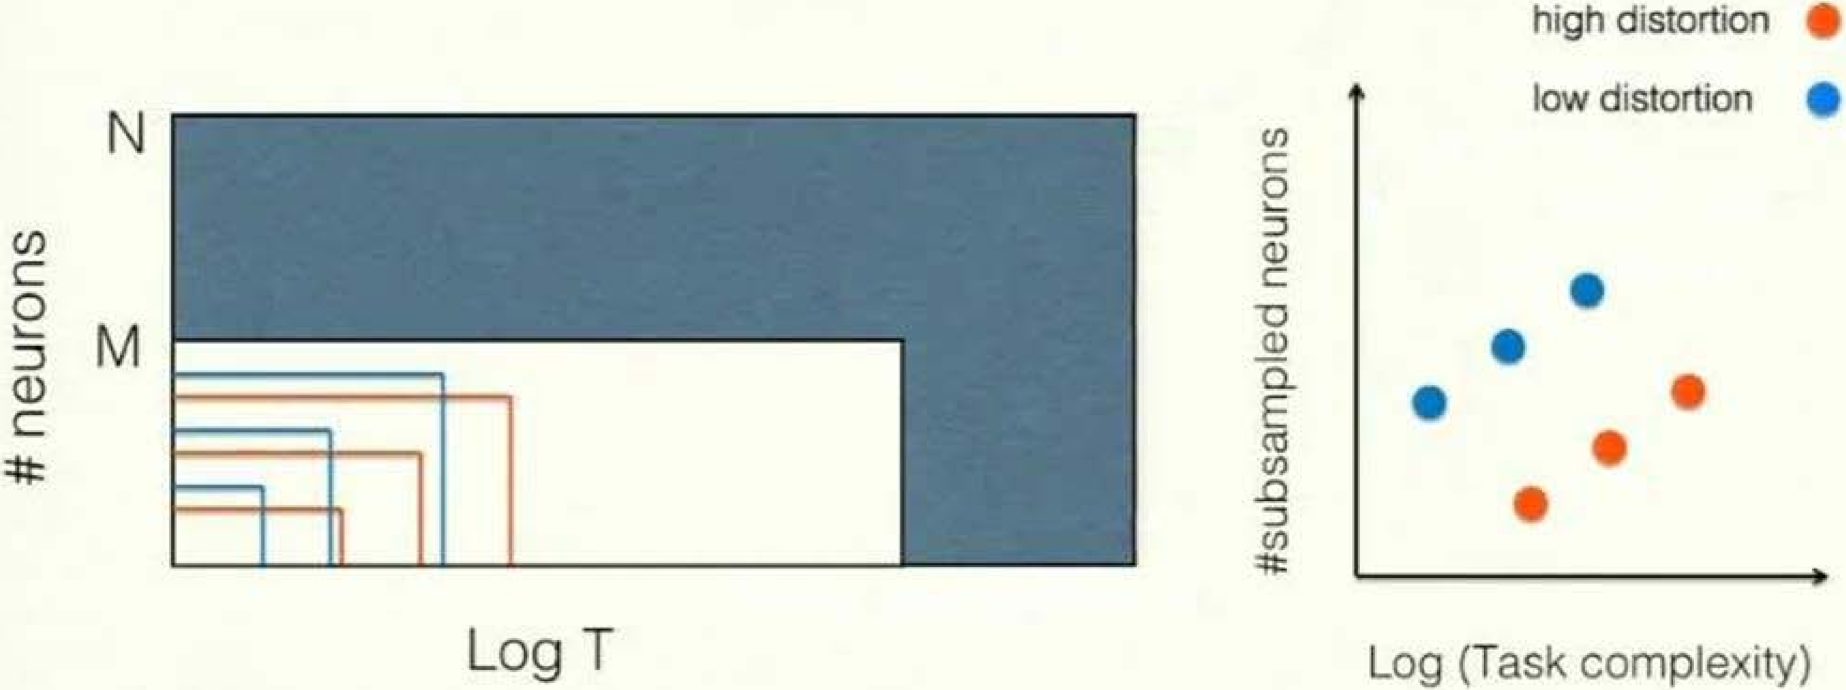

To keep the same level of desired distortion, # of neurons need only scale logarithmically with task complexity (good news!)

Static Decoding

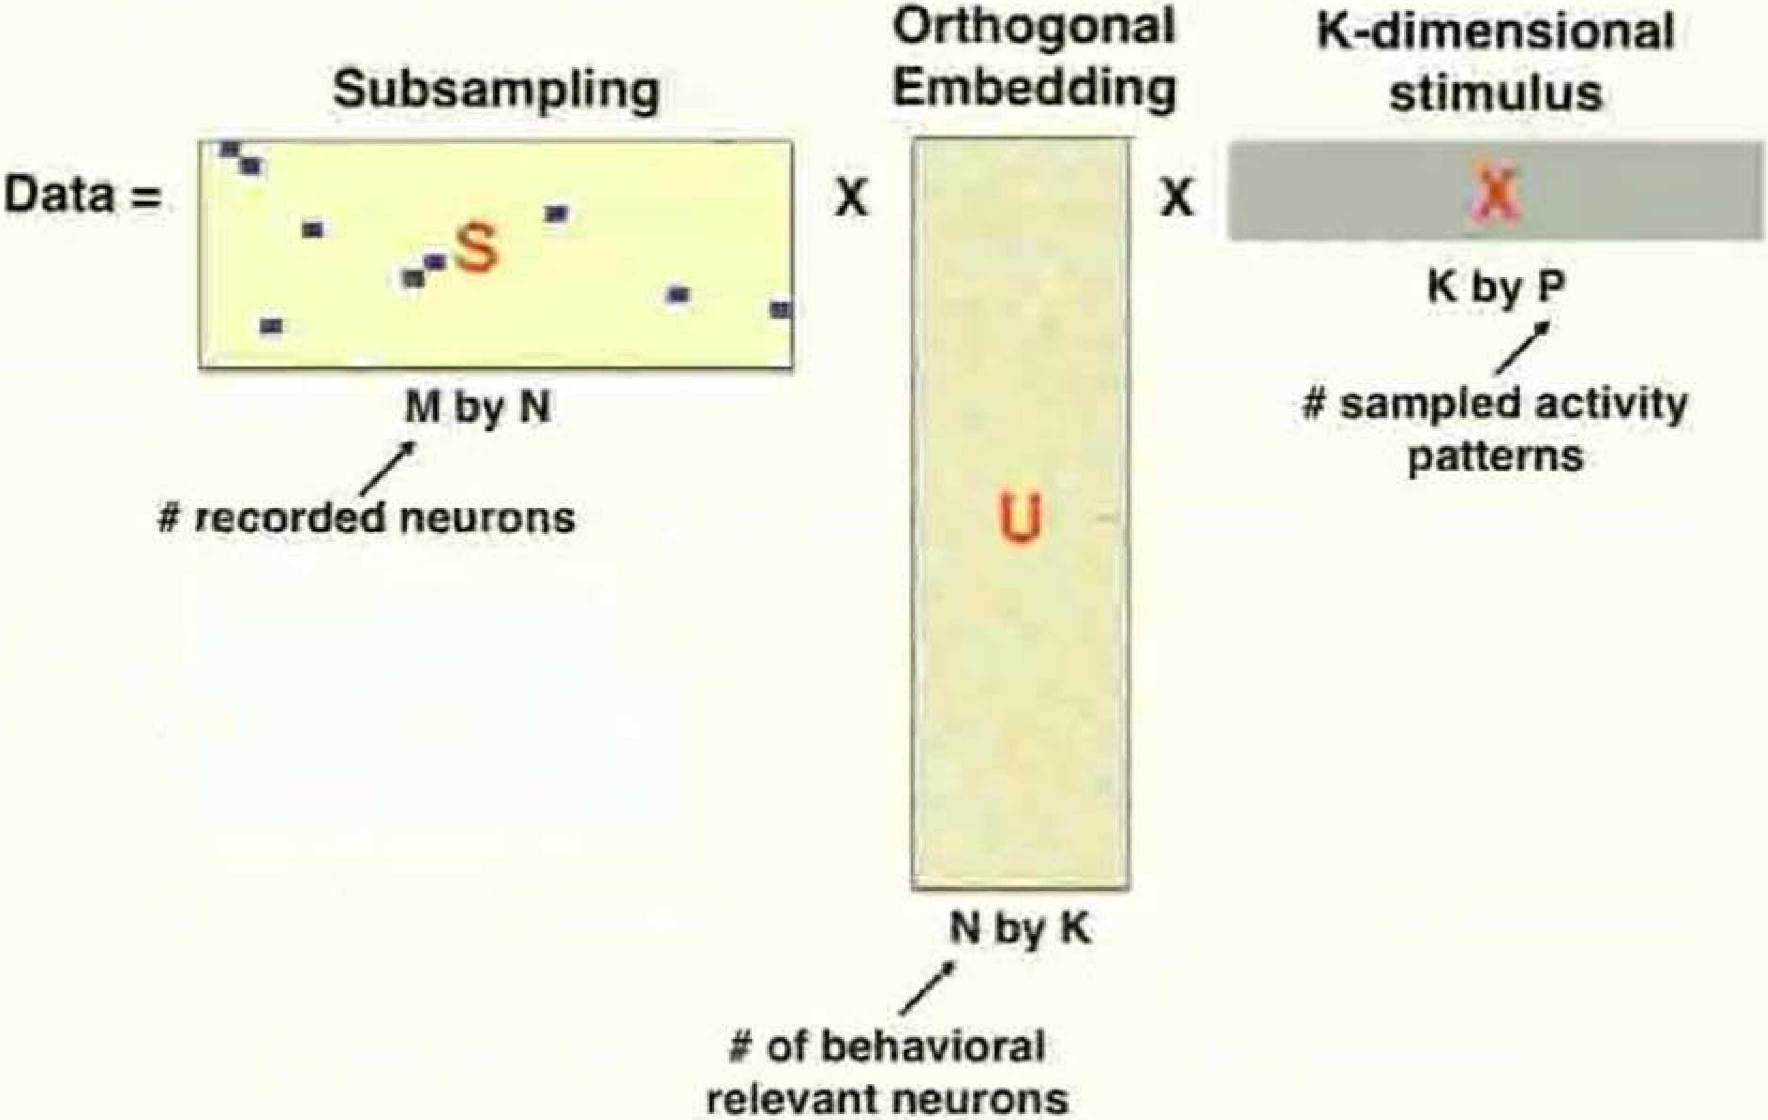

Subsampling

Orthogonal Embedding

K-dimensional stimulus

Data = S X U X X

M by N

recorded neurons

N by K

# of behavioral relevant neurons

K by P

sampled activity patterns

Subsampling

With partial observation (i.e. M < N):

$$R = SUX + Z$$

low-rank signal

high-dim noise

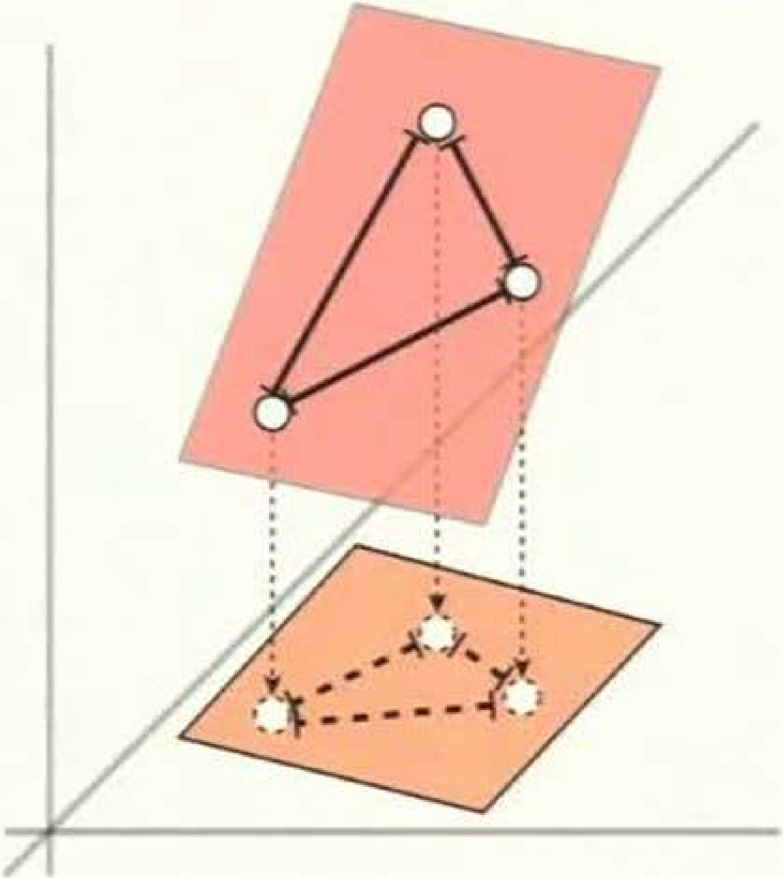

- K(=2)-dimensional stimulus space
- Embedded in N(=3)-dimensional neural space
- Subsampled to M(=2)-dimensional subspace
- Distance between sampled activity patterns are compressed
- Compressions are different depending on orientations
- Compression determined by the K singular values of SU

Static Decoding - Recovering Dimensionality

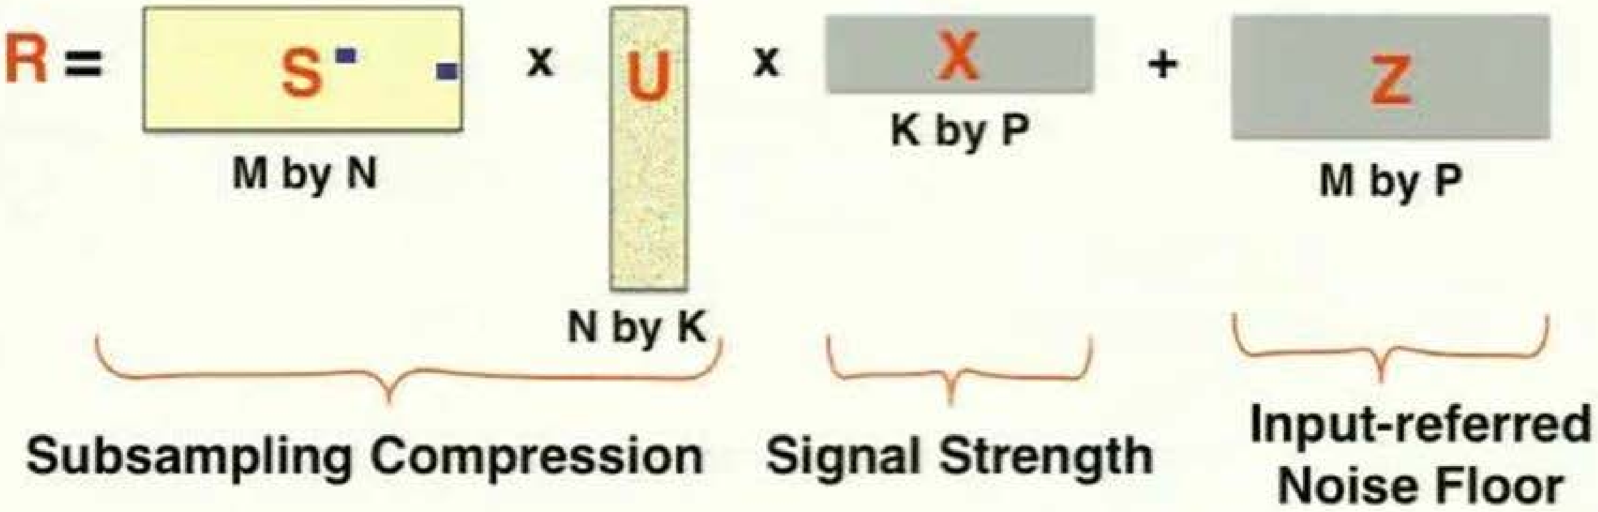

$$R = \underbrace{\underset{M \text{ by } N}{S} \times \underset{N \text{ by } K}{U}}_{\text{Subsampling Compression}} \times \underbrace{\underset{K \text{ by } P}{X}}_{\text{Signal Strength}} + \underbrace{\underset{M \text{ by } P}{Z}}_{\text{Input-referred Noise Floor}}$$

Static Decoding - Recovering Dimensionality

Subsampling Compression $\star$ **Signal Strength** $>$ **Input-referred Noise Floor**

(worst-case) (worst-case)

Simulations with N = 5000 and K = 20

Inferred dimensionality as # singular values > noise floor

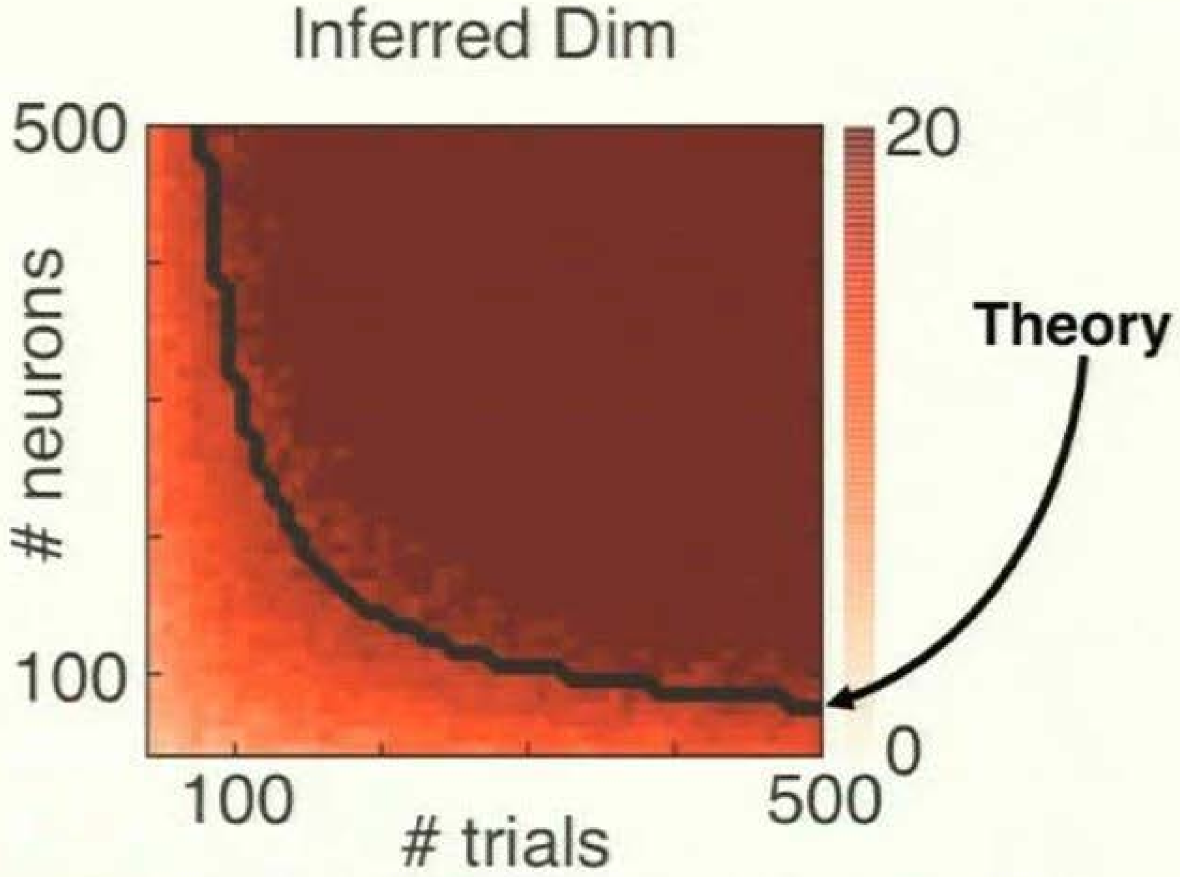

Static Decoding

Subsampling Compression * **Signal Strength** > **Input-referred Noise Floor**
(worst-case) (worst-case)

Simulations with N = 5000 and K = 20
Linear decoding using recovered signal in inferred subspace

Gavish & Donoho 2013

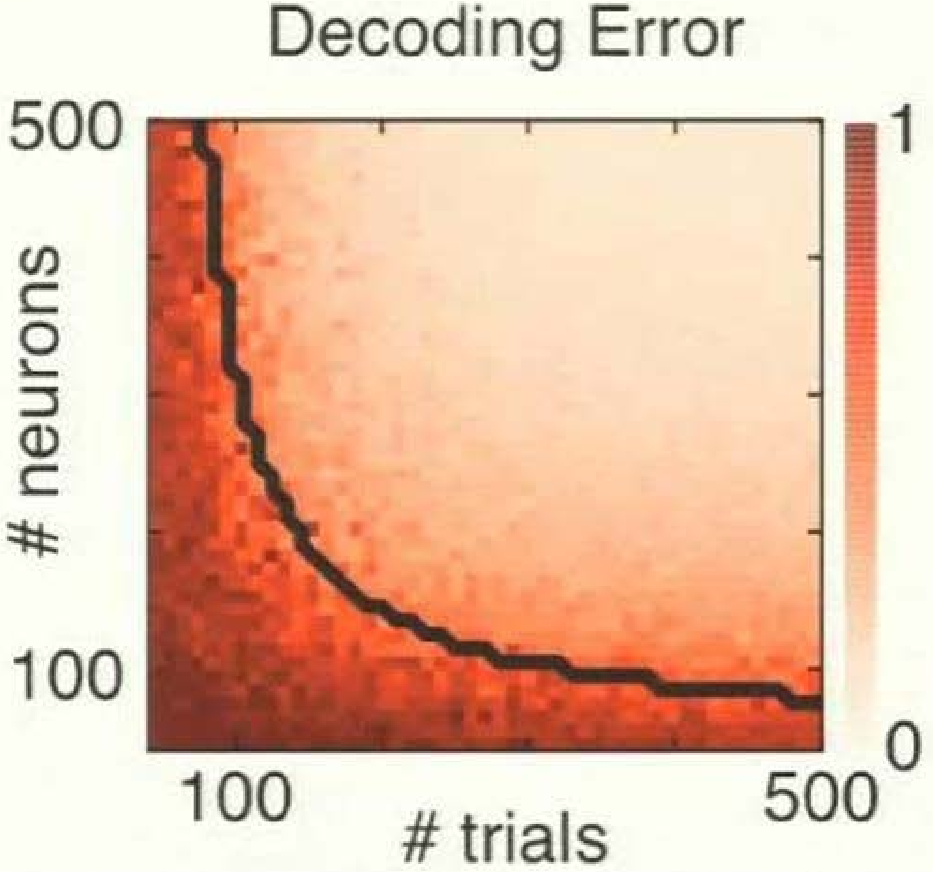

To understand the spectrum of the covariance matrix, we factorize the data

M by N ← # recorded neurons

N by N

N by T ← # of behavioral relevant neurons; # of time points recorded

$X = \text{sampling} * (X_{slow} + X_{fast})$

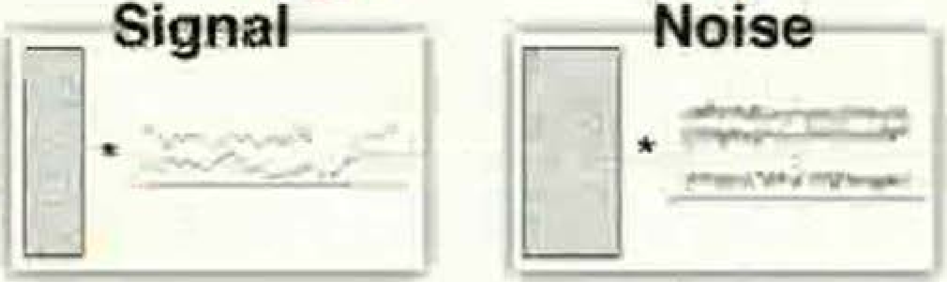

Similar setup to factor analysis, but the noise is correlated across time

Data can be thought of as a low-rank perturbation of a random noise matrix

$X_{slow} + X_{fast}$ Benaych-Georges & Nadakuditi, 2012

Eigenvalue spectrum of correlated noise deviates from the Marchenko-Pastur law

Theoretical eigenvalue spectrum for N = 1000, T = 2000 noise matrix

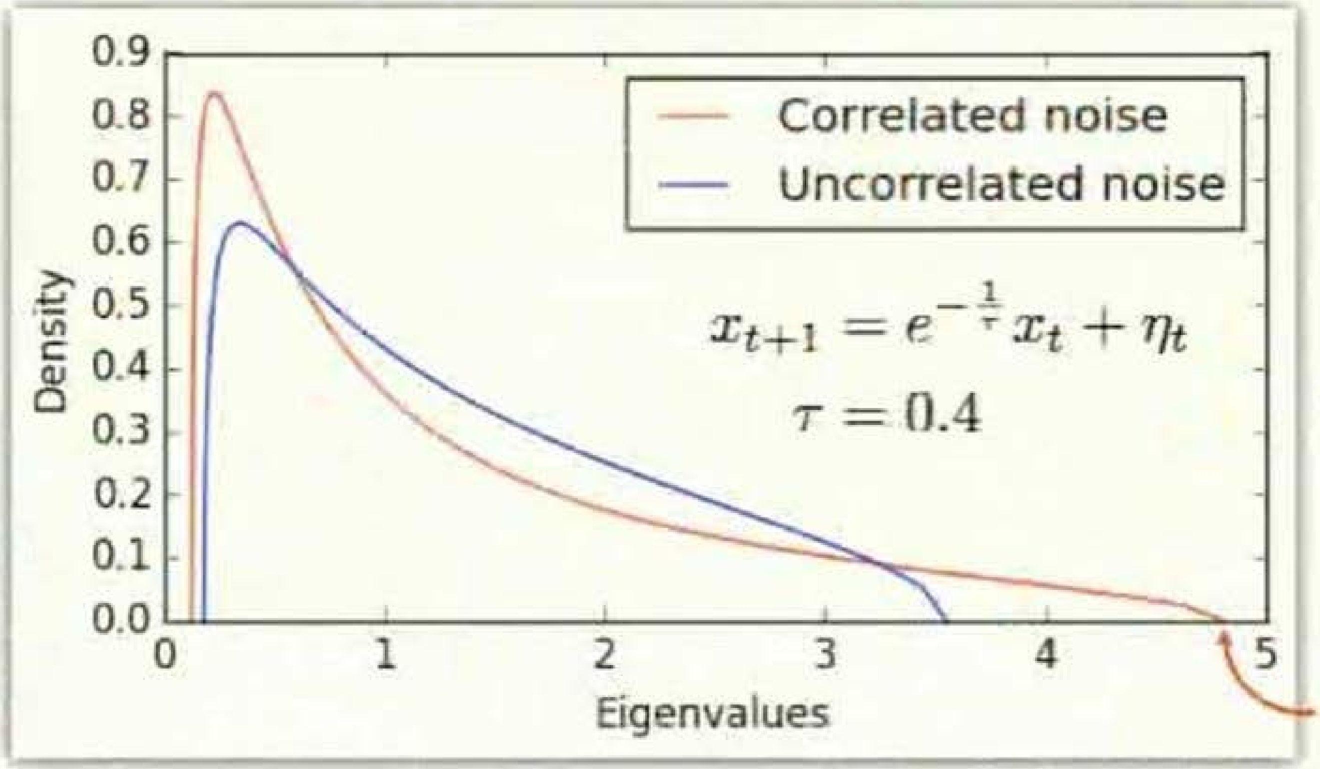

Marchenko & Pastur, 1967
Bai et al., 2008, Yao, 2014

b(N, T, τ)
noise floor

What does it take to get random neural manifolds?

A sufficient condition: every neuron has complex tuning for every task parameter.

Old paradigm:
Single units

The bane of existence for those thinking along the lines of single unit neurophysiology.

We cannot easily understand and classify single neurons ☹.

New paradigm:
Collective behavior

The saving grace of our ability to understand the brain!

With random trajectories, we can record from a relatively small number of neurons and infer the correct state space description of neural data!

Understanding what individual neurons do becomes the wrong question. We should focus instead on the collective.

Acknowledgements

Ganguli Lab

Peiran Gao
Subhaneli Lahiri
Jascha Sohl-Dickstein
Niru Maheswaranathan
Madhu Advani
Ben Poole
Jay Sarkar
Kiah Hardcastle

Shenoy Lab

Krishna Shenoy
Eric Trautmann
Byron Yu
Gopal Santhanam
Stephen Ryu

Funding

Bio-X Neuroventures
Burroughs Wellcome
Genentech Foundation
James S. McDonnell Foundation
National Science Foundation
Office of Naval Research
Simons Foundation
Sloan Foundation
Swartz Foundation