Low Dimensional Manifolds for Neural Dynamics

Sara A. Solla
Northwestern University

Dynamical Principles of Biological and Artificial Neural Networks
Banff International Research Station
January 10, 2022

Multi Electrode Arrays (MEAs)

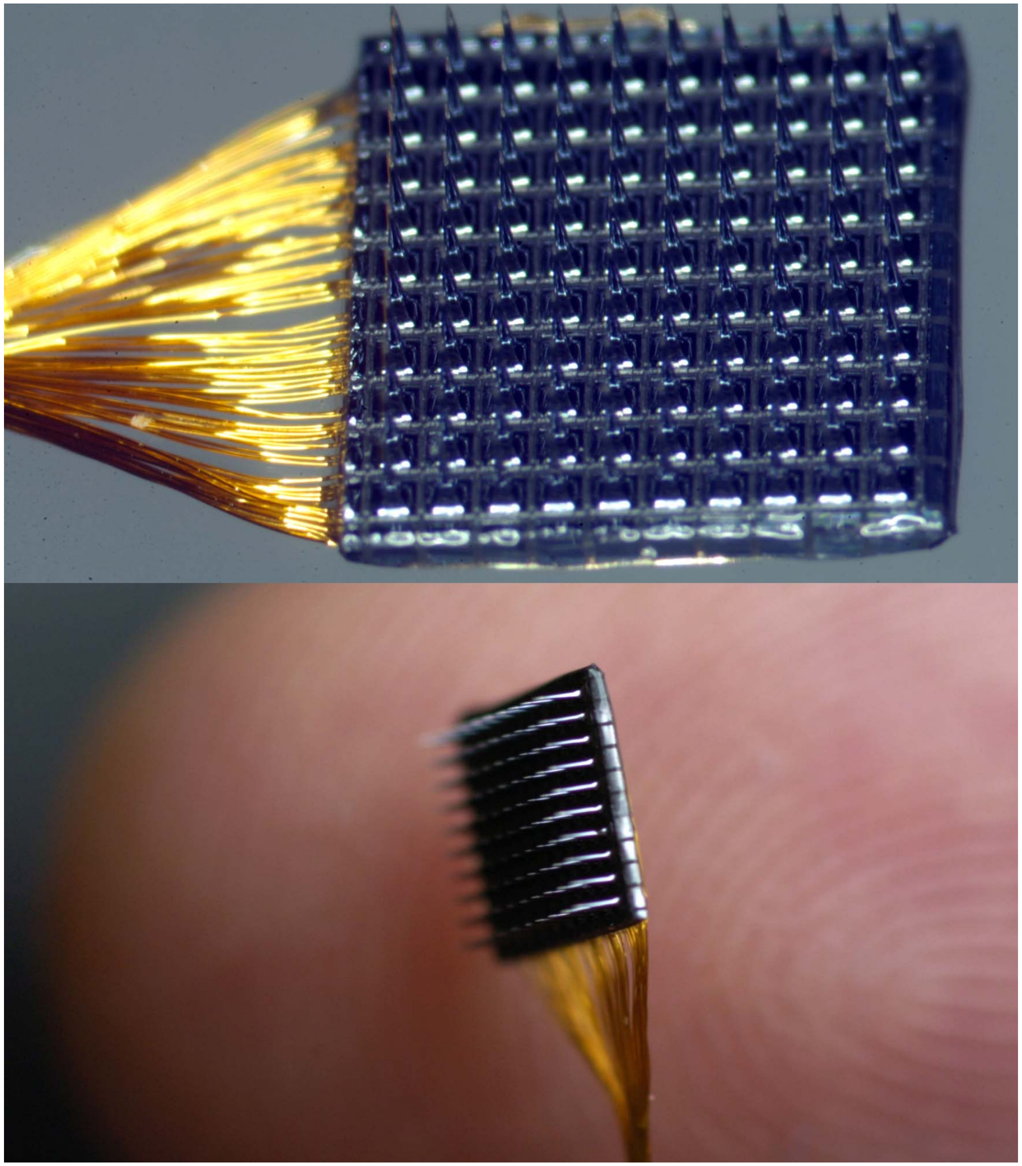

Population activity

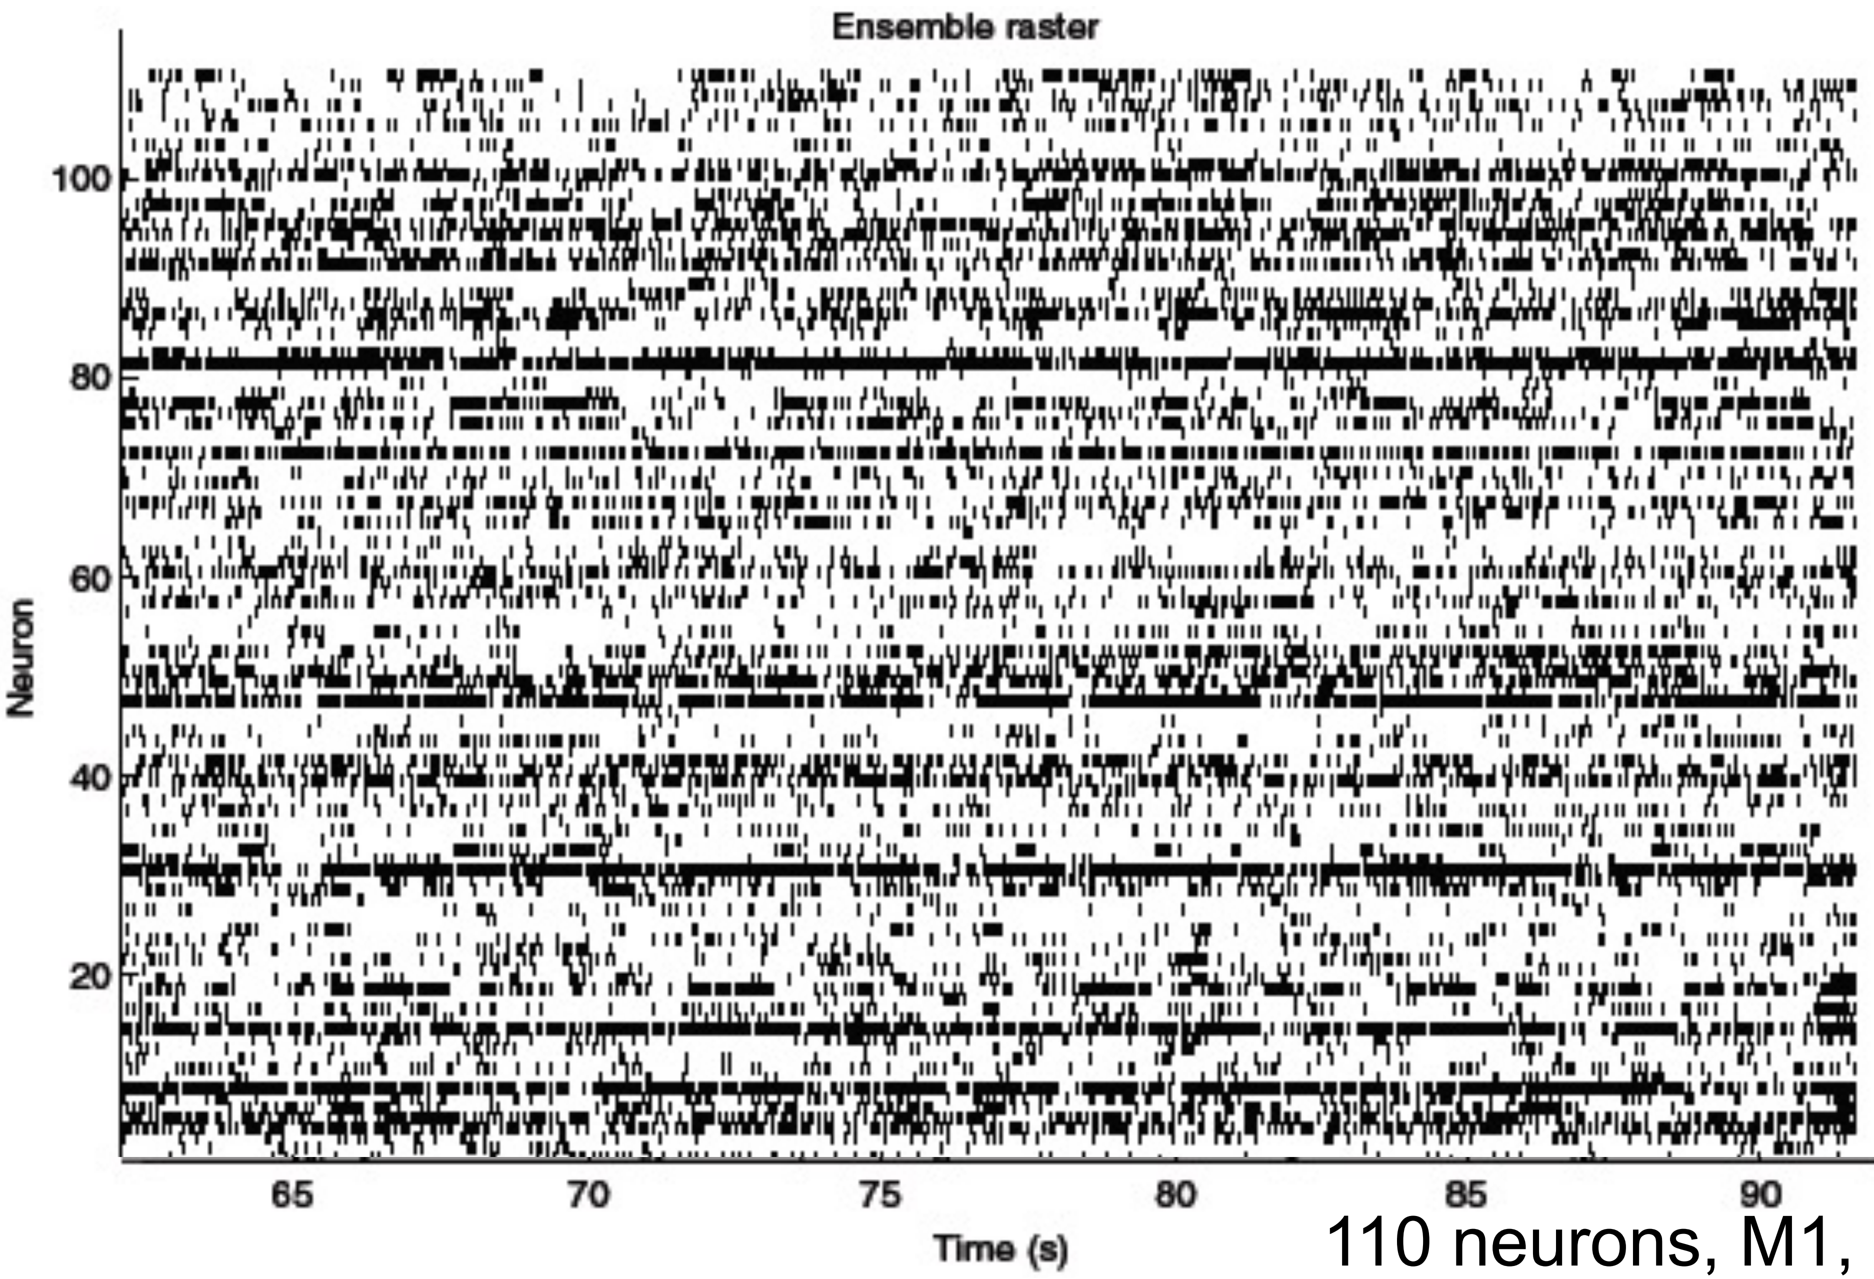

110 neurons, M1, hS3

Analysis of population activity

- Consider a population of N neurons whose spiking activity is observed during a time interval $(0,T]$.
- The interval is divided into K bins of size $\Delta = T/K$, labeled by an index $1 \leq k \leq K$.
- In each time bin k we observe the number of spikes $n_i(k)$ emitted by neuron i, for all $1 \leq i \leq N$.

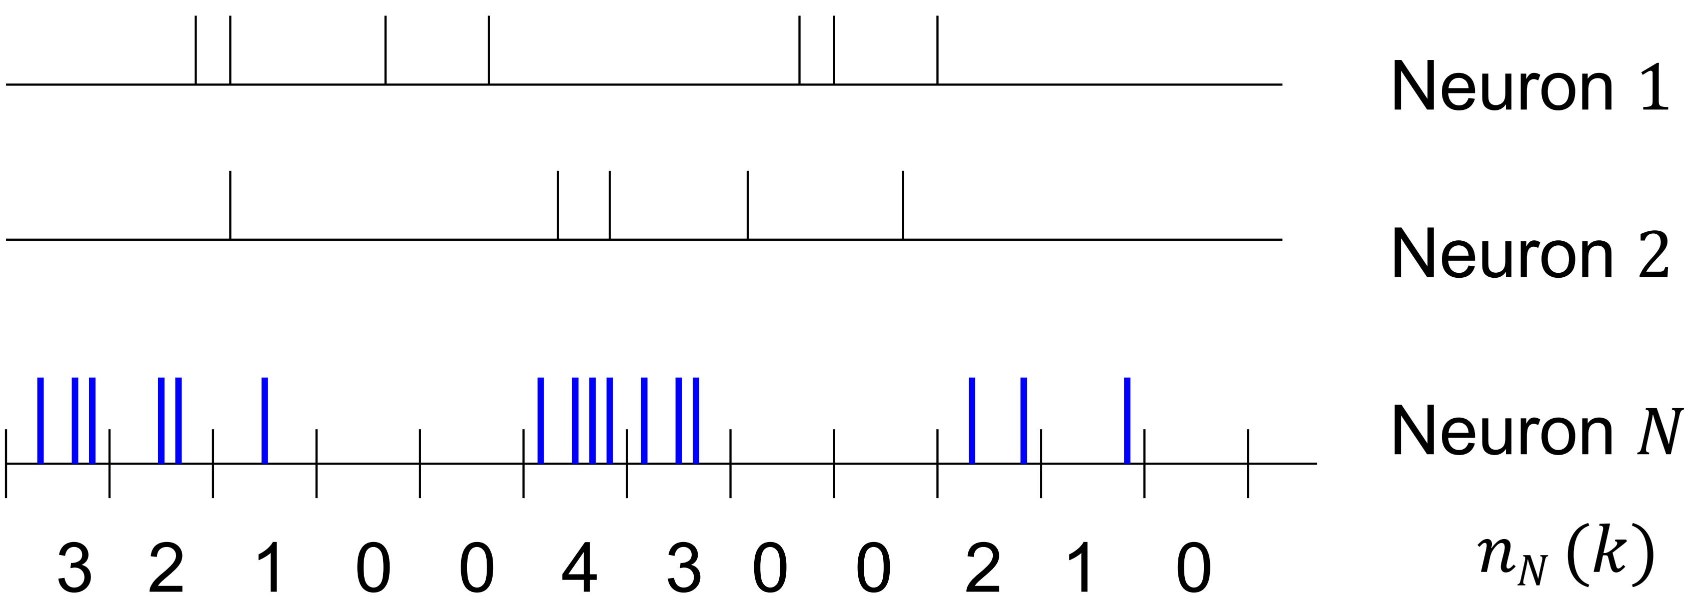

Population activity

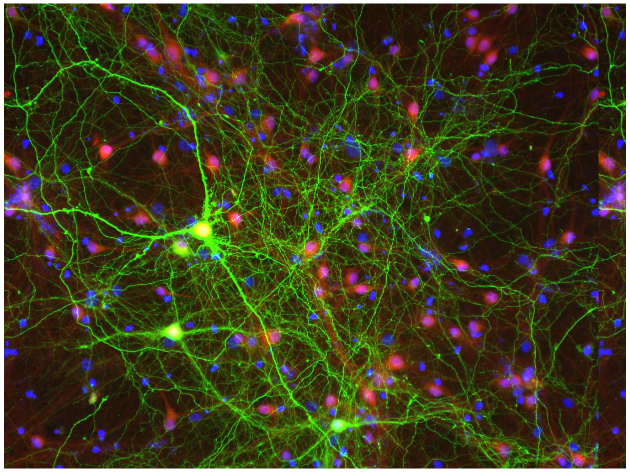

$$X_\infty = \begin{bmatrix} N_1^{t+1} & N_1^{t+2} & & N_1^{t+T} \\ N_2^{t+1} & N_2^{t+2} & \cdots & N_2^{t+T} \\ \vdots & \vdots & & \vdots \\ N_\infty^{t+1} & N_\infty^{t+2} & \cdots & N_\infty^{t+T} \end{bmatrix}$$

Data matrix X_∞ has N_∞ rows and T columns

T is the duration of the experiment in units of bin size Δ

For a reach movement, $N_\infty \approx 10^6$ in M1

Neural population activity

- A simple motor task

- Neural manifolds for motor control

- The unreasonable effectiveness of linear methods

A simple motor task: center-out reaches

Instructed delay center-out reaching task

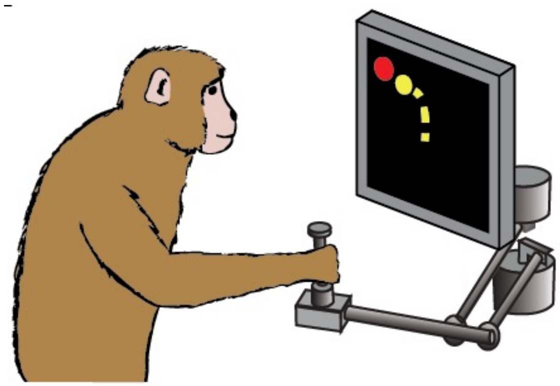

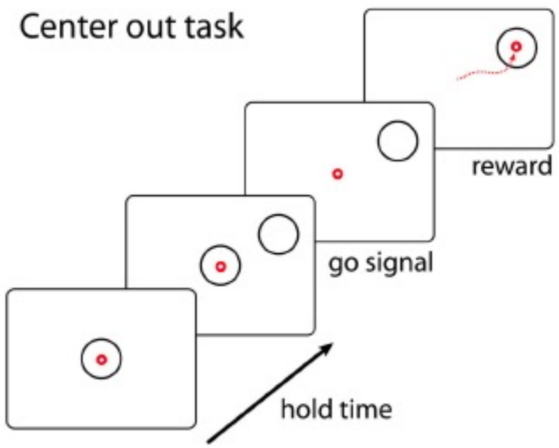

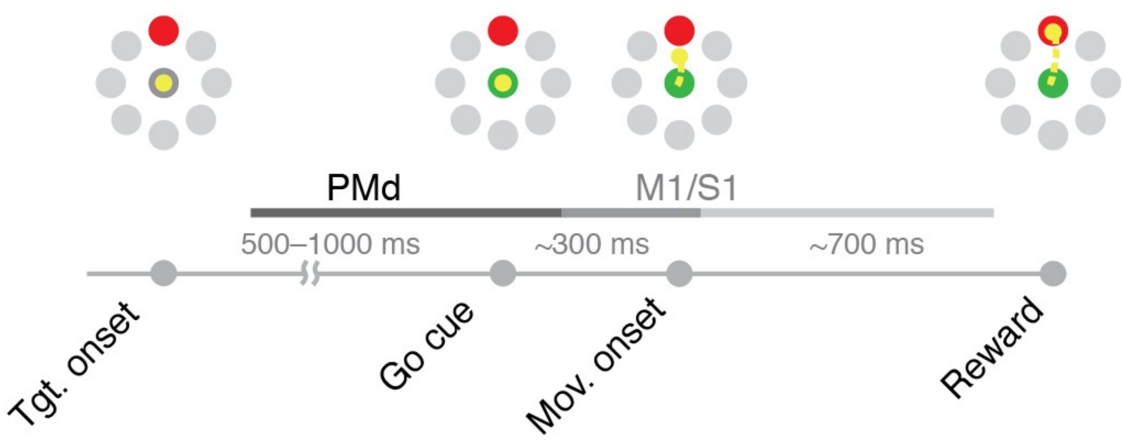

Neural activity: variability and specificity

Single neuron recordings

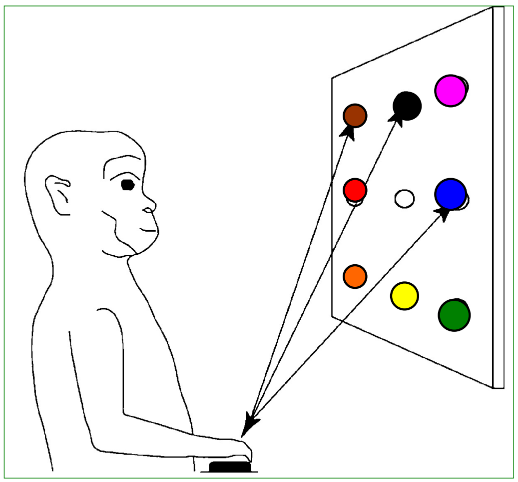

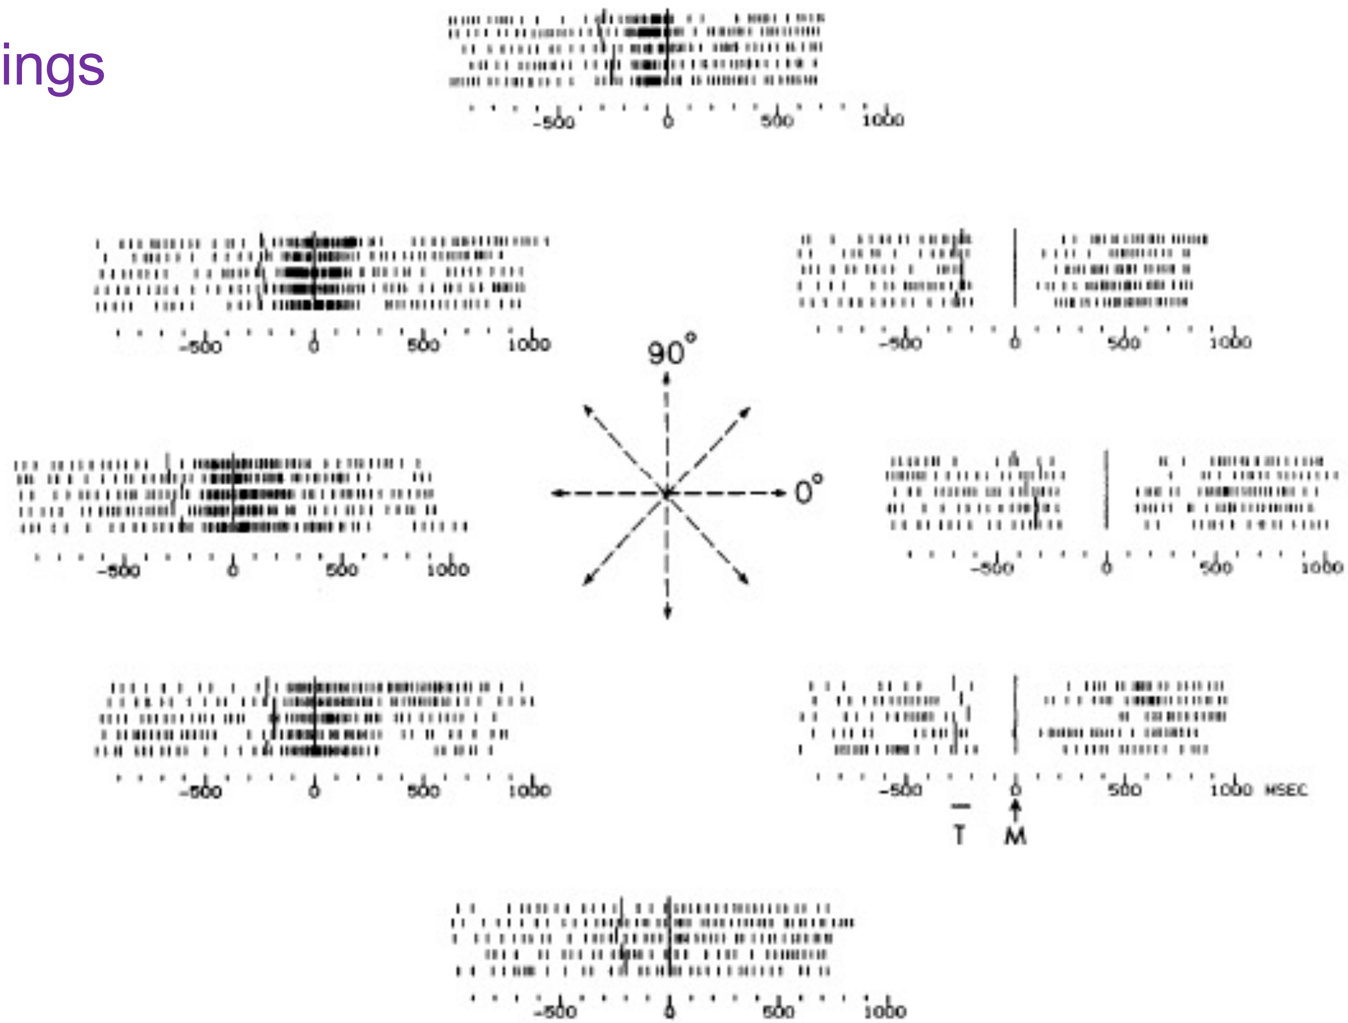

Georgopoulos, Kalaska, Caminity, Massey, *J of Neurosci* (1982)

Population activity : multiple targets

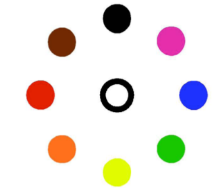

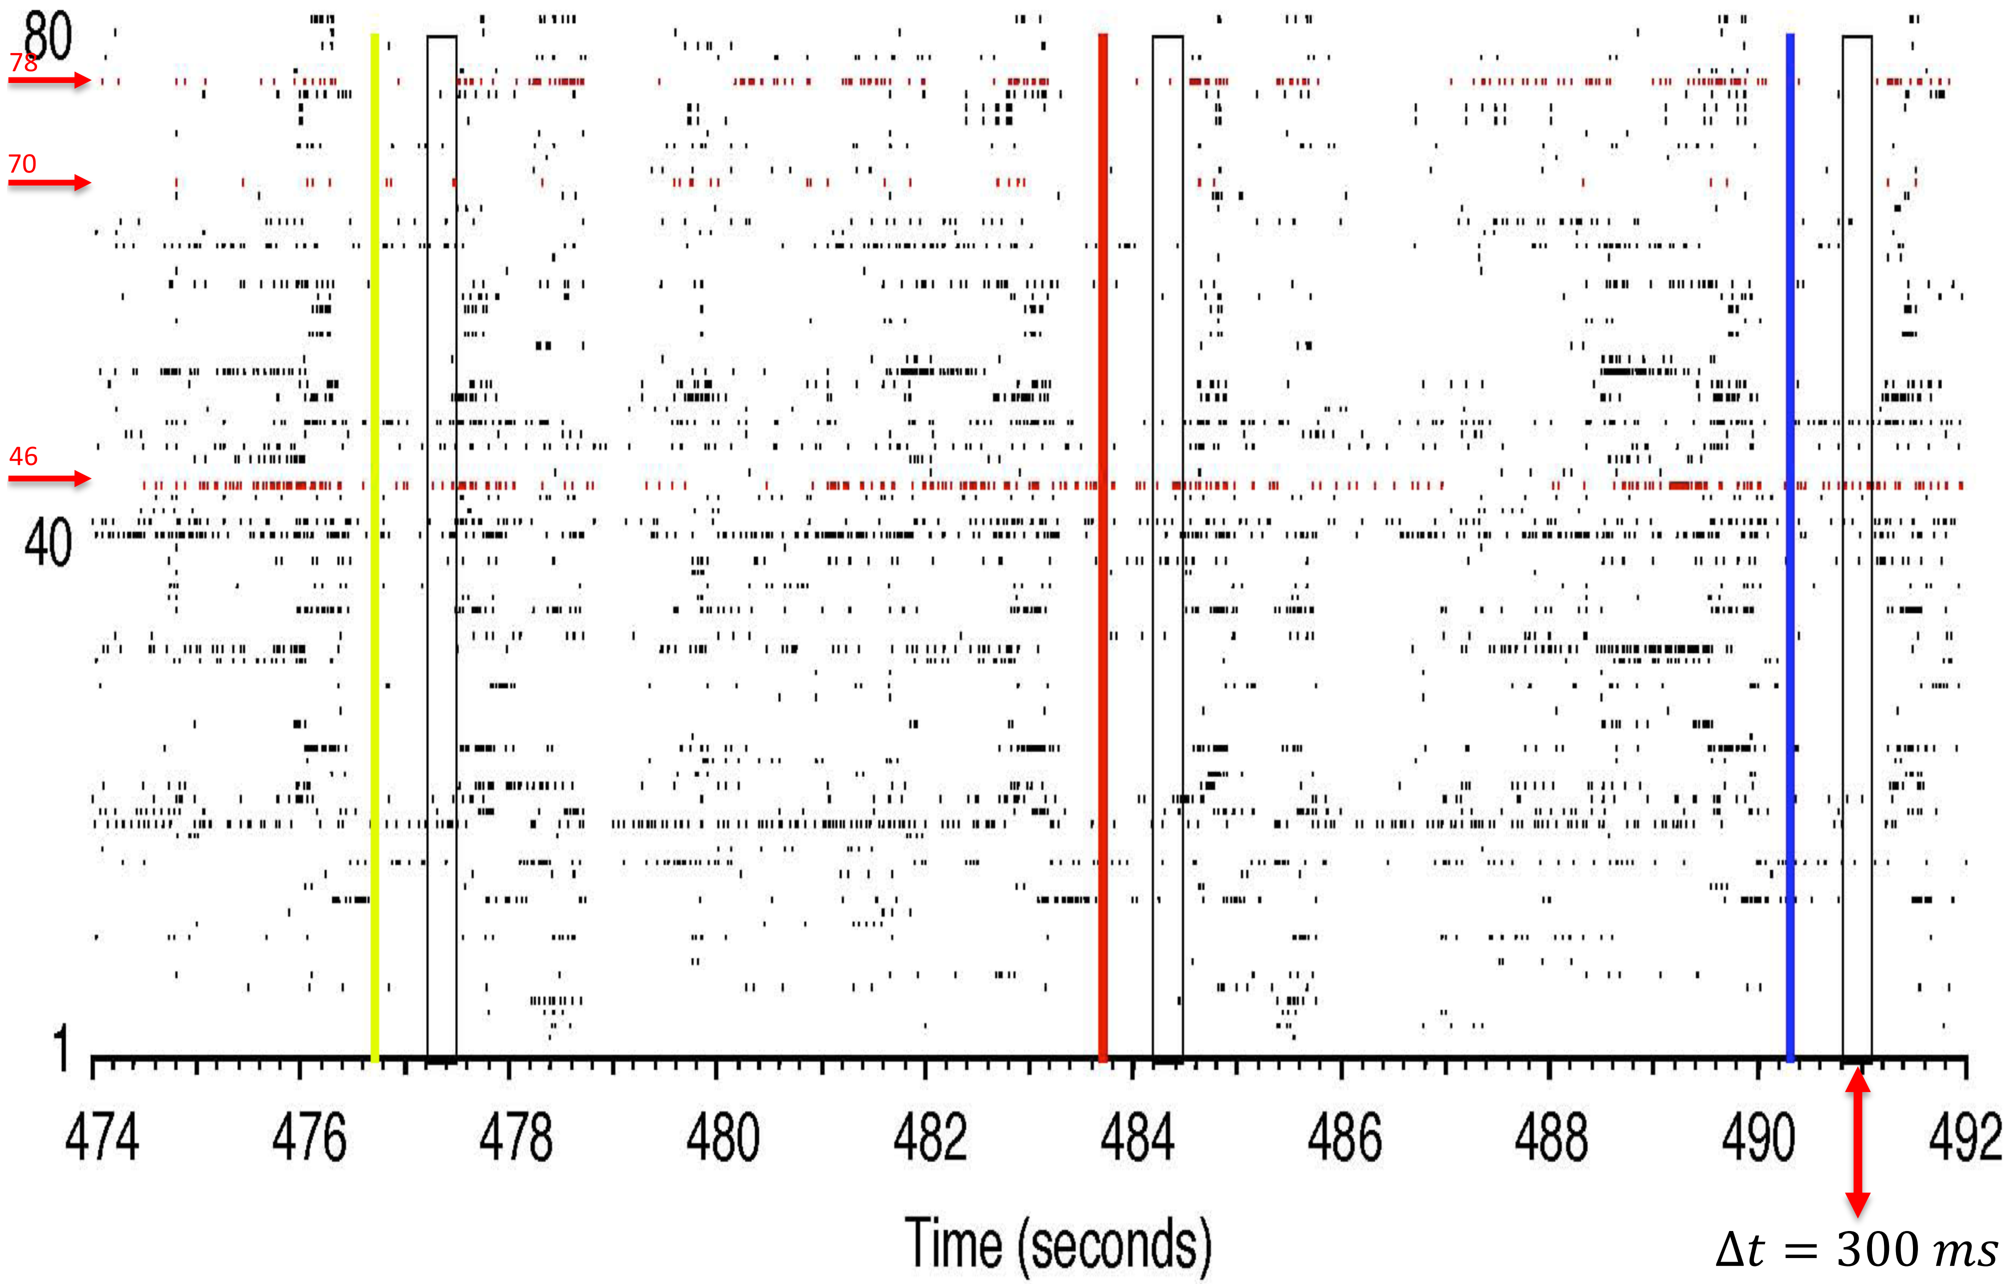

Target-dependent population activity

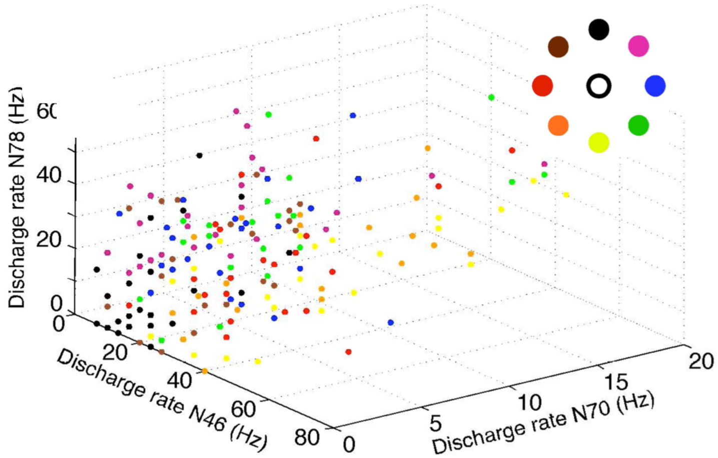

Neural modes: directions in neural space

Specific patterns of populations activity

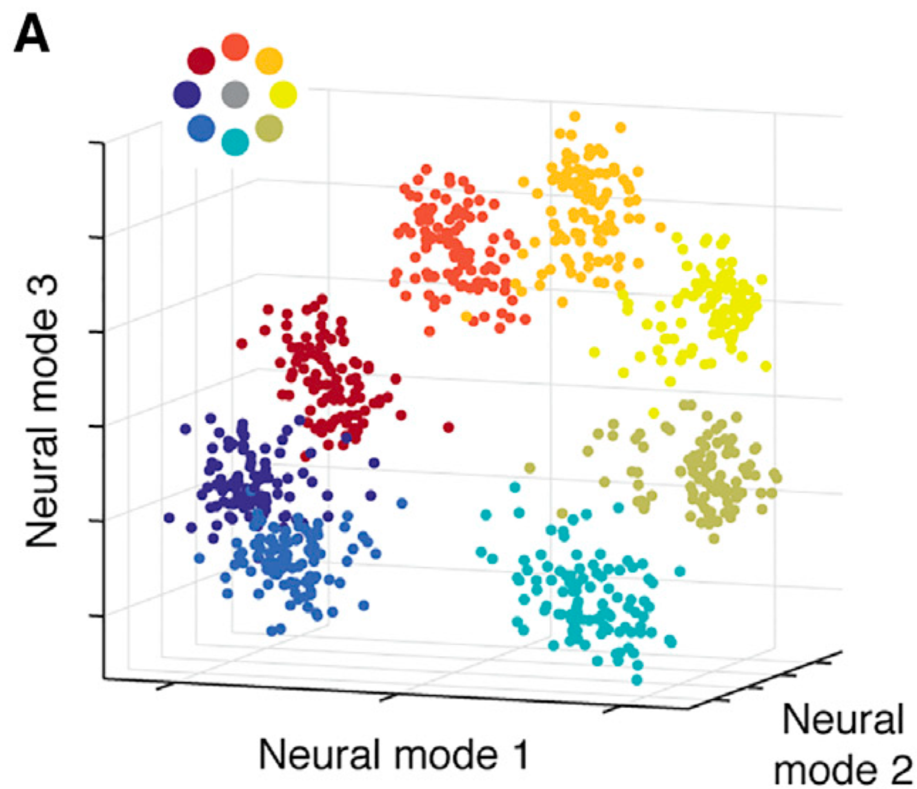

Santhanam, Yu, Gilja, Ryu, Afshar, Sahani, Shenoy, *J Neurophysiol* (2009)

Neural population activity

- A simple motor task
- Neural manifolds for the control of movement
- The unreasonable effectiveness of linear methods

Population activity: Subsampling

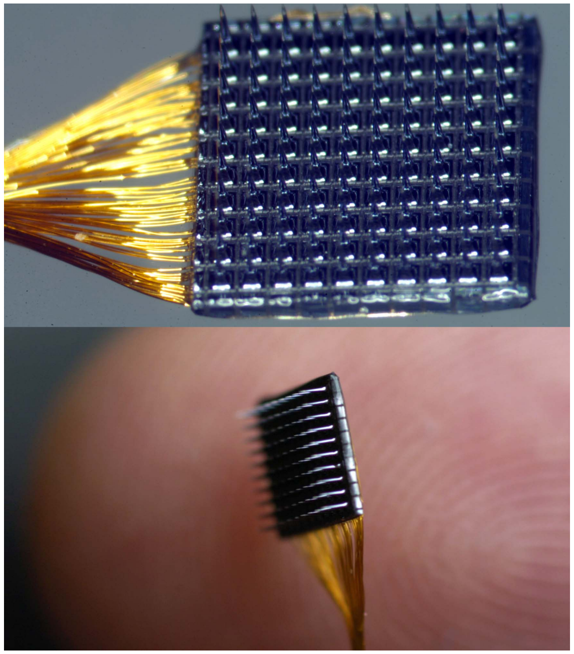

$$X_D = \begin{bmatrix} N_1^{t+1} & N_1^{t+2} & & N_1^{t+T} \\ N_2^{t+1} & N_2^{t+2} & \cdots & N_2^{t+T} \\ \vdots & \vdots & & \vdots \\ N_D^{t+1} & N_D^{t+2} & \cdots & N_D^{t+T} \end{bmatrix}$$

Data matrix X_D has N_D rows and T columns

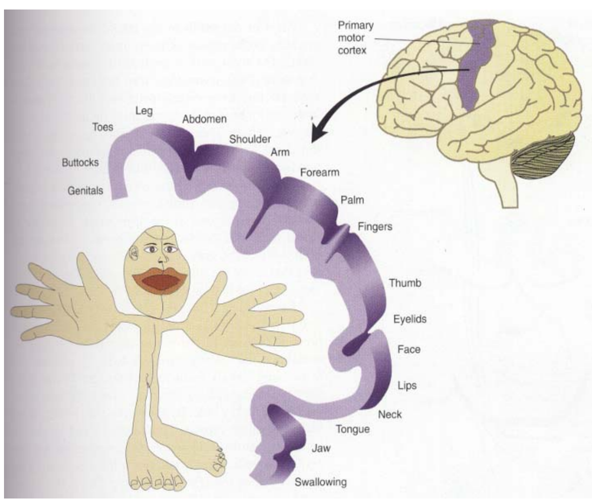

$D \approx 10^2$ for Multi-Electrode Arrays (MEAs)

$D \approx 10^3$ for Neuropixels

D is the ambient dimension

Population dynamics: the empirical neural space

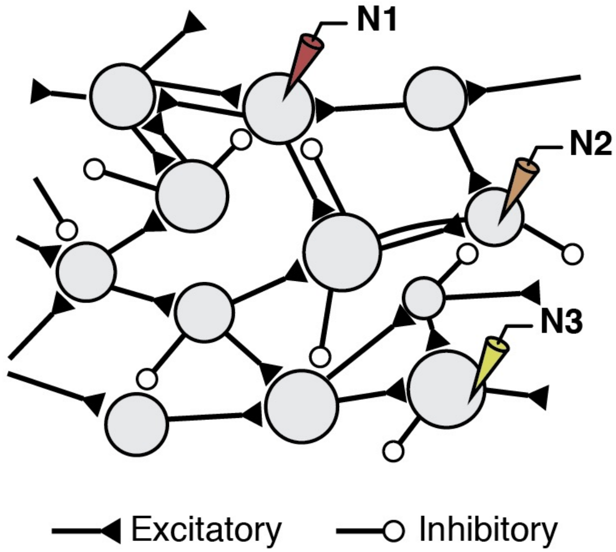

Ambient dimension D:
Number of recorded neurons

Population dynamics: the empirical neural space

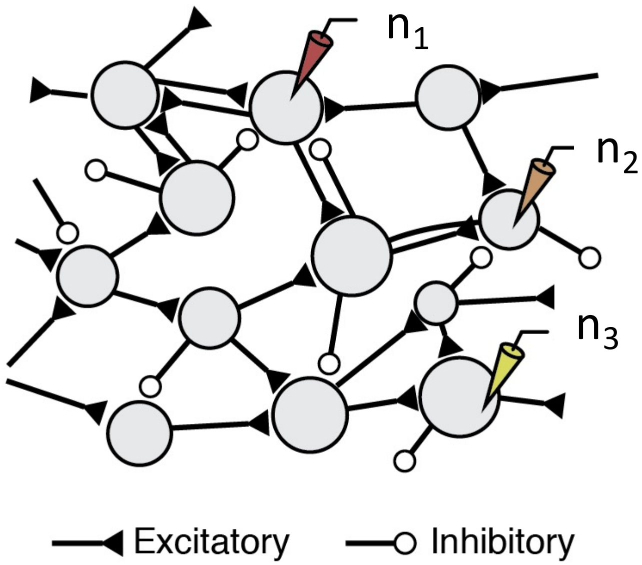

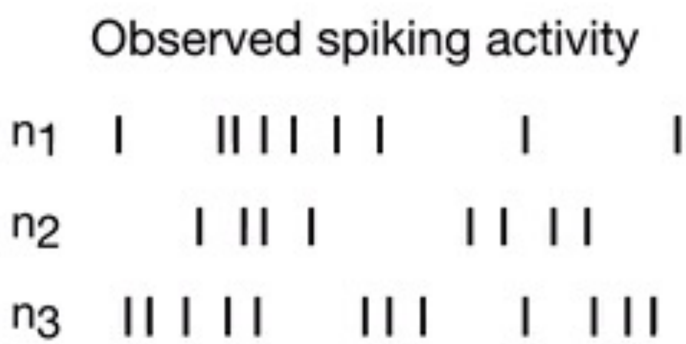

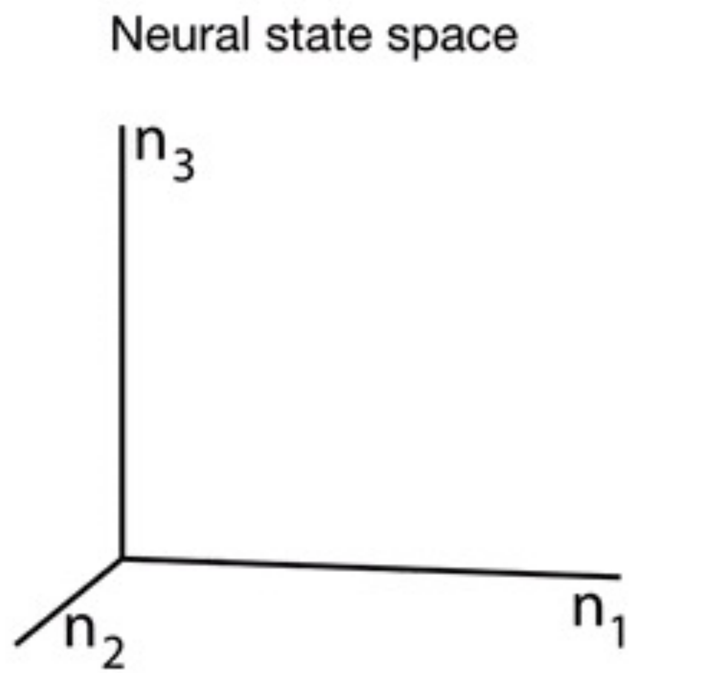

Cunningham & Yu, *Nat Neuro* 2014

Dimensionality reduction: neural modes and latent variables

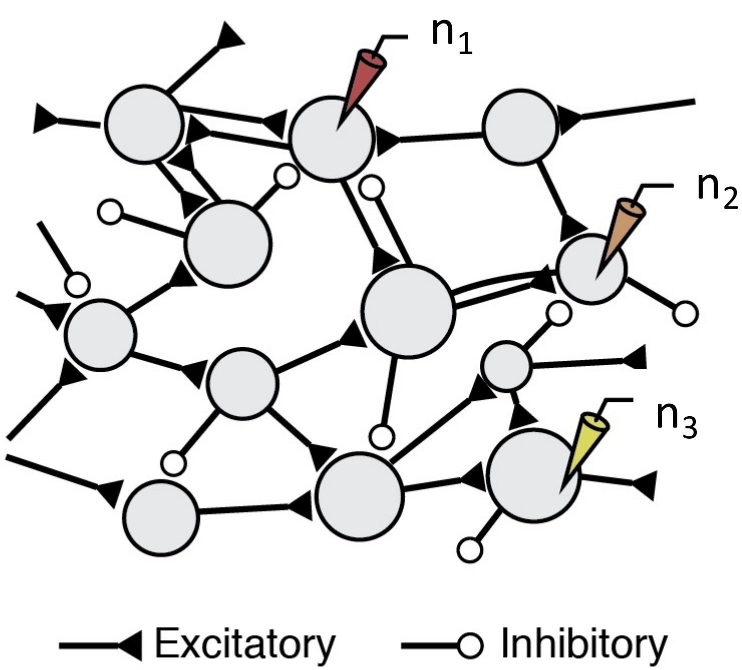

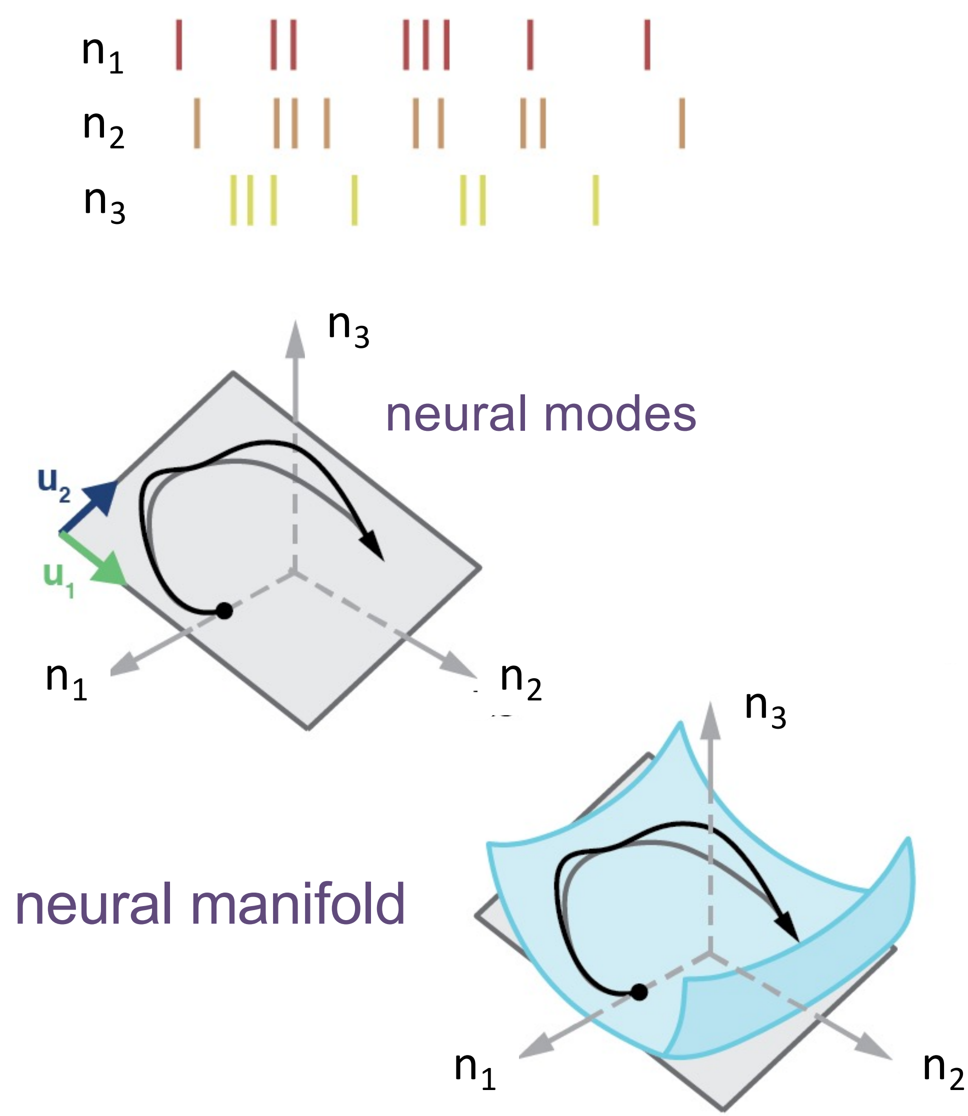

DIMENSIONALITY REDUCTION
linear or nonlinear?

Dimensionality reduction: neural modes and latent variables

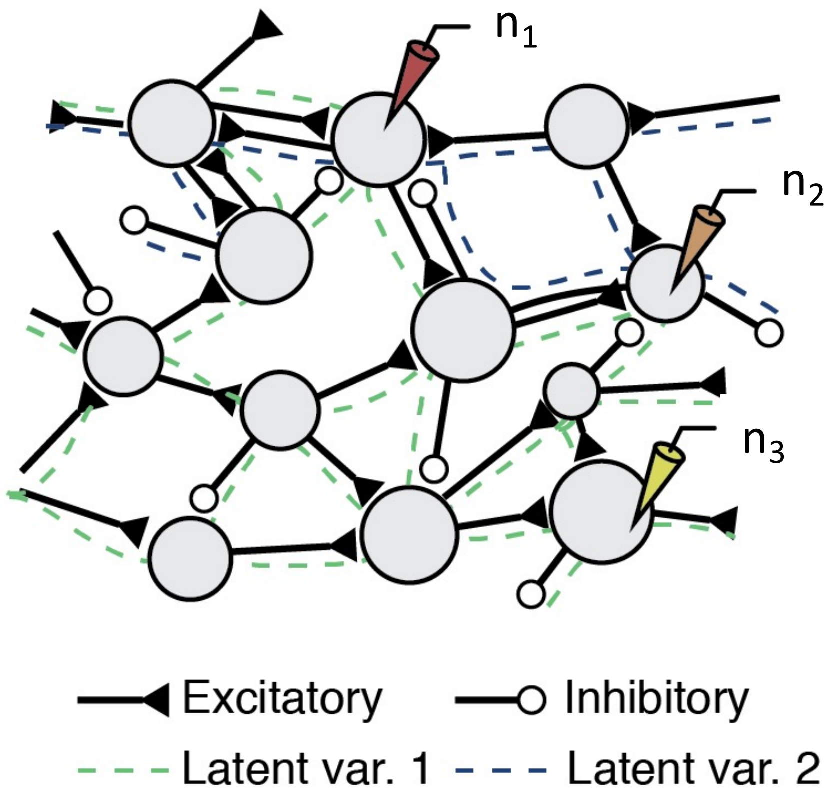

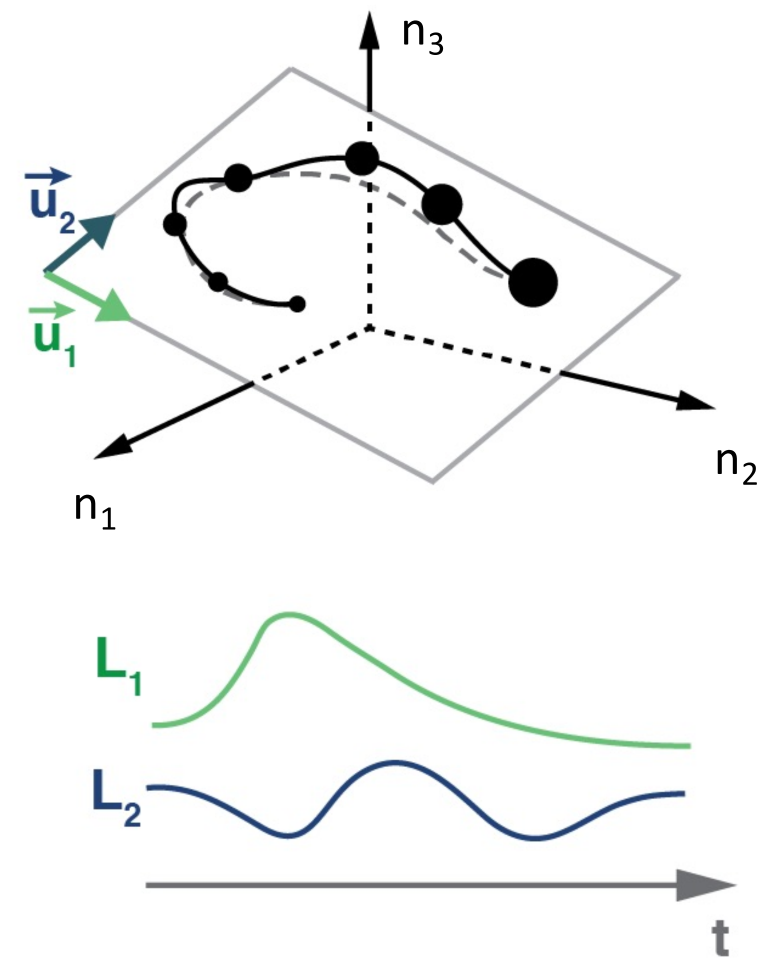

Population dynamics: latent variables as a generative model

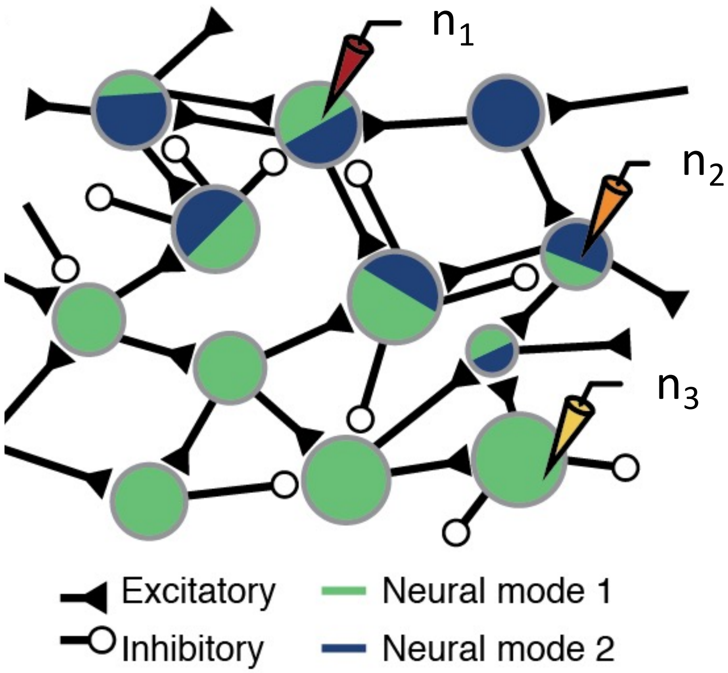

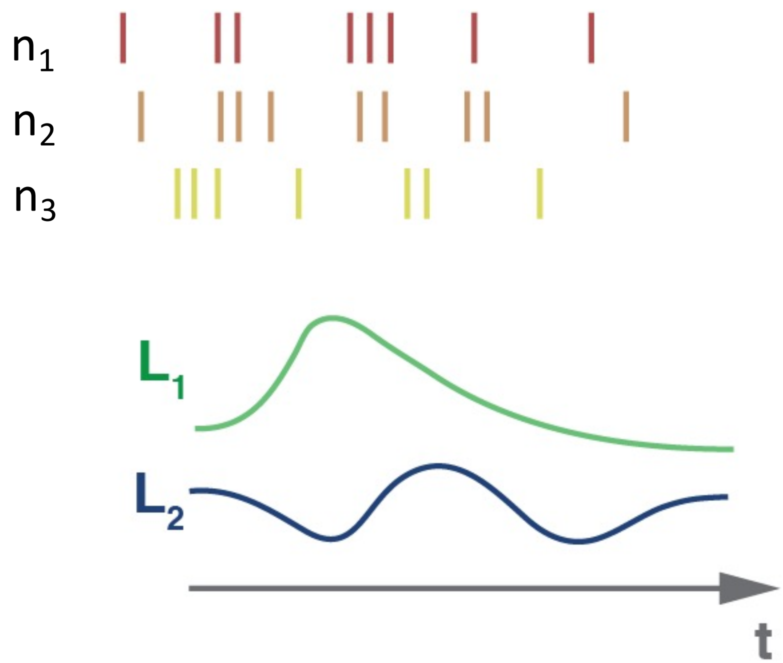

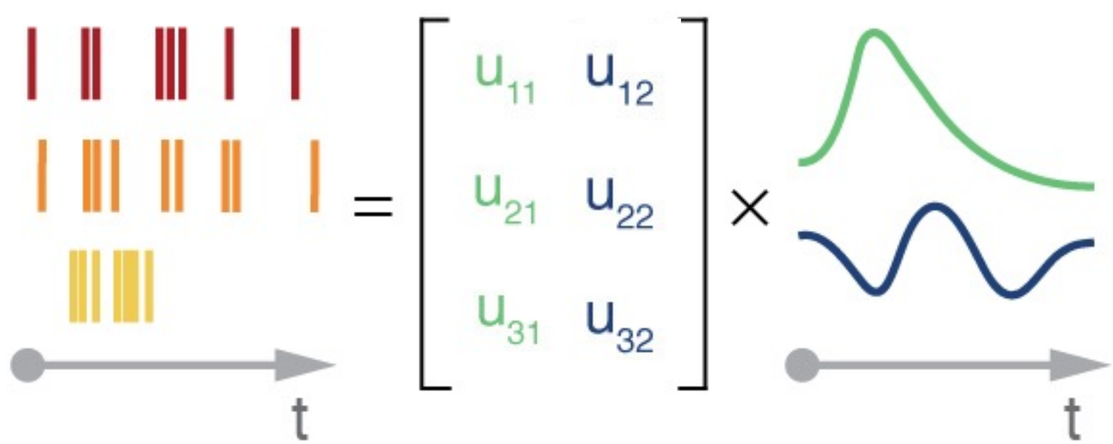

Neural population activity

- A simple motor task
- Neural manifolds for the control of movement
- The unreasonable effectiveness of linear methods

Behavioral stability

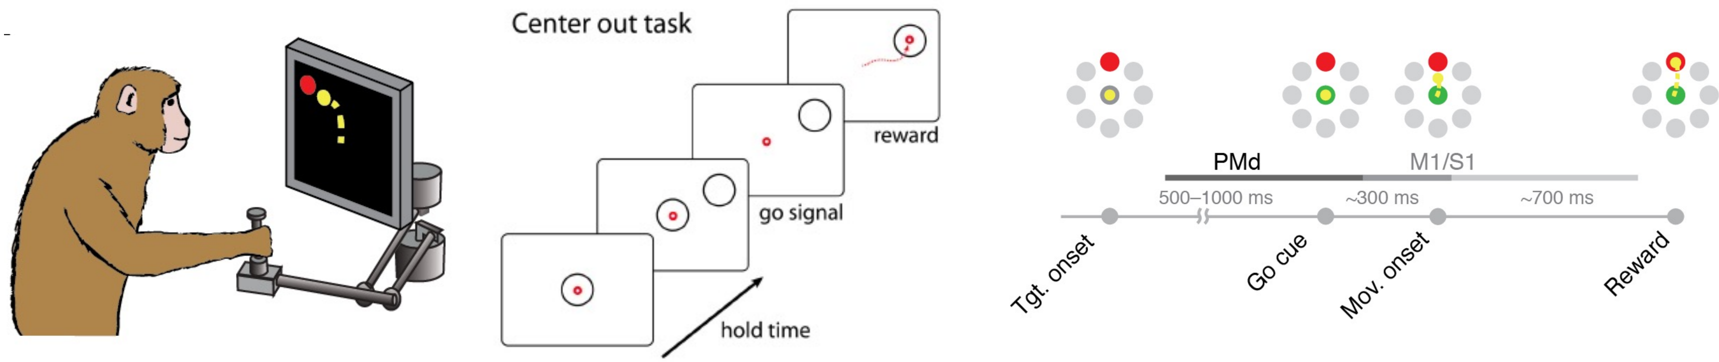

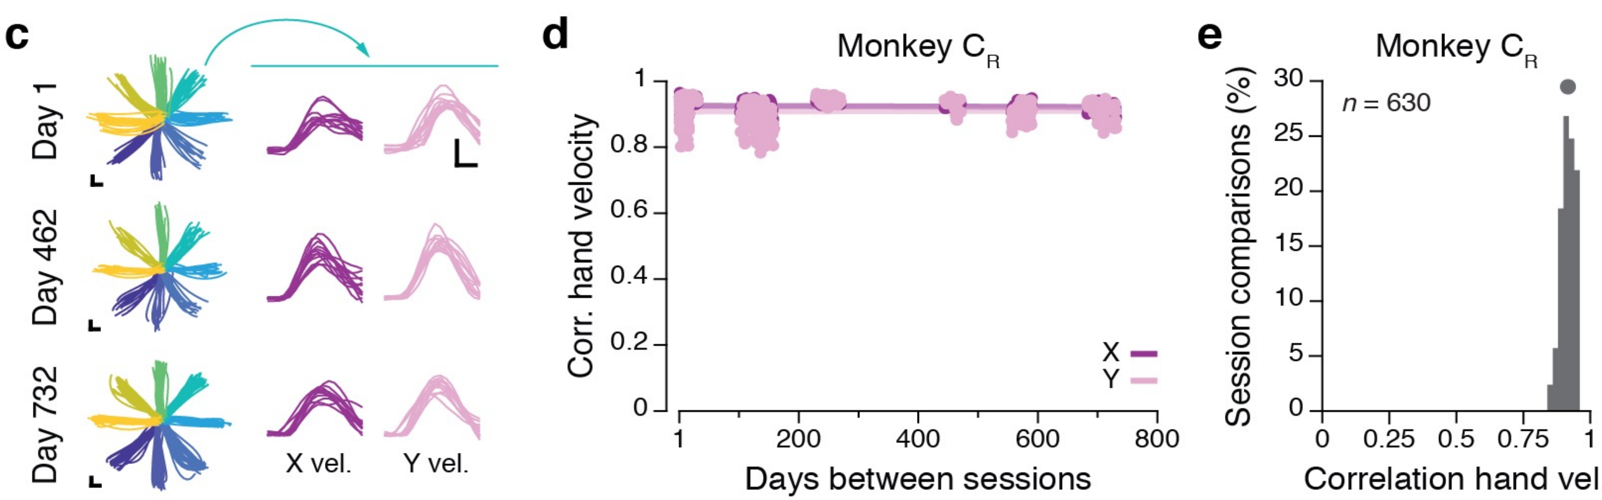

Stable neural dynamics underlying consistent behavior?

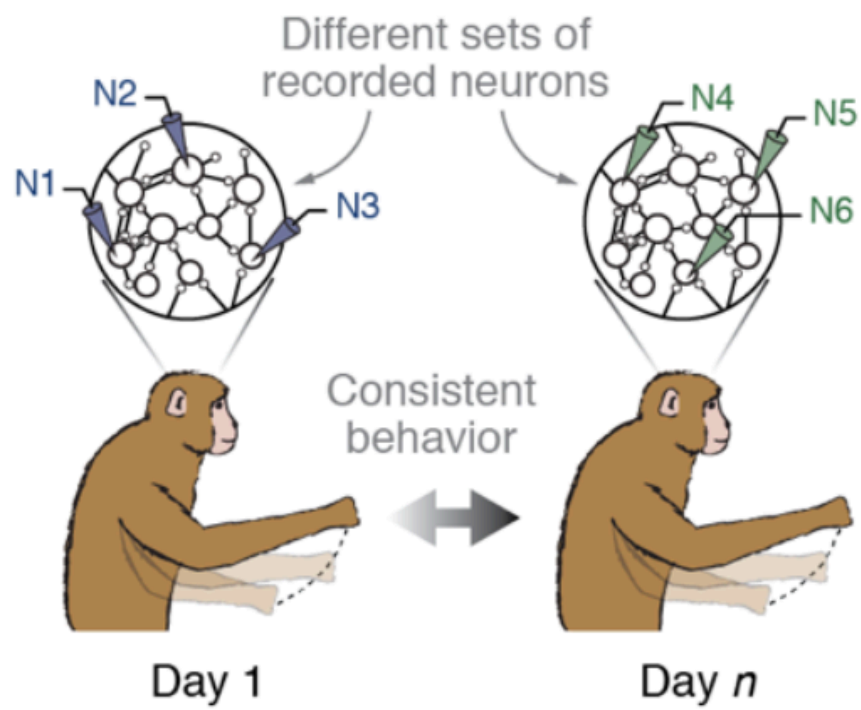

- Subjects consistently perform the same behavior over days, months, and years.
- Hypothesis: the true latent dynamics associated with consistent behavior should be stable.
- But: The same neurons cannot be recorded over this period.

In order to verify this hypothesis, we need to compensate for the fact that the true latent dynamics is being projected onto different empirical manifolds on different days.

Alignment of latent dynamics

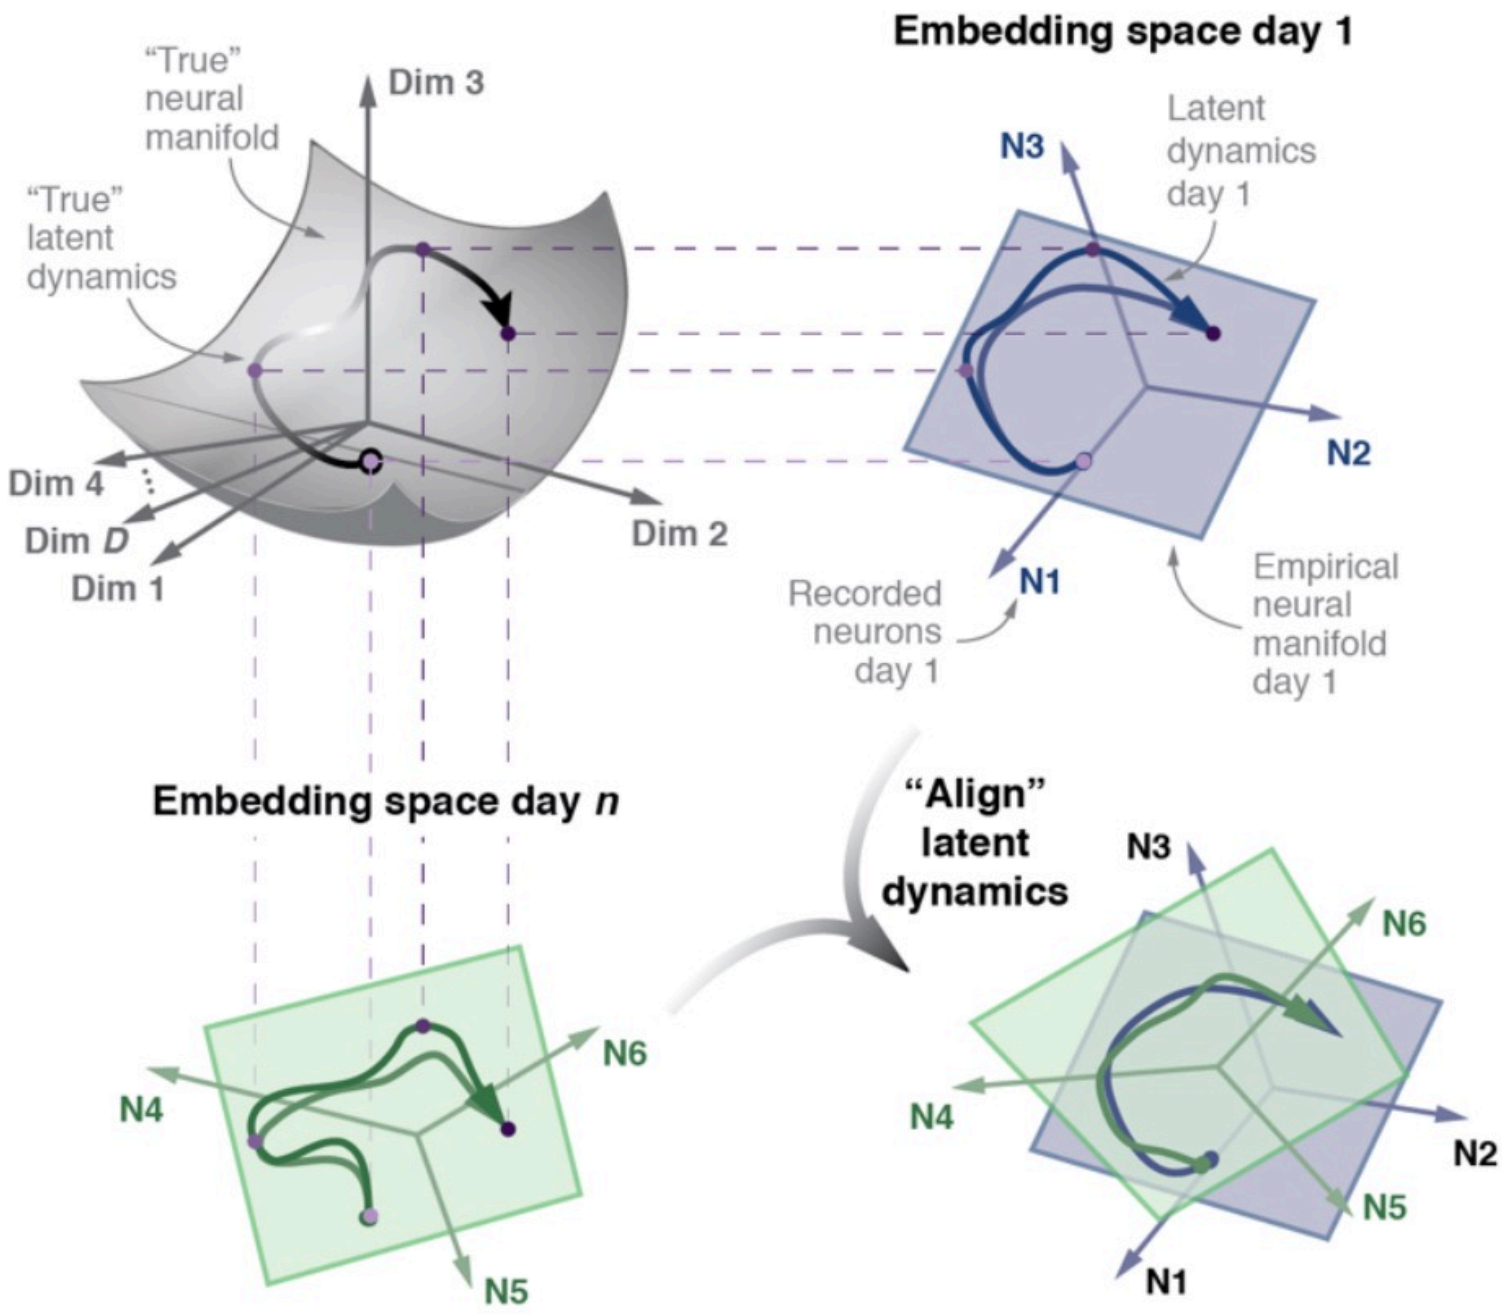

Day-specific neural modes and manifolds

$$X_D = \begin{bmatrix} N_1^{t+1} & N_1^{t+2} & \cdots & N_1^{t+T} \\ N_2^{t+1} & N_2^{t+2} & & N_2^{t+T} \\ \vdots & \vdots & & \vdots \\ N_D^{t+1} & N_D^{t+2} & \cdots & N_D^{t+T} \end{bmatrix}$$

Use Singular Value Decomposition (SVD) on data matrices X:

Data matrix for day n $\qquad X_n = U_n \, \Sigma_n \, V_n^T$

Data matrix for day m $\qquad X_m = U_m \, \Sigma_m \, V_m^T$

Both data matrices X are of dimension D by T, where the ambient dimension D is the cardinality of the union set of neurons recorded on days n and m, and T is the duration of the experiment.

Neurons unrecorded on a given day are assigned zero activity.

Day-specific neural modes and manifolds

Keep the first d columns of the matrices U_n and U_m, to obtain $\widetilde{U}_n$ and $\widetilde{U}_m$.

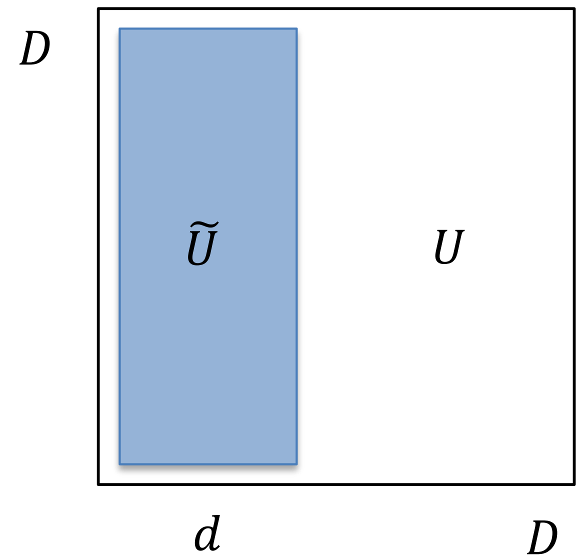

The day-specific low-dimensional manifolds are two hyperplanes:

the d-dimensional hyperplane spanned by the columns of $\widetilde{U}_n$
the d-dimensional hyperplane spanned by the columns of $\widetilde{U}_m$

These column vectors are the day-specific neural modes

d is the flat dimension of the day-specific manifolds

Canonical Correlation Analysis (CCA)

The data matrices X_n and X_m are projected onto the corresponding d-dimensional manifolds spanned by the neural modes using $\widetilde{U}_n$ and $\widetilde{U}_m$ to obtain the latent variables L_n and L_m:

$$L_n = \widetilde{U}_n^T X_n \qquad \text{and} \qquad L_m = \widetilde{U}_m^T X_m$$

These data matrices are of dimension d by T, where:

d: flat manifold dimensionality

T: duration of the experiment

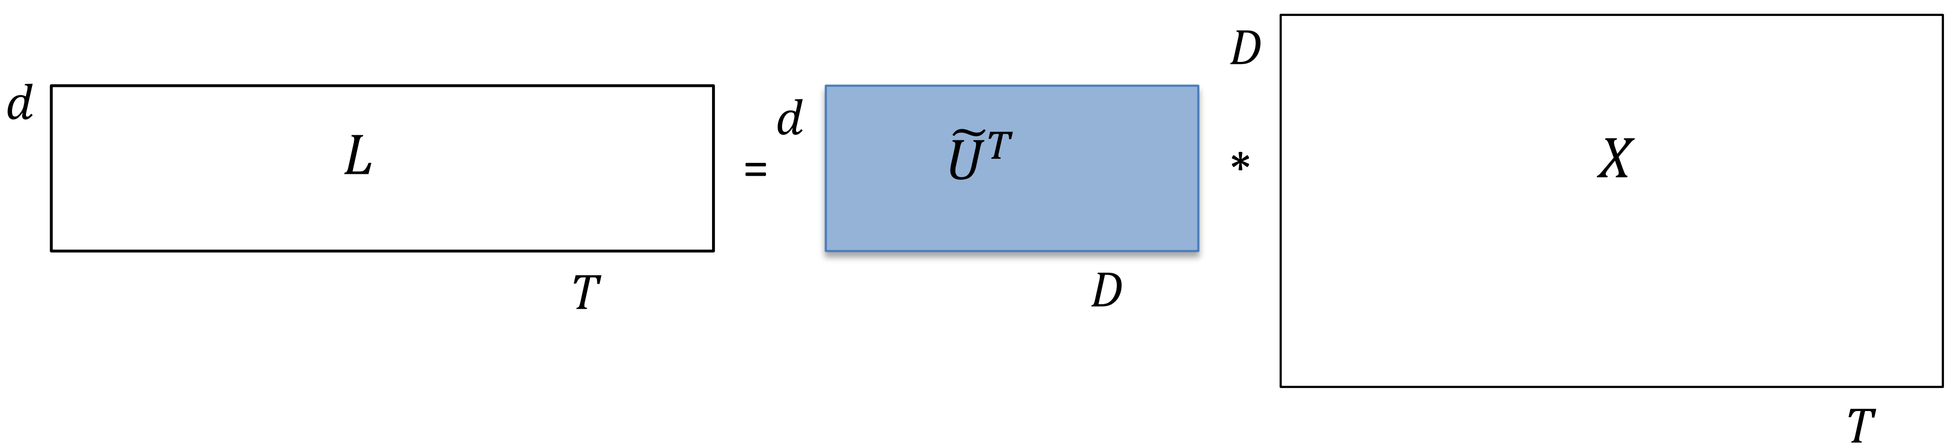

Canonical Correlation Analysis (CCA)

The CCs between the unaligned latent dynamics are the pairwise correlations between the rows of L_n and L_m: $L_n \, L_m^T$

CCA starts with a QR decomposition of the transposed latent variable matrices L_n and L_m,

$$L_n^T = \; Q_n \, R_n \qquad \text{and} \qquad L_m^T = \; Q_m \, R_m$$

The d column vectors of each matrix Q provide an orthonormal basis for the column vectors of the corresponding matrix L^T. The d by d inner product matrix of Q_n and Q_m yields a SVD:

$$Q_n^T \, Q_m = U \, S \, V^T$$

The elements of the diagonal matrix S are the canonical correlations (CCs), sorted from largest to smallest. They quantify the similarity in the aligned latent dynamics.

Canonical Correlation Analysis (CCA)

CCA yields new manifold directions that maximize the pairwise correlations between latent dynamics across the two days. The linear transformations that align the latent variables are effected by d by d matrices M_n and M_m:

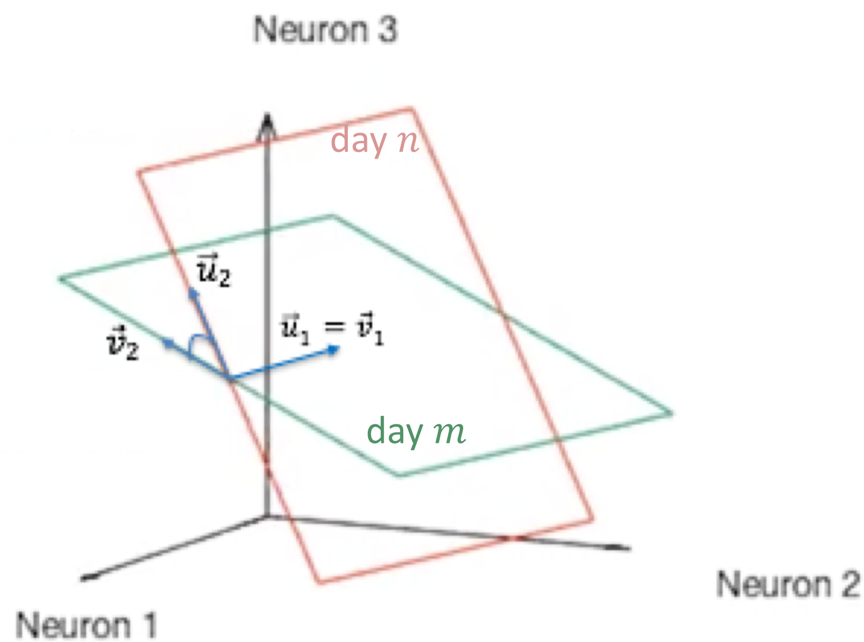

$$M_n = R_n^{-1} U$$

$$M_m = R_m^{-1} V$$

$$L_n \Longrightarrow M_n^T \; L_n$$

$$L_m \Longrightarrow M_m^T \; L_m$$

$$L_n \, L_m^T \Rightarrow (M_n^T \; L_n) \, (M_n^T \; L_n)^T = S$$

Stability of M1 latent dynamics

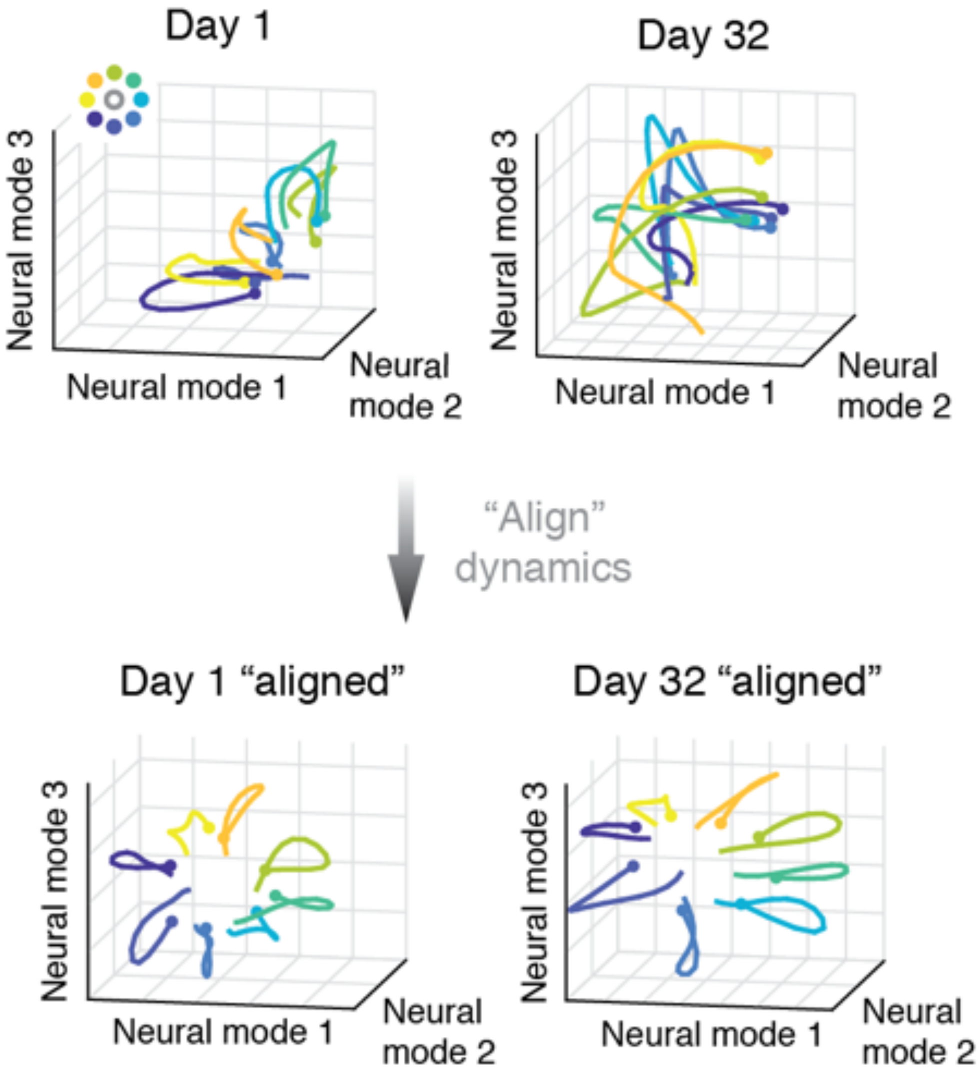

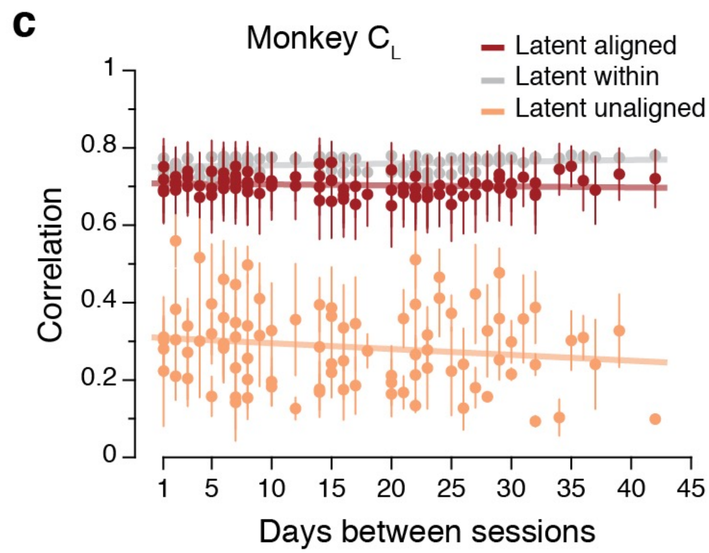

Alignment of latent dynamics

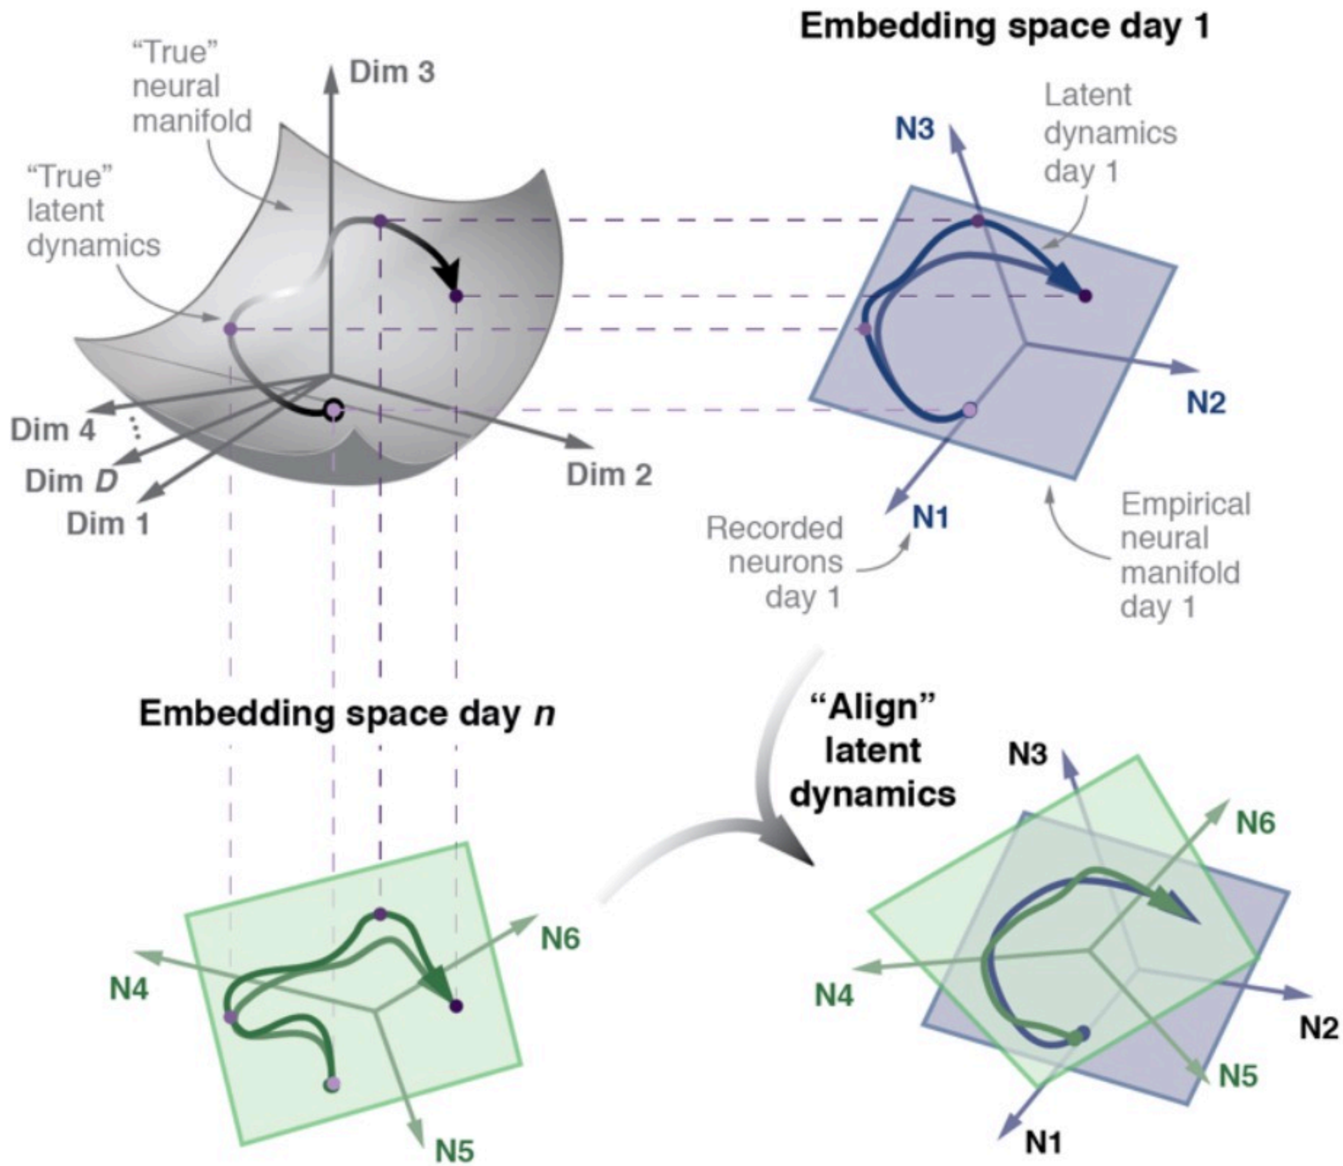

Gallego, Perich, Chowdhury, Solla, Miller, *Nature Neurosci* (2020)

Neural manifolds for the control of movement

Juan A. Gallego
Matthew G. Perich
Raeed H. Chowdhury
Lee E. Miller

Northwestern University

The future: natural behavior

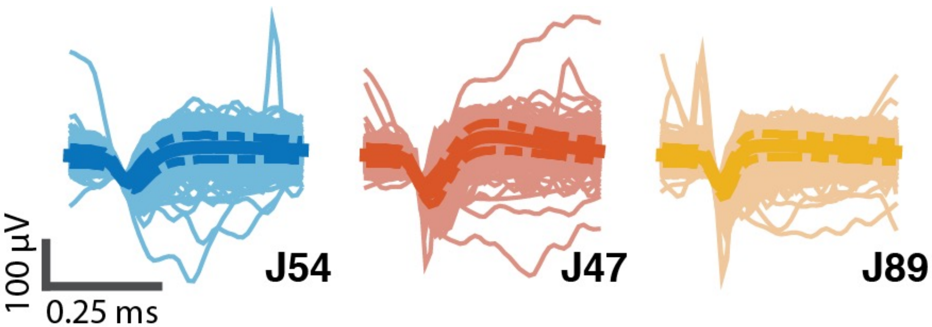

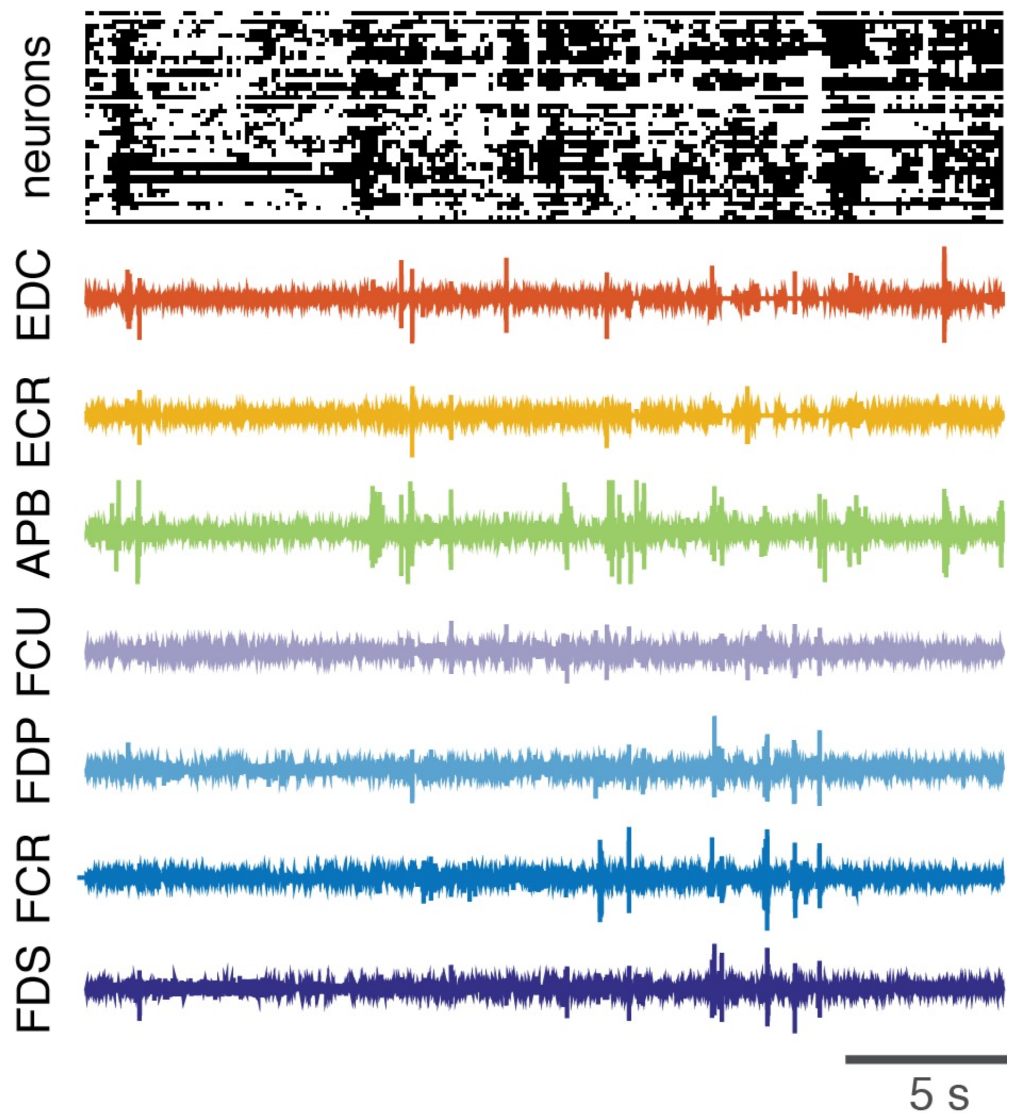

Stable prediction of movement kinematics

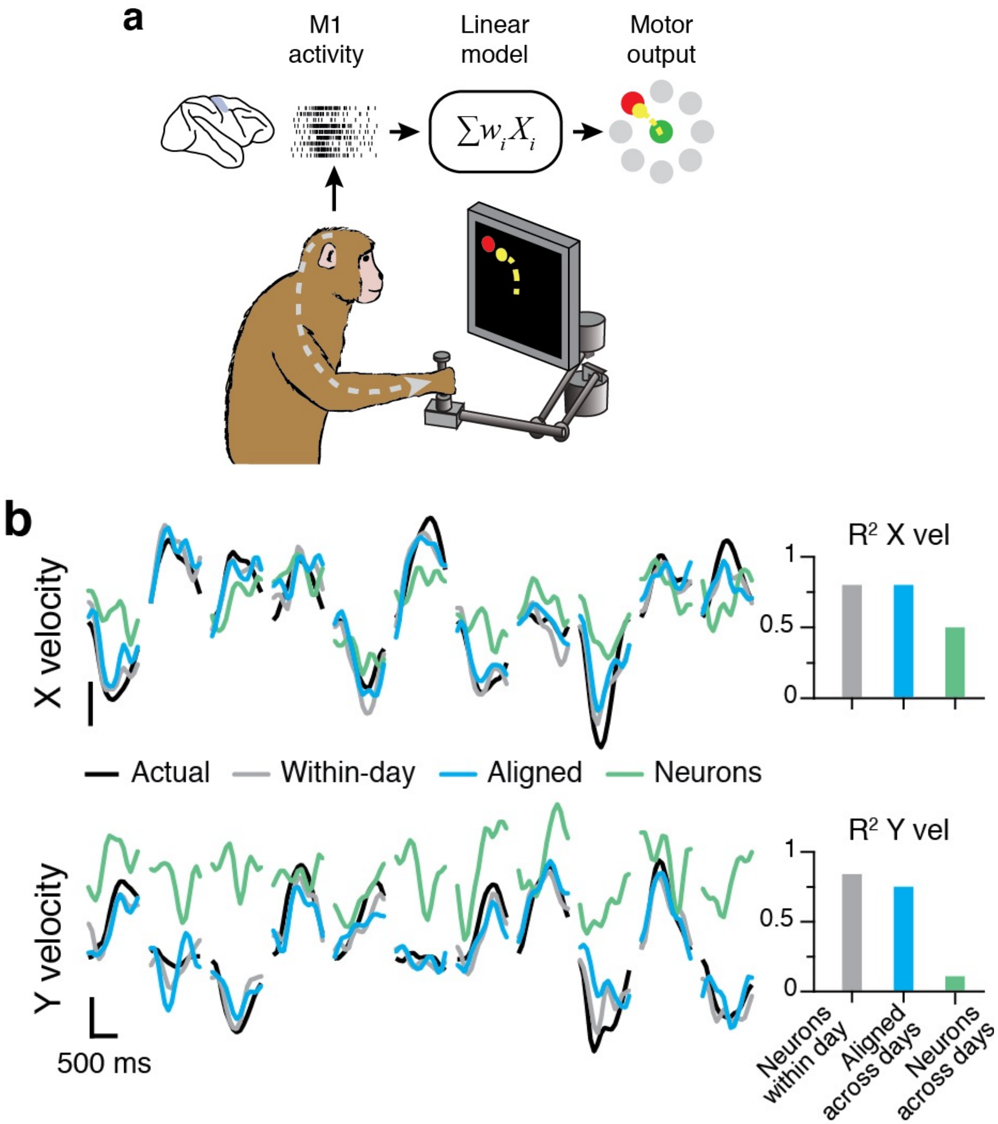

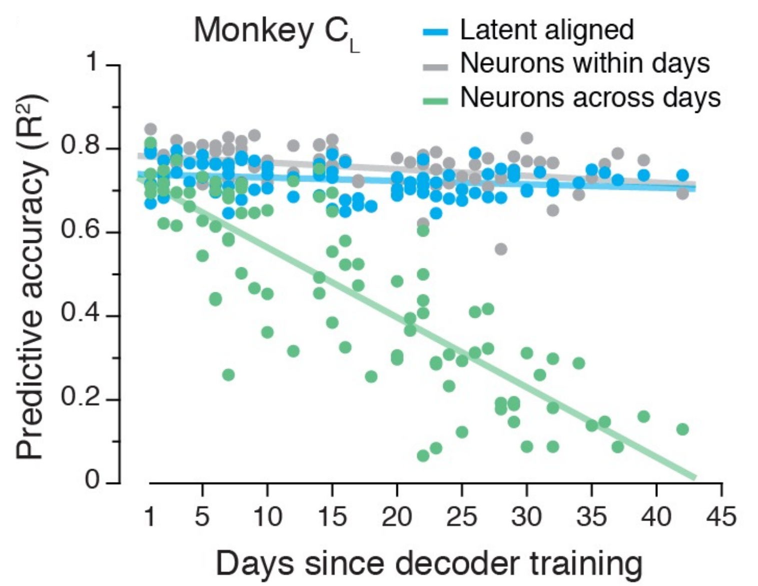

Prediction of muscle activity