Low dimensional manifolds for neural dynamics

Sara A. Solla
Northwestern University

Mathematics, Physics & Machine Learning
Técnico Lisboa
March 2, 2023

Manifolds for neural population activity

- A simple motor task

- Neural manifolds for the control of movement

- The unreasonable effectiveness of linear methods

Manifolds for neural population activity

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

A simple motor task: center-out reaches

Instructed delay center-out reaching task

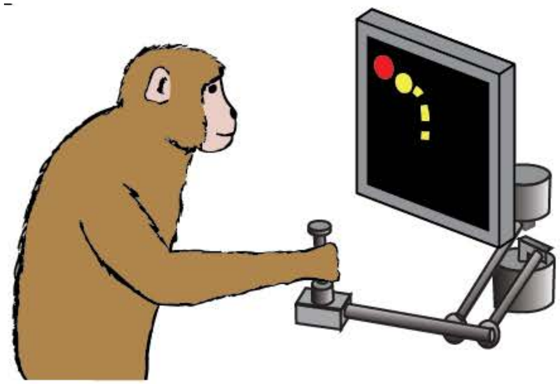

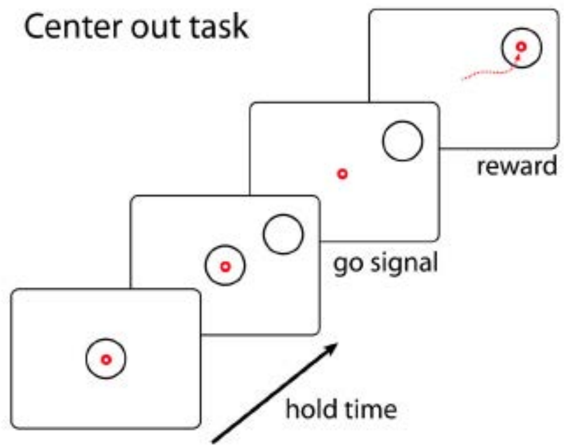

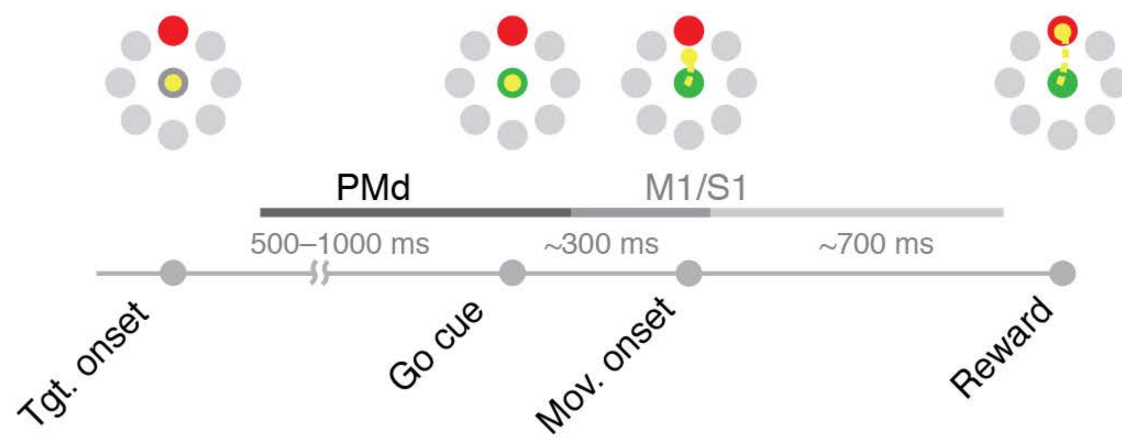

Neural activity: variability and specificity

Single neuron recordings

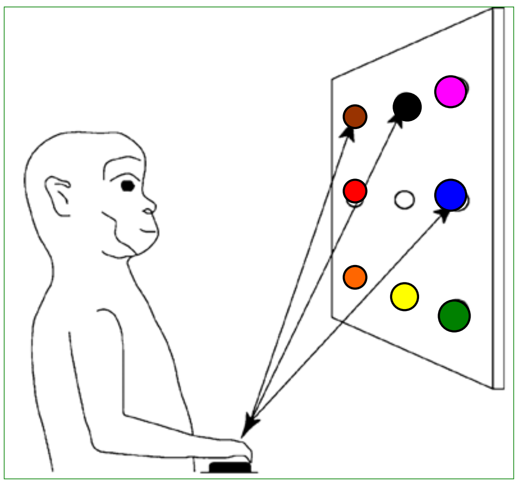

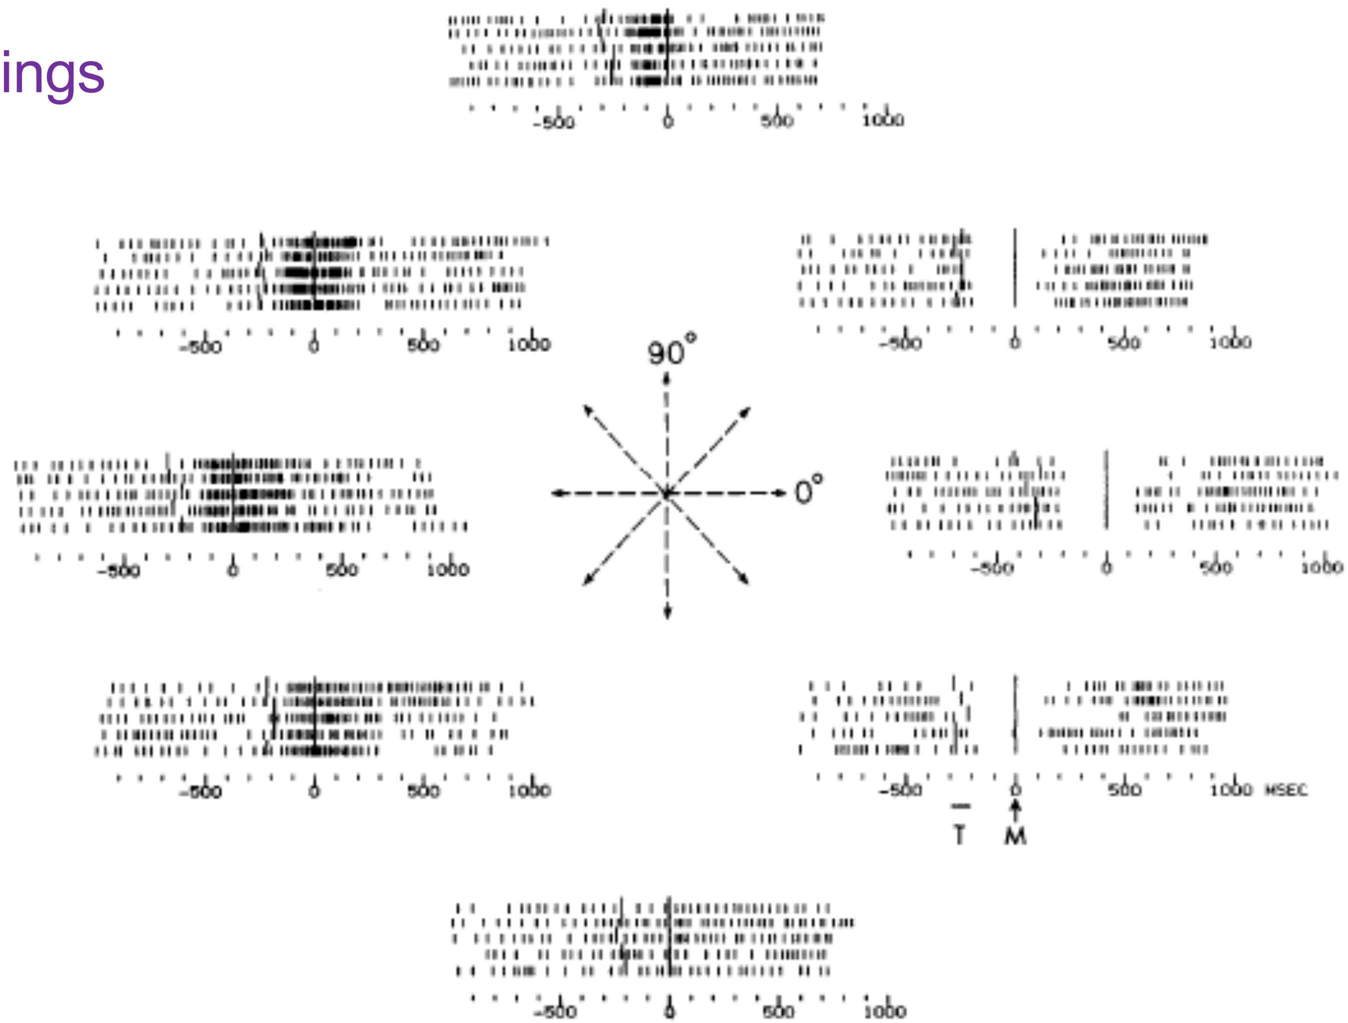

Georgopoulos, Kalaska, Caminity, Massey, *J of Neurosci* (1982)

Tuning curves in motor cortex

tuning curve $r = f(s)$

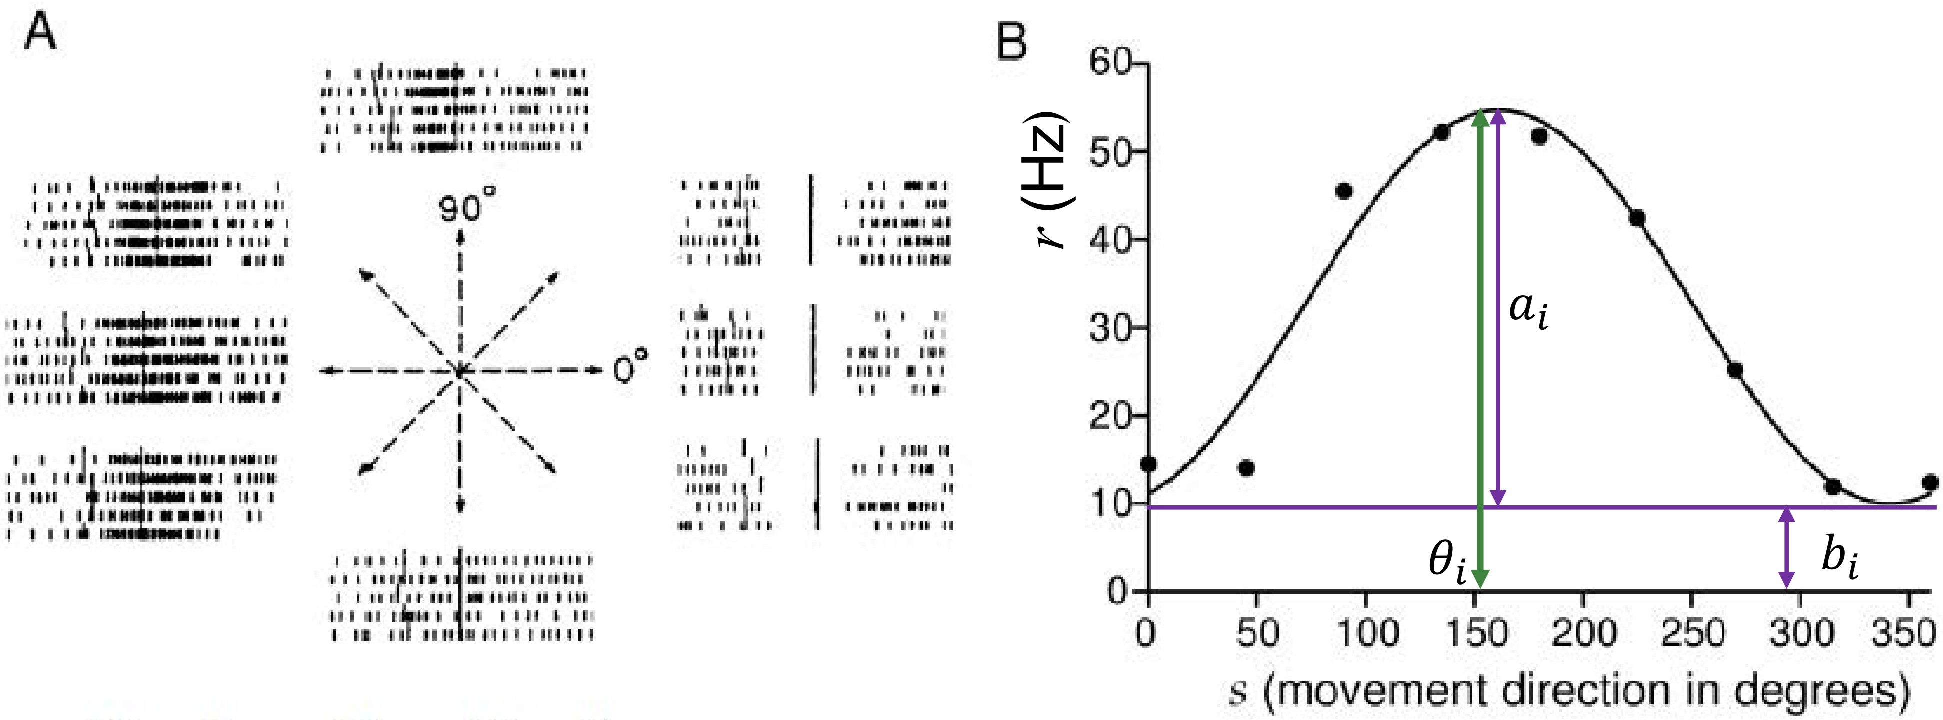

Hand reaching direction

reach (r, θ) implies $f_i = b_i + r a_i (1/2)[1 + cos(\theta - \theta_i)]$

Cosine tuning curve of a motor cortical neuron

Georgopoulos, Kalaska, Caminiti, Massey, *J Neurosci* (1982)

Multi Electrode Arrays (MEAs)

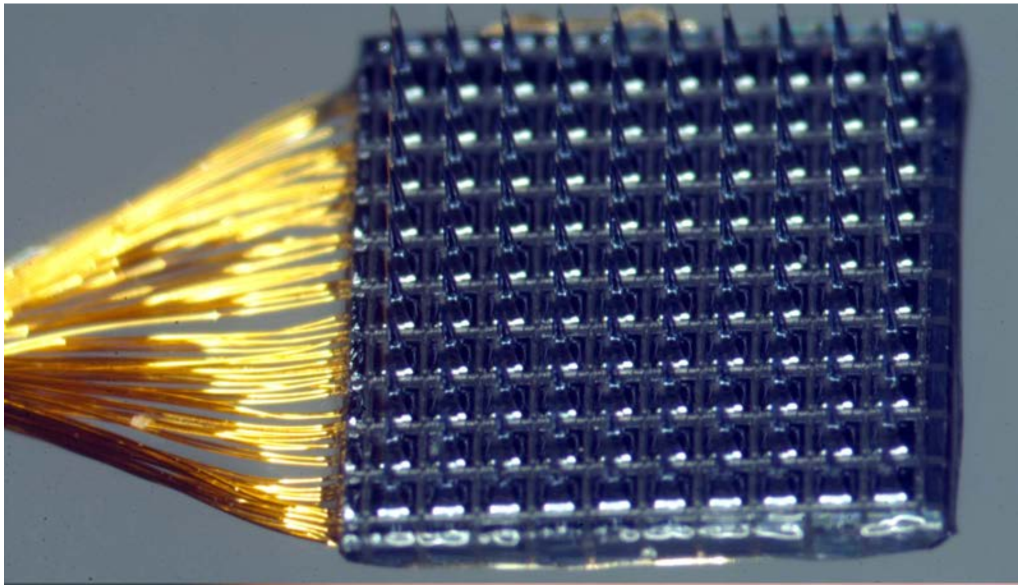

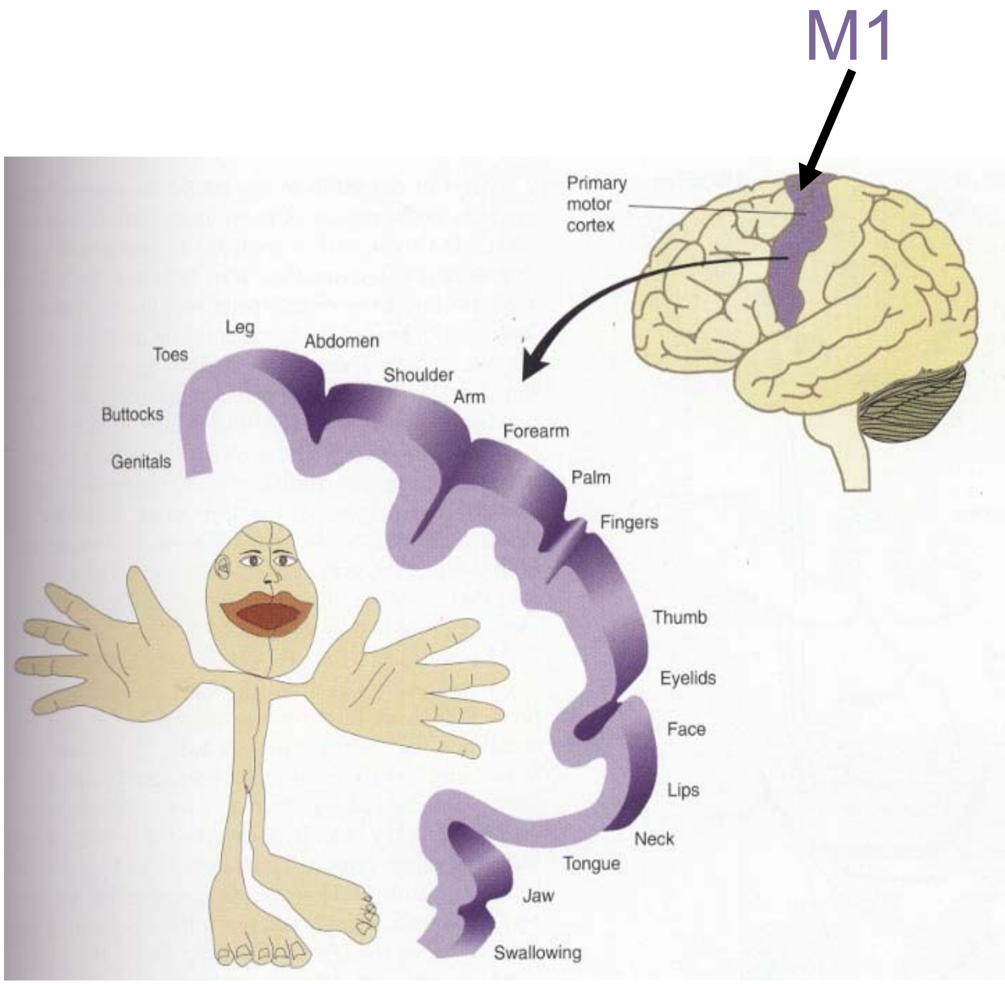

Neural recordings: center-out task

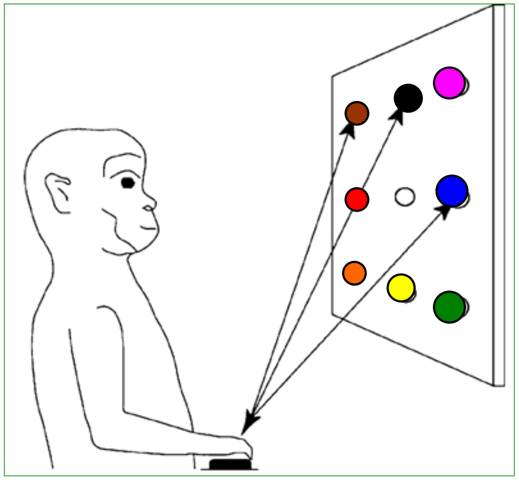

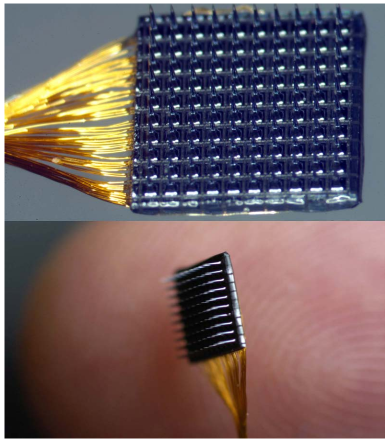

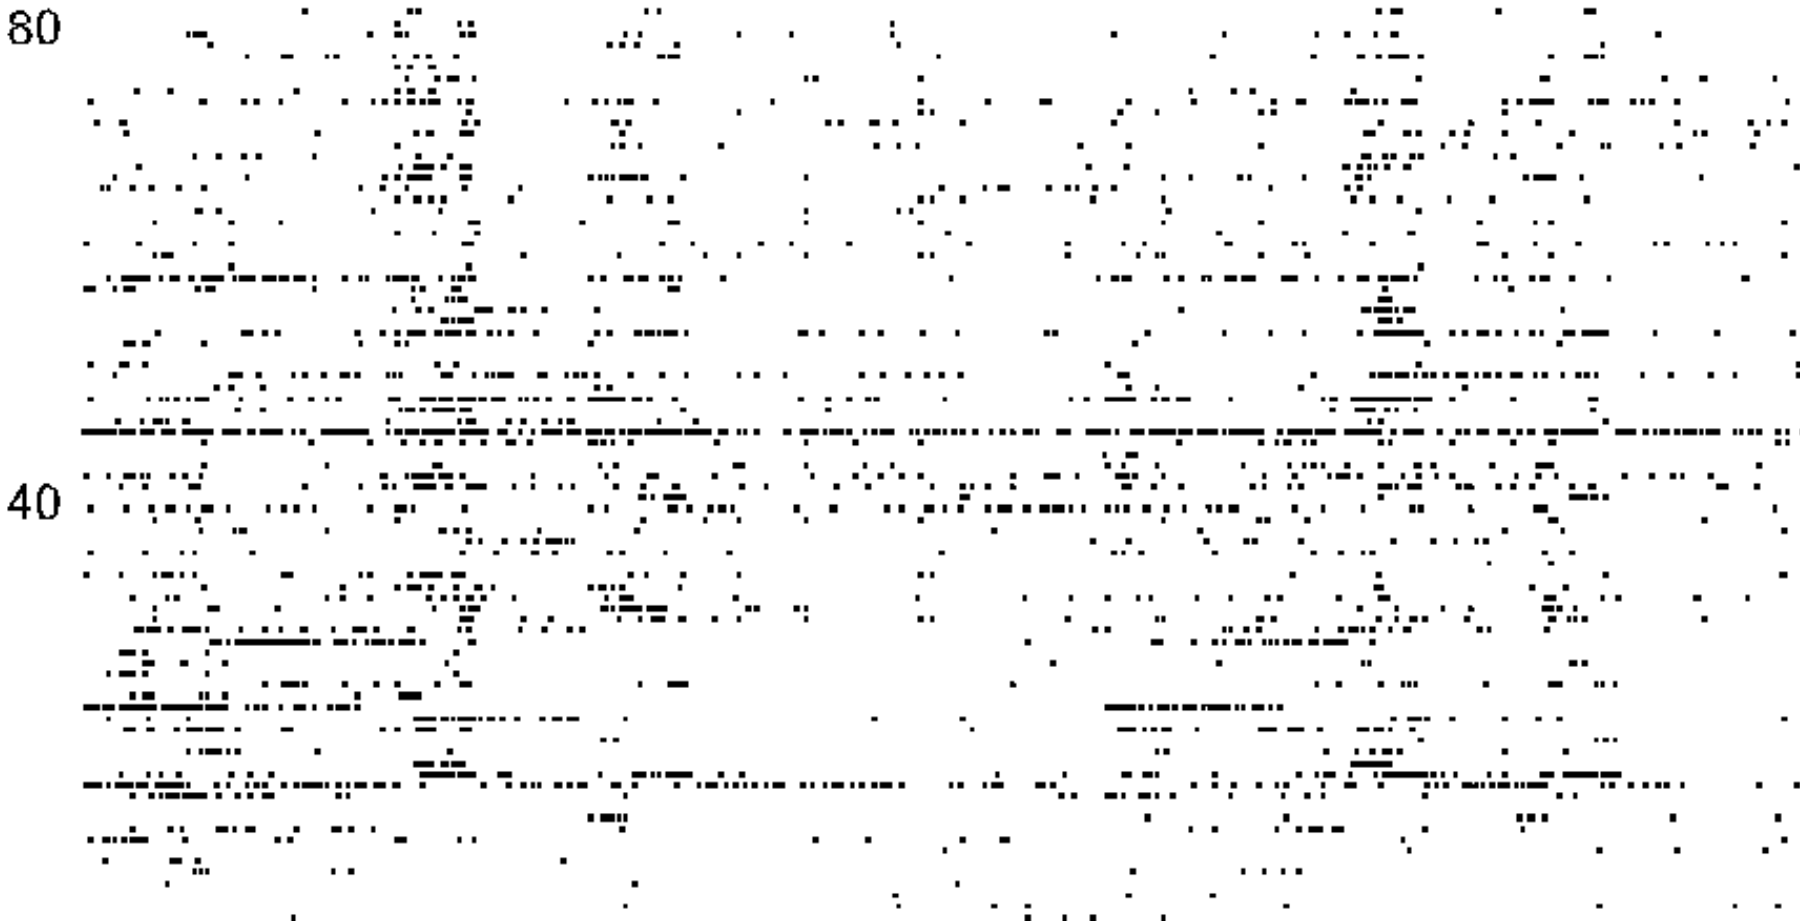

Neural recordings: center-out task

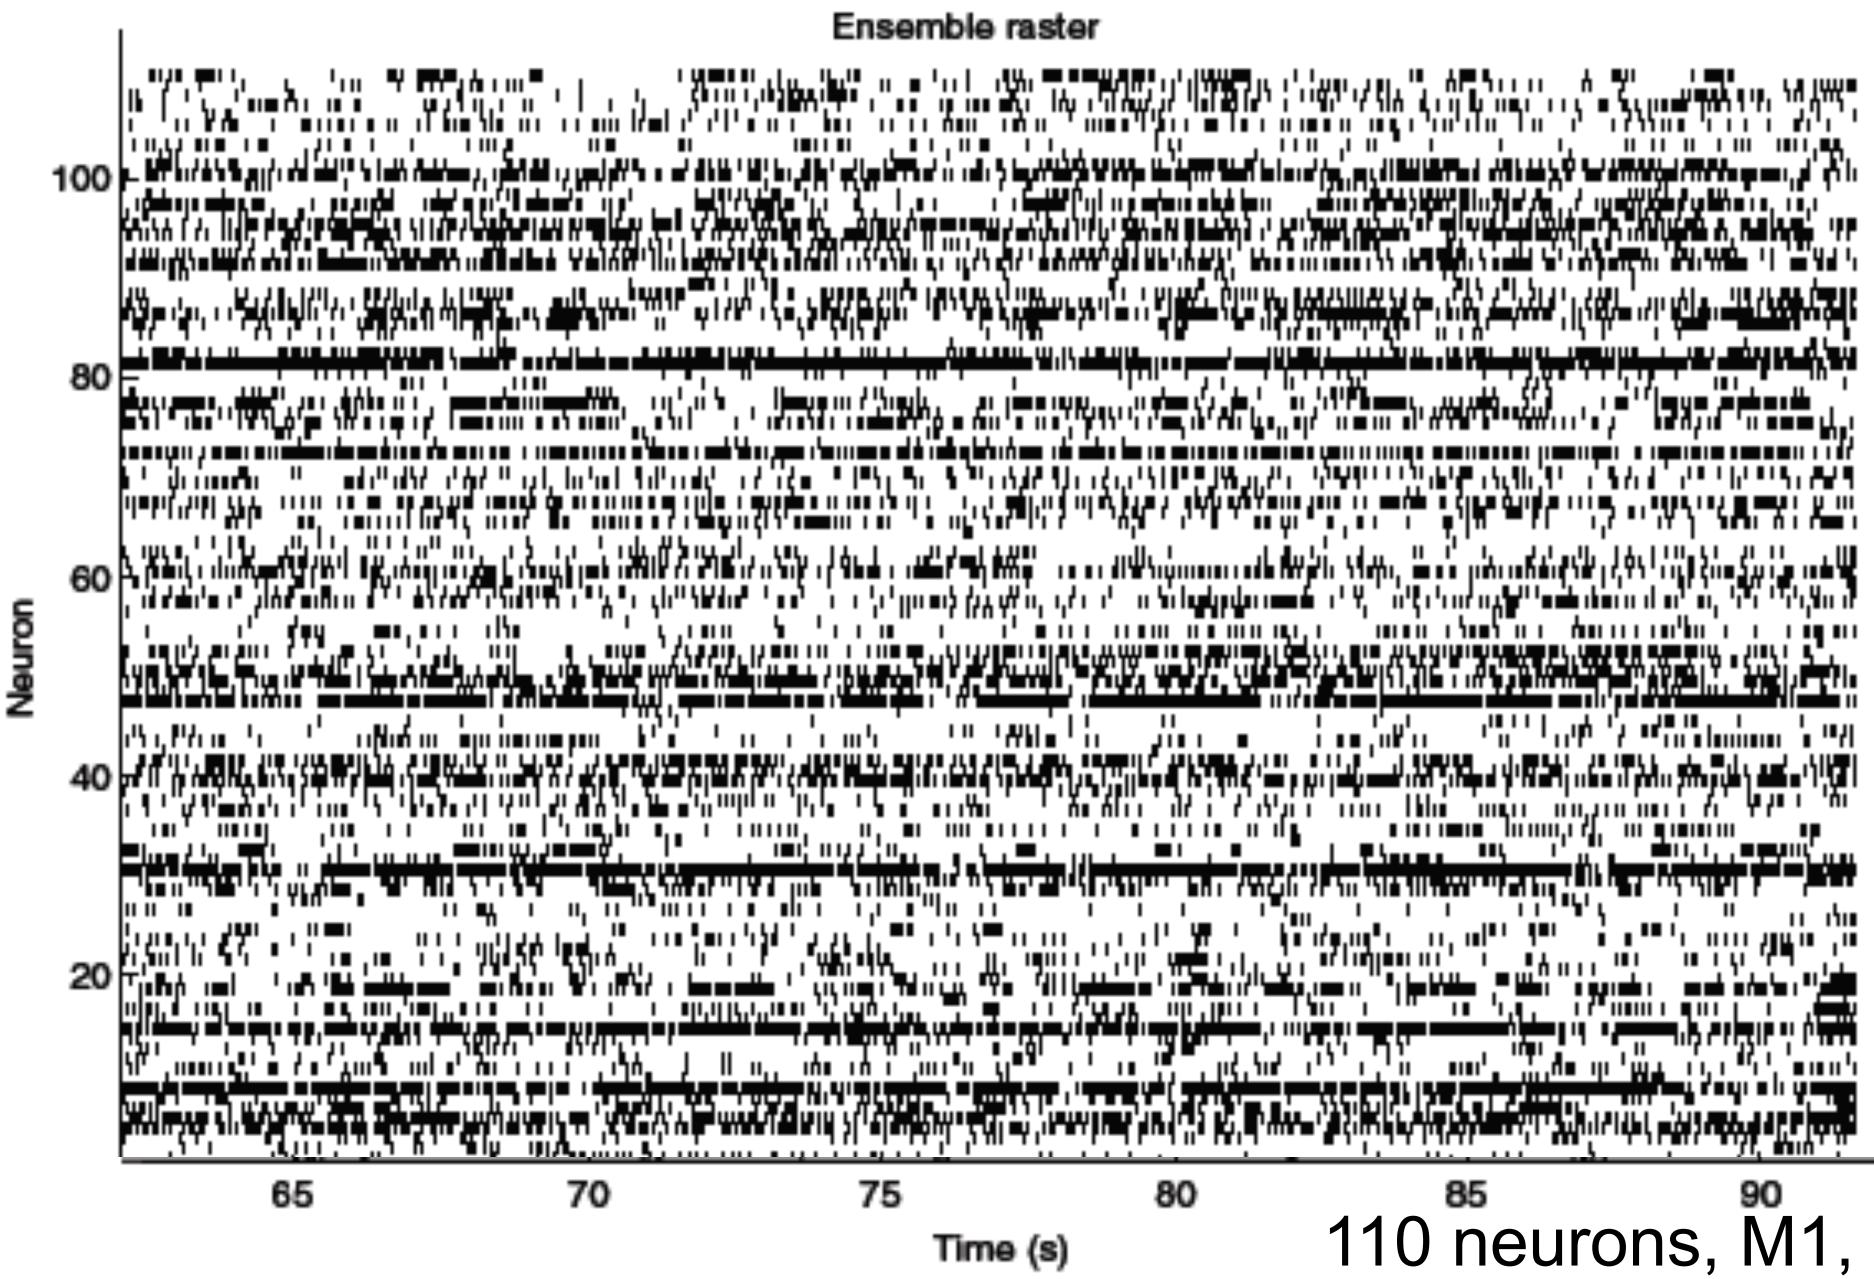

110 neurons, M1, hS3

Population activity : multiple targets

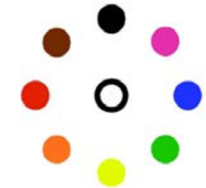

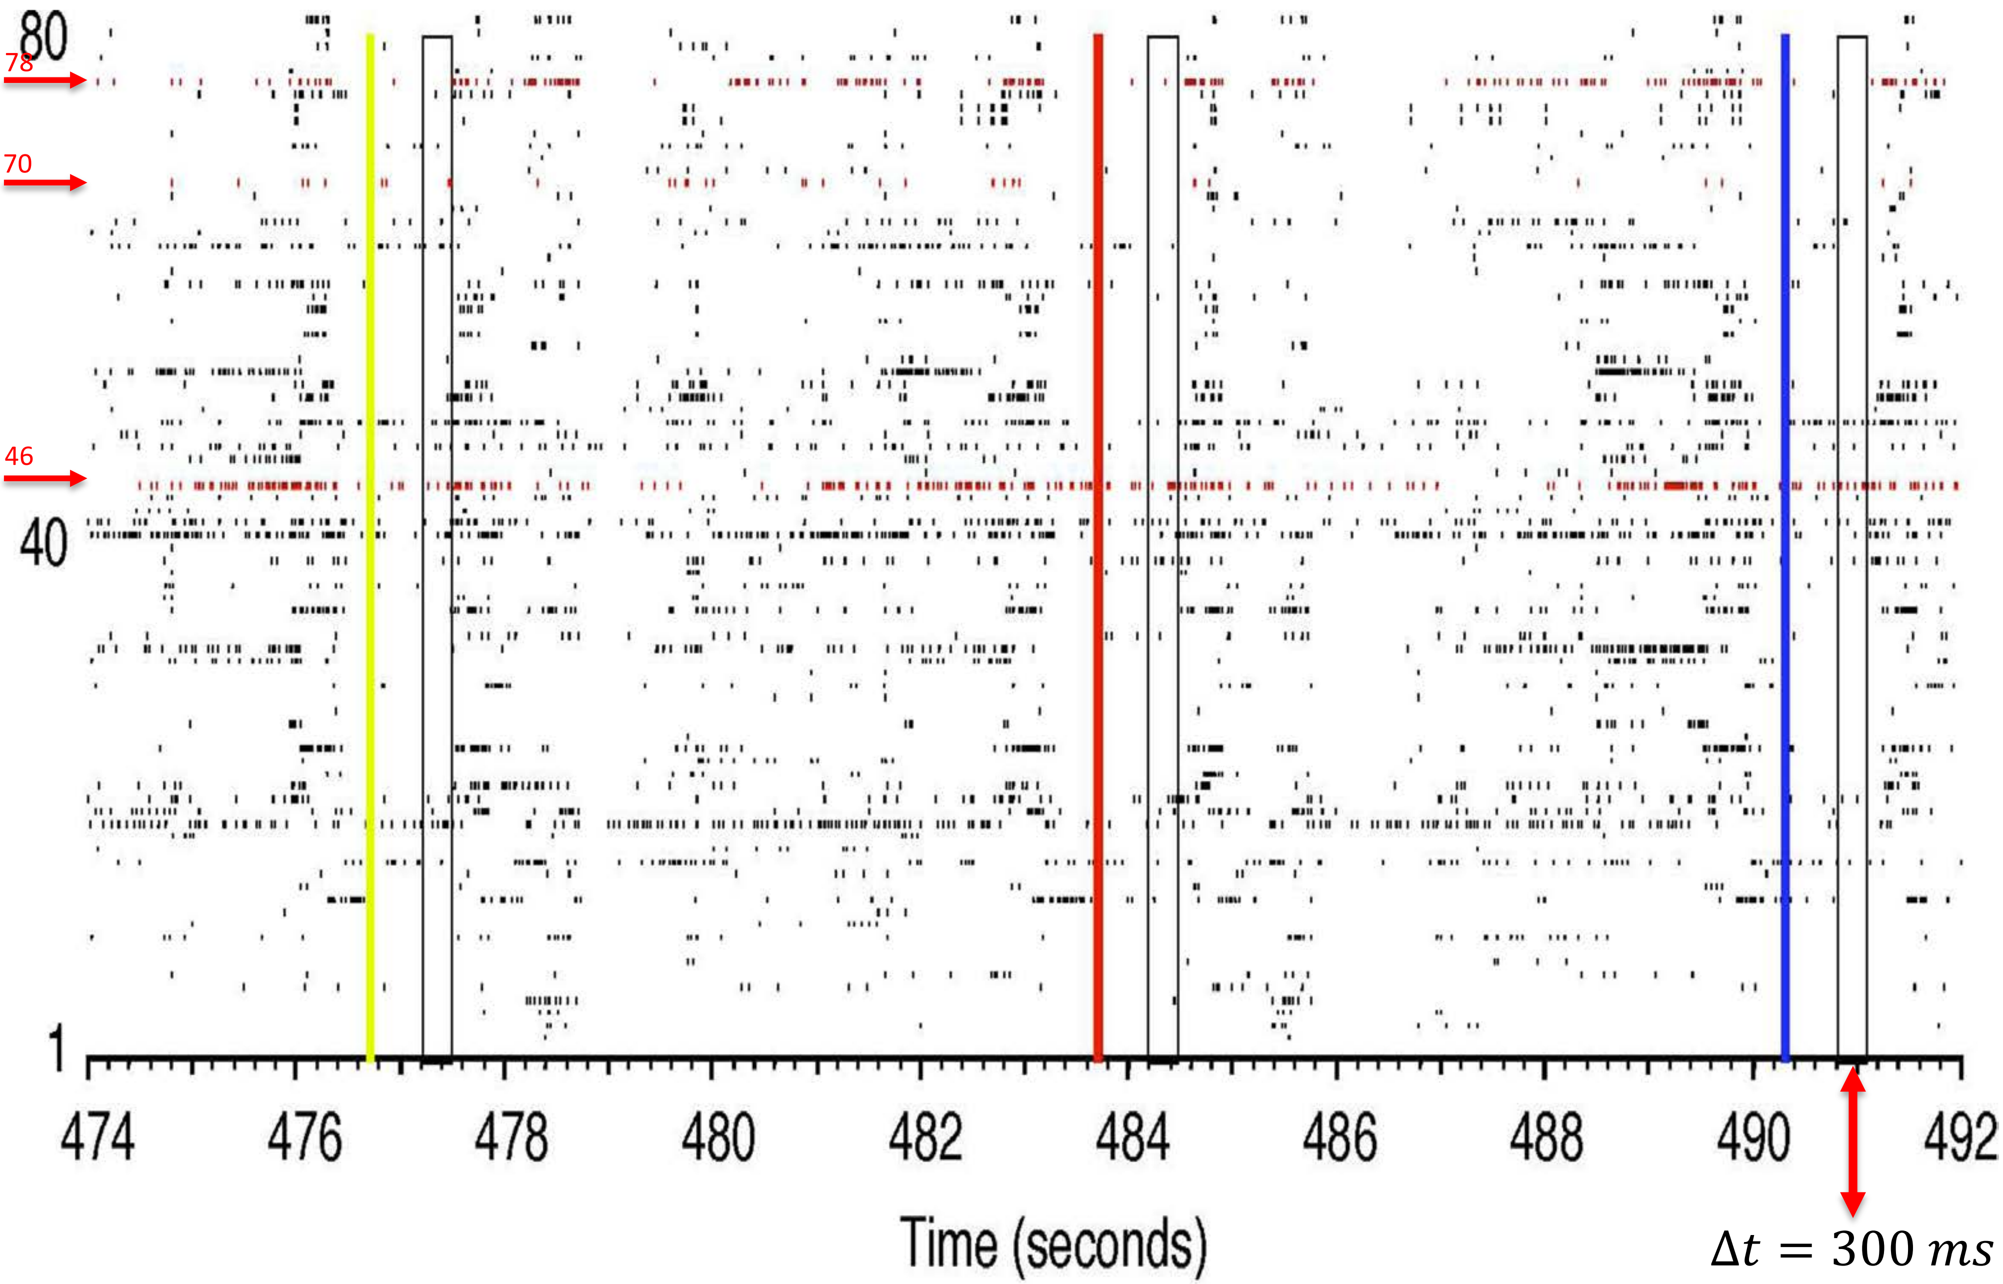

Target-dependent population activity

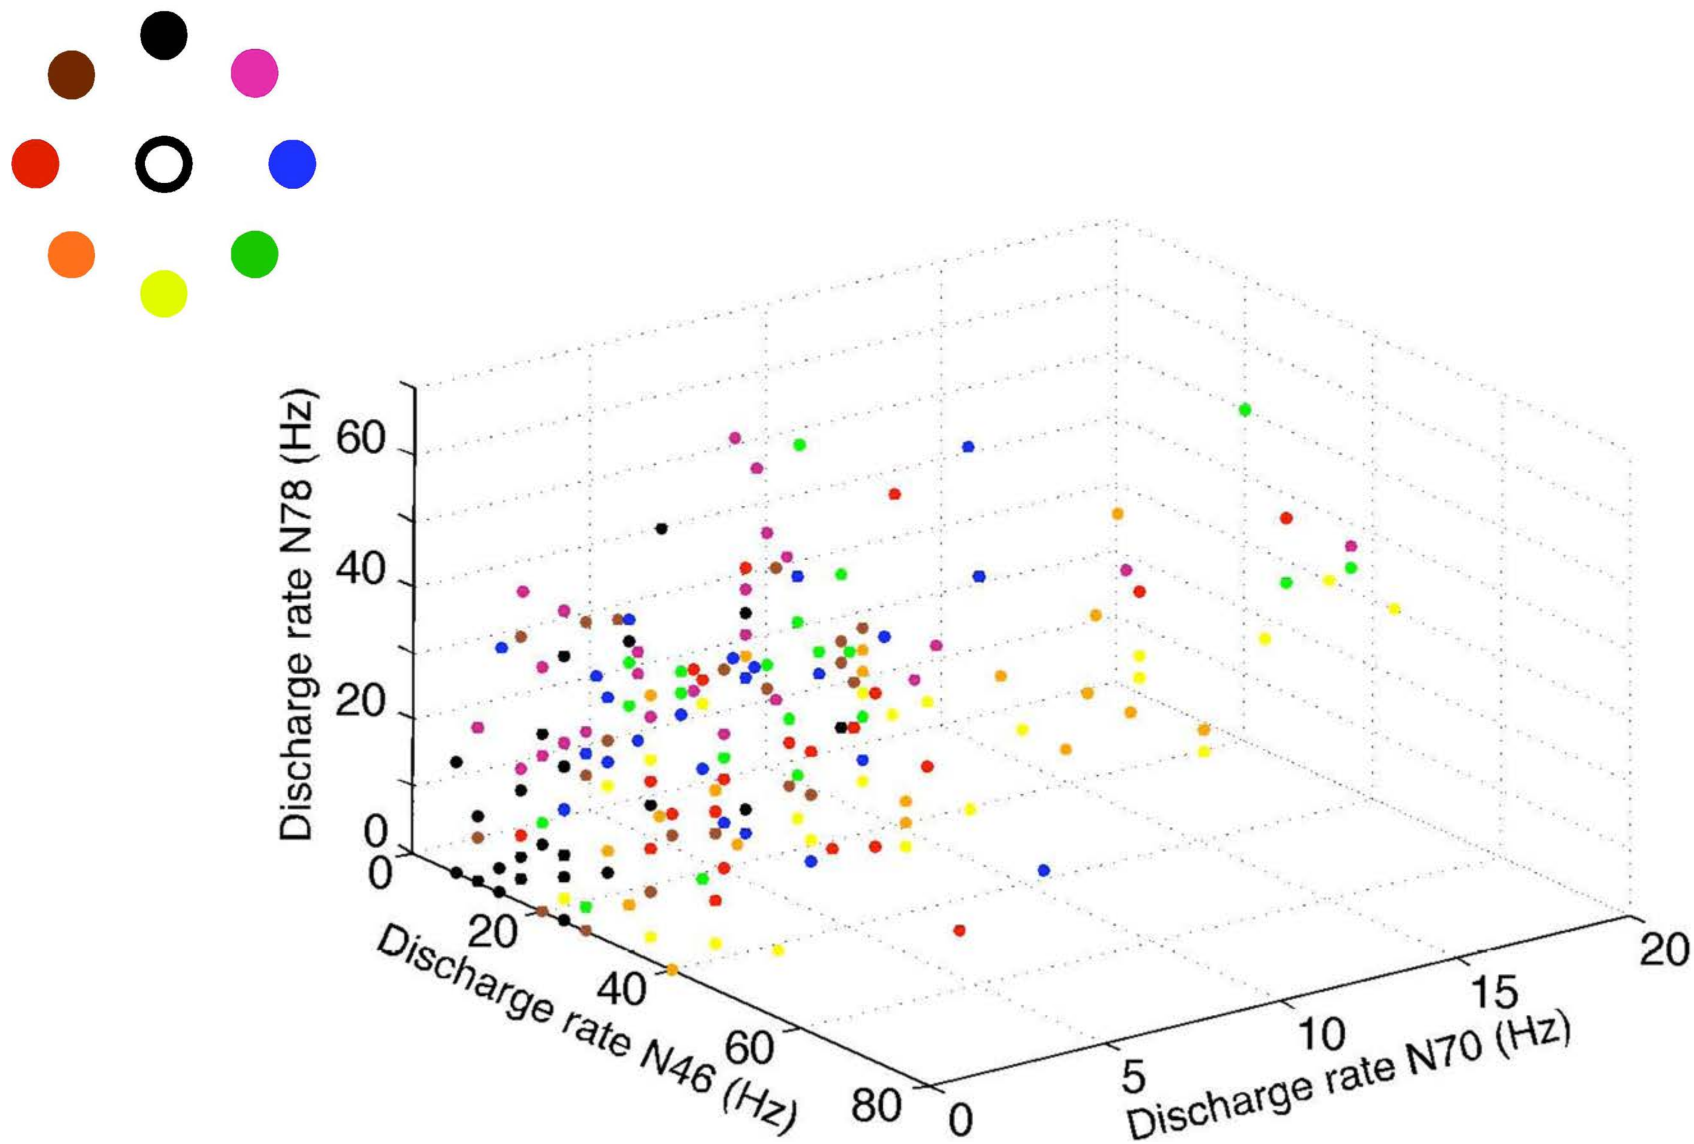

Cosine tuning for direction of motion

reach (r, θ) implies $f_i = b_i + r a_i (1/2)[1 + \cos(\theta - \theta_i)]$

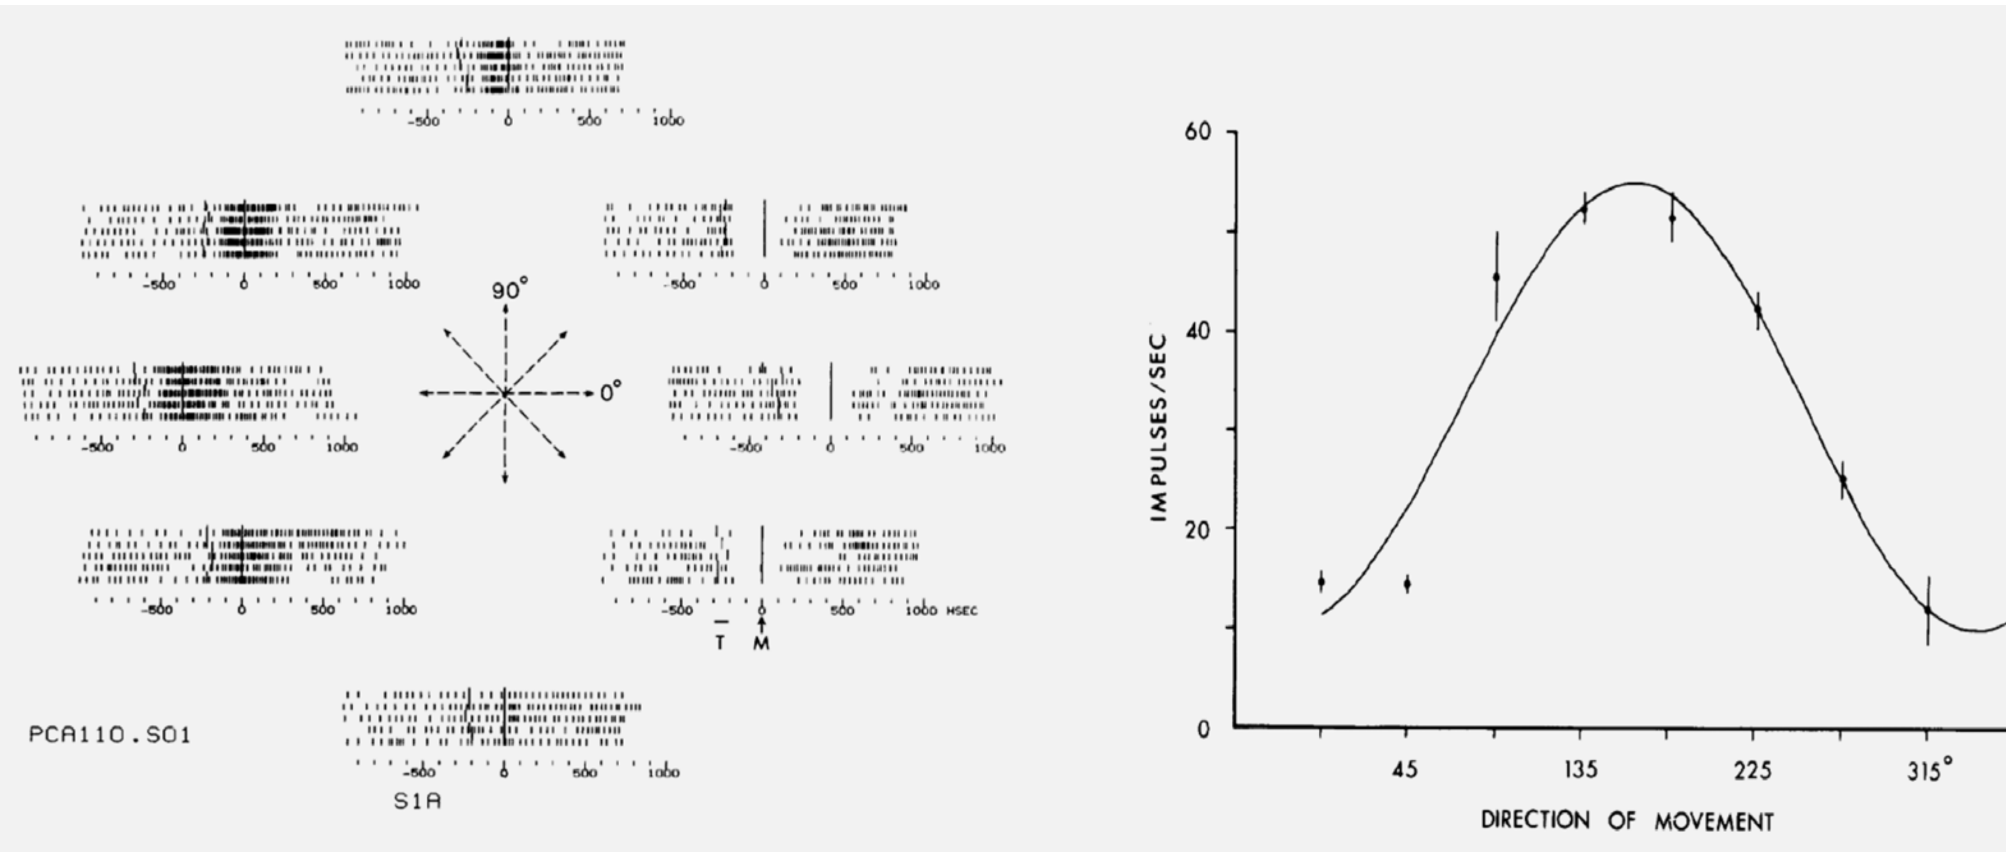

Georgopoulos, Kalaska, Caminiti, Massey, *J Neurosci* (1982)

Population of cosine tuned neurons

$$\vec{f} = \vec{b} + (1/2)r[\vec{a} + \cos(\theta)\vec{\phi_x} + \sin(\theta)\vec{\phi_y}]$$

$$\vec{b} = (b_1, b_2, \ldots, b_N) \quad \vec{\phi_x} = (a_1\cos(\theta_1), \ldots, a_N\cos(\theta_N))$$

$$\vec{a} = (a_1, a_2, \ldots, a_N) \quad \vec{\phi_y} = (a_1\sin(\theta_1), \ldots, a_N\sin(\theta_N))$$

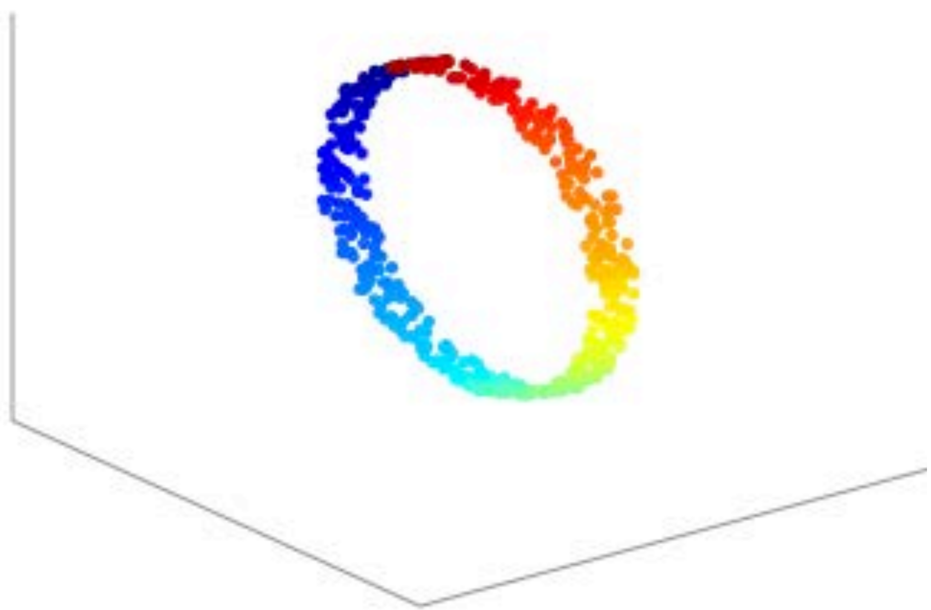

N=100

$r \in [10, 11]$

Target-dependent population activity

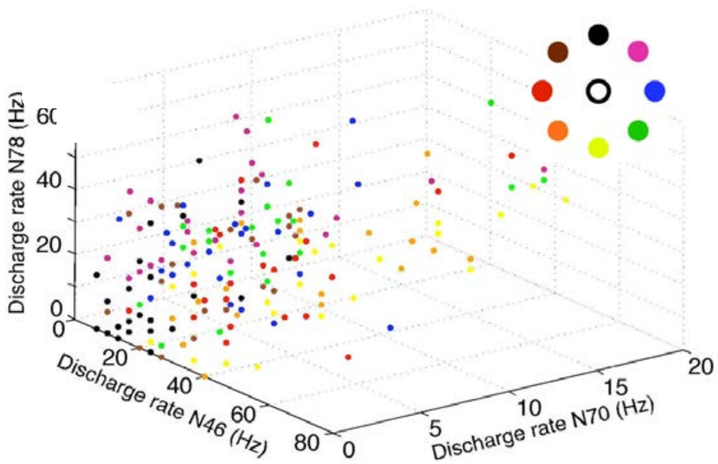

Neural modes: directions in neural space

Specific patterns of populations activity

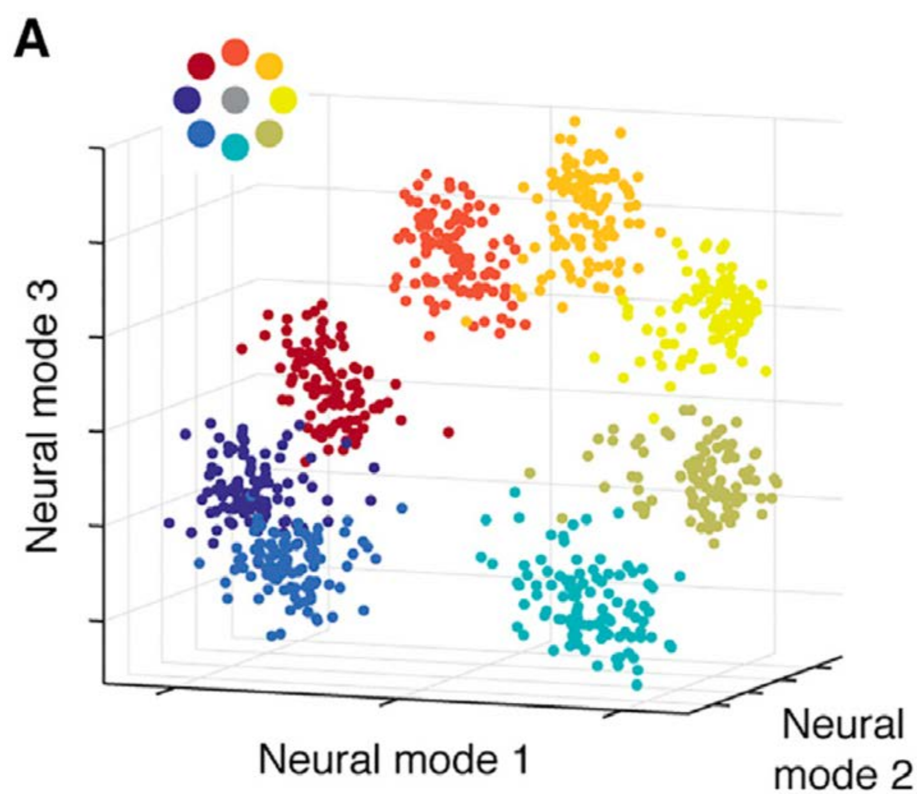

Santhanam, Yu, Gilja, Ryu, Afshar, Sahani, Shenoy, *J Neurophysiol* (2009)

Manifolds for neural population activity

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

Population activity

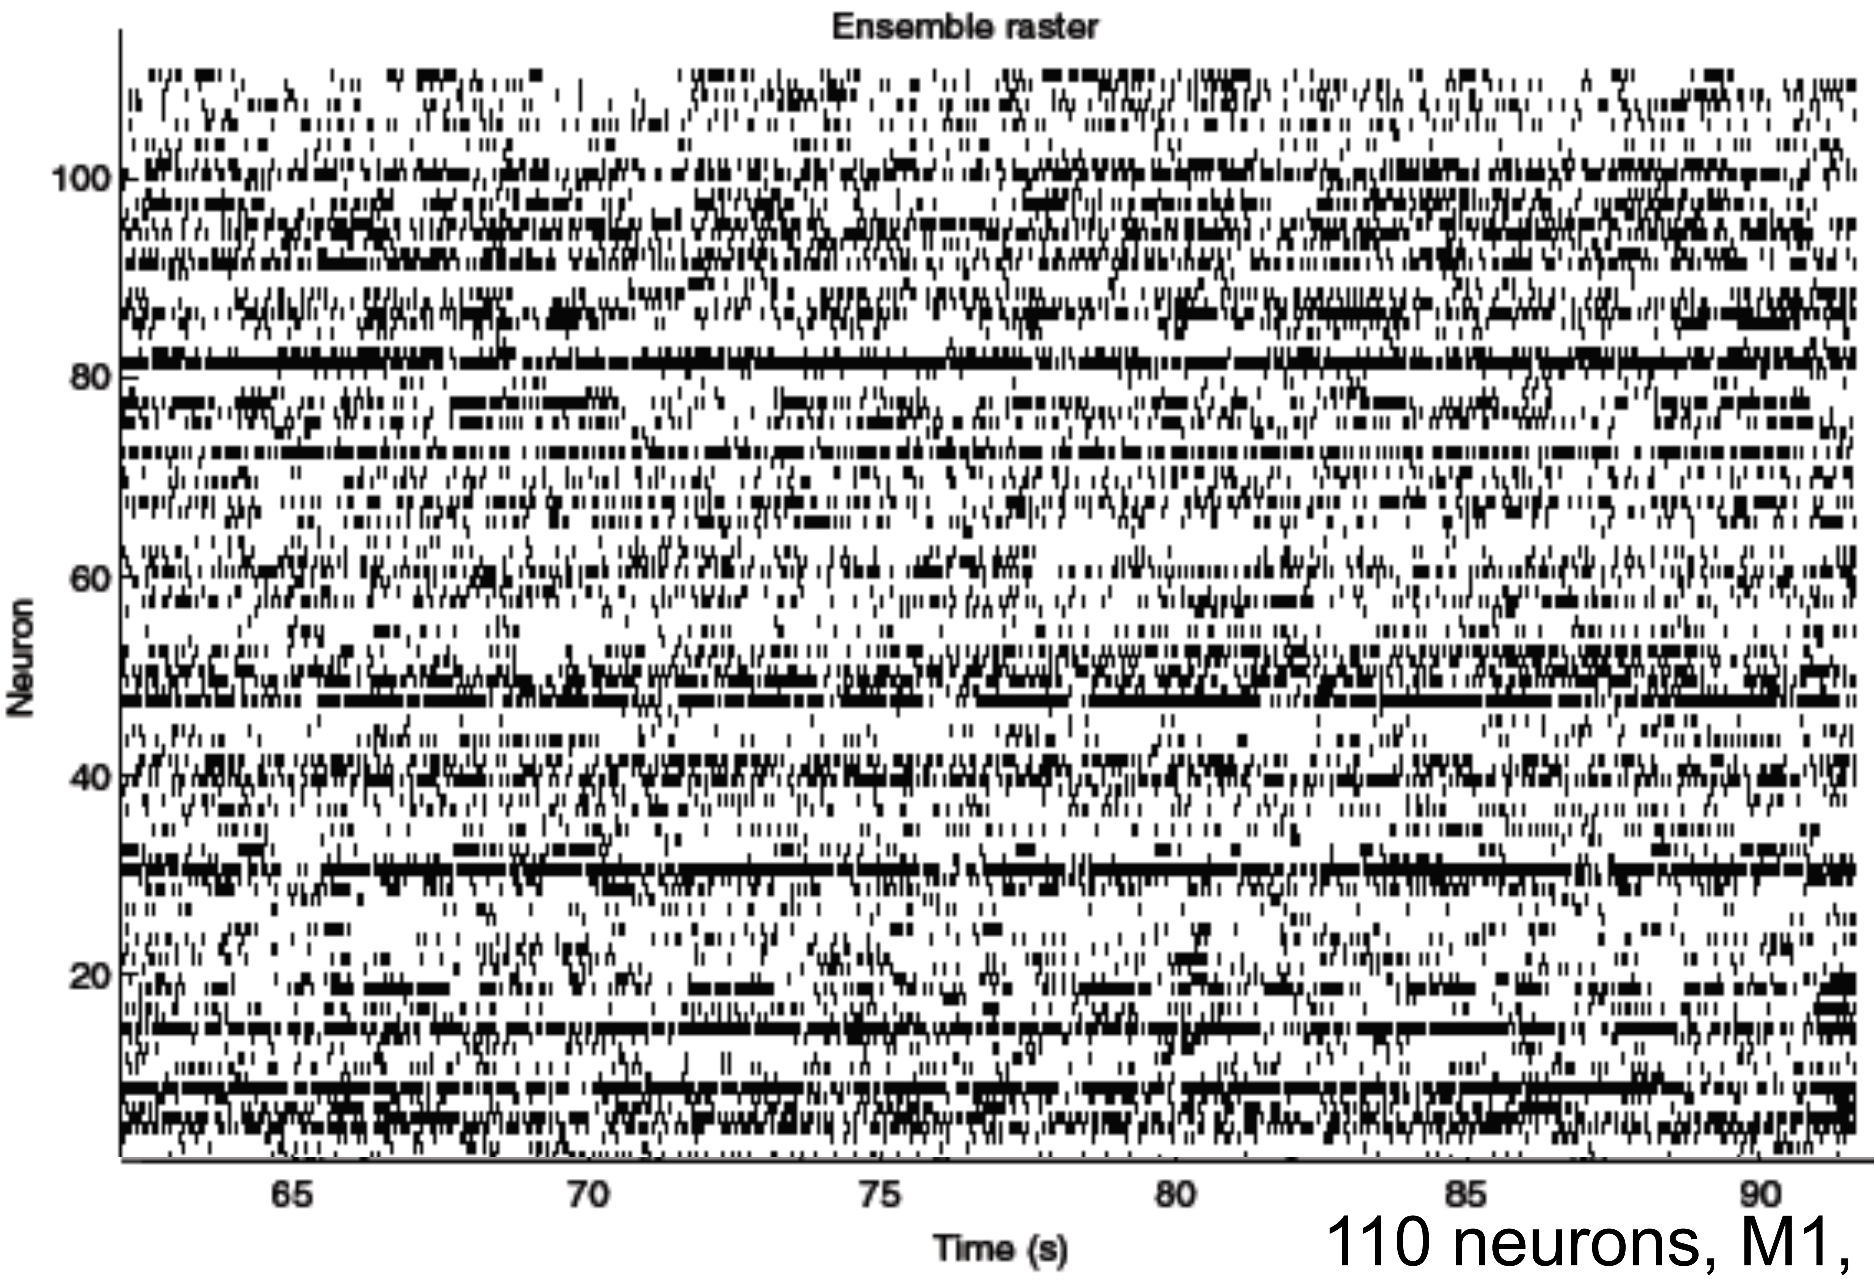

110 neurons, M1, hS3

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 $1 \leq i \leq N$.

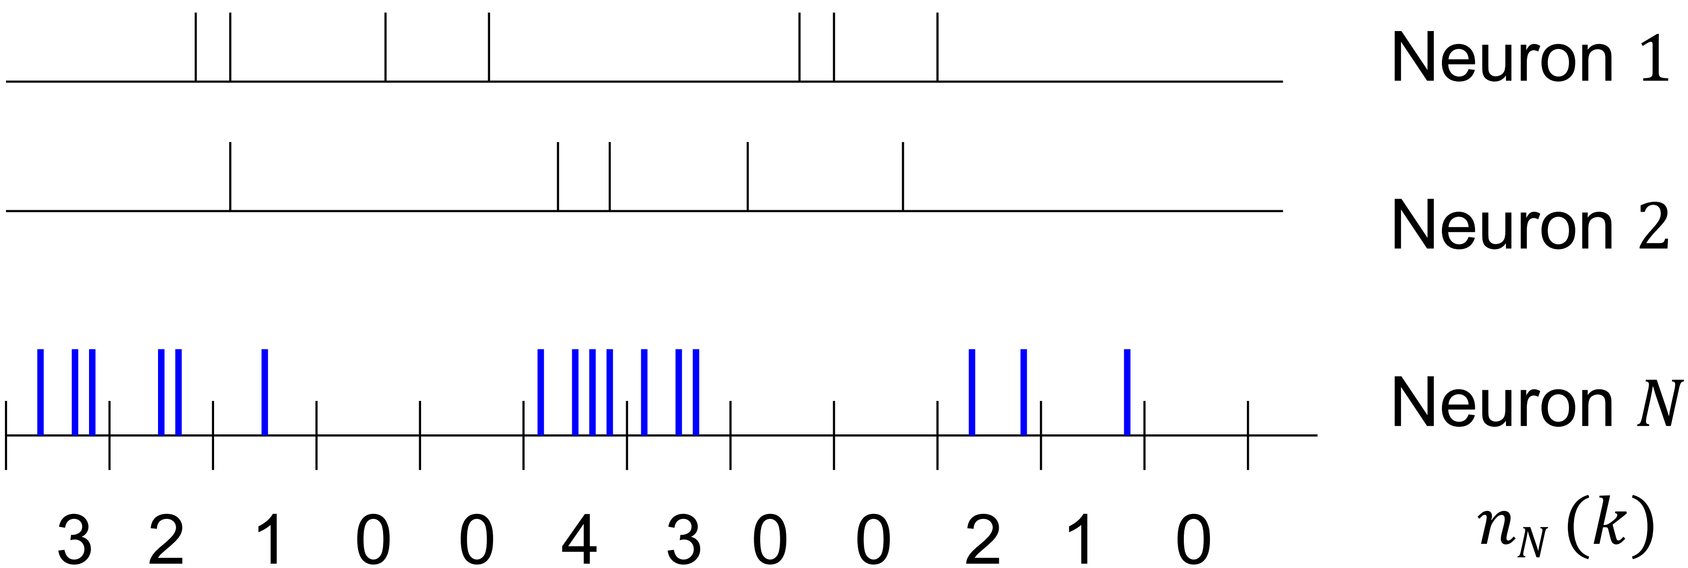

Population activity

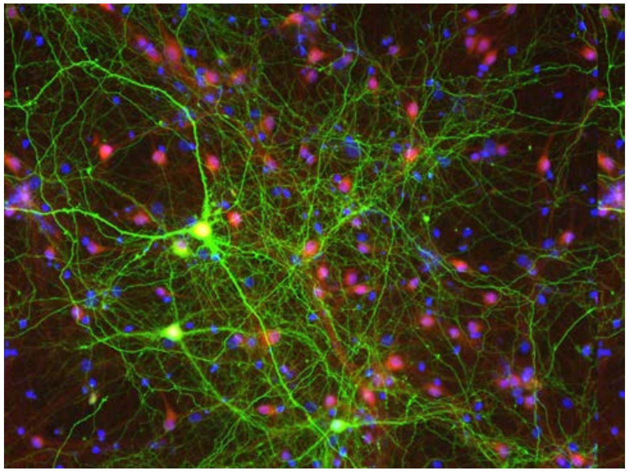

$$X_\infty = \begin{bmatrix} N_1^{t+1} & N_1^{t+2} & & N_1^{t+T} \\ N_2^{t+1} & N_2^{t+2} & \cdots & N_2^{t+T} \\ \vdots & \vdots & & \vdots \\ N_\infty^{t+1} & N_\infty^{t+2} & \cdots & N_\infty^{t+T} \end{bmatrix}$$

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

Stringer, Pachitariu, Steinmetz, Carandini, Harris, *Nature* (2019)

Population activity: Subsampling

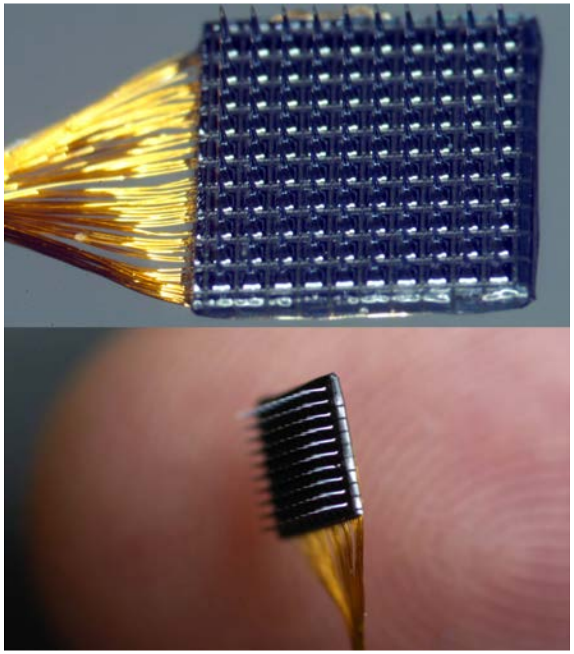

$$X_D = \begin{bmatrix} N_1^{t+1} & N_1^{t+2} & \cdots & N_1^{t+T} \\ N_2^{t+1} & N_2^{t+2} & & N_2^{t+T} \\ \vdots & \vdots & & \vdots \\ N_D^{t+1} & N_D^{t+2} & \cdots & N_D^{t+T} \end{bmatrix}$$

Data matrix X_D has N_D rows and T columns

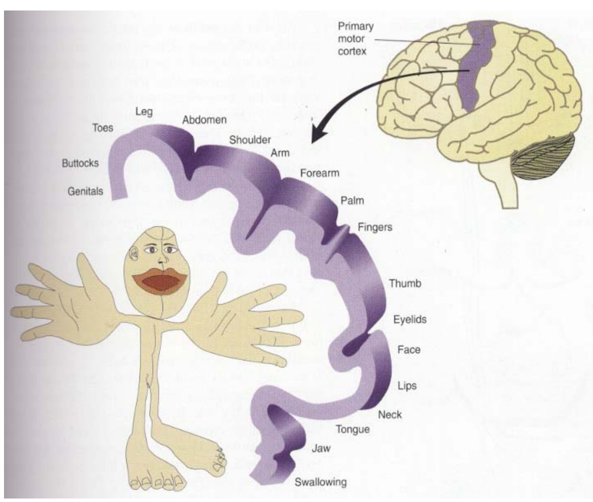

$D \approx 10^2$ for Multi-Electrode Arrays (MEAs)

$D \approx 10^3$ for Neuropixels

D is the ambient dimension

Population dynamics: the empirical neural space

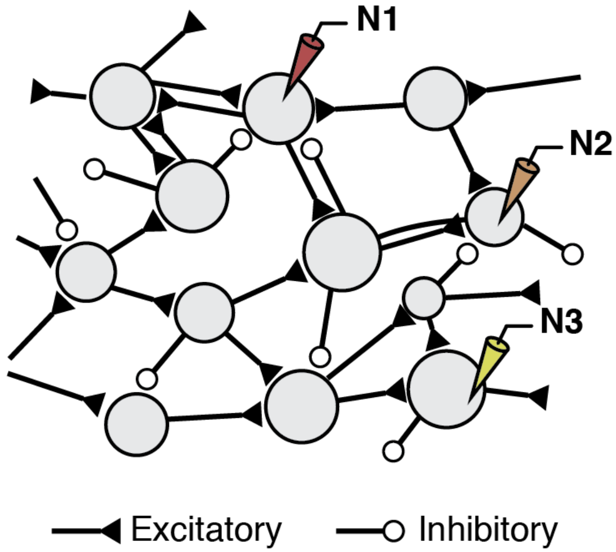

Ambient dimension D:
Number of recorded neurons

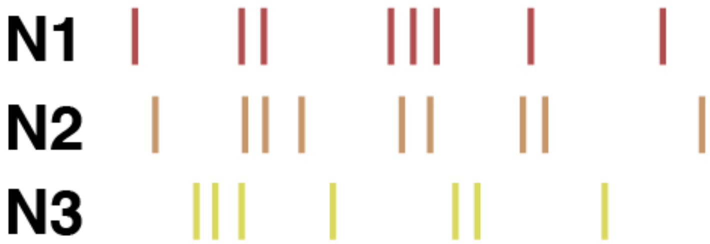

Population dynamics: the empirical neural space

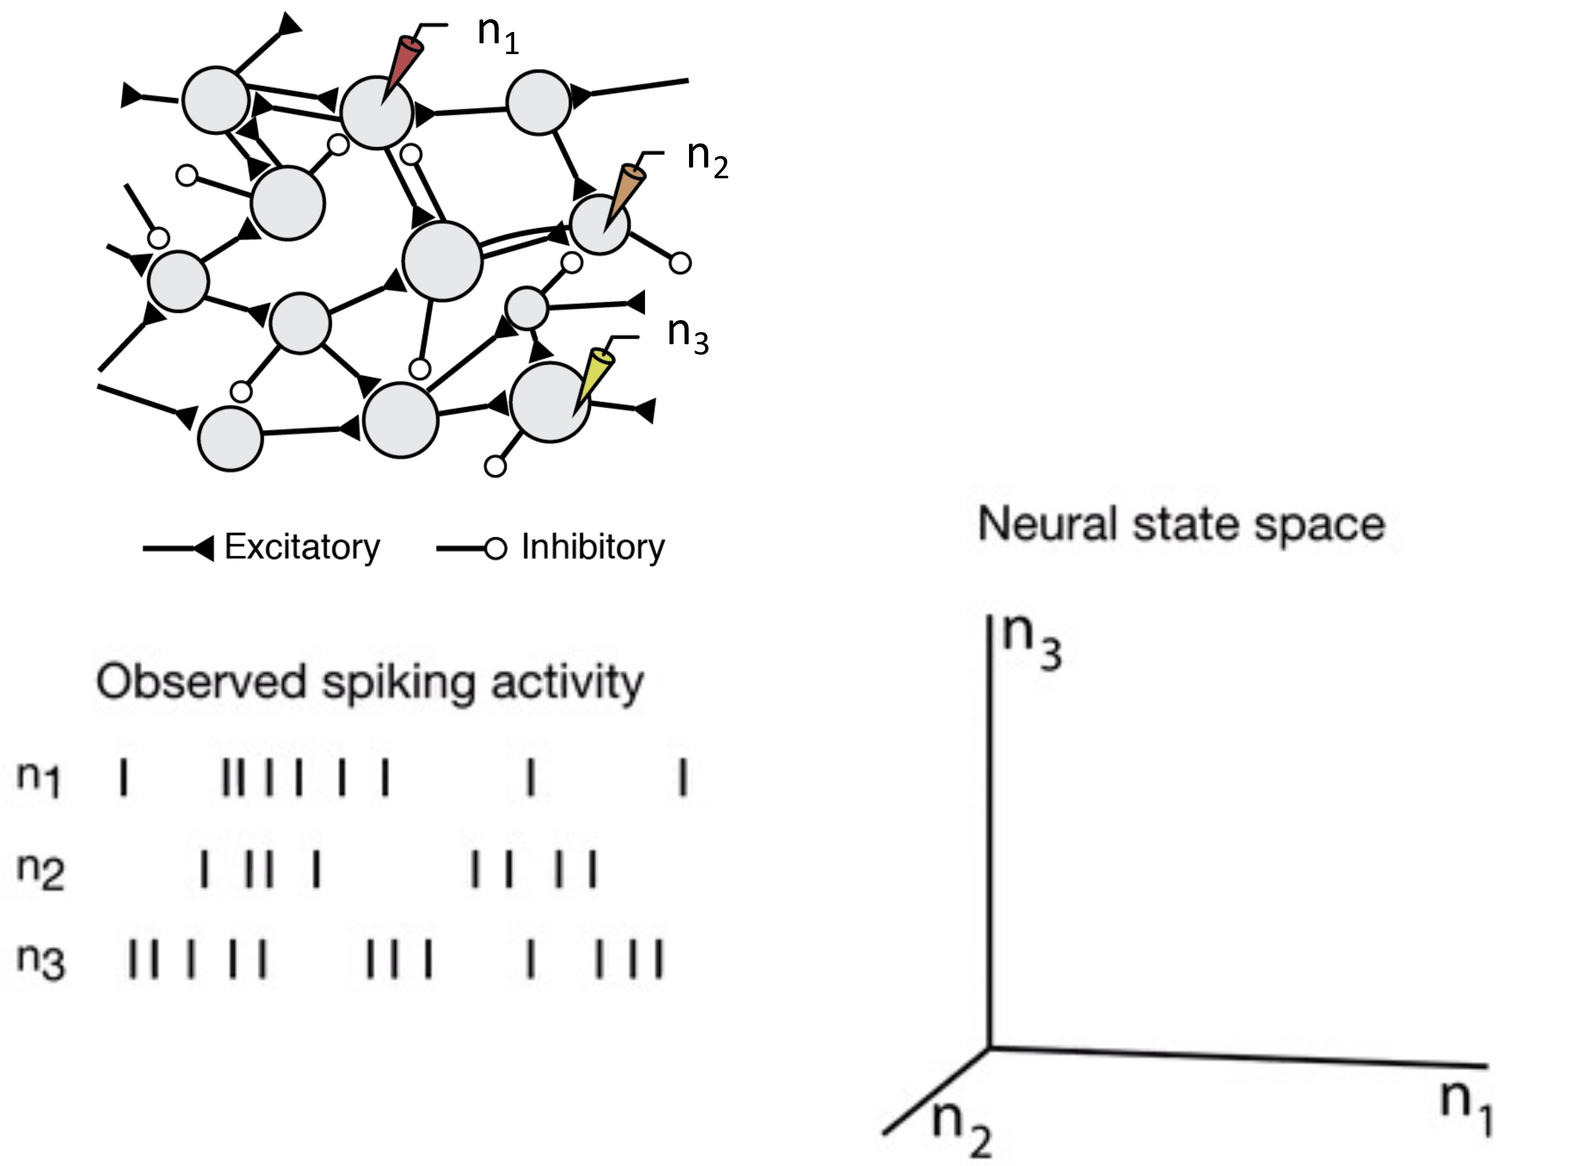

Cunningham & Yu, *Nat Neuro* 2014

Dimensionality reduction: neural modes and latent variables

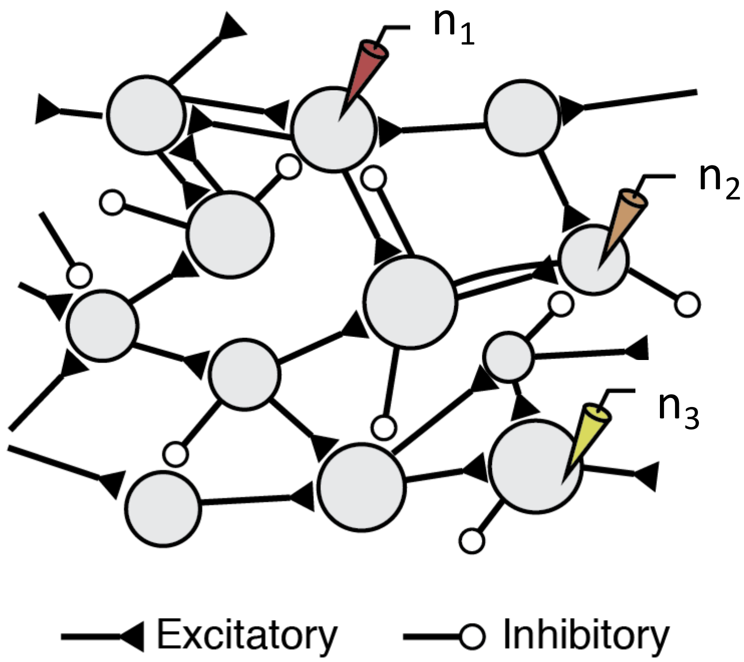

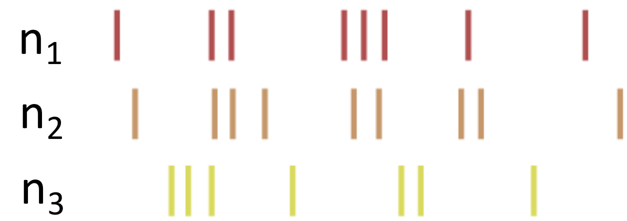

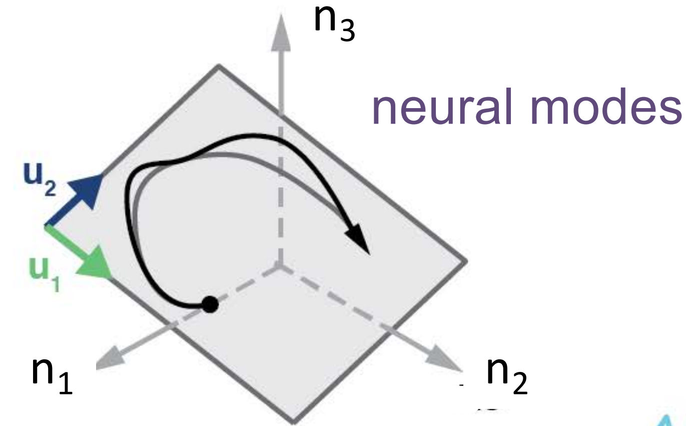

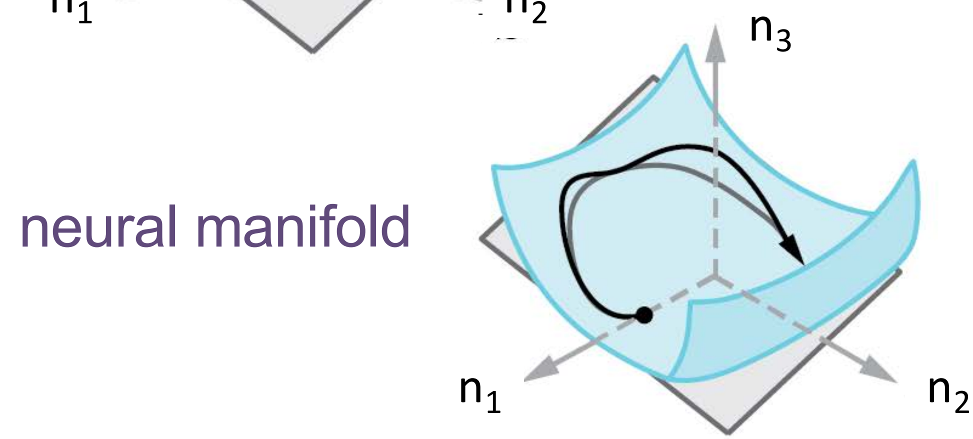

DIMENSIONALITY REDUCTION
linear or nonlinear?

Neural manifolds: linear or nonlinear?

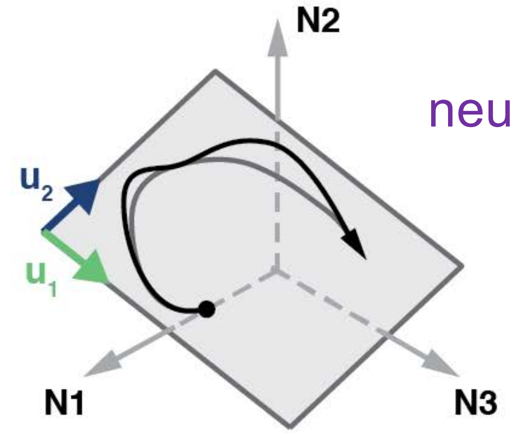

neural modes

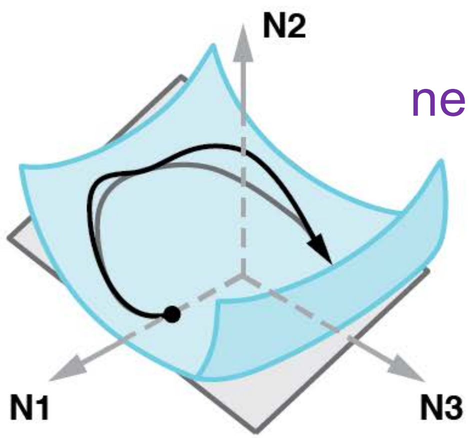

neural manifold

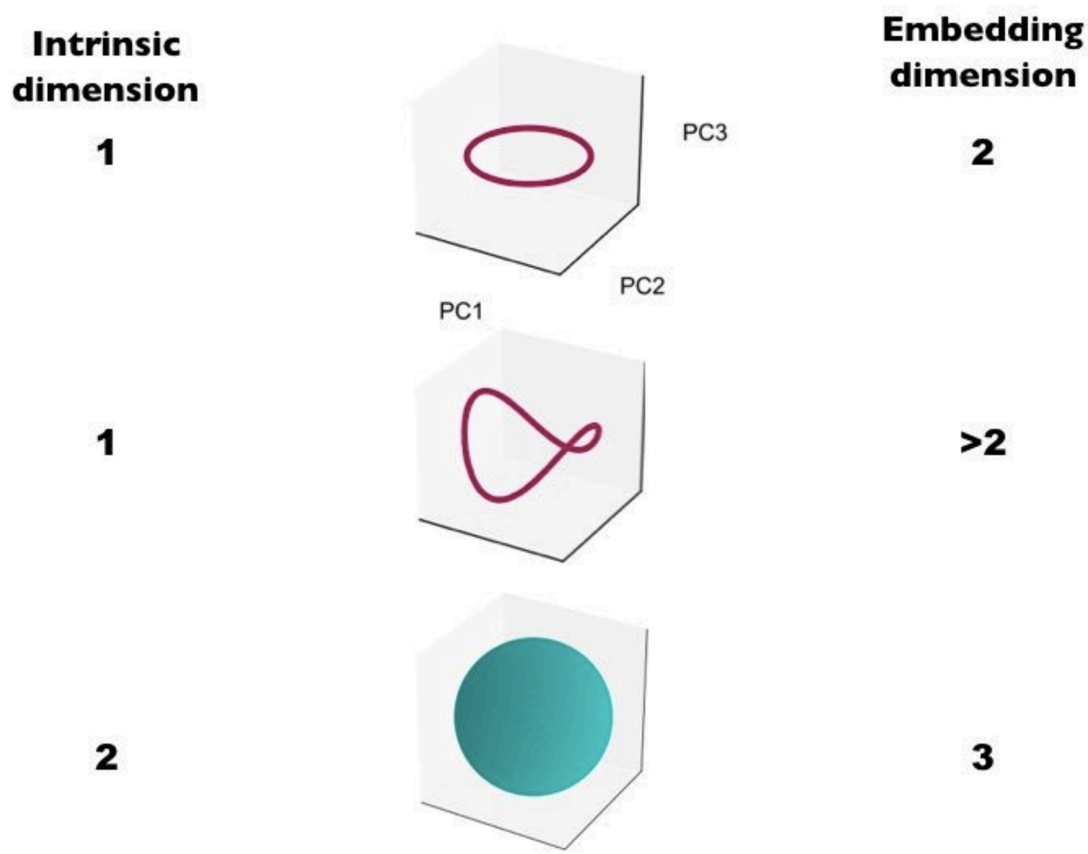

intrinsic

flat

Jazayeri, Ostojic, *Curr Opin Neurobiol* (2021)

Dimensionality reduction: neural modes and latent variables

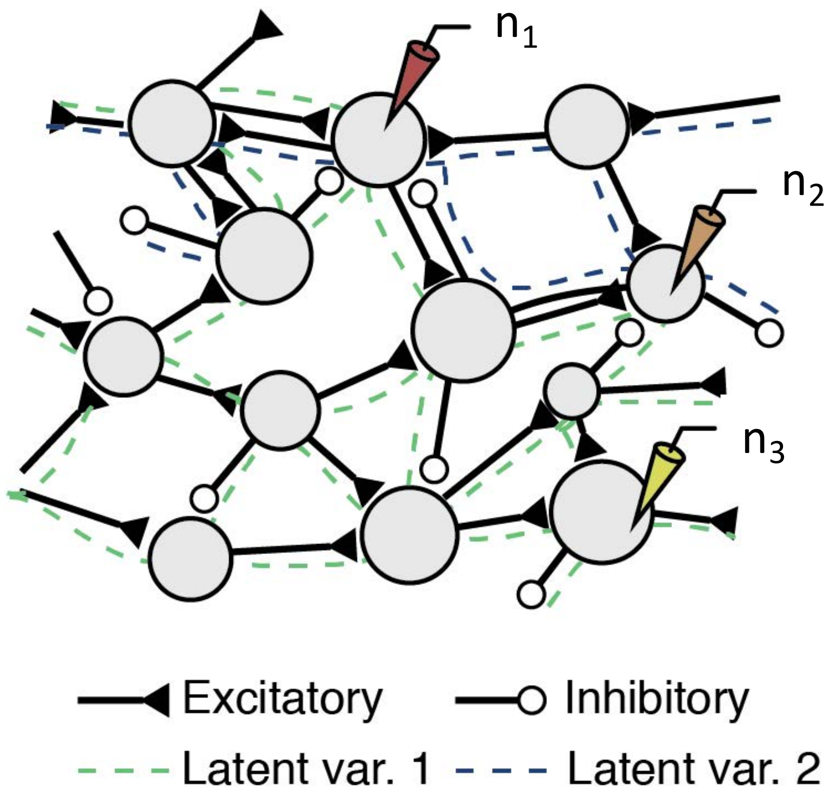

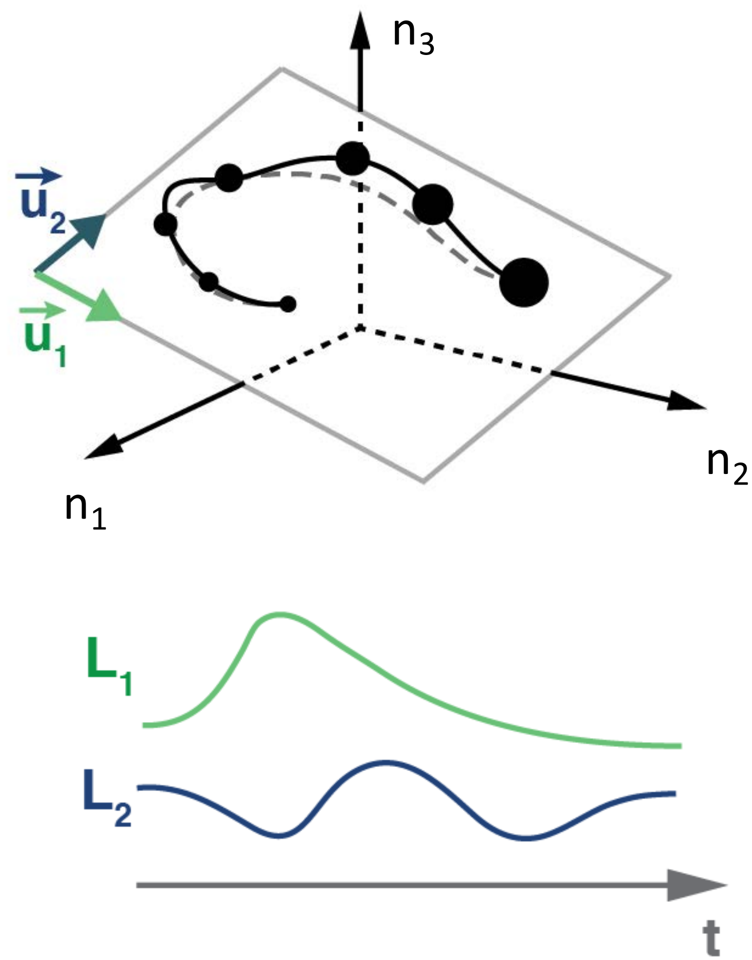

Population dynamics: latent variables as a generative model

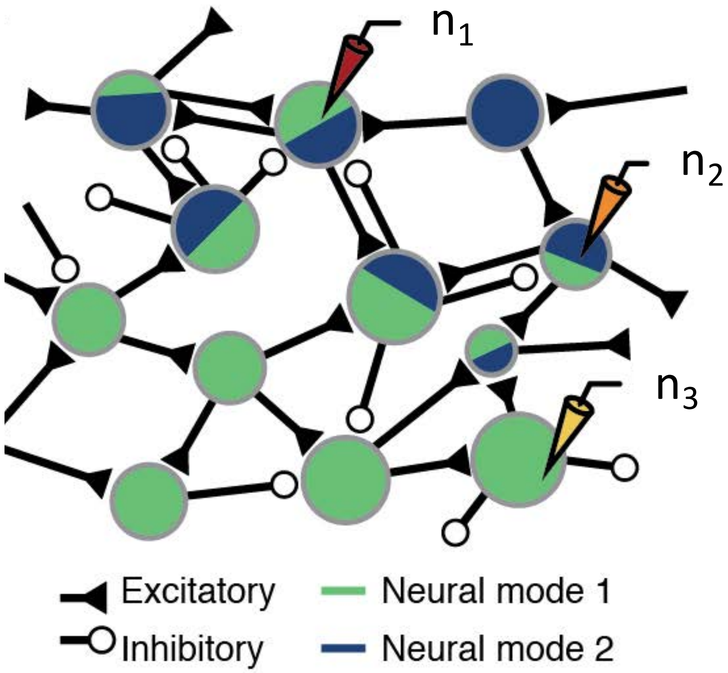

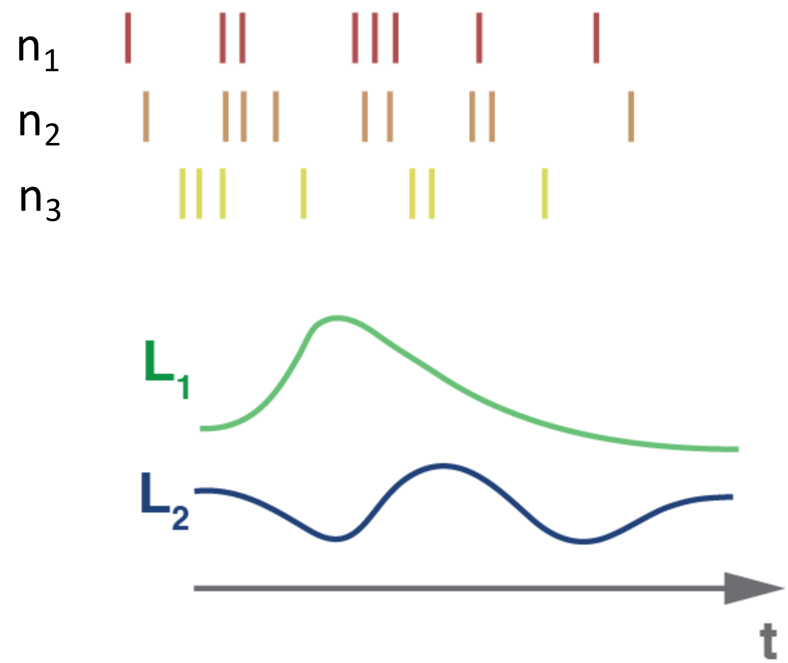

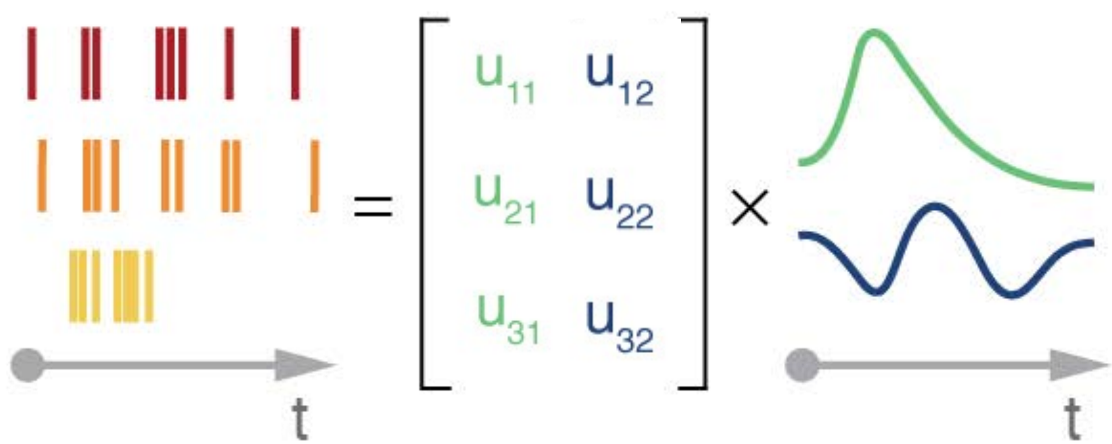

The neural manifold and latent dynamics

Neuron

Perspective

Neural manifolds for the control of movement

Juan A Gallego[1], Matthew G Perich[2], Lee E Miller[1,2,3], Sara A Solla[1,4]

[1] Department of Physiology, Northwestern University

[2] Department of Biomedical Engineering, Northwestern University

[3] Department of Physical Medicine and Rehabilitation, Northwestern University

[4] Department of Physics and Astronomy, Northwestern University

The analysis of neural dynamics in several brain cortices has consistently uncovered low dimensional manifolds that capture a significant fraction of neural variability. These neural manifolds are spanned by specific patterns of correlated neural activity, the "neural modes." We posit a model for neural control of movement in which the time varying activation of these neural modes is the generator of motor behavior. This manifold-based view of motor cortex may lead to a better understanding of how the brain controls movement.

Neuron 94, 978 (2017)

Latent dynamics: the time dependent activation of neural modes

The ubiquity of neural manifolds

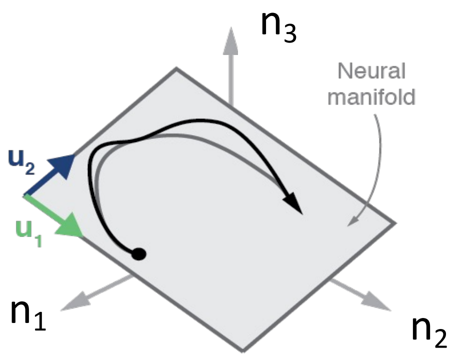

Neural modes have also been identified in other brain areas: frontal[25], prefrontal[26–29], parietal[30,31], visual[32–34], auditory[35], and olfactory[36] cortices.

25. Wang, J., Narain, D., Hosseini, E. & Jazayeri, M. Flexible control of speed of cortical dynamics . *bioRxiv* (2017). doi:10.1101/155390
26. Machens, C. K., Romo, R. & Brody, C. D. Functional, but not anatomical, separation of 'what' and 'when' in prefrontal cortex. *J. Neurosci.* **30,** 350–60 (2010).
27. Markowitz, D. a, Curtis, C. E. & Pesaran, B. Multiple component networks support working memory in prefrontal cortex. *Proc Natl Acad Sci USA* **112,** 11084–11089 (2015).
28. Mante, V., Sussillo, D., Shenoy, K. V & Newsome, W. T. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature* **503,** 78–84 (2013).
29. Durstewitz, D., Vittoz, N. M., Floresco, S. B. & Seamans, J. K. Abrupt transitions between prefrontal neural ensemble states accompany behavioral transitions during rule learning. *Neuron* **66,** 438–448 (2010).
30. Raposo, D., Kaufman, M. T. & Churchland, A. K. A category-free neural population supports evolving demands during decision-making. *Nat. Neurosci.* **17,** 1784–1792 (2014).
31. Harvey, C. D., Coen, P. & Tank, D. W. Choice-specific sequences in parietal cortex during a virtual-navigation decision task. *Nature* **484,** 62–8 (2012).
32. Churchland, M. M. *et al.* Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat. Neurosci.* **13,** 369–78 (2010).
33. Cohen, M. R. & Maunsell, J. H. R. A neuronal population measure of attention predicts behavioral performance on individual trials. *J. Neurosci.* **30,** 15241–53 (2010).
34. Cowley, B. R., Smith, M. A., Kohn, A. & Yu, B. M. Stimulus-Driven Population Activity Patterns in Macaque Primary Visual Cortex. *PLoS Comput. Biol.* **12,** e1005185 (2016).
35. Luczak, A., Barthó, P. & Harris, K. D. Spontaneous Events Outline the Realm of Possible Sensory Responses in Neocortical Populations. *Neuron* **62,** 413–425 (2009).
36. Kobak, D. *et al.* Demixed principal component analysis of neural population data. *Elife* **5,** 1–37 (2016).

Manifolds for neural population activity

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

Task: Behavioral stability

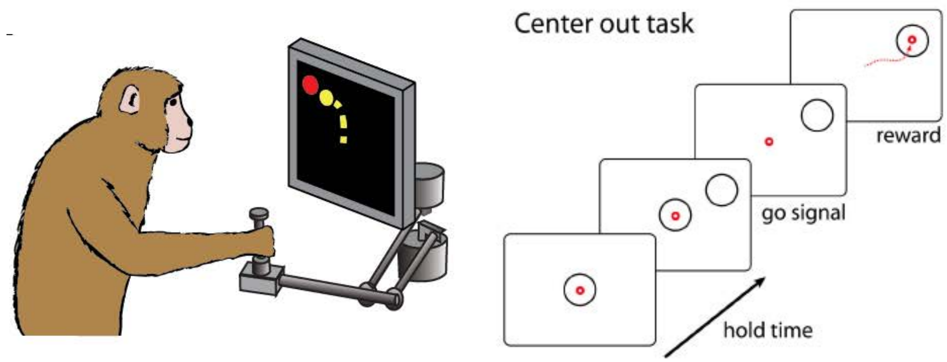

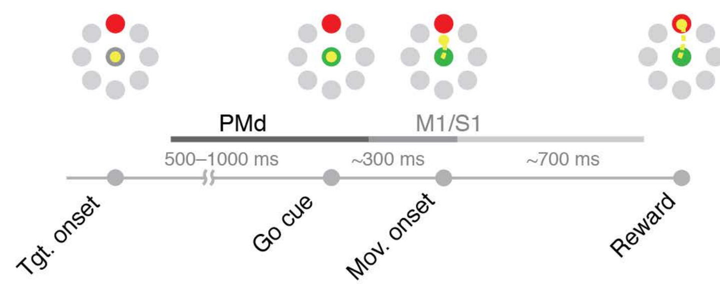

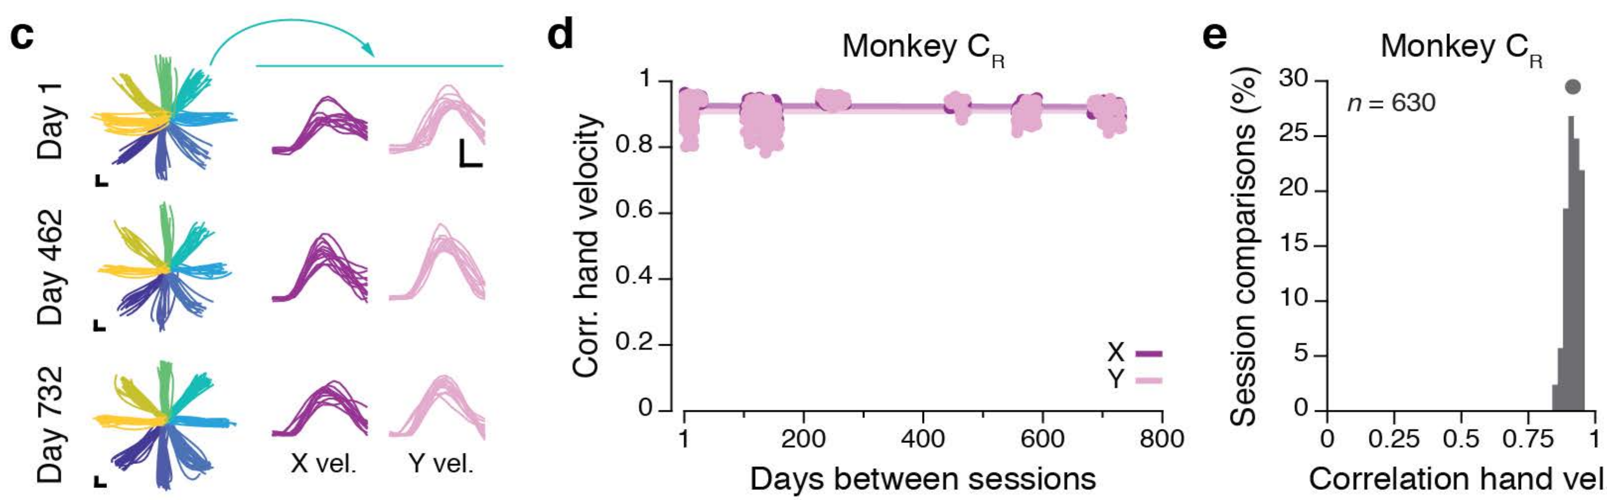

Stable neural dynamics underlying consistent behavior?

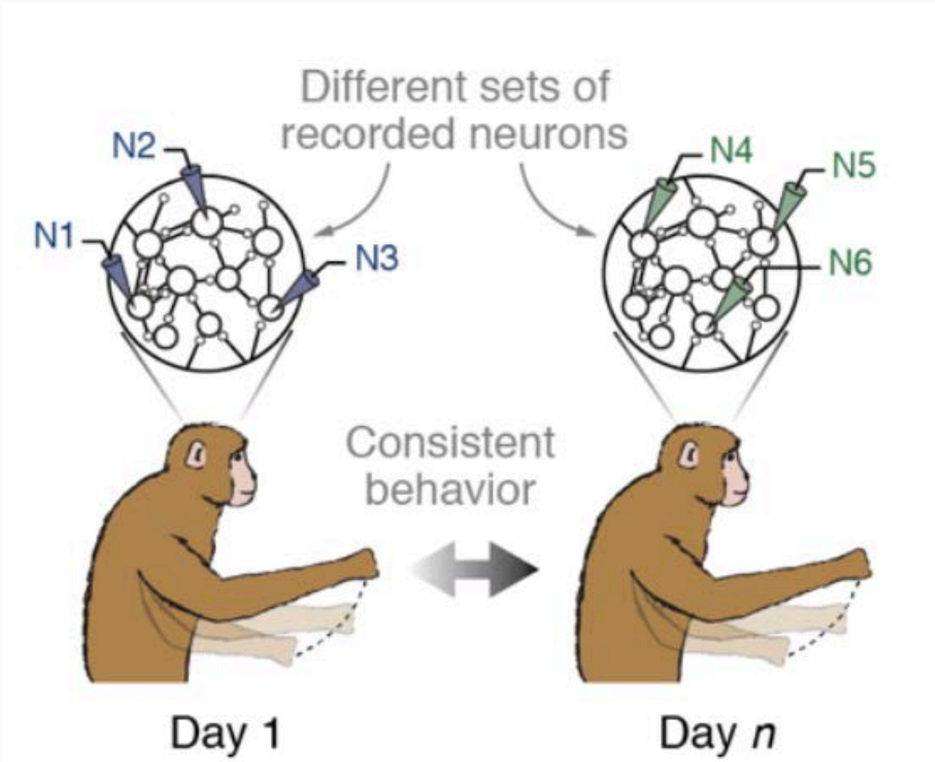

- Subjects consistently perform the same behavior over days, months, and years.
- Hypothesis: the true latent dynamics associated with consistent behavior should be stable.
- But: The same neurons cannot be recorded over this period.

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

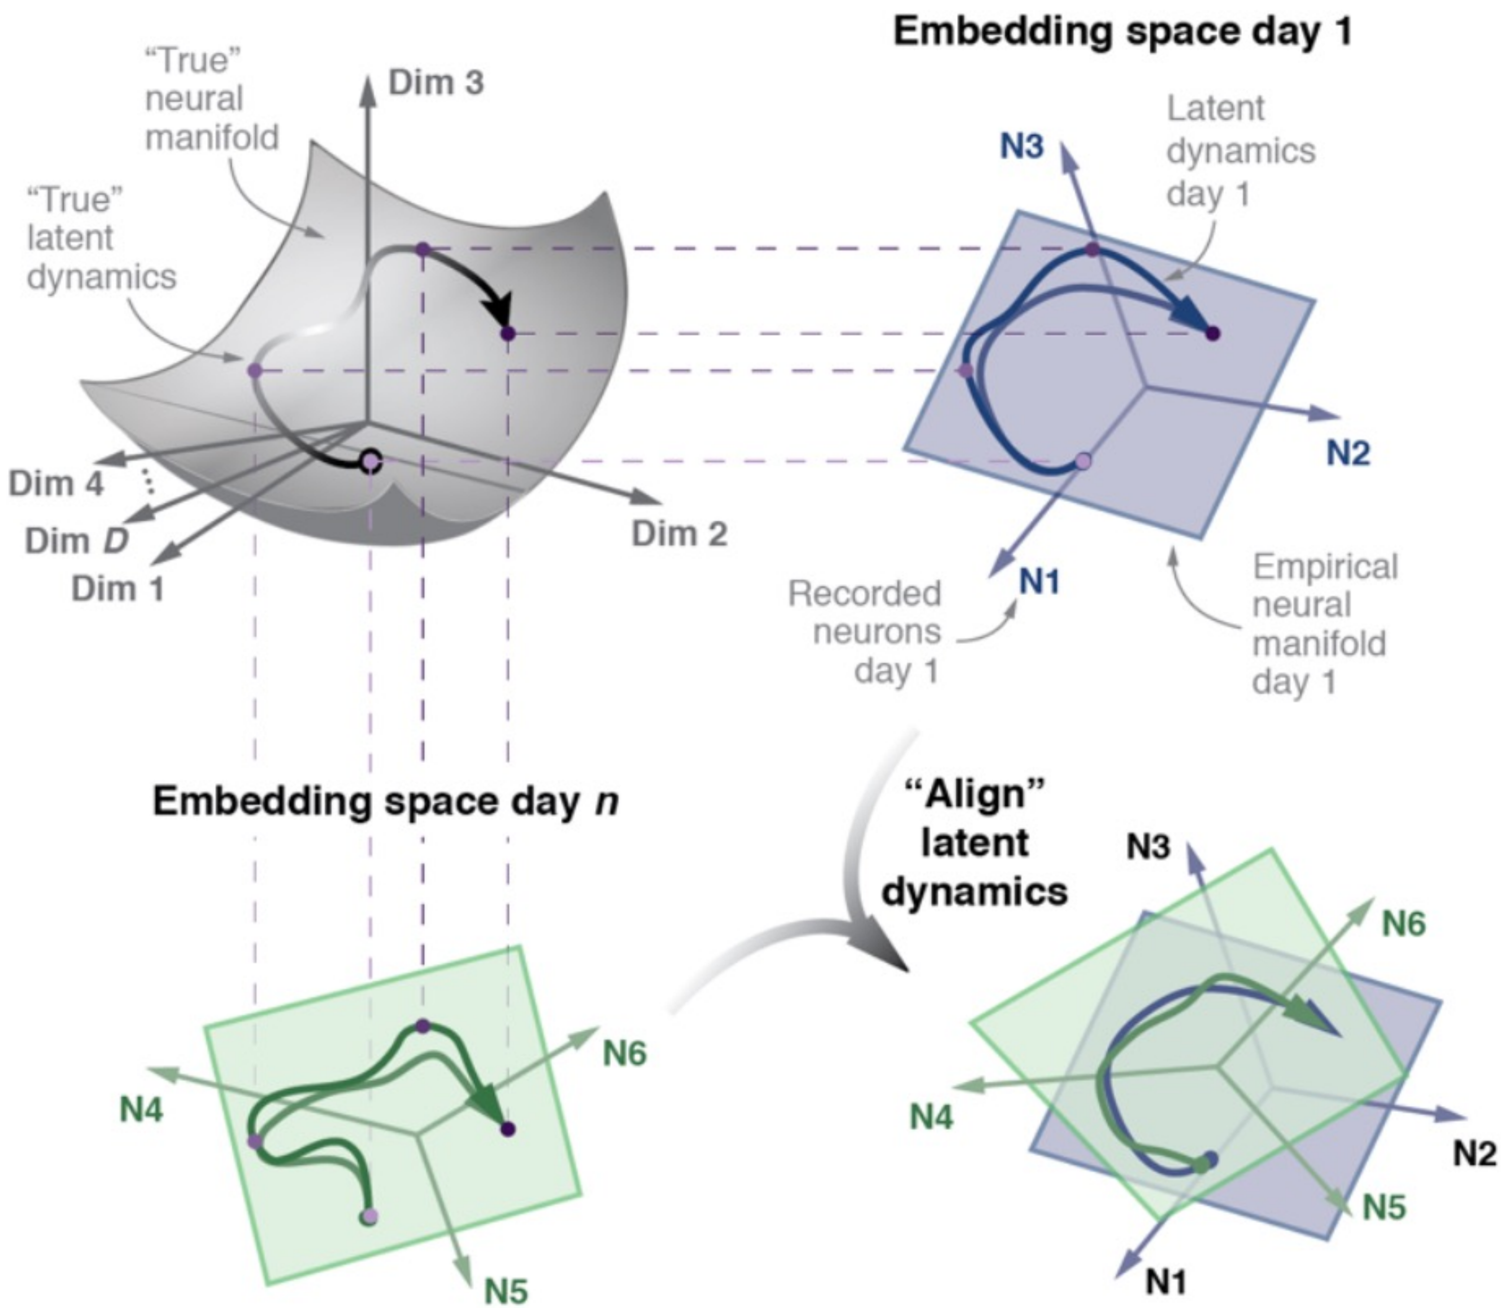

Day-specific neural modes and manifolds

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

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

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

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

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 $\widetilde{U}_m$.

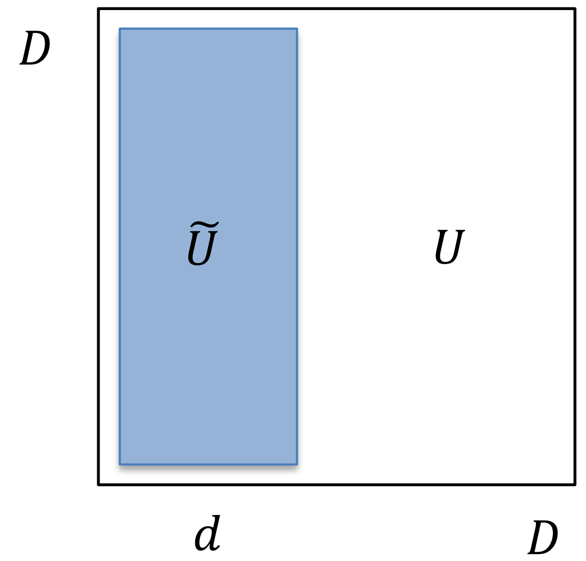

The day-specific low-dimensional manifolds are two hyperplanes:

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

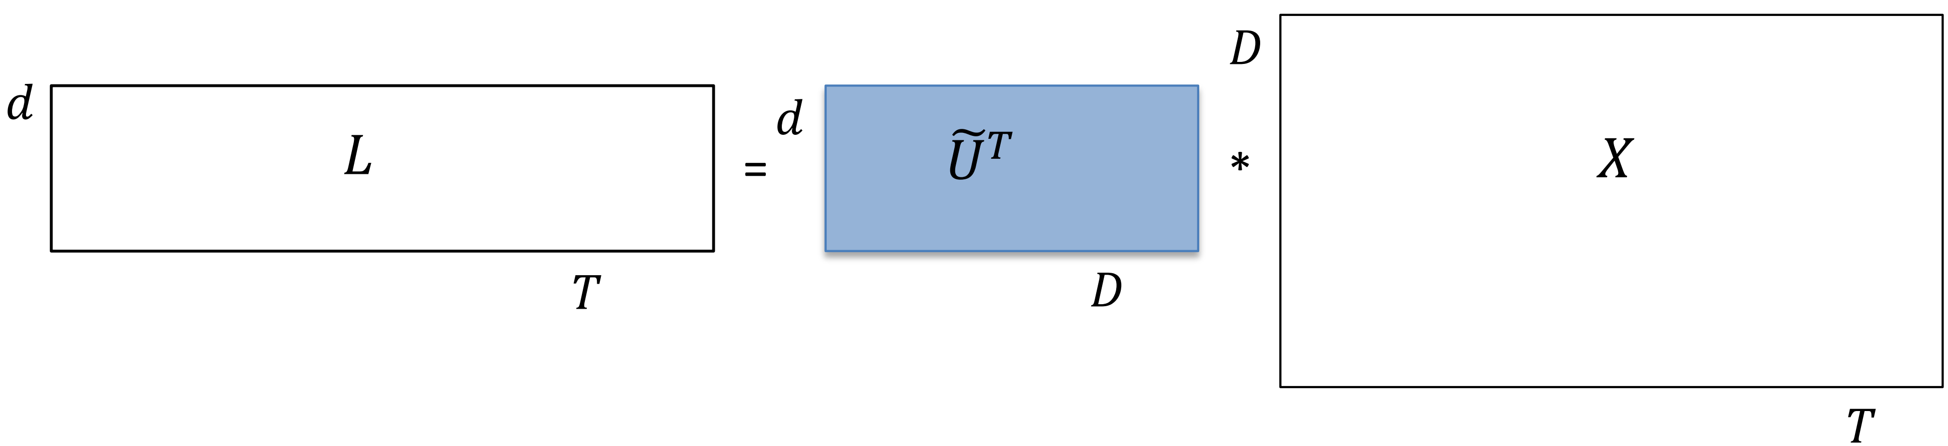

Canonical Correlation Analysis (CCA)

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,

$L_n^T = \; Q_n \, R_n$ and $L_m^T = \; Q_m \, R_m$

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:

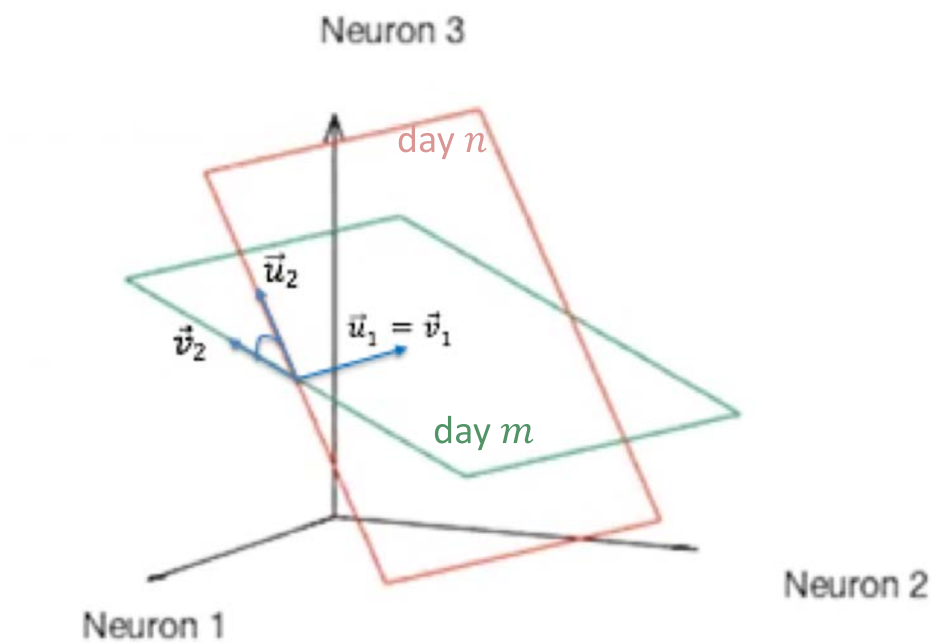

$$M_n = R_n^{-1} U$$

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

$$L_n \implies M_n^T L_n$$

$$L_m \implies M_m^T L_m$$

Bach, Jordan, *J Mach Learn Res* (2002)

Stability of M1 latent dynamics

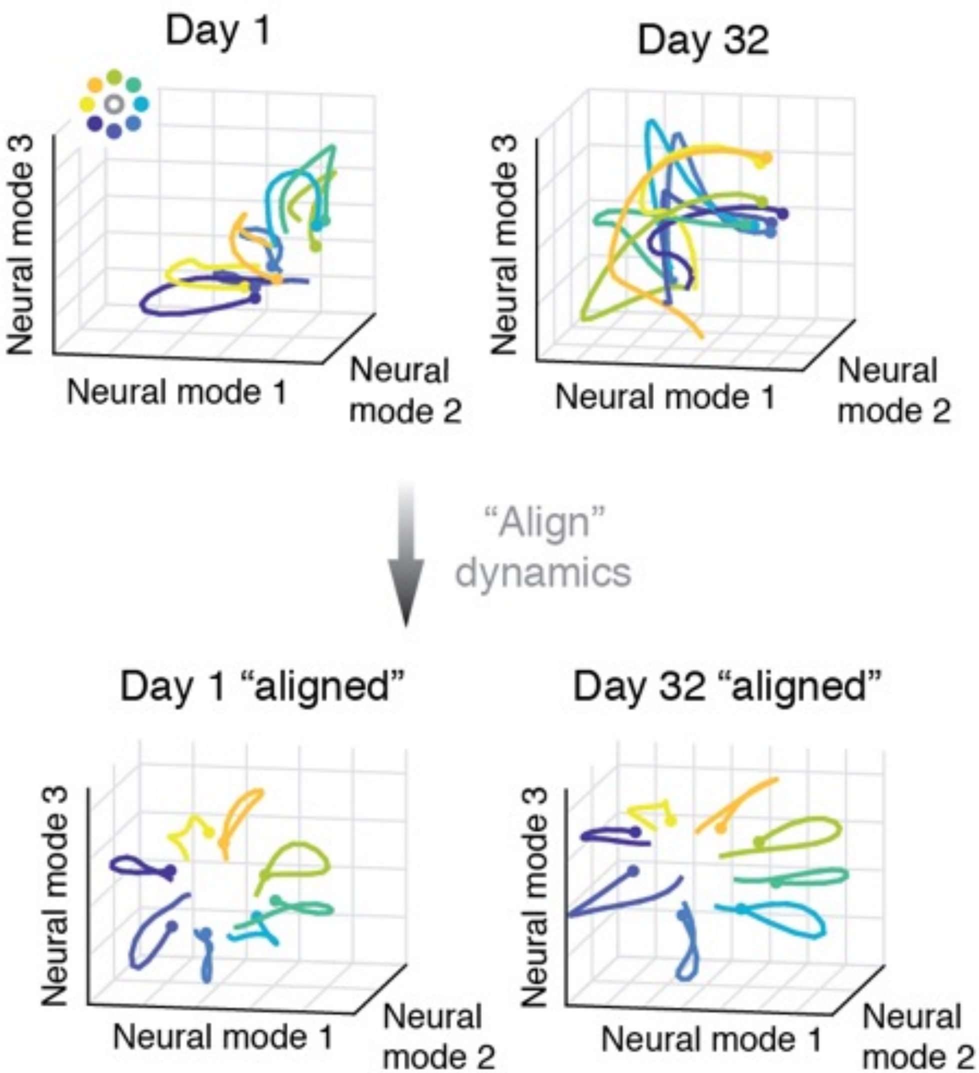

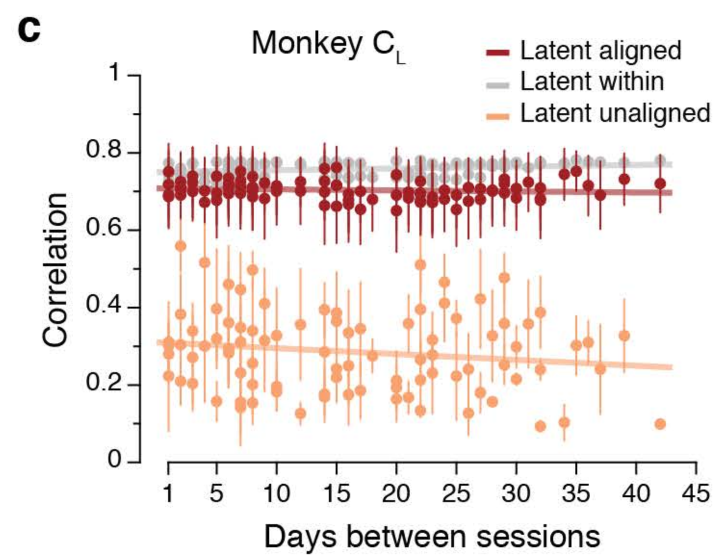

Gallego, Perich, et al, *Nat Neurosci* (2020)

Stability of M1 latent dynamics

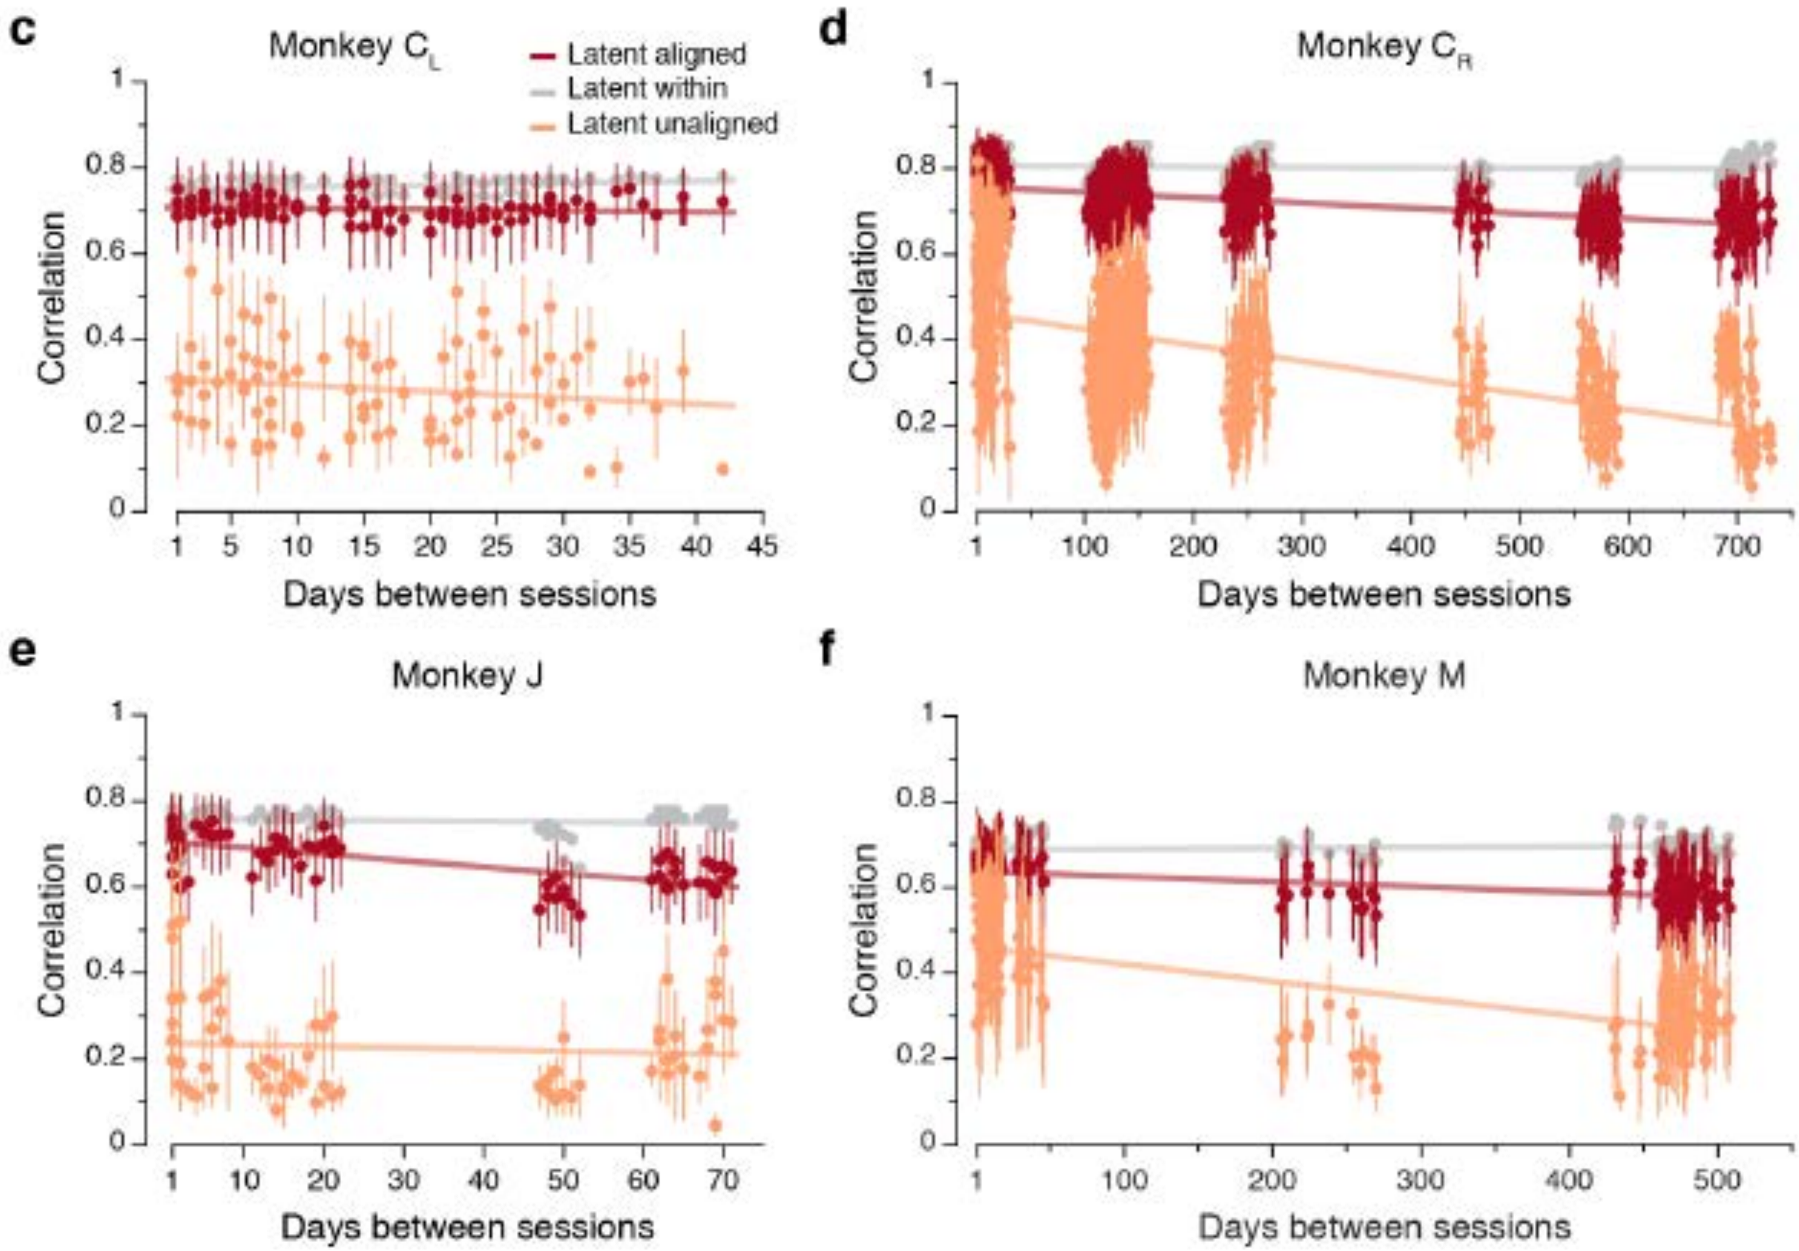

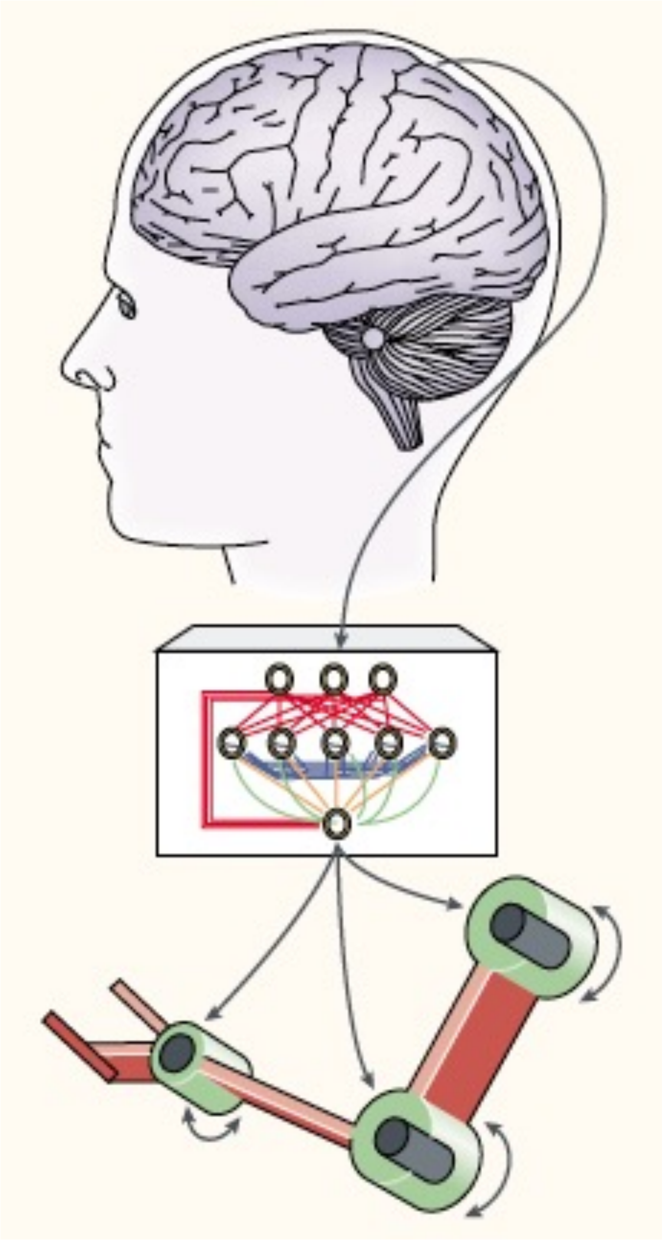

Population activity: relation to motor output

Stable prediction of movement kinematics

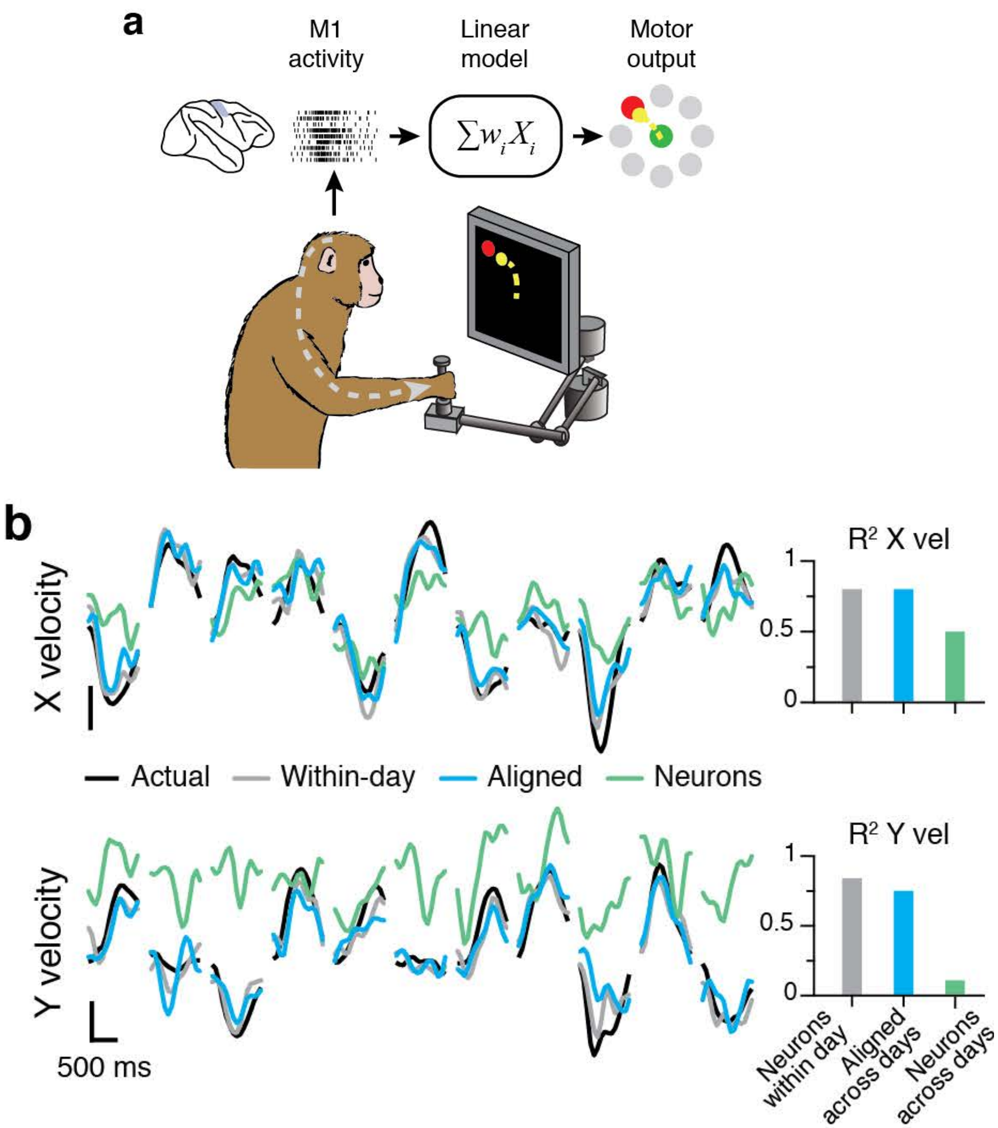

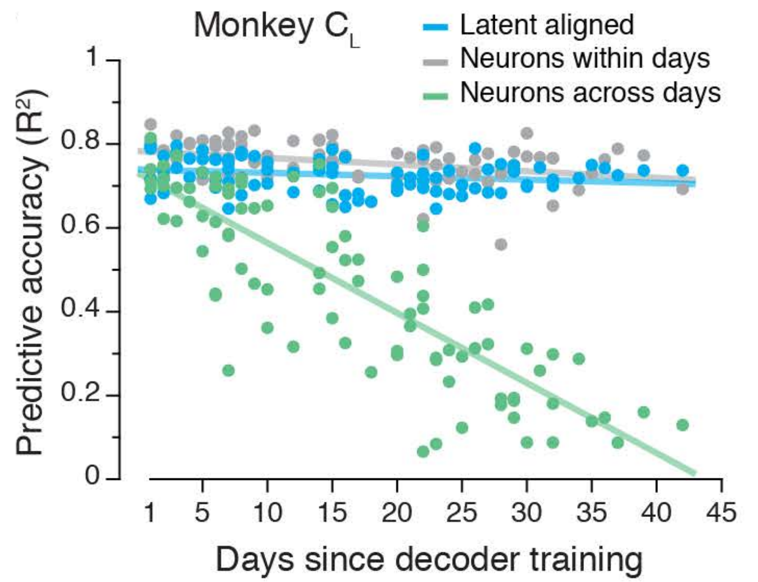

Prediction of muscle activity

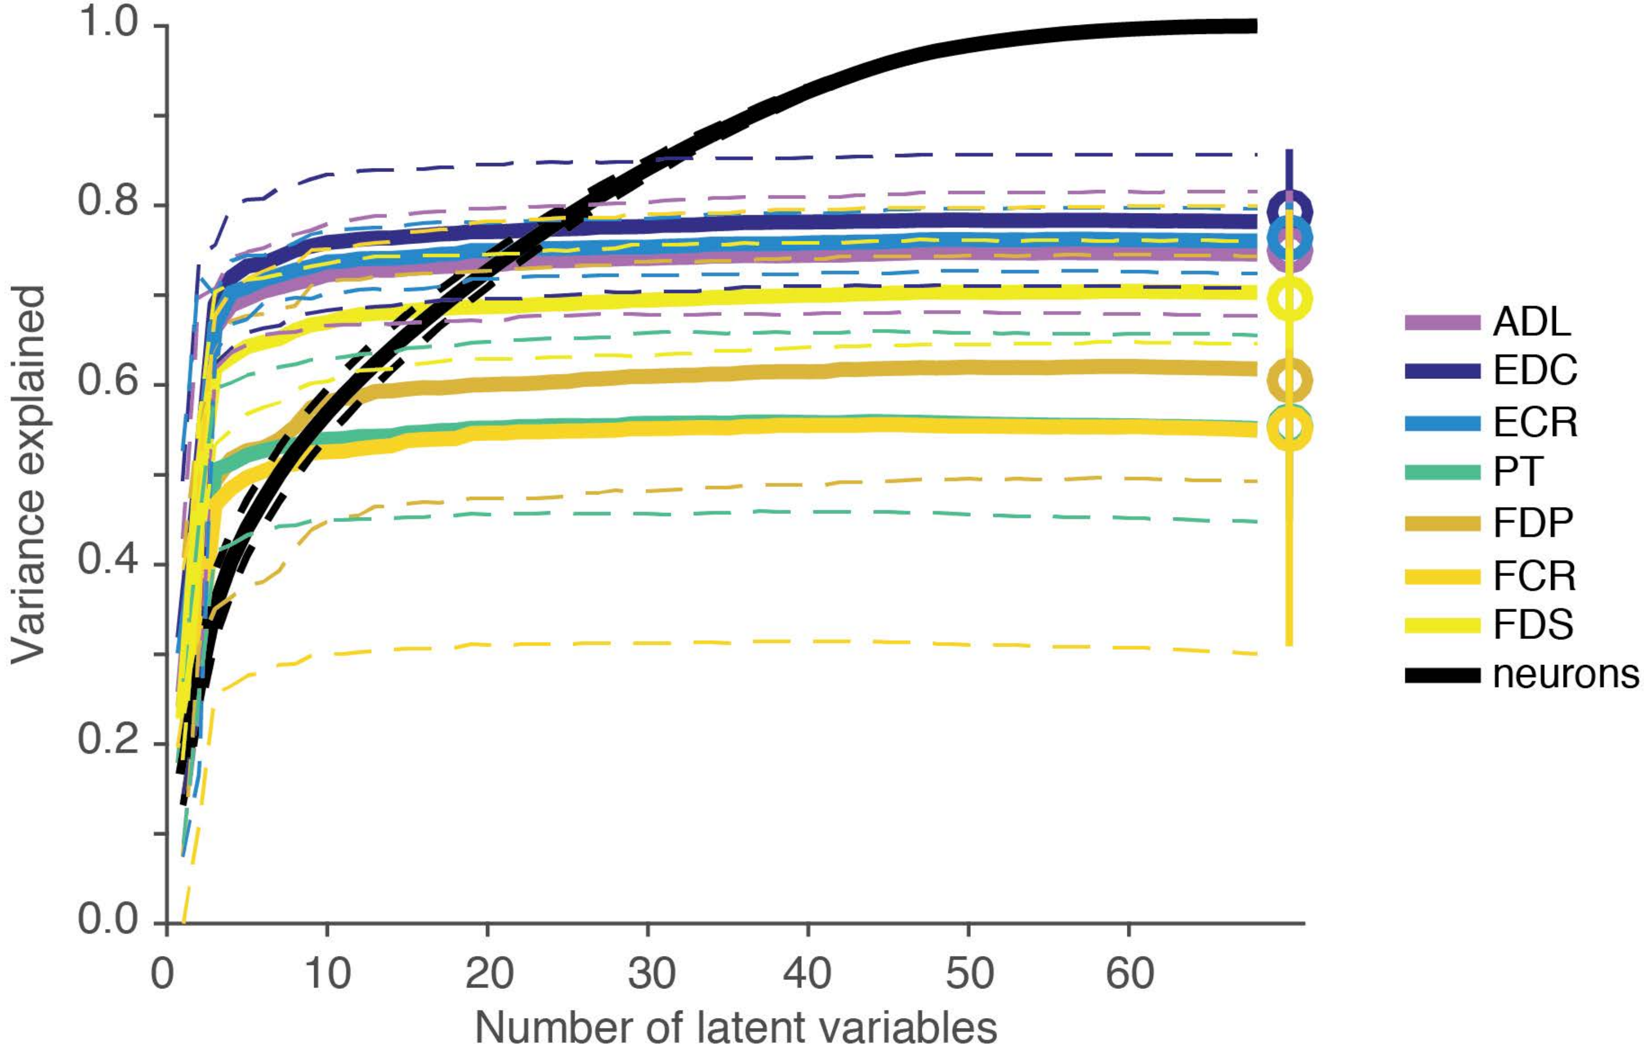

The future: natural behavior

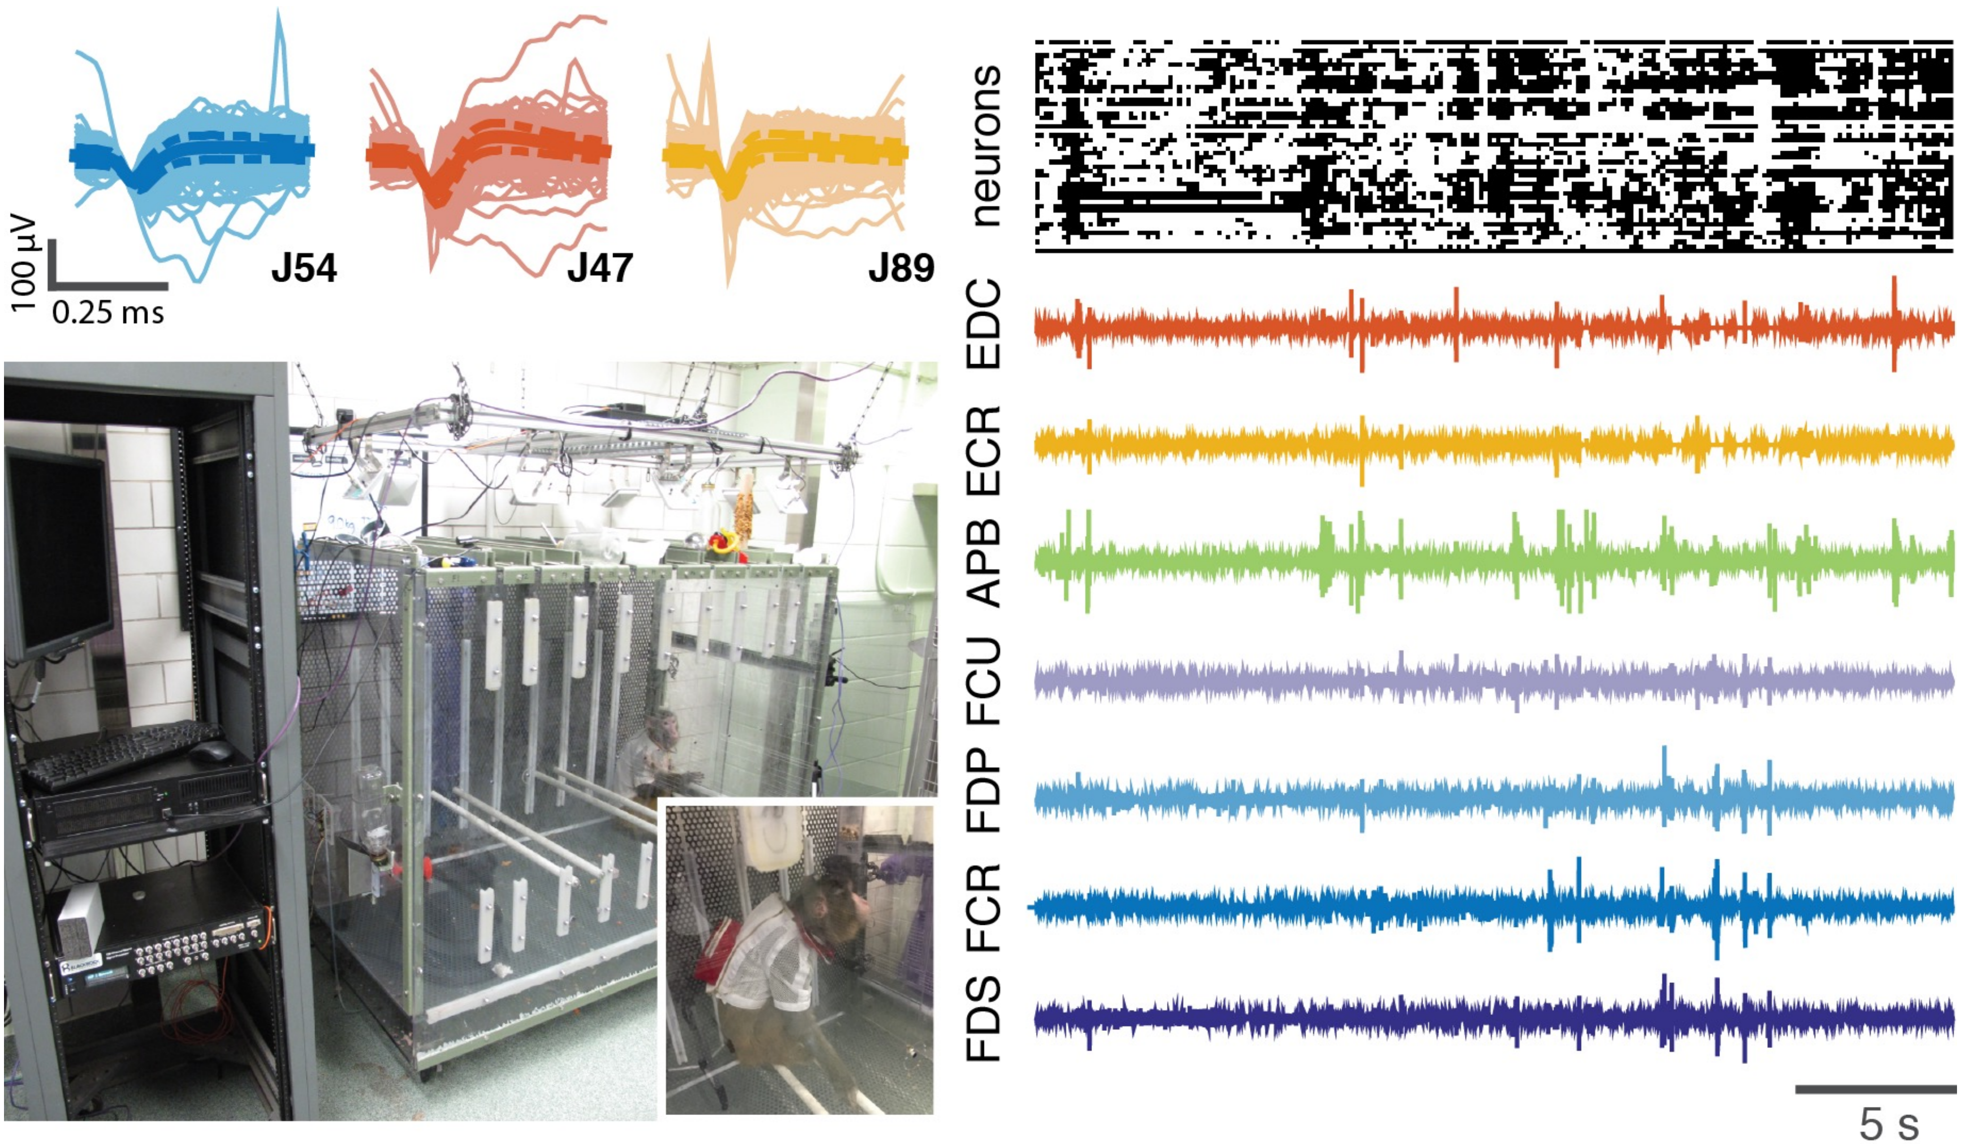

Alignment of latent dynamics

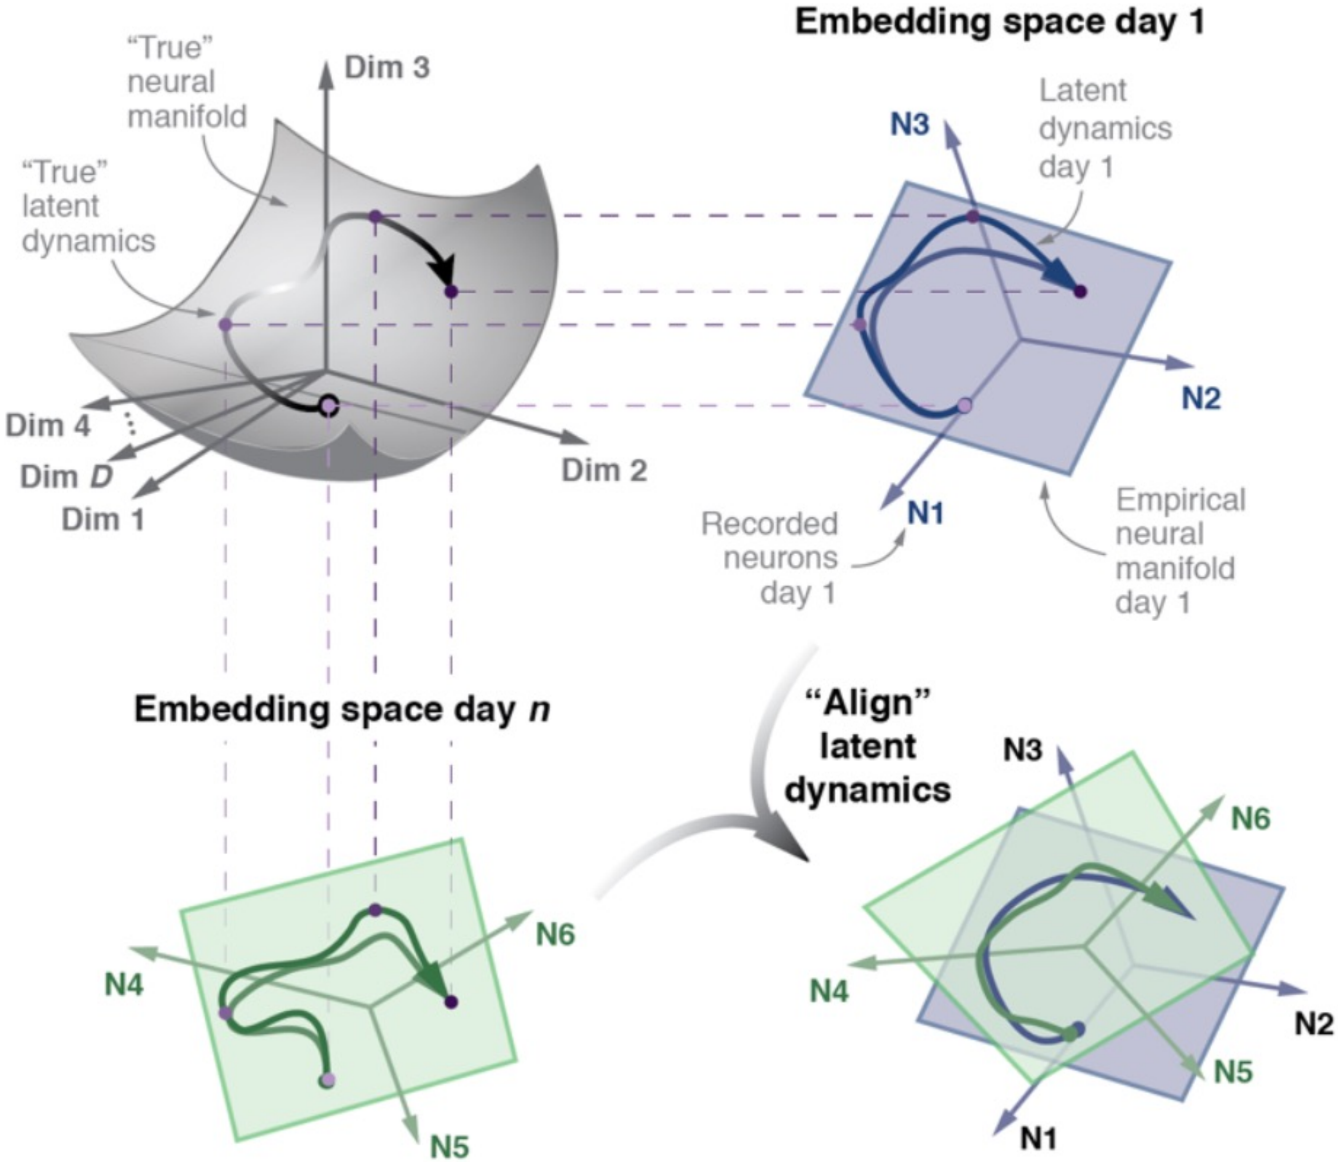

Gallego, Perich, Chowdhury, Solla, Miller, *Nature Neurosci* (2020)

Neural manifolds for the control of movement

Juan Gallego

Matthew Perich

Raeed Chowdhury

Lee Miller

Northwestern University