UCLA

Dynamic Random Partitions: Applications, Opportunities, and Challenges

Michele Guindani
Department of Biostatistics, University of California, Los Angeles

Nonparametric Bayesian Inference - Computational Issues
January 13th, 2026

Collaborators*

UCLA

Federica Zoe Ricci (Swarthmore College, Mathematics & Statistics)

Jaylen Lee (UCI Center for Statistical Consulting)

Erik Sudderth (UCI Computer Science and Statistics)

Megan Peters (UCI Cognitive Sciences)

Marina Vannucci (Rice University)

Alice Giampino (University of Milano Bicocca)

Bernardo Nipoti (University of Milano Bicocca)

Laura D'Angelo (University of Milano Bicocca)

Francesco Denti (University of Padova)

Antonio Canale (University of Padova)

* AKA "Those you should have invited instead of me"

Dynamic Partitions

UCLA

- In many applications, it is of interest to study the evolution of clusters of observations over time:

Dynamic Partitions

UCLA

➡️ In many applications, it is of interest to study the evolution of clusters of observations over time:

👉 **Environmental Sciences**, to see how **locations** cluster differently from other locations over time

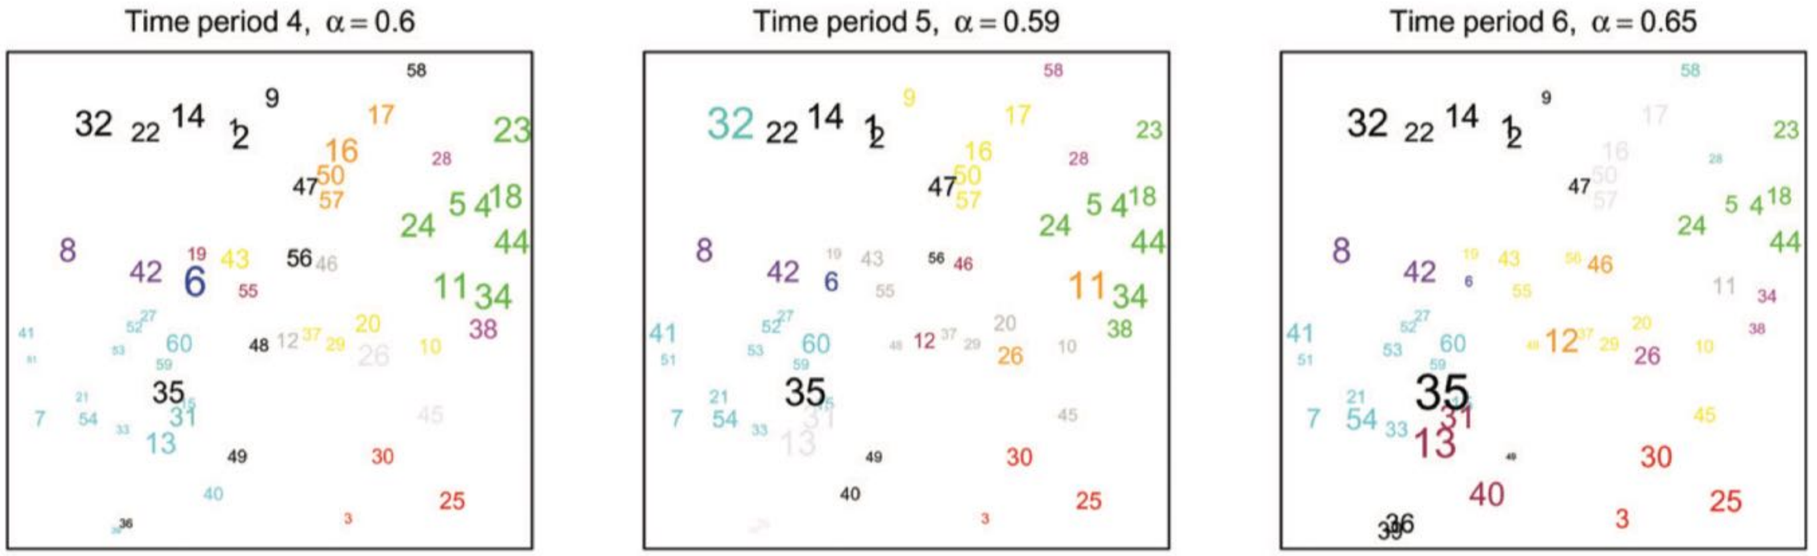

Page, Quintana, Dahl, JCGS, 2022

Dynamic Partitions

UCLA

➡️ In many applications, it is of interest to study the evolution of clusters of observations over time:

👉 **Mobile data, sports sciences,** to study the **coordination of movements** in time to improve sports performance

Dynamic Partitions

UCLA

➡️ In many applications, it is of interest to study the evolution of clusters of observations over time:

👉 **Mobile data, sports sciences,** to study the **coordination of movements** in time to improve sports performance

Dynamic Partitions

UCLA

➡️ In many applications, it is of interest to study the evolution of clusters of observations over time:

👉 **Brain imaging data**, to study how the activity of different brain regions **clusters over time and across different subgroups of subjects**,

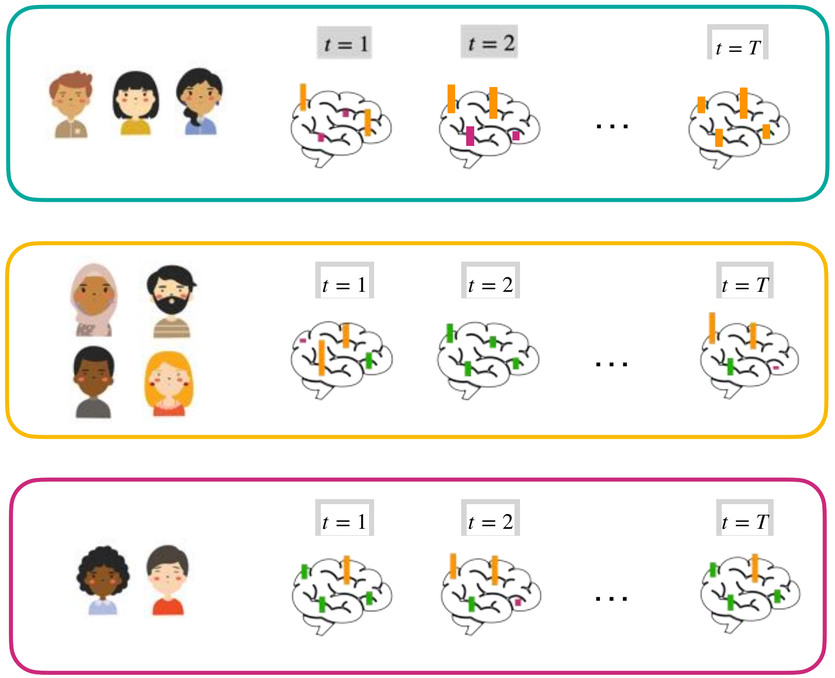

Dynamic Partitions

UCLA

➡️ In many applications, it is of interest to study the evolution of clusters of observations over time:

👉 **Brain imaging data**, to study how the activity of different brain regions **clusters over time and across different subgroups of subjects**, or to study how neuronal activations are differently clustered **as a function of behavior**

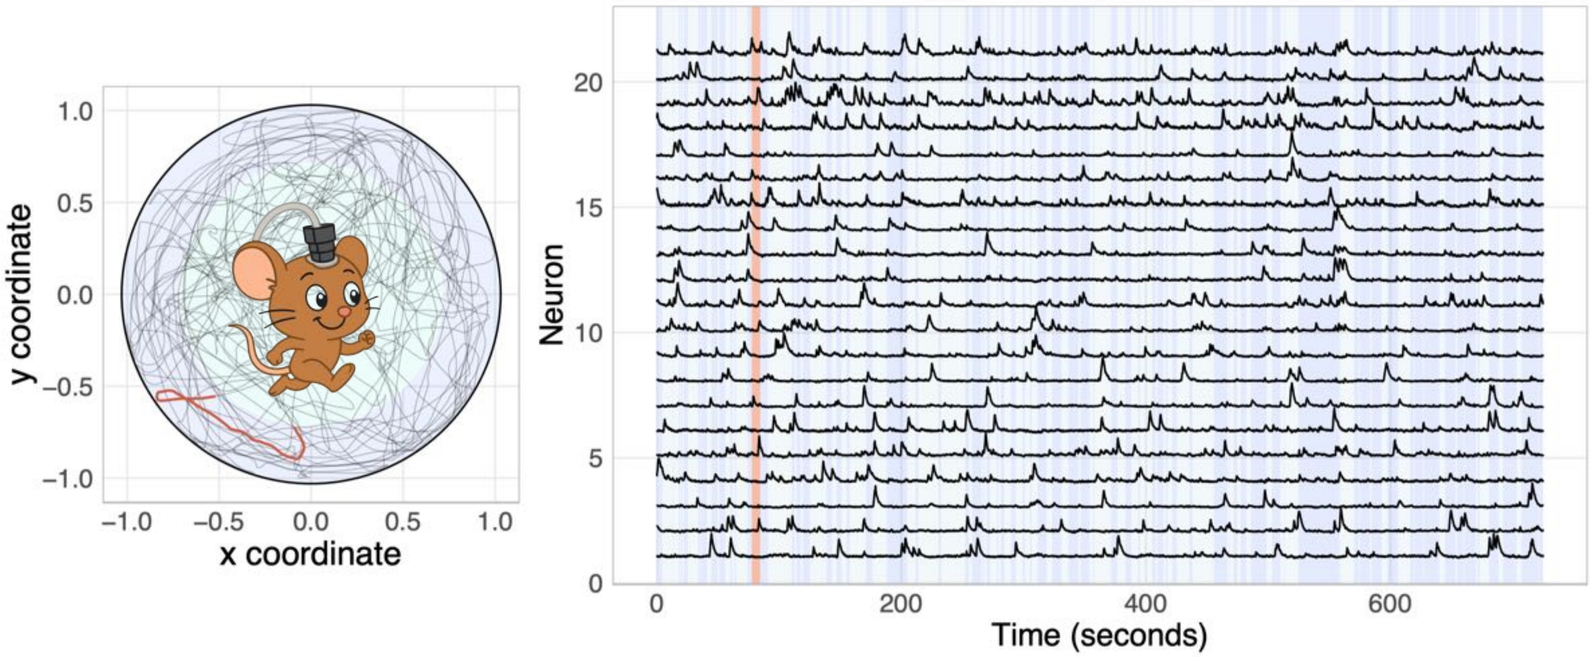

Dynamic Partitions

UCLA

➡️ In many applications, it is of interest to study the evolution of clusters of observations over time:

👉 **Brain imaging data**, to study how the activity of different brain regions **clusters over time and across different subgroups of subjects**, or to study how neuronal activations are differently clustered **as a function of behavior**

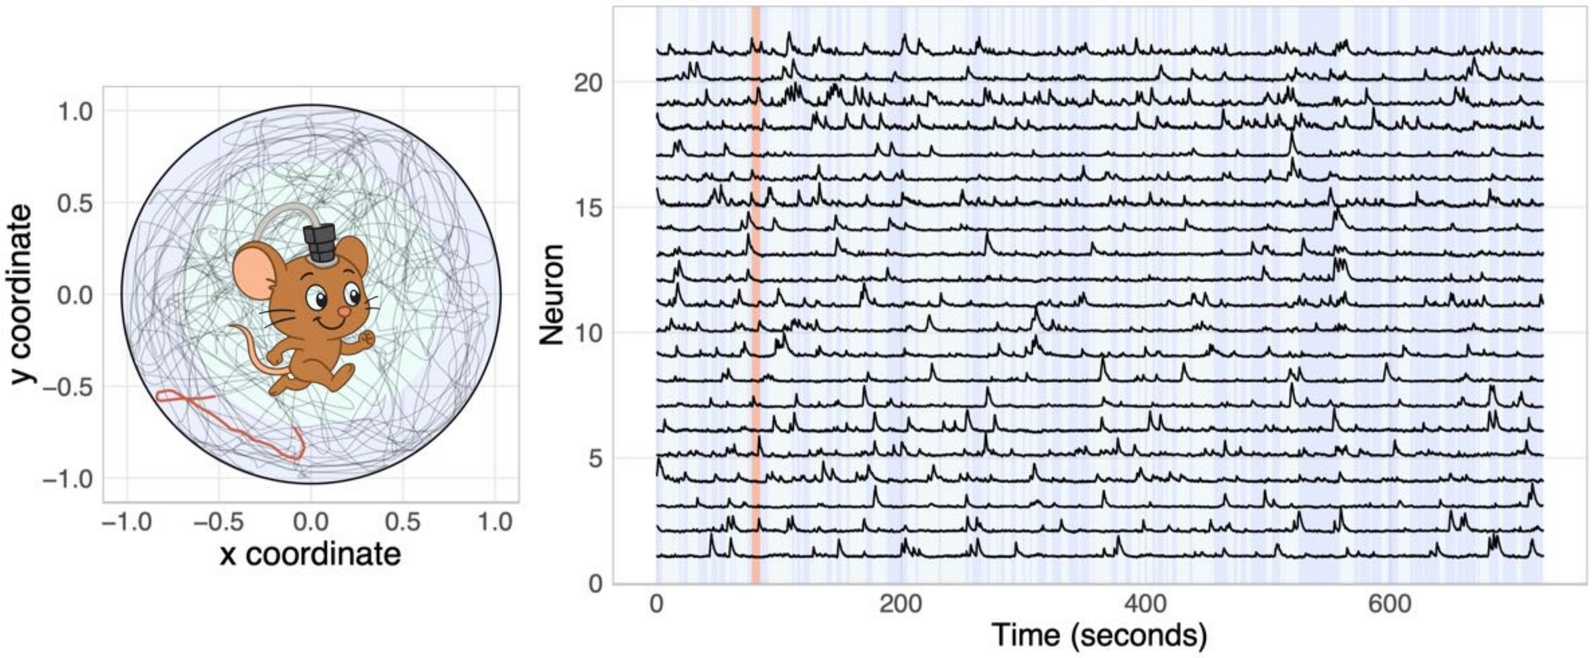

Every application may require a different modeling of the temporal dynamics

Modeling of dynamic evolution of partition parameters

UCLA

- **Time-varying Dirichlet process mixture models**
  (Caron, Davy, Doucet 2007)

$$\mathbb{G}_t = \sum_{k=1}^{\infty} V_{k,t} \delta_{U_{k,t}}$$

Modeling of dynamic evolution of partition parameters

UCLA

- **Time-varying Dirichlet process mixture models**
  (Caron, Davy, Doucet 2007)

$$\mathbb{G}_t = \sum_{k=1}^{\infty} V_{k,t}\delta_{U_{k,t}} \qquad U_{k,t} \sim \begin{cases} p\left(U_{k,t} \mid U_{k,t-1}\right) & \text{if } k \in \mathcal{I}\left(\mathbf{m}_{t-1}^t\right) \\ \mathbb{G}_0 & \text{otherwise} \end{cases}$$

Modeling of dynamic evolution of partition parameters

UCLA

Time-varying Dirichlet process mixture models
(Caron, Davy, Doucet 2007)

$$\mathbb{G}_t = \sum_{k=1}^{\infty} V_{k,t}\delta_{U_{k,t}} \qquad U_{k,t} \sim \begin{cases} p\left(U_{k,t} \mid U_{k,t-1}\right) & \text{if } k \in \mathcal{I}\left(\mathbf{m}_{t-1}^t\right) \\ \mathbb{G}_0 & \text{otherwise} \end{cases}$$

Other contributions (purposely non-exhaustive):

Fox, E., Sudderth, E.B., Jordan, M.I., and Willsky, A.S. (2011) - Switching DLM
Nieto-Barajas & Contreras-Cristan (2014) - Hierarchical LMM w. clustering of parameters
Cassese, Zhu, G., Vannucci (2019) - Dynamic model selection

Modeling of dynamic partitions

UCLA

- **Dynamic temporal allocation of the units** (Page, Quintana, Dahl, JCGS 2022)

Modeling of dynamic partitions

UCLA

➡ **Dynamic temporal allocation of the units** (Page, Quintana, Dahl, JCGS 2022)

‼ Introduce an auxiliary variable that identifies which of the experimental units at time *t-1* will be considered for possible cluster reallocation at time *t*

$$\gamma_{it} = \begin{cases} 1 & \text{if unit } i \text{ is not reallocated when moving from time } t-1 \text{ to } t \\ 0 & \text{otherwise} \end{cases}$$

across *i=1,…, m* units

Modeling of dynamic partitions

UCLA

➡ **Dynamic temporal allocation of the units** (Page, Quintana, Dahl, JCGS 2022)

‼ Introduce an auxiliary variable that identifies which of the experimental units at time *t-1* will be considered for possible cluster reallocation at time *t*

$$\gamma_{it} = \begin{cases} 1 & \text{if unit } i \text{ is not reallocated when moving from time } t-1 \text{ to } t \\ 0 & \text{otherwise} \end{cases}$$

across *i=1,…, m* units

See also Paganin S., Page G., Quintana, F.A. (2024): Informed Random Partition Models with Temporal Dependence.

Temporal biclustering

in multi-subject neuroscience studies

fMRI studies

UCLA

Measure small changes in blood flow related to brain activity

Experiments can be designed to study brain activity during a task

Multiple subjects often undergo the same experiment

fMRI studies

UCLA

Measure small **changes in blood flow** related to brain activity

Experiments can be designed to study brain activity during a task

Multiple subjects often undergo the same experiment

- What patterns of brain-region activation are there in any given moment of the experiment?

Brain region clusters

Within

fMRI studies

UCLA

Measure small **changes in blood flow** related to brain activity

Experiments can be designed to study brain activity during a task

Multiple subjects often undergo the same experiment

- What patterns of brain-region activation are there in any given moment of the experiment?
- How do patterns of brain-region activation change during the experiment?

Brain region clusters

Dynamic brain region clusters

Within

fMRI studies

UCLA

Measure small **changes in blood flow** related to brain activity

Experiments can be designed to study brain activity during a task

Multiple subjects often undergo the same experiment

- What patterns of brain-region activation are there in any given moment of the experiment?
- How do patterns of brain-region activation change during the experiment?
- How do patterns of brain-region activation vary across subjects?

Brain region clusters

Dynamic brain region clusters

Subject clusters

Within

Across

Multi-subject temporal biclustering

UCLA

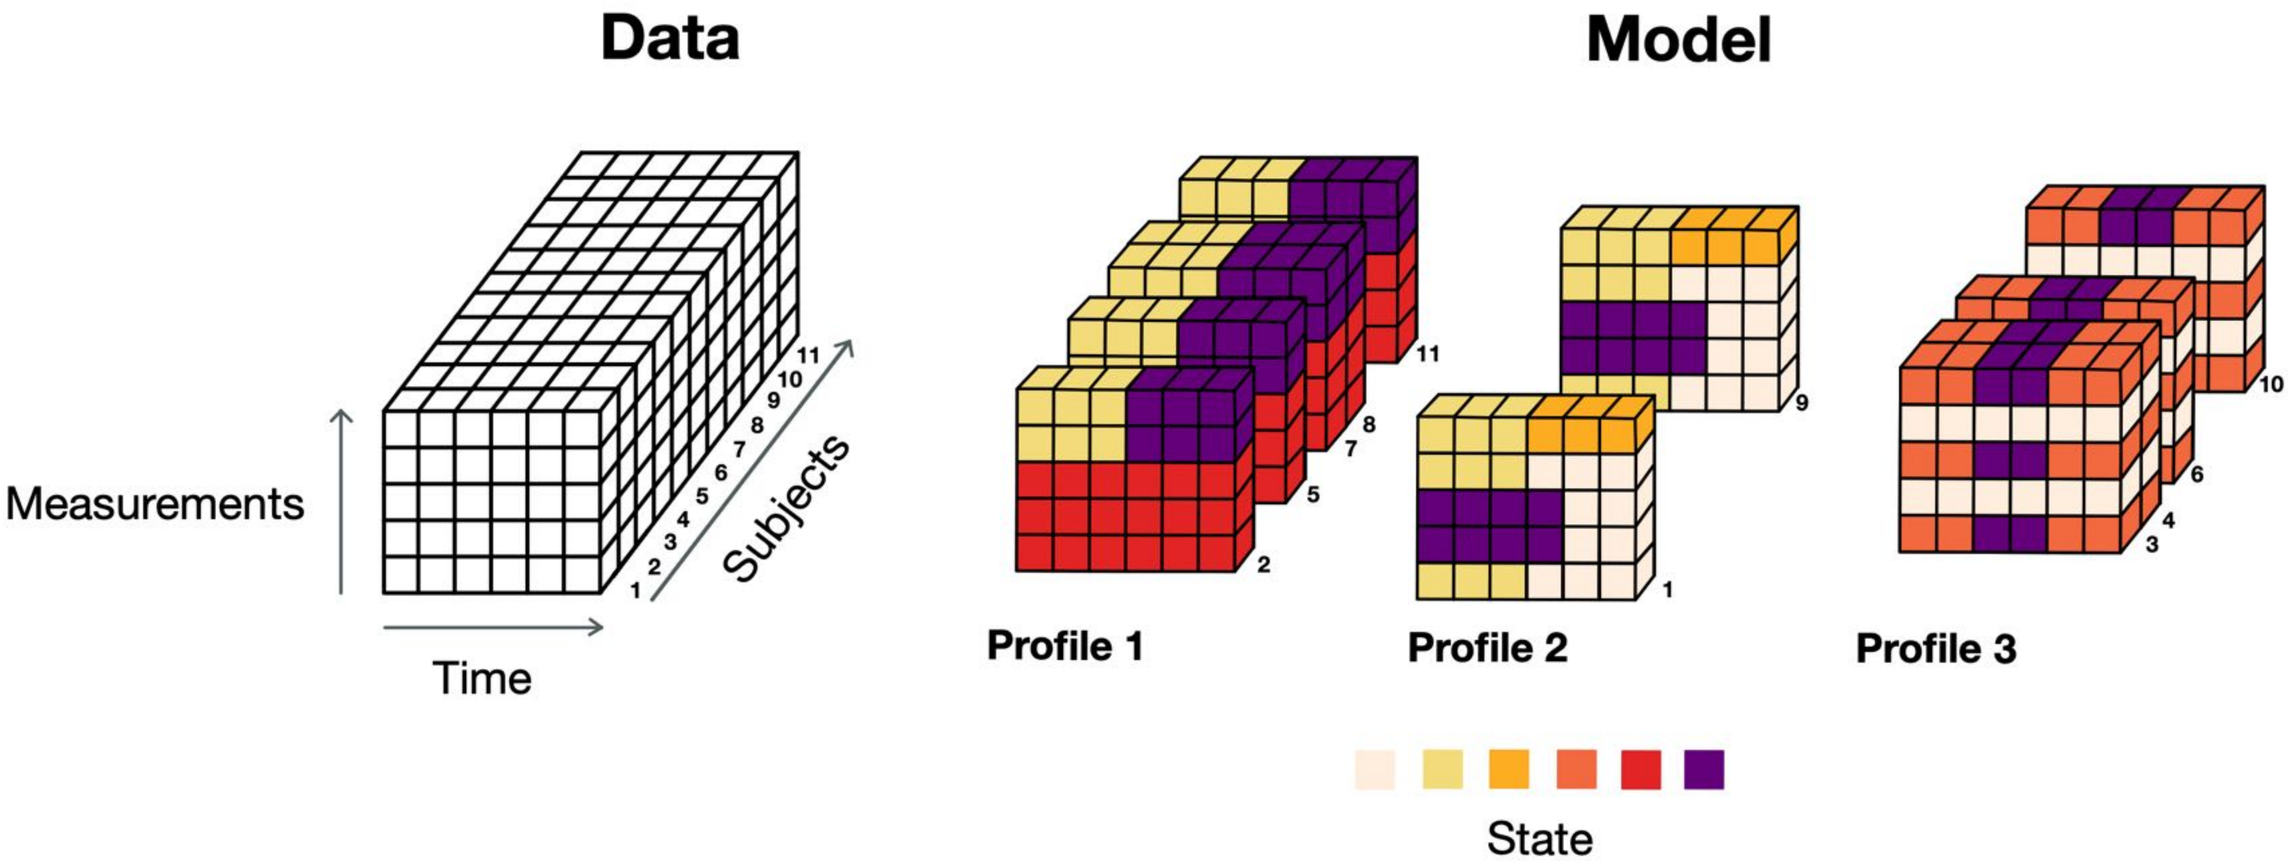

One subject model

UCLA

- 👉 **Basic model, time 0**

One subject model

UCLA

👉 **Basic model, time 0**

Data Y_r for ROI $r = 1, \ldots, R$

Statistical model

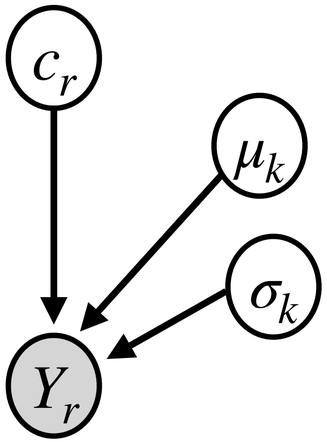

$$Y_r \mid c_r = k \sim \text{Student-t}(\mu_k, \sigma_k)$$

Inferred activation pattern

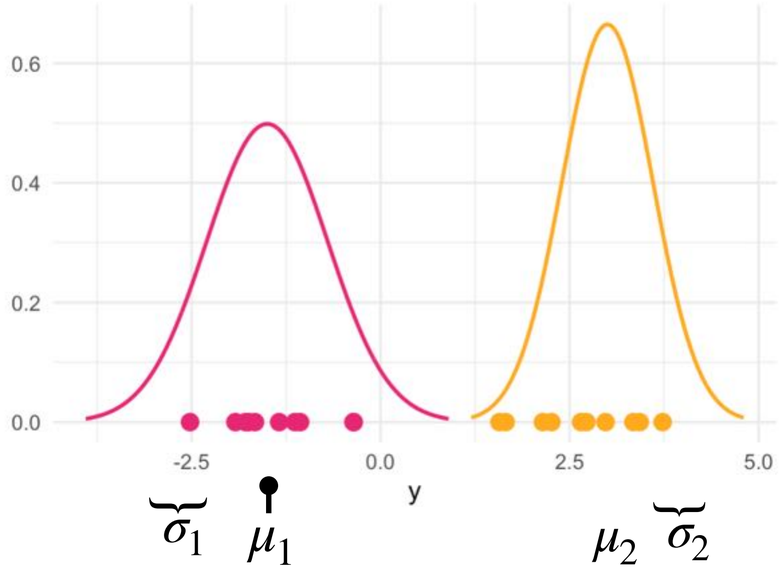

One subject model

UCLA

👉 **Dynamic Brain Region Clustering:**

$t = 1$

$i = 1$

One subject model

UCLA

👉 **Dynamic Brain Region Clustering:**

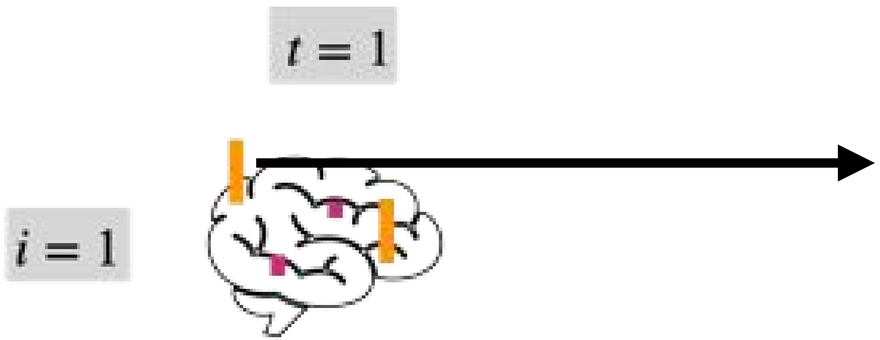

One subject model

UCLA

👉 **Dynamic Brain Region Clustering:**

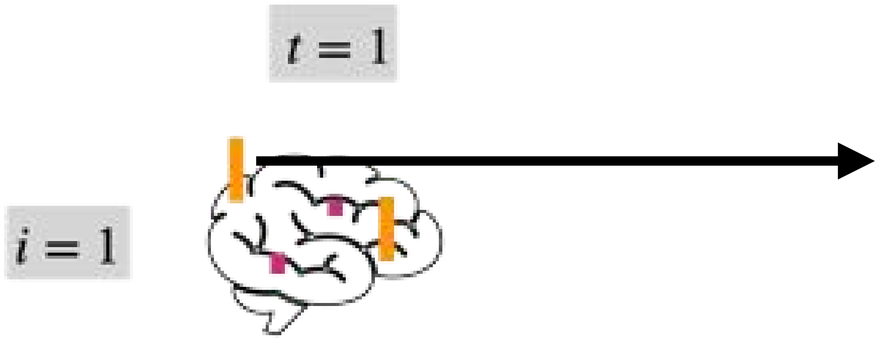

with probability α_2 : $c_{r2} = c_{r1}$

with probability $1 - \alpha_2$: $c_{r2} \sim \text{Categorical}(p_1, \ldots, p_K)$

$$Y_{r2} \mid c_{r2} = k \sim \text{Student-t}(\mu_k, \sigma_k)$$

One subject model

UCLA

👉 **Dynamic Brain Region Clustering:**

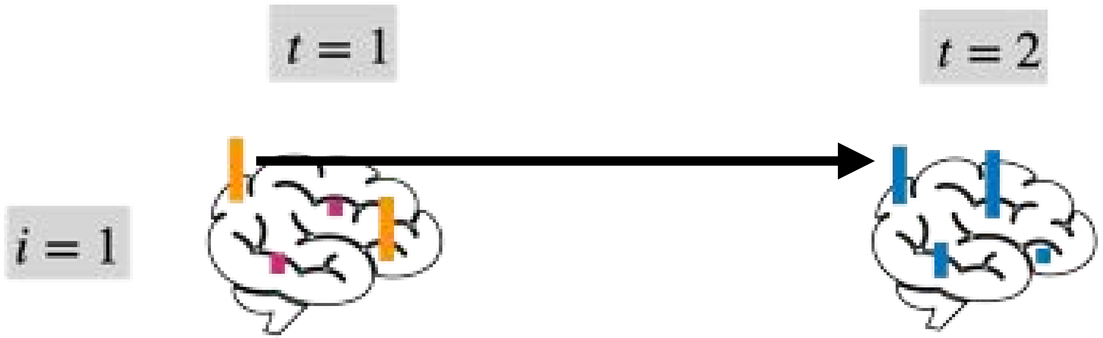

with probability α_2 : $c_{r2} = c_{r1}$

with probability $1 - \alpha_2$: $c_{r2} \sim \text{Categorical}(p_1, \ldots, p_K)$

$$Y_{r2} \mid c_{r2} = k \sim \text{Student-t}(\mu_k, \sigma_k)$$

Dynamic Temporal Allocation of the Units as in Page et al. 2022

One subject model

UCLA

👉 **Dynamic Brain Region Clustering:**

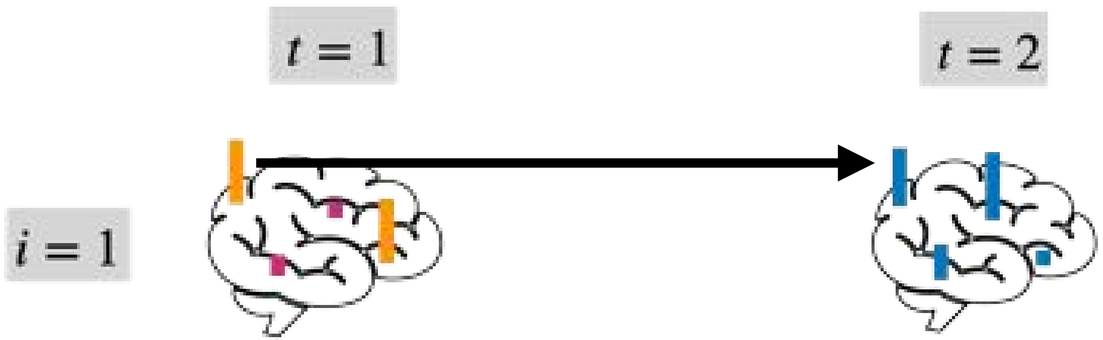

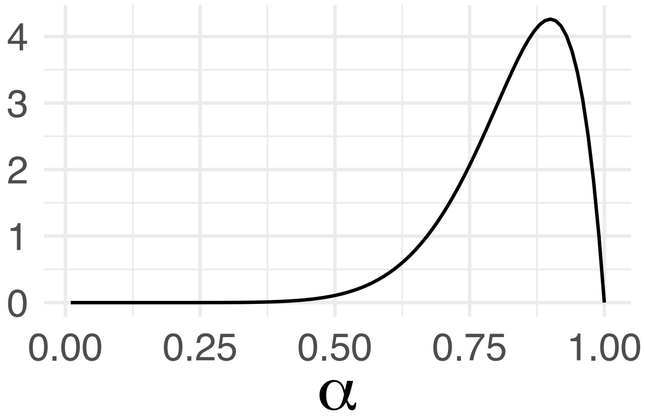

$$\alpha_2 \sim \text{Beta}(10,2)$$

with probability α_2 : $c_{r2} = c_{r1}$

with probability $1 - \alpha_2$: $c_{r2} \sim \text{Categorical}(p_1, \ldots, p_K)$

$$Y_{r2} \mid c_{r2} = k \sim \text{Student-t}(\mu_k, \sigma_k)$$

Large α_t encourages smooth dynamics!

Dynamic Temporal Allocation of the Units as in Page et al. 2022

One subject model

UCLA

👉 **Dynamic Brain Region Clustering:**

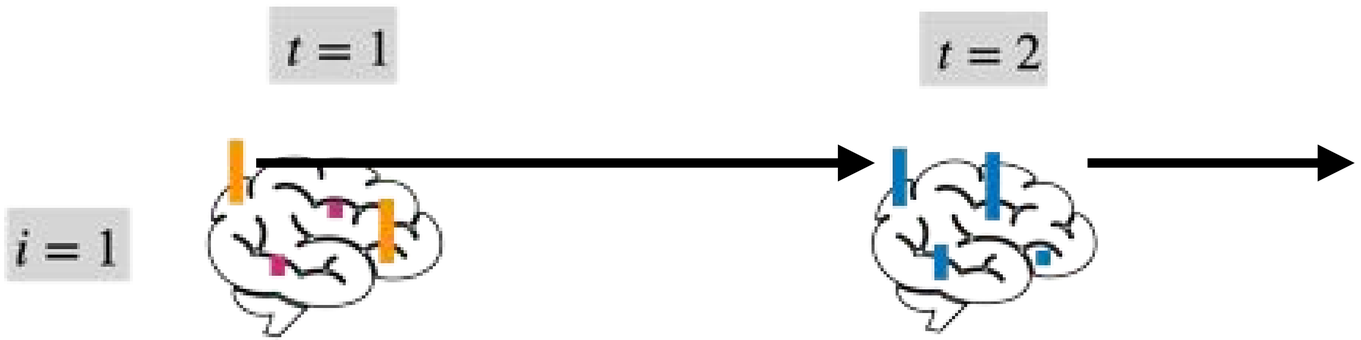

with probability α_3 : $c_{r3} = c_{r2}$

with probability $1 - \alpha_3$: $c_{r3} \sim \text{Categorical}(p_1, \ldots, p_K)$

$$Y_{r3} \mid c_{r3} = k \sim \text{Student-t}(\mu_k, \sigma_k)$$

Dynamic Temporal Allocation of the Units as in Page et al. 2022

One subject model

UCLA

👉 **Dynamic Brain Region Clustering:**

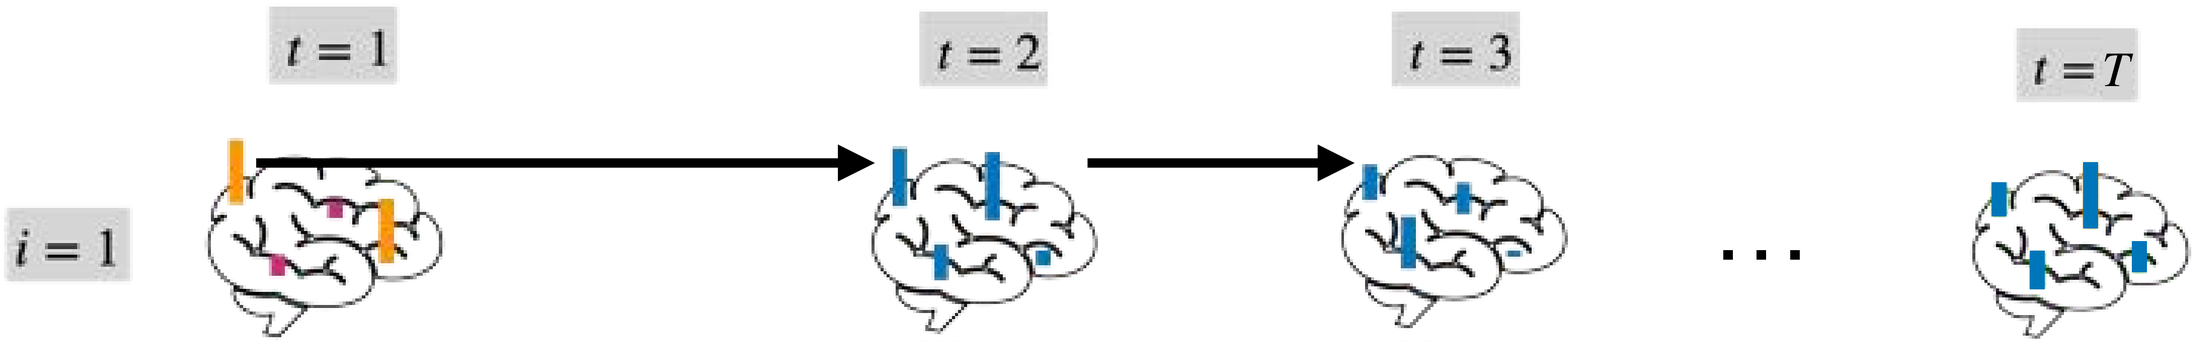

with probability α_3 : $c_{r3} = c_{r2}$

with probability $1 - \alpha_3$: $c_{r3} \sim \text{Categorical}(p_1, \ldots, p_K)$

$$Y_{r3} \mid c_{r3} = k \sim \text{Student-t}(\mu_k, \sigma_k)$$

Dynamic Temporal Allocation of the Units as in Page et al. 2022

One subject model

UCLA

👉 **Dynamic Brain Region Clustering:**

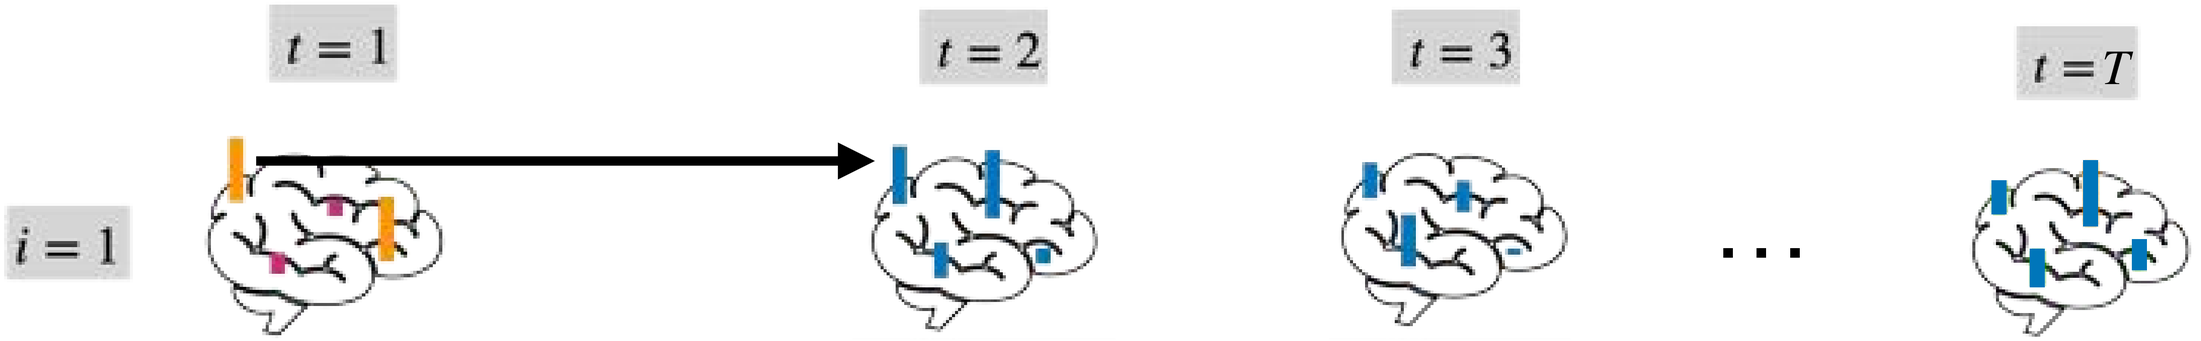

with probability α_t : $c_{rt} = c_{r,t-1}$

with probability $1 - \alpha_t$: $c_t \sim \text{Categorical}(p_1, \ldots, p_K)$

$$Y_{rt} \mid c_{rt} = k \sim \text{Student-t}(\mu_k, \sigma_k)$$

Prior on $\mathbf{p} = (p_1, \ldots, p_K)$ *such that active number of clusters can:*

- *be learned from data*
- *differ across time and subjects*

One subject model

UCLA

👉 **Dynamic Brain Region Clustering:**

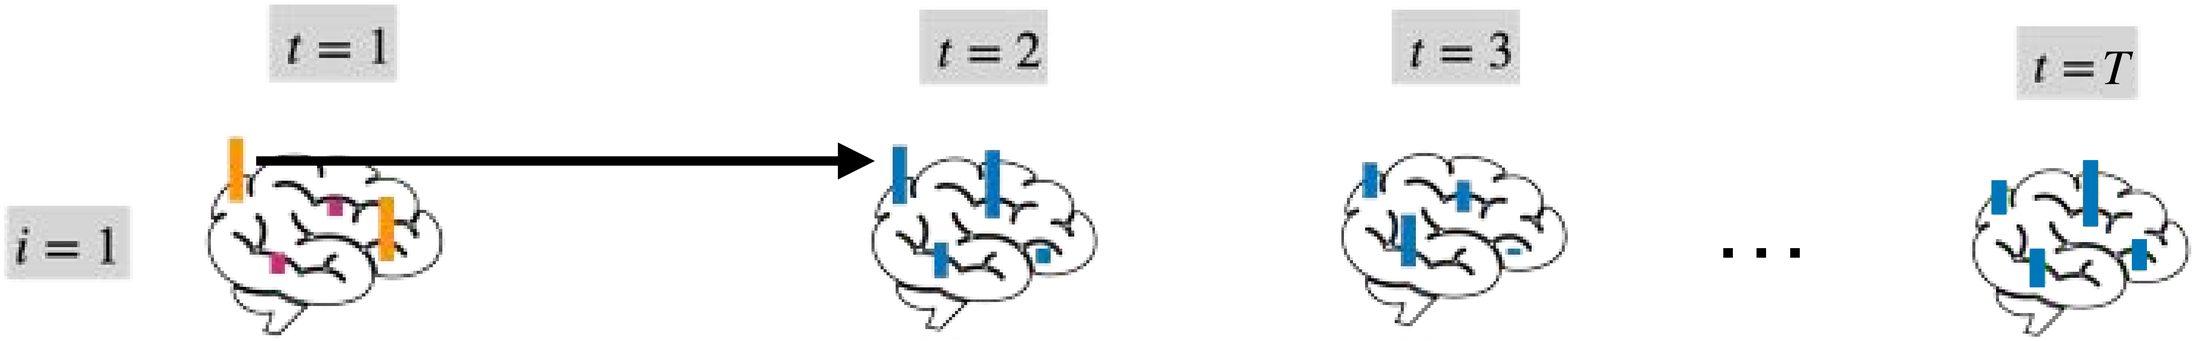

$$\boldsymbol{p} \mid \boldsymbol{\omega}_0 \sim \text{Dirichlet}(\phi\,\omega_{01}, \ldots, \phi\,\omega_{0K}),$$

$$\boldsymbol{\omega}_0 \mid \eta \sim \text{Dirichlet}\left(\frac{\eta}{K}, \ldots, \frac{\eta}{K}\right)$$

$$\eta \sim \text{Gamma}(d_1, d_2)$$

$$\boldsymbol{\omega}_0 = (\omega_{01}, \ldots, \omega_{0K})$$

Sparse Hierarchical Mixture of Mixtures Model
Malsiner-Walli et al. 2016 (Stat. Comput.)

Clustering Subjects

UCLA

From clustering brain regions over time in a single subject to **multiple subjects**:

Data + **Identify Temporal profiles (subgroups)** →

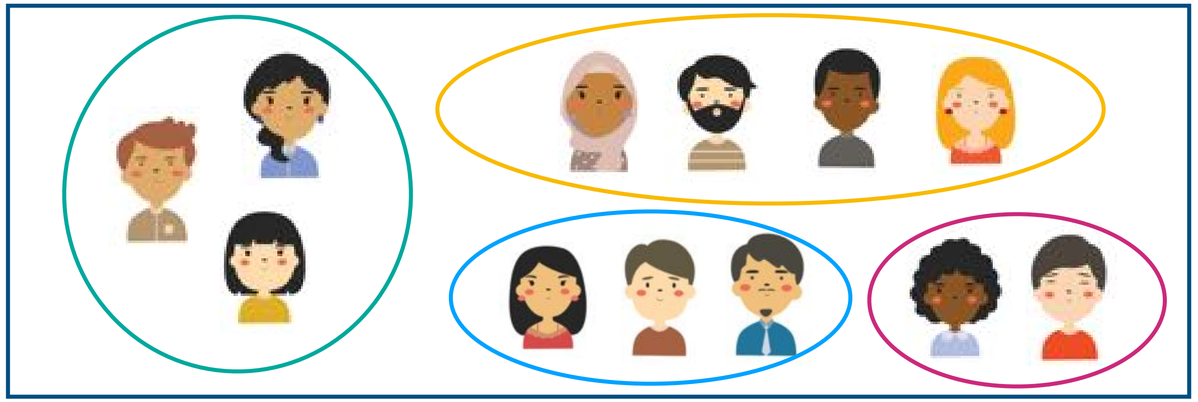

💡 **Profile**: specific sequence of brain-region clusters during the experiment

Subject clusters

UCLA

💡 **Profile**: specific sequence of brain-region clusters during the experiment

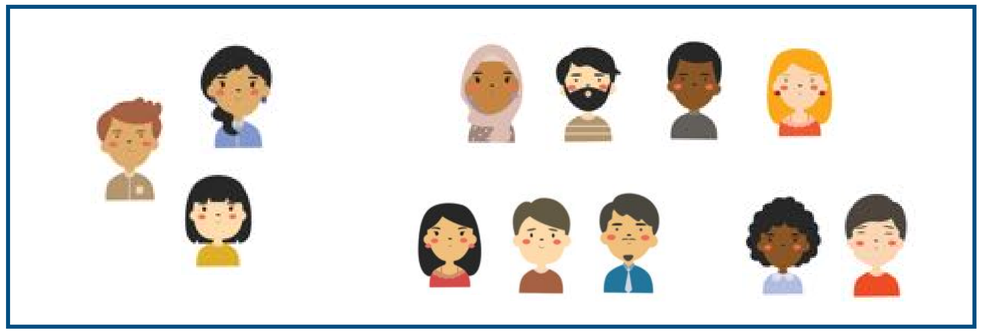

Subject clusters

UCLA

💡 **Profile**: specific sequence of brain-region clusters during the experiment

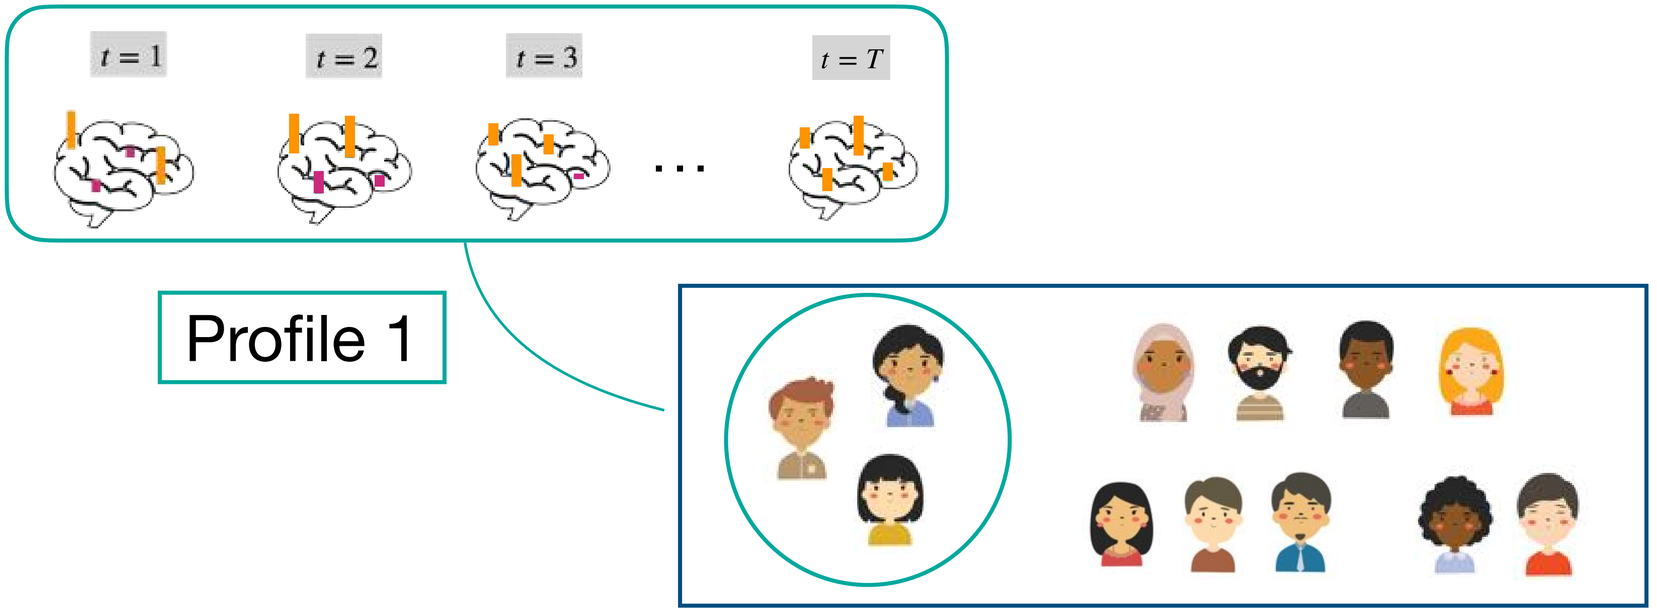

Subject clusters

UCLA

💡 **Profile**: specific sequence of brain-region clusters during the experiment

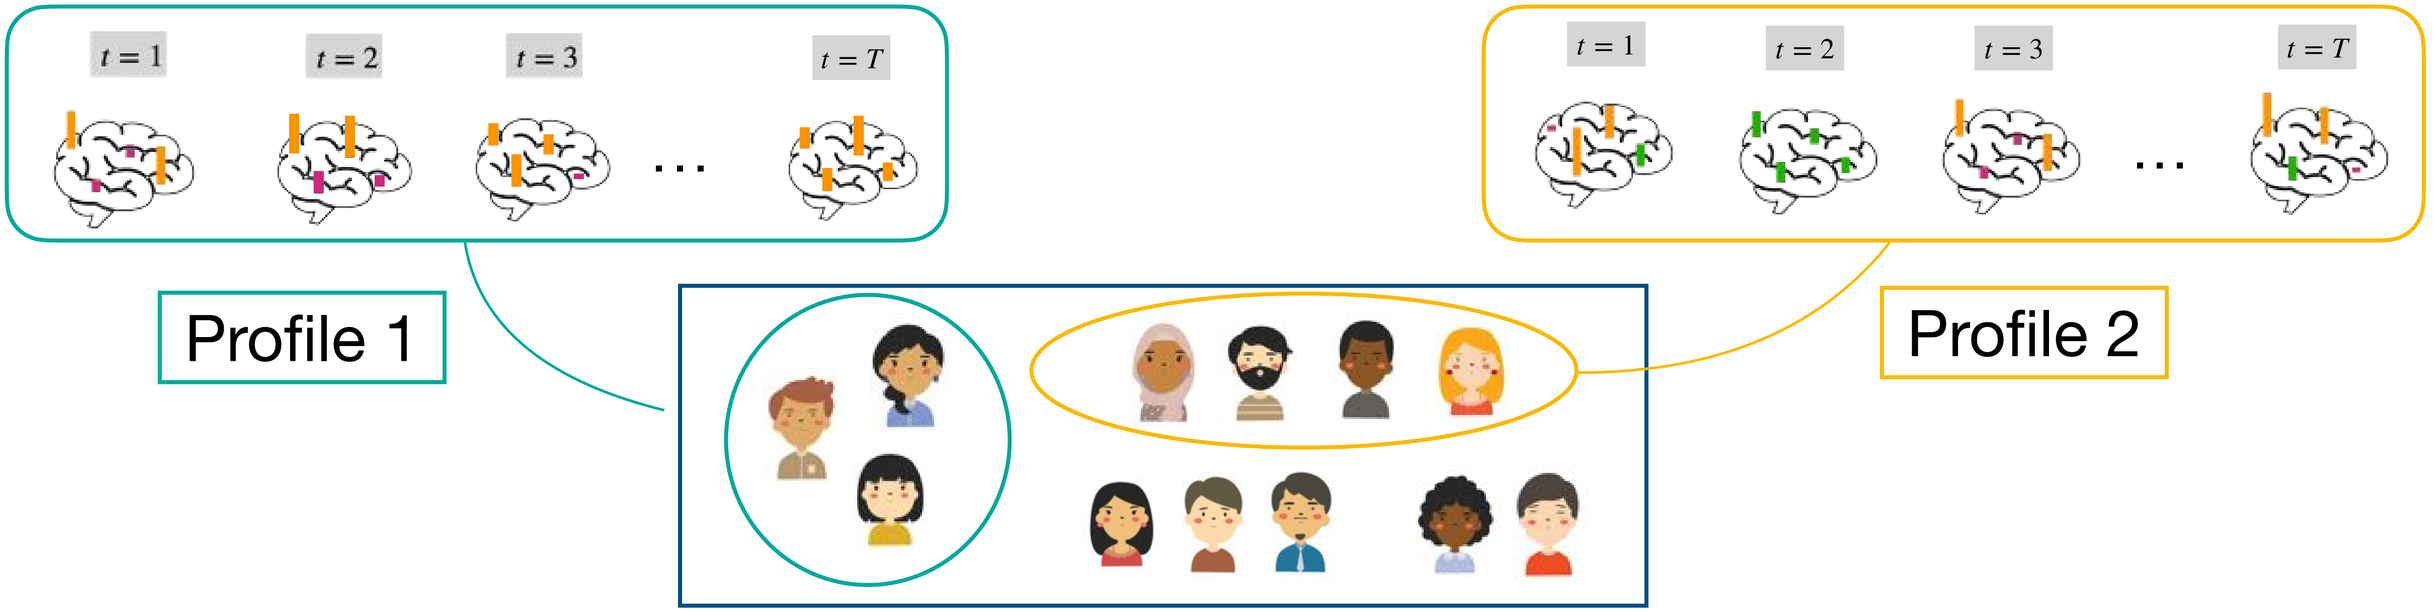

Subject clusters

UCLA

💡 **Profile:** specific sequence of brain-region clusters during the experiment

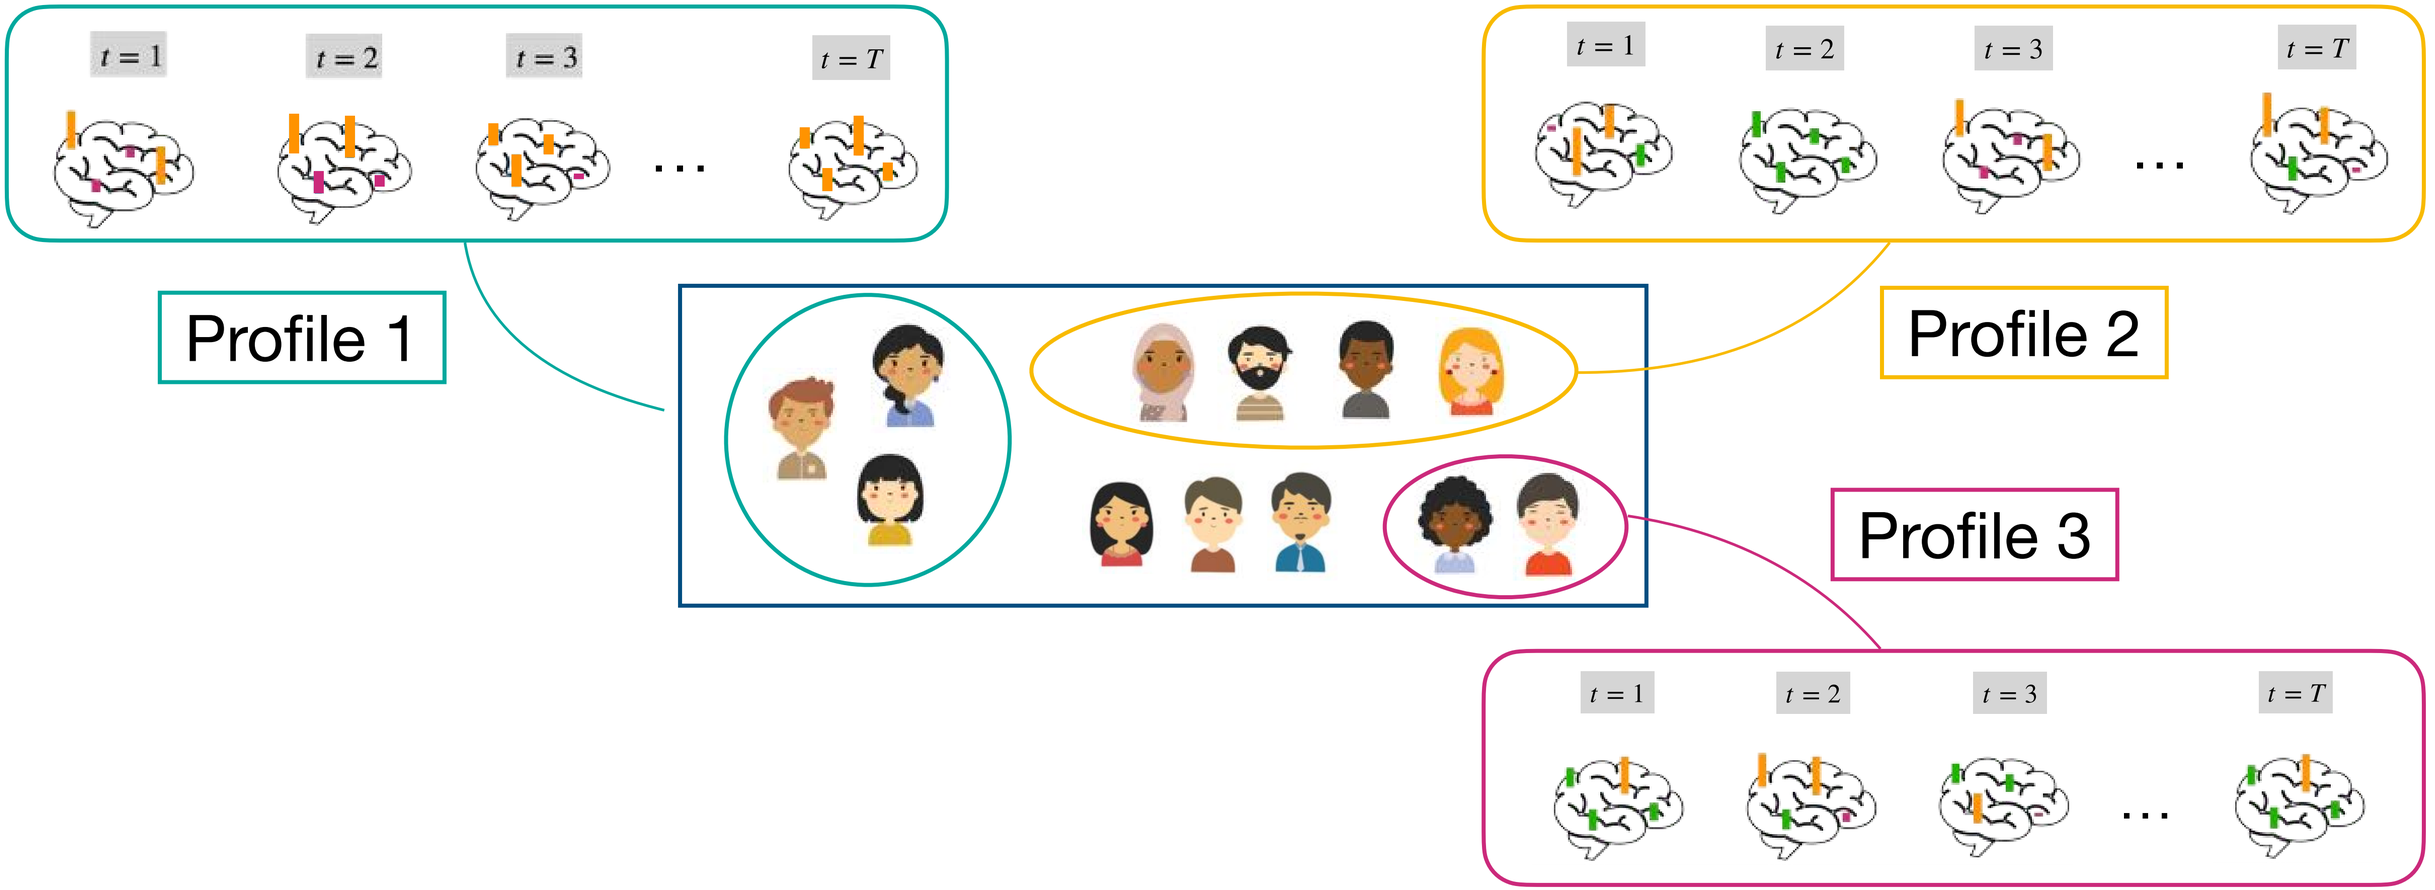

Subject clusters

UCLA

💡 **Profile**: specific sequence of brain-region clusters during the experiment

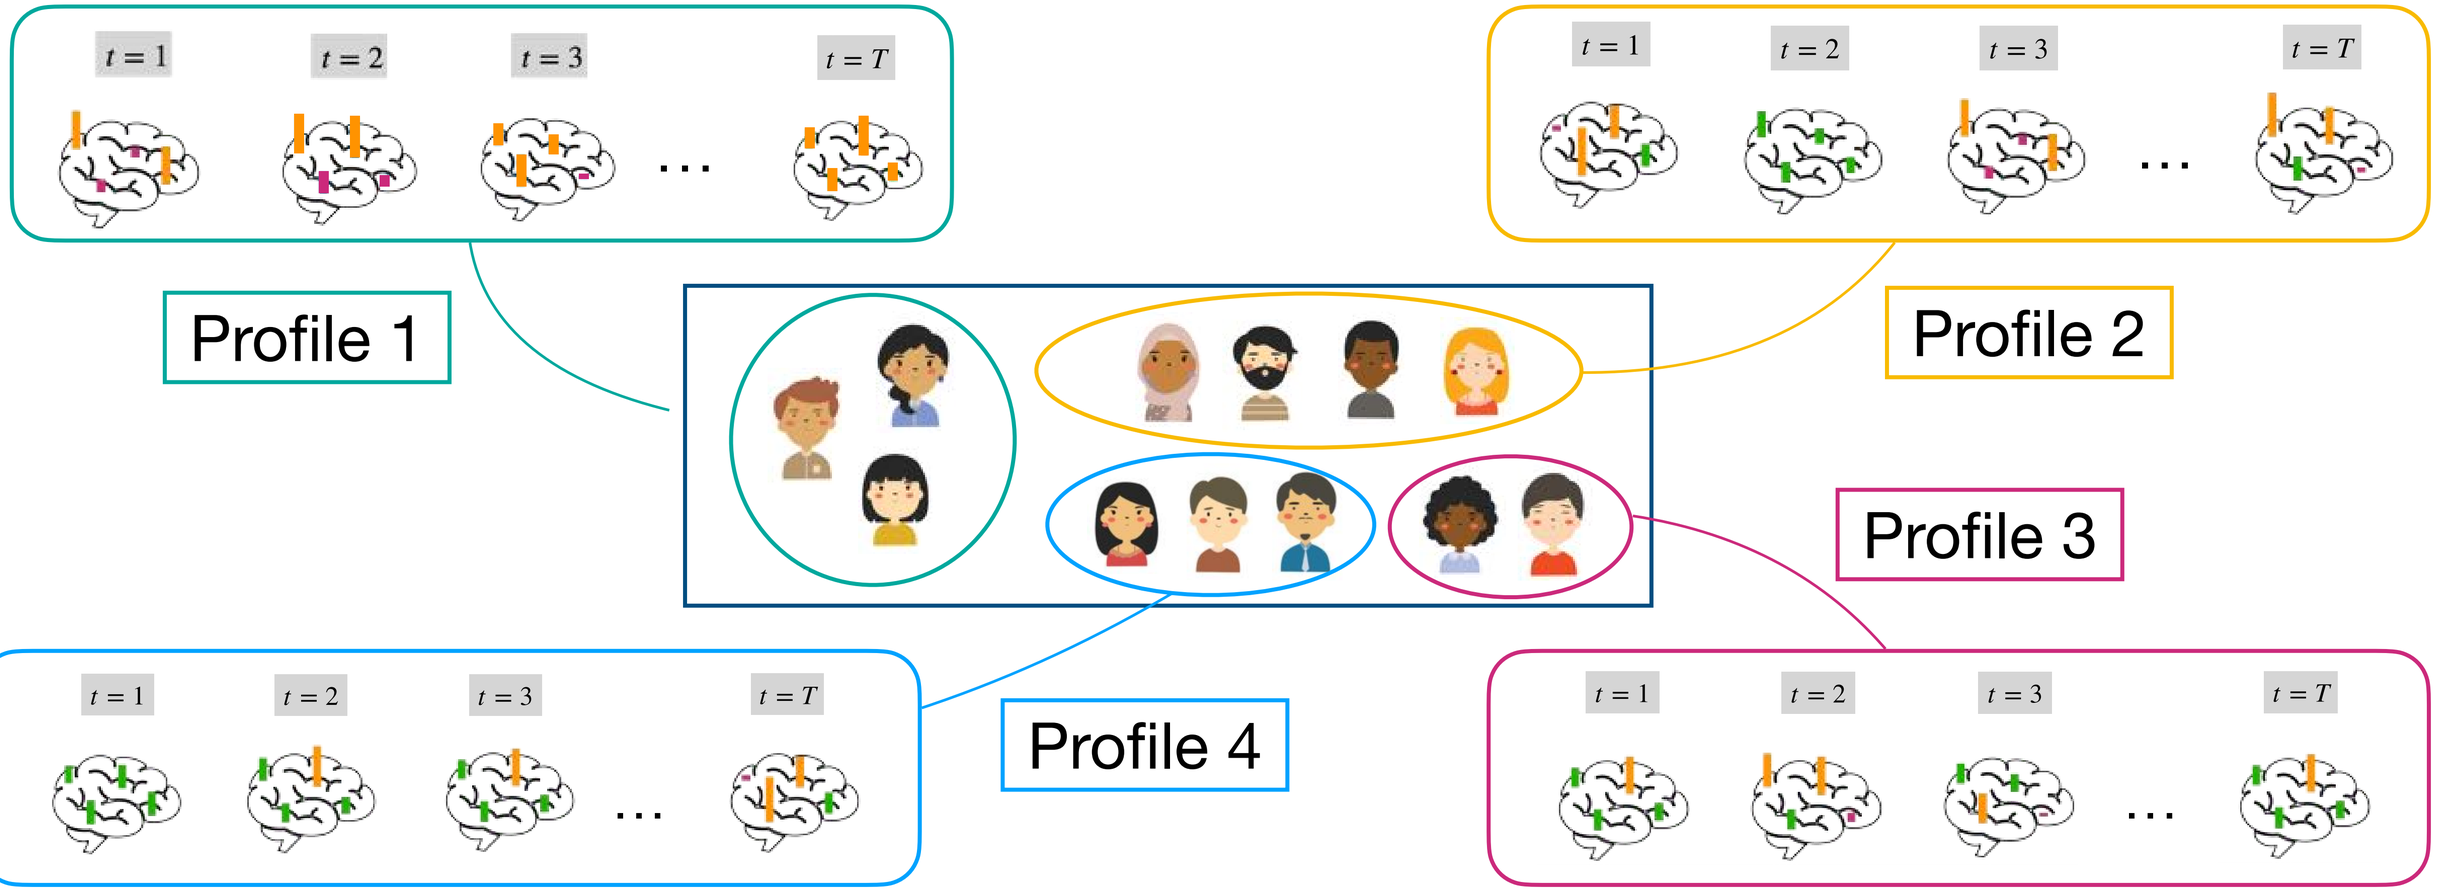

Full model

UCLA

For subject i, brain region r and time t

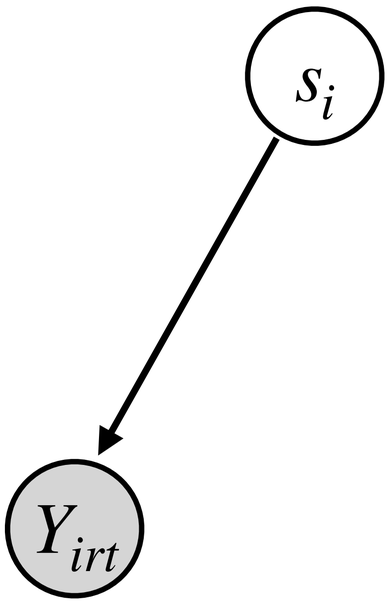

Profile of subject i

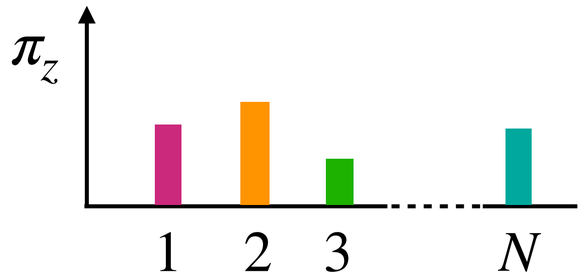

$$s_i \sim \text{Categorical}(\pi_1, \ldots, \pi_N)$$

$$\pi \sim \text{Dirichlet}\left(\frac{\varepsilon}{N}, \ldots, \frac{\varepsilon}{N}\right)$$

$$\varepsilon \sim \text{Gamma}(b_1, b_2)$$

Sparse Finite Mixture

Full model

UCLA

For subject i, brain region r and time t

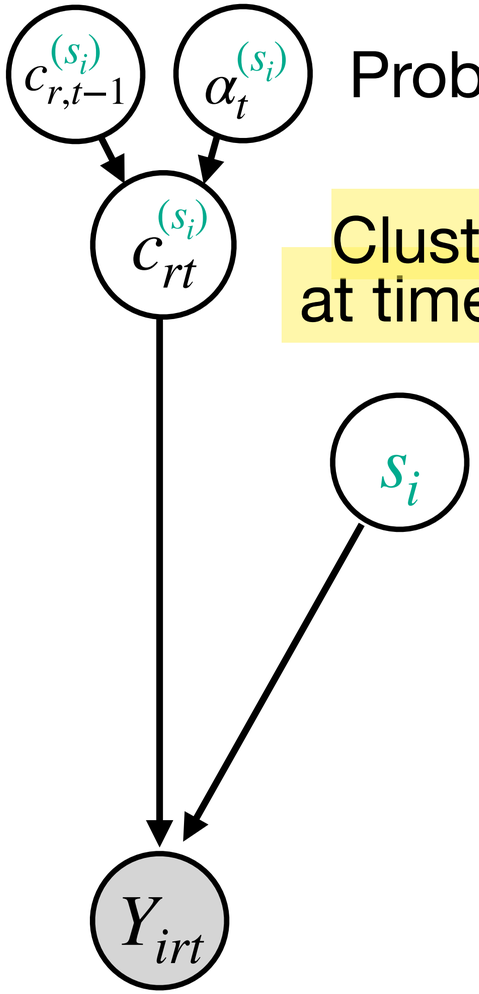

Probability that a cluster persists over time

Cluster of region r at time t for profile s_i

$$\begin{cases} \text{with probability } \alpha_t^{(s_i)} & c_{r,t}^{(s_i)} = c_{r,t-1}^{(s_i)} \\ \text{with probability } 1 - \alpha_t^{(s_i)} & c_{r,t}^{(s_i)} \sim \text{Categ}\left(p_1^{(s_i)}, \ldots, p_K^{(s_i)}\right) \end{cases}$$

Profile of subject i

$$\boldsymbol{p}^{(s_i)} \mid \boldsymbol{p}_0 \sim \text{Dirichlet}(\phi\,\omega_{01}, \ldots, \phi\,p_{0K}),$$

$$\boldsymbol{p}_0 \mid \eta \sim \text{Dirichlet}\left(\frac{\eta}{K}, \ldots, \frac{\eta}{K}\right)$$

$$\eta \sim \text{Gamma}(d_1, d_2)$$

[Finite approx. of Hierarchical Dirichlet Process]

Full model

UCLA

For subject i, brain region r and time t

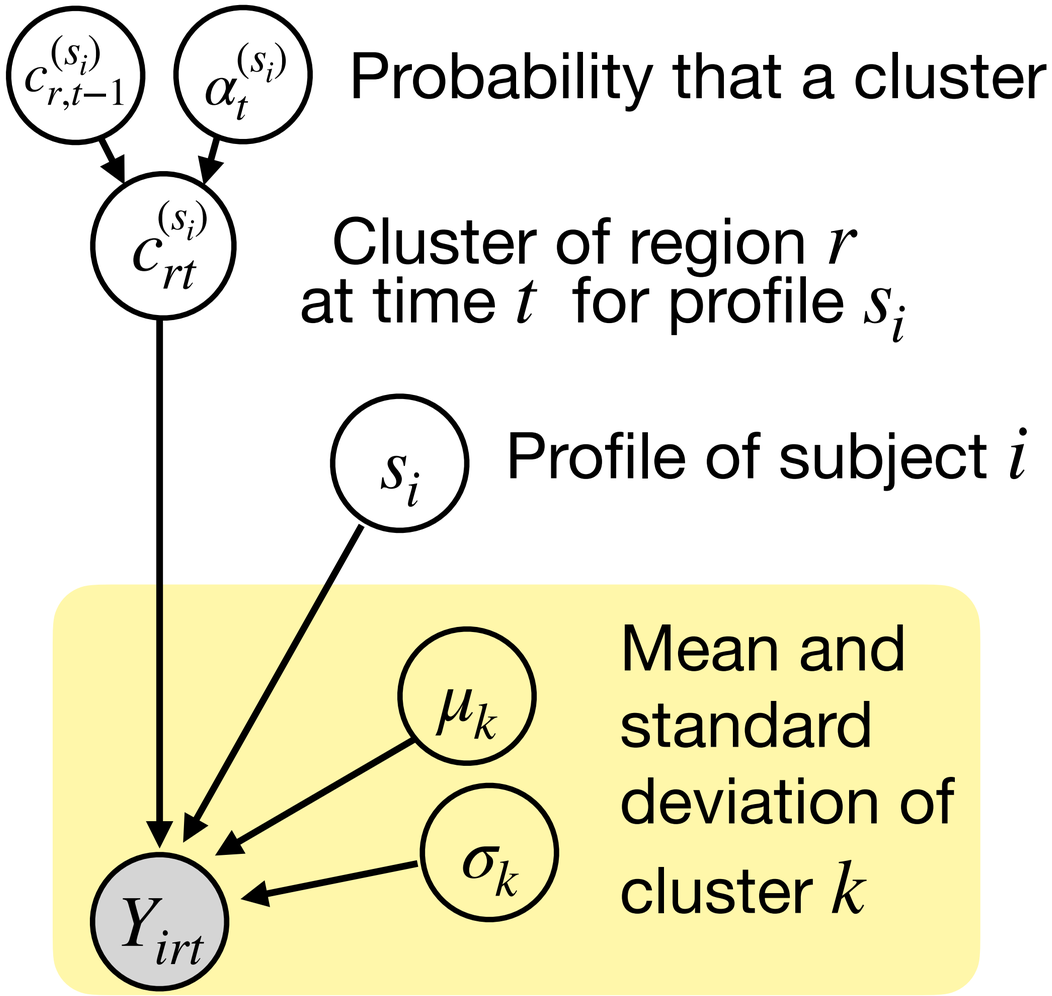

Probability that a cluster persists over time

Cluster of region r at time t for profile s_i

Profile of subject i

Mean and standard deviation of cluster k

$$Y_{irt} \mid s_i = z, c_{rt}^{(z)} = k, \mu_k, \sigma_k \sim \text{Student-t}(\mu_k, \sigma_k)$$

$$Y_{i,r,t} \mid c_{r,t}^{(s_i)} = k, \mu_k, V_{i,r,t} \overset{\text{ind}}{\sim} \text{Normal}(\mu_k, V_{i,r,t})$$

$$V_{i,r,t} \mid c_{r,t}^{(s_i)} = k, \sigma_k^2 \overset{\text{iid}}{\sim} \text{Inv-}\chi^2(\nu, \sigma_k^2)\,.$$

Posterior Inference

UCLA

- We design a MCMC for posterior inference, mostly using Gibbs updates
- Crucial step is the update of cluster-assignment sequence $(c_{r,1}^{(z)}, \ldots, c_{r,T}^{(z)})$ for each profile z and region r
- For the case with no profiles, Page et al. (2022) propose a marginal sampler
  - Let $\gamma_{r,t} = 1$ with probability α_t (so $\gamma_{r,t}$ is an indicator of cluster persistence)
  - Marginal updates are conditional on past, present and future persistence indicators and cluster assignments

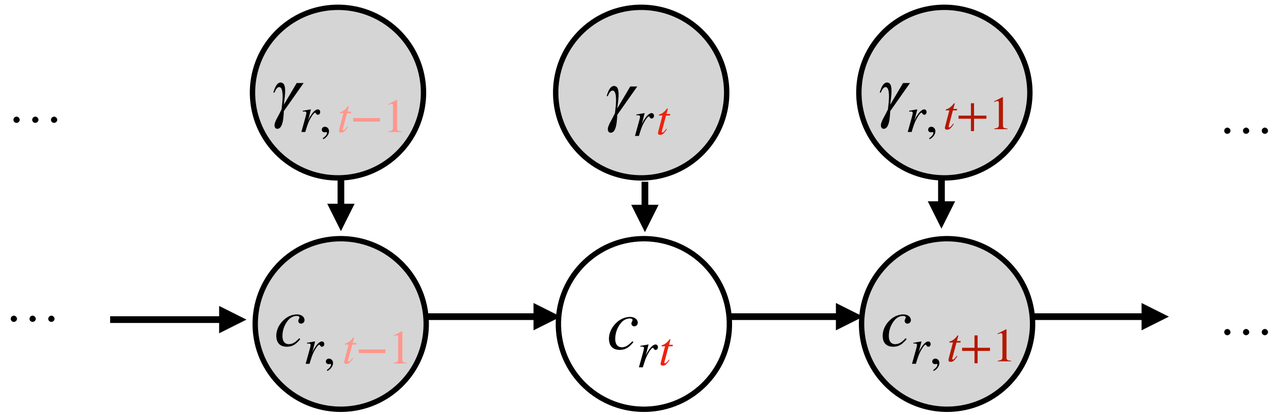

Block update of region clusters

UCLA

- We design update of cluster-assignment sequences *in block:*
  - Update persistence indicators and cluster assignments together and sequentially, only conditioned on the past

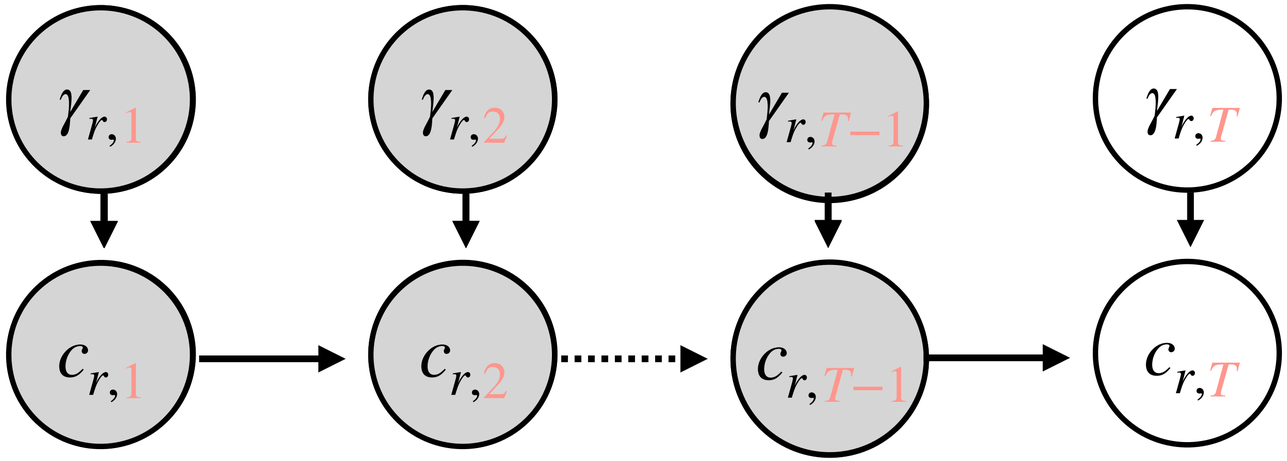

Blocked vs Marginal Sampler

UCLA

BL = Binder loss

MAE = Mean Absolute Error of state parameters

f-measure (F1, combining precision and recall) for changepoint detection accuracy

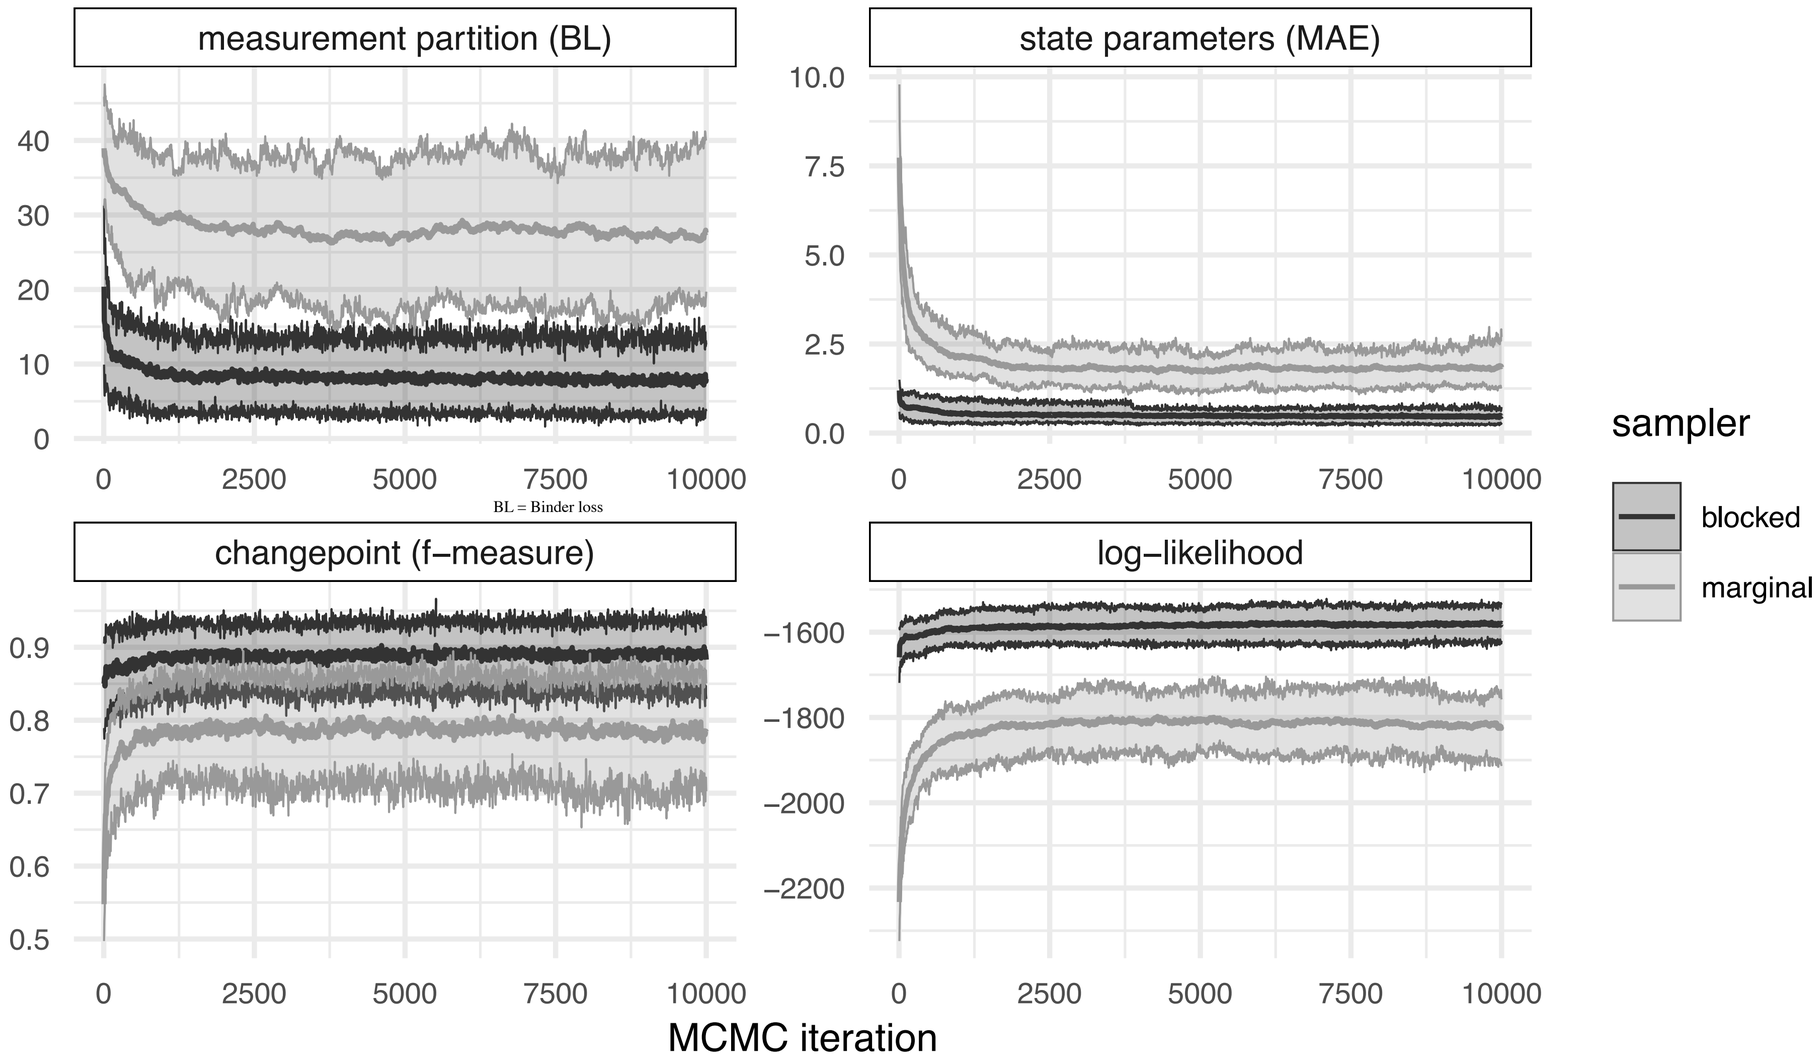

Application to fMRI study

UCLA

- **N = 23** Healthy Subjects **R = 11** ROIs known to be involved in

Default Mode Network (DMN): Posterior Cingulate Cortex (PCC), left/right Dorsolateral Prefrontal Cortex (L/R DLPFC), left/right posterior Inferior Parietal Lobule (L/R pIPL), left Parahippocampal Gyrus (L PHG), and left Inferolateral Temporal Cortex (L ITC)

Salience Network (SN): Left/right Insula (L/R Insula), left/right Lateral Parietal regions (L/R LP)

Image credit: Hussain, Sana, et al. (2023)

Results

UCLA

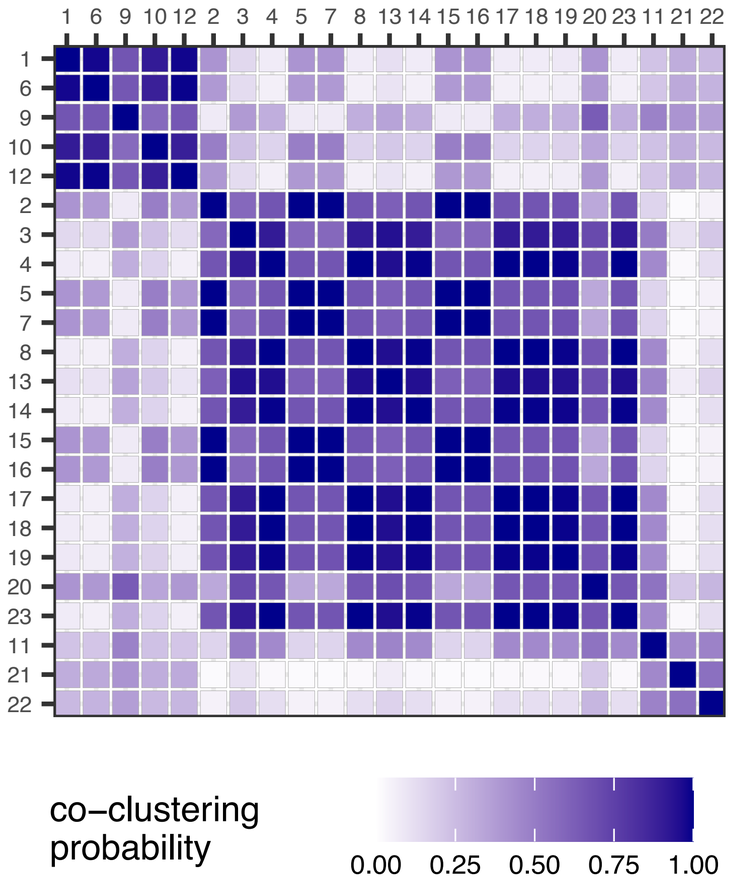

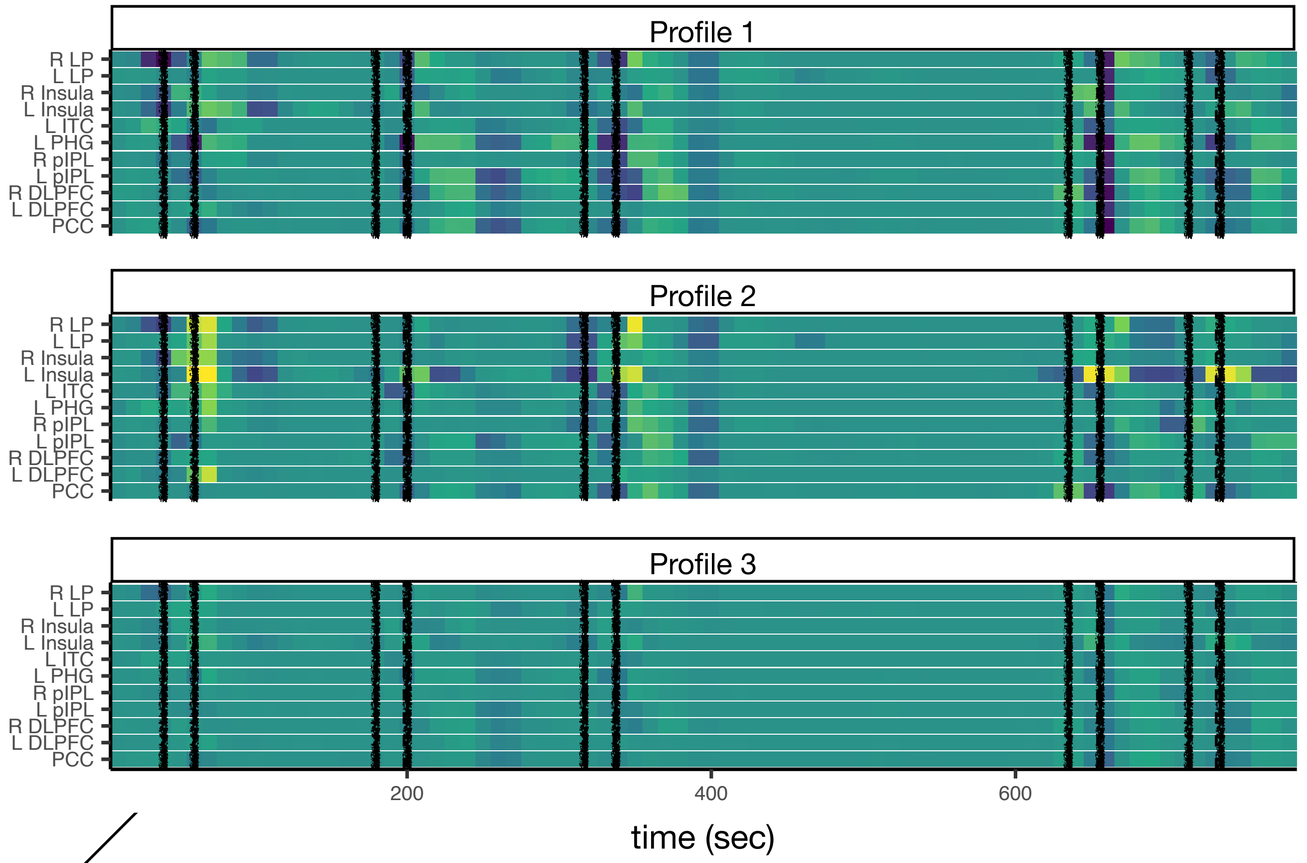

squeeze block

Allocation of Subjects to Profiles

Corresponding Temporal Evolution

Modeling of dynamic partitions

UCLA

Modeling of dynamic partitions

UCLA

➡️ Dynamic Partition Linear Model

‼️ Instead of looking at unit-specific allocations, look at the change in the partitions as a whole

Modeling of dynamic partitions

UCLA

➡️ Dynamic Partition Linear Model

‼️ Instead of looking at unit-specific allocations, look at the change in the partitions as a whole

Modeling of dynamic partitions

UCLA

➡️ Dynamic Partition Linear Model

‼️ Instead of looking at unit-specific allocations, look at the change in the partitions as a whole

Modeling of dynamic partitions

UCLA

➡️ Dynamic Partition Linear Model

‼️ Instead of looking at unit-specific allocations, look at the change in the partitions as a whole

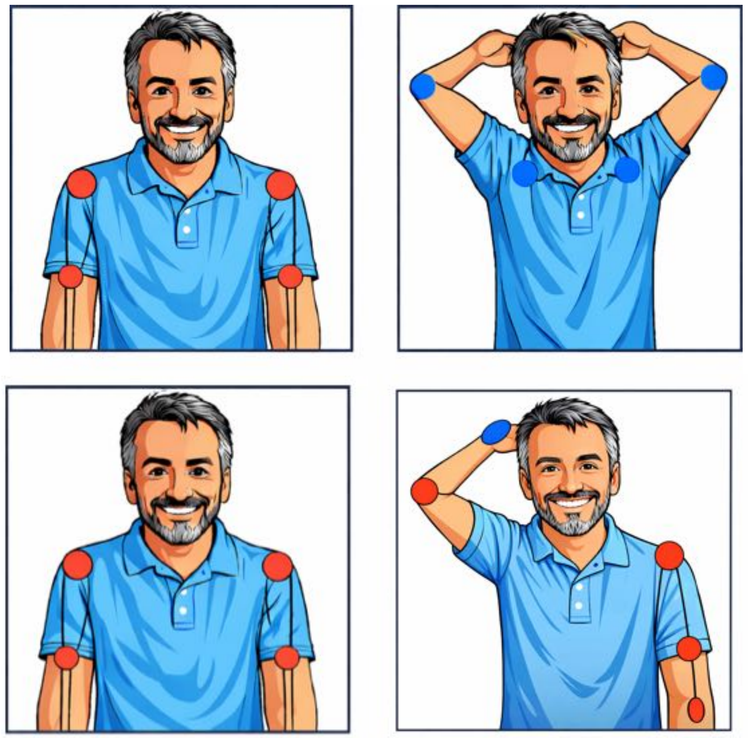

Dynamic Linear Models - A primer

UCLA

Dynamic linear models (DLMs) are commonly used for time-series data due to their flexibility and adaptability.

They define a class of **state-space** models

As such they are characterized by a system of two equations:

👉 an **observation equation**, which describes the observed data as a linear combination of latent state variables with noise,

👉 **state equation** that describes how latent states evolve over time, thereby tracking the underlying dynamics of the system.

A simple DLM: A Local Level model (LLM)

UCLA

Let $\boldsymbol{Y}_t = \{Y_{1,t}, \ldots, Y_{n,t}\}$ be an *n*-dimensional vector observed at $t = 1, \ldots, T$ time points

In an LLM, the observed data are composed of a level component plus a random noise:

$$Y_{i,t} = \beta_{i,t} + \varepsilon_i$$

with $\varepsilon_i \overset{iid}{\sim} N\left(0, \tau^2\right)$

👉 $\boldsymbol{\beta}_t = \{\beta_{1,t}, \ldots, \beta_{n,t}\}$ the underlying **level or trend** of the time series:

$$\beta_{i,t} = \beta_{i,t-1} + \omega_t$$

where $\omega_t \overset{iid}{\sim} N(0,1)$

Local Level Dynamic Random Partition Model (LLDPM)

UCLA

Local Level Dynamic Random Partition Model (LLDPM)

UCLA

For each unit (e.g., each arm) $i \in \{1, \ldots, n\}$, we have:

👉 a model for the **observation equation** (e.g., the Gaussian kernel):

$$Y_{i,t} \mid \beta_{i,t} \overset{ind}{\sim} p\left(y_{i,t} \mid \beta_{i,t}\right)$$

👉 The **state equation** governs the dynamics of the time-varying **partitions**

👉 Similarly to the LLM we want to tie the partition at time *t-1,* possibly with the partition at time *t*

LLM ➡ Random Walk
LLDPM ➡ partitions are discrete ➡ NDARMA-like formulation (Jacobs & Lewis, 1983)

Local Level Dynamic Random Partition Model (LLDPM)

UCLA

NDARMA describes a stationary discrete-valued analogue of ARMA built by randomly choosing whether the next value is copied from a past observation or taken from an innovation.

Local Level Dynamic Random Partition Model (LLDPM)

UCLA

NDARMA describes a stationary discrete-valued analogue of ARMA built by randomly choosing whether the next value is copied from a past observation or taken from an innovation.

Let π_t indicate a **partition** at time t

👉 Let $\left|\pi_t\right|$ indicate the number of clusters/blocks identified in the partition at time t among the n units

Local Level Dynamic Random Partition Model (LLDPM)

UCLA

NDARMA describes a stationary discrete-valued analogue of ARMA built by randomly choosing whether the next value is copied from a past observation or taken from an innovation.

Let π_t indicate a **partition** at time *t*

👉 Let $\left|\pi_t\right|$ indicate the number of clusters/blocks identified in the partition at time *t* among the *n* units

👉 Let γ_t indicate a (latent) **changepoint** (cgp) indicator (0=no cgp; 1=cgp)

➡️ detects changes in the partitions of units from time *t - 1* to time *t*

Local Level Dynamic Random Partition Model (LLDPM)

UCLA

NDARMA describes a stationary discrete-valued analogue of ARMA built by randomly choosing whether the next value is copied from a past observation or taken from an innovation.

Let π_t indicate a **partition** at time *t*

- Let $\left|\pi_t\right|$ indicate the number of clusters/blocks identified in the partition at time *t* among the *n* units

- Let γ_t indicate a (latent) **changepoint** (cgp) indicator (0=no cgp; 1=cgp)
  - detects changes in the partitions of units from time *t - 1* to time *t*

- $p^*\left(\pi_t\right)$ indicates the distribution of a **base** random partition model
  - a probability distribution that describes the probability of different clusters allocations at each time *t* ➡ e.g. CRP

Local Level Dynamic Random Partition Model (LLDPM)

UCLA

The partition-based state equation is characterized as a **mixture** over two partition models:

$$\pi_t \mid \pi_{1:(t-1)}, \gamma_{2:(t-1)} \sim \left(1-\gamma_t\right)\, \delta_{\pi_{t-1}}\left(\pi_t\right) + \gamma_t\, p^*\left(\pi_t\right)$$

Local Level Dynamic Random Partition Model (LLDPM)

UCLA

The partition-based state equation is characterized as a **mixture** over two partition models:

$$\pi_t \mid \pi_{1:(t-1)}, \gamma_{2:(t-1)} \sim \left(1-\gamma_t\right)\, \delta_{\pi_{t-1}}\left(\pi_t\right) + \gamma_t\, p^*\left(\pi_t\right)$$

Fully dependent case

The partition at time *t*
coincides
with the partition at
time *t-1*

Local Level Dynamic Random Partition Model (LLDPM)

UCLA

The partition-based state equation is characterized as a **mixture** over two partition models:

$$\pi_t \mid \pi_{1:(t-1)}, \gamma_{2:(t-1)} \sim (1-\gamma_t)\ \delta_{\pi_{t-1}}(\pi_t) + \gamma_t\, p^*(\pi_t)$$

Fully dependent case

The partition at time *t* coincides with the partition at time *t-1*

Innovation

The partition at time *t* is extracted from the **base** process **independently** of the partition at time *t-1*

Local Level Dynamic Random Partition Model (LLDPM)

UCLA

The partition-based state equation is characterized as a **mixture** over two partition models:

$$\pi_t \mid \pi_{1:(t-1)}, \gamma_{2:(t-1)} \sim \left(1-\gamma_t\right)\, \delta_{\pi_{t-1}}\left(\pi_t\right) + \gamma_t\, p^*\left(\pi_t\right)$$

Fully dependent case

The partition at time *t* coincides with the partition at time *t-1*

Innovation

The partition at time *t* is extracted from the **base** process **independently** of the partition at time *t-1*

$$\gamma_t \overset{iid}{\sim} \text{Bern}\left(\eta_t\right)$$

$$\eta_t \overset{iid}{\sim} \text{Beta}(a, b)$$

Local Level Dynamic Random Partition Model (LLDPM)

UCLA

The partition-based state equation is characterized as a **mixture** over two partition models:

$$\pi_t \mid \pi_{1:(t-1)}, \gamma_{2:(t-1)} \sim (1-\gamma_t)\ \delta_{\pi_{t-1}}(\pi_t) + \gamma_t\, p^*(\pi_t)$$

Fully dependent case

The partition at time *t* coincides with the partition at time *t-1*

Innovation

The partition at time *t* is extracted from the **base** process **independently** of the partition at time *t-1*

$$\gamma_t \overset{iid}{\sim} \text{Bern}(\eta_t)$$

$$\eta_t \overset{iid}{\sim} \text{Beta}(a, b)$$

One can be more **creative** in the definition state equation (NDARMA) and prior for γ_t

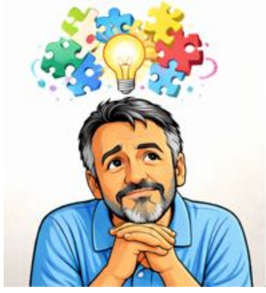

Local Level Dynamic Random Partition Model (LLDPM)

UCLA

Average lagged ARI for the pairwise comparison of *T=10* random partitions $\pi_{1:T}$ assuming a base process $p^*(\cdot) = p_{\text{CRP}}(\cdot)$.

For each matrix, the pixel in position (i, j) refers to the comparison of π_i and π_j.

For each value of η, values of the lagged ARI are averaged over a sample of 10,000 partitions.

The temporal dependence increases as the temporal dependence parameter η decreases.

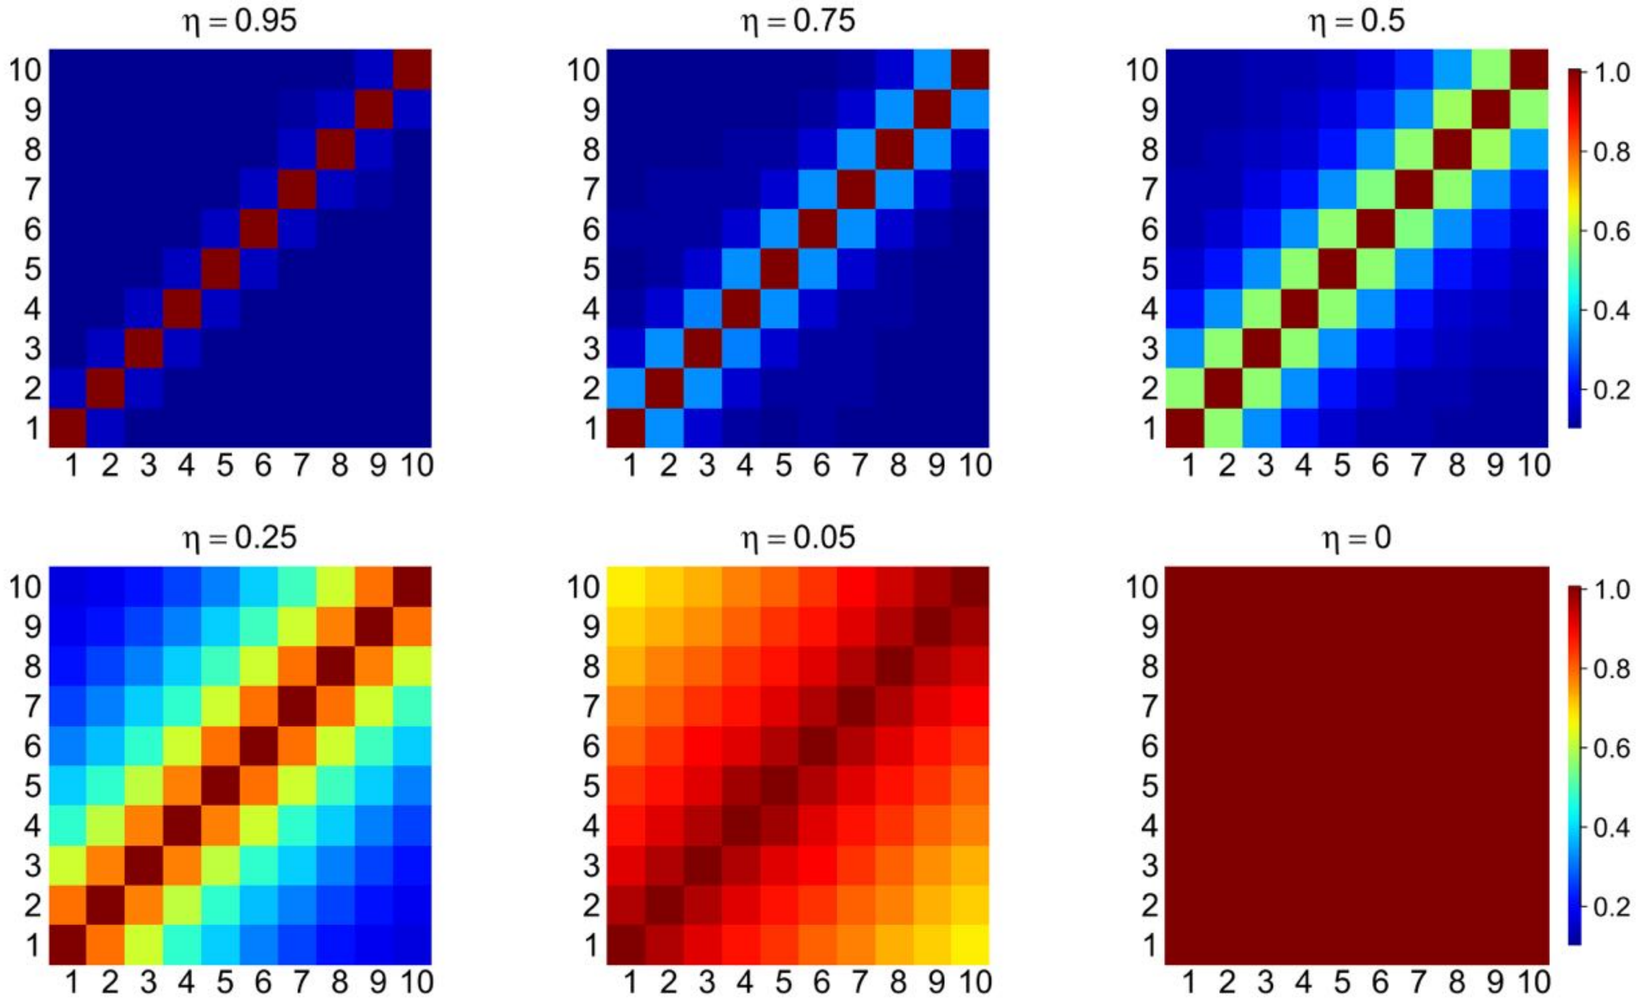

Gesture Phase Segmentation data

UCLA

Sensor data recordings of users recounting comic book stories (data from UCI ML repository)

The dataset provides scalar velocity and acceleration values over four sensors, placed on the left hand, right hand, left wrist, and right wrist

➡ $n = 8$ sensor measurements at regular time intervals

Gesture Phase Segmentation data

UCLA

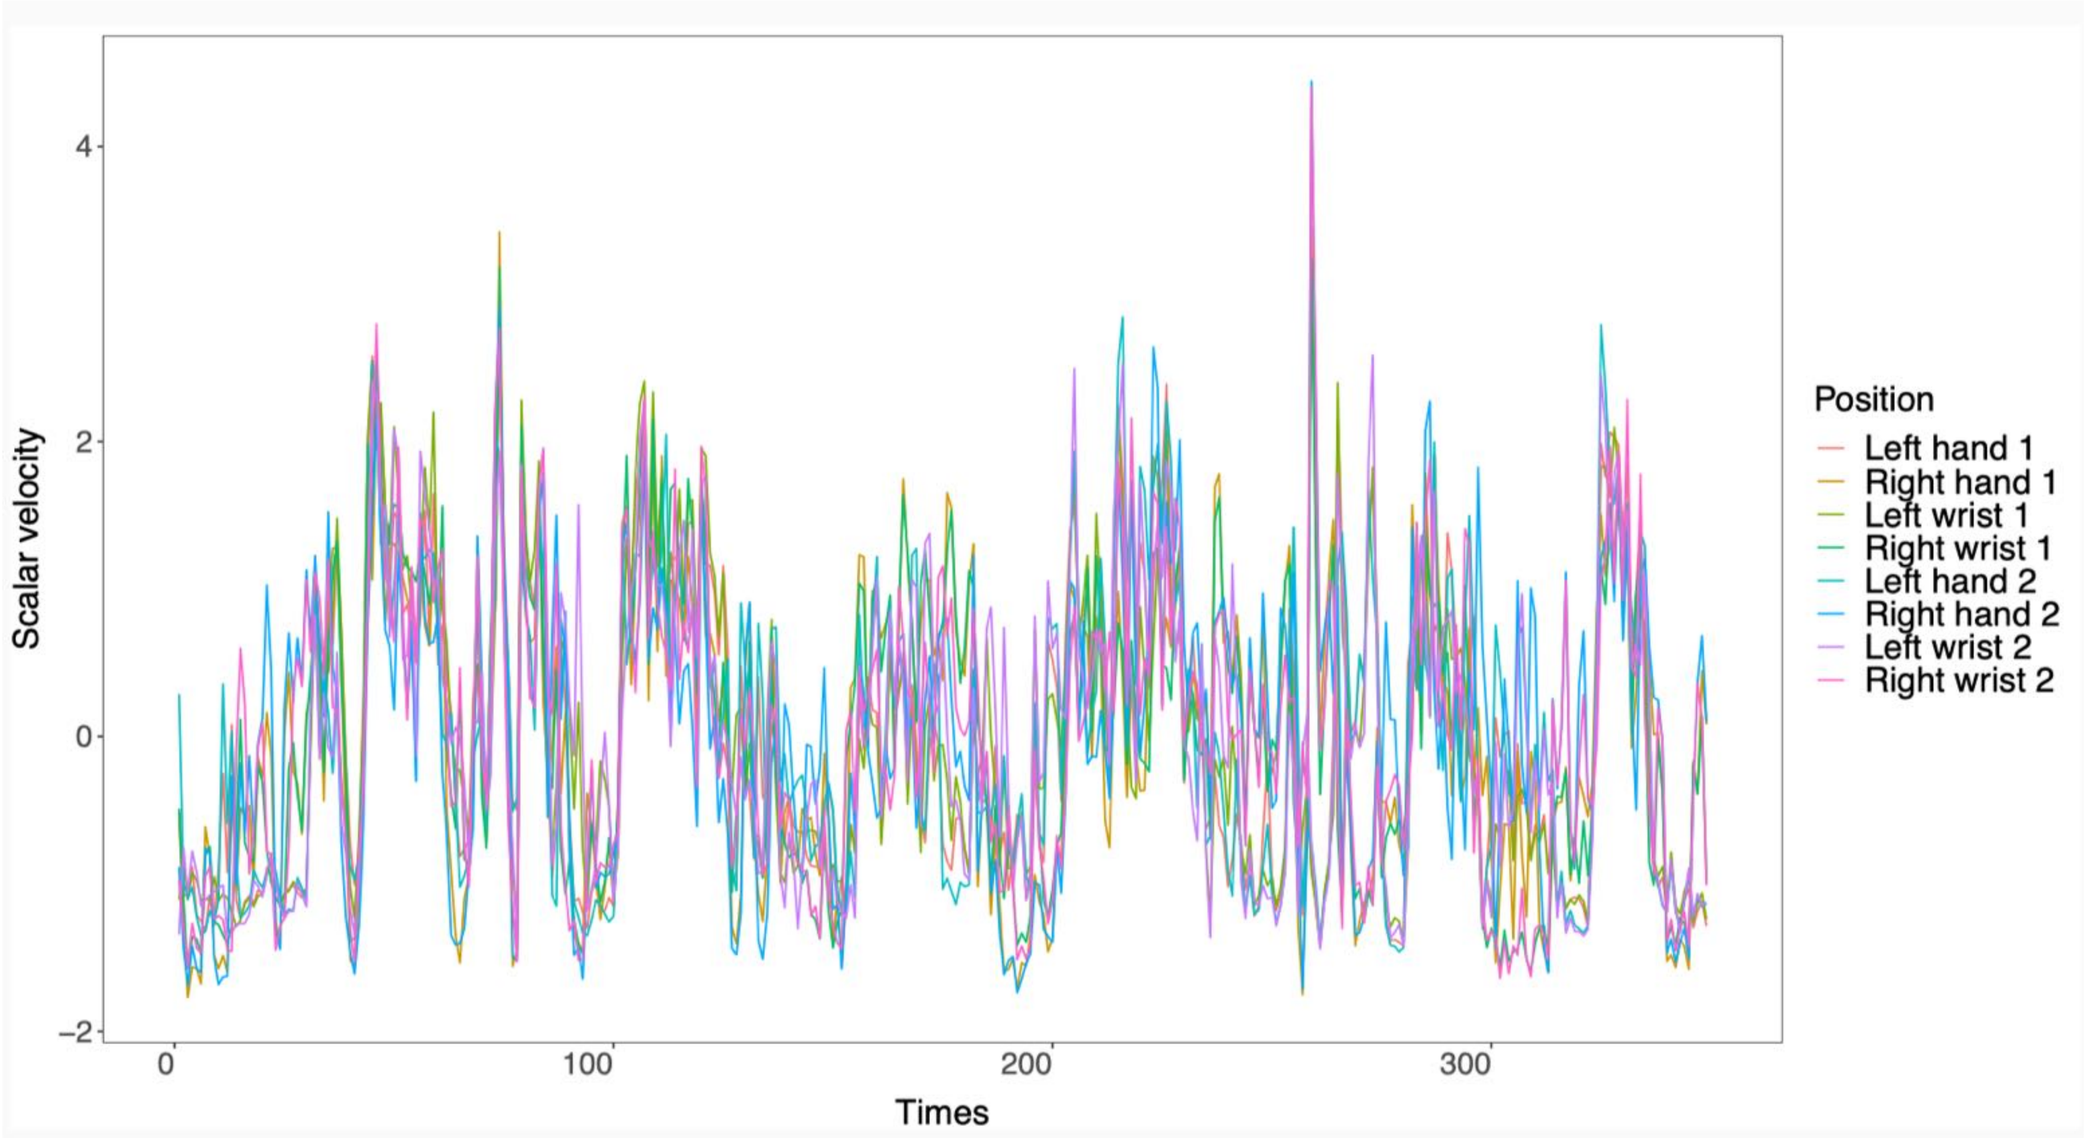

T = *349*

Changepoint detection

UCLA

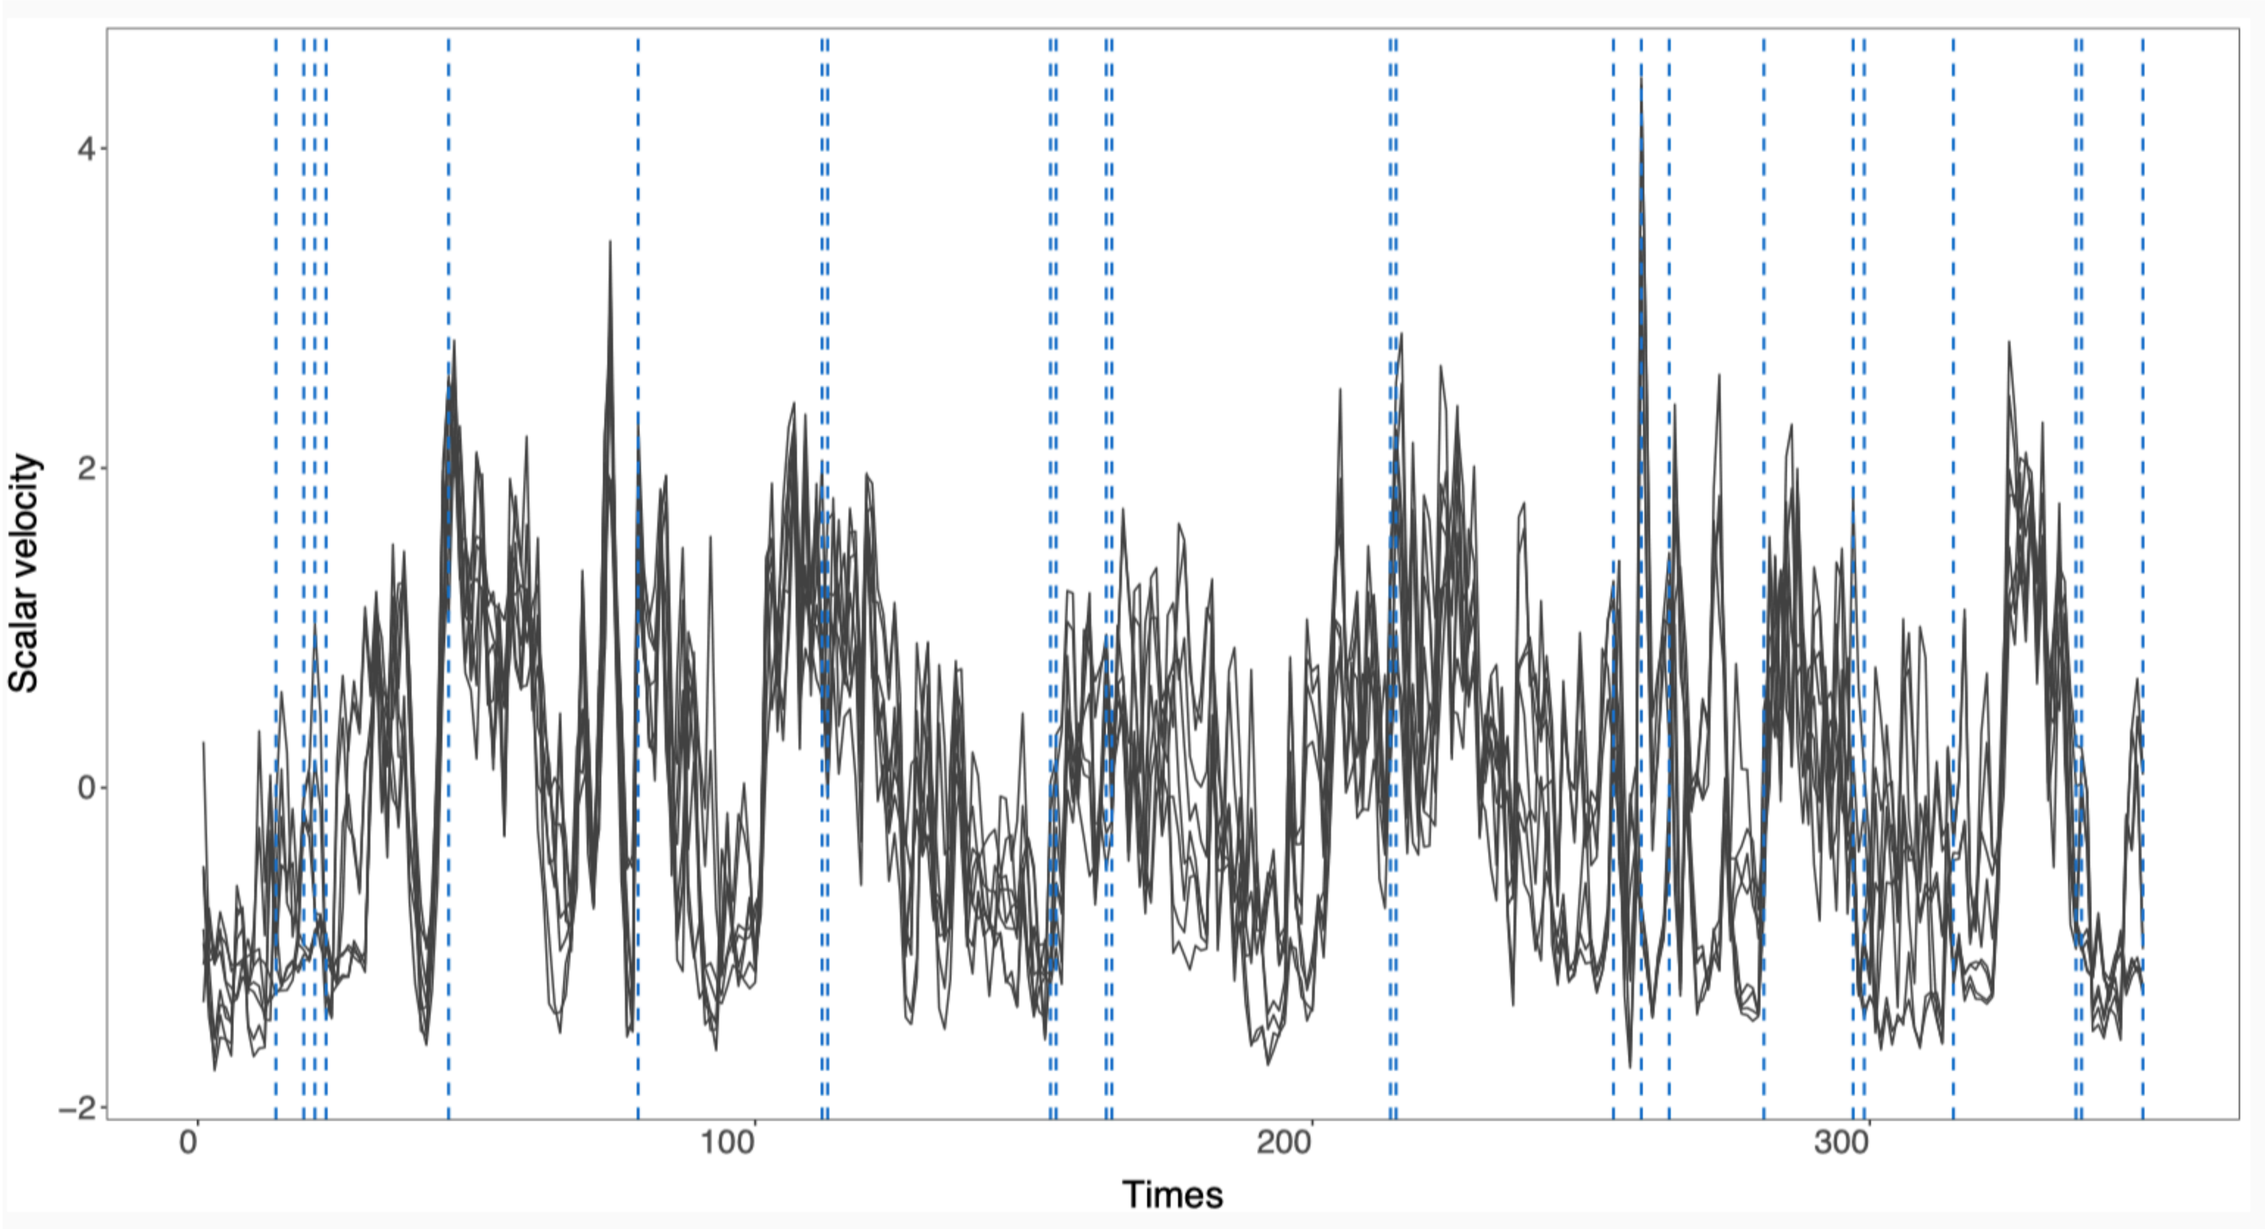

Zooming in on a time-window

UCLA

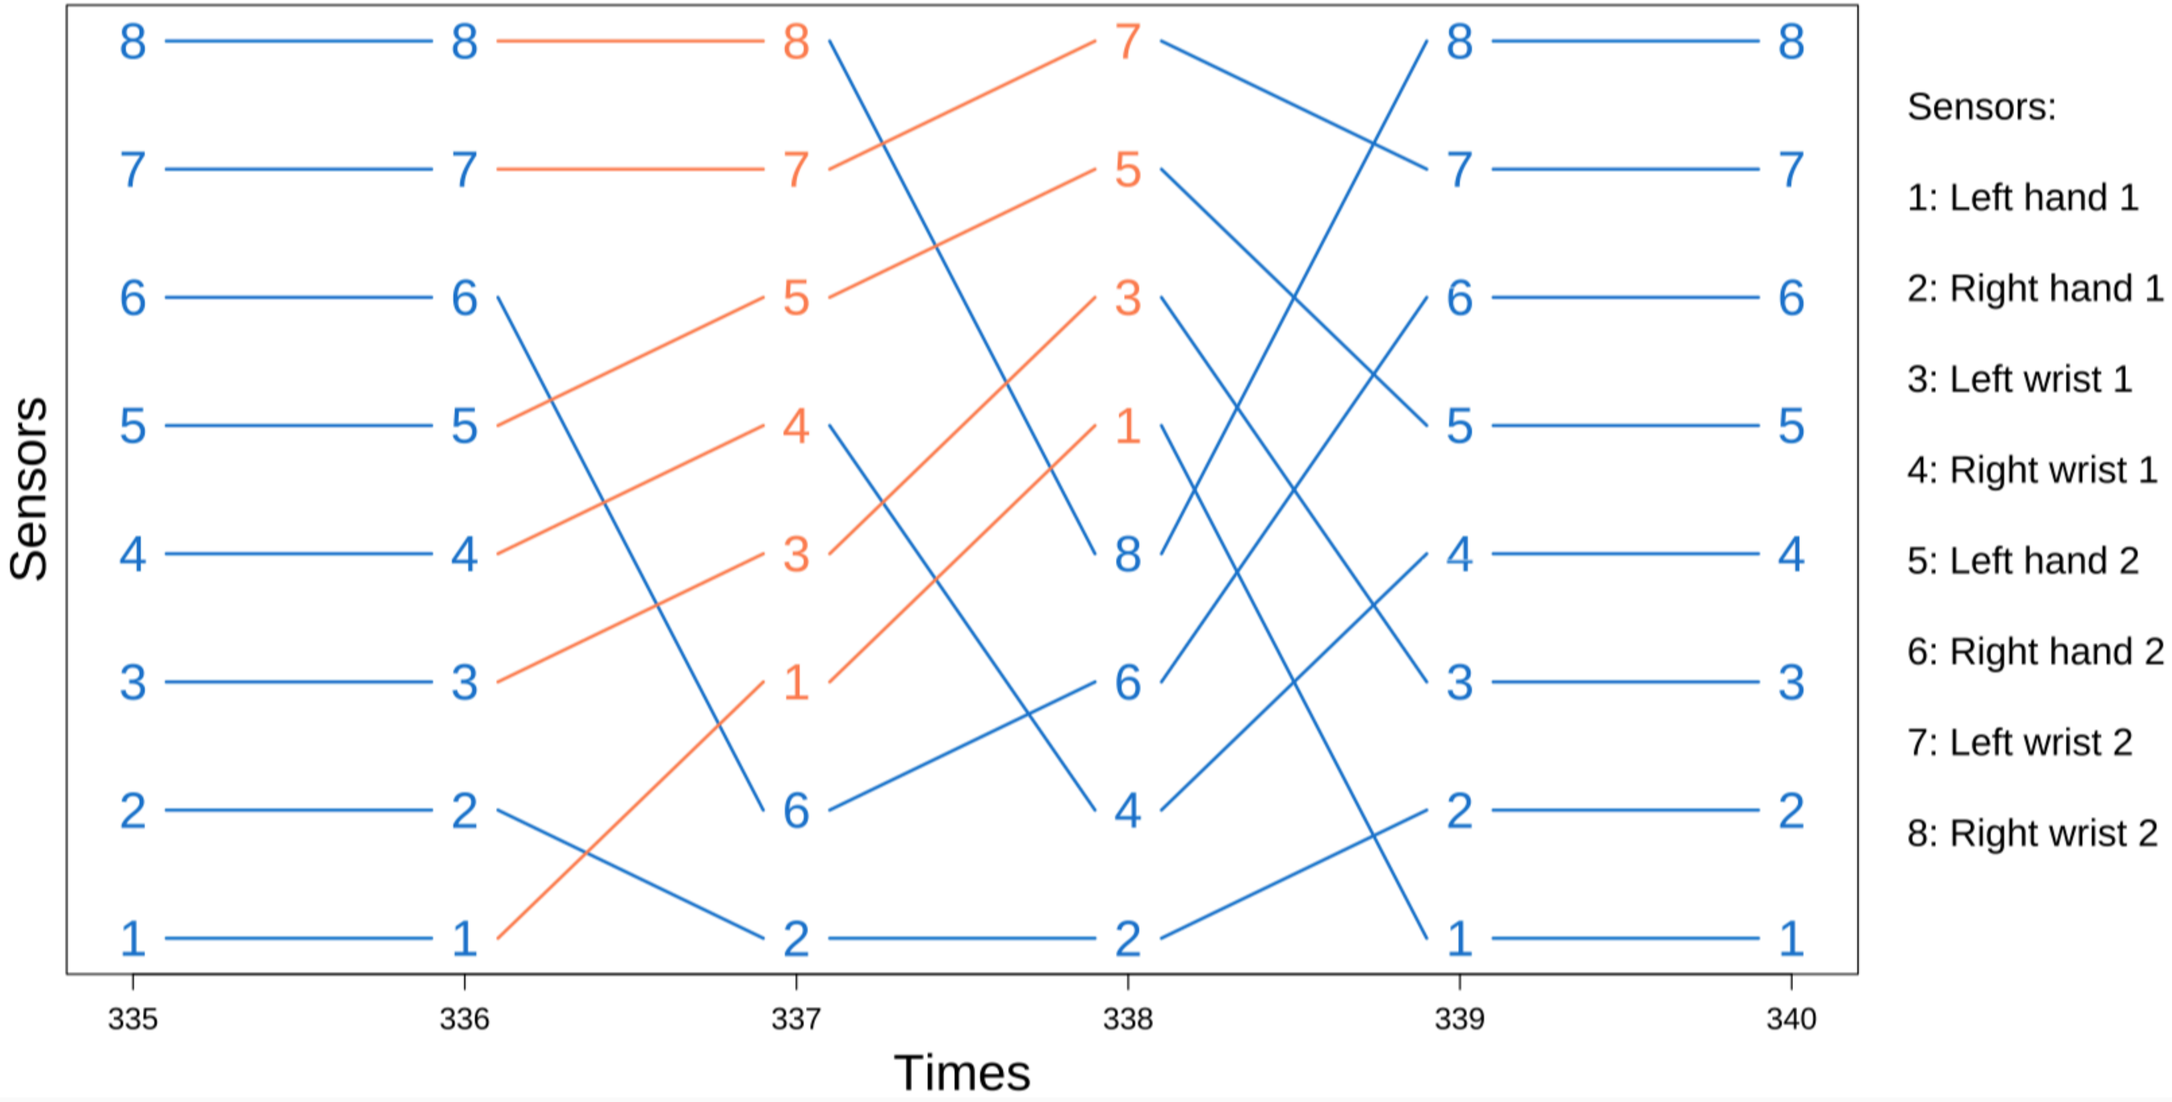

“Concluding” remarks: challenges in complex applications

UCLA

Scalability in all dimensions (units and time)

Incorporating available **information** (spatial dependence, covariates)

Combine BNP methods with **biological mechanistic models** over time (change points trigger different modeling)

Assessing **uncertainty** (in partitions, which is also a function of the inference on changepoints)

Interpretation of random partitions (e.g., association with measurable outcome

Decoding neuronal ensembles from spatially-referenced calcium traces

All the light that we can see: a new era in miniaturized microscopy

One major challenge in neuroscience is to uncover how defined neural circuits in the brain encode, store, modify, and retrieve information. Meeting this challenge comprehensively requires tools capable of recording and manipulating the activity of intact neural networks in naturally behaving animals. Head-mounted miniature microscopes are emerging as a key tool to address this challenge. Here we discuss recent work leading to the miniaturization of neural imaging tools, the current state of the art in this field, and the importance and necessity of open-source options. We finish with a discussion on what the future may hold for miniature microscopy.

Daniel Aharoni, Baljit S. Khakh, Alcino J. Silva and Peyman Golshani

Decades of neuroscience research have led to the development and refinement of diverse behavioral assays to probe the necessity and sufficiency of specific brain circuits and molecular pathways in a multitude of tasks[1,2]. In parallel, imaging techniques such as two-photon imaging emerged, allowing the structural and functional imaging of large neuronal networks at cellular resolution[3,4]. However, these imaging techniques required the animals to be head-fixed, thus limiting the behavioral repertoire and preventing researchers from using a large battery of behavioral tests that provide a wealth of associated information. At the same time, electrophysiological techniques were developed to probe the activity of large ensembles of neurons with single-cell resolution[5,6]. Yet these methods, for the most part, did not allow the identification of cell

Mapping neuronal activity in real time

UCLA

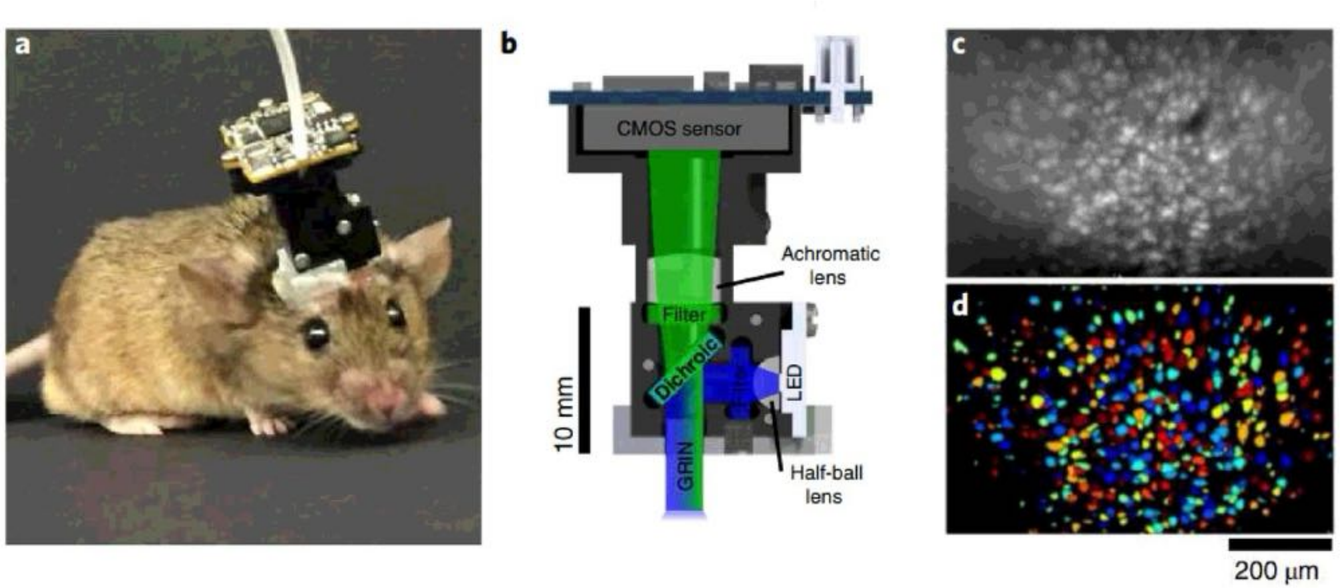

Fig. 1 | Open-source UCLA Miniscope. a, A mouse with a head-mounted Miniscope. **b**, Cross-sectional rendering of the Miniscope optical path. Blue, excitation path; green, emission optical path; GRIN, gradient-index lens. **c**, Maximum projection of a 10-minute motion-corrected Miniscope recording of hippocampal CA1 pyramidal neurons labeled with GCaMP6f. **d**, Spatial footprints of identified neurons from the recording in **c**. Scale bar in **d** applies to **c.**

- Calcium imaging is a microscopy technique to optically measure the intracellular calcium concentration of neurons in awake animals.

- The mechanism at the basis of calcium imaging is a physiological process of the cells: when a neuron fires, **calcium floods the cell** and produces a transient spike in its concentration

- Fluorescent Calcium Indicators bind to calcium ions during neuronal activation

👉 Outcome: **movie** of time-varying fluorescence intensities for each observable neuron in a targeted area.

Calcium imaging measurements

UCLA

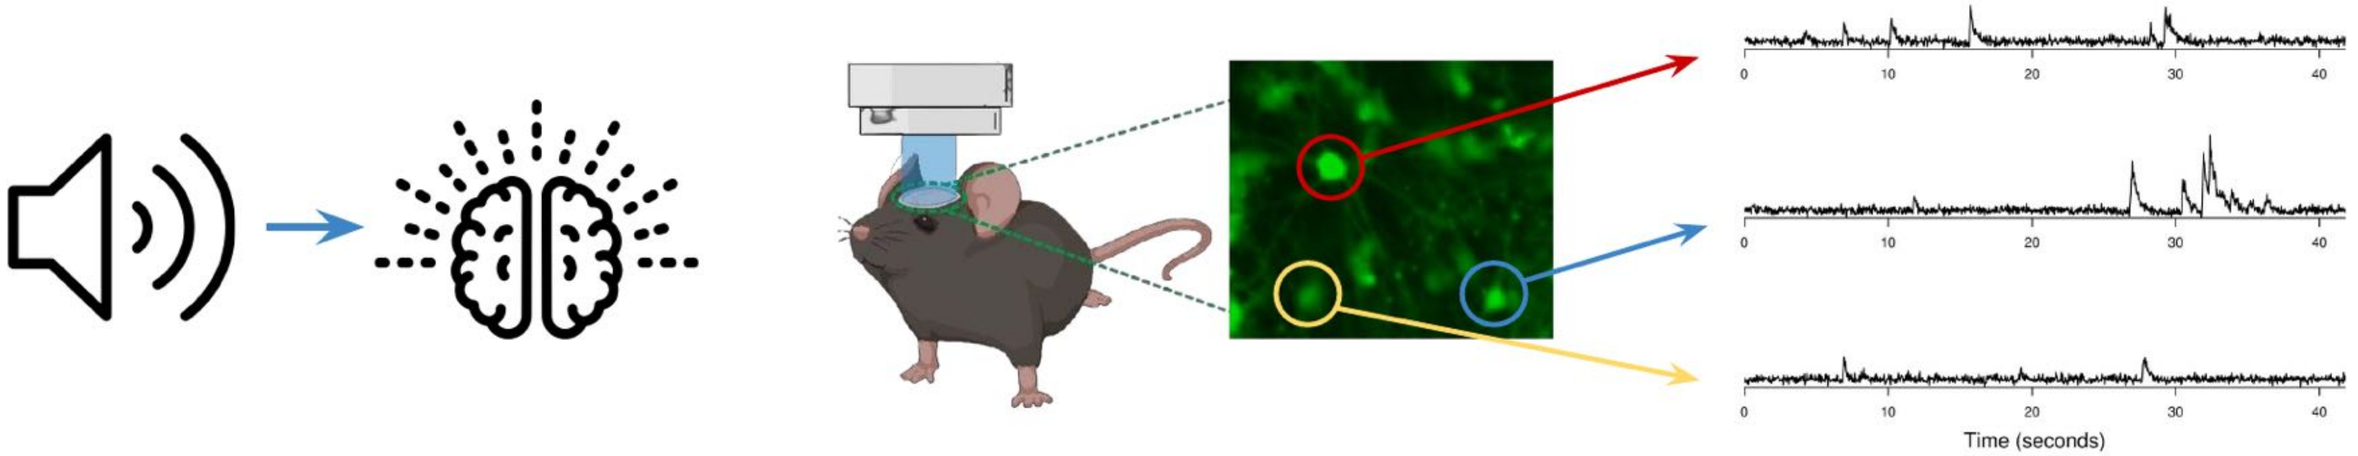

- Physiological process (often model mathematically) behind calcium imaging:
  - 👉 **External Stimulus:** Triggers a neuronal response
  - 👉 **Neuronal Activation:** Calcium floods the cell, causing a temporary increase in intracellular calcium concentration
  - 👉 **Return to Baseline:** calcium levels return to their normal state as the neuron goes back to rest
- Fluorescent calcium traces are **proxies of the activity** over time of individual neurons

Neuronal Data from a freely moving mice

UCLA

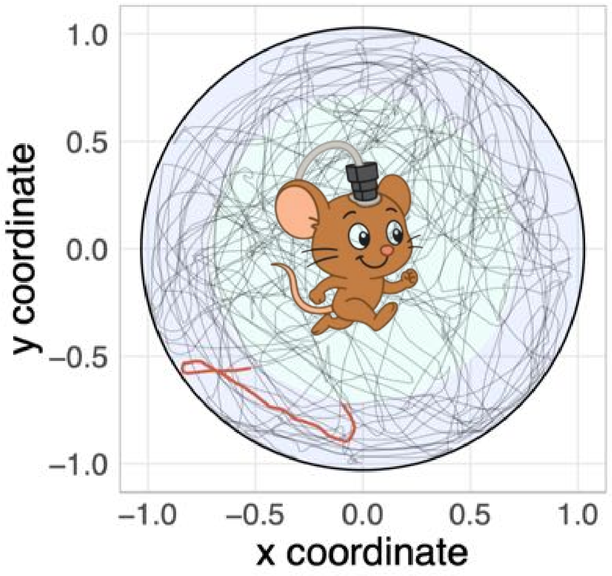

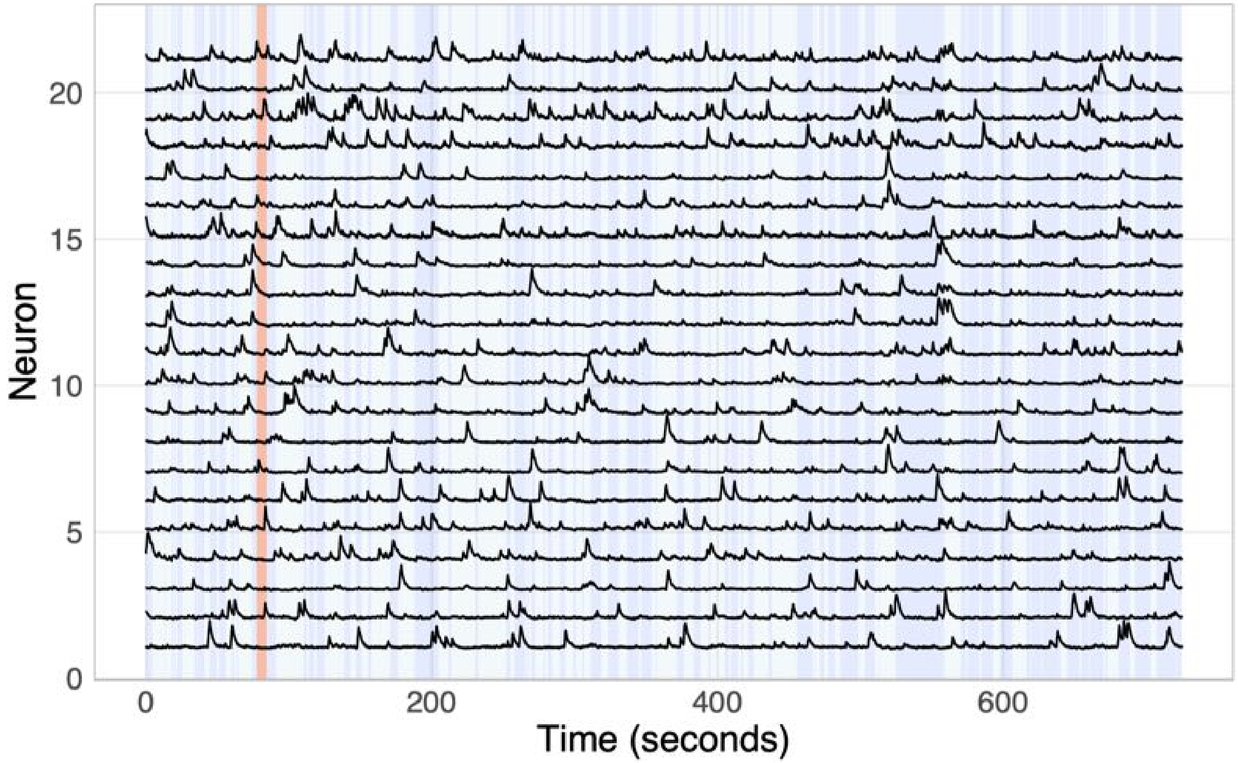

- Recorded hippocampal CA1 activity in a freely moving mouse exploring a circular arena (Chen et al, 2023)
- CA1 supports spatial navigation and episodic memory (implicated also in AD)
- **229 neurons** over **5,435 time points** across a 12-minute session
- The arena is split into a **center** and **outer ring**, and the time series is segmented into **position-defined windows**

Our Data

UCLA

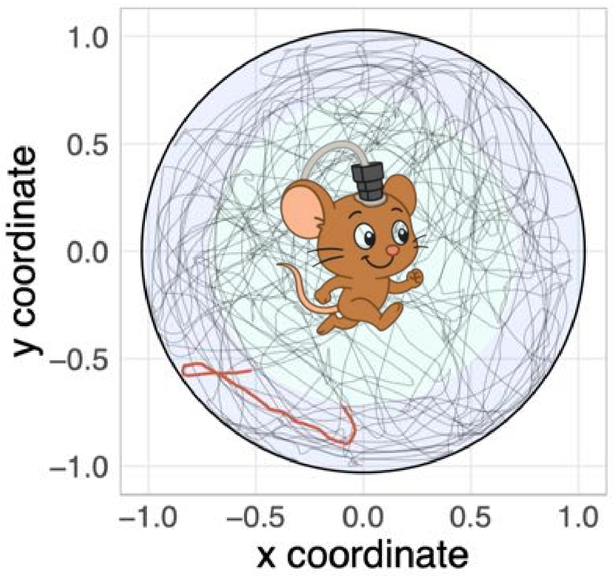

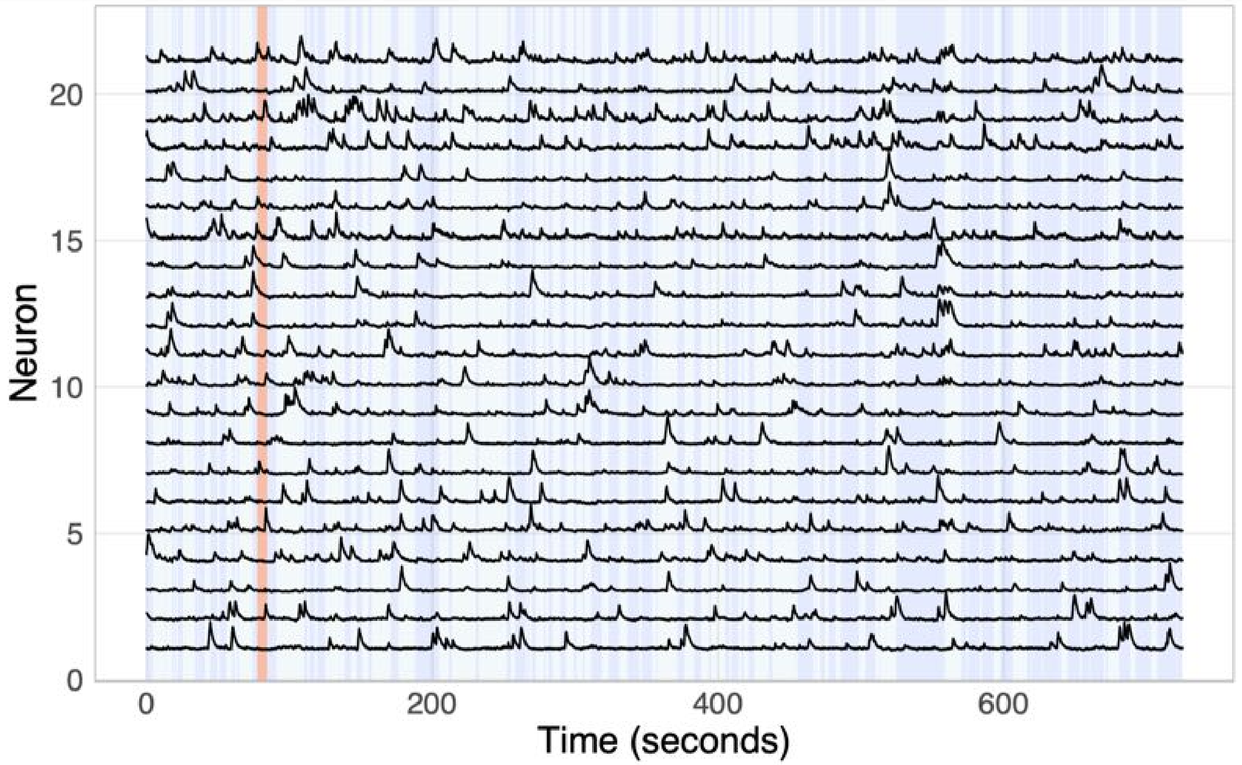

229 neurons, more than 5000 time points; scalability becomes important

1. **Jointly infer** spikes and cluster neurons from calcium imaging
2. Enforce **spatially coherent clustering**
3. Link neural ensembles to **behavior** & examine **context-dependent** shifts in clustering **(doubly-spatial problem)**

Clustering of neurons overe a time window

UCLA

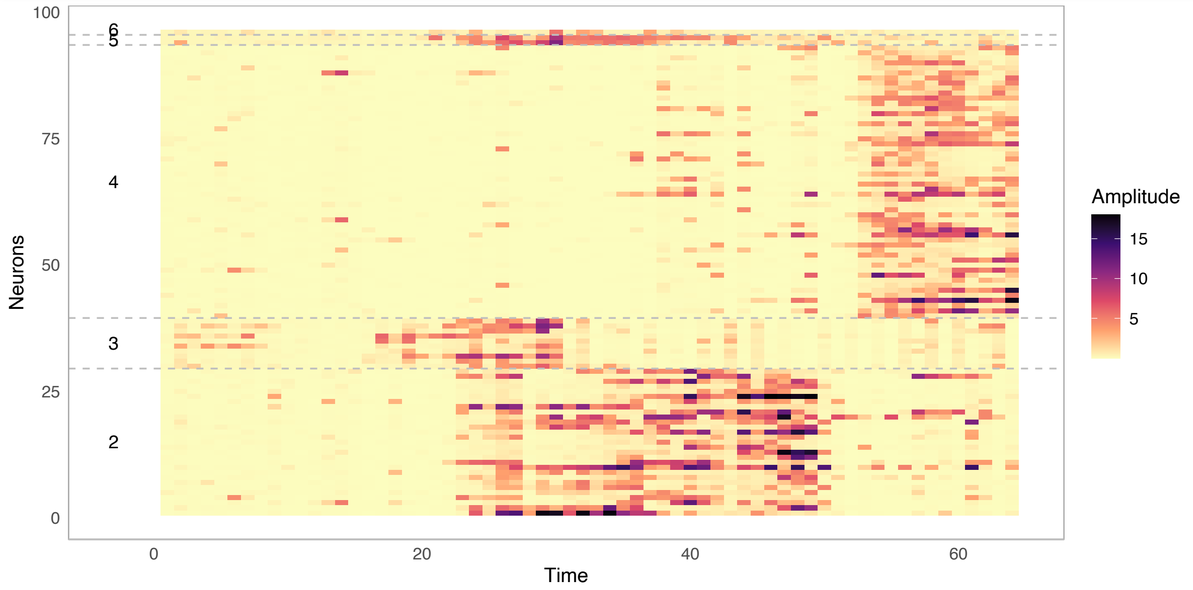

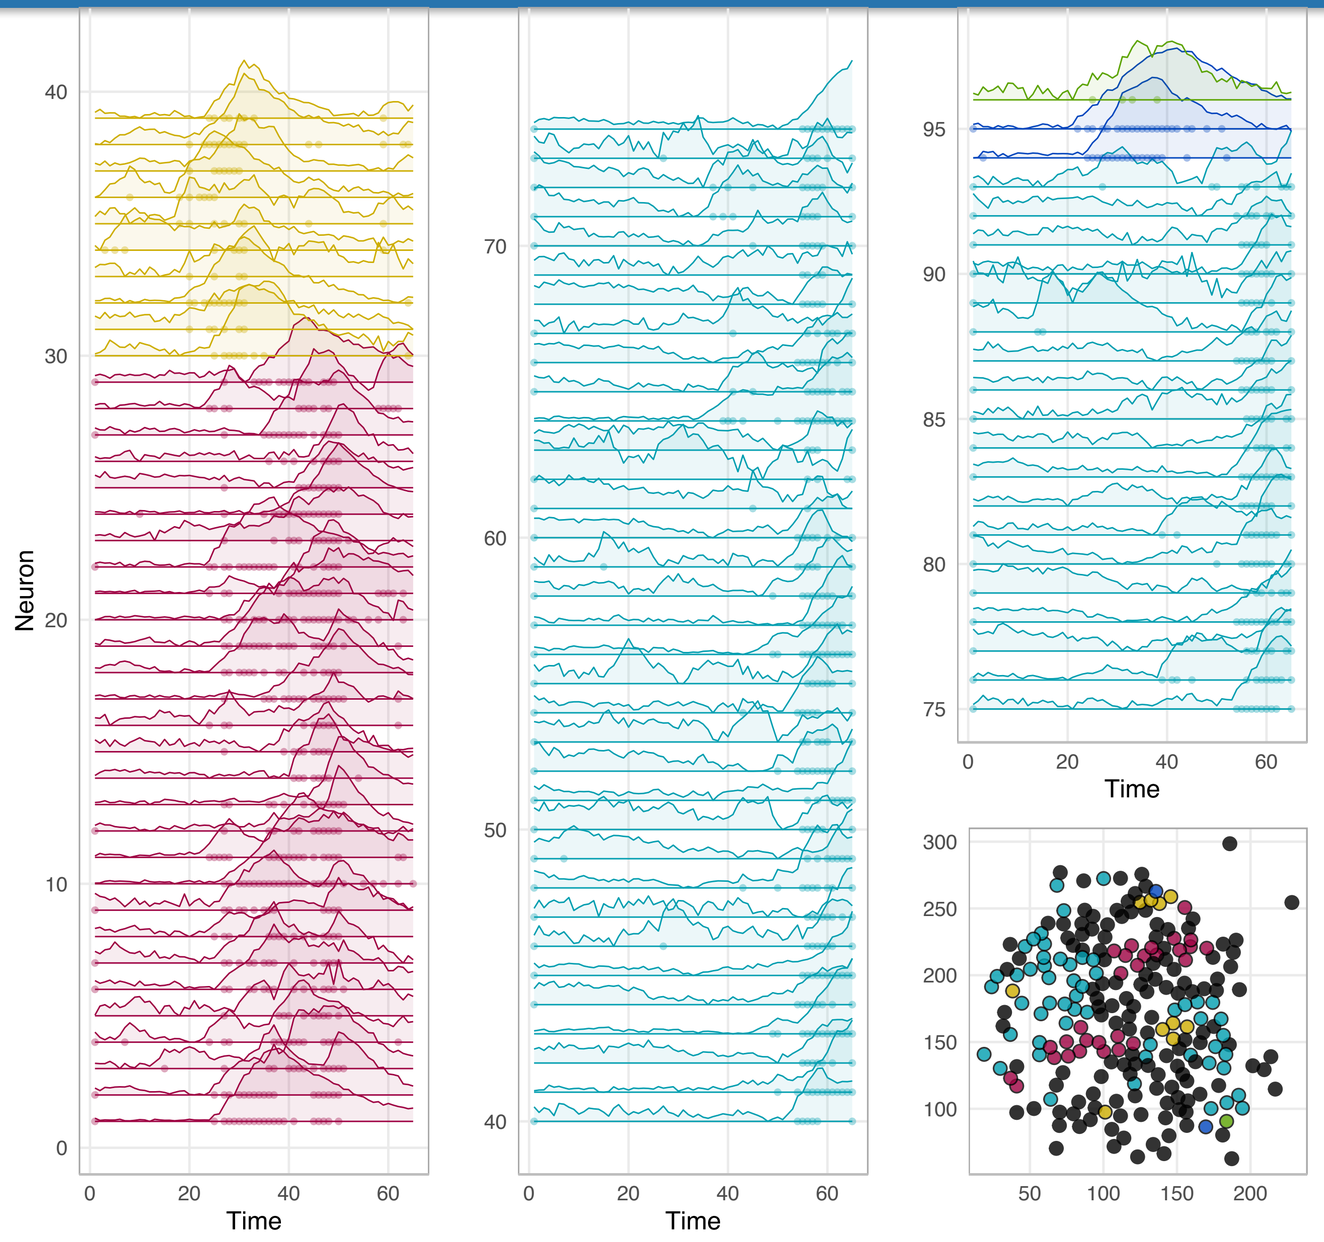

- 6 clusters
- Top: spike trains & amplitudes
- Right: time series and neurons' locations colored by cluster (cluster 1 with 122 inactive neurons not reported)

Neuronal responses to mouse position

UCLA

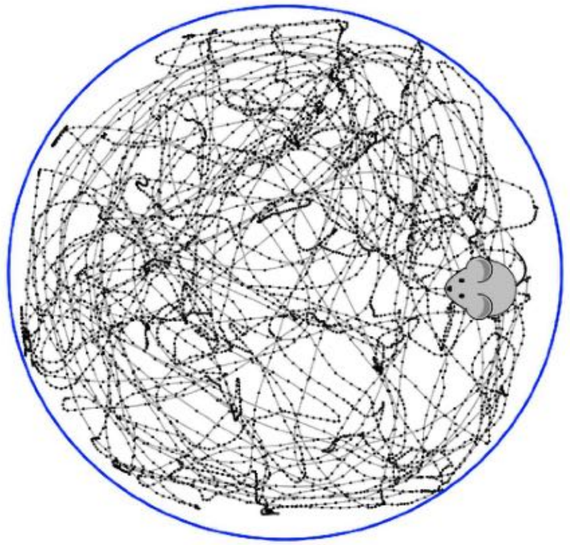

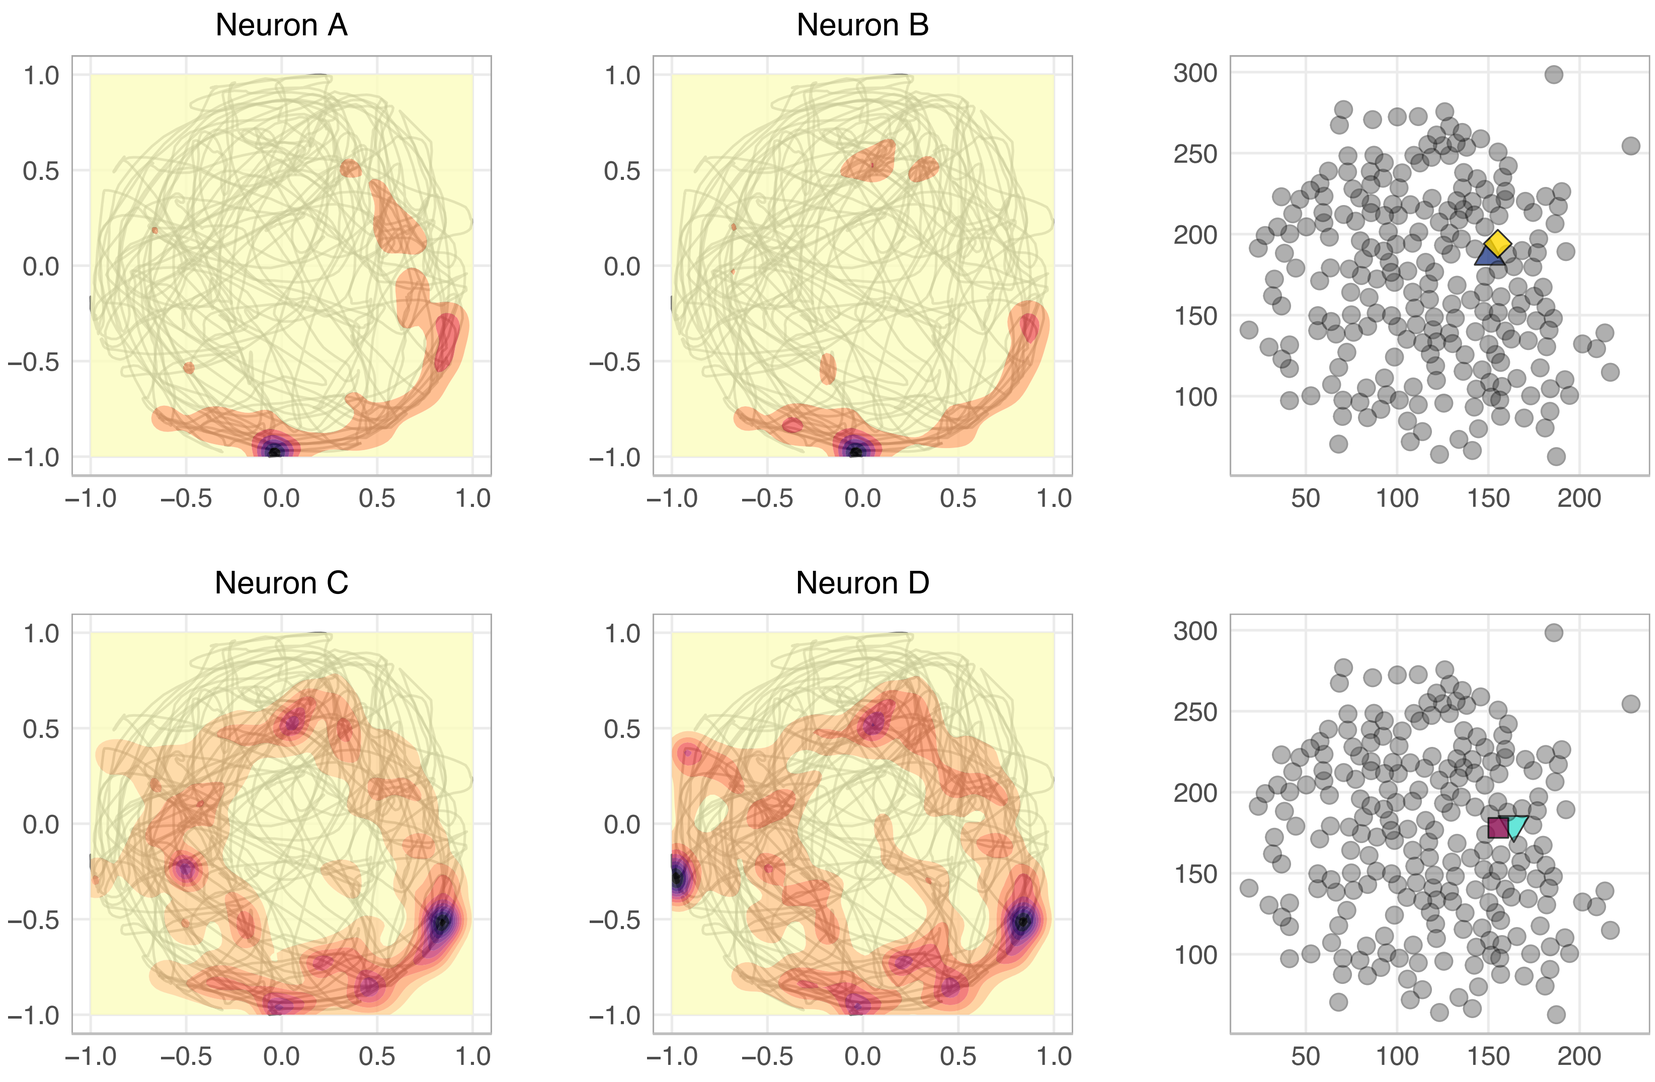

- **Co-clustering analysis:**

  select neurons frequently clustered together and try to understand their **spike patterns associated with the mouse position**

Spatial cluster variability

UCLA

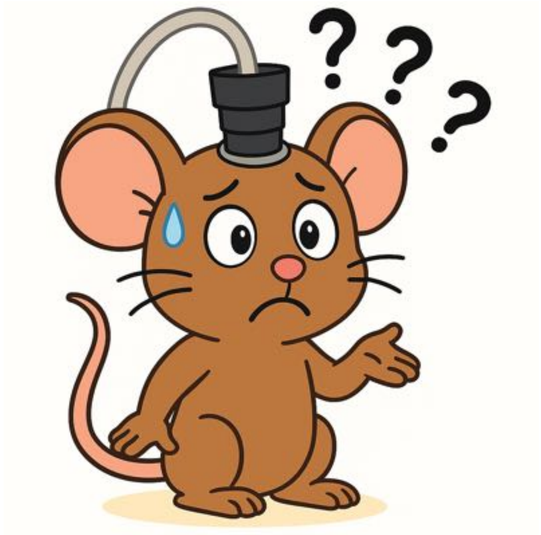

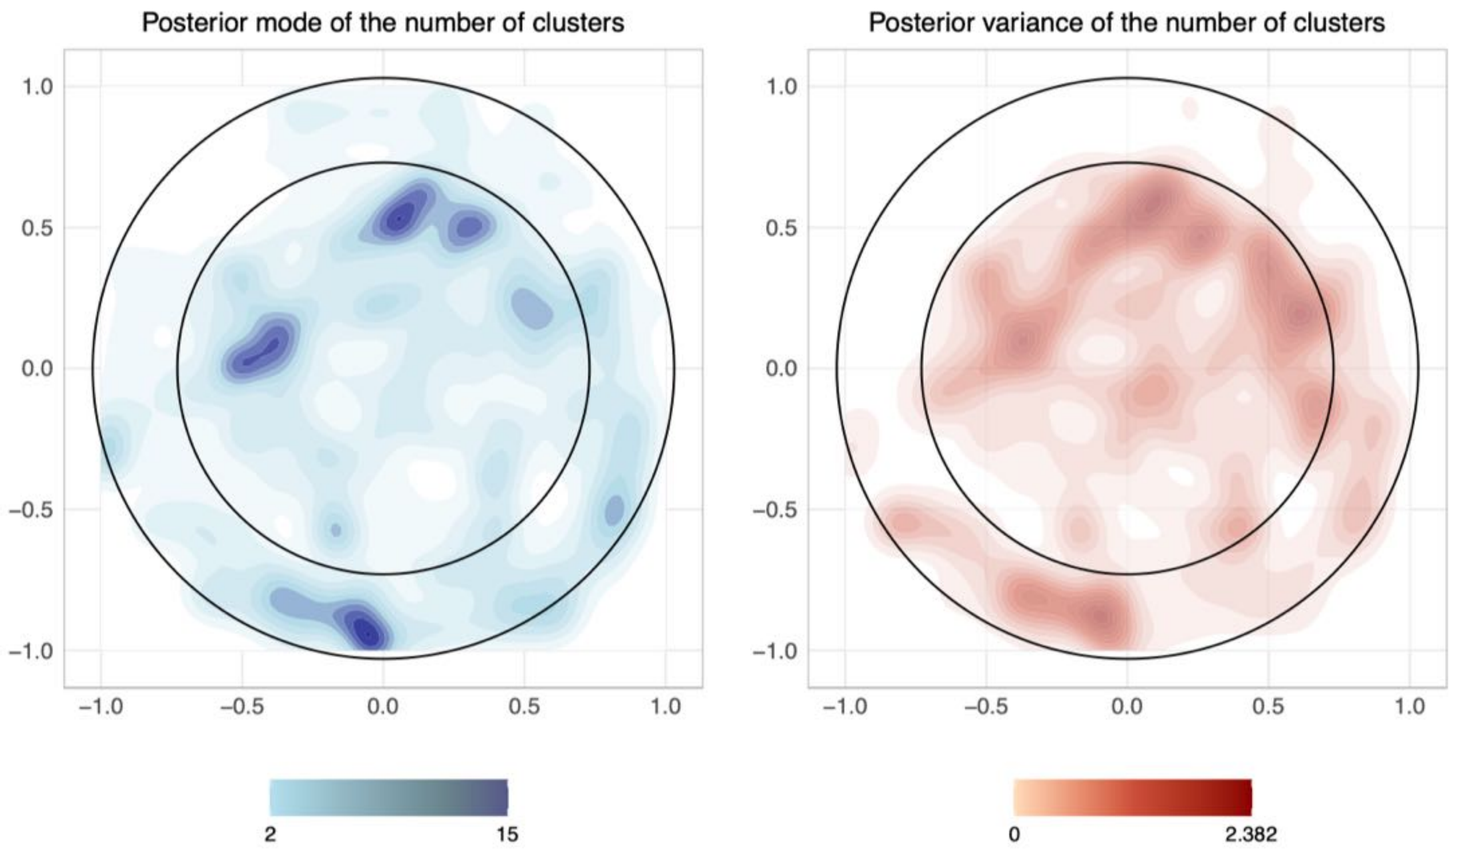

Heatmaps showing the spatial distribution of the clustering complexity and variability. Each point of the mouse trajectory is weighted by the mode (left panel) and variance (right panel) of the posterior distribution of the number of clusters in that location.

Similar data are also available on humans

UCLA

Article

Distinct neuronal populations in the human brain combine content and context

Marcel Bausch[1✉], Johannes Niediek[1,2], Thomas P. Reber[1,3], Sina Mackay[1], Jan Boström[4], Christian E. Elger[1] & Florian Mormann[1✉]

https://doi.org/10.1038/s41586-025-09910-2
Received: 10 August 2023
Accepted: 12 November 2025
Published online: 07 January 2026
Open access

Participants/recordings: 16 neurosurgical epilepsy patients; **3,109 single units** recorded across amygdala, parahippocampal cortex, entorhinal cortex, and hippocampus.

Neural data: microwire recordings from implanted intracranial electrodes in epilepsy patients.

Behavioral data: A **context question** (the rule), one of **five**: *Bigger? Last seen in real life? More expensive? or Older?* (depending on picture set), *Like better? Brighter?* ➡ Two pictures are shown sequentially (two of the four). ➡An answer prompt "1 or 2?"; the participant presses 1 or 2 to indicate which picture best fits the question.

Separate neurons into (mostly non-overlapping) sets based on whether their firing rates depend on:
Content (stimulus identity): neurons whose activity changes depending on *which picture* is shown
Context (task rule/question): neurons whose activity changes depending on *which question/rule* is active
Conjunctive (stimulus × context): neurons that respond specifically to a particular **combination** of picture and question

Bayesian temporal biclustering with applications to multi-subject neuroscience studies

Local Level Dynamic Random Partition Models for Changepoint Detection (BA, in press)

Decoding Neuronal Ensembles from Spatially-Referenced Calcium Traces: A Bayesian Semiparametric Approach

Thank you!