Dimensionality reduction of neural dynamics within and across trials by tensor decomposition

Alex H. Williams, Tony Hyun Kim, Forea Wang,
Saurabh Vyas, Krishna V. Shenoy, Mark Schnitzer,
Tamara G. Kolda, Surya Ganguli

Modern experiments capture a large range of timescales in neural data

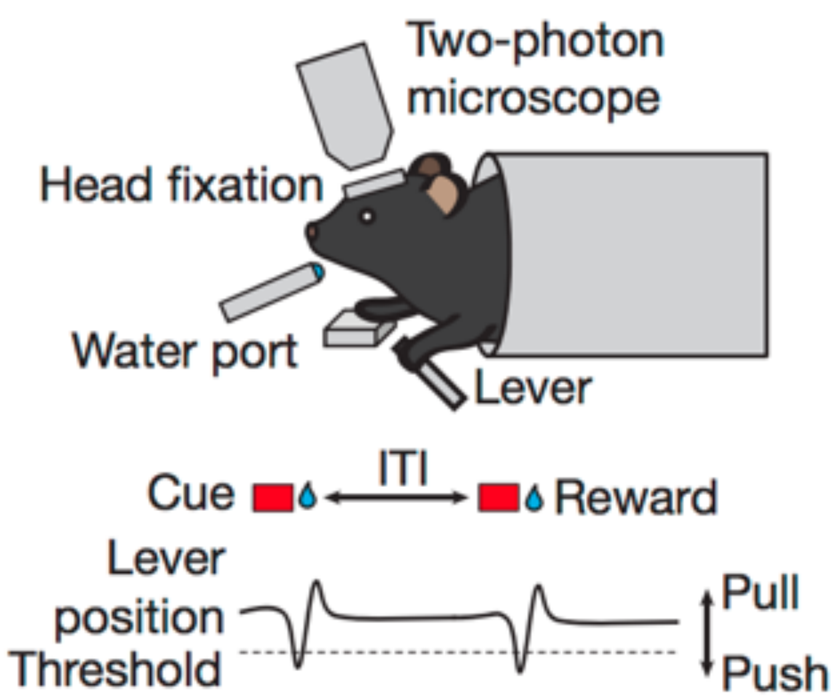

(Peters et al., 2014)

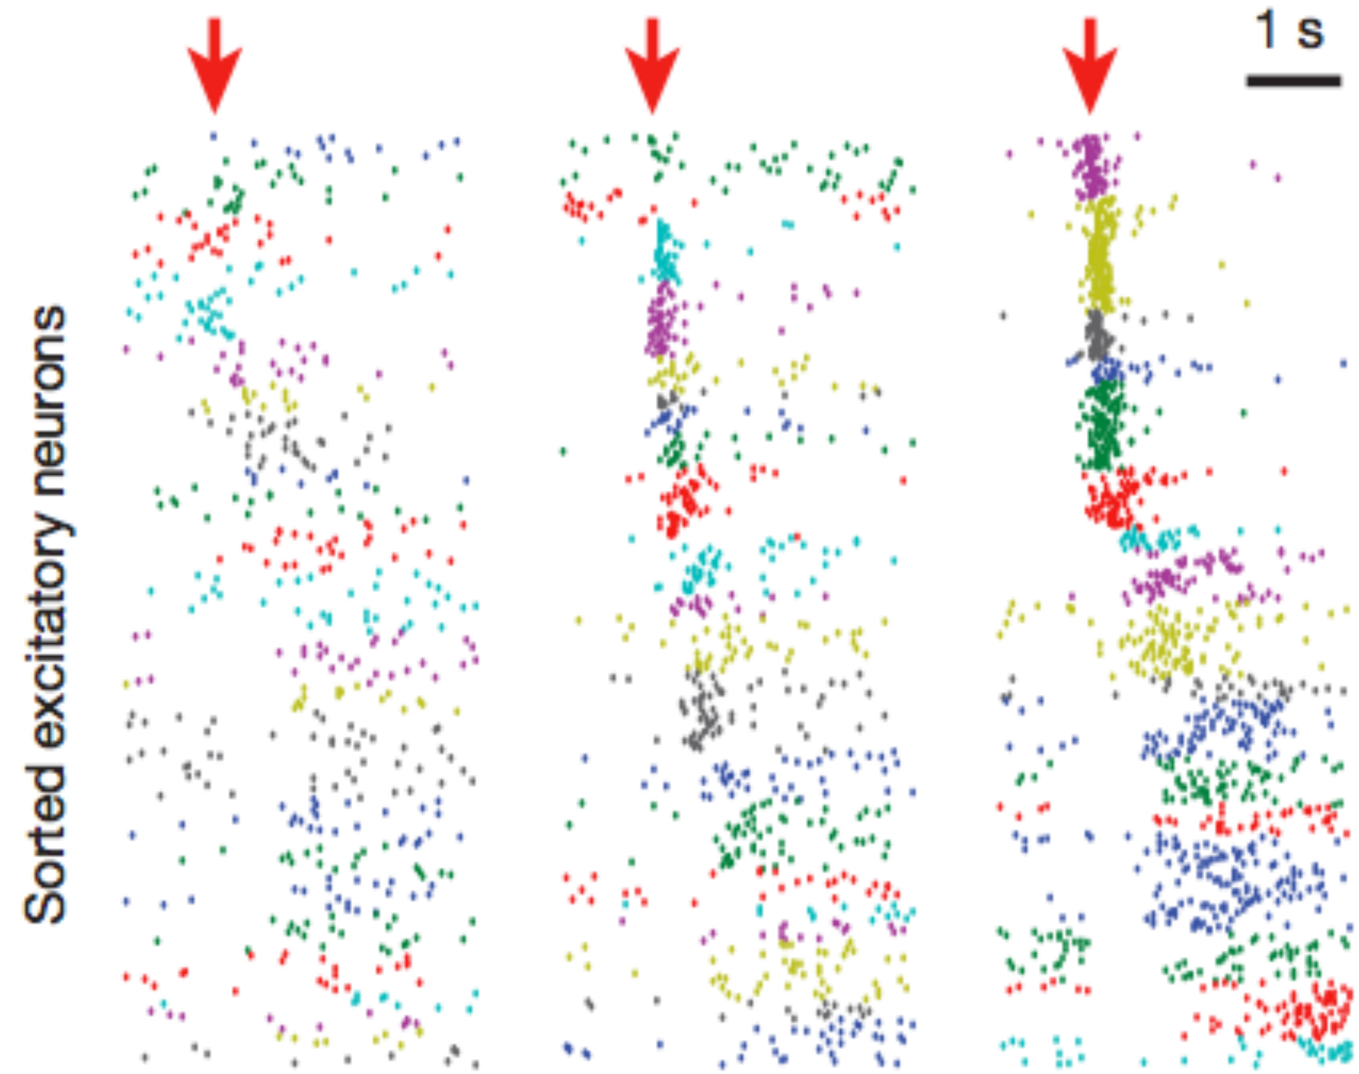

session A session B session C

Neural data is large, *along multiple dimensions*

- Ability to record from **thousands of neurons**, over **thousands of trials**.

Neural data is large, *along multiple dimensions*

- Ability to record from **thousands of neurons**, over **thousands of trials**.
- Existing dimensionality reduction techniques do not separate **fast, within-trial** dynamics from **slow, across-trial** components.

Neural data is large, *along multiple dimensions*

- Ability to record from **thousands of neurons**, over **thousands of trials**.
- Existing dimensionality reduction techniques do not separate **fast, within-trial** dynamics from **slow, across-trial** components.
- We propose **tensor decomposition** as a general-purpose framework to address this problem.

Neural data is large, *along multiple dimensions*

- Ability to record from **thousands of neurons**, over **thousands of trials**.
- Existing dimensionality reduction techniques do not separate **fast, within-trial** dynamics from **slow, across-trial** components.
- We propose **tensor decomposition** as a general-purpose framework to address this problem.
- We apply this method "off-the-shelf" to three datasets:

Neural data is large, *along multiple dimensions*

- Ability to record from **thousands of neurons**, over **thousands of trials**.
- Existing dimensionality reduction techniques do not separate **fast, within-trial** dynamics from **slow, across-trial** components.
- We propose **tensor decomposition** as a general-purpose framework to address this problem.
- We apply this method "off-the-shelf" to three datasets:

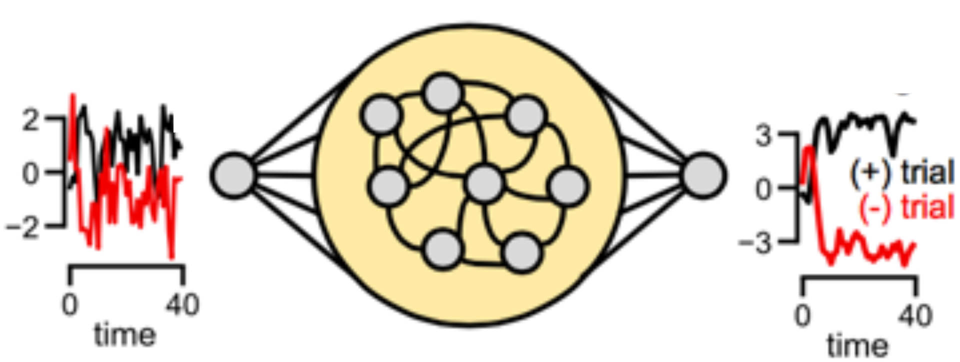

artificial RNN learning

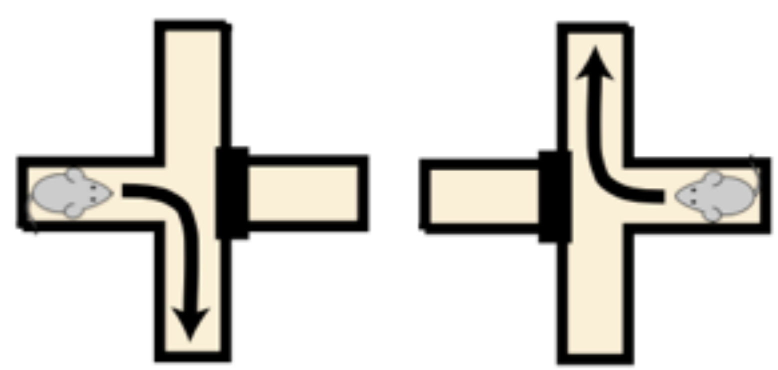

maze navigation

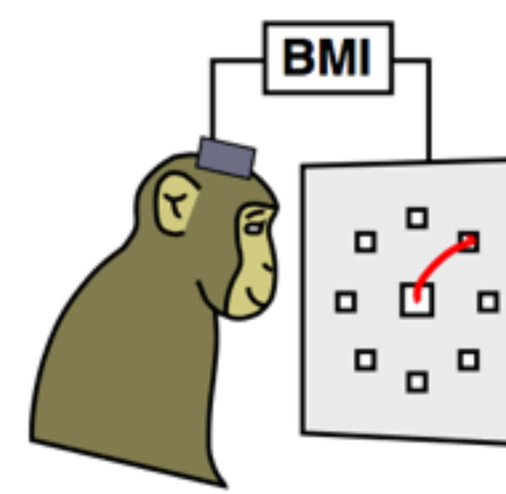

motor learning

Existing method #1: Trial-averaged PCA

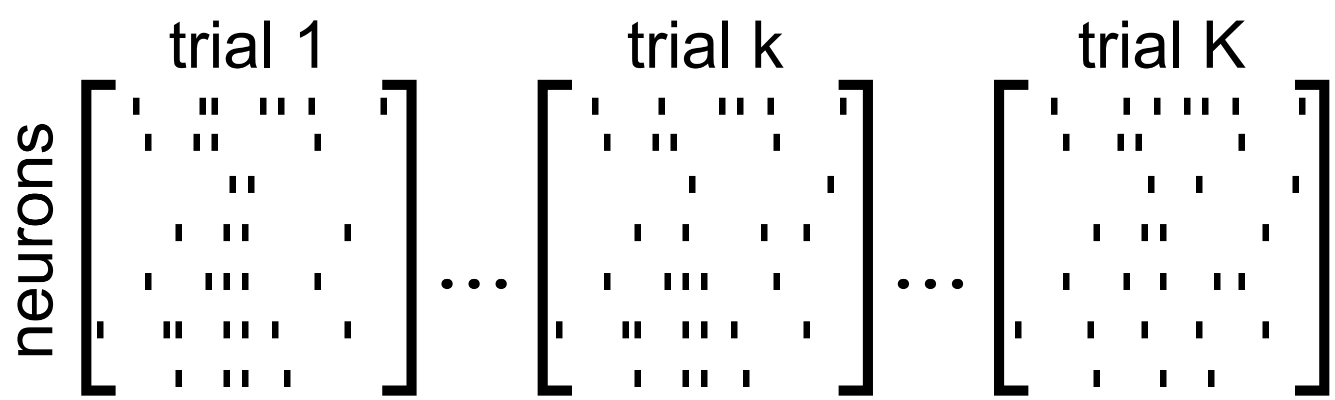

Existing method #1: Trial-averaged PCA

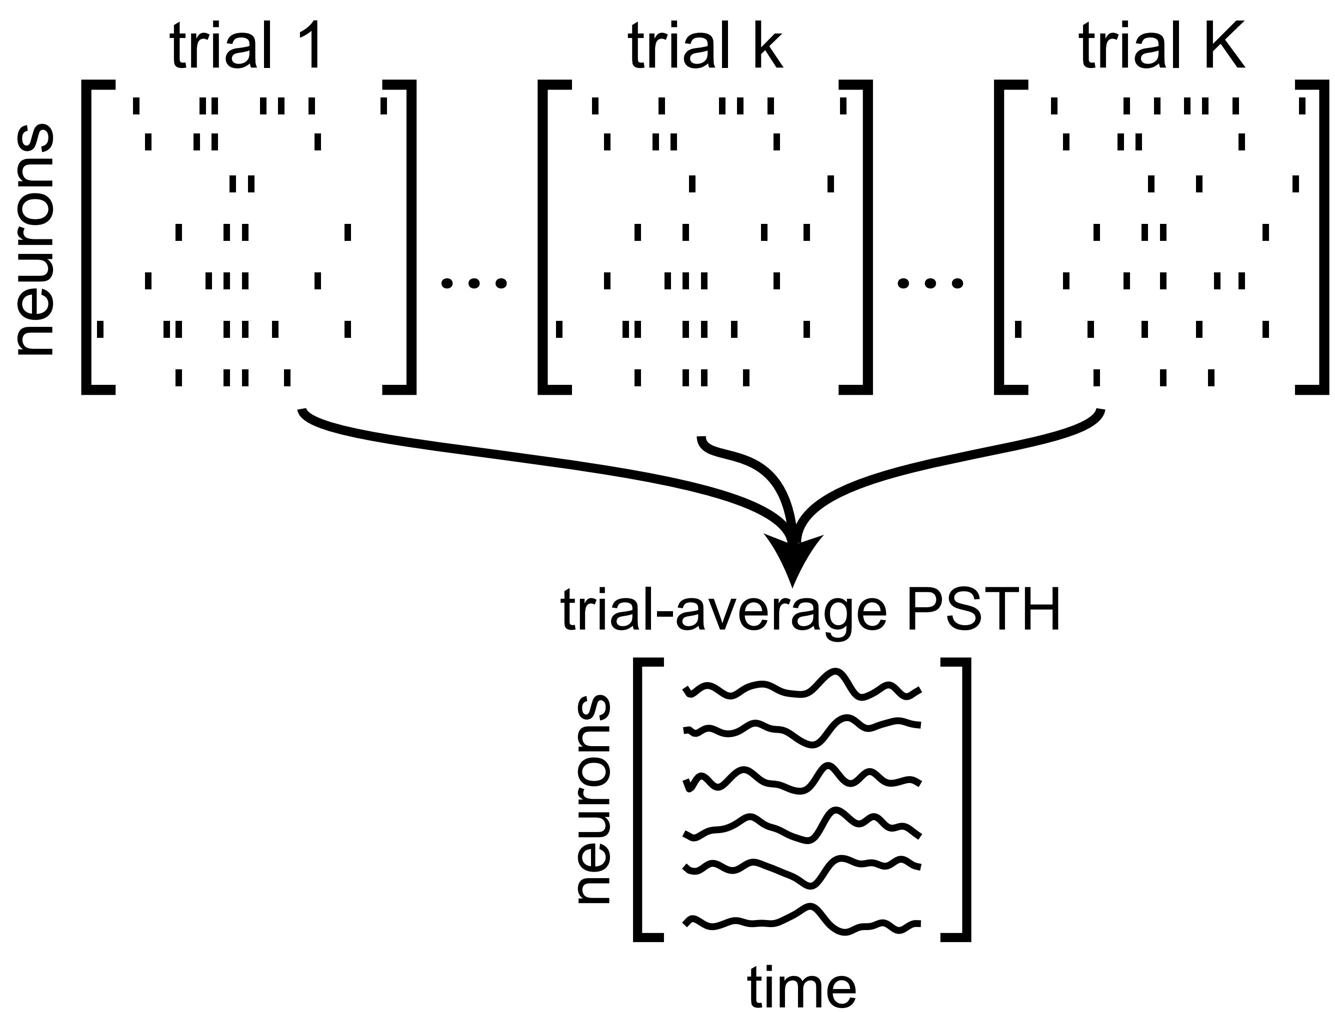

Existing method #1: Trial-averaged PCA

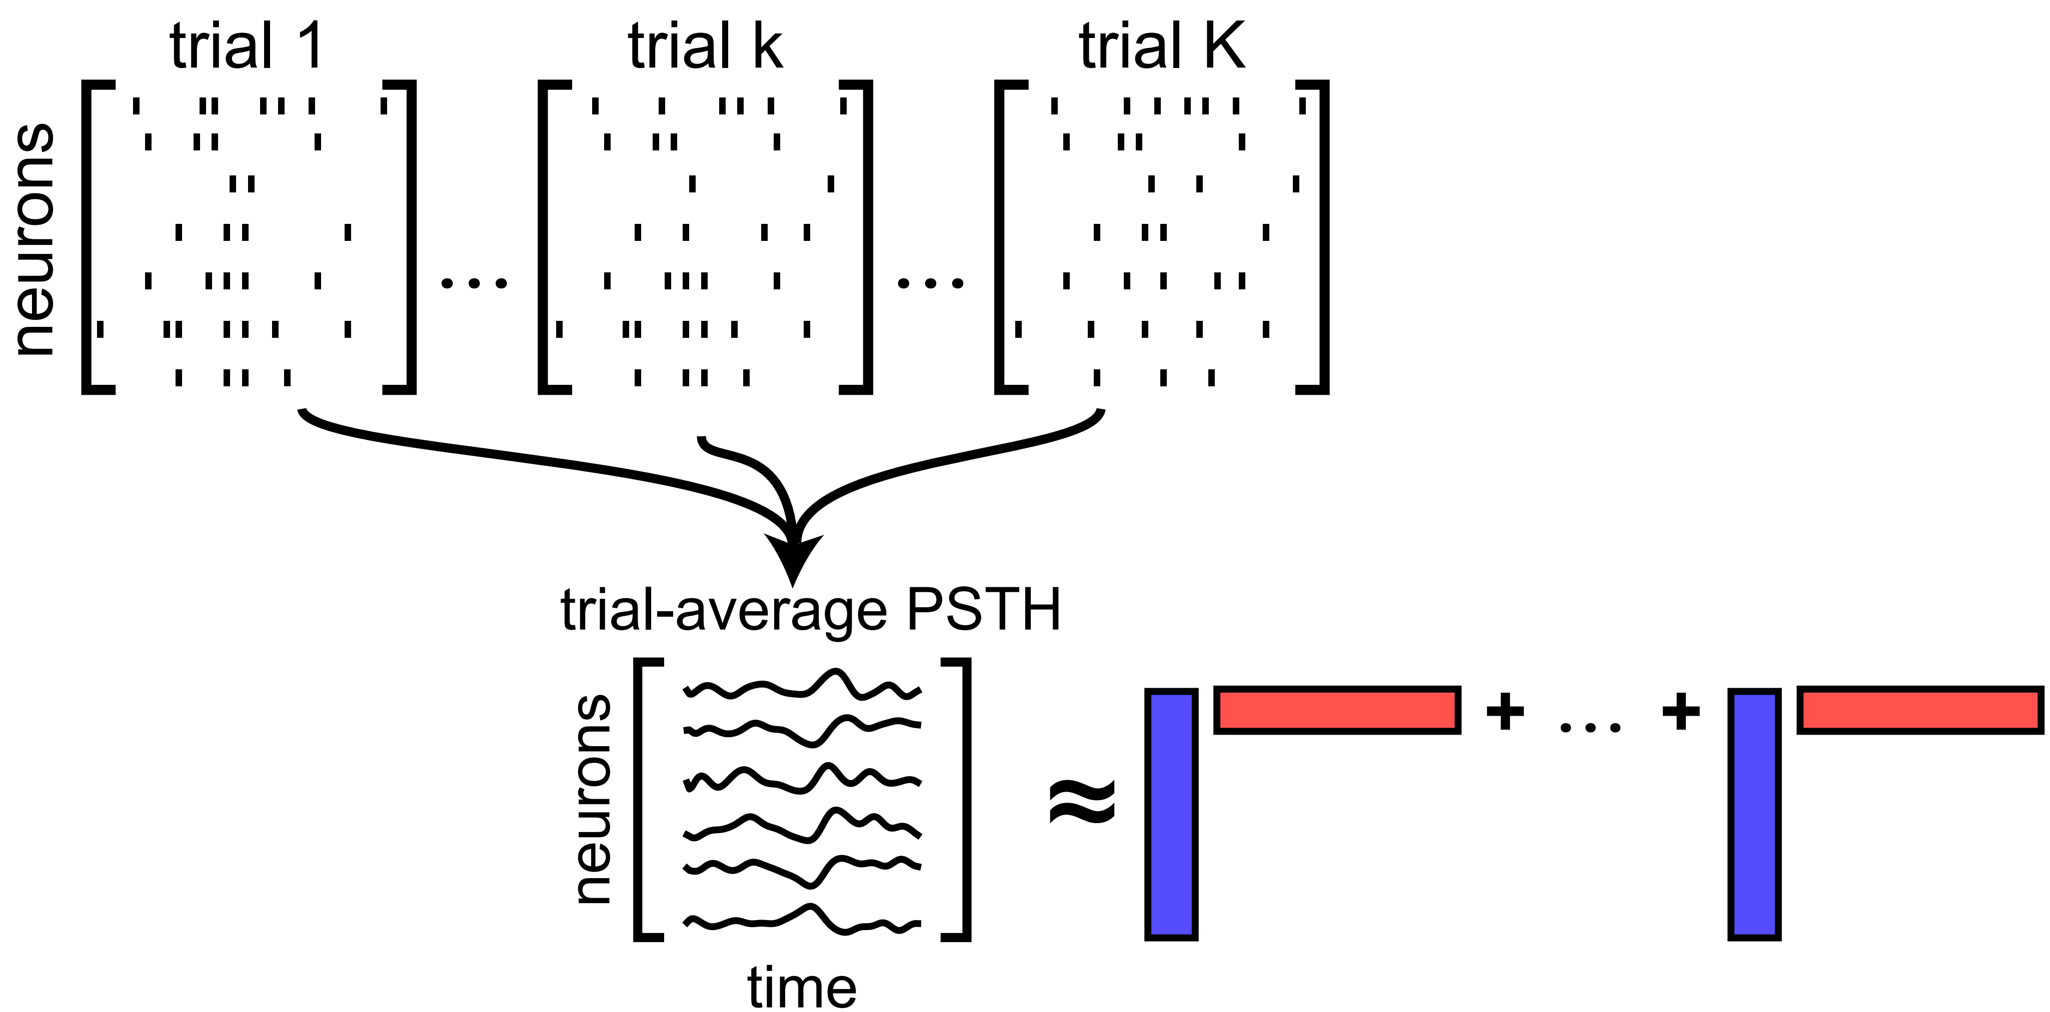

Existing method #1: Trial-averaged PCA

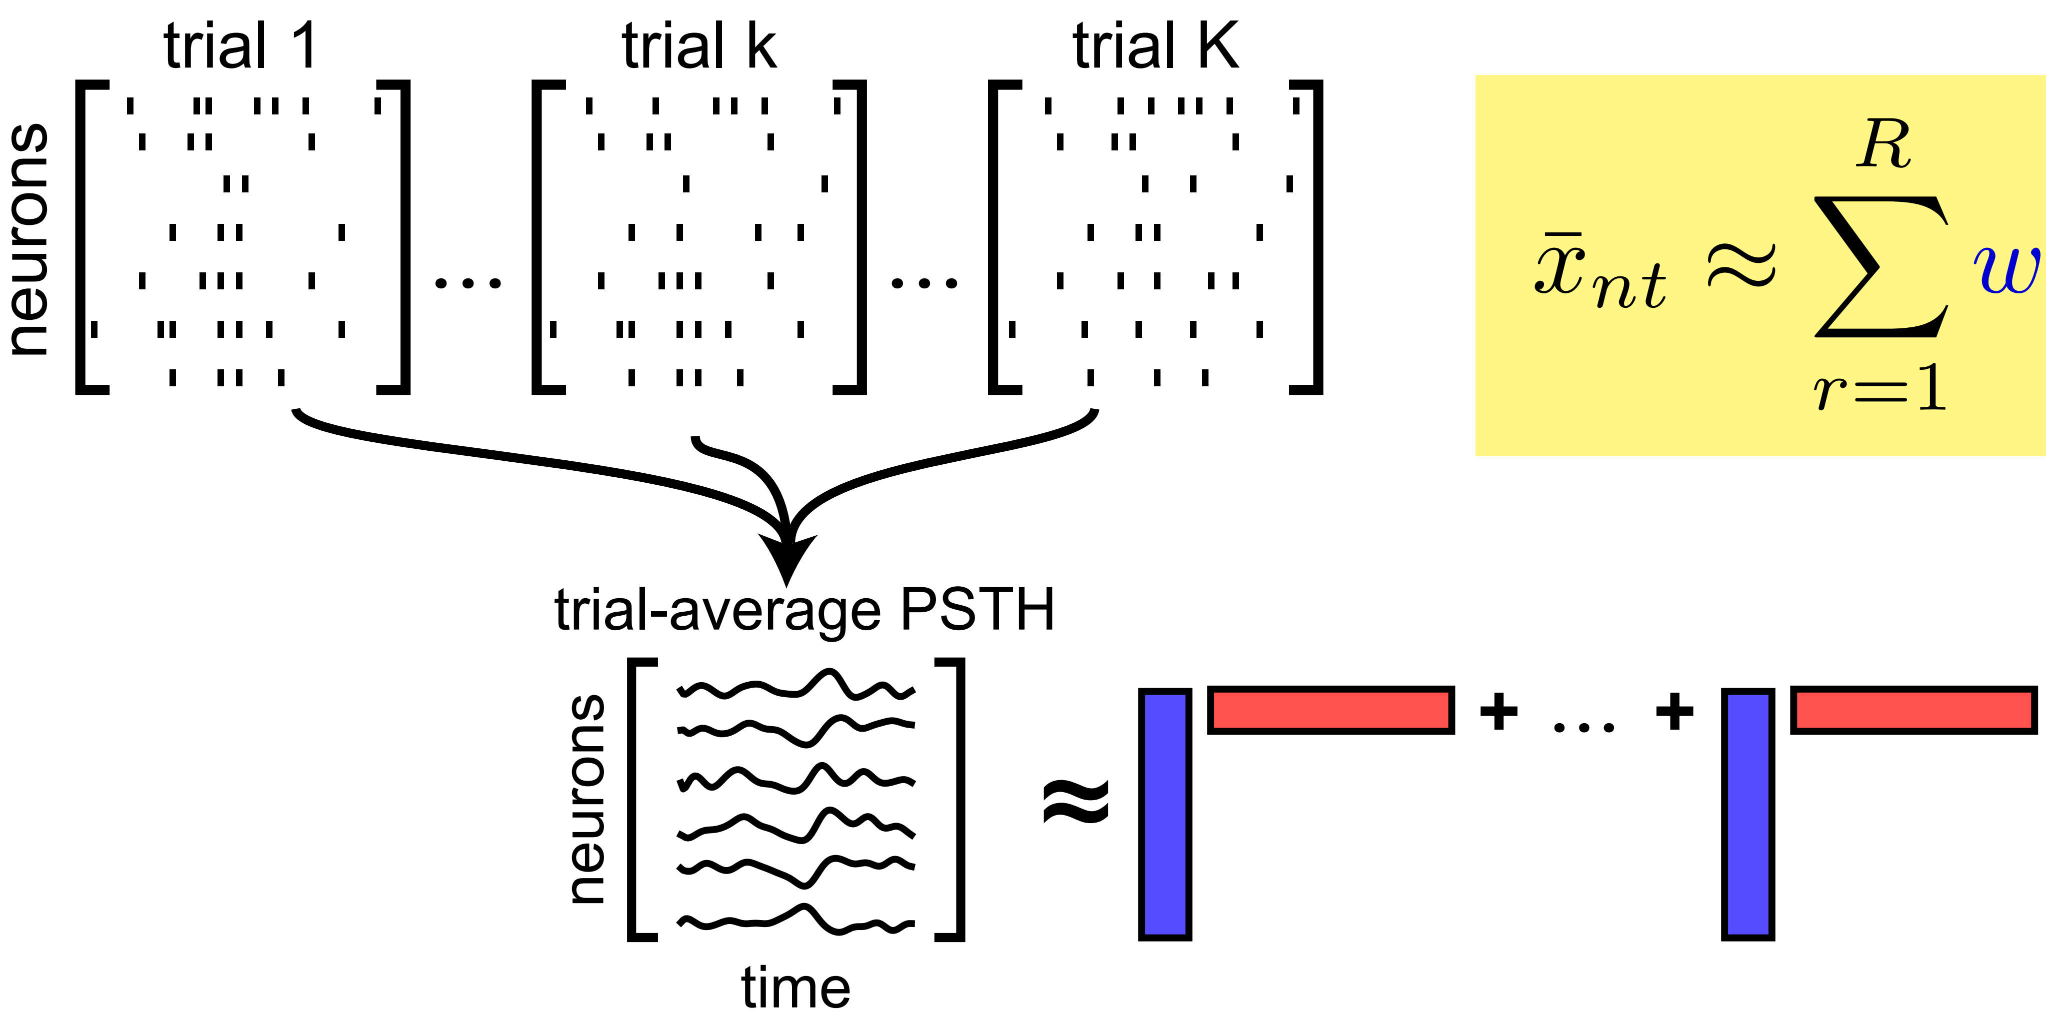

$$\bar{x}_{nt} \approx \sum_{r=1}^{R} w_n^r b_t^r$$

Existing method #1: Trial-averaged PCA

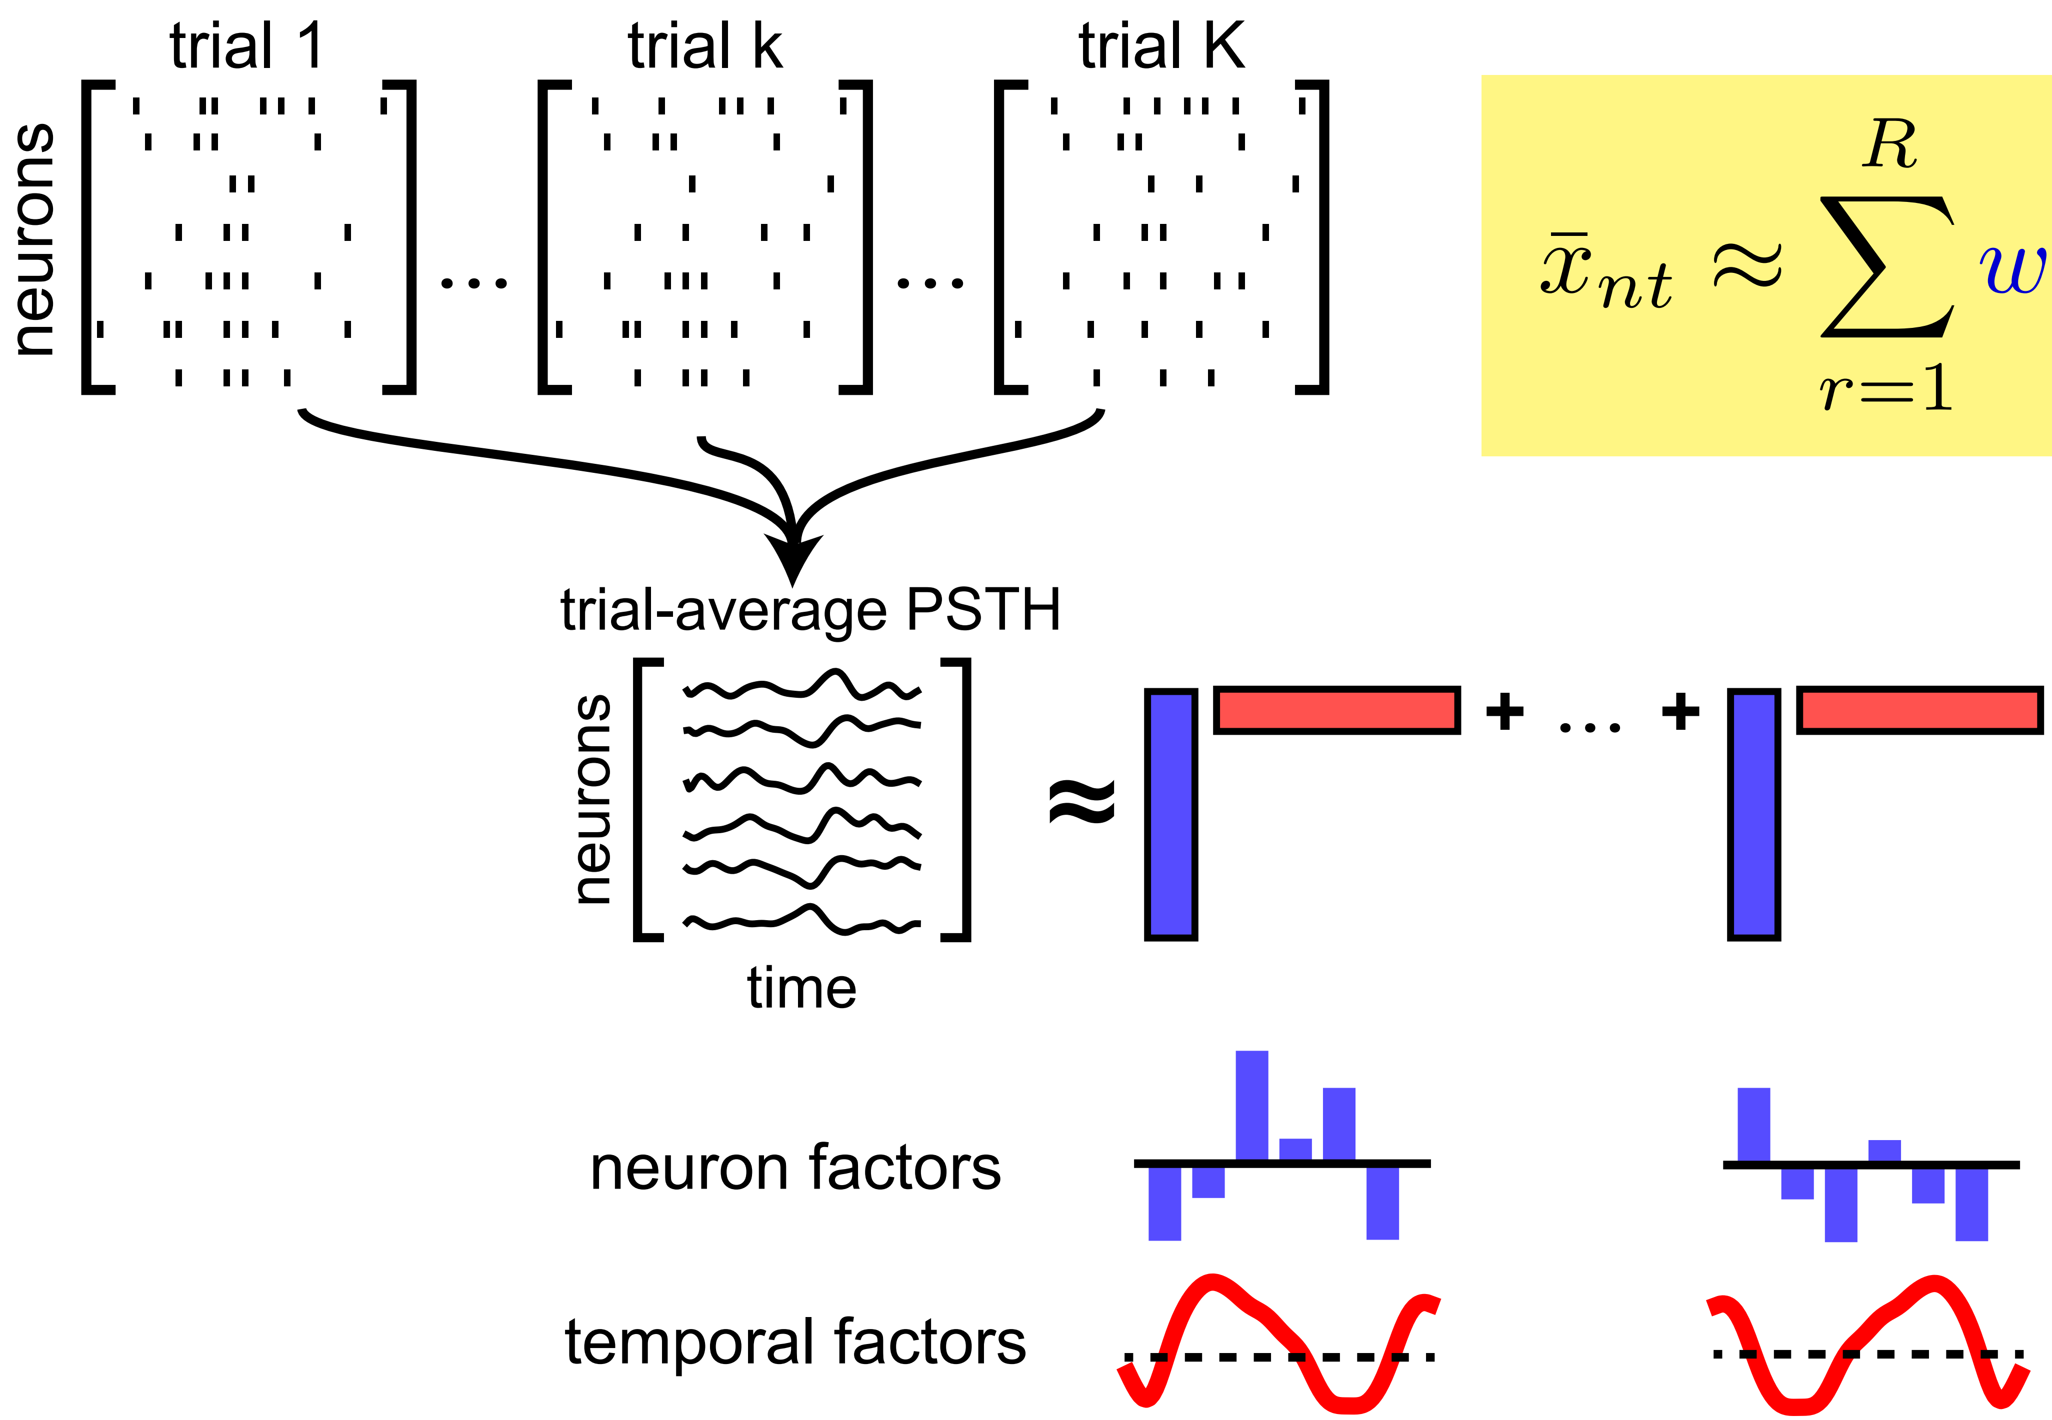

$$\bar{x}_{nt} \approx \sum_{r=1}^{R} w_n^r b_t^r$$

Existing method #2: Trial-concatenated PCA

(very similar to GPFA; Yu et al. 2009)

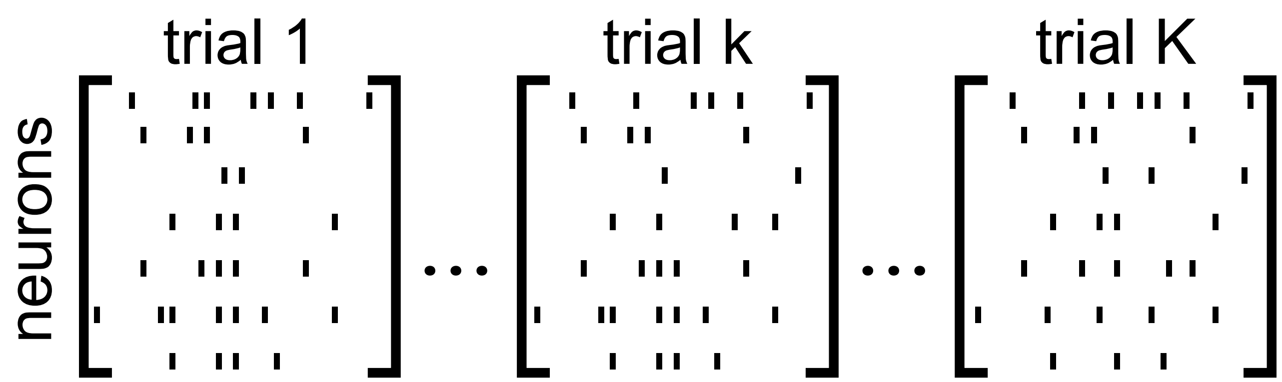

Existing method #2: Trial-concatenated PCA

(very similar to GPFA; Yu et al. 2009)

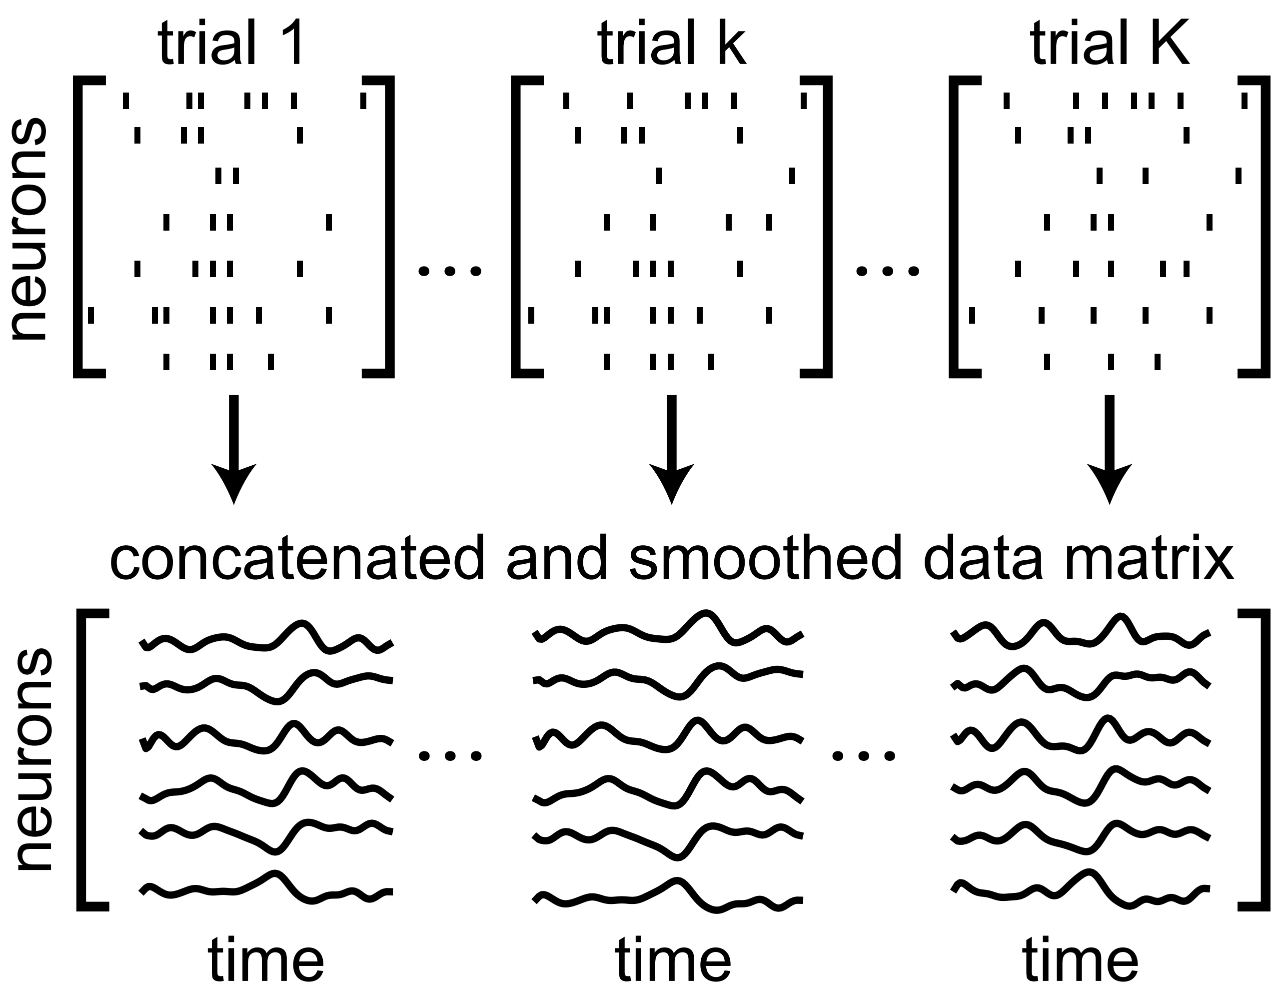

Existing method #2: Trial-concatenated PCA

(very similar to GPFA; Yu et al. 2009)

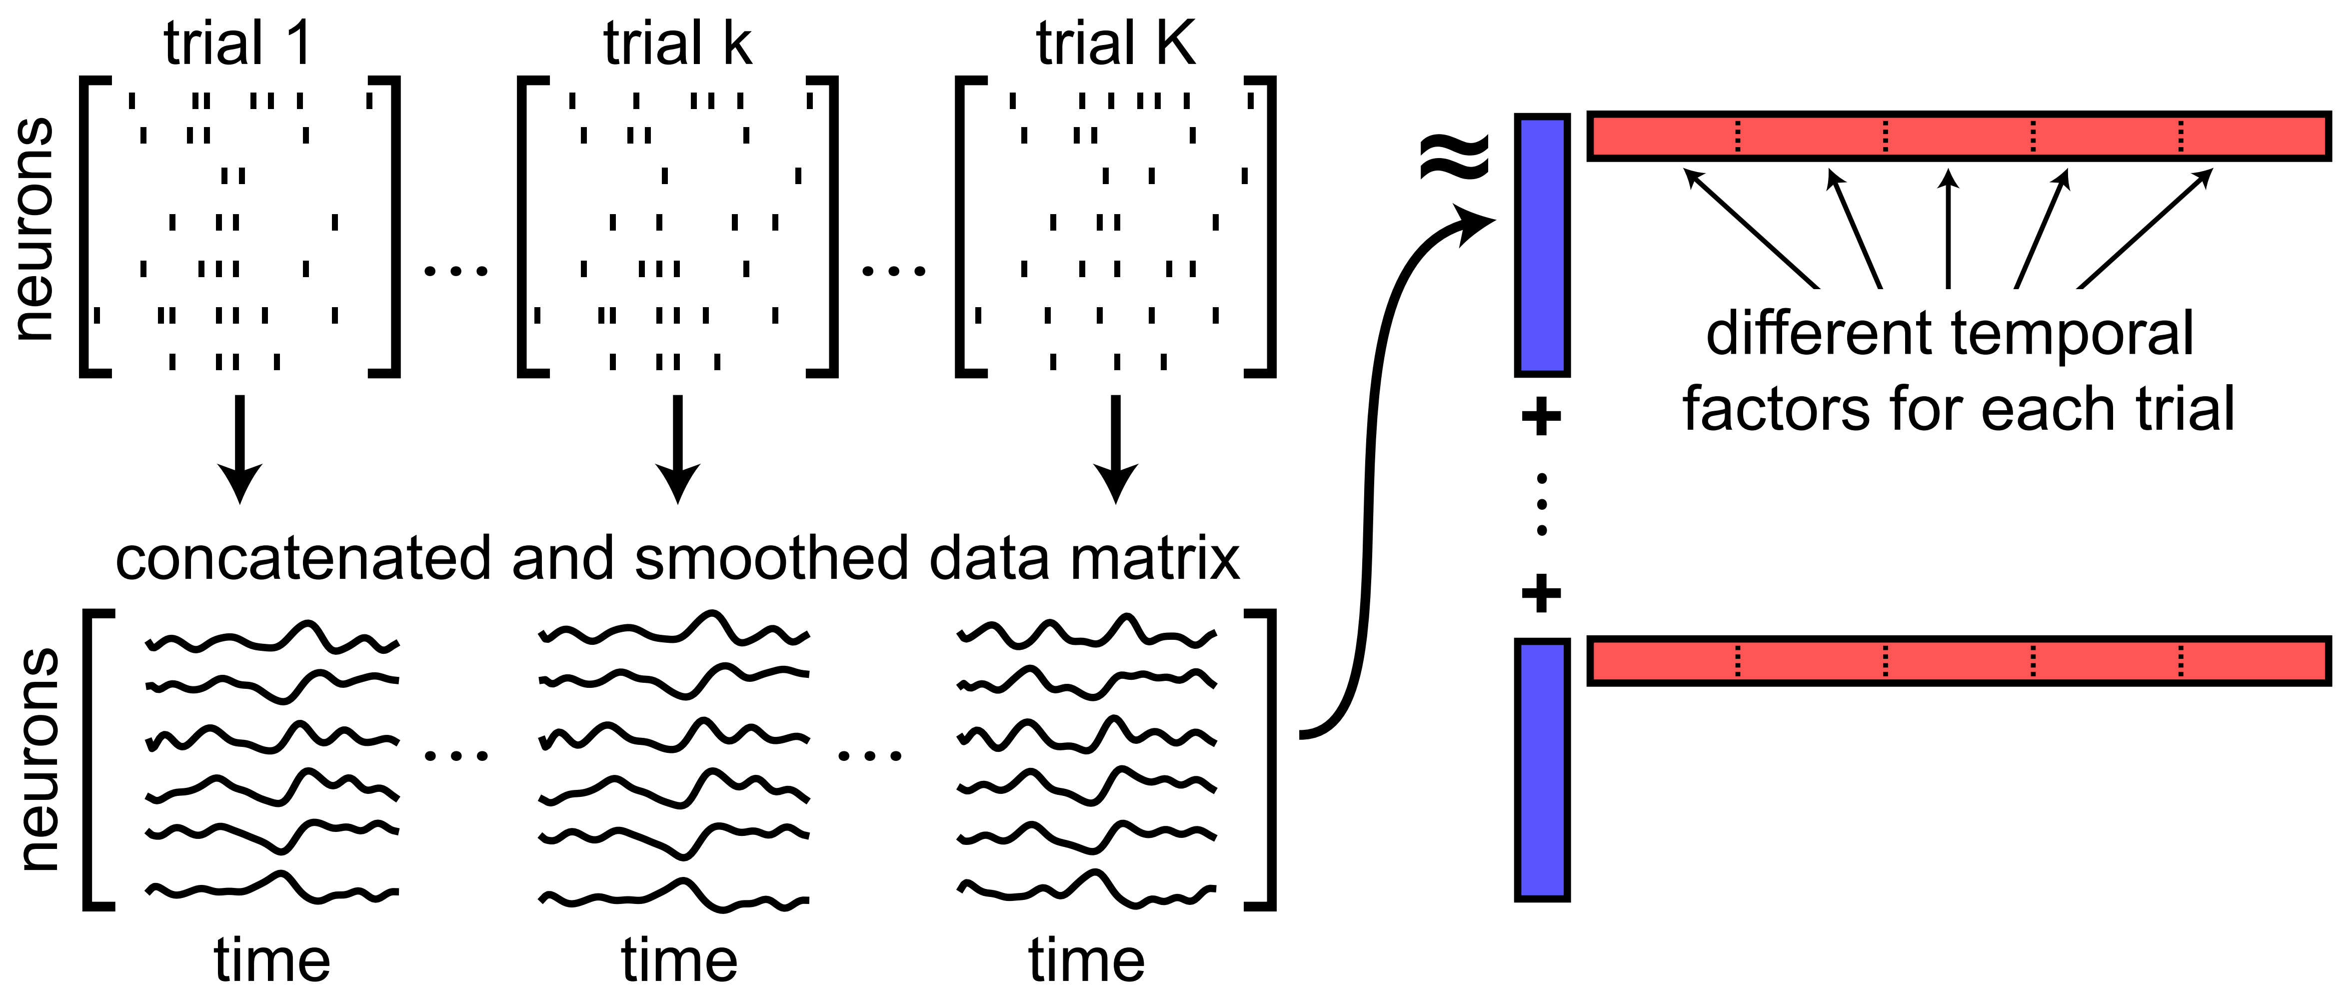

Our Goal: Find compact representation for within- and across-trial neural dynamics

neuron factors

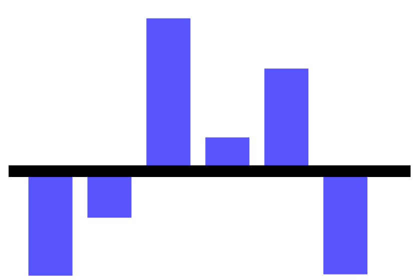

functional cell ensembles

temporal factors

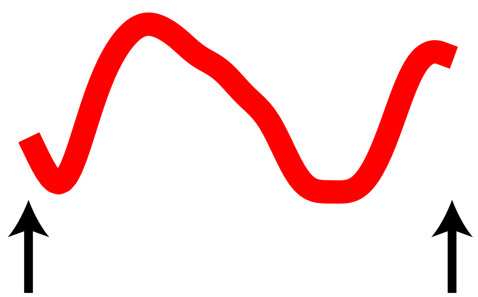

dynamics & cognition

trial factors

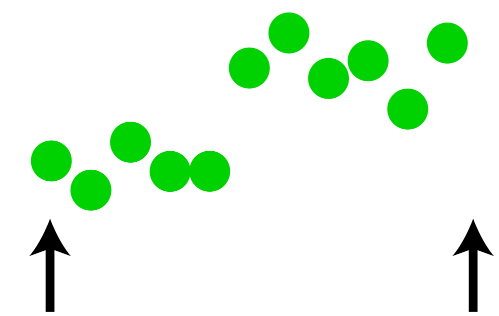

learning & stability

We apply standard tensor decomposition methods to extract these components

Tensor Components Analysis (TCA)

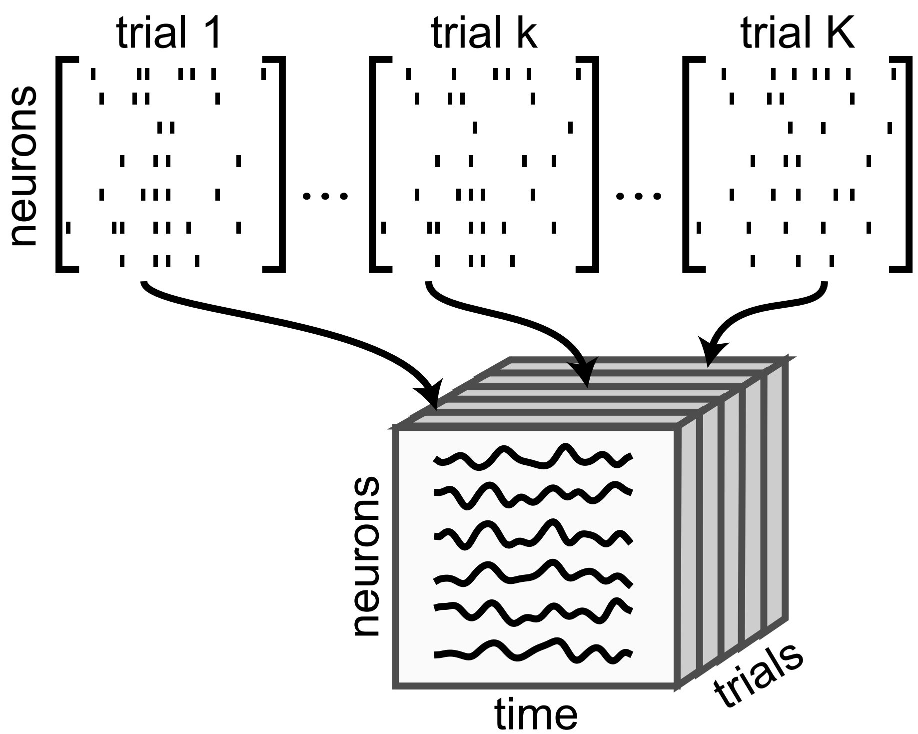

3rd-order tensor

We apply standard tensor decomposition methods to extract these components

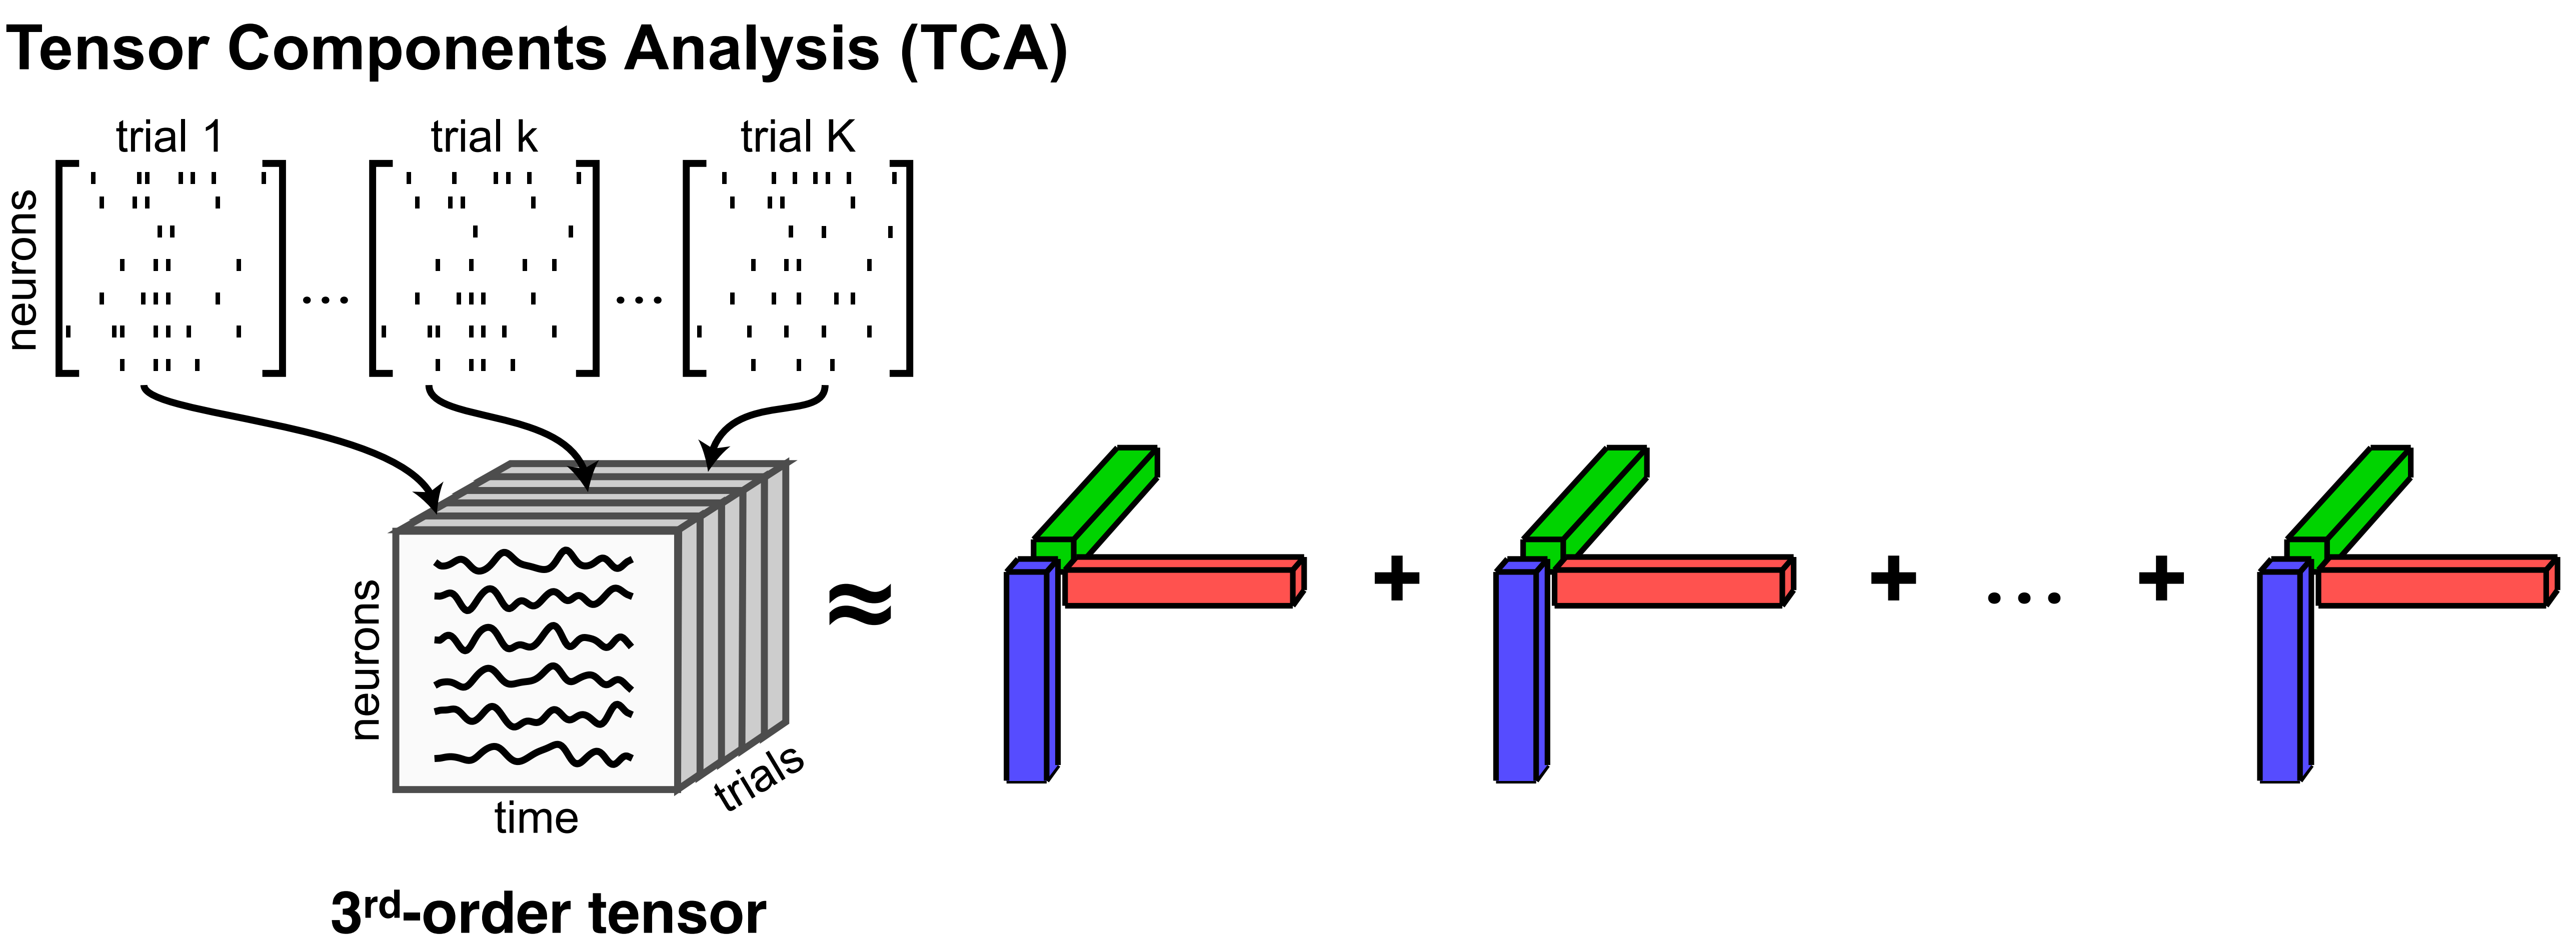

We apply standard tensor decomposition methods to extract these components

Tensor Components Analysis (TCA)

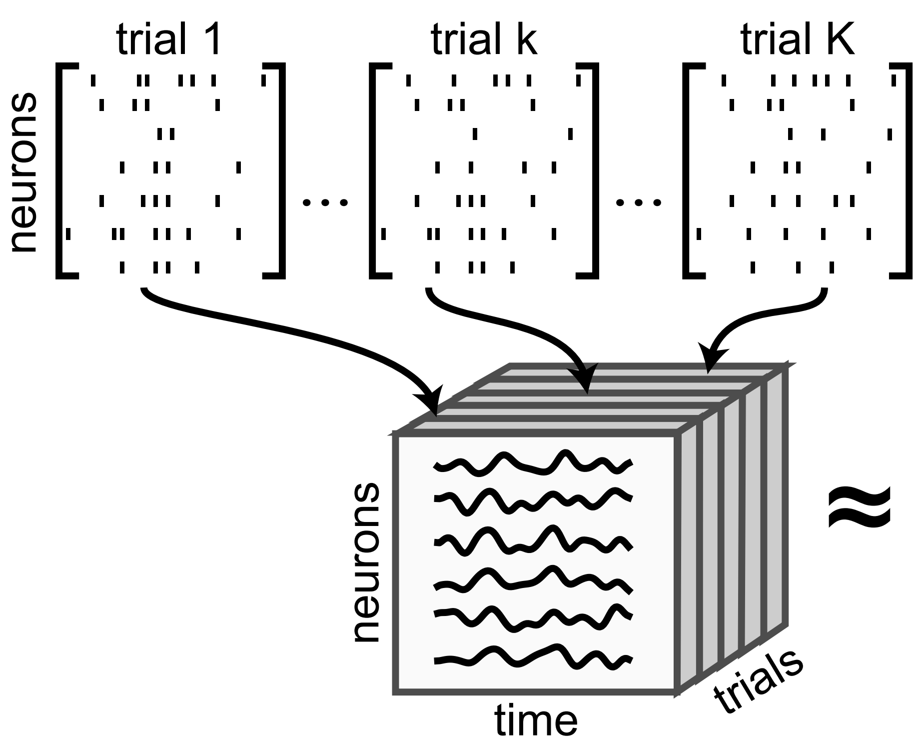

3rd-order tensor

$$x_{ntk} \approx \sum_{r=1}^{R} w_n^r b_t^r a_k^r$$

CANDECOMP/PARAFAC Decomposition

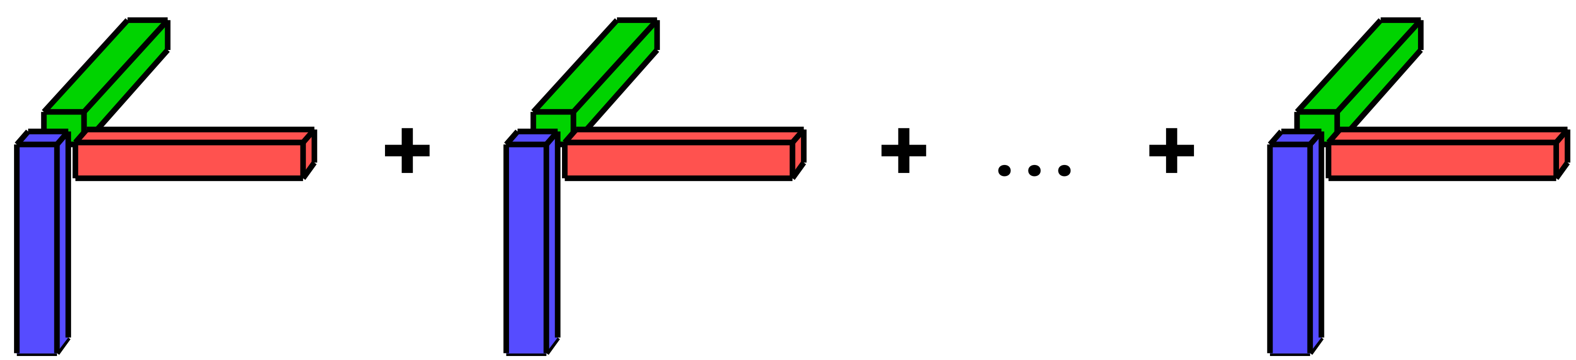

We apply standard tensor decomposition methods to extract these components

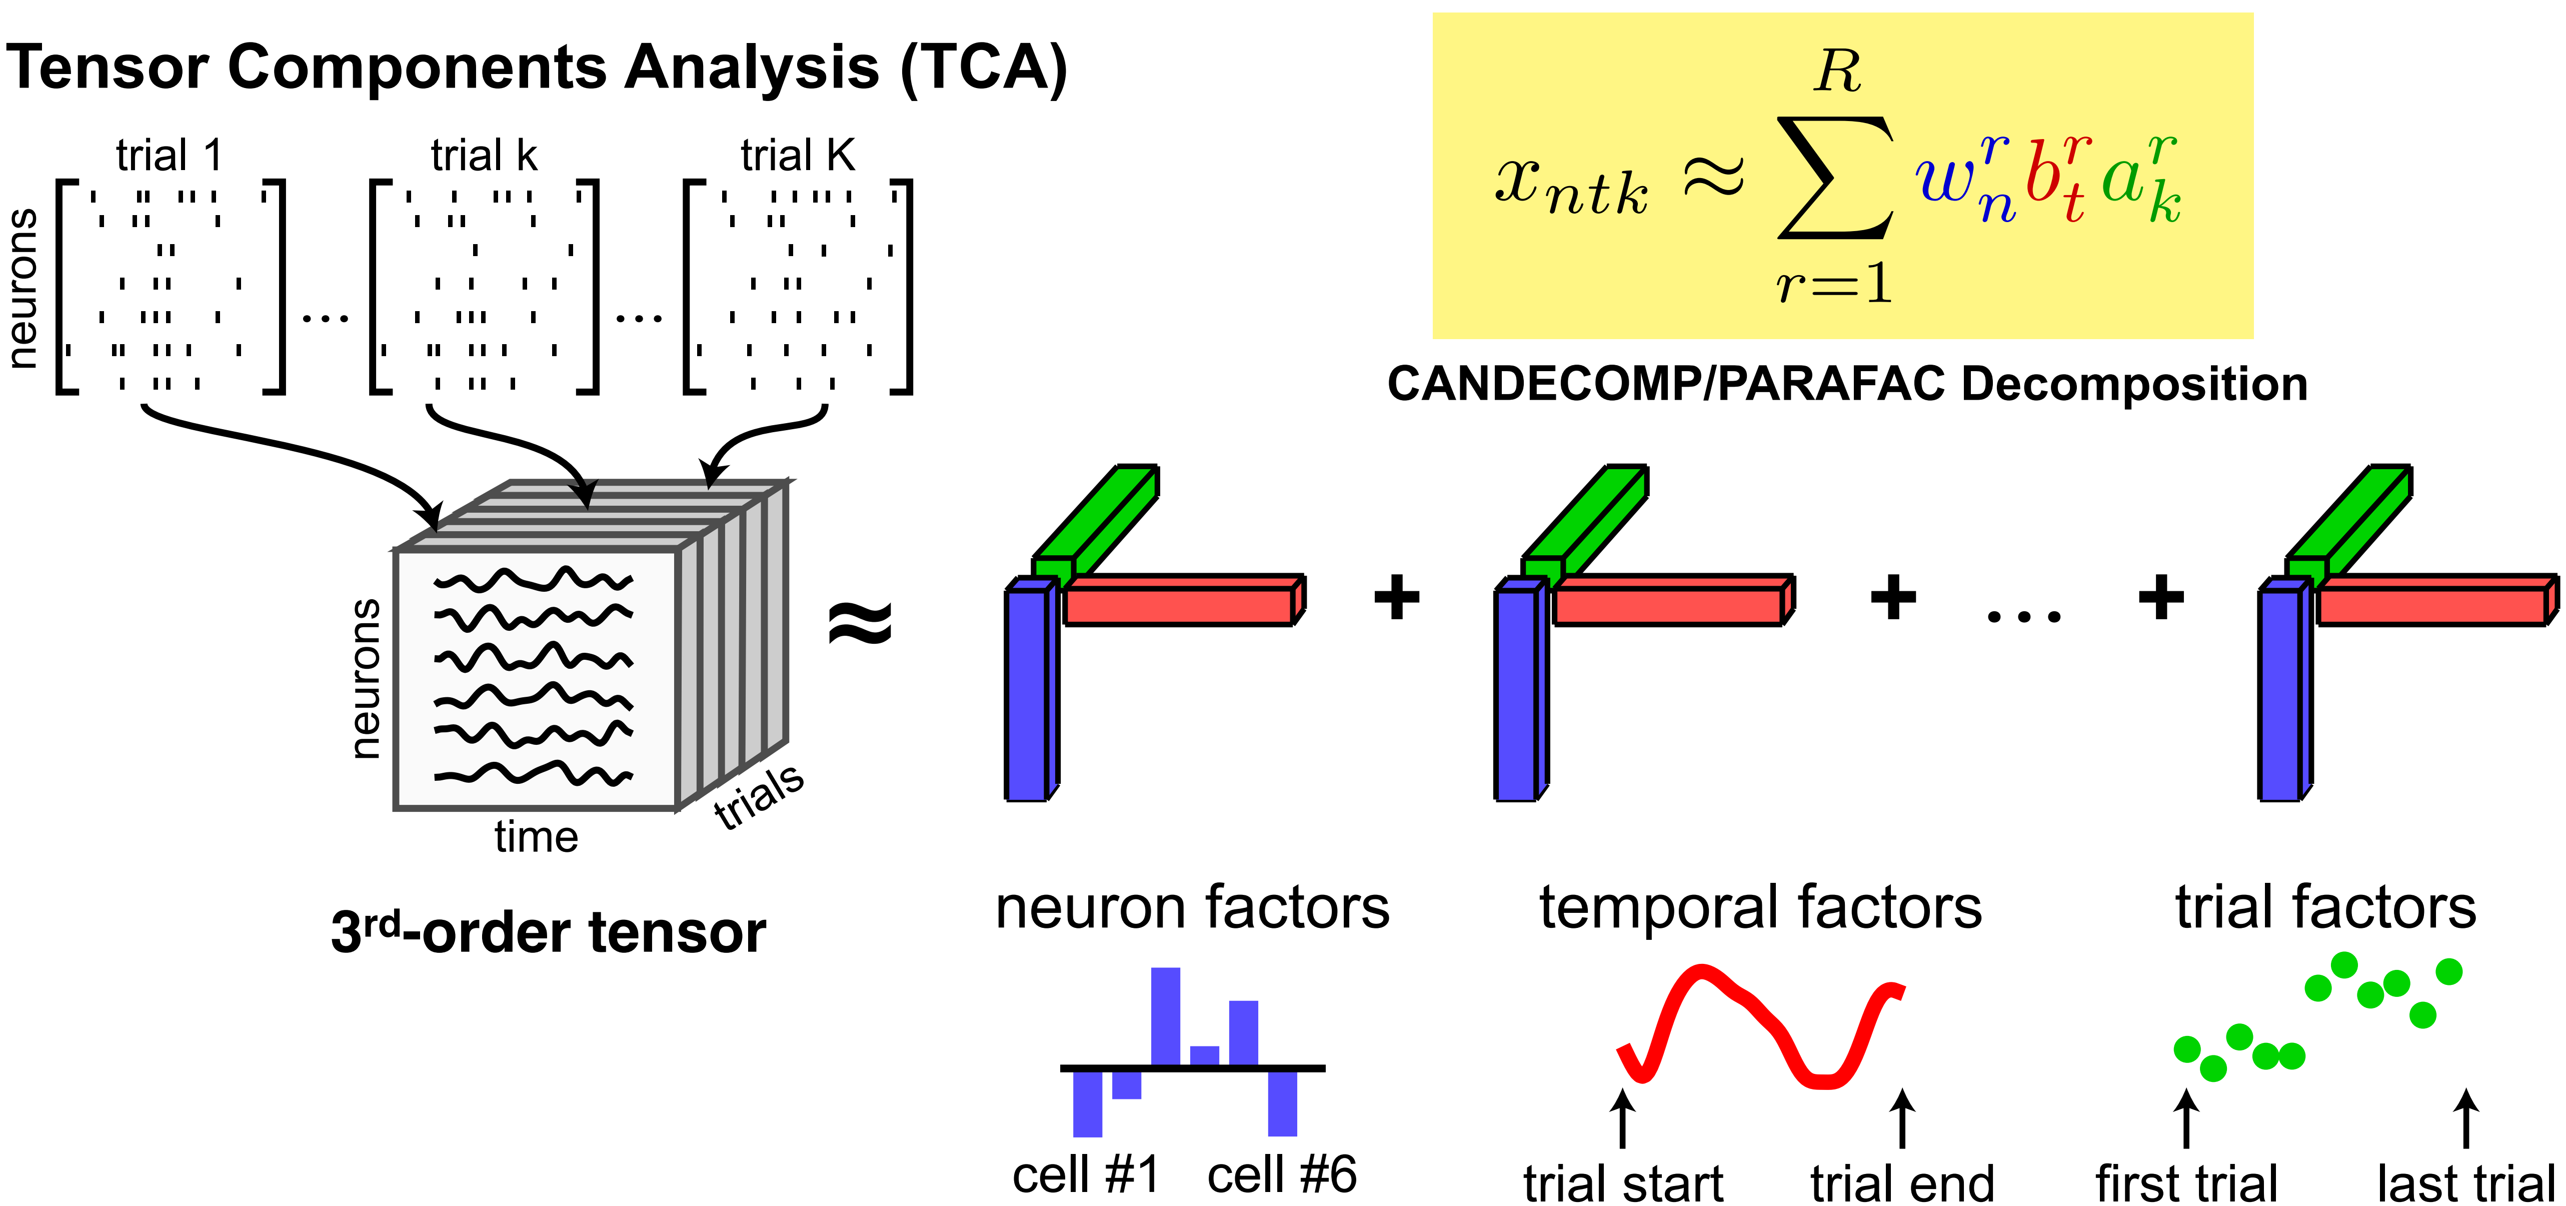

A neural circuit interpretation of TCA:

TCA is a linear network with gain modulation — an influential principle of cortical computation (e.g., Carandini and Heeger, 2012)

A neural circuit interpretation of TCA:

TCA is a linear network with **gain modulation** — an influential principle of cortical computation (e.g., Carandini and Heeger, 2012)

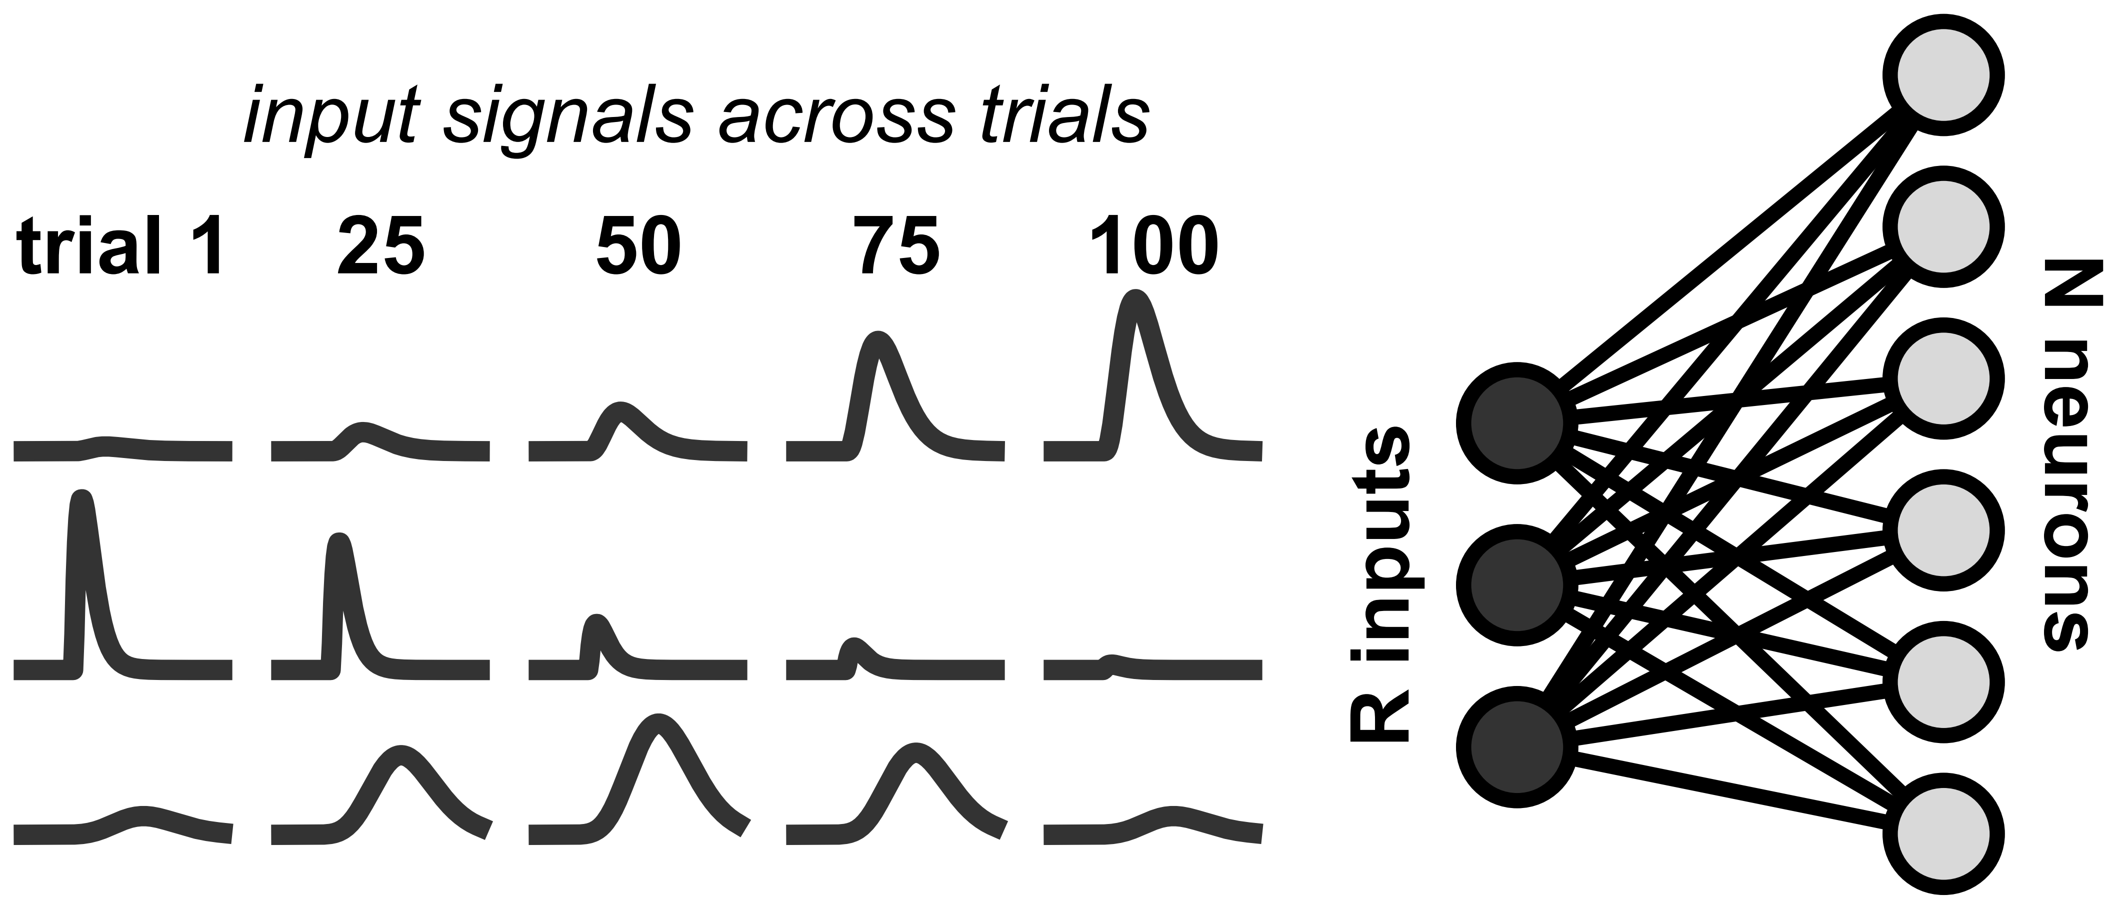

A neural circuit interpretation of TCA:

TCA is a linear network with **gain modulation** — an influential principle of cortical computation (e.g., Carandini and Heeger, 2012)

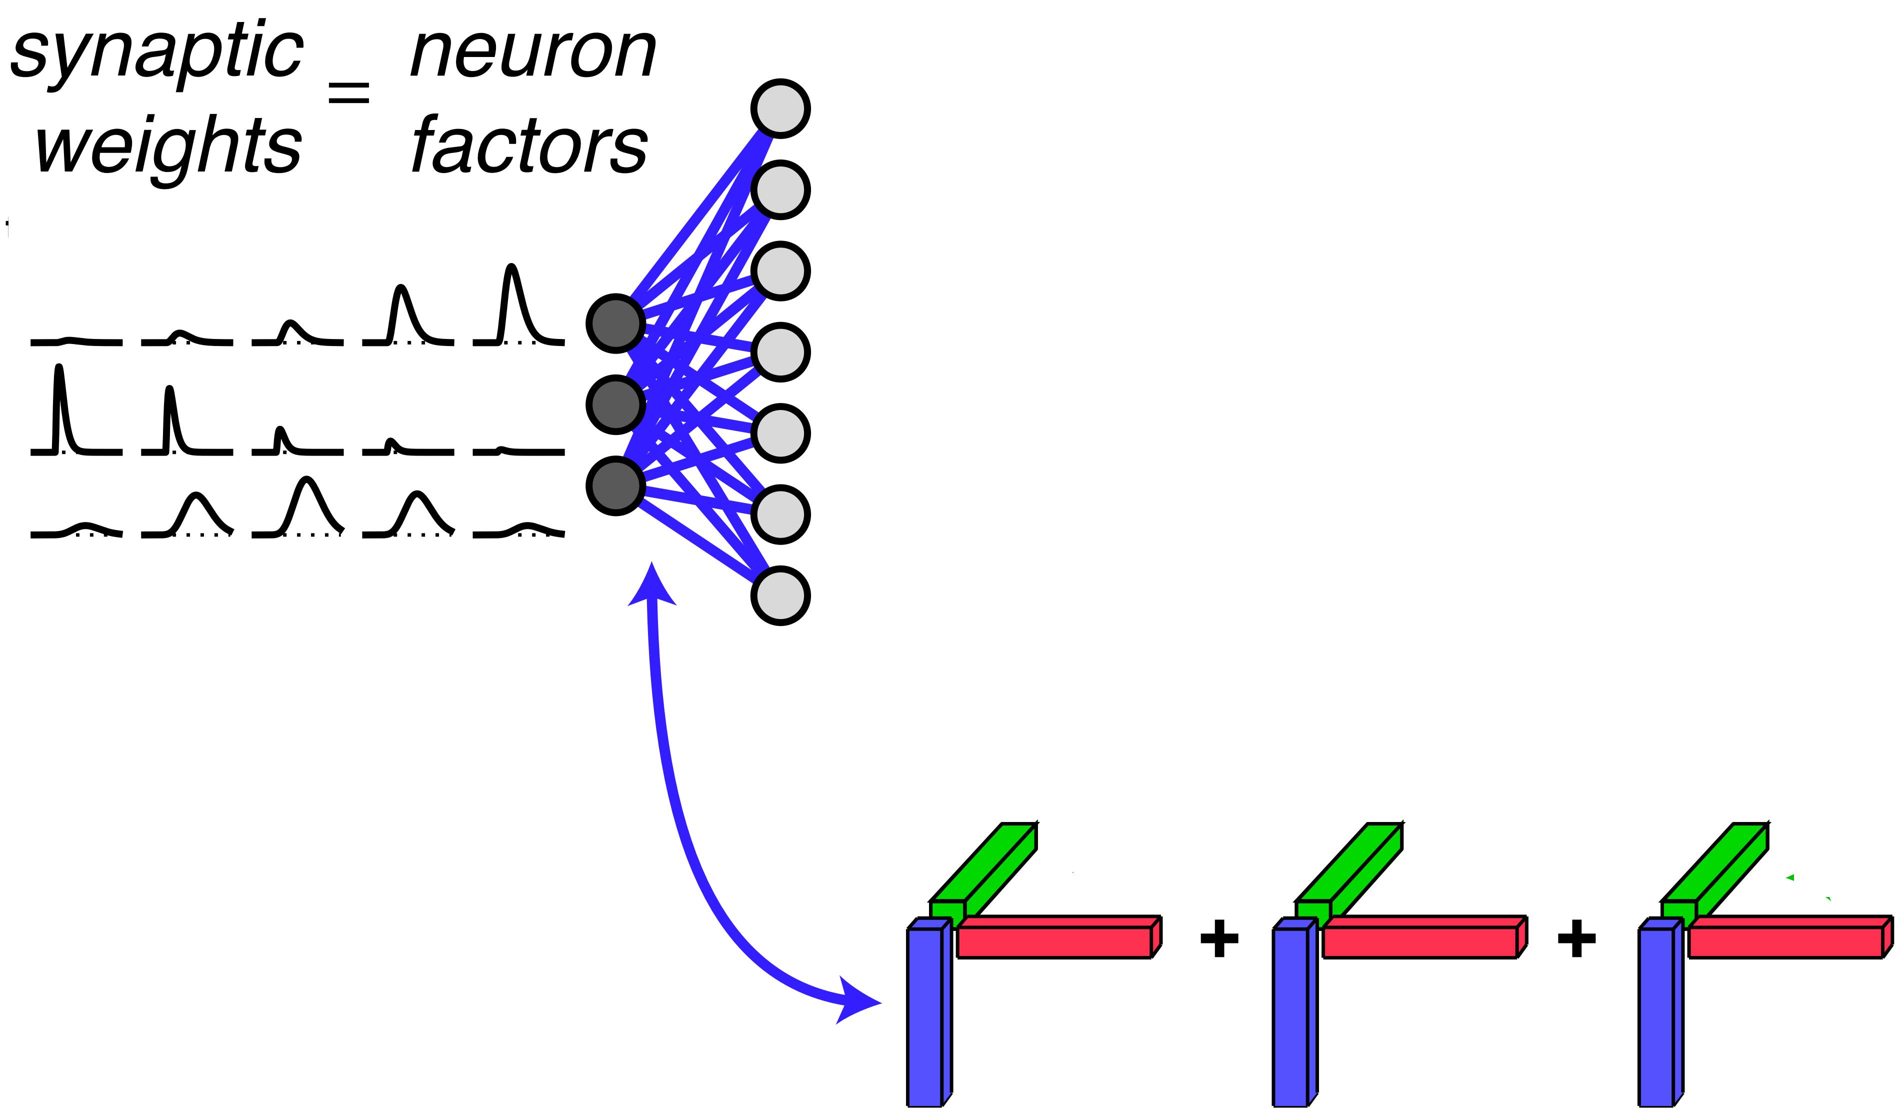

A neural circuit interpretation of TCA:

TCA is a linear network with **gain modulation** — an influential principle of cortical computation (e.g., Carandini and Heeger, 2012)

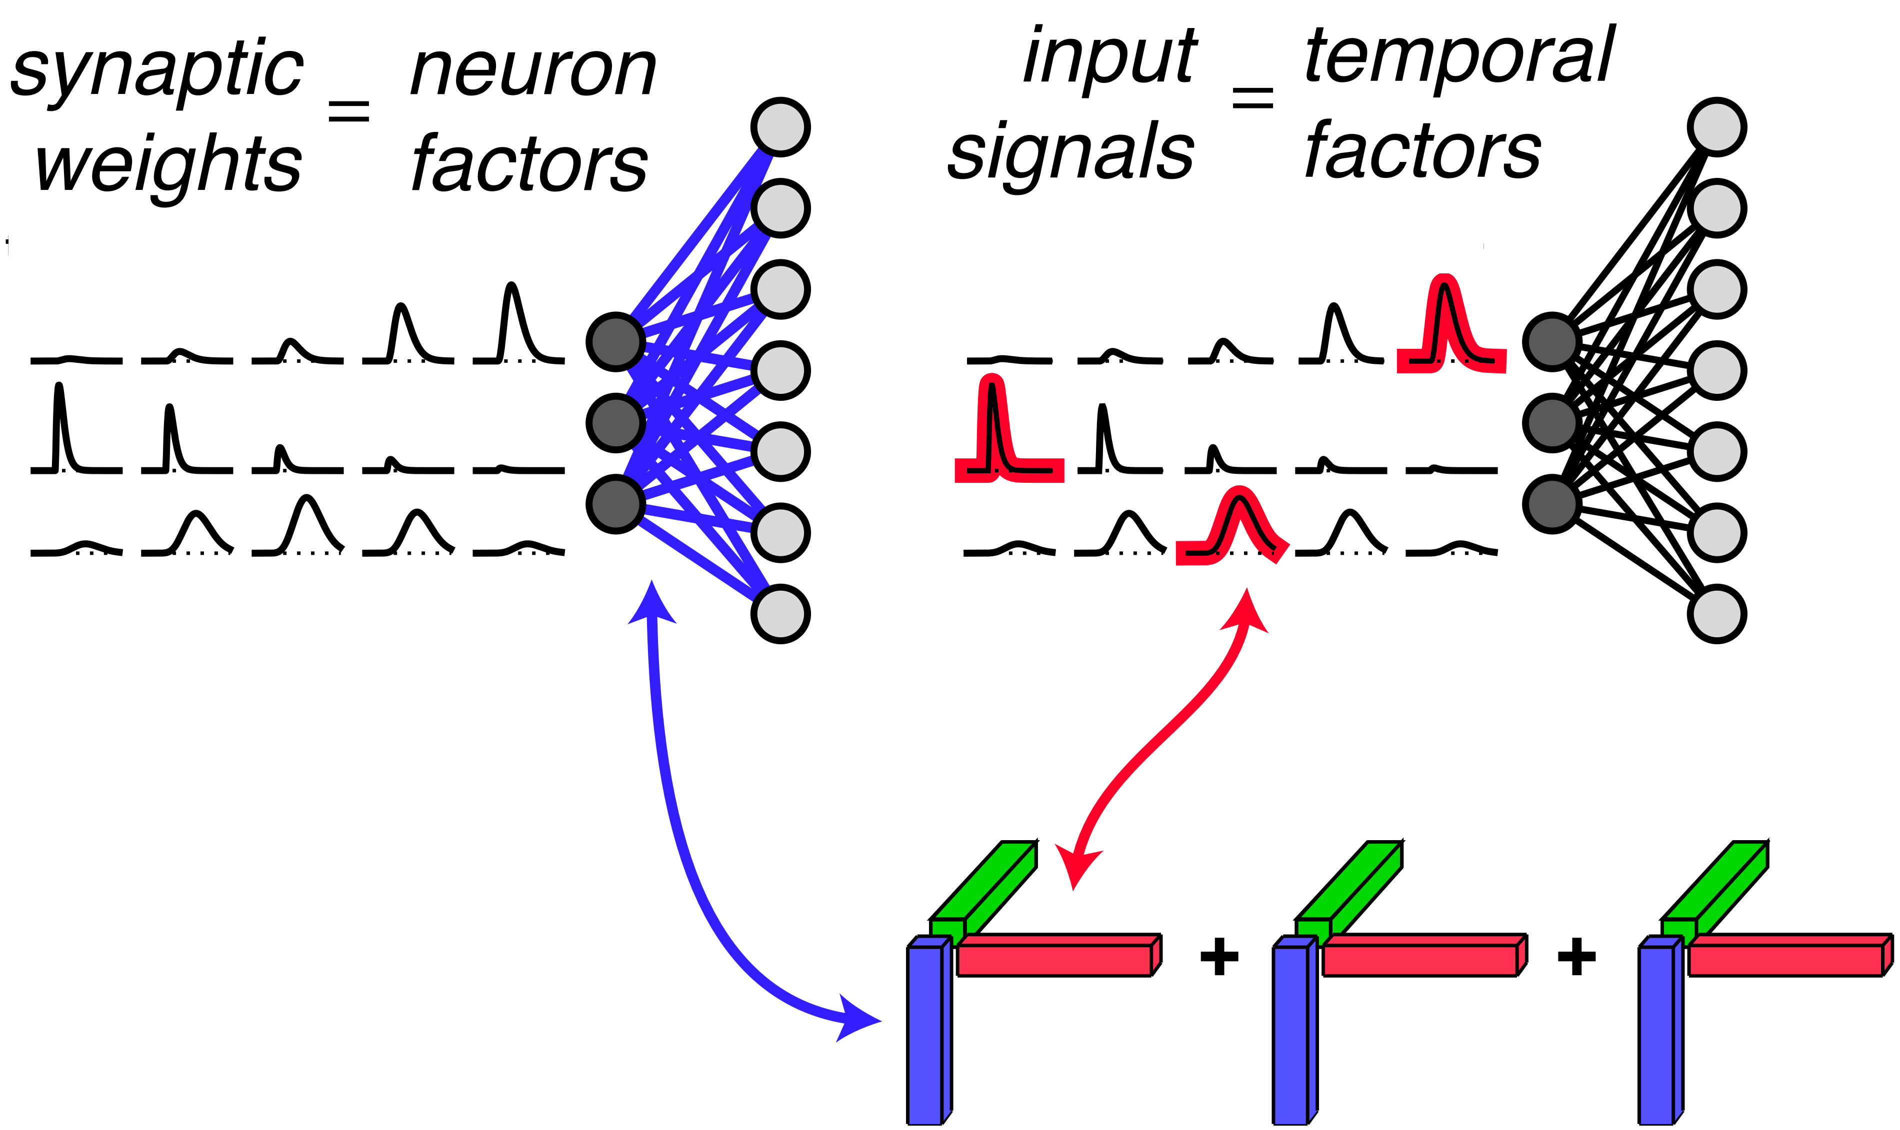

A neural circuit interpretation of TCA:

TCA is a linear network with **gain modulation** — an influential principle of cortical computation (e.g., Carandini and Heeger, 2012)

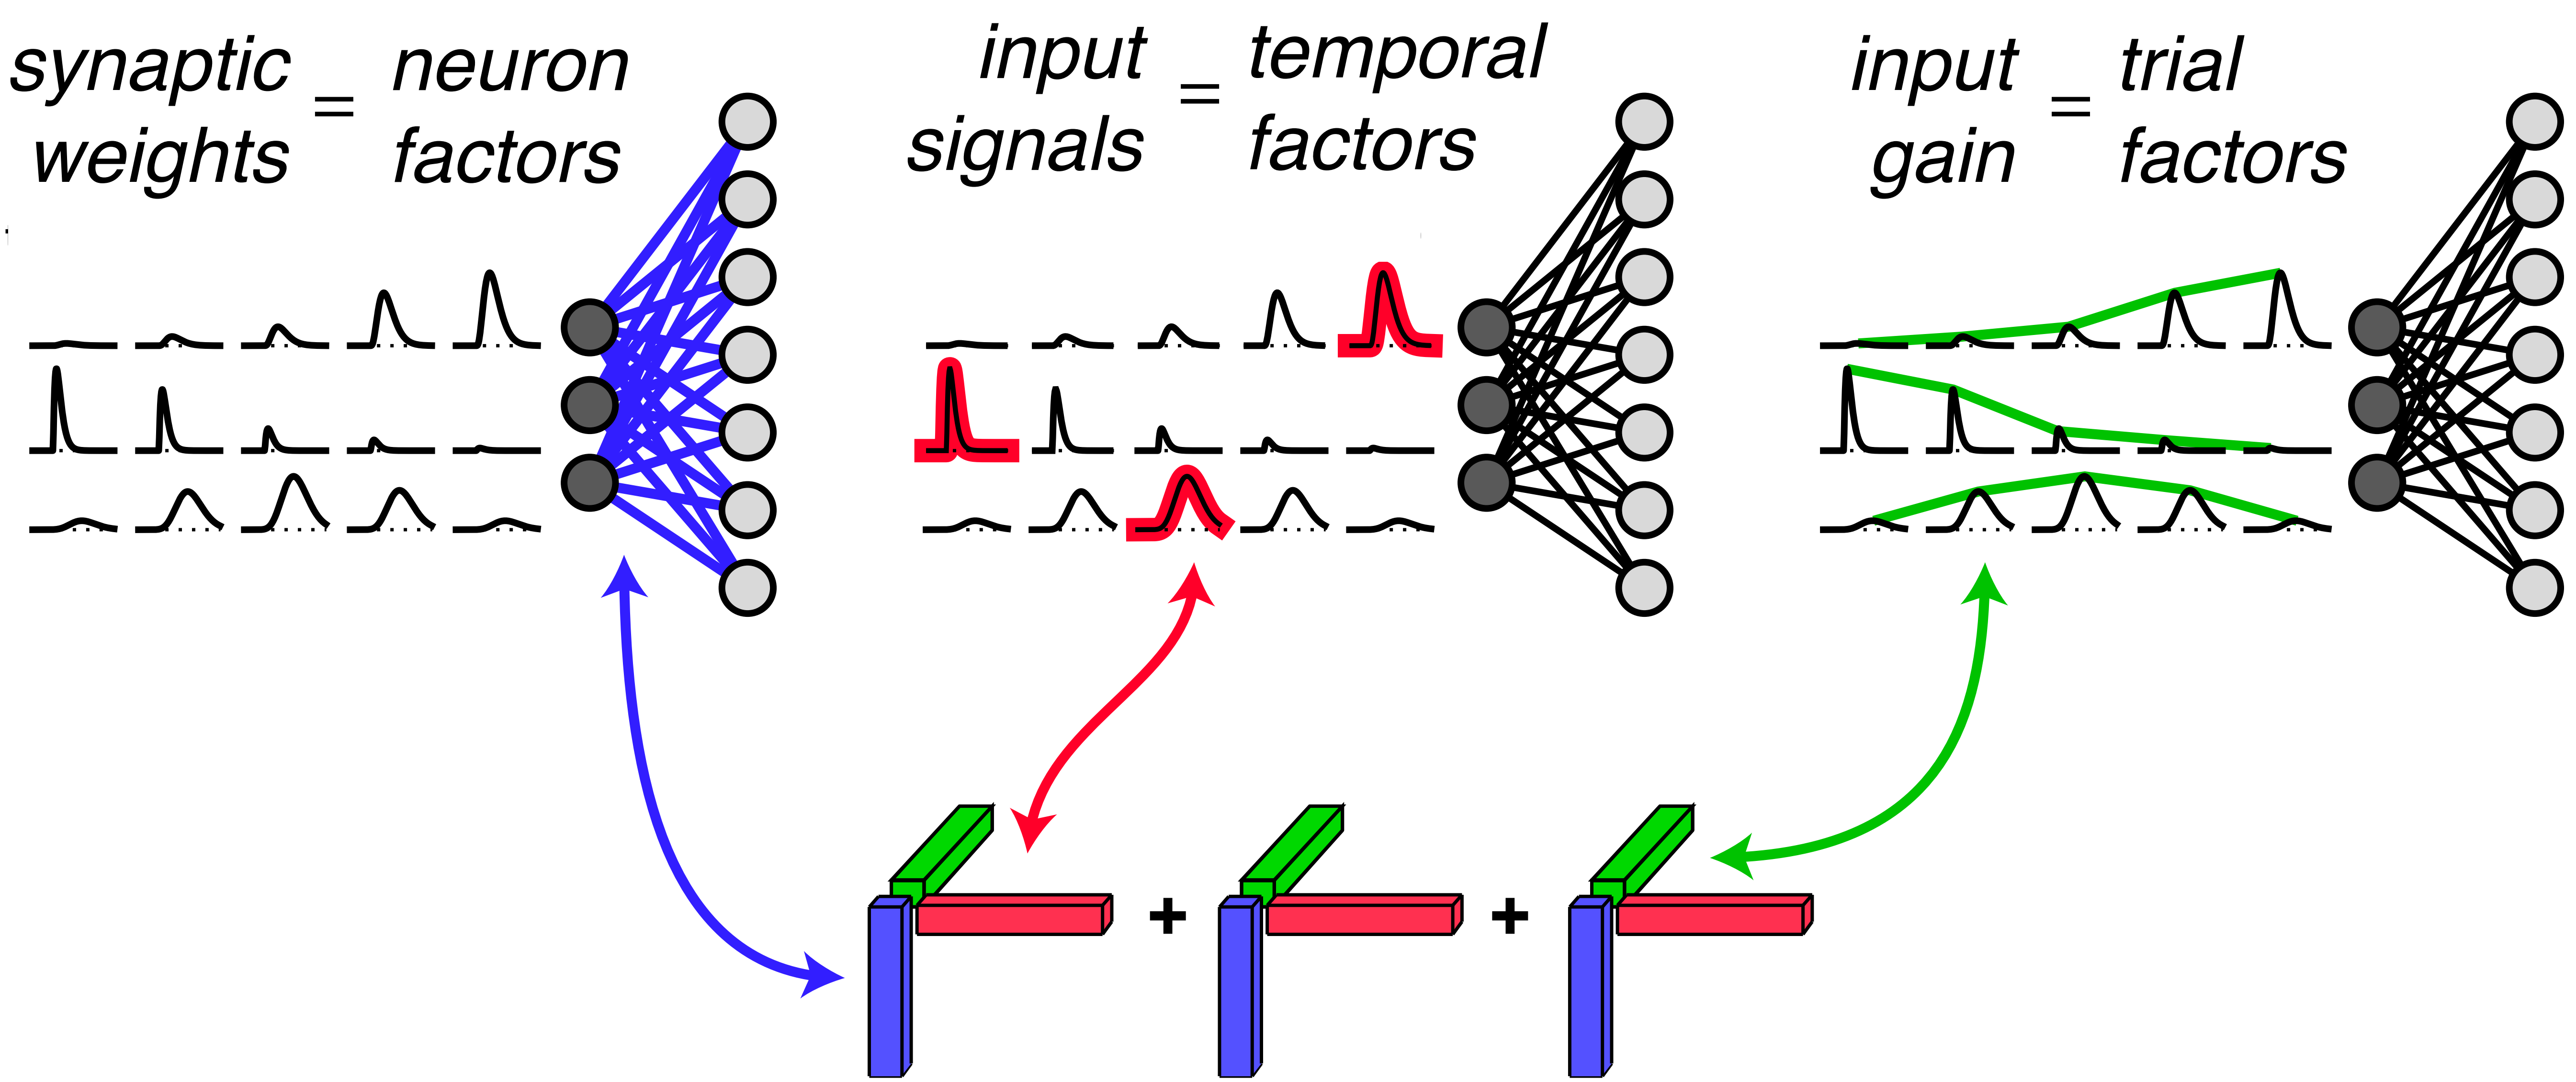

A demonstration of TCA using simulated data from the gain-modulated linear network

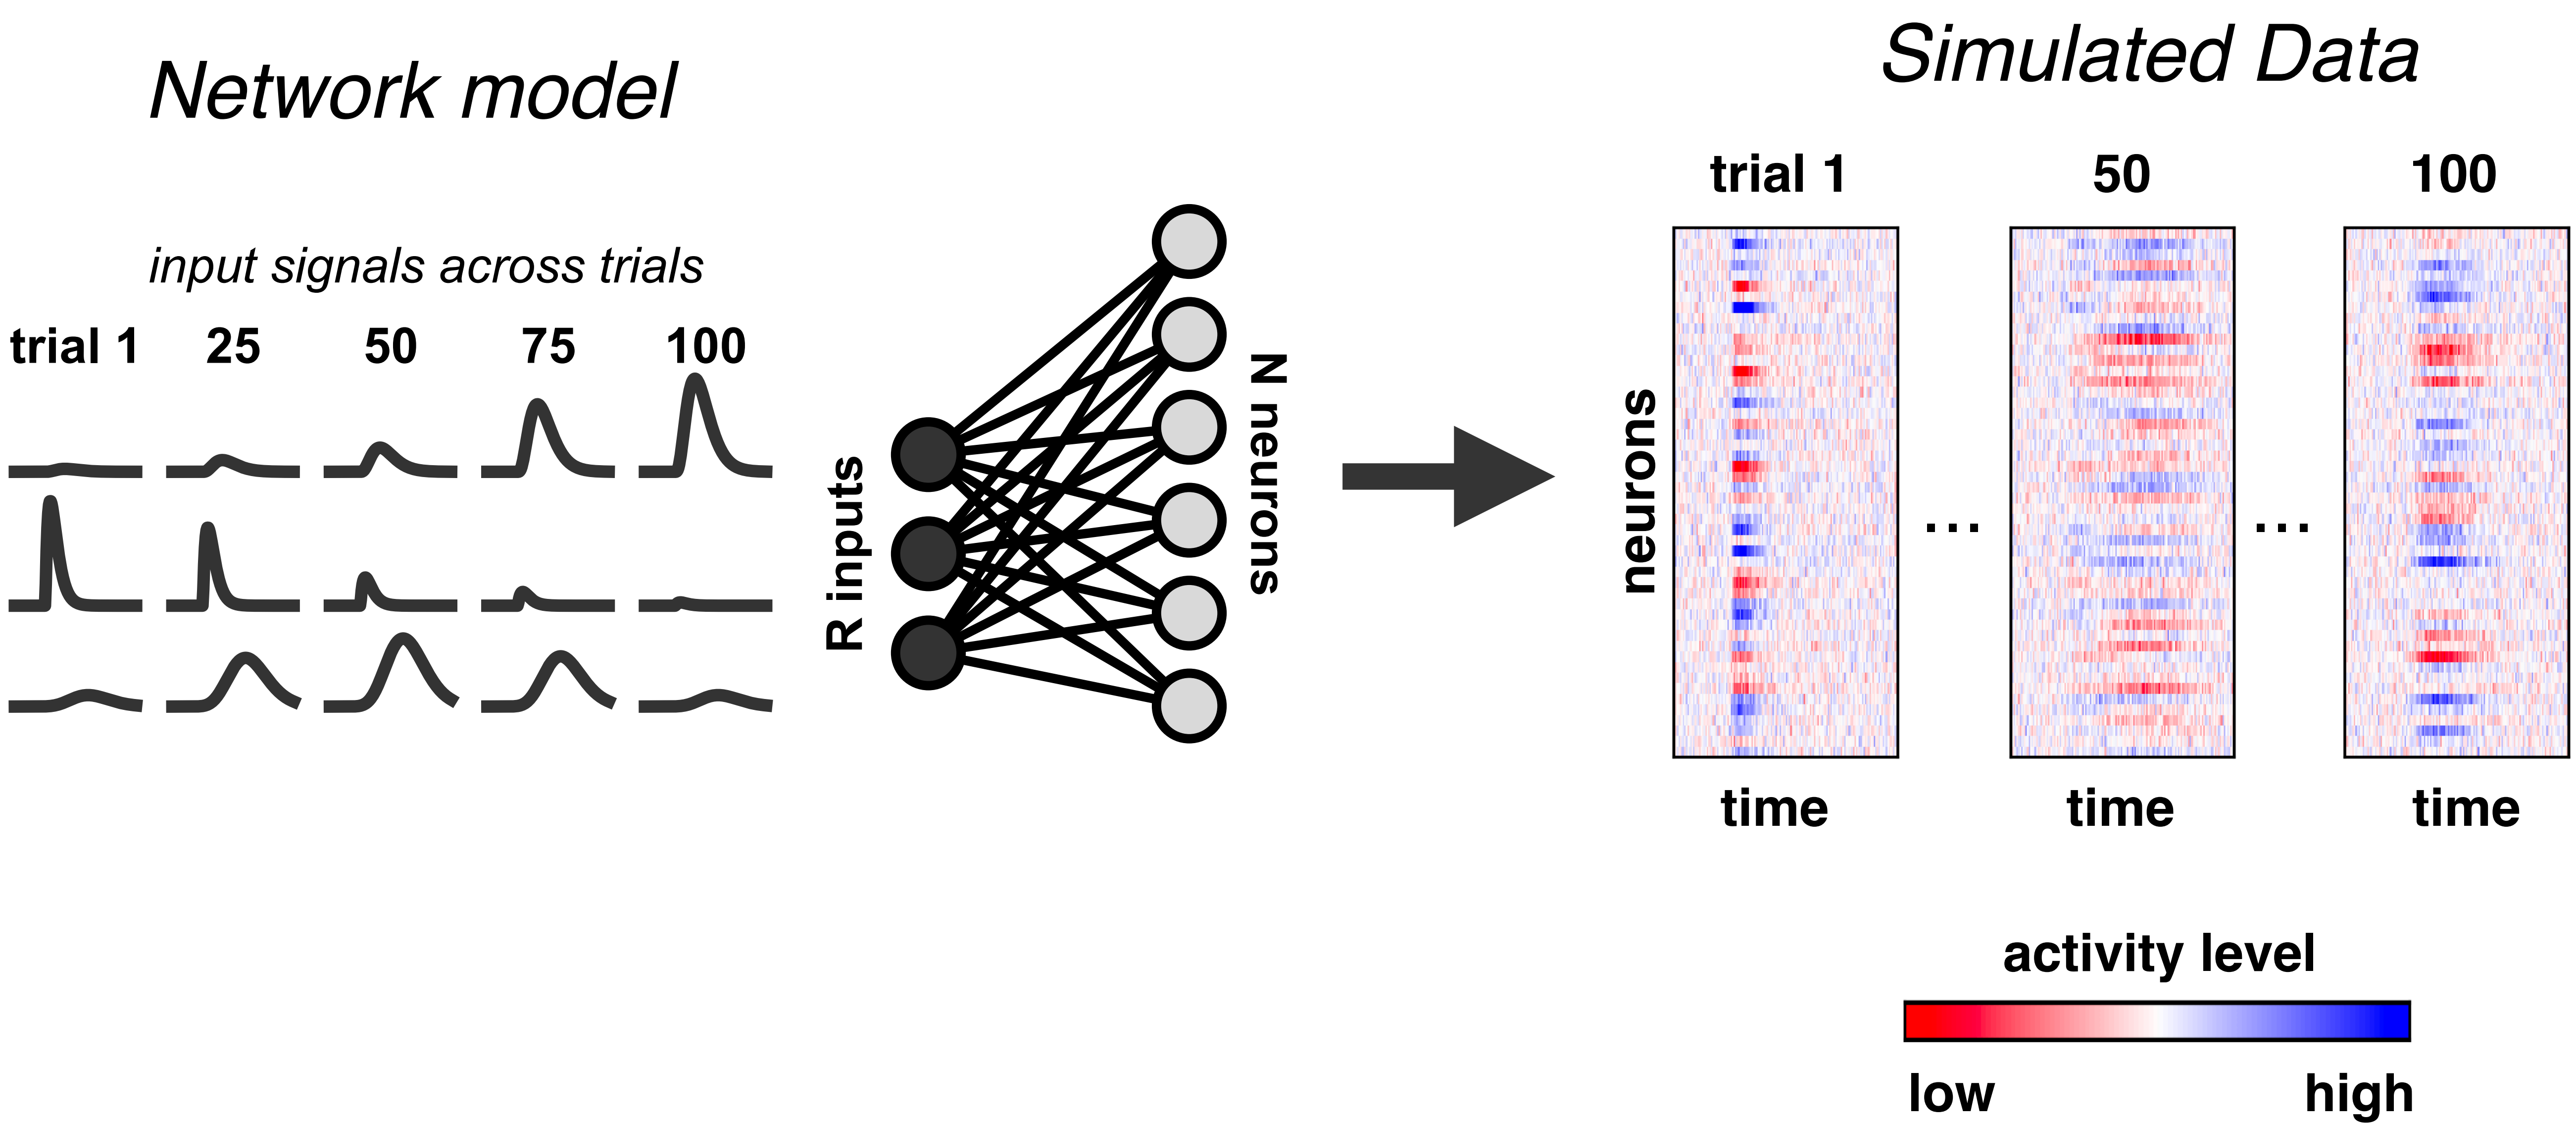

PCA fails to recover network parameters from simulated data

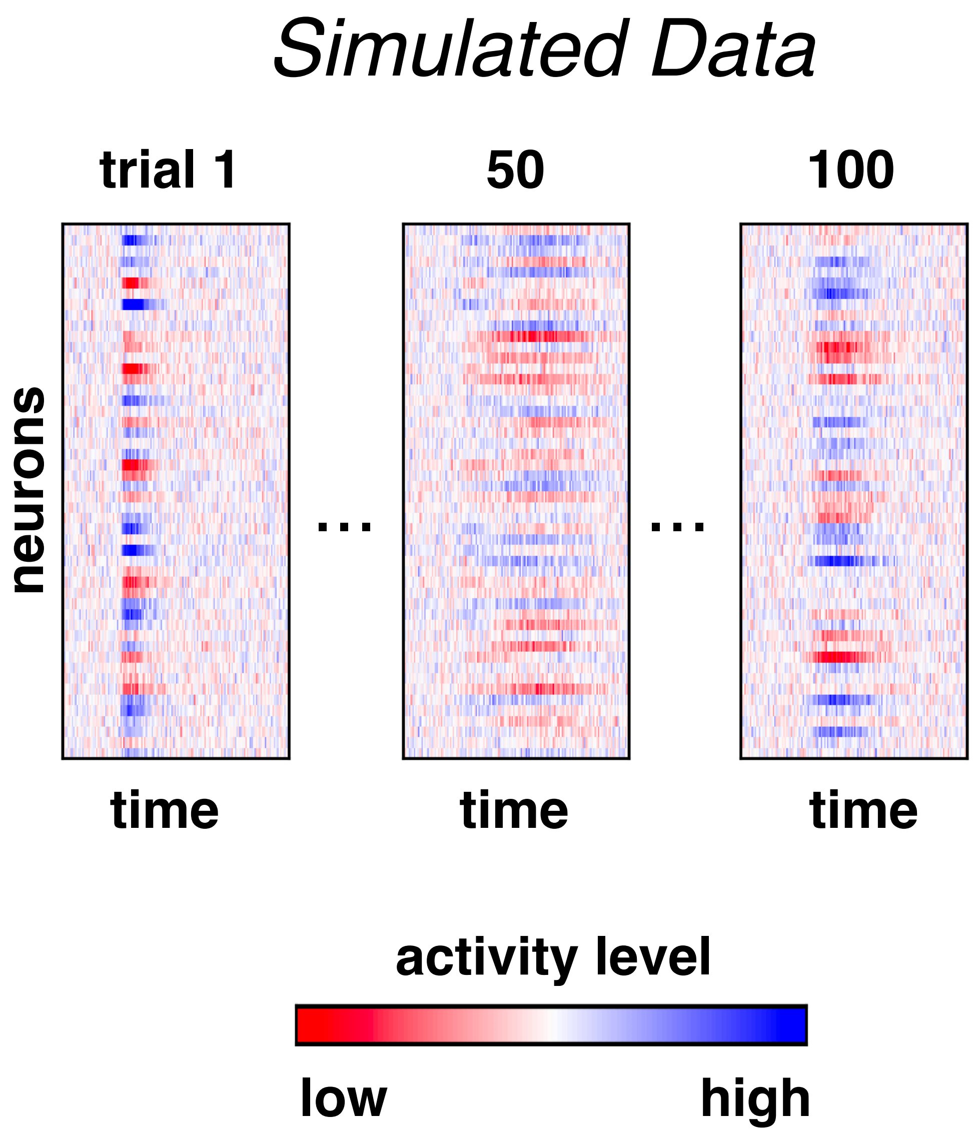

PCA fails to recover network parameters from simulated data

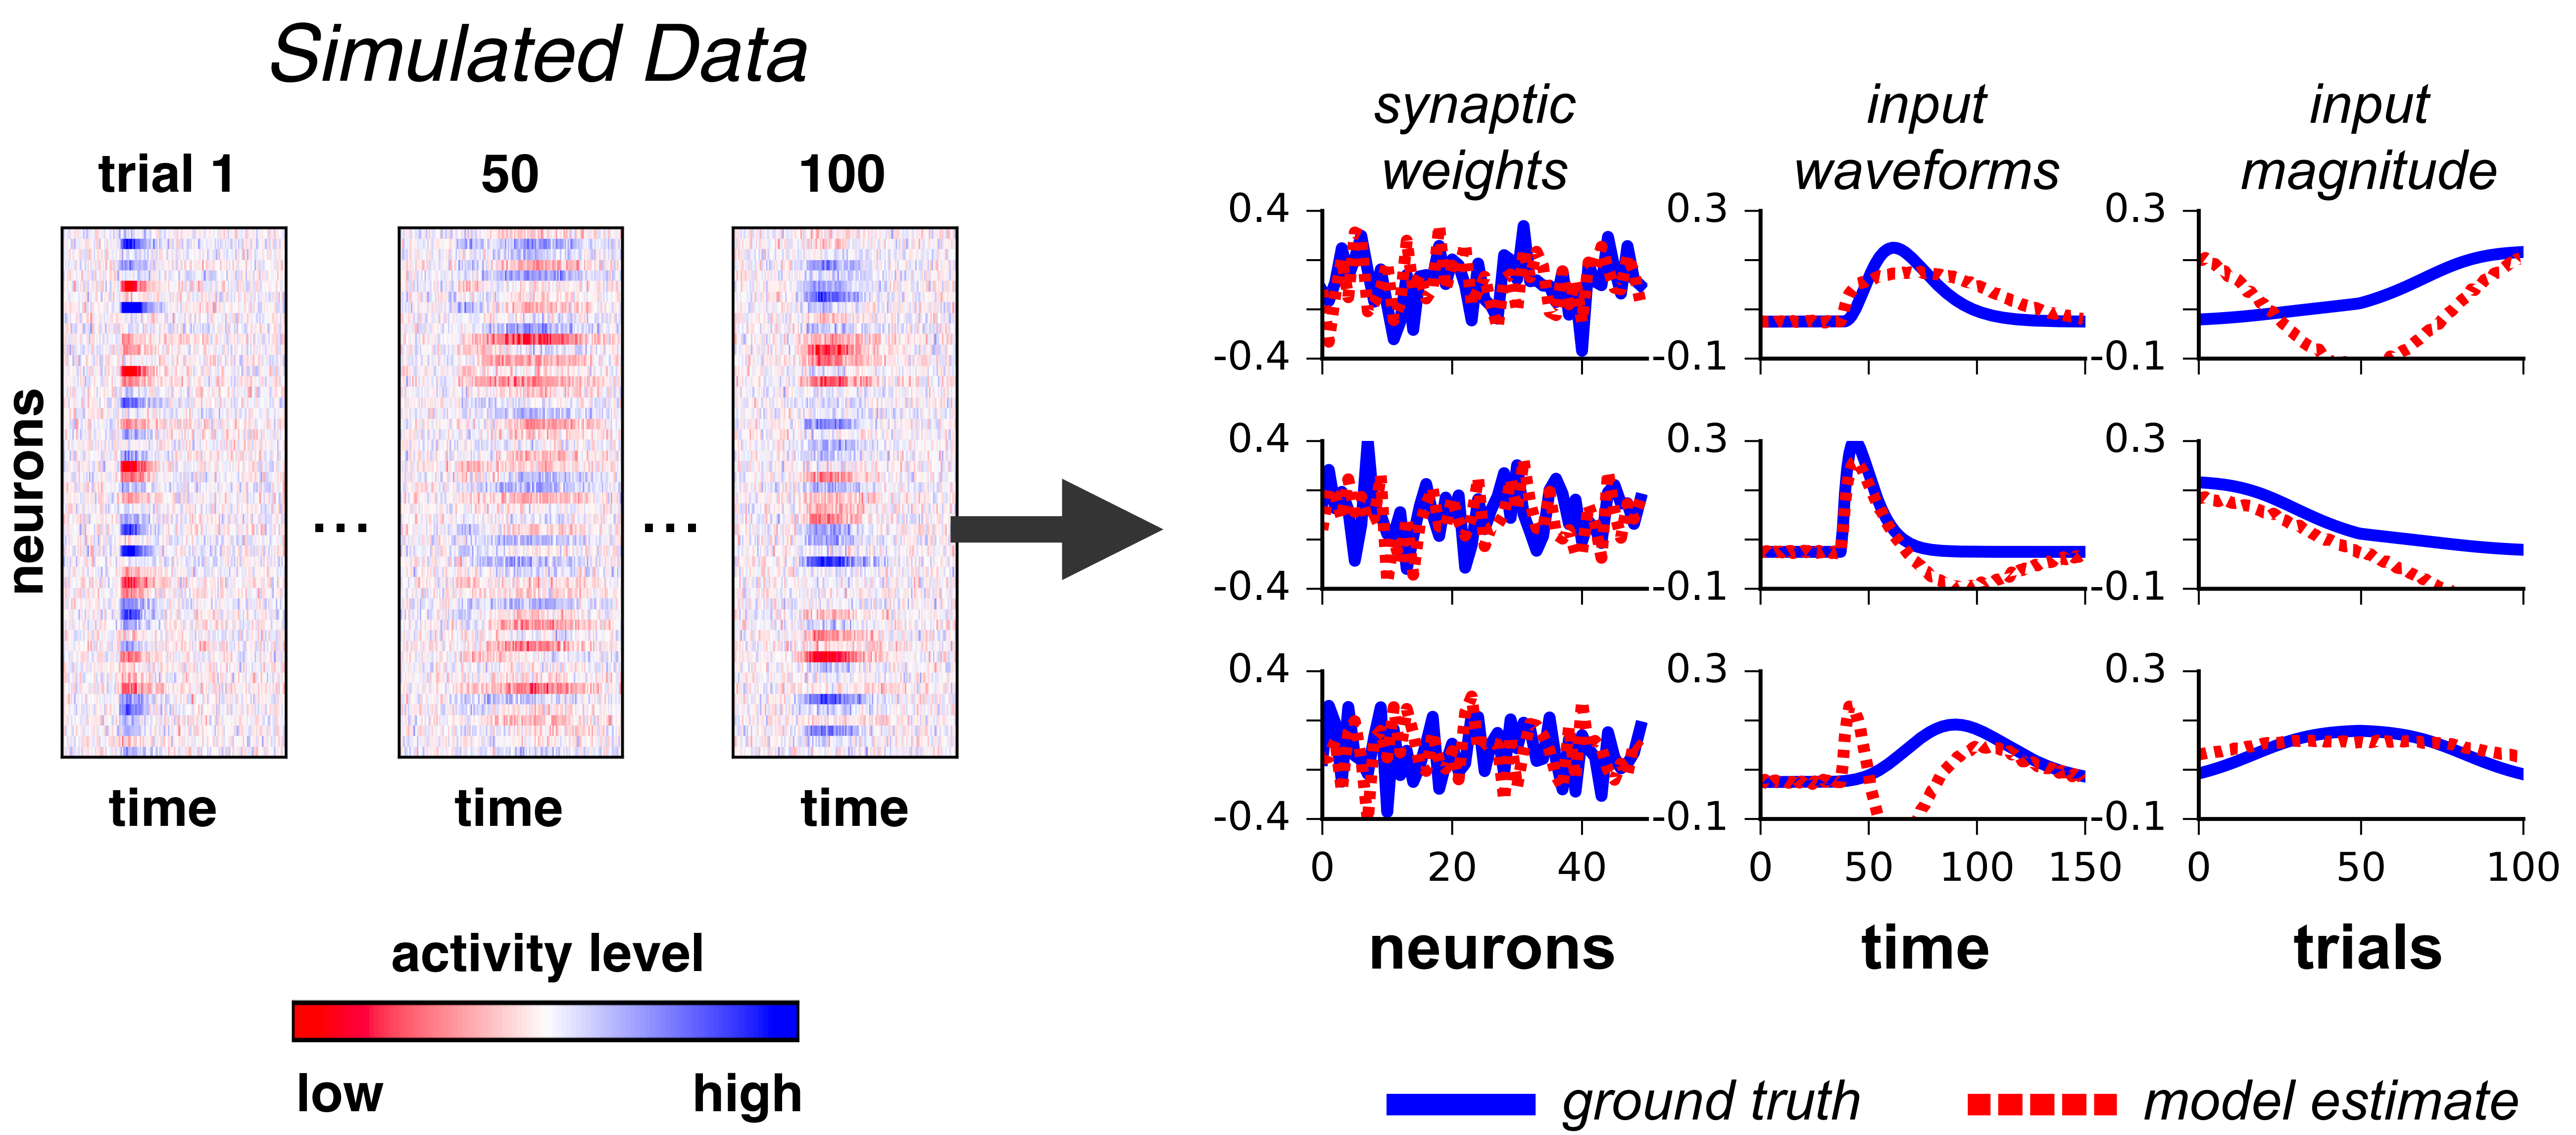

TCA precisely recovers all parameters

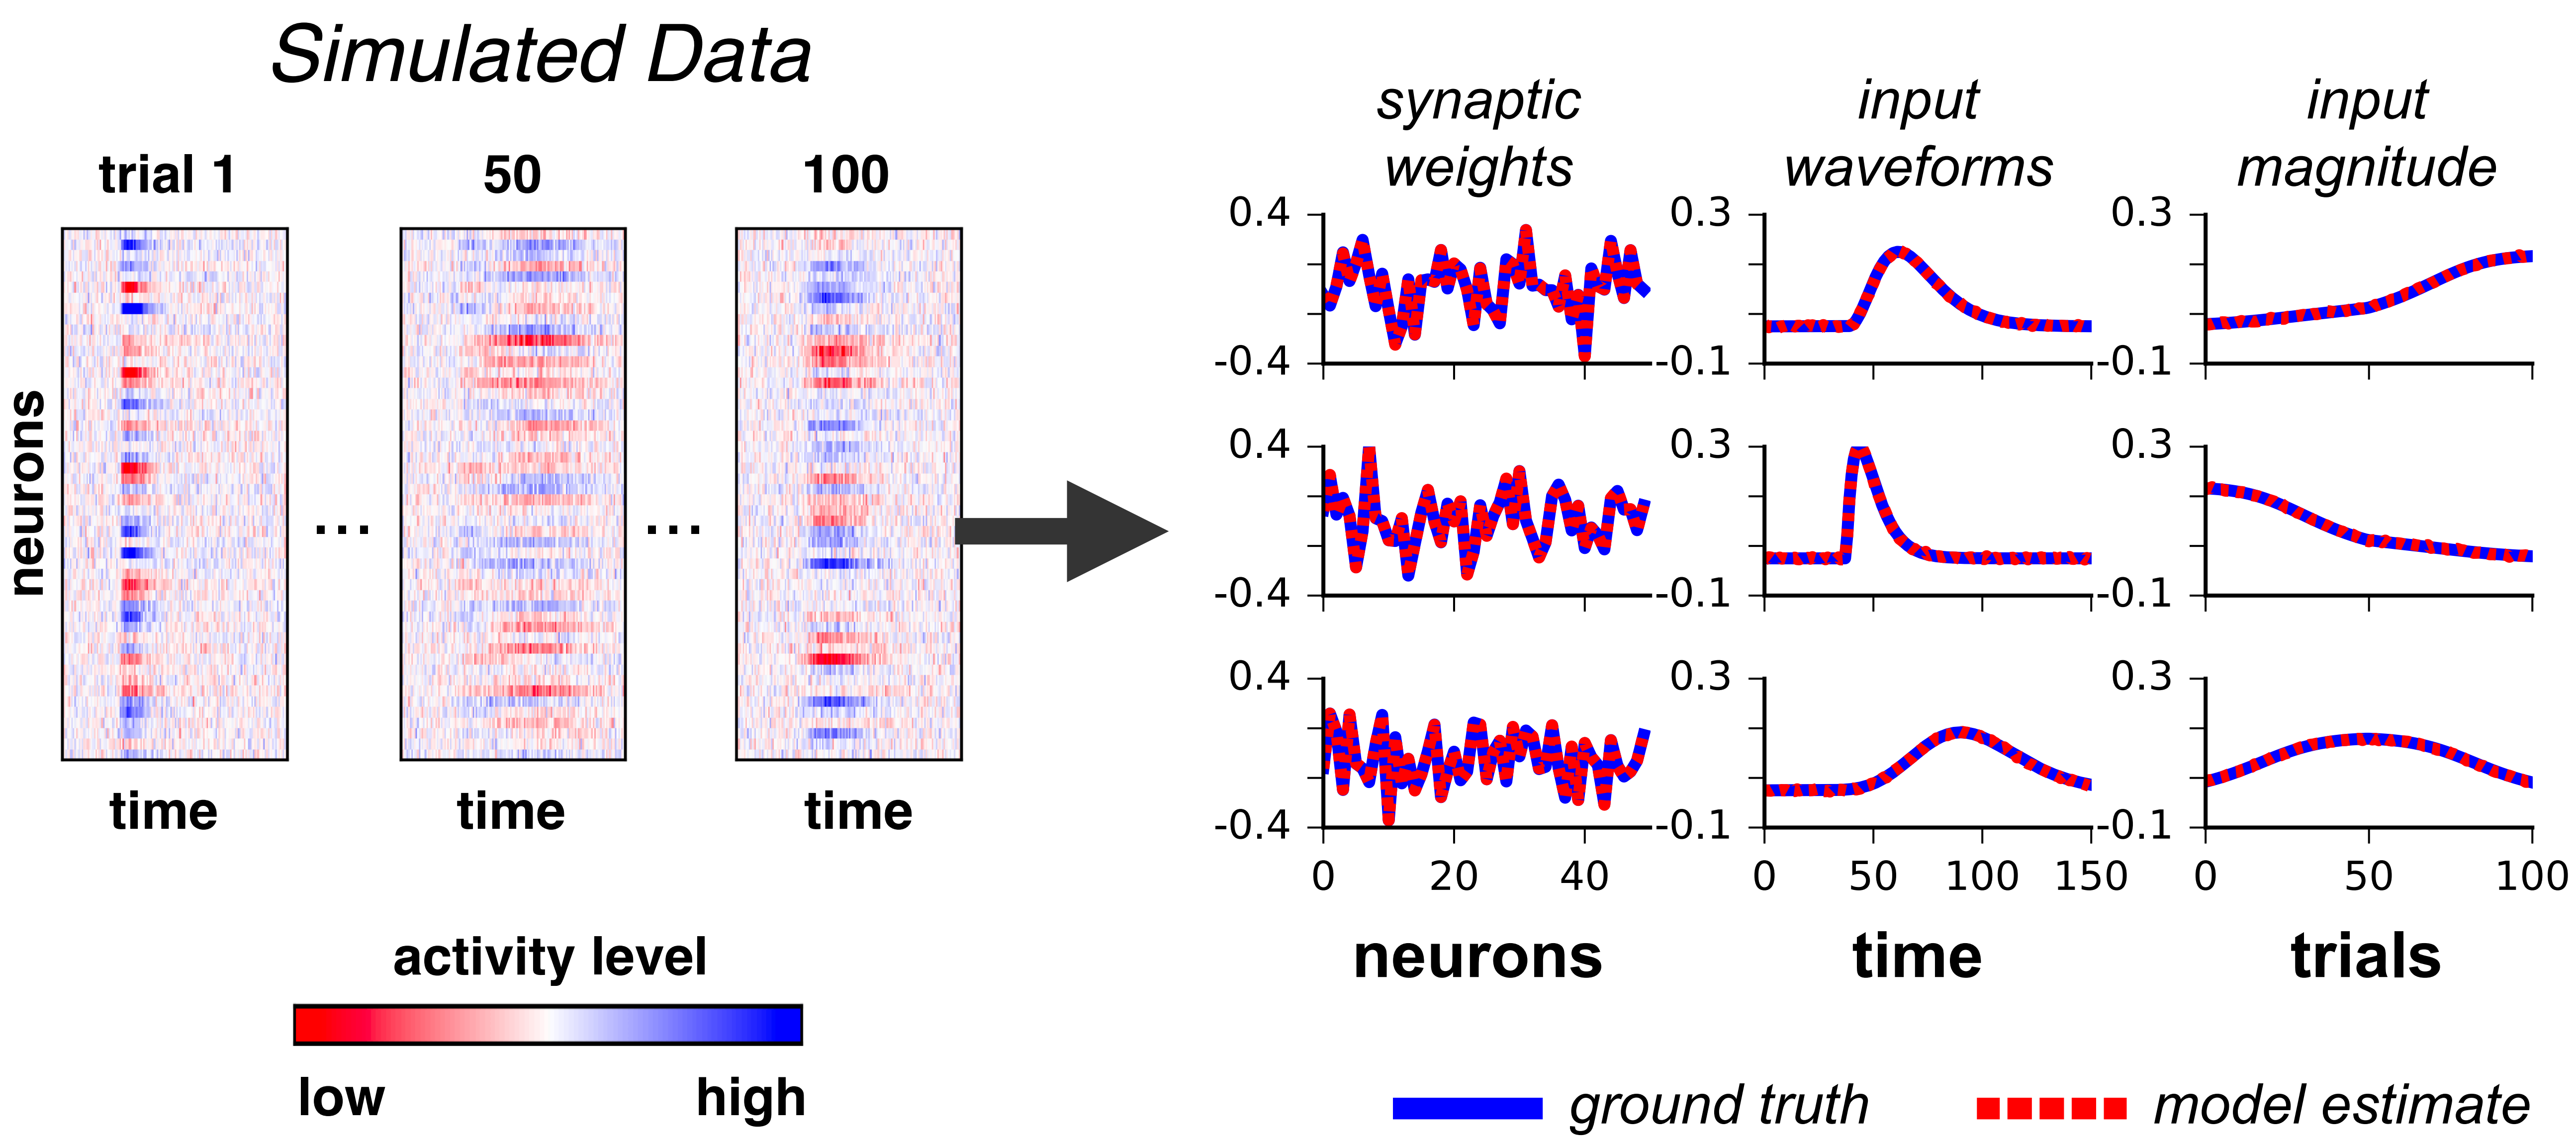

By design, PCA identifies a coordinate system instead of ground truth factors

Known as “rotational ambiguity” or the “rotation problem”

By design, PCA identifies a coordinate system instead of ground truth factors

Known as "rotational ambiguity" or the "rotation problem"

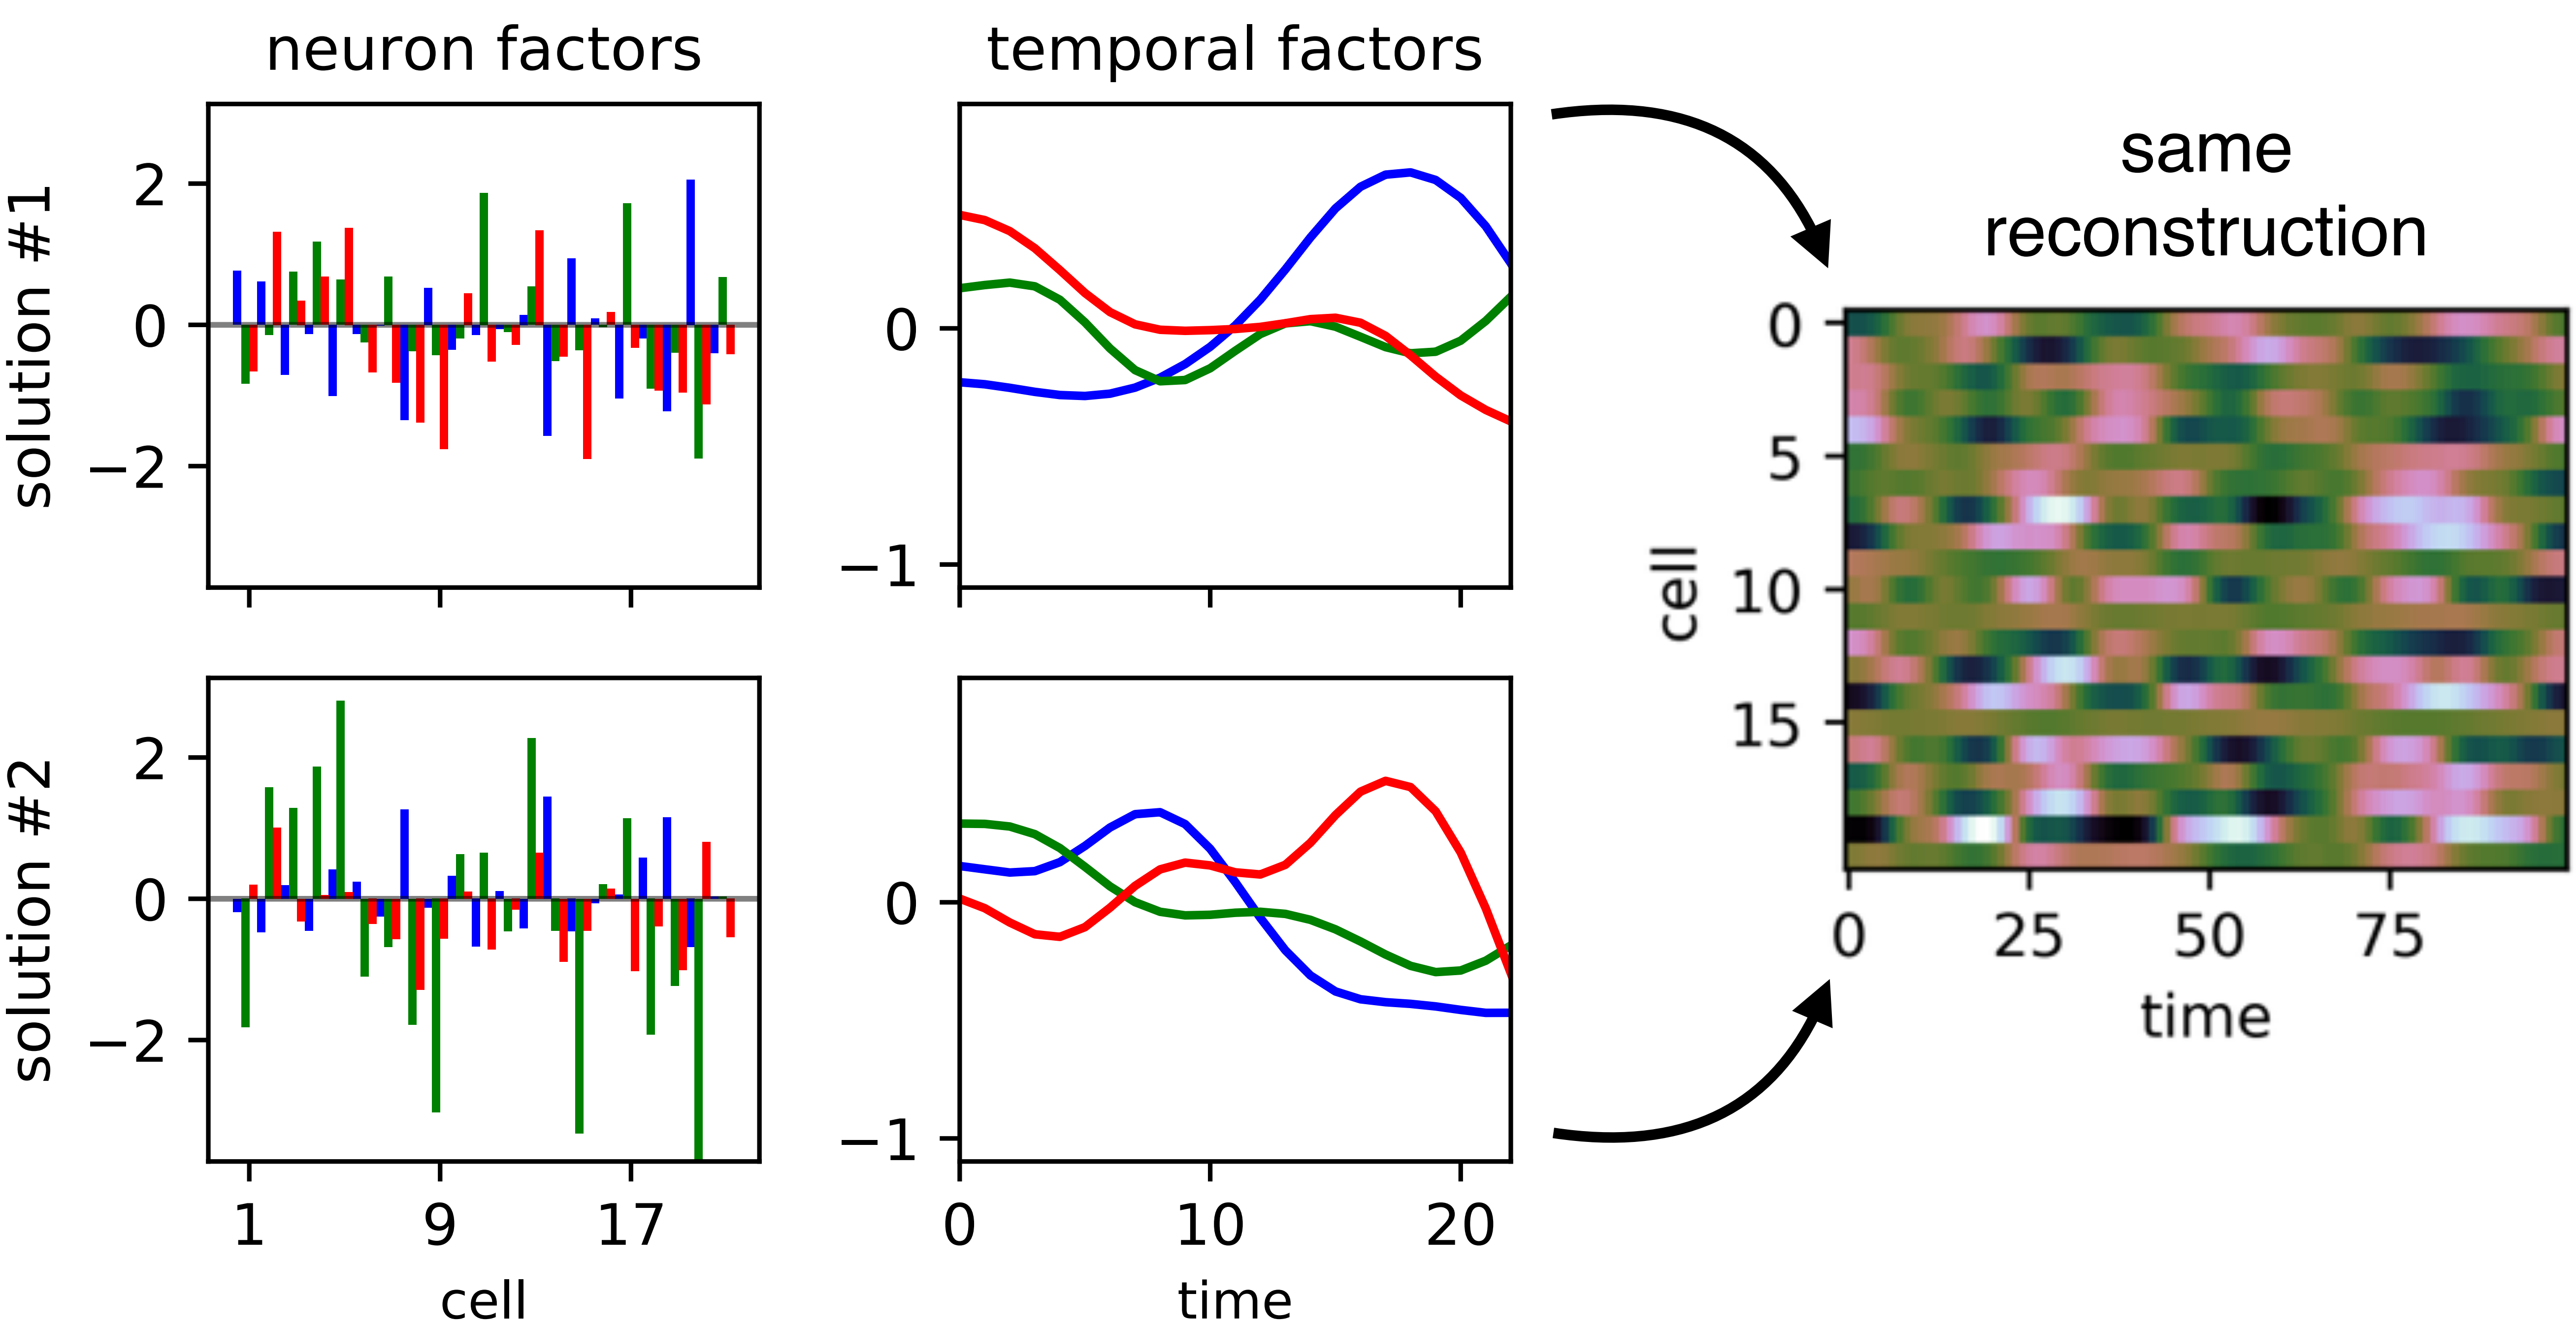

Seminal theorem (Kruskal, 1977) proves that linear independence is a sufficient condition for tensor decomposition identifiability

When PCA can recover ground truth:

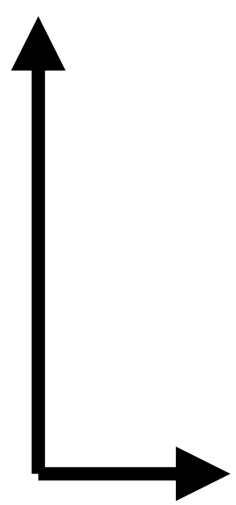

orthogonal factors, large eigengap

When TCA can recover ground truth:

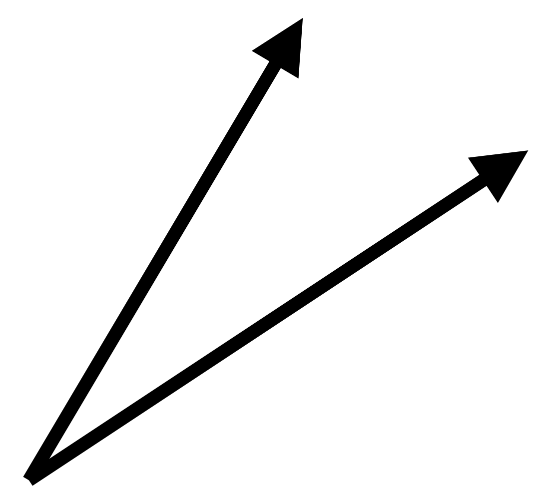

factors can be correlated and have similar magnitudes

Summary Thusfar

- TCA separates **fast, temporal factors** from **slow, across-trial factors**.
- Despite being a simple generalization of PCA, it has strikingly advantageous theoretical properties.
- Strong connection between TCA as a statistical model and the principle of gain modulation.

Summary Thusfar

- TCA separates **fast, temporal factors** from **slow, across-trial factors**.
- Despite being a simple generalization of PCA, it has strikingly advantageous theoretical properties.
- Strong connection between TCA as a statistical model and the principle of **gain modulation**.

Applications

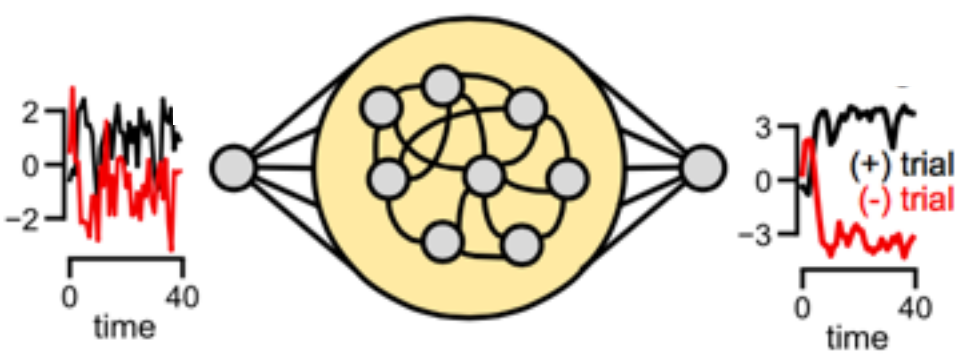

learning in artificial networks via backpropagation

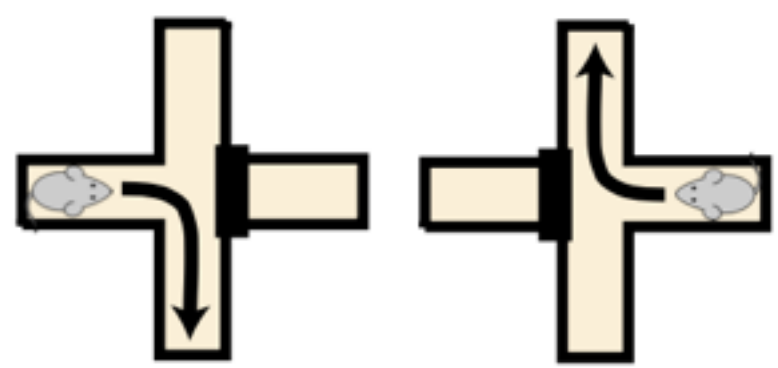

navigation with switching reward contingencies

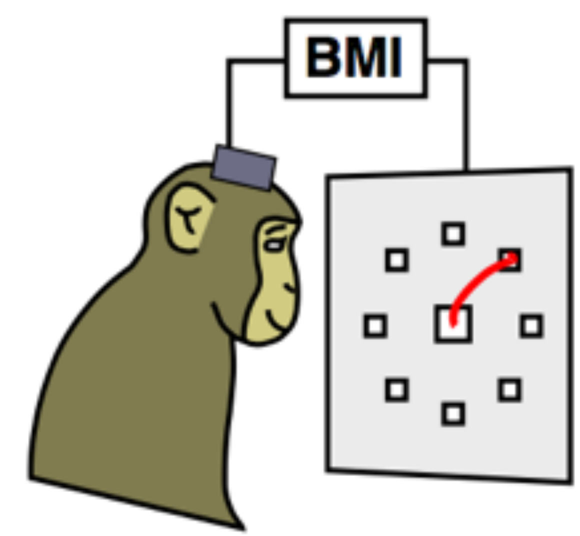

BMI learning and adaptation

Application #1: How does a model network learn a sensory discrimination task?

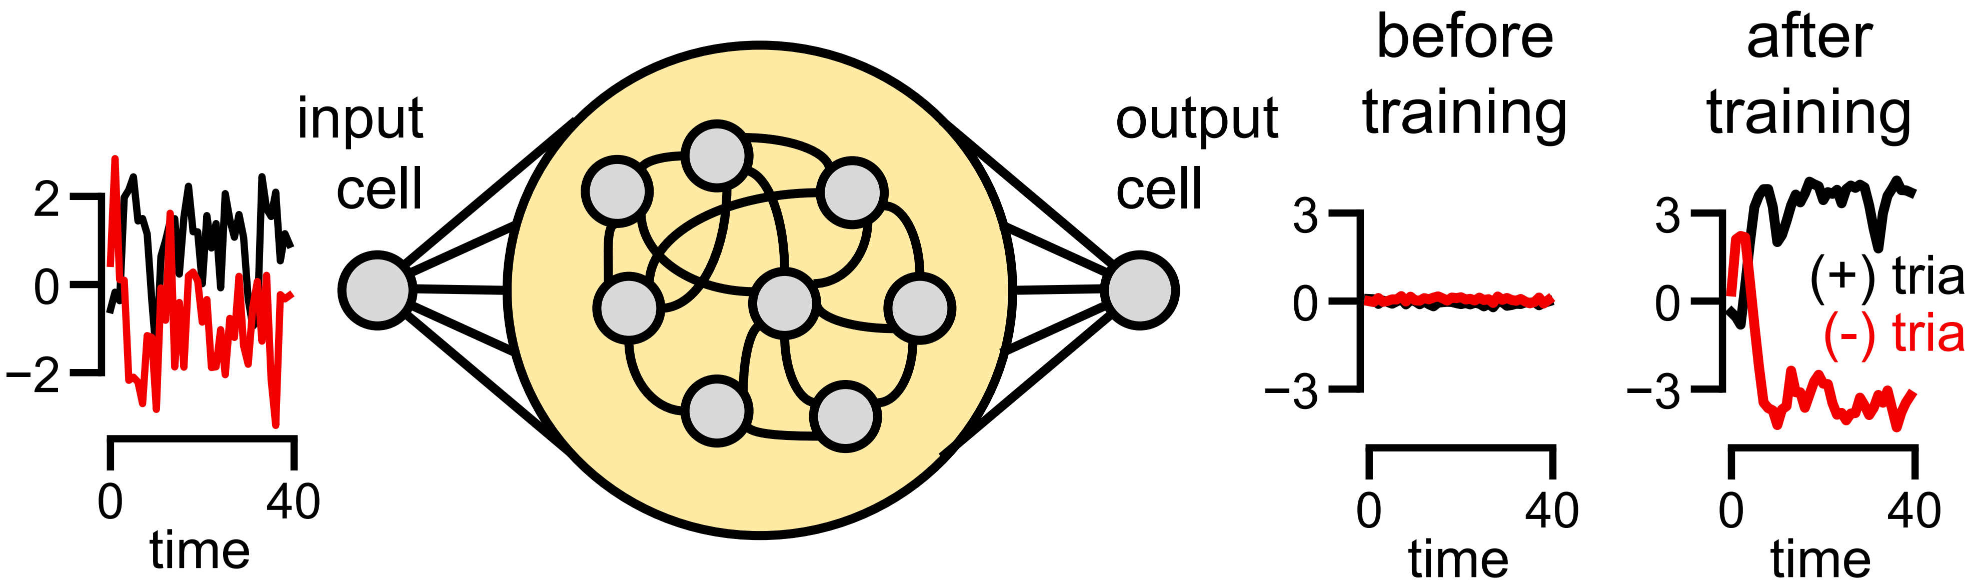

Application #1: How does a model network learn a sensory discrimination task?

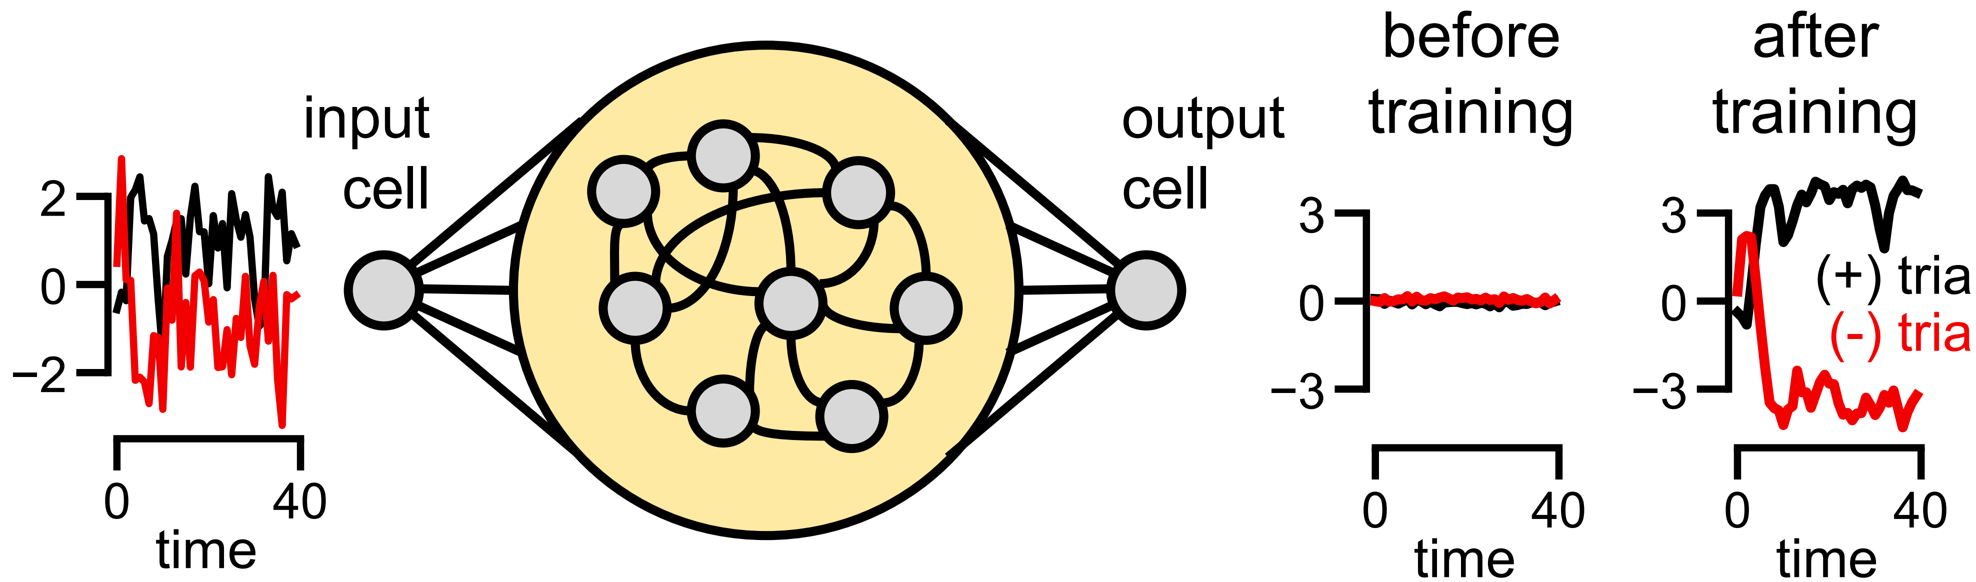

Analogous to classic experiments in primates

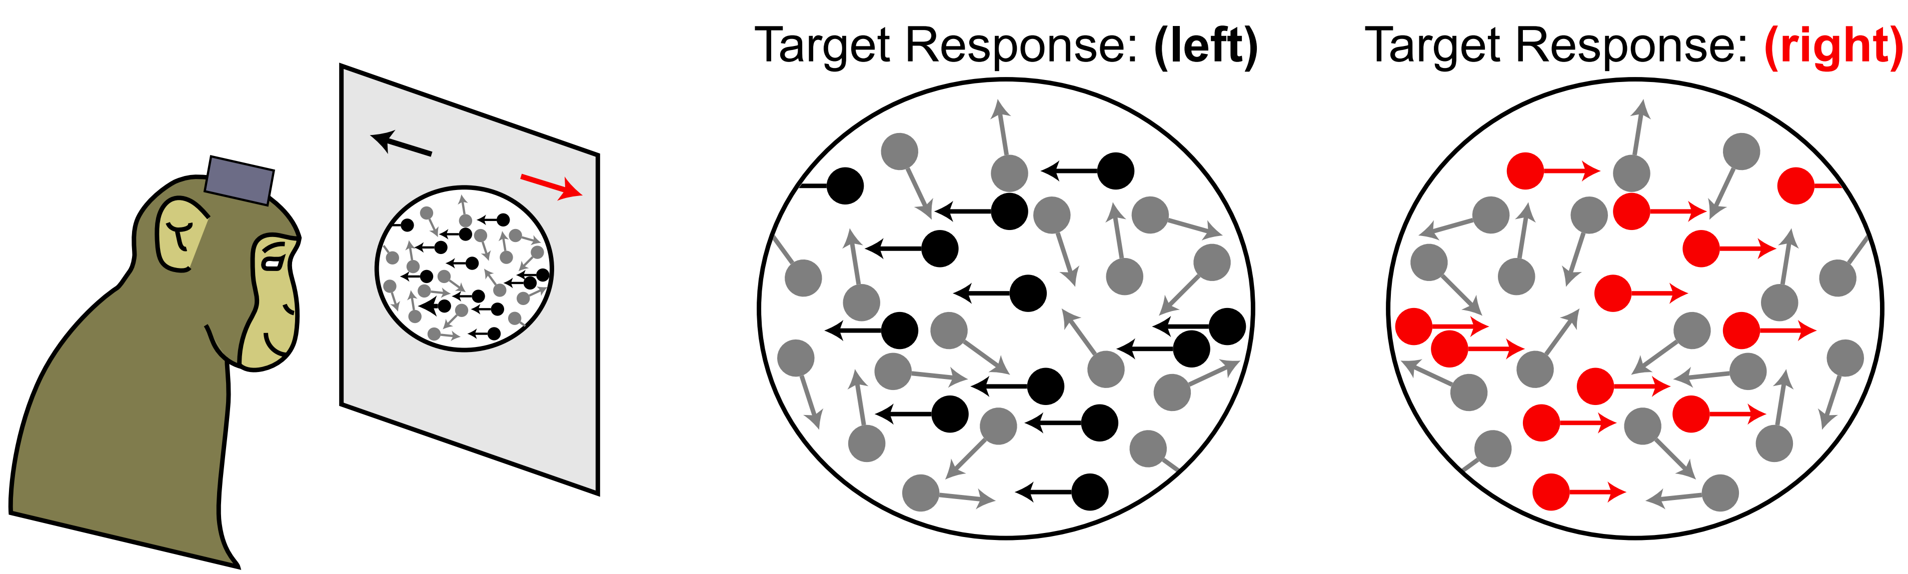

Application #1: How does a model network learn a sensory discrimination task?

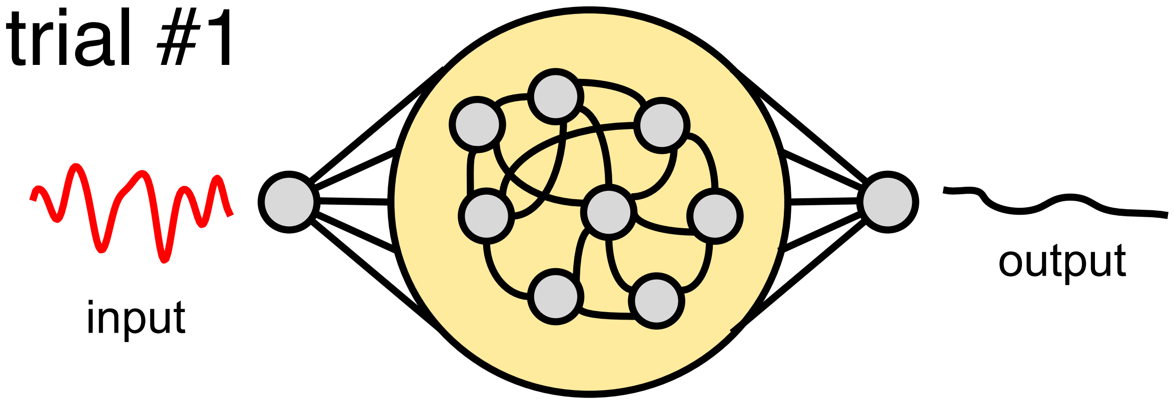

Application #1: How does a model network learn a sensory discrimination task?

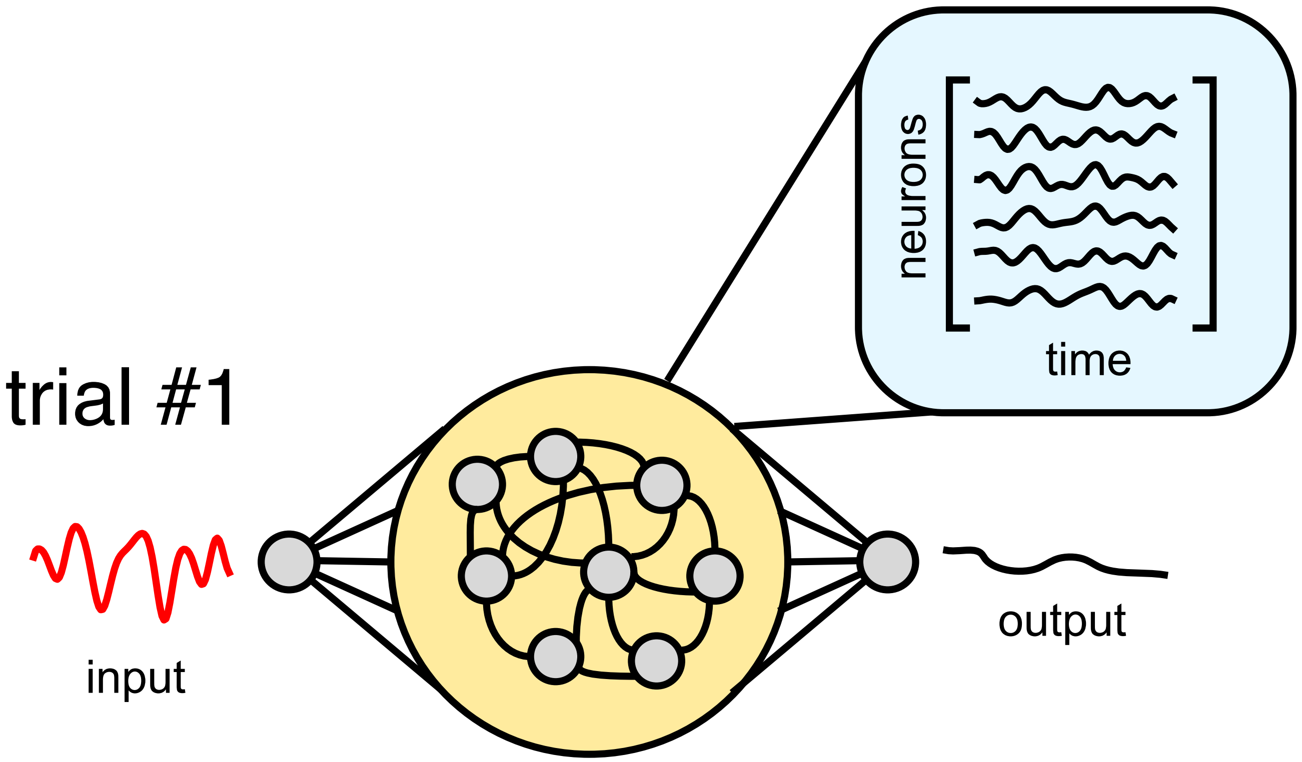

Application #1: How does a model network learn a sensory discrimination task?

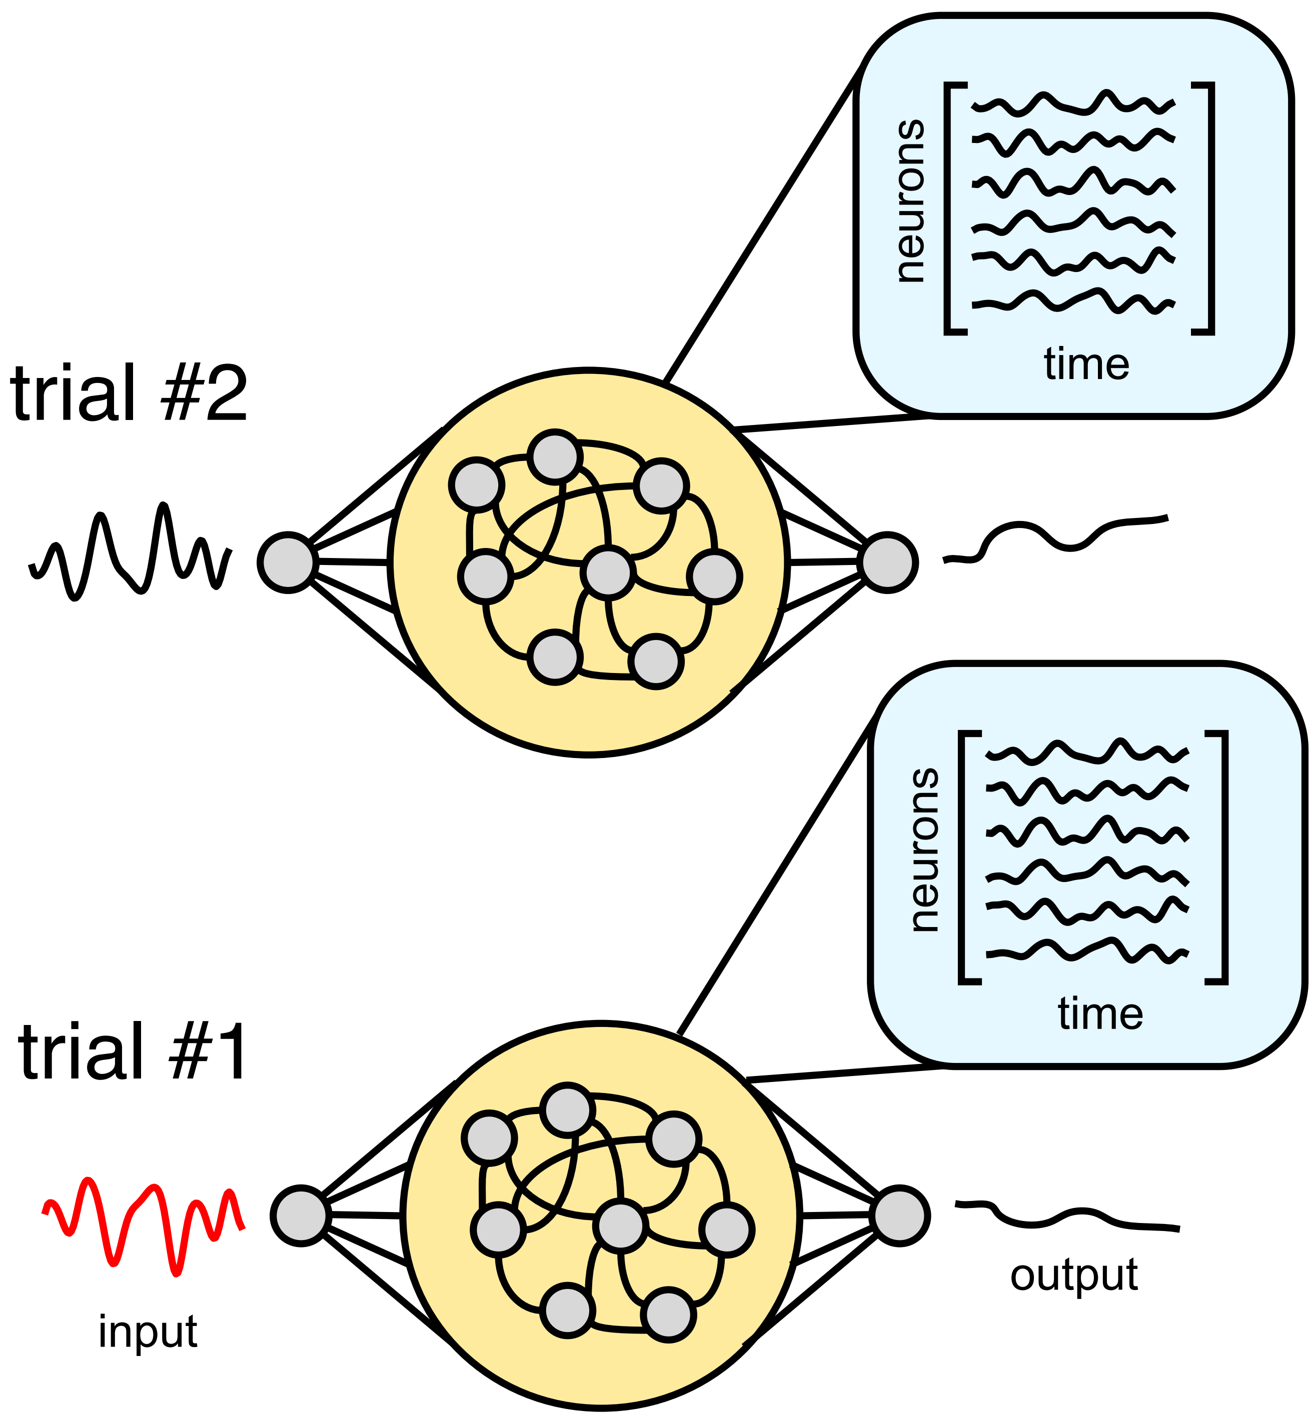

Application #1: How does a model network learn a sensory discrimination task?

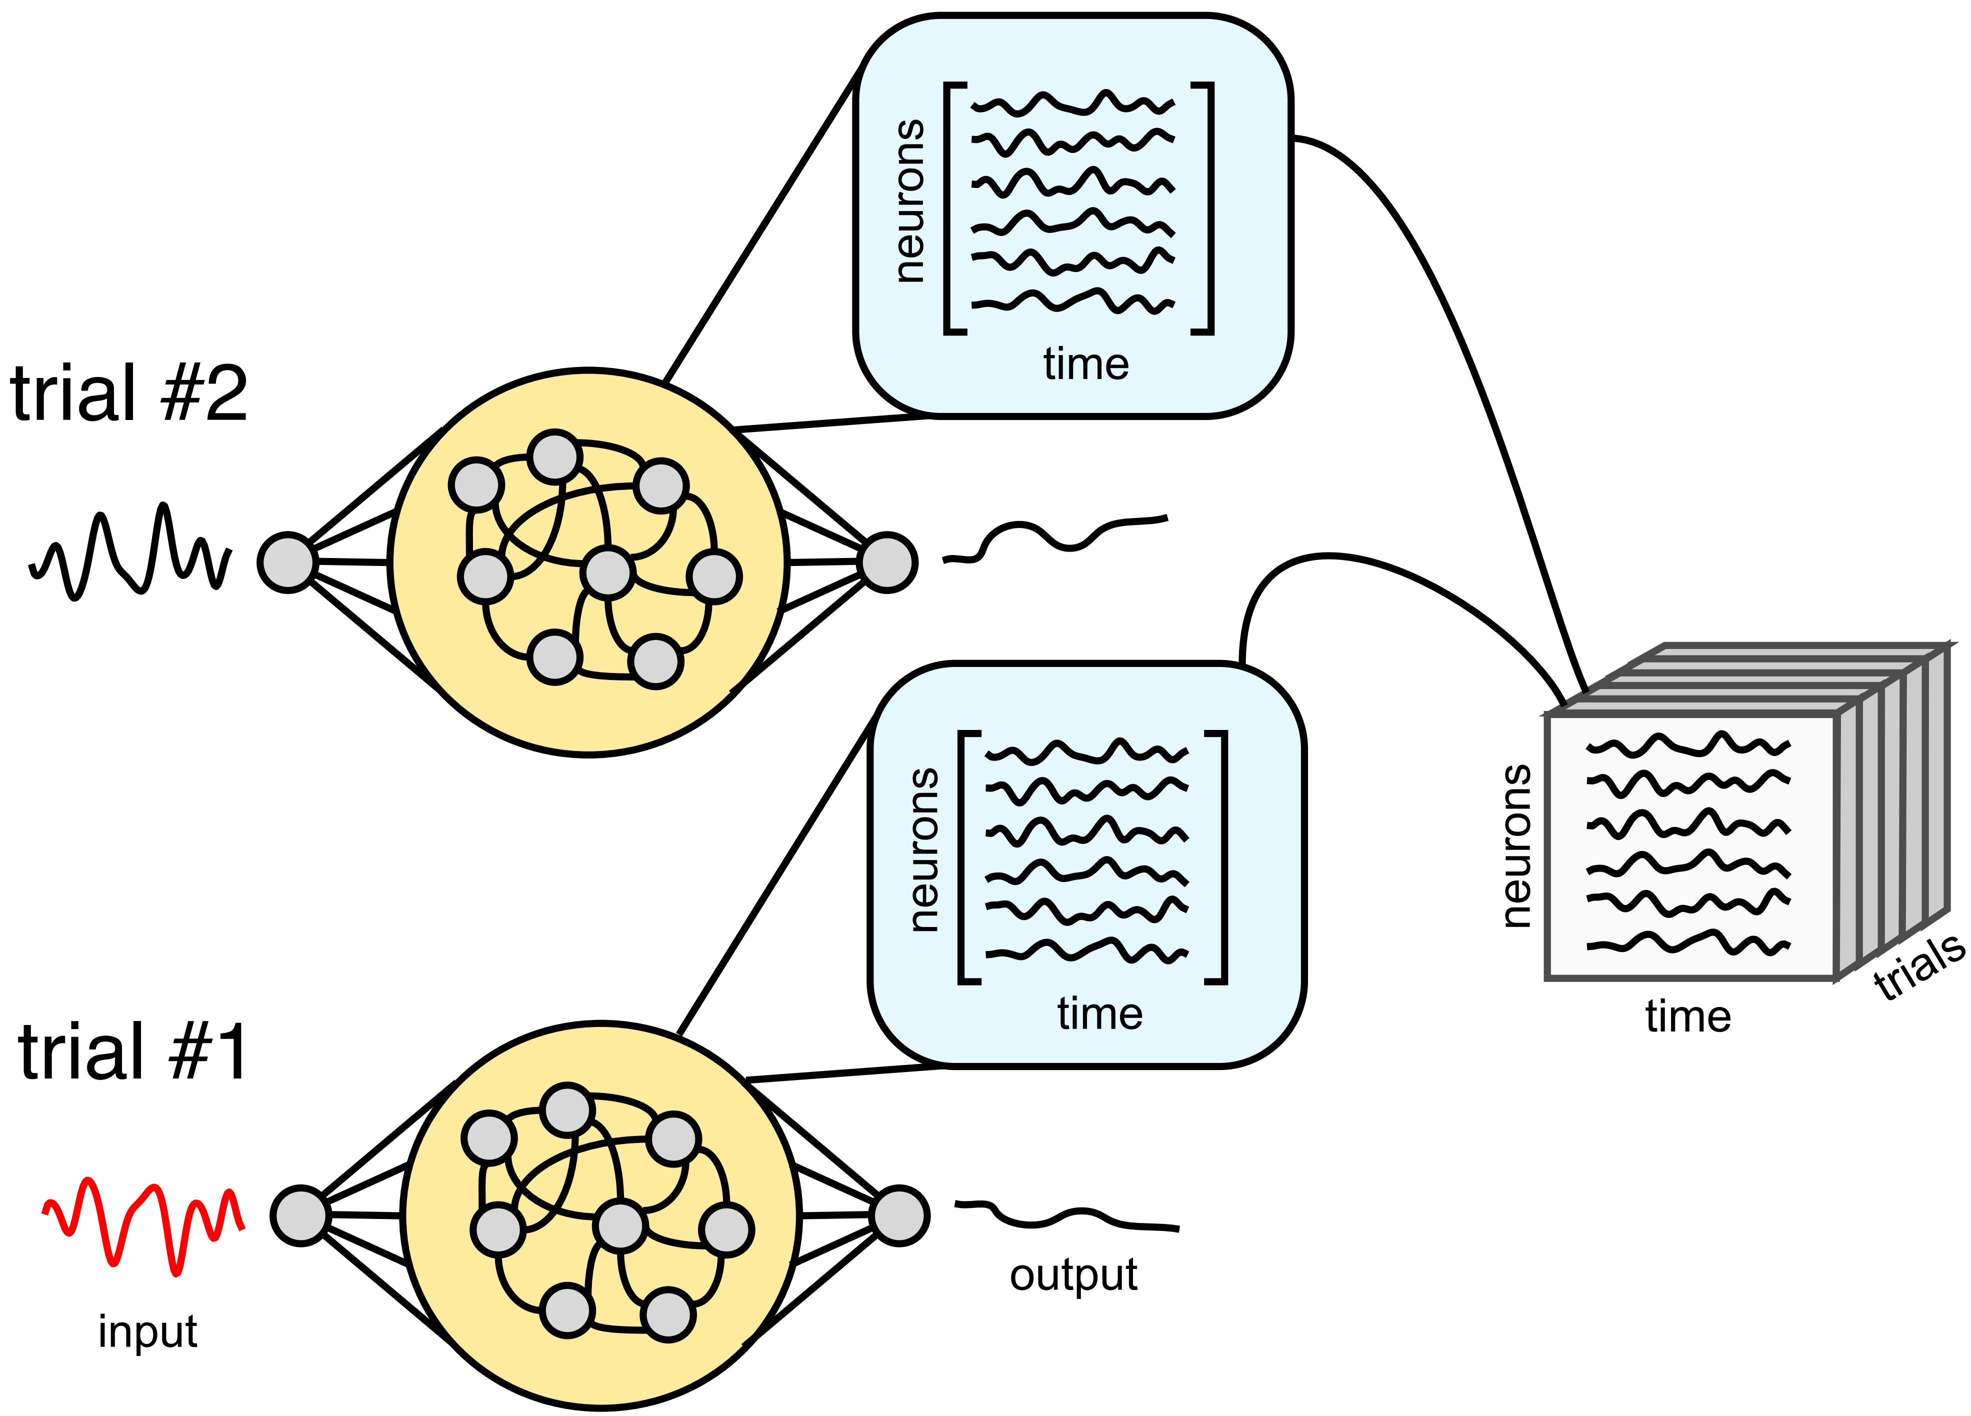

Application #1: How does a model network learn a sensory discrimination task?

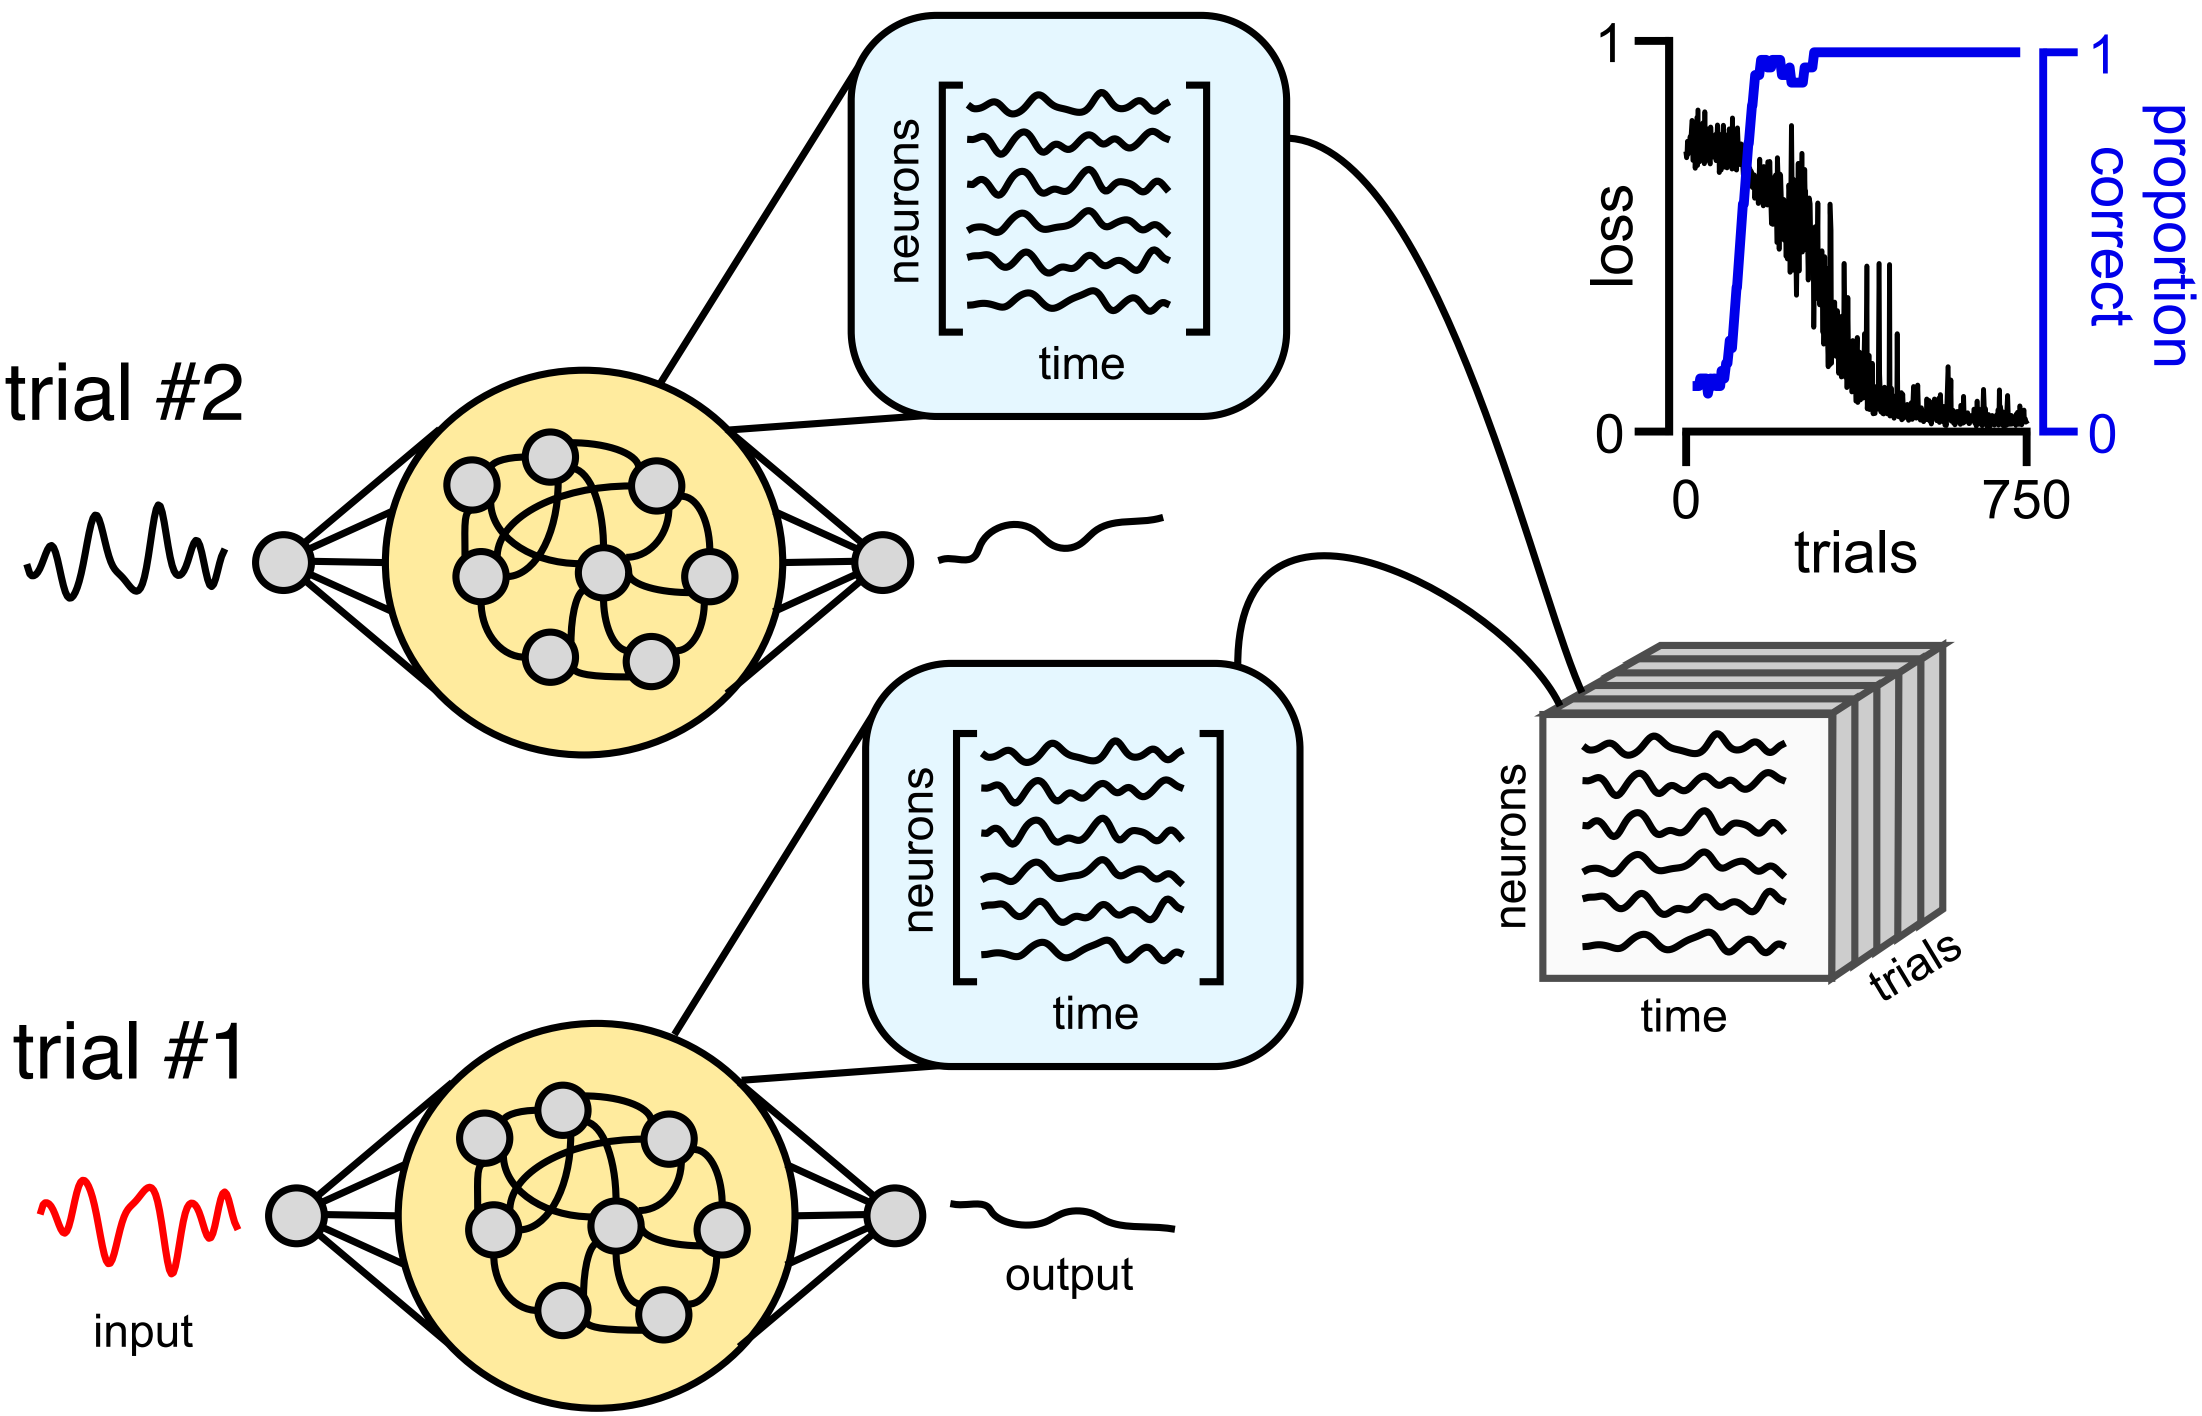

Gain modulation is a compact and accurate model of the network activity over all trials

Two example cells before and after training

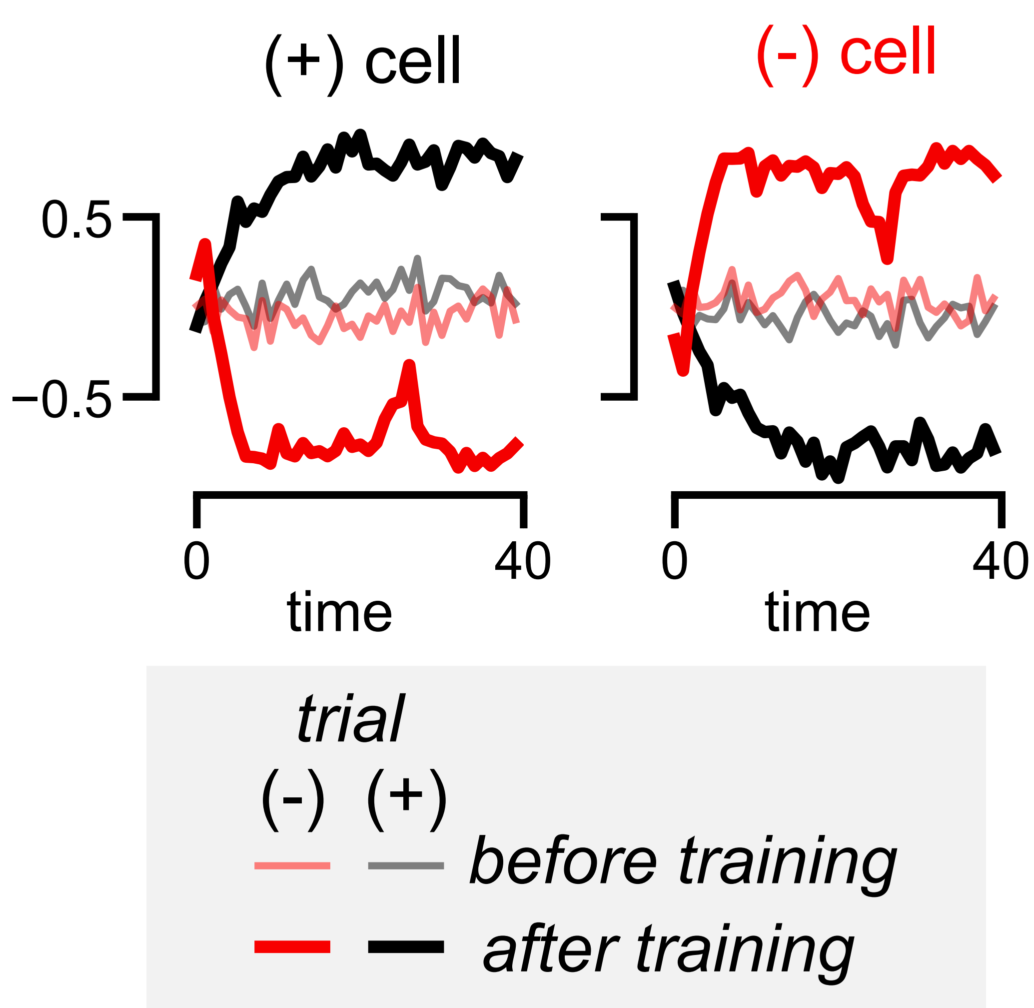

Gain modulation is a compact and accurate model of the network activity over all trials

Two example cells before and after training

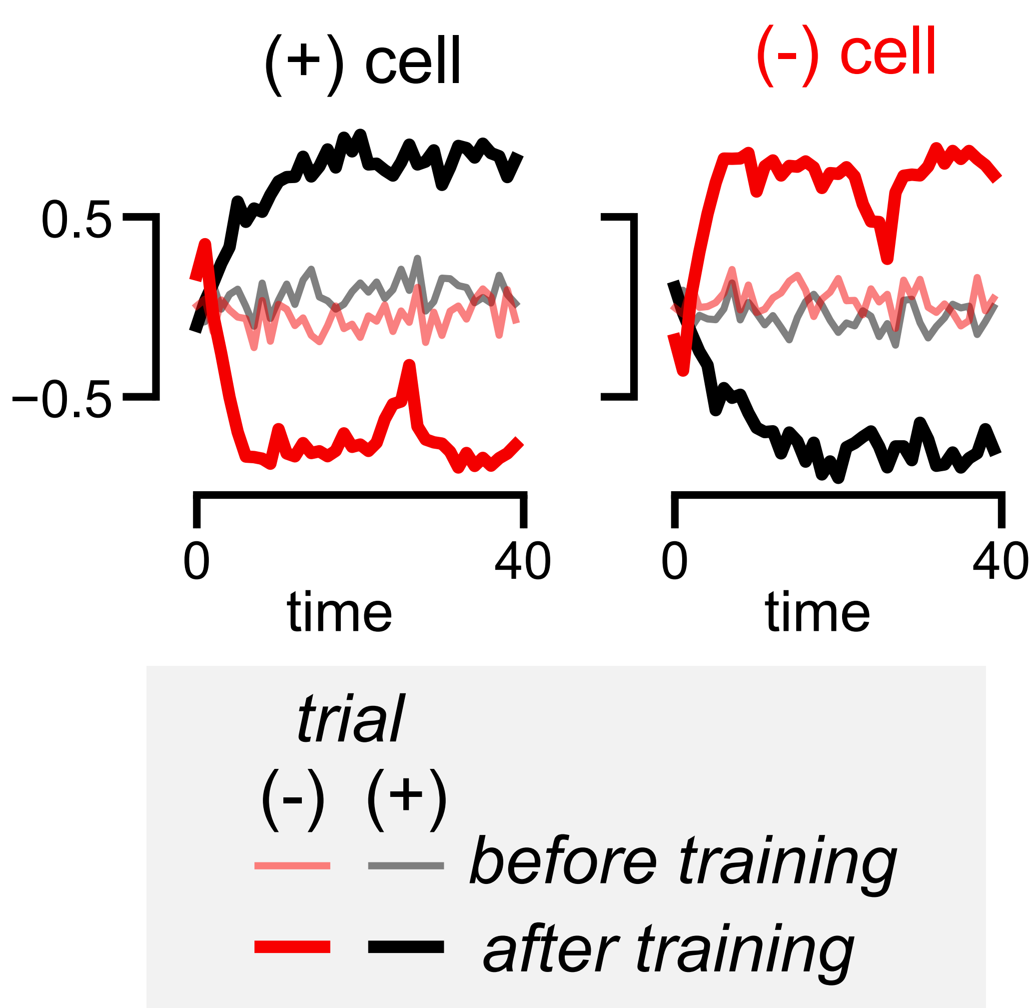

TCA with 1 component describes the vast majority of variance in firing rates

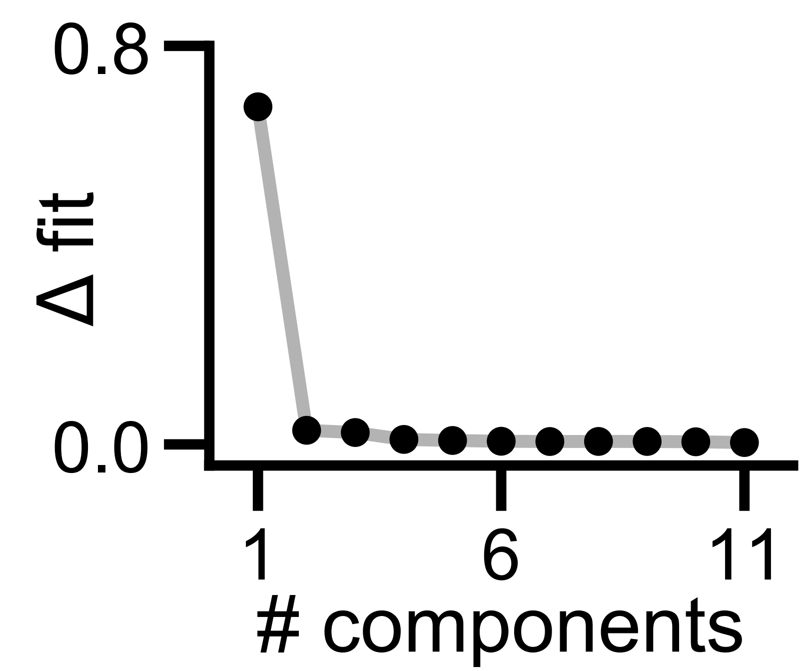

TCA with one component (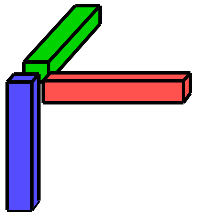) identifies:

TCA with one component (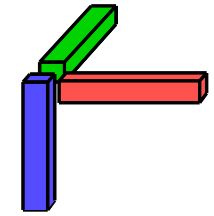) identifies:

1. Cells with opposing responses to the stimulus **(neuron factor)**

neuron factor

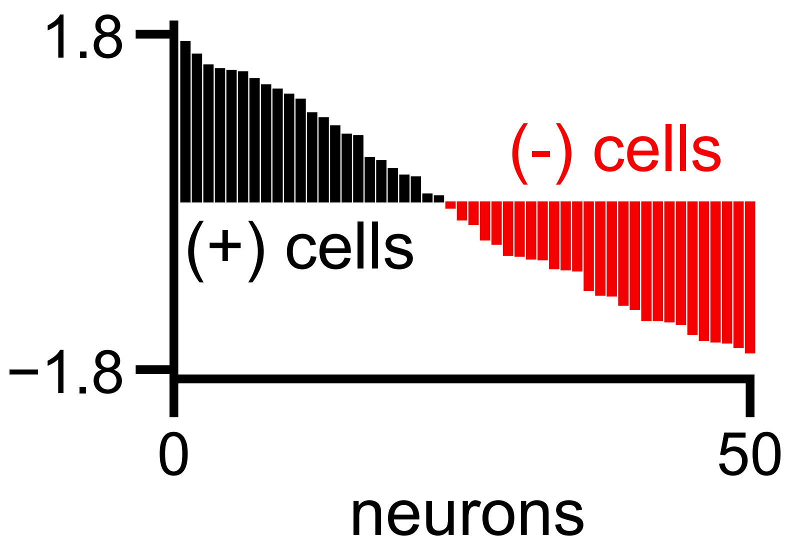

TCA with one component (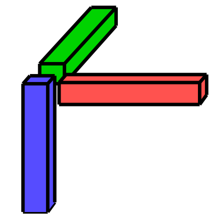) identifies:

1. Cells with opposing responses to the stimulus **(neuron factor)**
2. A decision timescale shared by all neurons **(temporal factor)**

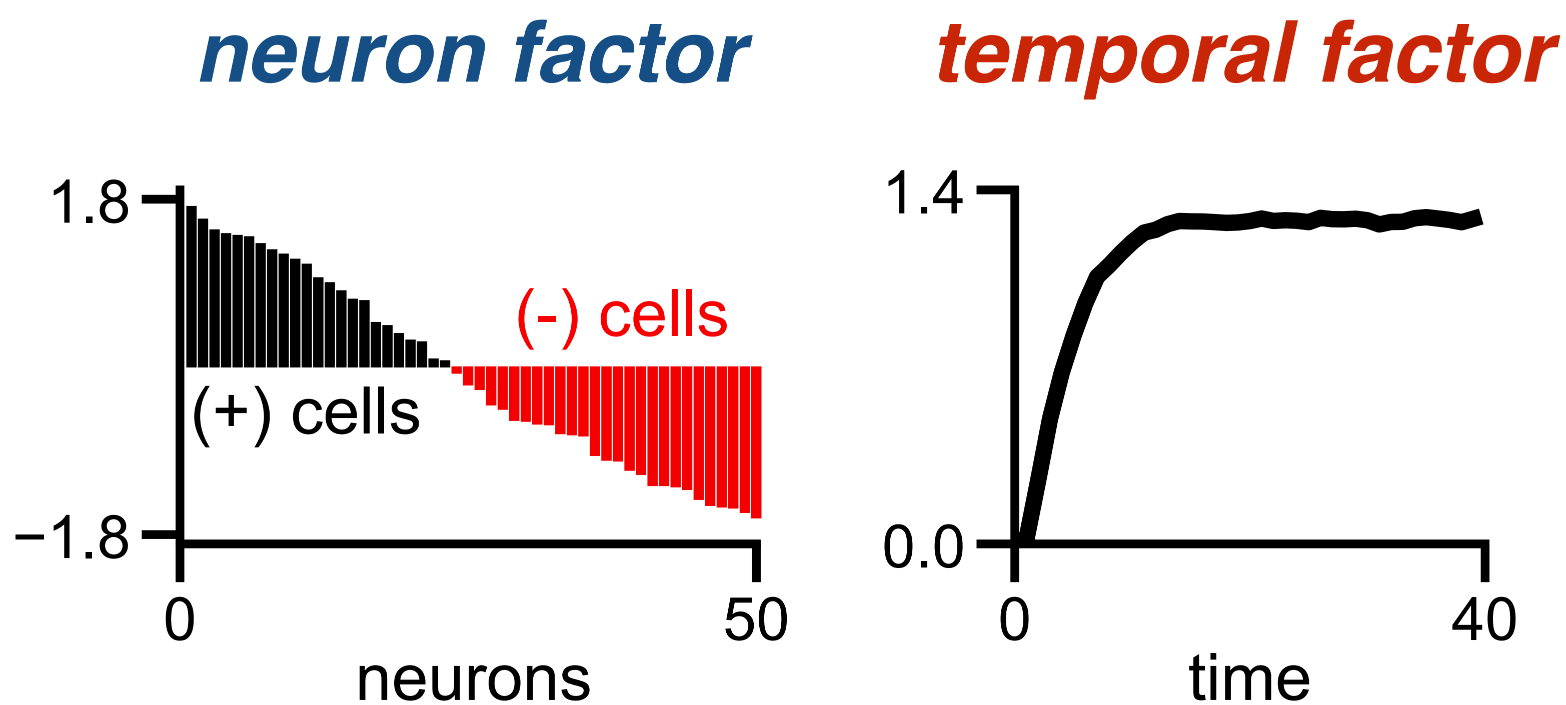

TCA with one component (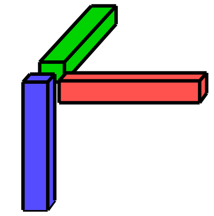) identifies:

1. Cells with opposing responses to the stimulus **(neuron factor)**
2. A decision timescale shared by all neurons **(temporal factor)**
3. Network decisions trial-by-trial and learning **(trial factor)**

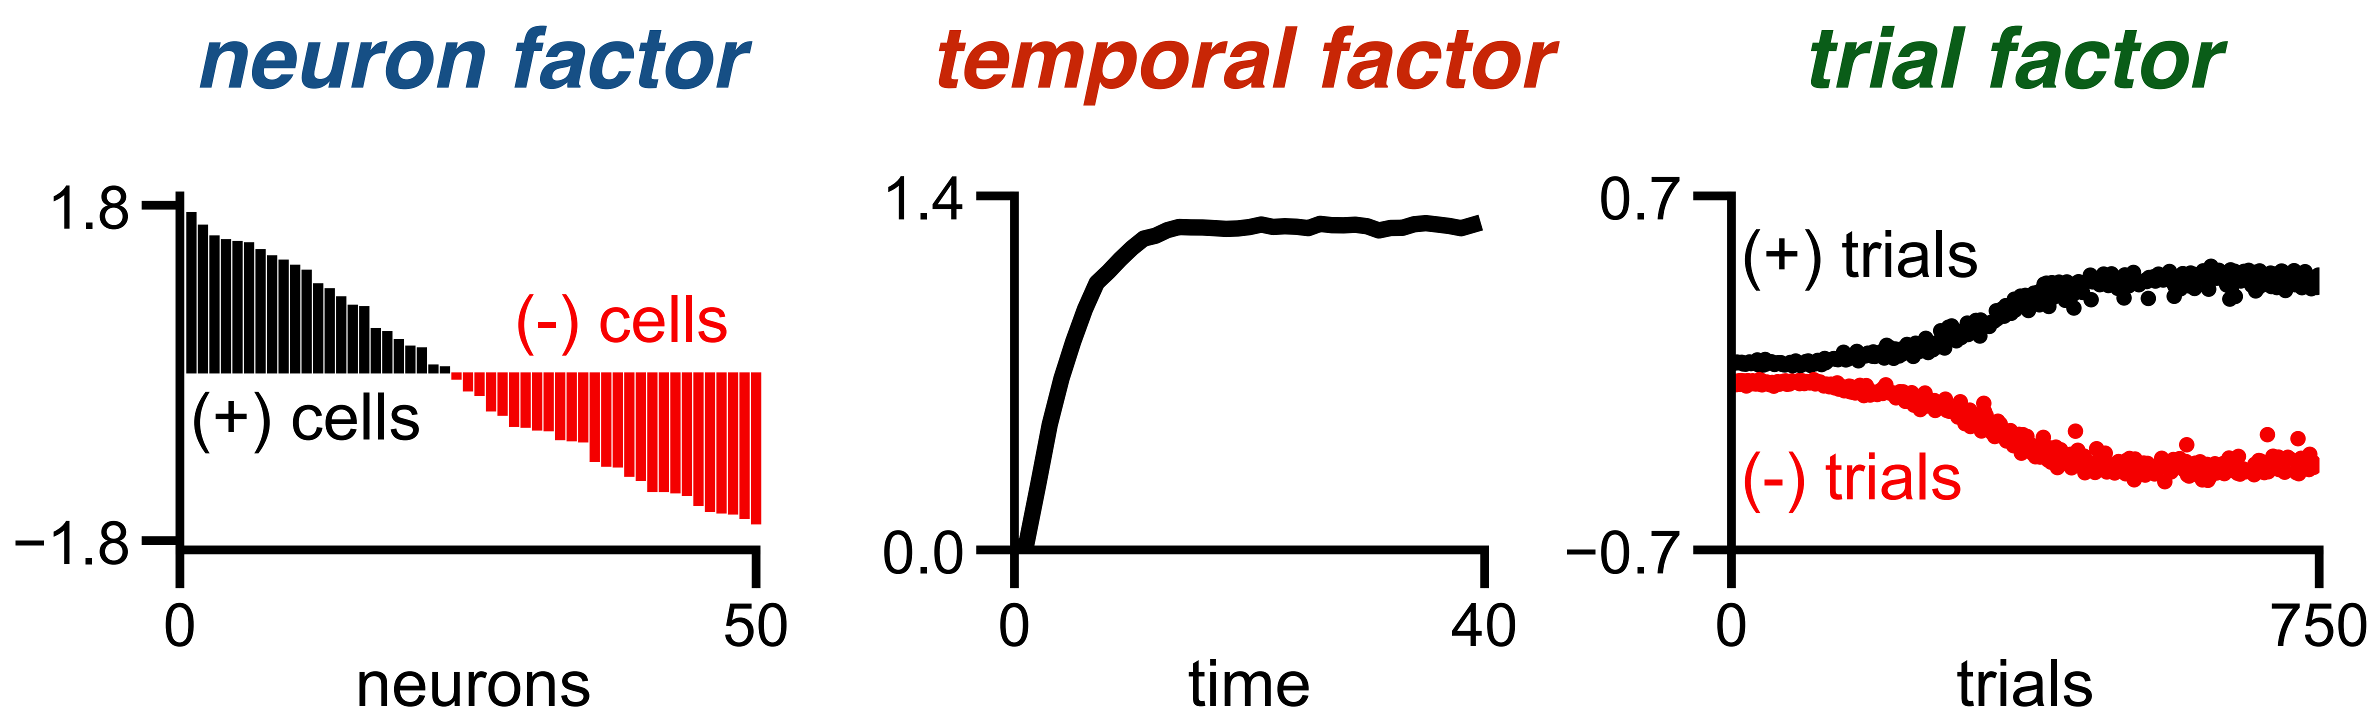

TCA with one component (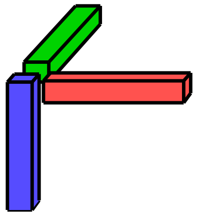) identifies:

1. Cells with opposing responses to the stimulus **(neuron factor)**
2. A decision timescale shared by all neurons **(temporal factor)**
3. Network decisions trial-by-trial and learning **(trial factor)**

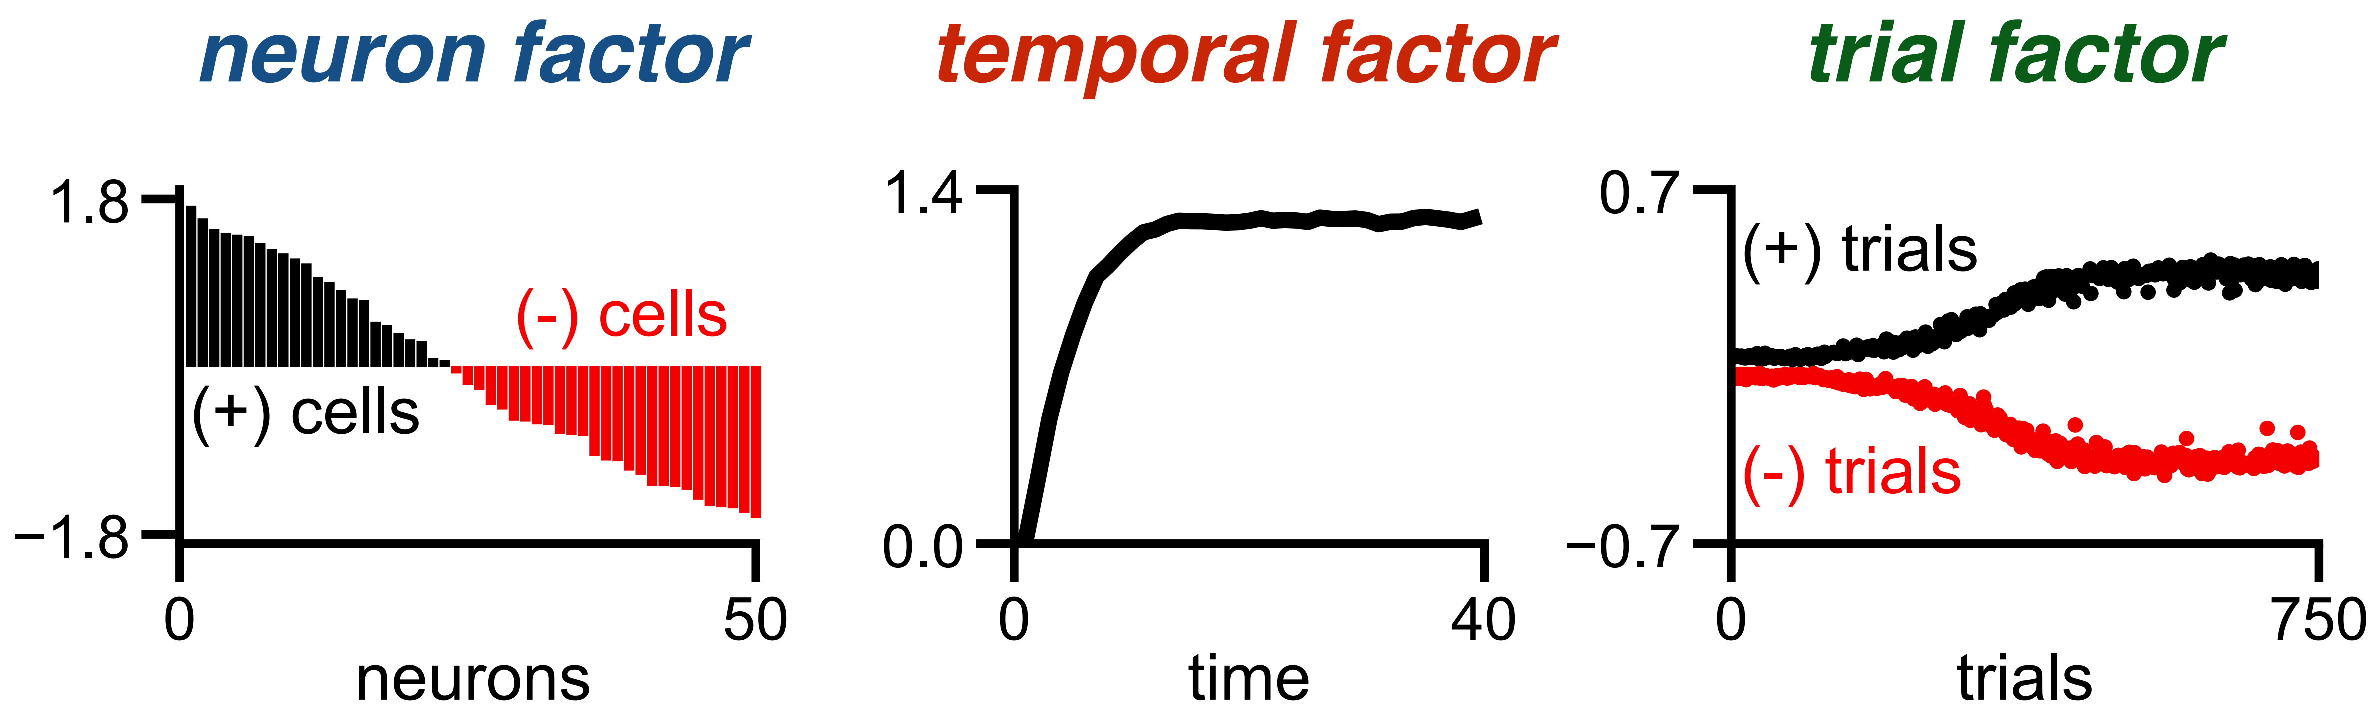

Resorting the network connectivity by the ***TCA neuron factor*** reveals winner-take-all structure

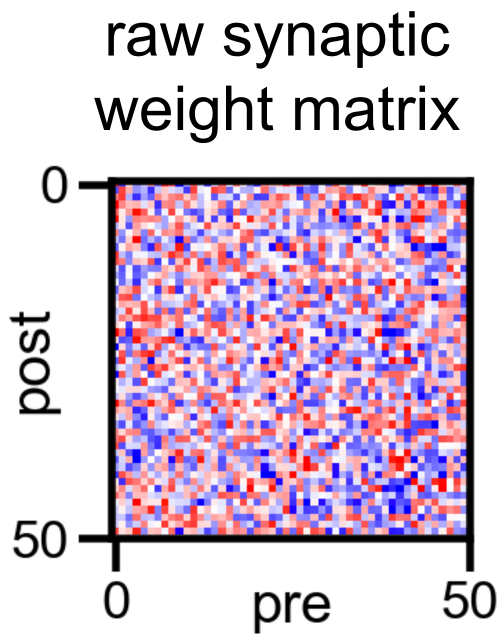

Resorting the network connectivity by the ***TCA neuron factor*** reveals winner-take-all structure

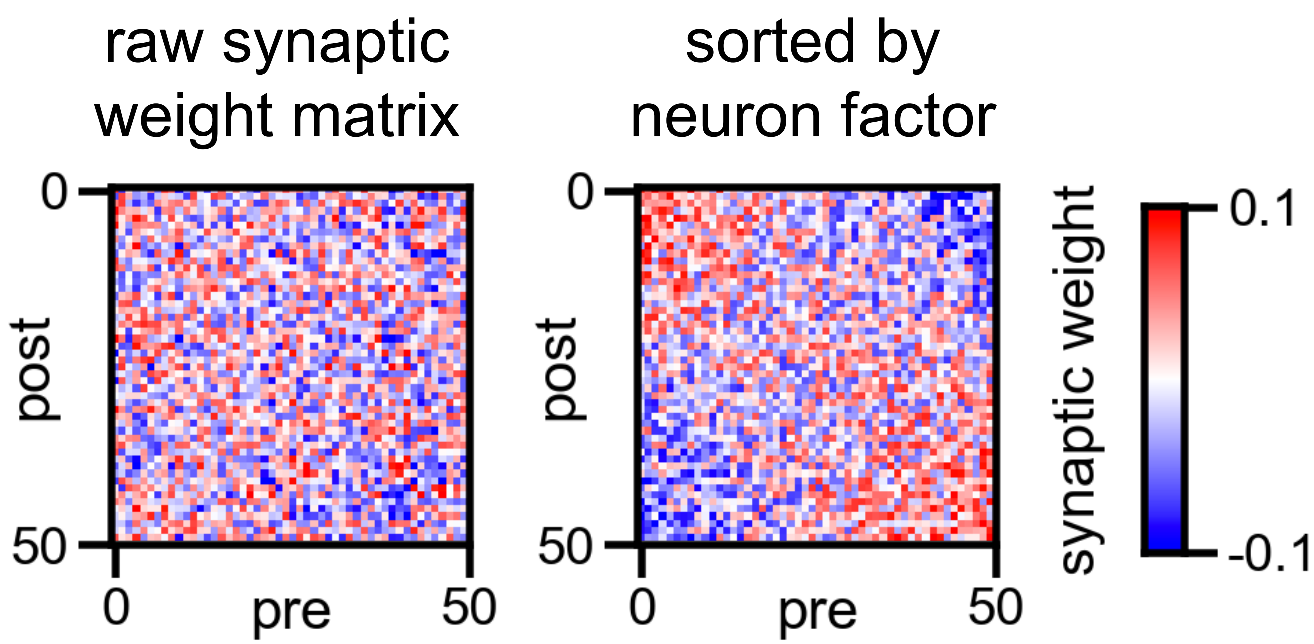

Resorting the network connectivity by the ***TCA neuron factor*** reveals winner-take-all structure

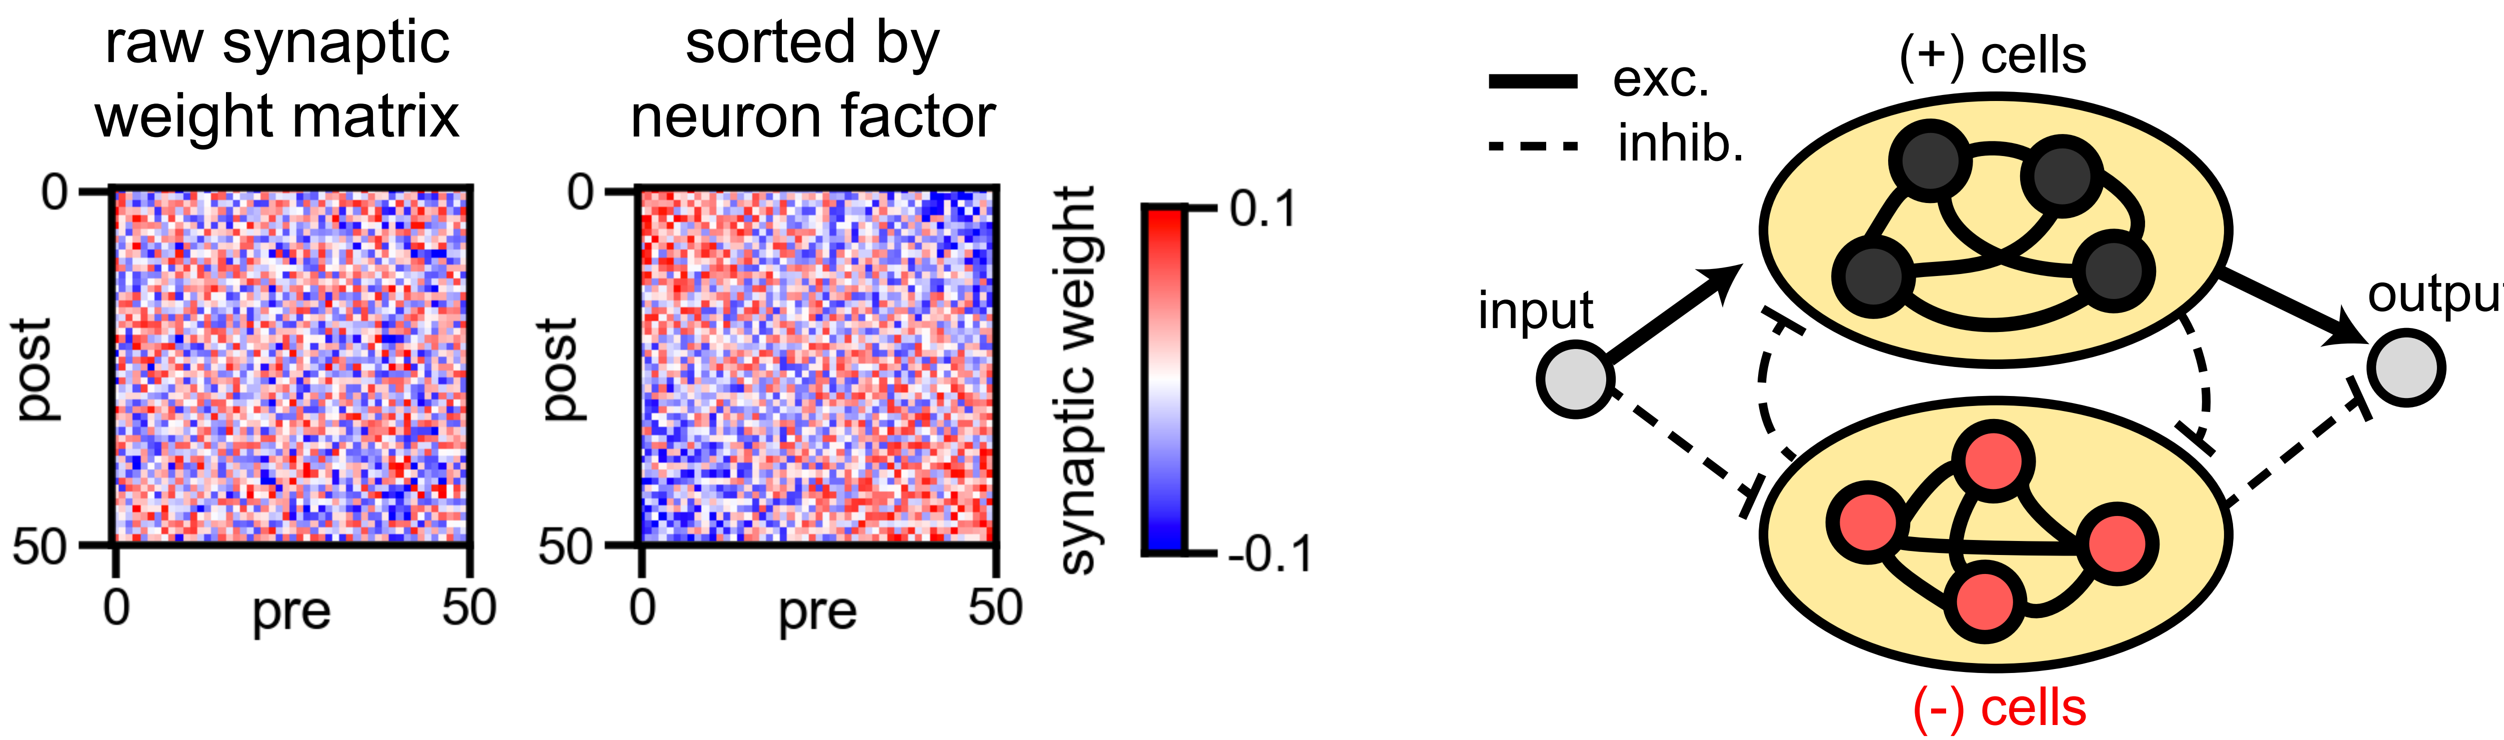

Application #2: How does prefrontal cortex encode place, actions, and rewards during maze navigation?

Application #2: How does prefrontal cortex encode place, actions, and rewards during maze navigation?

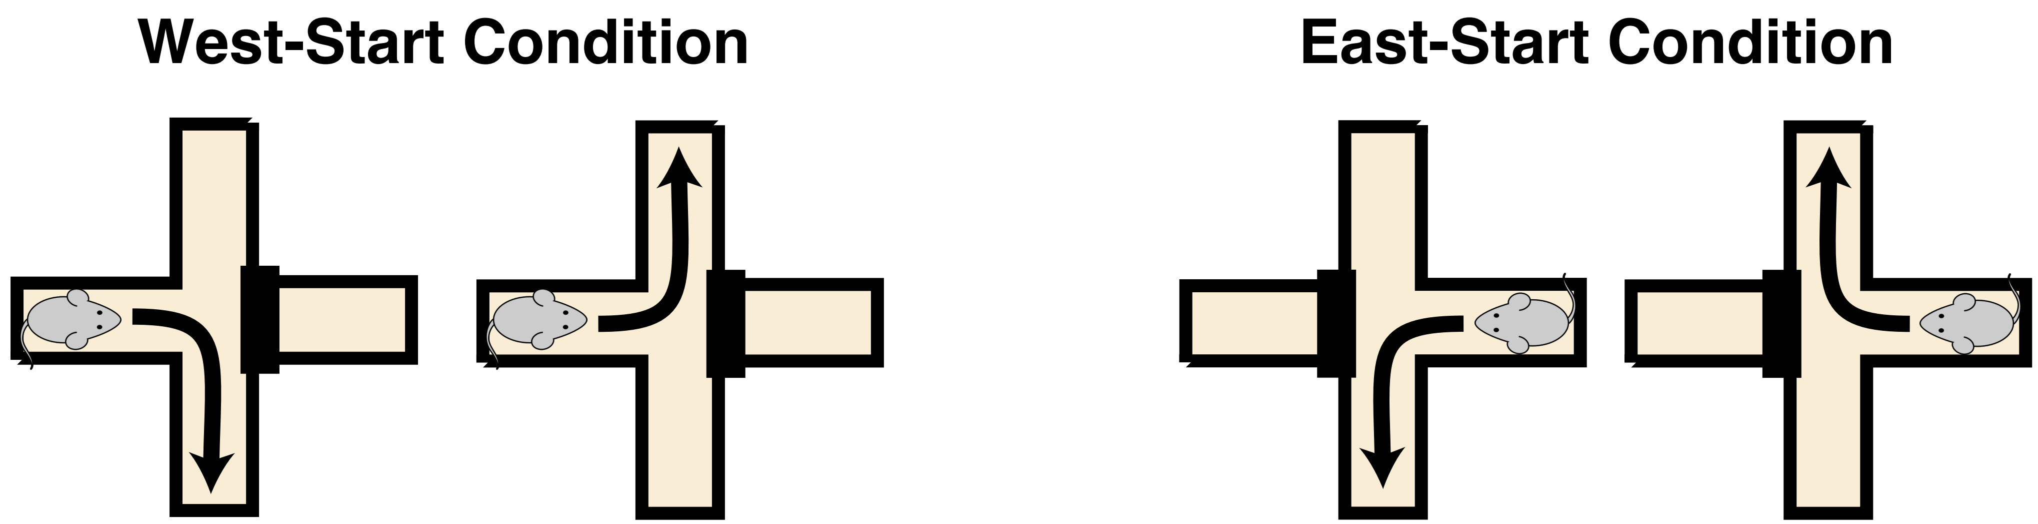

Trial Condition **Decision** **Trial Outcome**

East / West → North / South → Rewarded / Error

Application #2: How does prefrontal cortex encode place, actions, and rewards during maze navigation?

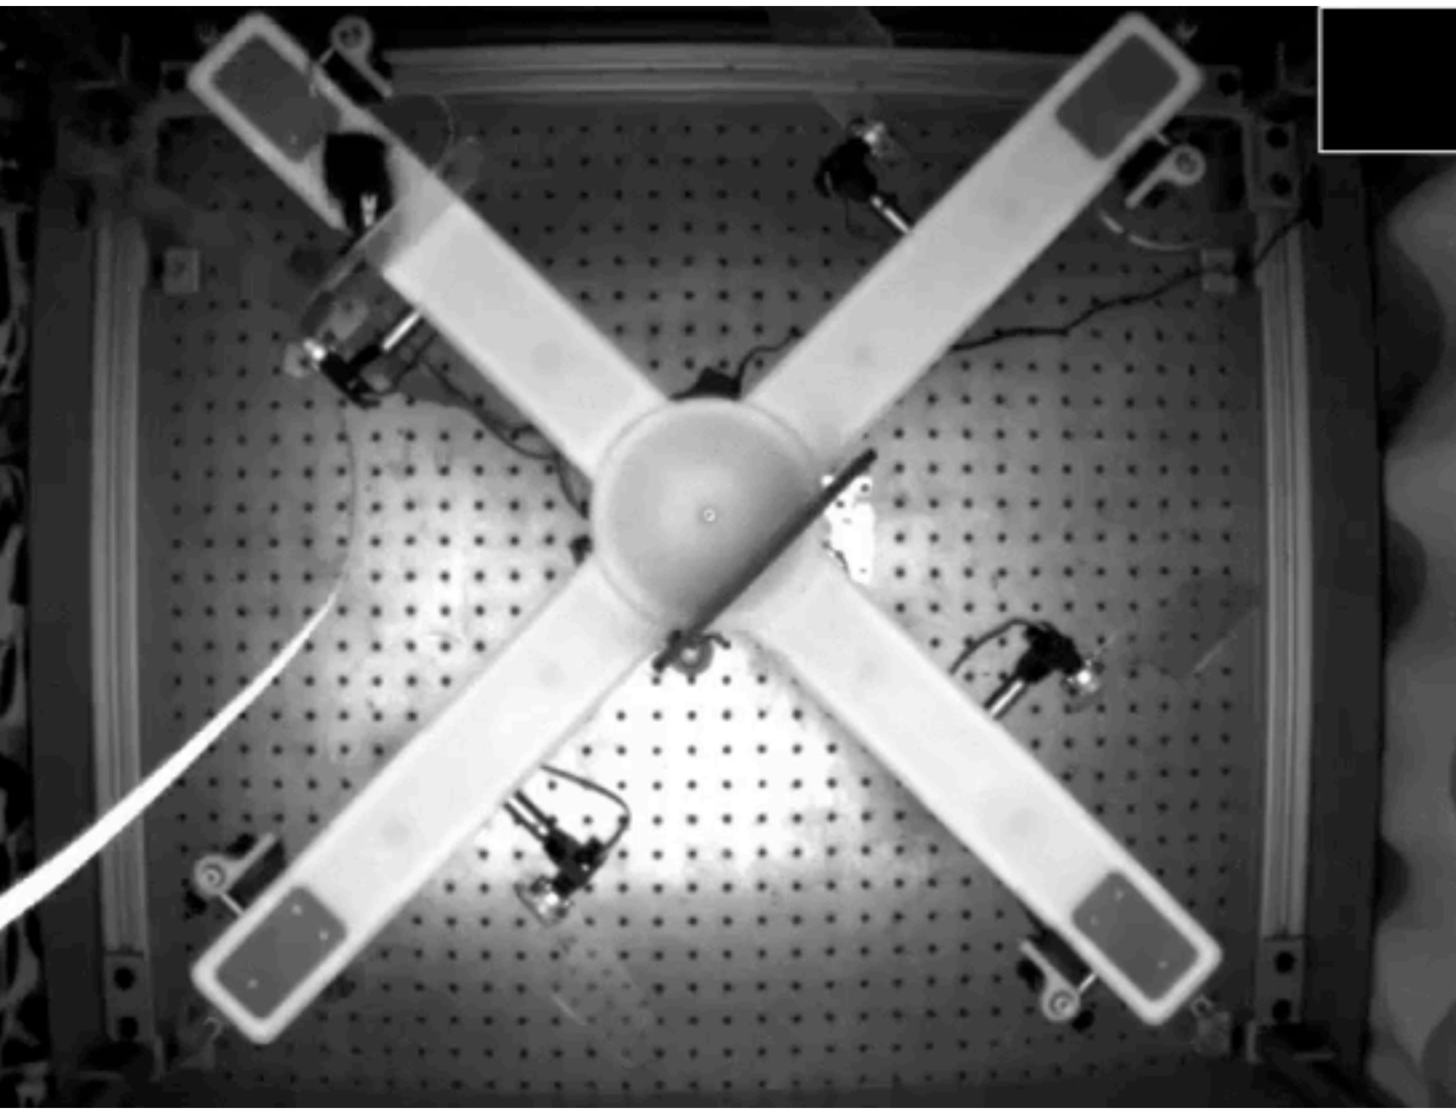

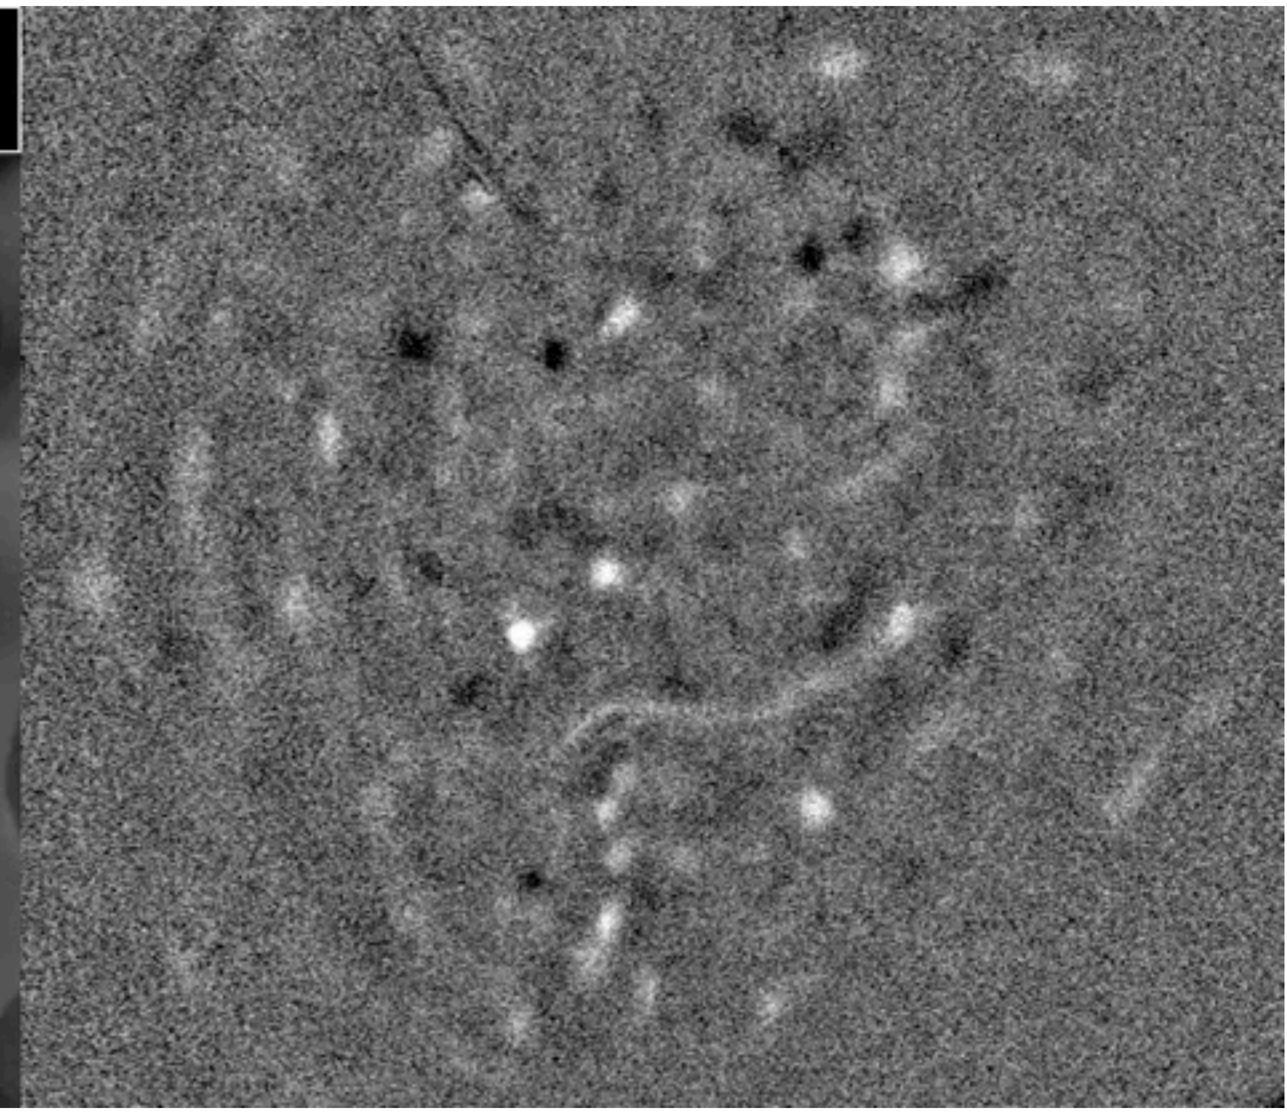

TCA (gain modulation) is a very compact and accurate model for trial-to-trial variability

two example cells encoding start location

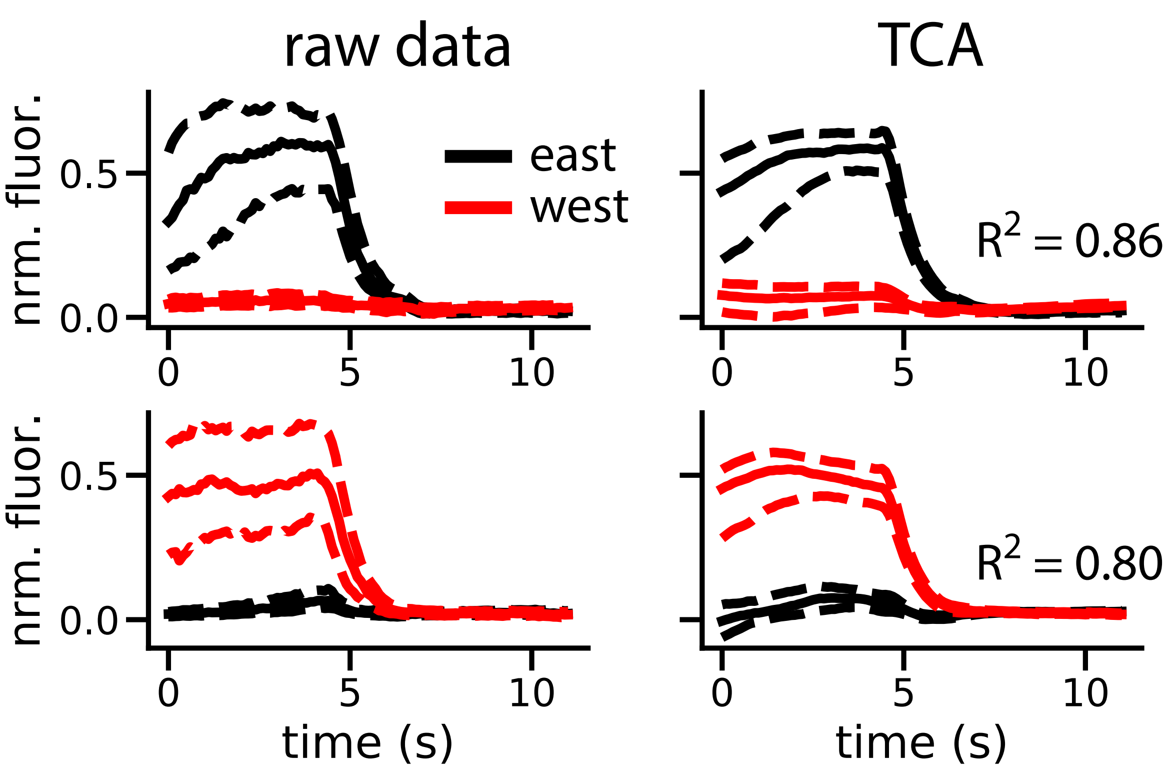

TCA (gain modulation) is a very compact and accurate model for trial-to-trial variability

two example cells encoding start location

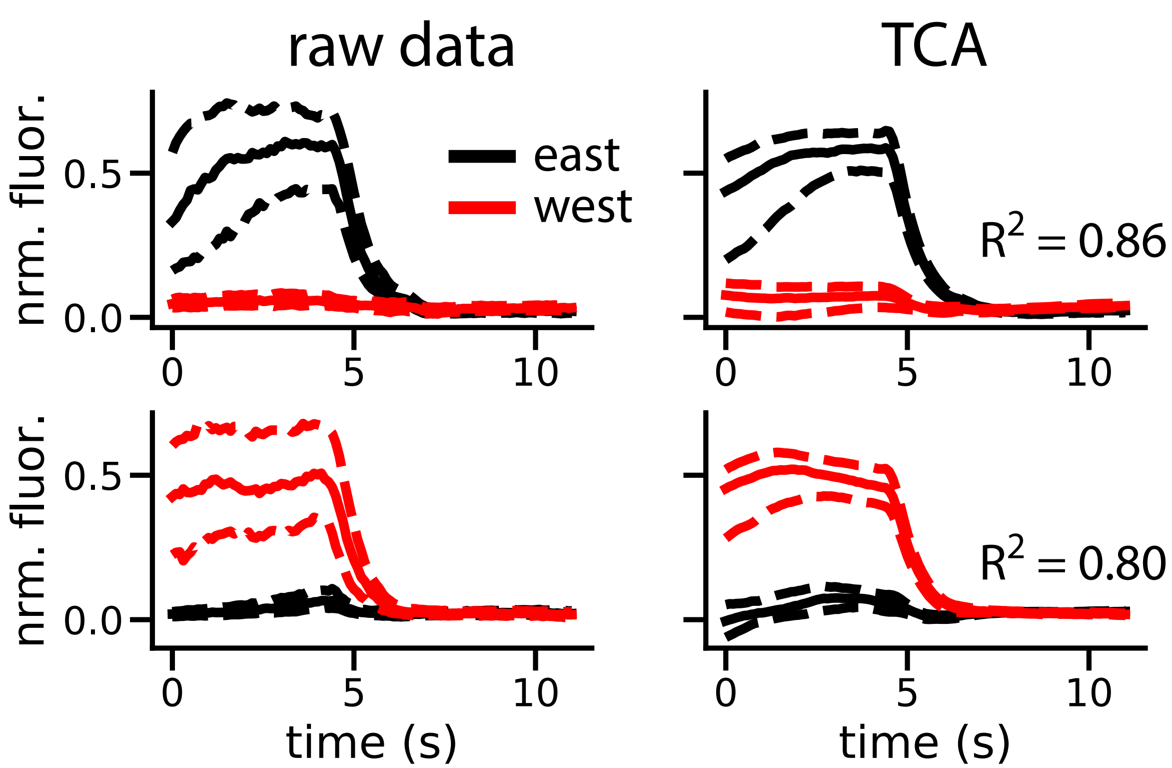

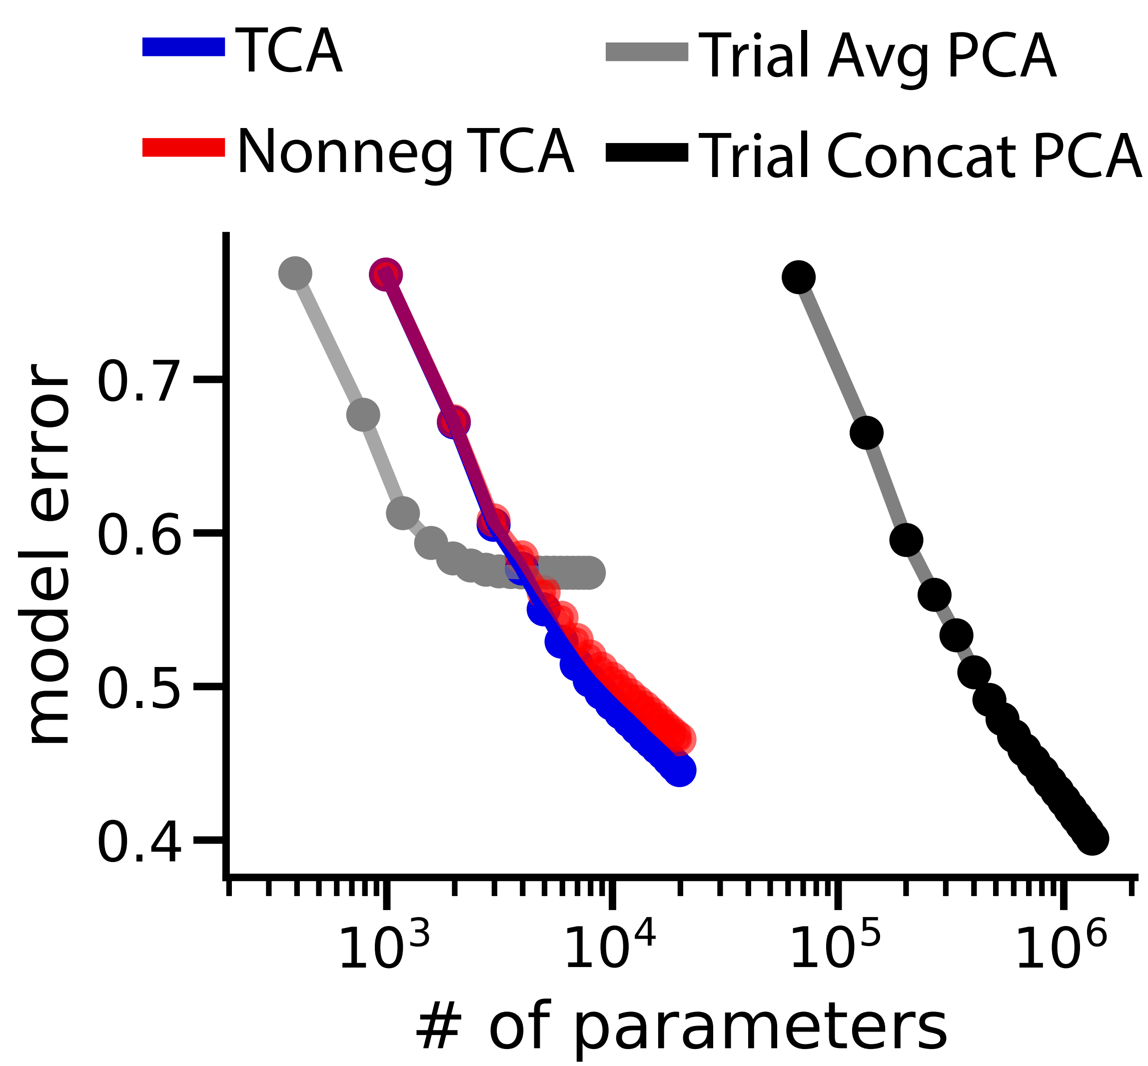

TCA factors map on to individual task variables

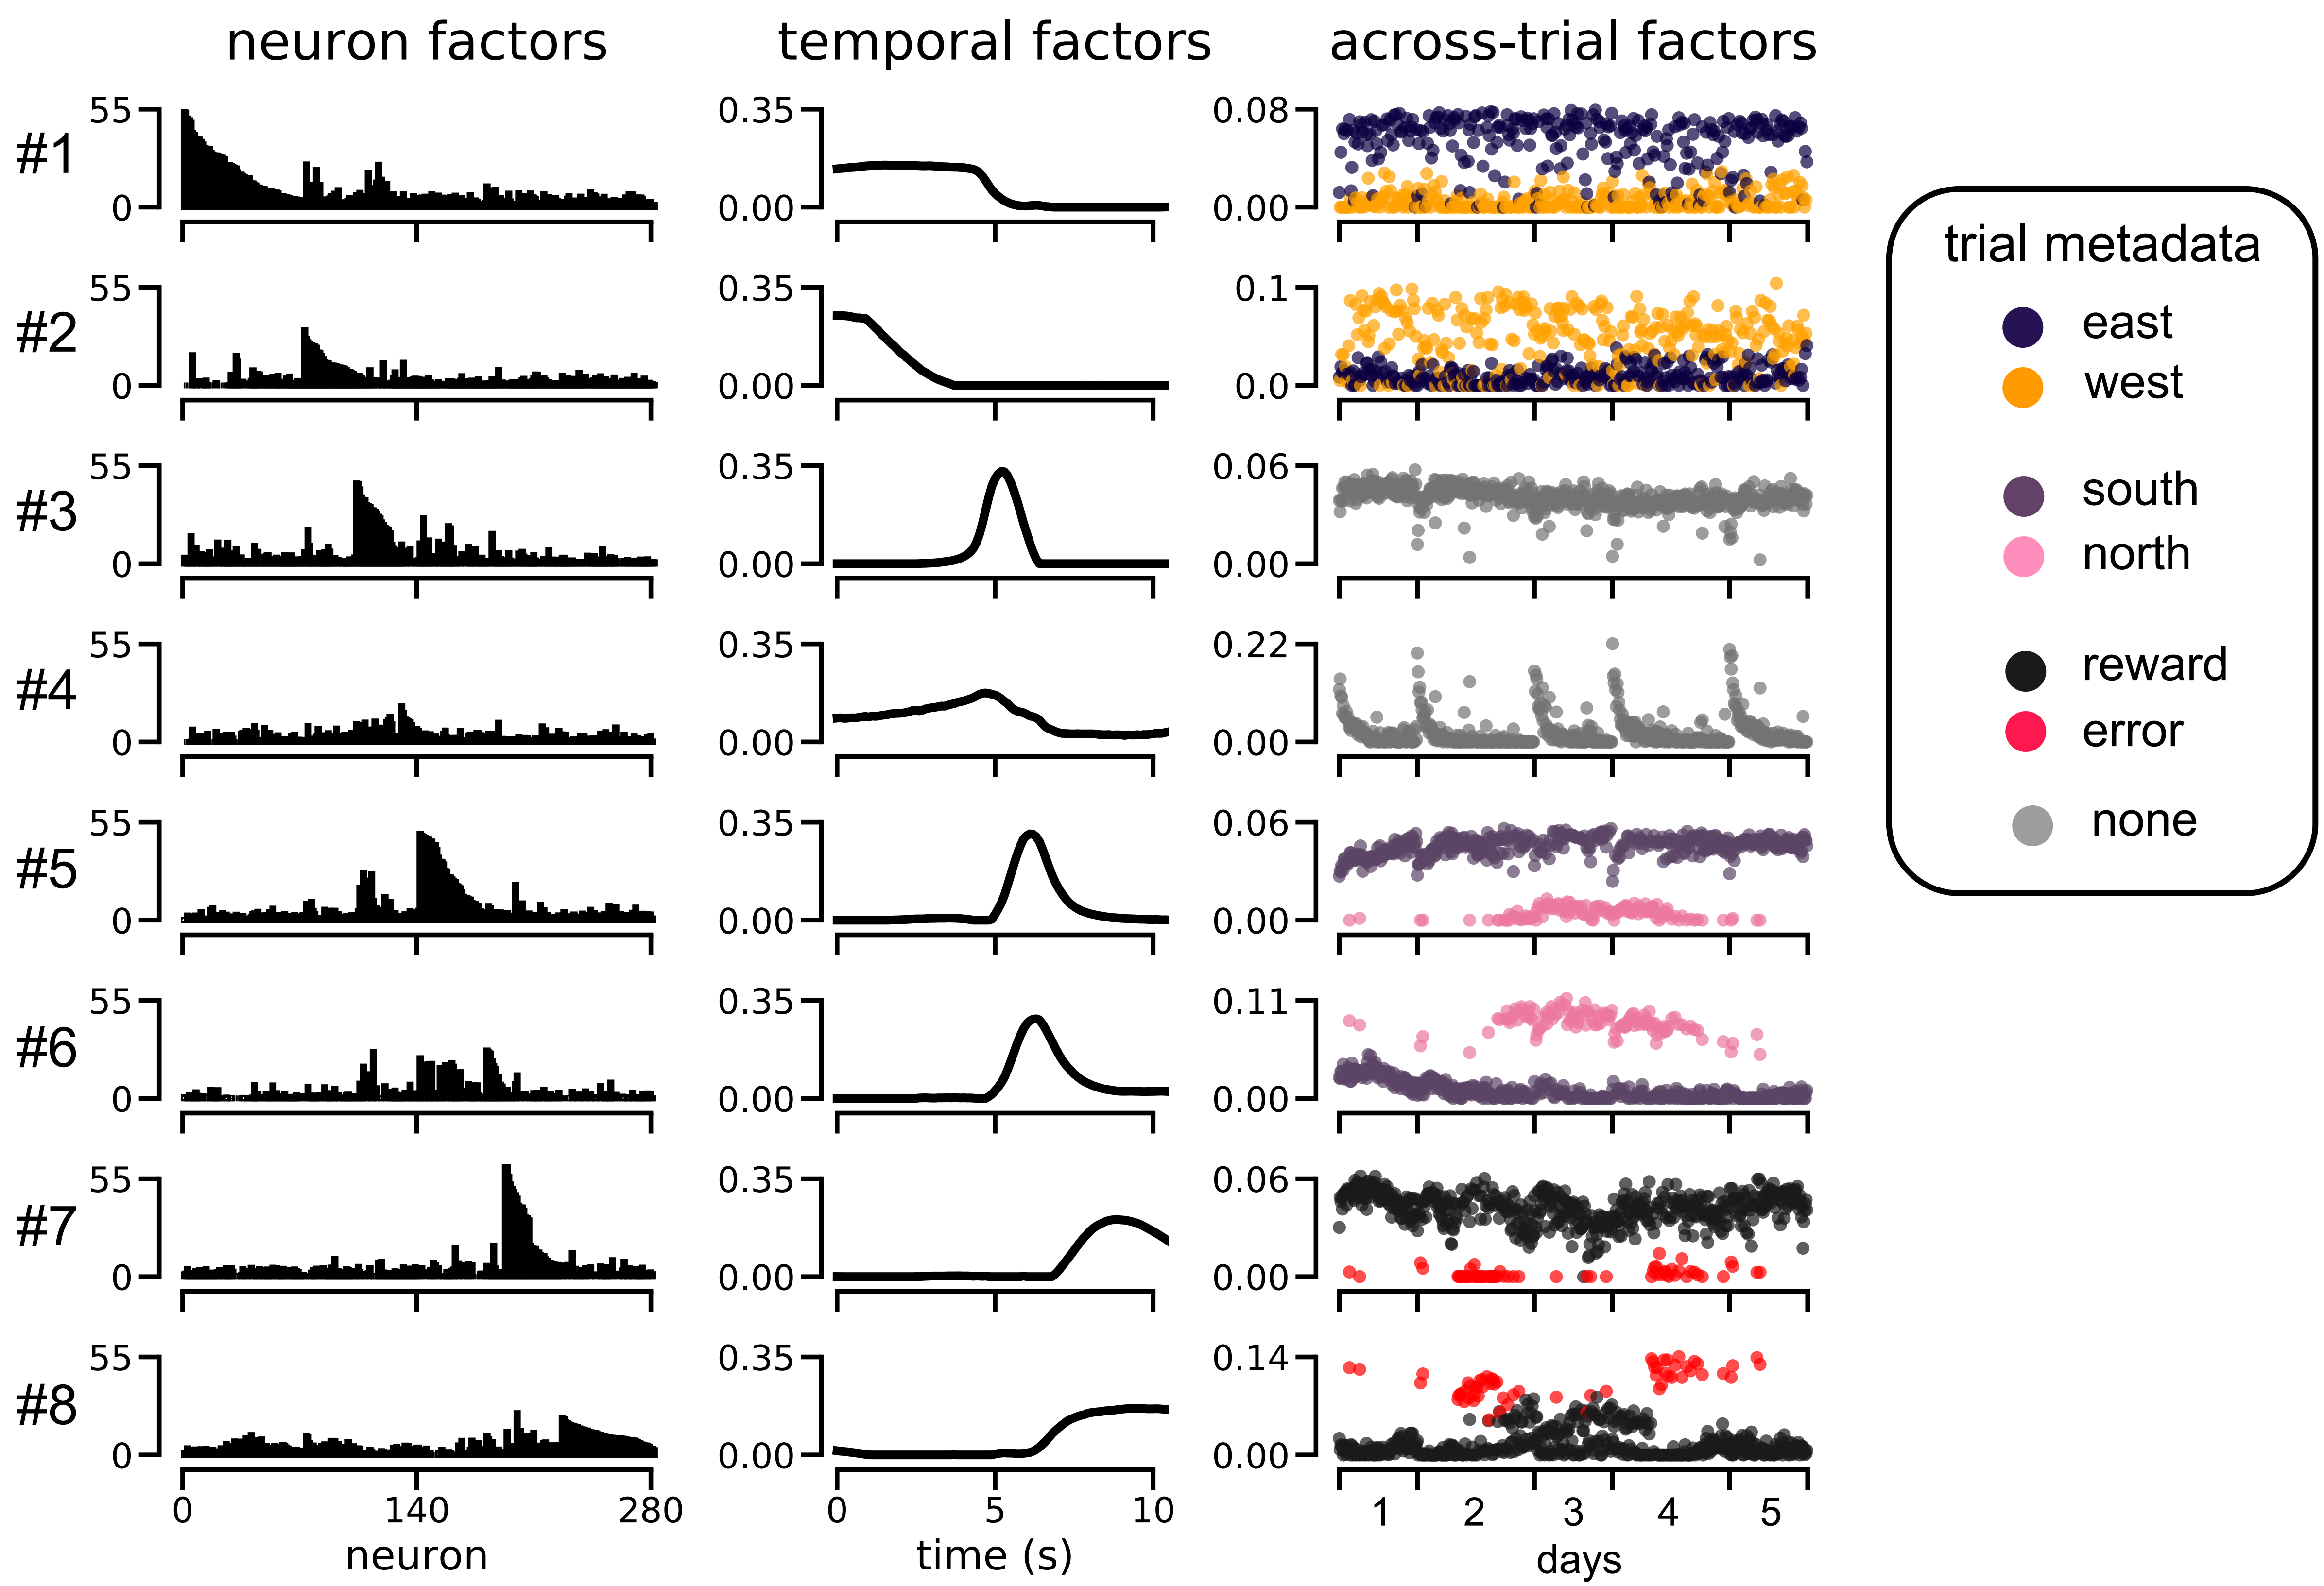

PCA components encode complex mixtures of task variables

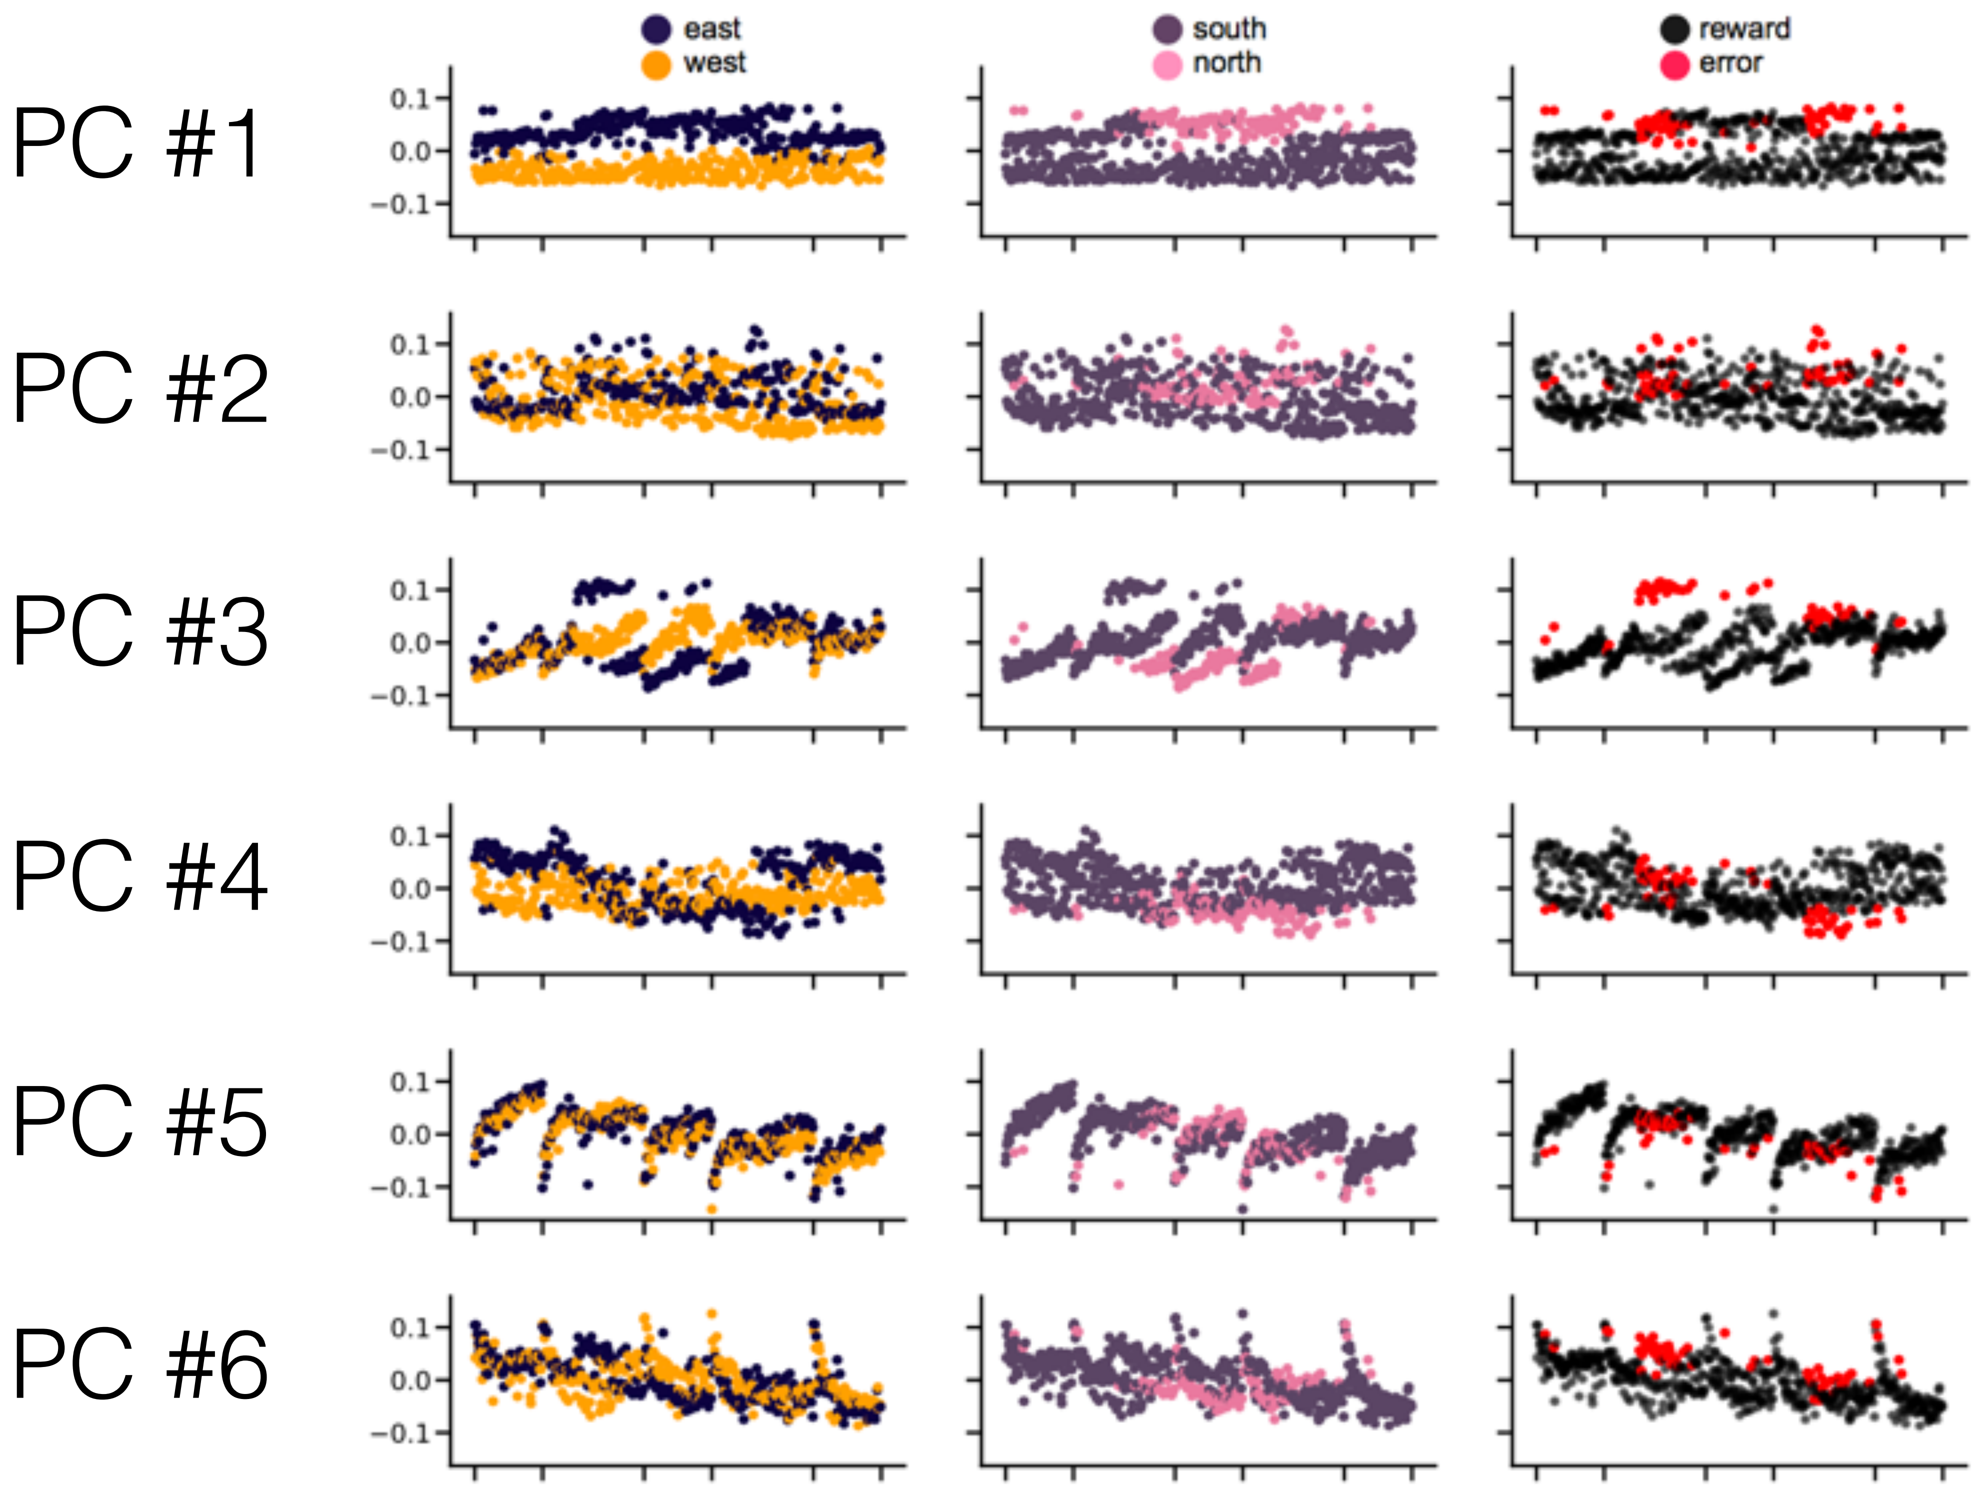

Application #3: How does motor cortex learn to control a cursor via a brain-machine interface?

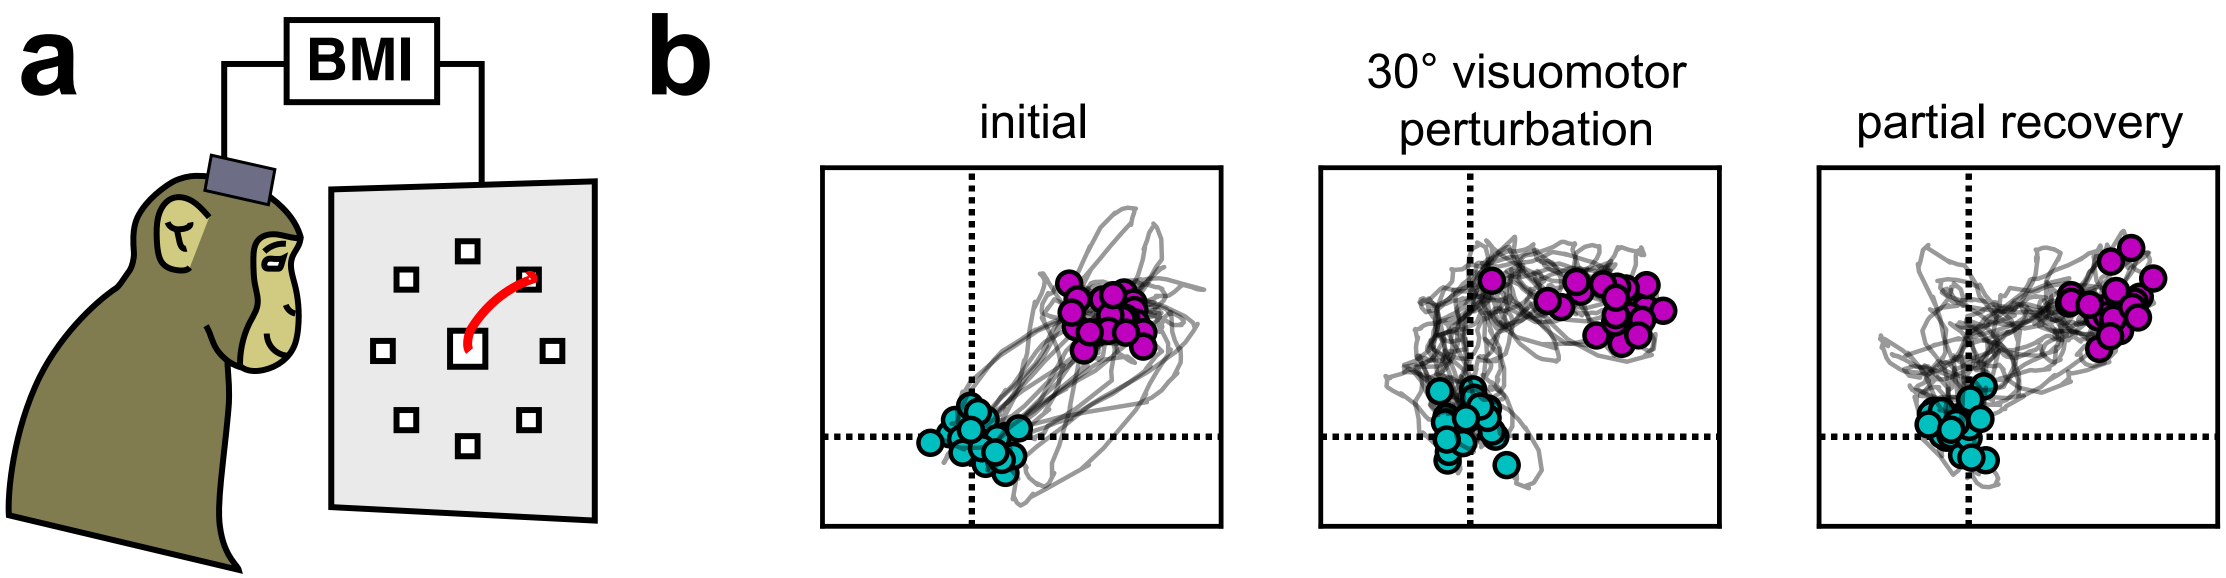

TCA identifies:

1. An **early component**, capturing the initial performance
2. A **compensatory component**, capturing within-trial corrections.
3. A **learned component**, capturing new neural dynamics that persist as the monkey adapts to the new BMI decoder.

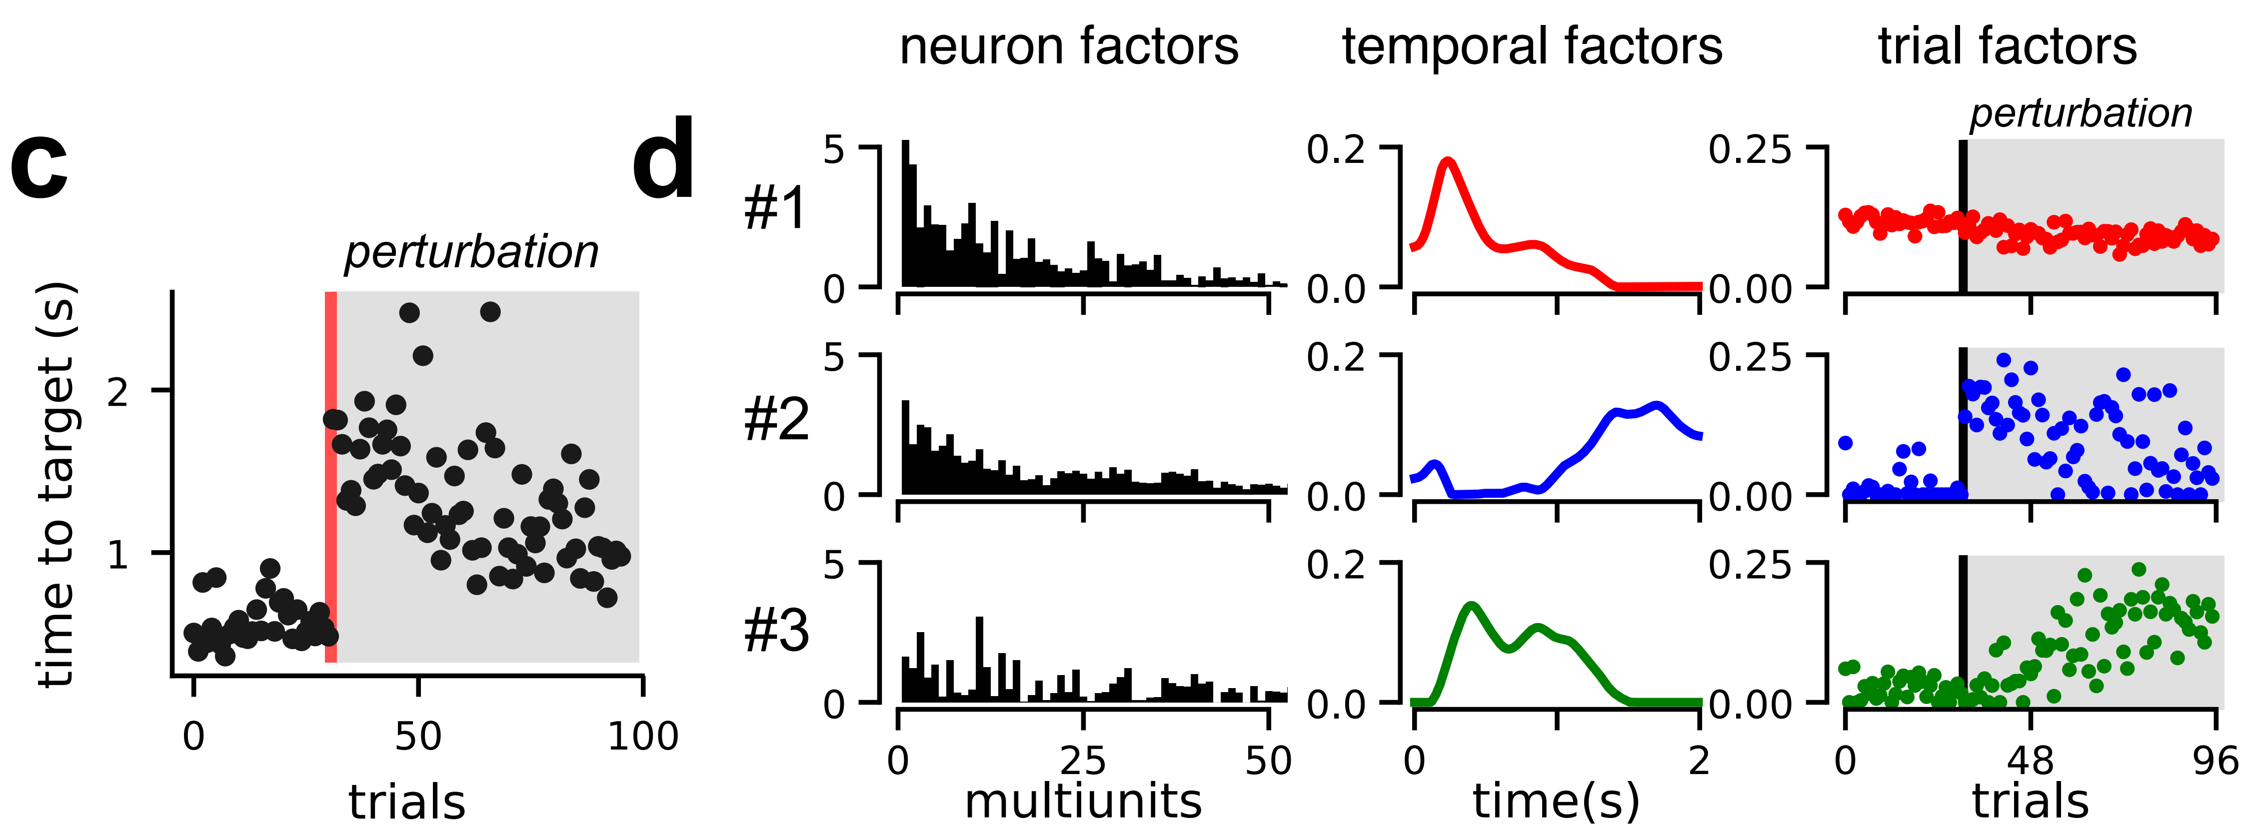

TCA identifies:

1. An **early component**, capturing the initial performance
2. A **compensatory component**, capturing within-trial corrections.
3. A **learned component**, capturing new neural dynamics that persist as the monkey adapts to the new BMI decoder.

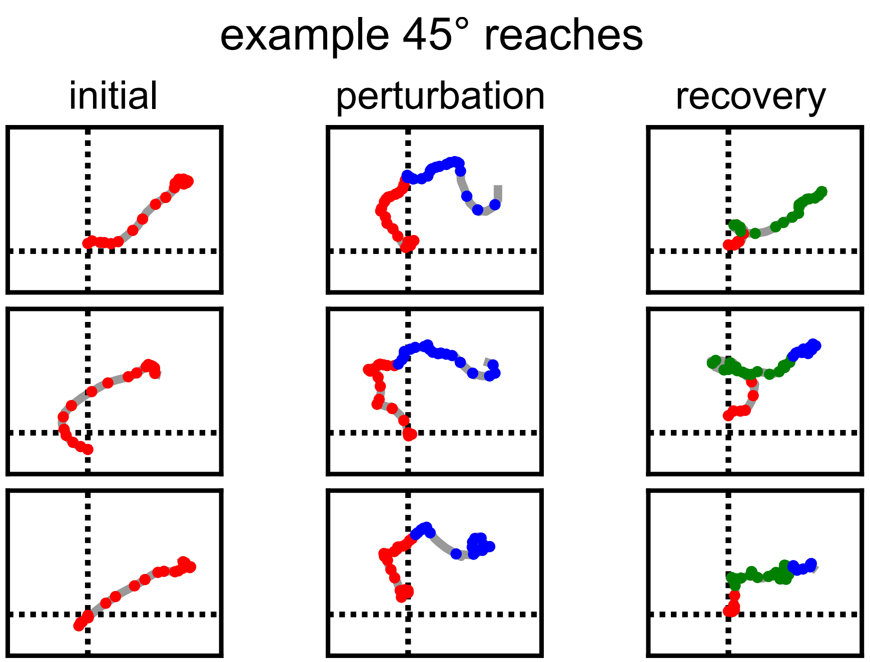

Tuning curves for each component

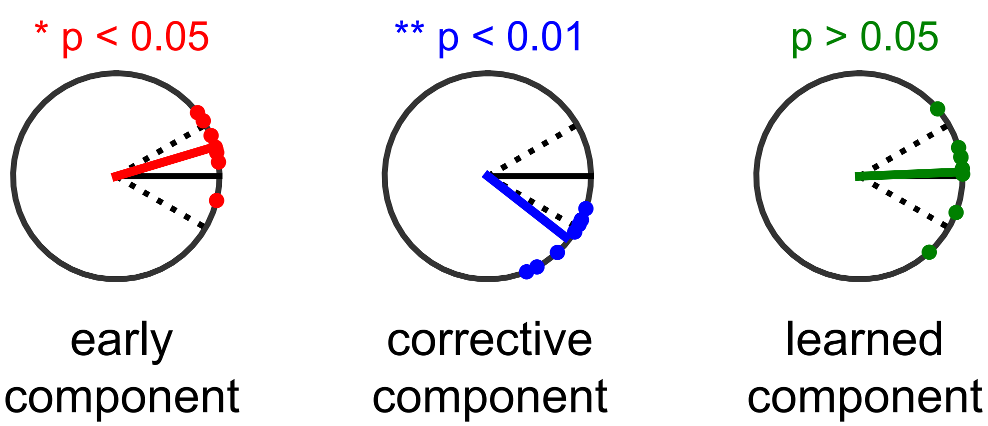

Summary

- There is growing interest in both *large-scale* and *long-term* recordings. TCA is can be used “*off-the-shelf”* on these data.

Summary

- There is growing interest in both *large-scale* and *long-term* recordings. TCA is can be used “*off-the-shelf”* on these data.
- TCA has nice theoretical properties (Kolda & Bader, 2009) and an intuitive interpretation as a gain-modulated network.

Summary

- There is growing interest in both *large-scale* and *long-term* recordings. TCA is can be used "*off-the-shelf*" on these data.
- TCA has nice theoretical properties (Kolda & Bader, 2009) and an intuitive interpretation as a **gain-modulated network**.

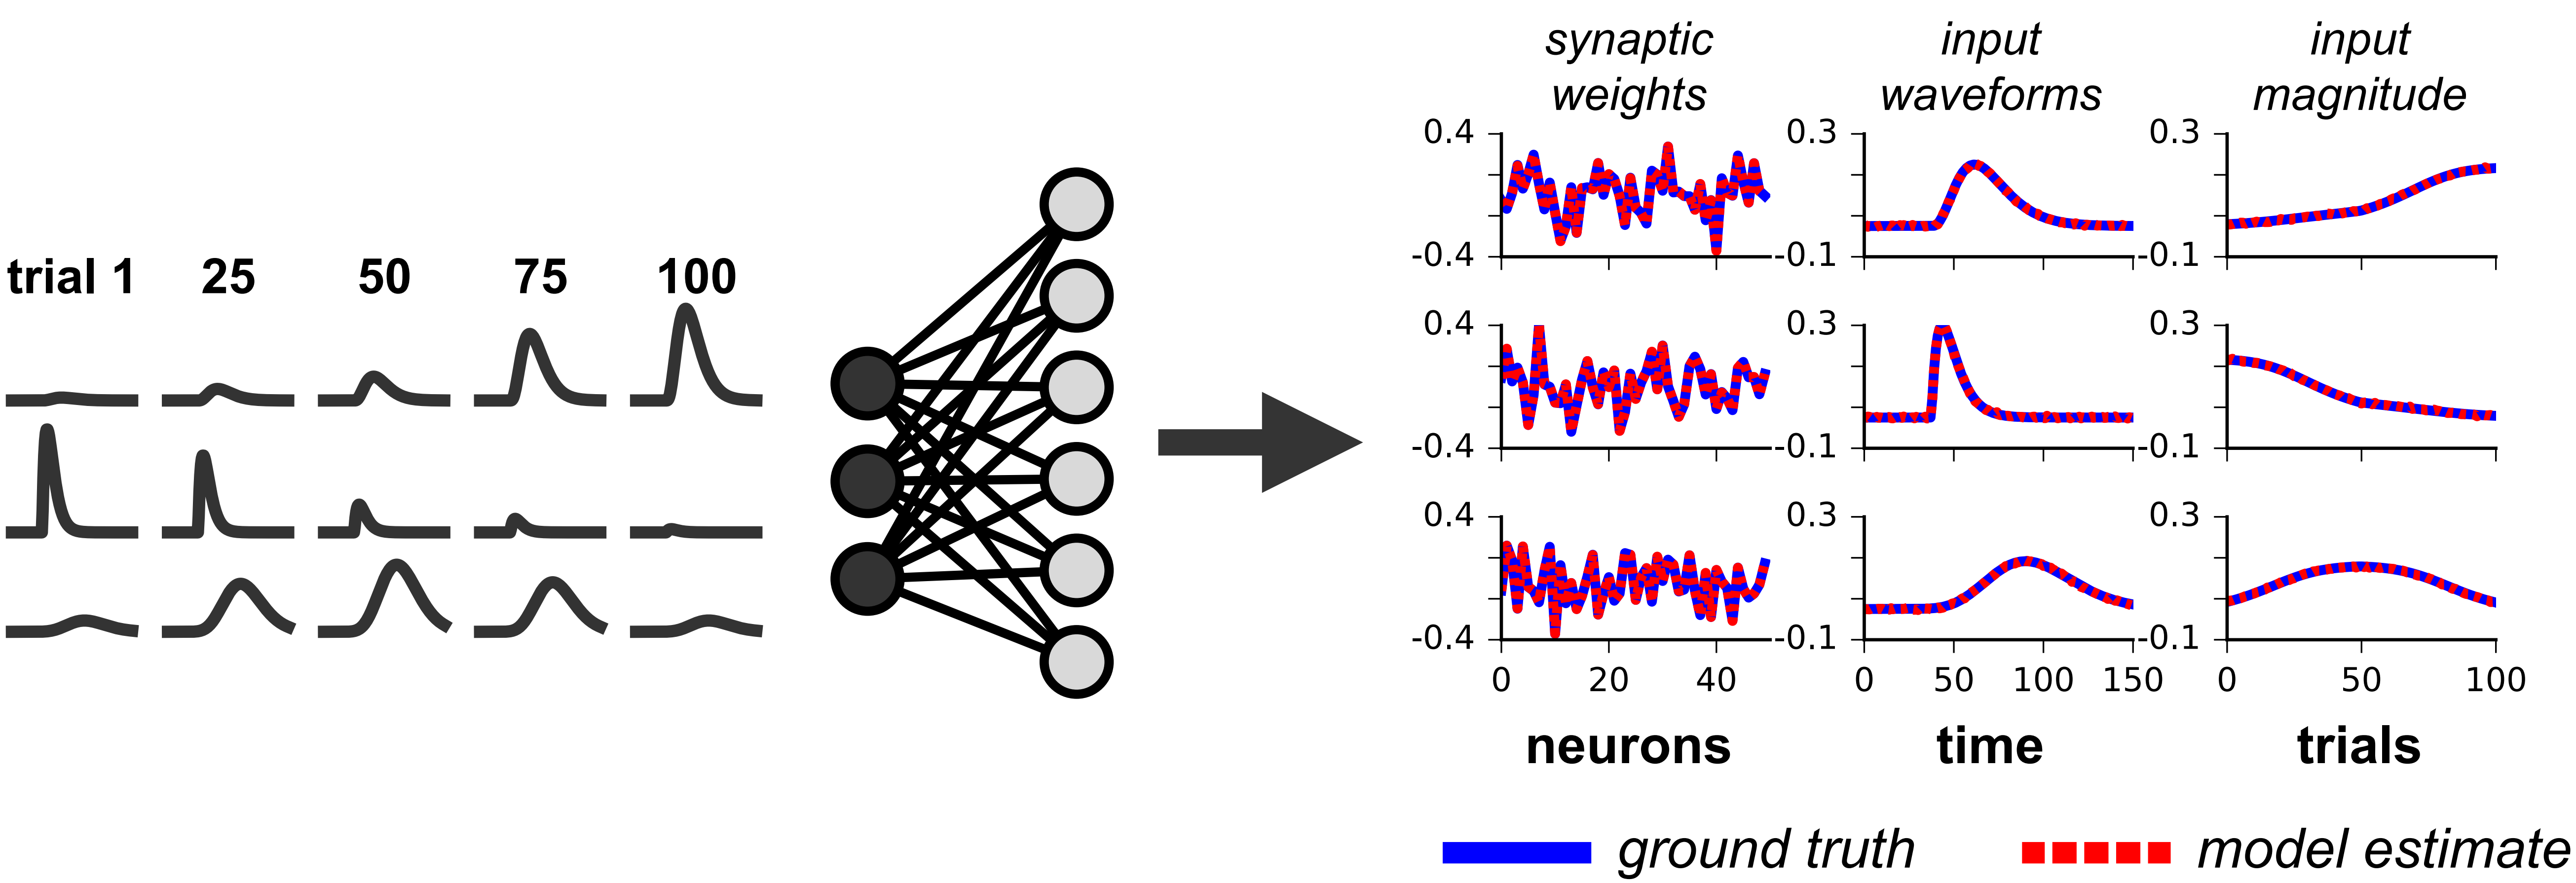

Summary

- There is growing interest in both *large-scale* and *long-term* recordings. TCA is can be used "*off-the-shelf*" on these data.
- TCA has nice theoretical properties (Kolda & Bader, 2009) and an intuitive interpretation as a **gain-modulated network**.
- Can reveal **discrete clusters** of neurons and single-trial dynamics.

discrete coding of place and reward in mPFC

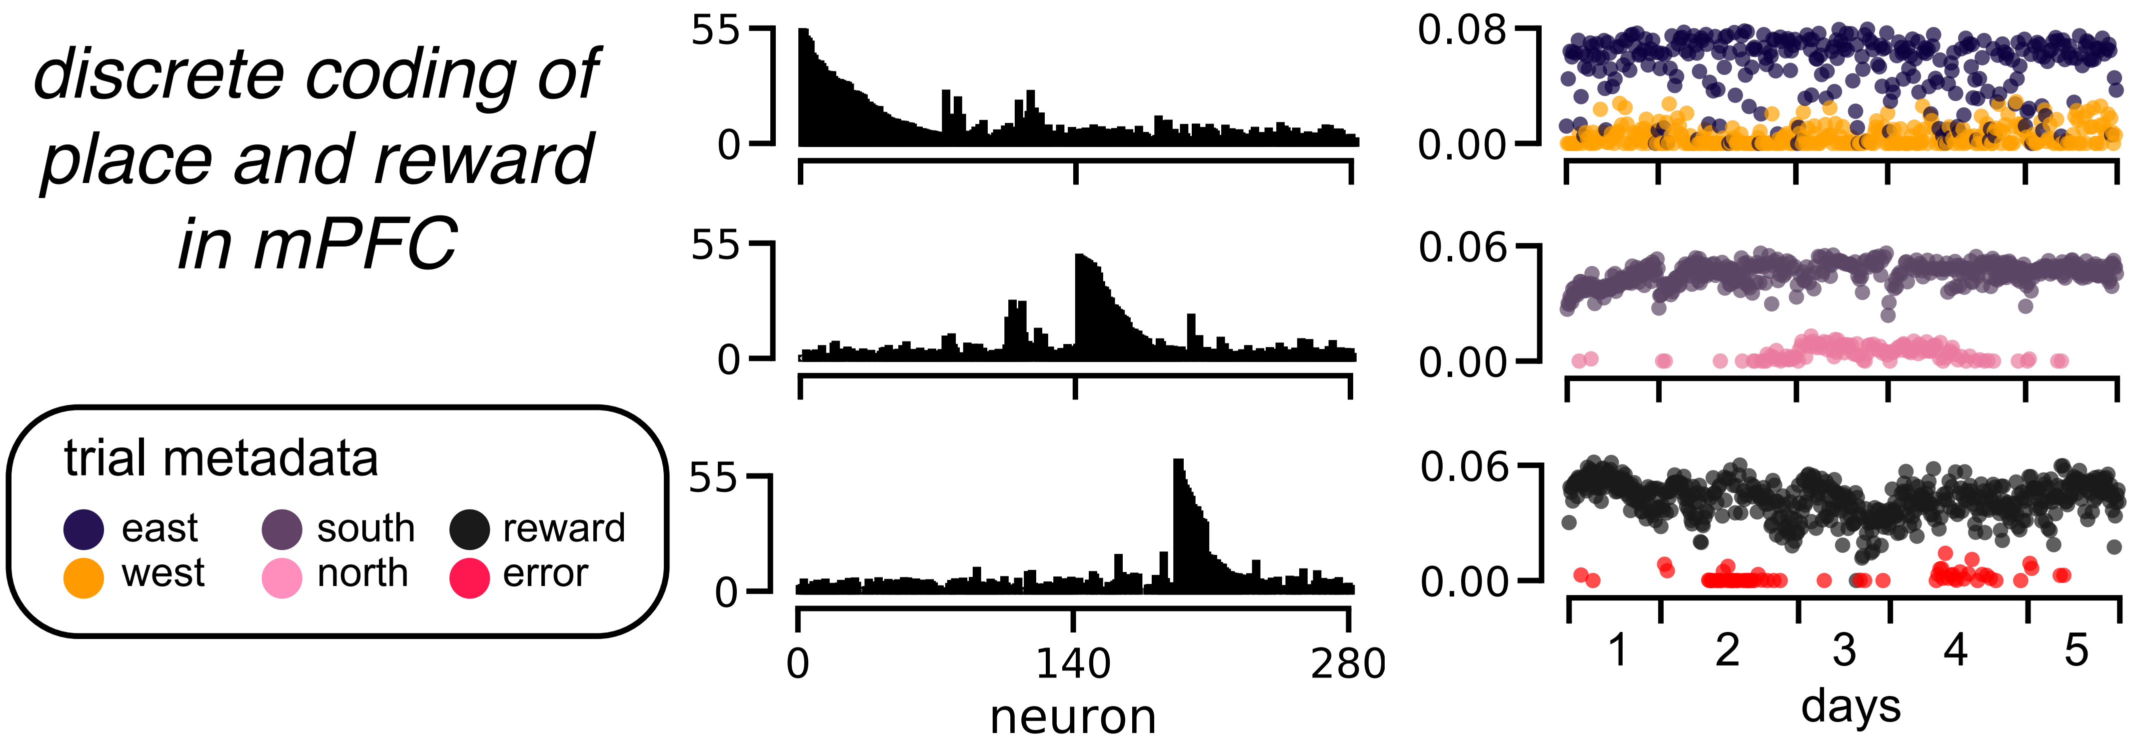

Summary

- There is growing interest in both *large-scale* and *long-term* recordings. TCA is can be used "*off-the-shelf"* on these data.
- TCA has nice theoretical properties (Kolda & Bader, 2009) and an intuitive interpretation as a **gain-modulated network**.
- Can reveal **discrete clusters** of neurons and single-trial dynamics.
- TCA can also pull out **gradual trends** within neural populations and across-trial learning dynamics

Summary

- There is growing interest in both *large-scale* and *long-term* recordings. TCA is can be used “*off-the-shelf”* on these data.
- TCA has nice theoretical properties (Kolda & Bader, 2009) and an intuitive interpretation as a **gain-modulated network**.
- Can reveal **discrete clusters** of neurons and single-trial dynamics.
- TCA can also pull out **gradual trends** within neural populations and across-trial learning dynamics

cell types in a RNN

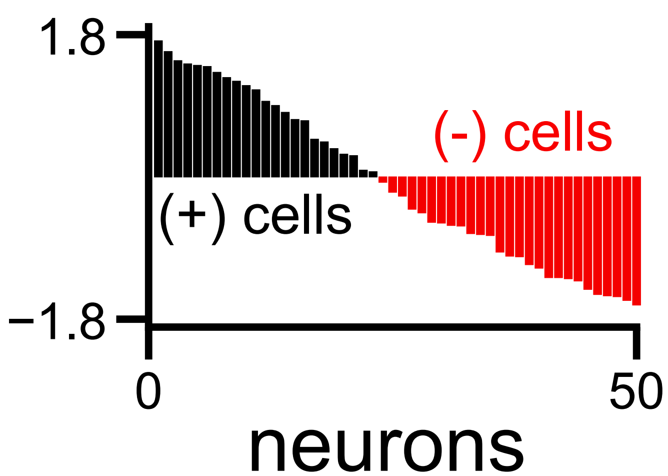

learning in motor cortex

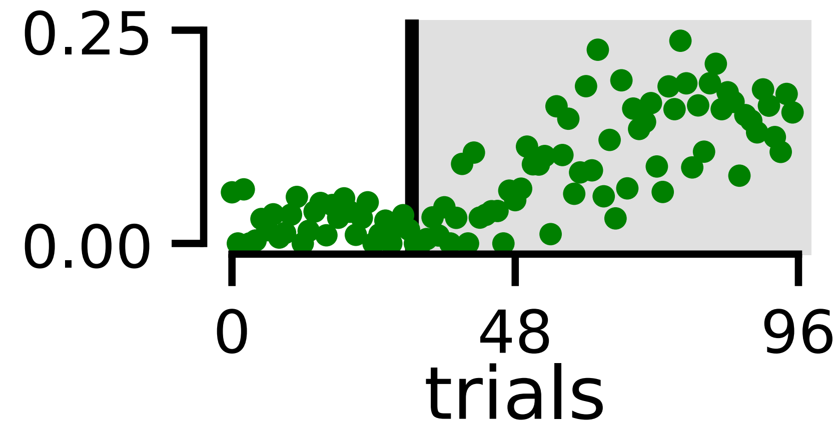

Summary

- There is growing interest in both *large-scale* and *long-term* recordings. TCA is can be used "*off-the-shelf"* on these data.
- TCA has nice theoretical properties (Kolda & Bader, 2009) and an intuitive interpretation as a **gain-modulated network**.
- Can reveal **discrete clusters** of neurons and single-trial dynamics.
- TCA can also pull out **gradual trends** within neural populations and across-trial learning dynamics

Overall, TCA is a well-motivated tool for identifying structure in some of the most challenging datasets in neuroscience involving large-scale & long-term neural recordings.

Try out TCA!

Python

https://github.com/ahwillia/tensortools

https://tensorly.github.io/

MATLAB

http://www.sandia.gov/~tgkolda/TensorToolbox/

https://www.tensorlab.net/

Julia

https://github.com/yunjhongwu/TensorDecompositions.jl

Contact : ahwillia@stanford.edu

Slides : alexhwilliams.info/pdf/nccd.pdf

Code : github.com/ahwillia/tensor-demo

STANFORD NEUROSCIENCE

Co-authors:

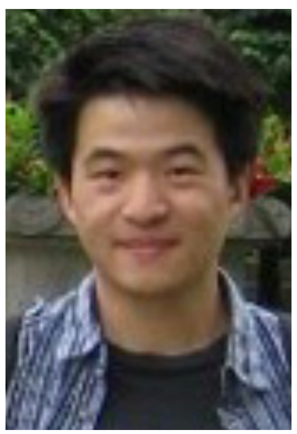

Tony Kim

Fori Wang

Mark Schnitzer

Saurabh Vyas

Krishna Shenoy

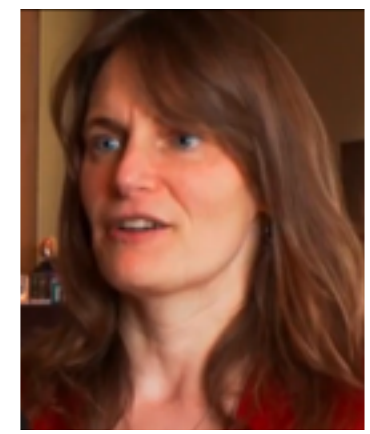

Tammy Kolda

Surya Ganguli

Funding & Support:

DOE CSGF

Sandia National Laboratories

James S. McDonnell Foundation

The McKnight Foundation

Simons Foundation