

# Beneficial roles for chaotic variability in learning systems

Matthew Farrell<sup>1-2</sup>, Stefano Recanatesi<sup>1</sup>, Timothy Moore<sup>1</sup>, Guillaume Lajoie<sup>3-4</sup>, and Eric Shea-Brown<sup>1-2</sup>

<sup>1</sup>*Computational Neuroscience Center, University of Washington*

<sup>2</sup>*Dept. of Applied Mathematics, University of Washington*


<sup>3</sup>*Mila—Québec AI Institute*

<sup>4</sup>*Dept. of Mathematics and Statistics, Université de Montréal*

**Abstract.** Neural responses are highly variable, even under identical task conditions. Significant efforts are being directed toward explaining how the brain copes with and may even leverage such variability to help learn the task and environment. Here we explore the issue in a recurrent neural network model that is trained to classify inputs. We find two potential beneficial roles for chaotic variability in these systems: (1) chaos can accelerate the flexible relearning of a task after it is modified; and (2) chaos can lift the network representation of data into a higher-dimensional space, which allows the network to classify inputs embedded in low-dimensional spaces.

Goal: investigate how the initial dynamics of a recurrent neural network (RNN) influence the solutions learned by the network when trained to do a task.

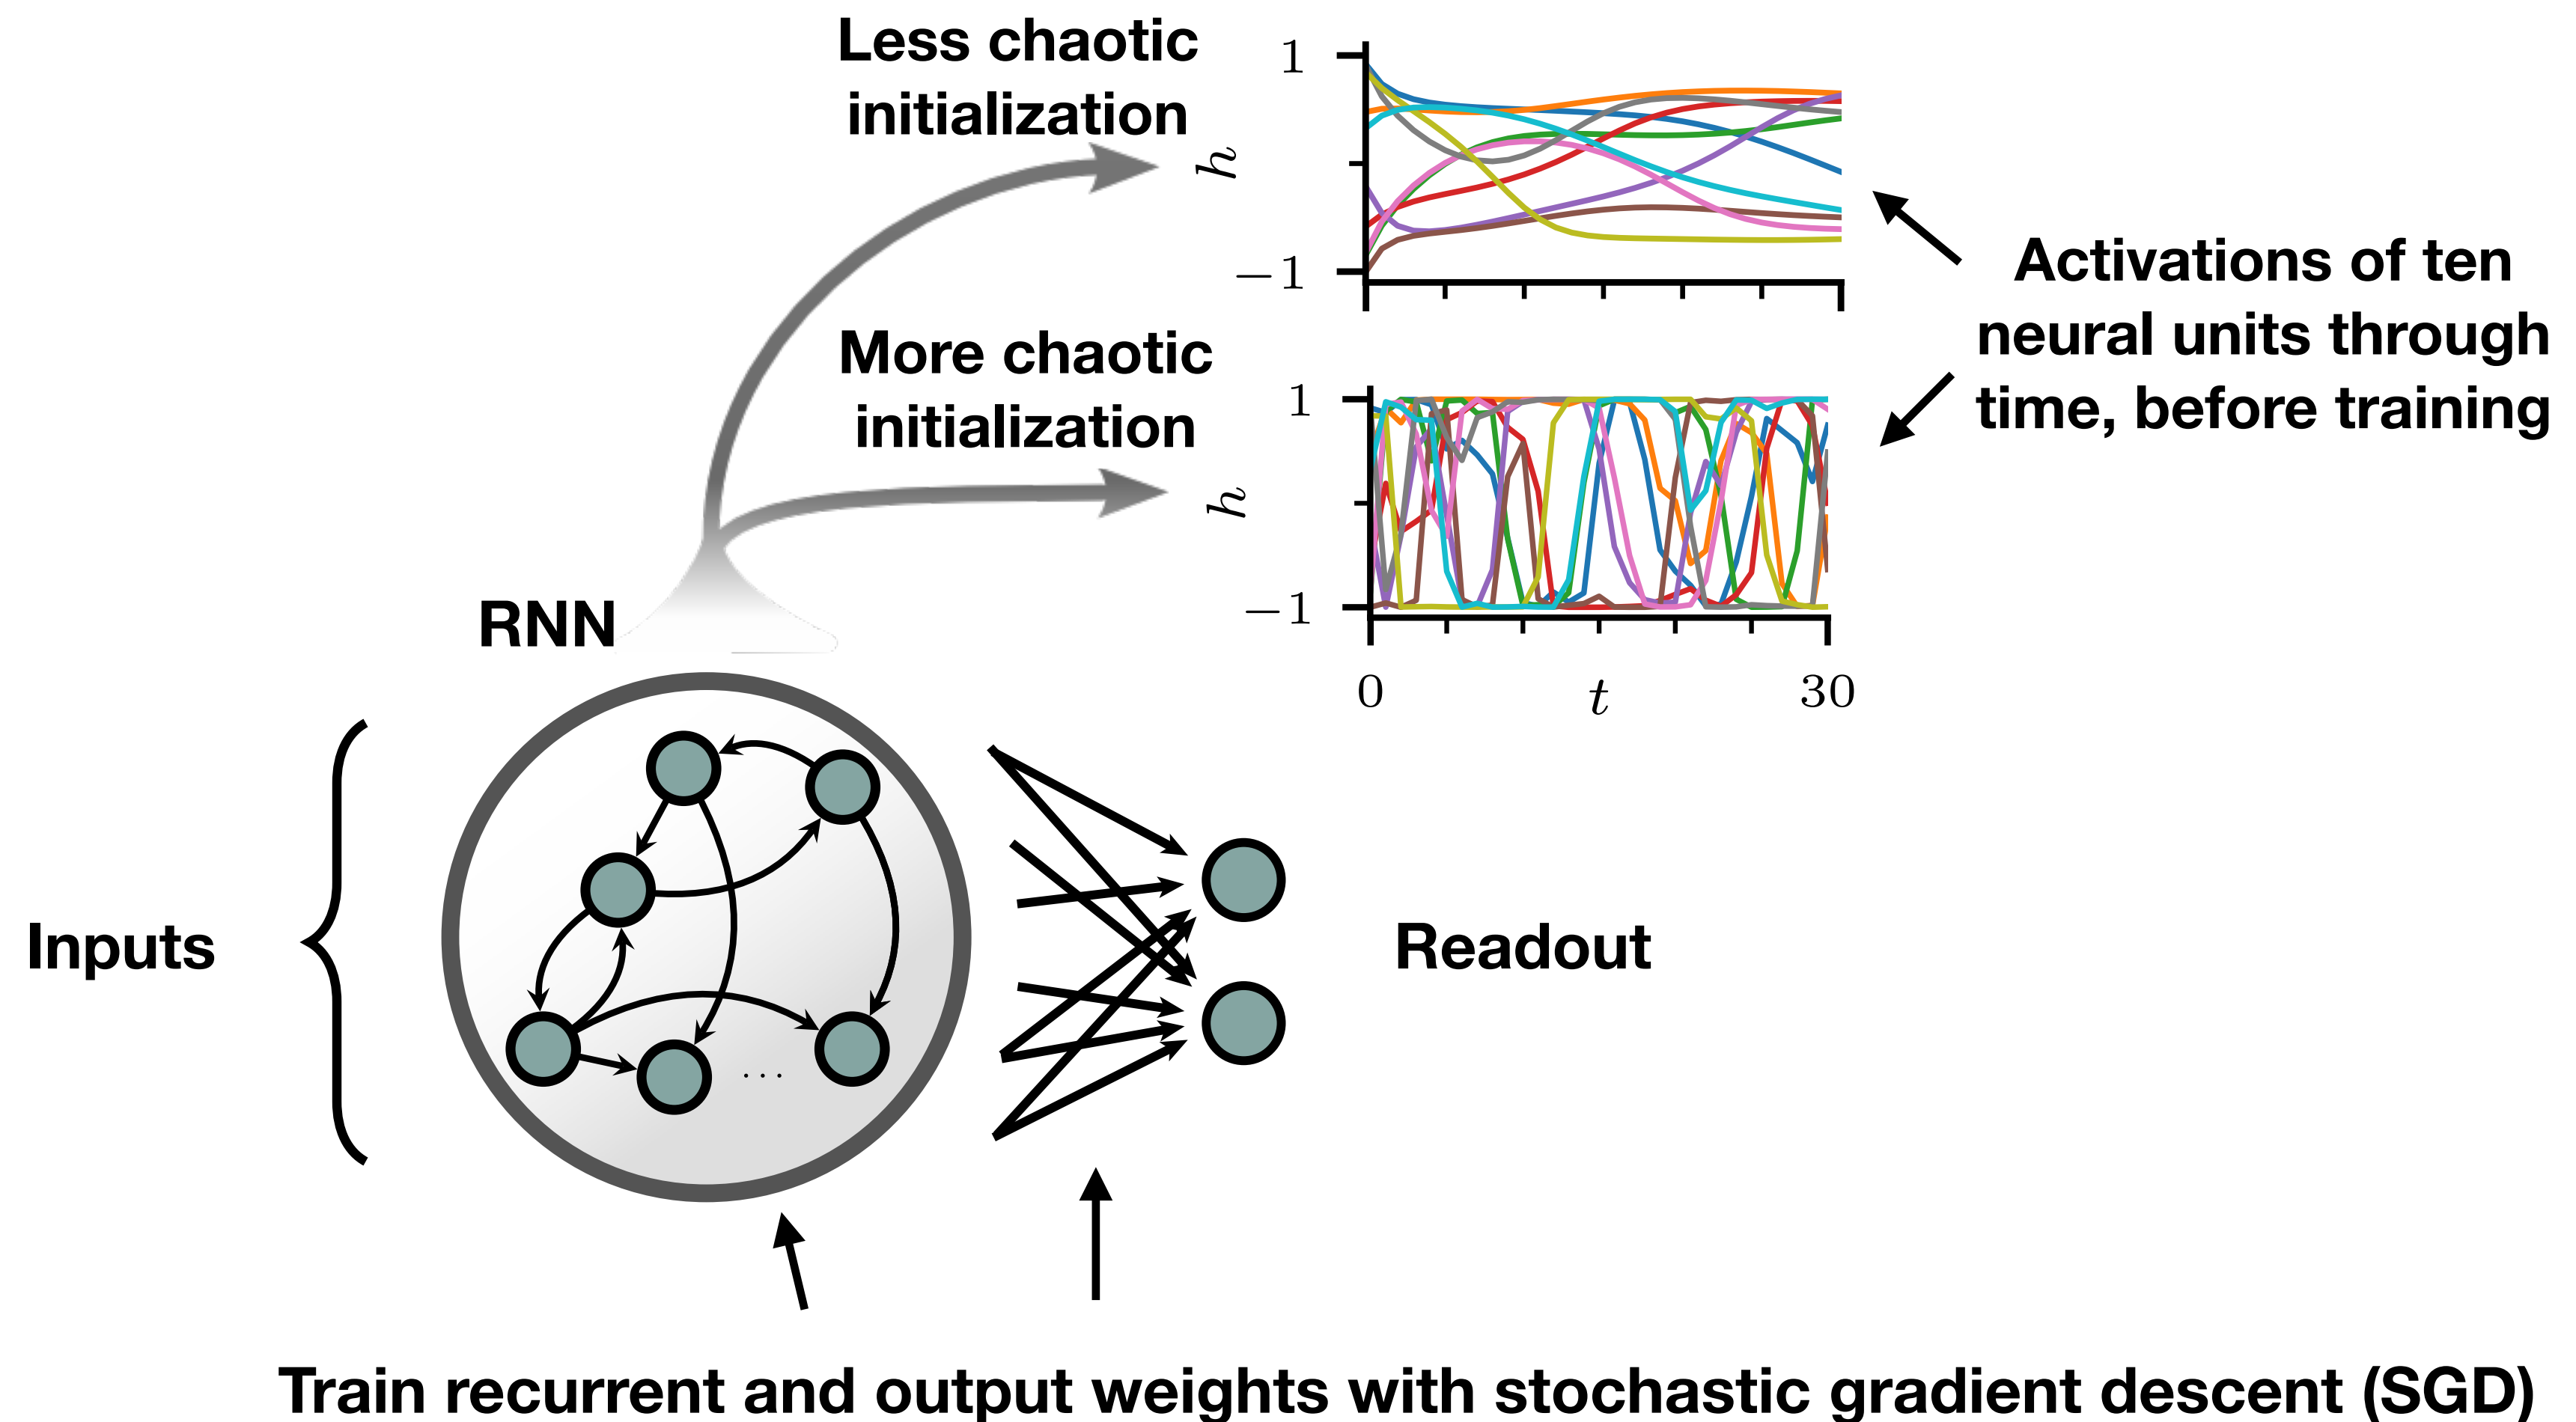
We characterize these dynamics by how chaotic the trajectories are.



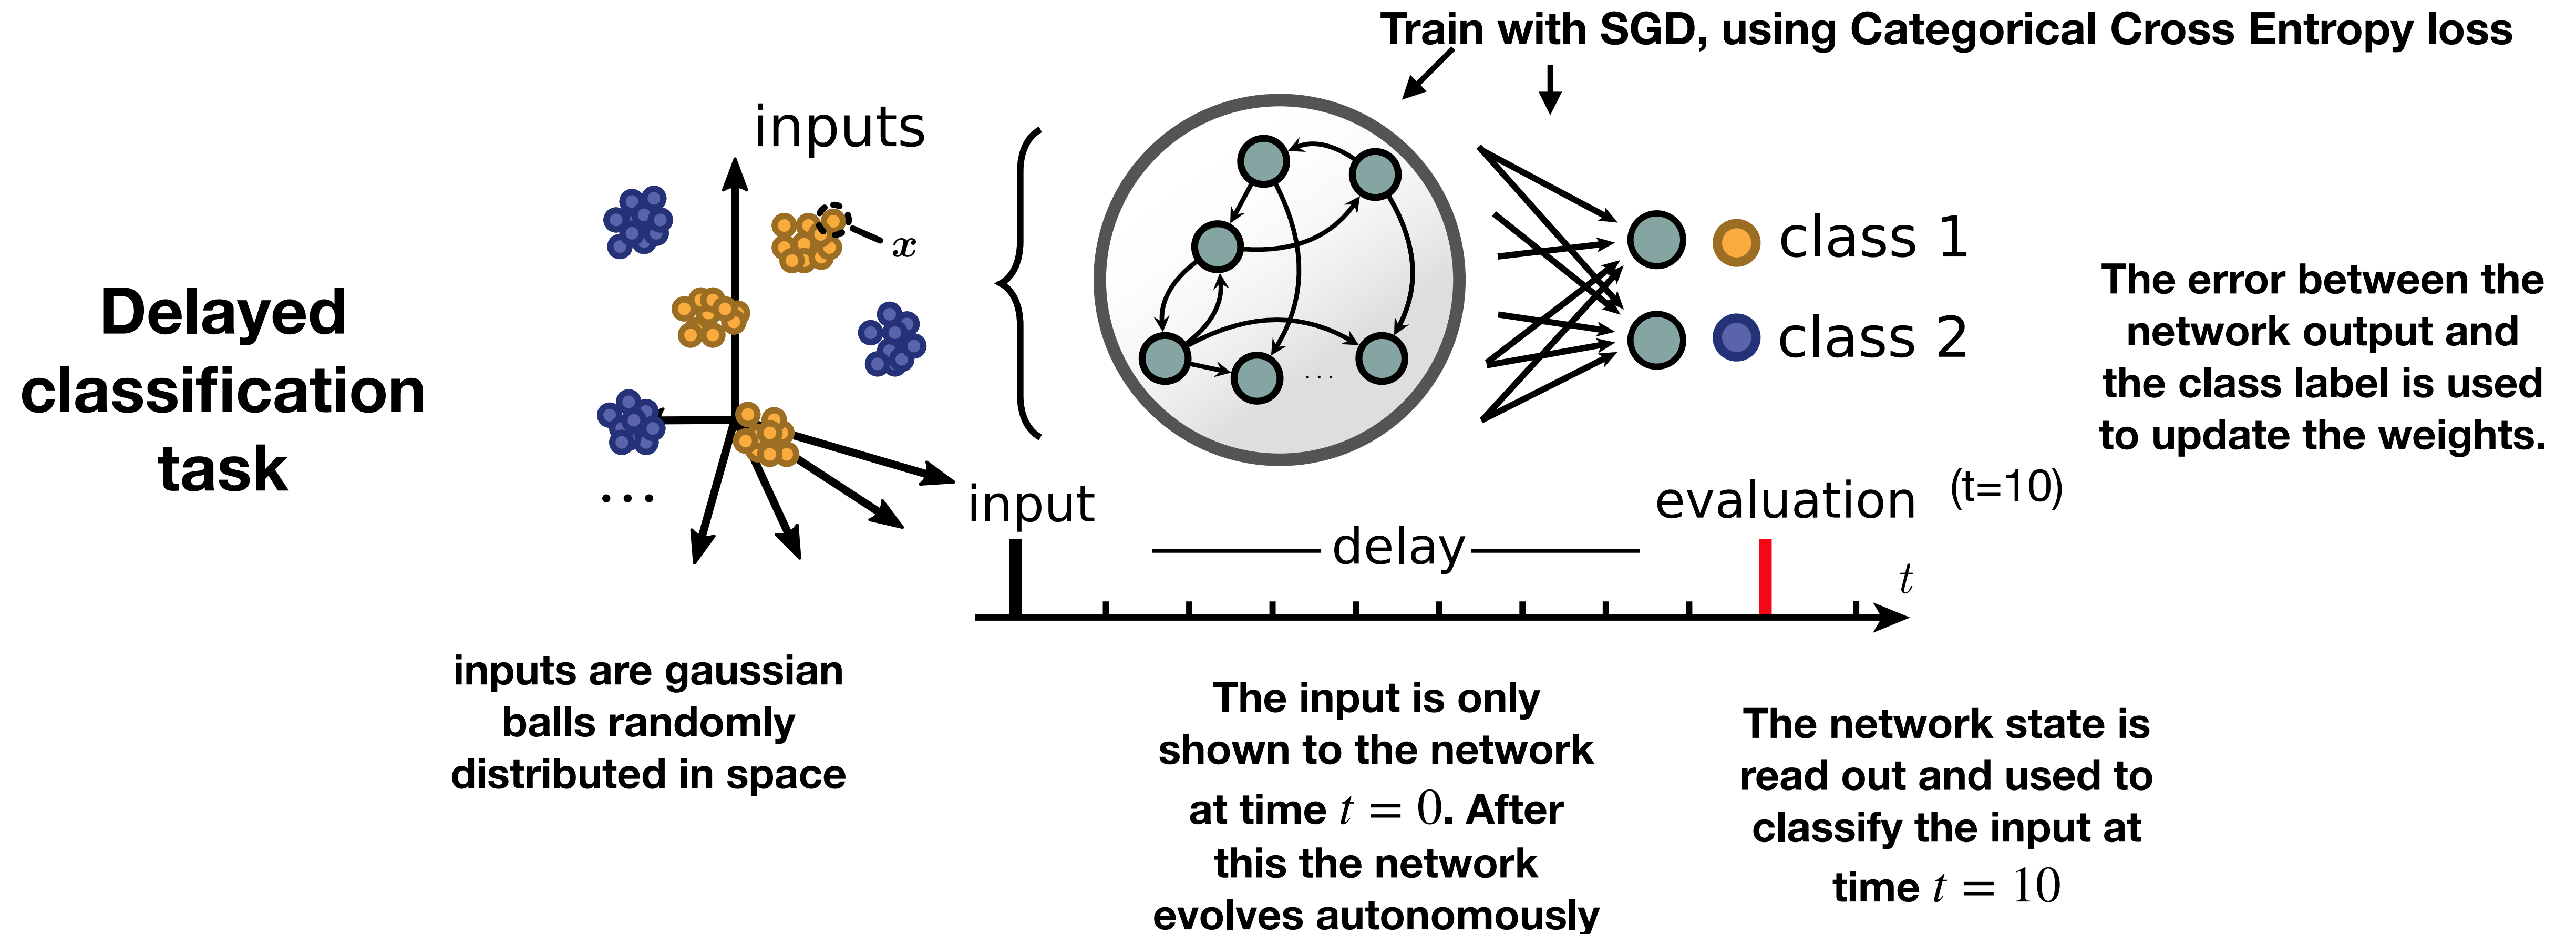
**Train recurrent and output weights with stochastic gradient descent (SGD)**

Goal: investigate how the initial dynamics of a recurrent neural network (RNN) influence the solutions learned by the network when trained to do a task.

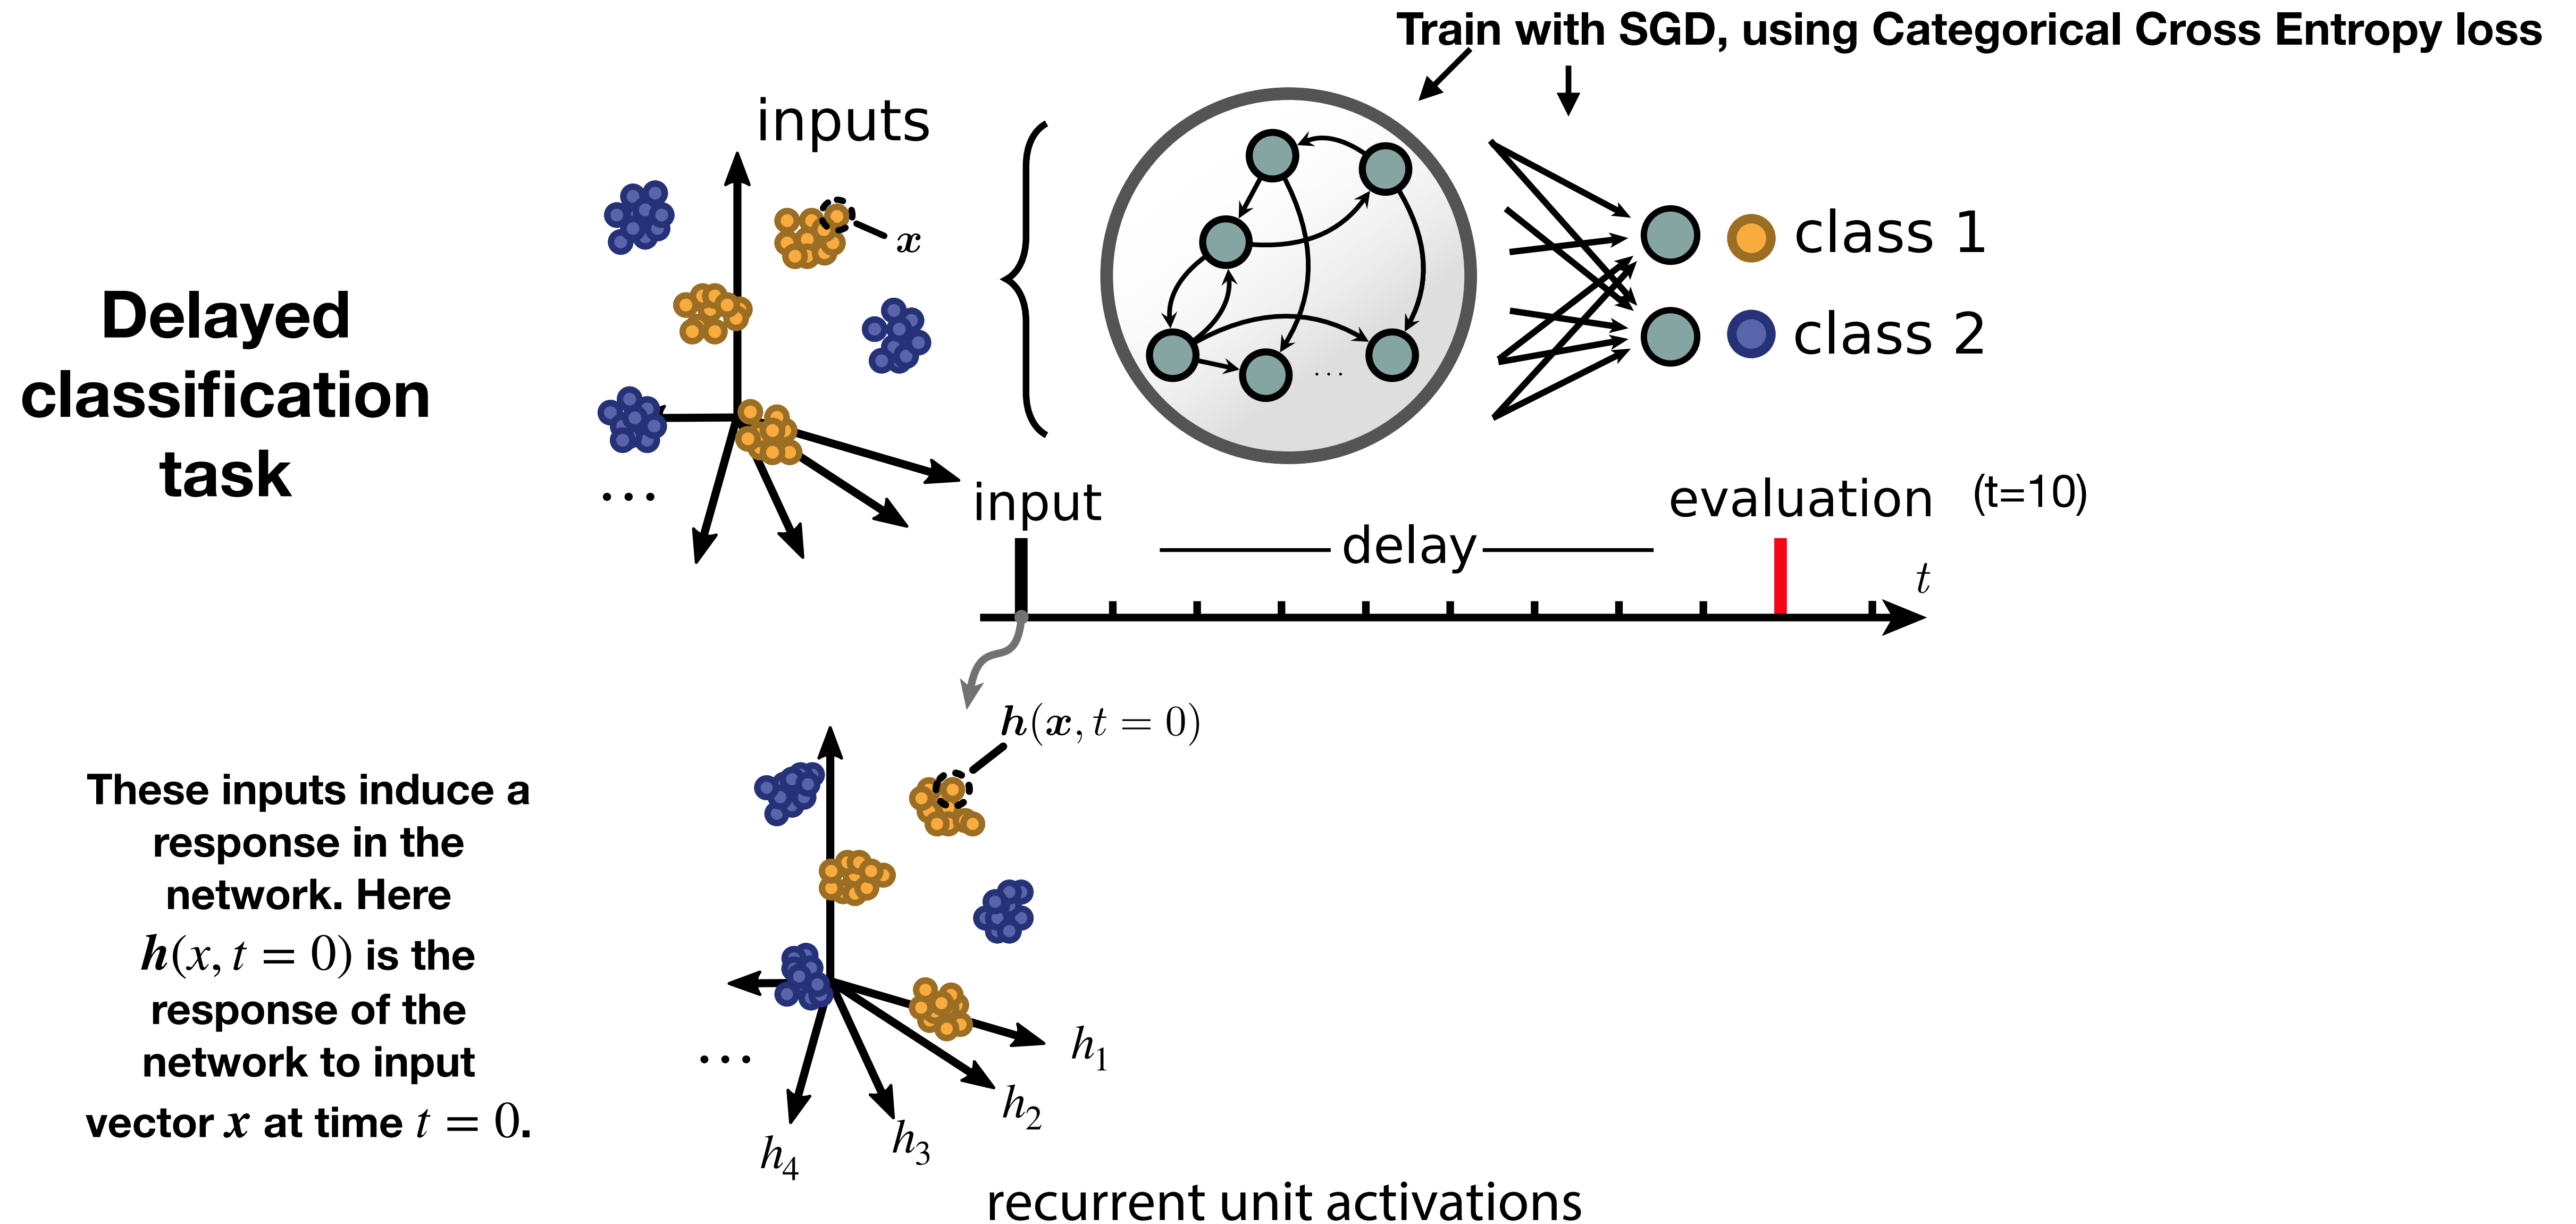
We characterize these dynamics by how chaotic the trajectories are.



Let's choose a task to train the network on




Let's choose a task to train the network on



Let's choose a task to train the network on

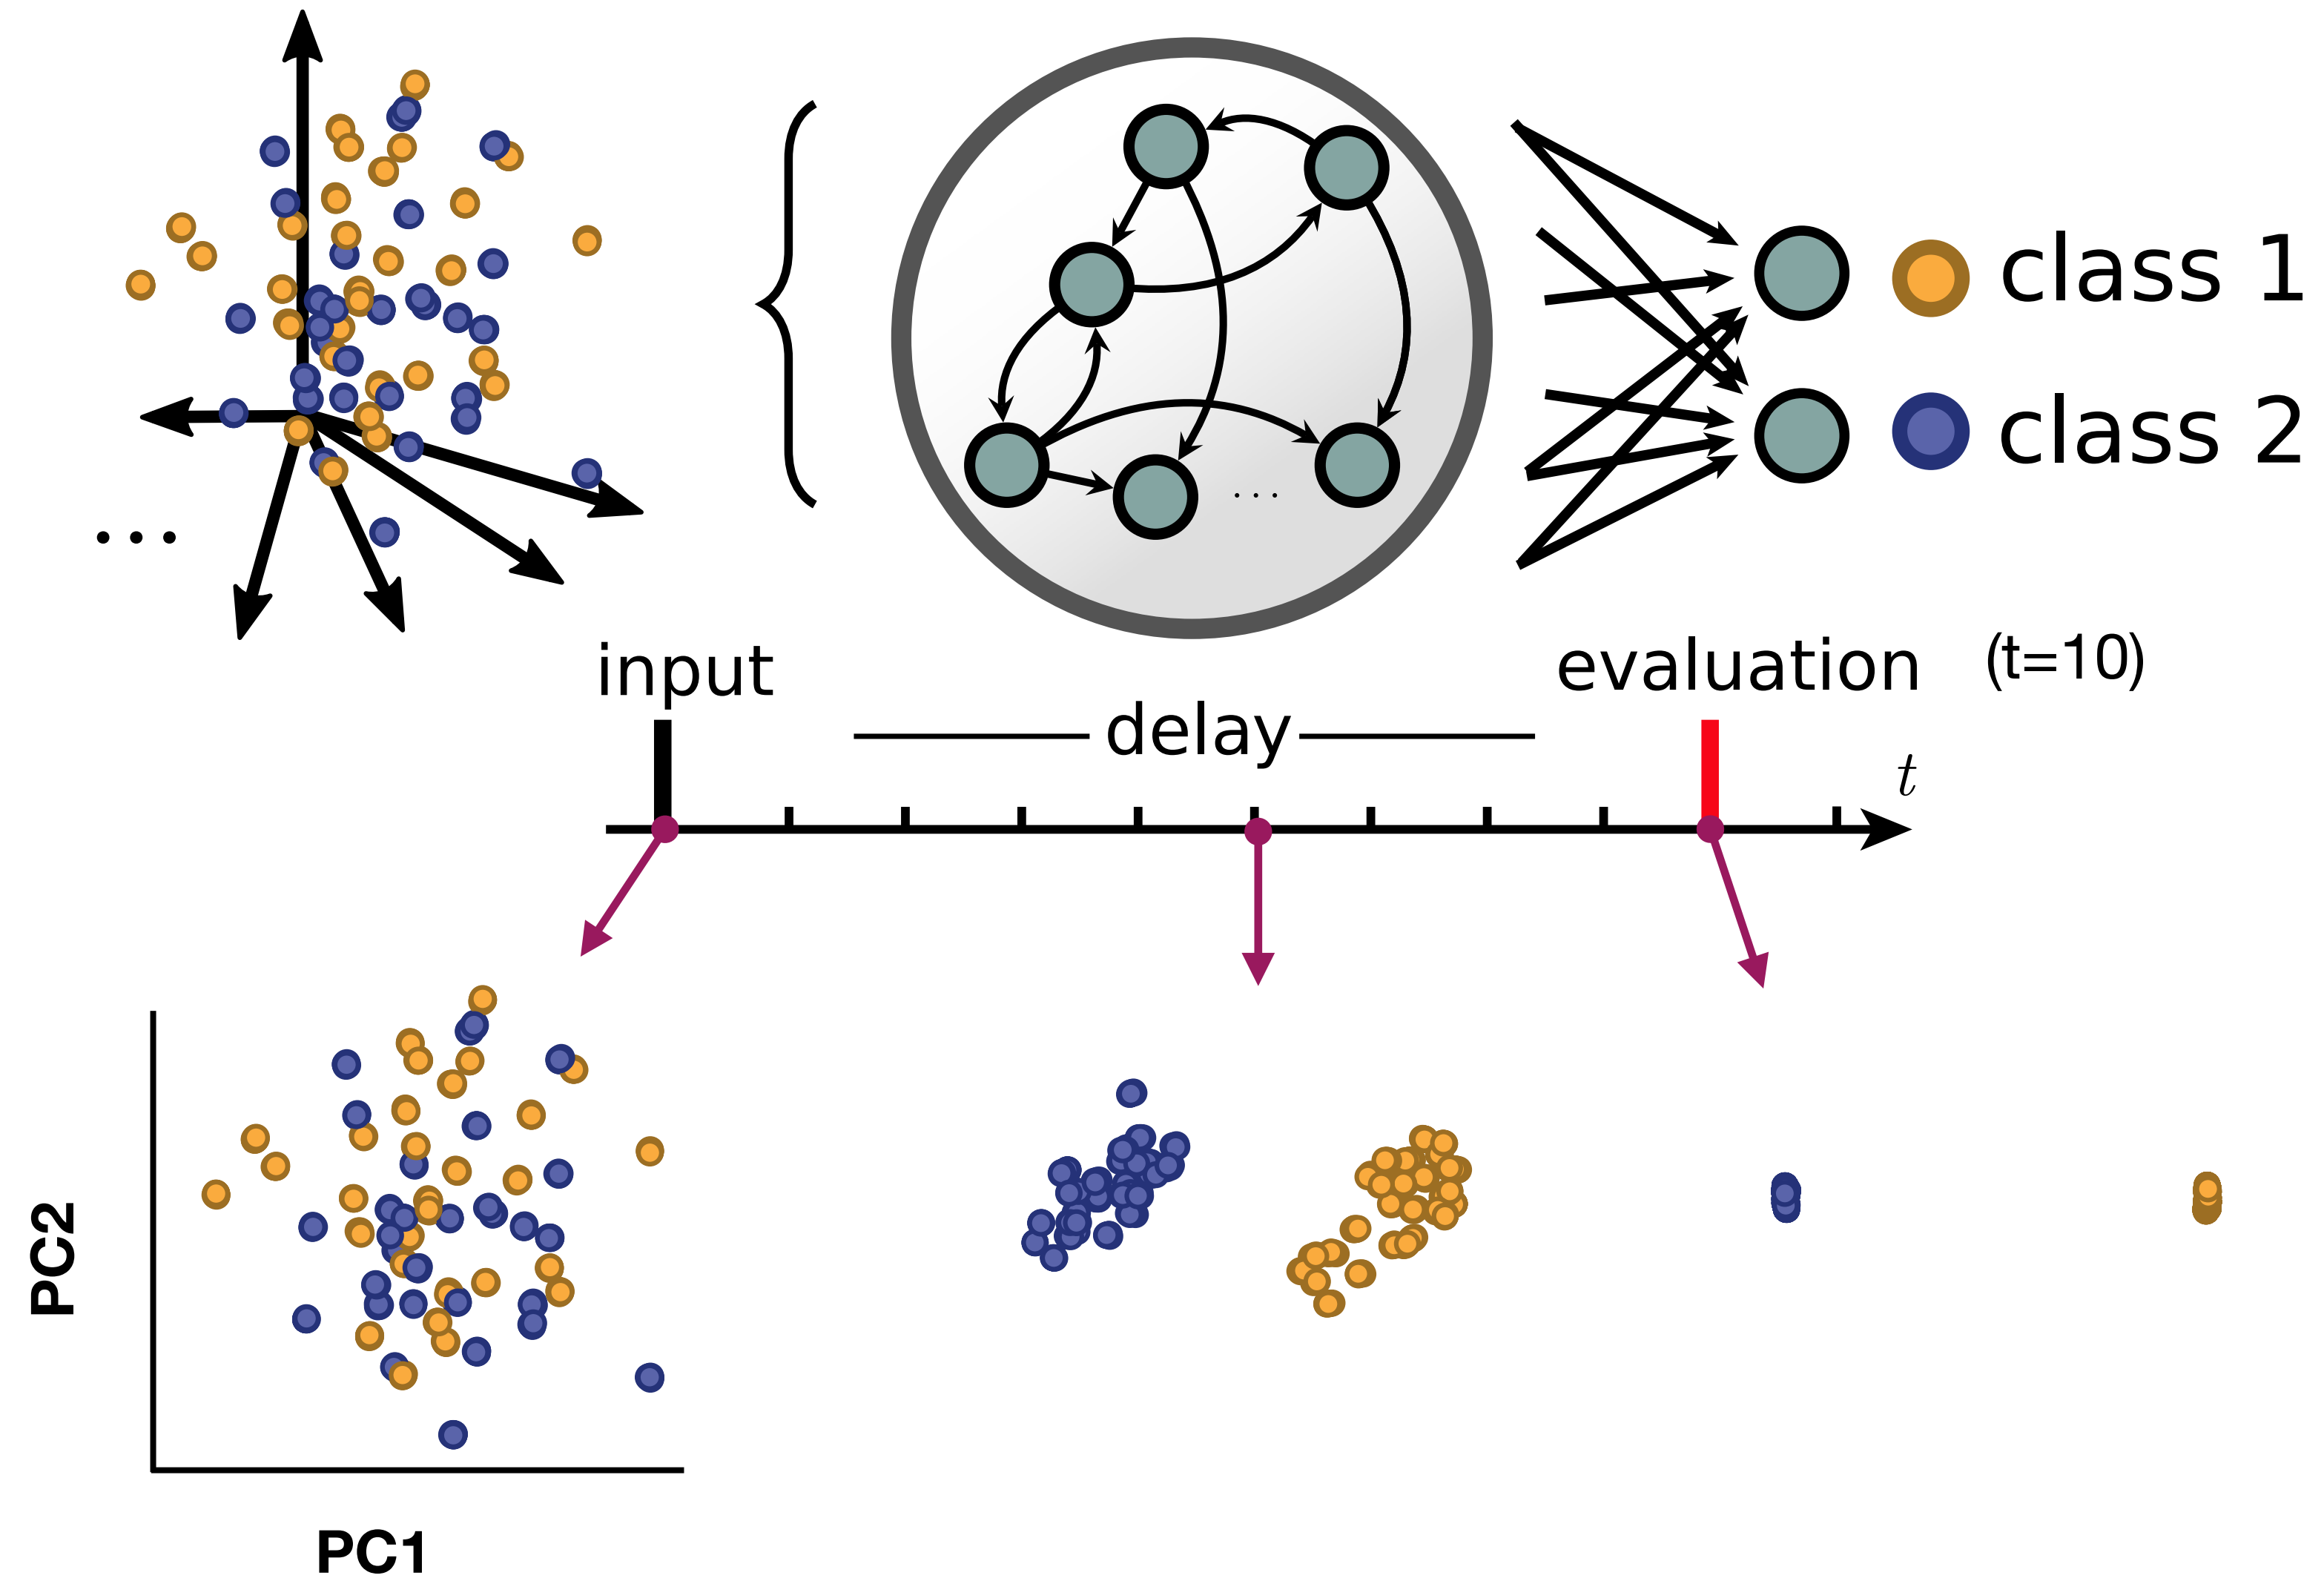
Train with SGD, using Categorical Cross Entropy loss



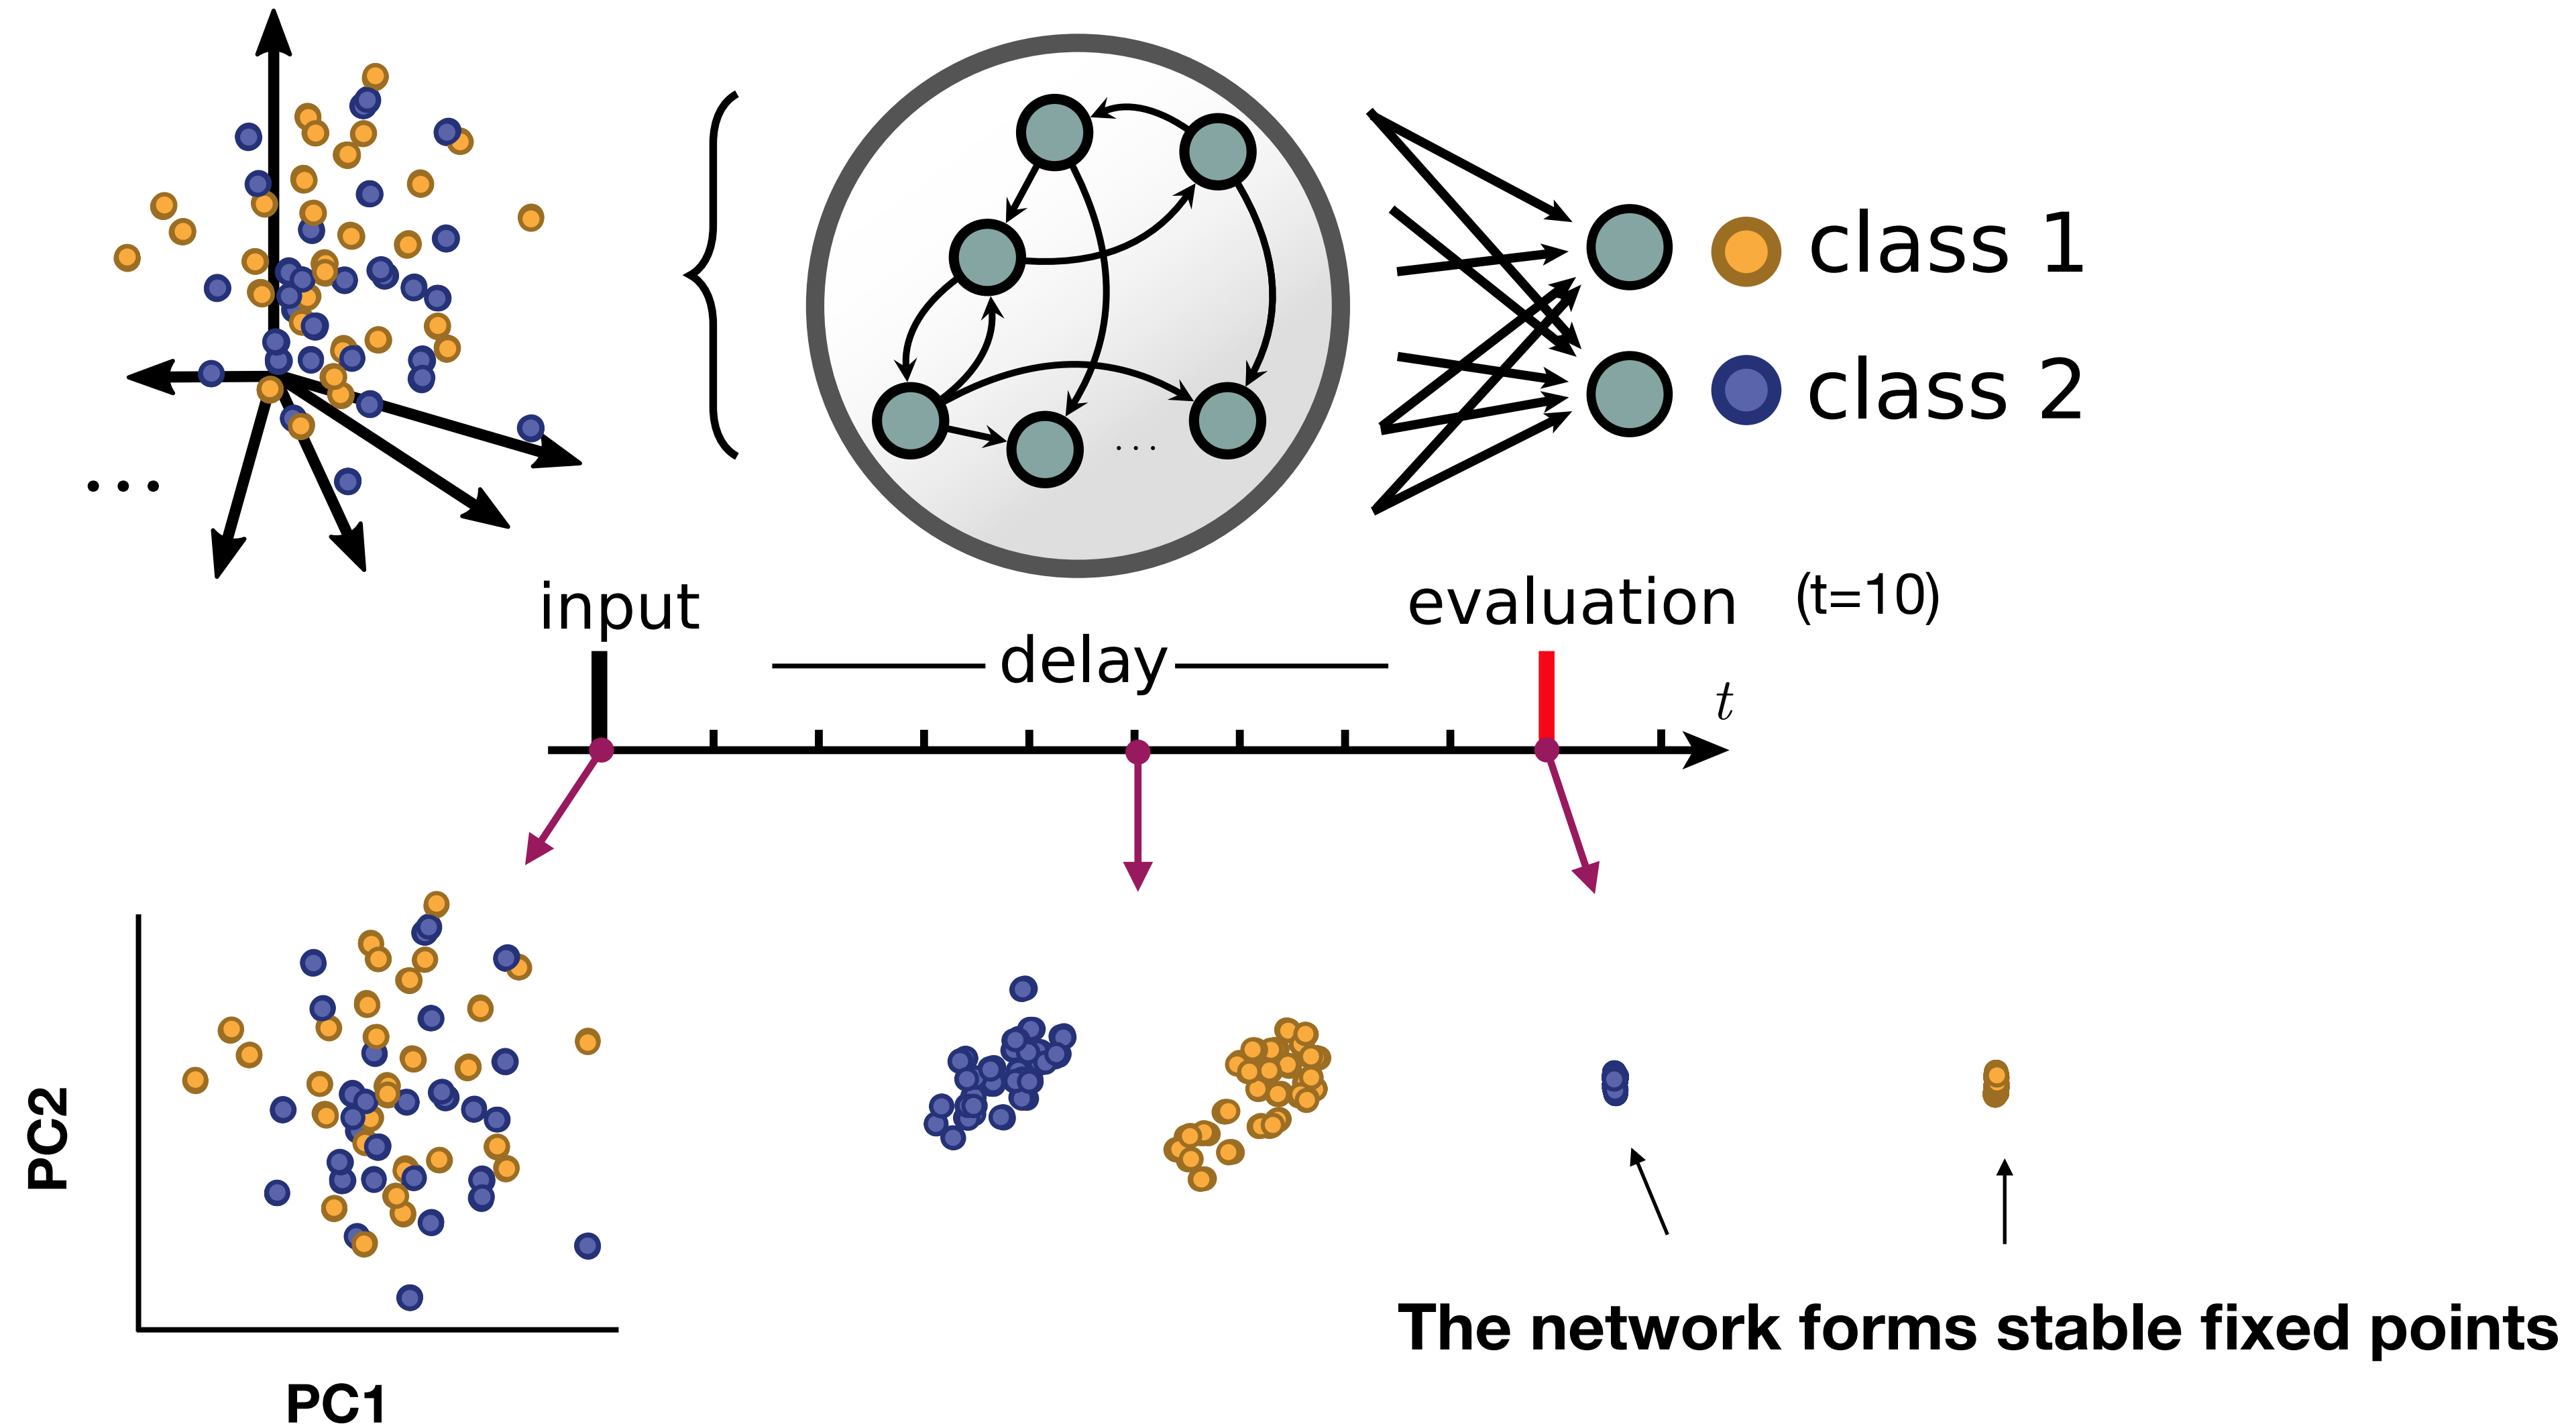
This response evolves through time. This is only a schematic though. Let's see what happens when we actually train the network.



High-dimensional,  
linearly separable input data



High-dimensional,  
linearly separable input data









How do we characterize the geometry of these?

Effective dimensionality (ED) is one approach

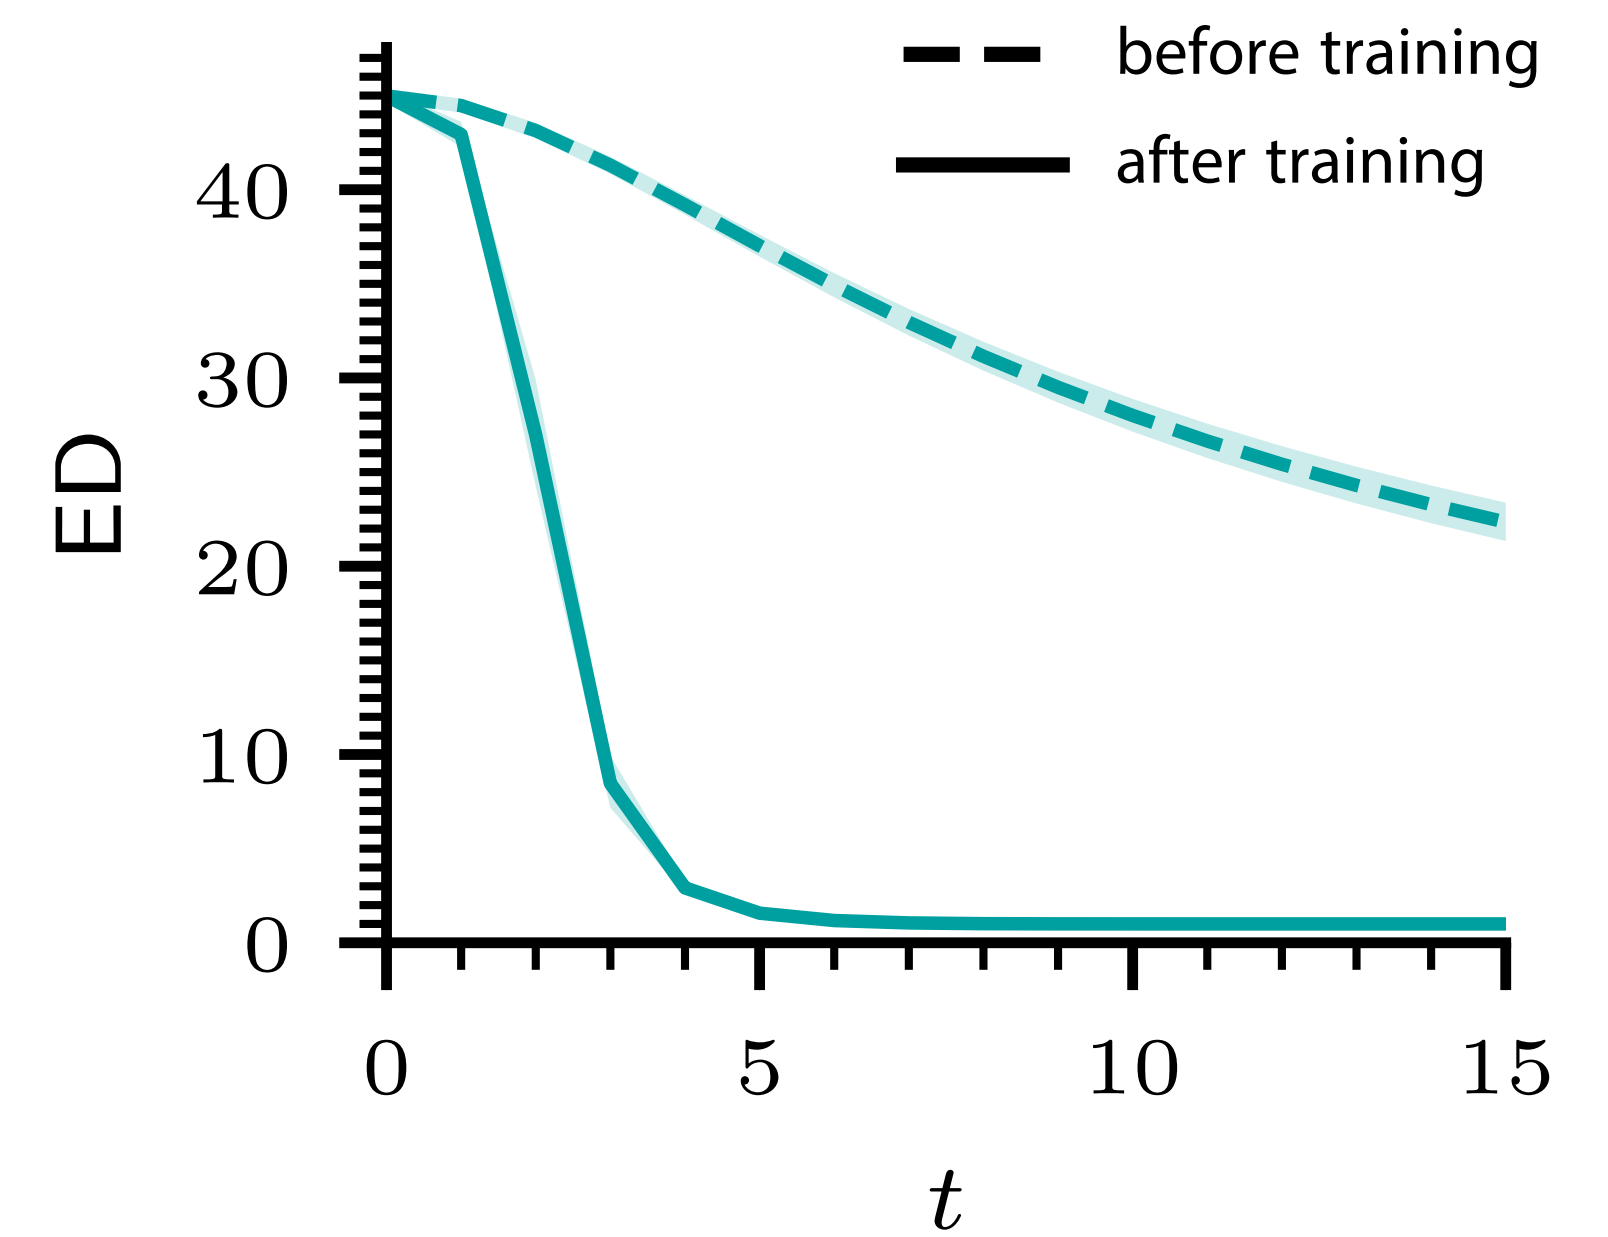


*Rajan et al., NeurIPS (2010)*



How do we characterize the geometry of these?

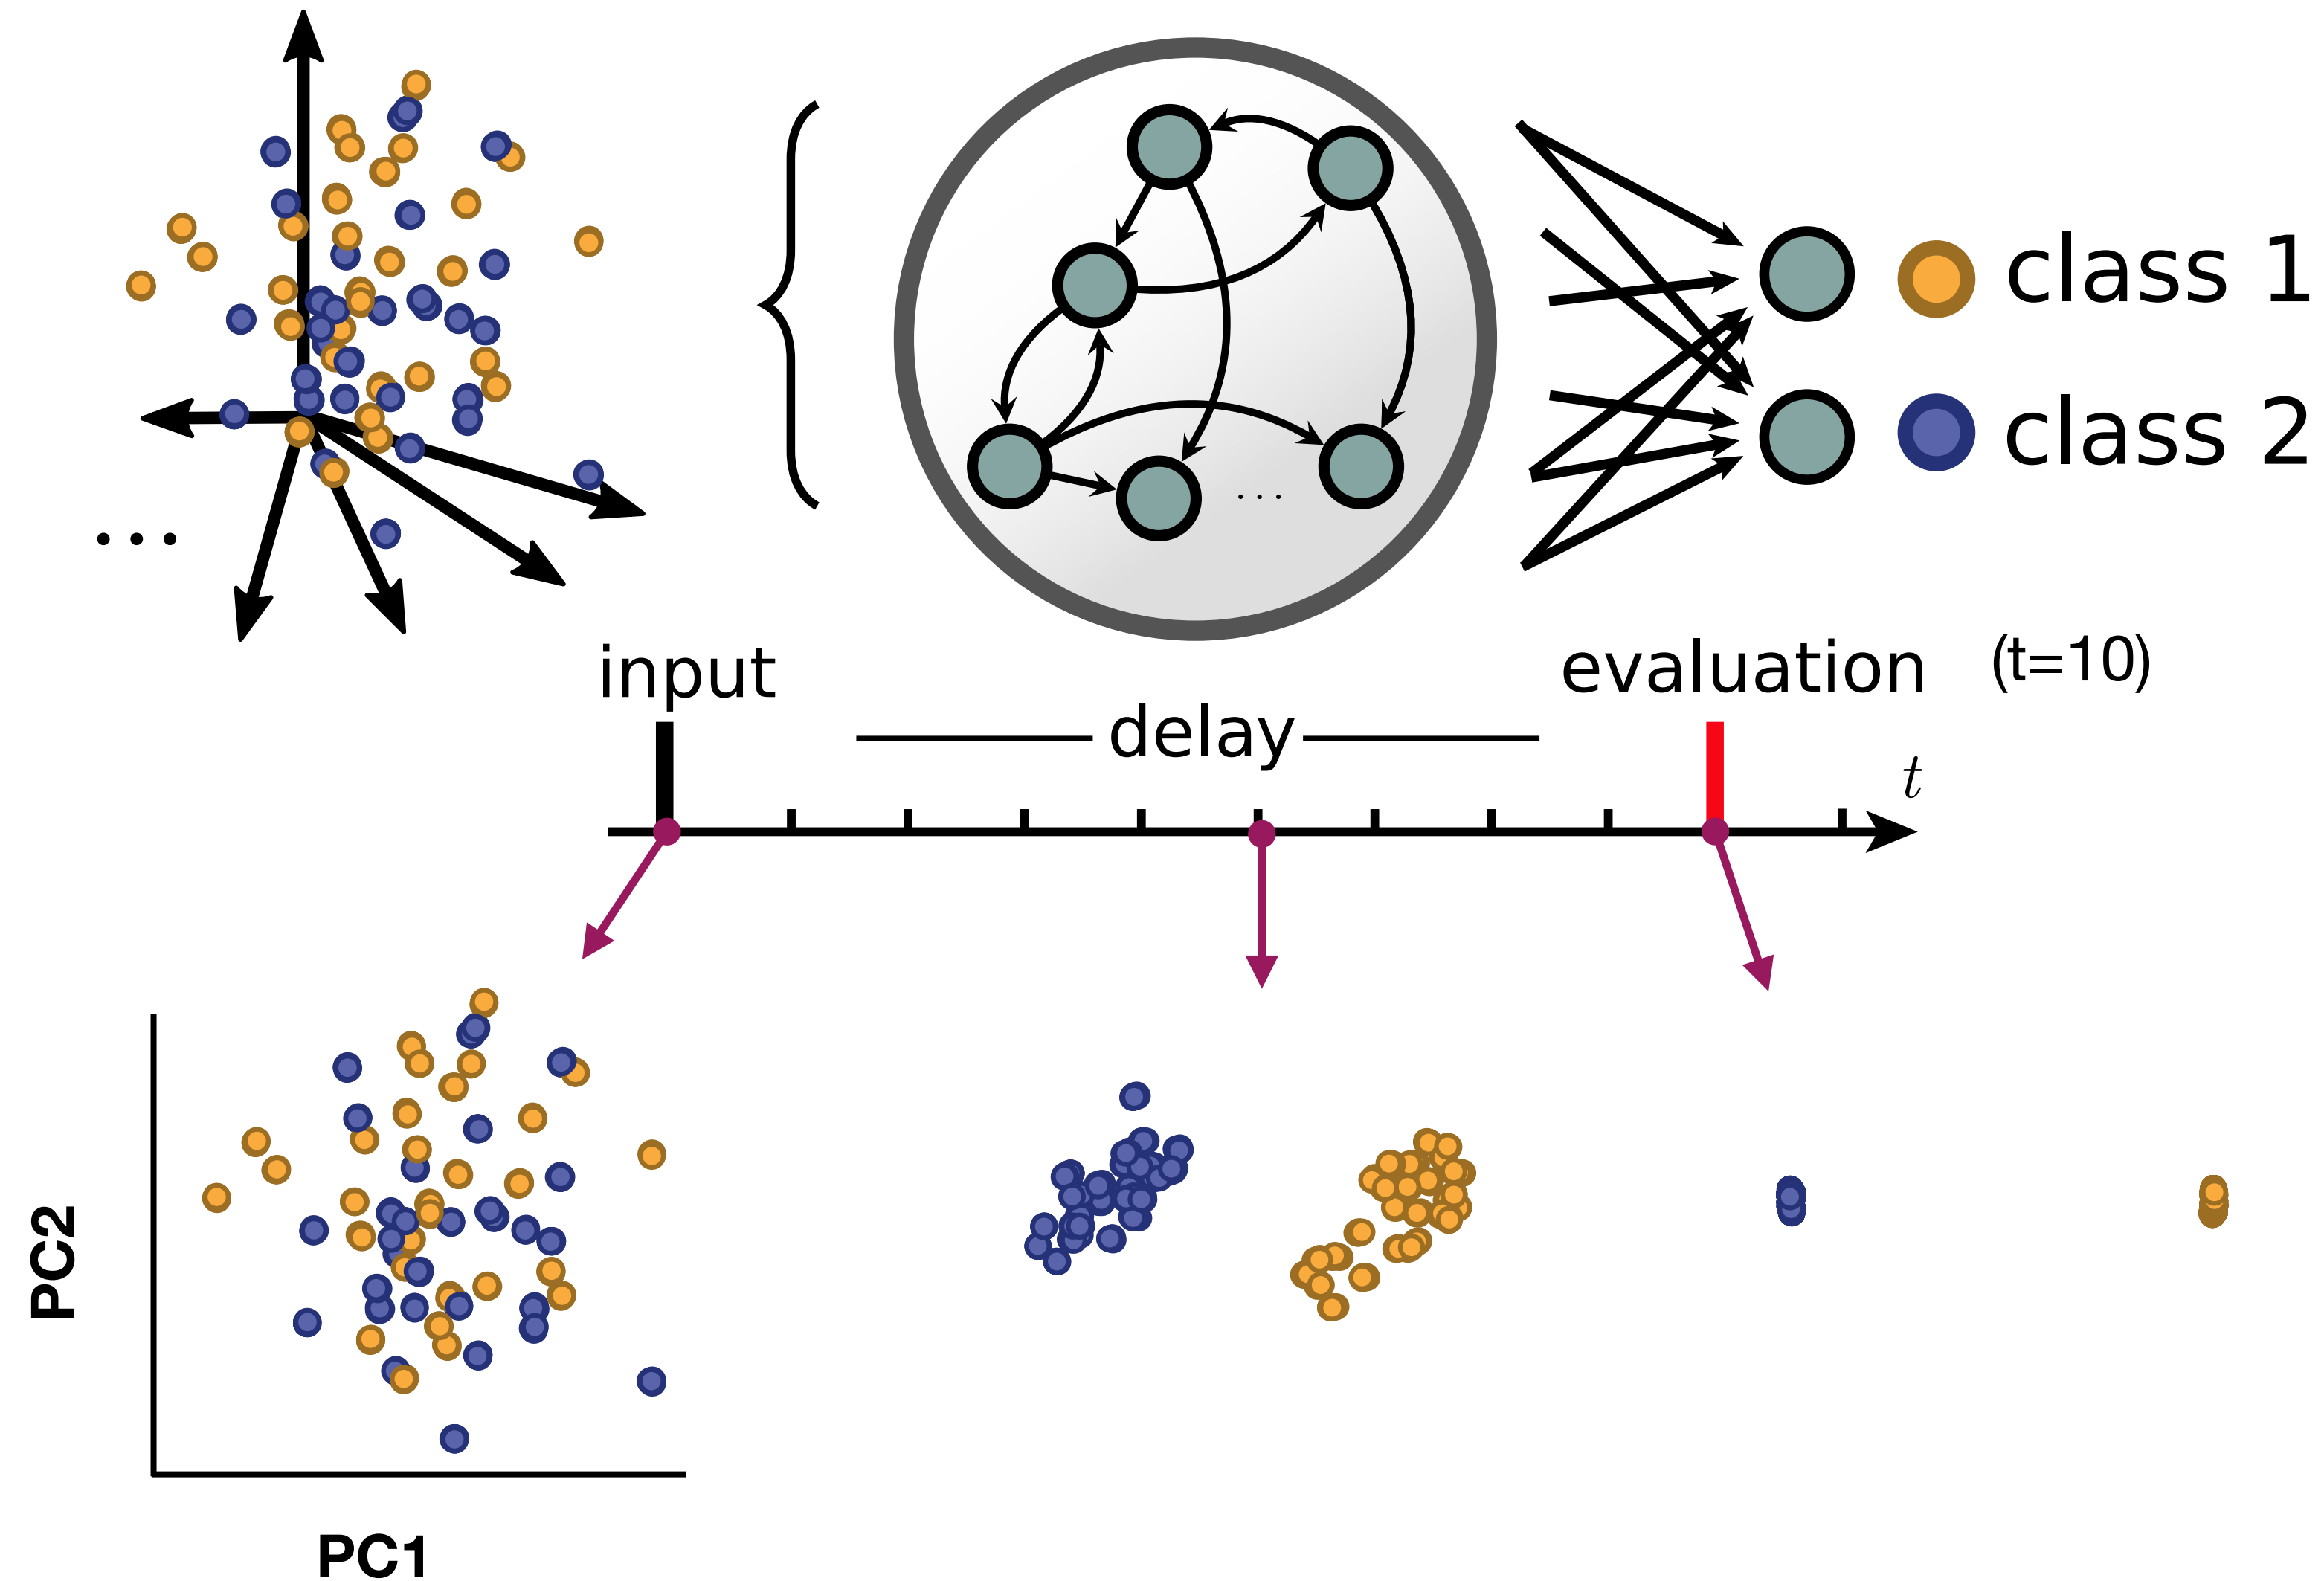
Effective dimensionality (ED) is one approach

*Rajan et al., NeurIPS (2010)*




High-dimensional,  
linearly separable input data



What happens if we initialize the network to  
be more chaotic?







High-dimensional,  
linearly separable input data



The network forms low-dimensional chaotic attractors

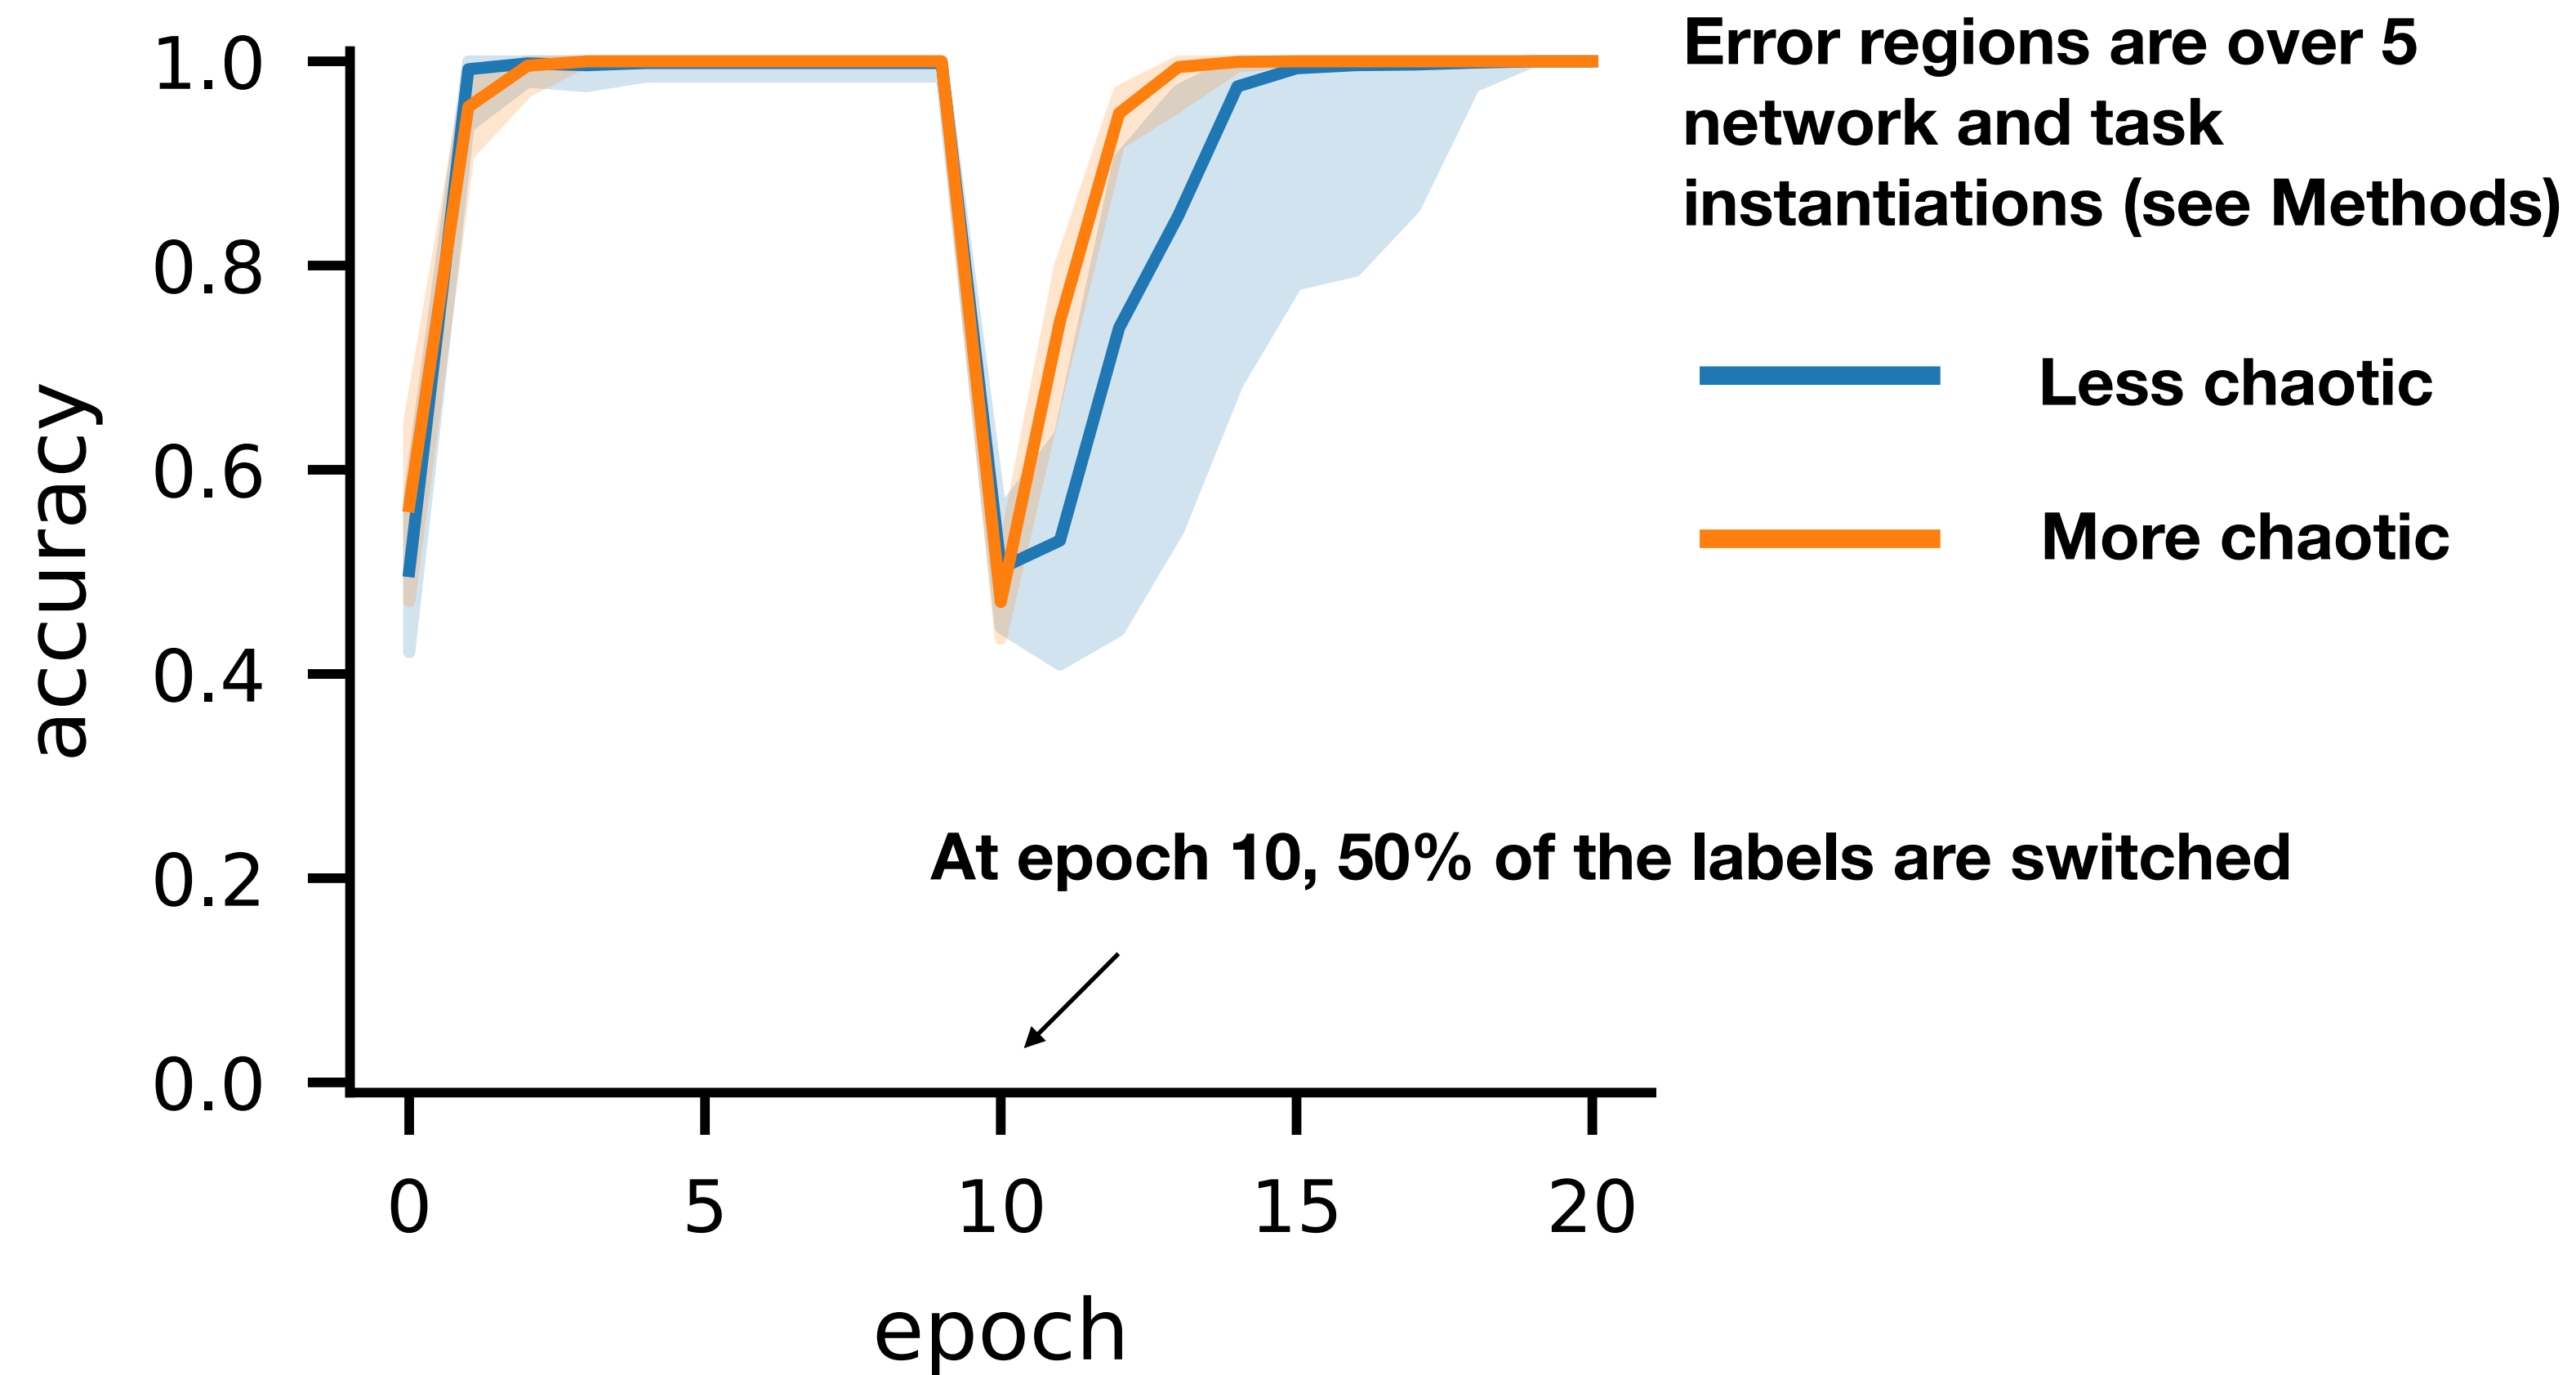


**If the labels are changed,  
this representation may be  
problematic due to tight  
overlap between points  
belonging to different  
classes**



**Both networks have perfect classification accuracy at the end of training.  
What happens if the task is changed in the middle of training?**


# Switching task labels



The more chaotic network is better at recovering after the label switch.

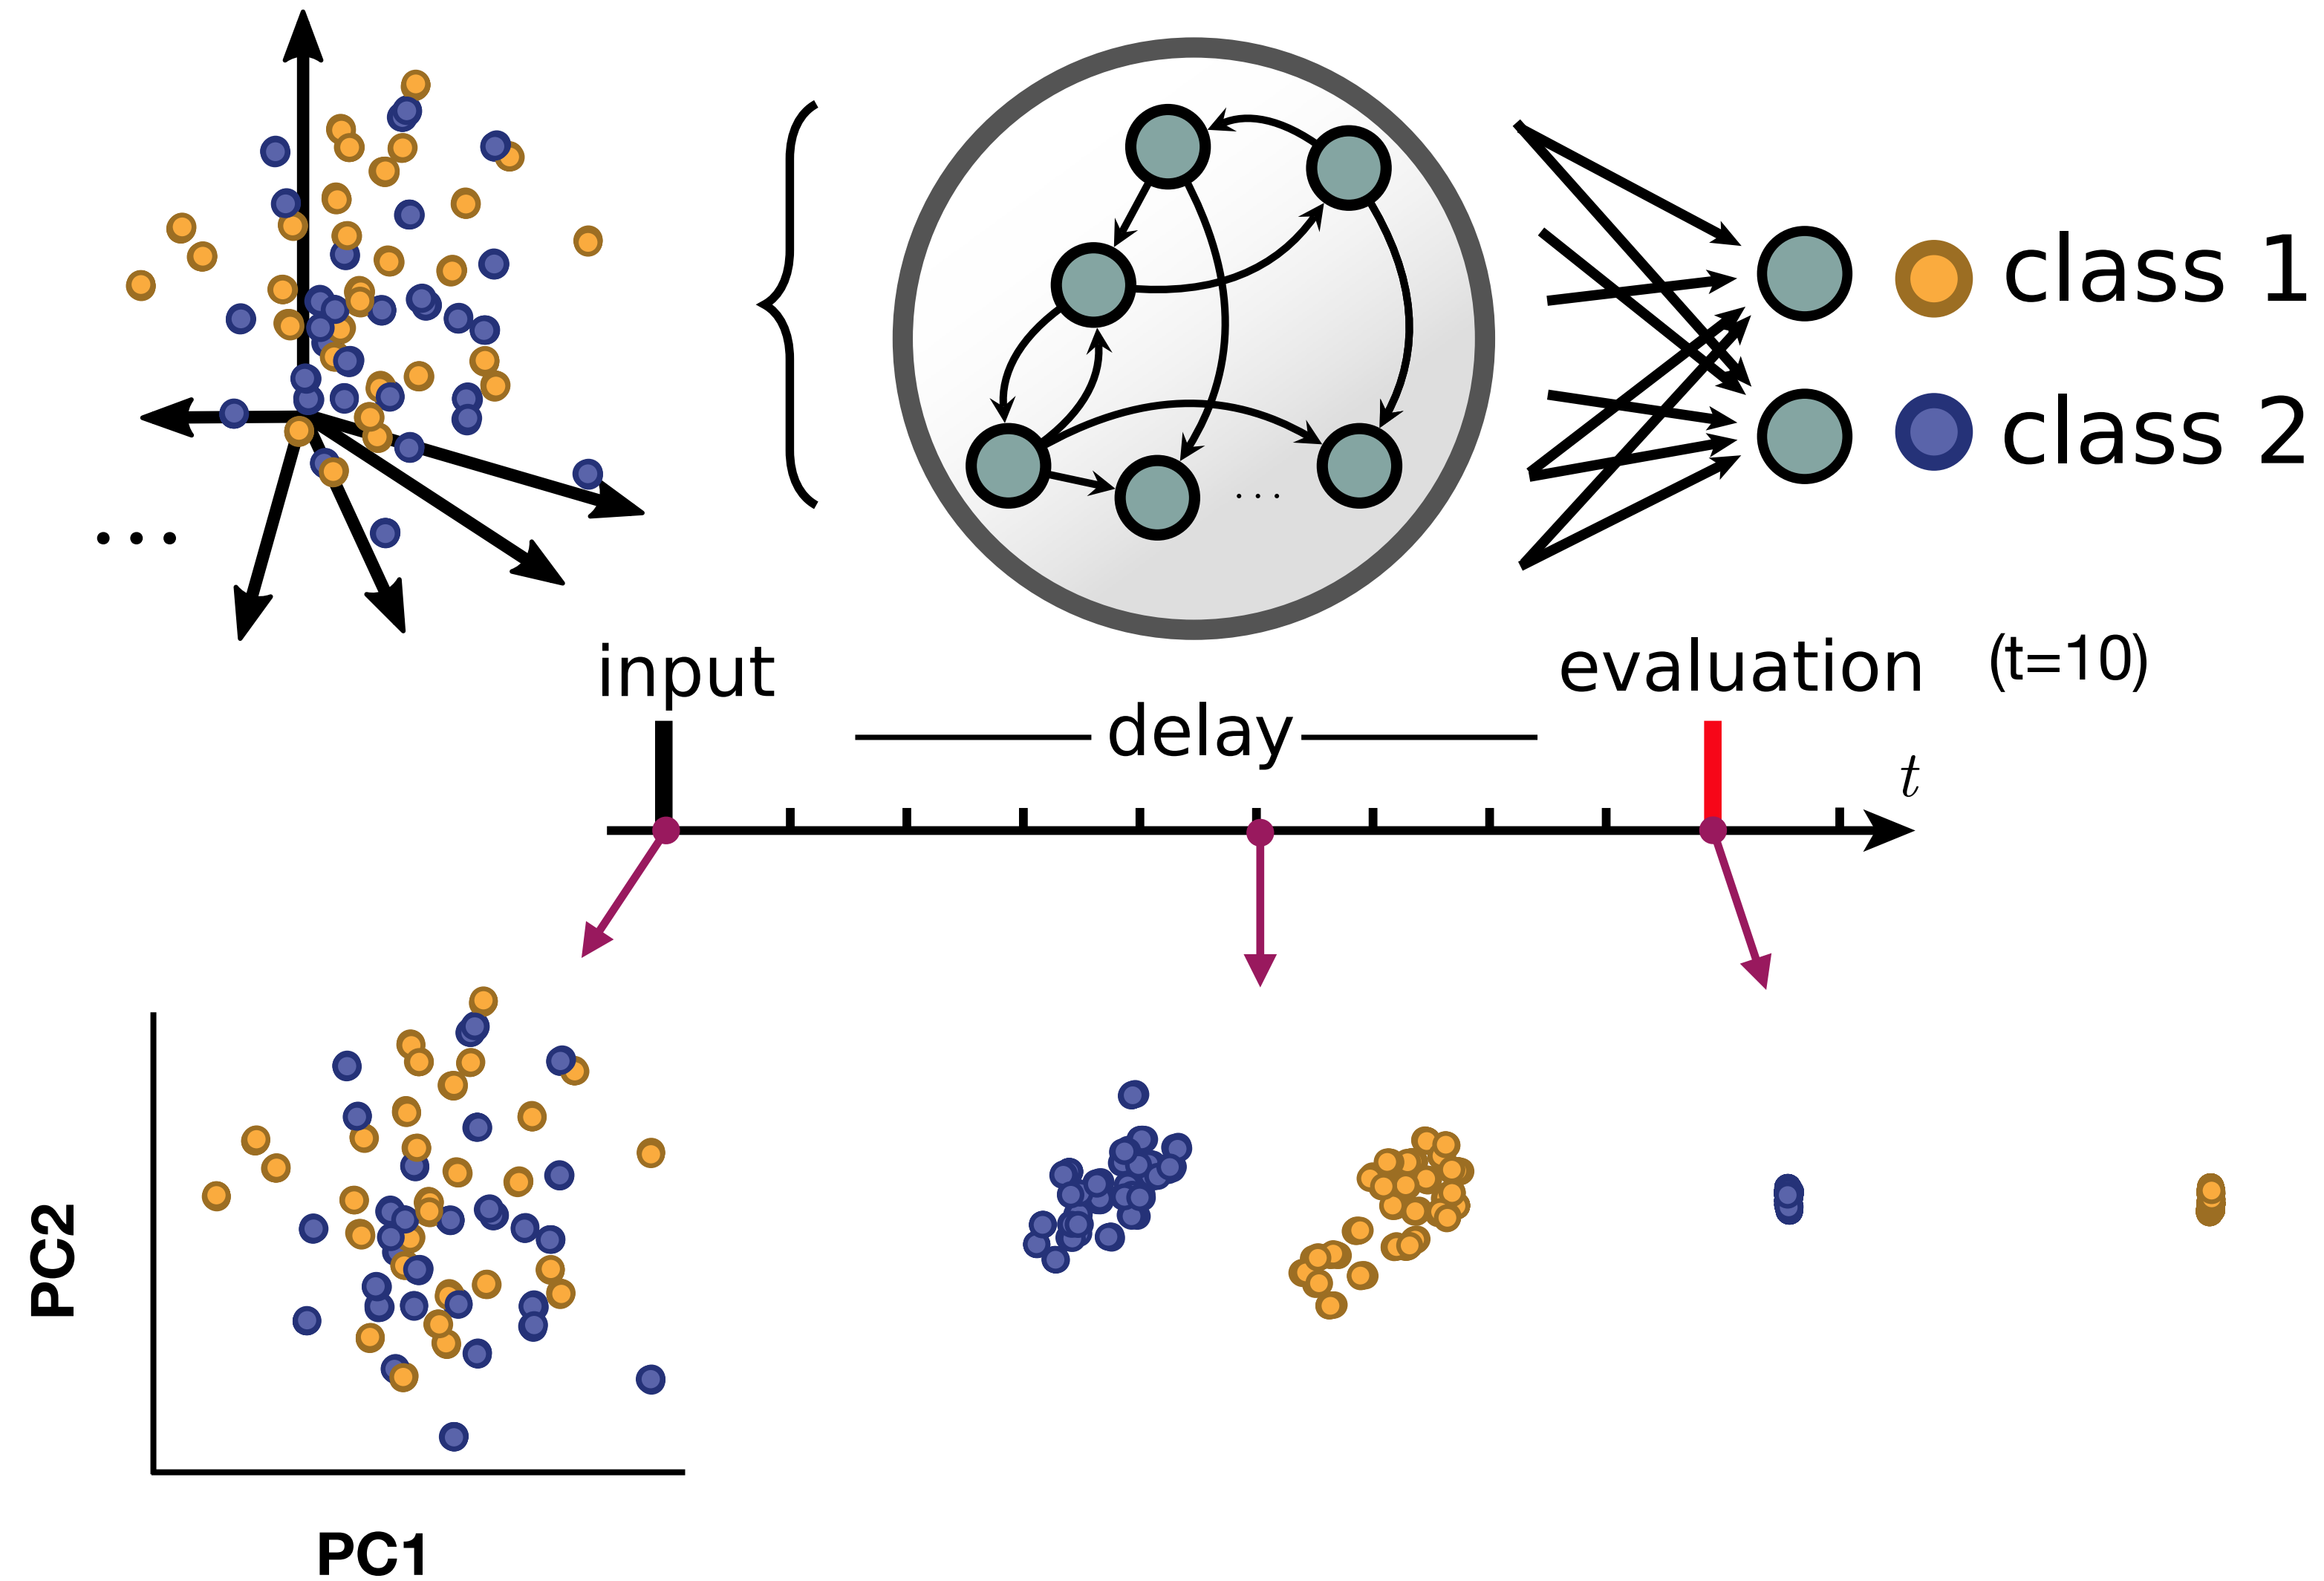


# Switching cluster locations




The more chaotic network is better at recovering after the label switch.

- Before we had high-dimensional, linearly separable input data.
- Now we'll look at the case of very low-dimensional input data



- Before we had high-dimensional, linearly separable input data.
- Now we'll look at the case of very low-dimensional input data


2-dimension inputs



- In this case, it is useful to expand dimensionality to help separate out the classes.
- Chaotic dynamics is one mechanism for this expansion (Legenstein and Maass, *Neural Networks* (2007)).


- Before we had high-dimensional, linearly separable input data.
- Now we'll look at the case of very low-dimensional input data

2-dimension inputs




- In this case, it is useful to expand dimensionality to help separate out the classes.
- Chaotic dynamics is one mechanism for this expansion (Legenstein and Maass, *Neural Networks* (2007)).

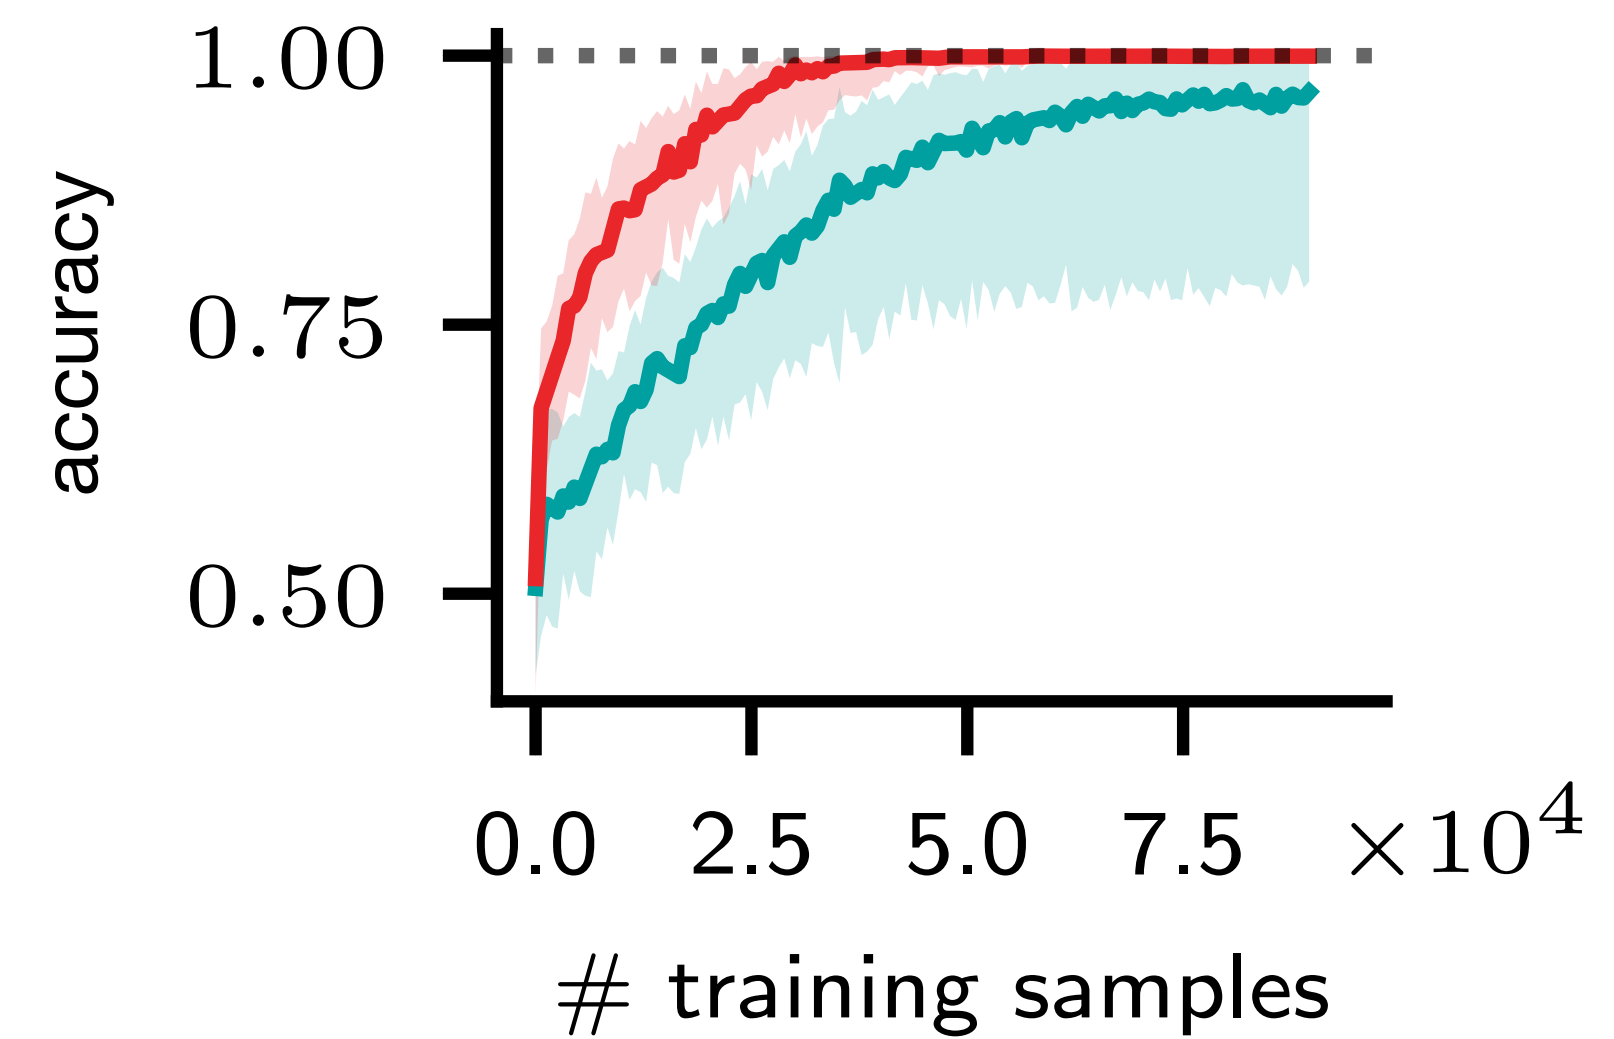
Chaotic divergence



- The more chaotic network expands dimensionality before training as well as after training.
- The less chaotic network is not able to learn this expansion.



- This allows the more chaotic network to do better at the task.



# Conclusions

- **Chaotic dynamics can have utility in recurrent neural network models:**
  - **Leads to solutions that can more quickly adapt to changing data.**
  - **Expands dimensionality, which can help when classifying low-dimensional data.**
- **This suggests beneficial attributes of the variability seen in biology, which may in part be generated by chaotic dynamics.**



# Methods

- Network equations

$$h_t = \tanh(\textcolor{red}{W}h_{t-1} + x_t + b) \quad \text{recurrent unit activations}$$

$$\hat{o}_t = \textcolor{blue}{R}h_t + b' \quad \text{readout}$$

- Chaos was modulated by increasing the variance of the connections between neurons at initialization:  $W = (1 - \epsilon)I + (g\epsilon/\sqrt{N})J$  where  $\epsilon = .01$  and  $J$  is a matrix with standard normal entries. For the less chaotic network  $g = 5$ , and for the more chaotic network  $g = 250$ .
- Error regions on plots are found by fitting a gamma distribution to the data, and shading the region containing 75% probability mass of the distribution. For the accuracy plots this distribution was reflected around the y-axis and shifted by +1 (since accuracy is bounded above by 1), and for the dimensionality plots the distribution was shifted by +1 (since the effective dimensionality is bounded from below by 1).