

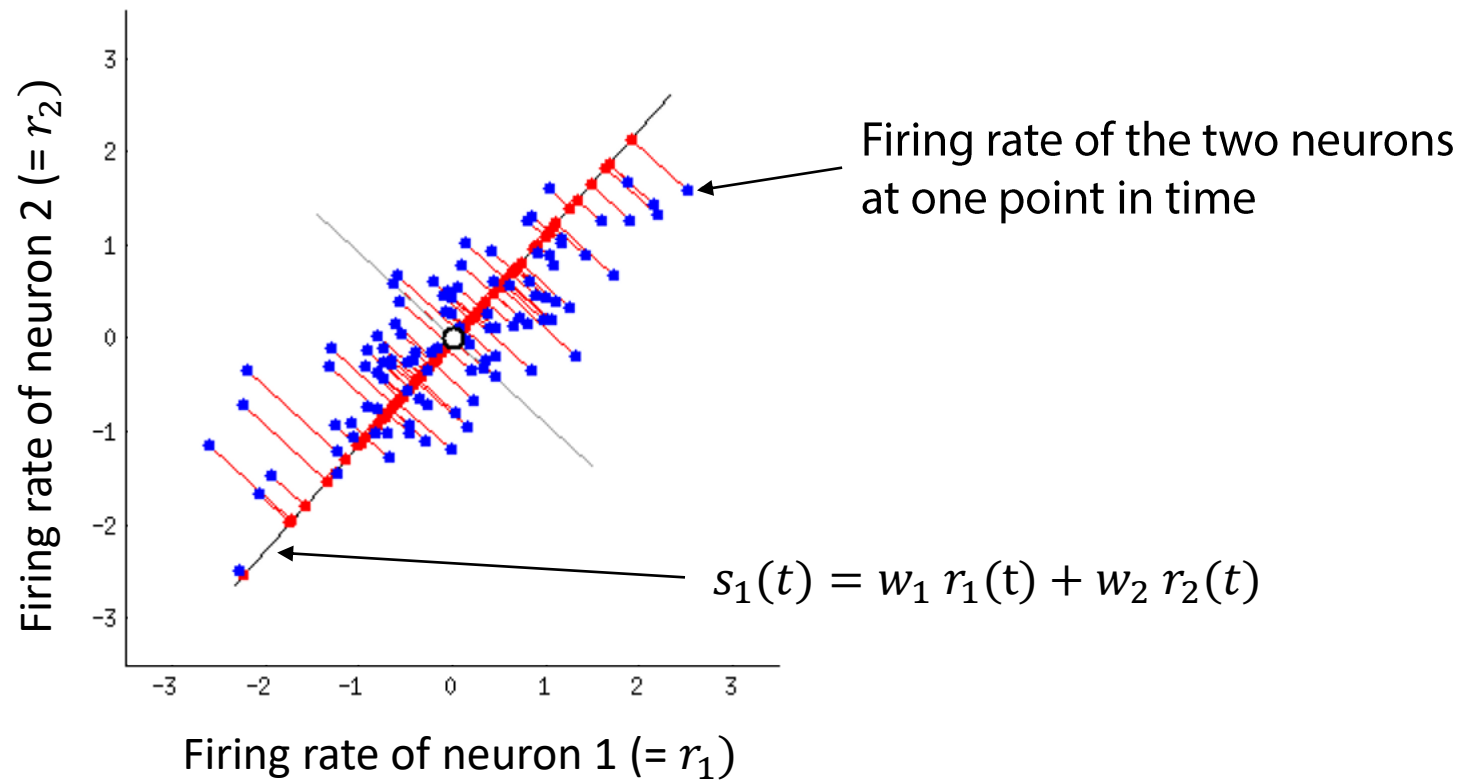
Dimensionality Reduction Part 2: Generative Models

Bi23: Methods in Neural Data Analysis

3/8/2019

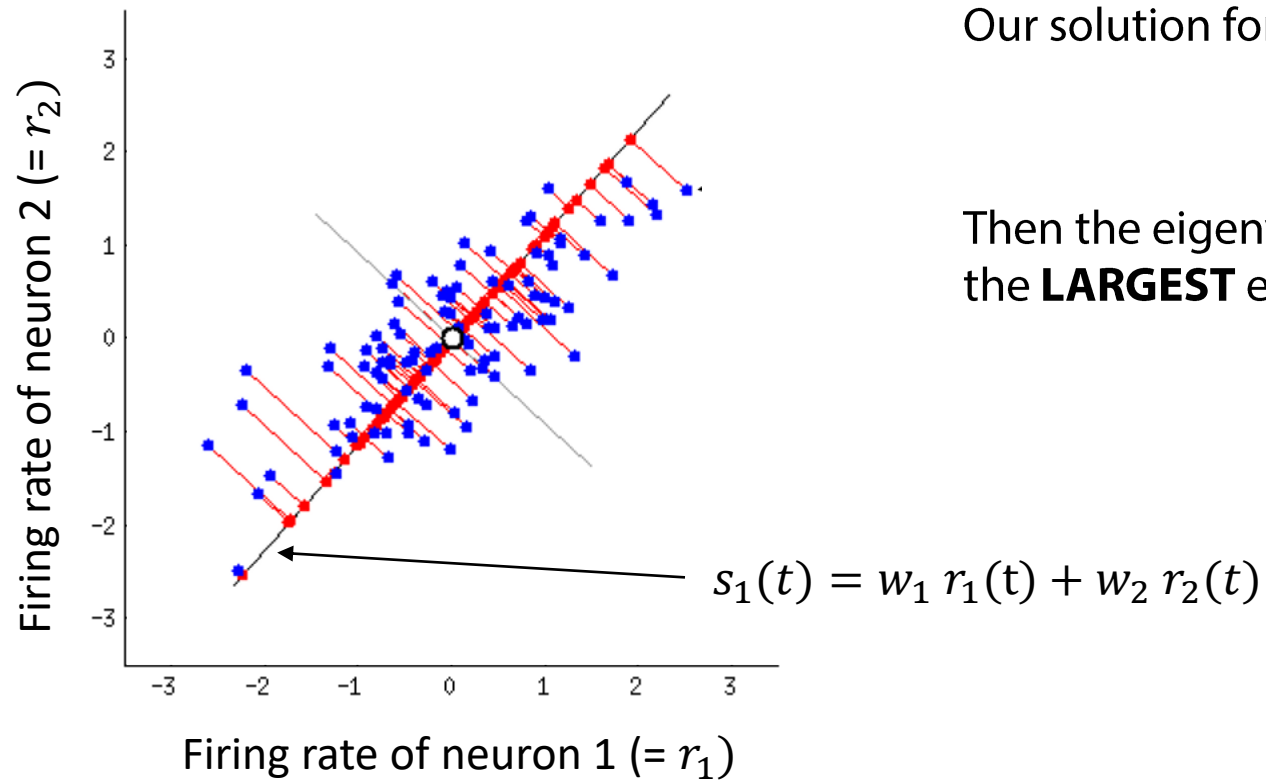
Last time: if we pick the right basis, we can reduce data dimensionality without losing much information

- Original data
- 1-dimensional "summary"



Last time: PCA finds the dimensions that capture the most variance in neural activity (ie the dimensions that lose the least information)



- Original data
- 1-dimensional "summary"




Our solution for w is an eigenvector of the **covariance matrix** Σ ,
$$\Sigma w^T = \lambda w^T$$

Then the eigenvector w_k that maximizes $w_k \Sigma w_k^T$ is the one with the **LARGEST** eigenvalue λ_k


So: PCA is computing the *eigenvectors* of the *covariance matrix* (Σ)



$$\Sigma = (\mathbf{r} - \bar{\mathbf{r}})(\mathbf{r} - \bar{\mathbf{r}})^T$$



Projection onto first
Principal Axis



A different perspective

Rather than projecting observed data into a low-dimensional space (eg with PCA), we want to construct a *hypothesis* about how the observed data was *generated*.

- Eg: the observed data (spikes) are noisy observations of a small number of underlying *latent variables*.

Generative models

A generative model describes the statistical process from which data arises.

- For example (GLM): **Underlying Spike rate** $\mu = \theta x$ **unknown parameter**
Predicted spike distribution $y \sim N(\mu, \sigma^2)$ **stimulus**

Generative modeling: drawing a picture of a car

Discriminative modeling: distinguishing a car from a motorcycle

Some other types of generative models:

- Hidden Markov Model
- Gaussian Mixture Model
- Variational Autoencoder
- Generative Adversarial Network (GAN)



AI-generated faces created with GAN 2.0

A Generative Modeling Take on Dimensionality Reduction

We observe the spiking of a population of N neurons, $\mathbf{y}(t) = \{y_1(t), y_2(t), \dots, y_N(t)\}$.

A simple model might be to assume that the firing rate of each neuron y_n is a function of a small number D ($D \ll N$) of latent variables, \mathbf{z} , as follows:

$$z_d(t) \sim N(0,1)$$

$$y_n(t) \sim N\left(\sum_{d=1}^D W_{n,d} z_d(t), \sigma^2\right)$$

Or, in matrix notation,

$$\mathbf{z}(t) \sim N(0, \mathbf{I})$$

Identity
matrix

$$\mathbf{y}(t) \sim N(\mathbf{W} \cdot \mathbf{z}(t), \sigma^2 \mathbf{I})$$

"factor loading" matrix ($D \times N$)

Or written another way:



$$\mathbf{y}(t) = \mathbf{W} \cdot \mathbf{z}(t) + \epsilon, \quad \epsilon \sim N(0, \sigma^2 \mathbf{I})$$

This is "Probabilistic PCA"!

Reconciling Probabilistic with ordinary PCA

$$\mathbf{z}(t) \sim N(0, \mathbf{I})$$

$$\mathbf{y}(t) = \mathbf{W} \cdot \mathbf{z}(t) + \epsilon, \quad \epsilon \sim N(0, \sigma^2 \mathbf{I})$$



We can make lots of models from this same general framework:

Probabilistic PCA:

$$\mathbf{z}(t) \sim N(0, \mathbf{I})$$

$$\mathbf{y}(t) = \mathbf{W} \cdot \mathbf{z}(t) + \epsilon, \quad \epsilon \sim N(0, \sigma^2 \mathbf{I})$$

Factor Analysis:

$$\mathbf{z}(t) \sim N(0, \mathbf{I})$$

$$\mathbf{y}(t) = \mathbf{W} \cdot \mathbf{z}(t) + \epsilon, \quad \epsilon \sim N(0, \Psi)$$

$$\Psi = \begin{bmatrix} \sigma_1^2 & 0 & \dots & 0 \\ 0 & \sigma_2^2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & \sigma_N^2 \end{bmatrix}$$

Linear Dynamical System:

$$\mathbf{z}(t) = \mathbf{A} \cdot \mathbf{z}(t-1) + \mathbf{v}, \quad \mathbf{v} \sim N(\mu, \sigma^2 \mathbf{I}_z)$$

$$\mathbf{y}(t) = \mathbf{W} \cdot \mathbf{z}(t) + \mathbf{b} + \epsilon, \quad \epsilon \sim N(0, \sigma^2 \mathbf{I}_x)$$

Gaussian Process Factor Analysis:

$$\mathbf{z}_i(t) \sim N(0, \mathbf{K}_i)$$

$$\mathbf{y}(t) = \mathbf{W} \cdot \mathbf{z}(t) + \mathbf{b} + \epsilon, \quad \epsilon \sim N(0, \Psi)$$

$$K_i(t_1, t_2) = \sigma_{f,i}^2 \exp\left(-\frac{(t_1 - t_2)^2}{2\tau_i^2}\right) + \sigma_{n,i}^2 \delta_{t_1, t_2}$$

...LOTS of models:

Hidden Markov Model:

$$\mathbf{z}(t) \in \{1, 2, \dots, D\} \quad (\text{discrete states!})$$

$$\mathbf{y}(t) \sim N(\mu_z, \Sigma_z)$$

Independent Component Analysis:

$$\mathbf{z}(t) \sim \text{maximally non - Gaussian}$$

$$\mathbf{y}(t) \sim N(\mathbf{z}, \sigma_z^2 \mathbf{I})$$

Switching Linear Dynamical System:

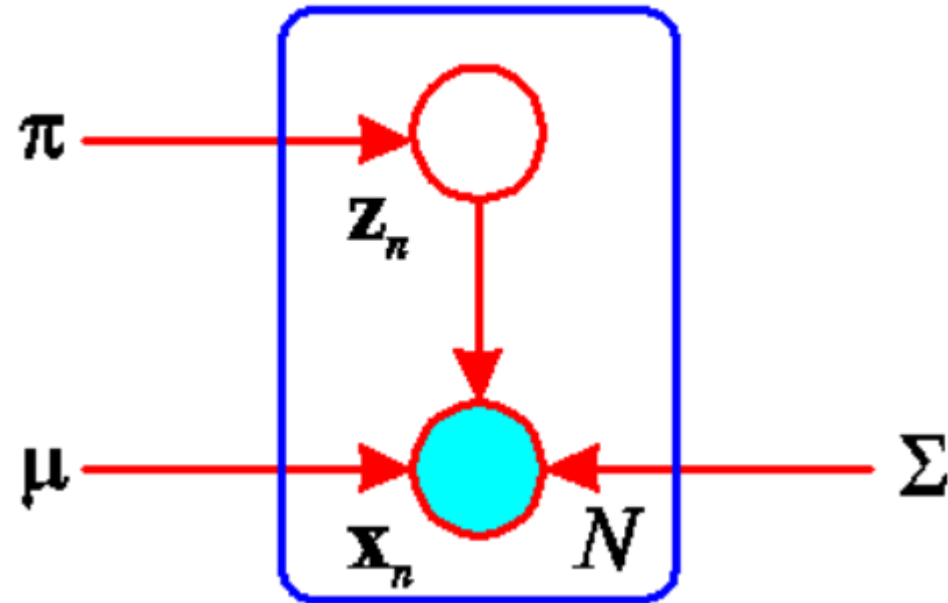
$$\pi(t) \in \{1, 2, \dots, D\}$$

$$\mathbf{z}(t) = \mathbf{A} \cdot \mathbf{z}(t-1) + v, \quad v \sim N(\mu_\pi, \sigma_\pi^2 \mathbf{I}_z)$$

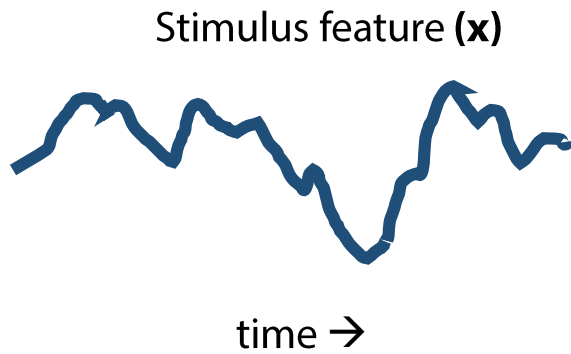

$$\mathbf{y}(t) = \mathbf{W} \cdot \mathbf{z}(t) + b + \epsilon, \quad \epsilon \sim N(0, \sigma^2 \mathbf{I}_x)$$

Sometimes you'll see these models represented graphically:

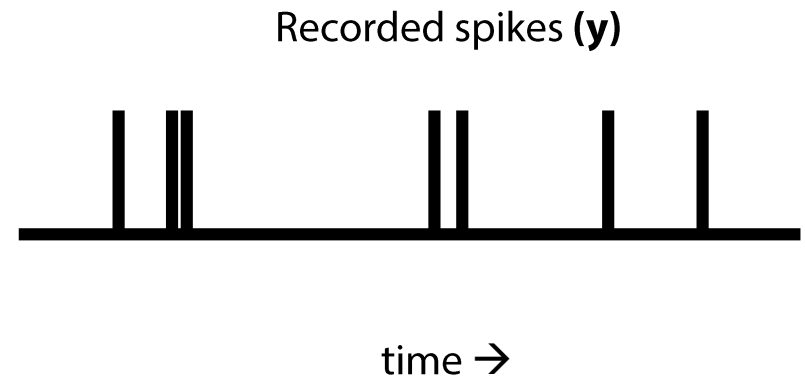
Graphical representation of a Gaussian Mixture Model



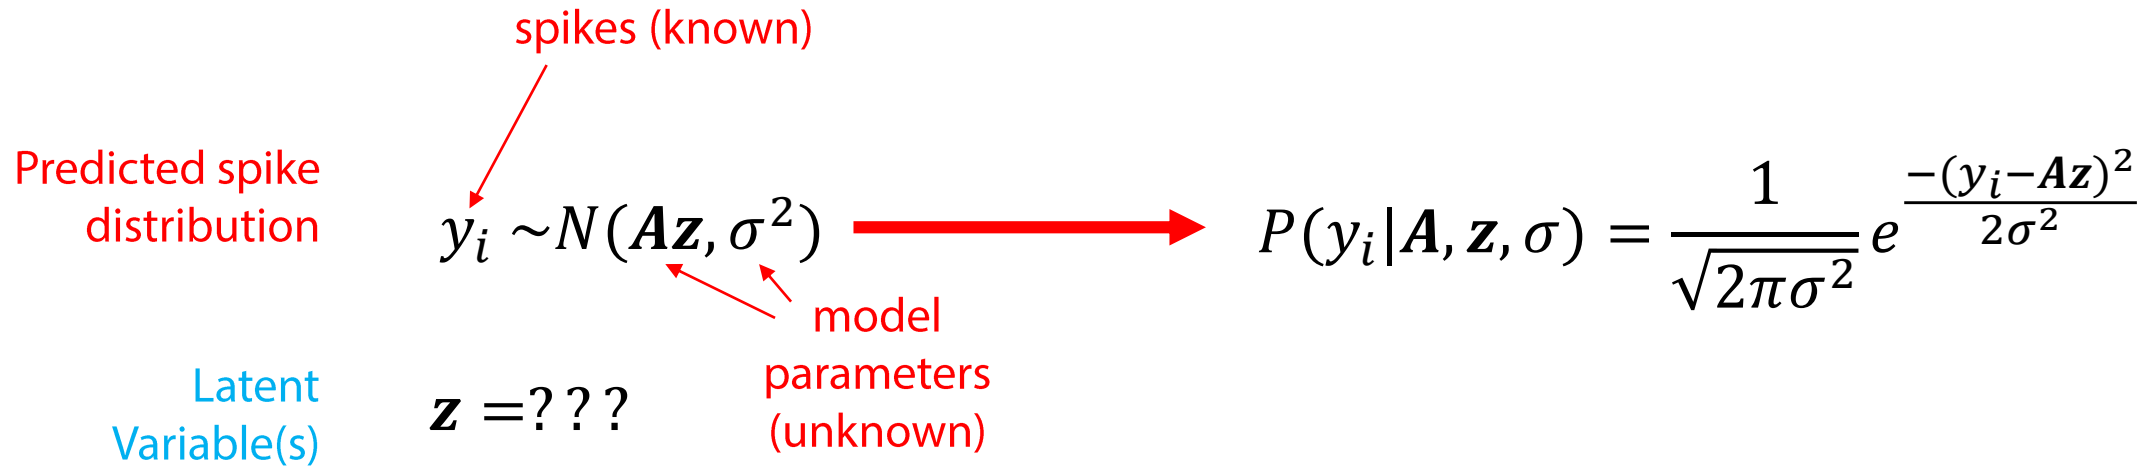
Recall how we fit the GLM...



Encoding model
 $P(\mathbf{y} | \mathbf{x})$



The Expectation Maximization algorithm converges to a local solution



- 1) **Initialization:** make a random guess for model parameters (here \mathbf{A} and σ).
- 2) **Expectation step:** compute the expected value of Z given y , \mathbf{A} , and σ .
- 3) **Maximization step:** re-estimate \mathbf{A} and σ given y and our new estimate of Z .
- 4) **Iterate** steps 2-3 until converged.

The Old Faithful Dataset

