Decoding stimulus information from neural activity: Naïve Bayes

Bi23: Methods in Neural Data Analysis

2/15/2019

Neural encoding

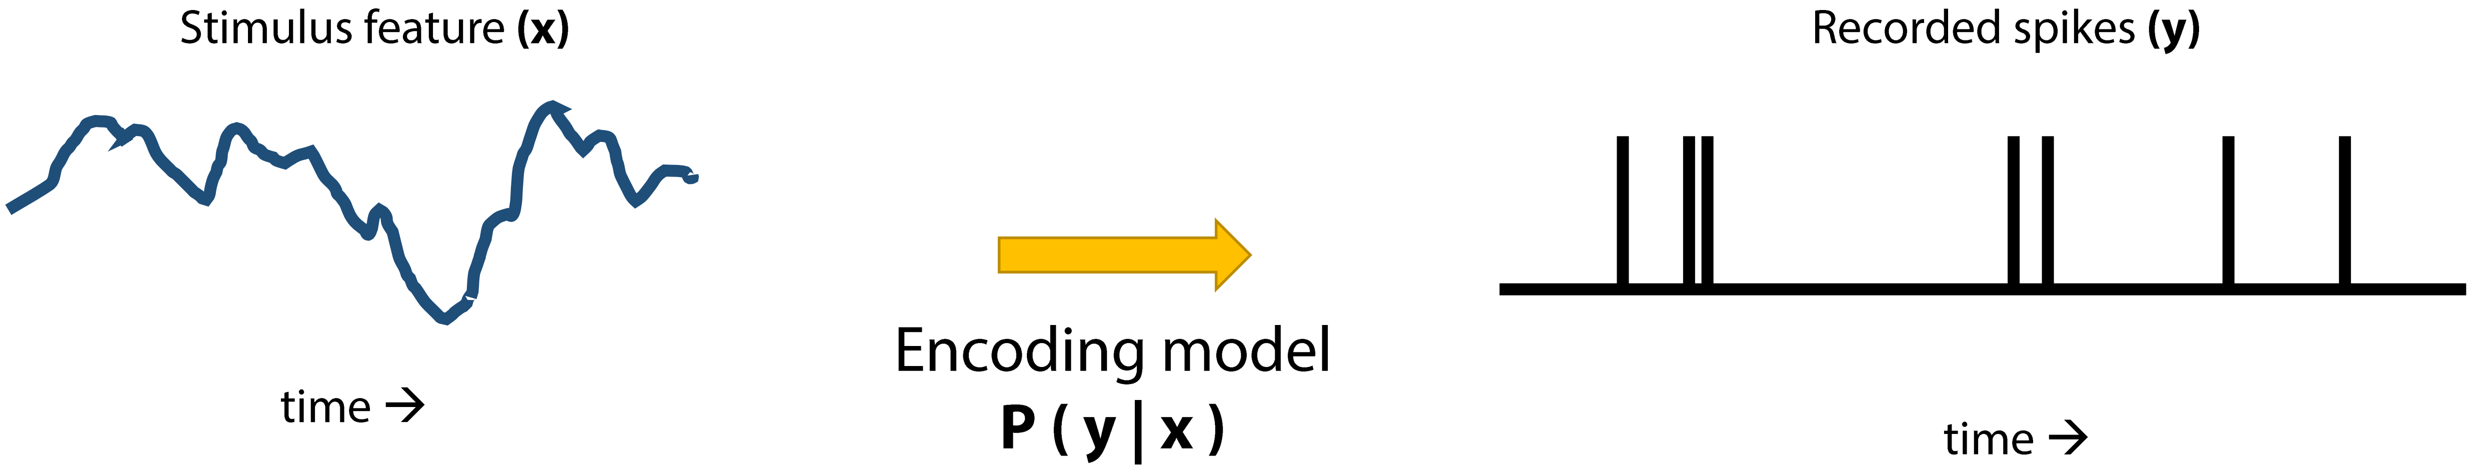

Stimulus feature **(x)**

time →

Encoding model

$P (y \mid x)$

Recorded spikes **(y)**

time →

Given that the animal saw stimulus “x”, what pattern of spiking “y” do we expect to observe in its brain?

Neural decoding

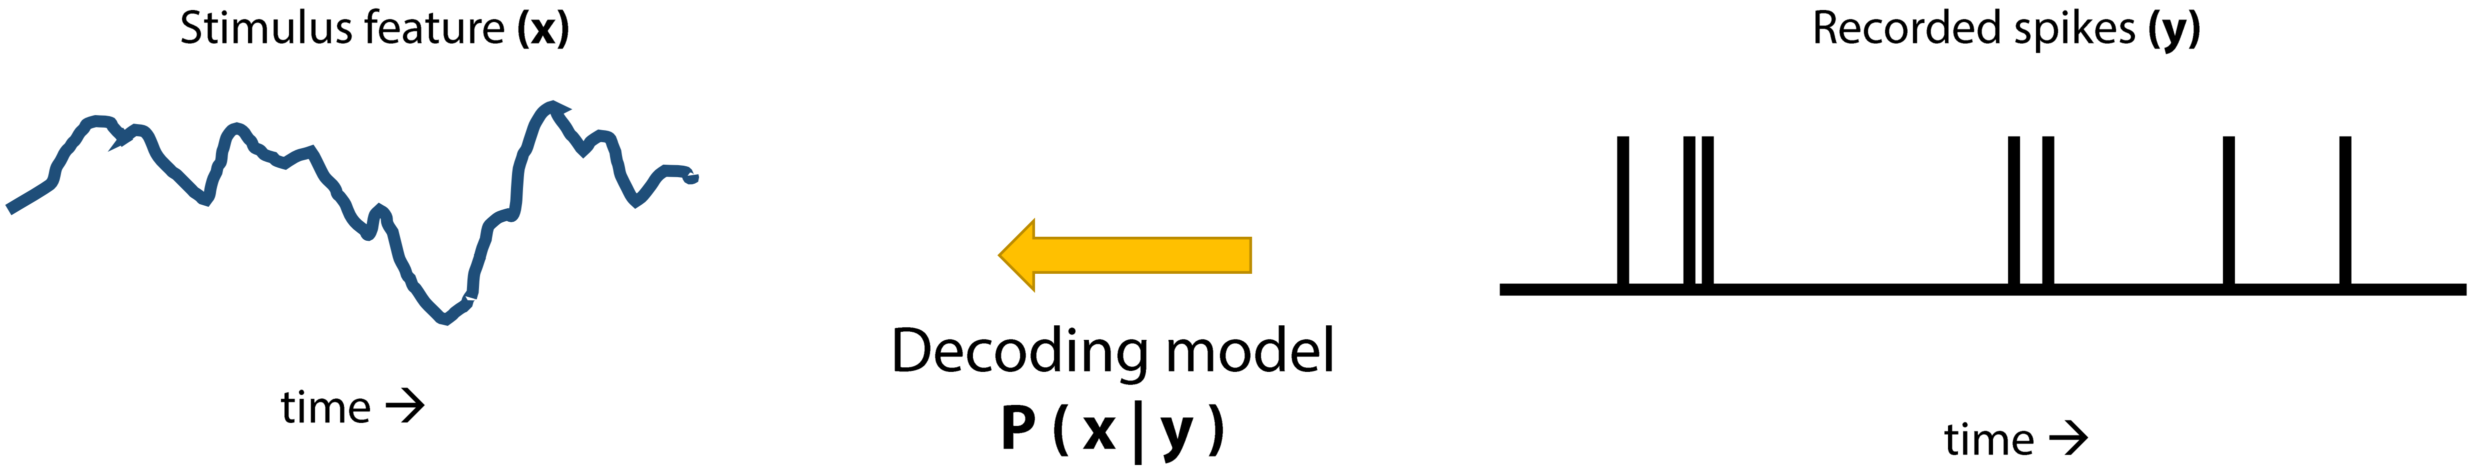

Stimulus feature **(x)**

time →

Decoding model

P (x | y)

Recorded spikes **(y)**

time →

Given the observed spiking “y”, what stimulus/behavior “x” is present?

Decoding from a single neuron

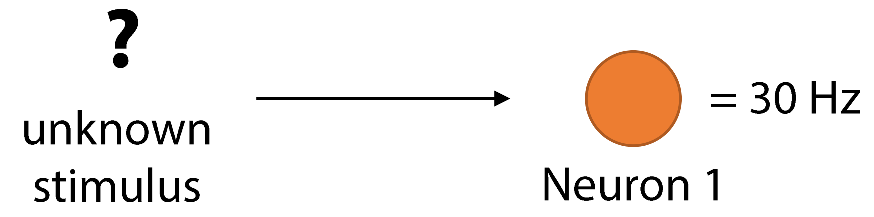

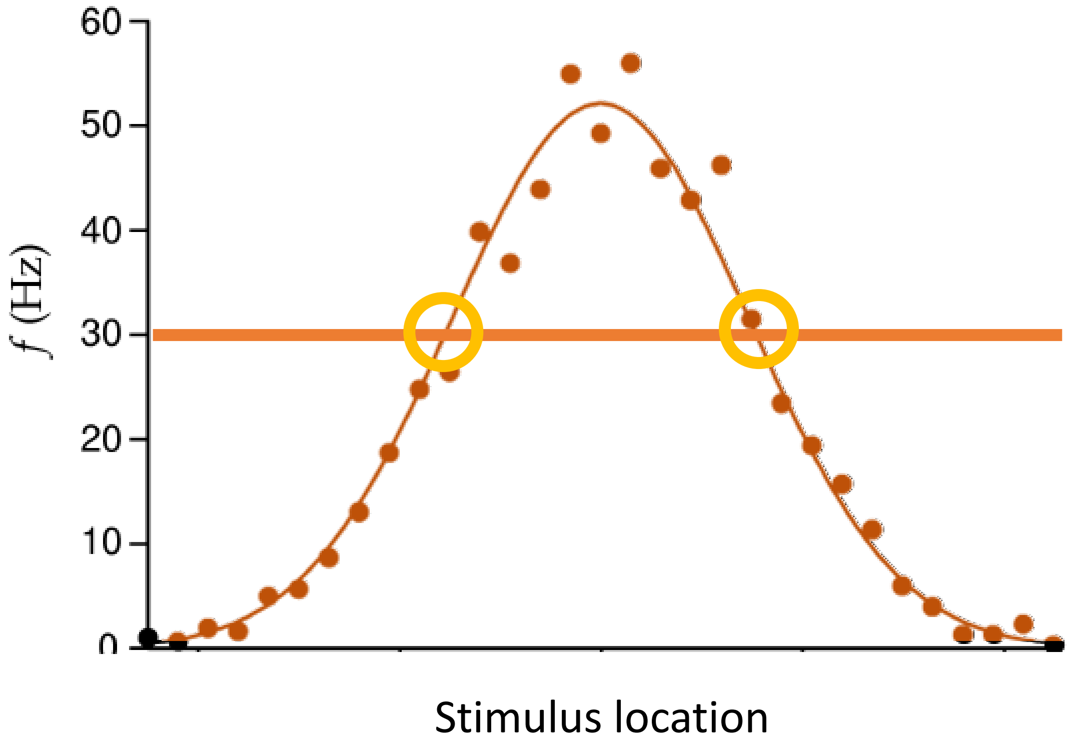

Decoding from multiple neurons

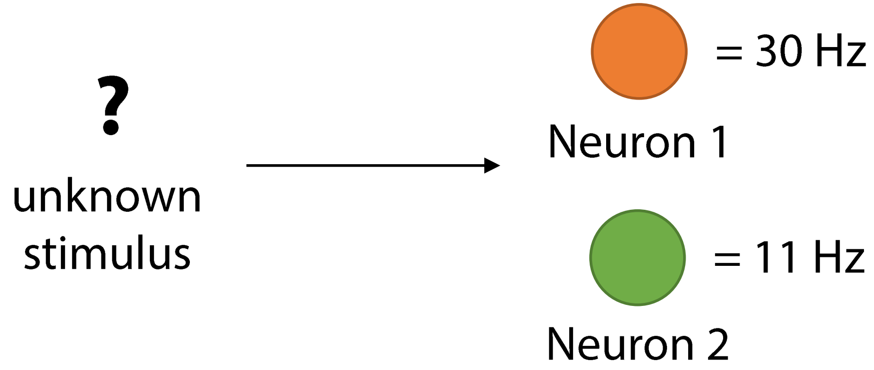

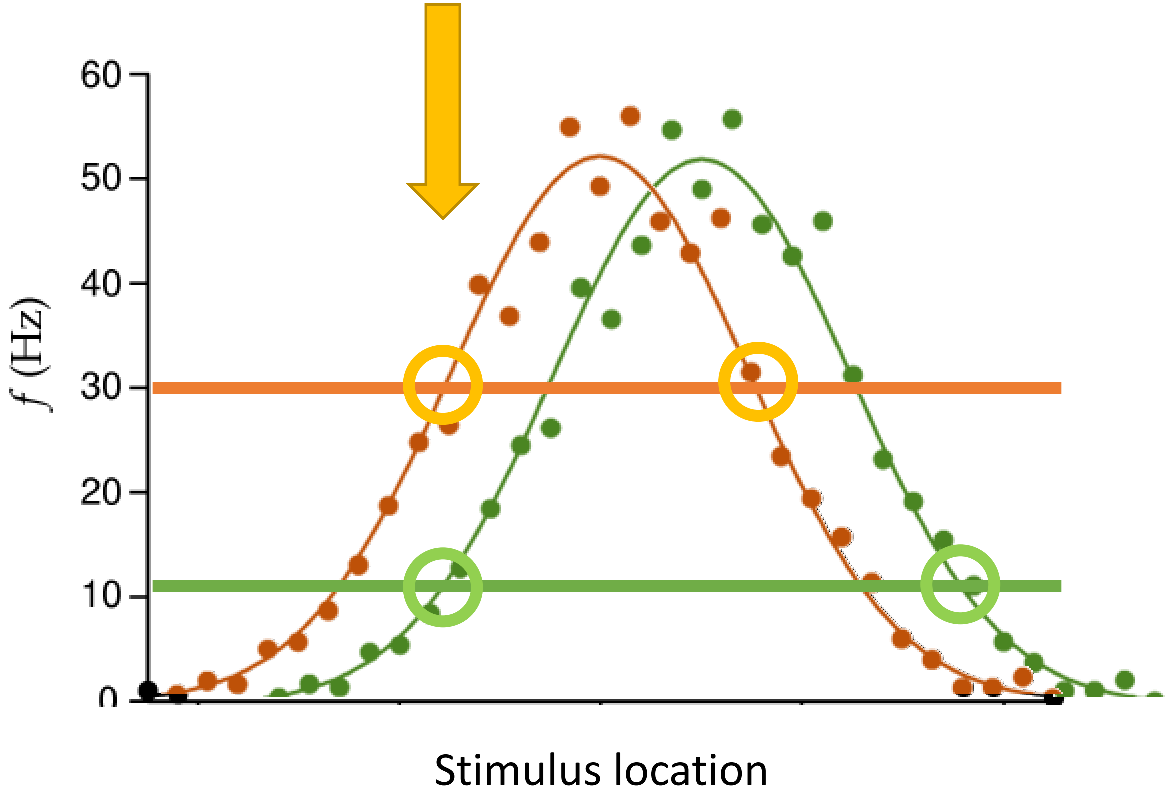

Probably the stimulus location?

A Bayesian perspective on decoding

Likelihood $P(y|x,\theta)$

stimulus "x"
spiking "y"
parameter(s) "θ"

"How likely it is the neuron fires *k* spikes given which stimulus the animal saw."

Posterior $P(x|y,\theta)$

"How likely it is the animal saw a certain stimulus given the observed neuron's spiking."

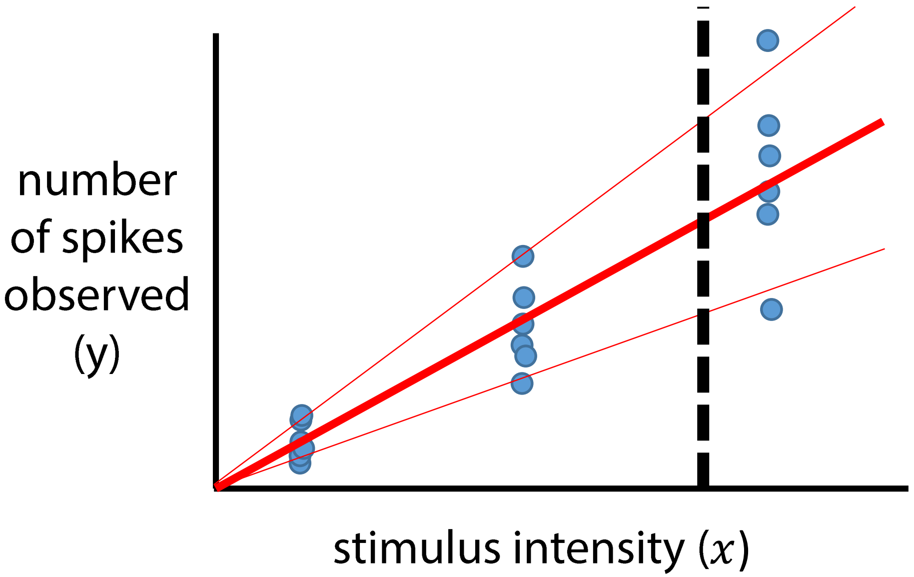

(This is our Poisson GLM from last time) :

Predicted spike distribution $y \sim Poiss(\lambda)$

Underlying spike rate $\lambda = \theta x$

A Bayesian perspective on decoding

stimulus “x”
spiking “y”
parameter(s) “θ”

Likelihood $P(y|x,\theta)$

“How likely it is the neuron fires *k* spikes given which stimulus the animal saw.”

Posterior $P(x|y,\theta)$

“How likely it is the animal saw a certain stimulus given the observed neuron’s spiking.”

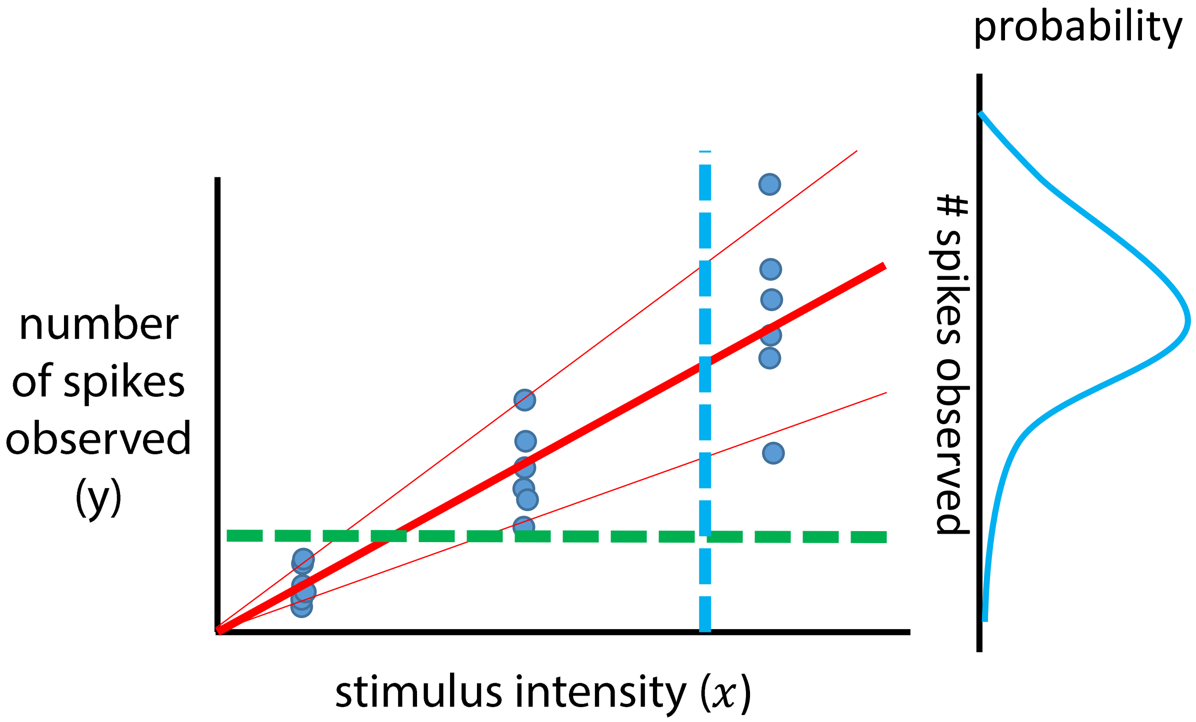

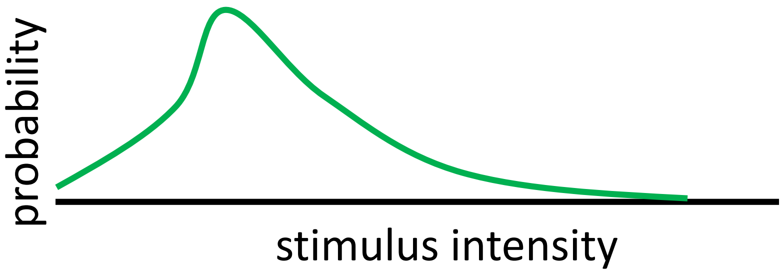

A Bayesian perspective on decoding

stimulus “x”
spiking “y”
parameter(s) “θ”

Likelihood $P(y|x,\theta)$

“How likely it is the neuron fires *k* spikes given which stimulus the animal saw.”

Posterior $P(x|y,\theta)$

“How likely it is the animal saw a certain stimulus given the observed neuron’s spiking.”

Bayes’ theorem

Prior over stimuli

$$P(x|y,\theta) = \frac{P(y|x,\theta)\,P(x)}{P(y)}$$

“Evidence”
(usually omitted)

Decoding from multiple neurons

$$P(x|y_1 y_2 \dots y_N, \theta) = \frac{P(y_1 y_2 \dots y_N|x, \theta)\, P(x)}{P(y)}$$

Hard to compute!

Likelihood $P(y|x, \theta)$

"How likely it is the neuron fires *k* spikes given which stimulus the animal saw."

Posterior $P(x|y, \theta)$

"How likely it is the animal saw a certain stimulus given the observed neuron's spiking."

stimulus "x"
spiking "y"
parameter(s) "θ"

Decoding from multiple neurons: *Naïve Bayes* approach

$$P(x|y_1 y_2 \dots y_N, \theta) = \frac{P(y_1 y_2 \dots y_N | x, \theta)\, P(x)}{P(y)}$$

Assume "conditional independence" (each neuron y is independent of each other neuron given the stimulus):

$$P(x|y_1 y_2 \dots y_N, \theta) = \frac{P(y_1|x,\theta)\, P(y_2|x,\theta) \dots P(y_N|x,\theta) P(x)}{P(y)}$$

Likelihood $P(y|x,\theta)$

"How likely it is the neuron fires *k* spikes given which stimulus the animal saw."

Posterior $P(x|y,\theta)$

"How likely it is the animal saw a certain stimulus given the observed neuron's spiking."

stimulus "x"
spiking "y"
parameter(s) "θ"

Now combine the Naïve Bayes model with a *decision rule:*

$$P(x|y_1 y_2 \dots y_N, \theta) = \frac{P(y_1|x,\theta)\, P(y_2|x,\theta) \dots P(y_N|x,\theta) P(x)}{P(y)}$$

Given observed spikes from each neuron at a given time point: $y_1, y_2 \dots y_N$

$$\hat{x} = \underset{m=1\dots M}{\operatorname{argmax}}\ P(x_m) \prod_{n=1}^{N} P(y_n|x_m)$$

Estimated stimulus

(this is called the Maximum A Priori (MAP) estimate of x)

Likelihood $P(y|x,\theta)$

"How likely it is the neuron fires *k* spikes given which stimulus the animal saw."

Posterior $P(x|y,\theta)$

"How likely it is the animal saw a certain stimulus given the observed neuron's spiking."

stimulus "x"
spiking "y"
parameter(s) "θ"

Data set for this week: rats running in mazes

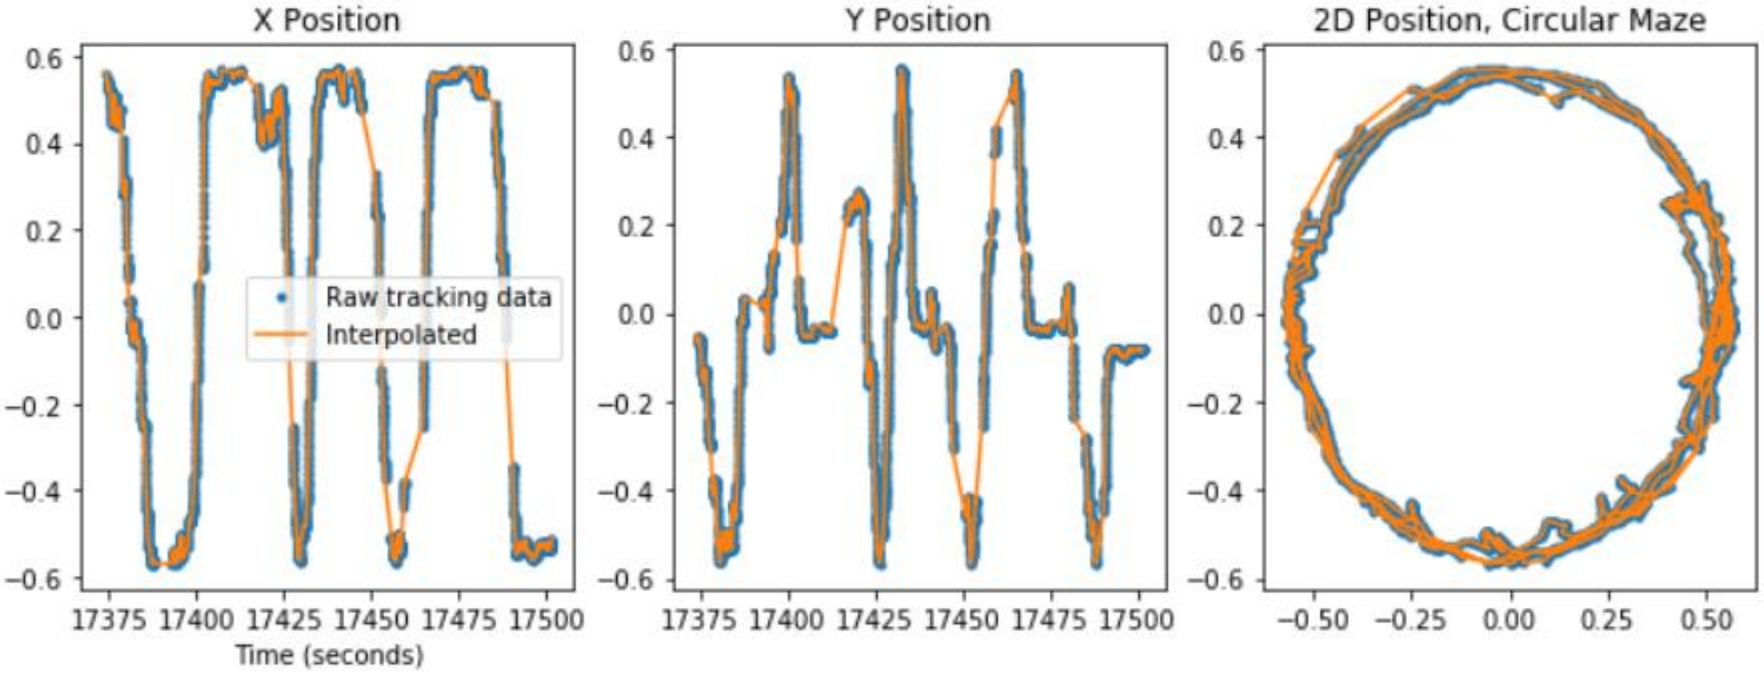

Evaluating decoder performance: confusion matrix

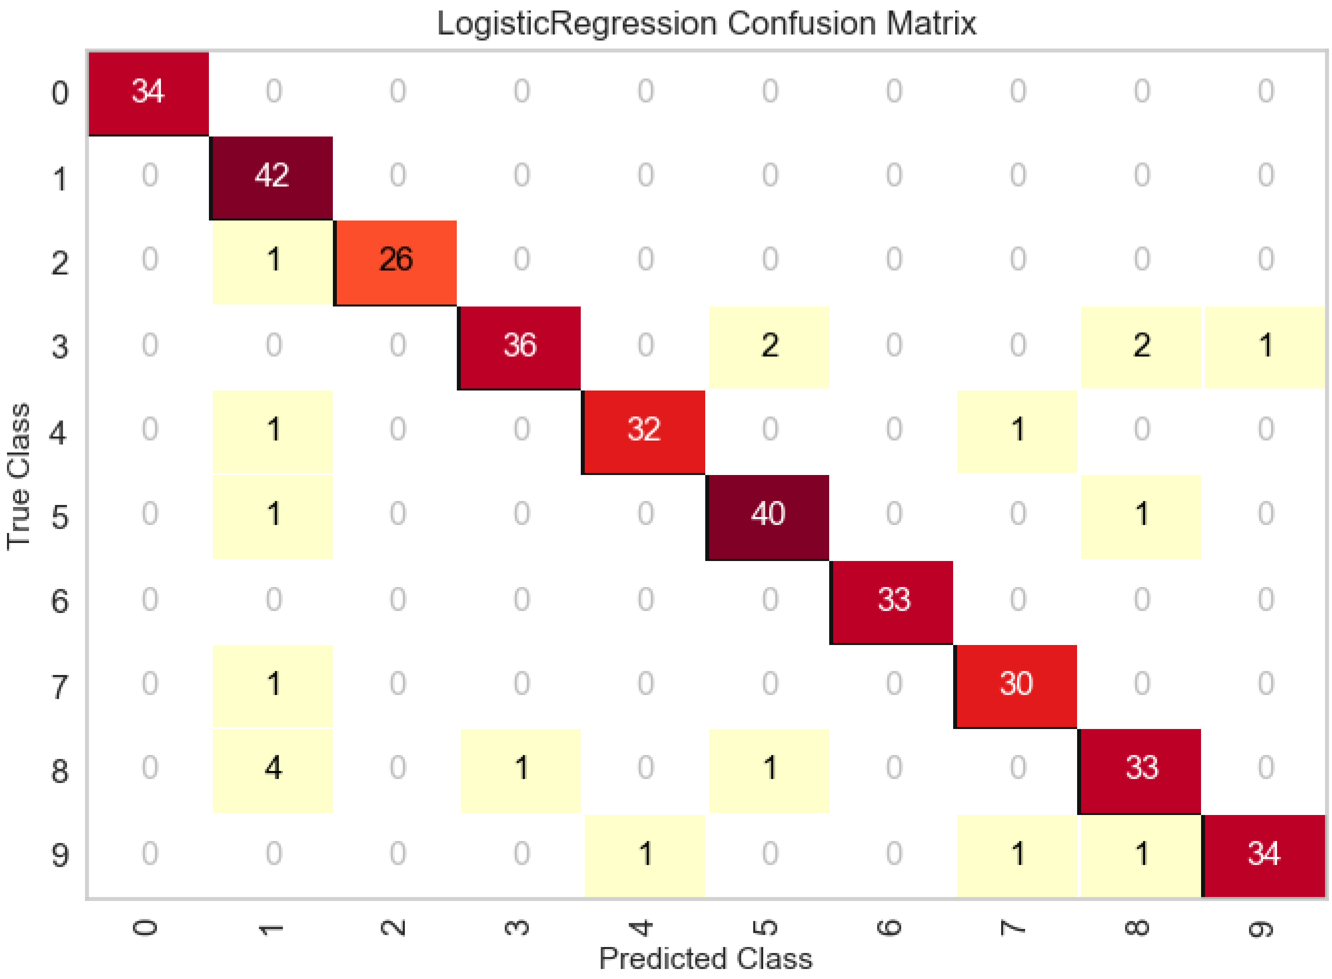

LogisticRegression Confusion Matrix

	0	1	2	3	4	5	6	7	8	9
0	34	0	0	0	0	0	0	0	0	0
1	0	42	0	0	0	0	0	0	0	0
2	0	1	26	0	0	0	0	0	0	0
3	0	0	0	36	0	2	0	0	2	1
4	0	1	0	0	32	0	0	1	0	0
5	0	1	0	0	0	40	0	0	1	0
6	0	0	0	0	0	0	33	0	0	0
7	0	1	0	0	0	0	0	30	0	0
8	0	4	0	1	0	1	0	0	33	0
9	0	0	0	0	1	0	0	1	1	34