Model-based fMRI analysis

Shih-Wei Wu

fMRI workshop, NCCU, Jan. 19, 2014

Outline: model-based fMRI analysis

I: General linear model: basic concepts

II: Modeling how the brain makes decisions: Decision-making models

III: Modeling how the brain learns: Reinforcement learning models

IV: Modeling response dynamics: Drift diffusion model

Model-based fMRI

Problem: Characterizing mental operations

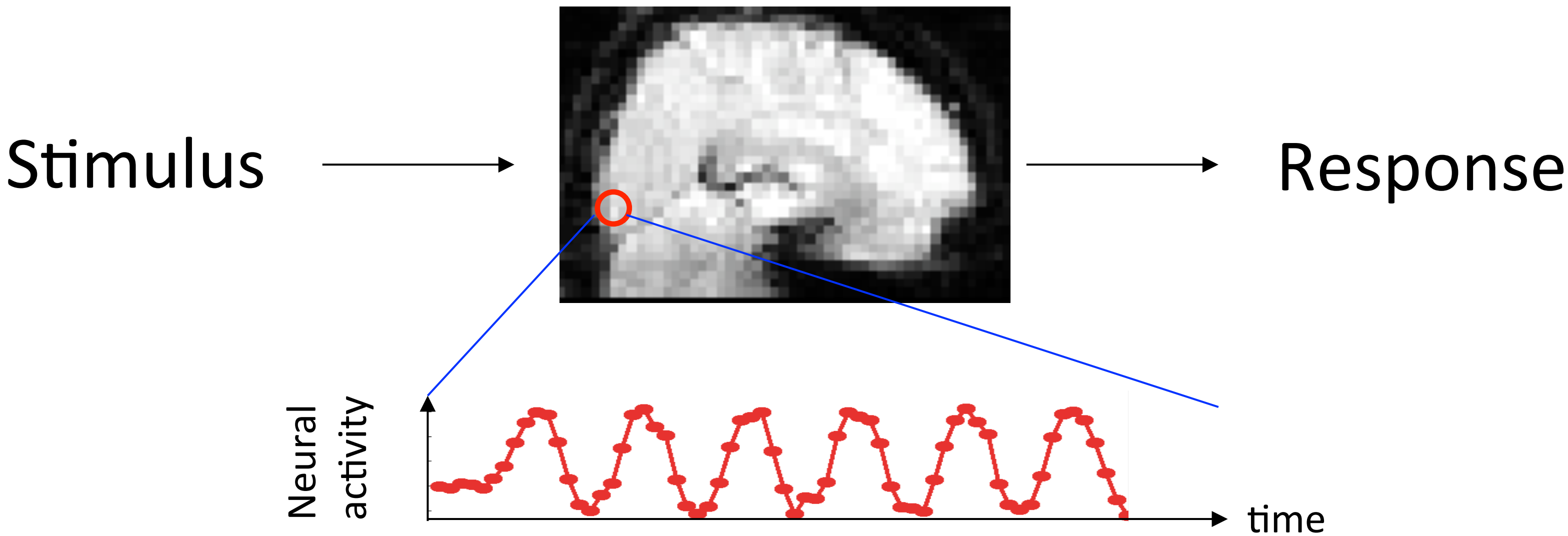

How do we characterize the mental operations involved in producing behavioral response given the stimulus?

Model-based fMRI

Problem: Characterizing mental operations

1. How do we characterize the mental operations involved in producing behavioral response given the stimulus?

Stimulus →  → Response

Mental operations cannot be simply represented by the observable: stimulus and response

Model-based fMRI

Problem: Characterizing mental operations

1. How do we characterize the mental operations involved in producing behavioral response given the stimulus?

2. Mental operations cannot be simply represented by the observable: stimulus and response

Stimulus →  → Response

Model-based fMRI

Problem: Characterizing mental operations

One option: Build or apply some computational model that

Stimulus →  → Response

↙ quantitatively describes the mental operations

↘ describes the behavioral performance

Univariate analysis

- Each voxel in the brain is analyzed *separately*

- Each voxel presents a time-series data

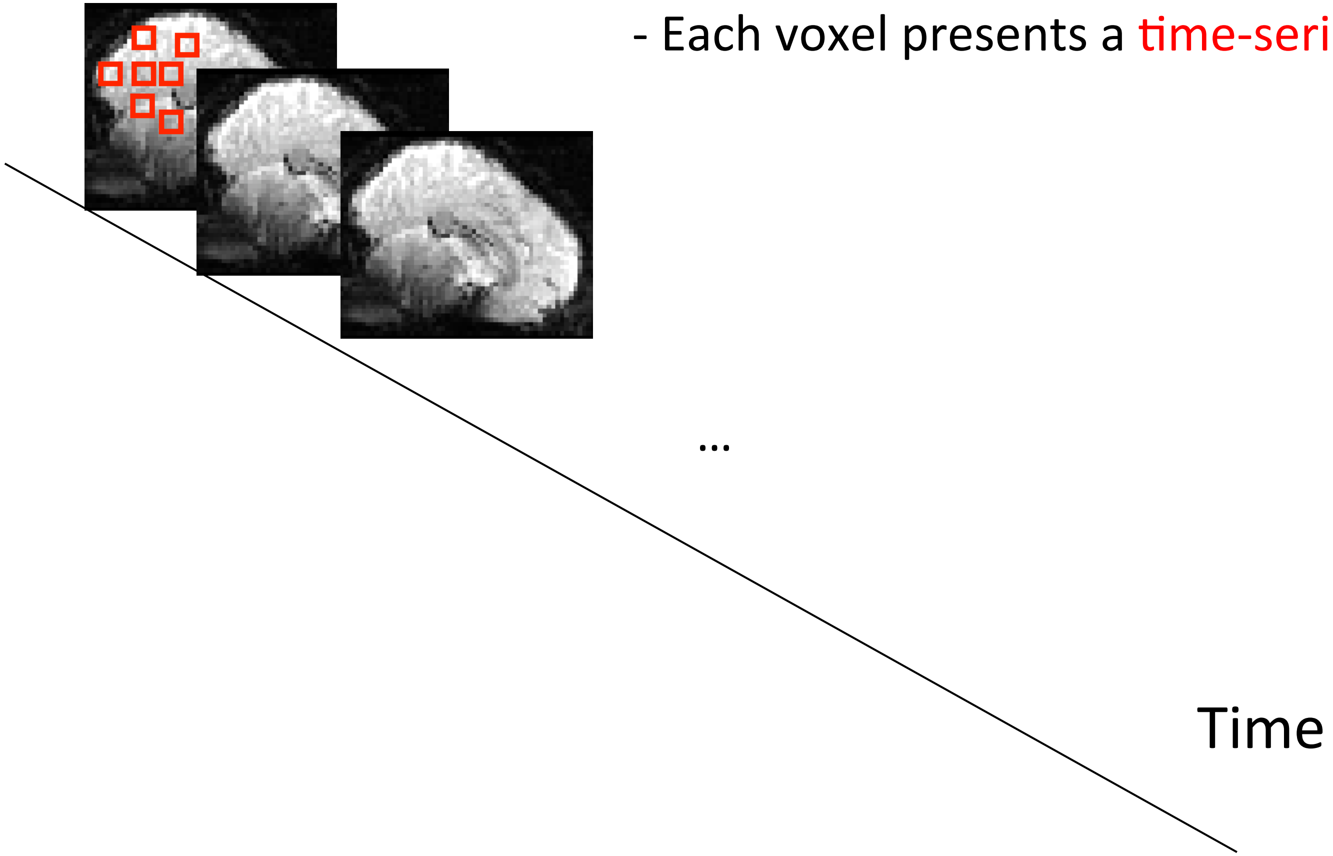

Time-series data

- Suppose you have the following experiment

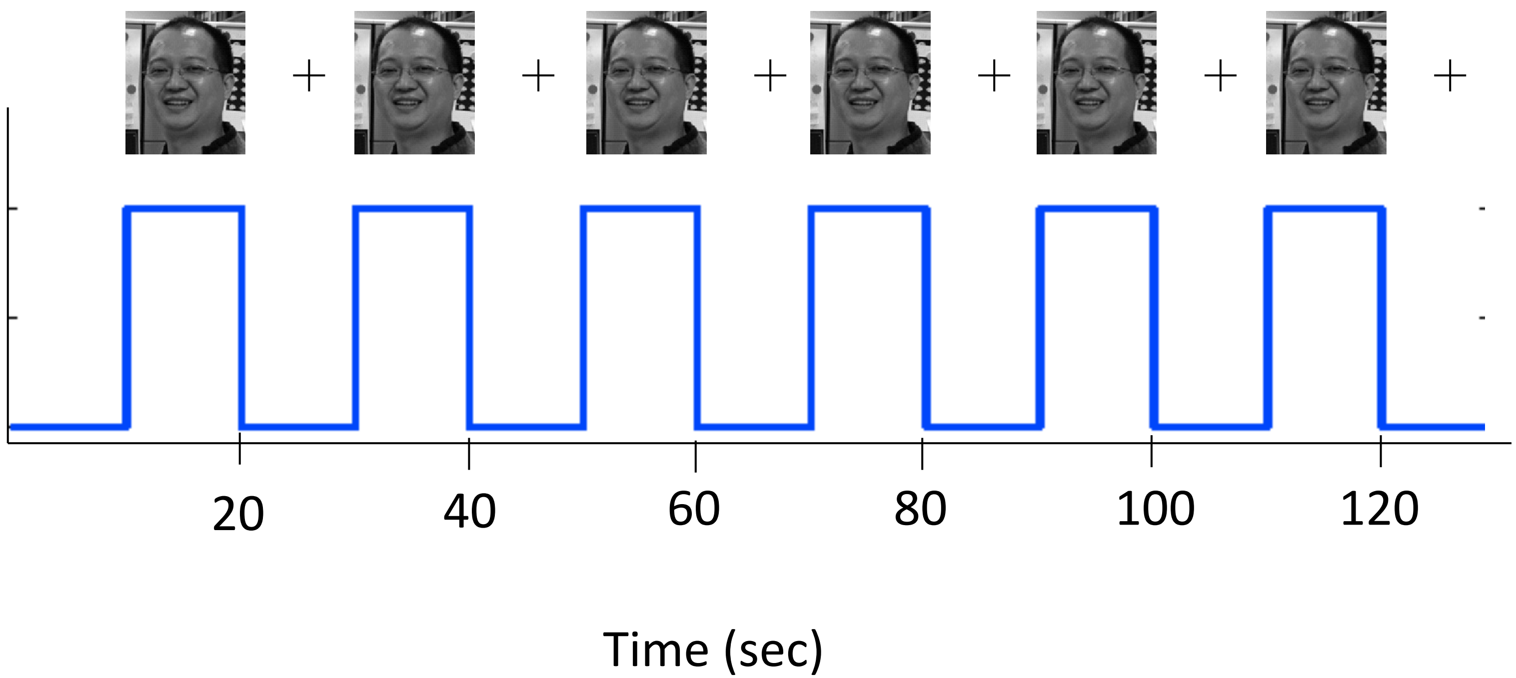

Time-series data

- This is the data you get (from a single voxel)

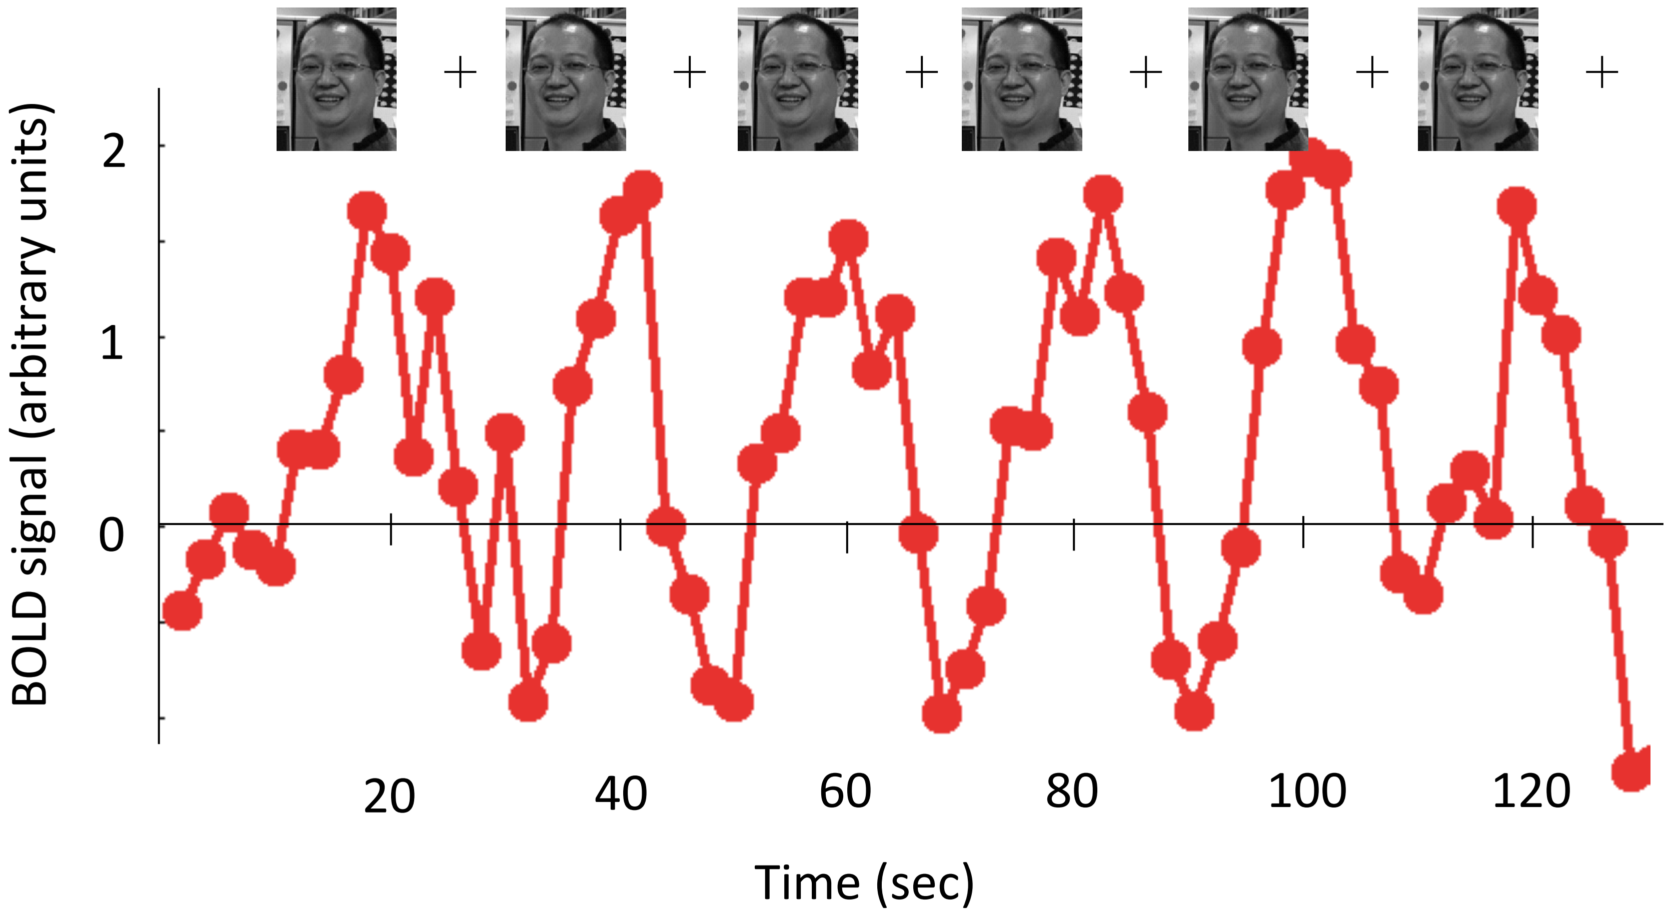

Time-series data

- When you compare prediction (based on your design) and data, you realize that there is somewhat a match, but not close

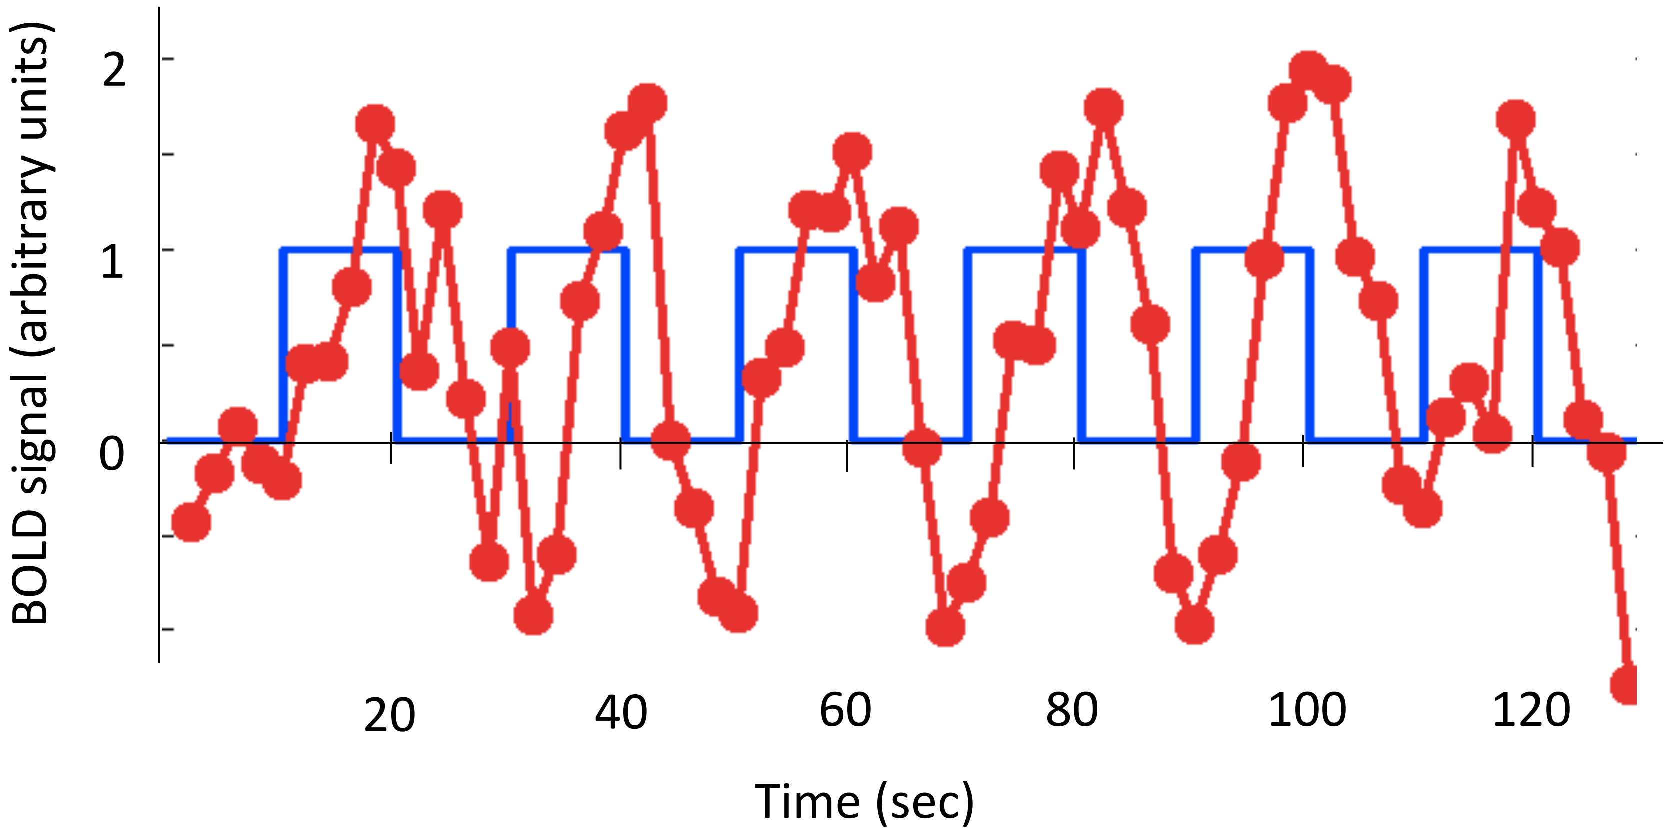

Time-series data

- What about this one? Which aspect of the comparison is the same, which aspect might be different?

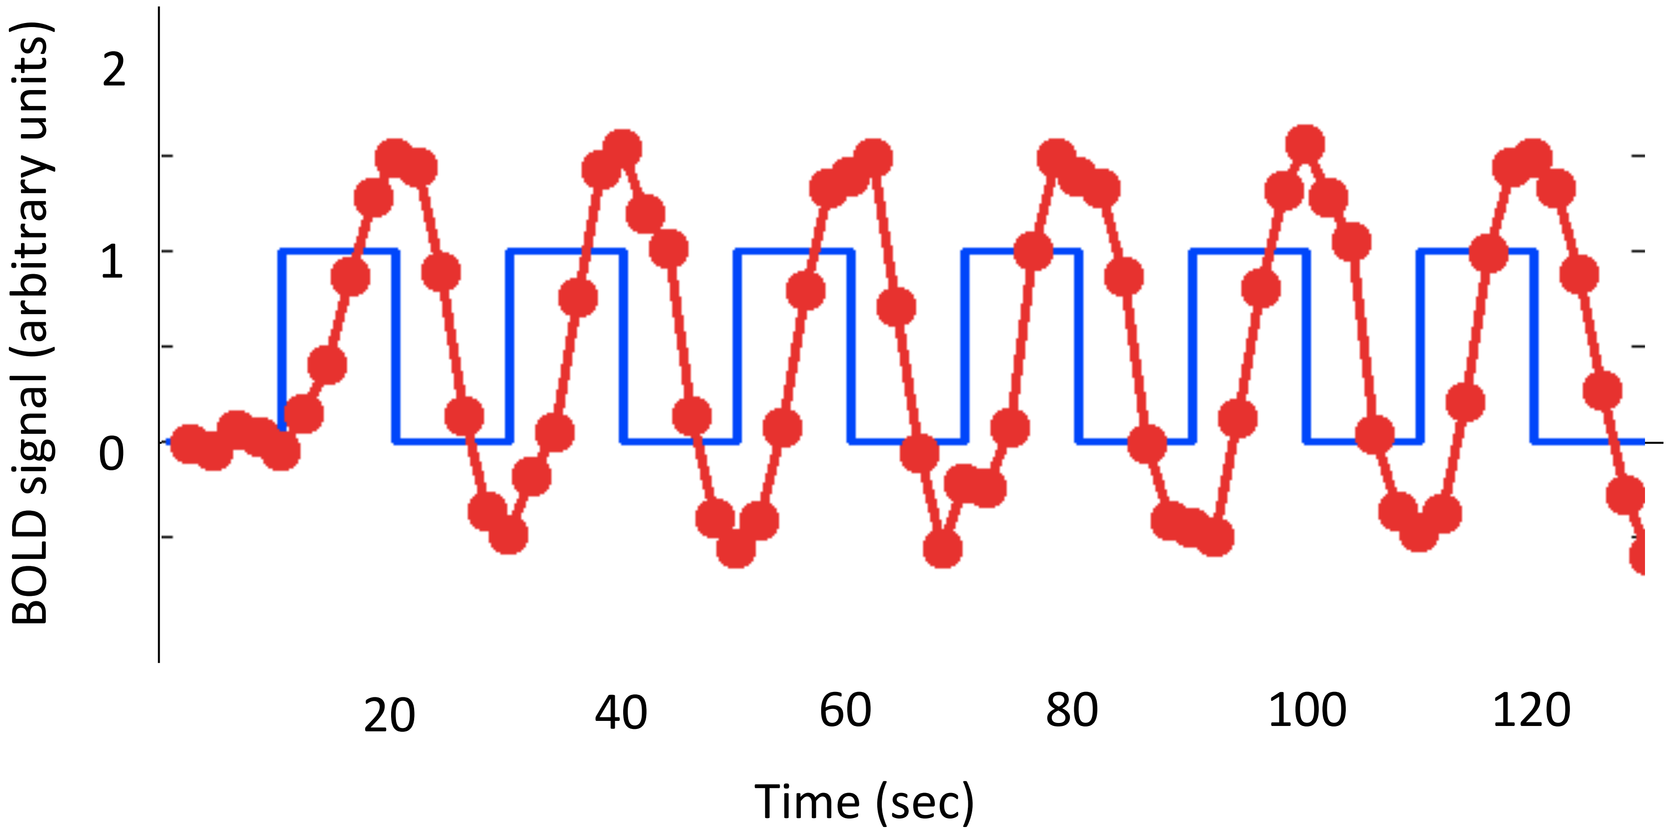

I. General linear model: basic concepts

BOLD signal

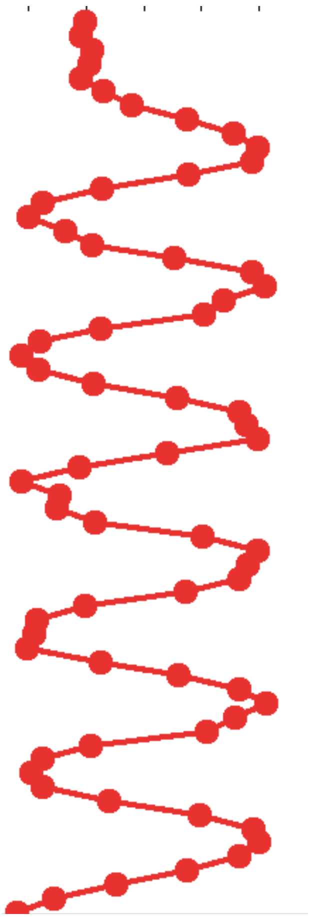

$= \beta$ X

Design matrix

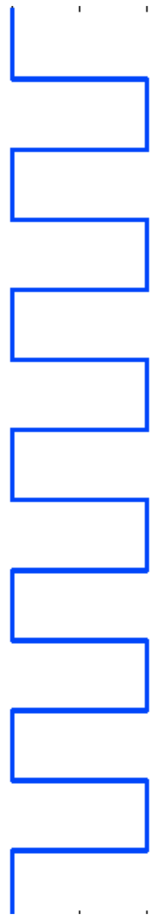

+ noise

Parameter estimate: this is what we are interested in

I. General linear model: basic concepts

$$Y = X\beta + \varepsilon, \varepsilon \sim N(0, \sigma^2)$$

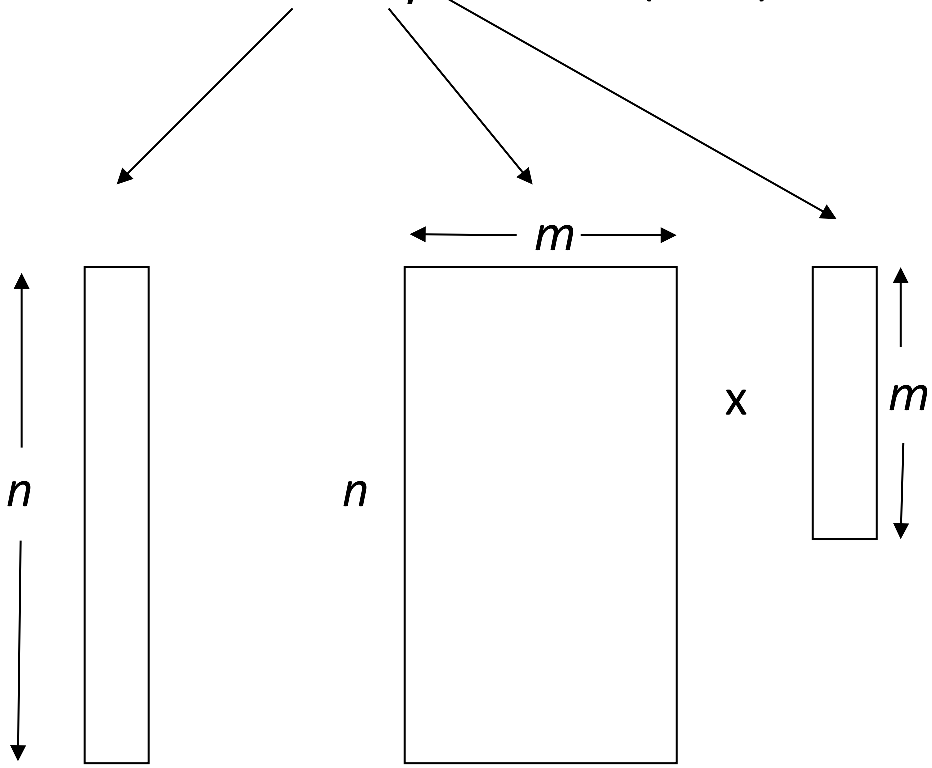

BOLD times series   Design matrix   Parameter vector

(BOLD: Blood Oxygenation Level Dependent)

Modeling how the brain makes decisions: Decision-making models

II. Modeling how the brain makes decisions

Losses loom larger than gains

贏 輸

- The psychological impact of a loss (or potential loss) is greater than the same-sized gain (or potential gain)

Prospect theory

- *Value function*

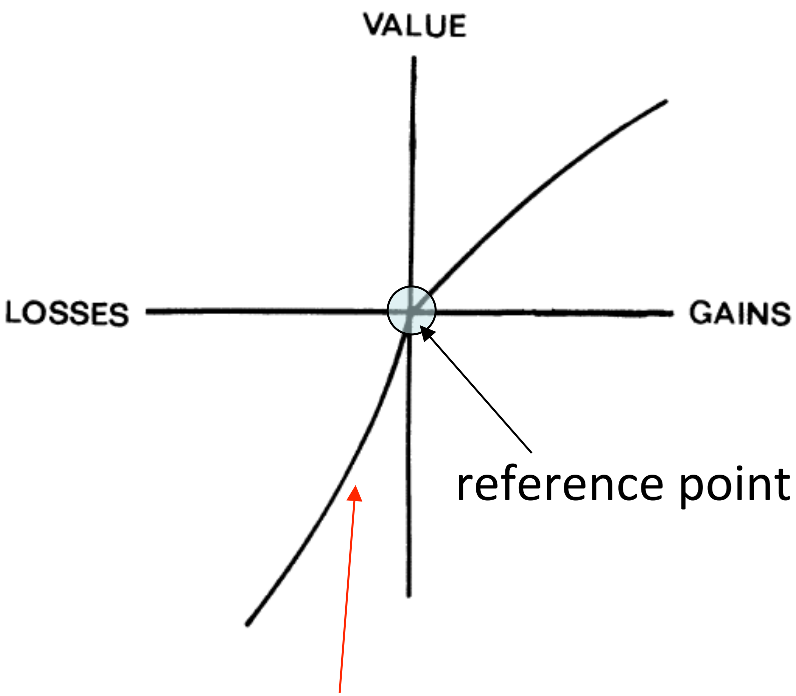

Loss aversion: Steeper slope
in the loss domain compared with gains

$$V(x) = \begin{Bmatrix} x^{\alpha}, x \geq 0 \\ -\lambda(-x)^{\beta} \end{Bmatrix}$$

λ : Controls the degree of loss aversion

Implication of loss aversion on choice behavior

Example: Is (gain $2000,50%; Lose $1000,50%) an attractive gamble?

Suppose

$$\alpha = 1, \beta = 1, \lambda = 2$$

Then the value of the gamble is

$$V(\$2000) \cdot 0.5 + V(-\$1000) \cdot 0.5$$
$$= 2000 \cdot 0.5 - 2 \cdot 1000 \cdot 0.5 = 0$$

This gamble is not attractive at all to the decision maker and hence it is not likely that s/he is going to bet on it

Prospect theory in the brain?

1. How does the brain represent gains and losses?

2. Is there a way to explain loss aversion from a neurobiological perspective?

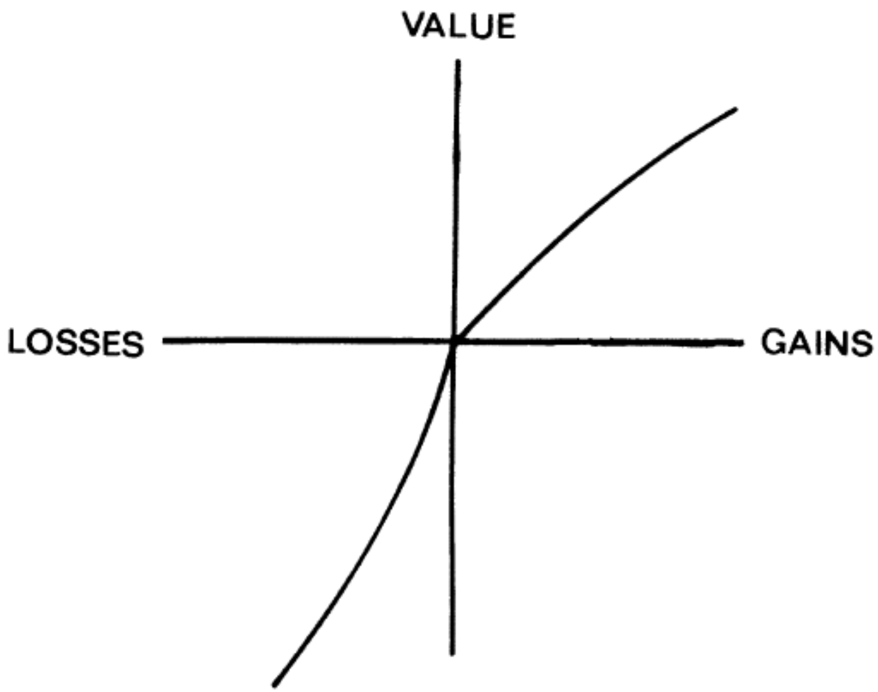

II. Modeling how the brain makes decisions

Neural basis of loss aversion

- Tom et al. (2007, Science): A decision-making experiment involving monetary gains and losses

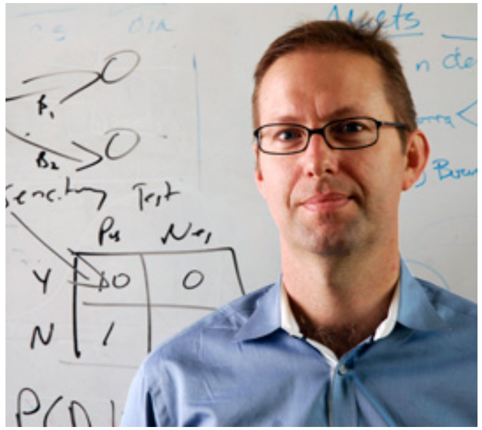

Russell Poldrack

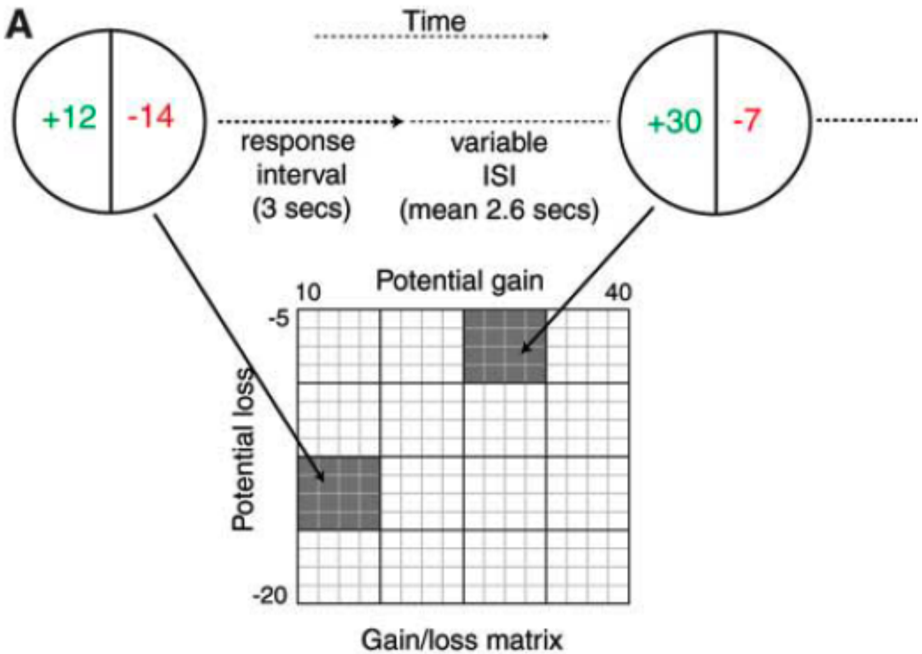

- Subjects in each trial had to decide whether to accept a gamble (Gain,50%; Loss,50%)

- Gain and loss in each trial were decided independently

Measuring loss aversion by choice behavior

- Suppose this is the data from a subject

Trial	Gain	Loss	Choice (yes/no)
1	$1000	$800	0
2	$2000	$1000	0
3	$350	$450	0
4	$500	$100	1
5	$1200	$380	1
6	$60	$55	0
7	$290	$148	0
.	.	.	.
.	.	.	.
.	.	.	.

Measuring loss aversion by choice behavior

- We can estimate how much loss averse a subject is based on his /her choice data

Trial	Gain	Loss	Choice (yes/no)
1	$1000	$800	0
2	$2000	$1000	0
3	$350	$450	0
4	$500	$100	1
5	$1200	$380	1
6	$60	$55	0
7	$290	$148	0
.	.	.	.
.	.	.	.
.	.	.	.

- Method: Logistic regression (a statistical method)

$$choice = \beta_G Gain + \beta_L Loss$$

β_G : how strong gains contribute to choice data

β_L : how strong losses contribute to choice data

- Degree of loss aversion

$$\lambda_{behavior} = \frac{-\beta_L}{\beta_G}$$

II. Modeling how the brain makes decisions

Measuring loss aversion by neural activity

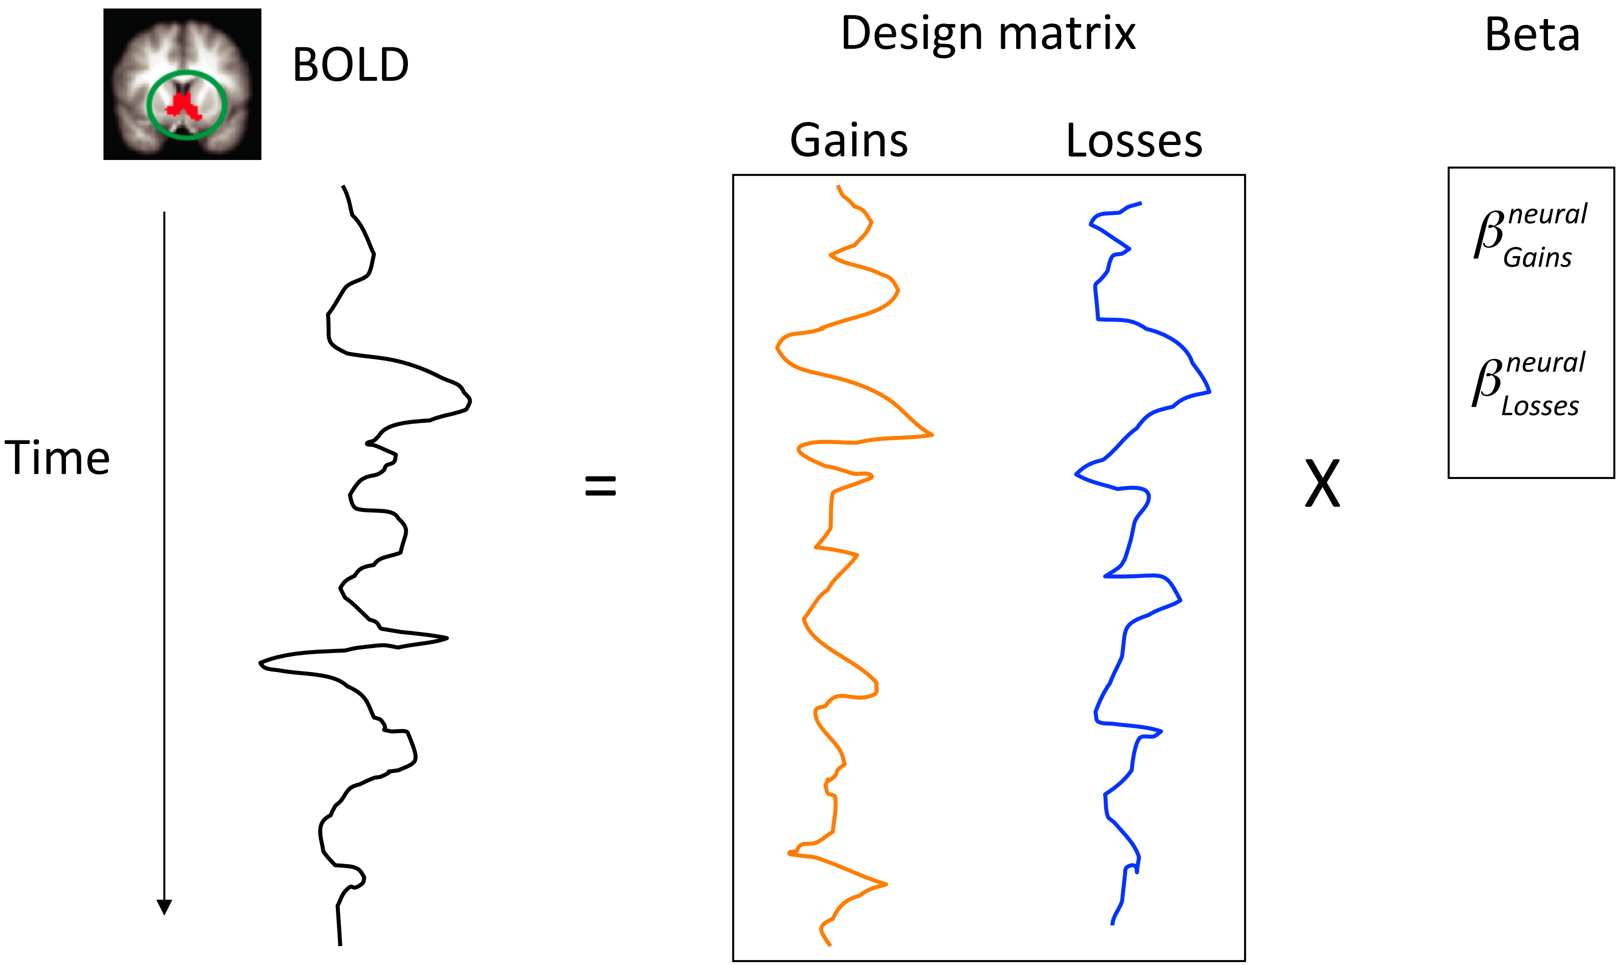

Neural measure of loss aversion: $\lambda_{neural} = -\beta^{neural}_{Losses} - \beta^{neural}_{Gains}$

Analysis focus

1. How does the brain represent gains and losses?

   - Prospect theory indicates positive correlation with gains, negative correlation with losses; if this is the case, then

$$\beta_{Gains}^{neural} : positive$$

$$\beta_{Losses}^{neural} : negative$$

Analysis focus

2. Neural basis of loss aversion

- An area driving (contributing to) loss aversion should exhibit a close match between loss aversion measured in behavior and loss aversion measured according to its neural activity (psychometric -neurometric match)

$$\lambda_{behavior} \propto \lambda_{neural}$$

Neural representation of gains and losses

- vmPFC and ventral striatum positively correlated with gains and negatively correlated with losses

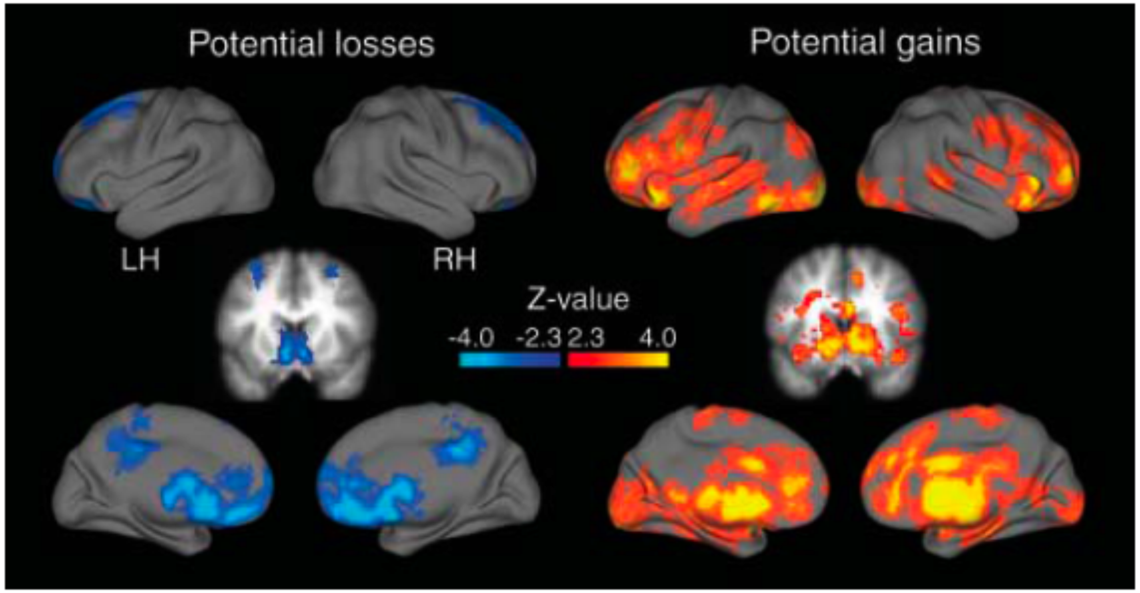

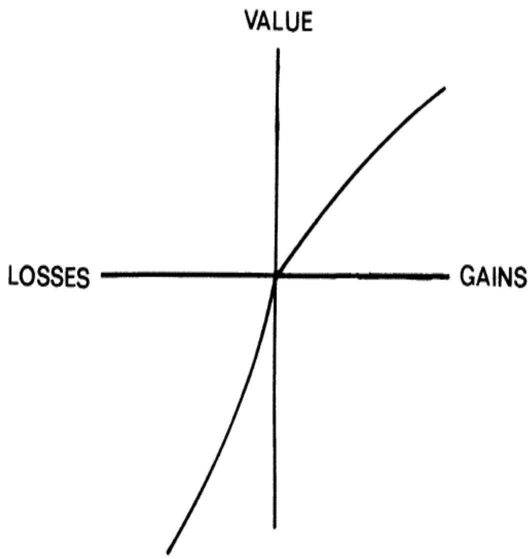

Neural basis of loss aversion

- Neural measure of loss aversion in ventral striatum strongly correlated with behavioral measure of loss aversion

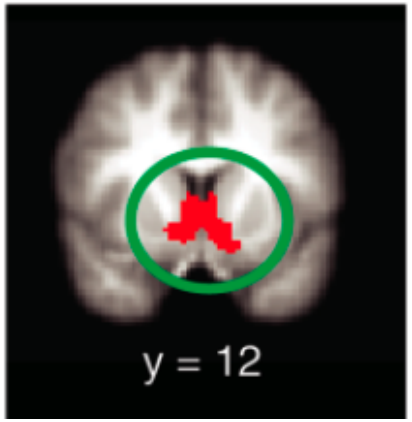

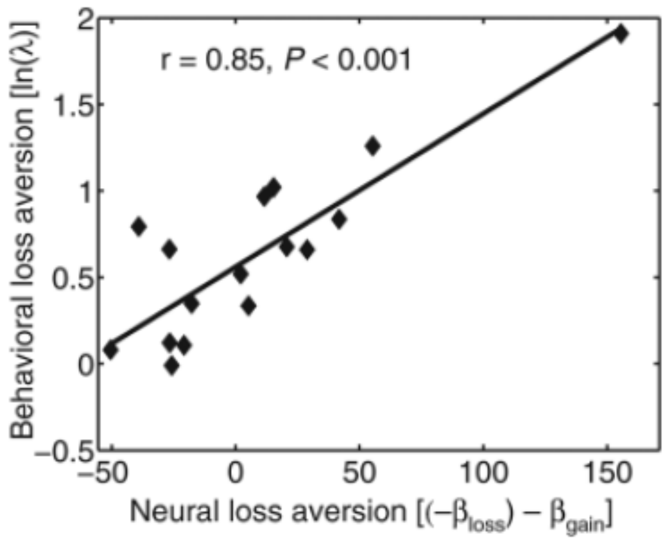

III. Modeling how the brain learns: Reinforcement learning models

III. Modeling how the brain learns

Example: O' Doherty et al. (2003)

Pavlovian conditioning task

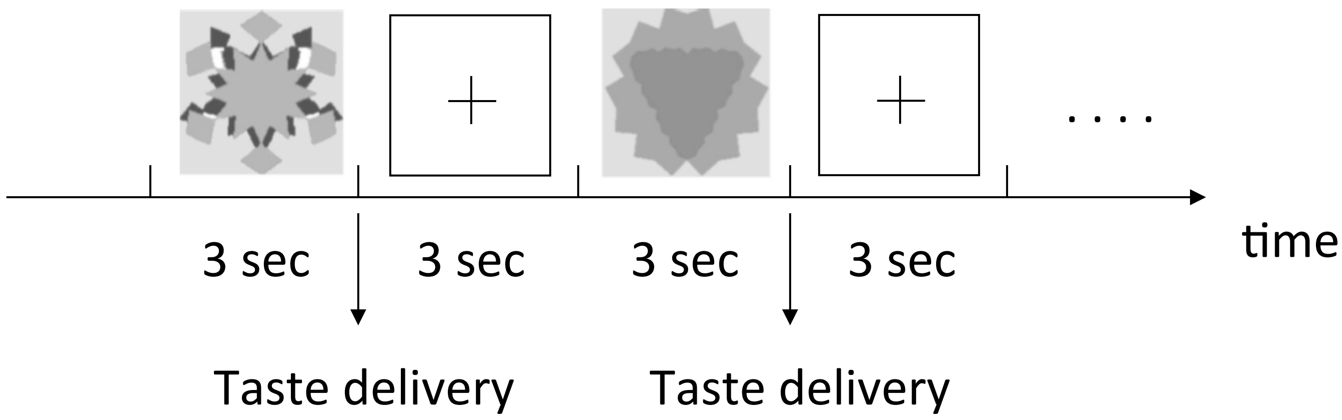

III. Modeling how the brain learns

Example: O’ Doherty et al. (2003) Pavlovian conditioning task

Stimulus-reward associations

	# trials delivery	# trials no delivery
	80	20
	80	20
	# trials = 80	

*Randomization on stimulus order and delivery/no delivery

III. Modeling how the brain learns

Example: O’ Doherty et al. (2003) Pavlovian conditioning task

Stimulus-reward associations

	# trials delivery	# trials no delivery
	80	20
	80	20
	# trials = 80	

*Randomization on stimulus order and delivery/no delivery

III. Modeling how the brain learns

Example: O' Doherty et al. (2003) Pavlovian conditioning task

- Behavioral results:

  The animals exhibit conditioned response (salivate when seeing the stimulus) after experiencing the stimulus-reward pairing

III. Modeling how the brain learns

- Activity of midbrain dopamine neurons (Schultz et al. 1997)

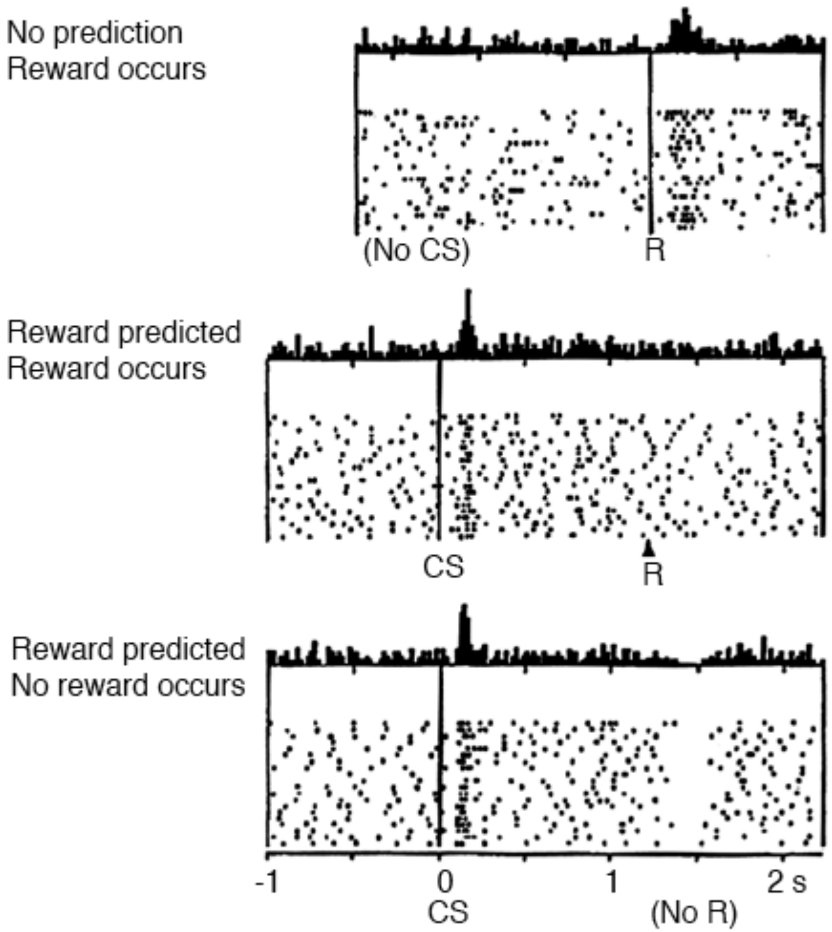

Burst of firing after reward delivery

Burst of firing after stimulus; activity remained at baseline at reward delivery

Burst of firing after stimulus; activity dipped down when no reward was delivered

III. Modeling how the brain learns

Question:

How do we characterize the learning process (learning the association between stimulus and reward) that takes place in the brain?

III. Modeling how the brain learns

Observations:

1. Midbrain DA neurons first showed an increase in firing in response to reward delivery

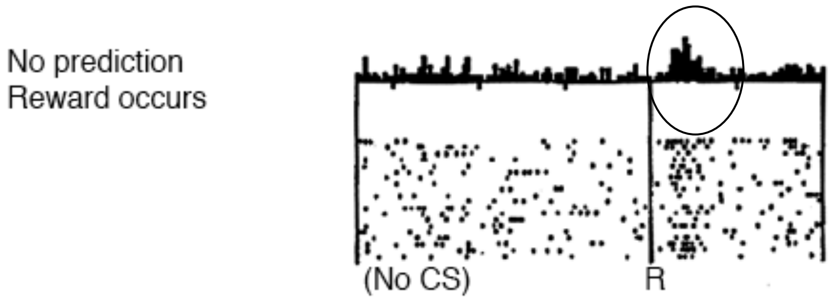

III. Modeling how the brain learns

Observations:

2. When a stimulus is paired with a reward, midbrain DA neurons gradually (over the course of experiment) 'shift' their responses to the time the stimulus is presented

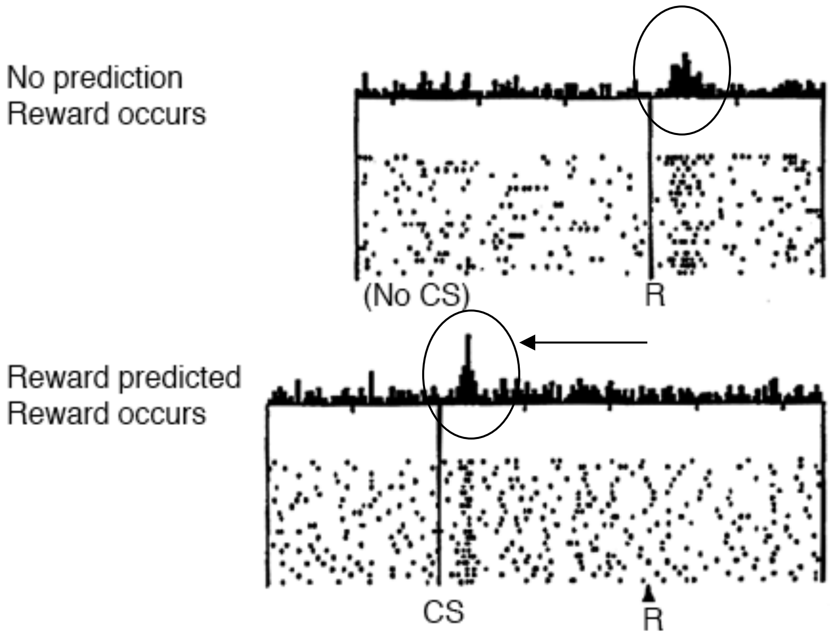

III. Modeling how the brain learns

Observations:

3. When a reward is expected after a stimulus is presented and when the reward is indeed delivered, no change in DA response

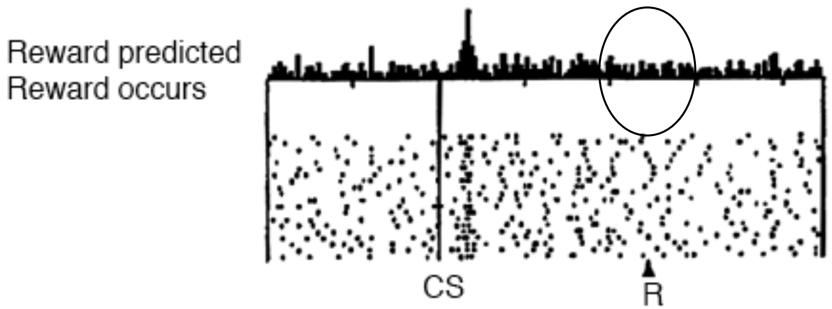

III. Modeling how the brain learns

Observations:

4. When a reward is expected after a stimulus is presented and when the reward is NOT delivered, there is a *decrease* in DA response

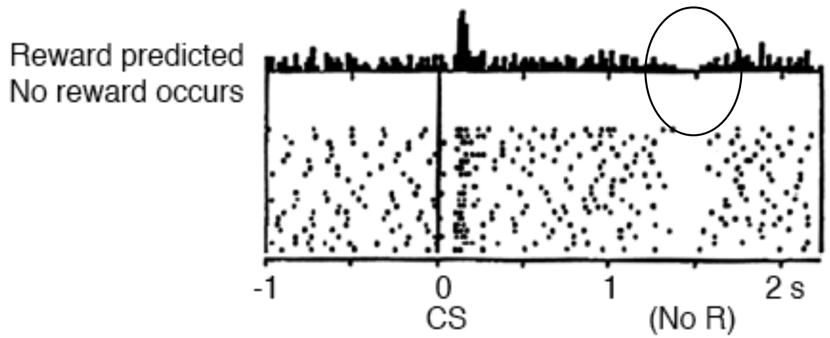

III. Modeling how the brain learns

Example: O' Doherty et al. (2003) Pavlovian conditioning task

Questions:

- How could we explain the 4 observations we just made about neural activity in midbrain DA neurons?

- How could we quantitatively describe or even predict the neuronal response profiles?

III. Modeling how the brain learns

Example: O' Doherty et al. (2003) Pavlovian conditioning task

A solution:

Apply a computational learning model: Temporal difference (TD) model (Sutton & Barto, 1990)

III. Modeling how the brain learns

Temporal difference (TD) learning

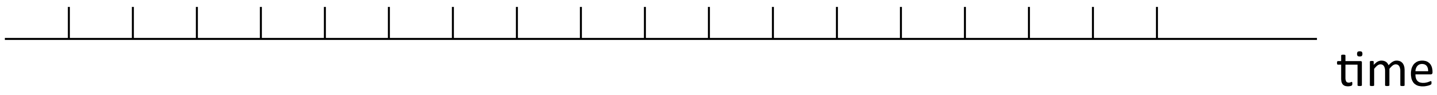

- Define a value for each moment in time separately, $V(t_i)$

$$v(t_i) = E\left[r(t_i) + \gamma r(t_i + 1) + \gamma^2 r(t_i + 2) + \ldots\right]$$

$0 \leq \gamma \leq 1$

discount parameter

The value at t_i is the sum of expected future rewards

III. Modeling how the brain learns

Temporal difference (TD) learning

$$V(t) = E\left[r(t) + \gamma r(t+1) + \gamma^2 r(t+2) + \ldots\right]$$

can be expressed as

$$V(t) = E\left[r(t) + \gamma V(t+1)\right]$$

Hence

$$E\left[r(t)\right] = \hat{V}(t) - \gamma \hat{V}(t+1)$$

III. Modeling how the brain learns

Temporal difference (TD) learning

Updating occurs by comparing the difference between

$$r(t) \qquad\qquad E\left[r(t)\right]$$

What actually occurs — What is expected to occur

$$V_{new}(t) = V_{new}(t) + \alpha \underline{\left[r(t) - E[r(t)]\right]}$$

Prediction error δ

III. Modeling how the brain learns

Temporal difference (TD) learning

Updating equation

$$V_{new}(t) = V_{new}(t) + \alpha\left[r(t) - E[r(t)]\right]$$

Given

$$E\left[r(t)\right] = \hat{V}(t) - \gamma\hat{V}(t+1)$$

We get

$$V_{new}(t) = V_{old}(t) + \alpha\left[r(t) + \gamma V_{old}(t+1) - V_{old}(t)\right]$$

(α ↓ learning rate)

$$0 \leq \alpha \leq 1$$

III. Modeling how the brain learns

Temporal difference (TD) learning

Updating equation

$$V_{new}(t) = V_{new}(t) + \alpha\left[r(t) - E[r(t)]\right]$$

$$V_{new}(t) = V_{old}(t) + \alpha\left[r(t) + \gamma V_{old}(t+1) - V_{old}(t)\right]$$

Prediction error δ

Temporal difference (TD) learning

<u>Trial 1:</u>

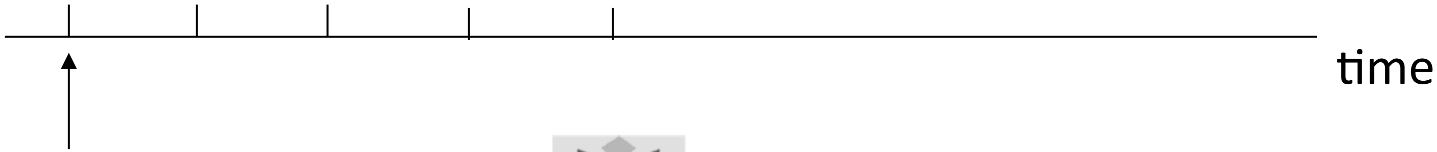

Suppose at t_1 you see 

$$V_{trial1}(t_1) = V_{trial0}(t_1) + \alpha\left[r_{trial1}(t_1) + \gamma V_{trial0}(t_2) - V_{trial0}(t_1)\right] = 0$$

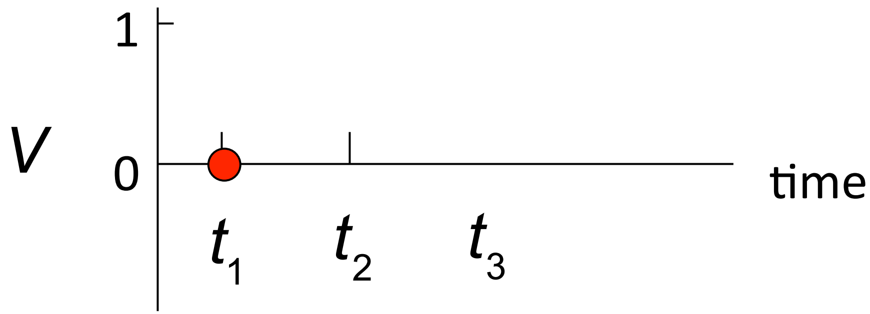

assume
$\alpha = 0.5, \gamma = 0.99$

III. Modeling how the brain learns

Temporal difference (TD) learning

<u>Trial 1:</u>

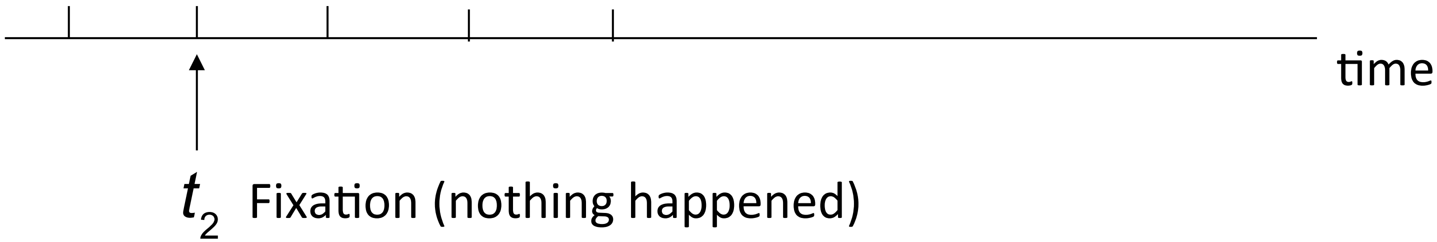

For t_2

$$V_{trial1}(t_2) = V_{trial0}(t_2) + \alpha \left[r_{trial1}(t_2) + \gamma V_{trial0}(t_3) - V_{trial0}(t_2) \right] = 0$$

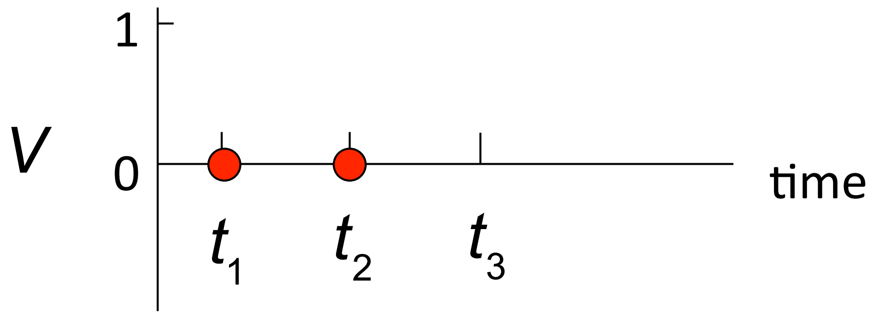

III. Modeling how the brain learns

Temporal difference (TD) learning

Trial 1:

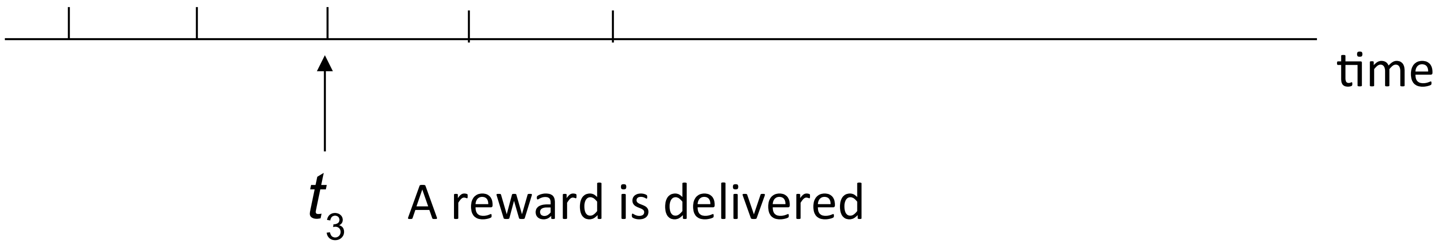

For t_3

$$V_{trial1}(t_3) = V_{trial0}(t_3) + \alpha\left[r_{trial1}(t_3) + \gamma V_{trial0}(t_4) - V_{trial0}(t_3)\right] = \alpha$$

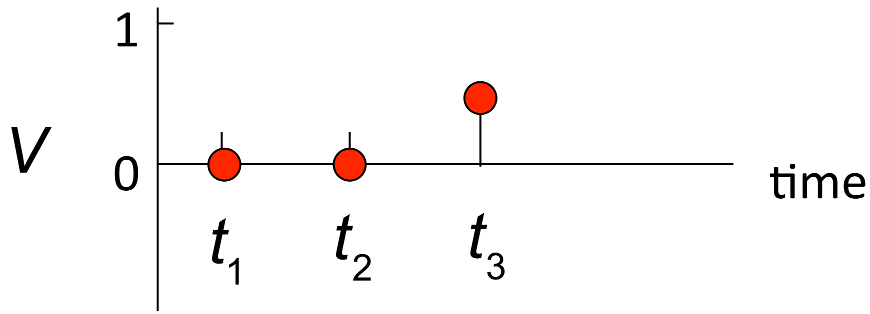

Temporal difference (TD) learning

<u>Trial 1:</u>

For t_4

$$V_{trial1}(t_4) = V_{trial0}(t_4) + \alpha\left[r_{trial1}(t_4) + \gamma V_{trial0}(t_5) - V_{trial0}(t_4)\right] = 0$$

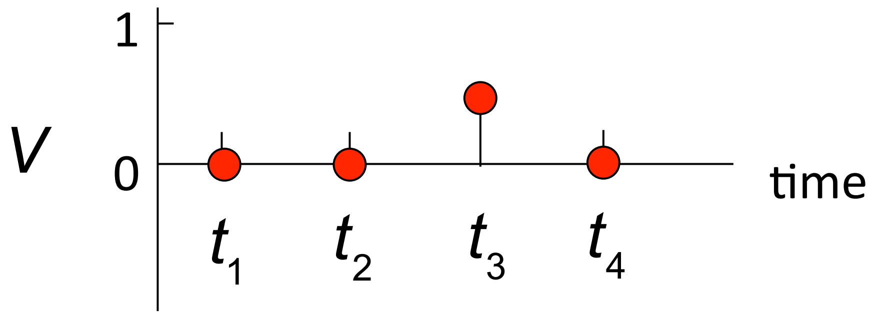

III. Modeling how the brain learns

Temporal difference (TD) learning

<u>Trial 2:</u>

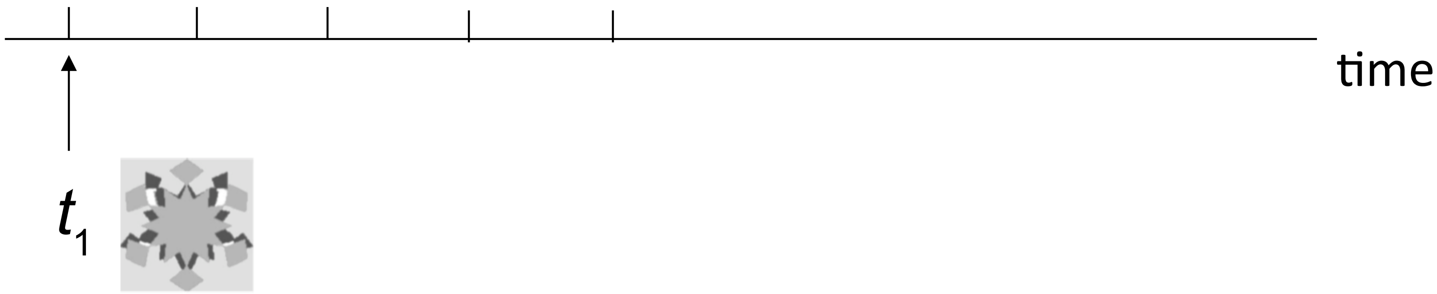

$$V_{trial2}(t_1) = V_{trial1}(t_1) + \alpha \left[r_{trial2}(t_1) + \gamma V_{trial1}(t_2) - V_{trial1}(t_1) \right] = 0$$

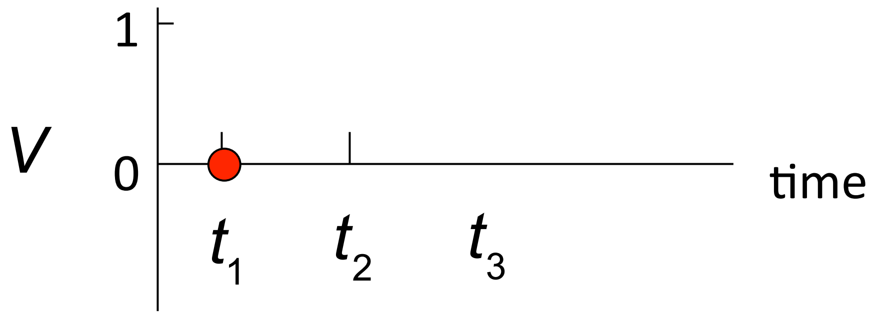

III. Modeling how the brain learns

Temporal difference (TD) learning

<u>Trial 2:</u>

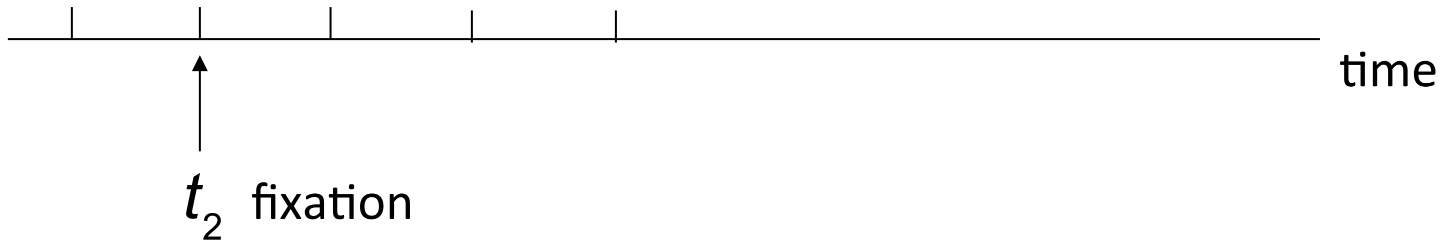

$$V_{trial2}(t_2) = V_{trial1}(t_2) + \alpha \left[r_{trial2}(t_2) + \gamma V_{trial1}(t_3) - V_{trial1}(t_2) \right] = \gamma\alpha^2$$

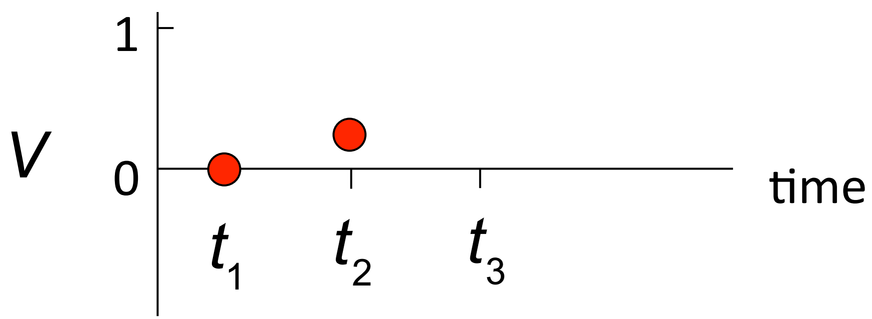

Temporal difference (TD) learning

Trial 2:

$$V_{trial2}(t_3) = V_{trial1}(t_3) + \alpha\left[r_{trial2}(t_3) + \gamma V_{trial1}(t_4) - V_{trial1}(t_3)\right] = 2\alpha - \alpha^2$$

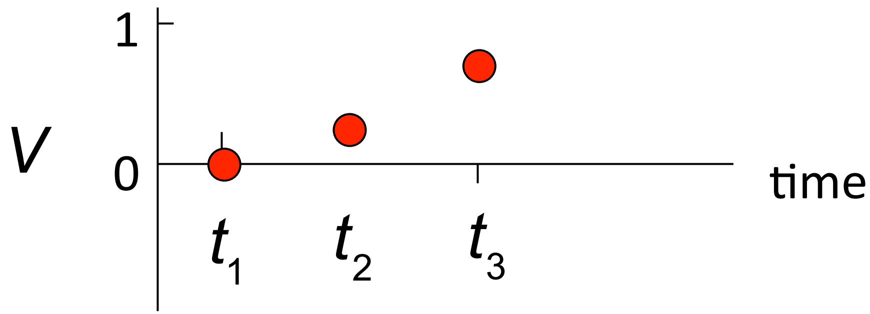

Temporal difference (TD) learning

Trial 2:

$$V_{trial2}(t_4) = V_{trial1}(t_4) + \alpha\left[r_{trial2}(t_4) + \gamma V_{trial1}(t_5) - V_{trial1}(t_4)\right] = 0$$

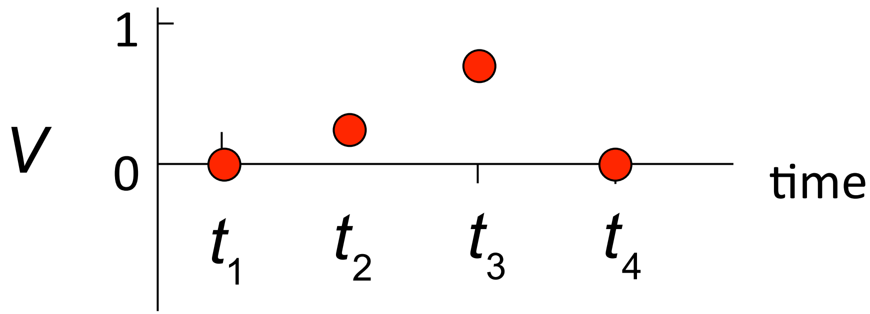

III. Modeling how the brain learns

Temporal difference (TD) learning

<u>Trial 3:</u>

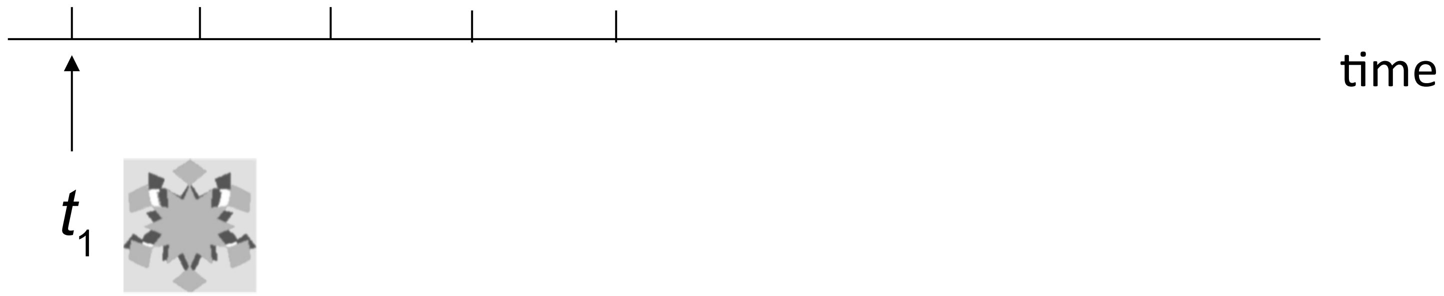

$$V_{trial3}(t_1) = V_{trial2}(t_1) + \alpha\left[r_{trial3}(t_1) + \gamma V_{trial2}(t_2) - V_{trial2}(t_1)\right] = \gamma^2\alpha^3$$

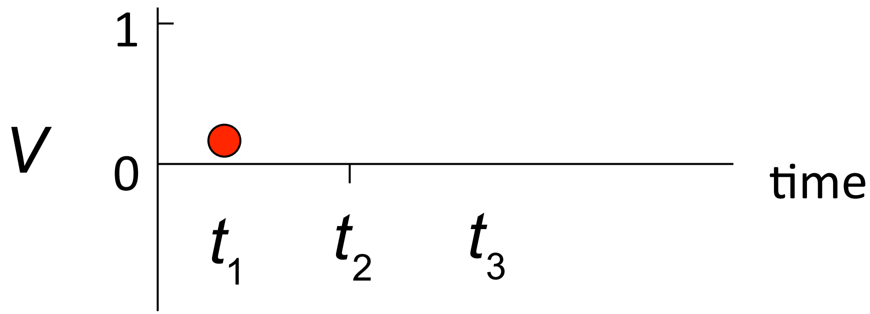

III. Modeling how the brain learns

Temporal difference (TD) learning

<u>Trial 3:</u>

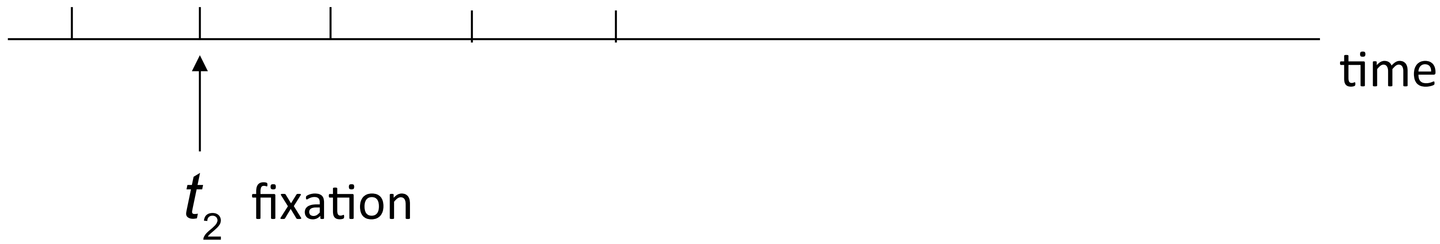

$$V_{trial3}(t_2) = V_{trial2}(t_2) + \alpha\left[r_{trial3}(t_2) + \gamma V_{trial2}(t_3) - V_{trial2}(t_2)\right]$$

$$= \delta(3\alpha^2 - 2\alpha^3)$$

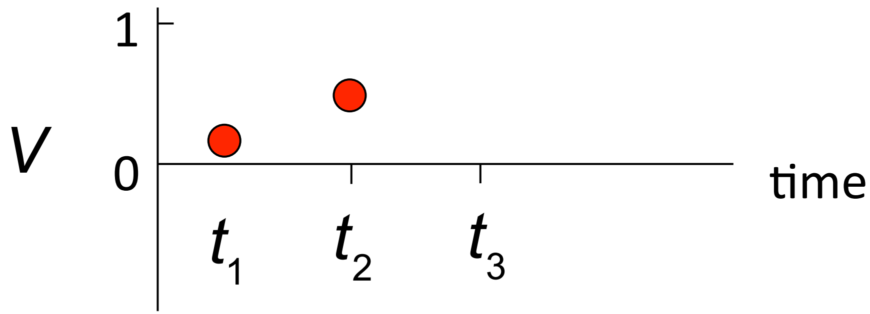

III. Modeling how the brain learns

Temporal difference (TD) learning

<u>Trial 3:</u>

$$V_{trial3}(t_3) = V_{trial2}(t_3) + \alpha\left[r_{trial3}(t_3) + \gamma V_{trial2}(t_4) - V_{trial2}(t_3)\right]$$
$$= 3\alpha - 3\alpha^2 + \alpha^3$$

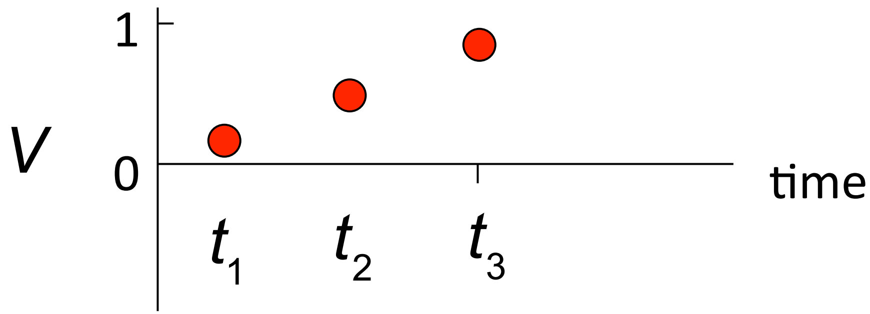

III. Modeling how the brain learns

Temporal difference (TD) learning

<u>Trial 3:</u>

$$V_{trial3}(t_4) = V_{trial2}(t_4) + \alpha\left[r_{trial3}(t_4) + \gamma V_{trial2}(t_5) - V_{trial2}(t_4)\right] = 0$$

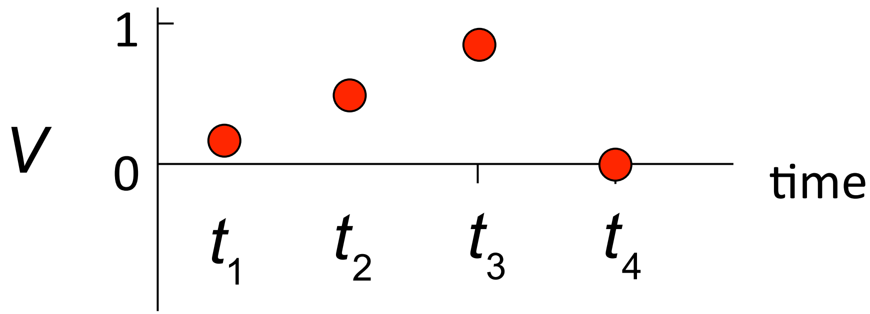

Temporal difference (TD) learning

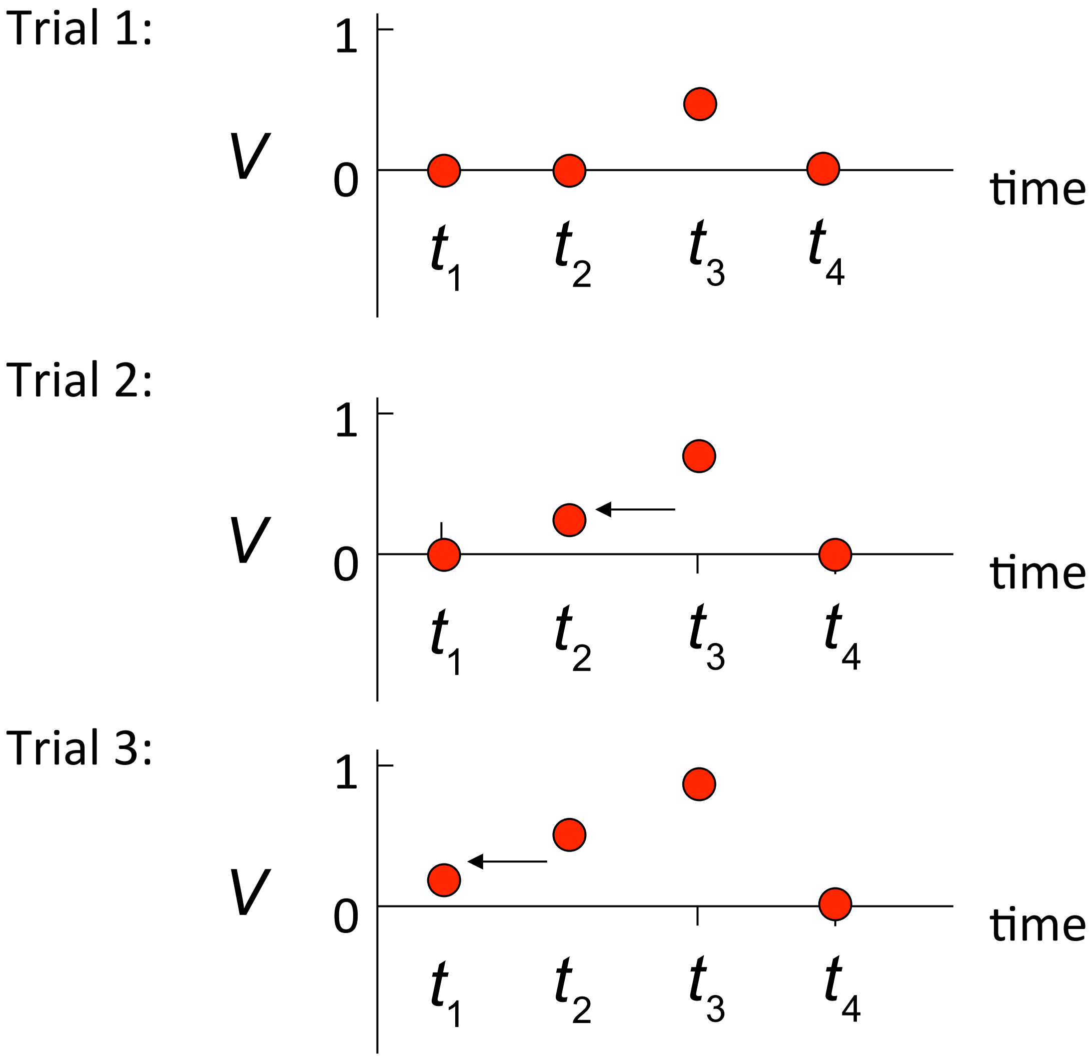

Temporal difference (TD) learning

Trial 4:

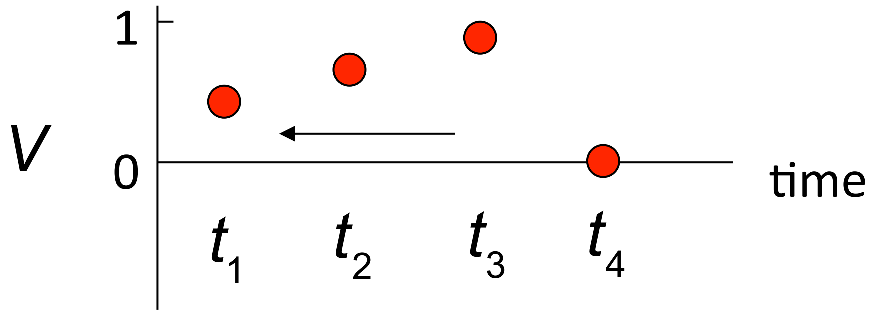

Trial 5:

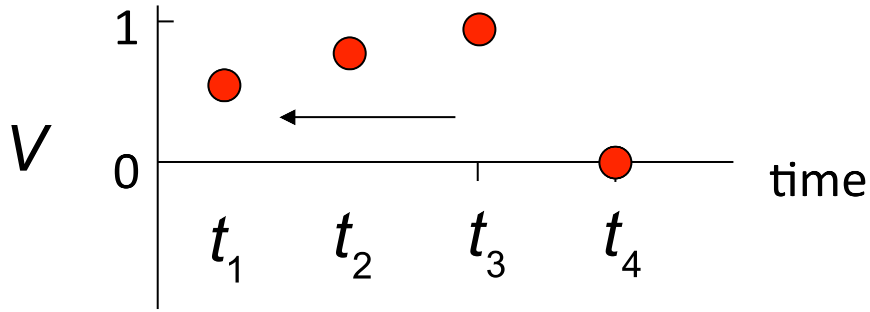

Trial 6:

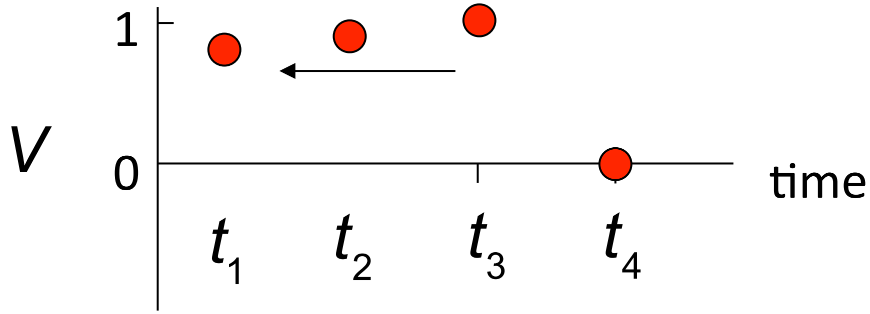

III. Modeling how the brain learns

Temporal difference (TD) learning

Let's look at prediction error δ

Recall that

$$V_{trial(X+1)}(t) = V_{trialX}(t) + \alpha \underbrace{\left[r_{trial(X+1)}(t) + \gamma V_{trialX}(t+1) - V_{trialX}(t) \right]}_{\delta}$$

III. Modeling how the brain learns

Just looking at t_3

Trial 1:

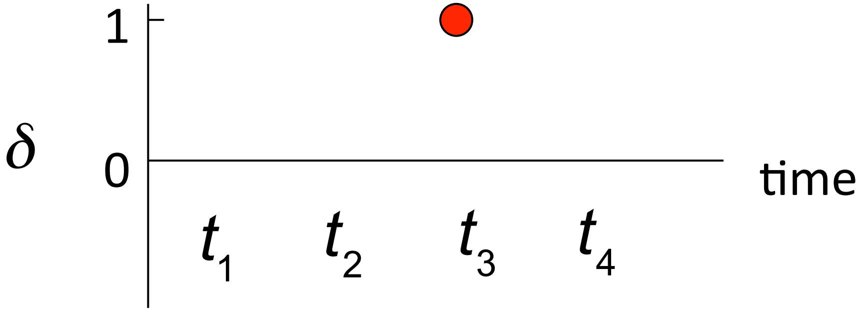

Trial 2:

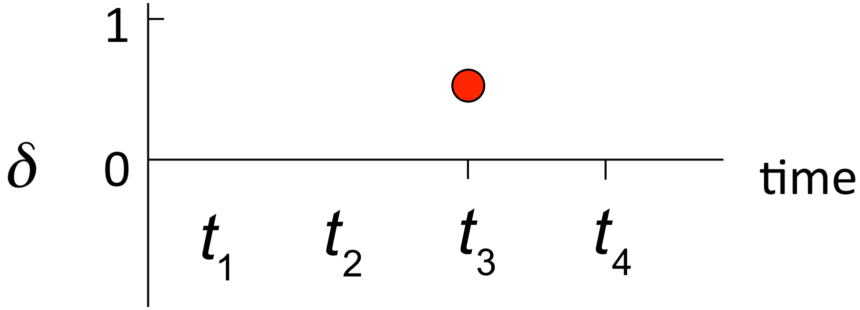

Trial 3:

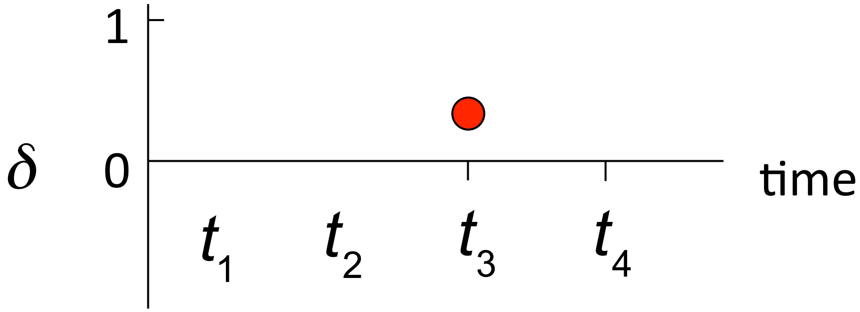

III. Modeling how the brain learns

Just looking at t_3

Trial 4:

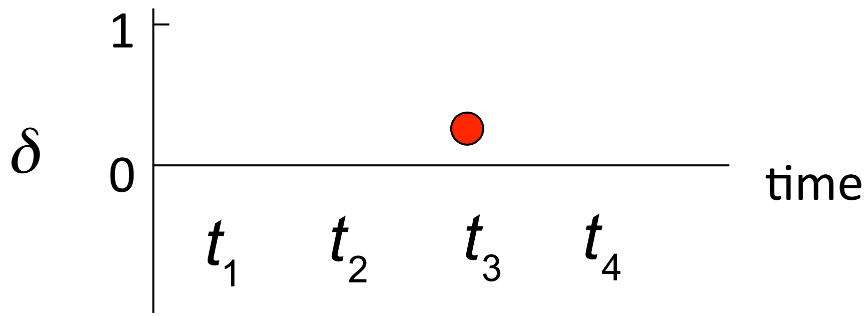

Trial 5:

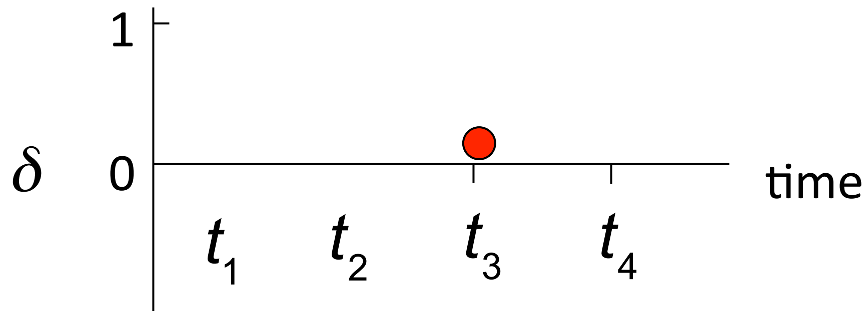

Trial 6:

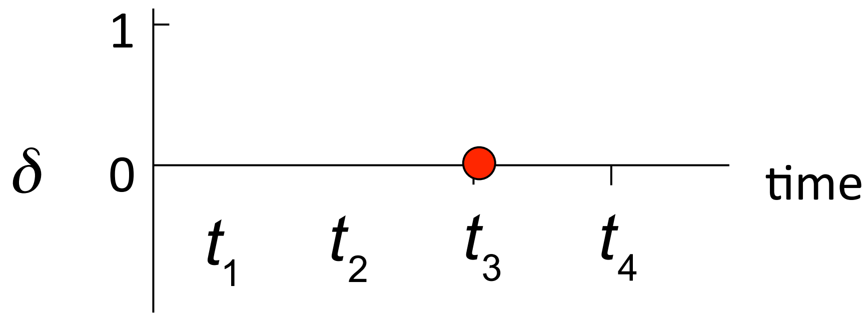

III. Modeling how the brain learns

Temporal difference (TD) learning

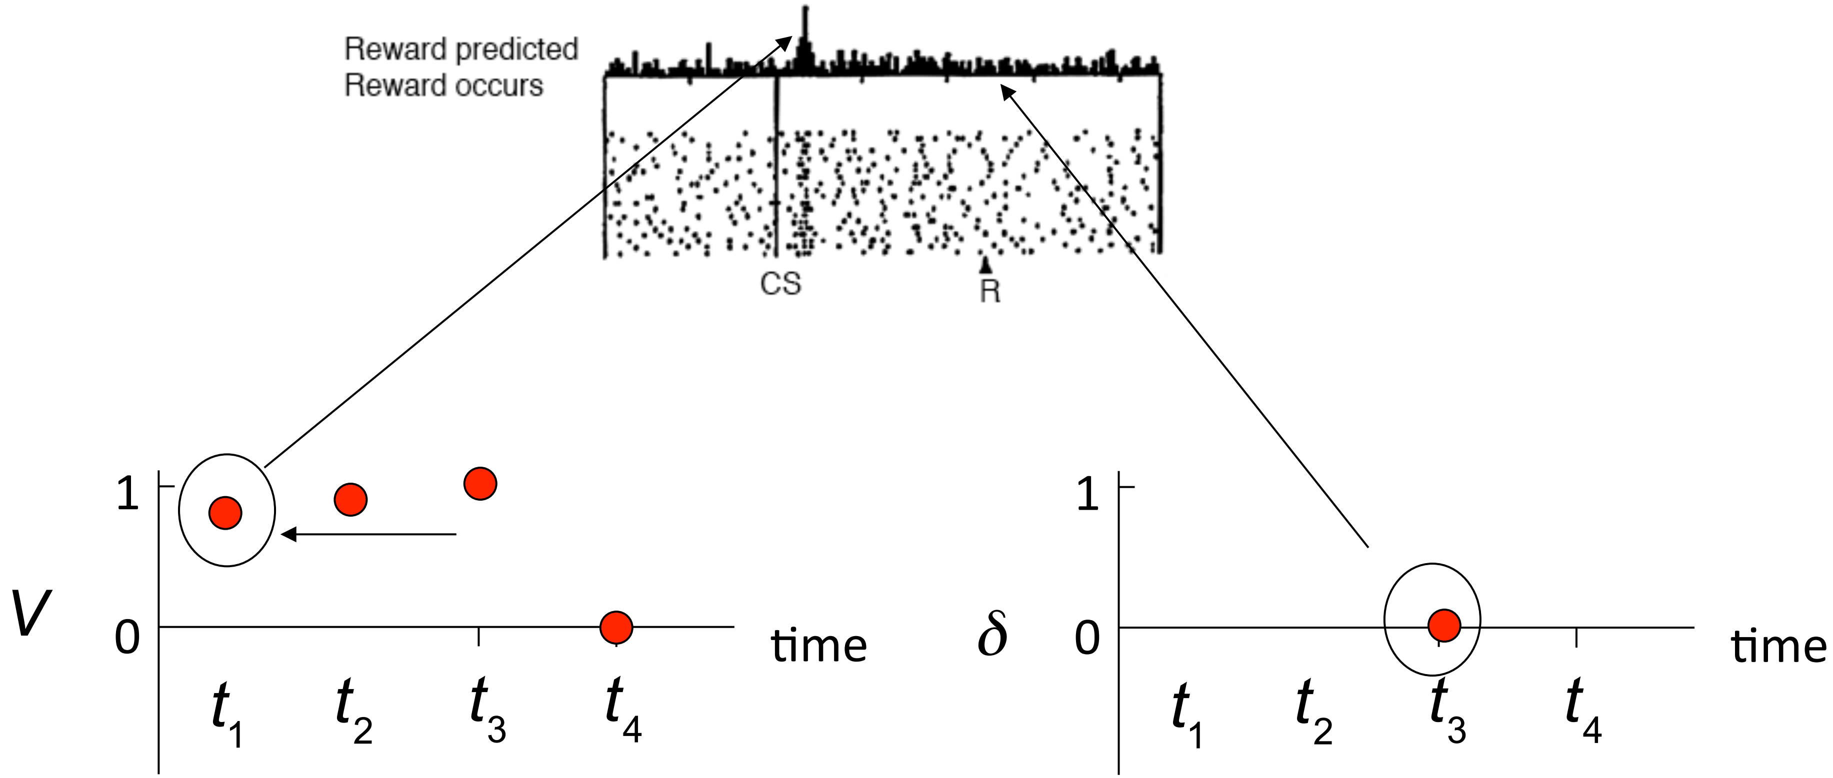

III. Modeling how the brain learns

Based on TD model, we can

- Construct a General Linear Model (GLM) to analyze data

$$Y(t) = \beta_0 + \beta_1 x_I(t) + \beta_V V(t) + \beta_\delta \delta(t)$$

Indicator variable

Value at t

Prediction error

TD model provides a quantitative prediction on the time course of data

Modeling response dynamics: Drift diffusion model

IV. Modeling response dynamics

Action selection is a dynamic process

- Multiple alternatives compete during this process

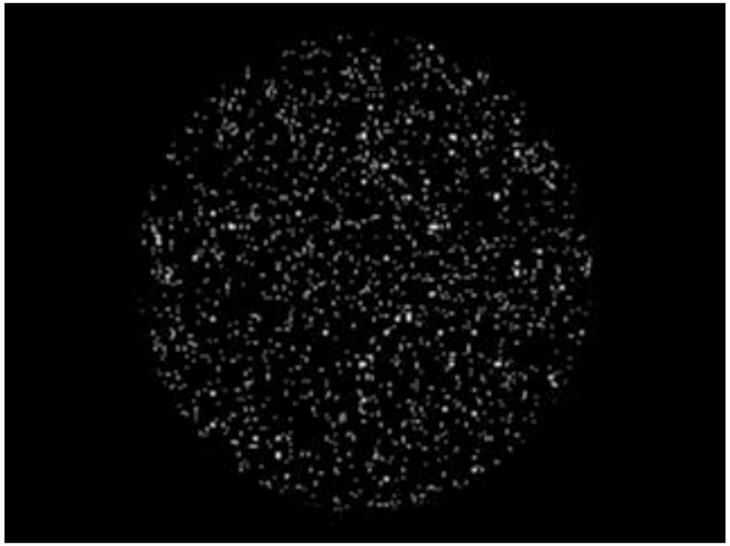

30% motion coherence

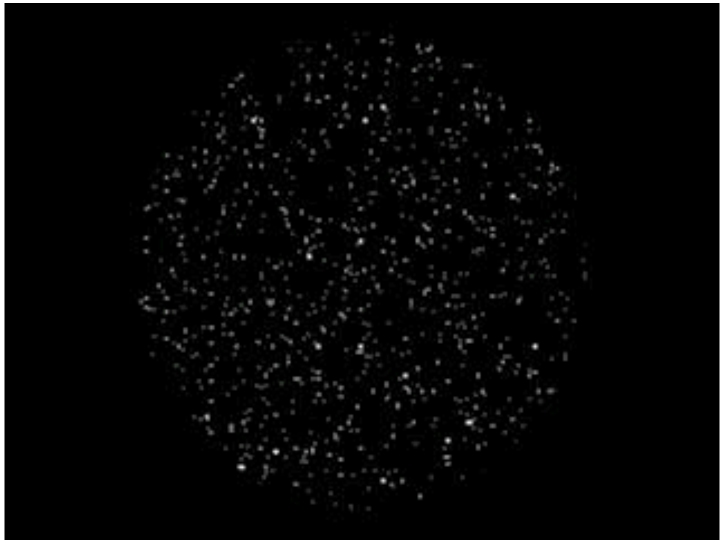

5% motion coherence

Question: how do we model the dynamics of neural activity during this process?

Courtesy to Bill Newsome

Dynamics of neural activity during stimulus presentation

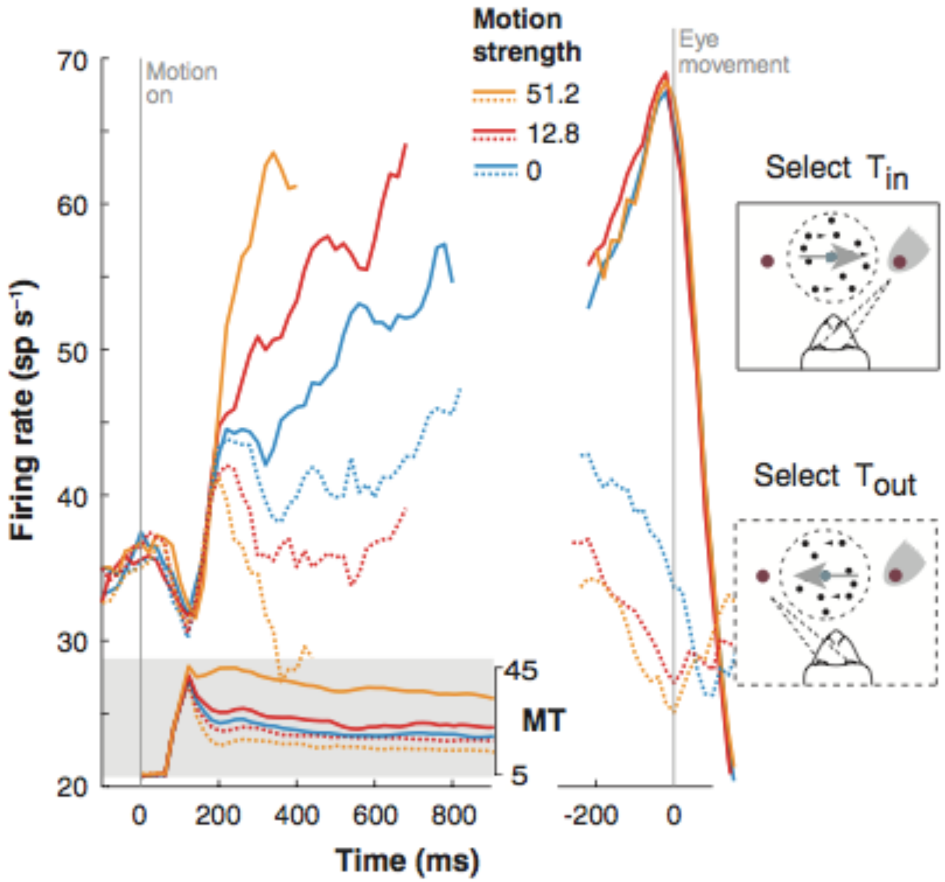

- Activity in area LIP rises up faster as motion coherence level increases

- Prior to eye movement, activity does not differ between different coherence trials

Golad & Shadlen (2007, *ARN*)

Modeling response dynamics as an evidence accumulation process

Firing rates behave as if neurons integrate *momentary evidence* over time

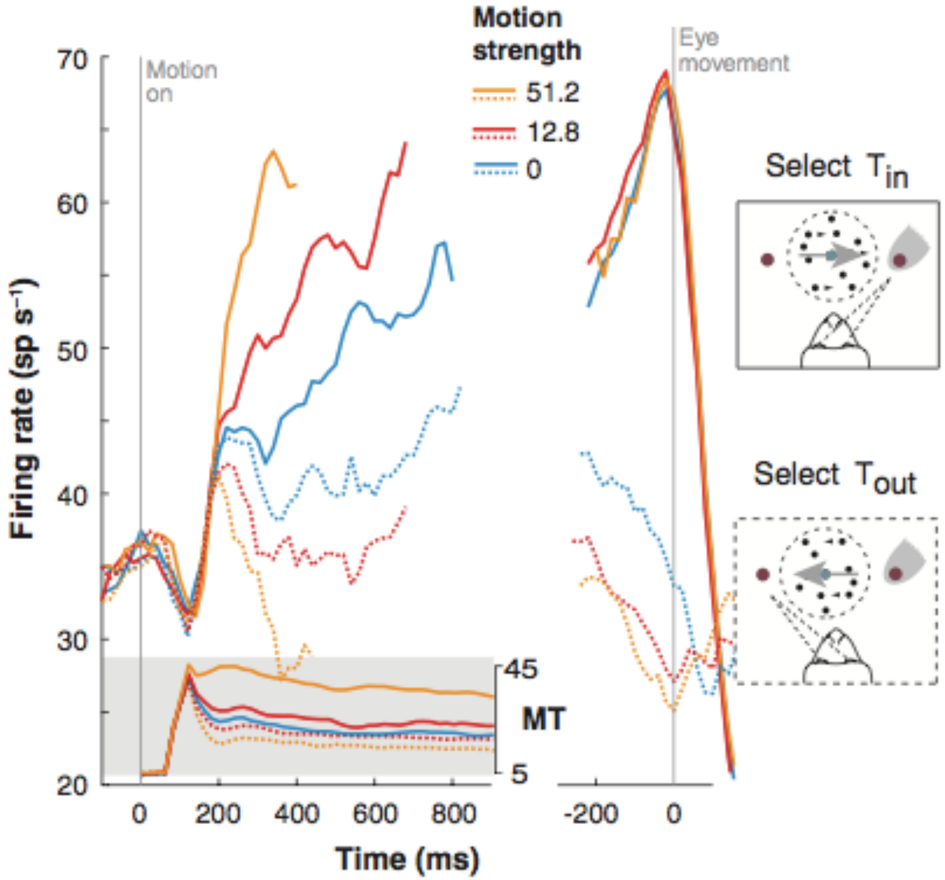

Each moment in time:

$$\log LR(t_i) = \log \frac{p(e(\theta,t_i)\,|\,L)}{p(e(\theta,t_i)\,|\,R)}$$

θ = motion coherence

Over time:

$$\log LR(t_1,\ldots t_k) = \sum_i \log LR(t_i)$$

Golad & Shadlen (2007, *ARN*)

Use Drift diffusion model to characterize evidence accumulation and action selection

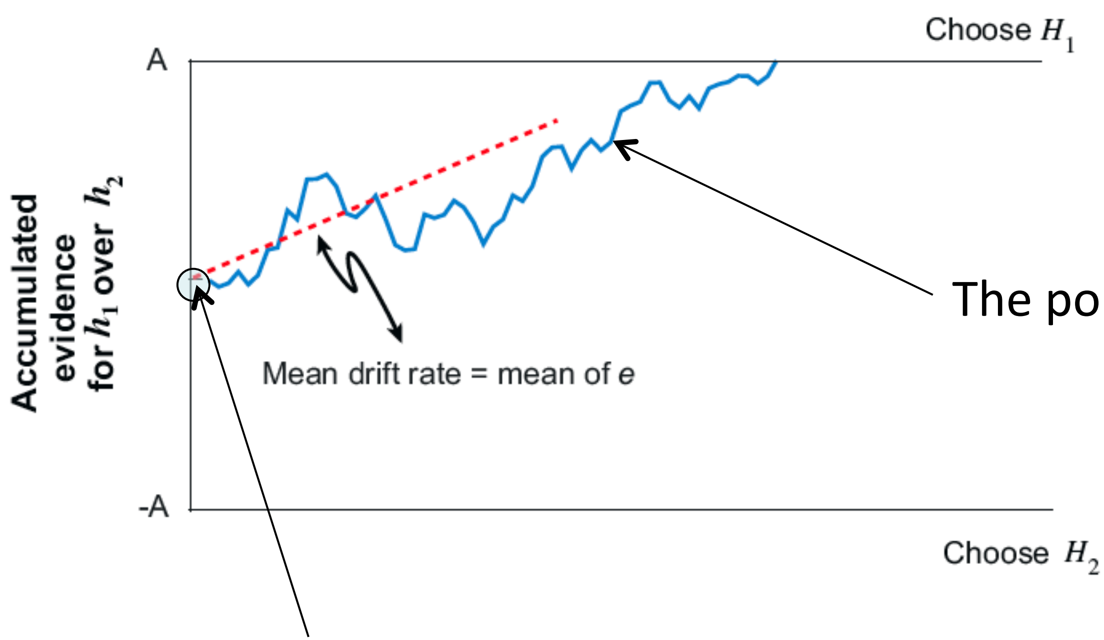

Decision bound
(e.g. leftward motion)

The point process ‘drifts’ as time goes by

Decision bound
(e.g. rightward motion)

Starting point
(unbiased)

Gold & Shadlen (2007, *ARN*)

IV. Modeling response dynamics

Use Drift diffusion model to characterize evidence accumulation and action selection

What determines the drift?
Ans: Momentary evidence is sampled from a Gaussian distribution to determine the next step

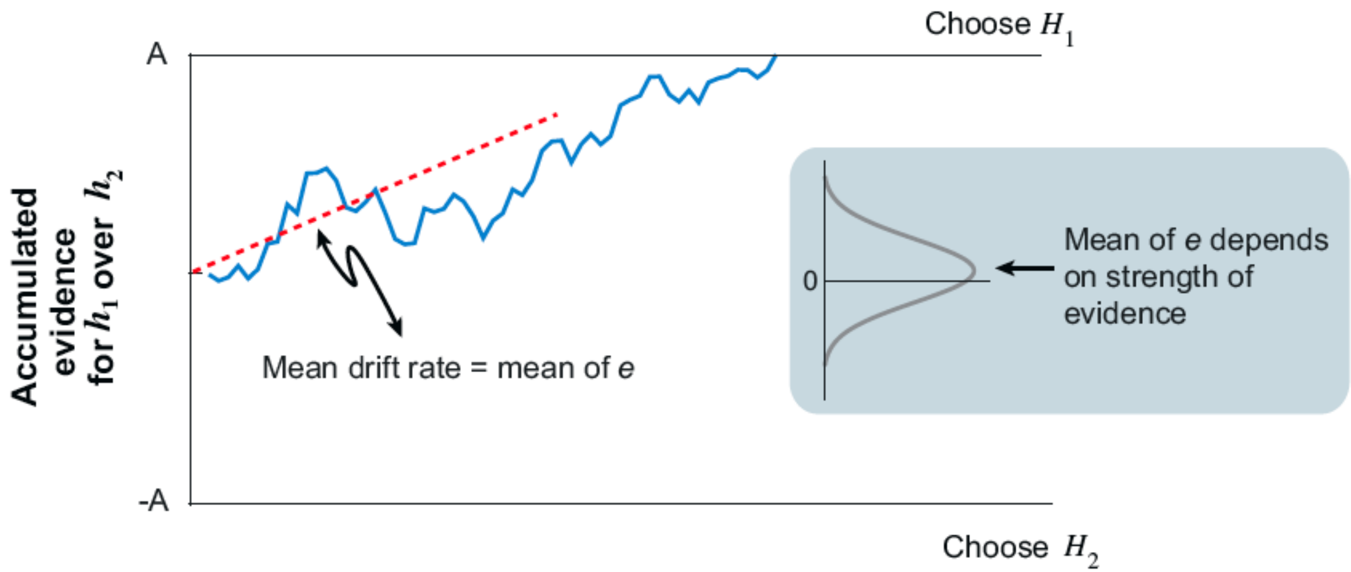

Gold & Shadlen (2007, *ARN*)