

1

## Combining Error-driven + Hebbian

Get benefits of both: Solve tasks, learn systematic representations, generalize to new stimuli.

What's left?...

Time!

Currently: networks learn *immediate* consequence of a given input.

- What if current input only makes sense as part of a *sequence* of inputs (e.g., language, social interactions)?
- What if the consequence of this input comes *later* (e.g., school/work, life)?

2

## Sequence Learning

How do we do it?

For example:

My favorite color is purple.

My favorite food is purple.

My favorite purple color is.

The girl picked up the pen.

The pig raced around the pen.

We represent the *context*, not just the current input.


in language, social interactions, driving (who should go at a 4-way stop?), etc.

3

## Representing Context for Sequence Learning

How does the brain do it?

How would we get our models to do it?




4

## Representing Context for Sequence Learning

How does the brain do it?


How would we get our models to do it?



Add layers to keep track of context (prefrontal cortex).

5

## Representing Context for Sequence Learning



Simple Recurrent Network (SRN; Elman, Jordan).

$$c_j(t) \approx h_j(t-1) + c_j(t-1) \quad (1)$$

$$c_j(t) = f_{m_{hid}} h_j(t-1) + f_{m_{prv}} c_j(t-1) \quad (2)$$

6

## An Example Task

BTXSE  
BPVPSE  
BTSXXTVVE  
BPTVPSE


Which of the following sequences are allowed?:

BTXXTTVVE  
TSXSE  
VVSXE  
BSSXSE

7

## An Example Task

BTXSE, BPVPSE, BTSXXTVVE, BPTVPSE, BTXXTTVVE  
TSXSE, VVSXE, BSSXSE




We implicitly learn such grammars (e.g., pressing buttons faster to letters that follow grammar).

8

## An Example Task

BTXSE, BPVPSE, BTSXXTVVE, BPTVPSE, BTXXTTVVE  
TSXSE, VVSXE, BSSXSE



Sims.

9

## Combining Error-driven + Hebbian

Get benefits of both: Solve tasks, learn systematic representations, generalize to new stimuli.

What's left?...

Time!

Currently: networks learn *immediate* consequence of a given input.

- What if current input only makes sense as part of a *sequence* of inputs (e.g., language, social interactions)? We represent the *context*, not just the current input.
- What if the consequence of this input comes *later* (e.g., school/work, life)?