Perceptrons

The basic model of synaptic learning

What are perceptrons?

First, they are the simplest possible model of a circuit:

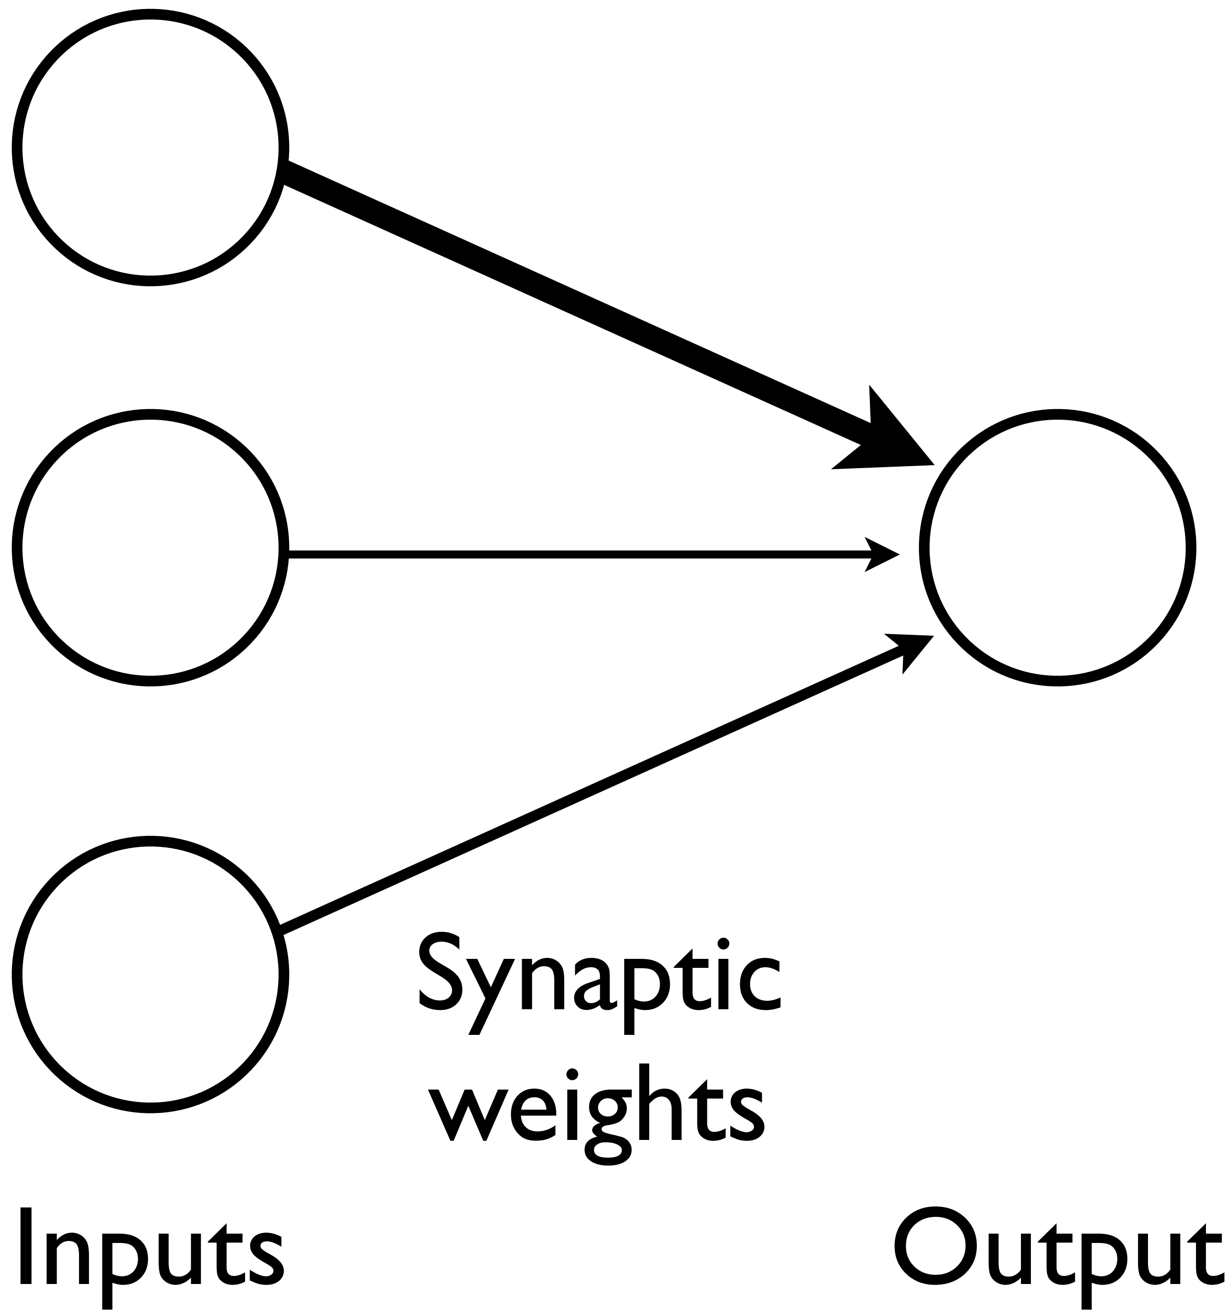

What are perceptrons?

What do they do? What are the dynamics?

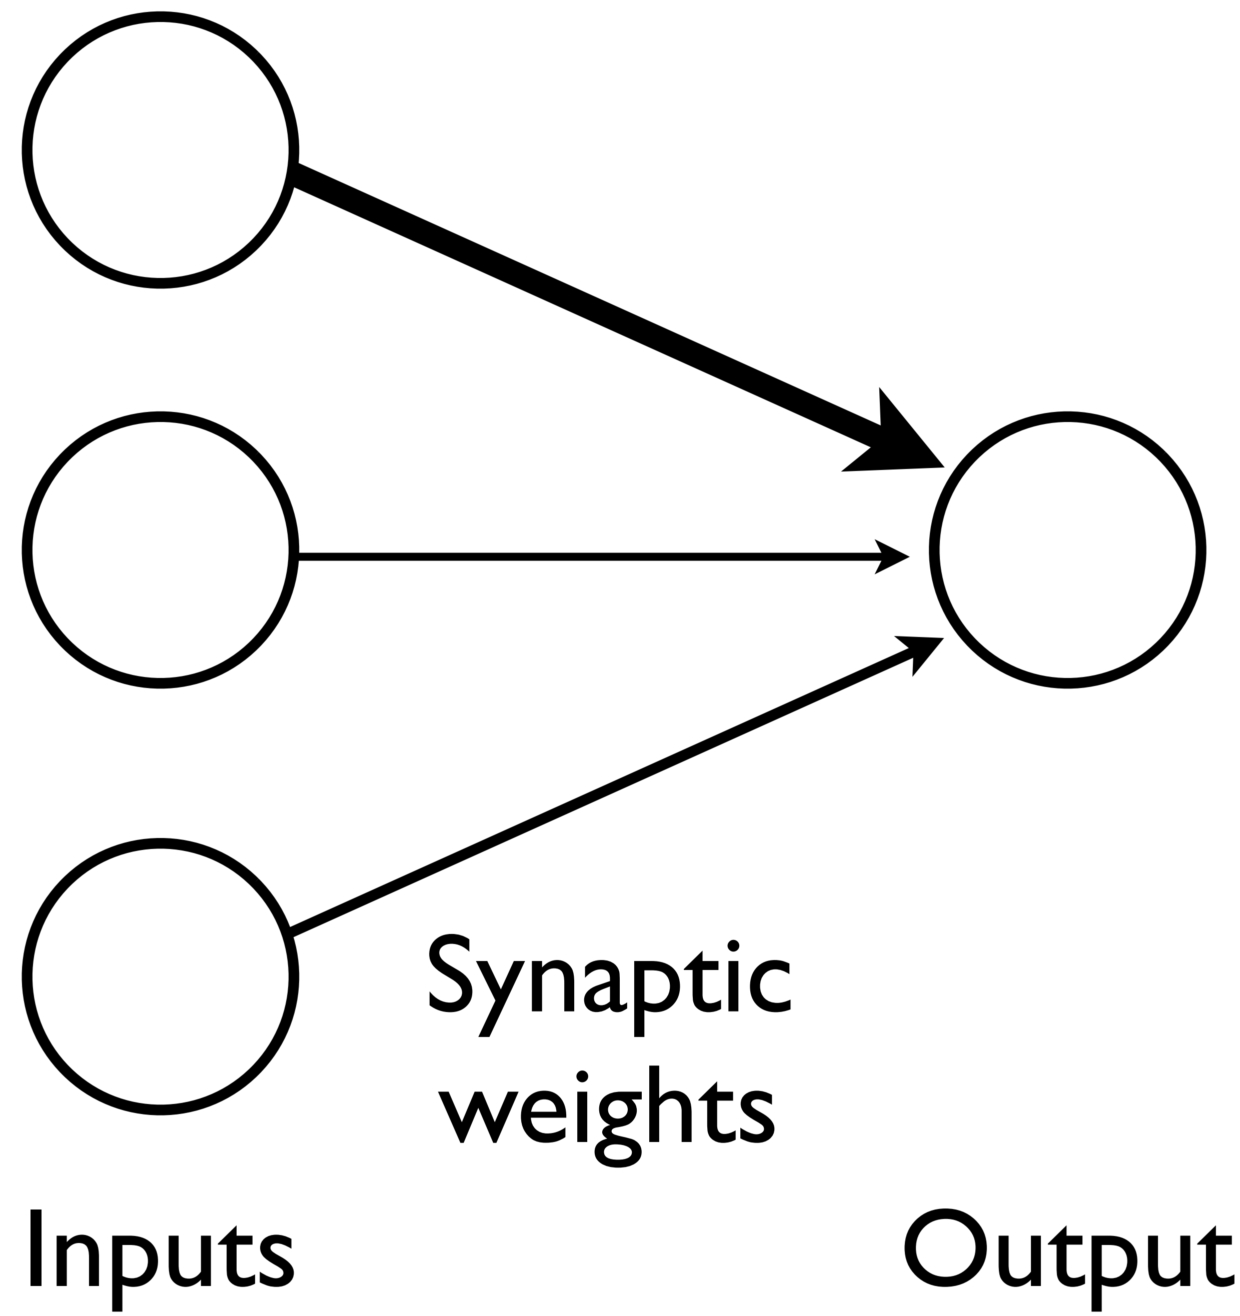

What are perceptrons?

Sums inputs according to synaptic weight and output is active if input is greater than threshold

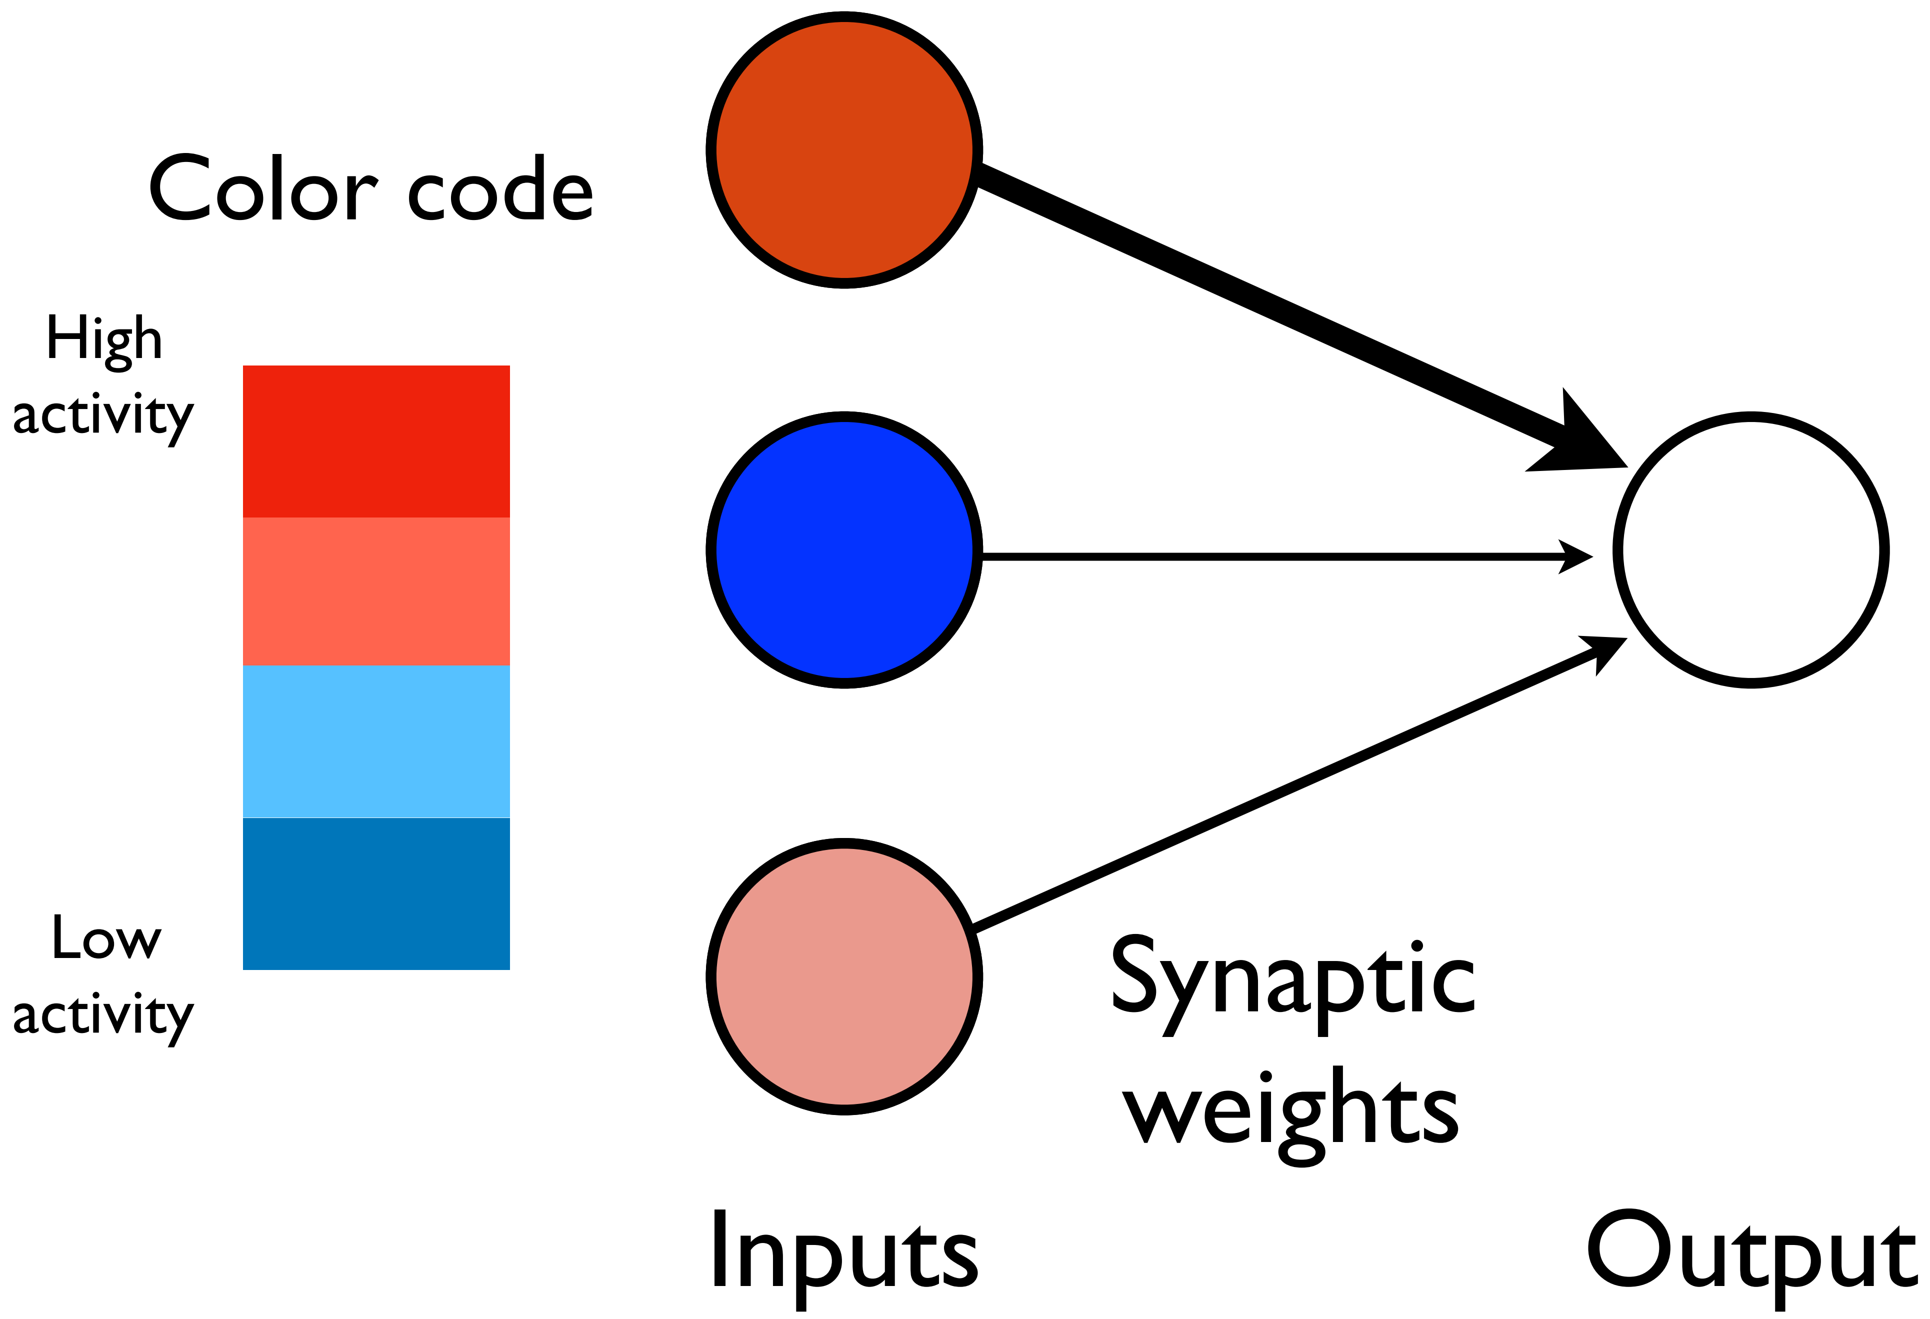

What are perceptrons?

Sums inputs according to synaptic weight and output is active if input is greater than threshold

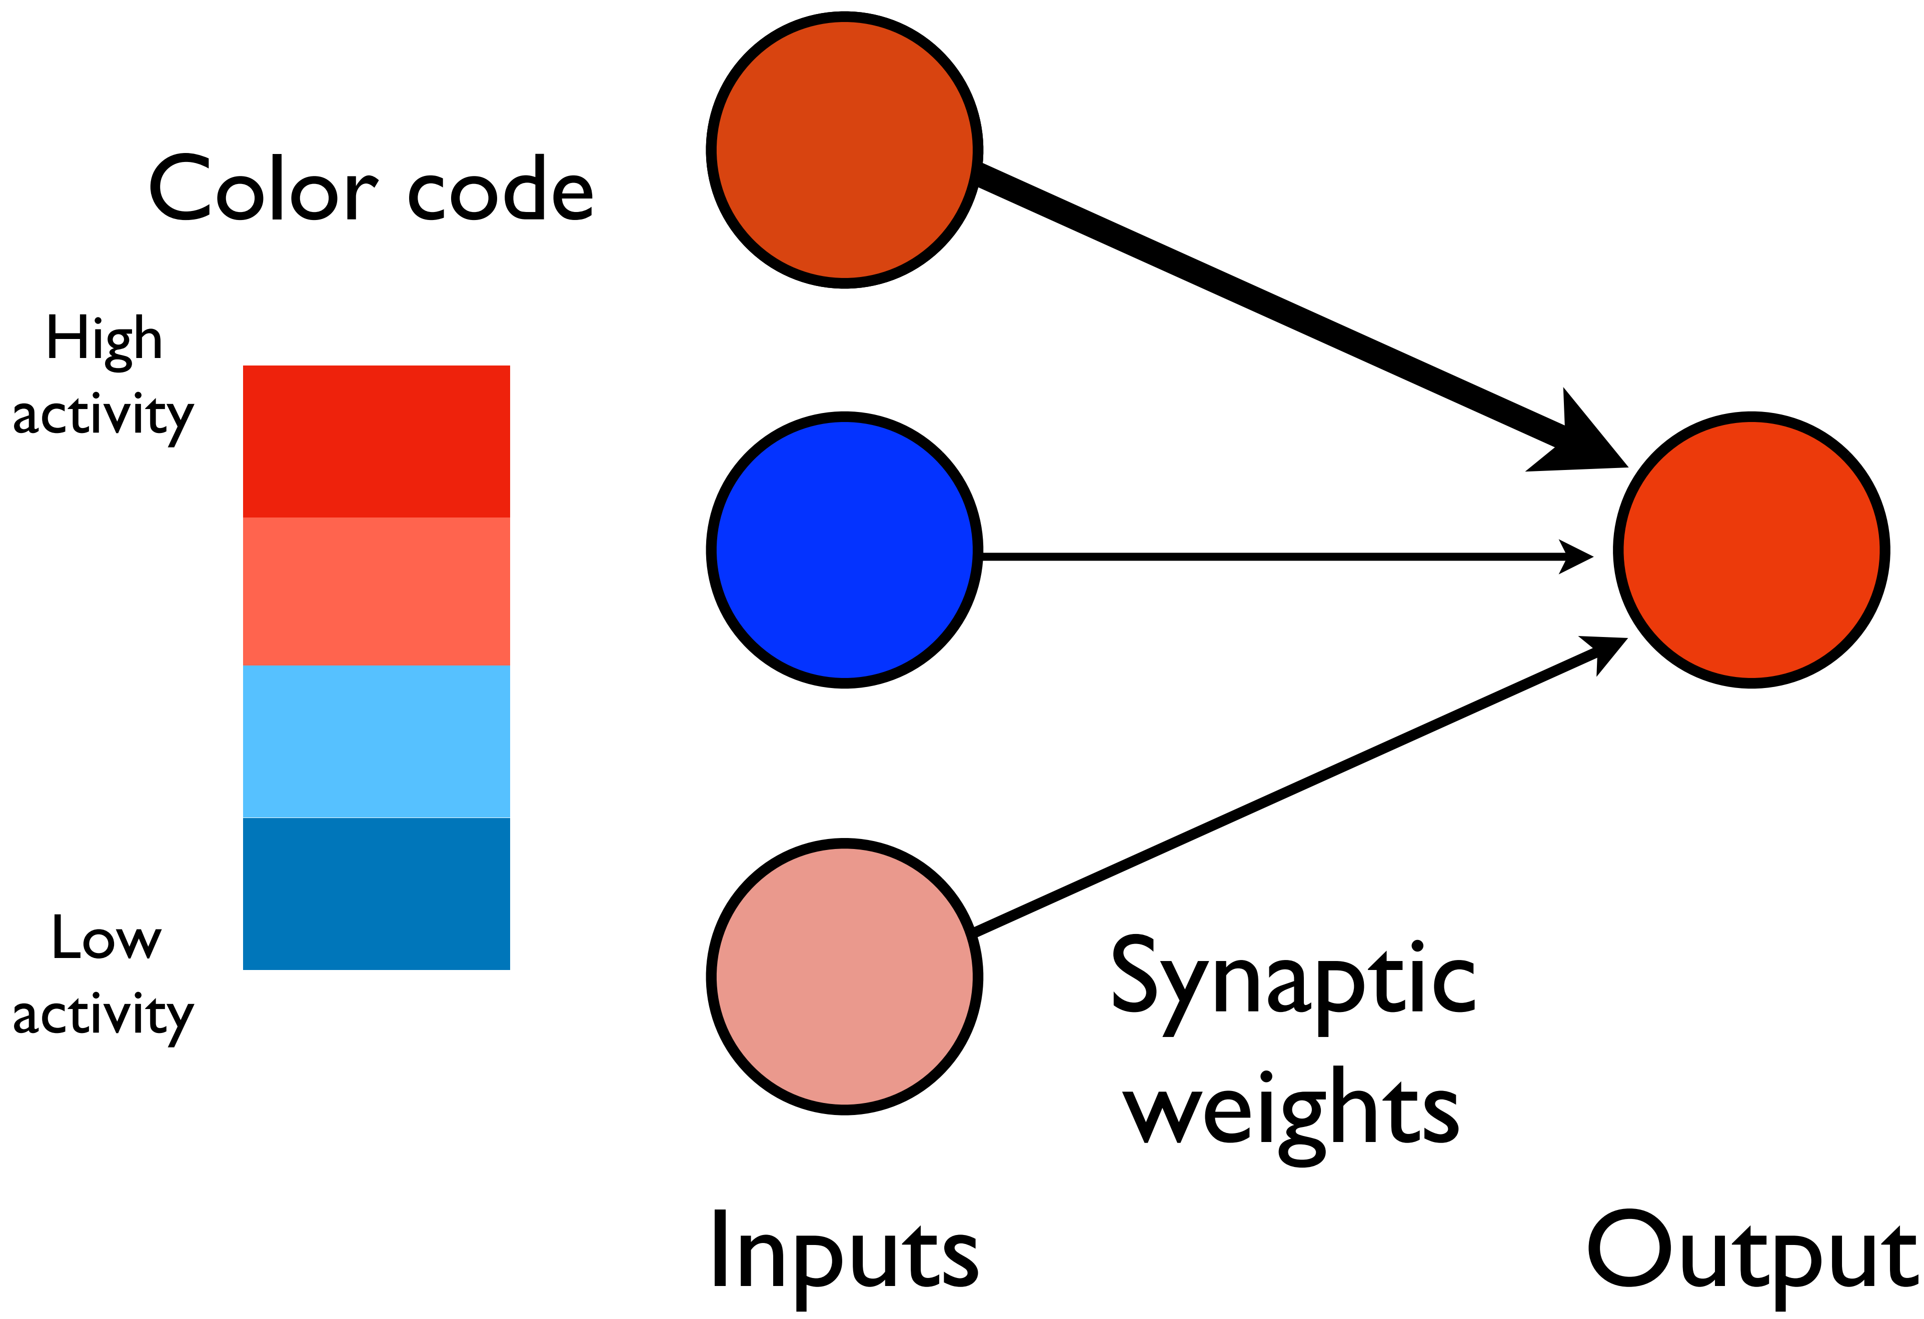

What are perceptrons?

Sums inputs according to synaptic weight and output is active if input is greater than threshold

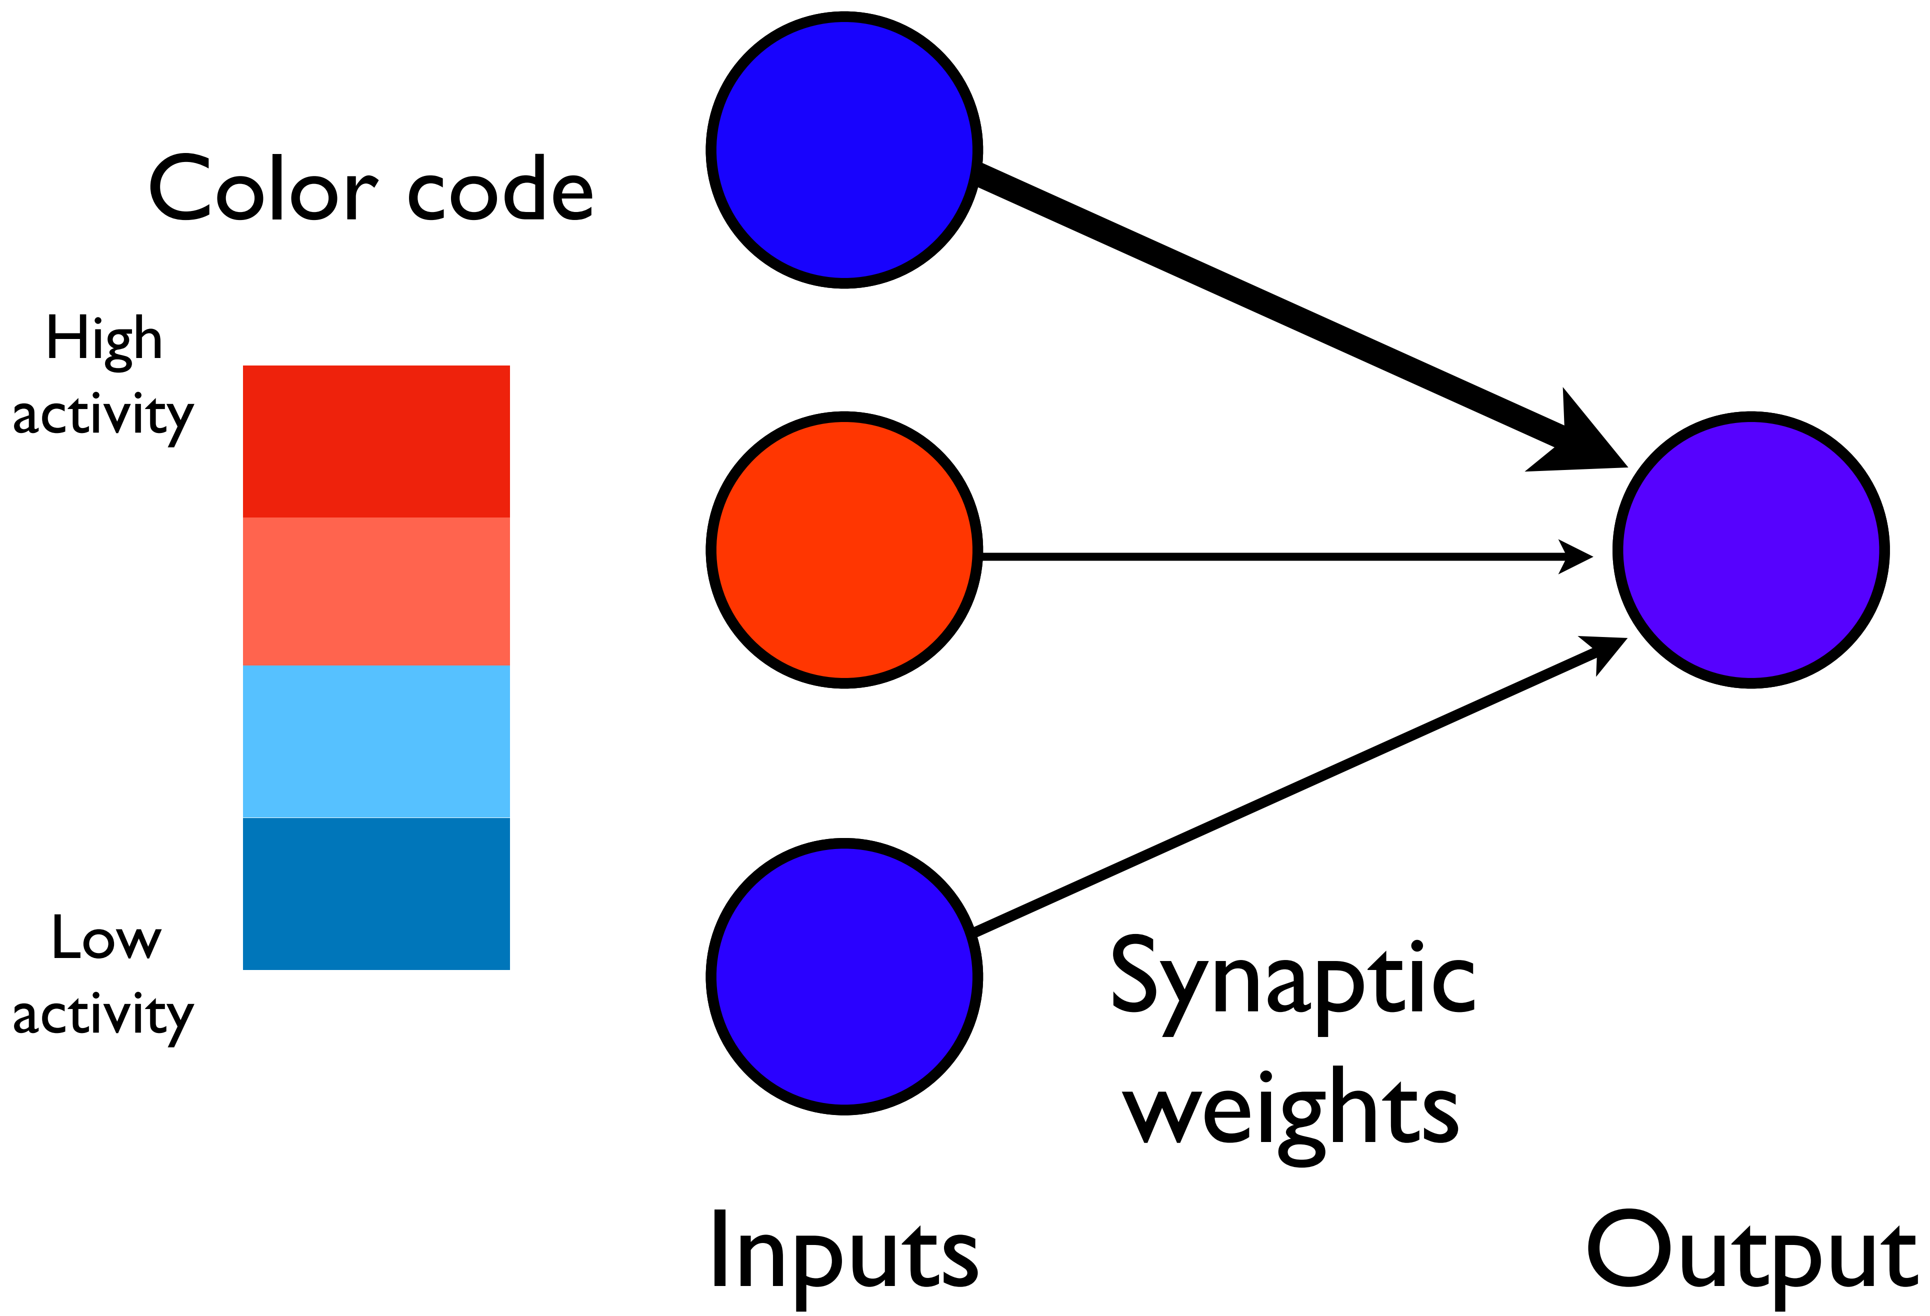

The interesting part is learning in the synaptic weights

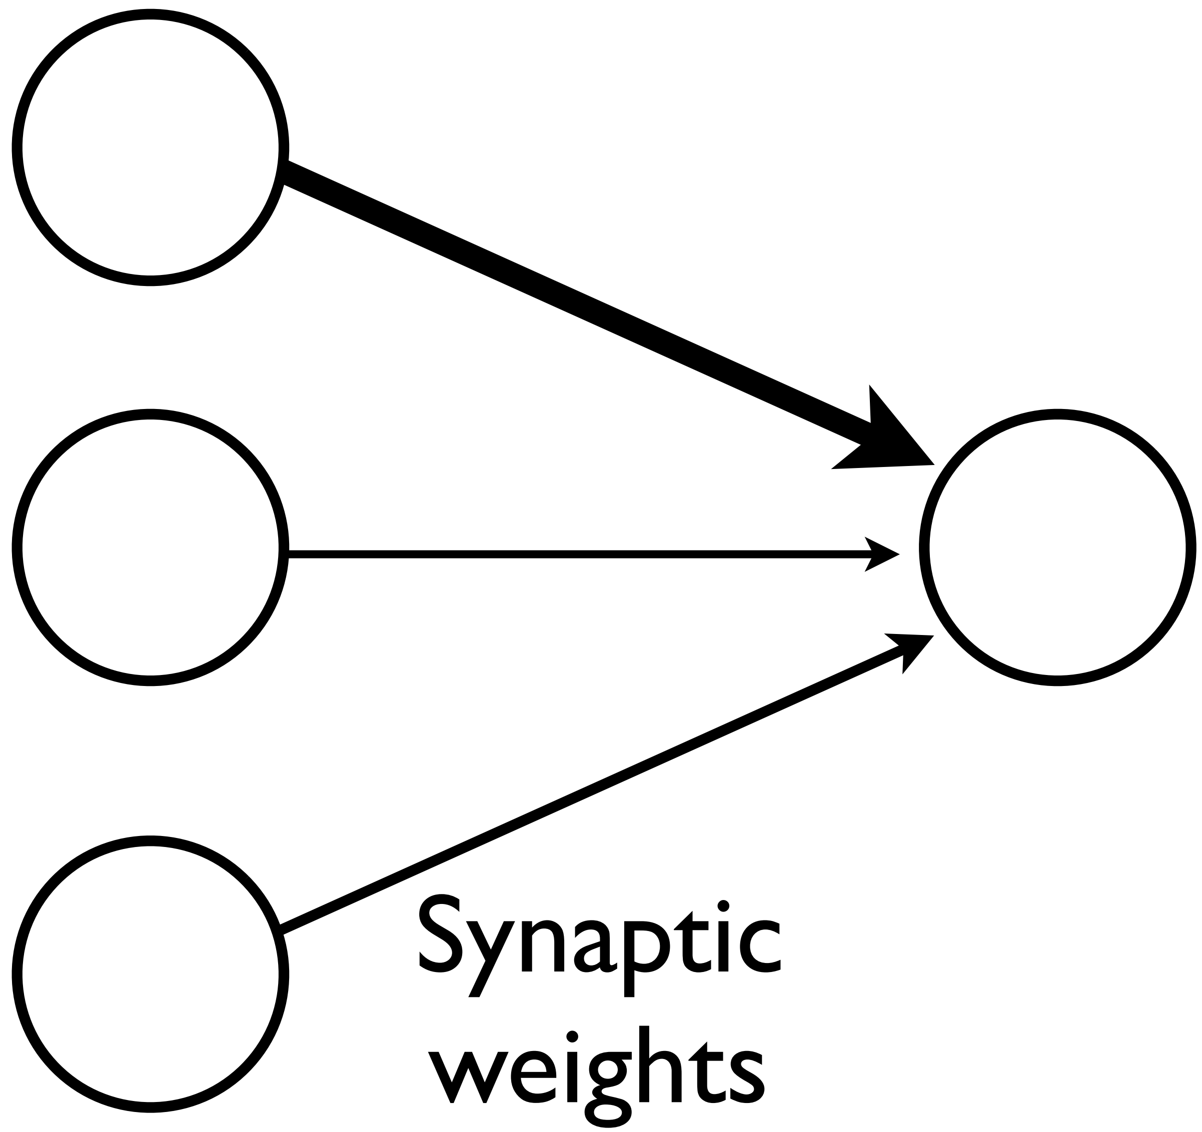

Similar to learning in Drosophila olfactory system

White board

Similar to learning in Drosophila olfactory system

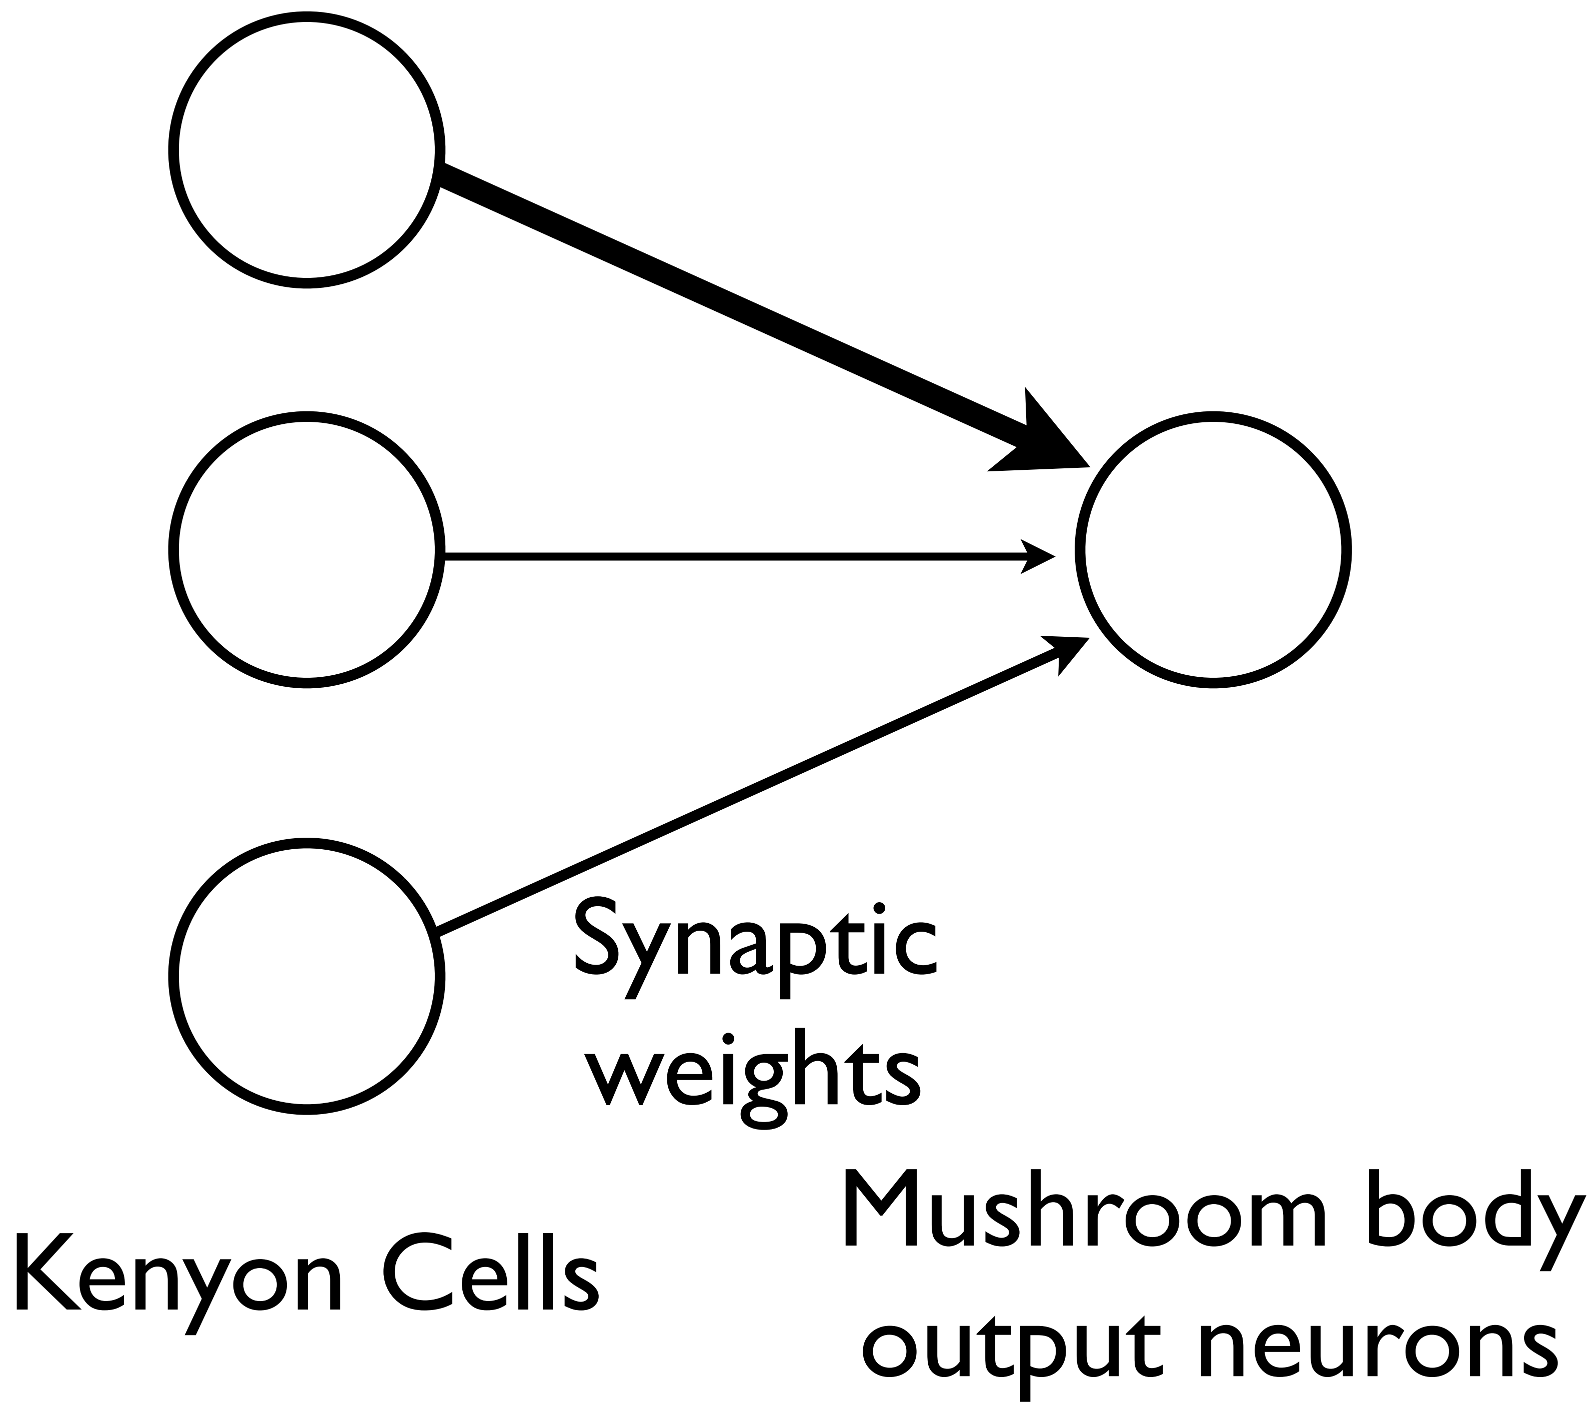

Perceptron learning rule

Learning is guided by teaching signal (e.g., Dopamine)

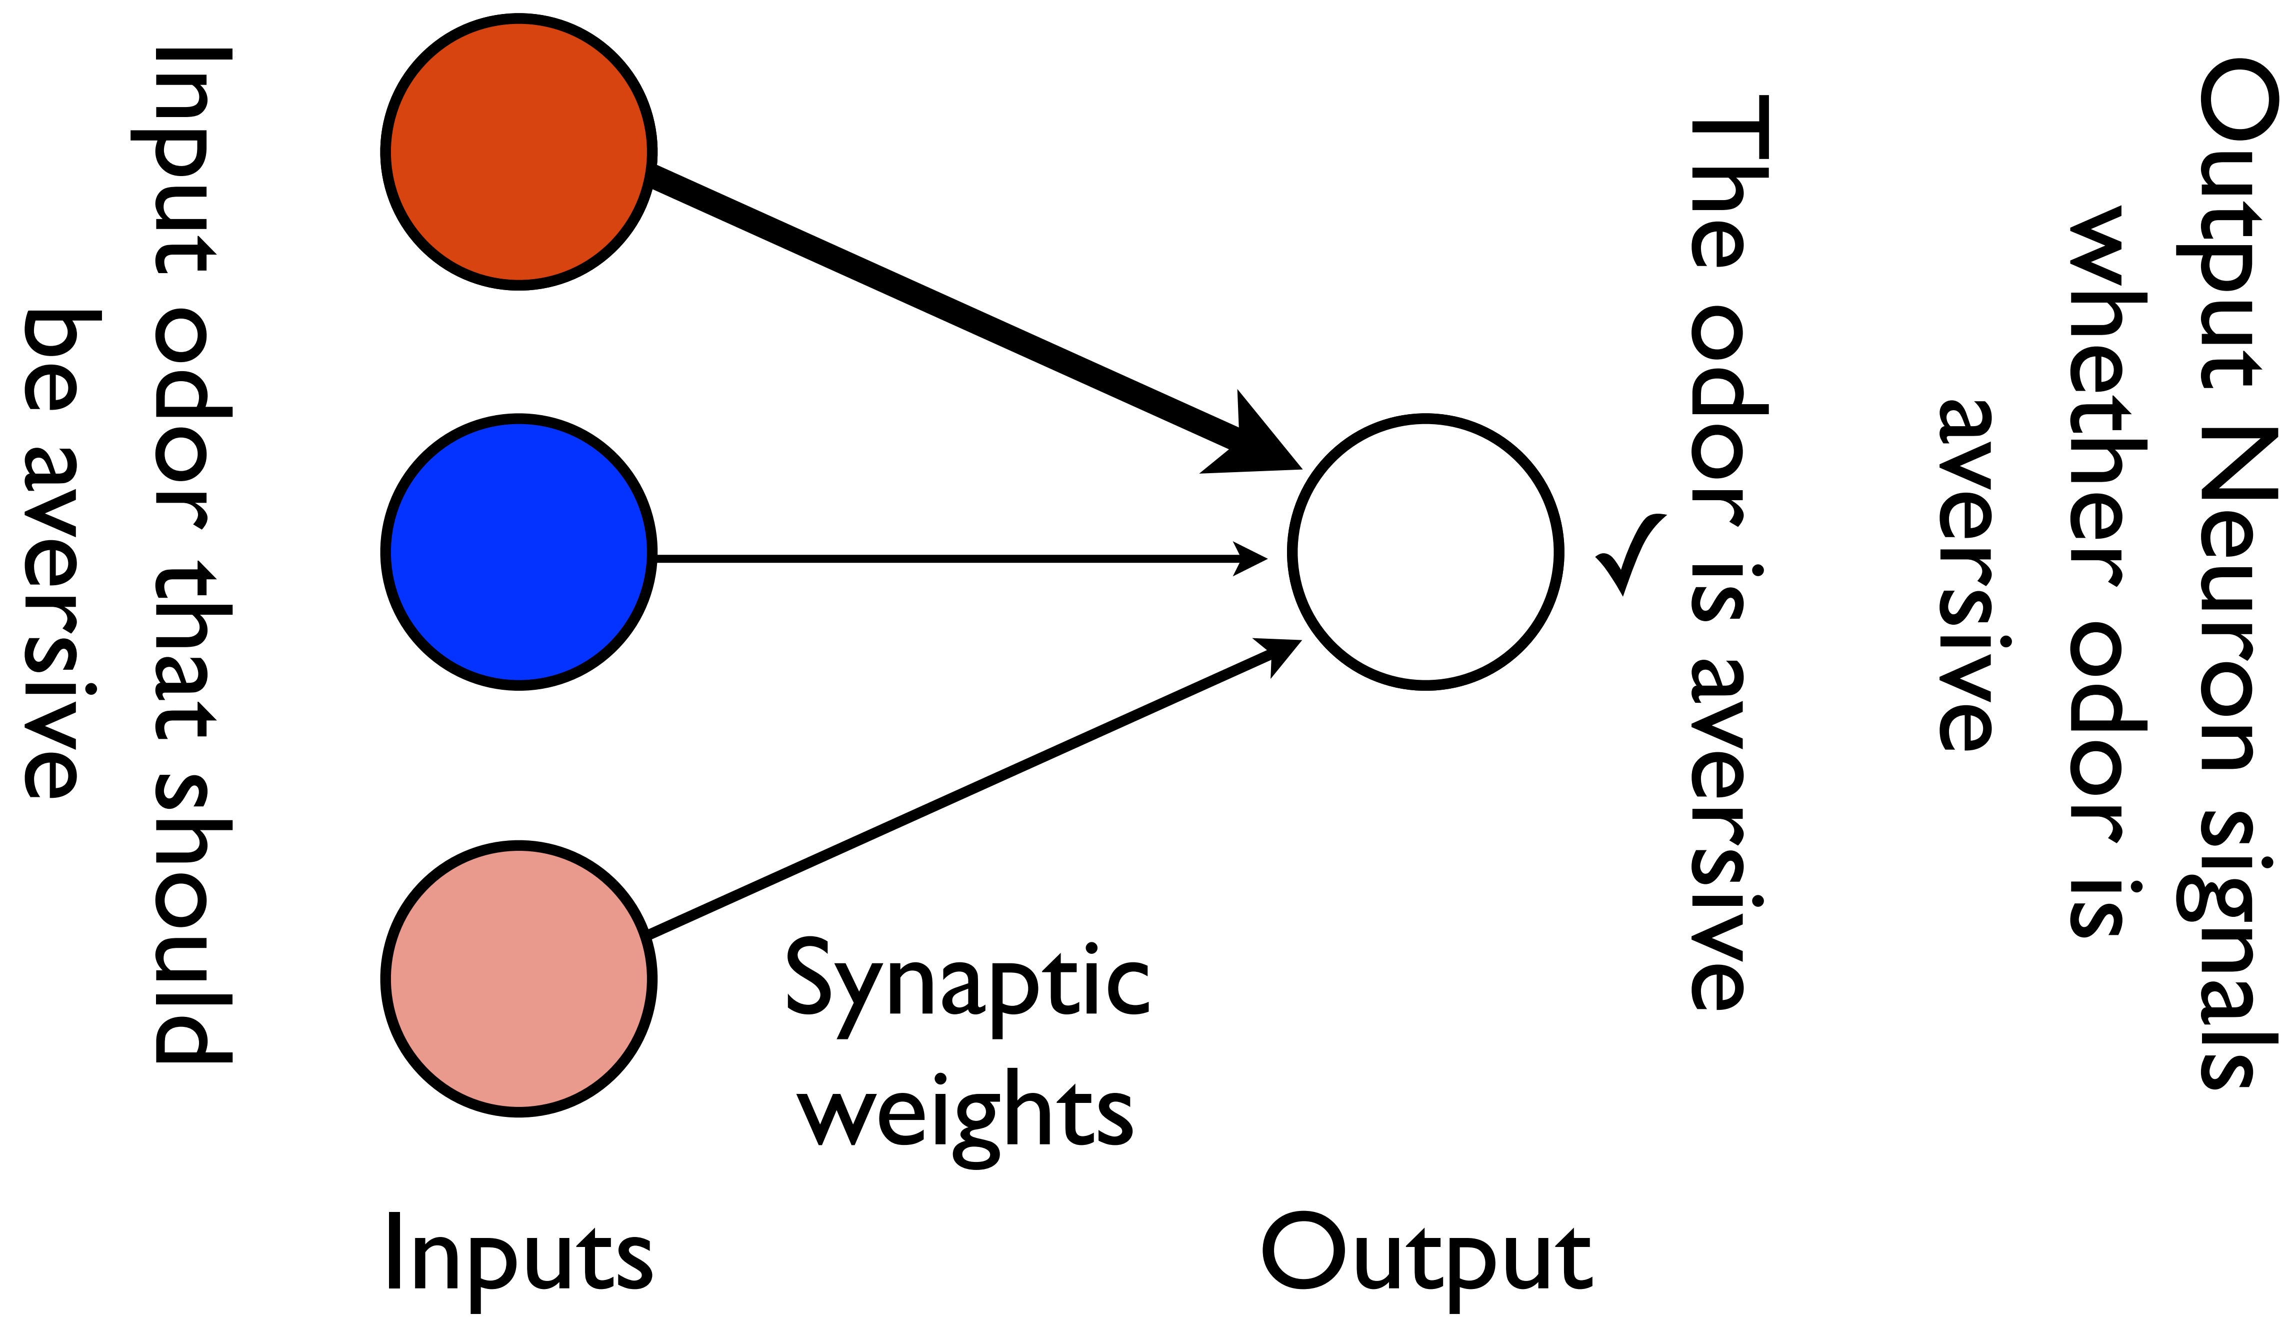

Perceptron learning rule

Learning is guided by teaching signal (e.g., Dopamine)

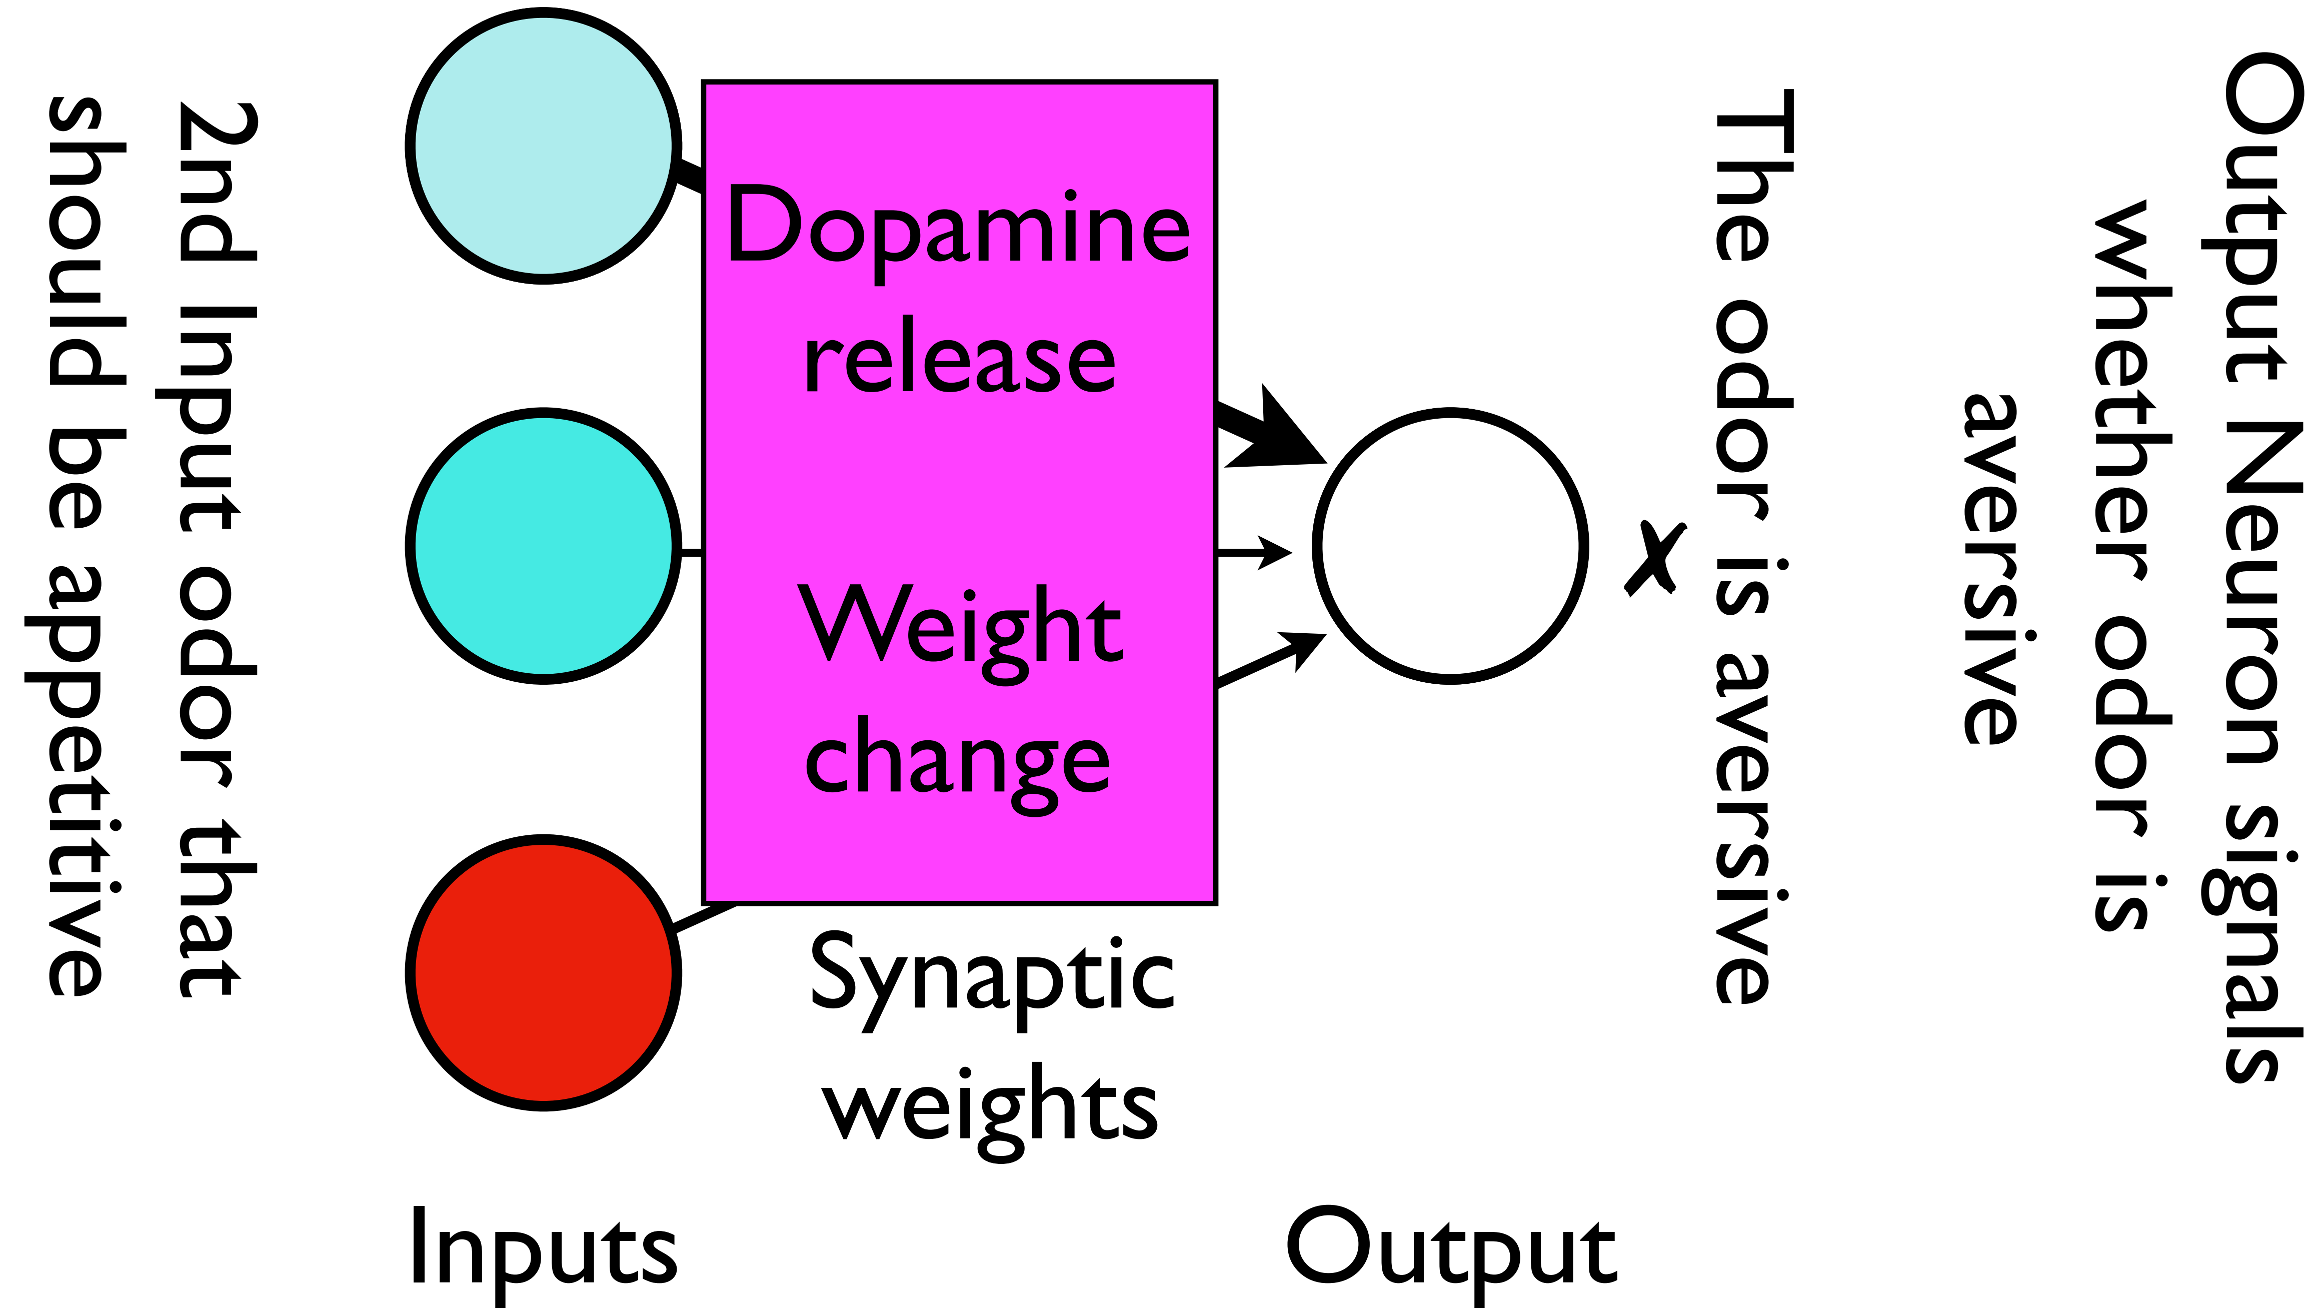

Perceptron learning rule

Learning is guided by teaching signal (e.g., Dopamine)

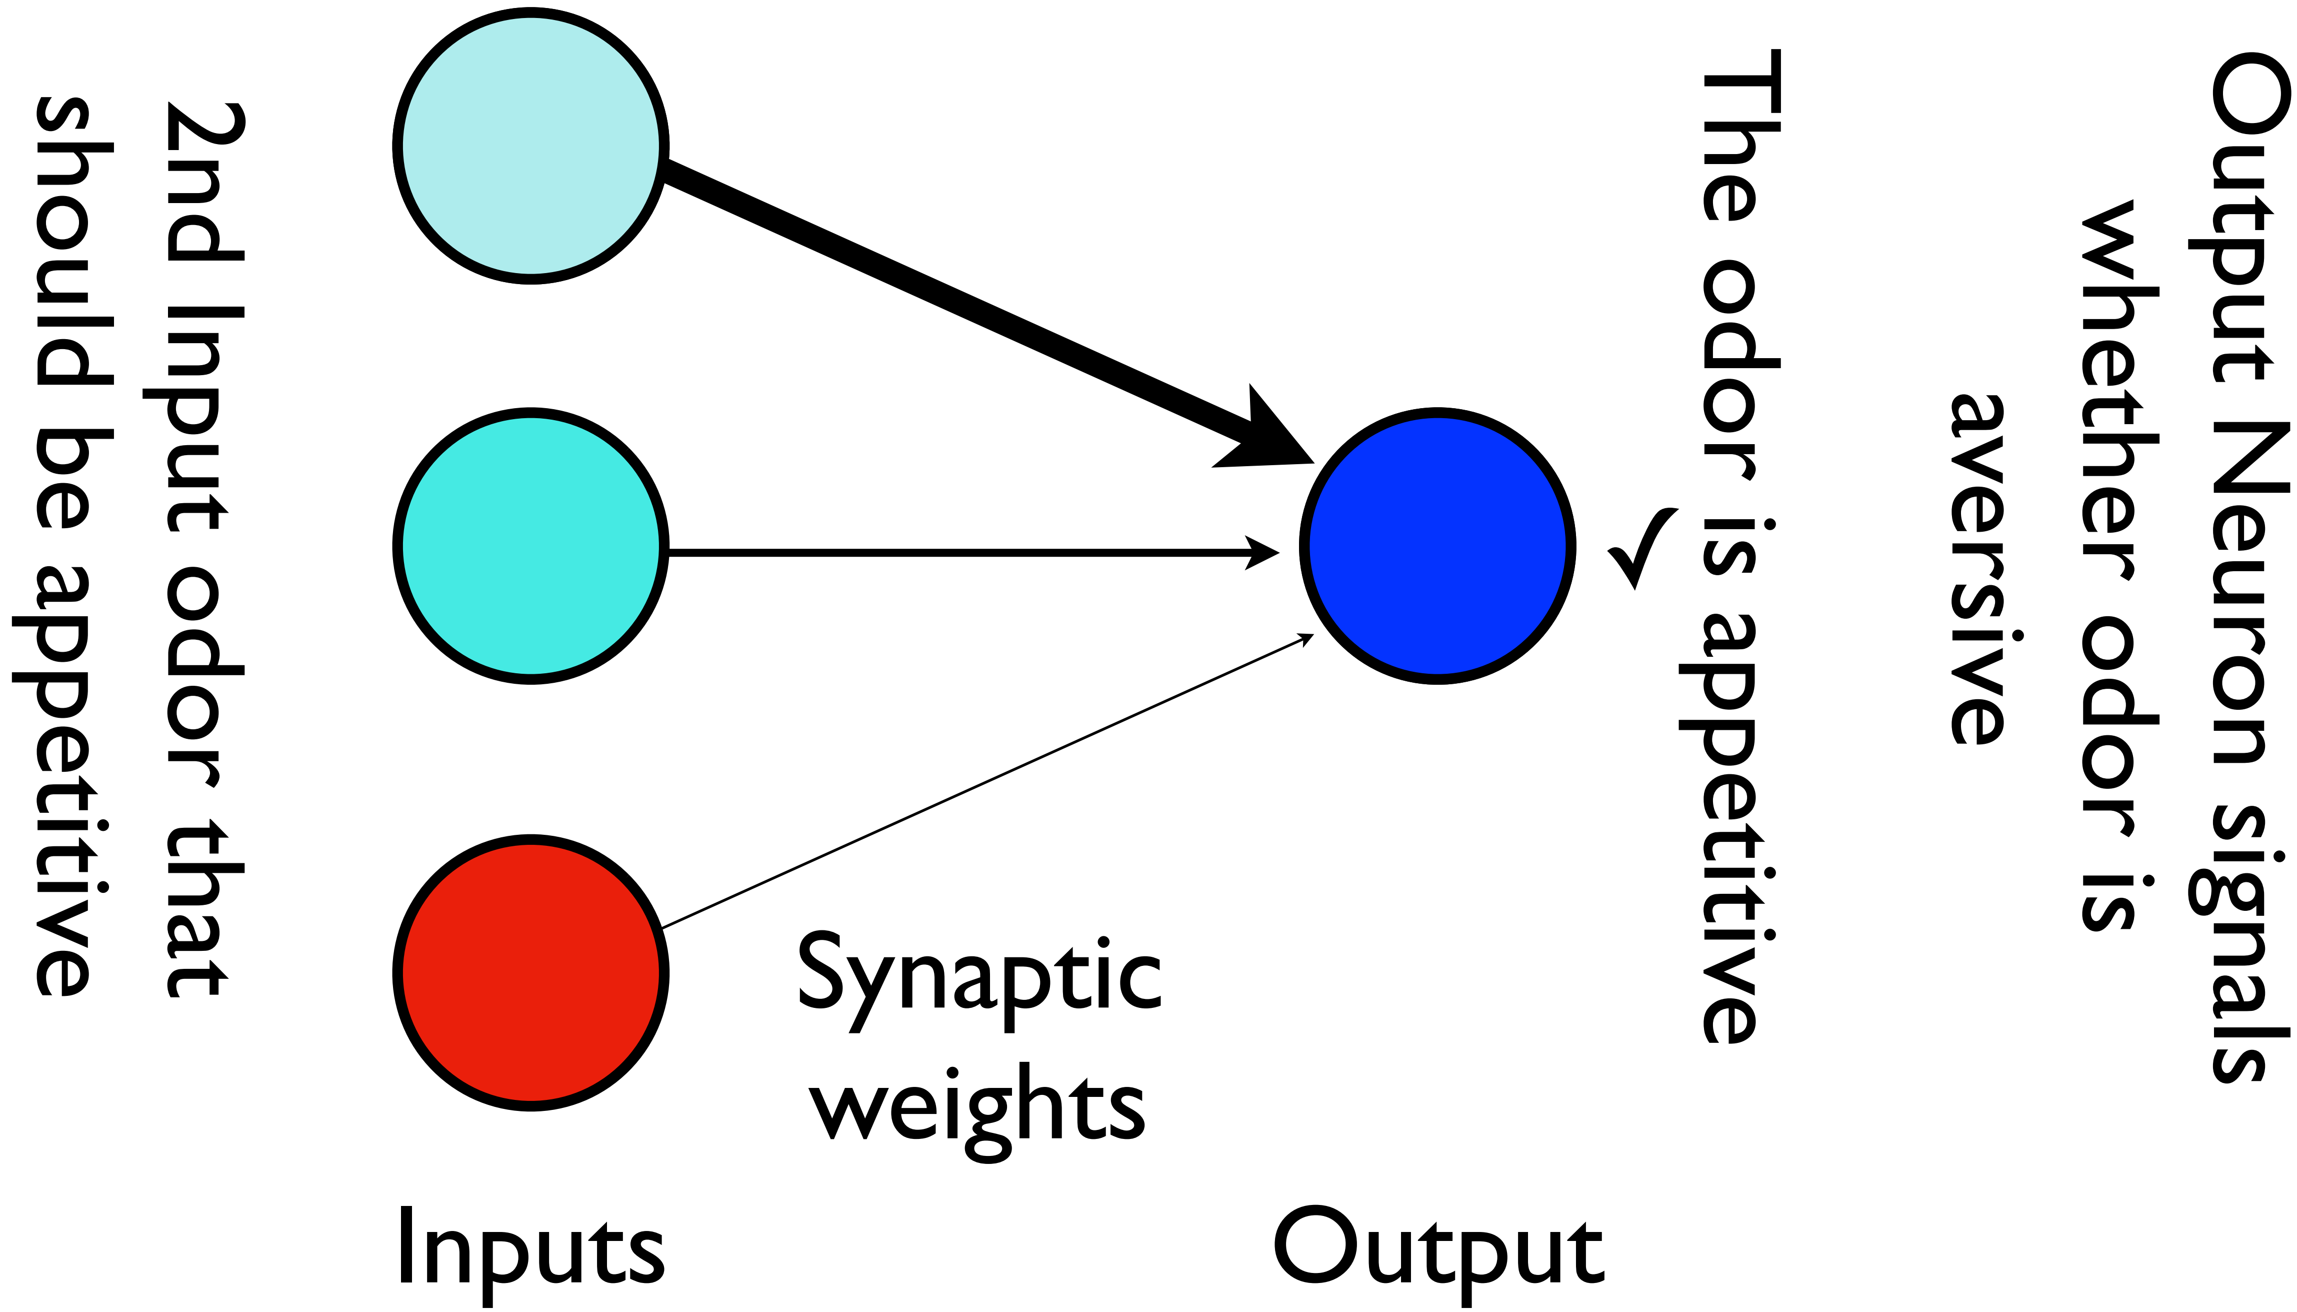

Perceptron learning

The circuit should change synaptic weights to have desired outputs for specific inputs

Aversive associated odors

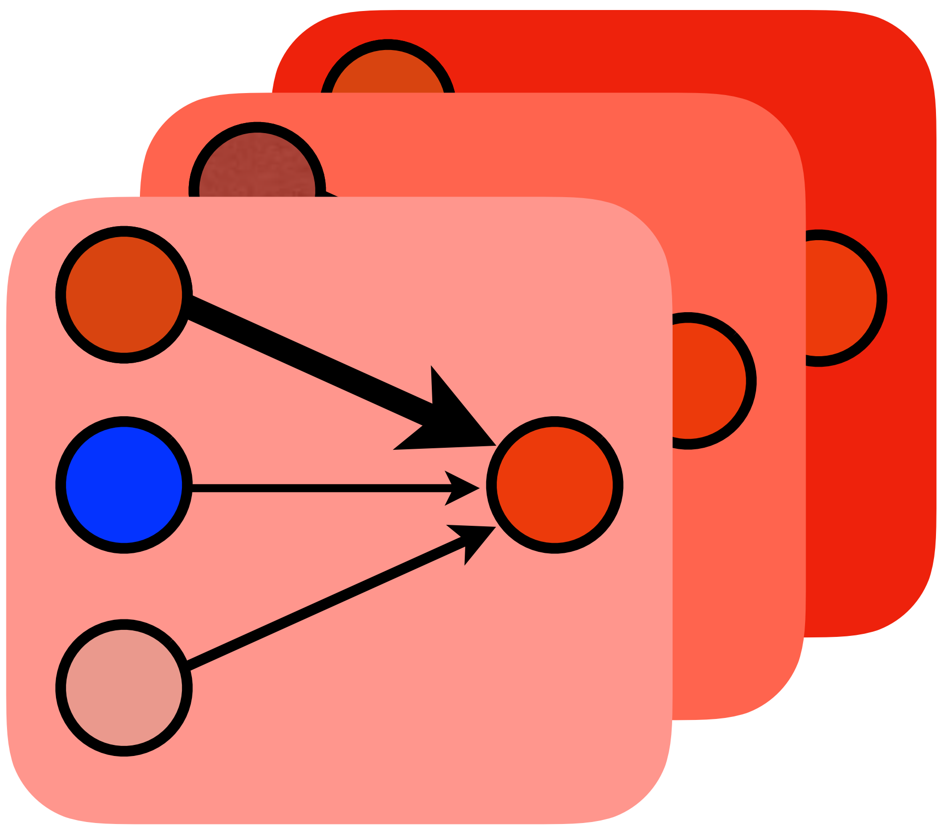

Appetitive associated odors

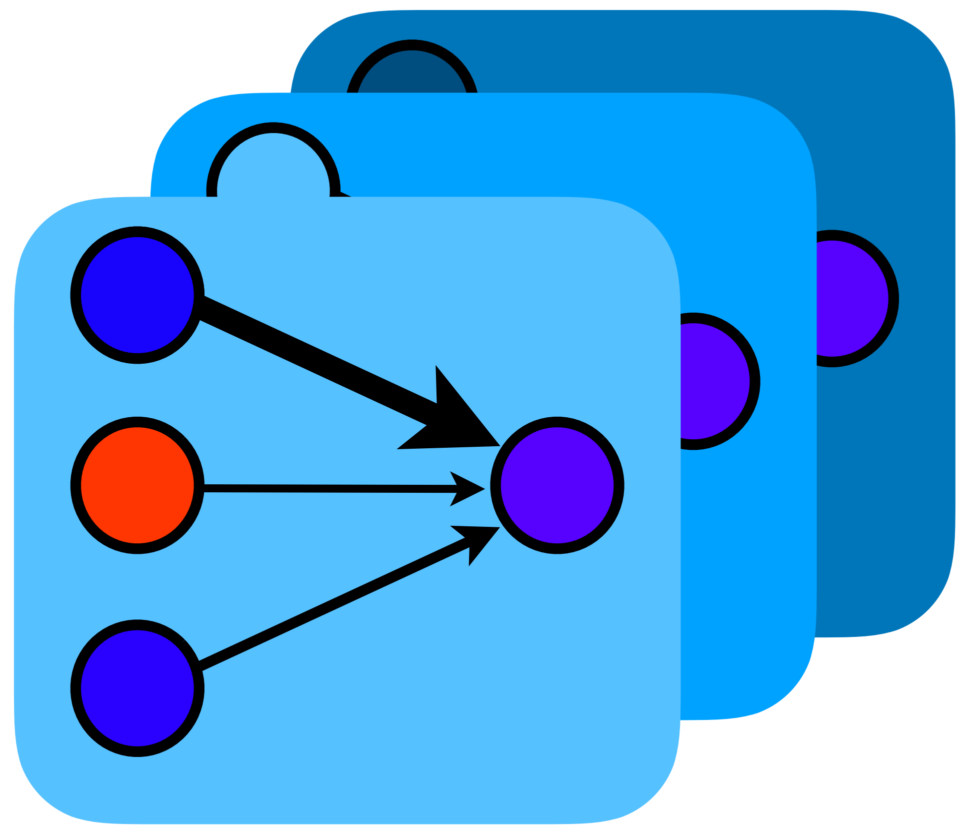

Perceptron demo

Perceptrons as model of learning

Anything a perceptron can possibly do,
it can learn to do,
just by (repeatedly) applying a simple learning rule

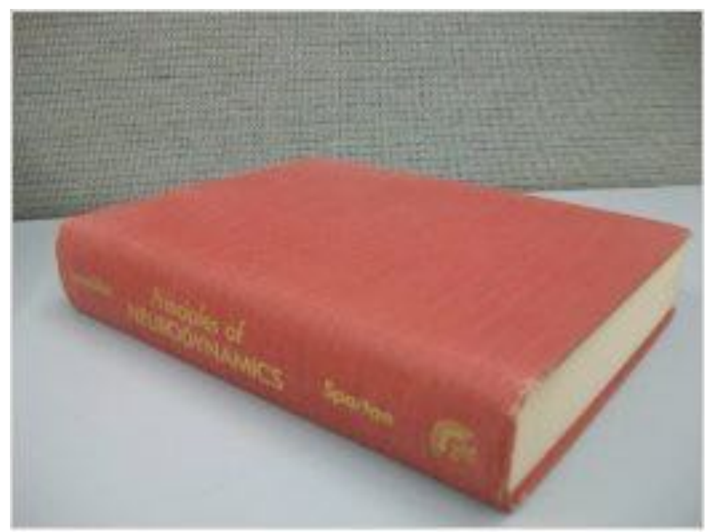

Best seller rank: 669,003 in books

Shakespeare's King Lear: #435,140

#1: The heroes of Olympus
(book five)

Rosenblatt (1962), Principles of Neurodynamics

Perceptrons as model of learning

Anything a perceptron can possibly do, it can learn to do

Inputs - percepts   Outputs - actions

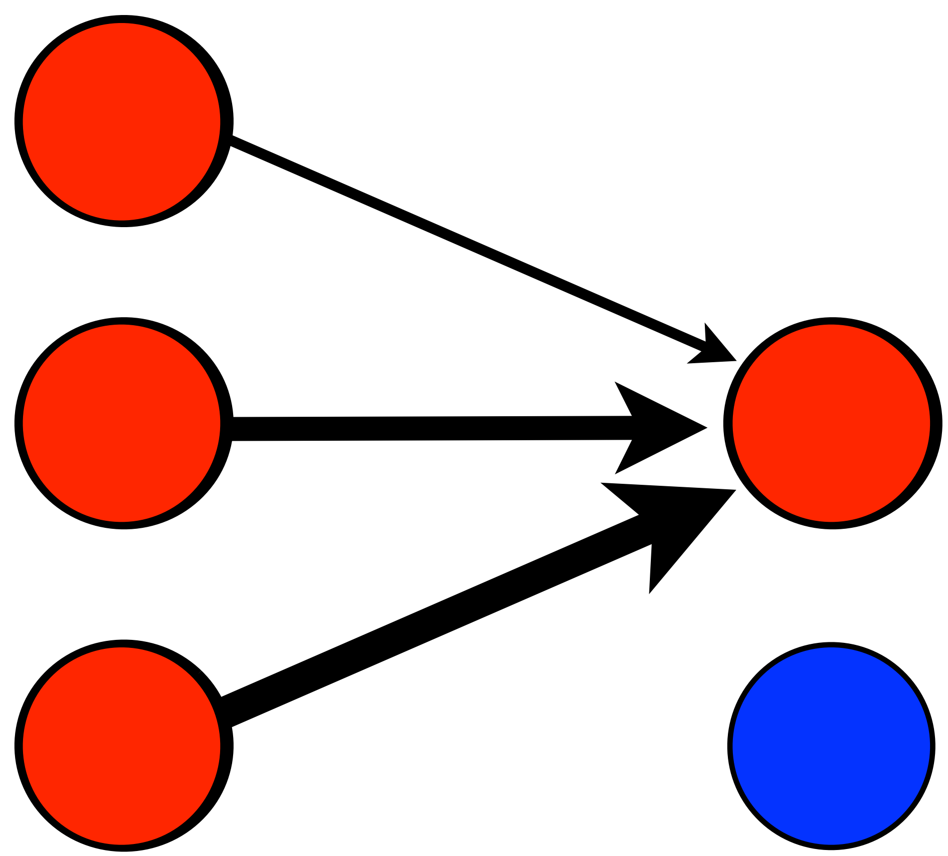

Rosenblatt (1962), Principles of Neurodynamics

A claim too far

Anything a perceptron can possibly do,
it can learn to do,
just by repeatedly applying a simple learning rule

Claim 2:
With enough perceptrons you can compute (almost) anything

We have a learning machine that can learn anything

Building a brain is now just a matter of implementation

Rosenblatt (1962), Principles of Neurodynamics

A claim too far

Building a brain is now just a matter of implementation

Rosenblatt (1962), *Principles of Neurodynamics*

The perceptron wars

Rosenblatt:
With enough perceptrons you can compute (almost) anything

Minsky and Papert:
(Simple, reasonable) perceptrons have fundamental limitations on what they can learn

Building a brain is **not** just a matter of implementation (yet)

Minsky and Papert (1969), Perceptrons (208,066 in books)

Perceptrons as model of learning

Lecun, Hinton and many others

With enough perceptrons in enough layers you can learn to compute (almost) anything

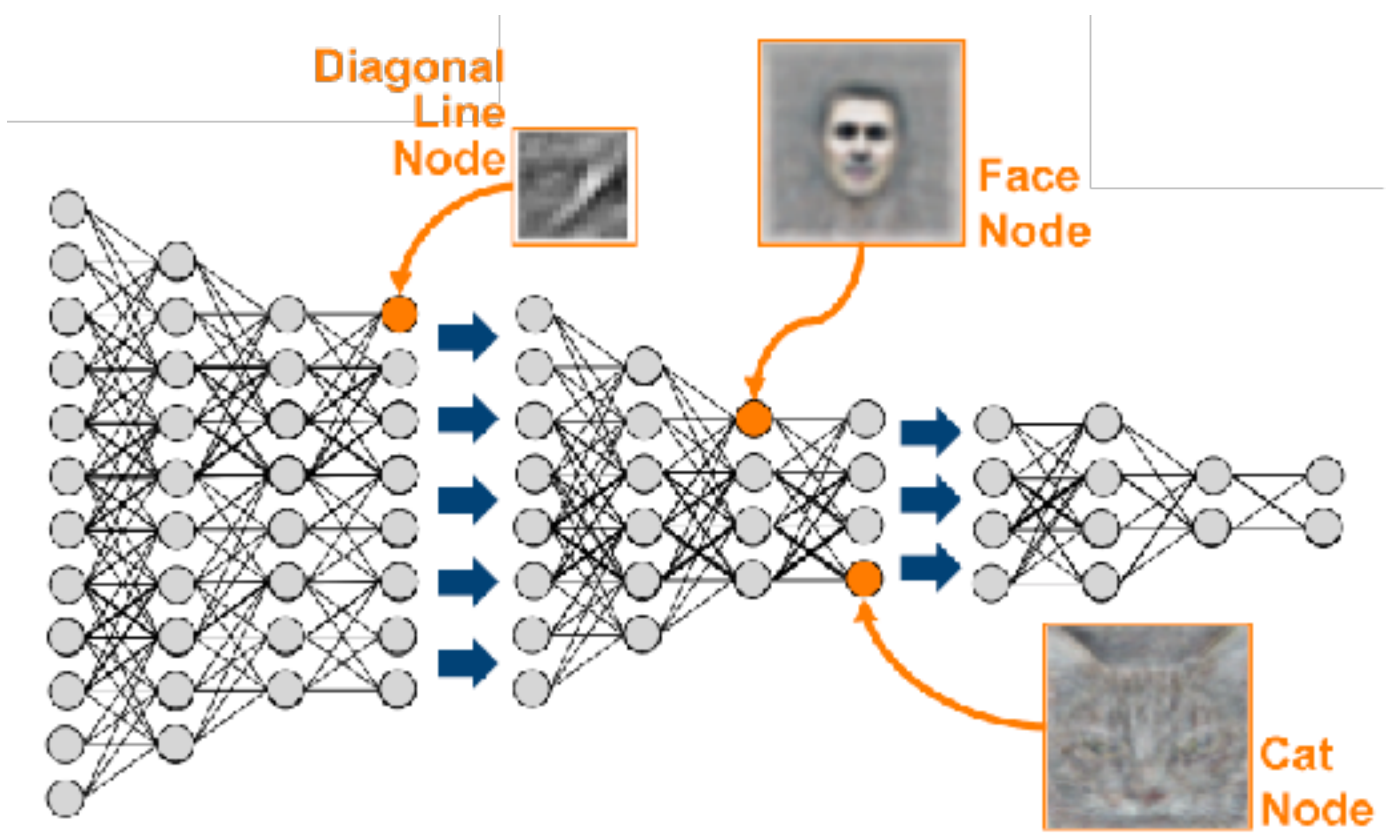

Perceptron demo

Interim summary

Perceptrons can learn to associate arbitrary input patterns with arbitrary output patterns

They can do so with a semi-biologically reasonable learning rule

They may be useful to understand associative learning

Lets do some math

Models make strong, often non-biologically realistic assumptions in order to simplify and solve systems

If you have a better sense of the math you will be able to better judge the impact of each assumption

The more domain-knowledge experts that can do this, the more neuroscience will improve