



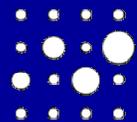
Cognitive Modelling

Establishing Topologically Ordered Population Codes

Tom Hartley

t.hartley@psychology.york.ac.uk

THE UNIVERSITY *of York*



VolkswagenStiftung

Hanse



Wissenschaftskolleg

Overview

- ❖ Examples of topological order in the brain
- ❖ Classification learning
 - ◆ supervised
 - ◆ unsupervised
- ❖ Establishing topologically-ordered representations
 - ◆ Kohonen, Self-Organizing Maps
- ❖ Demonstration using JavaNNS

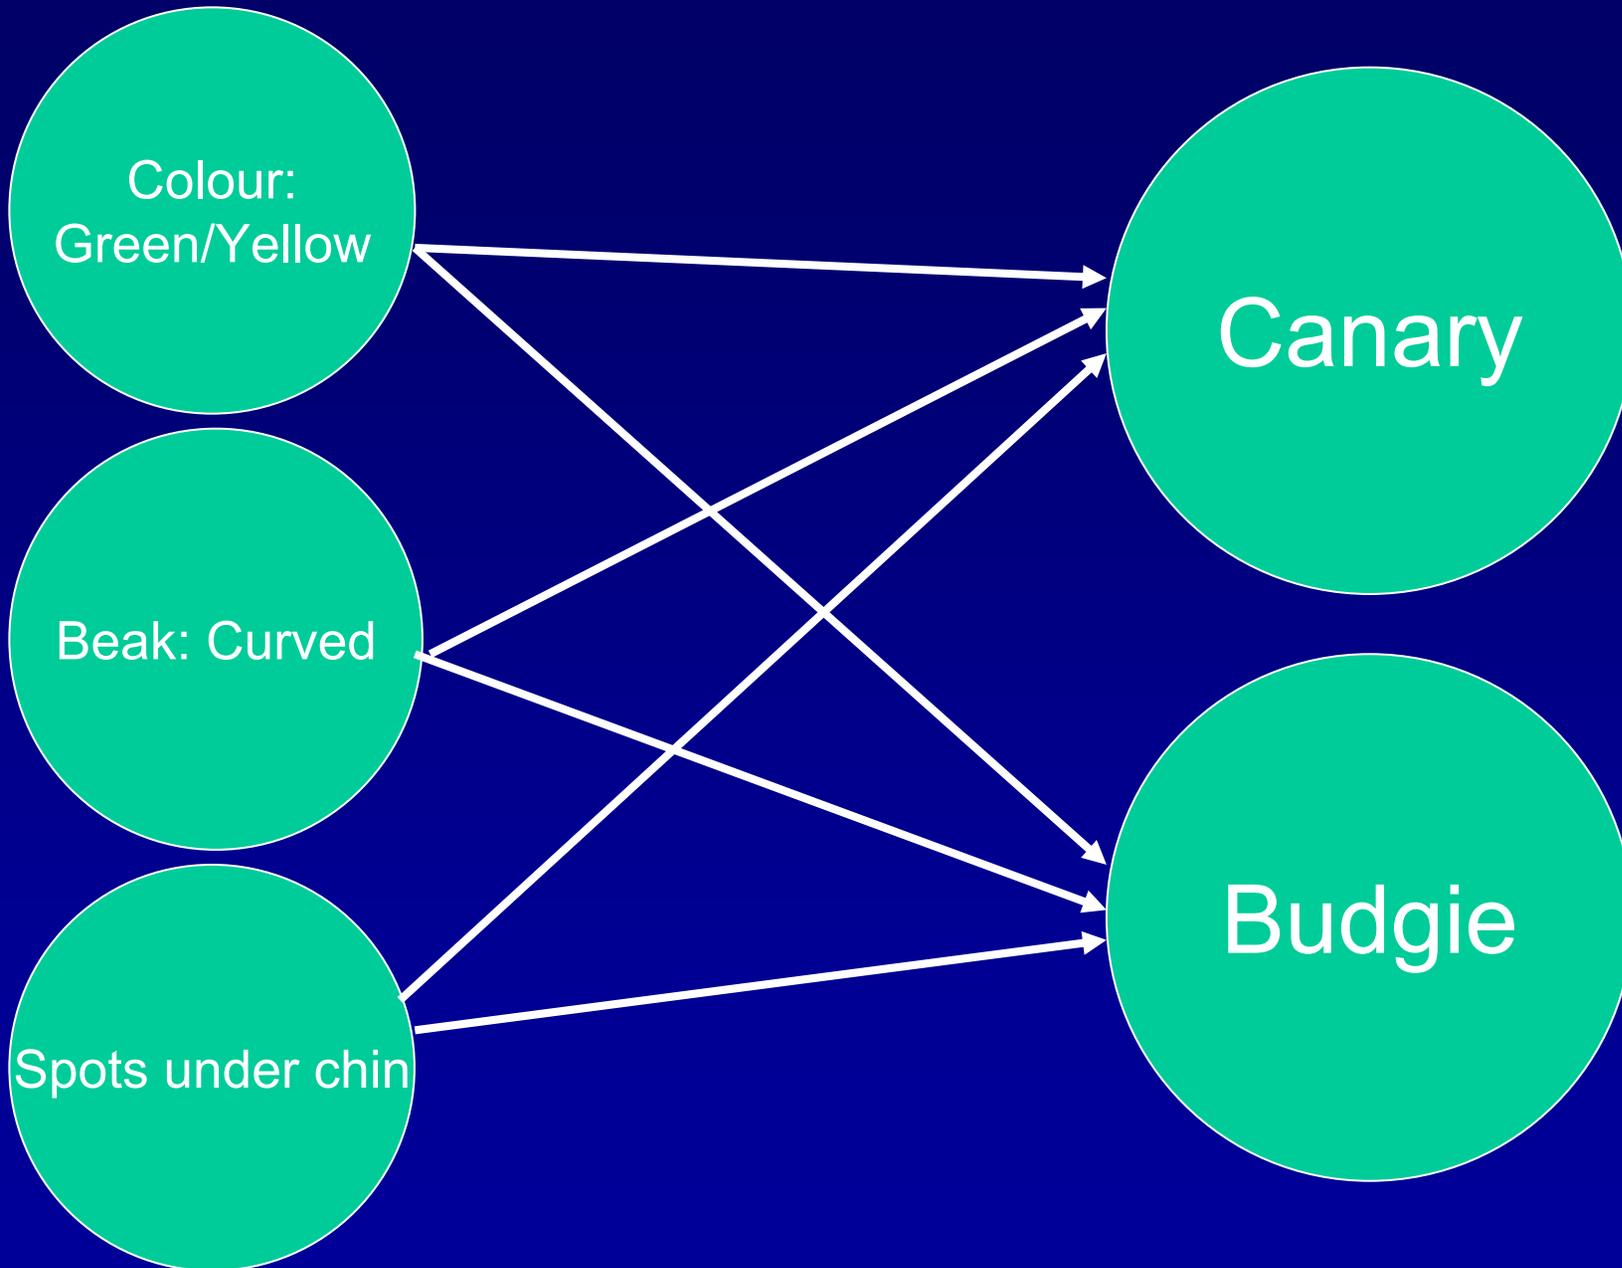
Learning

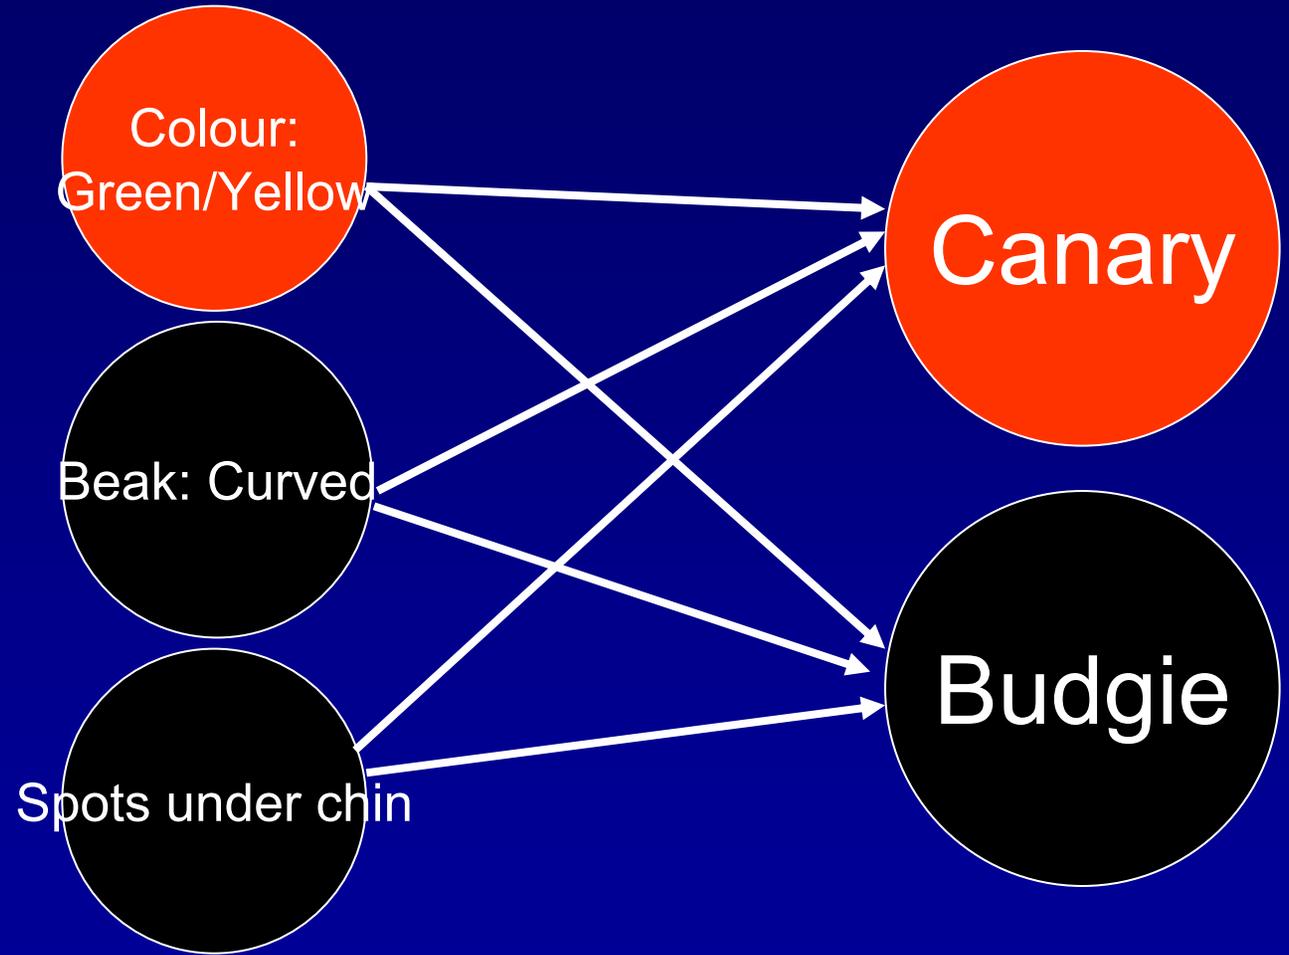
- ❖ Overview of how to adjust weights so that network performs a useful function
- ❖ Example: classifying canaries and budgies.

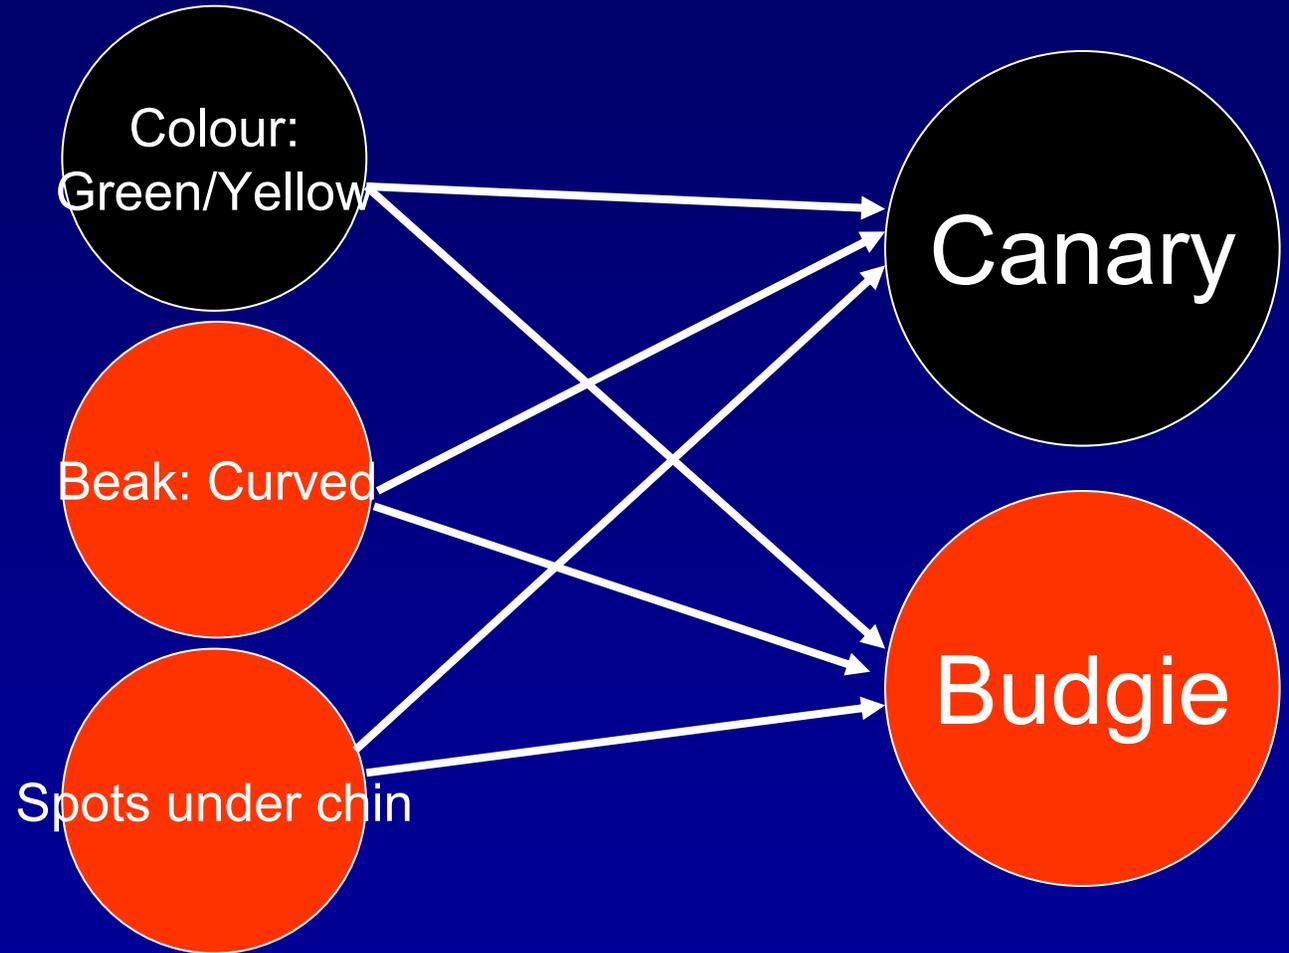


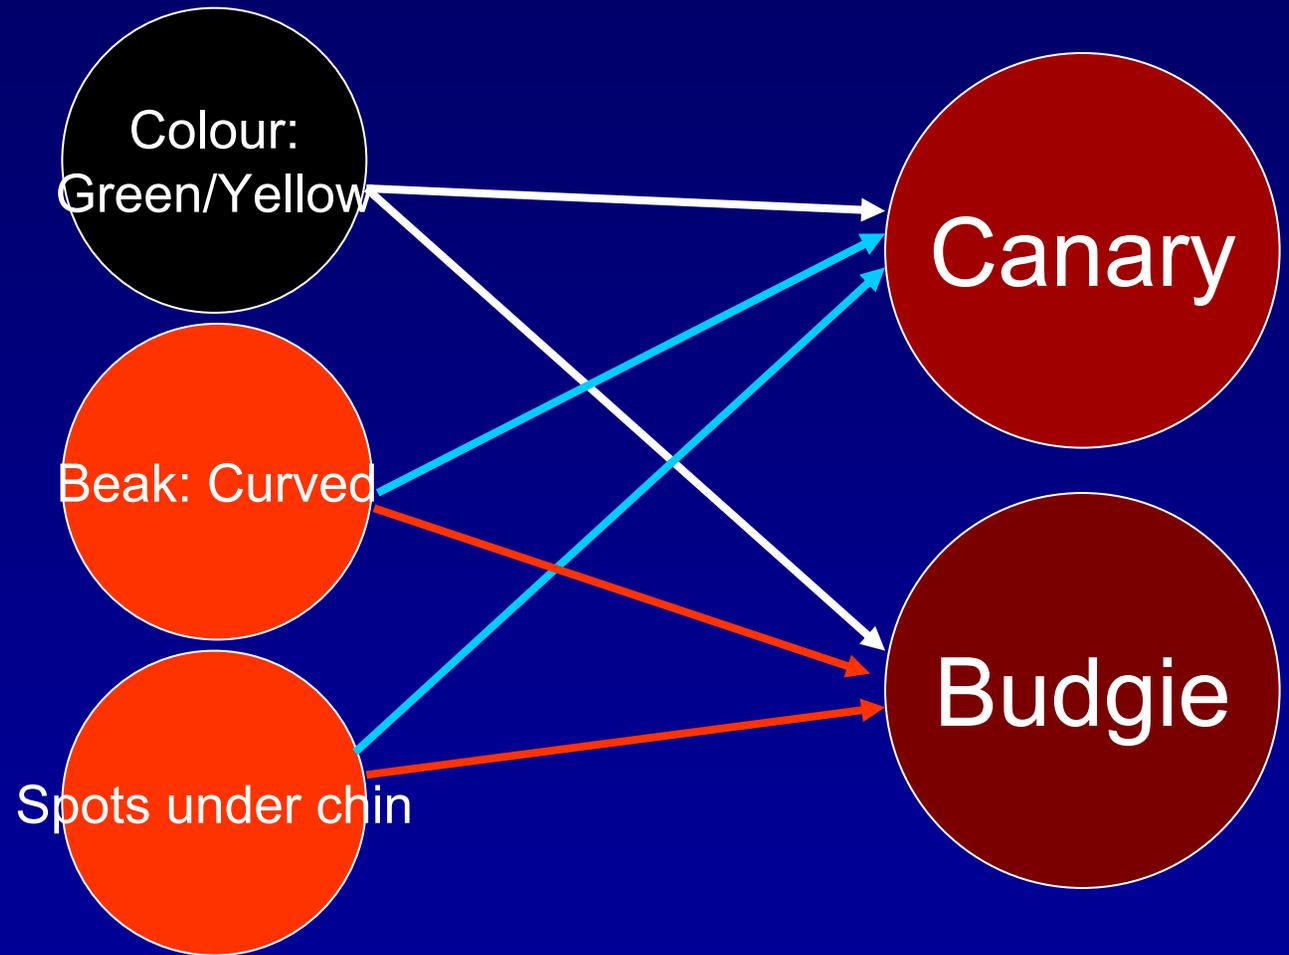
Feature

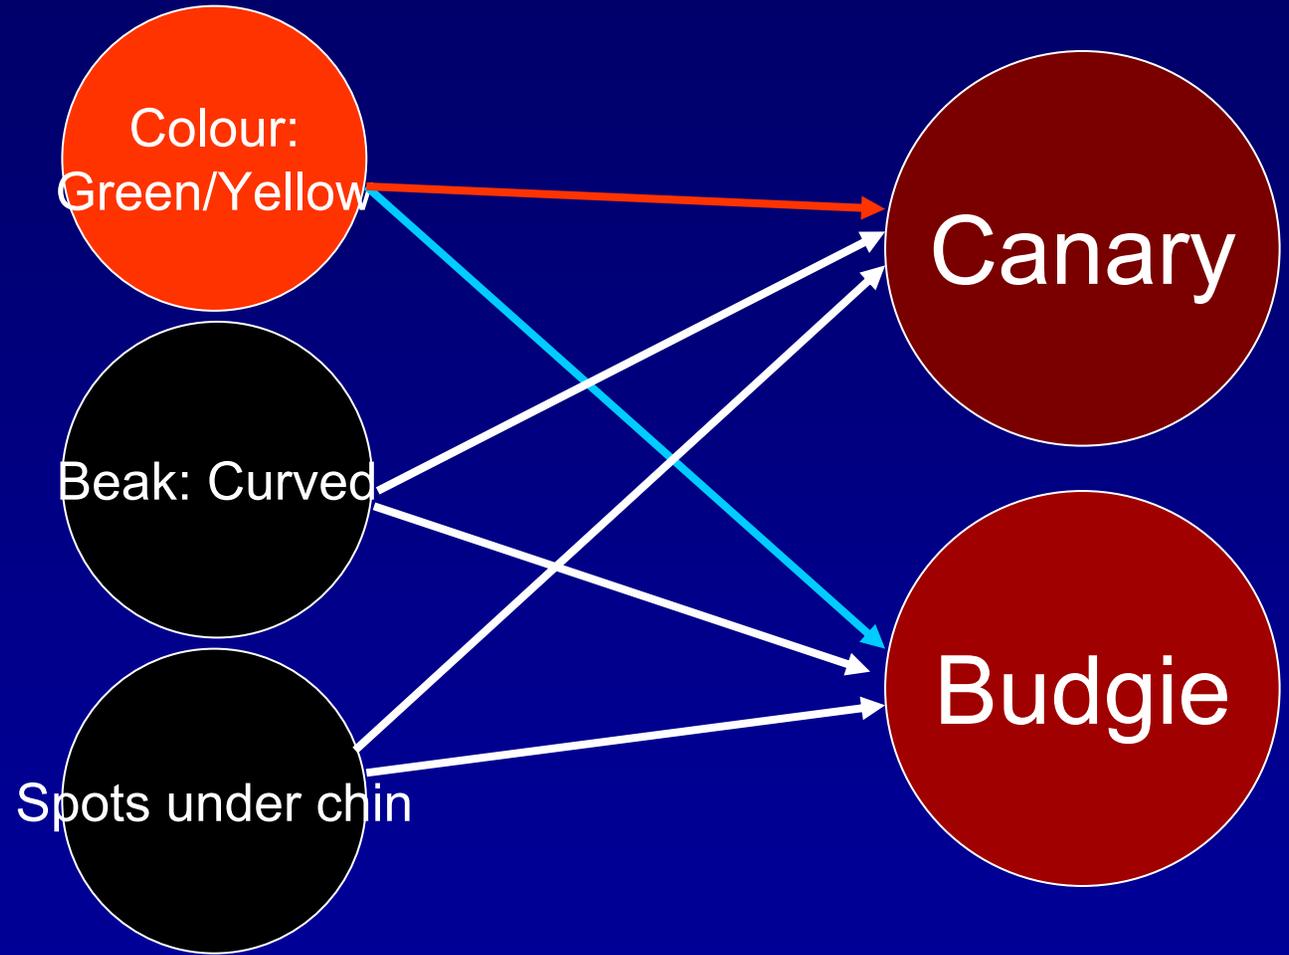
Category











Supervised Learning

- ❖ After each input pattern
 - ◆ compare output with correct/desired output.
 - ◆ weaken/strengthen each connection proportional to its contribution to the error.
- ❖ Repeat for many input patterns

this simplified description covers the supervised learning algorithms known as the delta rule for one layer of weights, and generalized delta rule (aka backpropagation of error) for two or more.

beak curve

color



Pattern Classification Limits

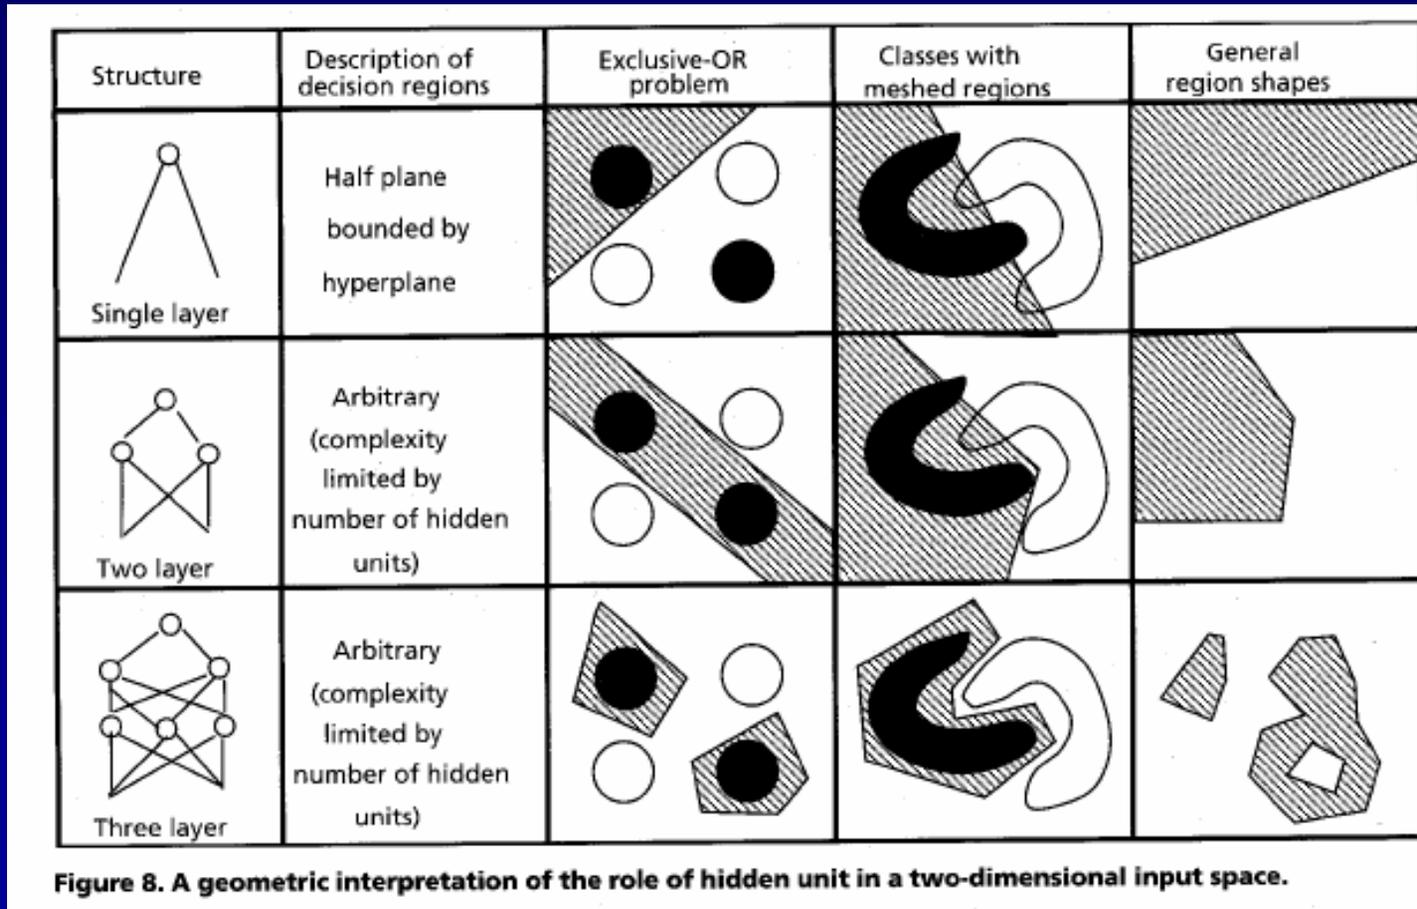
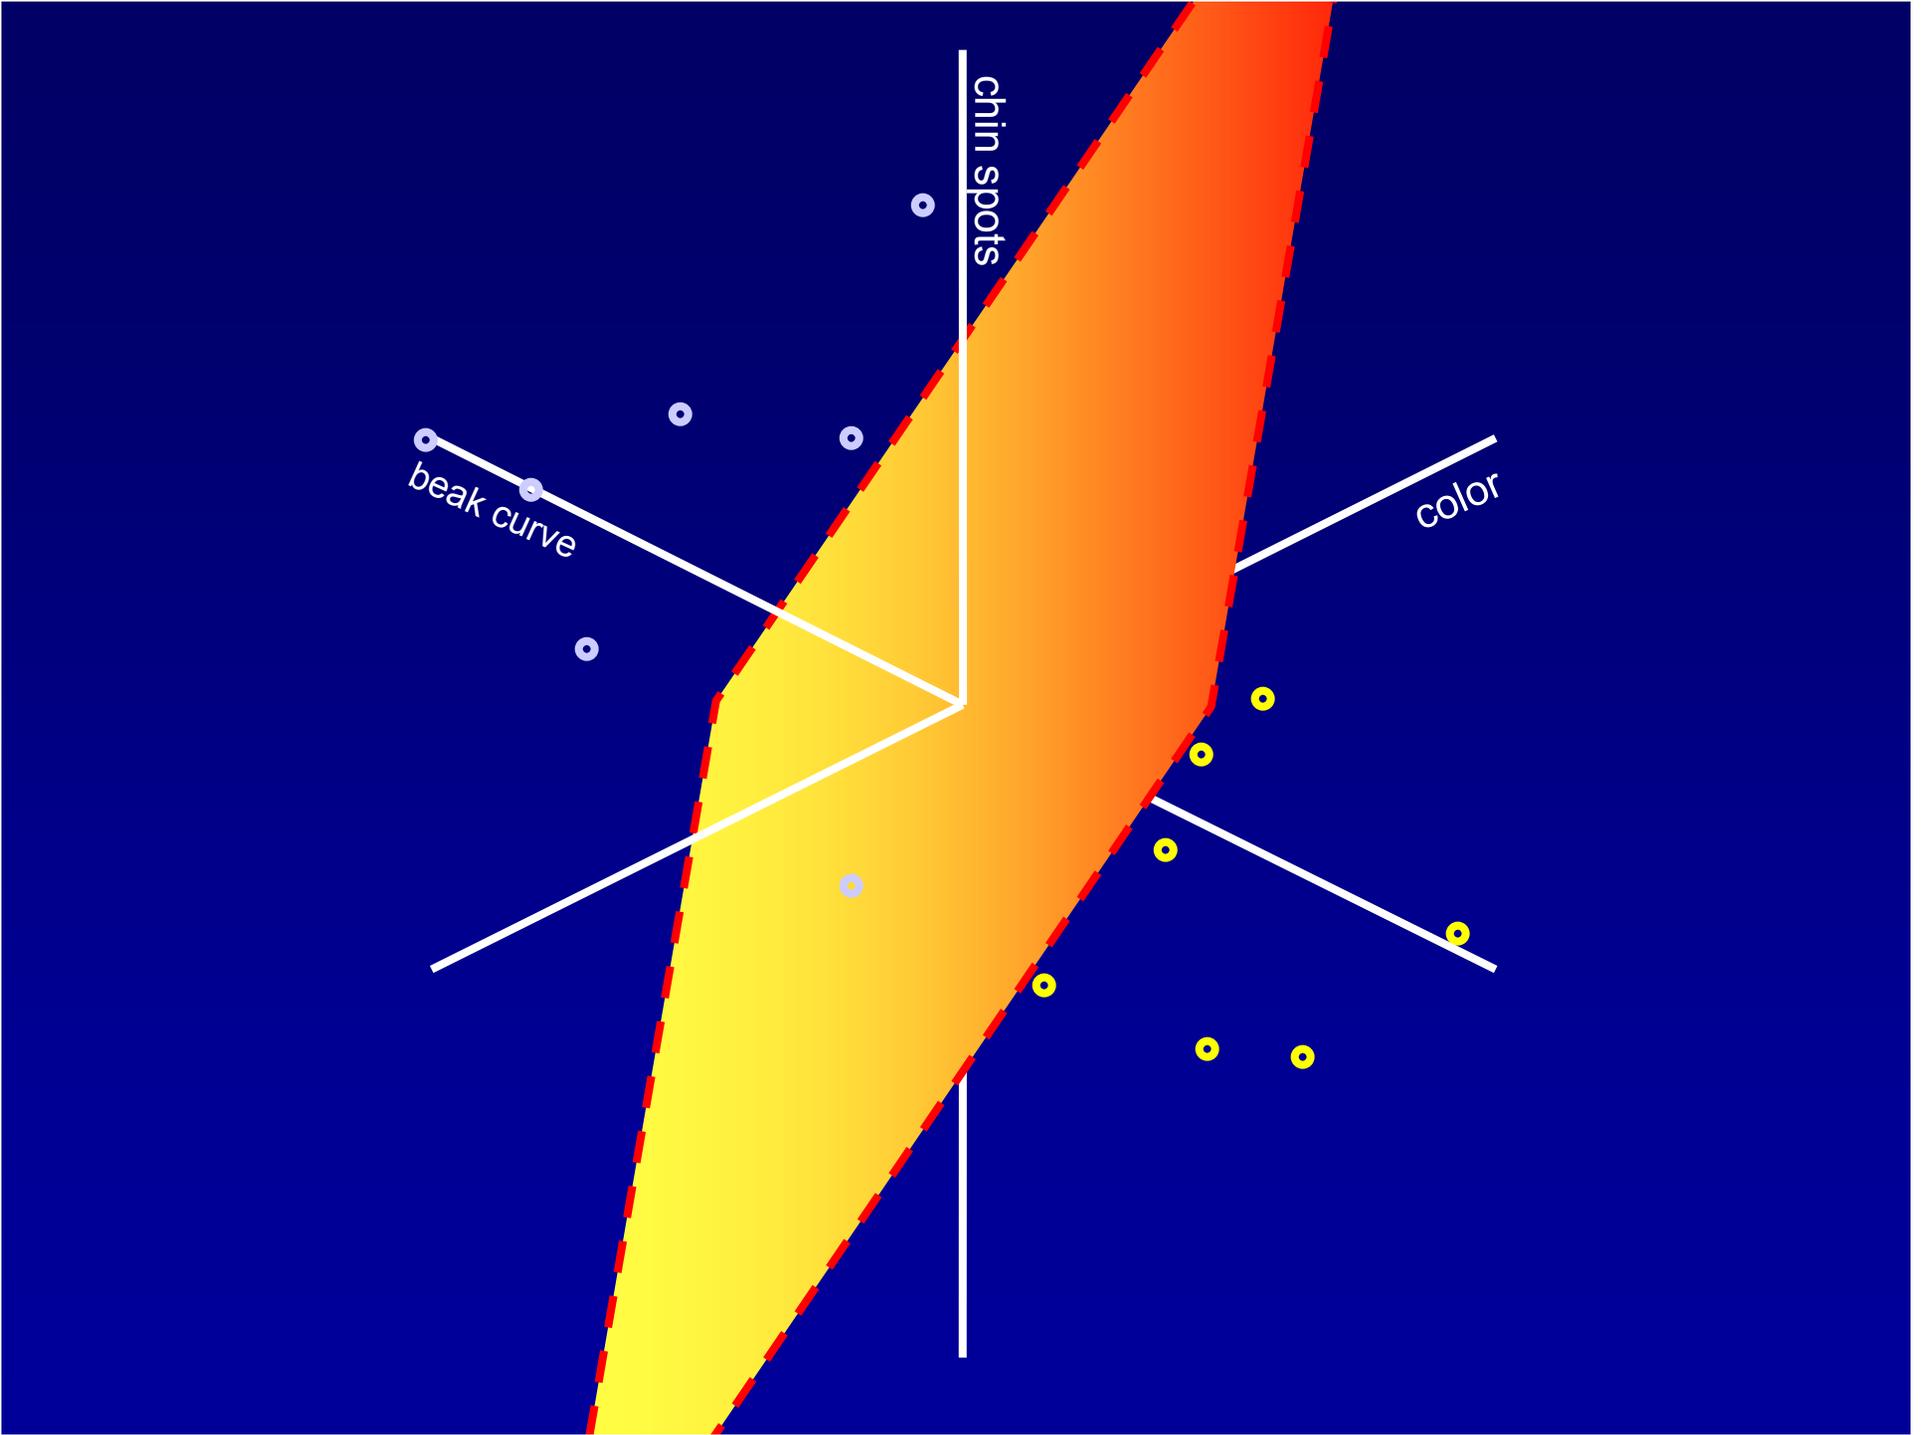


Figure 8. A geometric interpretation of the role of hidden unit in a two-dimensional input space.

For more complex problems, more layers are needed. Three layers of weights are sufficient to solve any classification problem (given enough neurons).

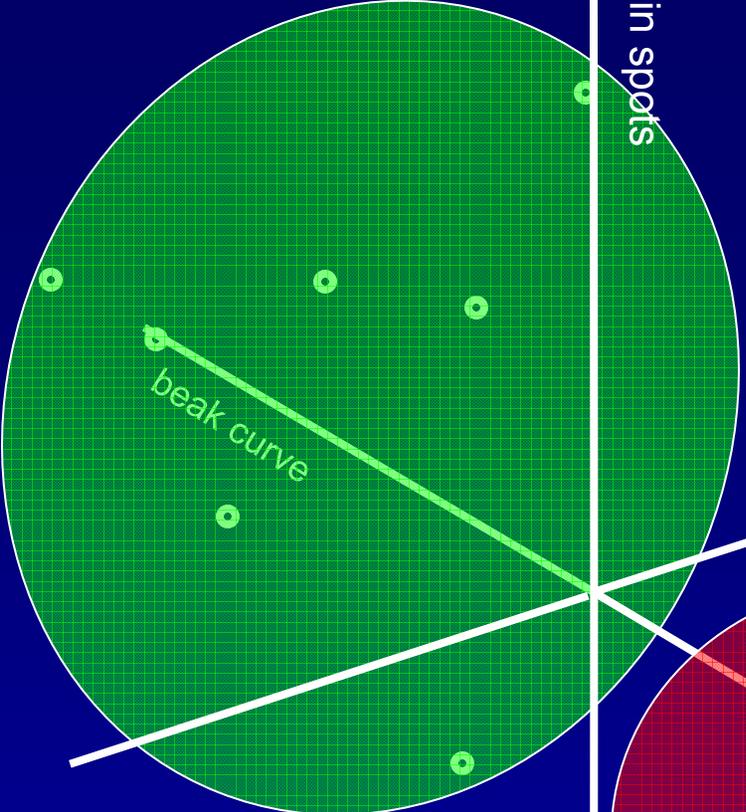


Supervised Learning

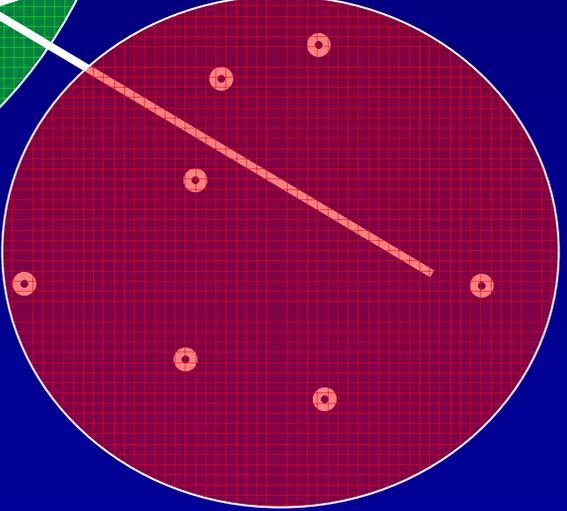
- ❖ Given the right weights, simple neural networks can solve complex problems.
- ❖ However the learning algorithms outlined in the previous example are not particularly realistic for many situations
 - ◆ Detailed “teaching” input is implausible for many tasks
 - ◆ Backpropagation of error (multilayer networks) is implausible
- ❖ Reinforcement learning is similar to supervised learning but avoids these issues.
 - ◆ No detailed teaching input – just global assessment of “correctness”.

Unsupervised Learning

- ❖ Works by discovering inherent structure in the input.
- ❖ Simple learning rules based on Hebb postulate:
 - ◆ Connections are strengthened between neurons that fire concurrently (Hebb, 1949)
 - ◆ More biologically plausible (LTP, LTD).



chin spots



color

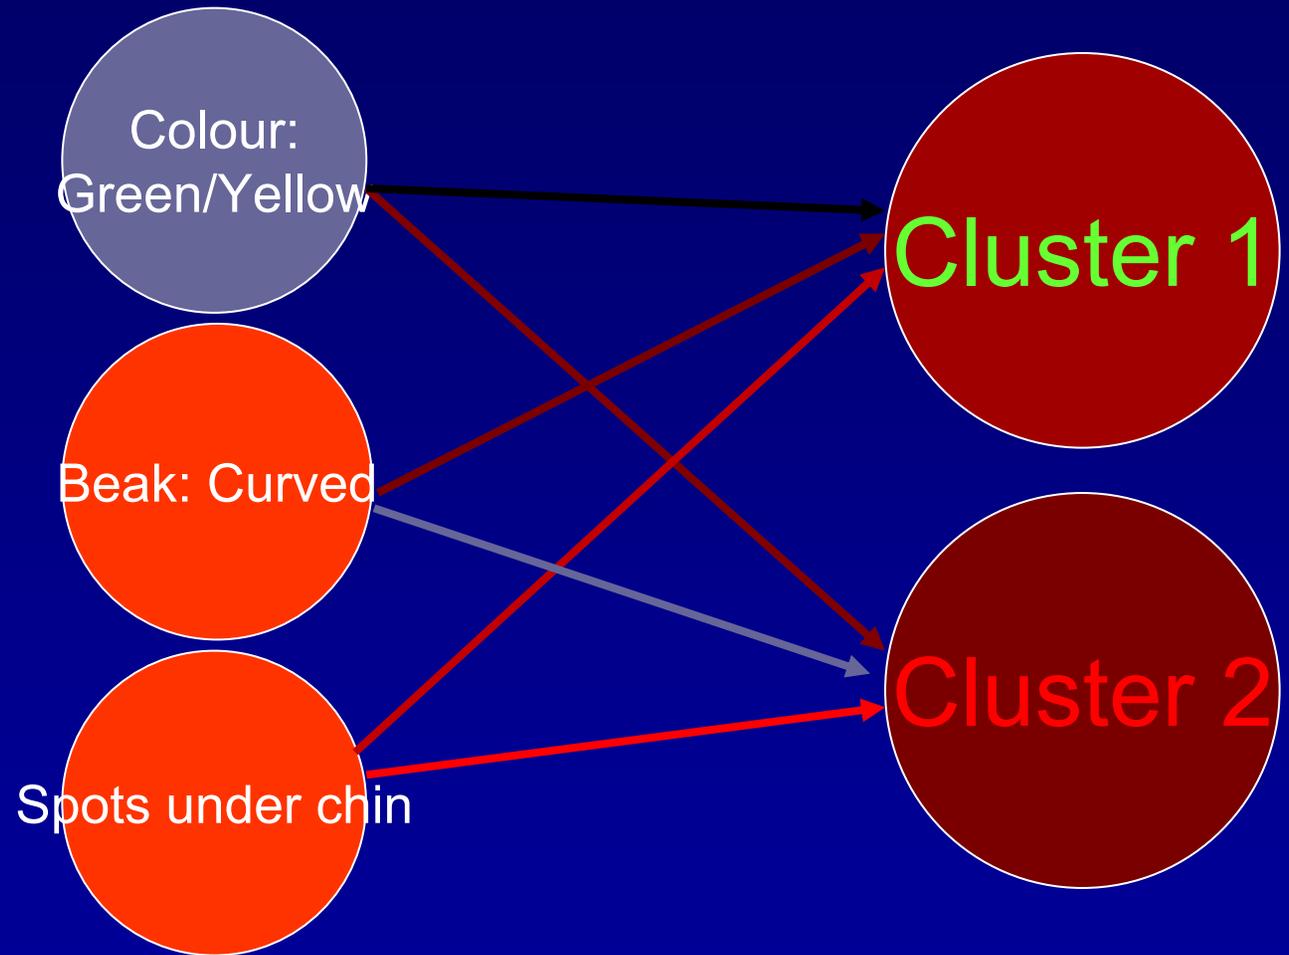
Competitive Learning

- ❖ Learns to identify clusters (and other statistical structure) in the input patterns
- ❖ The output neurons compete to respond to each input pattern
- ❖ Eventually each cluster will be represented by a different output neuron

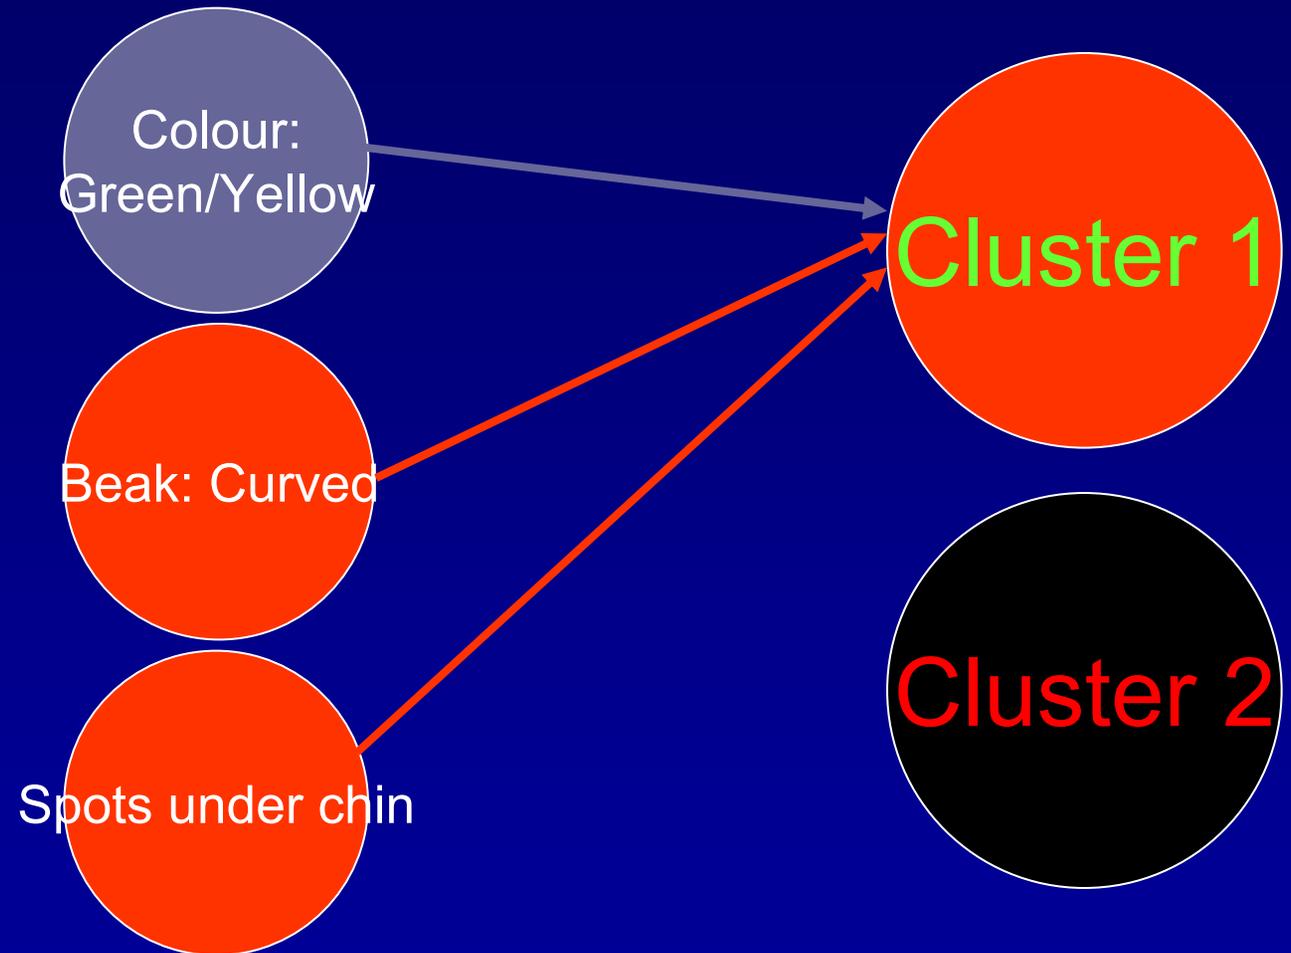
Competitive Learning

- ❖ *Output units compete*, so that eventually only one neuron (the one with the most input) is active in response to each output pattern.
- ❖ The total weight from the input layer to each output neuron is limited. If some connections are strengthened, others must be weakened.
- ❖ A consequence is that the winner is the output neuron whose *weights* best match the *activation pattern*...

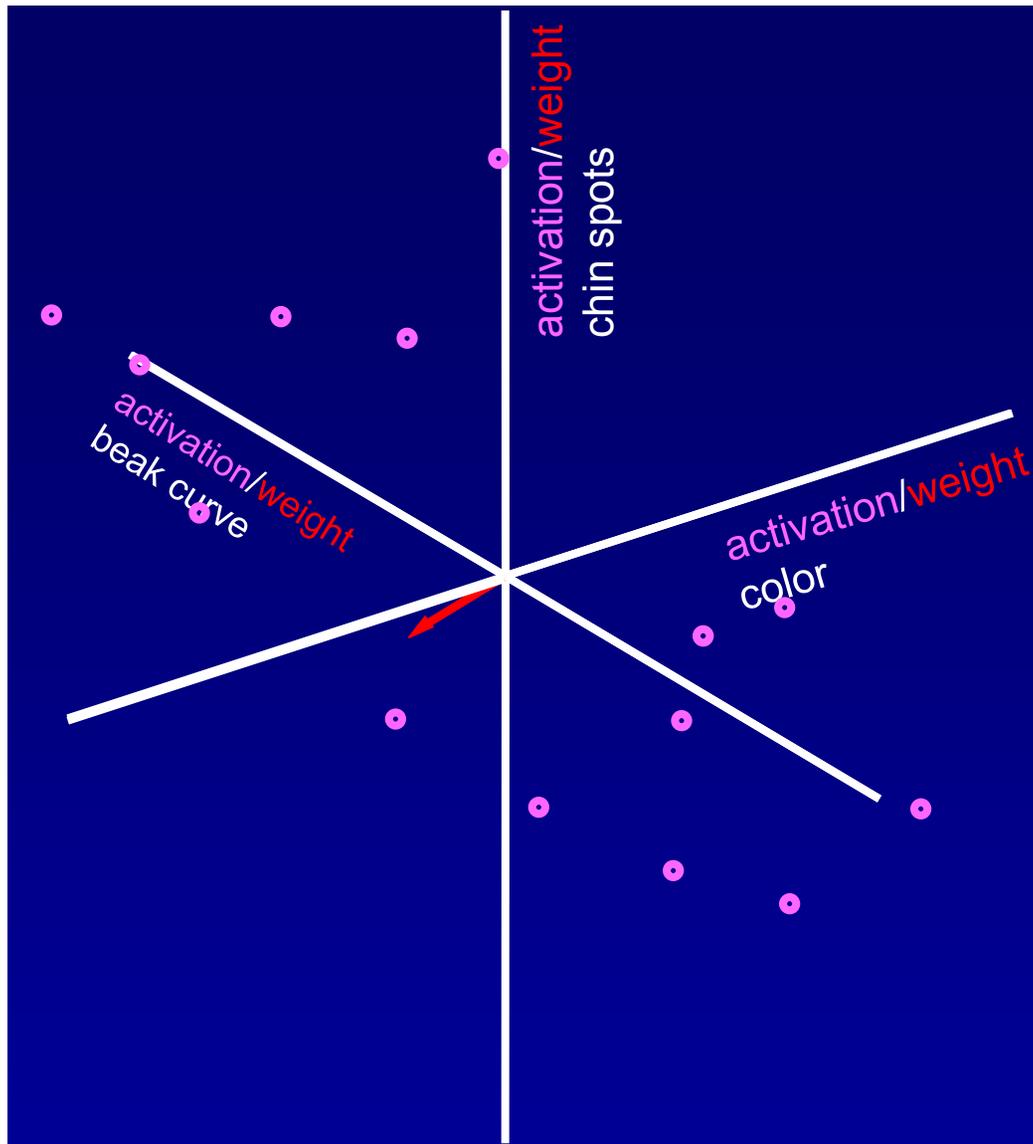
The output neuron whose weights are most similar to the input activations wins the competition.



The winning unit's input *weights* are changed to be a little more like the pattern of *activation* that made it the winner.



This means that when this pattern occurs again, the same output neuron will win (more easily). Similar input patterns will also tend to activate the same output unit.

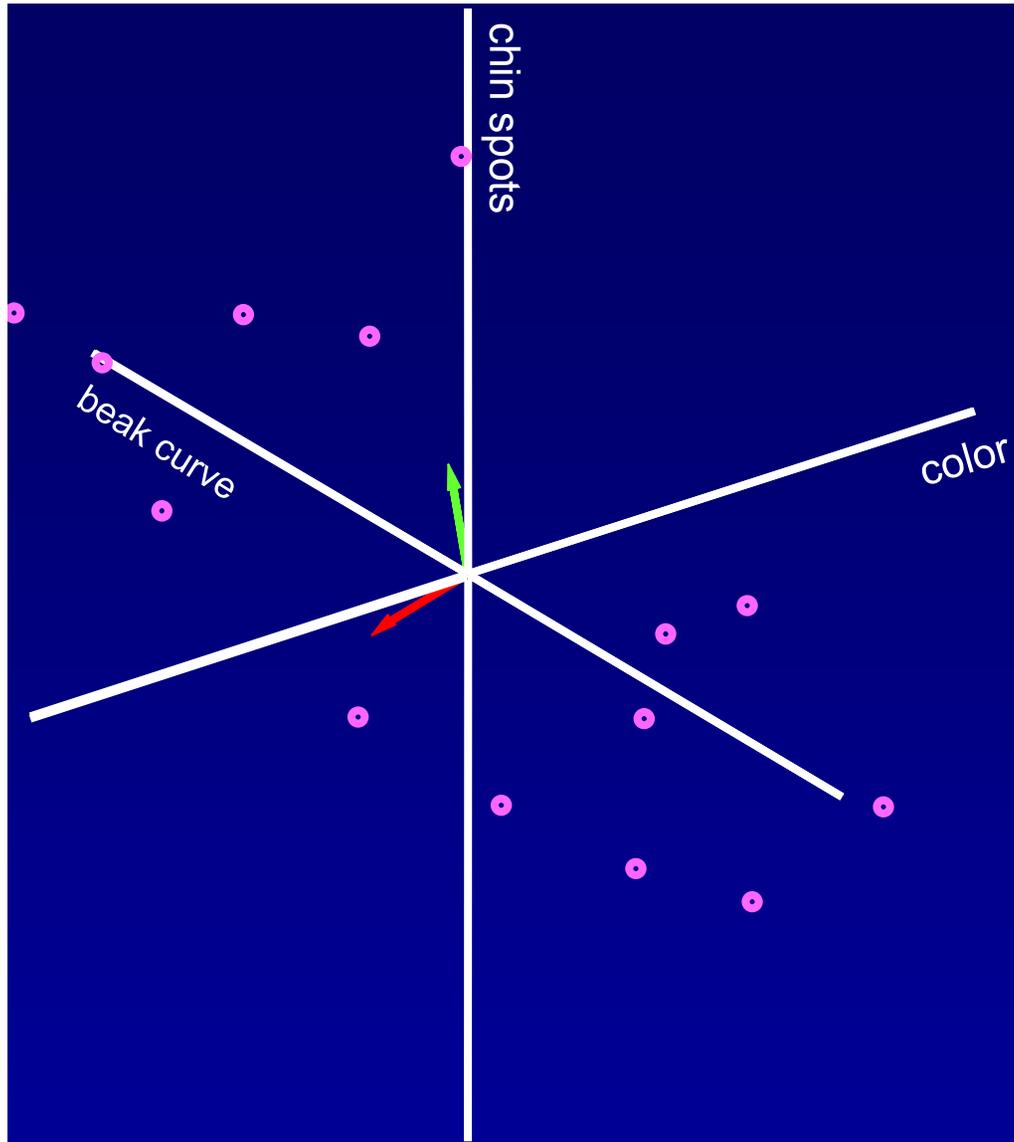


- ❖ The **arrow** represents weights going into an output unit
- ❖ The **circles** represent input activation patterns
 - ◆ As they each have the same components (color, beak curve and chin spots) we can represent them on the same graph
- ❖ The output unit's activation depends on how closely it points towards the current input pattern.

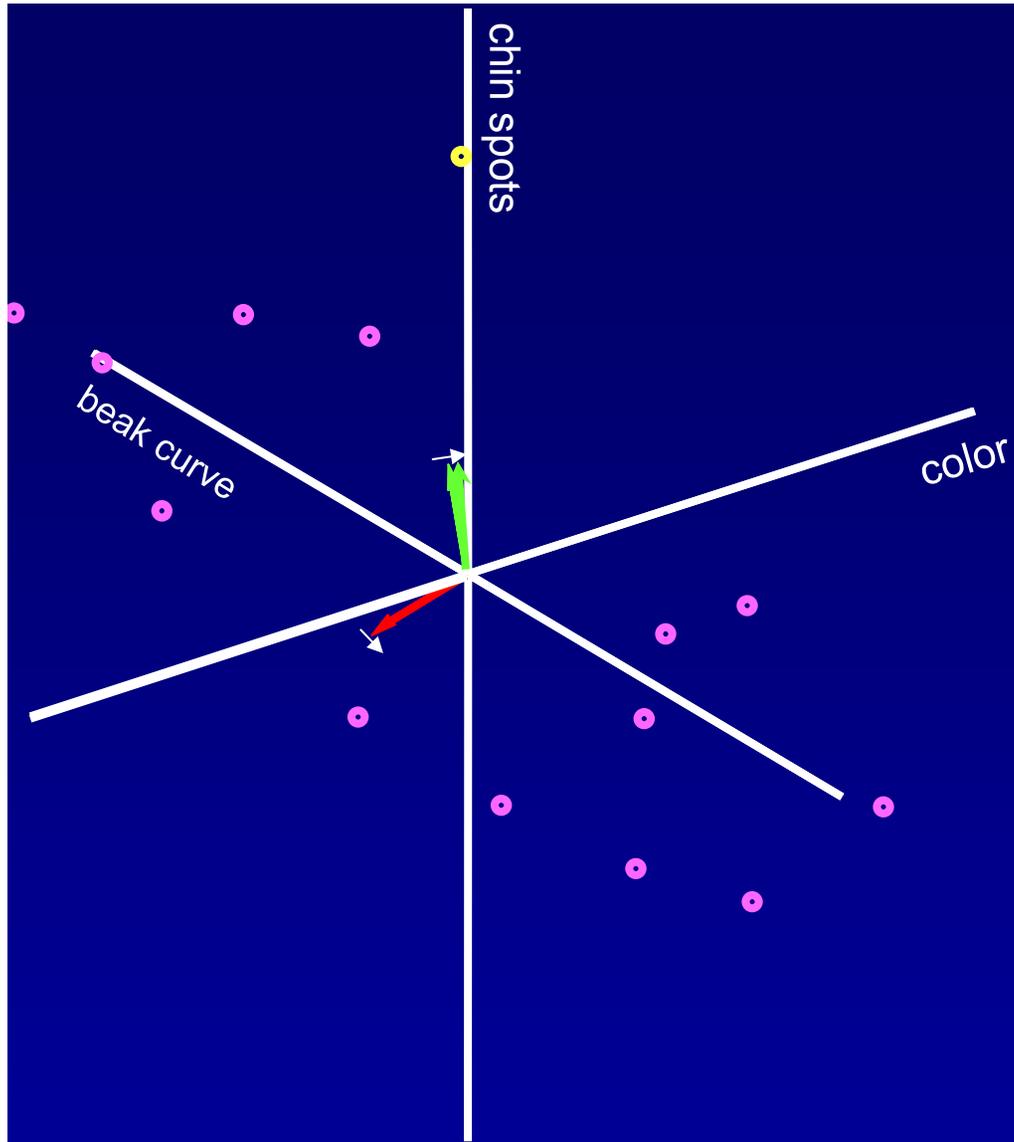
● input activation pattern for each bird

← weights into an output unit

- ❖ We need more than one output unit corresponding to the number of clusters that we need to find.



- input activation pattern for each bird
- ← weights (input) to cluster 2
- ← weights (input) to cluster 2

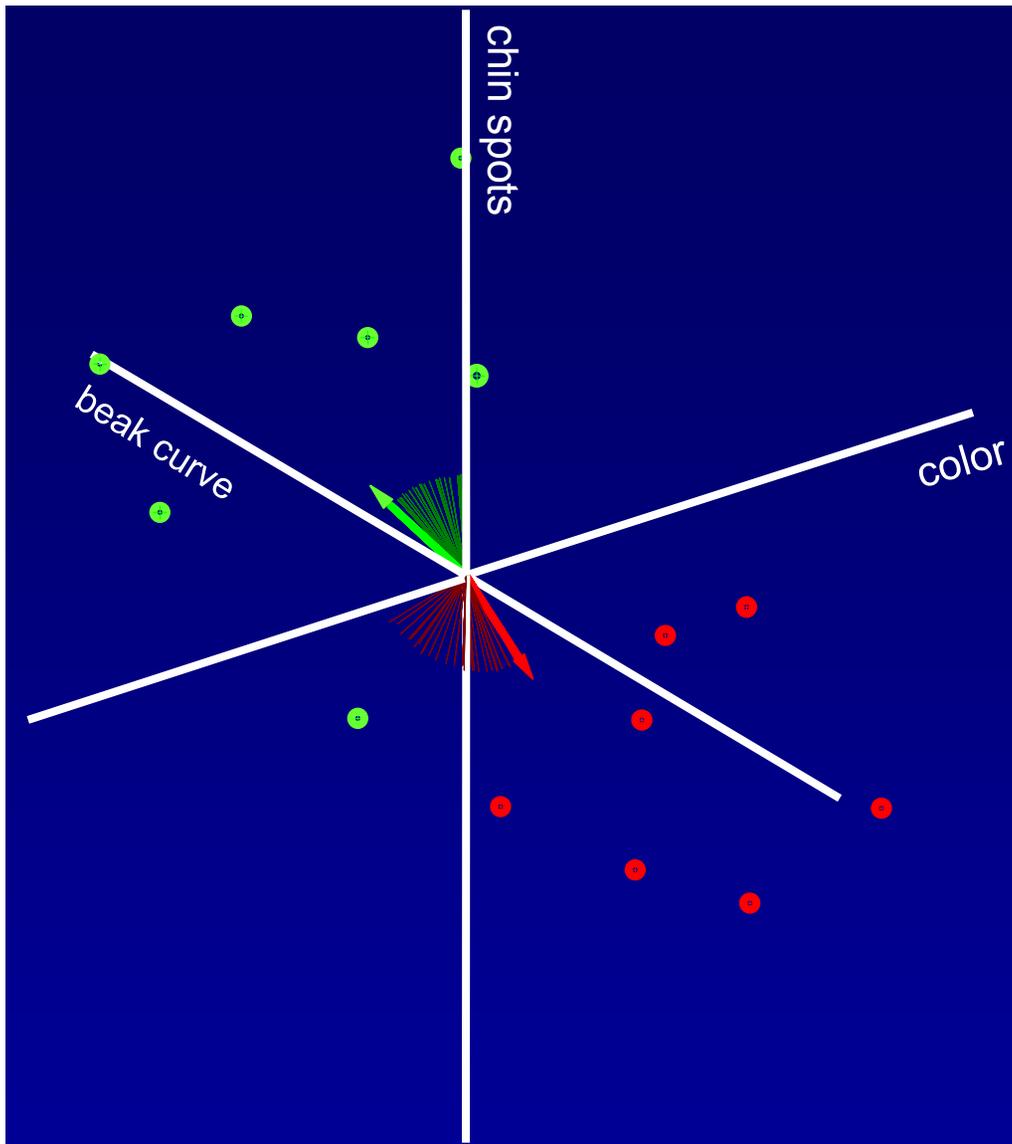


- ❖ After each input pattern the learning rule rotates the winning unit's arrow to point more nearly towards the current input pattern
- ❖ This means that when this pattern occurs again, the same output neuron will win (more easily). Similar input patterns will also tend to activate the same output unit.

● input activation pattern for each bird
 ← weights (input) to cluster 2
 ← weights (input) to cluster 2

Competitive Learning Animations

- ❖ version 1 (~1min)
- ❖ version 2 (~1min)
- ❖ version 3 (~1min)



- ❖ Arrows represent weights going in to each output unit
- ❖ They stay same length (because total weight is limited) but increasingly point toward the centres of the clusters.
- ❖ When a given input pattern appears, the arrow which points more nearly towards it “wins”. In the example, one output neuron would learn to fire when a **budgie** is seen, the other when a **canary** is seen.

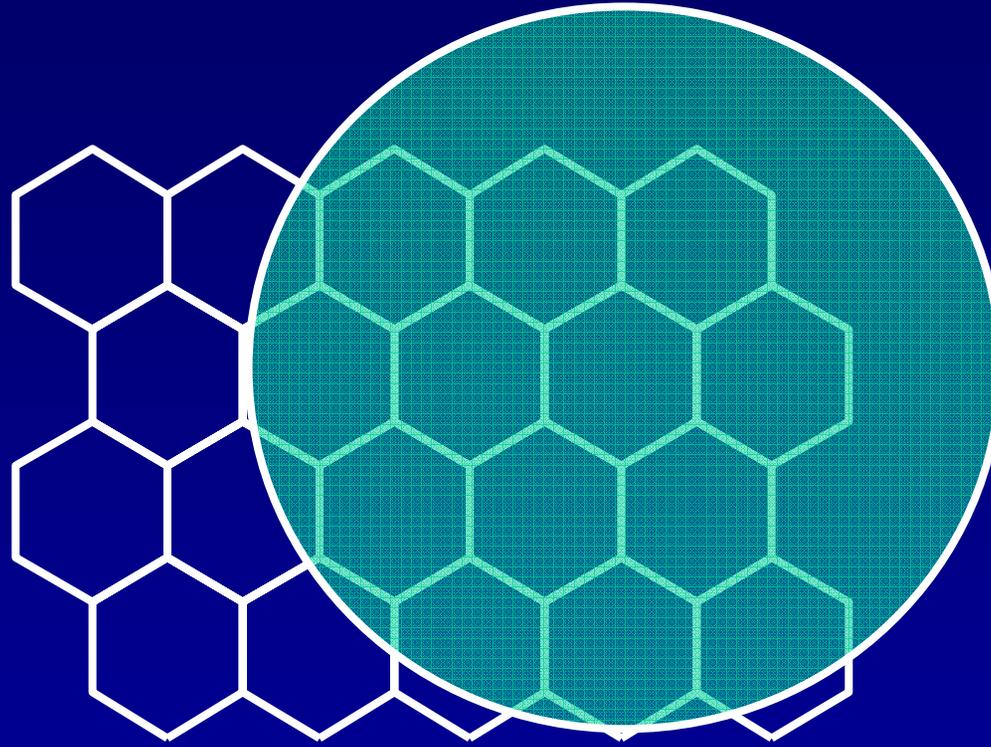
● input patterns: greatest activation in cluster 1 after learning
● input patterns: greatest activation in cluster 2 after learning
← weights to cluster 1 after learning — earlier iterations
← weights to cluster 2 after learning — earlier iterations

Self Organizing Maps

- ❖ Self Organizing Maps (Kohonen, 1982) are an extension of competitive learning which
 - ◆ takes into account the topology of the neurons
 - ◆ spontaneously develops a map of the input patterns

Self Organizing Maps

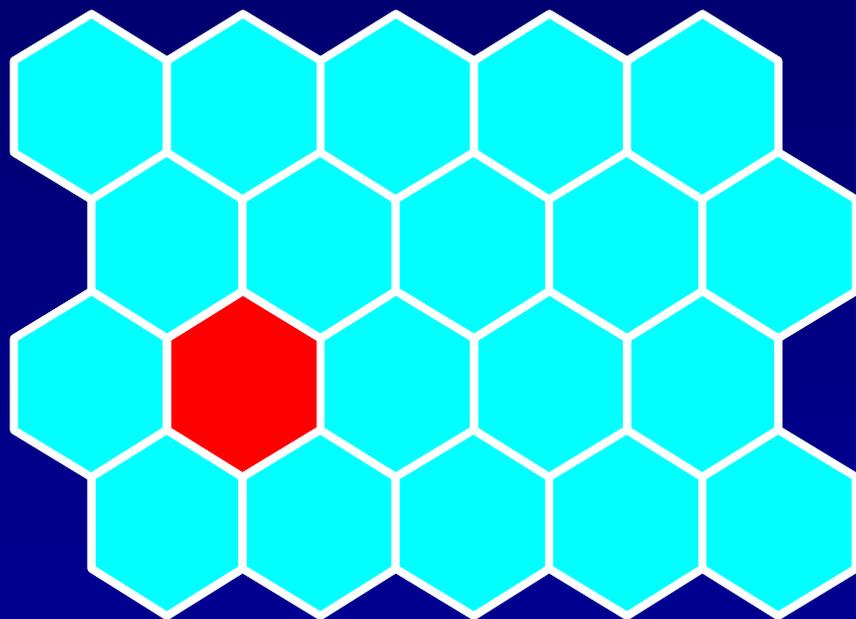
- ❖ Every neuron is defined as having a position as well as its weighted connections from the input.
- ❖ Units are typically laid out on a 2D grid (often hexagonal).
- ❖ The grid location defines a neighbourhood around each neuron (e.g., other neurons within a radius of 3 grid units)

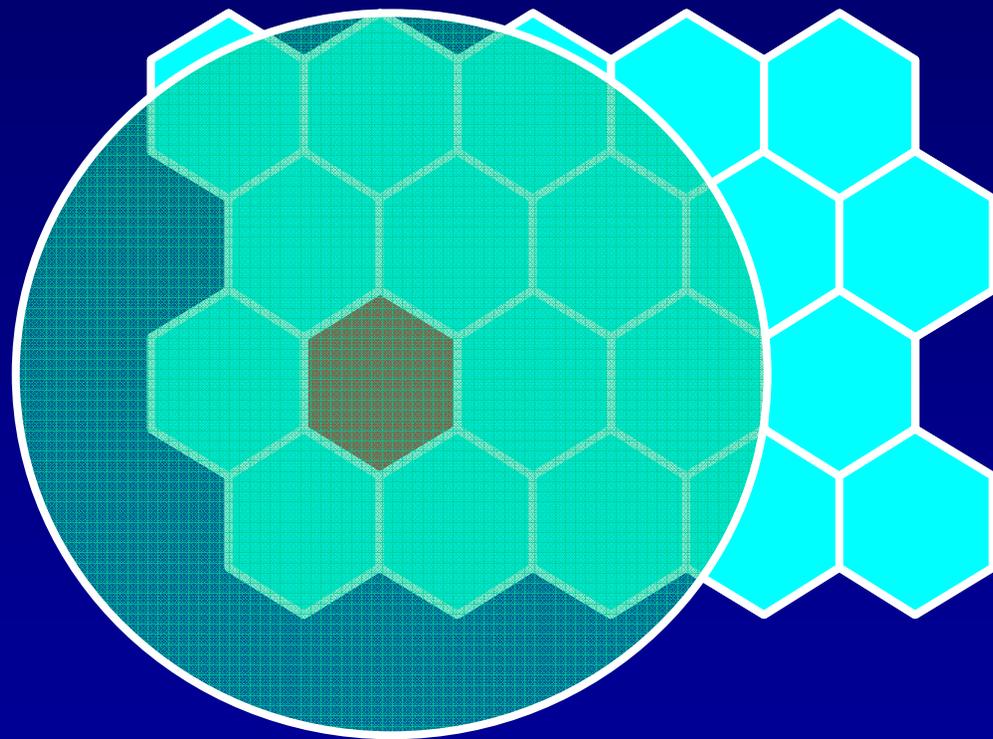


Self Organizing Maps

- ❖ Each time an input pattern is presented, neurons in the map compete, as in competitive learning.
- ❖ The winning neurons has its weights updated as in competitive learning
- ❖ *Neighbouring neurons also have their weights updated.* This step means that the topology of the network affects the learning.

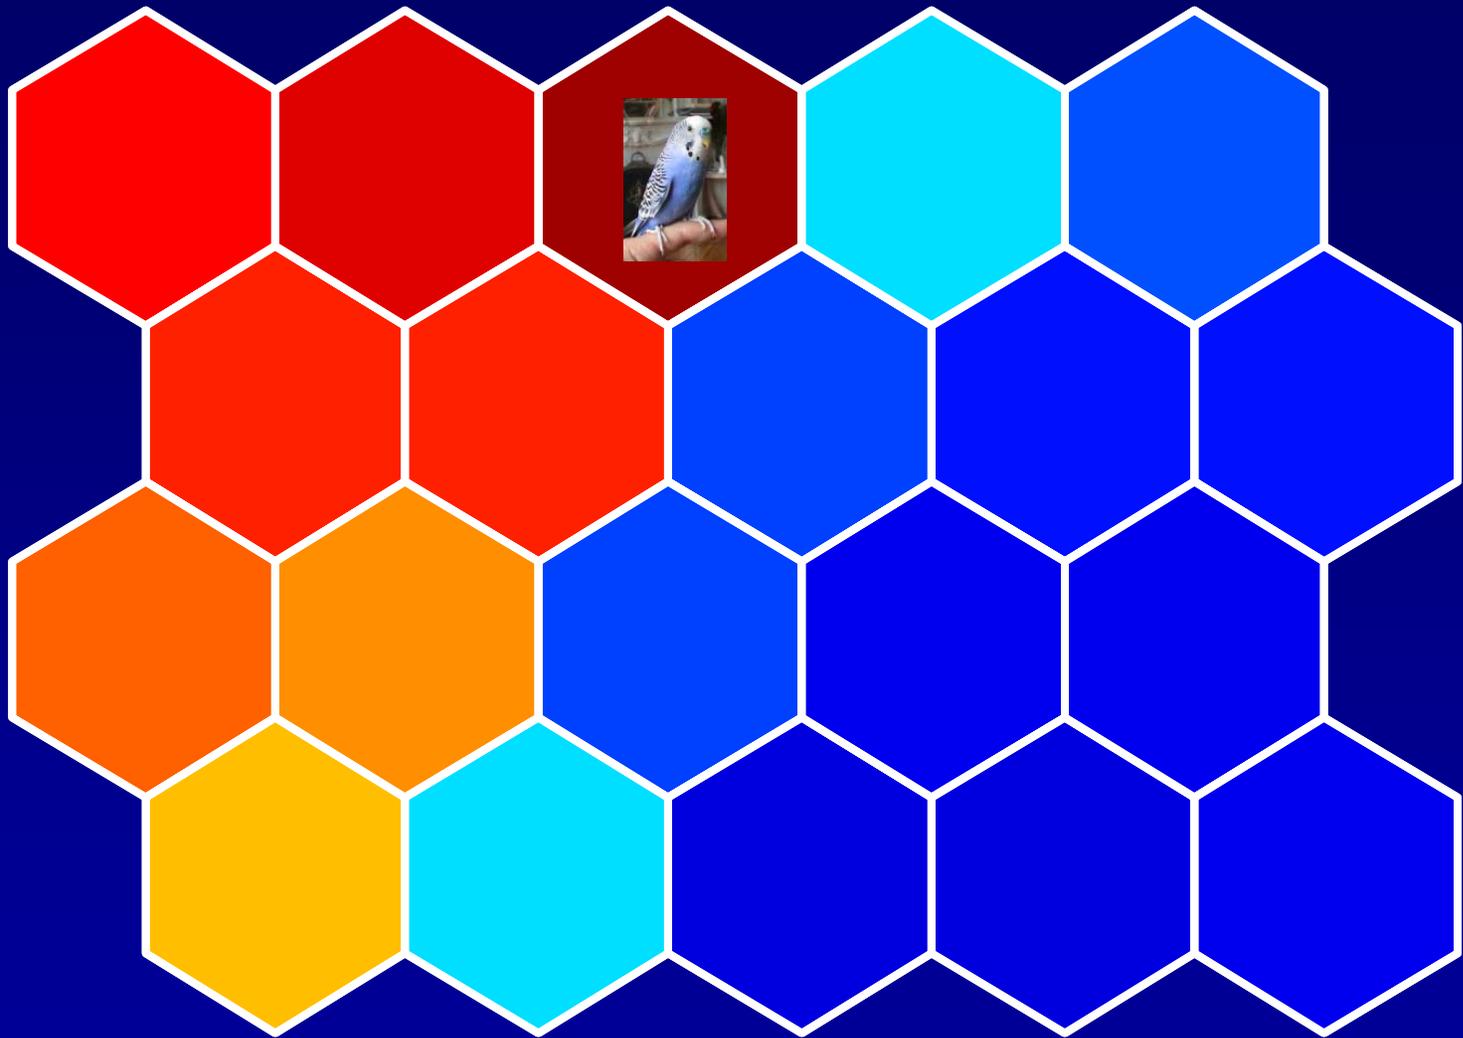






Self Organizing Maps

- ❖ Neighbouring neurons have similar weights (the weight vectors point in similar directions)
- ❖ Neurons' tunings spread evenly across the space of the input
 - ◆ Similar input patterns that occur more frequently are accorded more space in the map.
 - ◆ Unusual patterns are accorded less space





Self Organizing Maps

❖ [Video](#)



