Cognitive Modelling Examples and Applications

Spatial Representations and Memory
Modelling Learning Effects in fMRI

Tom Hartley

t.hartley@psychology.york.ac.uk

Overview

- Two disparate examples/applications
  - Hippocampal place cells, spatial learning and memory in humans (my own work)
  - Application of classification learning to predict neural activity in fMRI (someone else's)

Egocentric Reference Frames

Allocentric Reference Frame

world-centred
independent of subject's location/viewpoint (like a map)

Allocentric Reference Frames (hippocampus)

Place Cells

(O'Keefe & Dostrovsky, 1971)

Each cell fires in a particular region of the environment ("place field").

Different cells fire at different locations.

Place field location does not depend on the presence of individual cues, but their configuration (especially geometric cues – e.g., walls).

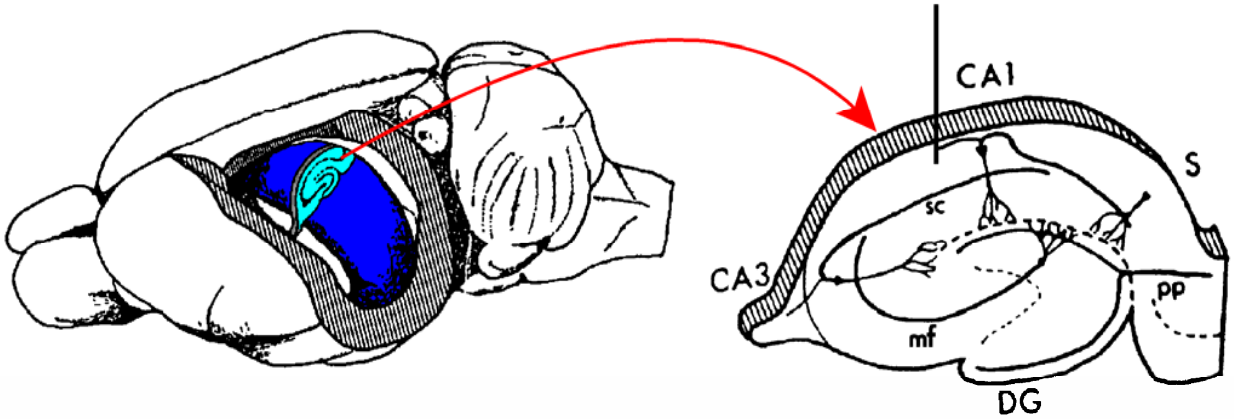

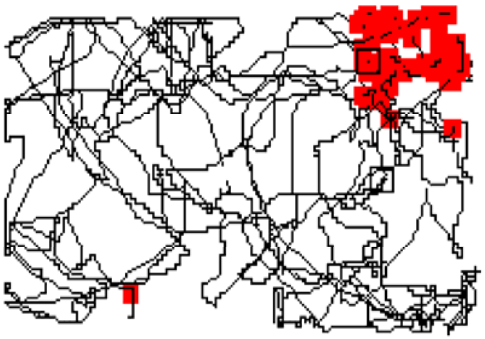

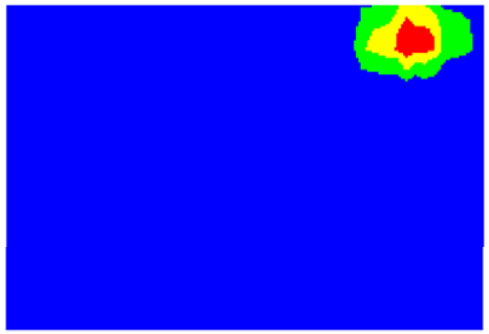

Allocentric Reference Frames (grid cells)

Grid cells

(Hafting et al, 2005)

Each cell has a grid-like firing field.

Different orientations and scales of grid are found in different cells. Scale appears to be organized topologically.

These cells are found in part of the entorhinal cortex inputting directly to the hippocampus.

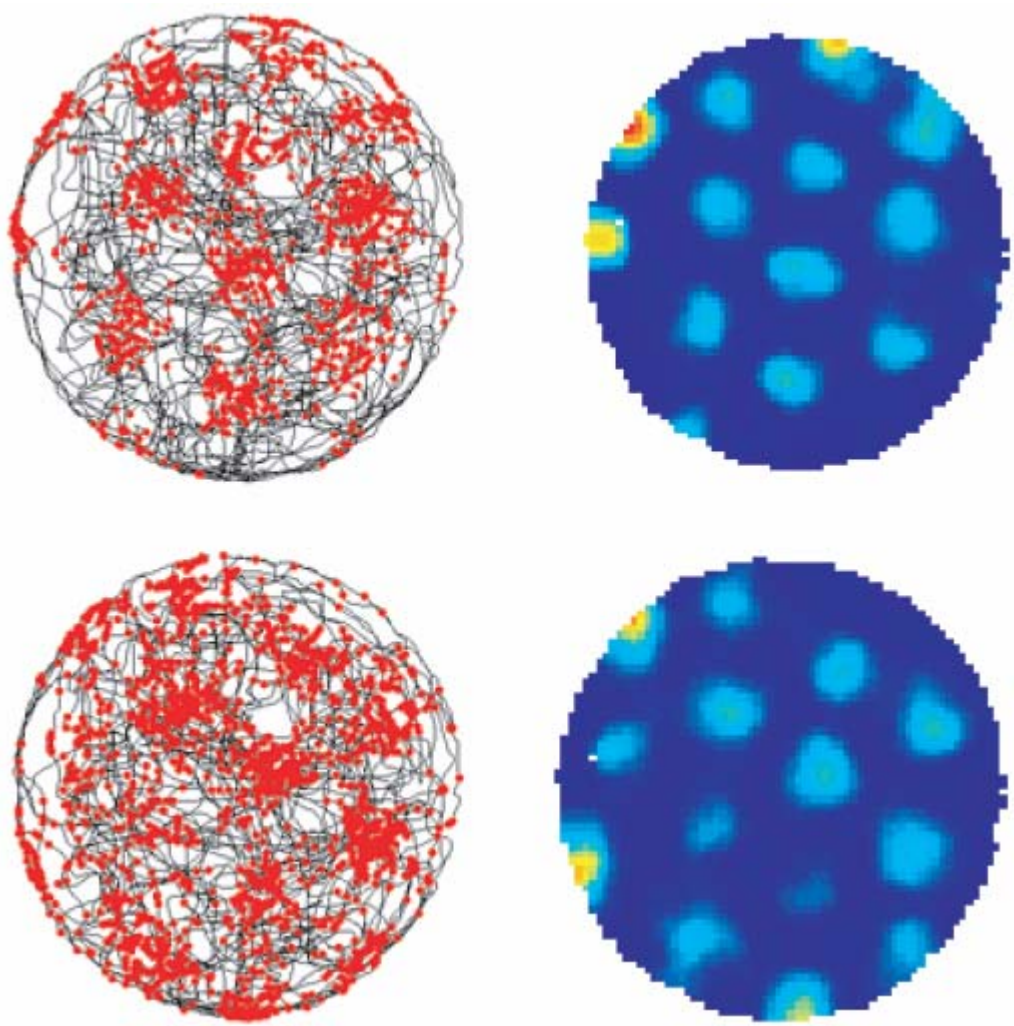

Allocentric Reference Frames (head direction system)

Head-direction cells

(Taube, Muller & Ranck, 1990)

Each cell fires when the animal is heading in a particular direction, regardless of where it is located.

Together the head direction cells act like a neural compass.

Changes in the orientation of the head direction system affect place fields.

Found in several brain regions including the subicular complex which is part of the hippocampal formation.

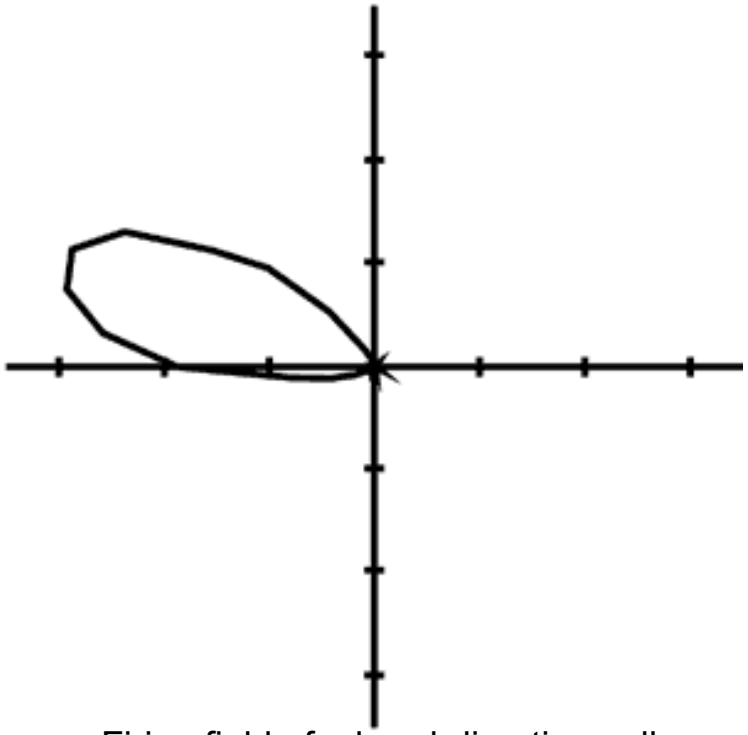

Firing field of a head-direction cell (source J.Knerim website). Radius corresponds to firing rate, orientation indicates heading.

Morris Watermaze

Morris (1981)
Morris *et al.* (1981)
Morris *et al.* (1990)

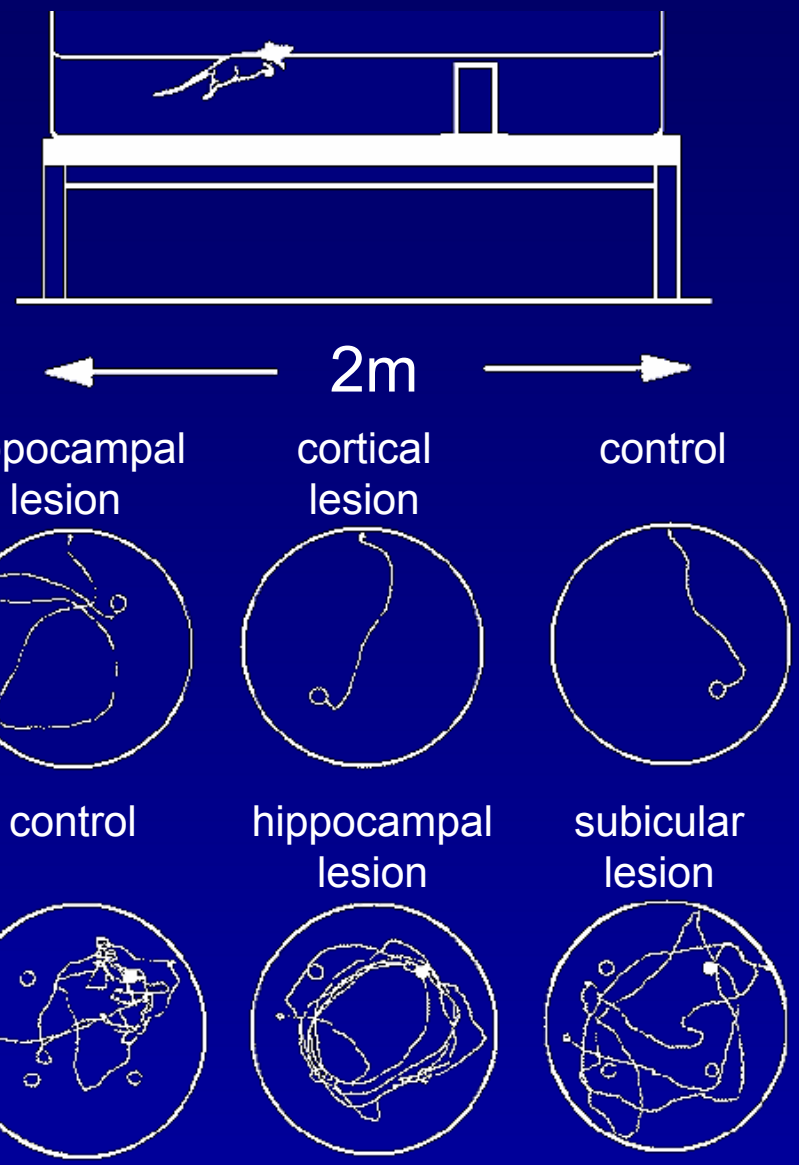


NEW
NEWS

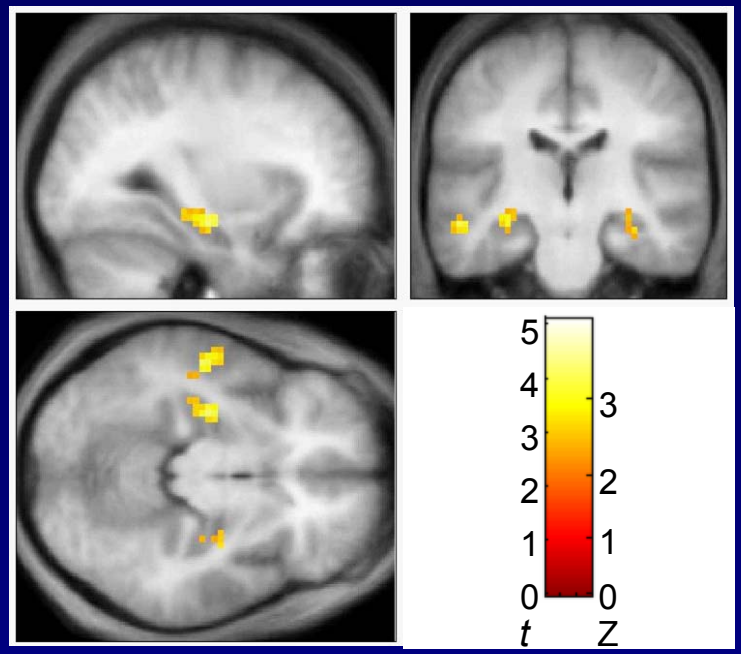
5
4
3
2
1
0
t
3
2
1
0
Z


1
2
3
0
10
20
30
40
50
60
virtual meters

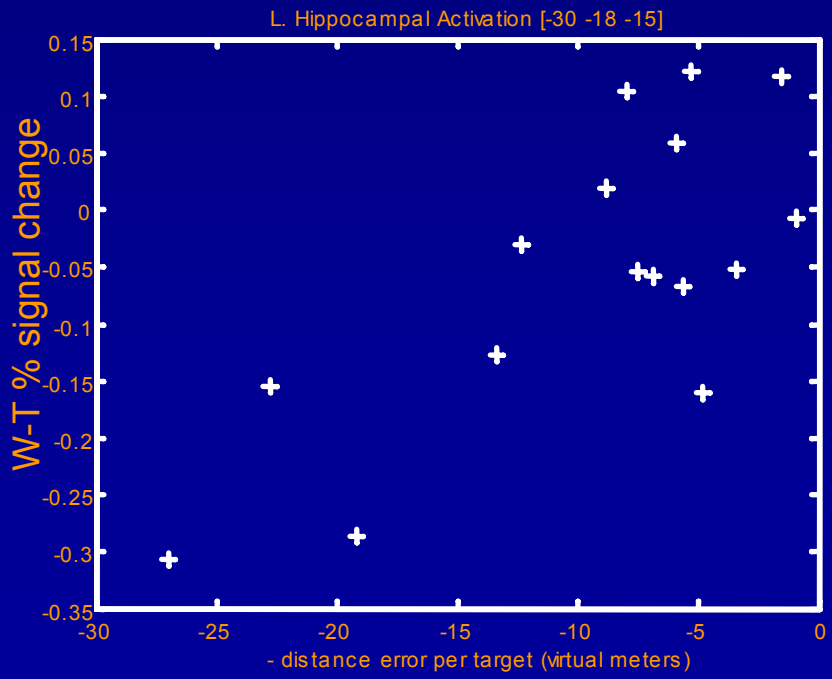
L. Hippocampal Activation [-30 -18 -15]
W-T % signal change
0.15
0.1
0.05
0
-0.05
-0.1
-0.15
-0.2
-0.25
-0.3
-0.35
-30
-25
-20
-15
-10
-5
0
- distance error per target (virtual meters)

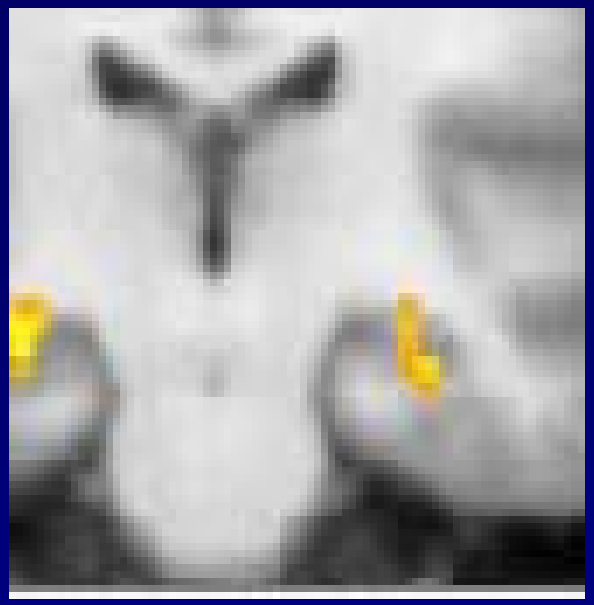

Performance-related activations in our fMRI navigation study (Hartley, Maguire, Spiers & Burgess, 2003)

Single unit recording in human hippocampus (Staba, Wilson, Fried and Engel, 2002)

electrode

Place Cells in Humans (?) (Ekstrom et al., 2003)

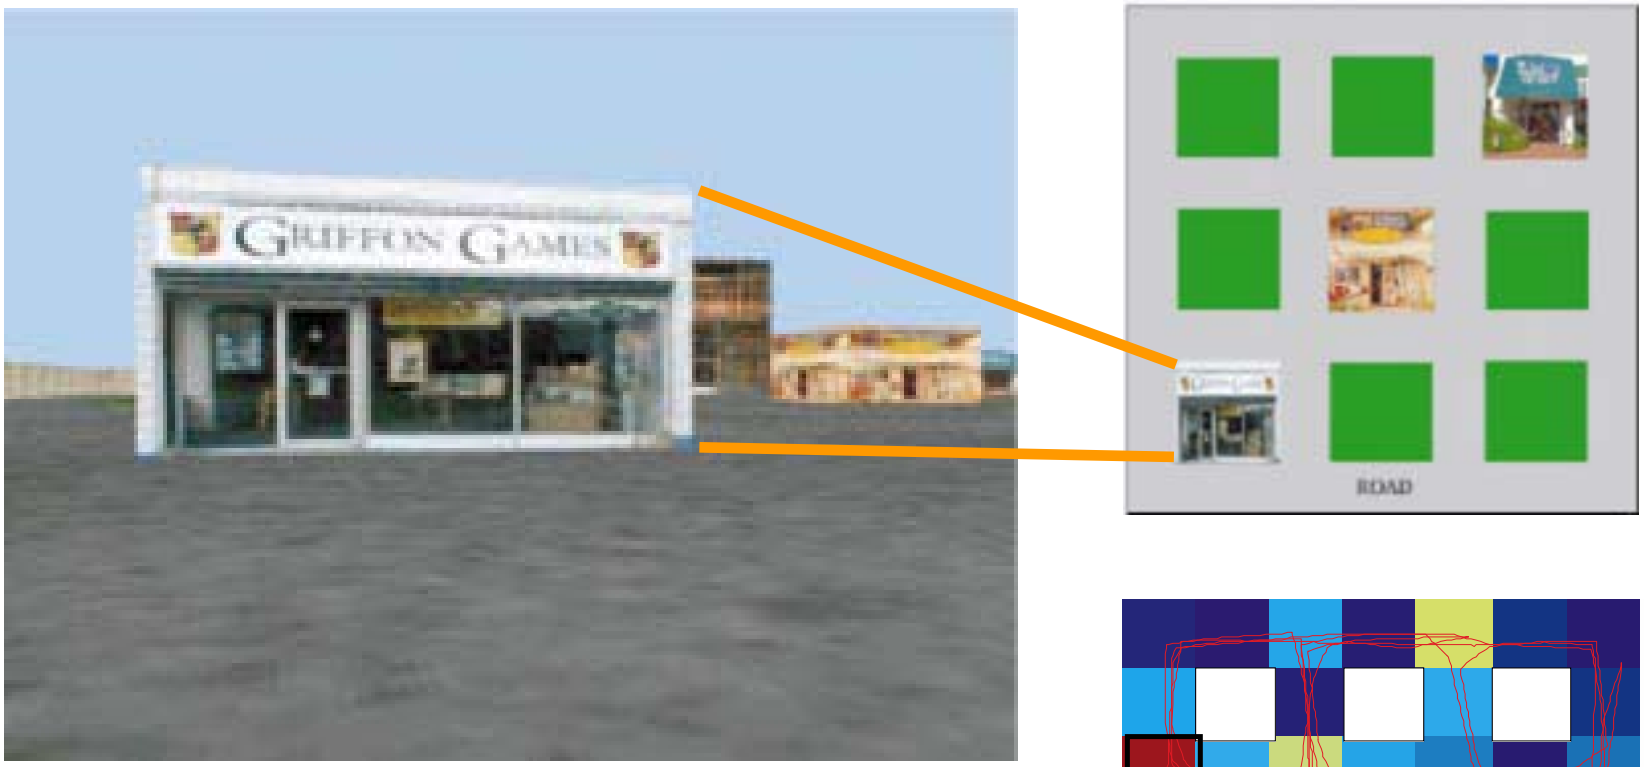

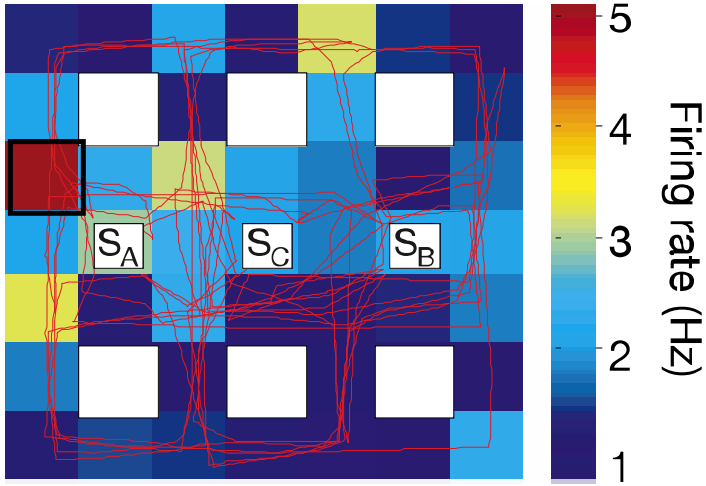

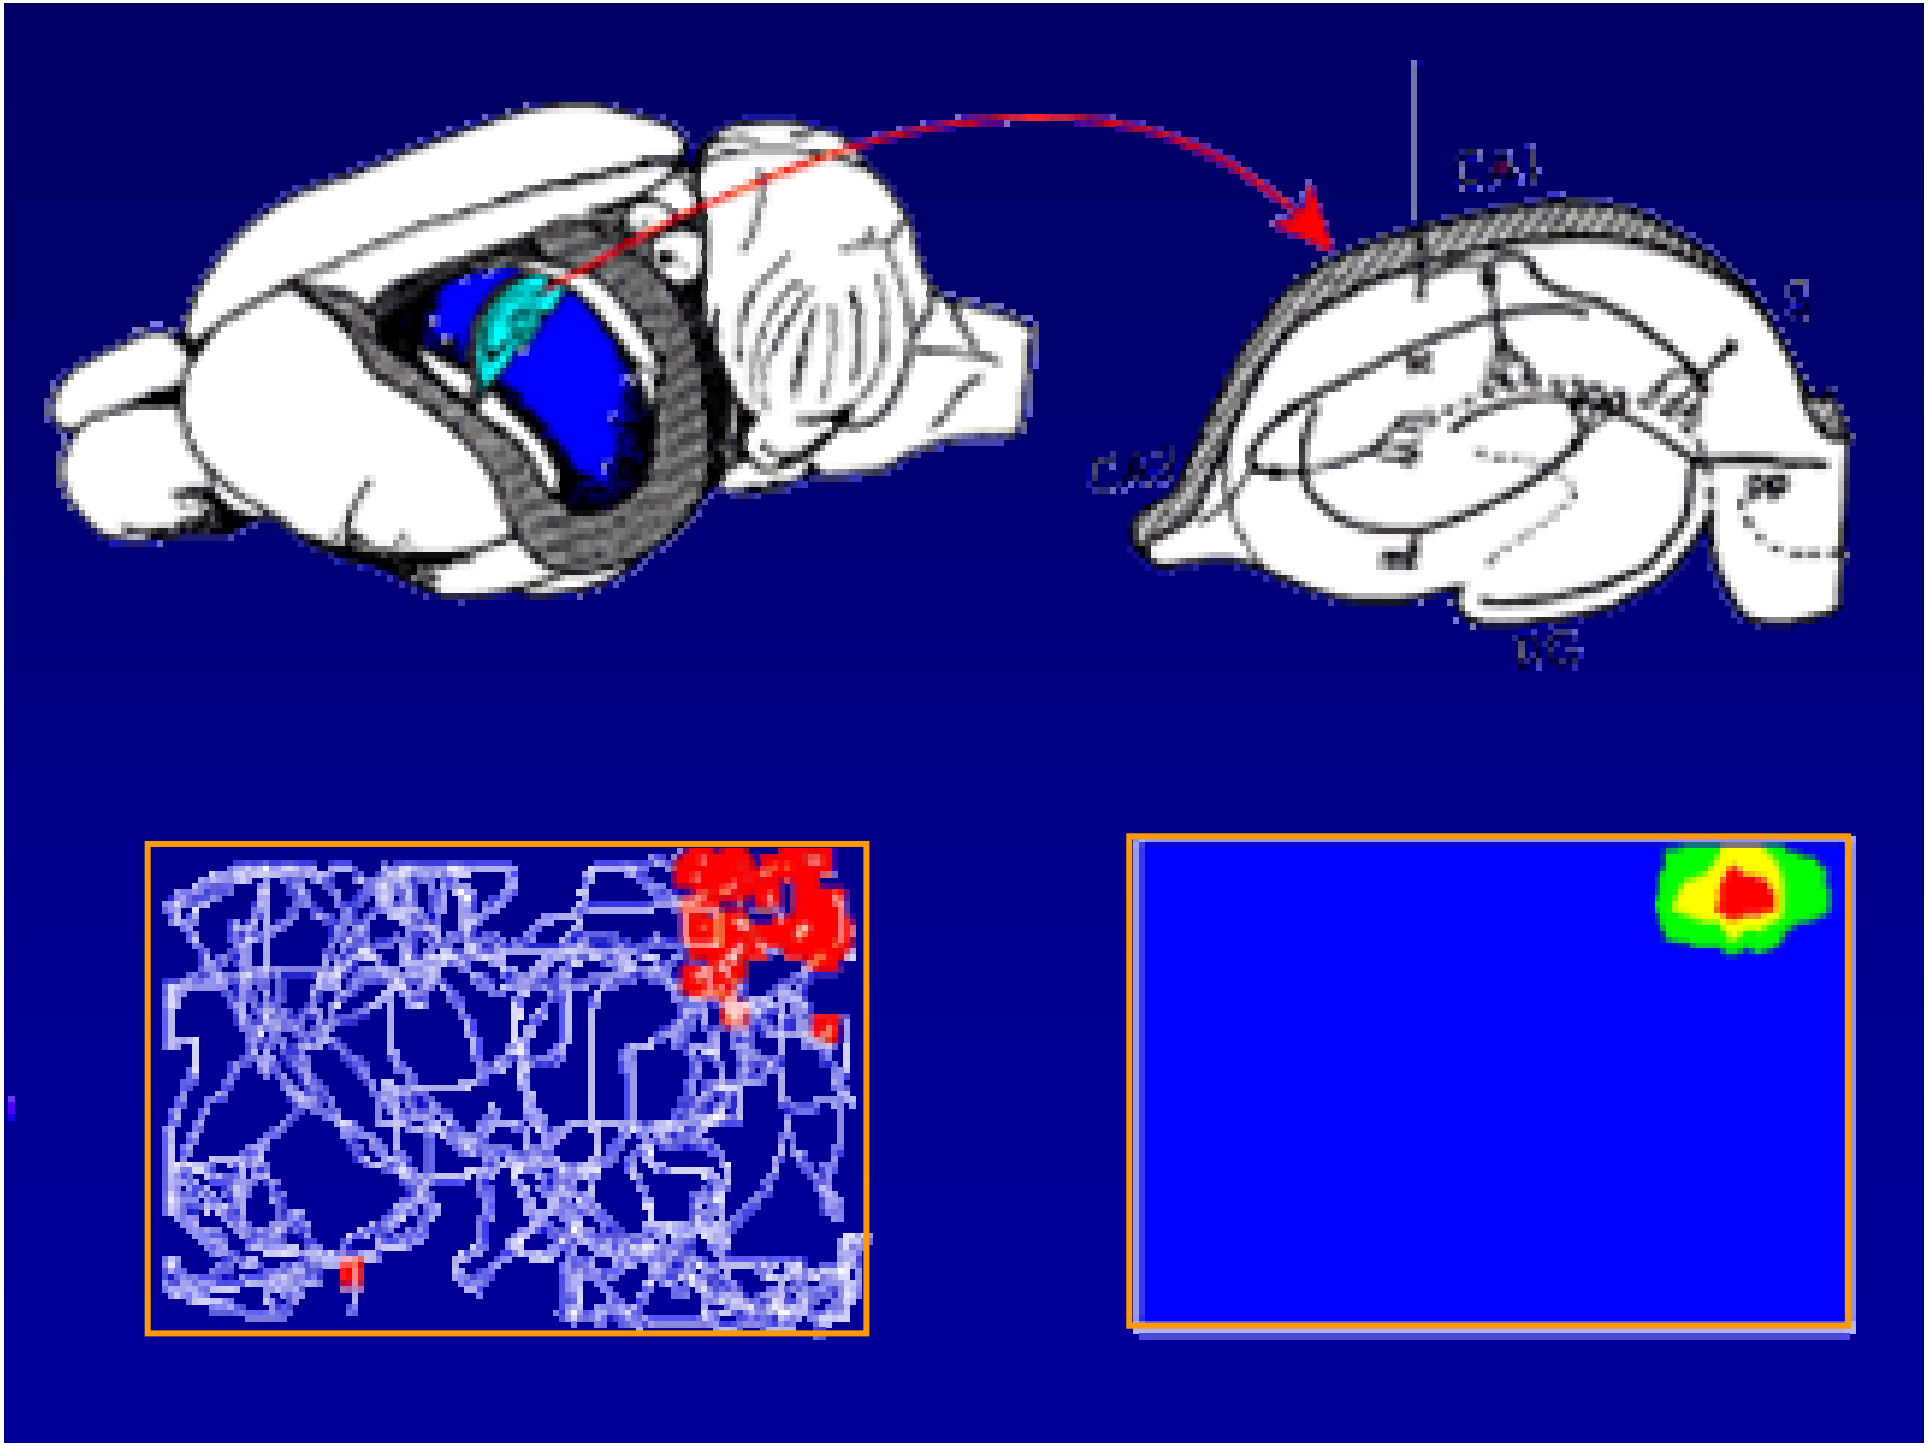
CA1
CA2
DG

Place Cells

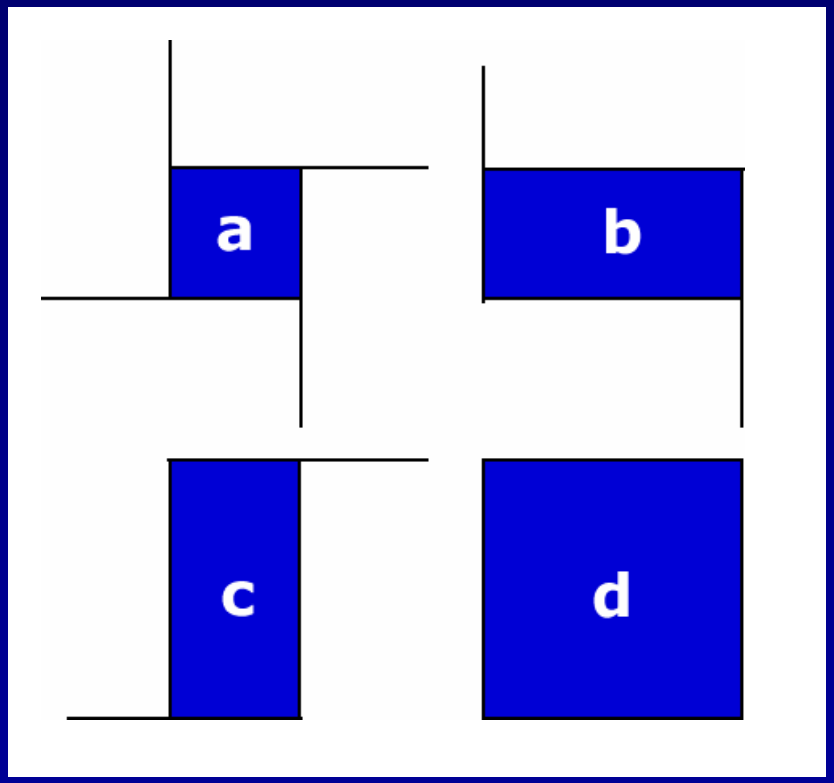

4 areas made from the same movable partitions

- The location of peak firing in different sized boxes tends to remain at fixed distances from two or more of the walls (O'Keefe & Burgess, 1996).

Place Cells

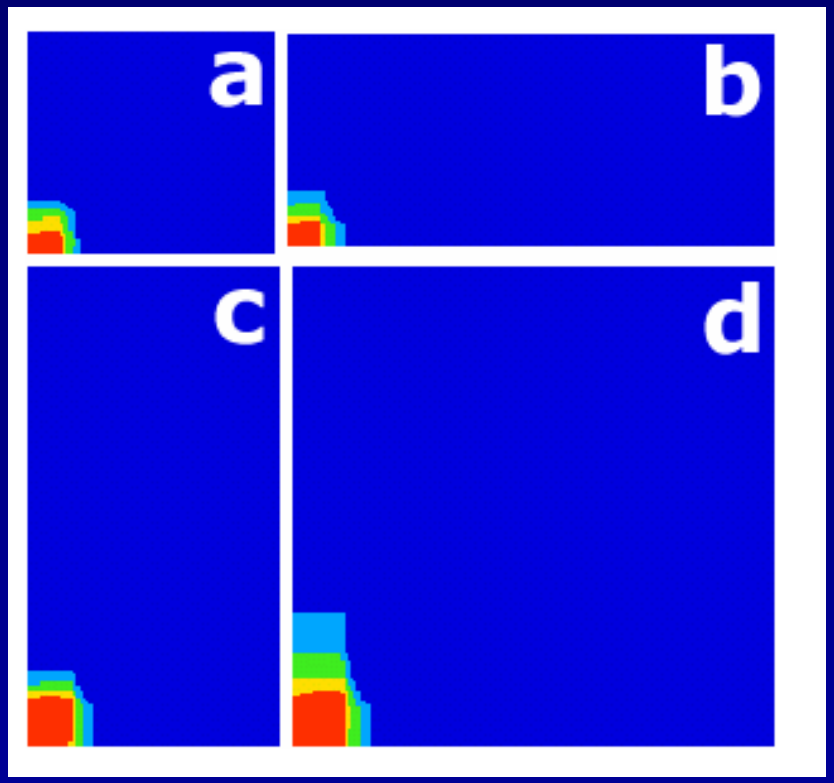

- The location of peak firing in different sized boxes tends to remain at a fixed distance from two or more of the walls (O'Keefe & Burgess, 1996).

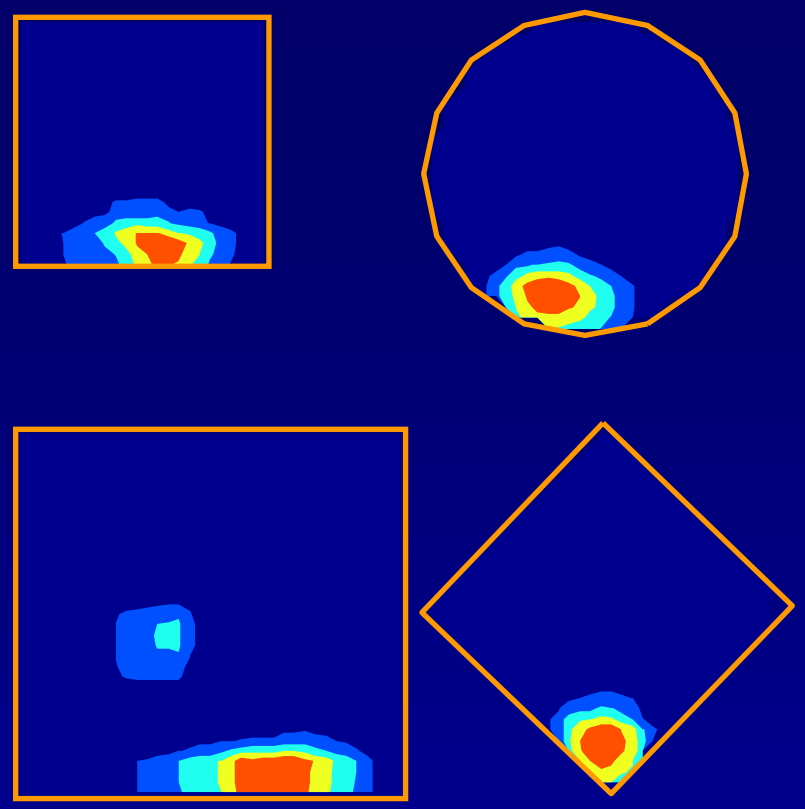

Place Cell Data (Lever *et al.,* 1999, 2002)

Phenomenon: place fields recorded from the same cell tend to “stick” to part of the boundary of the environment when it changes size or shape.

Model

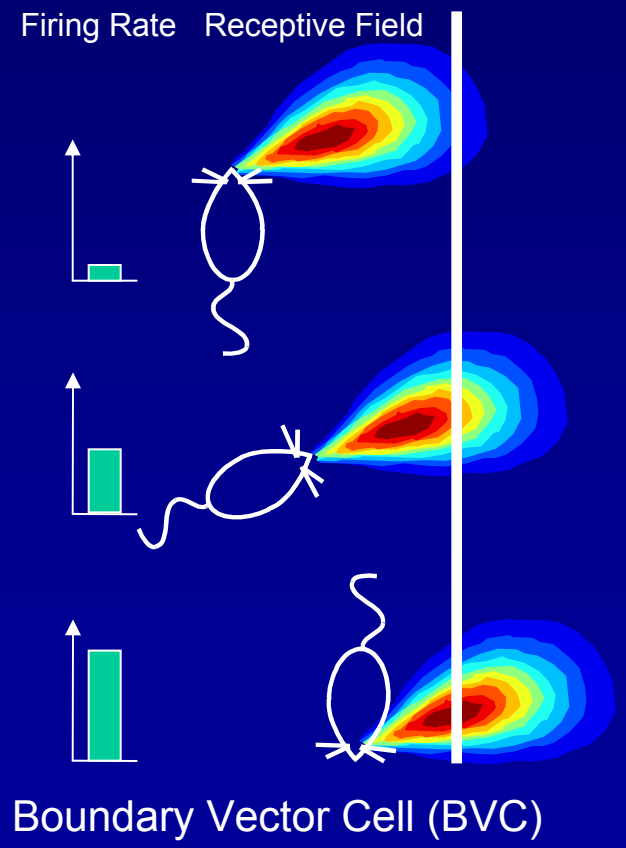

- We assume place cells receive inputs from "*boundary vector*" *cells (BVCs)* which respond to barriers in the environment.

Model

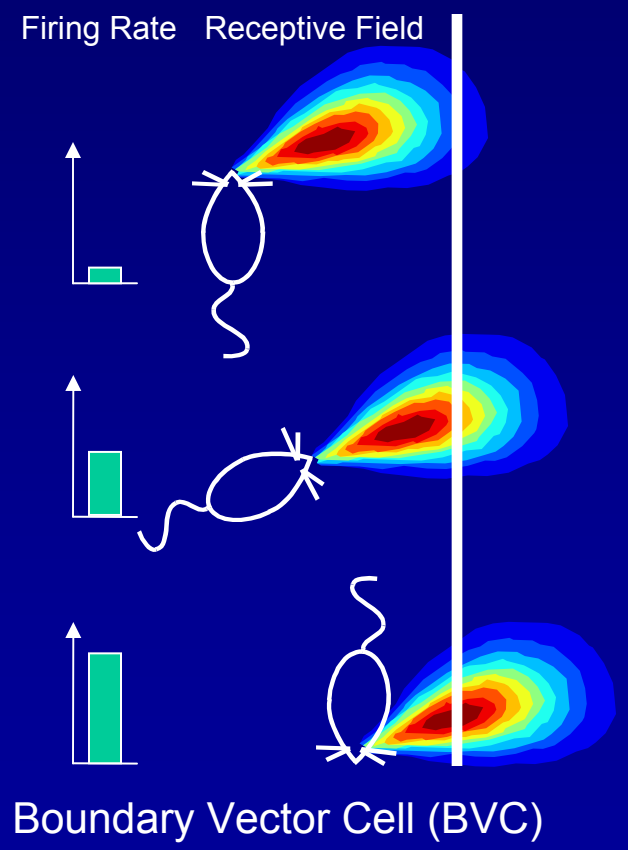

- Each BVC responds maximally when a barrier lies at a specific distance from the rat in a particular *compass* direction.

Model

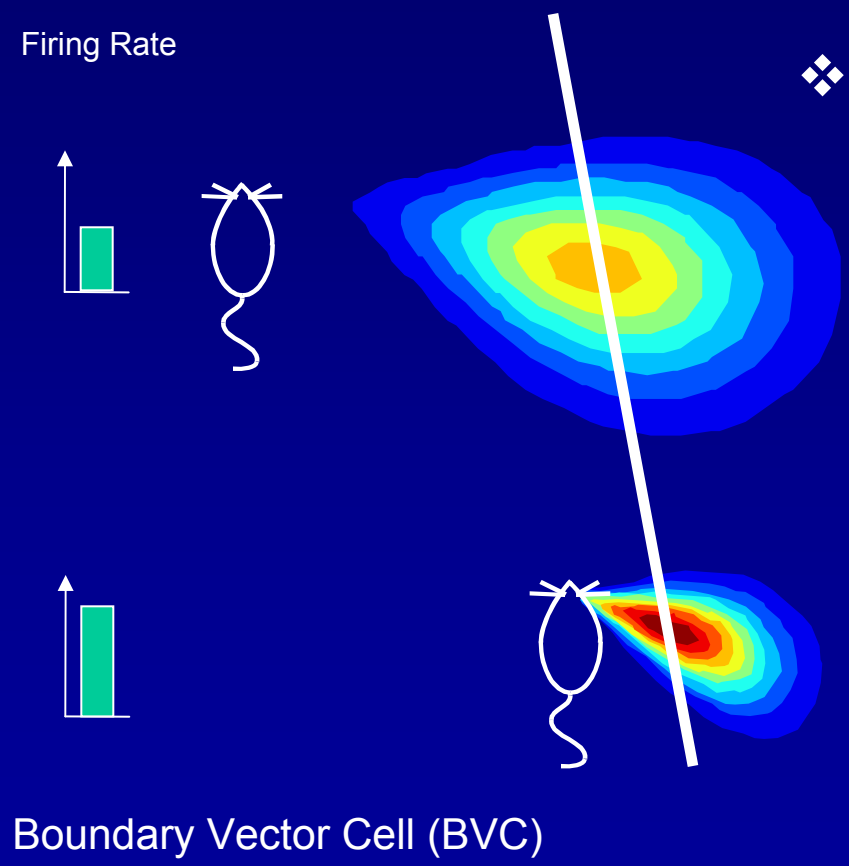

- BVCs tuned to respond to more distant walls have broader tunings

Model

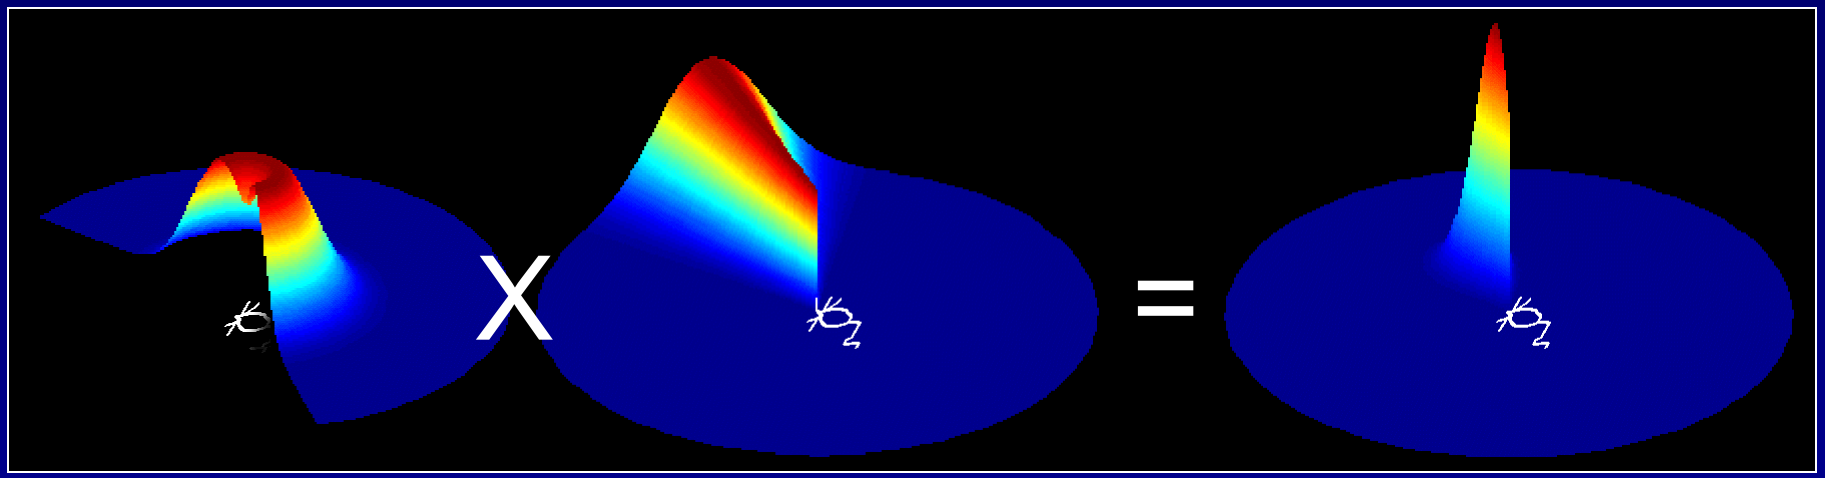

- The receptive fields of the BVC input cells are the product of Gaussian functions of distance and direction.

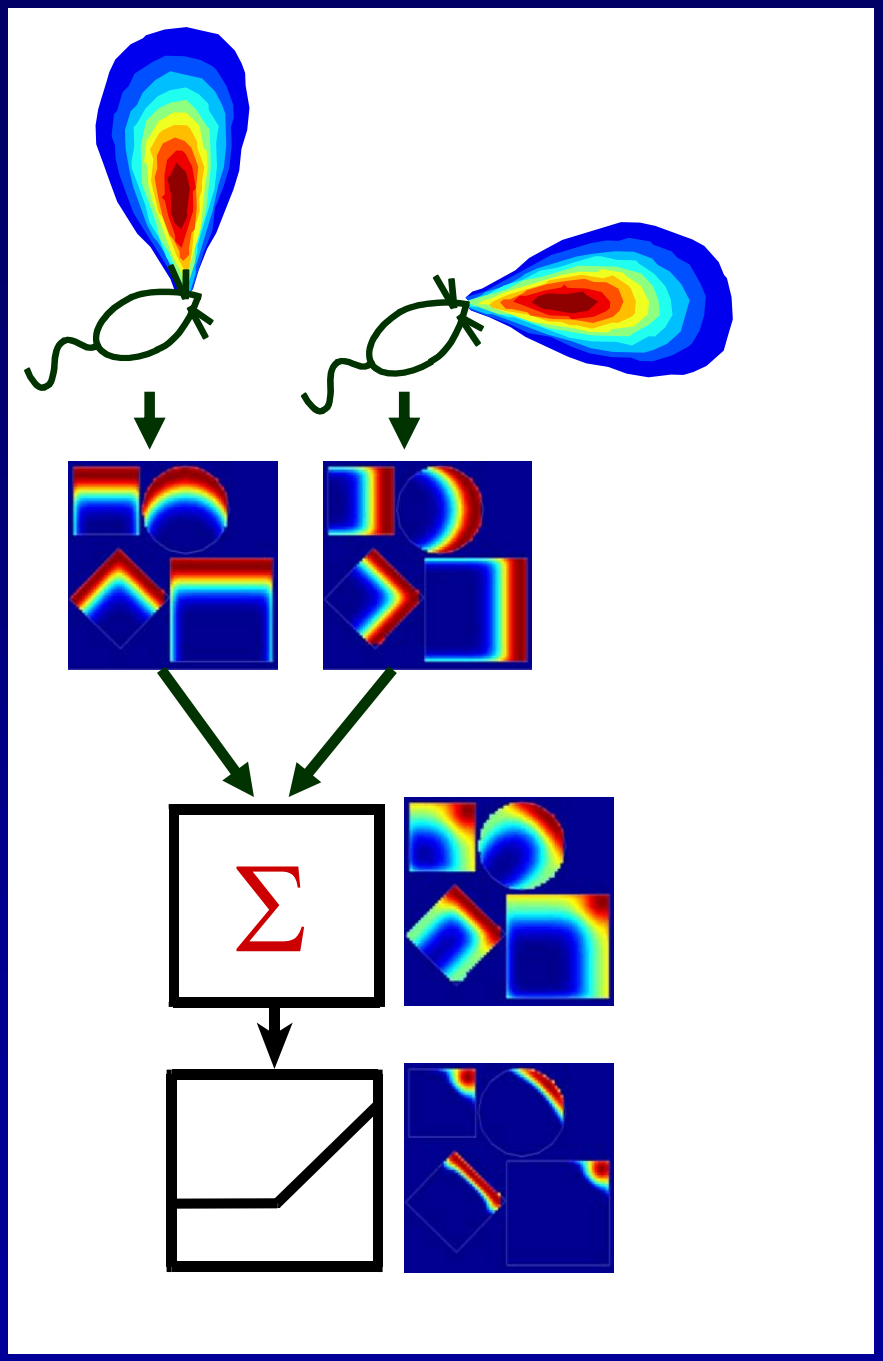

How we simulate Place Fields:

- BVCs have firing fields that follow the walls of the environment.
- Place fields are modelled as the thresholded sum of 2 or more BVC firing fields.

Fitting Place Cell Data

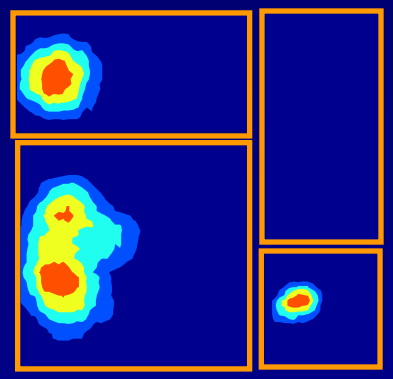

Place Cell Data (O'Keefe & Burgess, 1996)

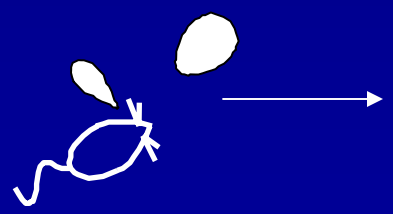

Candidate BVCs

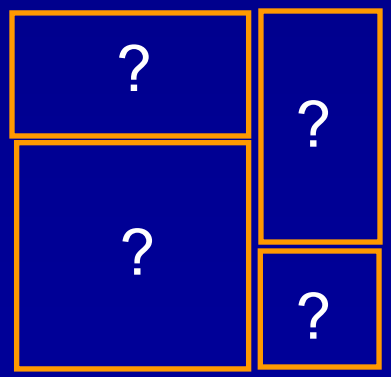

- To fit the experimental data we try to find a *small set* of BVCs which, when summed and thresholded, can account for the observed firing rate maps in different arenas.

Example

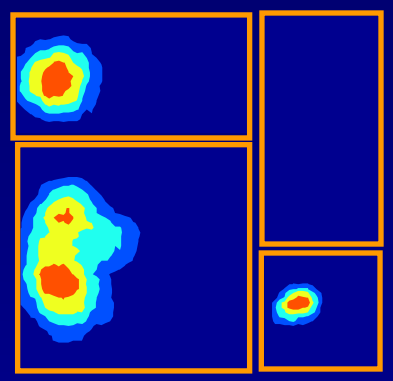

Place Cell Data (O'Keefe & Burgess, 1996)

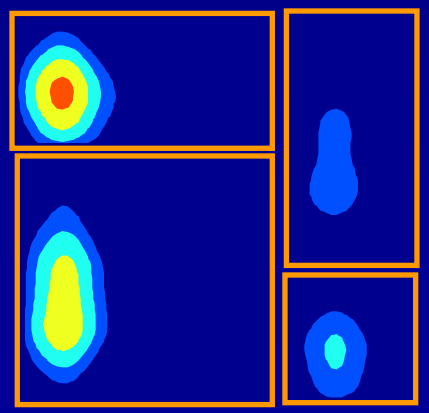

Simulated Place Fields

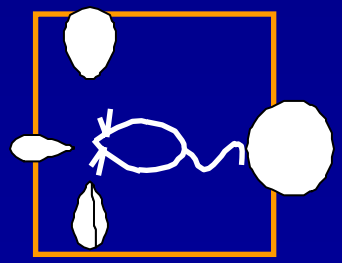

Best fitting BVCs

- 2-4 BVCs orientated at right angles to one another are sufficient to fit most fields satisfactorily.

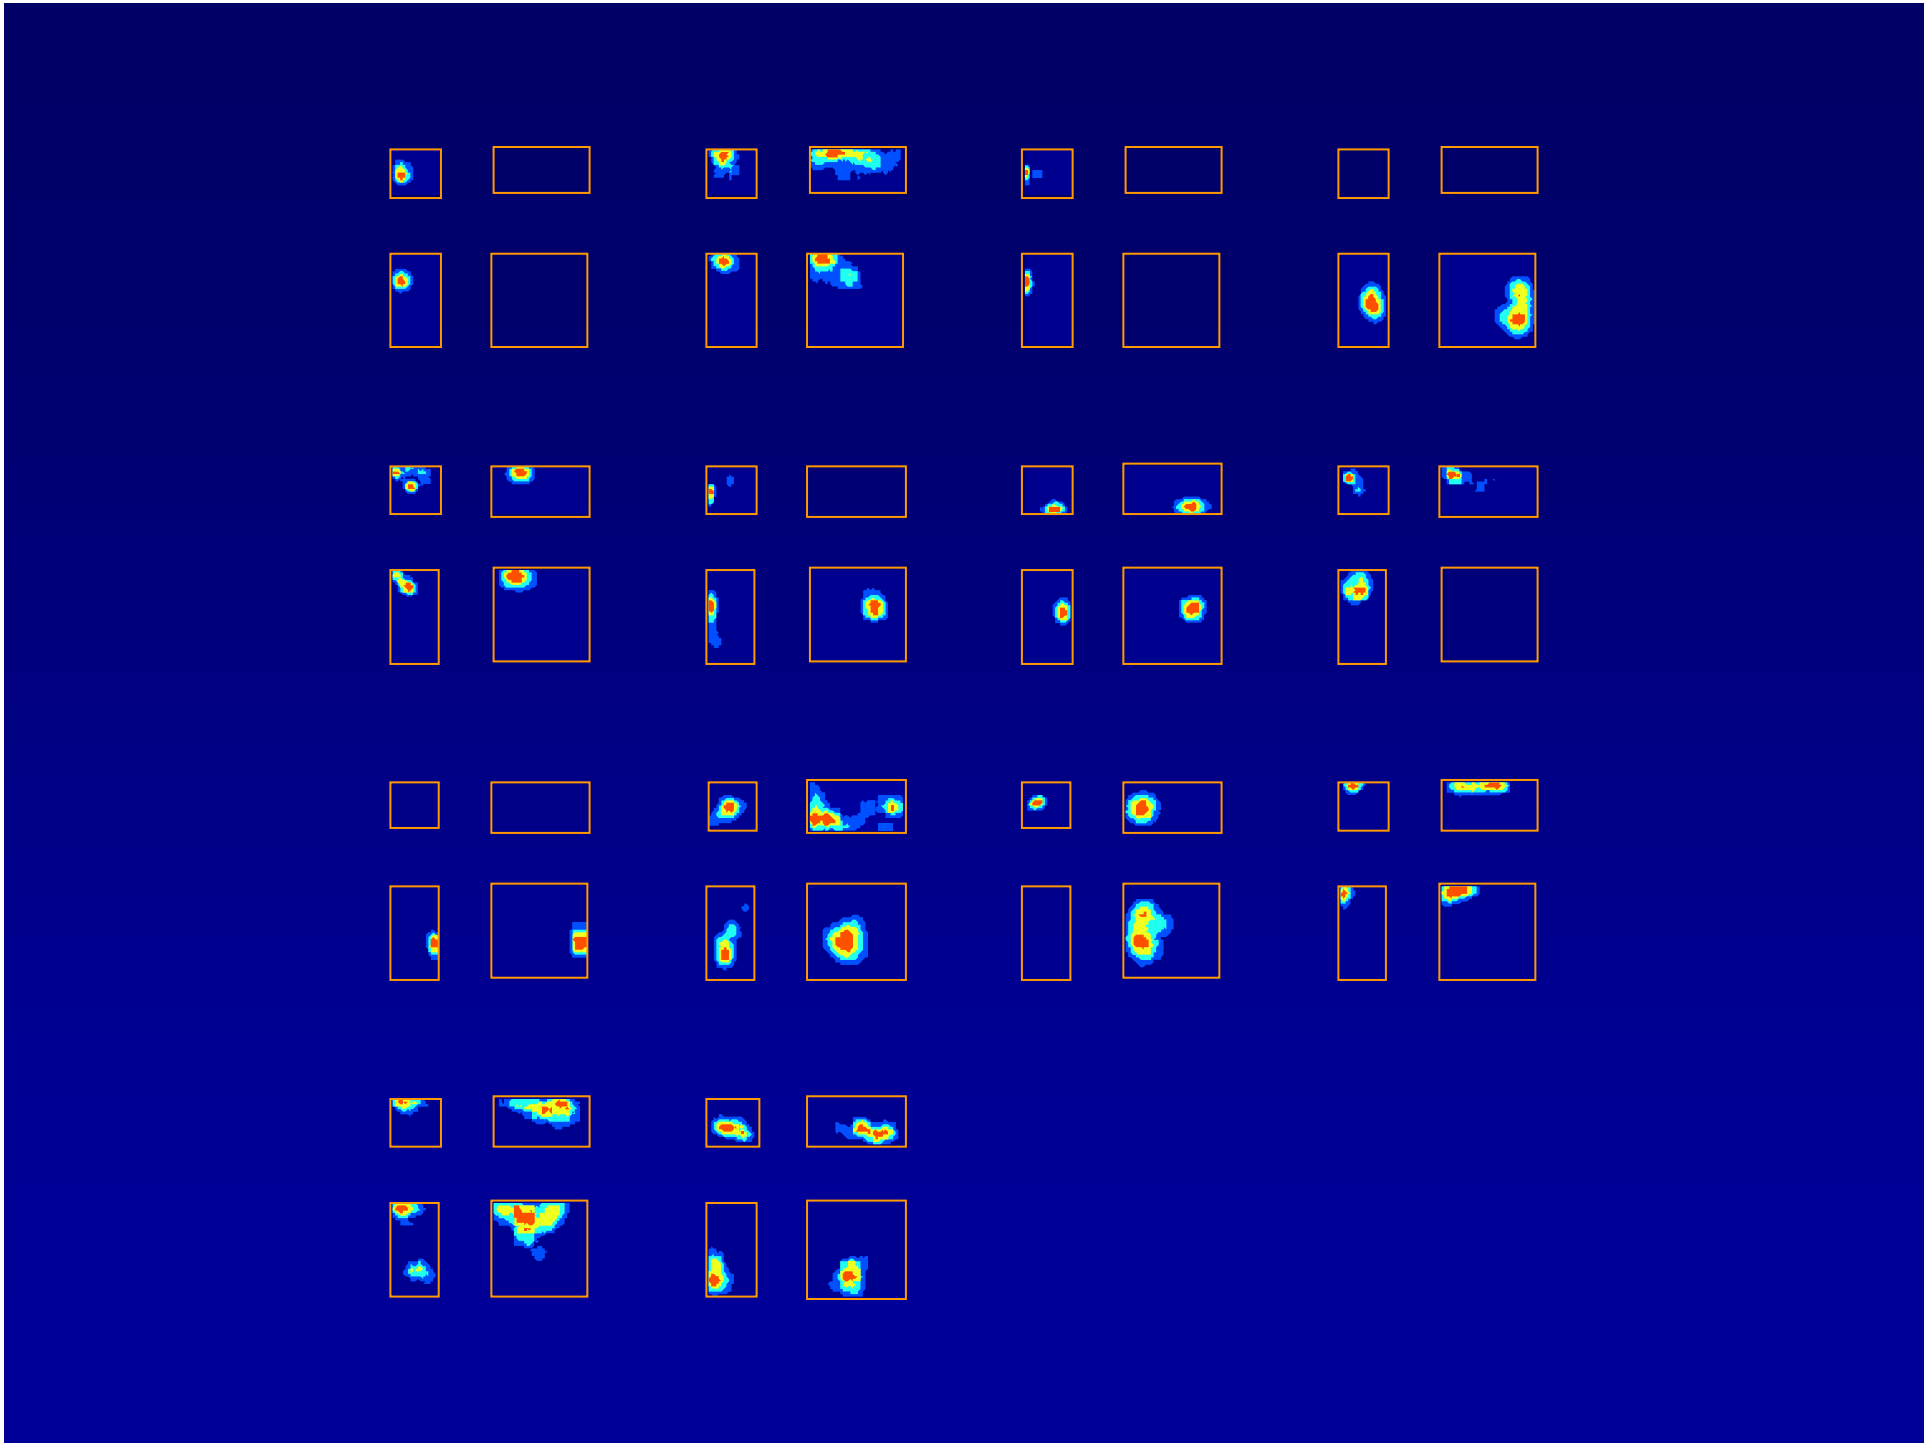

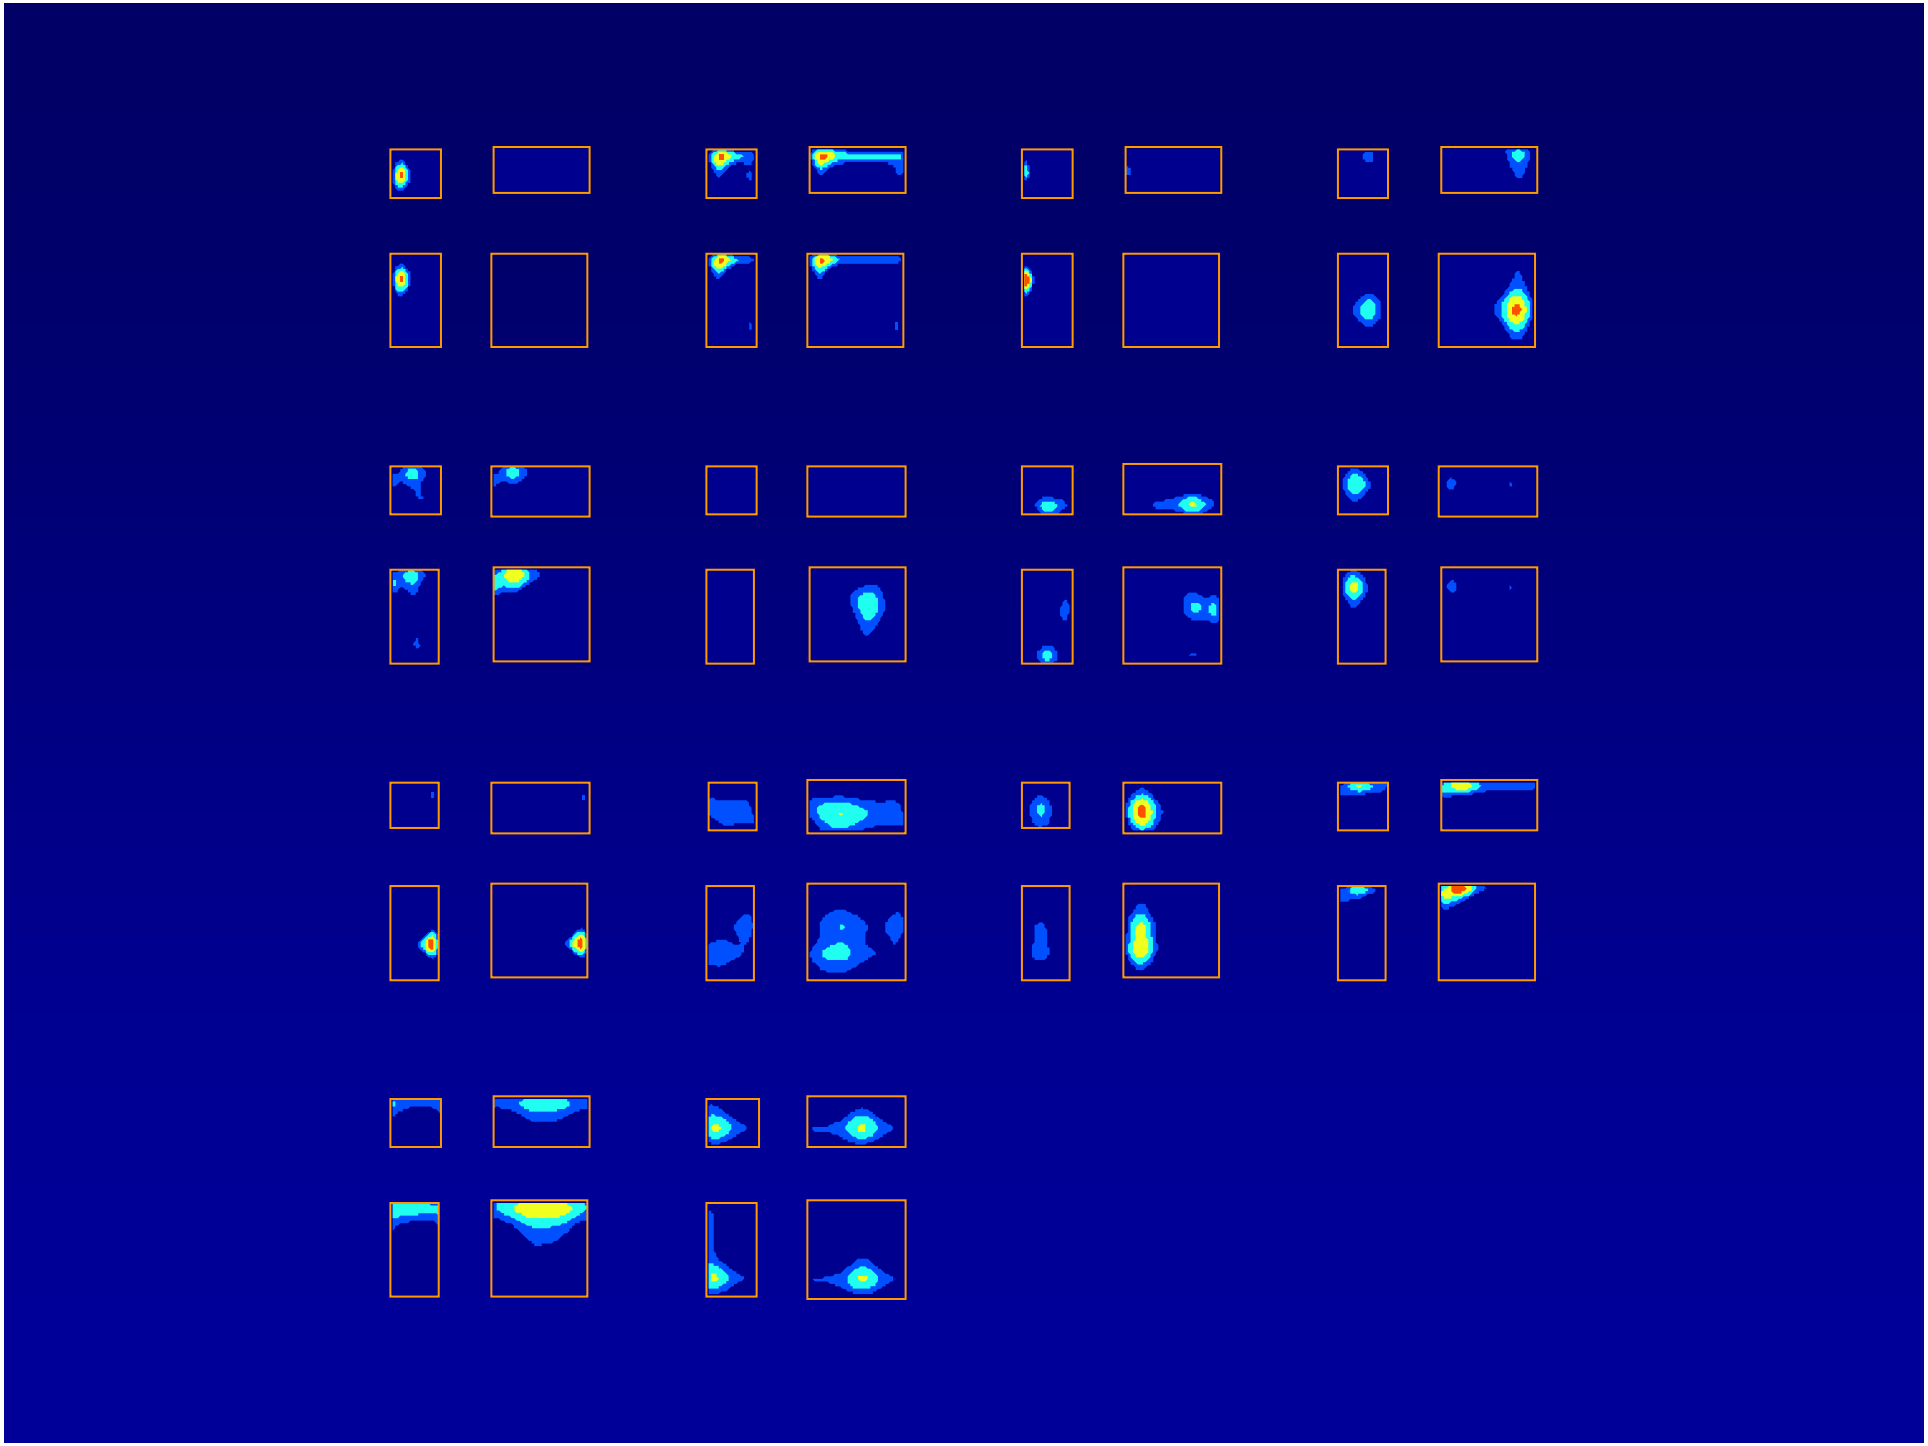

Physiological Prediction

Take the solution obtained by fitting a cell in one set of environments, and apply it to a novel set.

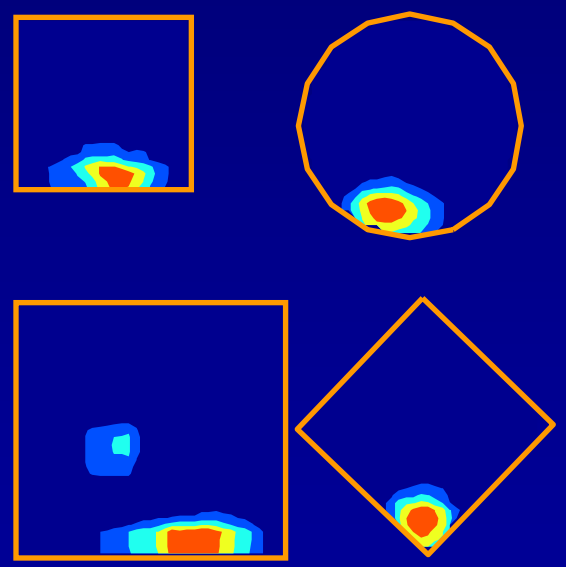

Place Cell Data
(Lever *et al*, 1999)

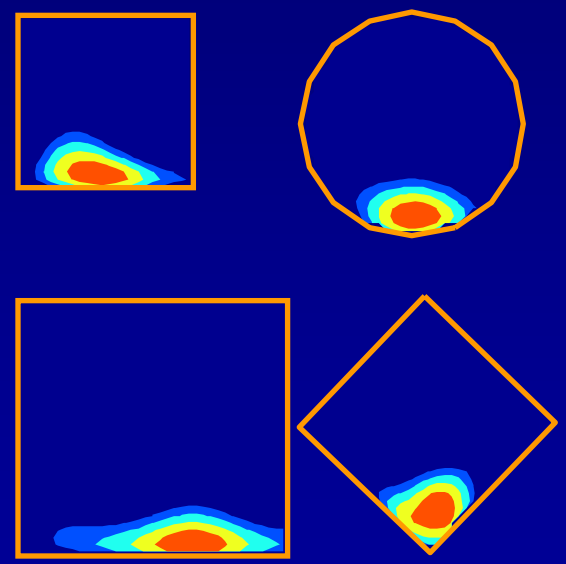

Simulated Place Cell

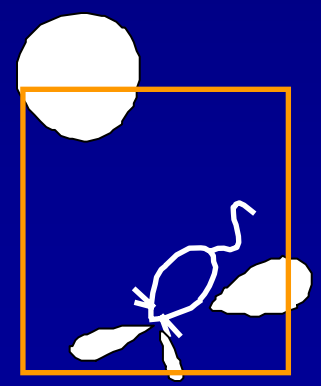

Best fitting BVCs

Physiological Prediction

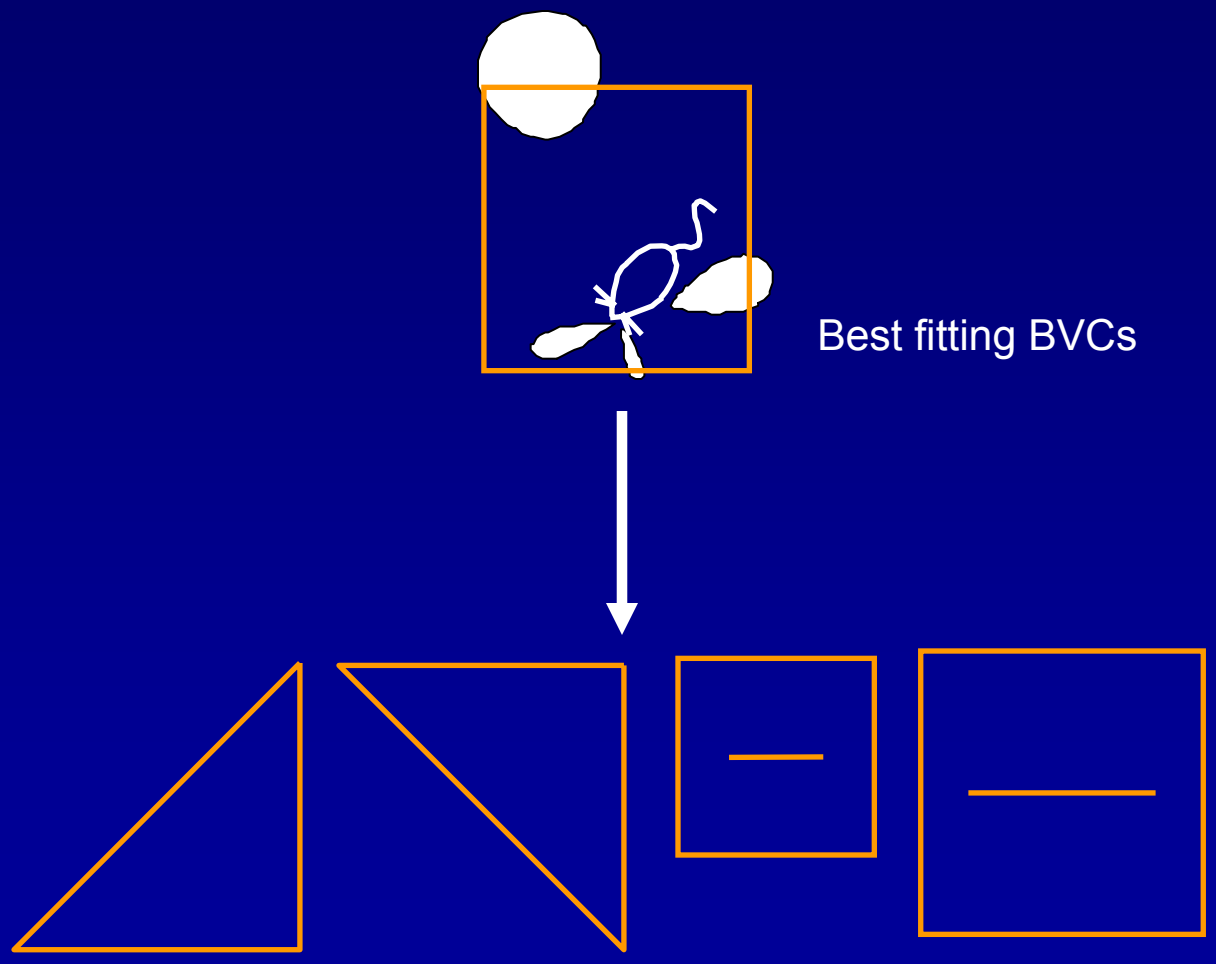

Physiological Prediction

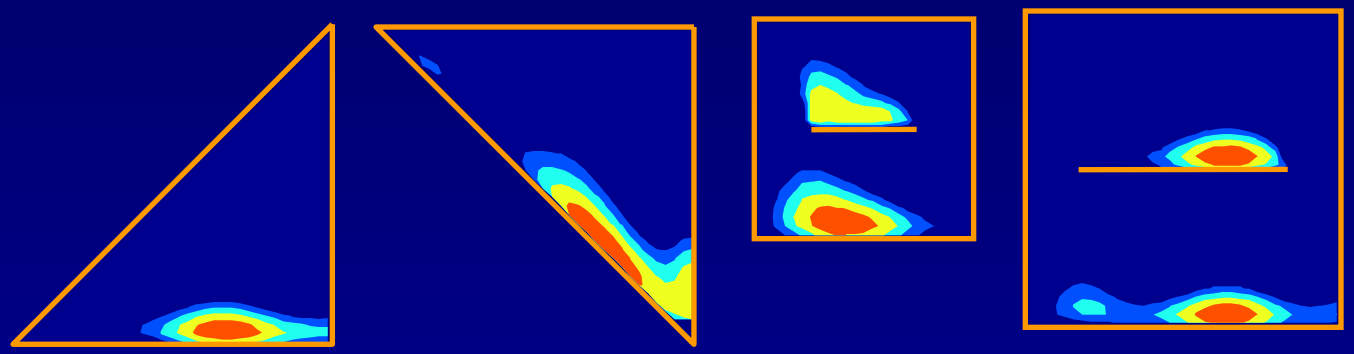

Simulation (Prediction)

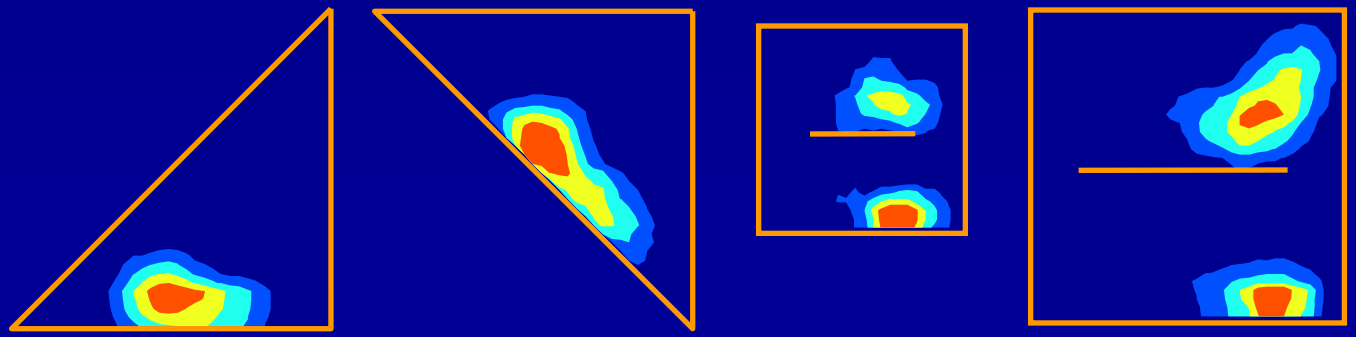

Experimental Data (Lever et al, 1999)

BVC Model
Hartley et al. (2000)

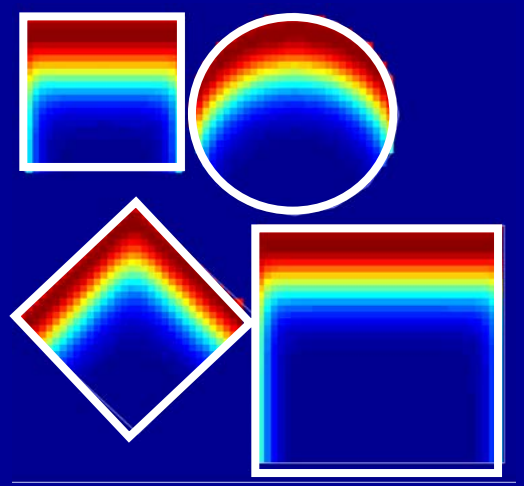

Examples of subicular cells
(Lever, unpublished; Barry et al.,in press)

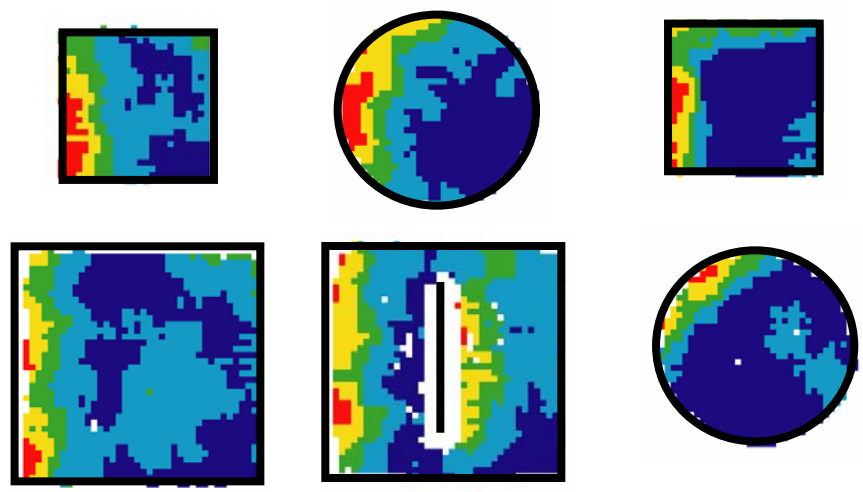

Cell A **Cell B**

Behavioural Model I

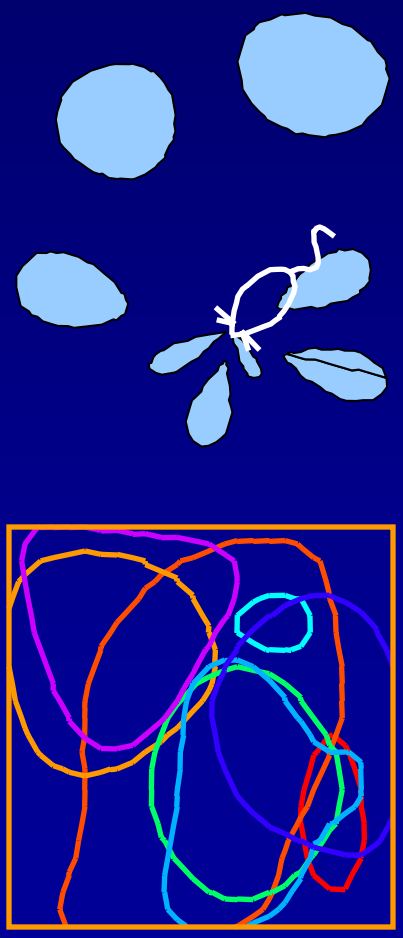

- Each place cell has a random set of BVC inputs.
- This results in a set of place fields which overlap and cover the space of an enclosure.

Behavioural Model I

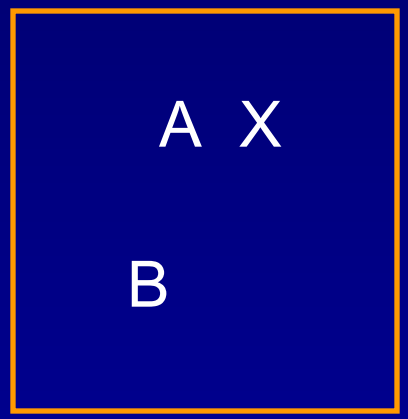

X, A and B: three points in an arena

Place cells' firing rates at X

Similar pattern at A (nearby)

Different pattern at B (distant)

Behavioural Model I

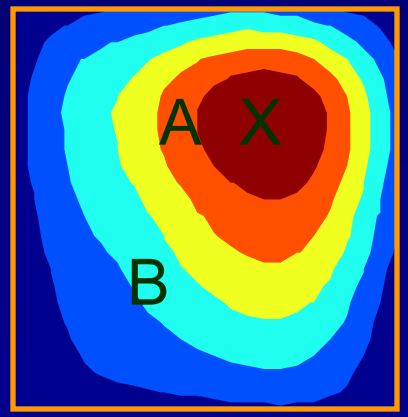

Similarity of points to X
(dot product of firing rate vectors for 100 simulated cells firing at X)

- To return to X (or to search for an object you last saw at X), go to the place where the pattern of place cell firing is most similar to that at X.

Behavioural Model I

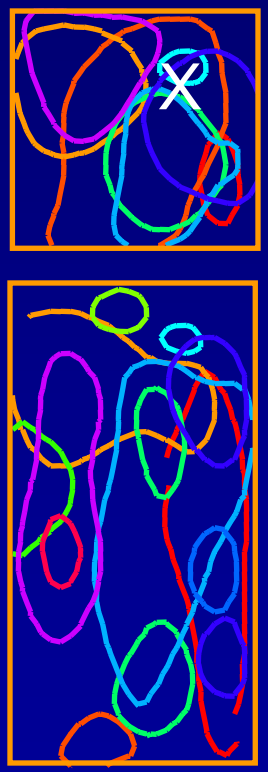

- The same cells fire in different environments. So we can predict where e.g., search activity would be directed following a distortion of the arena.

Behavioural Model I

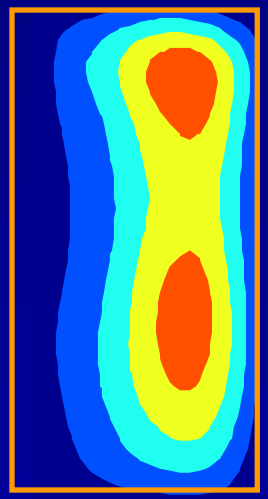

Similarity of points in a novel arena to X (based on dot product)

- The same cells fire in different environments. So we can predict where e.g., search activity would be directed following a distortion of the arena.

1. Store pattern of activity at the flag

2. Calculate change in activity pattern relative to stored pattern for each location in new environment

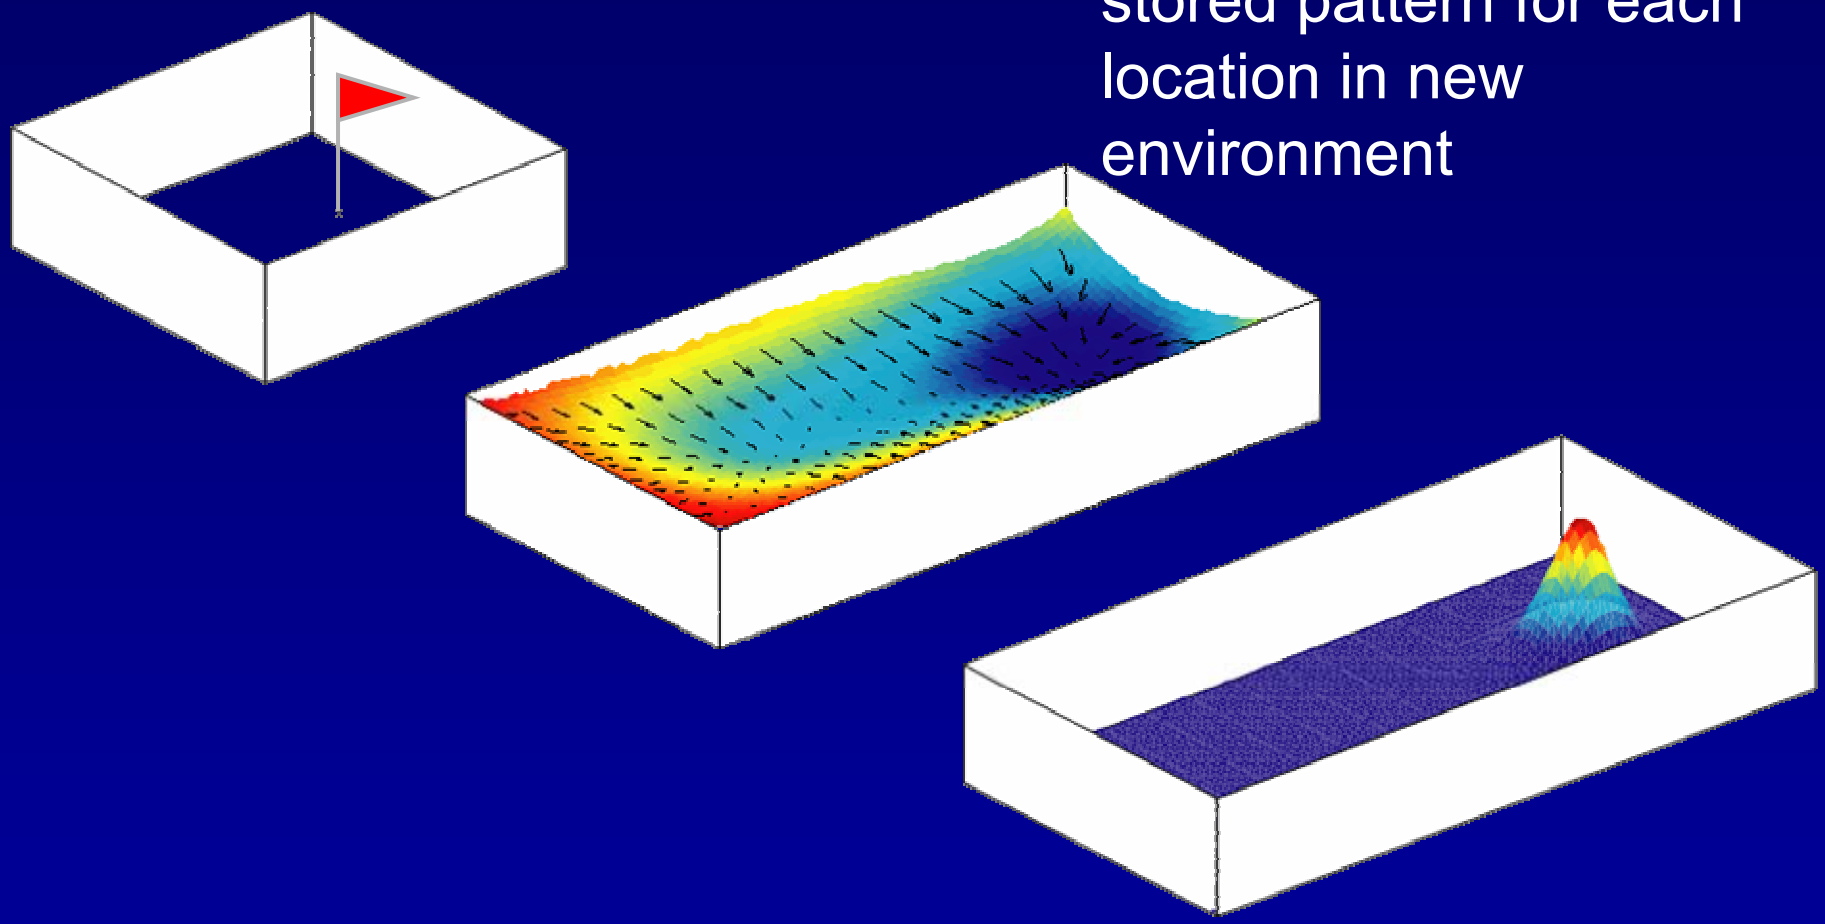

3. Probability of a response depends on the similarity of pattern

Human Behavioural Experiment

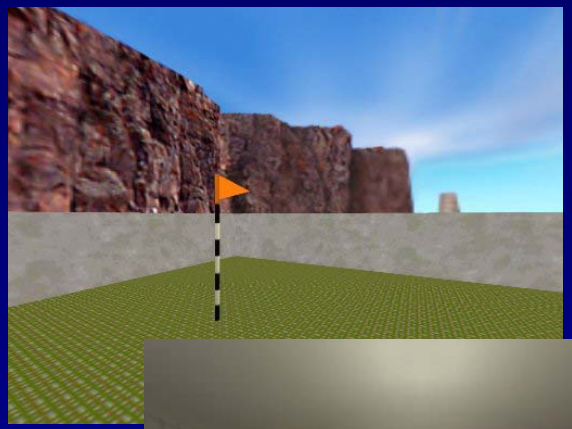

Presentation

Freely move about the arena, viewing golf flag.

Briefly removed from arena…

… then returned at randomized location*.

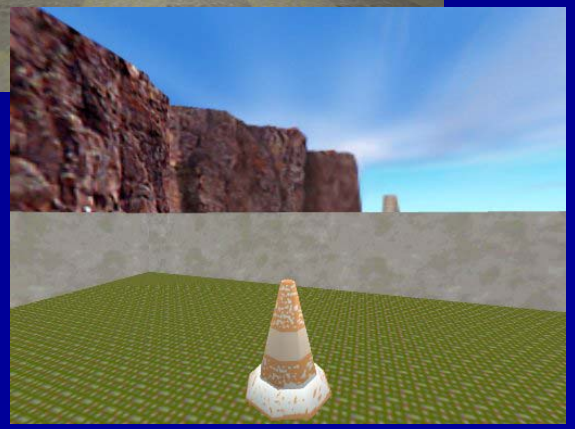

Testing

“Place marker where you think the golf flag was”

Repeat for several markers*.

background cues projected at ∞

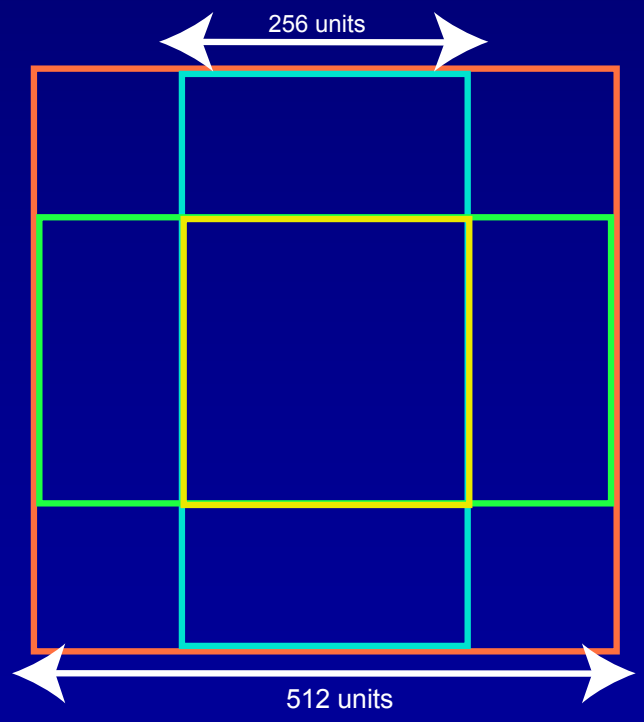

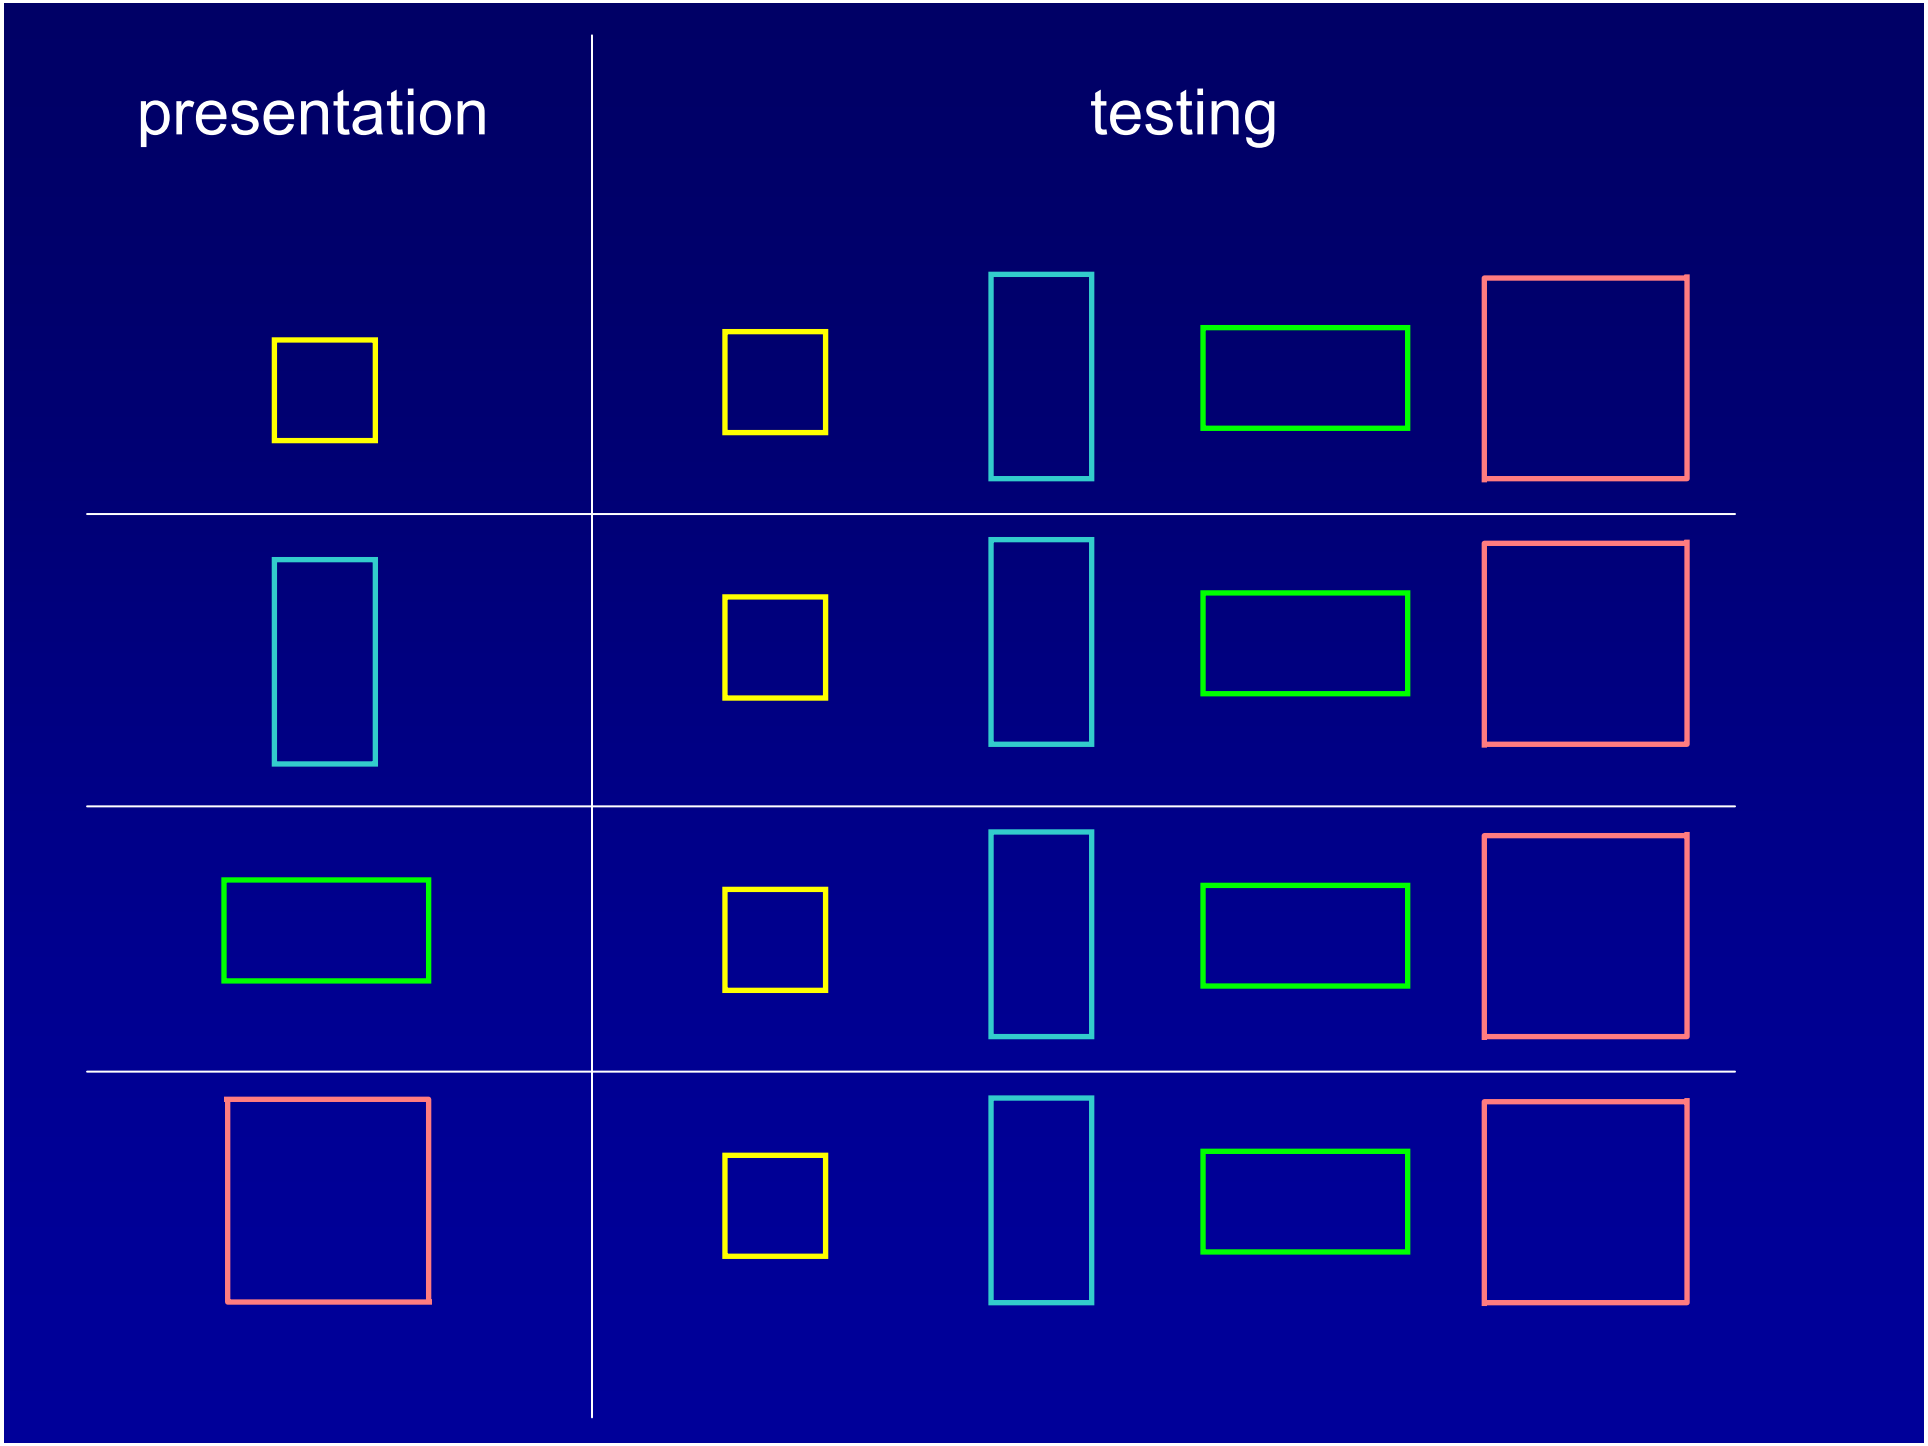
presentation
testing

Stretchy Room
DEMO

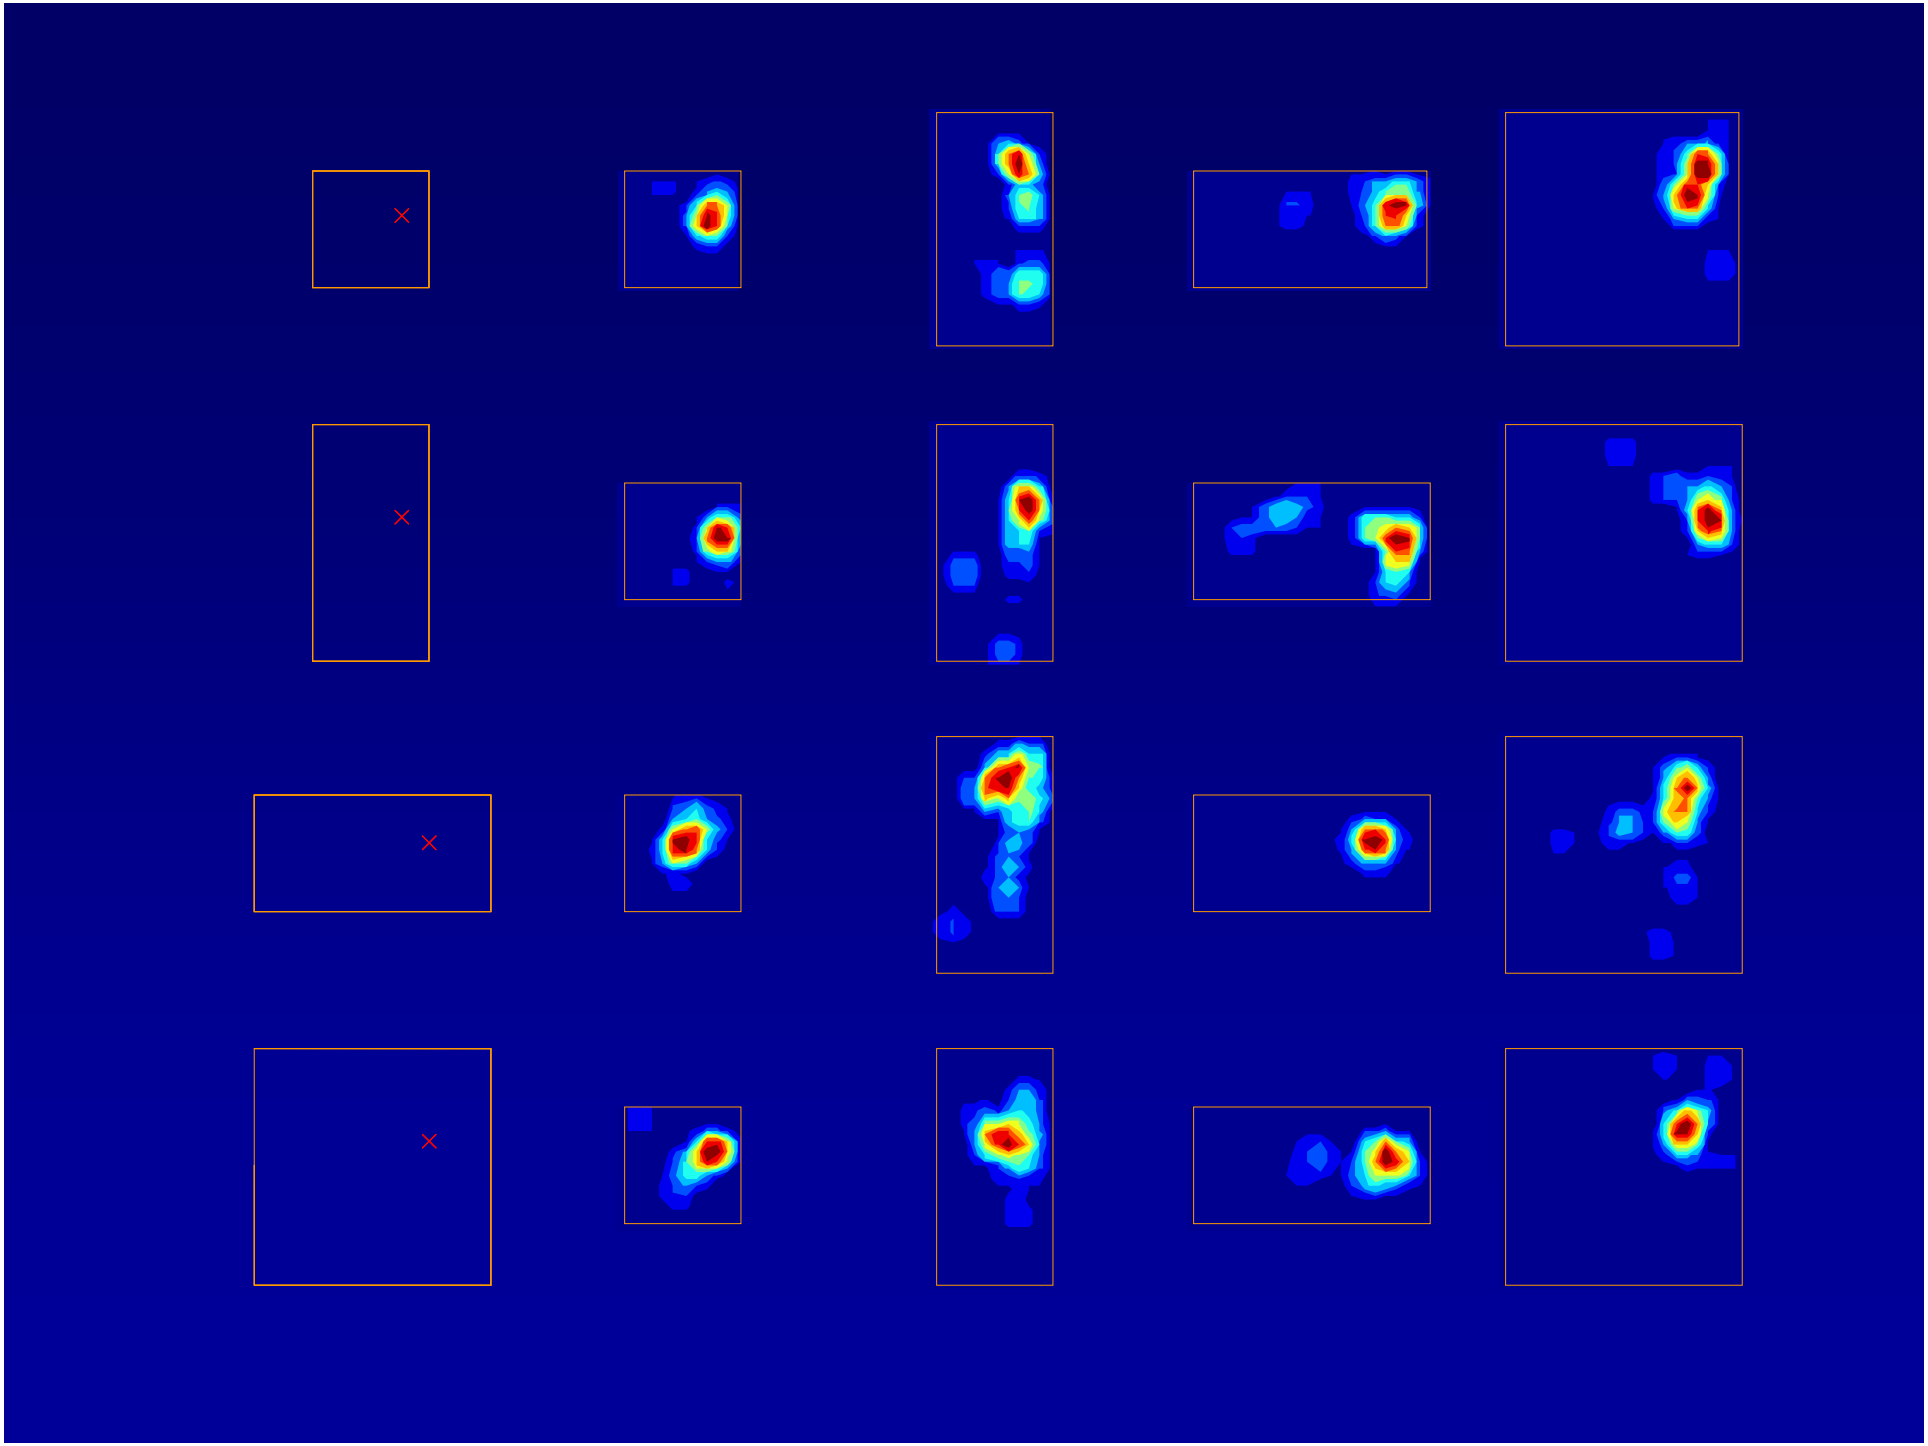

Fixed distances (to nearer walls) =△

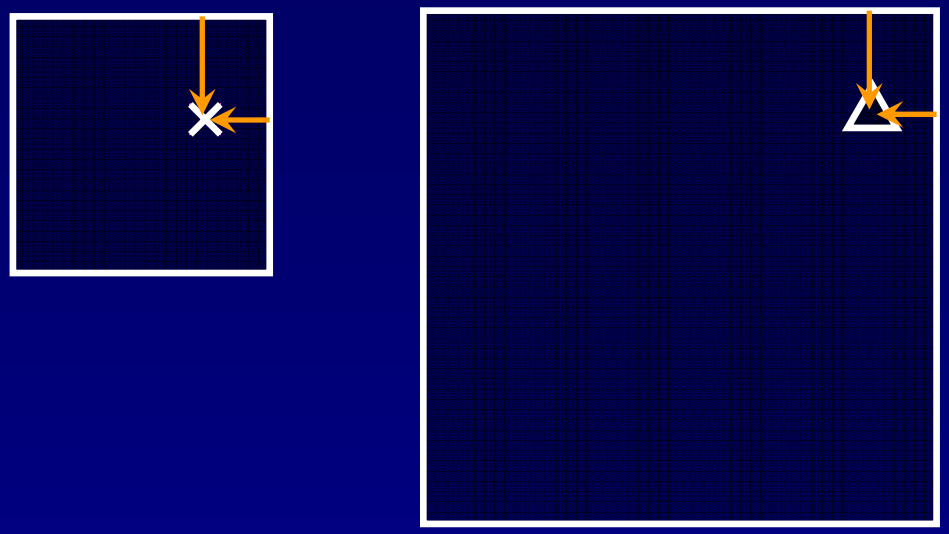

Fixed ratio (of distance between opposing walls) =○

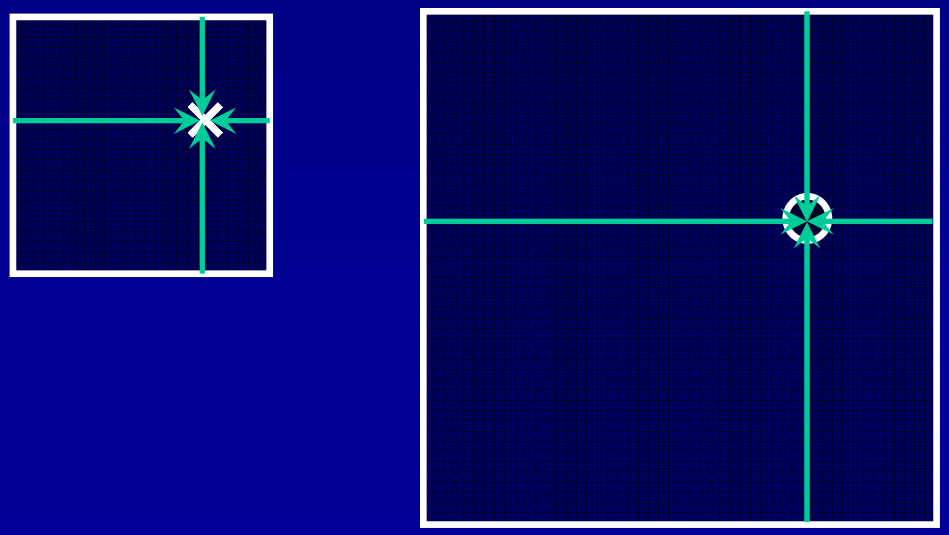

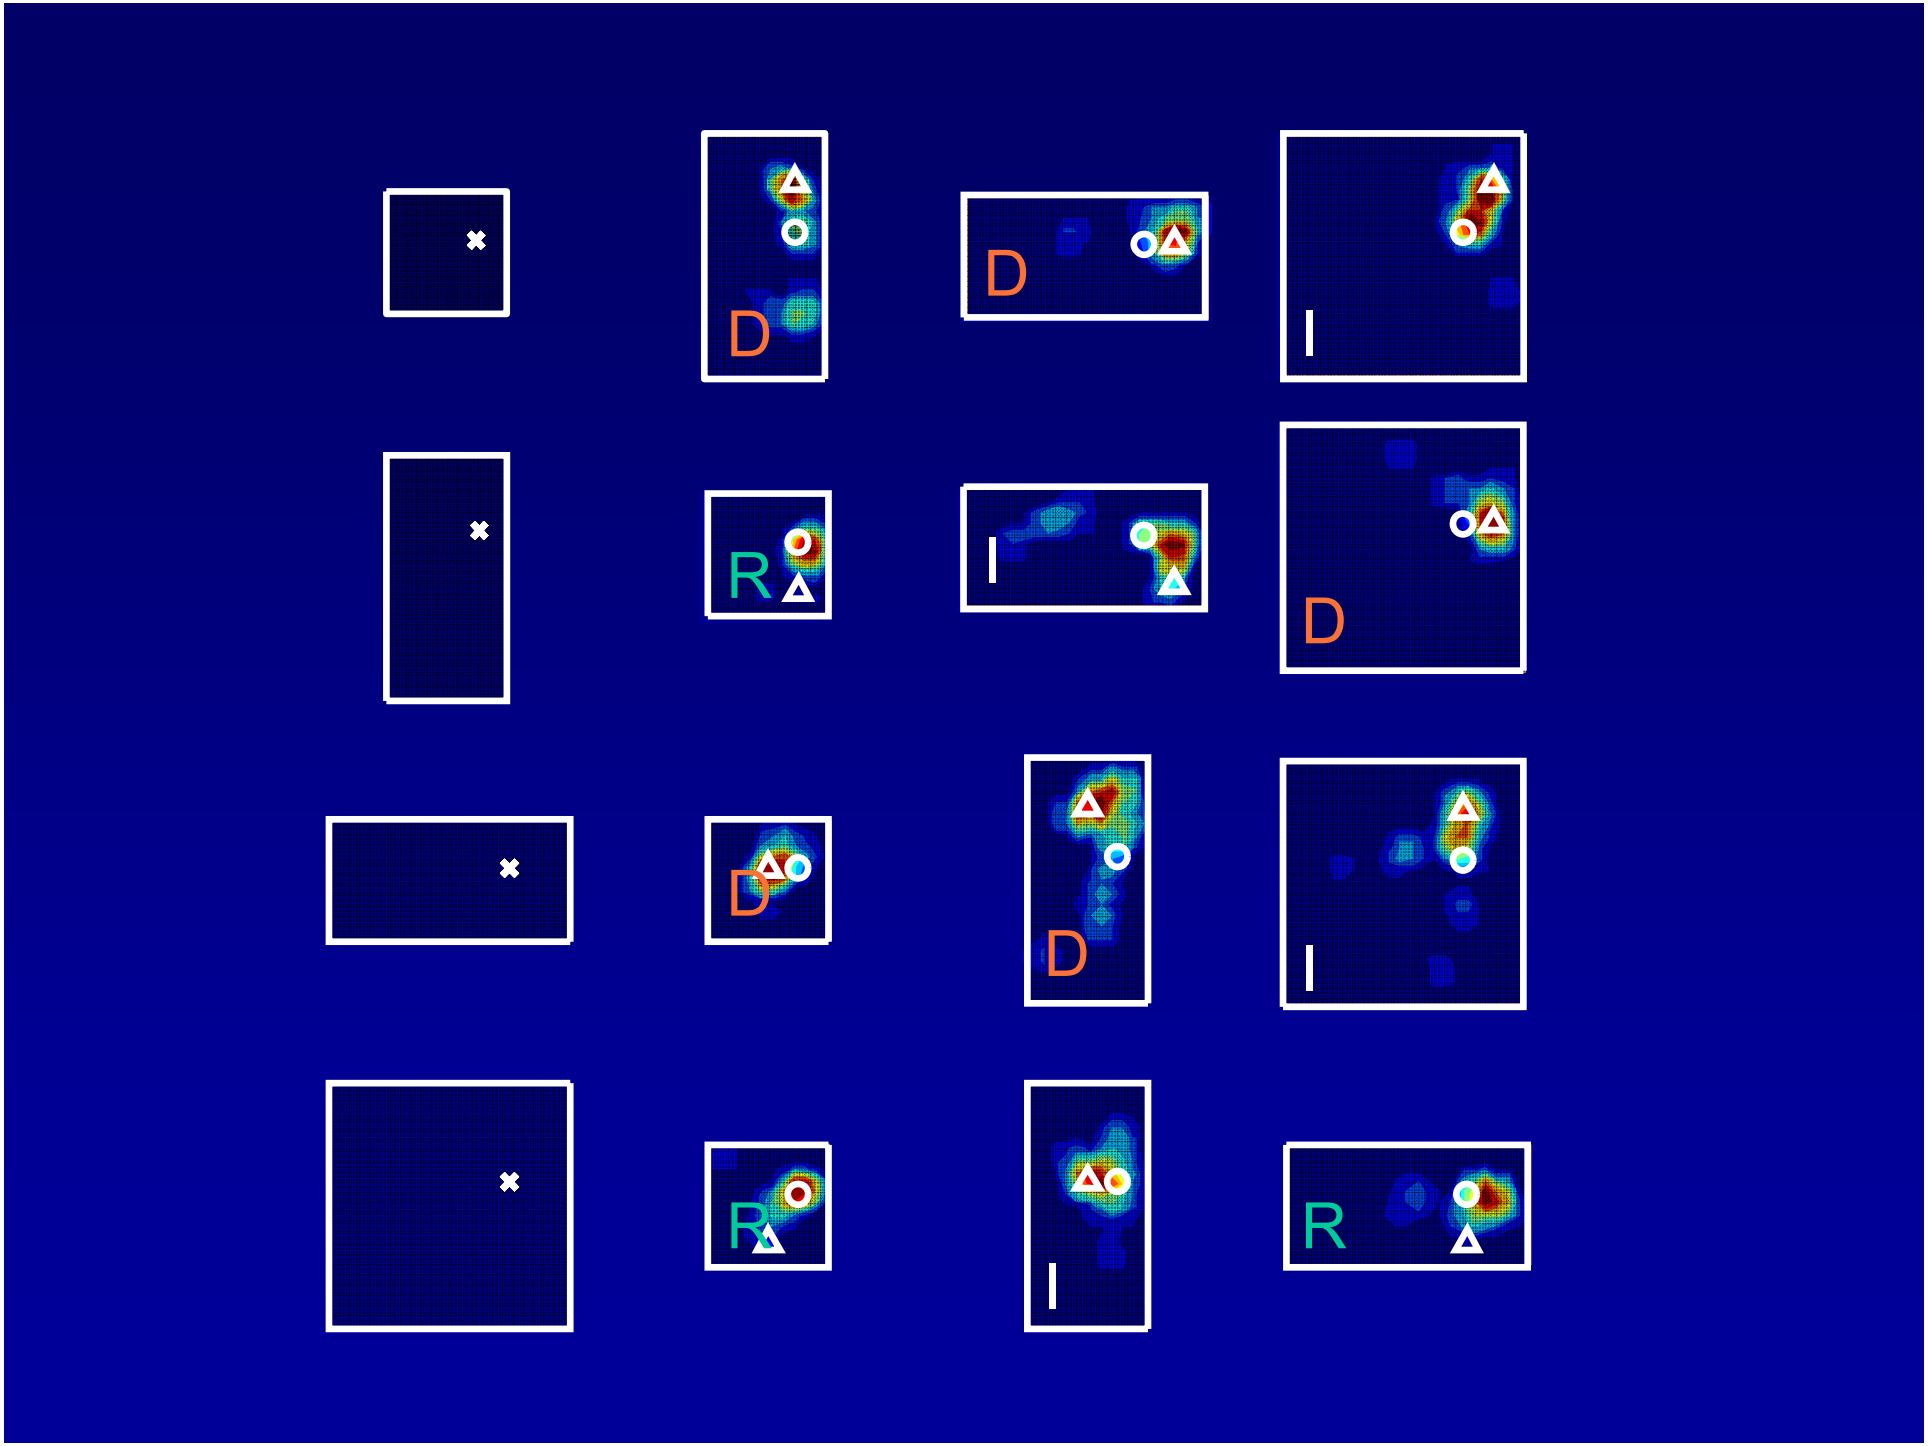
D
D
I
R
I
D
D
D
I
R
I
R

Interaction of geometric manipulation and cue location

- fixed distances in expansions of the arena and when the cued location was near to a wall

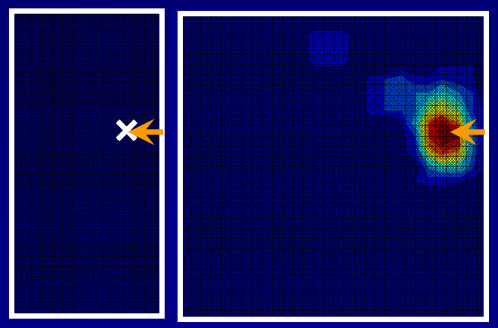

- fixed ratios more common in contractions of the arena and when the cued location was near the centre of the arena.

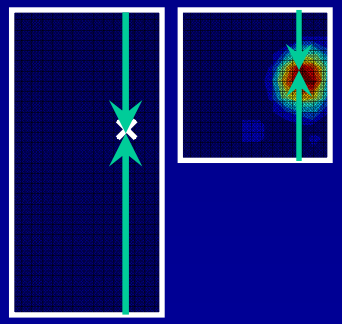

Problems with Behavioural Model I

- Predicts responses maintain a fixed distance to the nearer walls.
  - True, but only for expansions of arena, and for locations near the edge of the arena.
- Predicts bimodal response fields
  - Rather unusual in the data (but seen in some expansions).
  - However, some of this is due to between subject variation. Individuals responses within a given environment are usually unimodally distributed.
- Not really compatible with response distributions seen under contractions of the arena.

What might this mean for the model?

Readout: resolving ambiguity in the place cell representation.

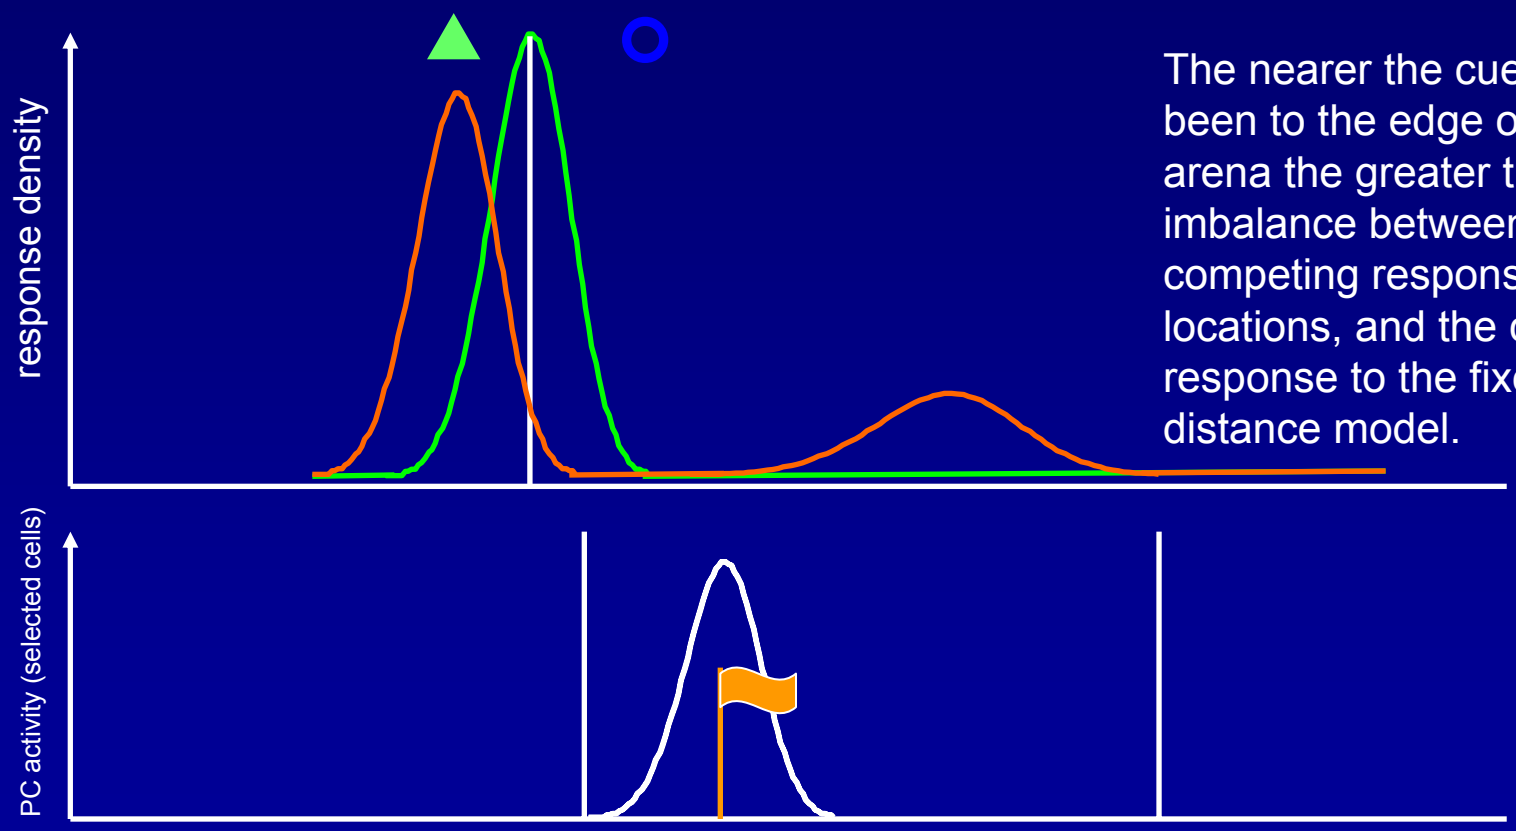

The nearer the cue had been to the edge of the arena the greater the imbalance between competing response locations, and the closer the response to the fixed distance model.

What might this mean for the model?

Readout: resolving ambiguity in the place cell representation.

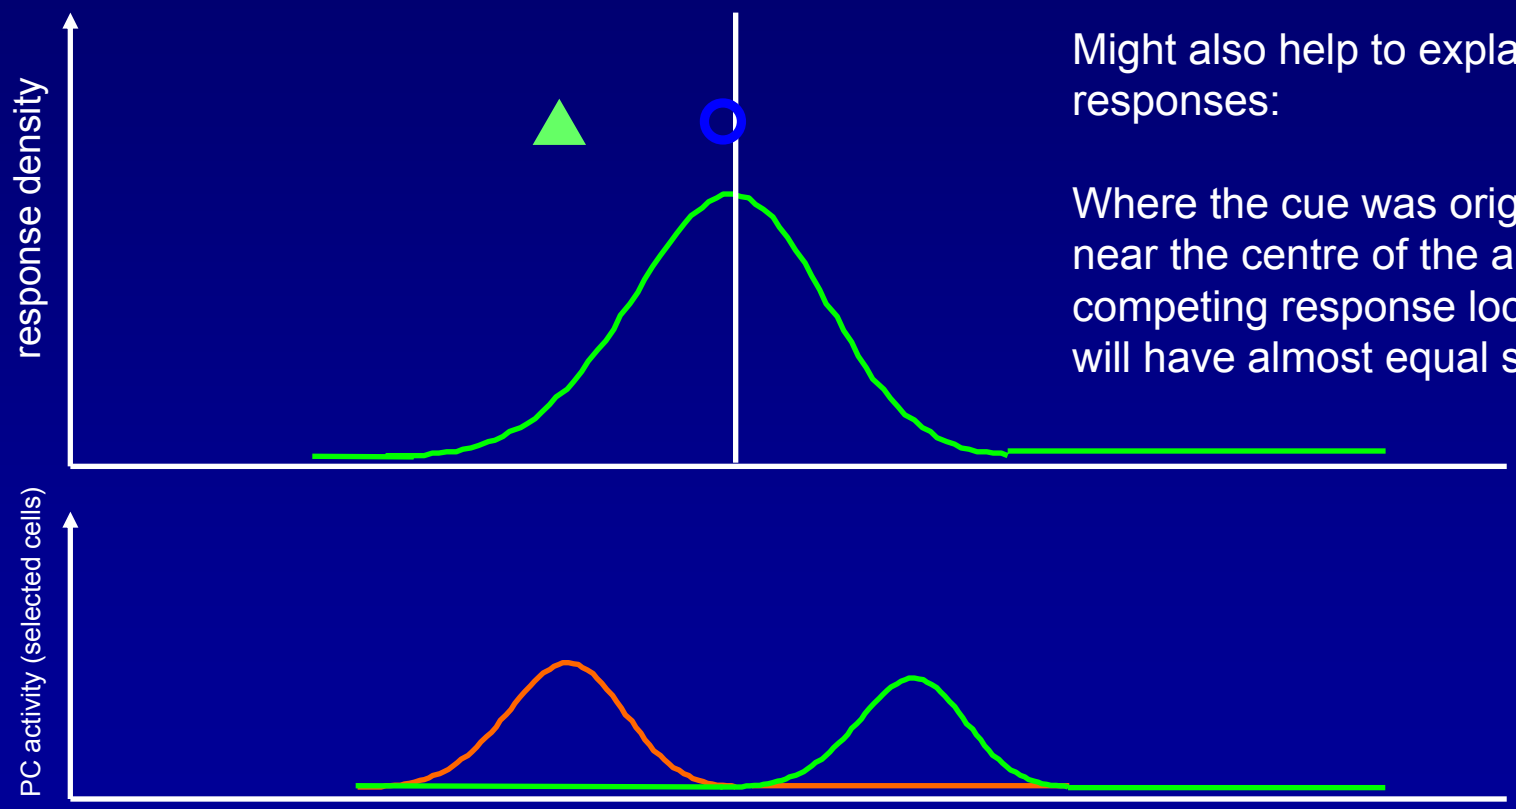

Might also help to explain ratio responses:

Where the cue was originally near the centre of the arena, the competing response locations will have almost equal strength.

Behavioural Model

- Approach
  - Use readout mechanism that resolves ambiguity due to cue conflict.
  - Examine effect of weighting the influence of geometric cues (boundaries) by their proximity $1/(d+c)$.
  - Compare with different possible models (i.e., cues that might be stored).

Behavioural Model II

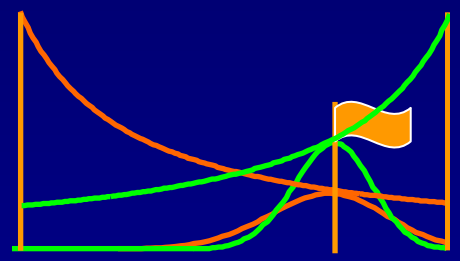

Instead of storing the pattern activation of a population of place cells at the cue location, we store the proximities *1/(d+c)* to each of the four walls.

This is equivalent to storing the peak firing rates of BVCs which are peaked at the cue location.

Learning: store proximities *1/(d+c)* of walls to cue (V_C)

Testing: find the place which minimizes the difference between stored and current (V_{Tx}) wall proximities in a least-squares sense.

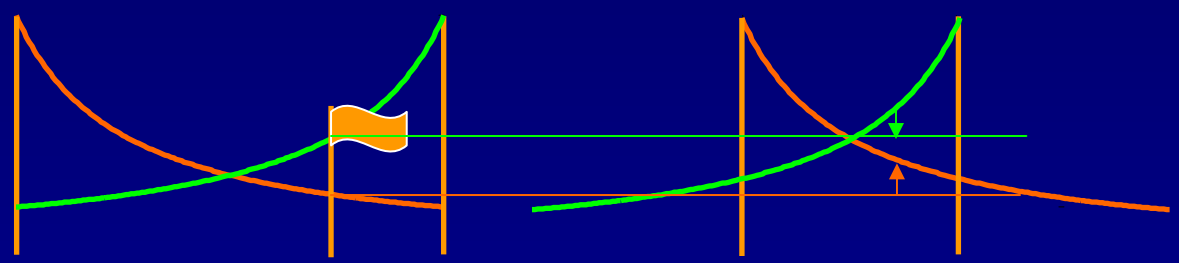

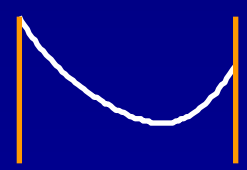

$s=|v_C-v_{Tx}|$

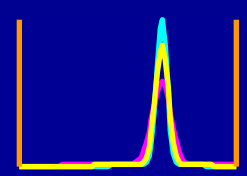

p(response)

Response distribution rotated to decrease fit to the current data set which sums to 1 over the area of the testing arena (SOFTMAX).

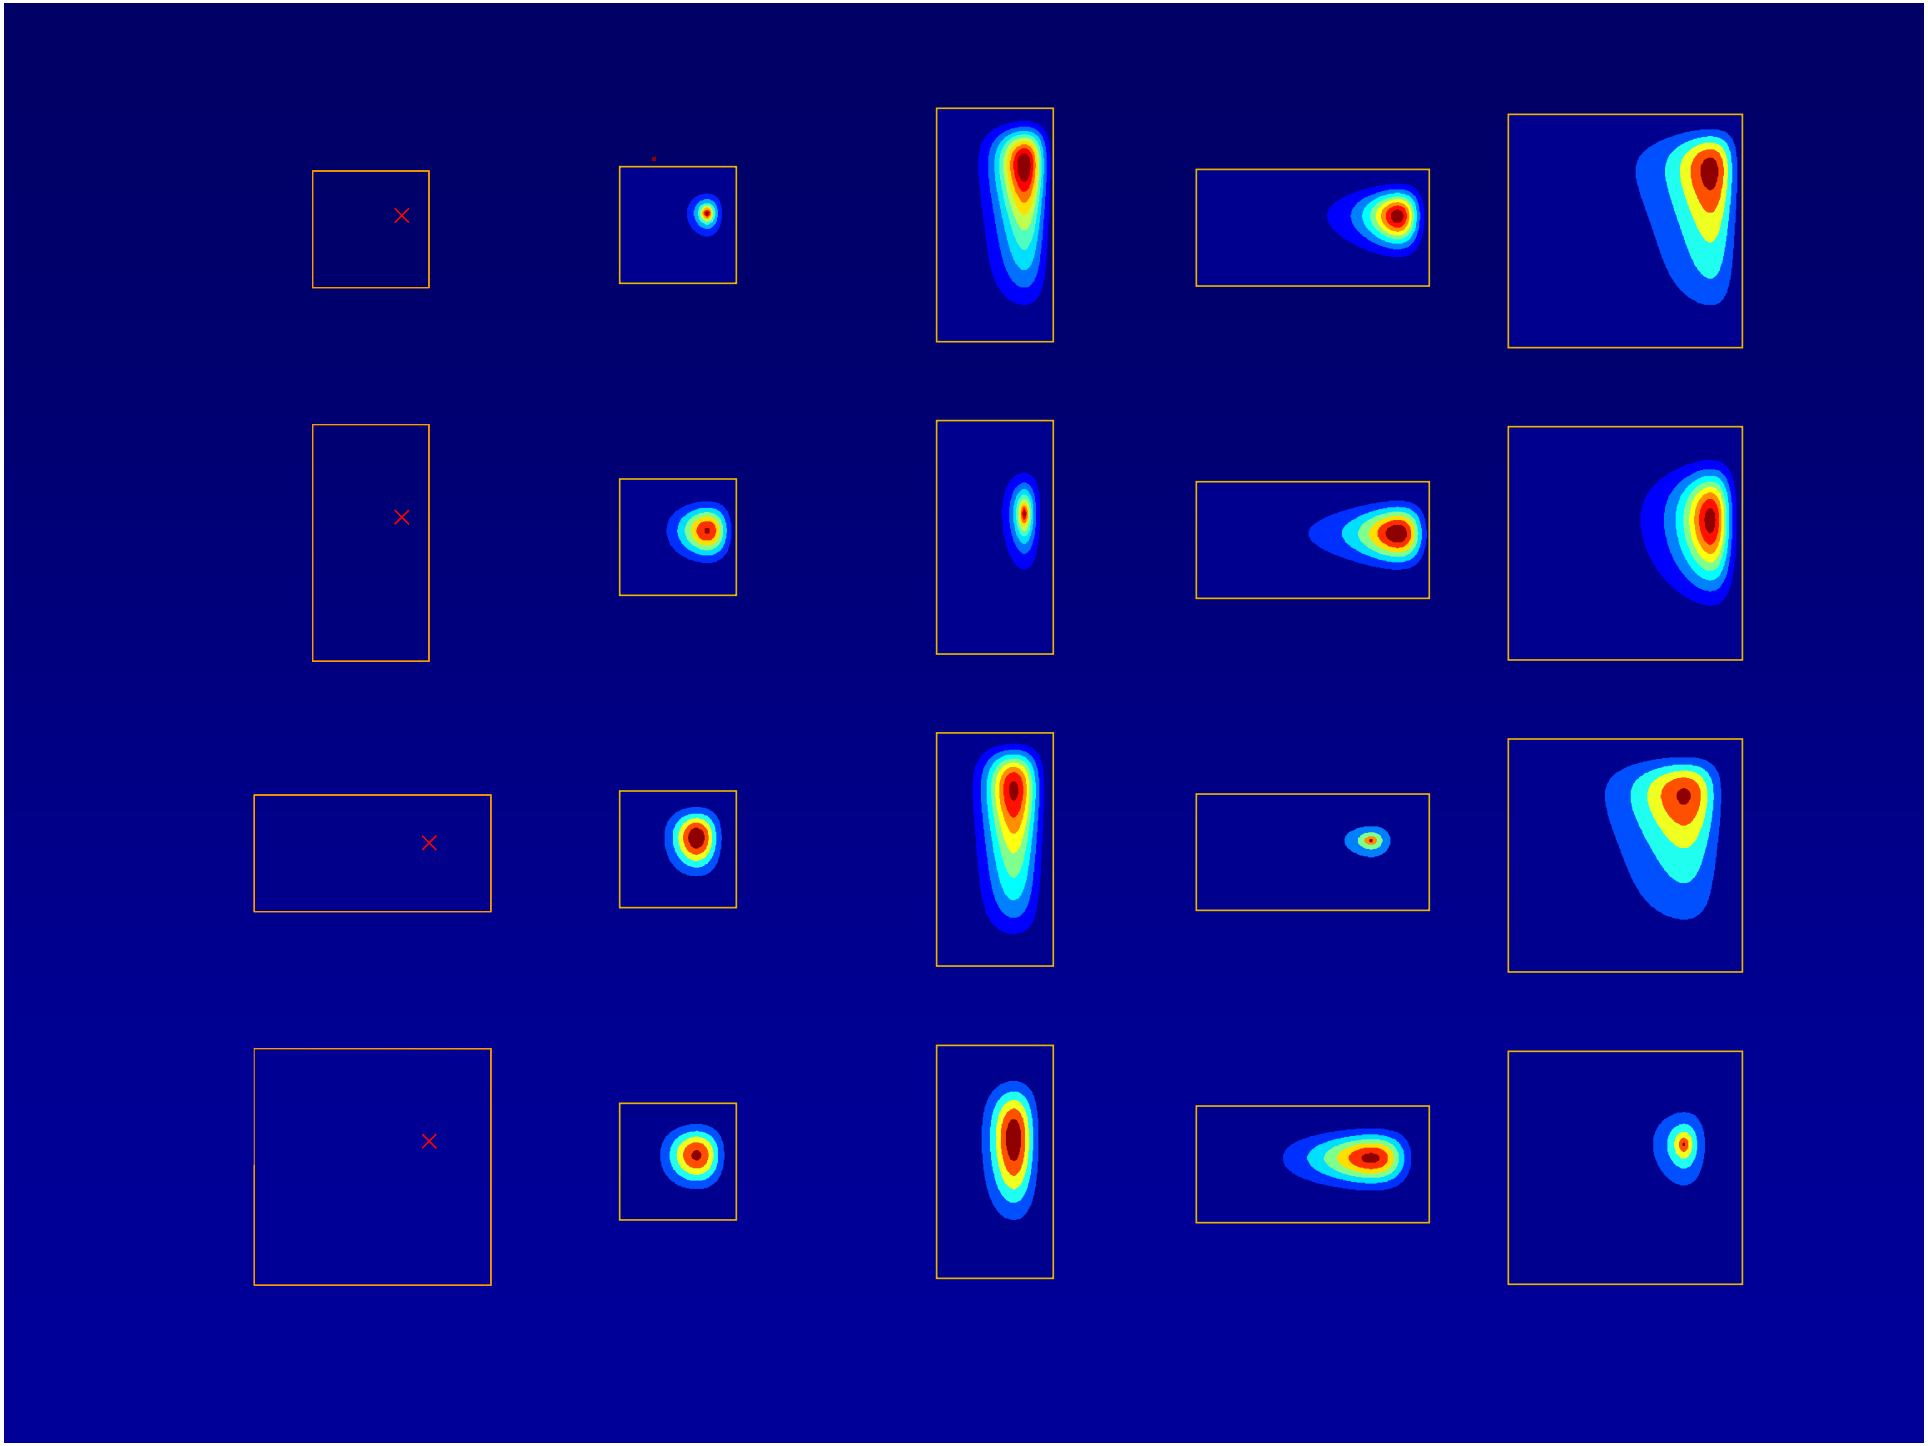

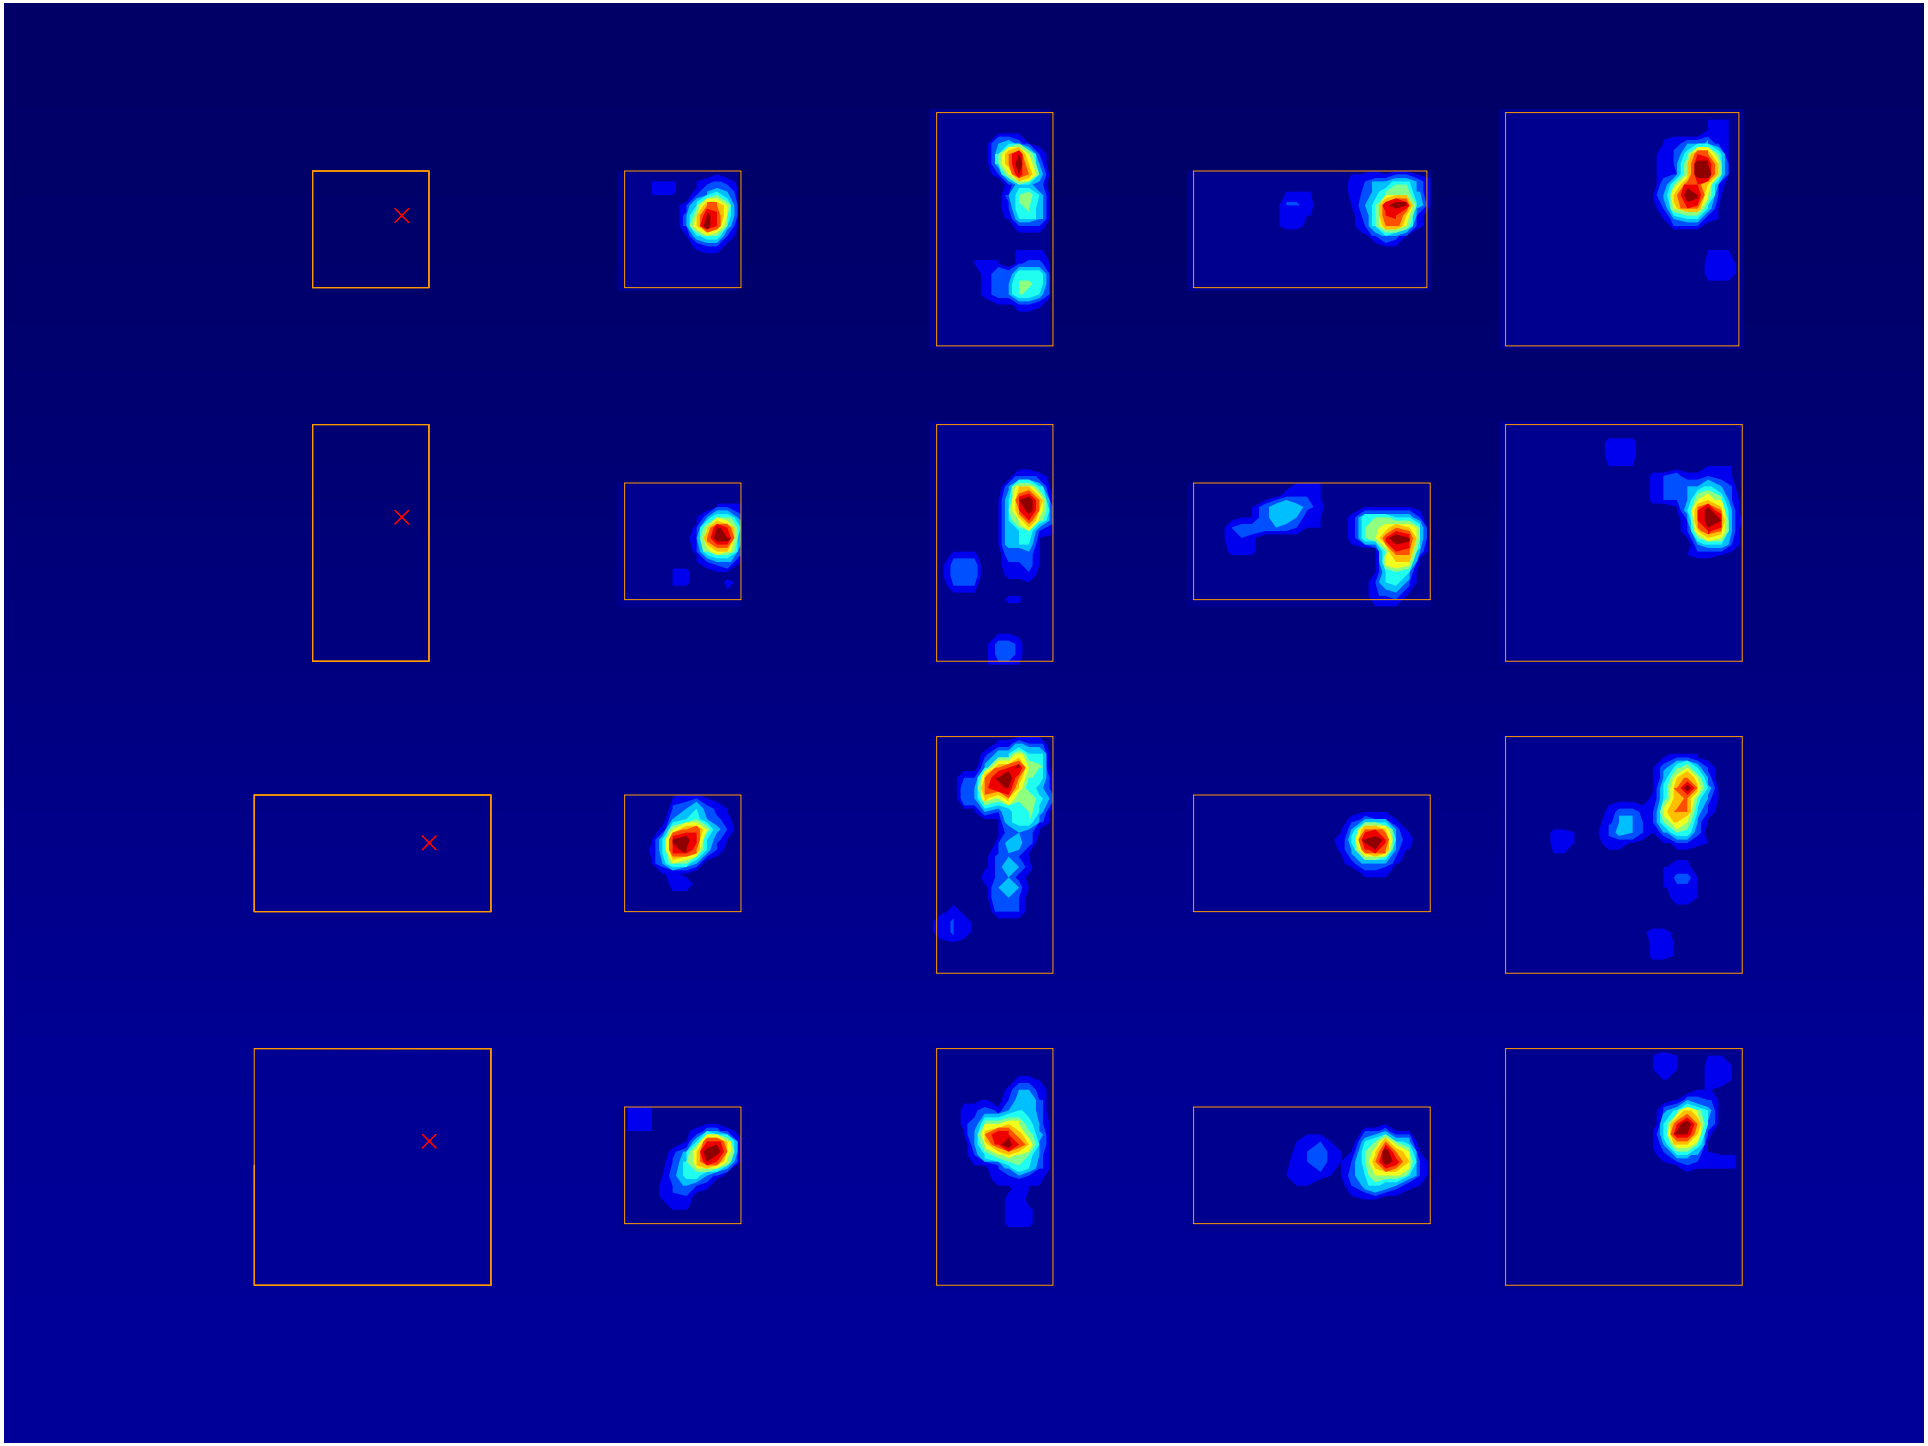

Data

Interaction between transformation type and cue location.

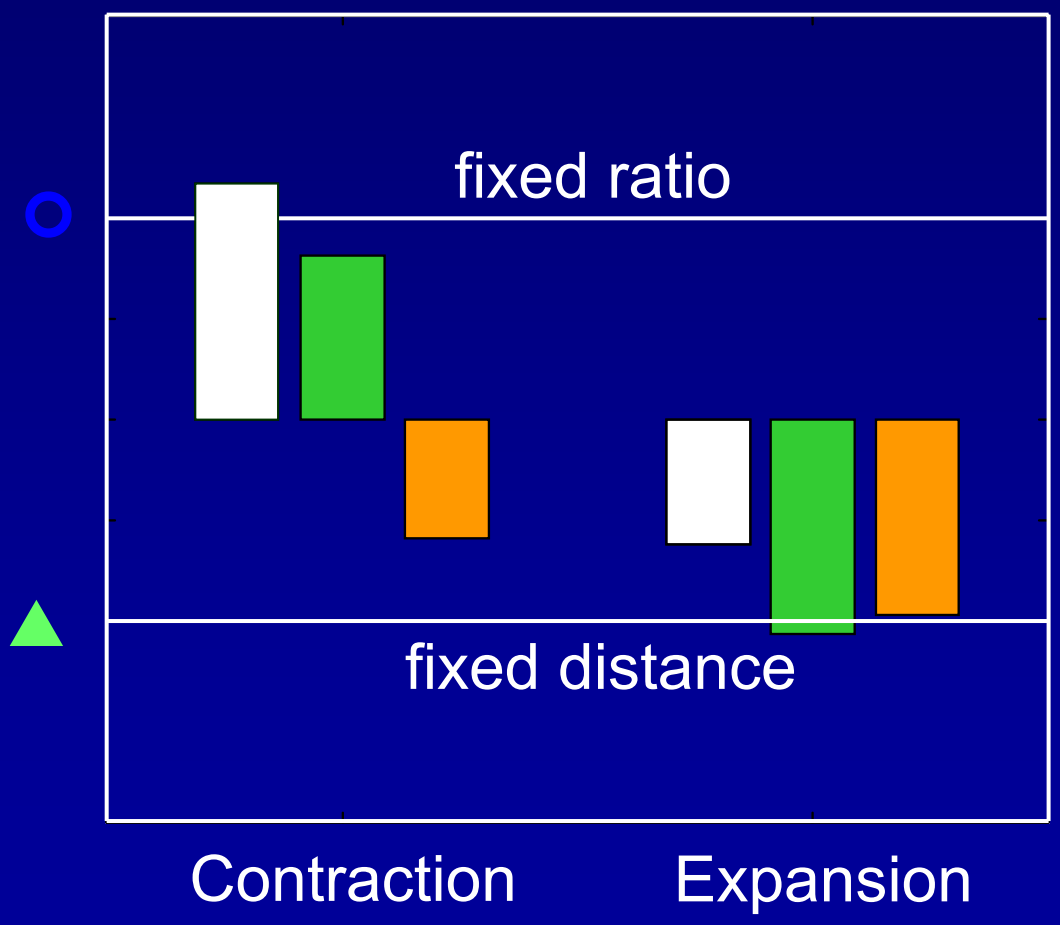

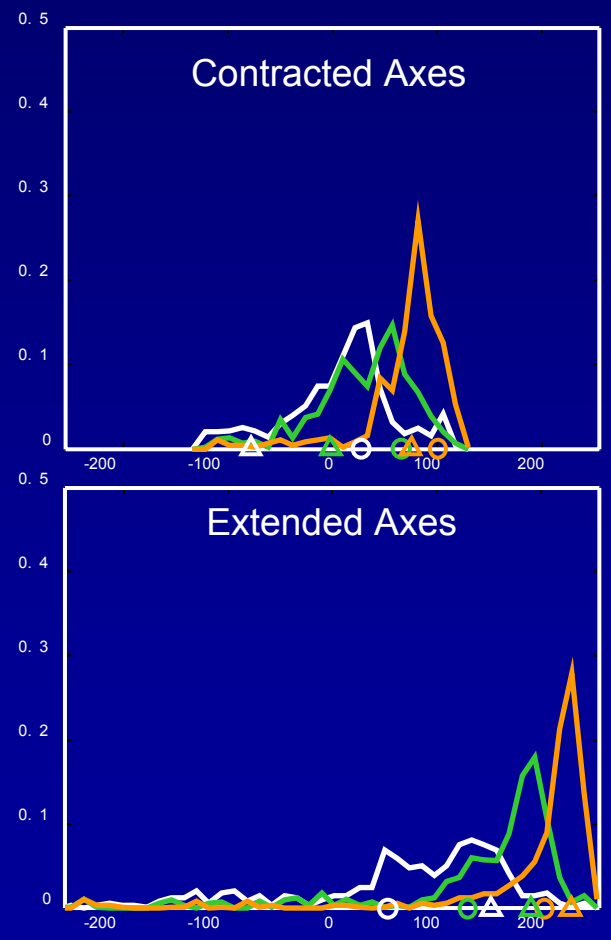

Boundary Proximity Model

Interaction between transformation type and cue location.

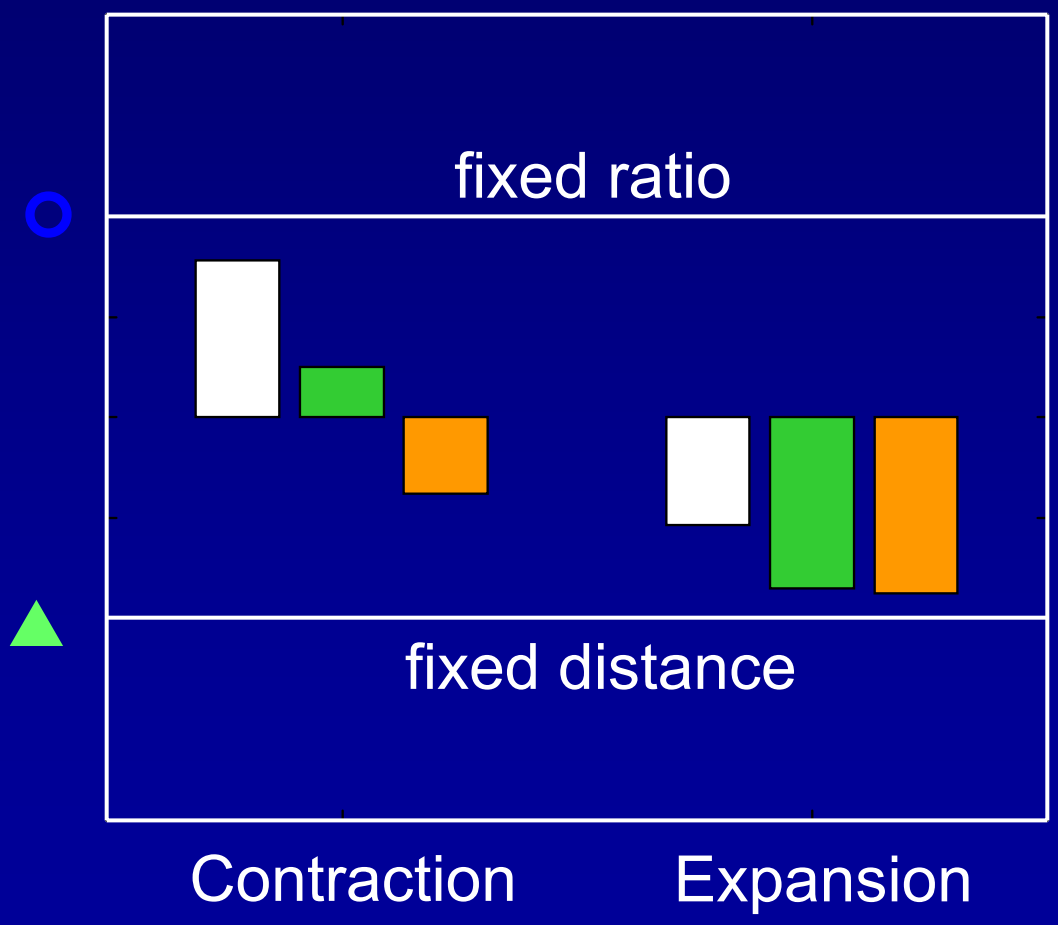

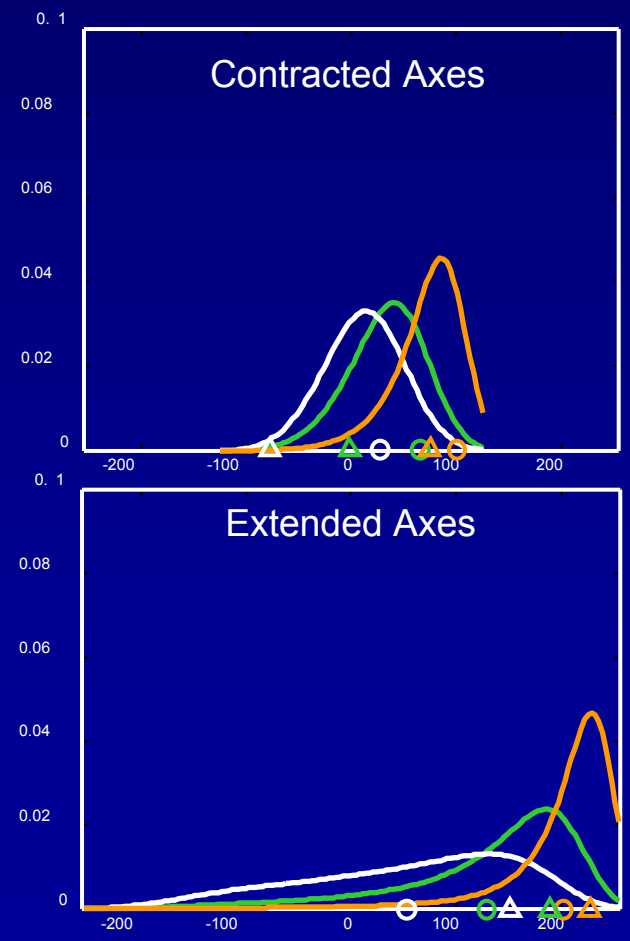

General Response Model

- PRESENTATION: Store some properties of the cue location (v_C) e.g., distance to walls in each direction.
- TESTING: Compare properties of each location in the testing arena (v_{Txy}) with the stored properties of the cue location:
  - Readout: minimize $s_{xy}=|v_C-v_{Txy}|$.
  - Response distribution: a monotonic decreasing function of s which sums to 1 over the area of the testing arena (SOFTMAX).
  - Sharpness of distribution varied to maximize fit to the entire dataset.

Some Alternative Models
(what are the elements of v_C and v_T)

- ▲ Distances to nearer walls
- ○ Ratios between opposing walls
- □ Distances to all four walls
- ♦ Bearings to corners
- ★ Proximity of all four walls $1/(d+c)$

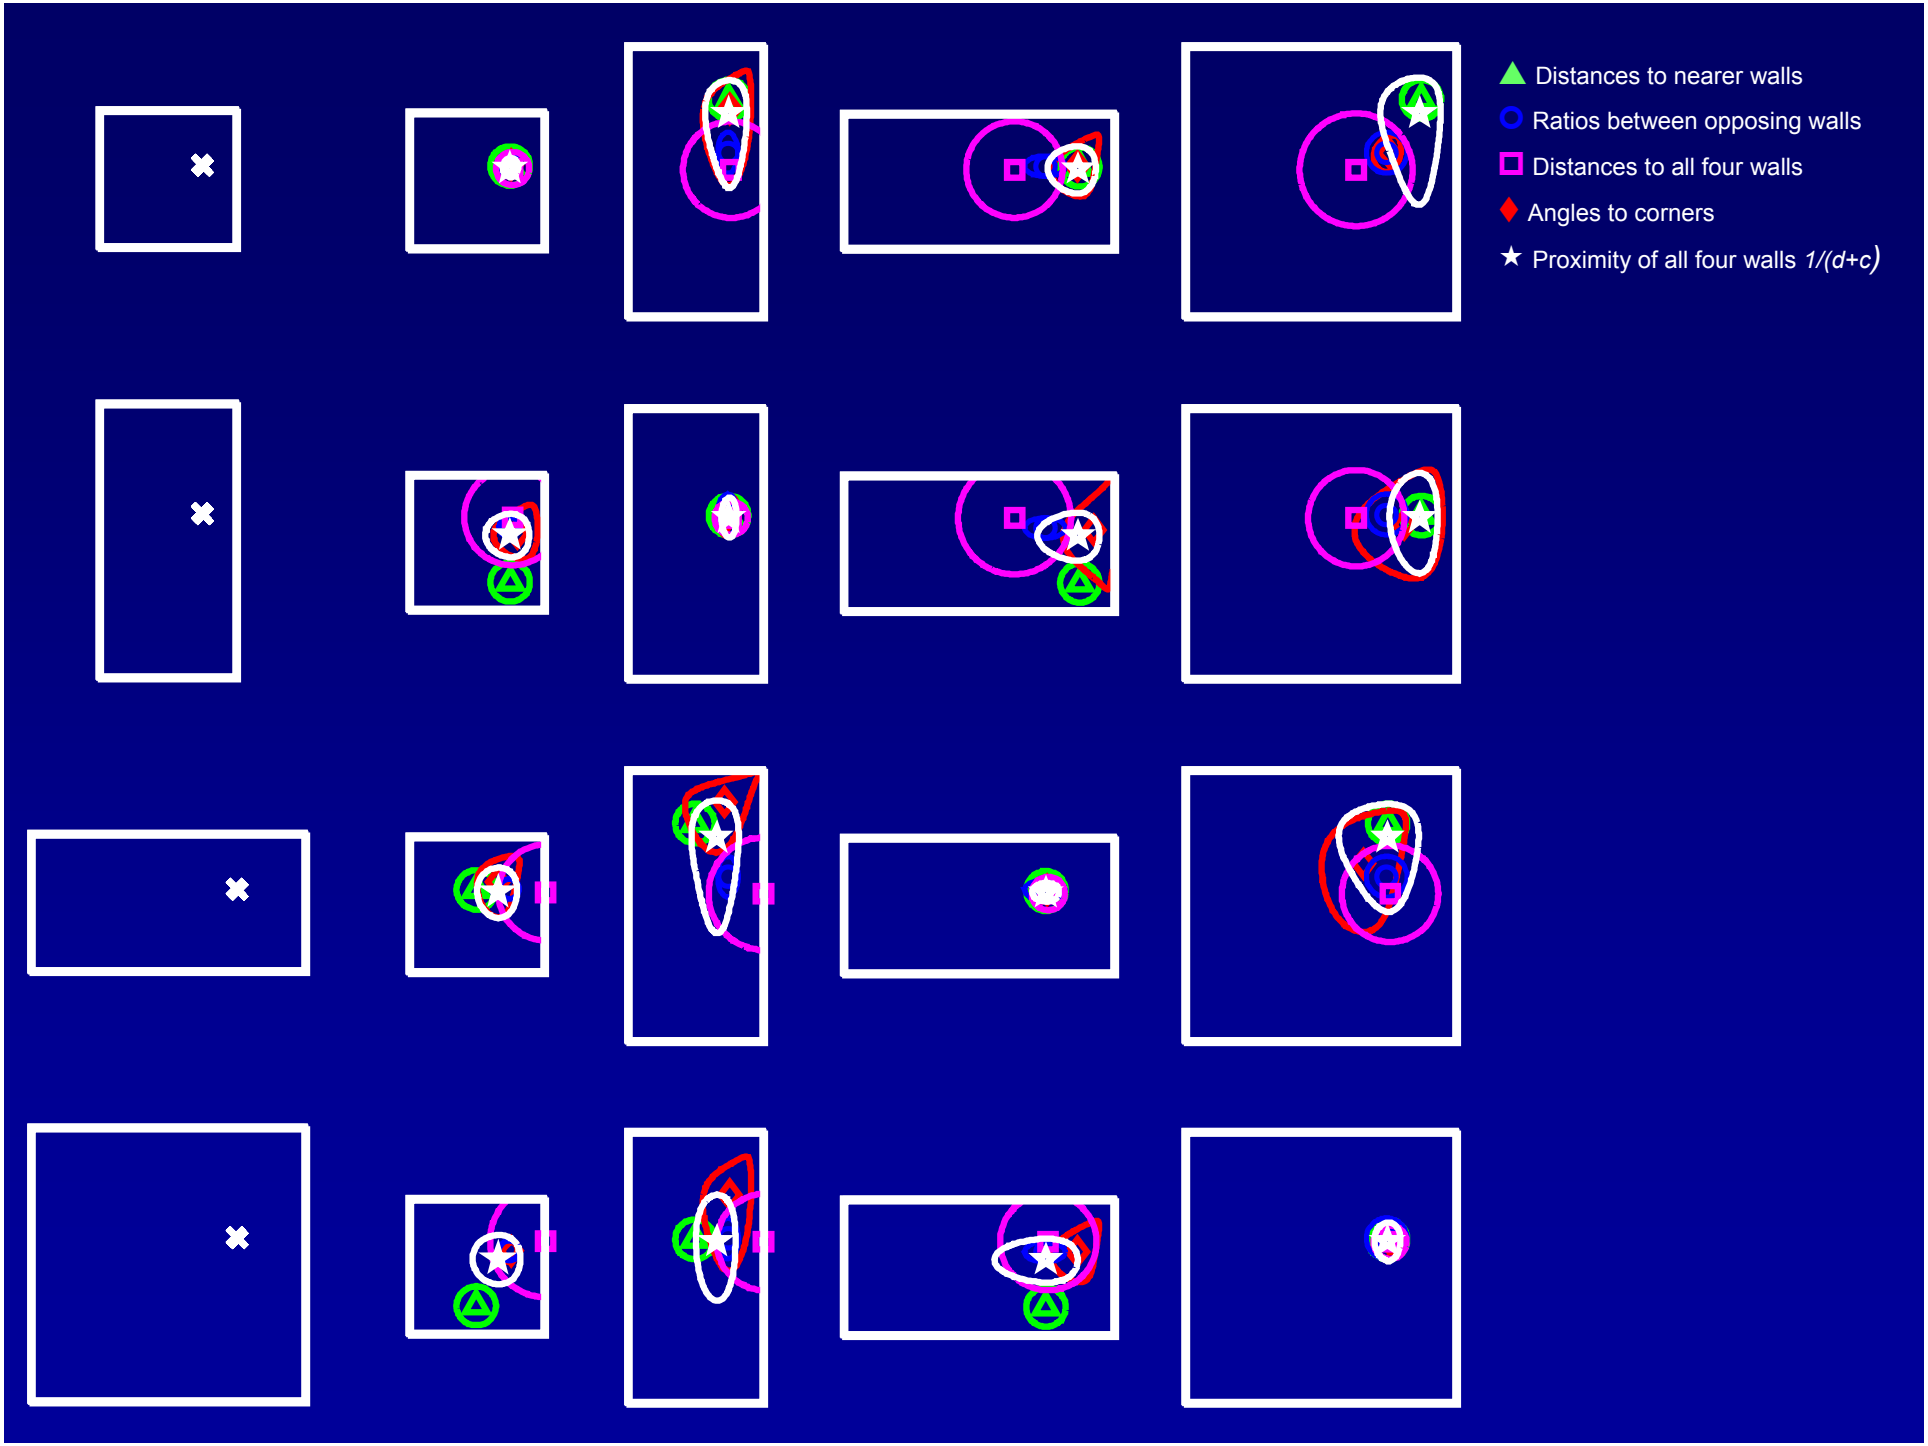
Distances to nearer walls
Ratios between opposing walls
Distances to all four walls
Angles to corners
Proximity of all four walls 1/(d+c)

Evaluating Models

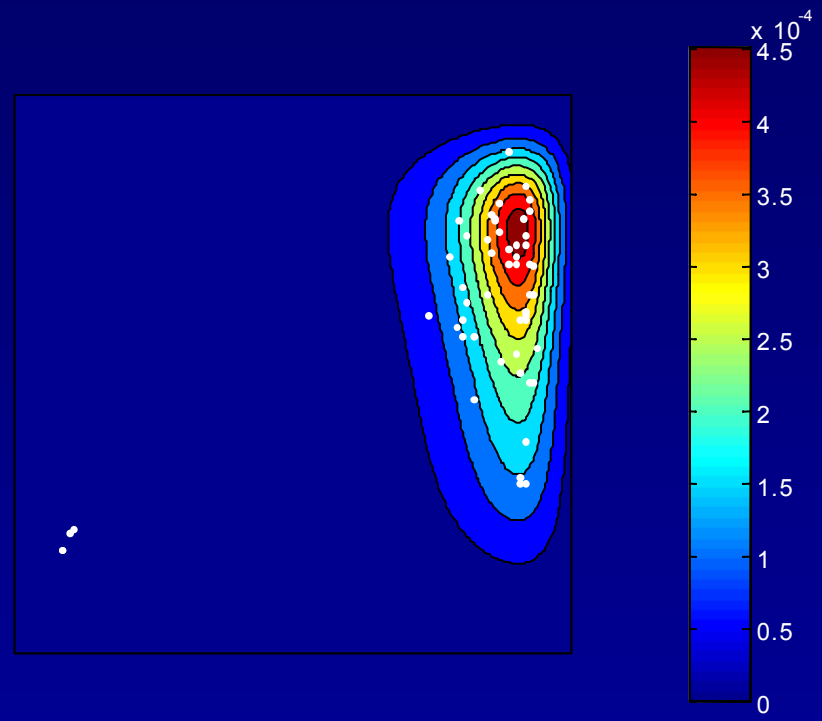

- For each model we calculated the log likelihood of the entire dataset (given the model) varying the sharpness of the distribution to maximize the fit.

Evaluation

Model	Fit log(p(data|model))
boundary proximity	**-2.3134×10^4**
corner angle	**-2.3357×10^4**
fixed-ratio	**-2.3763×10^4**
fixed-distance	**-2.4251×10^4**
absolute-distance	**-2.4981×10^4**
uniform distribution (chance)	**-2.72371×10^4**

Evaluation

Factor by which each data point is more likely under a given model than under a uniform distribution

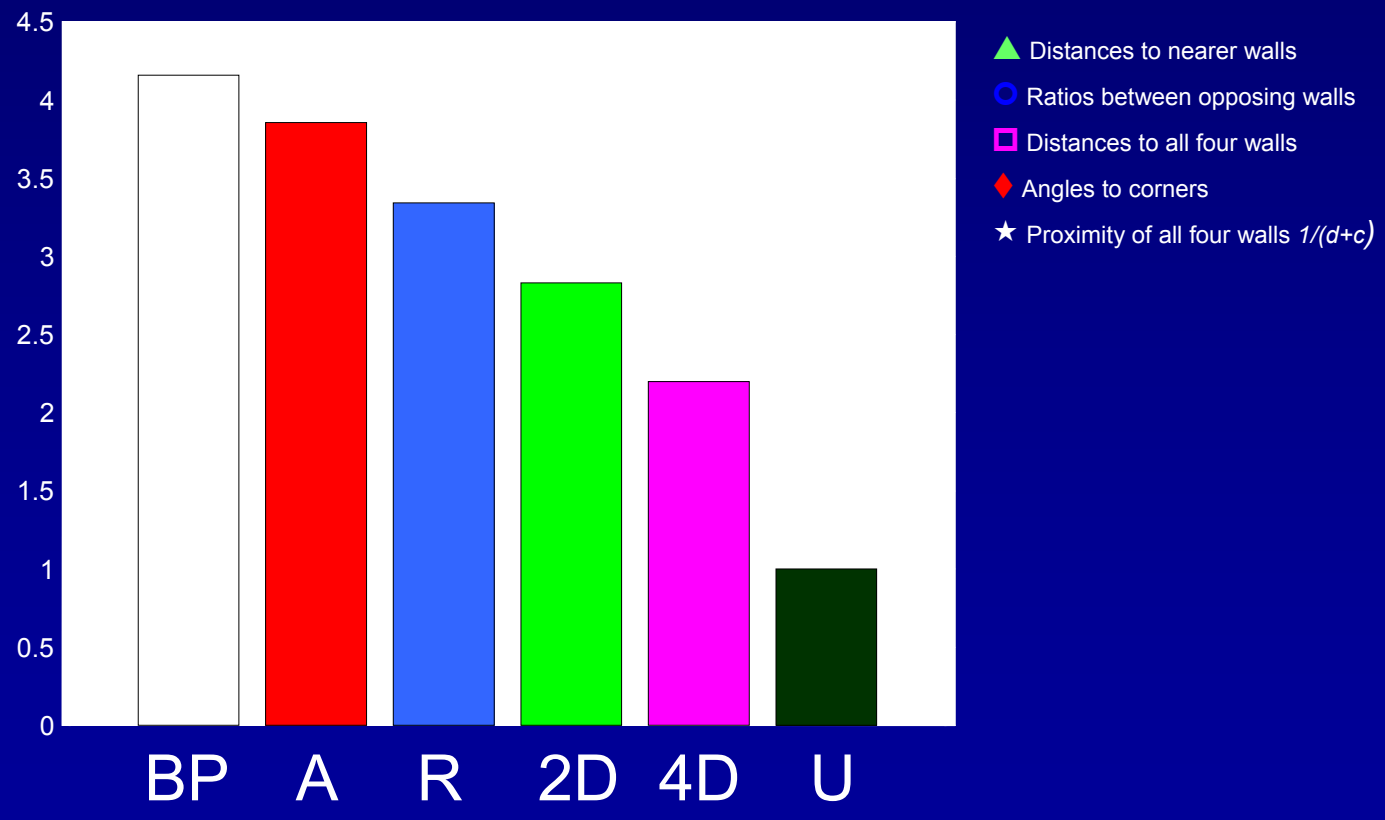

Conclusions

- A complicated pattern of behaviour in our experiment can be explained by assuming a representation based on boundary proximity.
- This was inspired by our model of place fields.
- Future work will be aimed at reconciling the readout mechanism used in the boundary proximity model with the place cell model.

1. Store pattern of activity at the flag

2. Calculate change in activity pattern relative to stored pattern for each location in new environment

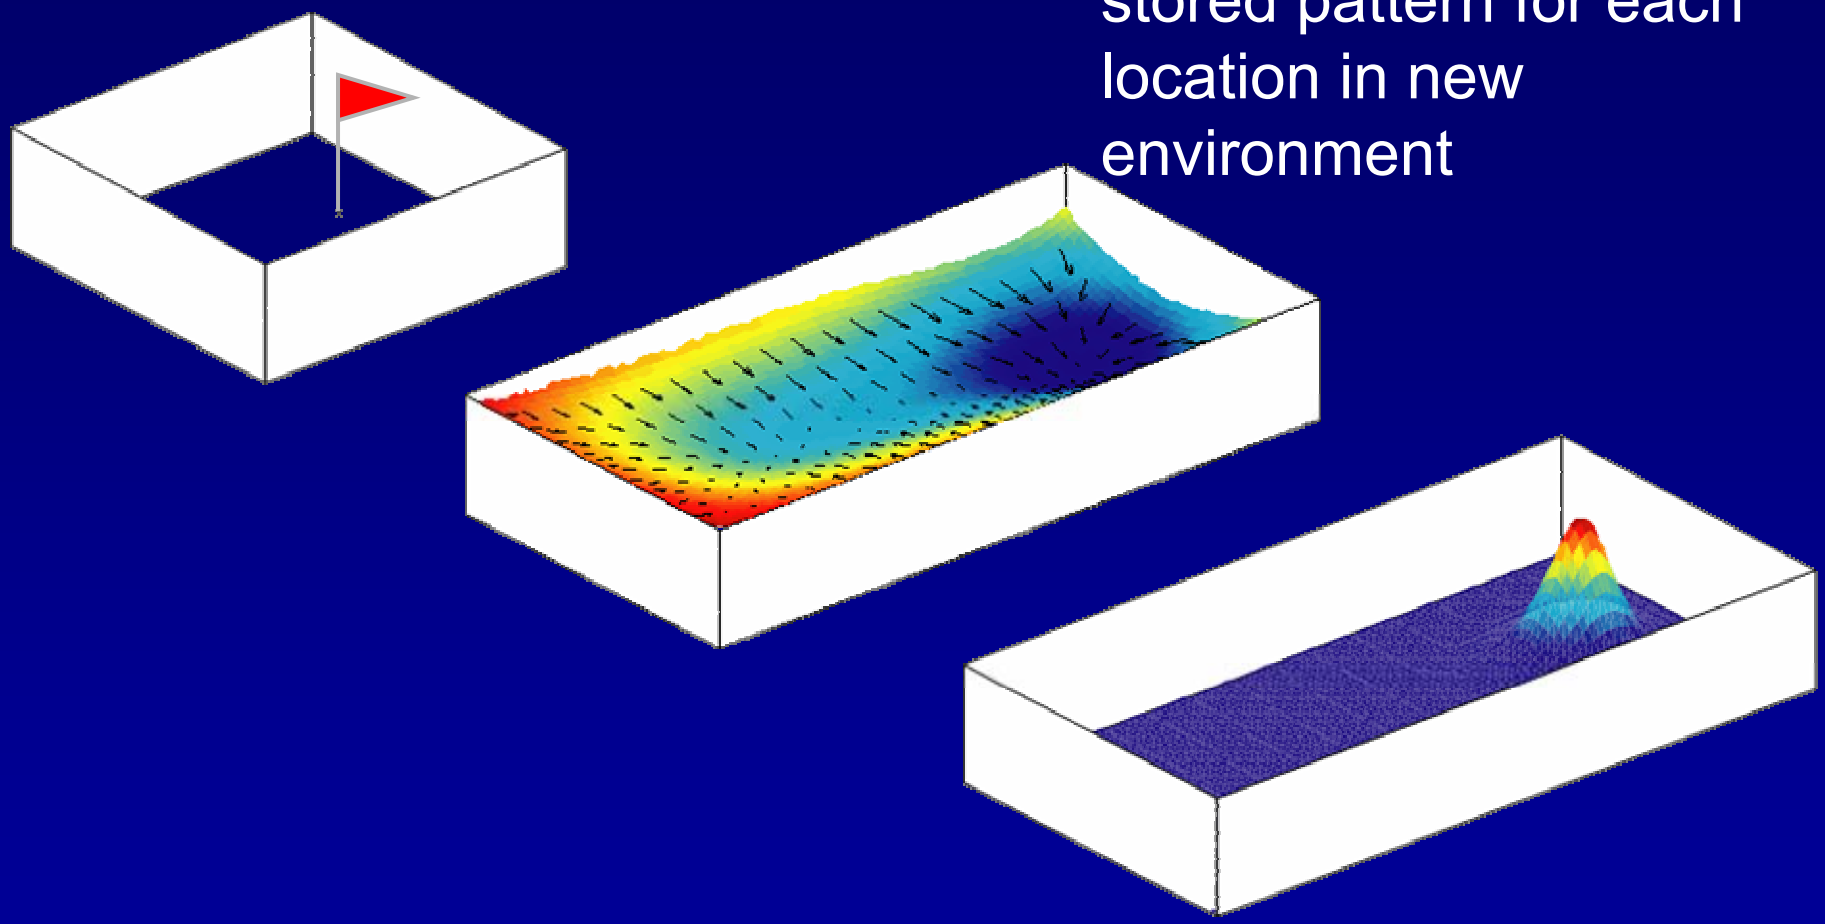

3. Probability of a response depends on the similarity of pattern

Future Directions

- The model now deals with environments of arbitrary shape
- Behavioural tests e.g.,
  - Barrier
  - Rotations
  - Forced Choice
- Neuroimaging tests
  - fMRI signal change as subject moves around
- Modelling navigation in complex large scale environments

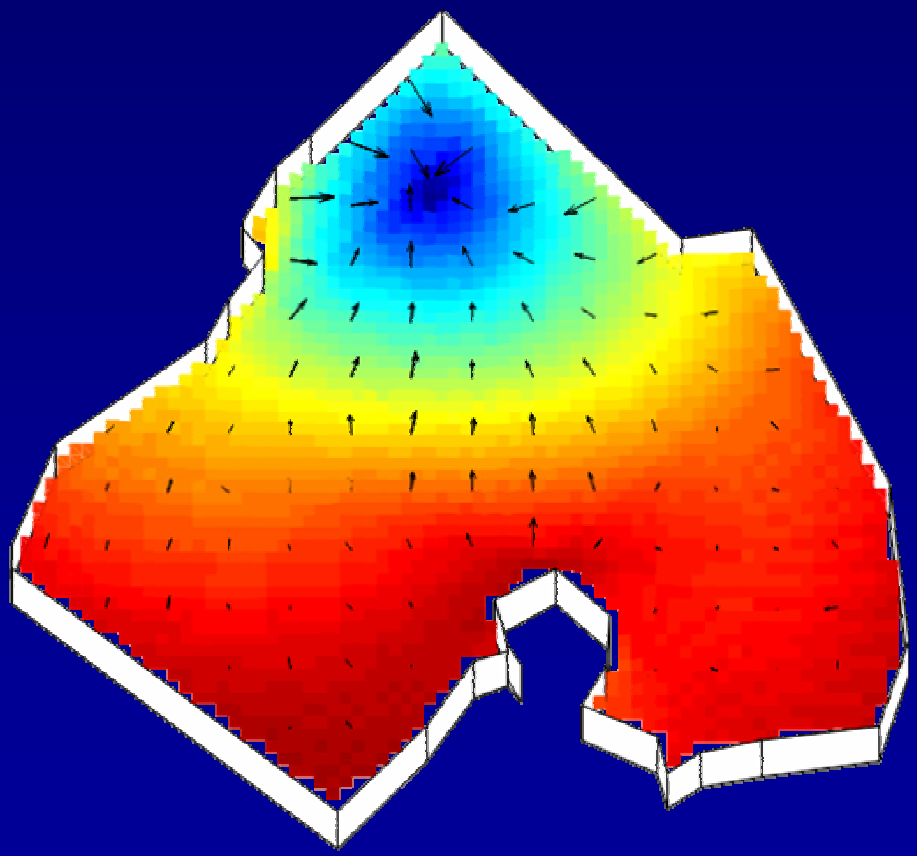

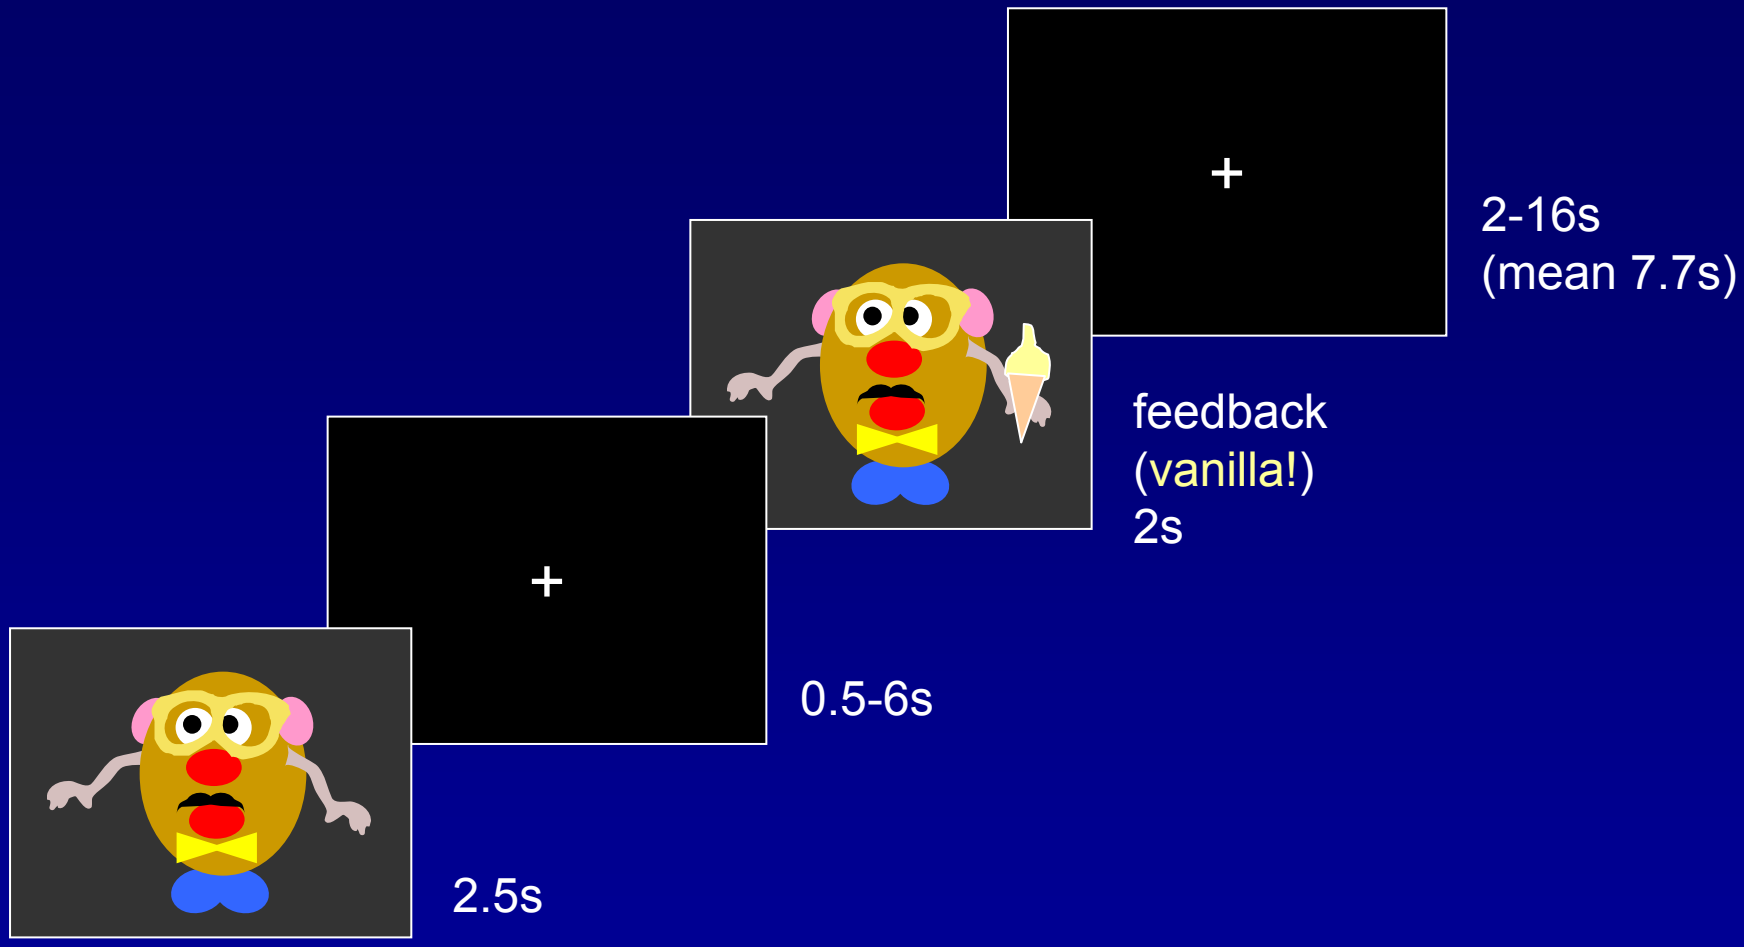

Rodriguez et al. (2005, Hum Brain Mapp.)

Table I. Stimuli as well as objective/empirical/model-based probabilities for the chocolate (as opposed to vanilla) outcome

	Features				P(choc)		
Stimulus	Moustache	Hat	Glasses	Bowtie	Objective	Subjects' responses	Model outputs
1	0	0	0	1	0.94	0.81	0.89
2	0	1	0	0	0.75	0.56	0.63
3	0	1	0	1	0.91	0.84	0.83
4	0	0	1	0	0.12	0.32	0.2[illegible]
5	0	0	1	1	0.88	0.72	0.70
6	0	1	1	0	0.50	0.30	0.49
7	0	1	1	1	0.94	0.62	0.75
8	1	0	0	0	0.07	0.13	0.17
9	1	0	0	1	0.50	0.43	0.50
10	1	1	0	0	0.17	0.11	0.29
11	1	1	0	1	0.71	0.32	0.60
12	1	0	1	0	0.07	0.14	0.20
13	1	0	1	1	0.60	0.38	0.40
14	1	1	1	0	0.14	0.12	0.23

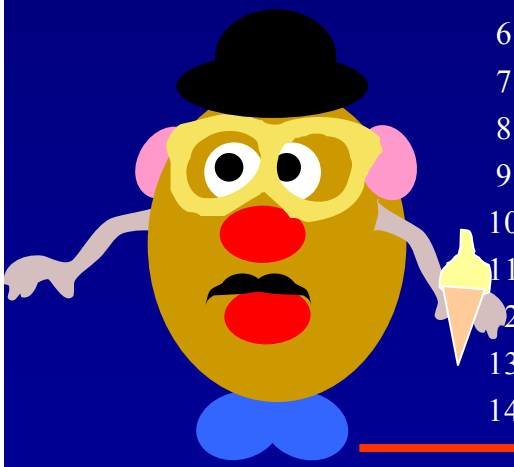

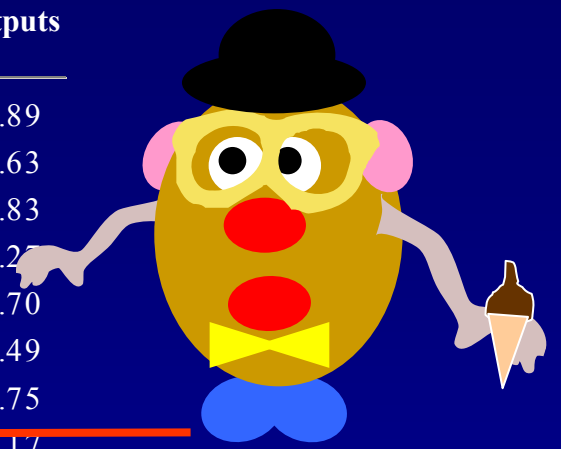

Classification Learning: Behavioural Data

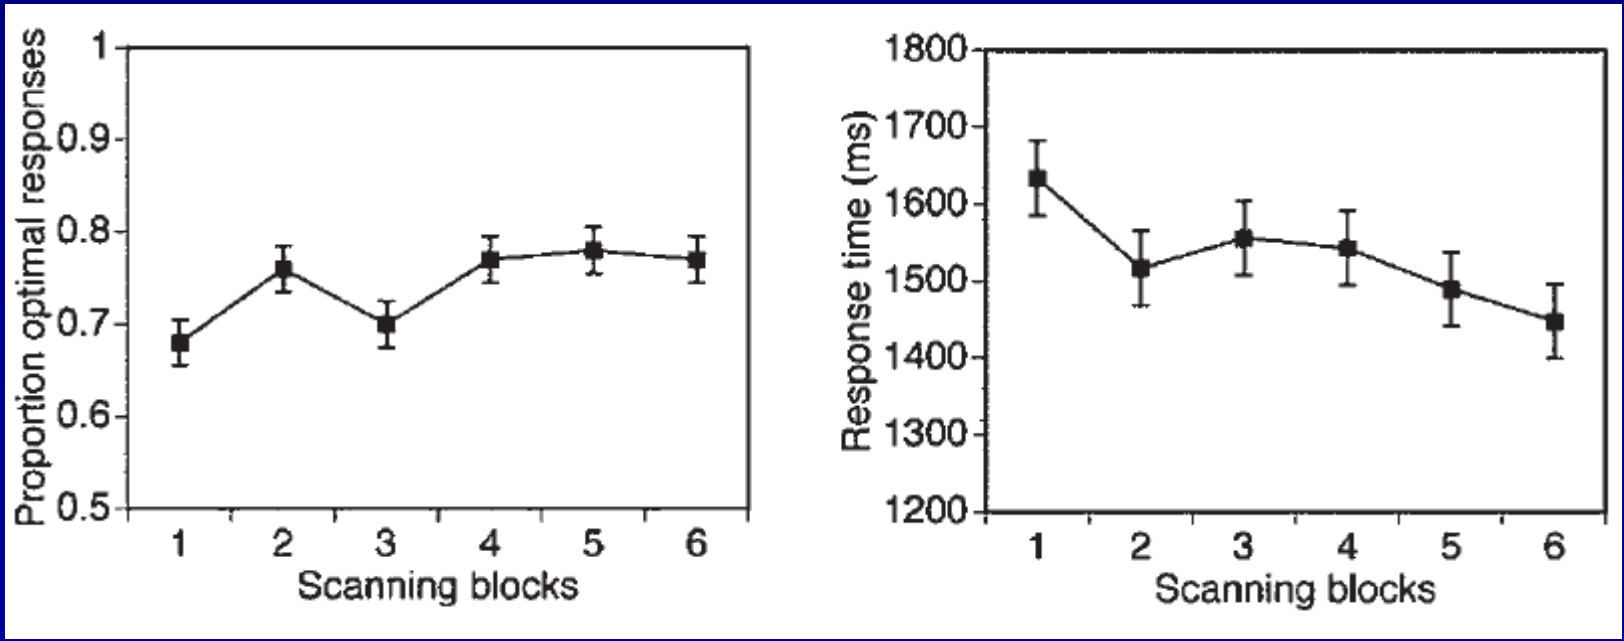

25 trials per block with feedback

Classification Learning
Simple Model

input units
representing features

output units
representing responses

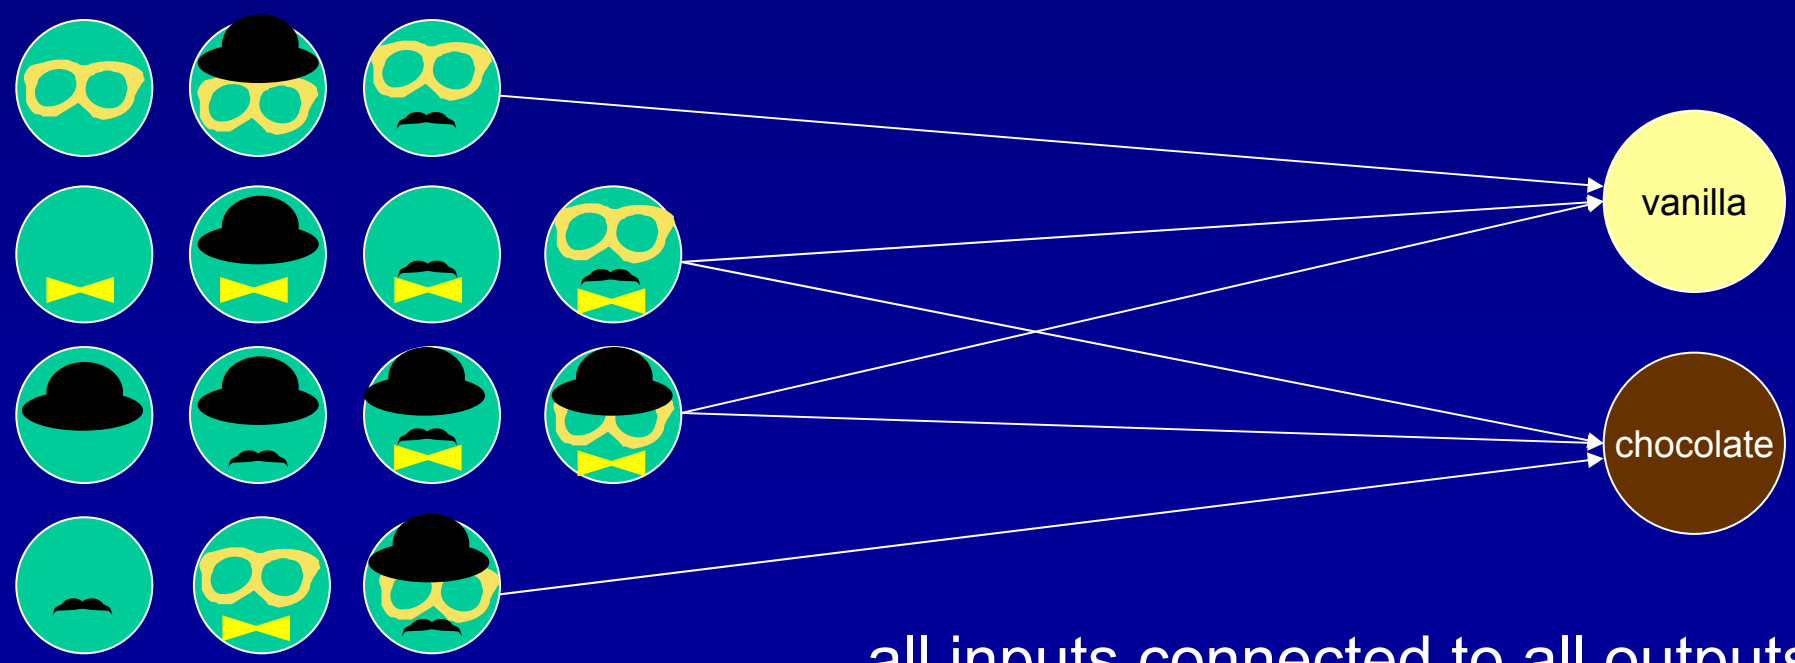

all inputs connected to all outputs

Here (at the beginning of learning) the pattern of features present in the stimulus activates one response (chocolate) slightly more than the other (vanilla).

However, the correct answer turns out to be vanilla.

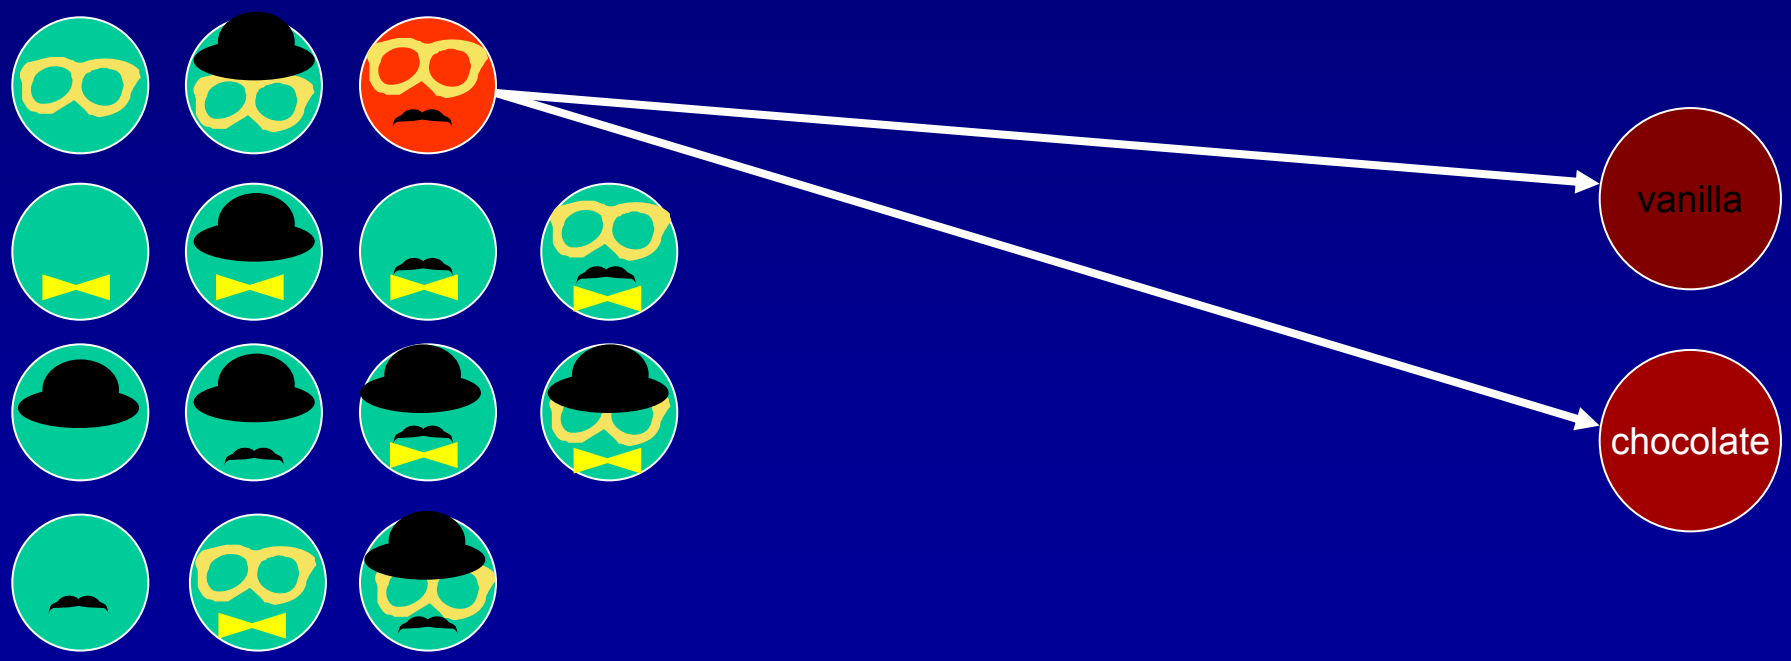

In this example if the correct answer was vanilla the learning rule would strengthen the connection from  to the vanilla response

and weaken the connection from  to the chocolate response.

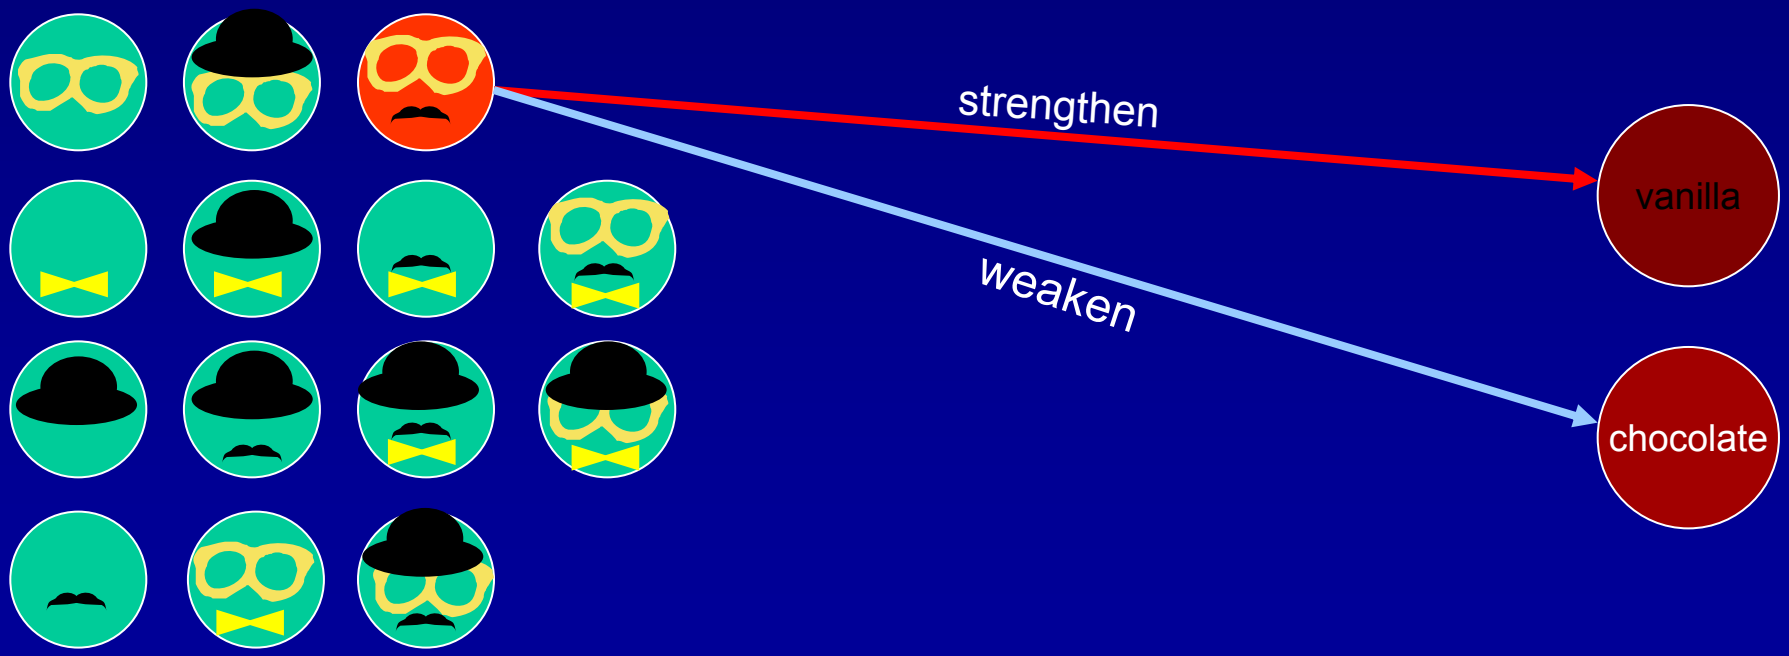

Classification Learning Simple Model

- After every pattern, the response is compared to the correct answer.
- The connections are strengthened or weakened according to how much each contributes to the error.
- Responses get more and more accurate.

Brain basis of reinforcement learning

- It is hypothesized that the error signal used in this kind of learning originates in midbrain dopiminergic neurons.
- These neurons relay signals to structures in the ventral striatum
- The signal is thought to convey the discrepancy between the observed and expected outcomes of an action.

Rodriguez et al. (2005)

- Used a parametric analysis to see whether ventral striatal activity was predicted by the error signal in the model.

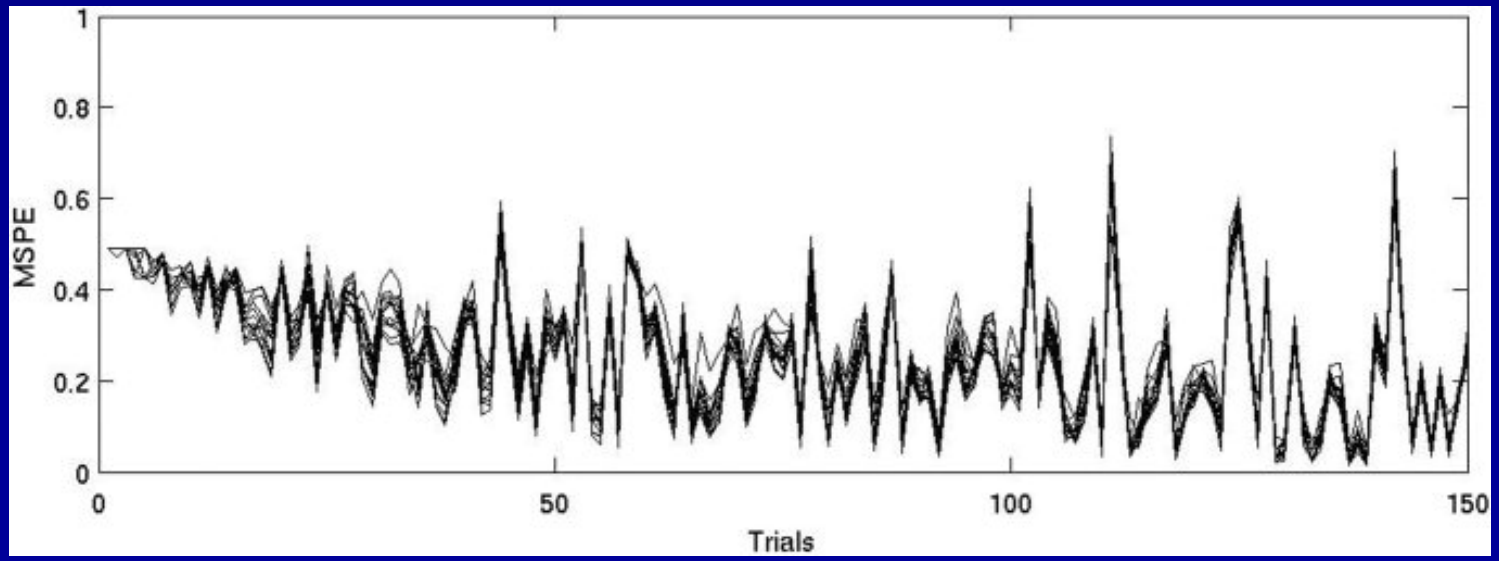

Feedback events: parametric modulation

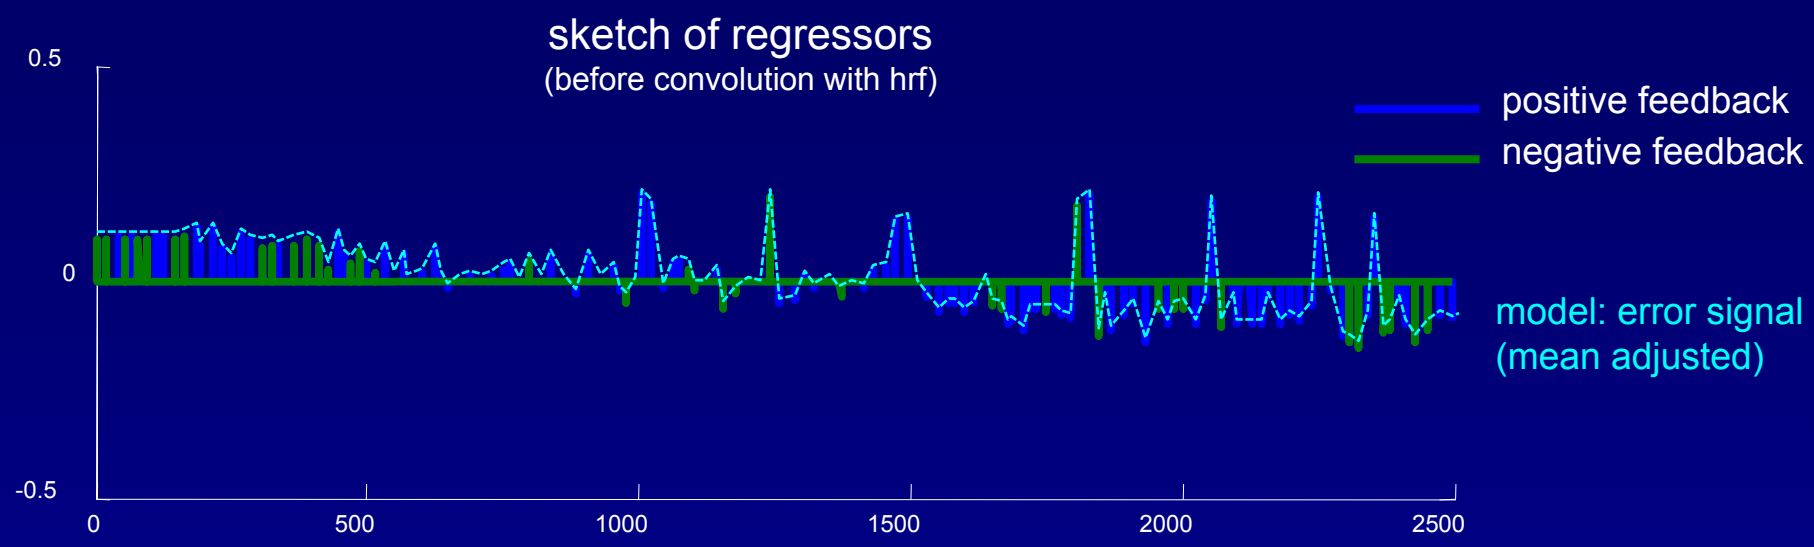

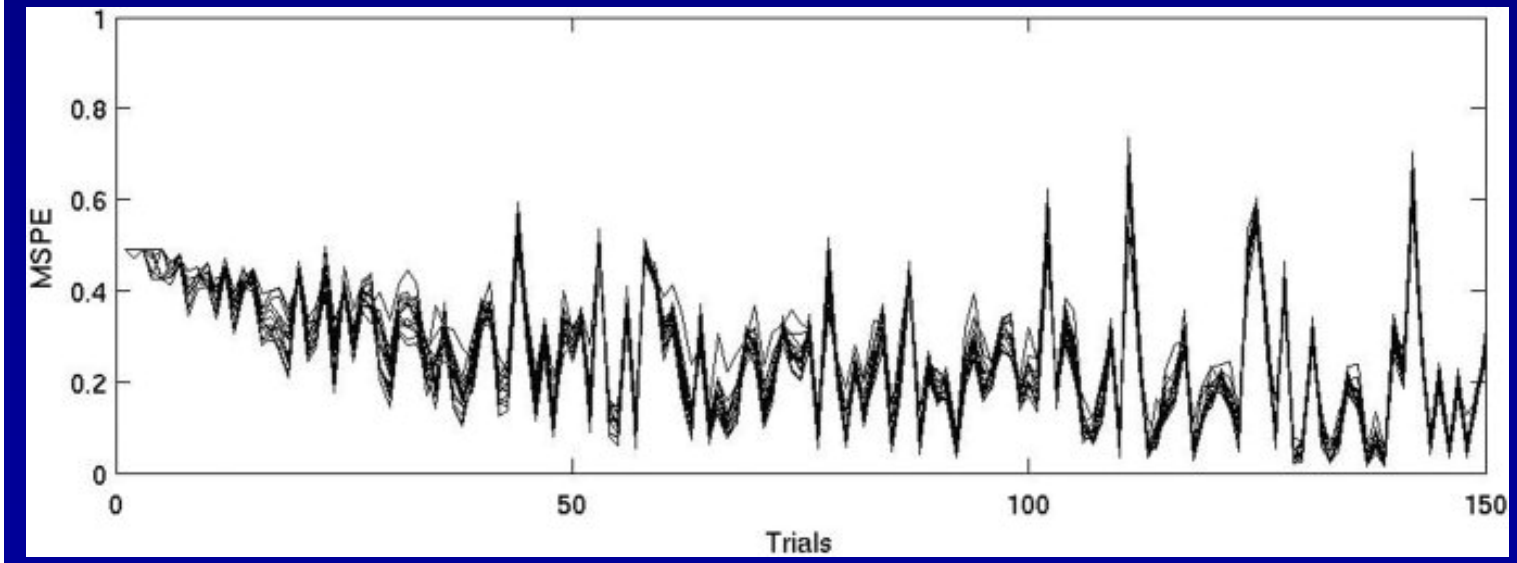

actual model predicted error signal (per subject)

The error signal gets lower on average.

There is also more positive feedback.

The odd surprise results in high error signal.

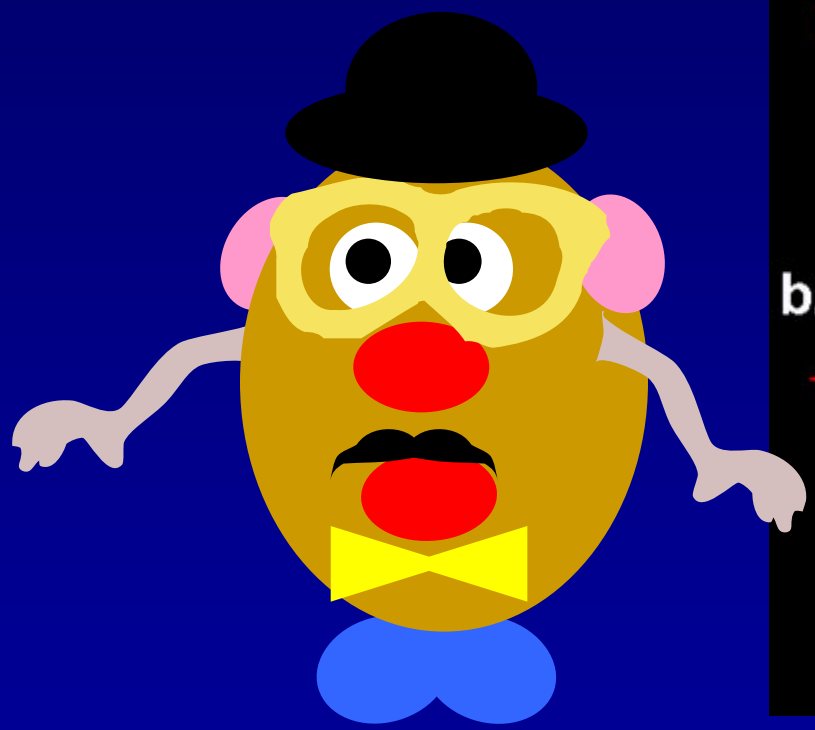

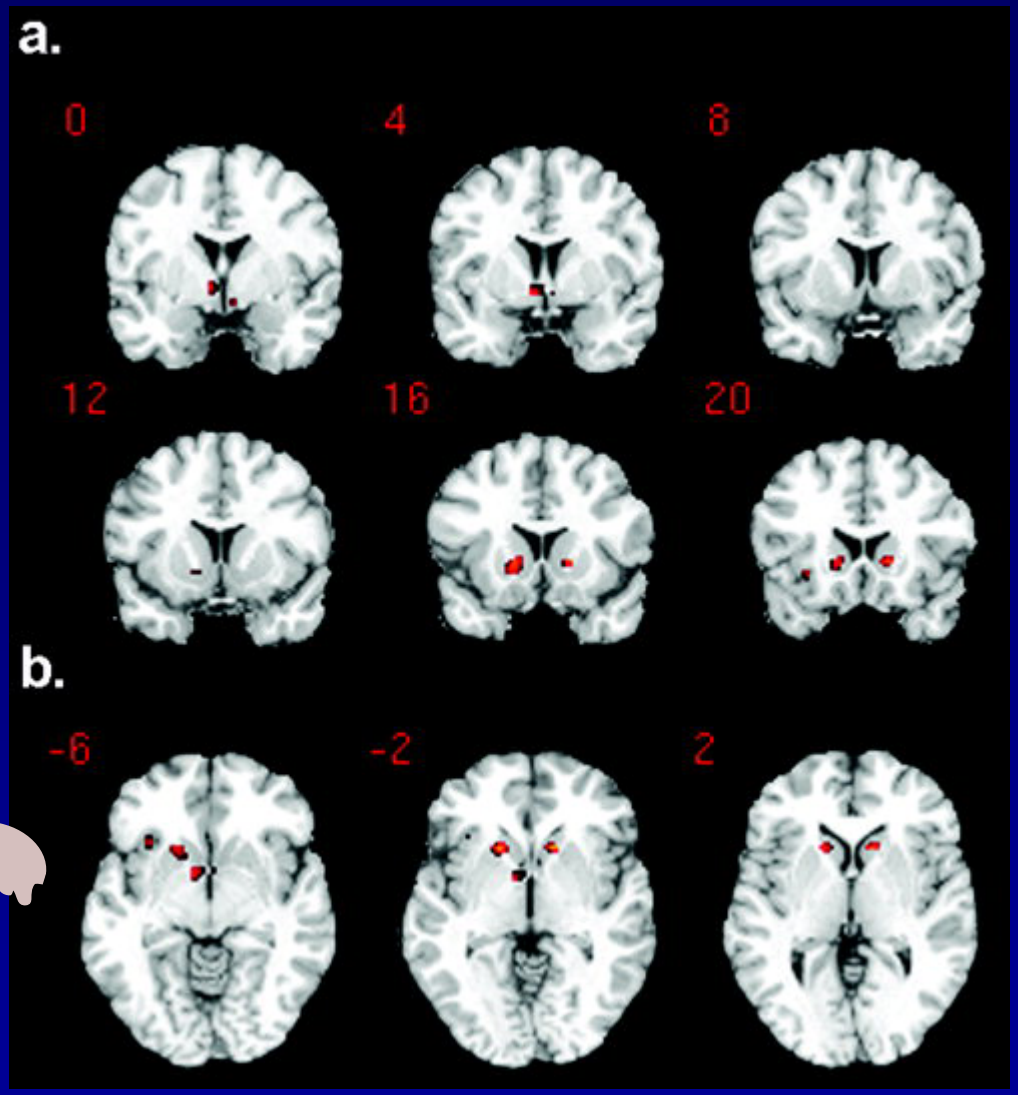

Nucleus accumbens, caudate/putamen activity correlated with model's prediction error on negative feedback trials

Overall

- Experiments inspired by classic neurophysiological paradigms combined with simple models of neural representation can provide useful insights into the processes underlying complex human behaviours.
- Computational models can address neuroimaging data.
- Modelling can provide a quantitative link between neuroscience and psychology.

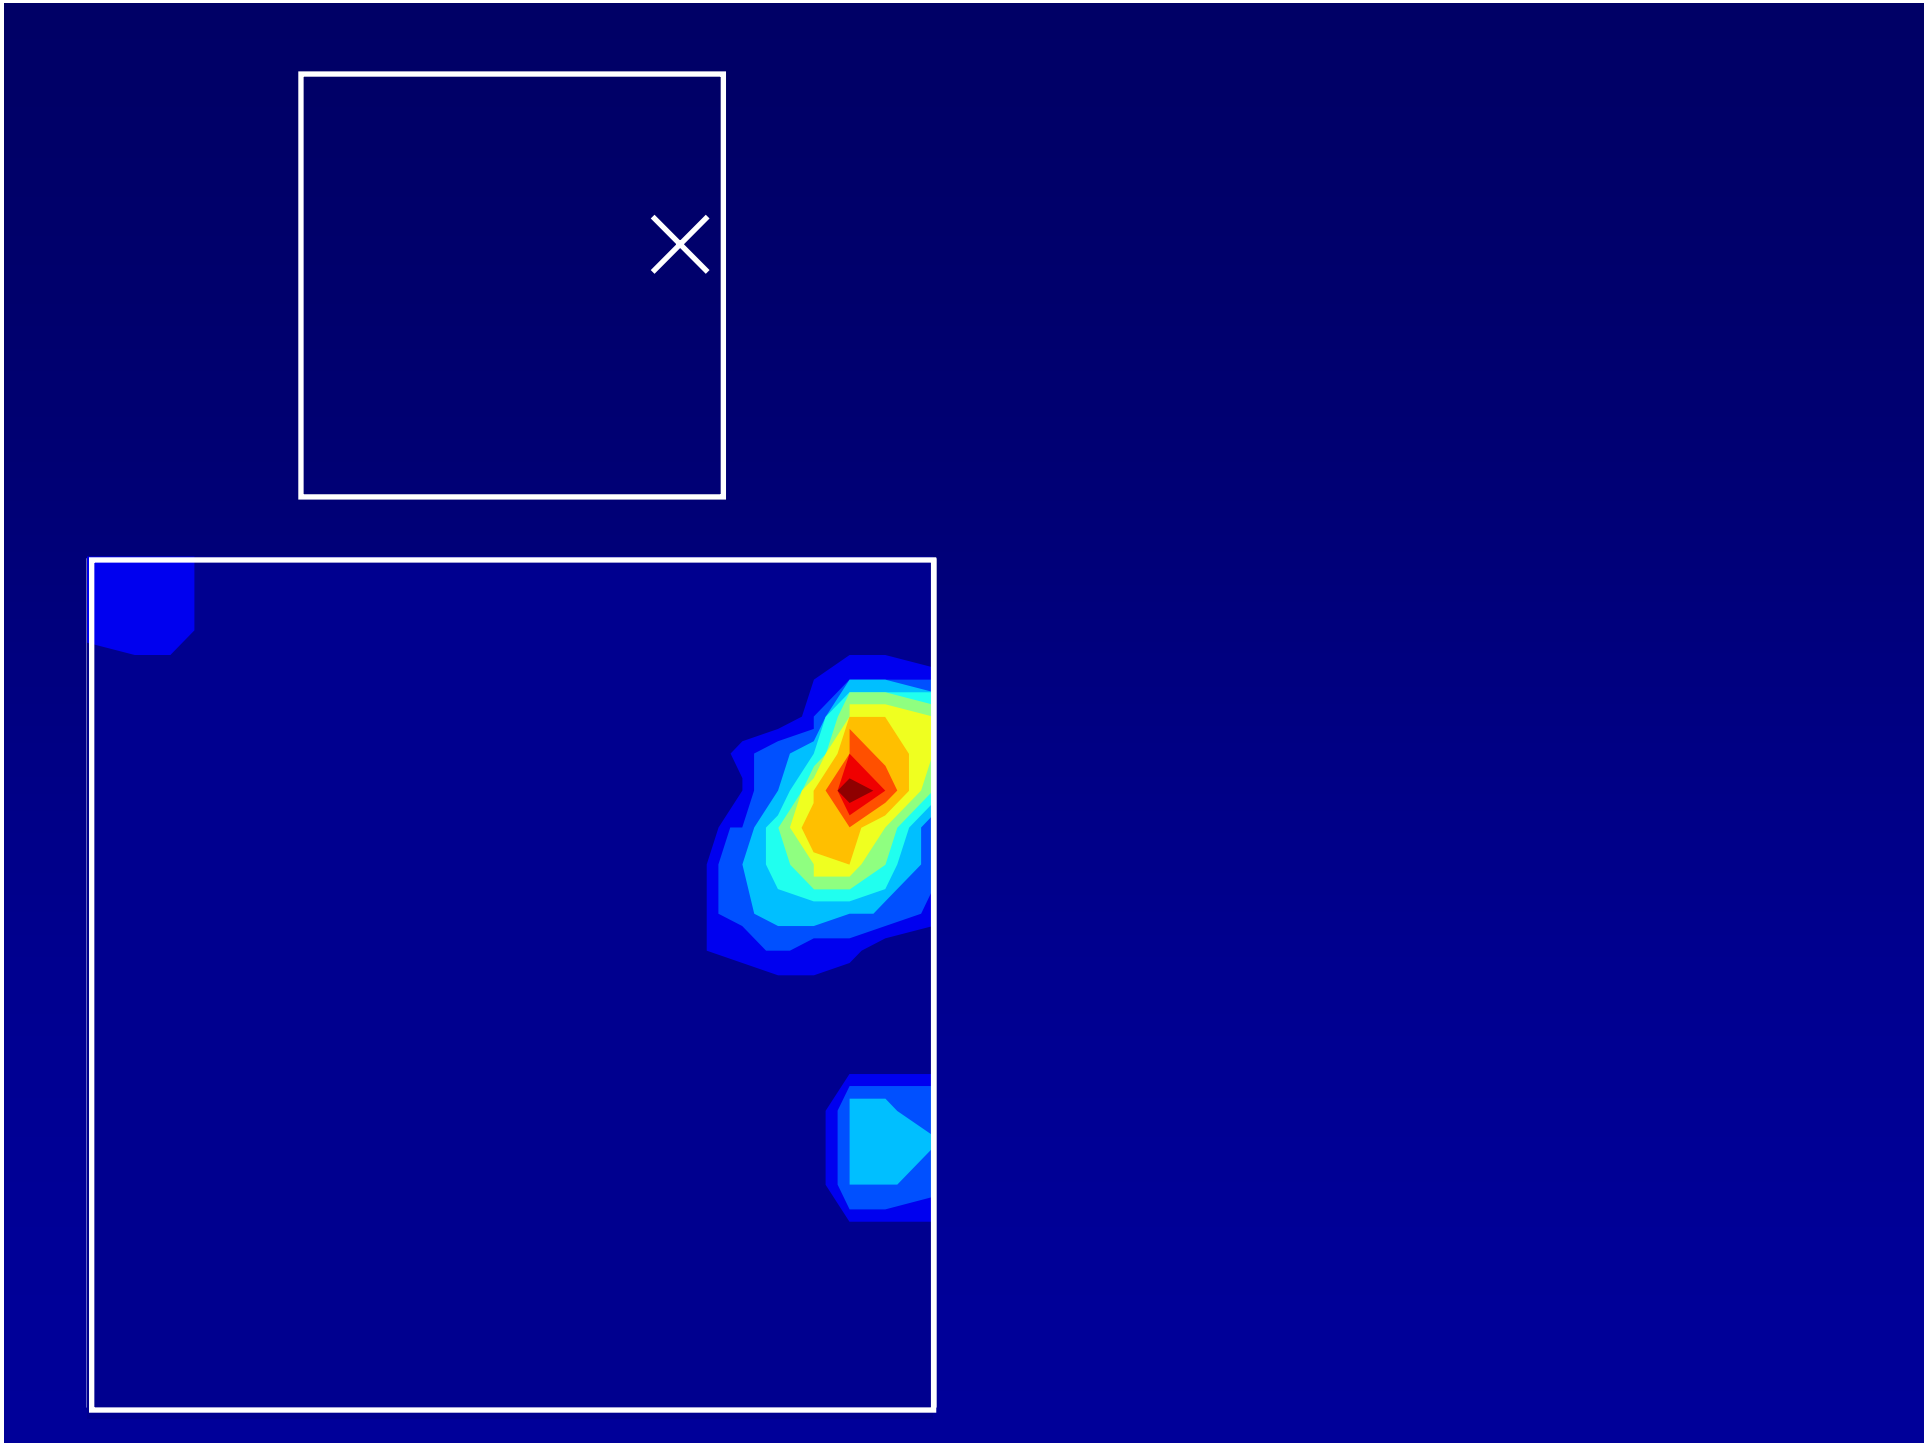

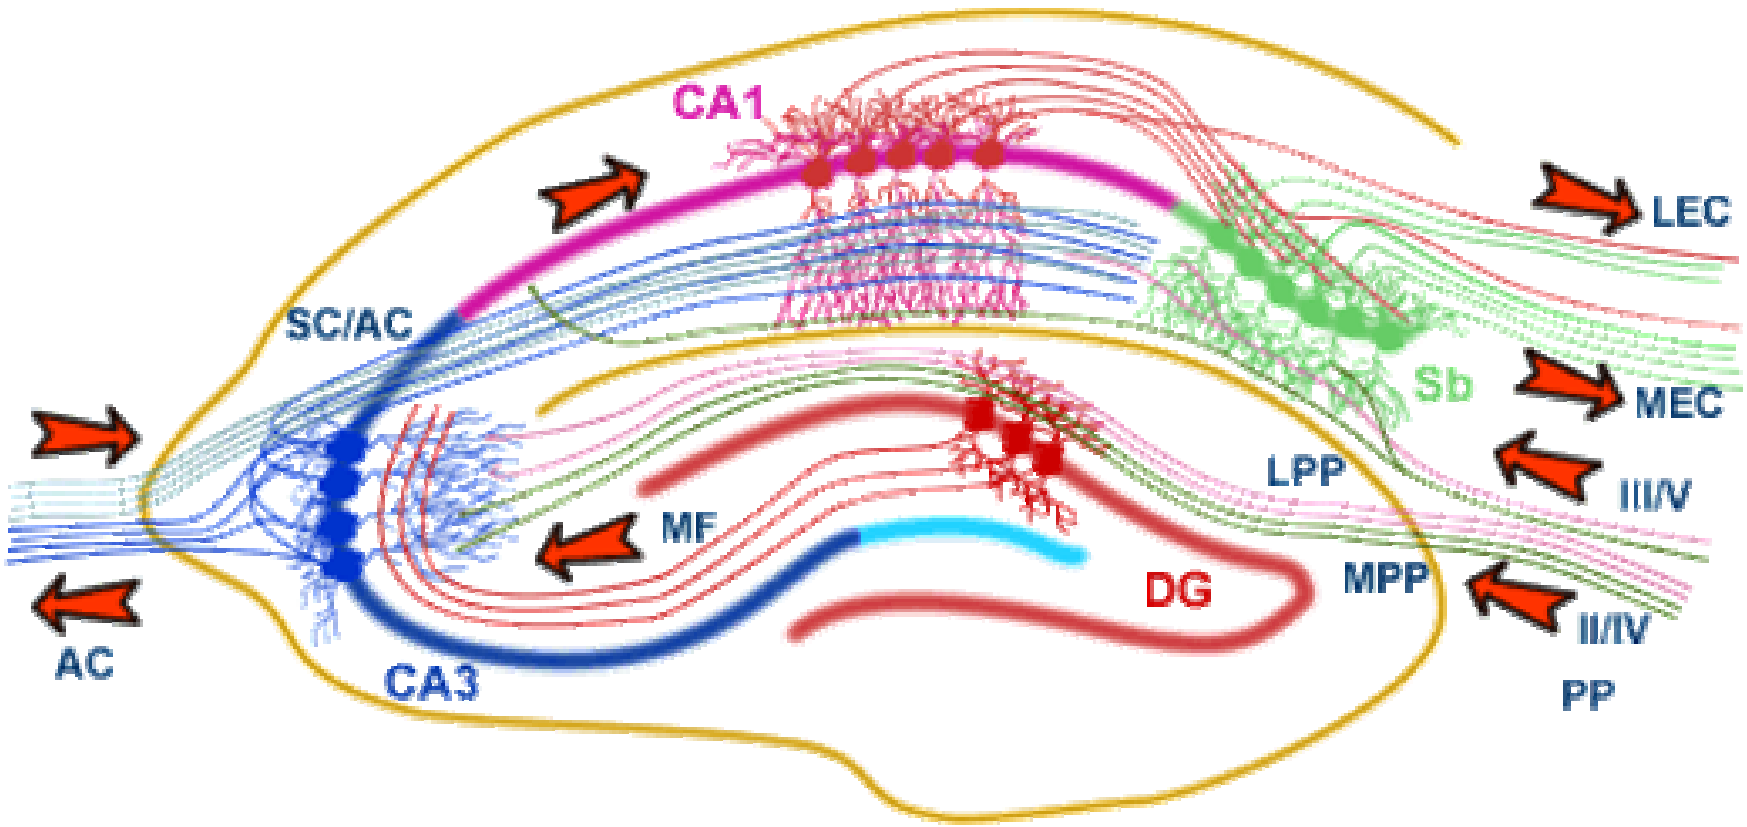

The Hippocampal Network: The hippocampus forms a principally uni-directional network, with input from the Entorhinal Cortex (EC) that forms connections with the Dentate Gyrus (DG) and CA3 pyramidal neurons via the Perforant Path (PP - split into lateral and medial). CA3 neurons also receive input from the DG via the mossy fibres (MF). They send axons to CA1 pyramidal cells via the Schaffer Collateral Pathway (SC), as well as to CA1 cells in the contralateral hippocampus via the Associational Commissural pathway (AC). CA1 neurons also receive input directly from the Perforant Path and send axons to the Subiculum (Sb). These neuron in turn send the main hippocampal output back to the EC, forming a loop.

Boundary Proximity Model

Contraction

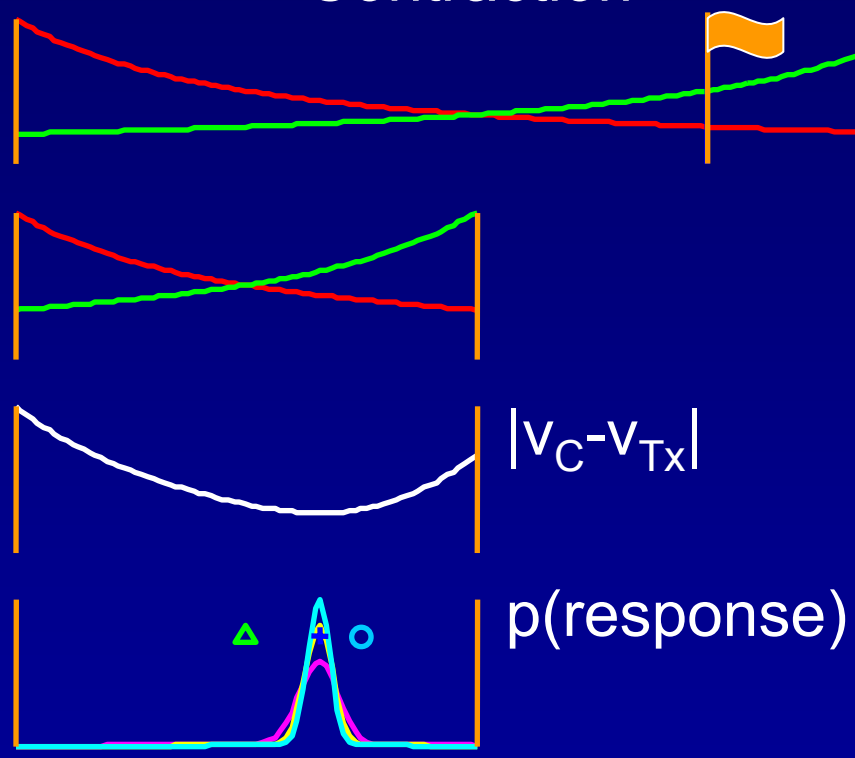

Boundary Proximity Model

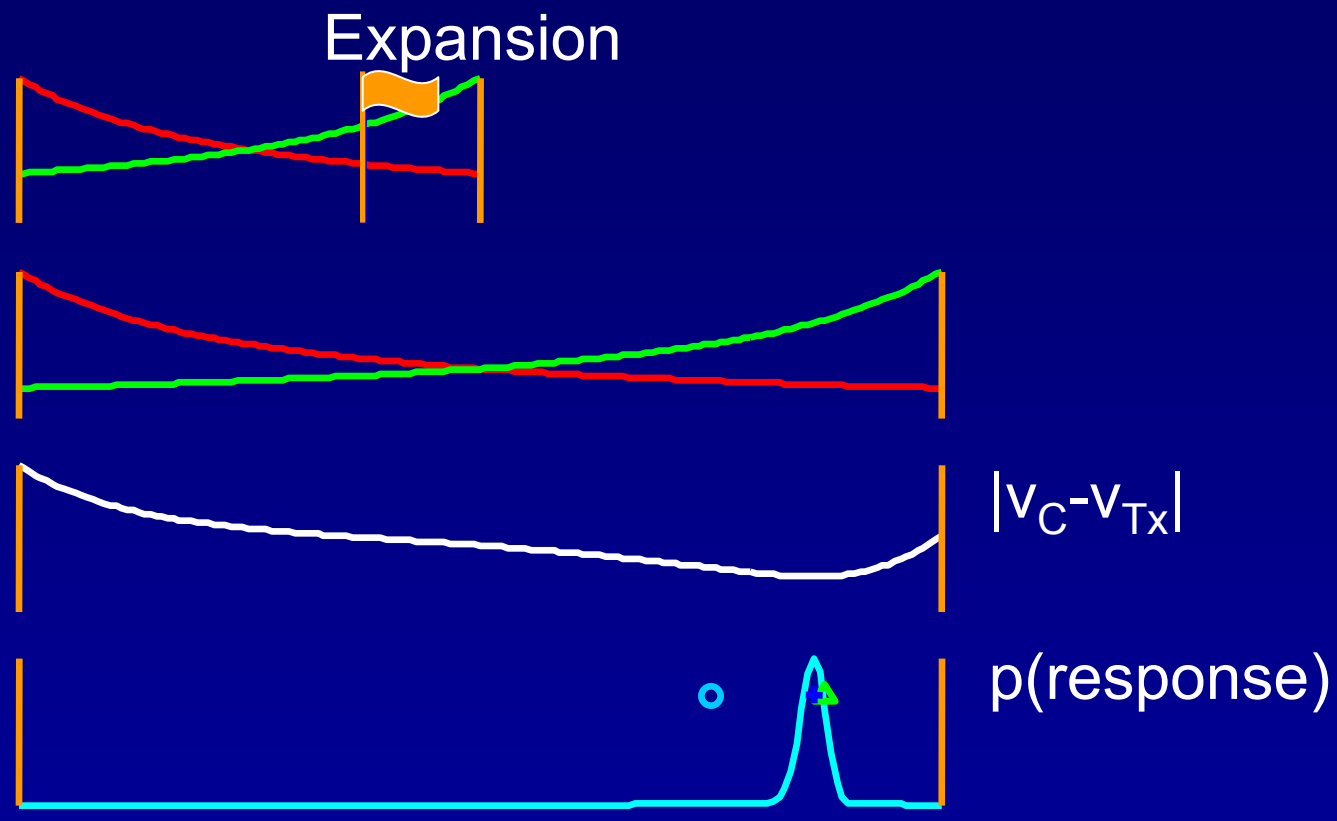

Boundary Proximity Model

Cue Near Centre - Contraction

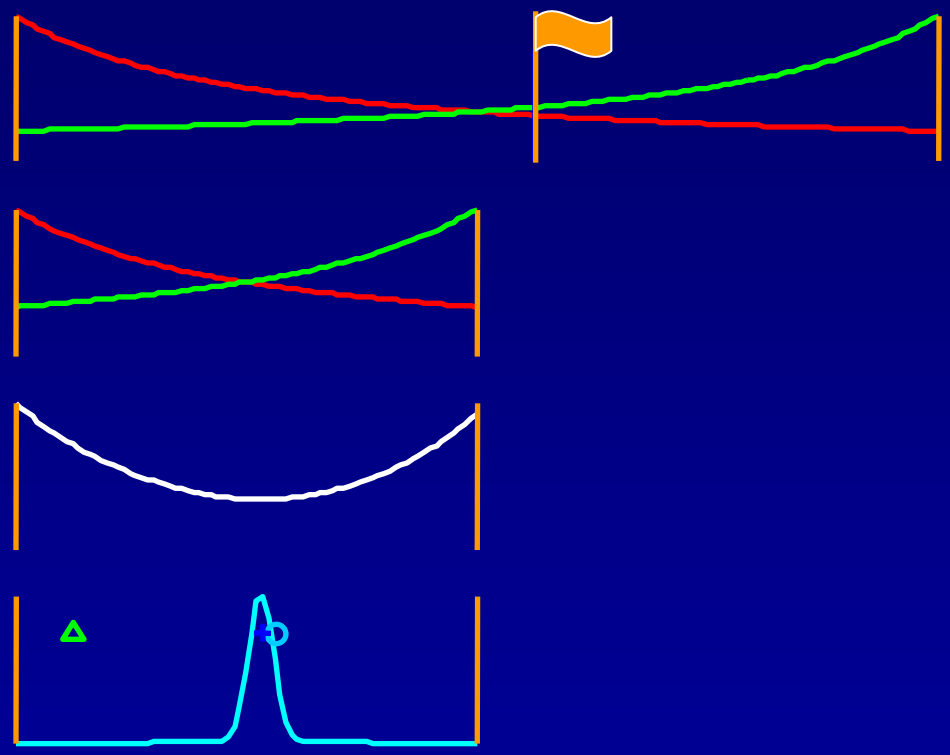

Boundary Proximity Model

Cue Near Edge - Contraction

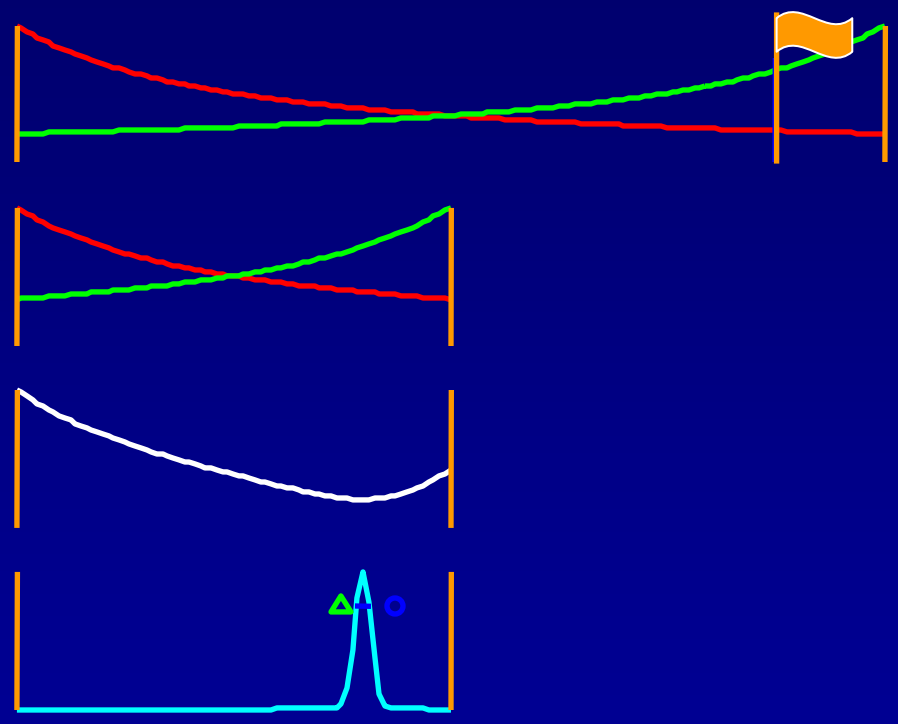

Boundary Proximity Model

Cue Near Centre - Expansion

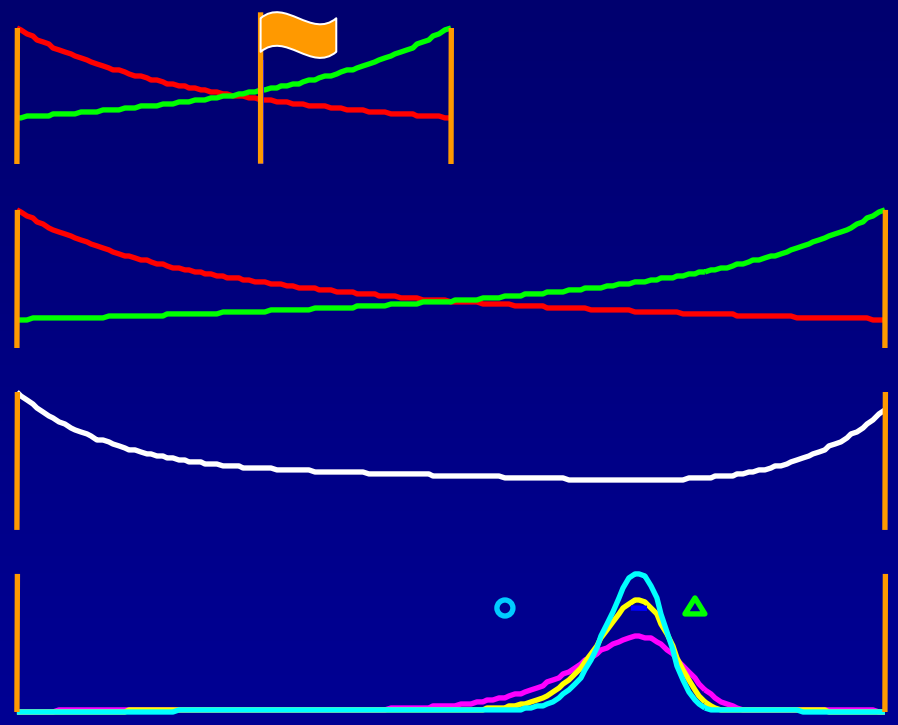

Boundary Proximity Model

Cue Near Edge - Expansion