Making Sense from Models

Modelling the Senses with Multi-net Systems

Matthew Casey
27 November 2007

www.cs.surrey.ac.uk

Motivation

Can knowledge of biological systems produce a step-change in the capability of artificial systems?

Can using artificial systems improve our knowledge of biological systems?

Current Focus

- Develop models of sensory processing
  - Task adaptation in human vision
  - Multi-sensory integration in the superior colliculus
- Explore levels of abstraction
  - How complex does a model have to be?
  - How biologically plausible?
  - Simple behavioural models (rate-coded) or more biologically plausible systems (pulse-coding)
- Develop strong theoretical foundations
  - For modelling with multi-net systems (partially ordered sets – with Mike Shields)
  - Understanding multiple classifier performance (game theory – with Richard Zanibbi)

Overview

- Modelling visual categorical perception
  - Work with Dr Paul Sowden (Psychology, Surrey)
  - Help improve our understanding of natural cognitive systems
  - Explore levels of abstraction: simple behavioural models (rate-coded)
- Categorical perception (CP)
  - Traditionally, once developed, low level vision is thought to be static
  - However experiments suggest that low level vision adapts on a task by task basis

Visual CP

Do these belong to the same or a different category?

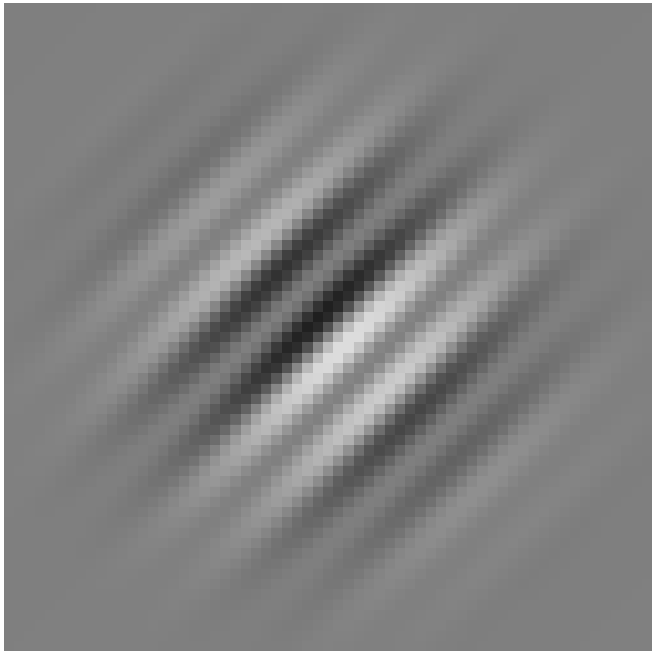

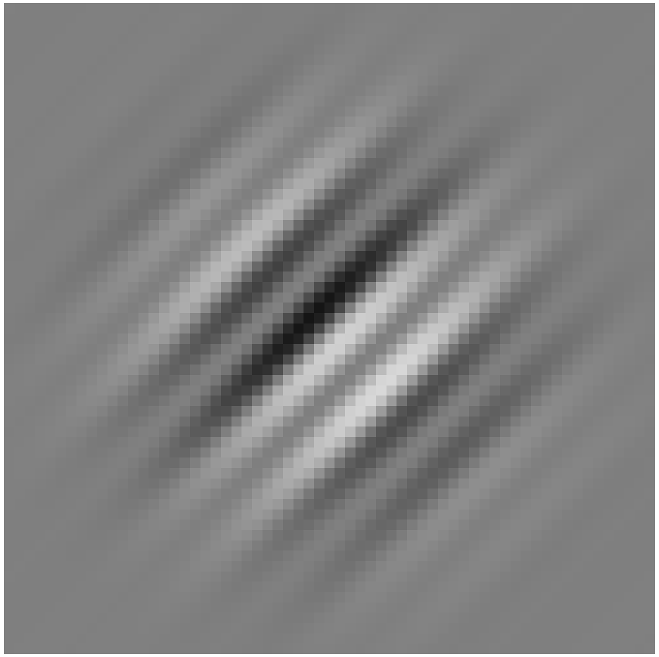

Visual CP

Do these belong to the same or a different category?

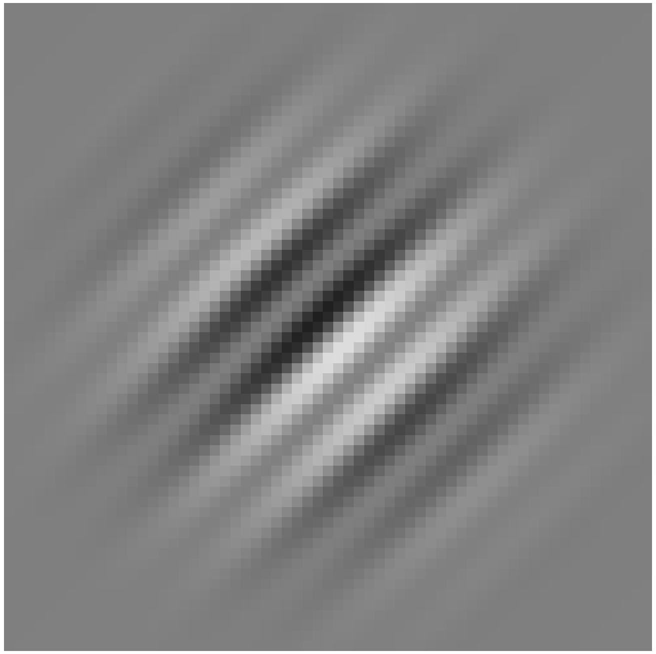

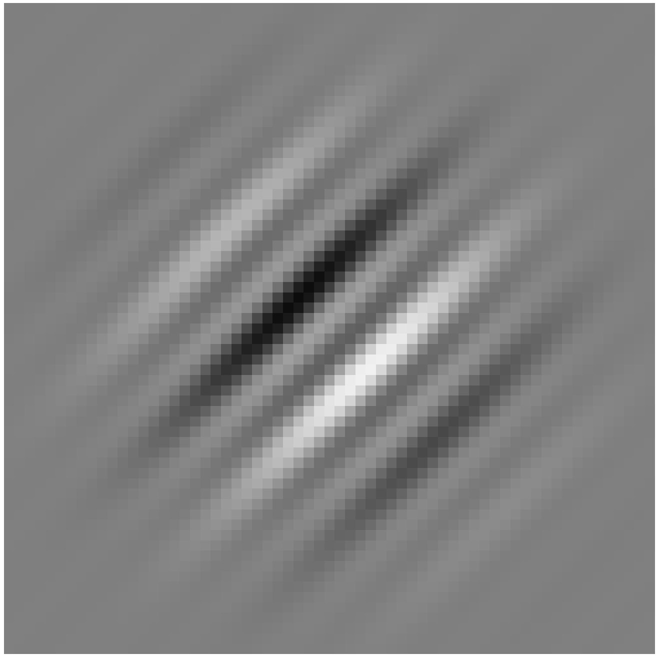

Visual CP

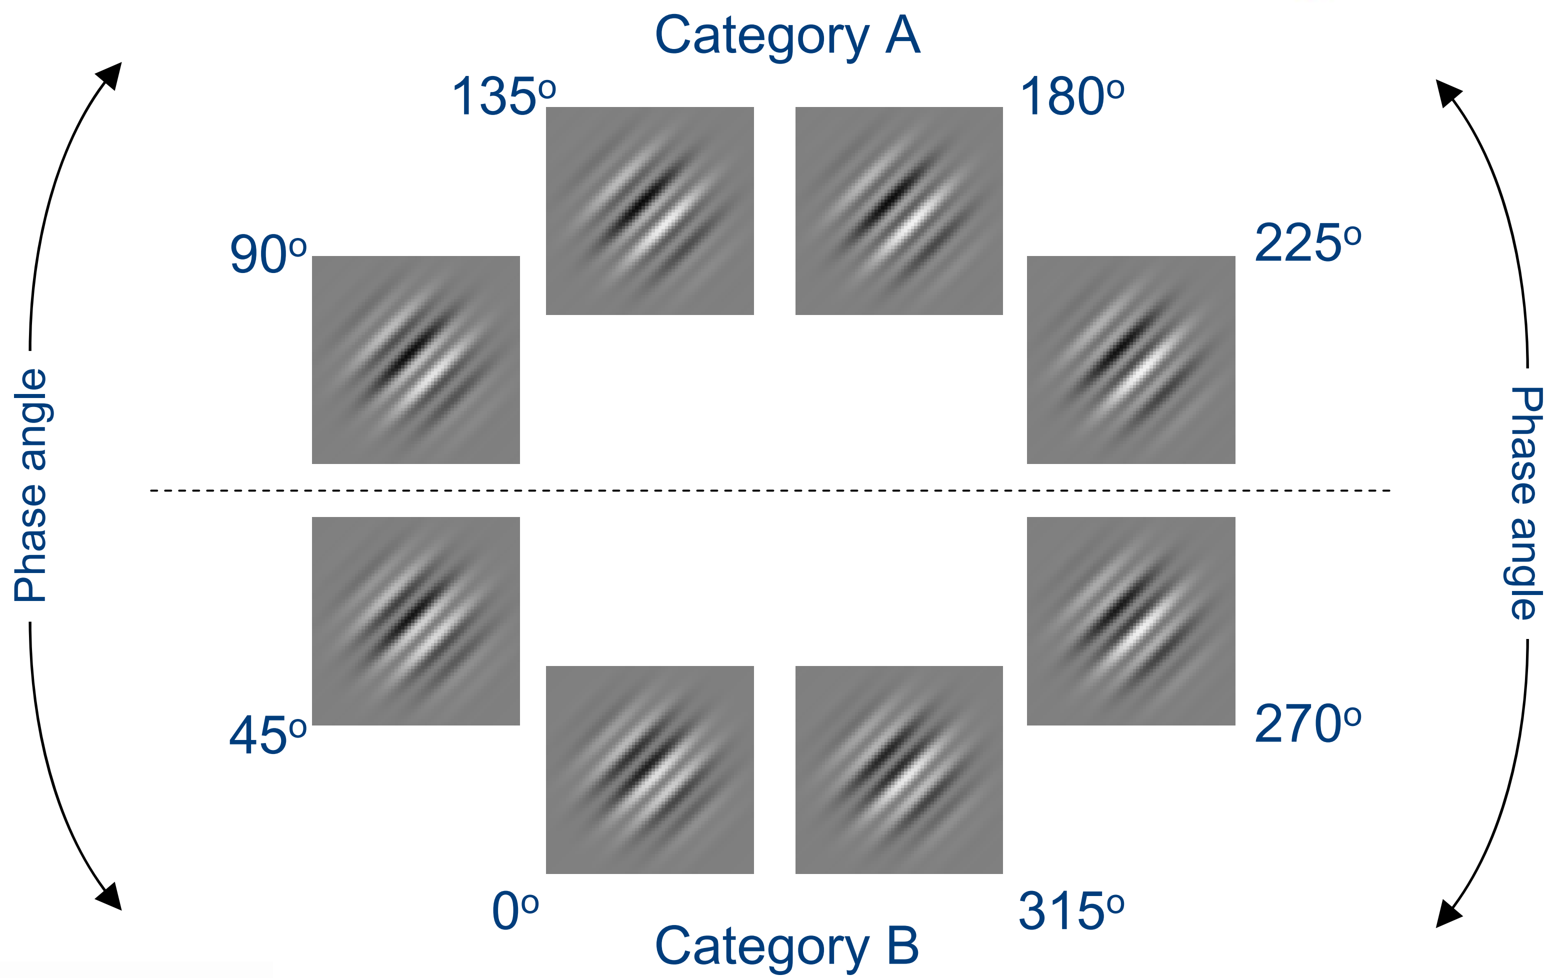

Notman, L.A., Sowden, P.T. & Özgen, E. (2005). The Nature of Learned Categorical Perception Effects: A Psychophysical Approach. *Cognition*, vol. 95(2), pp. B1-B14.

Visual CP

- Categorical perception effect:
  - Discriminating between categories becomes easier with training

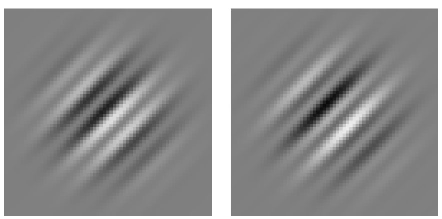

0° 135°

Between

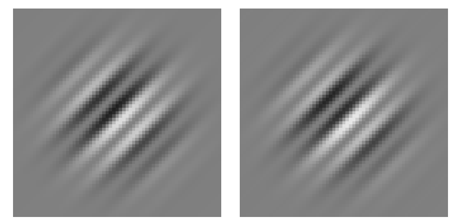

0° 315°

Within

Findings

- Evidence that the behaviour of (at least) the primary visual cortex (V1) appears to adapt
  - Changes to the way receptive fields interact
  - On a task by task basis
  - Different to the traditional view of vision
- However
  - How can we confirm that this really is the case?
  - What psychophysical experiments could explore this further?
- Computational modelling may help
  - Are such models useful?

Computational Modelling

- Prototyping vision with a rate-coded model
  - Abstract, behavioural representation
- Initial supervised approach
  - Naïve use of backpropagation for receptive fields
  - Good image conversion but poor generalisation
- Unsupervised approach with Hebbian learning
  - Based on work by Armony et al’s (1997) successful model of auditory processing in rats for fear conditioning
  - Pavlovian-like conditioning of a model on an auditory stimulus
  - Led to new understanding of the amygdala and further biological experiments

Armony, J.L., Servan-Schreiber, D., Cohen, J.D. & LeDoux, J.E. (1997). Computational Modeling of Emotion: Explorations Through the Anatomy and Physiology of Fear Conditioning. *Trends in Cognitive Sciences*, vol. 1(1), pp. 28-34.

Input Representation

- Simplified input to model phase receptive fields
  - Model of firing rate as discrete Gaussian values
  - Based upon assumptions about human processing

Image Generation

Orientation ±45°

Constant *f* frequency
(0.32 cycles per degree)

Constant *3f* frequency
(0.96 cycles per degree)

Varying combination phase

Example

+45°

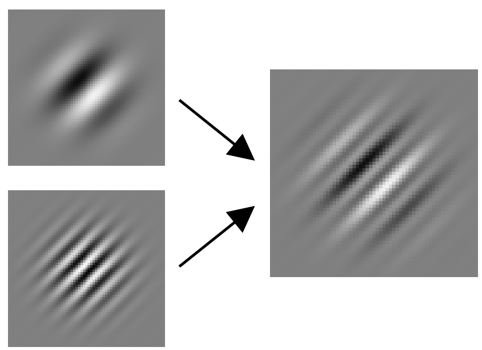

180°

Input

0

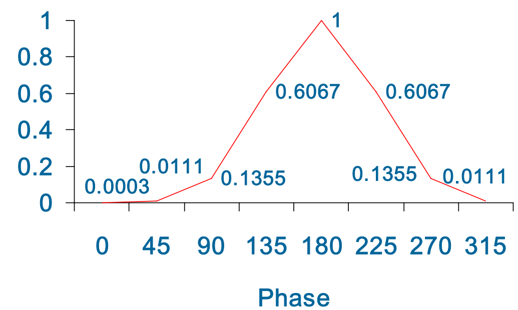

Architecture

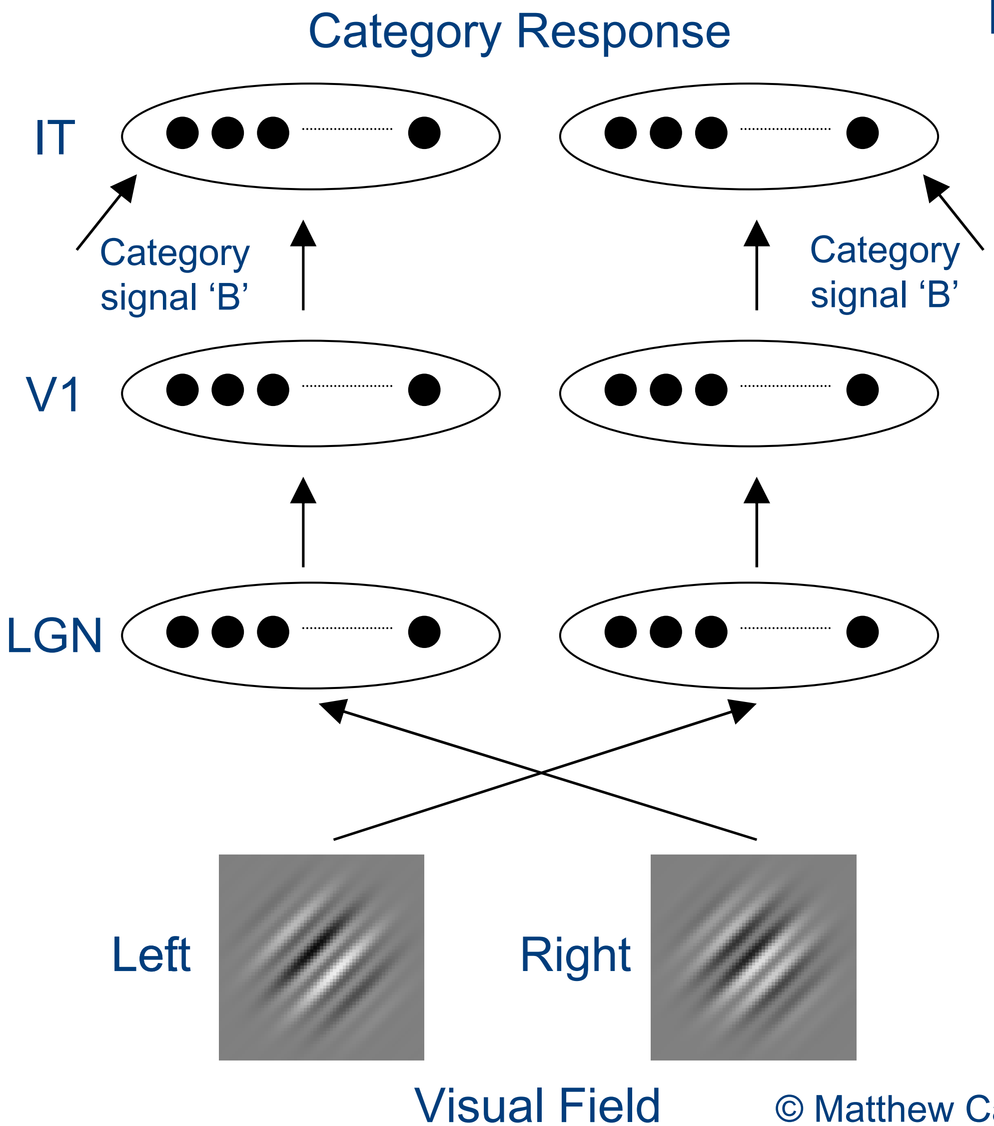

Based on Armony et al's (1997) model

$$u_r = \sum_s a_s w_{rs}$$

Weighted summation

$$a_r = \begin{cases} f(u_r) & \text{if } r = \arg\max_j f(u_j) \\ f(u_r - \mu a_{win}) & \text{otherwise} \end{cases}$$

Neuron activation

$$f(u) = \begin{cases} 1 & u \geq 1 \\ u & 0 < u < 1 \\ 0 & u \leq 0 \end{cases}$$

s	sending units
r	receiving units
win	winning unit – highest activation
a_s	activation of sending unit
w_{rs}	weights from s to r
μ	inhibition rate (I 0.2, V 0.6, L 0.4)
ε	learning rate (0.1)
λ	update factor (1)

Learning Algorithm

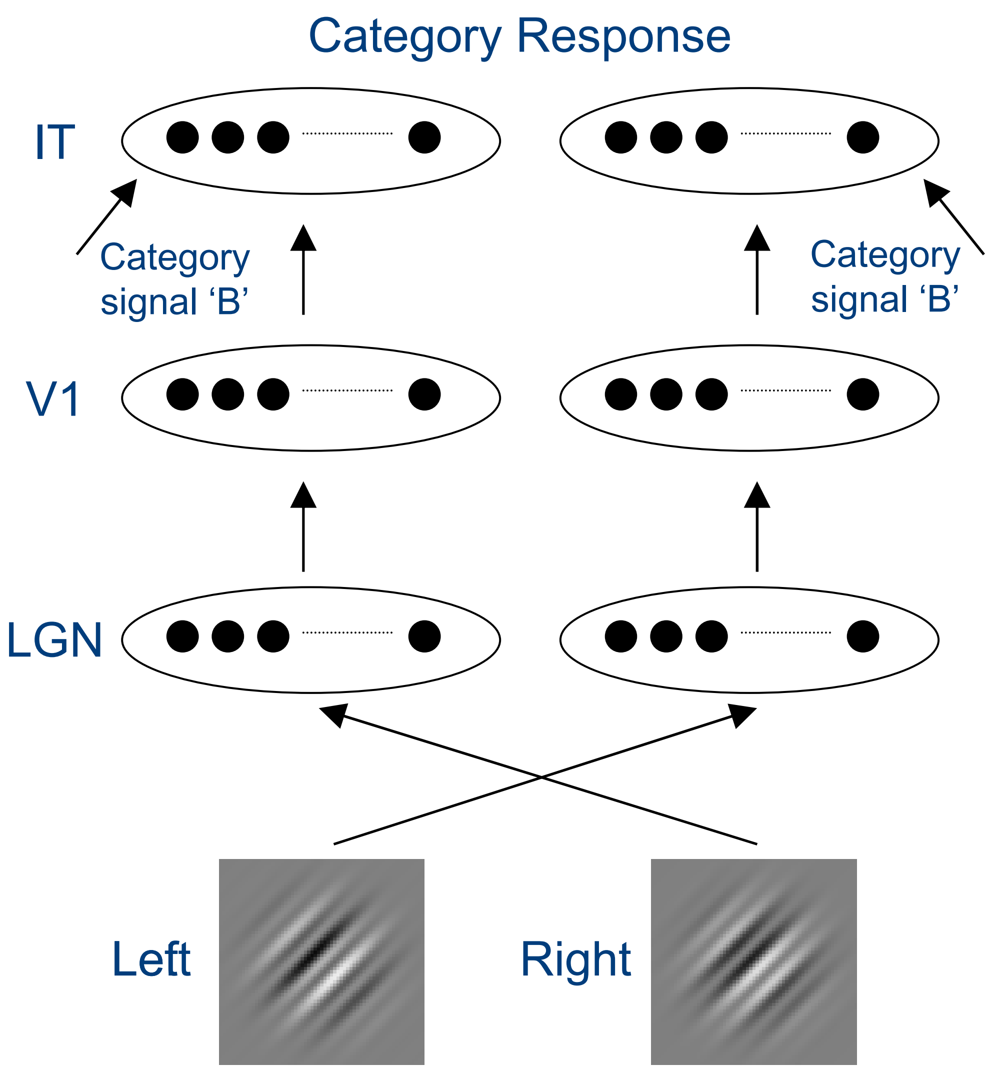

Visual Field

$$w'_{rs} = \begin{cases} w_{rs} + \varepsilon a_r a_s & \text{if } a_s > \lambda a_{avg} \\ w_{rs} & \text{otherwise} \end{cases}$$

Competitive weight change

$$w_{rs} = \frac{w'_{rs}}{\sum_s w'_{rs}}$$

s	sending units
r	receiving units
win	winning unit – highest activation
a_s	activation of sending unit
w_{rs}	weights from *s* to *r*
μ	inhibition rate (I 0.2, V 0.6, L 0.4)
ε	learning rate (0.1)
λ	update factor (1)

Pre-training

- Assumption
  - Human experiments carried out on adults
  - Fully functional visual system
  - Model therefore needs appropriate receptive field firing prior to category training
- Achieved by a period of pre-training
  - Single image (not a pair)
  - All phases 0, 45, 90, 135, 180, 225, 270, 315
  - Presented to a random visual field (left or right)
  - Trained until activation is stable (10000 epochs)
  - No category input

Category Training

- Assumption
  - That the category training matches the human tasks as closely as possible
  - Double and single training tasks with category feedback
- Category training
  - 1 x all possible pairs – within and between category
  - 3 x all single images
  - Presented to a random visual field (left or right)
  - Category input
  - Repeated for just 11 epochs

Testing

- Assumption
  - That the testing matches the human tasks as closely as possible, including the metrics
  - Single and sequential testing with no category feedback
- Testing
  - 1 x all single images – same to both visual fields
  - 1 x all sequential pairs (0, 45; 45, 90; …)
  - Presented to a random visual field (left or right)
  - No category input
  - Repeated 10 times
  - Average responses recorded over trials

Metrics

- Discrimination performance (within and between pairs)
  - Measured using A' (signal detection theory)
- Hits vs. misses
  - Outputs from each module summed per pattern and normalised
  - Difference between left and right IT responses
  - Same image vs. different image
  - Threshold used to determine difference in responses: $\sigma = 0.2$
- When left and right response differ > σ
  - Hit (within) if image pair in the same category: H_w
  - Hit (between) if image pair in different categories: H_b
  - False alarm otherwise: F

$$A' = \begin{cases} \frac{1}{2} & \text{if } \Pr(H) < \Pr(F) \\ \frac{1}{2} + \frac{(\Pr(H) - \Pr(F))(1 + \Pr(H) - \Pr(F))}{4\Pr(H)(1 - \Pr(F))} & \text{otherwise} \end{cases}$$

	Same	Different
Same	Hit (H)	Miss (M)
Different	False Alarm (FA)	Correct Reject (CR)

Model Parameters

- Number of neurons
  - To determine appropriate layer sizes
  - Left and right LGN and V1 varied together: 1 to 16
  - Left and right IT varied together: 1 to 16
  - In each module, 7 gives maximum average A’ difference
- Pre-training
  - To determine stability in pre-training
  - Varying epochs from 20 to 10000
  - No change after 10000 epochs

UNIVERSITY OF SURREY

Results

- Model demonstrates the CP effect
  - Learnt as a result of category training

Human Results

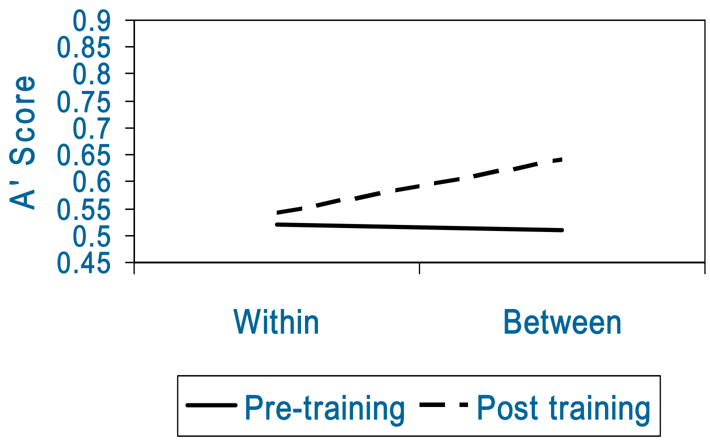

Model Results

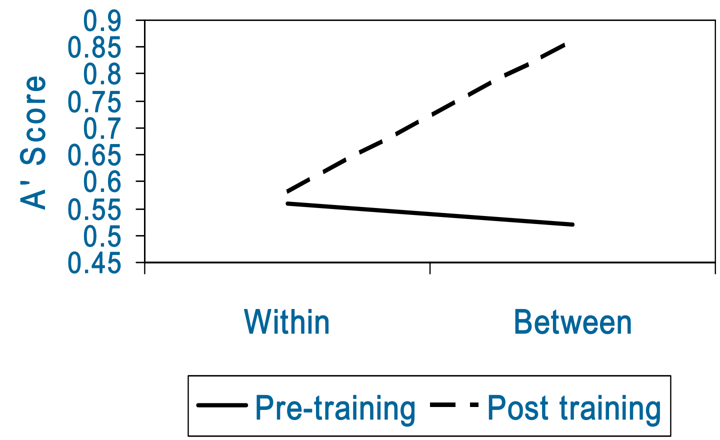

Averaged from 16 participants

Averaged over 100 trials

But…

- Human hypothesis:
  - Evidence from psychophysics
  - Learning to categorise changes visual processing (possibly V1 and LGN)
- Model disagrees
  - Same (abstract) experimental set-up
  - Category feedback only at the discriminating stage (IT) and not earlier
  - Task successfully completed without adaptation in earlier discriminating stages

	Baseline	Rolled Back LGN	Rolled Back V1	Rolled Back LGN and V1
A’ Within	0.52	0.53	0.52	0.52
A’ Between	0.86	0.84	0.86	0.84

More Buts…

- Is the model strong enough to really make conclusions?
  - Human experiments also looked at generalisation of CP to orientations
  - Findings showed that CP is generalised to images within a 6.5° orientation bandwidth
- Model extended
  - Input representation extended to model orientation vs. phase firing
  - Phase bandwidth 106°
  - Orientation bandwidth 30°

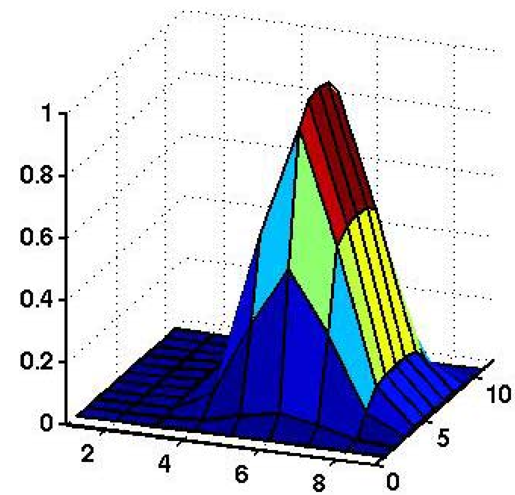

Orientation Generalisation

- Pre-training
  - On selected orientations in pairs
  - 0, 15, 30, 40, 43, 45, 47, 50, 60, 75, 90
- Category training
  - No change: still conducted on 45^{o} orientation only
- Difference recorded for:
  - Within average A' for pre- and category training
  - Between average A' for pre- and category training
- Model parameters
  - Varying update factor λ to determine scale of orientation generalisation
  - Varied from 0 (adjust all weights) to 3 (only those with strong activation)

Results

- Model demonstrates orientation generalisation
  - Similar profile to human data
  - Bandwidth tuned via update factor

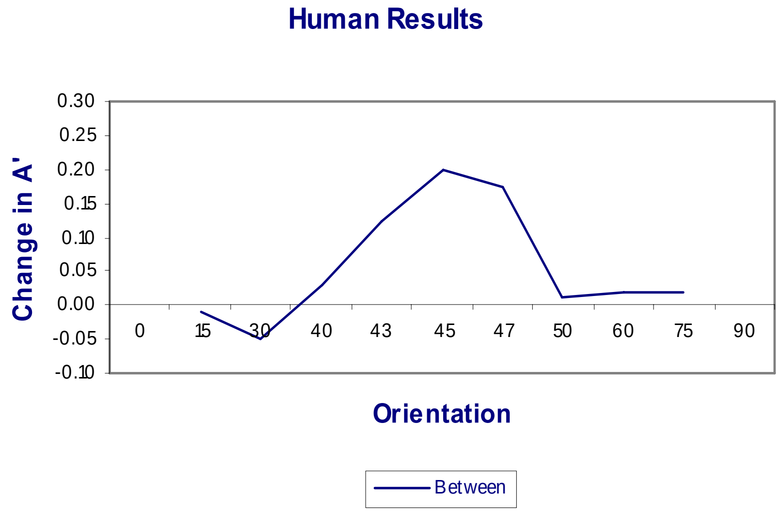

Averaged from 16 participants

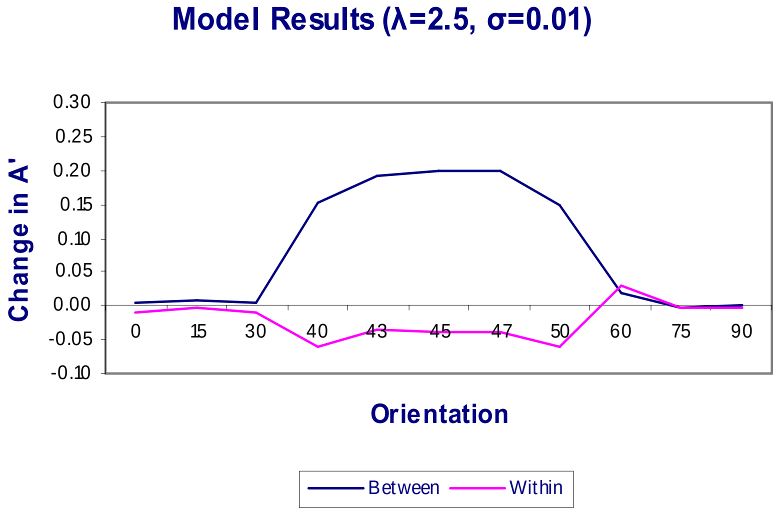

Averaged over 100 trials

Limitations

- It's a very simple model
  - Assumptions regarding suitability of the abstract representation to model visual processing
  - Arbitrary labels for the modules?
  - Arbitrary parameters – lots of tuning?
- Is that really how we should represent visual inputs?
  - No, but actually dealing with visual inputs is a complex question all on its own
- However
  - Demonstrated that the CP effect can be learnt by a simple model
  - Behaviour matches human results closely
  - Demonstrated that feedback is not needed in earlier processing stages

What Next?

- More psychophysics
  - The model has allowed us to test the human hypothesis
  - Now we need to more tightly explore earlier stages of visual processing
  - Image masking is an established method
  - Prevent visual stimuli from reaching higher processing (V1, IT)
  - Human experiments already underway
- Matching this, we need an improved model
  - Can we model the biology better?
  - What about rate-coding vs. pulse coding?
  - What effect does noise have?
  - Evidence suggests that noise is important sub-cortex, but isn't so apparent in cortex
  - Can we develop receptive field processing to replace pre-processing (Linsker 1986?)

Linsker, R. (1986). From Basic Network Principles to Neural Architecture: Emergence of Spatial-opponent Cells. *Proceedings of the National academy of Science*, vol. 83, pp. 7508-7512.

Related Work

- Modelling the superior colliculus
  - Modelling visual and auditory integration for motor output
  - Prototype SOM and Hebbian model developed (with Thanos Pavlou)
  - Plan to put the prototype in a robot (with Constantinos Kroustis)

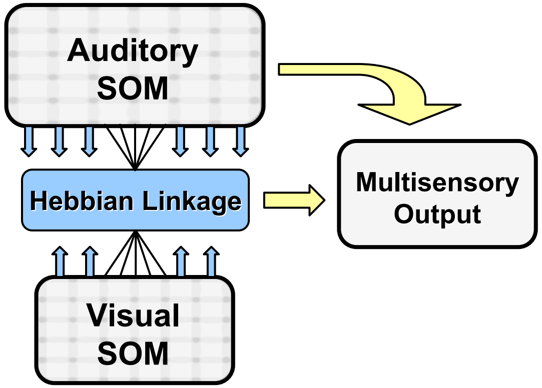

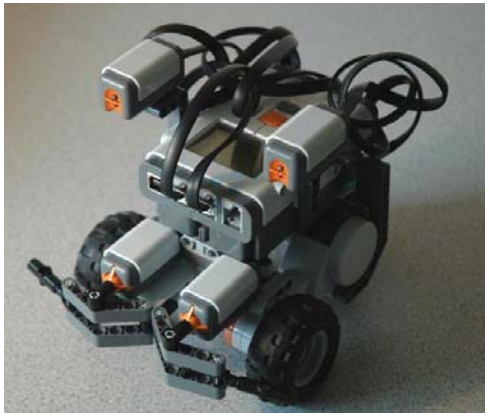

- Modelling visual fear conditioning
  - Cortical and sub-cortical pathways (with Thanos Pavlou)

Where's the Step Change Then?

- What have we learnt so far?
  - That simple models can be used to inform neuroscience and design human experiments
  - But they are simple and suffer from assumptions regarding real-world interaction (pre-processing)
- What next?
  - Build incrementally larger models
  - Put the models in the real world
  - Get them to interact with real environmental stimuli (senses)
- Complex models interacting with the real world
  - The step change will come when we start putting larger models in real environments (Grand Challenges – Denham 2002)

Denham,M. (2002). The Architecture of Brain and Mind: Integrating Low-level Neuronal Brain Processes with High-level Cognitive Behaviours. http://www.nesc.ac.uk/esi/events/Grand_Challenges/proposals/ArchitectureOfBrainAndMind.pdf. UK Computing Research Committee (UKCRC).

What's Needed?

- An interdisciplinary understanding
  - Computer scientists, neuroscientists (biology and psychology) and engineers (electronic and robotic) need to work together
  - Put the models in the real world, in hardware (silicon)
- Theoretical underpinnings
  - Use models we can understand in order to extract the computational principles
  - Partially ordered sets (Mike Shields)
  - Game theory (with Richard Zanibbi)
- Large-scale interdisciplinary project proposal
  - Computer scientists (Austin, Wermter, Murray, Furber, O'Keefe): 3d cameras, cognitive robotics, chip design/processors
  - Computational modellers (Bednar, Casey, Smith): rate and pulse-coding
  - Neuroscientists (Platt, Spence, Stein): animal biology and human psychophysics

People

Dr Paul Sowden, Psychology

Dr Richard Zanibbi, RIT

Dr Mike Shields, Computing

Thanos Pavlou, Computing

Constantinos Kroustis, Computing

Thank you

Questions?

T: +44 (0)1483 689635
F: +44 (0)1483 686051
m.casey@surrey.ac.uk

Abstract

- Our understanding of both natural and artificial cognitive systems is an exciting area of research that is developing into a multi-disciplinary subject with the potential for significant impact on science, engineering and society in general. There is considerable interest in how our understanding of natural systems may help us to apply biological strategies to artificial systems, whilst there are the more traditional opportunities for using computational modelling as a tool to aid understanding of natural cognitive systems. In this talk, we will look at an example of a simple computational model of human vision built to explore categorical perception. The model successfully replicates human behaviour within a simulated set of psychophysical experiments. The work demonstrates that such simple models can give insight into how natural systems operate, whilst providing knowledge to target further human experiments. We will also relate this work to the wider context of artificial cognitive systems, including work being carried out on modelling larger brain structures such as the superior colliculus, reflecting upon the need for a paradigm shift in computational techniques, and hinting at where such a shift may come from.

Varying Neurons

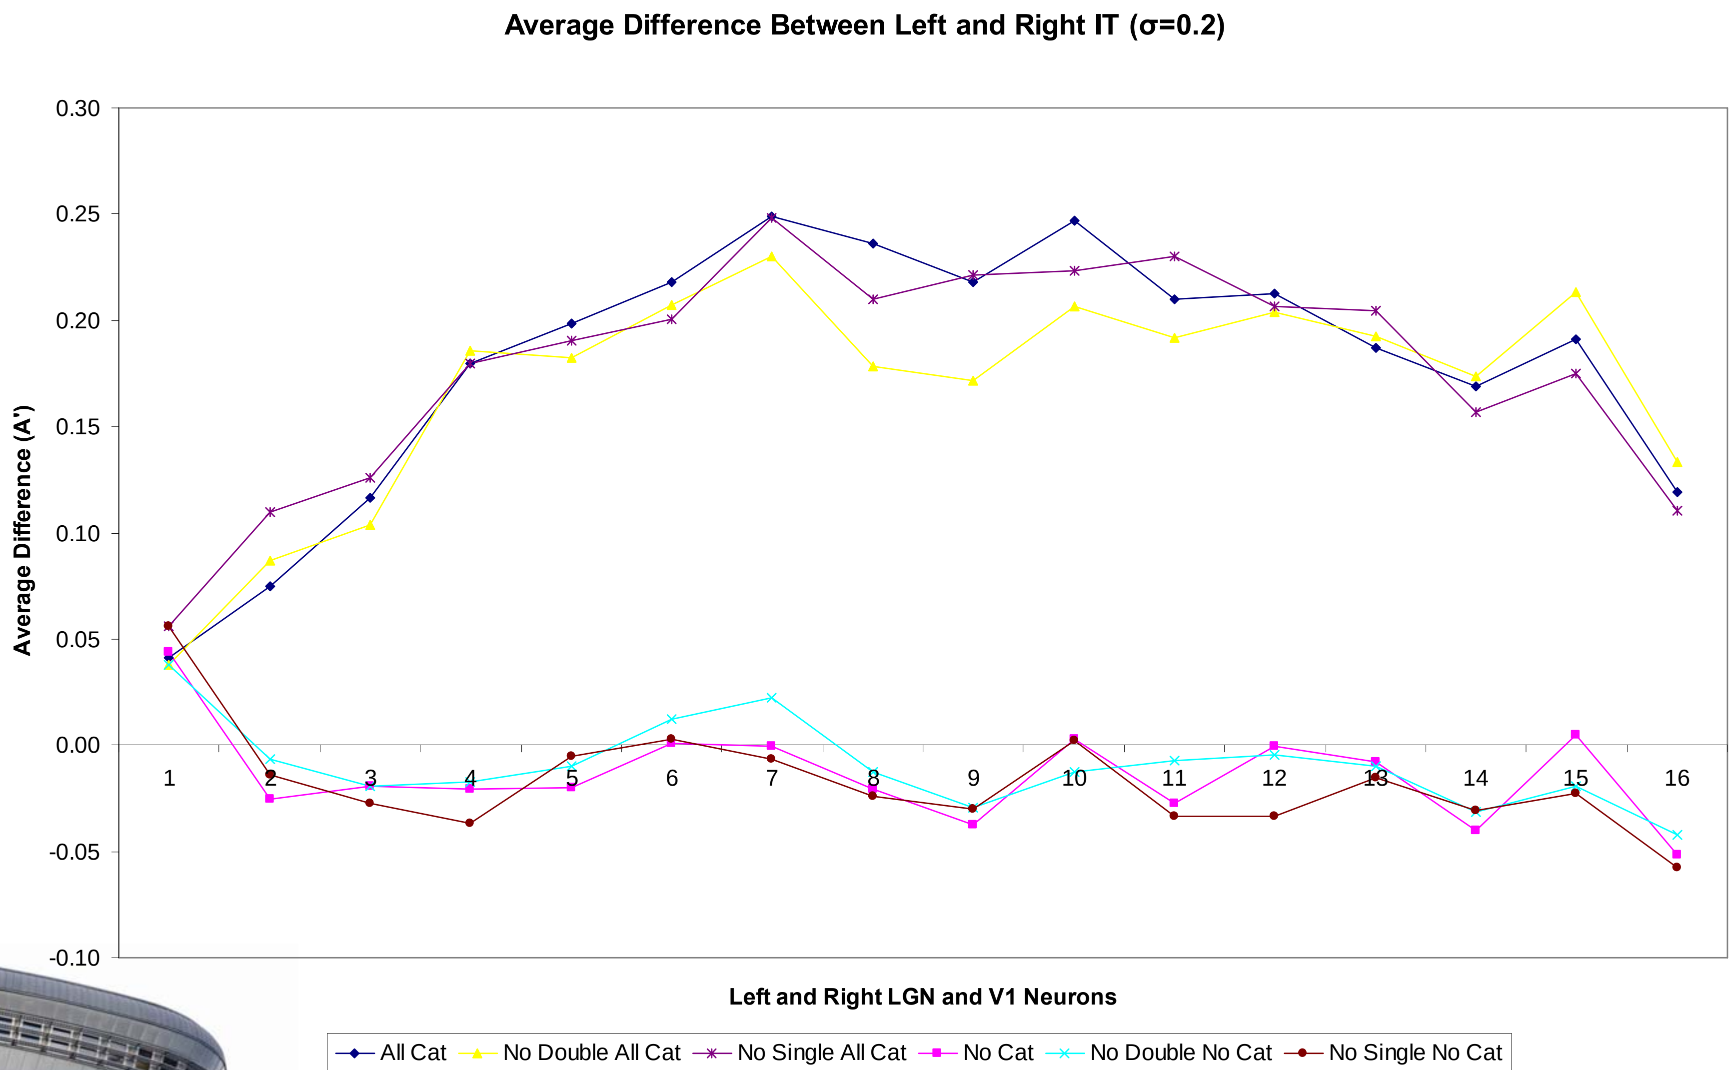

Varying Neurons

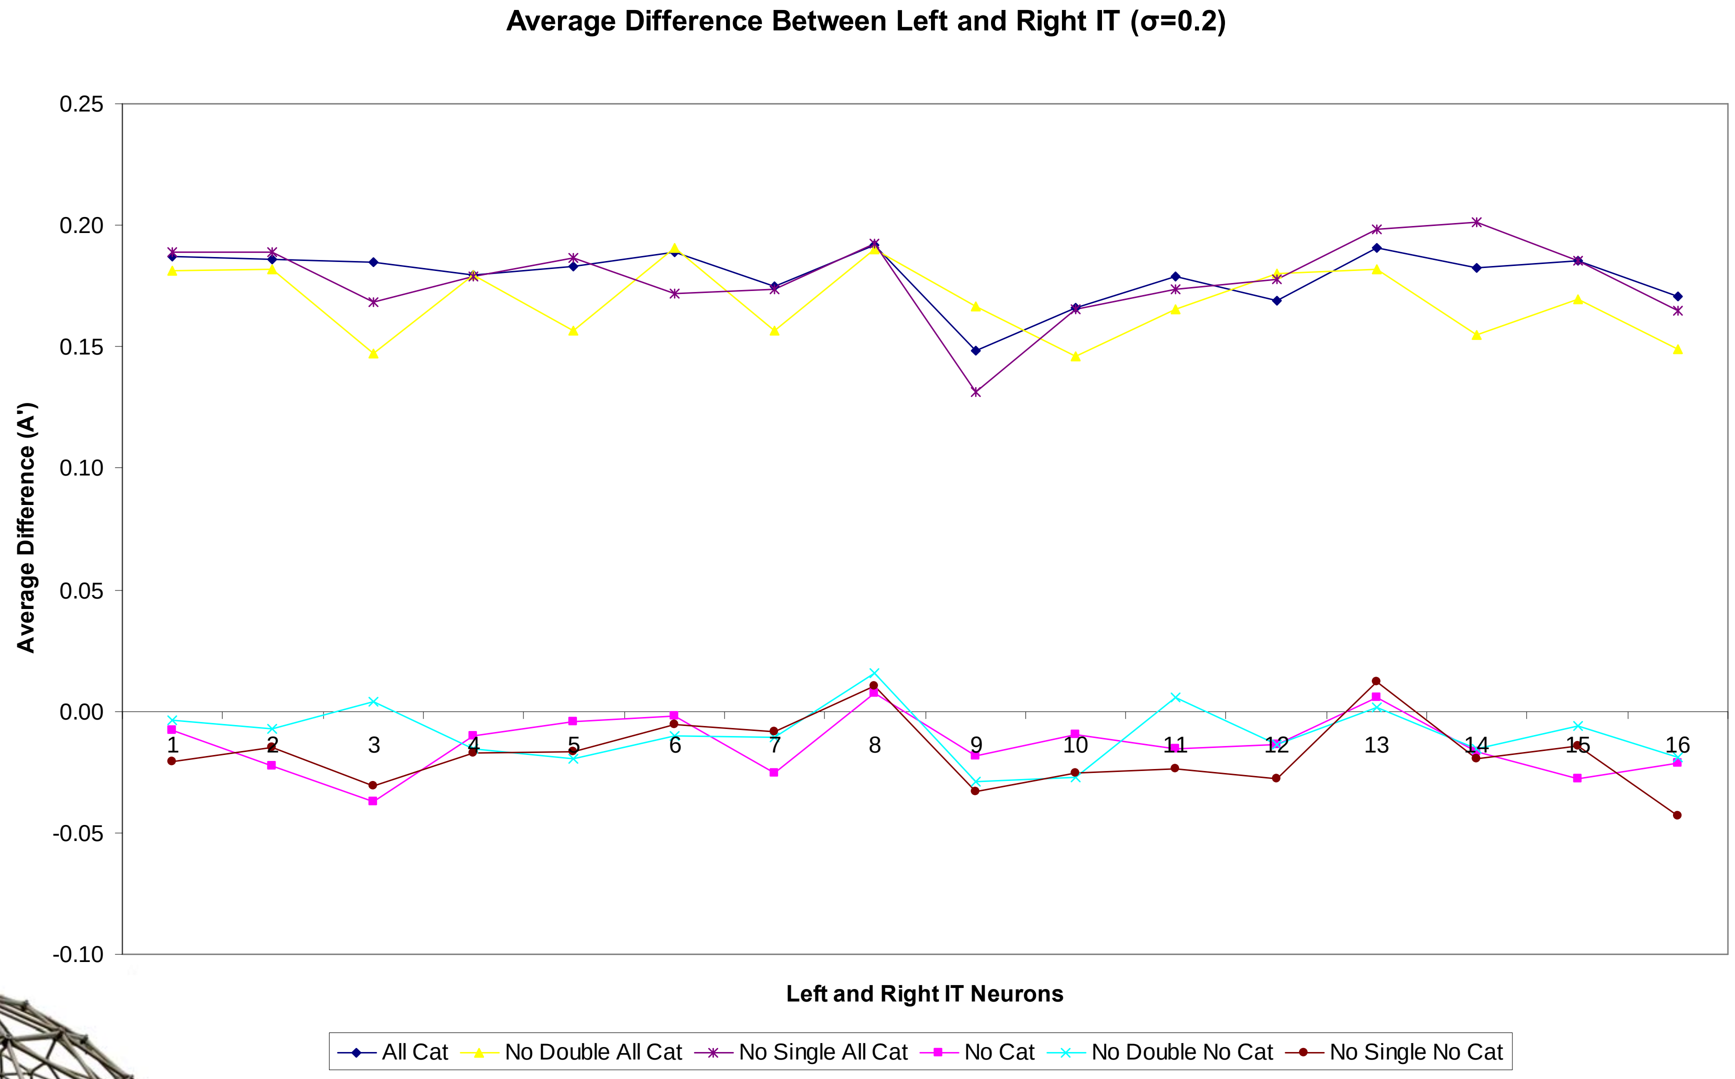

Varying Epochs

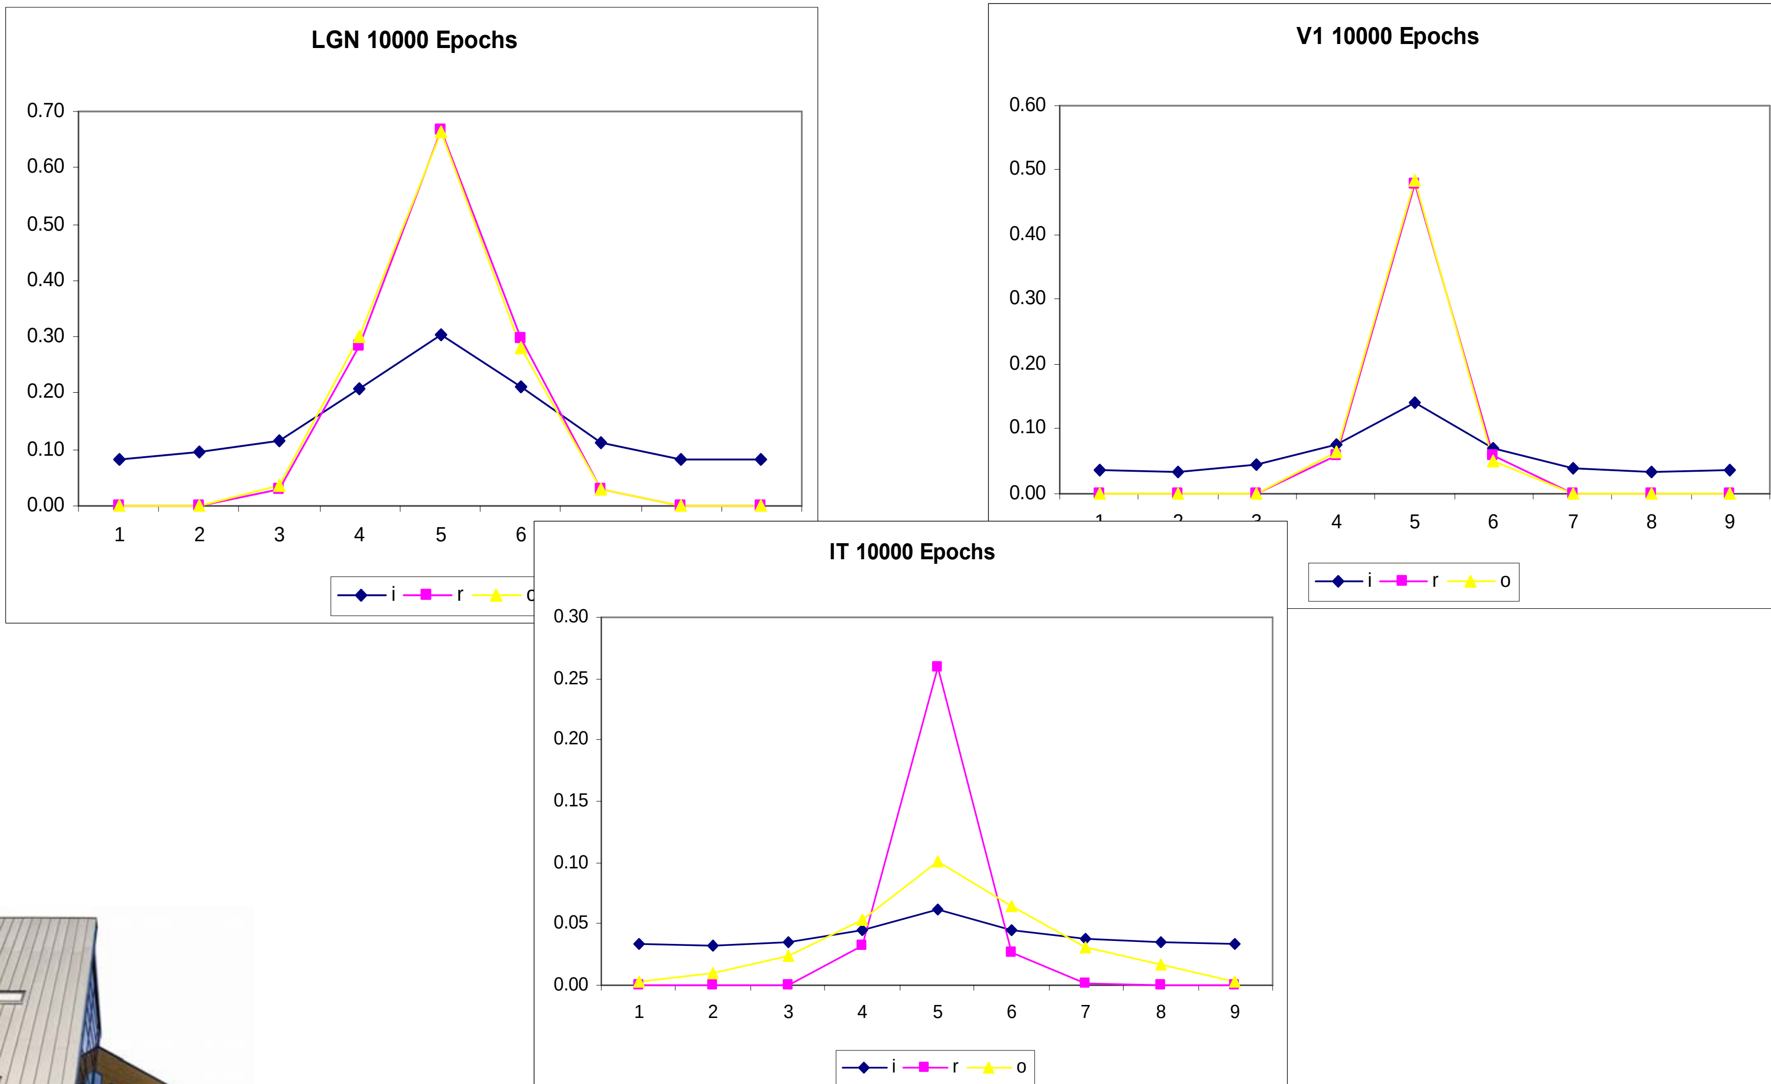

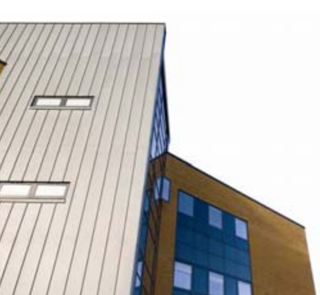

Sensory Maps

Auditory Map Test: 2500 inputs

Red : 225
Dark red : 311
Green : 676
Dark green : 710
Blue : 337
Dark blue : 229

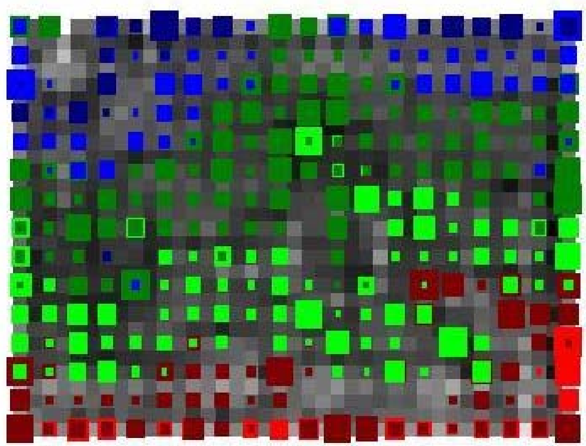

[-180 180],[-90 -65]
[-180 180],[-64.9 -30]
[-180 180],[-29.9 -1]
[-180 180],[-0.9 29]
[-180 180],[30 64.9]
[-180 180],[65 90]

Visual Map Test: 900 inputs

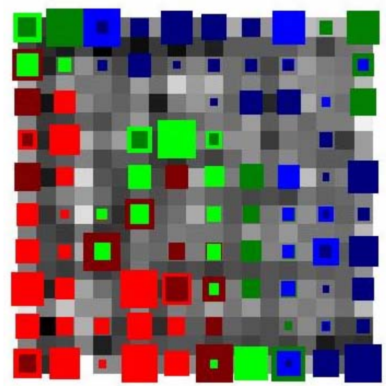

Red : 237
Dark red : 66
Green : 142
Dark green : 156
Blue : 85
Dark blue : 208

[-90 90],[-65 -30]
[-90 90],[-29.9 -15]
[-90 90],[-14.9 0]
[-90 90],[0.1 15.9]
[-90 90],[16 30.9]
[-90 90], [31 65]

Multisensory Integration

Visual to Auditory Link Test: 900 inputs

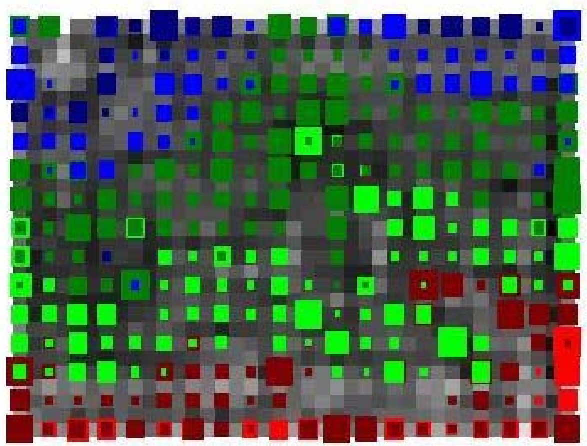

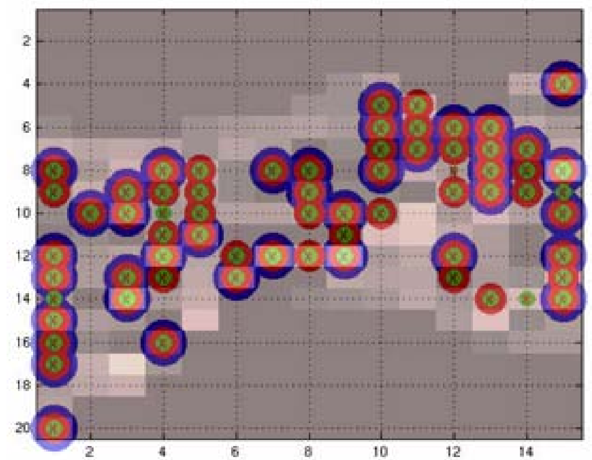

- Perfect hits (302,26.93%)
- Hits with radius=1 (687,61.15%)
- Hits with radius=2 (806,71.73%)
- Hits with radius=3 (934,83.11%)

Not shown in the figure:

Hits with radius=4 (977,86.93%)
Hits with radius=5 (990,88.08%)
Hits with radius=6 (1009,99.77%)
Hits with radius=7 (1012,90.04%)
Hits with radius=8 (1028,91.46%)
Hits with radius=9 (1036,92.17%)

Multisensory Integration

Example Coincident Stimuli test

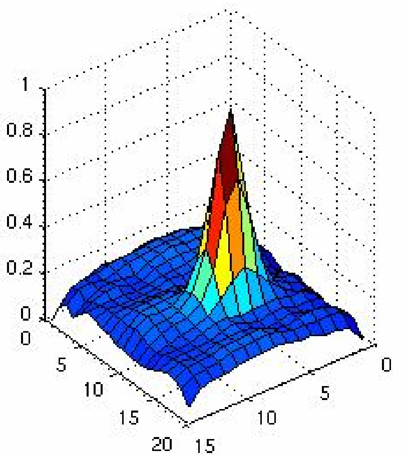

Auditory

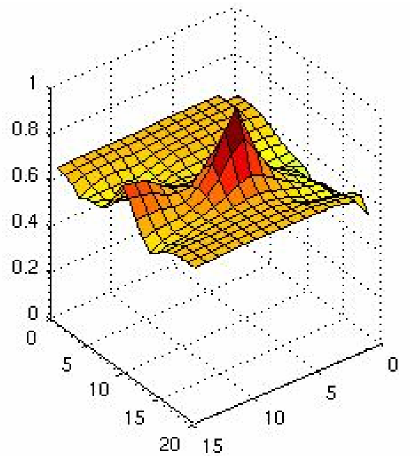

Visual Link

Mulisensory Input_index = 500 (of 1125)
location [12 7] mutisensory_max = 2.00
location [12 7] auditory = 1
location [12 7] link = 1

location [12 7] auditory_max = 1
location [12 7] link_max = 1

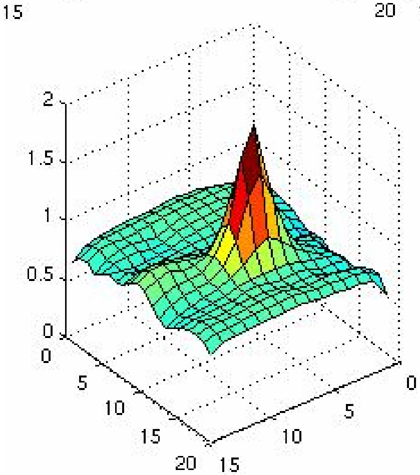

Multisensory