Putting Intelligence into Computational Intelligence

Leslie S. Smith

Computing Science and Mathematics, University of Stirling

Stirling FK9 4LA, Scotland, UK

UNIVERSITY OF STIRLING

Overview

- **What is Computational Intelligence?**
  - **Computational? Intelligence?**
  - What's intelligence
    - And have we captured it?
- Context and Invariance
  - How and why these matter for everyday intelligence.
- Time
  - Temporal context
- Concluding remarks

Overview

- **What is Computational Intelligence?**
  - **Computational? Intelligence?**
  - What's intelligence
    - And have we captured it?
- Context and Invariance
  - How and why these matter for everyday intelligence.
- Time
  - Temporal context

What is Computational Intelligence?

“Fuzzy Computation, Evolutionary Computation and Neural Computation” (IJCCI 2013 website)

“computational intelligence, neural networks, learning, fuzzy logic, fuzzy systems, machine learning, biologically-inspired computing, evolutionary computing, evolutionary algorithms, genetic programming, differential evolution, memetic algorithms, swarm intelligence, ant colony optimisation, artificial immune systems, multi-agent systems, games, data mining, web intelligence, computational intelligence and self-awareness, computational intelligence in bioinformatics, computational intelligence in digital economy, computational intelligence in healthcare, computational intelligence for security, computational intelligence in finance, computational intelligence in robotics, computational intelligence in manufacturing, computational intelligence in energy” (UKCI call, 2012)

Computational Intelligence…

- Generally, solving problems that don’t have obvious algorithmic solutions
  - Dimensionality reduction
  - Classification of high dimensional data
  - Prediction of time series
  - Optimisation of complex processes
  - Finding signals in noise
- … using techniques that might have a biological inspiration,
  - Neural networks, Genetic/Evolutionary Algorithms, Ant Colony Optimisation, Immune systems
- … or might not
  - GOFAI, Bayesian Inference, Decision trees, Fuzzy systems

But what's this got to do with *Intelligence*?

- Difficult to answer: what's *Intelligence*?
  - And what are the (neural) mechanisms that underlie intelligence in animals and humans?
- Is Intelligence in animals necessarily different from Intelligence in CI?
  - If so we should be looking at hybrid living/silicon systems

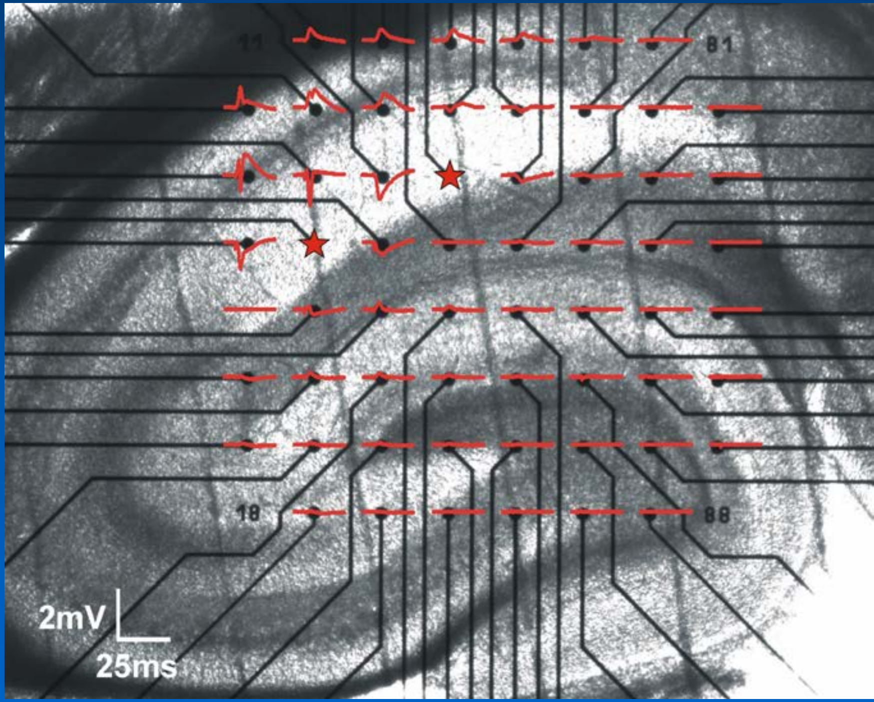

An aside on Intelligence

- Difficult to define, and source of many conflicts
  - Not just in CI!
- Eysenck and Intelligence tests
  - What do intelligence tests measure?
  - Is intelligence defined by what intelligence tests measure?
- Racial and cultural issues
  - Cultural stereotyping in tests: test including questions on cars or planes, given to groups who hardly ever drive or fly
    - Has been a major issue …

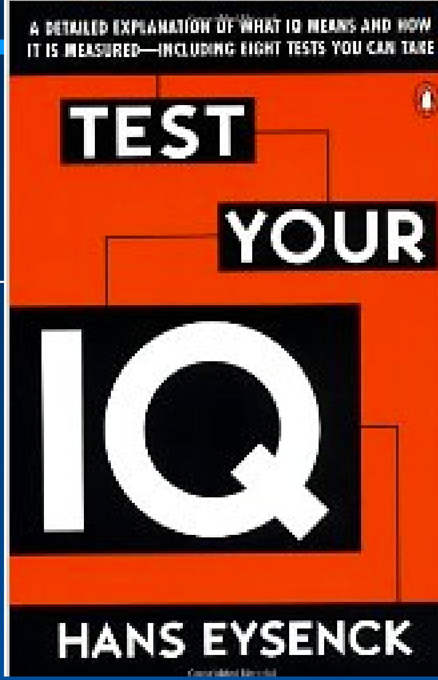

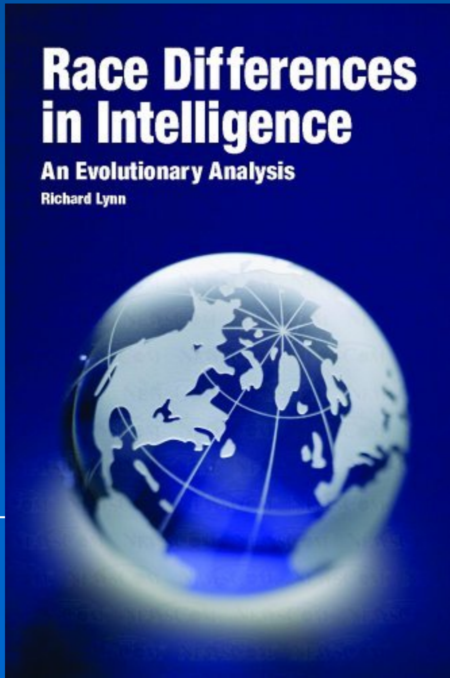

Easier question:

- What problems require intelligence in order to solve them?

Intelligence and Computational Intelligence

- Generally, in CI:
  - the ability to generalise, not just to recall;
  - the ability to use context to influence decisions, to bring a great deal of information to bear on the interpretation of some current data;
  - the ability to discard the irrelevant and to use the information that matters at the time it's needed…
- And not
  - purely functional approximation, classification, WTA networks, optimisation, …
- But…

Intelligence is a moving target

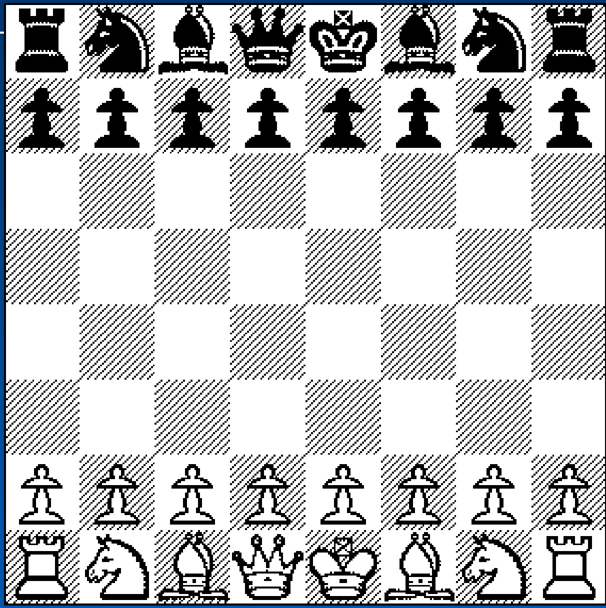

- As techniques for solving problems that were believed to require human-like intelligence have developed
  - We have redefined what we mean by intelligence
  - Chess-playing?
    - Brute-force and clever searching techniques
  - Character and face recognition?
    - Data reduction, followed by a recogniser
  - Other game playing?
    - E.g. Tesauro's TDGammon
- As we understand the techniques used in these systems, we have tended to move away from considering them as intelligent.
  - Or at least as requiring human-like intelligence

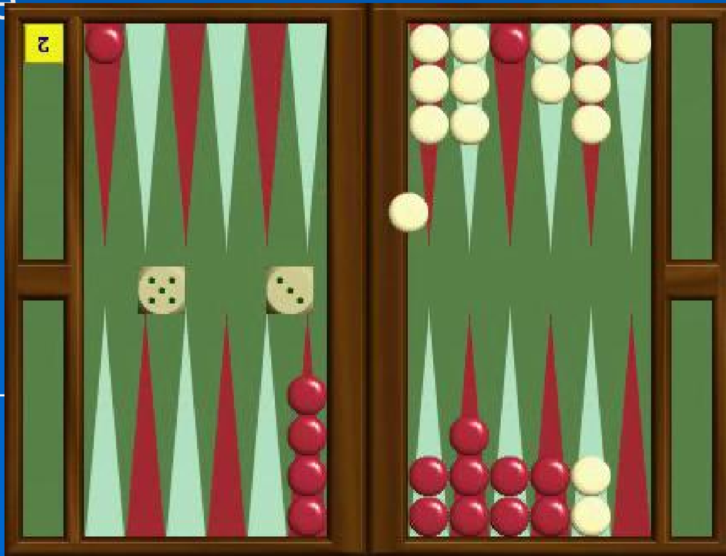

So what's left?

- Self-awareness, consciousness, motivation, …
  - No-one is claiming that current CI techniques show these capabilities
    - As opposed to *modelling* these capabilities
- But: is there an area that might reasonably be called intelligence
  - But which doesn't require 1st person science that seems implied by the capabilities above?
- What elements of intelligence can be dissociated from
  - The clever techniques of non-algorithmic computing

  and

  - The 1st person aspects of self-awareness, consciousness and motivation

Possible answers

1. Nothing’s left.
   – Pursue awareness/consciousness research

Possible answers

1. Nothing's left.
   – Pursue awareness/consciousness research
2. There is something left.
   – Try to identify what it is
   – Pursue it!

Possible answers

1. Nothing's left.
   - Pursue awareness/consciousness research
2. There is something left.
   - Try to identify what it is
   - Pursue it!
3. Maybe there's something left & maybe not
   - See (2) above!

What is there that's not just (!) current CI, nor yet 1st person awareness/consciousness, but could help to implement/understand intelligence?

The Computationalist view

"Computationalism is the view that cognitive capacities have a computational explanation or, somewhat more strongly, that cognition is (a kind of) computation" *Piccinini, Oxford Handbook of Philosophy and Cognitive Science forthcoming.*

- We already know the answer:
  - Cognition can be implemented using …
    - A computer (1960's/70's)
    - A Perceptron-based machine (1960's)
    - A logical inference machine (Japan, 5th generation project , early 1980's)
    - A Back-propagated delta-rule based machine (mid 1980's)
    - A Reinforcement Learning Machine (1980's/90')
    - A Bayesian Inference machine (2000's: Hawkins & George, for example)
  - Perhaps it has aspects of all (or at least some) of these
    - And perhaps it is rather more than these.

Intelligence and interacting with the world

- Robots/computer systems are generally poor at this
  - Except in highly restricted areas
- Better interaction with the real world …
  - As opposed to better interaction with people
  - Or better interactions mediated by people
- … could lead to machines which might *appear* more intelligent
  - And would certainly be more capable and useful
- But how might this be achieved?
  - What’s missing?

Overview

- **What is Computational Intelligence?**
  - **Computational? Intelligence?**
  - What's intelligence
    - And have we captured it?
- Context and Invariance
  - How and why these matter for everyday intelligence.
- Time
  - Temporal context

Everyday interaction

- What makes animals appear intelligent?
  - Clue: it’s not chess-playing or problem solving!
- It’s the everyday “stuff”!
  - Walking through a doorway without colliding with it
  - Smoothly walking over a rough surface
  - Lifting up an egg without breaking it
  - Finding one’s way home
  - …etc.
- What underlies this everyday behaviour?

Invariance and Context, and time

- Invariance and Context
  - Interpretation (and interaction) taking into account nearby /recent events sensory data, signals, etc.
  - Invariance: the ability to interpret sensory data in an appropriate way independent of some specific changes in the sensory data
    - caused e.g. by illumination (vision) or reverberation (sound)

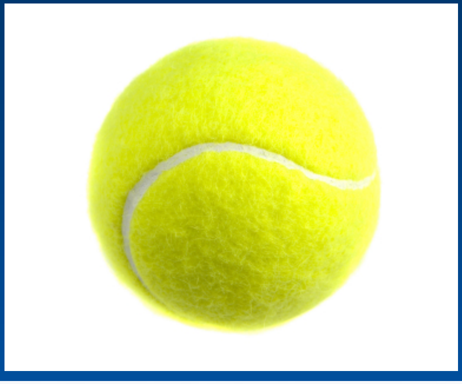

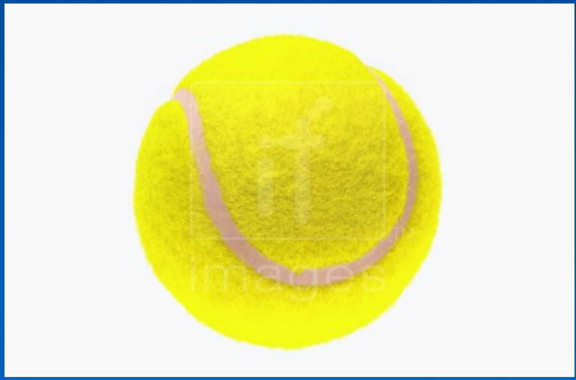

- Time: a much ignored dimension.

Invariance and context

- Nothing in the real world is interpreted on its own.
  - Low-level vision
    - Completely different energy spectra are interpreted as the same colour in different illuminations
      - Or the same energy spectra are interpreted as different colours
    - Or even in different parts of an image
      - Colour constancy.

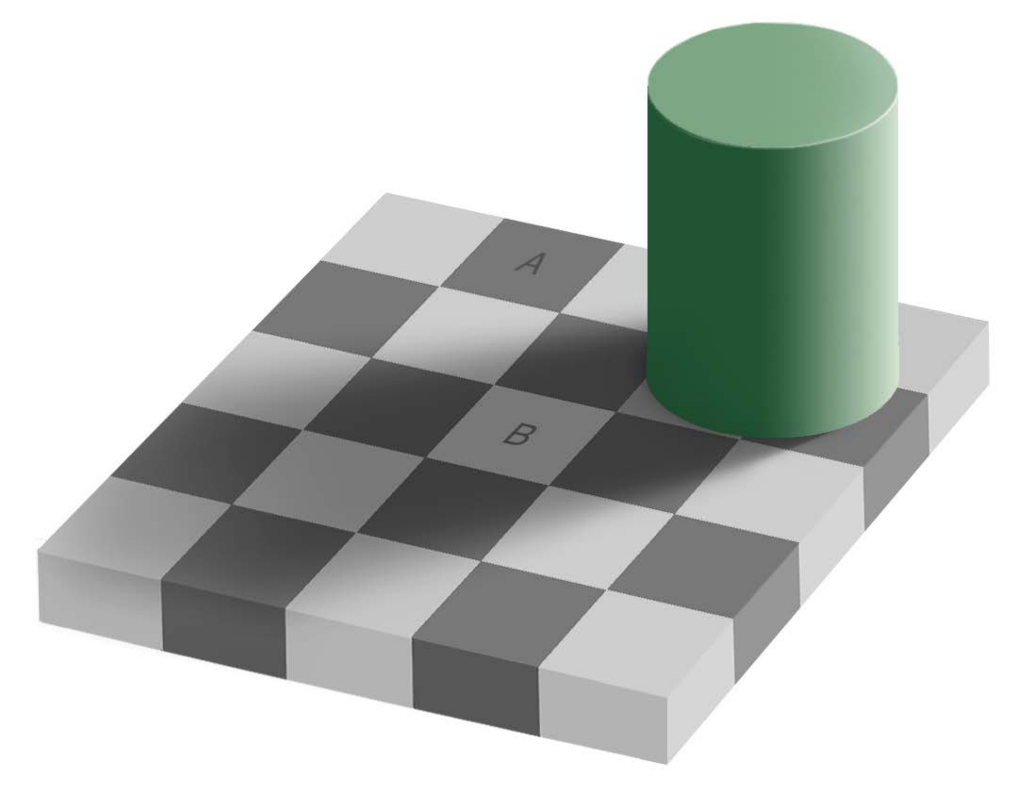

What is context?

- Sensory contexts / sensory fusion
  - Lifting a coffee-cup: locating it, moving a hand to it, gripping it hard enough not to drop it, but not so hard that it breaks
    - Visual, proprioceptual, tactile sensing all need to be fused.
- Usually thought of in terms of *motor programs* that have been learned
  - But is this a good description?
- Contextual modulation is clearly present at every level of neural description
  - E.g. Phillips and Kay at the synapse level
  - The McGurk effect at the sensory level
  - Difficulty recognising someone out of context

What is invariance?

- Adjusting perception to reality
  - Interpreting as the same, things which are the same, even when they appear different
  - Invariant visual perception
    - under varying illumination, varying distance (size), differing orientation
  - Invariant auditory perception
    - Varying loudness, varying reverberation levels, in the presence of background “noise”

Why are context & invariance so important?

- For real-world interaction.
  - For pure physics we'd want the opposite
- Objects produce characteristic reflectivity patterns, characteristic sounds, …
  - And we want to identify them correctly in varying visual and auditory environments
  - Critically important both for animal and synthetic intelligent systems
    - "characteristic" actually includes what matters for identifying the object, rather than what is physically the case.

How is invariance implemented?

- ***Invariance*** is tied to specific modalities
  - Invariance under …
  - In animal perception it is often at the sensor
  - Visually:
    - Retinal architecture
    - Lots of processing on the retina, before it is transmitted to the cortex
  - Invariance under varying illumination
    - Overall, and local changes

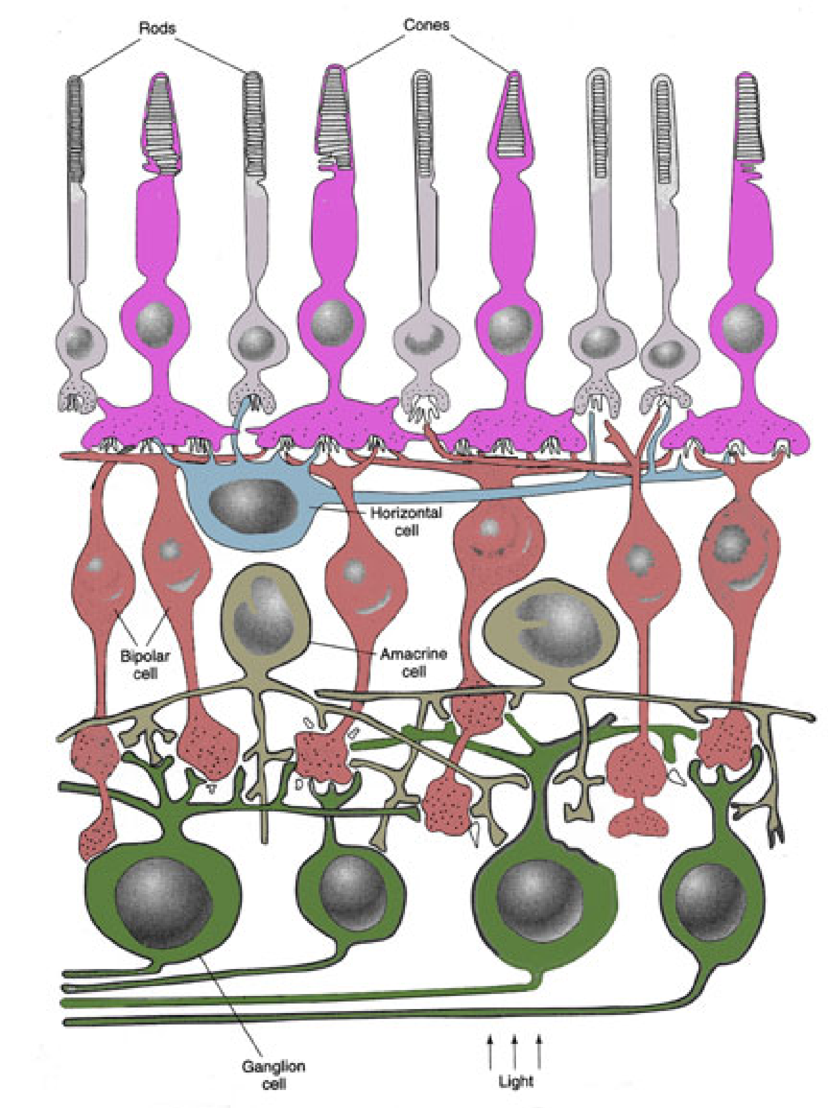

… and in synthetic visual systems

- Two levels of resistive networks
  - Aim is to enable local contrast enhancement
  - Allowing imaging where the local contrast remains visible even when the overall contrast is very large
  - E.g. where there is sunlight and shadow

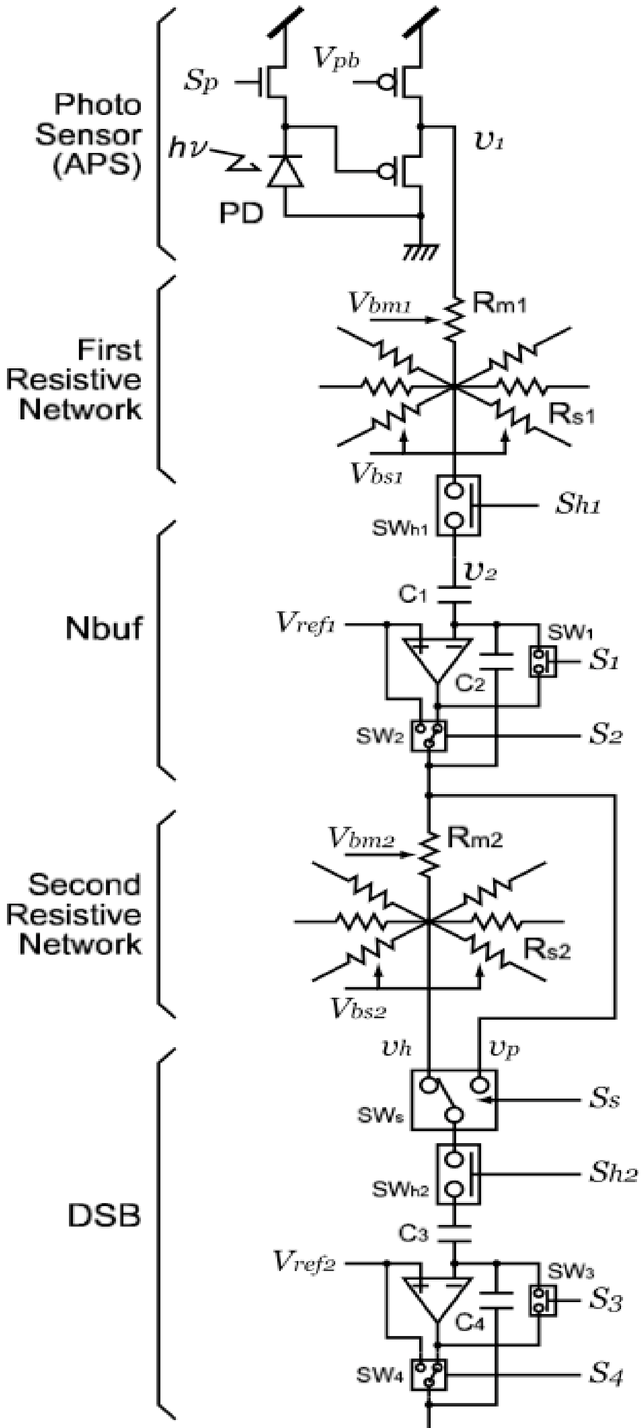

Delbruck's Dynamic Vision Sensor

- Asynchronous temporal contrast silicon retina
- Reports object movement
  - Not framestore based!

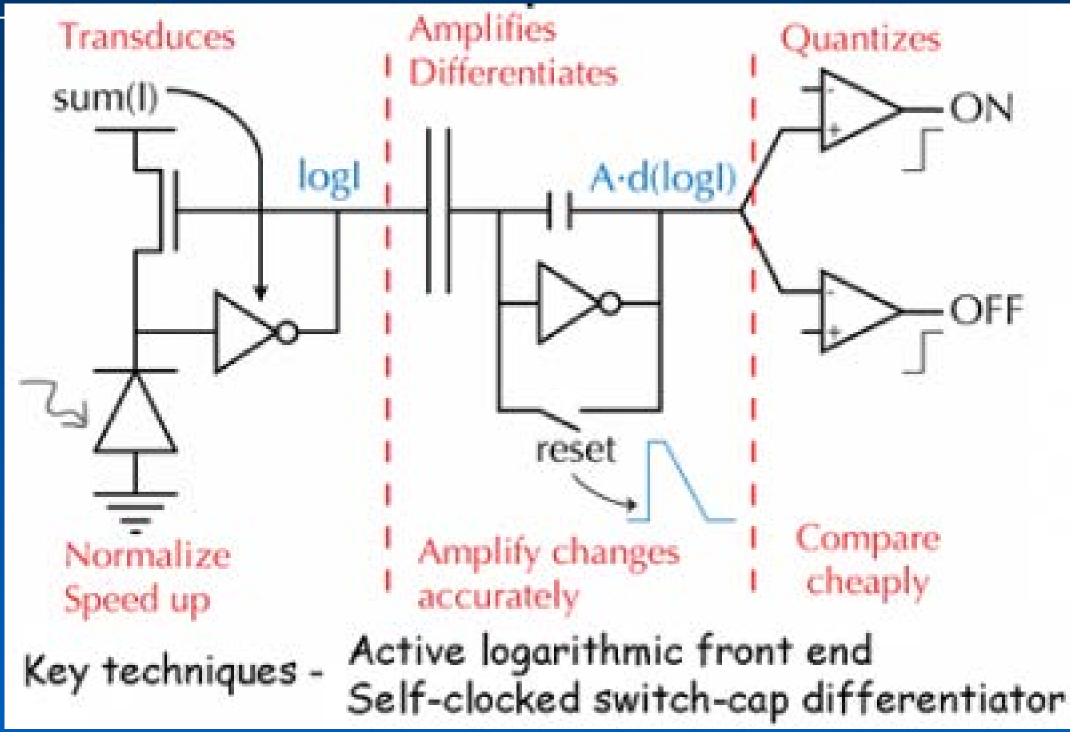

See http://siliconretina.ini.uzh.ch/wiki/index.php

Why invariance in audio perception?

- Speech recognition systems…
  - Work well in good signal: noise situations
  - Don’t work very well when there’s
    - Reverberation
    - Multiple sources
  - Yet the normal situation is that there are many sources, and lots of reflections.

What does the auditory system do to succeed in this situation?

Invariance in early audition

- Processing takes place in brainstem nuclei
- Before transmission to the auditory midbrain, and thence to the cortex
- Very similar across a large range of animals
- Invariance under level, reverberation changes,
  - Also sound source separation
  - …

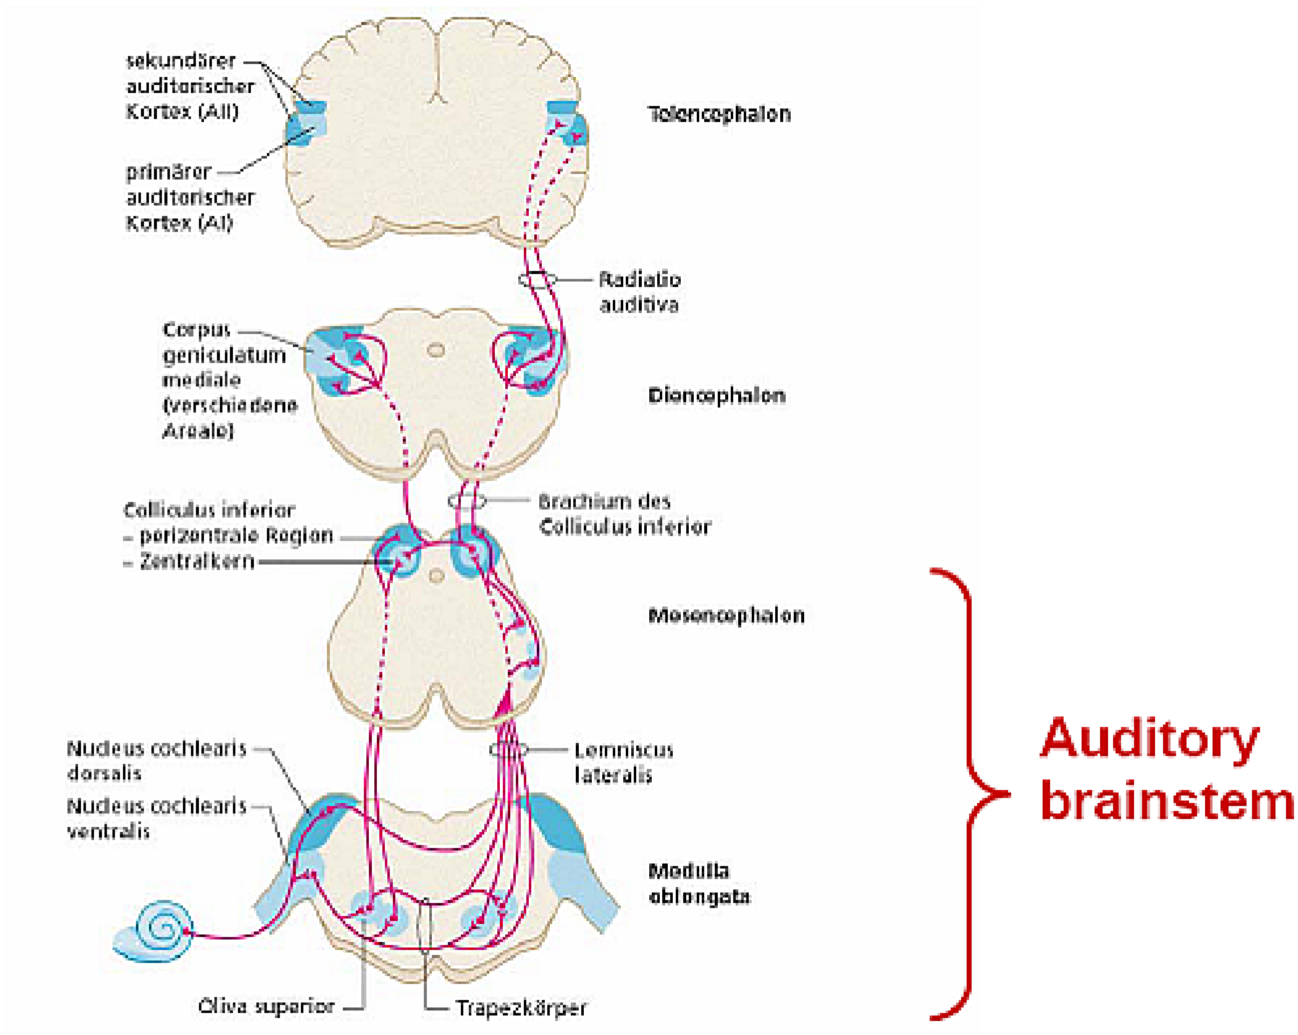

Fig. 1 The central auditory system is built-up by many centers and pathways

An aside: on onsets

- An interesting feature
  - "defined" as a sudden increase in energy in some part(s) of the spectrum
    - Many different ways of detecting them
  - Of interest in music recognition
    - Note starts, drumbeats
  - In speech segmentation
    - Voicing onset, sibilance onset
  - (reasonably) immune to degradation by reverberation
    - (i.e. invariant under reverberation)
  - Very clearly enhanced in the auditory brainstem
    - Multipolar and octopus cells in cochlear nucleus, and many others as well

Onsets and offsets

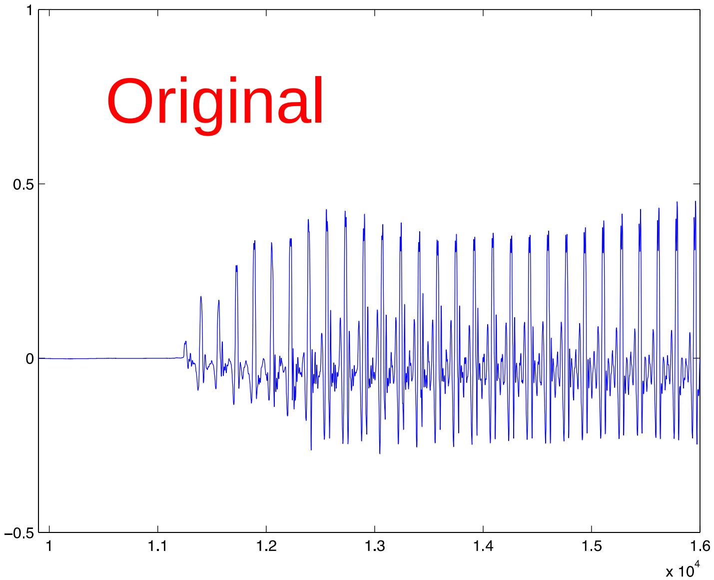

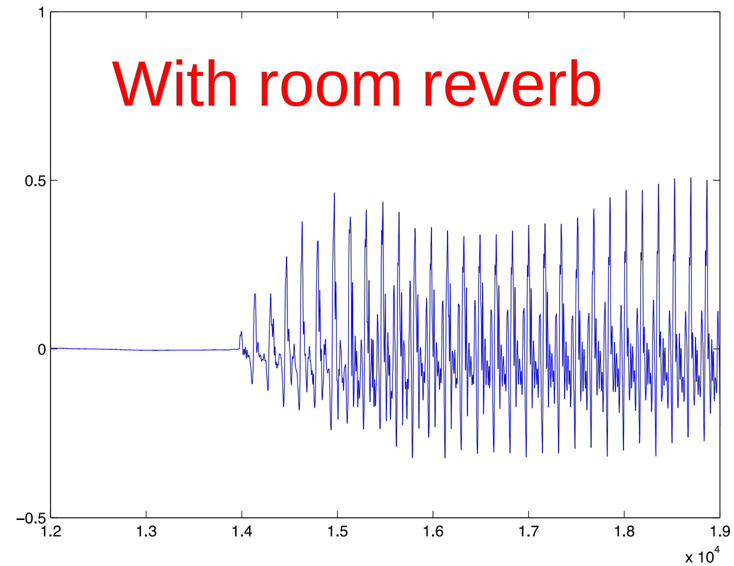

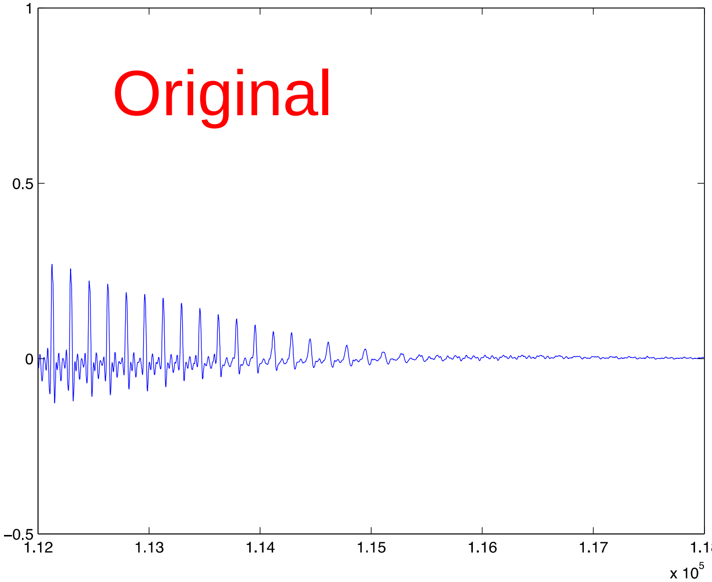

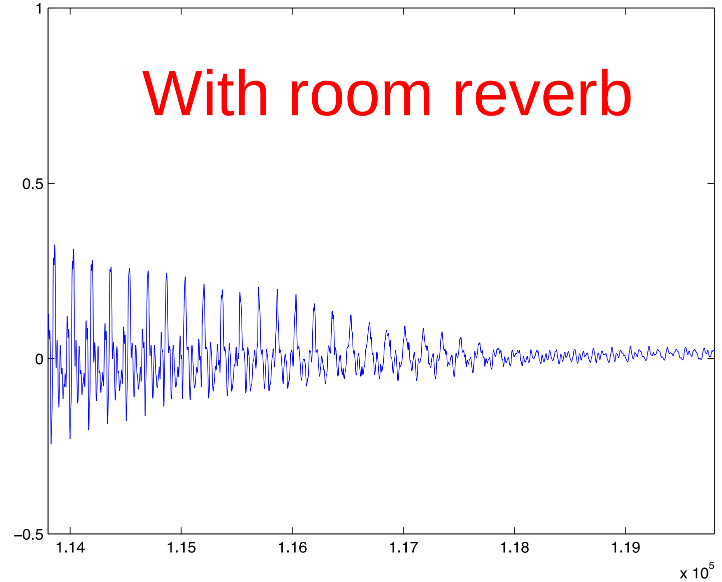

A bio-inspired front end

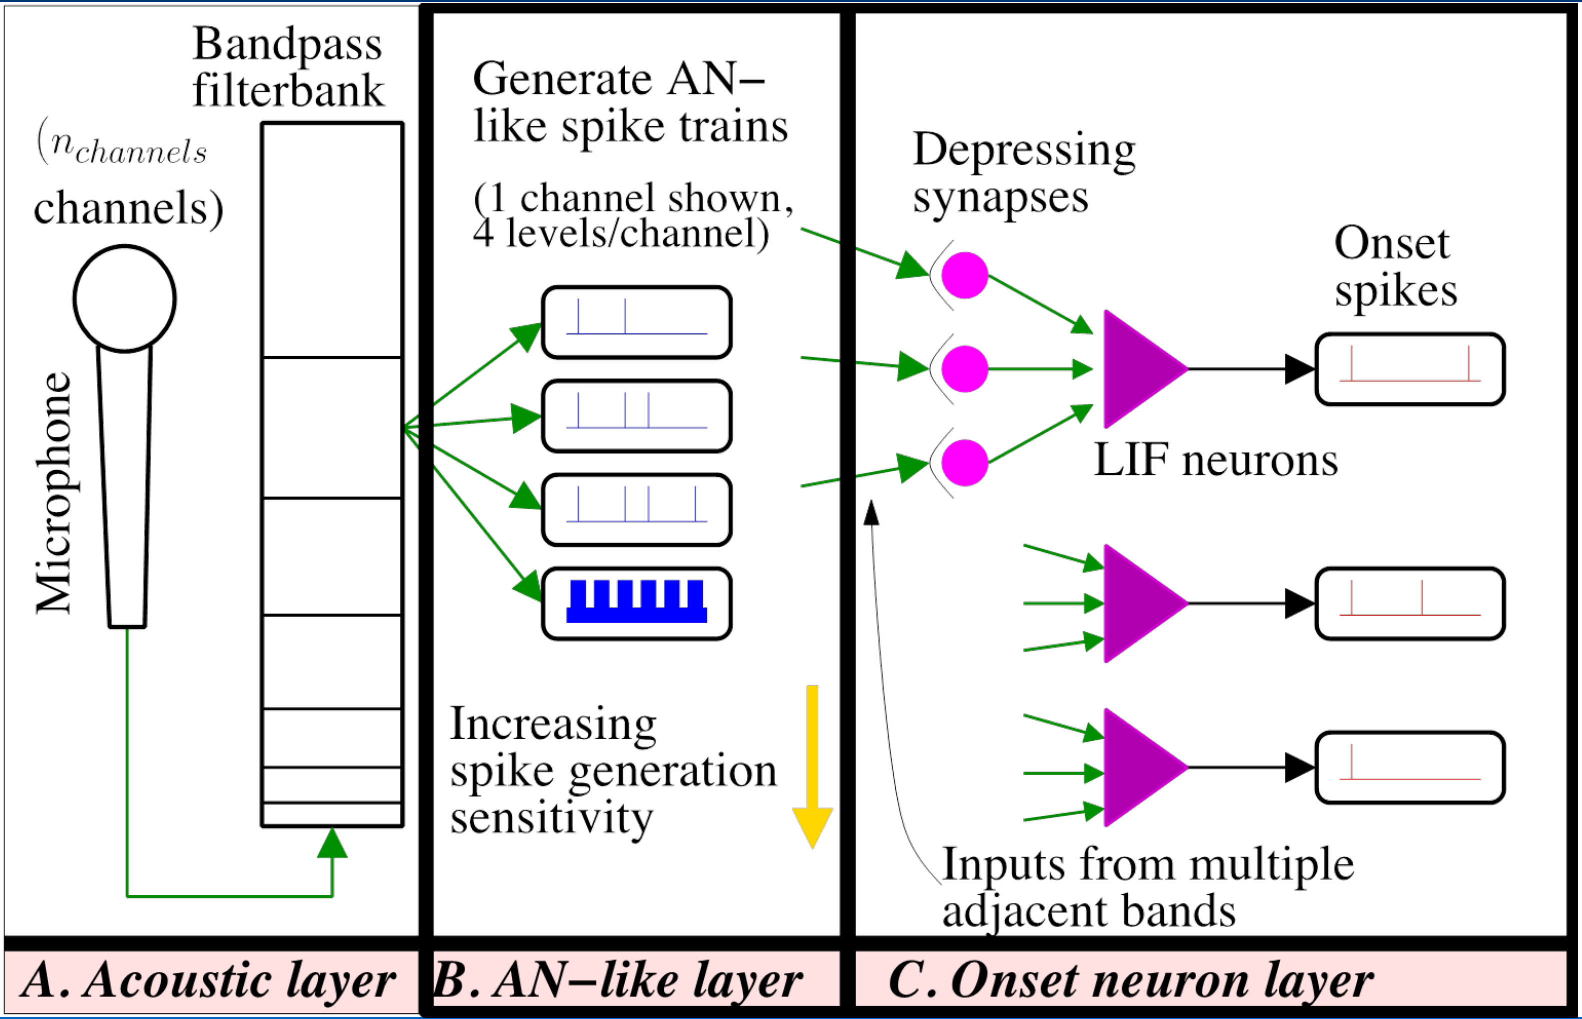

Musical Instrument class problem: detail

Brass

x 417

Reed

x 417

Bowed string

x 417

Plucked string

x 417

Struck string

x 417

Strategy 1: spiking auditory model

Strategy 2: cepstral coefficients

Sound descriptor (2085 examples in total)

Train classifier with 1460 examples (292 per class)

CLASSIFIER MODEL

Test classifier success using 625 unseen examples (125 per class)

Results (1: on the McGill dataset)

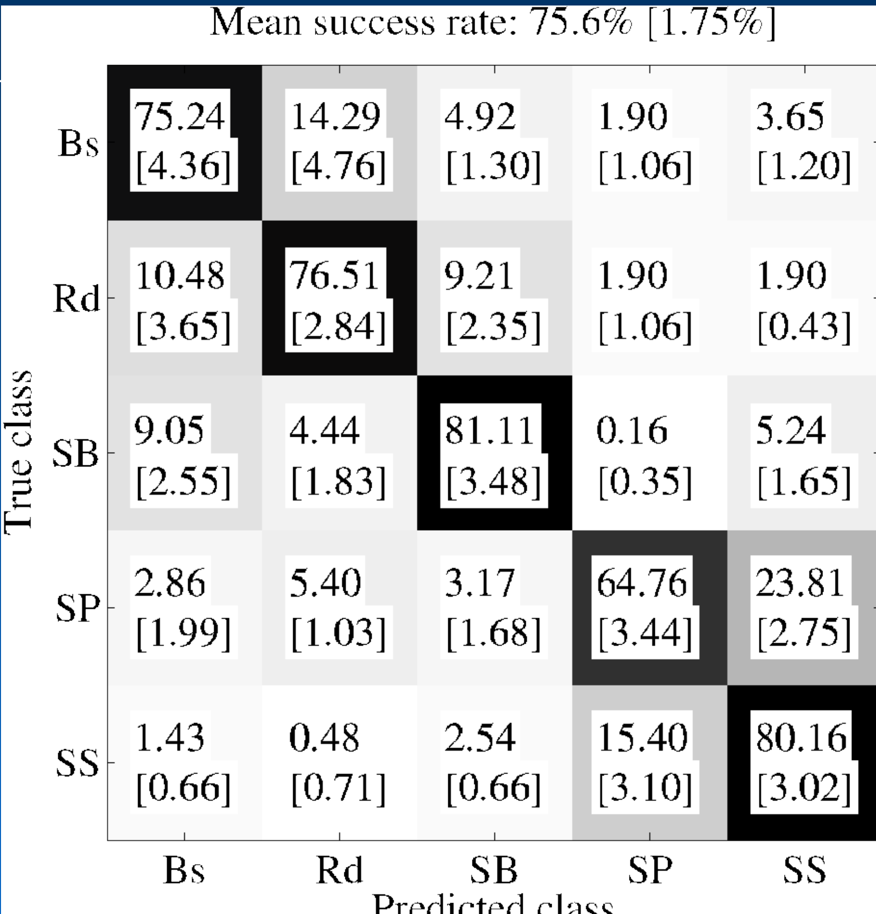

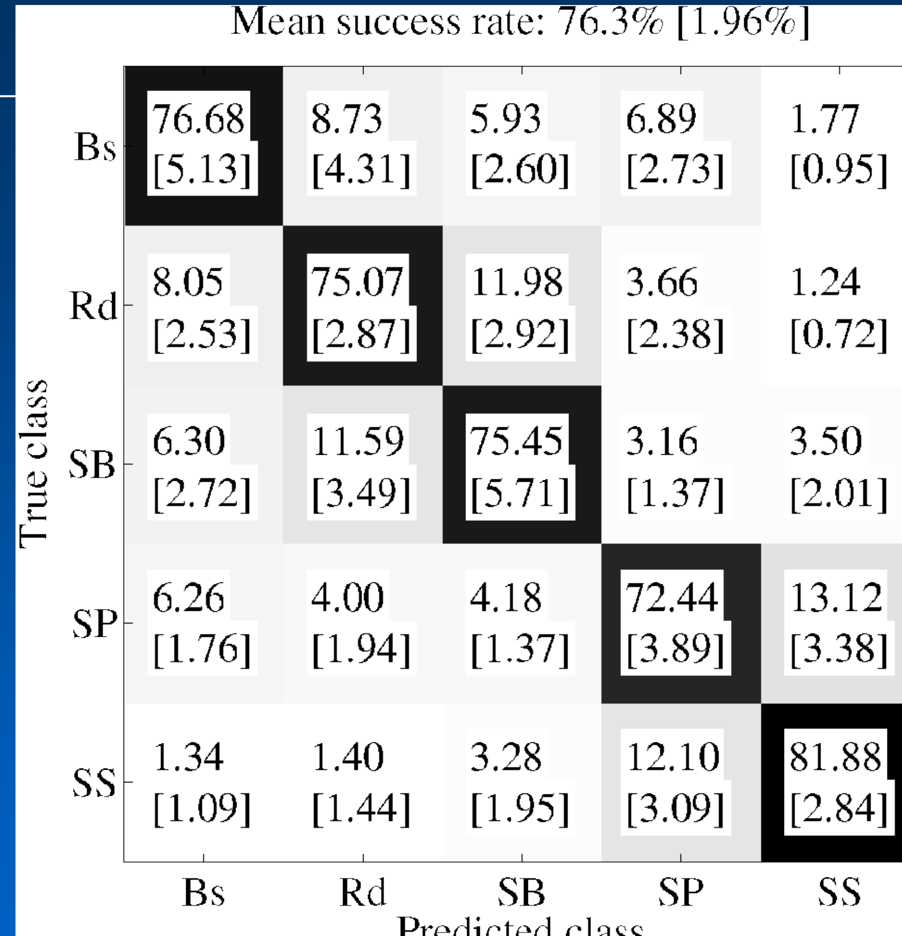

Onset fingerprint coding — MFCC coding (whole note)

Results are very comparable: but the confusions differ.

Results 2: train on McGill + test on Iowa dataset

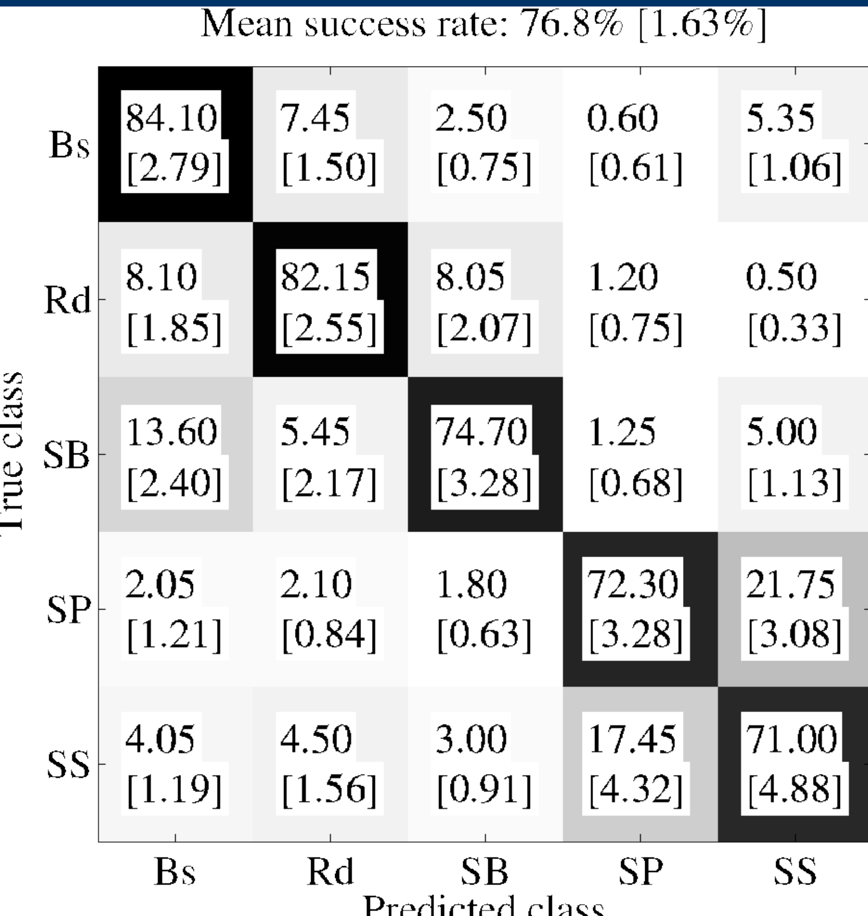

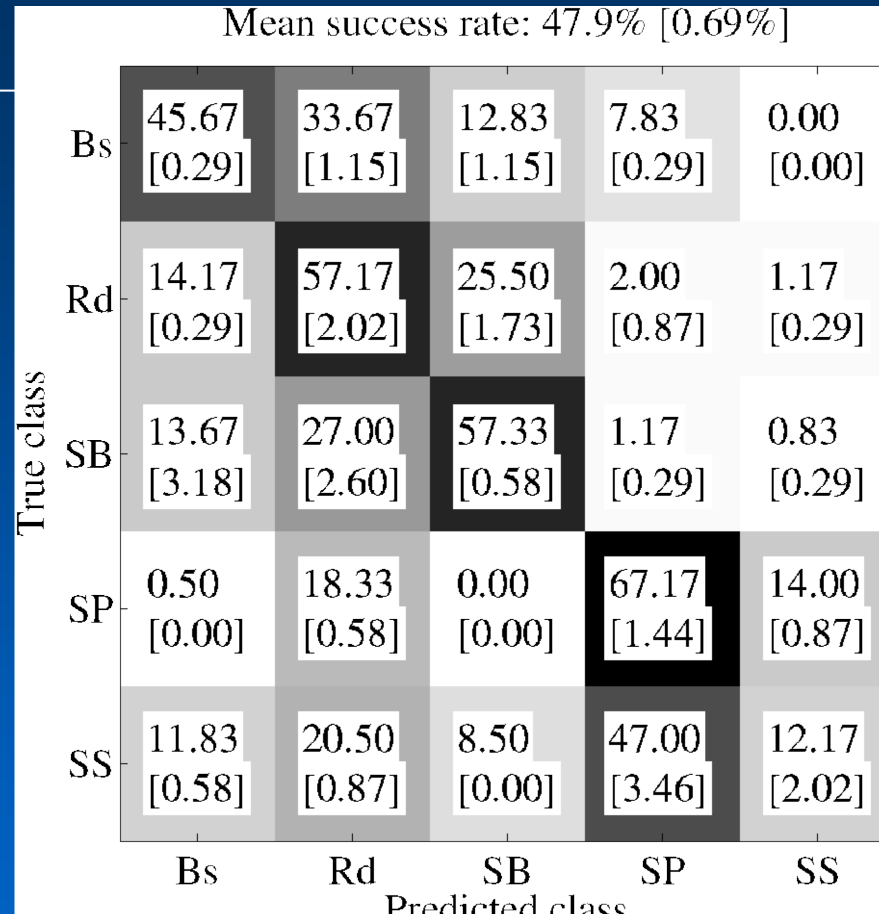

Onset fingerprint MFCC

Note that the onset technique holds up much better on a differently recorded dataset.

So?

- We argue that sensory preprocessing is important
  - And should be applied as early as possible
    - Don’t try to do everything all at once
  - Even at the sensor!
    - Delbruck’s camera
    - Multi-element bandpassing microphones
    - See http://siliconretina.ini.uzh.ch/wiki/index.php
    - Or adaptive multi-element microphone

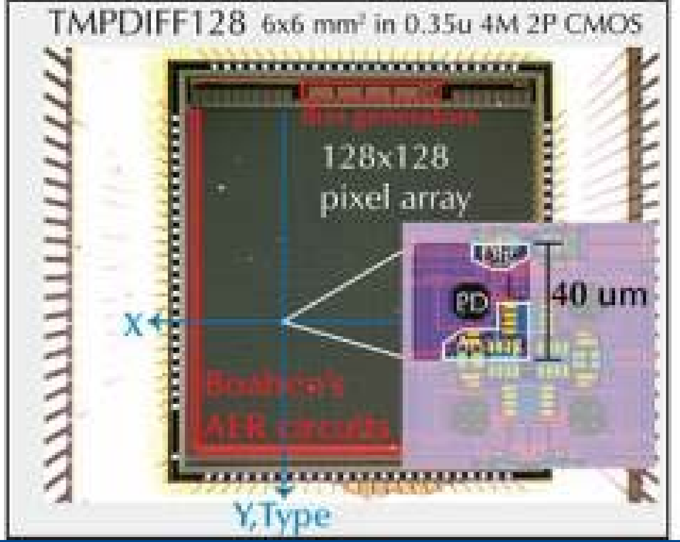

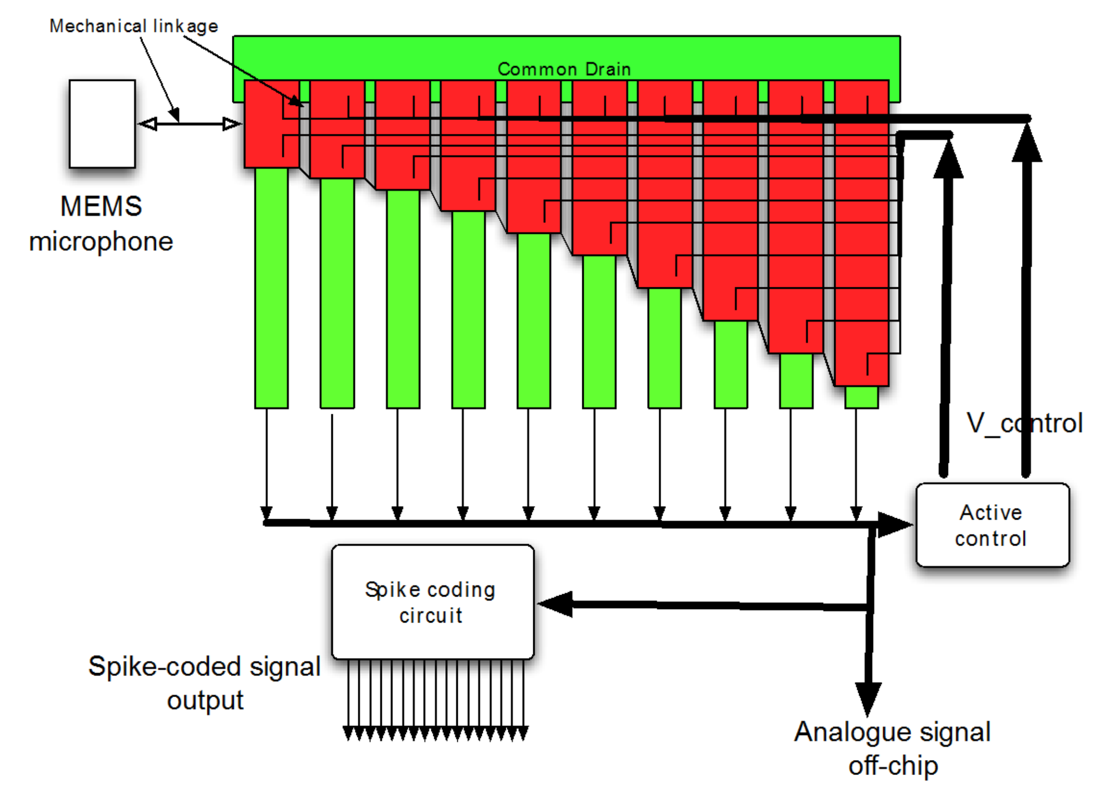

Context and CI?

- At what level should context be taken into account?
  - Within-sensory-modality context?
    - As early as possible (at sensor, brainstem)
  - Multi-sensory context?
    - Needs initial context resolved first
      - mid-brain and later at cortex level
  - Temporal context…

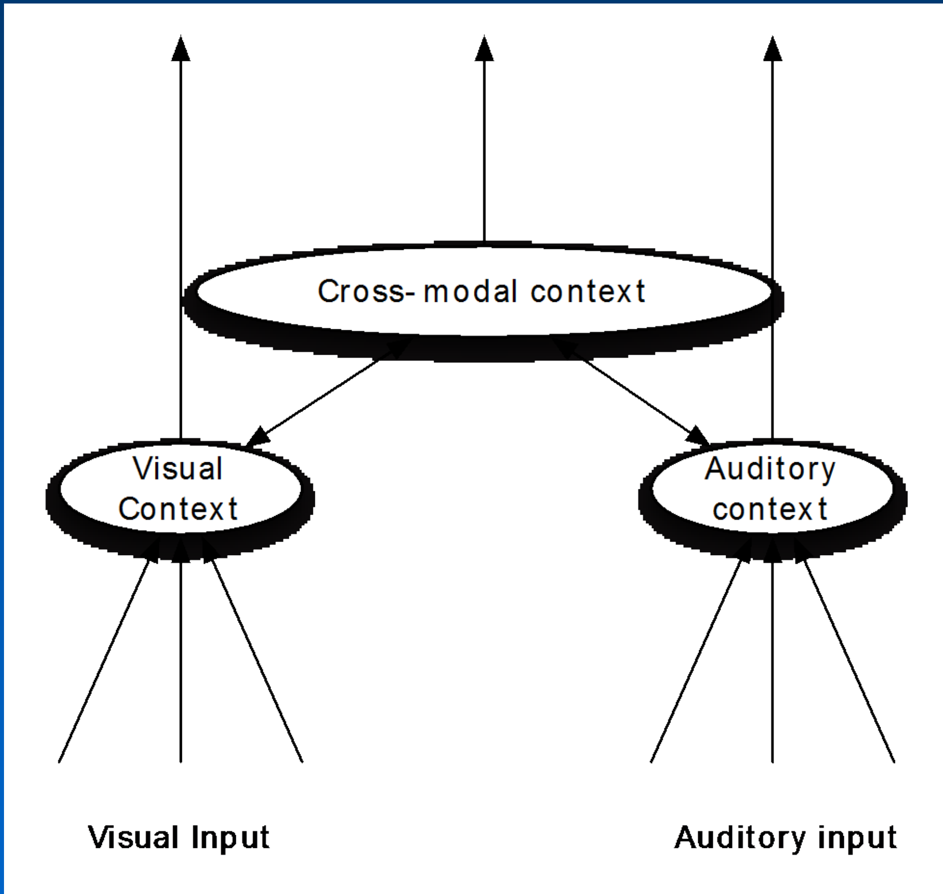

Overview

- **What is Computational Intelligence?**
  - **Computational? Intelligence?**
  - What's intelligence
    - And have we captured it?
- Context and Invariance
  - How and why these matter for everyday intelligence.
- Time
  - Temporal context

Time (temporal context)

- Perceptual time is different from physical time.
  - Not surprising:
  - perceptual images are different from the light that causes them, and
  - perceptual sound differs from the pattern of pressure waves that are the physical bases of perceptual sound
- Percepts are internal versions of physical entities
  - Usually (sometimes they are illusory or imaginary)

How shall we tell the minutes?
The time it takes to swipe
A lonely pint of Guinness
Or load a friendly pipe?
O. St. J. Gogarty

Temporal context and CI

- How does this relate
  - To Computational Intelligence
- In terms of plain categorisation or function approximation
  - Not at all

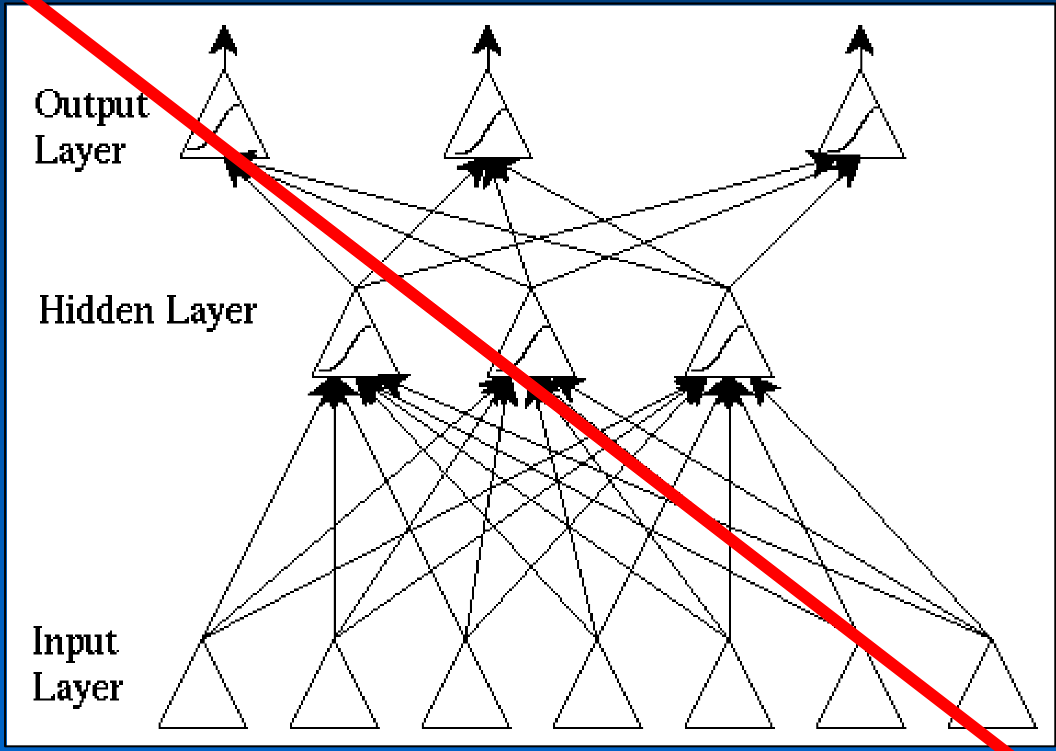

Temporal context

Clearly time matters for prediction and for time series classification

- But how should it be used
- Simple ordering?
  - OK when time is a sequence of integers
  - Usable when one simply takes a window of input for classification or prediction
  - Clearly a very impoverished view of time

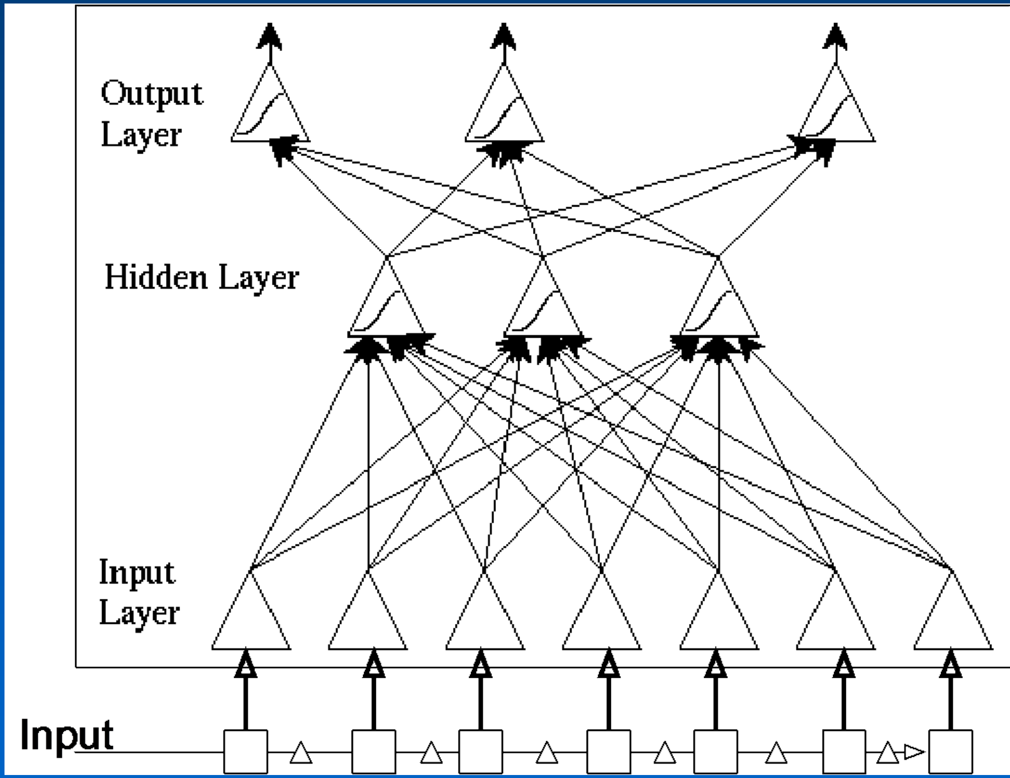

Time and CI systems

- Simple window-based/integer time systems
- Discounted time systems
  - Reinforcement Learning (RL), Time difference Learning (TDL)
- Using elements with time constants
  - Leaky integrate-and-fire (LIF) neurons
  - Dynamic synapses
- Reservoir networks
- Even in GOFAI systems, there's often a blackboard used to hold input over time
  - Though often indefinitely

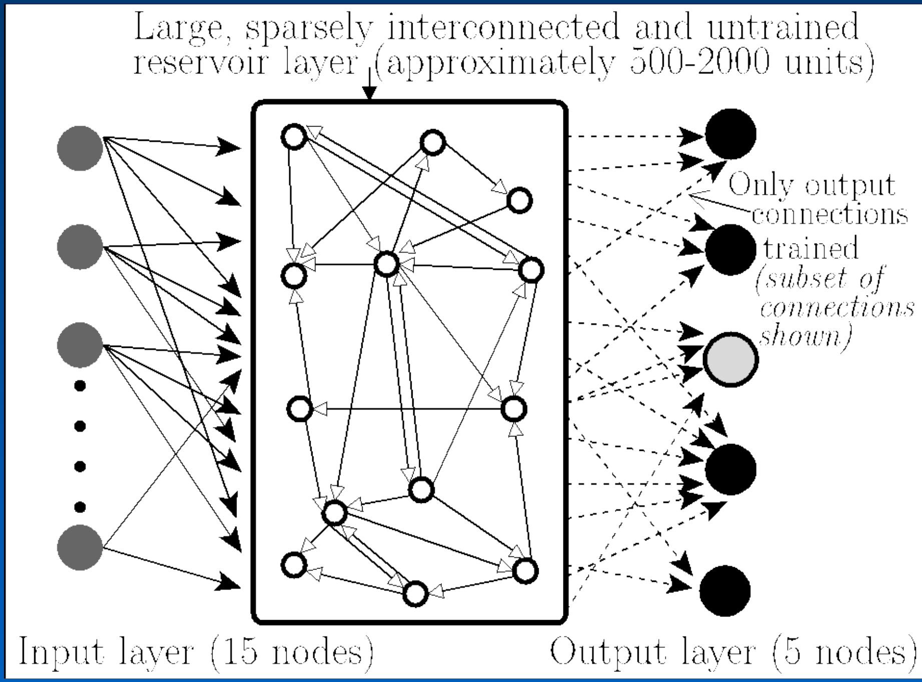

Reservoir network (of LIF neurons) for classification of time-varying signals.

Why does perceptual time matter?

- How we interact with everyday objects is strongly influenced by dynamic and physical constraints
  - Not moving too fast – yet moving fast enough
  - Dynamics of everyday objects: physical constraints
  - Speed of movement, resonance, oscillation, gravity
- So it makes good sense to imbue synthetic systems with a similar sense of time, if they are to interact with everyday objects

Perceptual time: from psychology

- Dunne 1925 (An experiment with time)

  "attention is never really confined to a mathematical instant. It covers a slightly larger period"

- But what is *the duration of the present instant*?
  - Also known as *the specious present* and the *mental present*
- Two different interpretations
  1. The period below which the present seems indivisible: set to 40 to 50ms by von Baer and Poppel
  2. The period for pre-sematic temporal integration, set at about 3 seconds by Poppel.

Perceptual time and features

- How long does a coherent percept take to form?
  - Visually? Auditory? Olfactory? Tactile? Or cross-modal?
  - Appears to be bounded below by about 40-50ms
    - And can certainly be longer
- Auditory example: increases in energy in different areas of the auditory spectrum
  - …onsets (including vowel onsets /s/ onsets, etc.).
  - Seem to need to be within 40-50ms if they are to be considered as a single entity
- And repeated pulses merge into a single sound at 18-20 pulses/second
- More difficult to perform experiment with other sensory modalities.
  - But consider movement percepts in video projection.

On the neural construction of perceptual time

- There are many different rhythmic oscillations in the brain
  - With time constants from about 0.1Hz to about 400Hz.
  - But energy spectrum from local field potentials is not evenly spread through the spectrum
- But what are the time constants for perceptual time?

Perceptual time and neural systems

- Cortical columns seem to be the basic unit of the cortex.
- They are made up of excitatory cells (mostly pyramidal cells) and inhibitory cells (of many different types)
- Excitation/Inhibition systems often oscillate

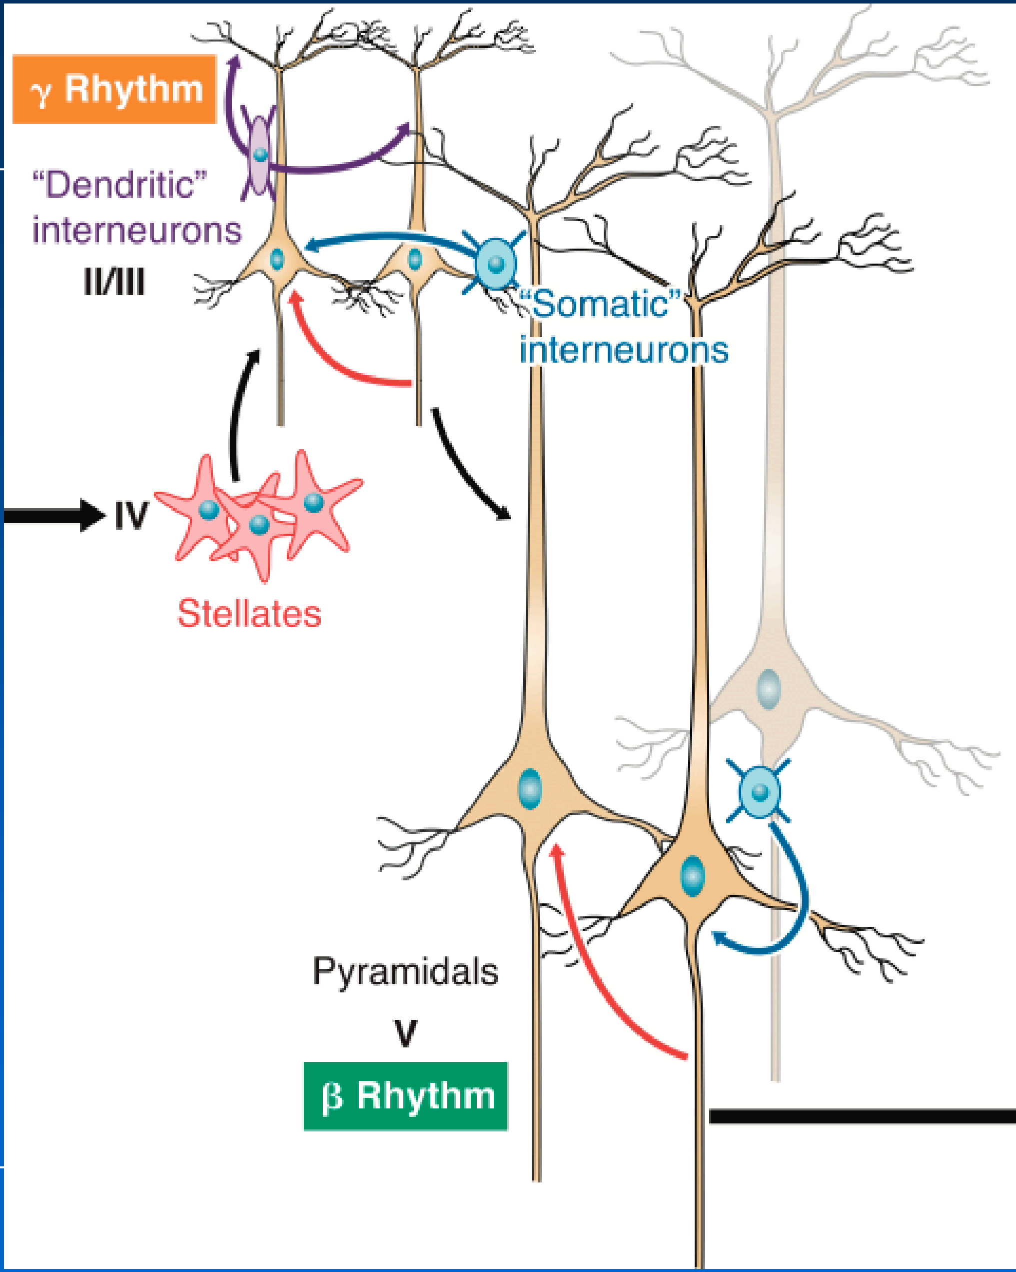

What do cortical columns do?

- Cortical column local field potentials seem to be in β and γ ranges (about 15 to 80Hz)
- Cortical columns appear to integrate a number of inputs
  - Both direct (external/sensory) and contextual
- And produce outputs regularly
  - At about 40ms intervals
- They seem to code their inputs, resetting every 40ms
  - Looks like a good candidate for a neural correlate of (fast) perceptual time
  - Time of neural firing within an oscillatory cycle may be important

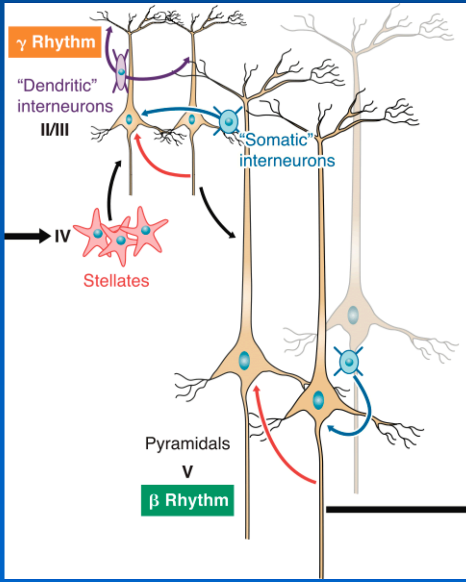

Why is the perceptual instant this length?

- Two factors
  - What matters in everyday situations
  - What can be achieved neurally
- Everyday situations:
  - In 40 ms,
    - a sound travels about 13 metres.
      - And we often need to fuse what we can see with what we can hear
    - A ball drops about 8mm (drops 3cm in 80ms)
  - Consonant duration starts at about 40ms.
  - Running fast (15 mph) one moves 25cm
  - In other words, 40ms fits well with everyday movement timings.
- Neurally?
  - Given the speed of a neuron, 40ms is a reasonable length of time for some parallel processing
    - Much faster timings do exist, for specialised purposes, like sound source direction finding
    - But even there, the sounds tend to be at least 40ms long!

What about longer perceptual instants?

- Multiple cycles for longer perceptual integration
  - Integrating across a number of columns
- Longer times? 1-3 seconds?
  - Perceptual interpretation of a number of events
    - Each bounded below by 40ms
  - Neural basis is unclear
    - But there are candidates
    - E.g. synaptic changes, plus intra-cortical loops.

Concluding on time

If we want CI systems to display everyday intelligence

- We need to match the “perceptual time” of a CI device to everyday event timings
  - To be able to interact with everyday objects in an everyday way
- We don’t need to do so in the same way that neural systems do this
  - But we do need to make the integrating interval appropriate to the environment.

Putting this together

- What makes us accredit an animal with intelligence?
  - Appropriate behaviour
    - Really rapid reactions sometimes
    - More deliberative at other times
    - Using context appropriately
  - Learning
    - Not making the same mistake twice
      - Evolutionally critical: no second chance!

One view: unified operation.

- Intelligence comes from systems working together.
  - Very fast reactions (brainstem: single modality)
  - Slower integration across some modalities, immediate time (mid-brain)
  - Still slower, deliberative, integrating over longer periods (cortical)
  - Very slow, long-term operation

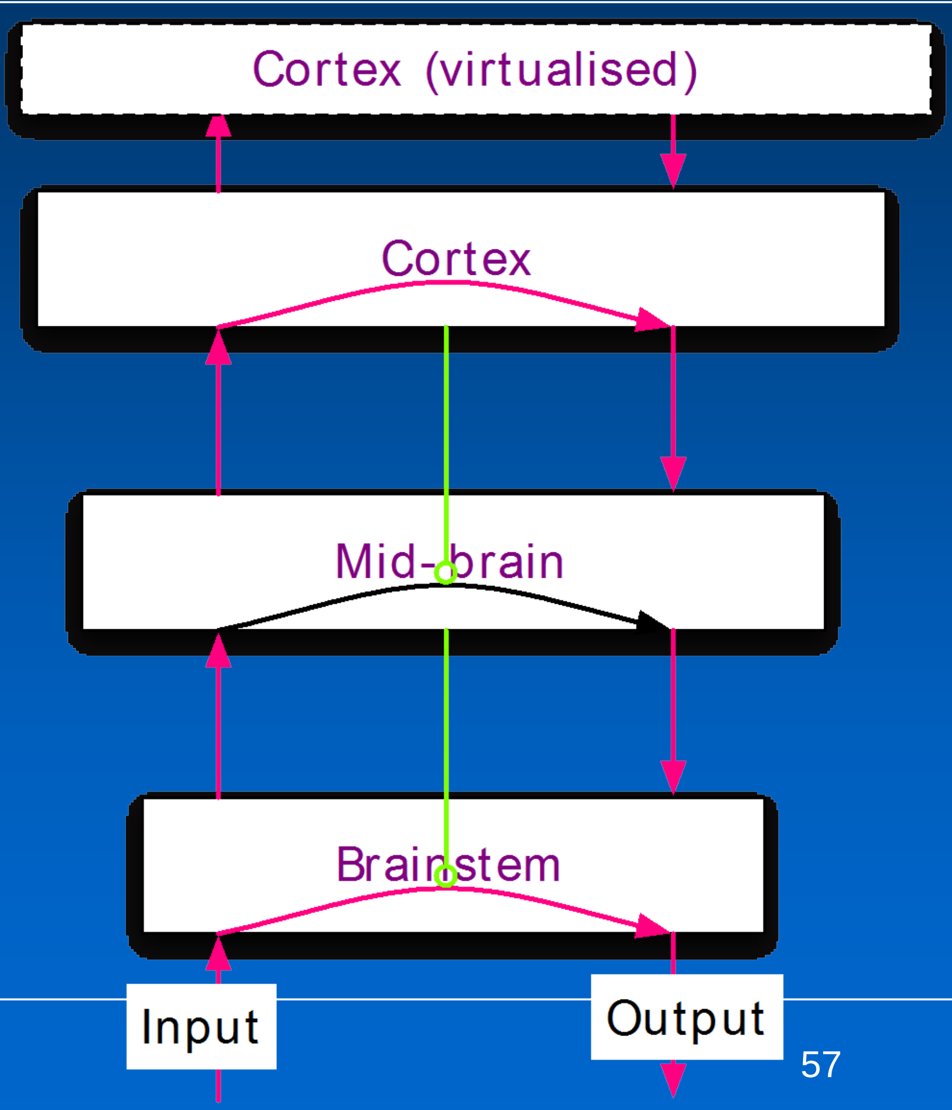

Context at more than one level

- Cross-modal context
  - Initially at midbrain level
  - Collicular
- Temporal context
  - At midbrain level
  - At cortical level
- Learning?
  - Brainstem and midbrain in early life
    - And evolutionary as well
  - Cortical throuhout life

Going out on a limb…

A mathematics for intelligence

- The INBIOSA project
  - Link to white paper
  - http://www.cs.stir.ac.uk/~lss/recentpapers/INBIOSAWhitePaper(BookVersion).pdf
- A general theory of living systems
  - Seeking a new mathematics for describing and understanding living/intelligent systems
- Interested in Category Theory
  - As a mathematics that includes mathematics
  - Self-reflexive.

Category theory for learning

View model induction (learning with or without a teacher)

as creation of a structure that represents what's out in the world.

Category theory view:

Sensory system induces functors between categories

(and we could add arrows the opposite way to include motor capabilities to give us active sensing)

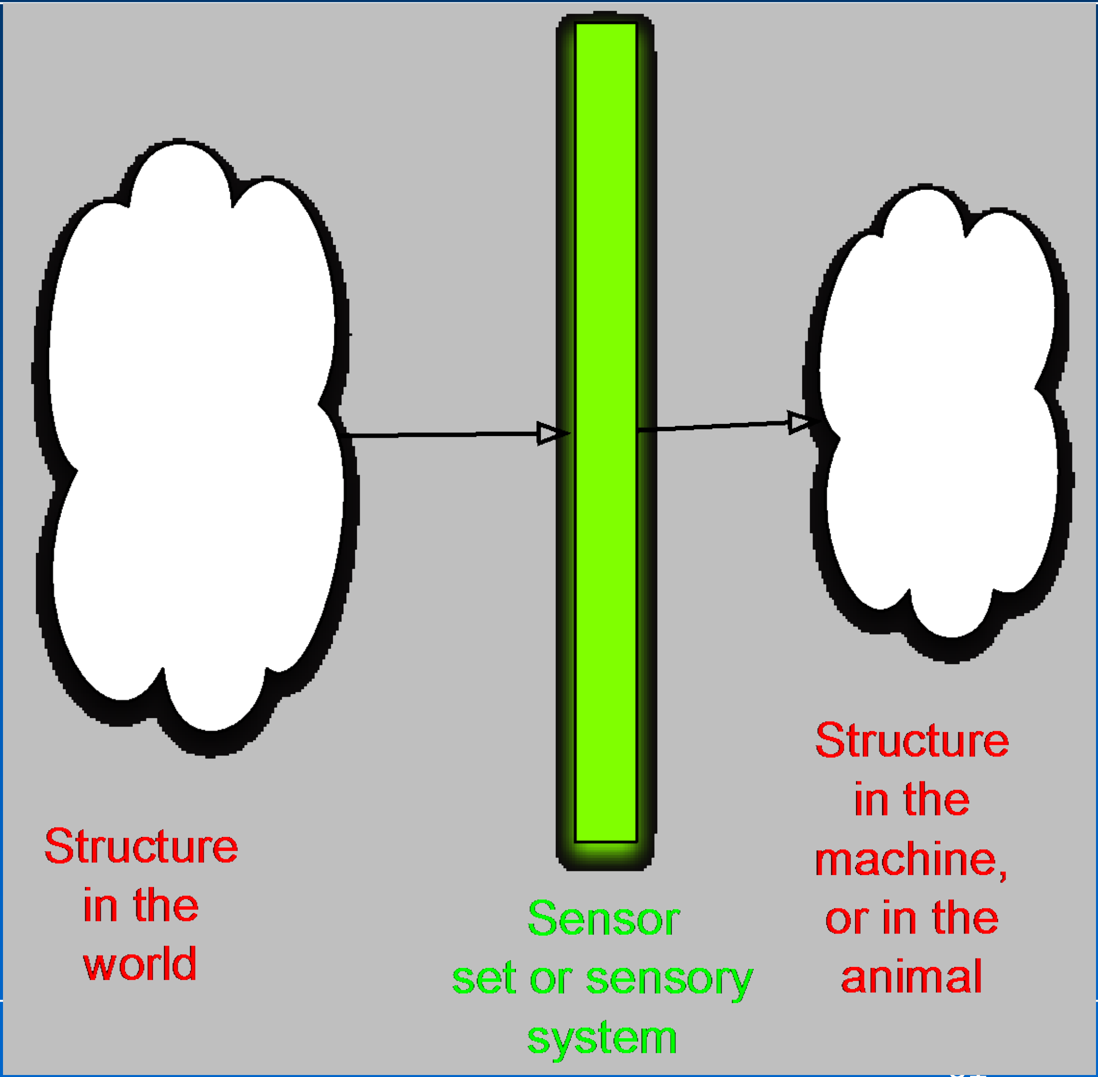

The Computationalist view

"Computationalism is the view that cognitive capacities have a computational explanation or, somewhat more strongly, that cognition is (a kind of) computation" *Piccinini, Oxford Handbook of Philosophy and Cognitive Science forthcoming.*

- We already know the answer:
  - Cognition can be implemented using …
    - A computer (1960's/70's)
    - A Perceptron-based machine (1960's)
    - A logical inference machine (Japan, 5th generation project , early 1980's)
    - A Back-propagated delta-rule based machine (mid 1980's)
    - A Reinforcement Learning Machine (1980's/90')
    - A Bayesian Inference machine (2000's: Hawkins & George, for example)
  - Perhaps it has aspects of all (or at least some) of these
    - And perhaps it is rather more than these.

Replacing Computationalism: A new Pythagoreanism?

- F. Zalamea's *Synthetic philosophy of contemporary Mathematics*, (based on earlier work, e.g. Lautmann)
  - (2012) translation of "Filosofía sintética de las matemáticas contemporáneas",
- Zalamea lights a way towards a new philosophy of Mathematics that brings together
- the constructive imagination of what he calls *eidal* Mathematics
- the Physically based *quiddital* Mathematics,
- and the idea of Mathematics of mathematics in *archeal* Mathematics

Pythagoreanism (2)

- In Zalamea’s quiddital Mathematics, I see the handiwork of God, whether in biology, or physics, or any other branch of science.
- But in eidal or archeal Mathematics, I see possibilities that might or might not be in the actual Universe. I see connections between the possible workings of the Universe: perhaps we see into the mind of this Pythagorean God.

And why this matters…

- A Pythagorean God is not a deity that helps us directly to live our lives. It’s not a God in the usual sense in Abrahamic religions. A Pythagorean God is more in the background, more about the unity of the Universe, more about the underlying structure.
- Perhaps a view of the world that brings together Science and God
  - Perhaps we can then build machines that are aware, alive, in an image of god
  - And not try to bring aware and intelligent systems down to the level of clockwork

General Conclusions

- We argue that there are levels of everyday intelligence that can be achieved without recourse to 1st person awareness/ consciousness (etc.) research
  - Not arguing (here!) for hybrid silicon/living systems (cyborgs): different issue!
  - At the sensory level, invariances are important
  - Context and time are critically important
    - Context: across modalities, over time,
    - Time: physical time and simple ordering are different from perceptual time
      - Perceptual time underlies much of how we perceive the world
- But if we want to go further we may need to go beyond Computationalism.

…and thanks to …

- Bill Phillips
- Plamen Simeonov
- Bruce Graham
- Michael Newton
- Andrew Abel
  - For useful discussions
  - … but the responsibility for the ideas (and the blame too) lies with me alone!