Cognitive science for machine learning 2: Empirical Methods

Nick Chater

OVERVIEW

1. OBSERVATIONS AND EXPERIMENTS
2. THE SPECTRUM OF EXPERIMENTAL METHODS
3. CAN YOU PLAY 20 QUESTIONS WITH NATURE AND WIN?

OBSERVATIONS AND EXPERIMENTS

THE CONTRAST

Observations

- Observe human performance in some interesting scenario
- May merely use verbal reports (introspection)
- Or, more usually,turn reports into a task (cf. Signal detection theory, Felix, next session)
- But the most interesting observations are often very general and qualitative

Experiments

- Manipulate one or more independent variables
- Measure one or more independent variables
  - E.g., responses,reaction times (Felix, next session)
  - Body or eye movements
  - Activities of individual neurons
  - Pupil dilation
  - Galvanic skin response
  - EEG waves
  - Cerebral blood flow (fMRI)

OBSERVATIONS, EVEN WHEN INFORMAL, ARE DATA, TOO

Example: Human category learning

- The observation: people can learn to classify the same objects along different dimensions, very flexibly
- But traditional instance-based psychological models, based on a fixed similarity metric don't…

Katherine Heller

Adam Sanborn

Heller, K., Sanborn, A. N., & Chater, N. (2009). *NIPS*

CATEGORIES ARE DEFINED BY DIFFERENT *TYPES* OF DIMENSION

(e.g., object vs substance, Linda Smith et al)

For objects: Shape matters, colour doesn’t

BUT FOR SUBSTANCES...



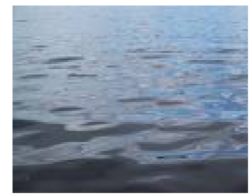
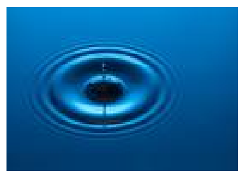




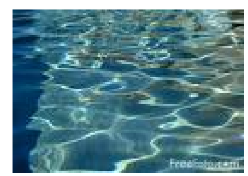

Shape irrelevant; colour, texture, movement matters

Generalization for just two instances can be powerful---but puzzling for conventional learning accounts

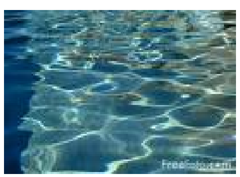

?

?

A Bayesian approach to learning this category structure

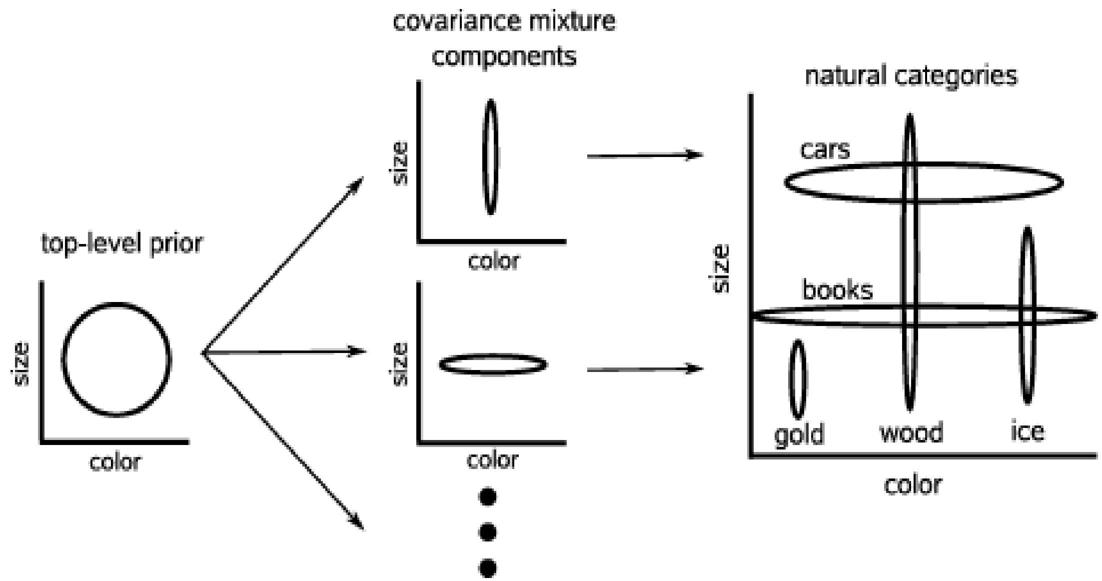

NOW COMPARE WITH EXPERIMENTS...

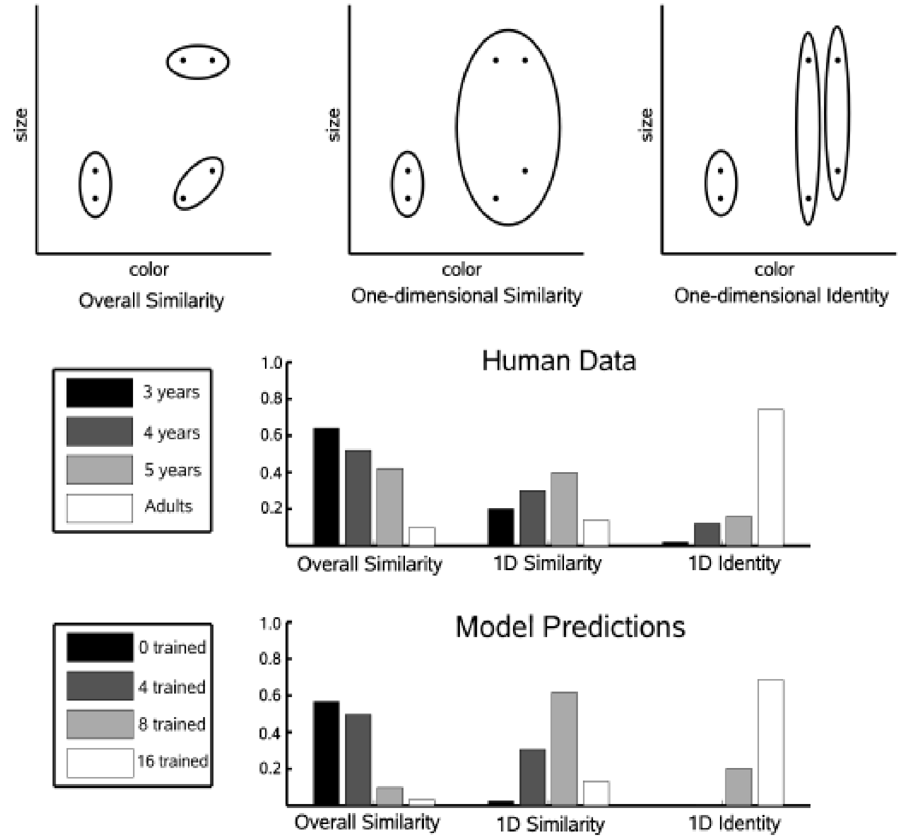

Linda Smith (Psychological Review, 1989) showed that children gradually learn to categories more locally, and less globally...

...AND CONDUCT NEW EXPERIMENTS

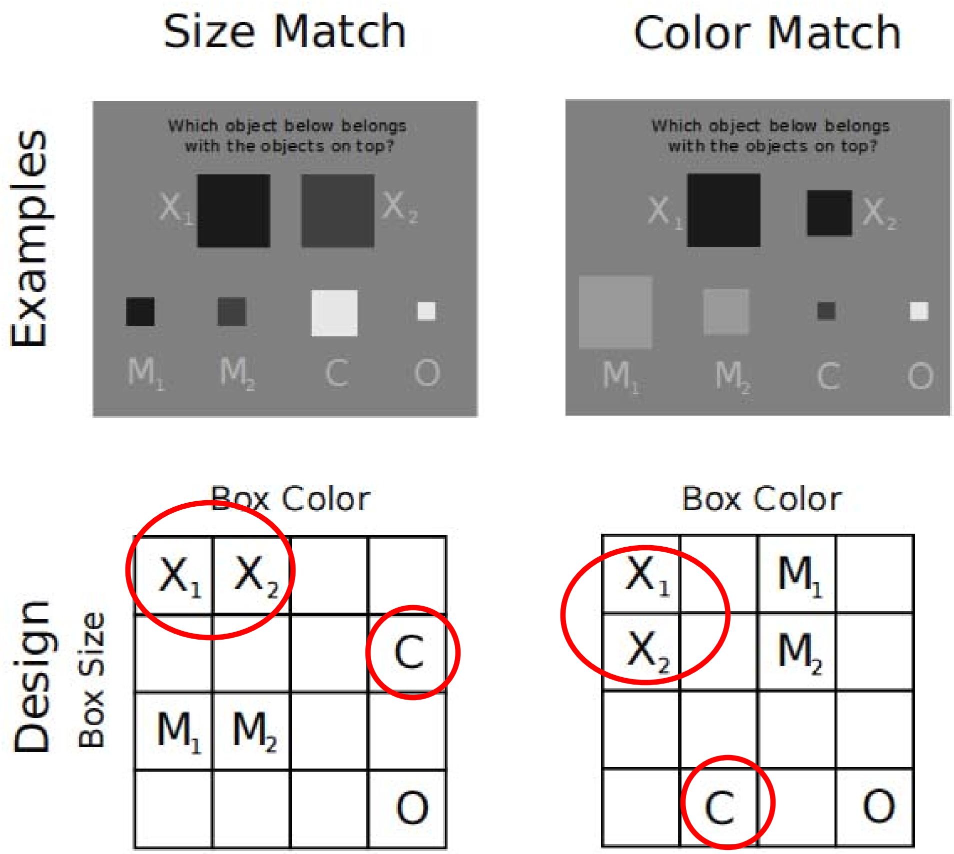

Sanborn, Heller & Chater, under review

AND QUALITATIVE OBSERVATIONS THAT PEOPLE ROUTINELY…

- Infer causality
- Infer agency
- Infer communicative intent

Inferring causality from observation

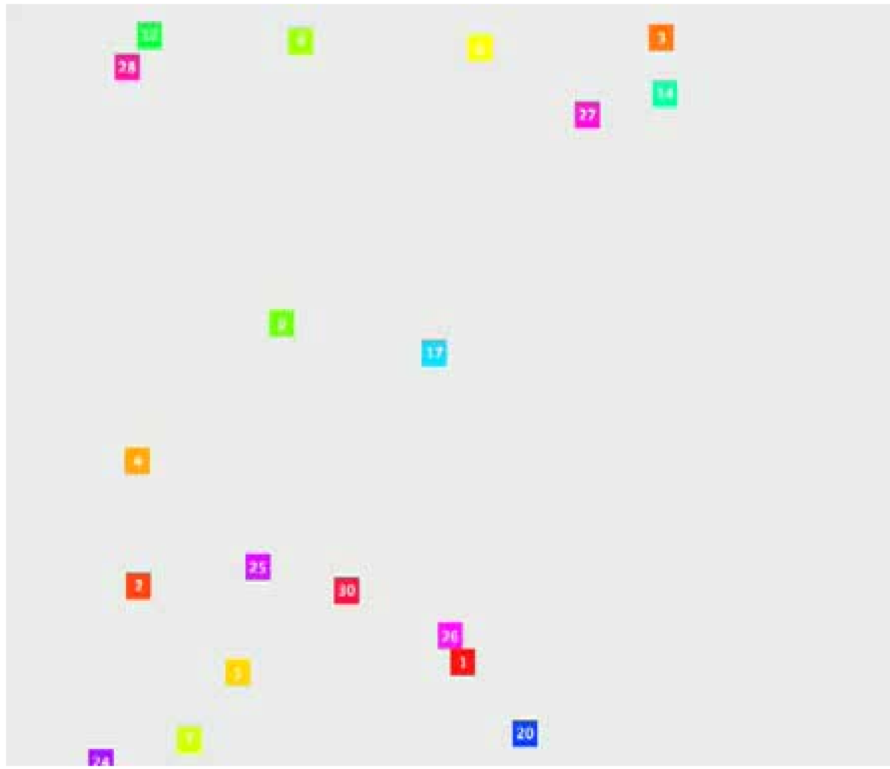

Cf. Pearl (2000):

$$\boldsymbol{a} = \boldsymbol{a}(F,\mathrm{m}) = F/m$$

Not $\mathbf{F} = \mathbf{F}(m,a) = ma$

Inferring agency from observation

From Heider & Simmel, 1944:

Inferring “theory of mind” (with, or without?) powerful innate constraints)

Inferring communicative intentions from observation

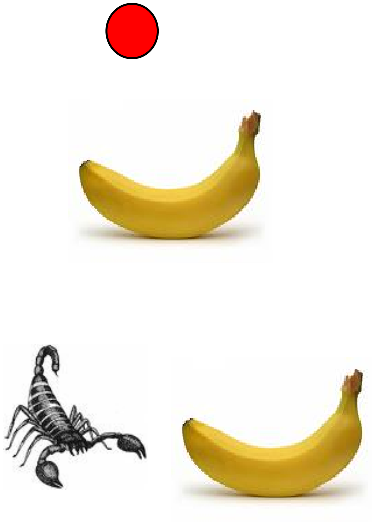

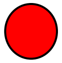




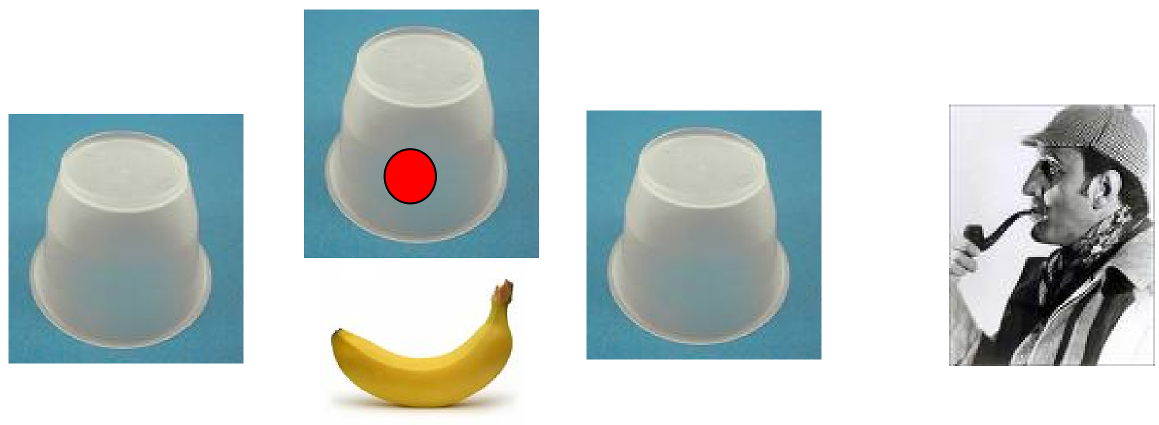

Common knowledge is crucial

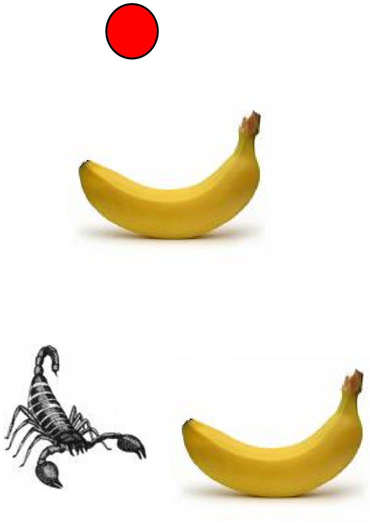







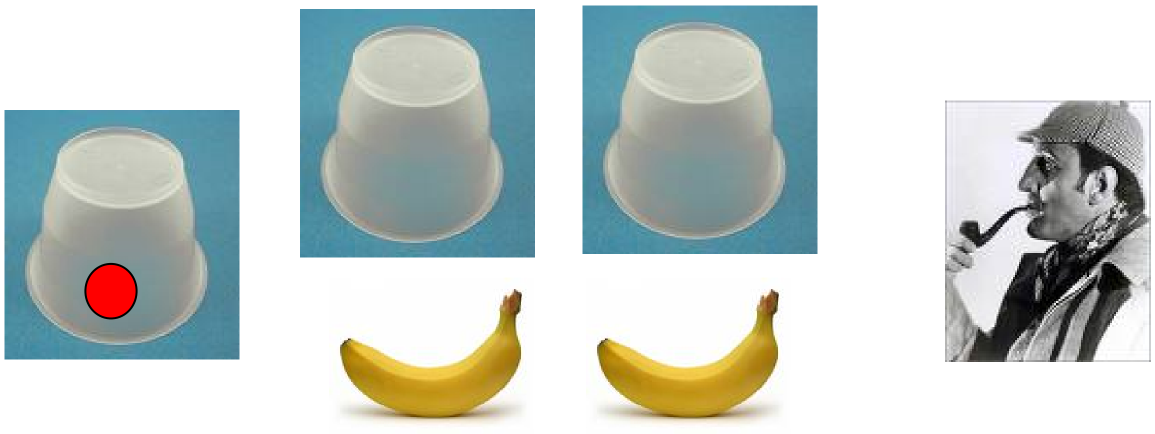

Communicative, rather than mere signalling, requires common knowledge of joint intentions; and “team reasoning” (Bacharach, Sugden). How might this be learned???

THE STUDY OF VISION IS HEAVILY BASED ON OBSERVATIONS, NOT MERELY EXPERIMENTS

- Kanizsa triangles

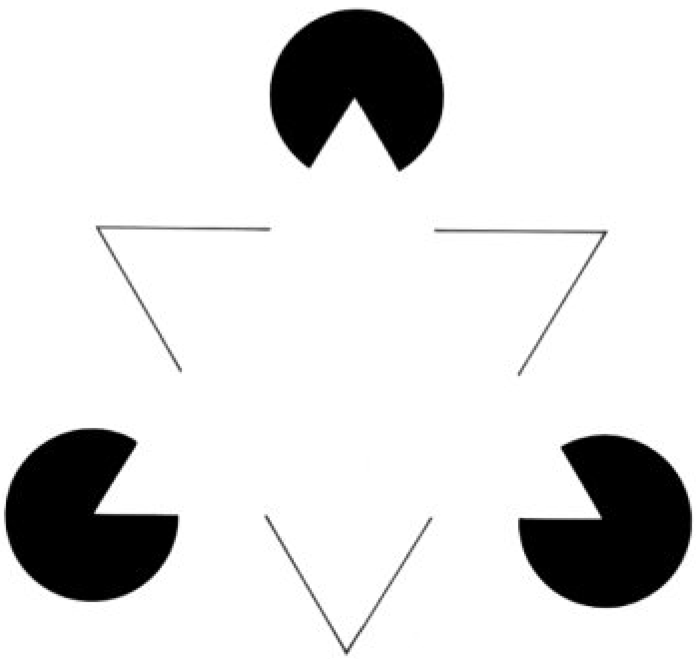

- What is the origin of the lightness changes on the edges of the 'virtual' triangle?

A 3D variation…

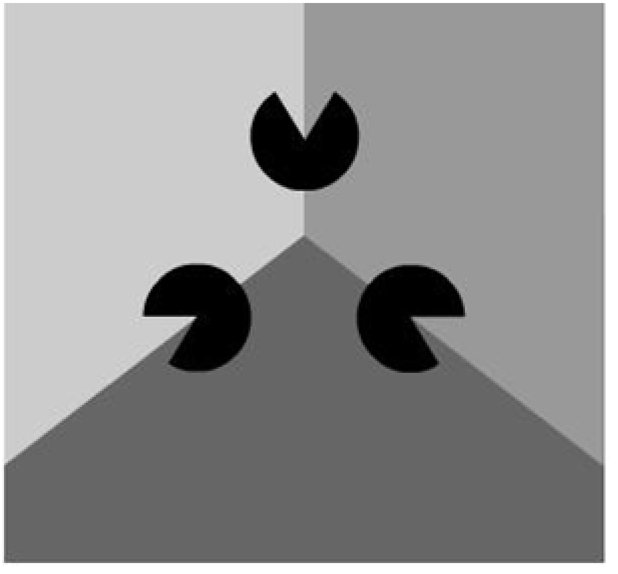

Note the difference... beginnings of an experiment

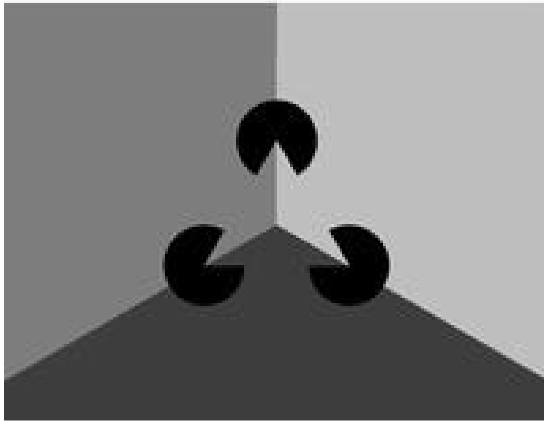

Pietro Guardini & Luciano Gamberini: and they have a lovely *moving* 3D version

OBSERVATIONS MAY SUGGEST GENERAL PRINCIPLES – E.G., FAVOUR THE SIMPLEST EXPLANATION

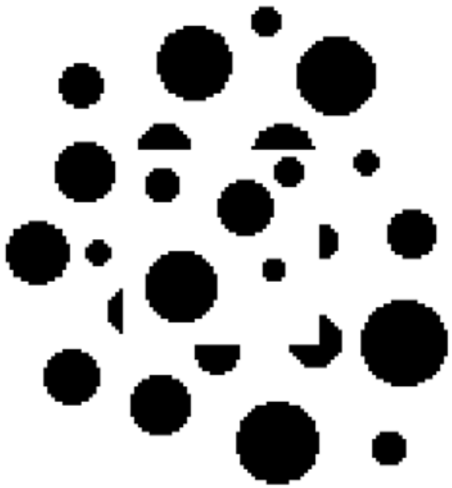

Kanizsa

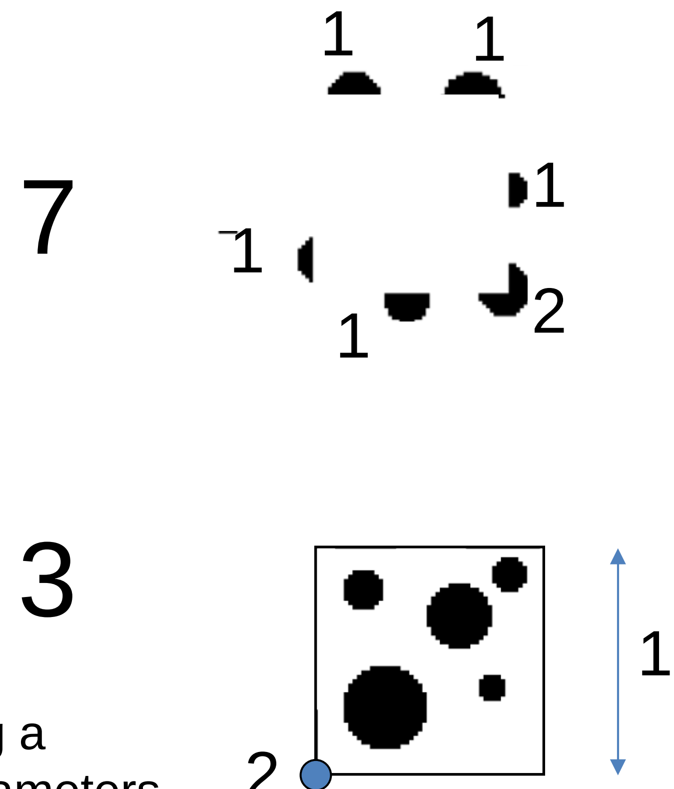

Find simple abstract patterns… e.g., postulating a square needs 3 parameters; simpler than 7 parameters for accounting for ‘cuts’ in circles separately

IF SO, CAN ONLY SHARE INFORMATION WITHIN A SINGLE OBJECT

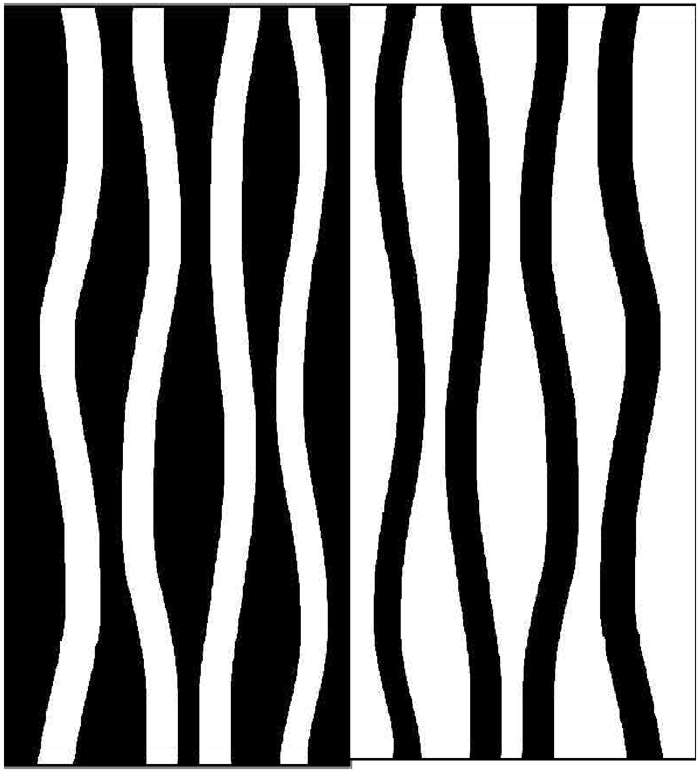

- Figure-ground separation is different in the two halves

IF SO, CAN ONLY SHARE INFORMATION WITHIN A SINGLE OBJECT

- 1 vs 2 codes for parallel curves
- Such observations reveal quite a lot about visual structure
- Relation to image processing by machine?

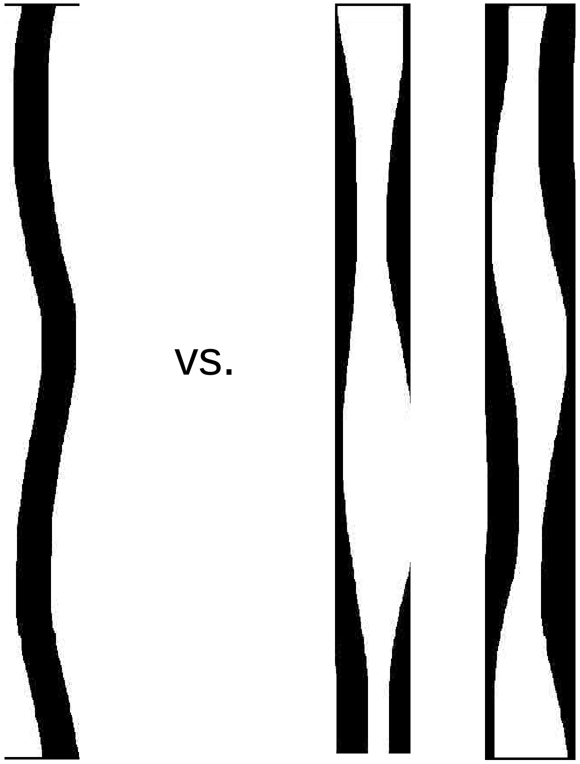

WE CAN ALSO OBSERVE, IN A LOT OF DETAIL, THE PATTERNS IN LANGUAGE

- Methodology in linguistics
  - Thesystematisation of intuitions about well formedness

*John Mary likes

John likes Mary

pif vs *fpi

- Or semantics (e.g., the binding contraints):

- John_j saw himself_j
- John_j saw him_k
- He_i saw John_j
- John_j saw a picture of Mary and John_k

...OR OBSERVATION MAY INVOLVE DISPLAYING THE STRUCTURE IN LINGUISTIC INPUT

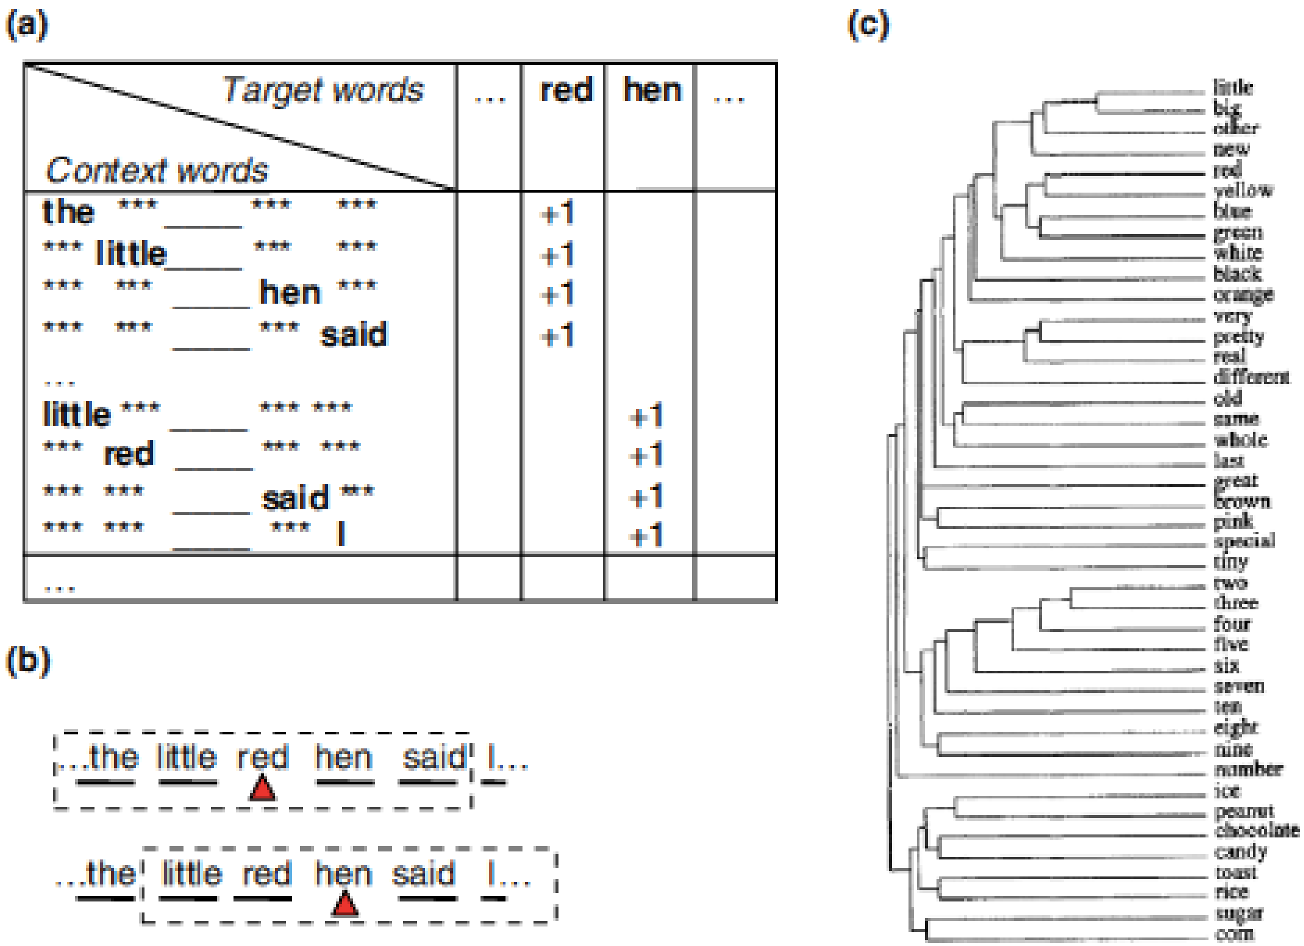

DISTRIBUTIONAL ANALYSIS TO RECONSTRUCT SYNTACTIC CATEGORIES (REDINGTON, CHATER & FINCH, 1998)

OR APPLY BAG OF WORDS TOPICS MODELS TO FIND SEMANTIC CATEGORIES

217	274	126	63	200	209
INSECT	SPECIES	GENE	STRUCTURE	FOLDING	NUCLEAR
MYB	PHYLOGENETIC	VECTOR	ANGSTROM	NATIVE	NUCLEUS
PHEROMONE	EVOLUTION	VECTORS	CRYSTAL	PROTEIN	LOCALIZATION
LENS	EVOLUTIONARY	EXPRESSION	RESIDUES	STATE	CYTOPLASM
LARVAE	SEQUENCES	TRANSFER	STRUCTURES	ENERGY	EXPORT
42	2	280	15	64	102
NEURAL	SPECIES	SPECIES	CHROMOSOME	CELLS	TUMOR
DEVELOPMENT	GLOBAL	SELECTION	REGION	CELL	CANCER
DORSAL	CLIMATE	EVOLUTION	CHROMOSOMES	ANTIGEN	TUMORS
EMBRYOS	CO2	GENETIC	KB	LYMPHOCYTES	HUMAN
VENTRAL	WATER	POPULATIONS	MAP	CD4	CELLS
112	210	201	165	142	222
HOST	SYNAPTIC	RESISTANCE	CHANNEL	PLANTS	CORTEX
BACTERIAL	NEURONS	RESISTANT	CHANNELS	PLANT	BRAIN
BACTERIA	POSTSYNAPTIC	DRUG	VOLTAGE	ARABIDOPSIS	SUBJECTS
STRAINS	HIPPOCAMPAL	DRUGS	CURRENT	TOBACCO	TASK
SALMONELLA	SYNAPSES	SENSITIVE	CURRENTS	LEAVES	AREAS
39	105	221	270	55	114
THEORY	HAIR	LARGE	TIME	FORCE	POPULATION
TIME	MECHANICAL	SCALE	SPECTROSCOPY	SURFACE	POPULATIONS
SPACE	MB	DENSITY	NMR	MOLECULES	GENETIC
GIVEN	SENSORY	OBSERVED	SPECTRA	SOLUTION	DIVERSITY
PROBLEM	EAR	OBSERVATIONS	TRANSFER	SURFACES	ISOLATES
		109	120		
		RESEARCH	AGE		
		NEW	OLD		
		INFORMATION	AGING		
		UNDERSTANDING	LIFE		
		PAPER	YOUNG		

Griffiths & Steyvers (2004, PNAS).

SUCH ANALYSIS OF THE STRUCTURE OF THE ENVIRONMENT LOOKS A LOT LIKE MACHINE LEARNING

- And non-rigid boundary between
  - characterising the information in the input vs
  - providing a computational model of cognition

- Link with empirical data
  - E.g., Mintz et al 2002, Cognitive Science re: distributional information
  - E.g., Griffiths, Steyvers, Tenenbaum, 2007, Psych Review, on word associations

OBSERVATIONS, WHETHER EVERYDAY OR CLEVERLY CONSTRUCTED, ARE VITAL CLUES

- Vision (interpretations of visual images, illusions)
- Language (intuitions underlying linguistics, semantics)
- Common-sense knowledge and reasoning

- But observations are not enough---in particular to establish *causal* relationships between variables, it is often vital to conduct experiments

THE SPECTRUM OF EXPERIMENTAL METHODS

Simple psychophysics - detection

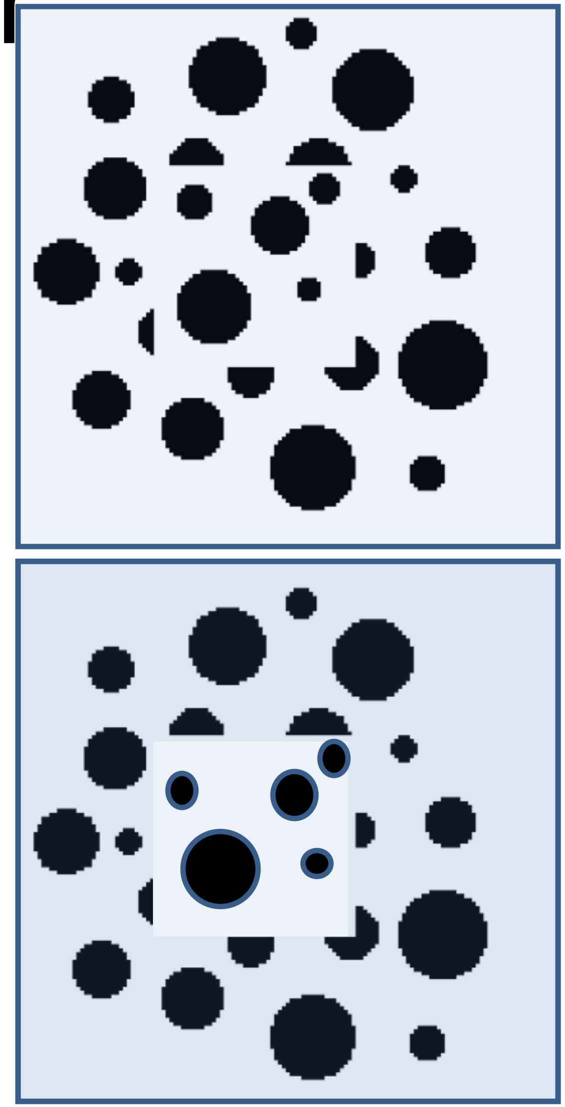

- On half the trials, there really is a faintly brighter square
- 2AFC response
- Manipulate, e.g., whether/how many blobs are cut by the square/size of blobs
- Measure detection rate (signal detection theory)

Simple psychophysics - detection

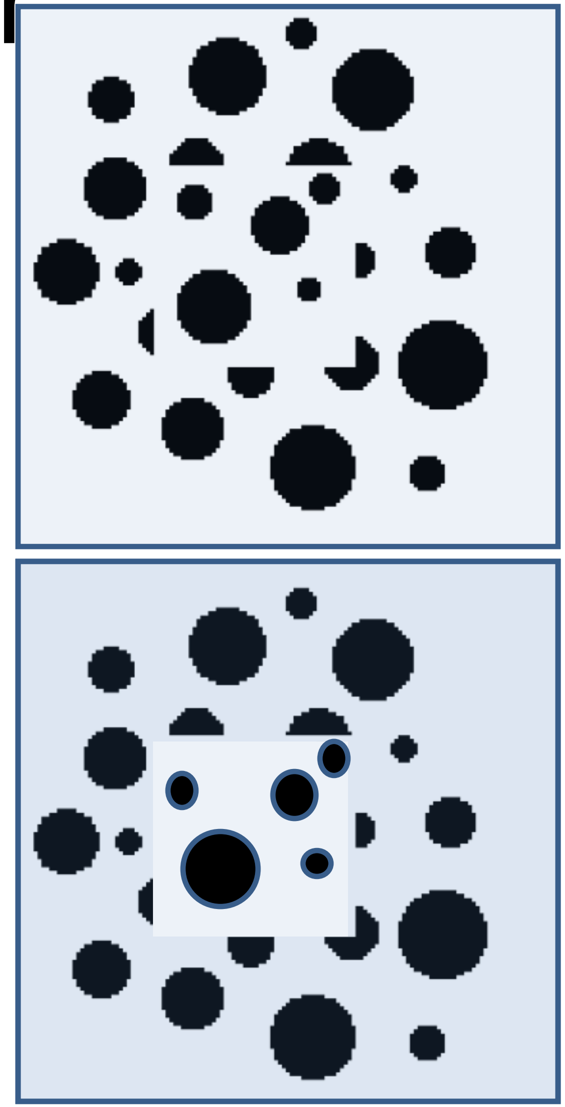

- We might manipulate display duration
- Response deadline
- Density of dots
- Size of the stimulus
- Side of the square...

- But we can also consider a richer dependent variable, e.g., RTs

FROM RESPONSES TO REACTION TIMES

- Helmholtz measure of speed of nervous conduction (roughly 27m/s)

FROM RESPONSES TO REACTION TIMES

- Helmholtz measure of speed of nervous conduction (roughly 27m/s)

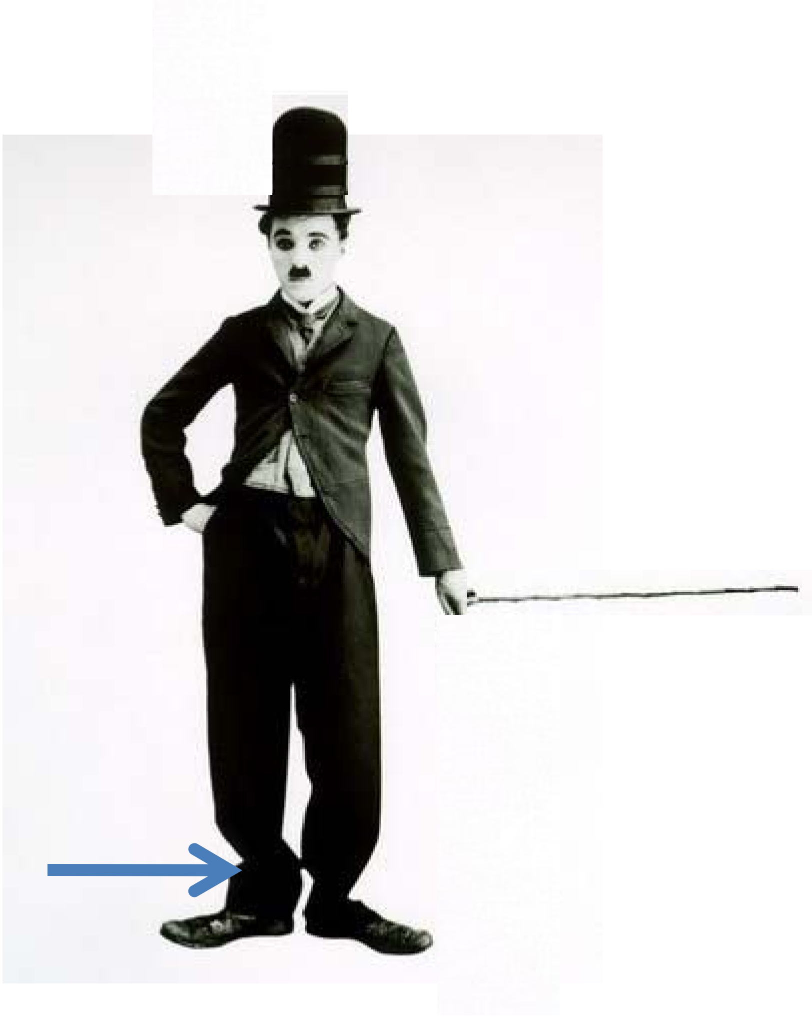

FROM RESPONSES TO REACTION TIMES

- Helmholtz measure of speed of nervous conduction (roughly 27m/s)

FROM RESPONSES TO REACTION TIMES

- Helmholtz measure of speed of nervous conduction (roughly 27m/s)

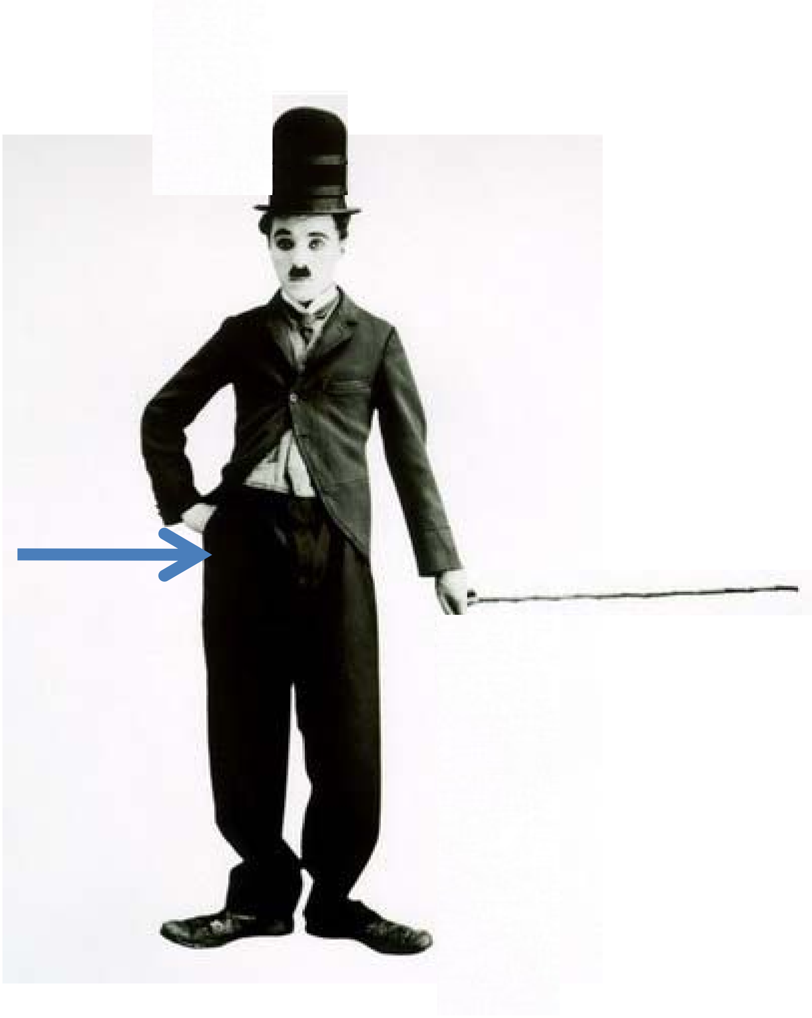

FROM RESPONSES TO REACTION TIMES

- Helmholtz measure of speed of nervous conduction (roughly 27m/s)

FROM RESPONSES TO REACTION TIMES

- Helmholtz measure of speed of nervous conduction (roughly 27m/s)

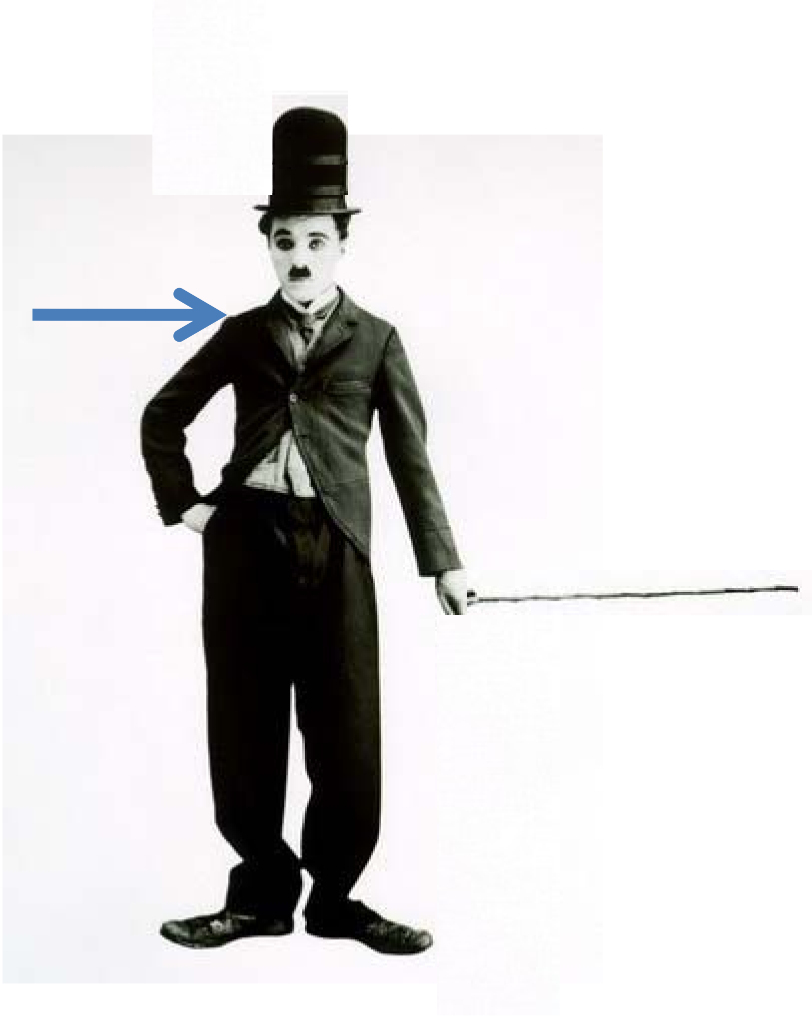

BUT OFTEN IT IS VALUABLE TO MEASURE INCIDENTAL VARIABLES

– e.g., eye-tracking reveals incremental interpretation in reading (e.g., Traxler, Pickering et al) – slow down on the word that doesn't fit...

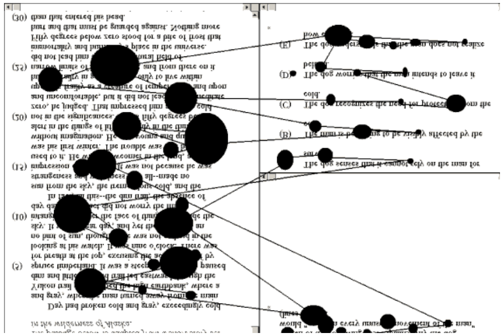

AND WE MAY MEASURE THE BRAIN, NOT JUST THE EYE...

- Event-related potentials (ERP) have characteristic patterns for syntactic and semantic anomaly (e.g., Kutas)

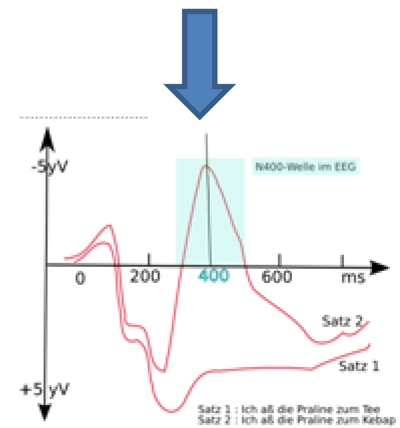

AND WE MAY WONDER *WHERE* IN THE BRAIN (PET, FMRI)

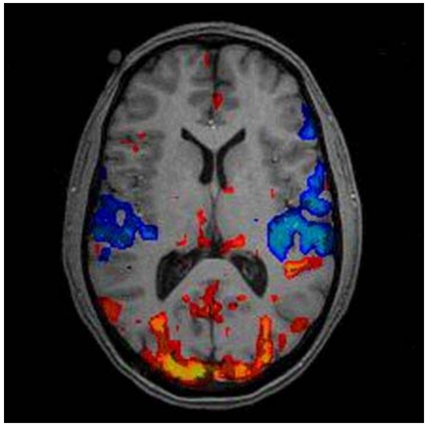

Specific regions appear associated with fear, surprise, different locations in the visual field, motor planning, social emotions, and so on

So a rich new set of dependent variables... though always a danger of their interpretation being circular

Imaging technologies provide another role for machine learning

- What can we really infer?
- Potentially horrible problems of multiple comparisons

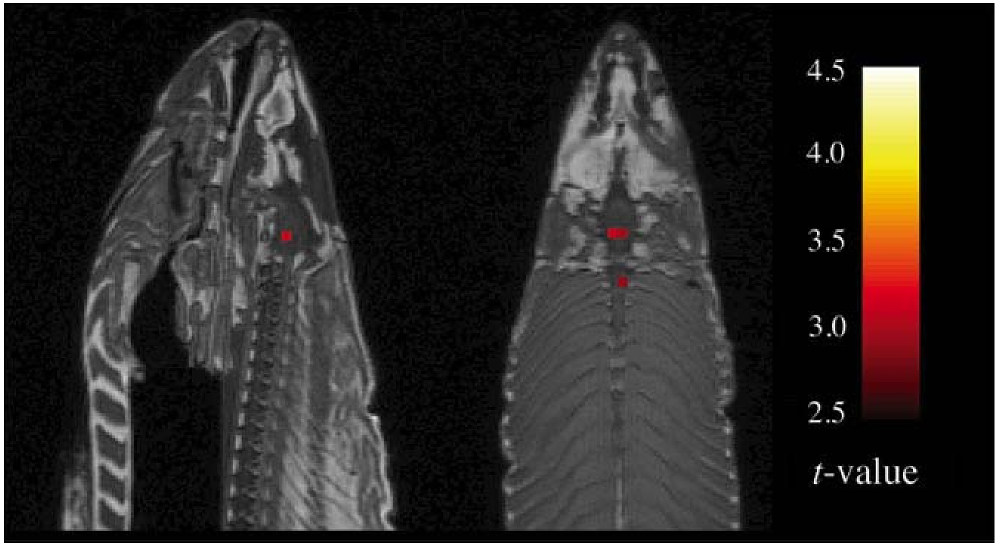

Cause celebre: fMRI of social emotions in a dead salmon (Craig Bennett, 2009)

MACHINE LEARNING CAN HELP

- But the results may be rather demoralizing...???

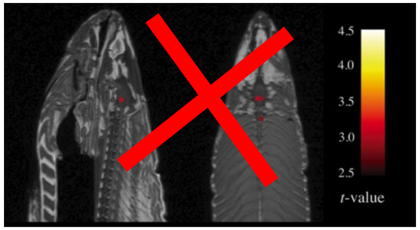

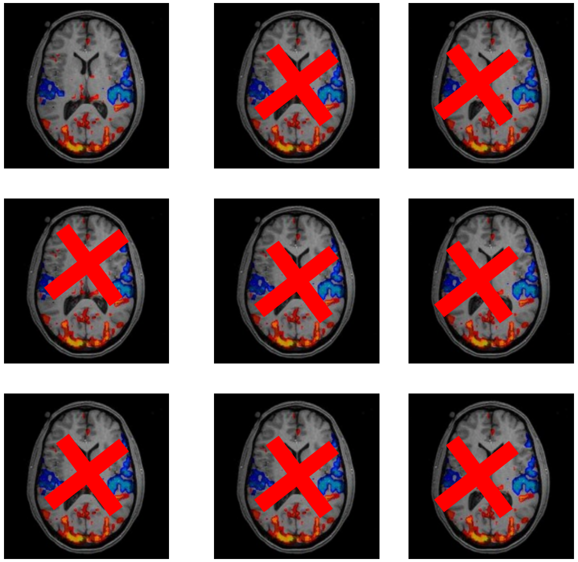

M Liou, H-R Su, J-D Lee, A Savostyanov, JAD Aston, C-H Chuang and PE Chen. Beyond p-values: Averaged and Reproducible evidence in fMRI experiments. (2009) *Psychophysiology* , 46: 367-378.

CAN YOU PLAY 20 QUESTIONS WITH NATURE AND WIN?

Alan Newell 1927-1992

EXPERIMENTAL PSYCHOLOGY AND THE SEARCH FOR BINARY DIVISIONS

- Experiments can at best answer questions posed to choose between broad classes of option
  - Face processing system vs general visual perception?
  - Semantic vs episodic memory
  - Short vs long term memory
  - Hot vs cold cognitive processes
  - ...
  - ...
  - ...

BUT EXPERIMENTAL PSYCHOLOGY CAN NEVER SUCCEED ALONE

- Newell (1973) argued that most such debates are inconclusive
- And even if conclusive 20 bits of information is not enough to specify a theory of the mind
- Observations/intuitions are vastly richer, and provide very rich constraints (as we have seen)
- And perhaps the biggest constraint of all is functional: good reverse engineering must be good engineering!

BUT WE NEED OUR 20 QUESTIONS NONETHELESS

- ...unless there is one solution to the problem of intelligence

- Cognitive science and machine learning need to ask the right questions, splitting along genuinely critical issues

  - Discriminative vs model-based perception???
  - Modular vs non-modular processes???
  - Exemplar vs rule-based process???
  - Symbolic vs connectionist cognition???

FINDING THESE QUESTIONS IS A CHALLENGE FOR THE REST OF THIS WORKSHOP, AND BEYOND!