Ontologies for Baby Animals and Robots.

Aaron Sloman

School of Computer Science, University of Birmingham

`http://www.cs.bham.ac.uk/~axs/`

These PDF slides are available in my 'talks' directory:
`http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#prague09`

NB: This is still a draft version which will be clarified, tidied up and extended when I have time (and better ideas!).

This is part of a large and growing collection of presentations on a battery of related topics about requirements and designs for intelligent systems, natural and artificial, available here:
`http://www.cs.bham.ac.uk/research/projects/cogaff/talks/`

Feel free to suggest improvements or send me criticisms: A.Sloman@cs.bham.ac.uk

High level aims and non-aims

My research is more science and philosophy than engineering: I am not trying to build a useful robot or a useful machine vision system, nor an ontology-based interface to the internet.

Rather, I am trying to understand what design requirements biological evolution had to address in producing types of animal that can perceive, interact with and manipulate a complex and changing 3-D environment, that includes large scale mostly static structures and smaller scale, more dynamic, structures and processes changing on different time scales, some under the control of the animal, some not, and some involving other information users.

Including design requirements met by these types of animal:
human, orangutan, hunting mammal, elephant, nest-building bird, octopus.

My aim is also to get some ideas about how those design problems were solved.

The main output of this research comprises: descriptions (mostly informal still) of both requirements and (still very sketchy) partial designs.

Making progress includes trying to test ideas about both requirements and designs by building working systems – though these are still very limited.

High level aims and non-aims (2)

My research is more science and philosophy than engineering

The requirements analysis and the designs for biological agents and human-like robots may not be relevant to very precisely specified practical vision systems and robots, with restricted and unchanging functionality, e.g. systems used in motor manufacturing – welding car bodies and other strongly constrained practical problems.

However the work is relevant to future ambitious robotic projects, e.g. designing general purpose domestic robots to help people who may be blind or have reduced mobility – subsuming most of the functions of a guide dog.

The work is also relevant to understanding human visual development and may be relevant

- to various clinical practical applications

- and to future educational policies.

But I shall say nothing about those topics here.

See the Cognition and Affect web site

    http://www.cs.bham.ac.uk/research/projects/cogaff/

my talks directory

    http://www.cs.bham.ac.uk/research/projects/cogaff/talks/

and the CoSy project papers

    http://www.cs.bham.ac.uk/research/projects/cosy/papers/

What are ontologies, and why are they important?

In order to acquire, manipulate, reason with, test, revise, store or use information about anything, it is necessary to have information components from which more complex information structures can be constructed.

What are ontologies, and why are they important?

In order to acquire, manipulate, reason with, test, revise, store or use information about anything, it is necessary to have information components from which more complex information structures can be constructed.

These information components may be about types of location, types of "stuff", types of motion, types of relationship, types of surface feature, types of extended object, types of interaction, or even types of mental event and mental process...

The ontology used by a perceiver, thinker, reasoner, active agent depends on which of these elements are represented in the ontology and how they can be combined to form more complex information structures representing more entities, processes, etc.

We can think of principles for composition of information fragments to form larger information structure as parts of the ontology or as parts of the mode of representation.

The most well known principles of composition are those in formal languages (logic, algebra, programming languages) and human natural languages.

However there are others, e.g. maps, diagrams of many kinds, pictures, sign languages, computer data-structures, neural nets, etc.

It seems that visual systems need forms of representation that combine some of the features of maps and pictures and some of the features of the more formal languages (e.g. supporting inference).

But current formalisms used in vision research tend to be too mathematically precise, and too lacking in practically useful information, to explain animal competences.
(As explained later.)

More on semantic composition

In order to acquire, manipulate, reason with, test, revise, store or use information about anything, it is necessary to have information components from which more complex information structures can be constructed.

We also need principles for composition of information fragments to form larger information structures as parts of the ontology or as parts of the mode of representation.

Example of a **linguistically composed** complex information structure involving several components of a spatial ontology: Two fairly flat roughly parallel surfaces facing each other are moving together with a cylindrical object between them oriented with its axis roughly parallel to the surfaces.

This is fairly abstract information, omitting many details, which could in principle be added, about the precise locations, orientations, colours, textures, rigidity, elasticity, temperature, etc.; yet it may suffice for certain tasks such as planning, predicting, explaining.

Something closer to a **pictorial** form of representation is required to relate particular visual episodes to the more abstract linguistic or logical representation, with information structures in the abstract representation partly in registration with the optic array.

The ontology available for constructing such percepts in a learning developing animal or robot could change over time – including the primitive components and the modes of composition.

E.g. it seems that the ontology of a very young child does not include 'boundary alignment', required for inserting an irregular shape into its recess, e.g. in jigsaw puzzles.

But the child's ontology does grow and most will include such things later.

Some requirements for an animal or baby ontology

- My concern is with animals or robots that need to acquire and use information about, reason about, and interact with rich and complex 3-D structures and processes in the physical environment.

  Artificial systems could include automated design, inspection and repair of complex machinery; automated rescue systems; domestic aids for disabled people; and robots performing tasks in remote and humanly uninhabitable environment, e.g. on space platforms and other planets.

- The ontology will not refer only to abstract structures, as an "internet ontology" system might.

- Instead the visual/spatial ontology would need to include spatial structures and processes, causal interactions, assembly or disassembly of objects of varying degrees and kinds of complexity

- including perceiving and interacting with process that involve changes of
  - material properties (e.g. becoming brittle),
  - spatial relations (including shape changes),
  - causal relations (e.g. producing obstructions, or loosening a grip)
  - functional relations (e.g. modifying a structure to serve a new purpose)

- Perceiving and thinking about other agents requires the ontology to have meta-semantic components,

  e.g. beliefs, goals, of others, or of oneself in the past or future, or in some hypothetical state.
  (This raises problems of "referential opacity".)
  Show video of pre-verbal child with ontology including meta-semantic ontology.

Some of what current systems cannot do

Familiarity with roles of low level pictorial cues in representing 3-D edges, orientation, curvature of surfaces, joins between two objects or surfaces, etc., allows you to use compositional semantics to see 3-D structure, and some causal and functional relationships, in pictures (even static, monocular pictures) never previously seen.

How many features, relationships (topological, semi-metrical, metrical, causal) can you see in these?
How many items (including substructures) can you identify in more than one of the scenes?
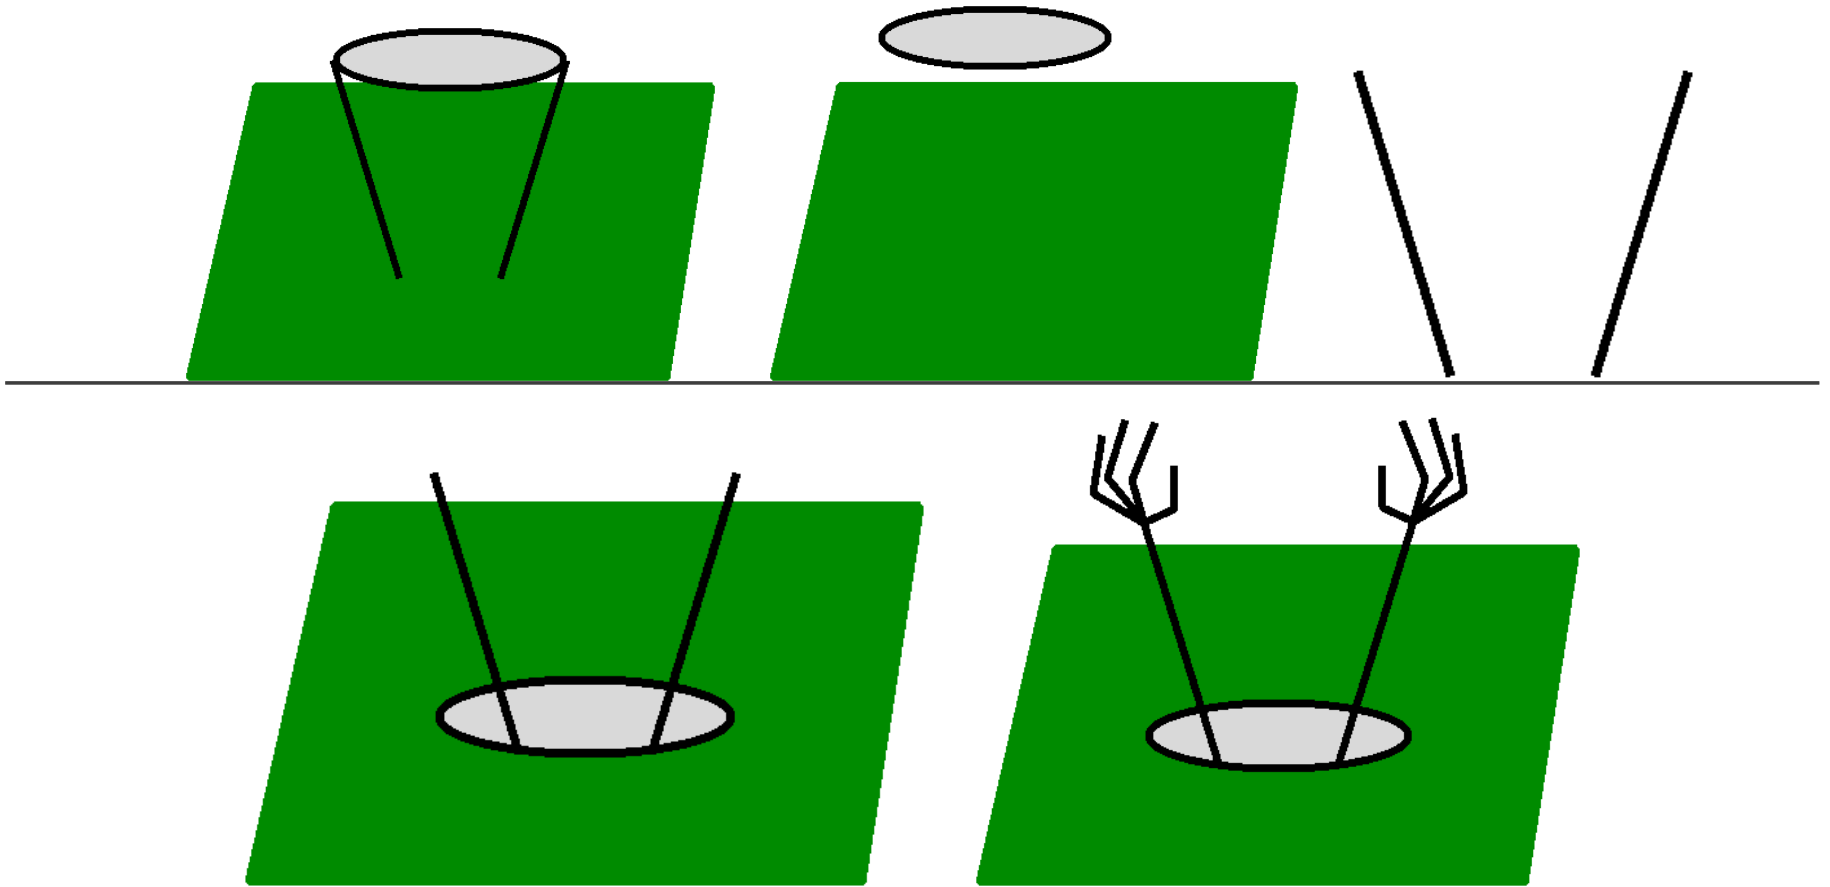

http://www.cs.bham.ac.uk/research/projects/cosy/photos/crane/
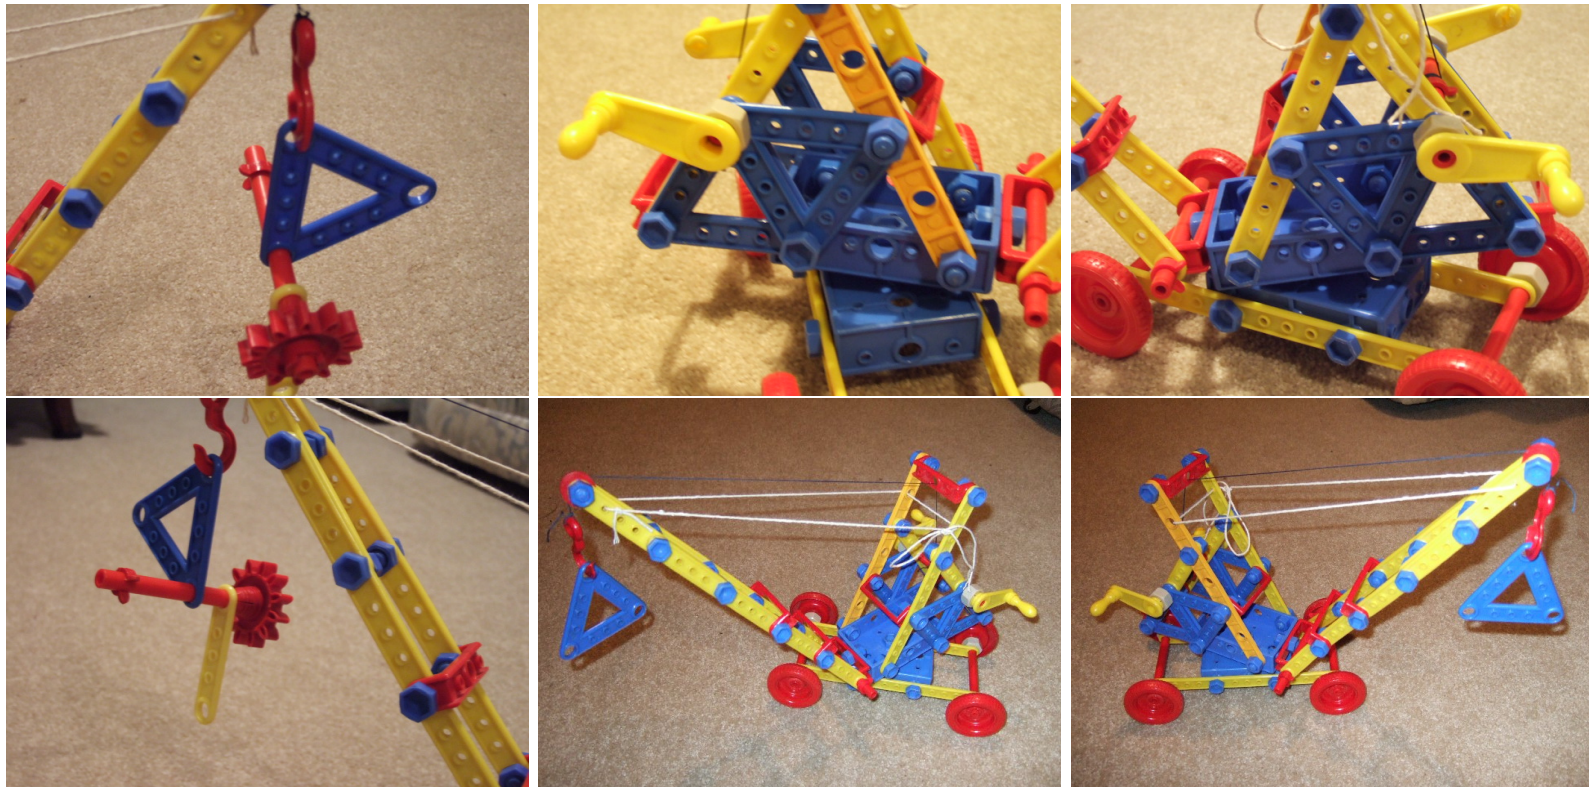
No AI vision system comes close to being able to do that – yet.

Different combinations of the same elements

What do you see in these pictures? Only 2-D configurations?

Notice how context can influence interpretation of parts.
Perceptual compositional semantics is highly context-sensitive.

Words can add more context: Strong worm catches early bird?
What about: Shark-infested sewer?

Seeing is not primarily recognising objects
Though recognising objects can help
# with the main tasks – e.g. disambiguation.

You can do many things with something you see but do not recognize:

- Climb over it, kick it, prod it, hit it, lick it, put it in your mouth, push it out of your way, ...

- pick it up in different ways,

- adjust your motion and your grip depending on where you decide to grasp it and with which hand,

- see, without having to try, that some grasps are impossible (grasping the back of an arm chair with one hand),

- seeing in advance how your hand will need to rotate when you grasp a particular object in a particular way

  and working out what will happen to the cup of coffee you are holding if you do that.

I can point at various locations on surfaces of these objects and if you see the shape (despite noise and low resolution) you will be able to work out roughly how two fingers need to be oriented to grasp at those locations.

These competences do not require **high precision** spatial information.
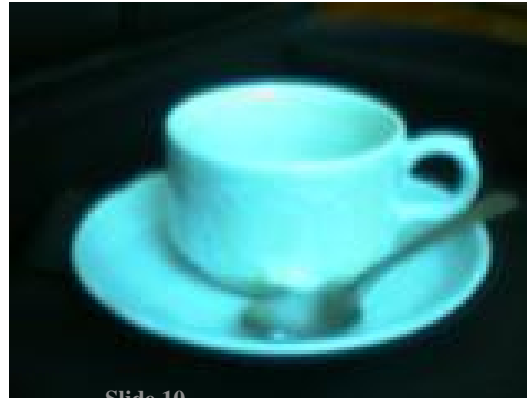
Ontologies for intelligent agents (natural & artificial)

Ontologies are for living, not just for reasoning and communicating.

- Active information processors need to be able to acquire, construct, and use novel information structures.

- This requires some (possibly growing) set of primitive information building blocks, e.g.:
  - Some way of encoding the primitives and ways in which they can be combined to deal with novelty.
  - Ways of deriving them from sensory input, or sensory-motor interactions.
  - Ways of using them to formulate goals, questions, specific beliefs, general theories, plans, actions and communications.
  - Currently we know very little about how humans, other intelligent mammals intelligent birds, intelligent cephalopods do such things.

- Researchers have focused on using forms of representation we already understand, rather than trying to develop requirements, and then possibly new forms.

  (Including new forms of composition and manipulation)

- We need to understand this if machines are to be able to communicate effectively with humans.

- There has been much misguided emphasis on embodiment as concerned with body morphology rather than with the nature of the environment. See:
  `http:`
  `//www.cs.bham.ac.uk/research/projects/cogaff/misc/embodiment-issues.html`

What sort of initial ontology?

Many theorists assume that the initial ontology includes only sensory and motor contents and patterns relating them, a somatic, multi-modal, ontology) – I claim that will not suffice for children, chimps, or crows.

CONJECTURE: From the start the learner will also use, and attempt to extend, an exosomatic, amodal ontology (about what's going on outside – not just the shadows on Plato's cave wall), including:

- bits of stuff (of various kinds) that can occur in the environment

- bits of surface of bits of stuff, in various shapes, locations, orientations

- bits of process (of various kinds) that can occur in the environment

- ways of combining them to construct larger structures and processes in the environment (not necessarily with global consistency)

- at various levels of abstraction: metrical, semi-metrical, topological, causal, functional....

  Semi-metrical representations include things like: "W is further from X than Y is from Z", orderings with gap descriptions, symmetries and partial symmetries. (And other things, still to be determined.)

  Semi-metrical distance and angle measures could include comparisons between distances and angles instead of use of global units, like 'cm' or 'degrees'.

  Kinds of curvature: spherical, cylindrical, conical, elliptical, wavy.... not implying mathematical precision.

  Instead of items in the environment being located relative to a single global coordinate frame, they could be embedded in (changing) networks of more or less local relations of the above types.

How should spatial properties/relations be represented?

The current default approaches among most vision researchers and roboticists have two serious flaws:

1. the use of a global coordinate system;

2. the use of precise metrics **globally** and **locally** (e.g. for curvature, orientation, shape, distance, area, volume, velocity, acceleration, etc.).

Systems that avoid problem (1) (e.g. by using geometric algebras?) still have problem (2).

The use of a high precision representation for information is a problem because:

- if a representation expresses high precision, and the sensory or other evidence available is noisy or of low precision, then no one representation is available, only a space of representations compatible with the sensory evidence – but we see one (roughly characterised) shape, not a probability distribution;

- researchers/designers try to deal with this by using probability distributions, but results are weak;

- there are many contexts where that is both unnecessary and counter-productive, e.g. reasoning about how some mechanism, such as an old clock, works: such machines are best treated as deterministic;

- in particular, more abstract **possibility distributions** provide a deeper and more general form of representation.

We need more abstract forms of representation that contain enough information to be useful for decision making, action control, etc. but not so much that they require unsupported precision.

Examples include ordering information and semi-metrical information, e.g. 'A is longer than B', 'the angle through which I have turned is not enough to ensure avoiding bumping into X', or 'my fingers are far enough apart to be able to straddle object Y'.

See: `http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0702` "Predicting affordance changes"

Representations for vision: some hints

Various people have suggested that visual systems should use models that can be combined in various configurations

E.g. Biederman proposes 'geons': various types of 3-D solids

   (cylinders, spheres, cones, blocks, etc.) that can be distorted, added, subtracted, merged, etc., to form a wide range of shapes with convexities and concavities.

      http://en.wikipedia.org/wiki/Geon_(psychology)

Suppose we replace geons with lower level surface fragments that can be composed to form a wide range of 3-D structures

   (e.g. bumps, dents, grooves, ridges, edges, hollows of various shapes, protuberances of various shapes, openings, rims, etc., – compare Barrow and Tennenbaum 1978 on "intrinsic images".)

and also add various process fragments that can be combined with various 3-D structures to form a wide range of 3-D processes

   (e.g. translations, rotations, relative movements, changes in alignments, alterations in curvature, changing gap sizes, of surface fragments, etc.)

that can be composed to form processes in which surfaces change their shapes and their relationships.

We could also add kinds of stuff with different properties

   (e.g. rigidity, impenetrability, plasticity, elasticity).

The hard part is to specify modes of representation supporting composition, using imprecise qualitative notions of contact, inclusion, overlap, alignment, ordering, etc.

Seeing possibilities

- J.J. Gibson criticised the notion that the main function of perception is to produce some sort of model or description of physical/geometrical external reality.

- He argued that organisms need to acquire information about affordances, not locations, categories, pose, etc. of objects.

- I suggest his ideas need to be generalised, by removing two of his three requirements for affordances:

    - Possibilities for and constraints on changes/movements that may not actually be occurring.
    - Relevance to the perceiver's actual or possible goals, preferences, etc.
    - Restricted to possible actions of the perceiver

- If we retain just the first of these we have proto-affordances: possible motions and constraints on motions, not necessarily relevant to the perceiver's goals or possible actions.

- Required for thinking about possible motions or constraints on motions of inanimate objects and other agents (trees in the wind, rocks falling, rivers flowing, waves breaking).

- A special case is perception of vicarious affordances
    (e.g. for prey, predators, conspecifics).

Gibson took some important steps along a new road.

But it is a much longer road than he realised.

# Our visual ontology isn't used just to build complete, consistent, metrically precise models of scenes
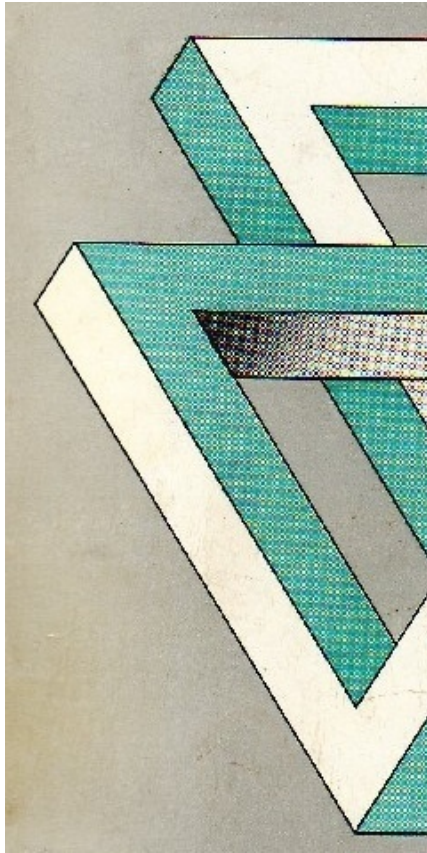
What can you say about 3-D distances and differences of orientation between the two short edge-fragments bounding the bottom right shaded part and the topmost visible edge in the scene?
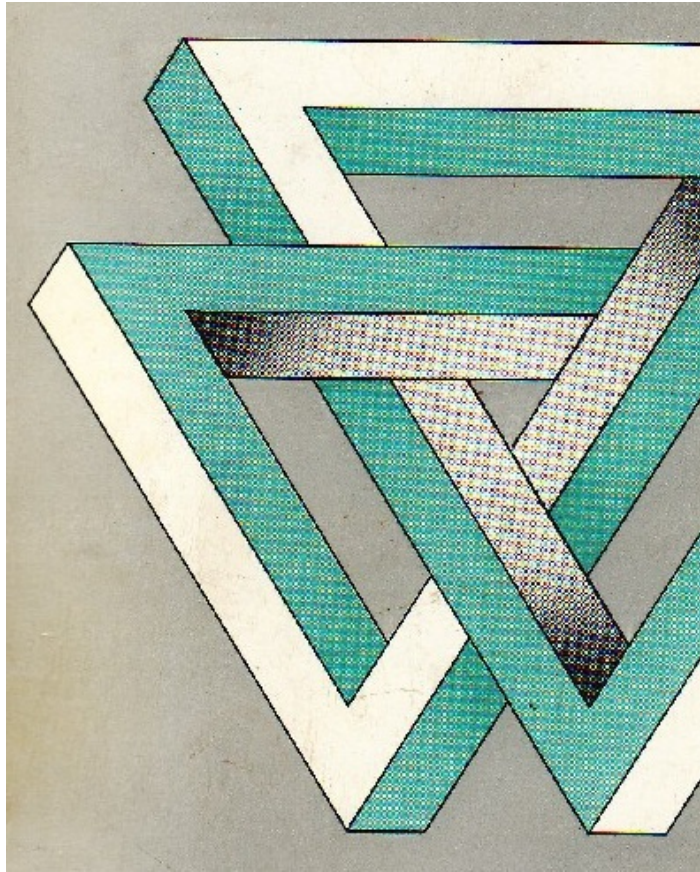
Does what you see now alter your interpretation of the previously visible edges?

All of it
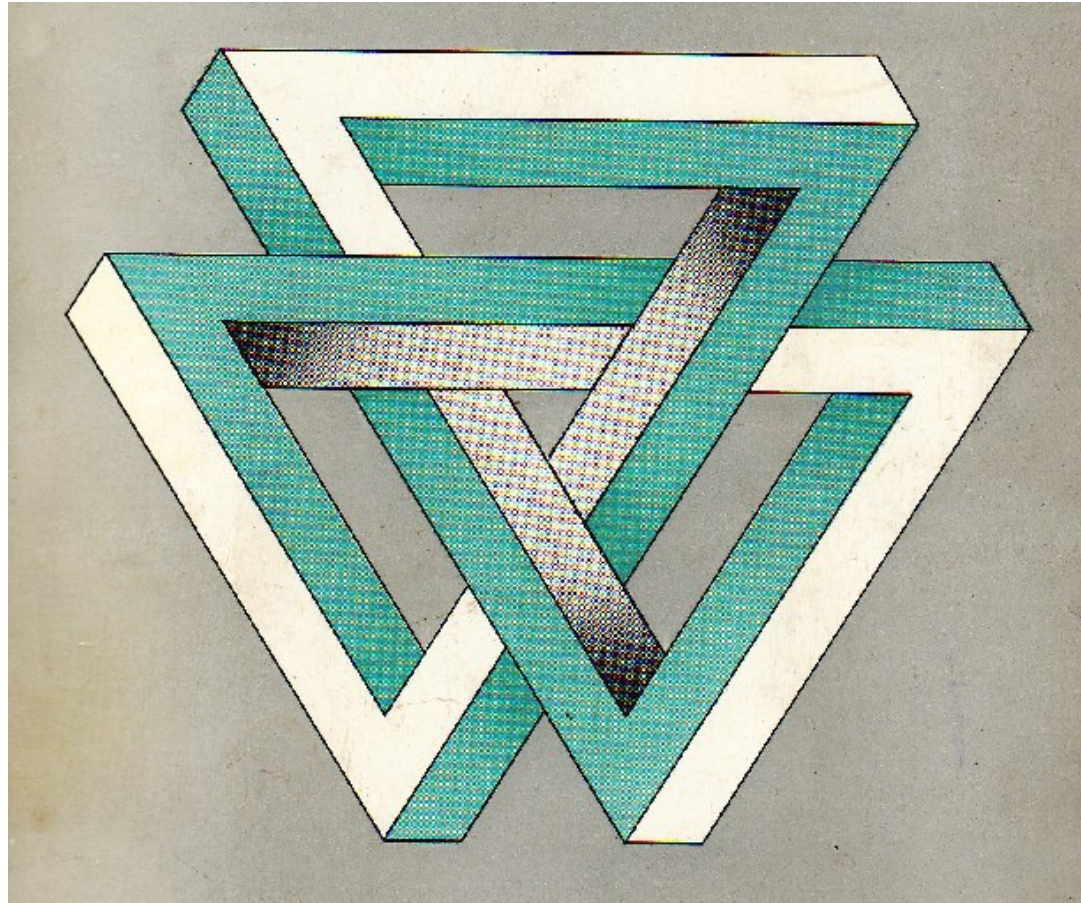
Does what you see now alter your interpretation of the previously visible edges?

Such a picture (the whole picture, but not previously shown parts) represents an impossible object whose impossibility could be easily missed
(and would be missed by a young child).

A precursor to Penrose
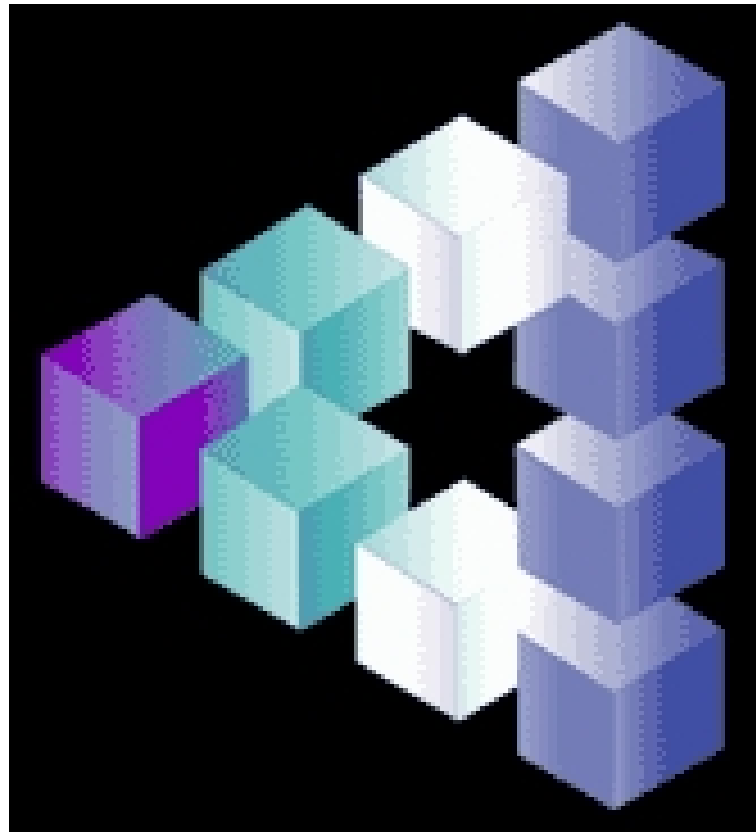
Picture by Swedish artist, Oscar Reutersvärd (1934)

The whole configuration is impossible, yet removing any of the blocks leaves a possible configuration.

Moreover, even as it is you can see possibilities for 3-D processes, e.g. moving your hand between two adjacent blocks, moving a block away fom the configuration, replacing it with another block, etc.

What you see supports a rich collection of clearly identifiable possible processes, even though the total configuration is impossible.

A 3-D apparently impossible object

The lower part looks impossible, but only when viewed from a certain direction.

Otherwise it comes apart and is no longer impossible.

What do such pictures of impossible objects show?

That we don't build only globally consistent models and we don't always check for global consistency in our percepts – we don't need to normally because nothing in the environment can be globally inconsistent, even if it looks globally inconsistent.

The mechanisms used in creating such percepts solve multiple constraint problems and sometimes these cause local information to be wrongly interpreted, as in the Ames room, where the preference for a rectangular shape for the room leads to serious perceptual distortions.

Likewise the hollow mask illusion. (Richard Gregory)
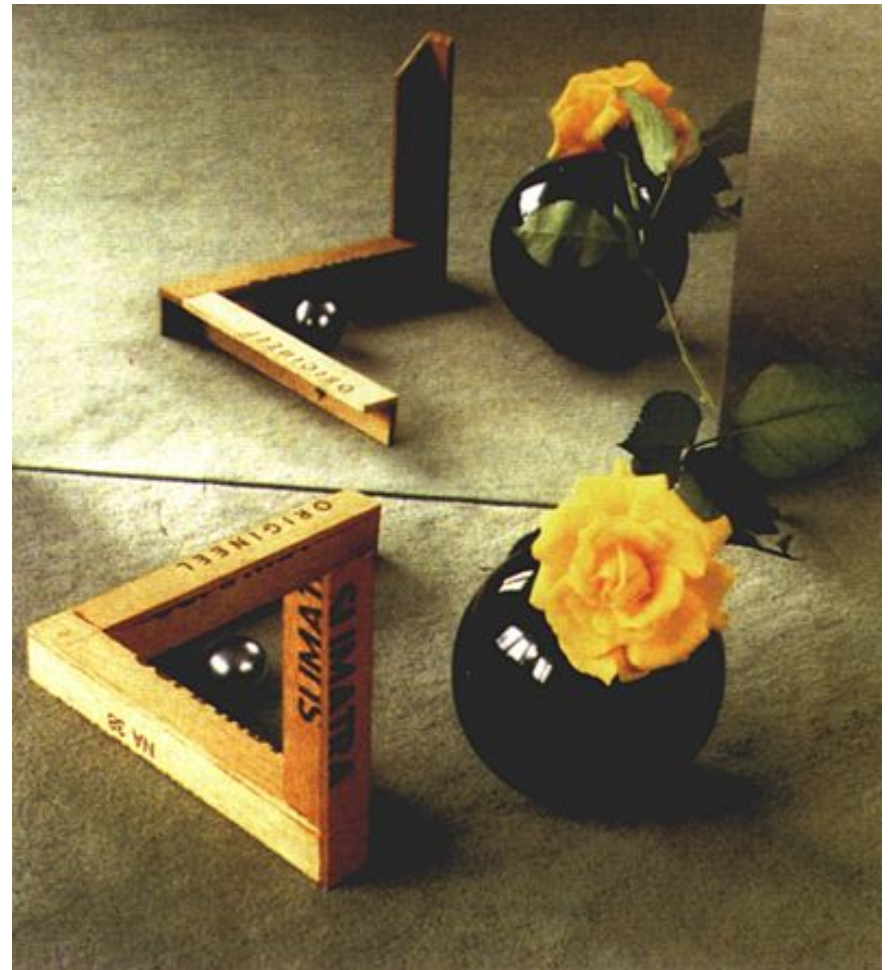
Picture by Bruno Ernst, after Richard Gregory.

How can perceptual systems cope with such ontologies?

Powerful multi-layer, extendable constraint-propagation mechanisms will need to be available for vision, haptic perception, reasoning, planning, predicting, etc. to work.

For more on this see, for example
`http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0804`
COSY-TR-0804 (PDF) Some Requirements for Human-like Robots: Why the recent over-emphasis on embodiment has held up progress. (To appear in a book based on a 2007 Honda research conference.)

The unsolved problem is: what forms of representation are required to support these processes?

See `http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#glang`
Evolution of minds and languages. What evolved first and develops first in children:
Languages for communicating, or languages for thinking (Generalised Languages: GLs)?

I have argued that both in our pre-linguistic evolutionary history, in our pre-verbal individual development and in some other non-verbal animals, there are "languages" that are not used for communication, but are used **internally** for perception, reasoning, goal formation, planning, plan execution, question formation, prediction, explanation, causal understanding....

and those languages include

structural variability (for dealing with novelty)

compositional semantics (context sensitive)

manipulability (for reasoning, planning, hypothesising, etc.)
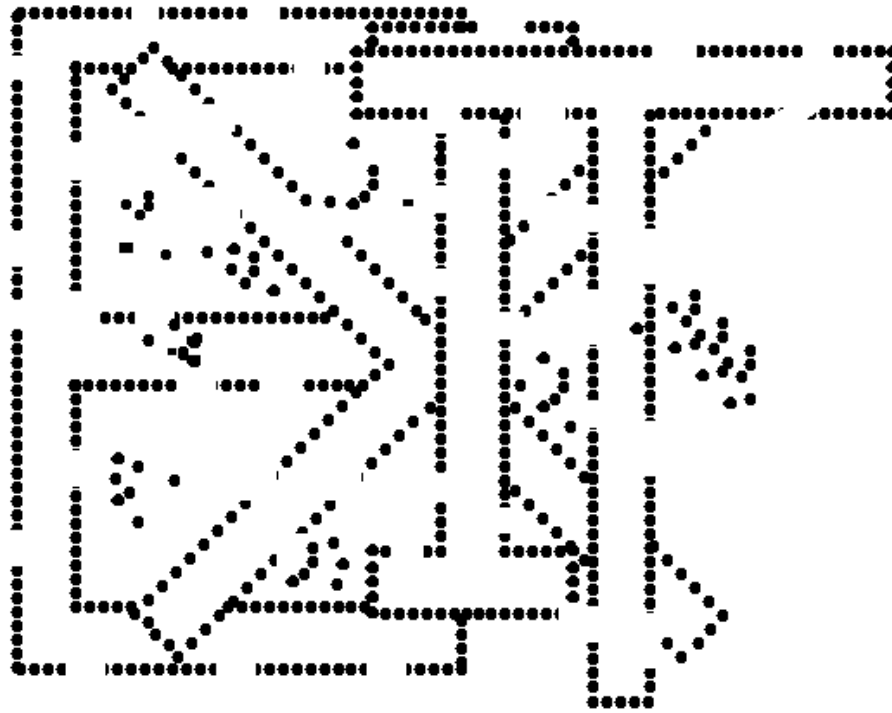
Interpreting multi-level noisy images

The next three slides come from Chapter 9 of *The Computer Revolution in Philosophy* (1978)

`http://www.cs.bham.ac.uk/research/projects/cogaff/crp/chap9.html`

The first picture is a noisy and cluttered image of a word made of outline capital letters.

Try looking at it for less than a second and see if you can tell what the word is.

Noisy and cluttered image of a word composed of outline capital letters.

The same cluttered image, but with noise removed.

The next slide shows how different levels of abstraction, using different ontologies are required to do the interpretation.

# An artificial example to illustrate some problems

This shows how several layers of interpretation may be involved in seeing letters in a dot-picture.

Each layer is a domain of possible configurations in which substructures may represent or be represented by features or substructures in other layers.
The following domains are illustrated:

    (a) configurations of dots, spaces, dotstrips, etc.,

    (b) configurations of 2-D line-segments, gaps, junctions, etc.,

    (c) configurations of possibly overlapping laminas (plates) in a 2.5D domain containing bars, bar-junctions, overlaps, edges of bars, ends of bars, etc.,

    (d) a domain of stroke configurations where substructures can represent letters in a particular type of font,

    (e) a domain of letter sequences,

    (f) a domain of words composed of letter sequences.

These can be processed in parallel using top-down, bottom-up and middle-out processing, concurrently, with much constraint propagation in all directions.
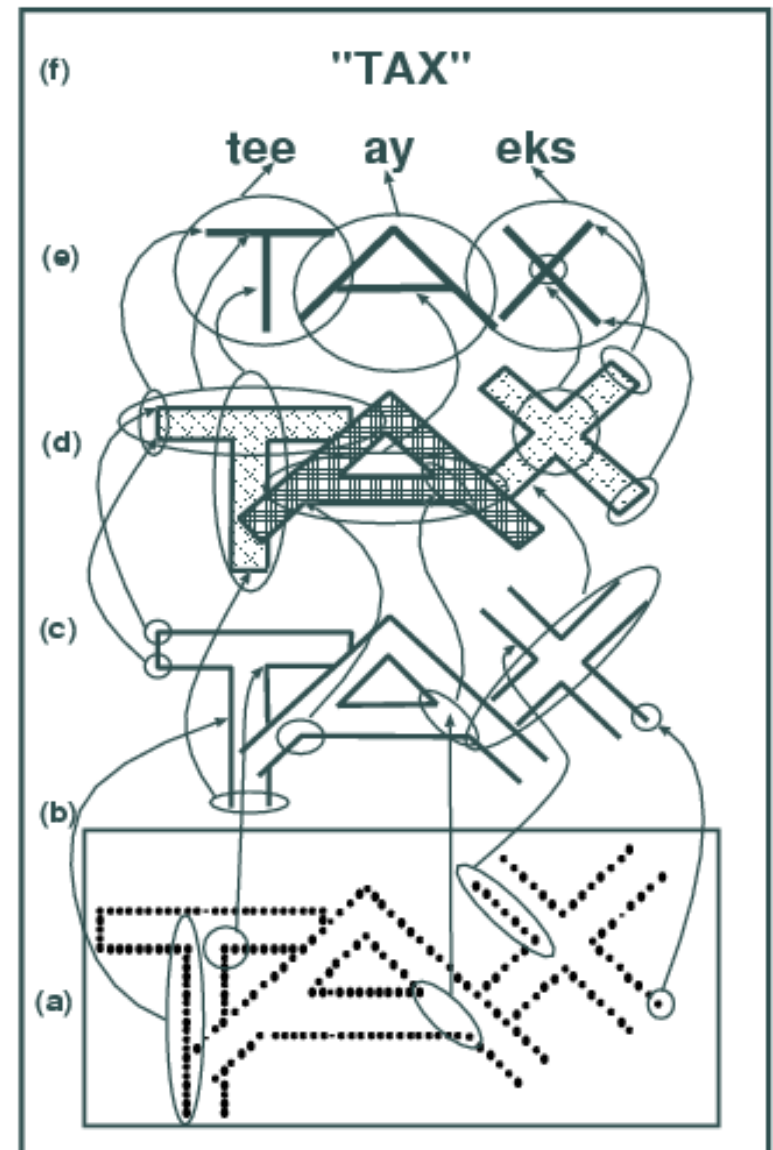
A familiar type of dynamical system
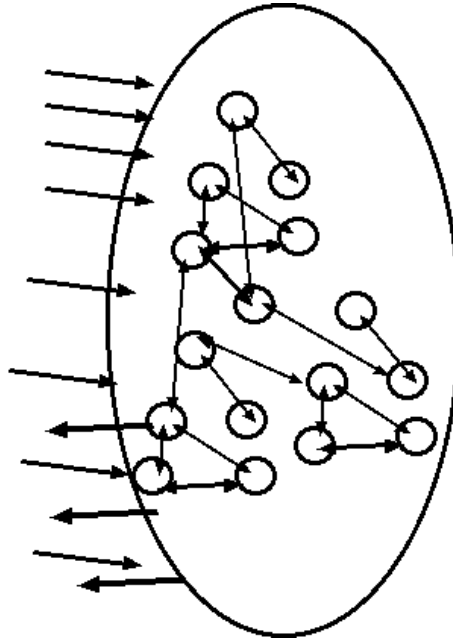
A multi-stable dynamical system closely coupled with the environment through sensors and effectors.

All the semantics may be somatic referring only to states of sensors and effectors and their relations.

Perhaps new sorts of dynamical system
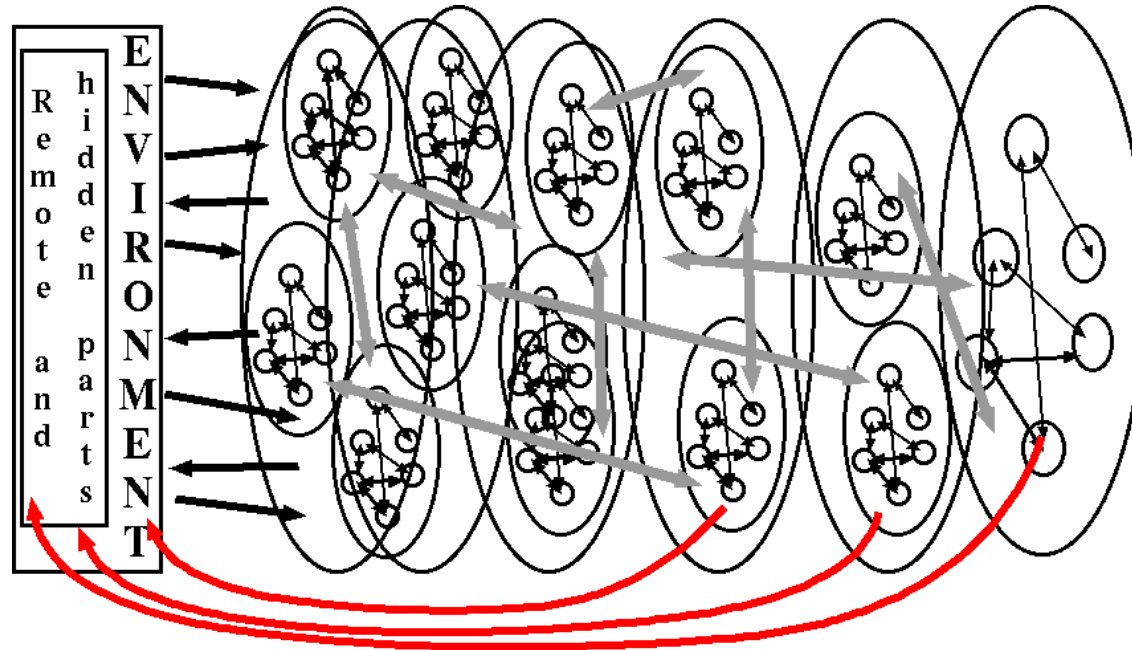
A multi-layered, multi-component dynamical system with most sub-systems not coupled to the environment, but some able to refer to things in the environment that cannot be sensed, e.g. past, future, remote, and hypothetical entities. This uses an exo-somatic ontology.

But there are many questions about how the external environment is represented.

E.g. a unified global, fully metric coordinate system?

Compare large numbers of topological and semi-metrical information structures, some of them dynamically changing.

The huge gap to be bridged

Some AI vision and robotic systems perform impressively in very restricted tasks, requiring little understanding of what they are doing or why. BUT ...

- Although there are mobile robots that are impressive as engineering products,

  e.g. BigDog – the Boston dynamics robot

  `http://www.bostondynamics.com/content/sec.php?section=BigDog`

  and some other mobile robots that are able to keep moving in fairly rough terrain, including in some cases moving up stairs or over very irregular obstacles.

- They lack understanding of what they are doing, what they have done, what they could have done, what goals they could achieve in different circumstances, why some goals should be abandoned, etc. though they can sometimes react *as if they* understood.

  e.g. sticking out a leg to prevent a fall sideways: a trained or programmed reflex.

- Existing robots that manipulate objects can be triggered to perform an action, but cannot perceive processes, notice new possibilities, or reason about what the result would be if something were to happen, except in very simple cases.

- Neither can they reason about why something is not possible.

- I.e. they lack the abilities underlying the perception of positive and negative affordances.

- They cannot wonder why an action failed, what would have happened if...; notice that their action might have failed if so and so had occurred part way through, etc.; realise that some information was available that they did not notice at the time.

Moreover, what they can see, represent or think about, in the environment, is too limited.

A newborn human infant cannot see or do those things
Why not?
And what has to change to produce those competences?

We must not forget that some newborns can do very sophisticated things very soon after birth (e.g. deer, chicks) so evolution **can** produce innate sophisticated competences.

If infant humans, orangutans, corvids, ... lack behavioural competences some other species have, perhaps that is because they have something more powerful.

Everyone assumes learning is that more powerful something: but what sort of learning? and from what starting point?

Often it is assumed that the learning is of a general kind, that can learn anything, provided that enough training data can be provided.

The designers of such systems don't bother to study the environment: they expect to leave that to their future learning systems – but that will not work.

In "The Well-Designed Child" (AIJ, Dec, 2008) John McCarthy wrote:

"Evolution solved a different problem than that of starting a baby with no a priori assumptions.
....
Instead of building babies as Cartesian philosophers taking nothing but their sensations for granted, evolution produced babies with innate prejudices that correspond to facts about the world and babies' positions in it. Learning starts from these prejudices. What is the world like, and what are these instinctive prejudices?"

Turing's mistake?

A major challenge for such an investigation is
- to understand the variety of possible starting points
- for an individual born or hatched in a particular sort of environment,
- after millions of years of evolution of the species

In his 1950 *Mind* article, "Computing machinery and intelligence", Turing wrote:

"Instead of trying to produce a programme to simulate the adult mind, why not rather try to produce one which simulates the child's? If this were then subjected to an appropriate course of education one would obtain the adult brain. Presumably the child brain is something like a notebook as one buys it from the stationer's. Rather little mechanism, and lots of blank sheets. (Mechanism and writing are from our point of view almost synonymous.) Our hope is that there is so little mechanism in the child brain that something like it can be easily programmed."

On this point (little mechanism and much space), Turing was uncharacteristically badly wrong, like all the AI researchers who try to find a small number (some hope **one** will suffice) of powerful, general, learning mechanisms that can learn from arbitrary data.

Evolution did not produce general-purpose data-miners. (see McCarthy – next slide)
- Most species produced by evolution start off with almost all the information they will ever need, leaving only scope for minor adjustments of parameters, e.g. for calibration and minor adaptations.

- A few species learn a lot using mechanisms that evolved to learn in a 3-D world of static and changing configurations of objects, including other intelligent agents:
    they start with powerful special-purpose mechanisms.

Evolution is a general-purpose data-miner, changing what it mines
But it needs something like a planet-sized laboratory, and millions of years, to produce things like us.

McCarthy does not agree with Turing

"Animal behavior, including human intelligence, evolved to survive and succeed in this complex, partially observable and very slightly controllable world. The main features of this world have existed for several billion years and should not have to be learned anew by each person or animal."

> http://www-formal.stanford.edu/jmc/child.html
> To be published in the AI Journal (December 2008)

McCarthy grasped an important point missed by Turing (and by many AI researchers).

McCarthy's own theories about requirements for a neonate are tempered by his goal of attempting to see how much could be achieved using logic.

We need to keep an open mind as to which forms of representation and modes of syntactic composition and transformation may be required, or may be useful at times.

> As argued in 1971 in: Interactions between philosophy and AI: The role of intuition and non-logical reasoning in intelligence.
>
> http://www.cs.bham.ac.uk/research/cogaff/04.html#200407
>
> Also Chapter 7 of The Computer Revolution in Philosophy (1978)
> http://www.cs.bham.ac.uk/research/projects/cogaff/crp/chap7.html

I am not arguing **against** the use of logic, but **for** a search for additional (new) forms of representation.

Evolution produced something richer

A logicist roboticist might think innate prejudices can be expressed as axioms and deployed through a logic engine.

However, studying the environment animals interact with, and learn in, suggests that we need a much richer theory, involving what McCarthy describes, and also

- An initial architecture, that can extend itself in certain ways including ontology extension.

- Initial (still unknown) forms or representation adequate for encoding specific sorts of information, and which support specific forms of information manipulation.

- Initial sensory, motor, and internal processing mechanisms

    including mechanisms for constructing new goals, for goal conflict resolution, and for detecting opportunities to learn.

- Initial behavioural dispositions that drive learning tailored to perceiving and producing 3-D structures and processes.

- An initial, mostly implicit, "framework theory" determining the type of ontology that is assumed and ways in which it can be used and extended. Compare Immanuel Kant (1780).

    E.g. implicit assumptions about the topology of space/time, kinds of stuff able to occupy and move around in space, modes of composition of structures and processes, kinds of process that can occur involving the stuff, kinds of causation, the differences between doing and passive sensing, ...

- **Delayed activation of an architectural layer that uses the combination of the environment and the early architecture as a new developmental "playground" in order to drive ever more sophisticated testing, debugging, and extensions.**

In order to make progress

- Instead of the normal AI strategy of thinking about how to extend our existing mechanisms, or how to deploy them in new ways,

- We need to engage in a deep study of features of the environment and ways of interacting with it

- Looking at examples of children and other animals doing that, and altering their competences as a result.

- Trying to derive constraints on the forms of representation and ontologies that can explain the detailed phenomena observed at different stages of development (which in children are partially, not totally ordered).

- In the light of all that, trying to design and test mechanism, architectures, robots that illustrate the theories.

NB: the problems will be different for different sorts of organisms and robots, e.g. depending on the complexity of their sensors and manipulators, the kinds of terrain they inhabit and the kinds of things they need to acquire and avoid.

See: http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0704
Diversity of Developmental Trajectories in Natural and Artificial Intelligence,
in Computational Approaches to Representation Change during Learning and Development. AAAI Fall Symposium 2007,
Technical Report FS-07-03, pp. 70–79,

Composition/Binding

These different aspects of reality can be composed/combined in many different ways.

Long before there was algebraic/functional/logical composition there was spatio-temporal composition.

Also auditory/temporal composition – music and many natural sounds.

We need to distinguish

- Composition in the spatio-temporal environment

    e.g. combining actions and things acted on, or sounds

- Composition in internal representations of things that can be spatio-temporally combined: i.e. composition in representations in virtual machines.

At present we have only a relatively small number of forms of information-composition that we can implement and use in computers.

By studying the environments of various sorts of intelligent systems very carefully we can derive new requirements for forms of representation and forms of composition and manipulation.

This may lead to the creation of new kinds of artificial information-processing systems.

Composition/Binding and Creativity

Have you ever spilt a bowl of porridge on a thick carpet?

How should the result be cleaned up?

Composition/Binding and Creativity

Have you ever spilt a bowl of porridge on a thick carpet?

How should the result be cleaned up?

- Using a vacuum cleaner?

- Using a broom?

- Using a dustpan and brush?

- Using a cloth – wet? or dry?

- Using a mop and a bucket of water? With/without detergent?

- Take the carpet out and shake it?

- Use a shovel, or a trowel?

- Find an animal that likes eating porridge and ....

  .....Other options?

If you have never encountered the problem how can you think about the solution?

Should an intelligent internet, or a semantically sophisticated domestic helper be able to give advice on such problems?

Further examples are here

`http://www.cs.bham.ac.uk/research/projects/cogaff/sloman-oii-2009.pdf`
Requirements for Digital Companions: It's harder than you think.

Life is information processing – of many kinds

The world contains: matter, energy, information

Organisms acquire and use information,
  in order to control how they use matter and energy
      (in order to acquire more matter, energy and information,
      and also reproduce, repair, defend against intruders, dispose of waste products...).

Somehow evolution produced more and more sophisticated information processors.
These pose challenges for science and engineering, namely:

- To understand that process.
- To understand the products.
- To replicate various aspects of the products.

We need to understand
- the structure of design space
- the structure of niche space
- the many design tradeoffs linking them
- the possible trajectories in design space,
- the possible trajectories in niche space,
- the many complex feedback loops linking both.

Development of environment and cognition 1

The cognitive system, including sensory mechanisms, motor control systems, learning systems, motivational mechanisms, memory, forms of representation, forms of reasoning, etc. that an organism (or robot) needs will depend both on

- **what is in the environment**
  and
- **what the physical structure and capabilities of the organism are.**

For a micro-organism swimming in an ever changing chemical soup it may suffice to have hill-climbing mechanisms that sense and follow chemical gradients, perhaps choosing different chemical gradients according to the current needs of the organism.

As the environment becomes more structured, more differentiated with more enduring objects and features (e.g. obstacles, food sources, dangers, shelters, manipulable entities) and the organisms become more articulated, with more complex changing needs, the information-processing requirements become increasingly more demanding.

As more complex information processing capabilities develop, the opportunities to observe, modify and combine them in new ways also develop.

See: Diversity of Developmental Trajectories in Natural and Artificial Intelligence
AAAI07 Fall Symposium
`http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0704`

Development of environment and cognition 2

The cognitive system, including sensory mechanisms, motor control systems, learning systems, motivational mechanisms, memory, forms of representation, forms of reasoning, etc. that an organism (or robot) needs will depend both on

- what is in the environment
  and

- what the physical structure and capabilities of the organism are.

Many researchers who emphasise the importance of embodiment of animals and robots make a mistaken assumption:

they claim that embodiment and physical morphology solve the problems and reduce the burdens on cognition, by producing required results "for free" when movements occur.

However, the point I am making is that

As bodies become more complex, with more parts that can be moved independently to cooperate with one another in performing complex actions on complex, changeable structures in the environment, the cognitive demands (for perception, learning, planning, reasoning, and motor control, and the ontologies involved) increase substantially, requiring more powerful forms of representation and more complex information-processing architectures.

For more on this see `http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0804`

A. Sloman, "Some Requirements for Human-like Robots: Why the recent over-emphasis on embodiment has held up progress". in *Creating Brain-like Intelligence*
Eds. B. Sendhoff and E. Koerner and O. Sporns and H. Ritter and K. Doya, 2009, Springer-Verlag

Developmental psychologists vs Designers

Many developmental psychologists investigate what is and is not innate in newborn humans, and other animals.

Examples studying humans include (among many more):
E. Spelke, P. Rochat, E. Gibson & D. Pick, A. Karmiloff-Smith, and much earlier J. Piaget,

and studying animals:
N. Tinbergen, K. Lorenz, J. Goodall, W. Köhler, E.C. Tolman, I. Pepperberg, M. Hauser, A. Kacelnik (and colleagues), N. Clayton, S. Healey, F. Warneken, M. Tomasello,

Unfortunately not enough of these researchers have learnt to look at something done by a child, chimp, or chick and ask

How could **that** work? What else can the mechanisms do? **How** do they do it?

Instead most of them ask questions like

- Under what conditions does this happen?
- How can the task be made easier or more difficult for species X?
- Is this innate or learnt?
- If it is learnt what triggers the learning?
- Which other animals can and cannot do it?
- How early does it happen?
- Which additional tests can I perform to detect these and similar competences?

They don't adopt what McCarthy calls "the designer stance".

TO BE REVISED AND EXTENDED