First year lecture: Intro to AI, Birmingham 12 Dec 2007
(Based on Seminar in School of Psychology in October 2007)
Mind as Machine Weekend Course, Oxford 1-2 Nov 2008

What evolved first:

Languages for communicating?
or
Languages for thinking ?[*]

(Generalised Languages: GLs),

Aaron Sloman
http://www.cs.bham.ac.uk/~axs/

Ideas developed with **Jackie Chappell** (Biosciences, Birmingham).

These slides are accessible here:

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

[*] Where 'thinking' here loosely refers to any kind of internal information processing.

The Context of this Lecture

One use of these slides is as a contribution to the first year module:

Introduction to Artificial Intelligence

`http://www.cs.bham.ac.uk/~jab/Modules/IntroAI/07-08/`

In previous years I gave a lecture on AI and Philosophy, still available here

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

Further background information can be found in this high-level overview of the goals and methods of AI (AI as science and AI as engineering):

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

A comparison of the aims, methods and tools of AI and more conventional software engineering languages and development environments can be found here:

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

I am assuming that students will recall some of the previous lectures on AI, and will know that AI includes the study of, among other things:

- vision and other forms of perception
- planning and decision making
- learning of various kinds
- reasoning and problem solving of various kinds
- natural language processing
- emotions and other affective states and processes
- architectures, representations and algorithms for intelligence
  symbolic, logic-based, rules-based, neural, diagrammatic, dynamical systems, ...

(skip) These Slides (work in progress)

These slides were originally written for presentation at the Language and Cognition seminar, School of Psychology, University of Birmingham, 19th October 2007

The topic is far too large for a single seminar, so only a subset of the slides were presented, after some videos of animals and prelinguistic children doing things that demonstrated perception of structure in the environment, and in two cases interpretation of the intentions of an adult human, and apparently spontaneous actions to help achieve the adult's goals.

> Thanks to Felix Warneken for use of his videos, available here
> `http://email.eva.mpg.de/~warneken/video.htm`

Some of the videos I use in this context are available here (e.g. broom video):

> `http://www.cs.bham.ac.uk/research/projects/cosy/conferences/mofm-paris-07/sloman`

Some of the papers referred to in the talk are listed at the end, with URLs.

I thank the members of the seminar for their patience and for interesting comments and questions which have led to some changes to these slides.

NOTE:
My slides are written so as to be readable online,
so they contain a lot more detail than most presentations.

This makes them less suitable for live presentations.

Some references are provided at the end.

Some background assumptions

Things we have learnt from AI research over the last 60 years or so include the following:

- An animal or machine acting in the world needs an information-processing architecture including different components capable of performing different sorts of tasks concurrently
- This is a virtual-machine architecture, not a physical architecture
- The various components need information of different sorts, which has to be represented or encoded in an appropriate way
- There is no one right form of representation: different tasks have different requirements, e.g.
  - for collections of symbolic facts (particular and general)
  - for structures representing spatial relationships e.g. maps, 3-D models
  - for algorithms that can be executed to produce internal or external behaviours
  - for doing statistical processing (e.g. building and using histograms)
  - for doing fast pattern recognition (using statistical or symbolic mechanisms)
  - for representing control information, including goals, preferences, partially executed plans, future intentions, etc.
- SOME information processing requirements BUT NOT ALL can be very dependent on the contents of the environment and on the body of the robot or animal (its structure, the materials used, etc.).
- Some animal information-processing architectures are mostly genetically determined, allowing only minor adaptations, whereas others are grown during individual development and are strongly influenced by interactions with the environment

How can information be used?

For an animal or robot to **use** information there are various requirements that need to be satisfied

- There must be some way of acquiring the information e.g. through the genes, through the senses, by reasoning or hypothesis formation, or some combination

  NOTE: acquisition through senses is not a simple matter of recording sensory signals: A great deal of analysis, interpretation, abstraction, and combination with prior knowledge may be involved.

- There must be some way in which the information can be encoded or represented so that it can be used. This may be
  - transient and used once

    (e.g. information used in continuous control)
  - enduring and reusable

    for short or long periods e.g.
    * percepts and immediate environment
    * generalisations,
    * geographical (extended spatial) information,
    * motives, preferences, values, ...
    * intended or predicted future events/actions, ... etc.

- There must be mechanisms for selecting relevant information from large stores.

- The form of representation must allow information-manipulations that derive new information or construct new hypotheses or goals

- Some animals (e.g. humans) and some robots need ways of representing novel information about things never previously encountered.

Some videos/demos

- **Parrot**

  Video of parrot scratching back of head with feather

- **Crow**

  Betty, the New Caledonian Crow makes hooks to lift a basket of food.

- **Infant helper**

  Show Warneken video of child spontaneously opening cupboard door to help researcher. (On web site mentioned above)

  The researchers ask: can a child spontaneously decide to help someone?

  I ask: if a child can do anything of that sort — what are the representational and architectural requirements?

  Remember the deliberative sheepdog: Demo 5 here:

  `http://www.cs.bham.ac.uk/research/projects/poplog/figs/simagent`

- **Broom pusher**

  Toddler steers broom down corridor and round corner

Non-human animals and and pre-verbal children must have information represented using **internal** languages.

See Sloman 1979

The primacy of non-communicative language

`http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#43`

What representations are NOT

It is often stated, e.g. in AI text-books, and in research publications, that representations “stand for” or “stand in for” or “resemble” the things they represent: **this is a serious confusion.**

- Information about X and X itself will generally be used for quite different purposes.

  A recipe for a type of cake gives information about ingredients and actions to be performed to make an instance.

  – If you mix and cook the ingredients properly (eggs, flour, sugar, etc.) you may get a cake.

  – Compare trying to mix and cook bits of the recipe (e.g. the printed words for the ingredients)

- A 2-D representation of a 3-D object cannot be used as a replacement for the object.
- If X is some type of physical object, then information about X might be used to work out how to make X, to decide whether to make X, to reason about the cost of making X, to work out how to destroy X, how to produce a better X, ...
- If X is a type of action, then information about X can be used to decide whether to perform X, to work out how long X will take, to work out risks in doing X, to decide how to perform X, to produce a performance of X, to modulate the performance of X, to evaluate the performance of X, to teach someone else how to perform X.
- If X represents a generalisation, e.g. “All unsupported objects near the surface of the earth fall”, then there is no object X refers to that can be used, manipulated, modified etc.

AI and Evolutionary Theory

- The study of biological evolution has many facets.
- Evolution of physical forms has been studied most, partly because that is what is most directly evident about differences between animals, and partly because much fossil evidence is about physical form.
- In recent years, study of the evolution of genetic makeup in DNA has accelerated.
- Some people try to understand the evolution of behaviour and intelligence by attempting to draw conclusions from physical form of animals, and from evidence of products of behaviour.
- In the case of micro-organisms and some insects it is also possible to observe behavioural changes across generations, e.g. evolution of resistance to disease in some plants and resistance to medical treatments in some antigens, but not for animals that evolve more slowly, like humans.
- In any case, observing changes in behaviour is different from understanding what produces those changes.

Observing what animals do, when they do it, and which don't do it leaves open the deeper questions: How do they do it, and how did that competence evolve or develop?

A complementary approach, based on AI

Many aspects of biological evolution look different from the standpoint of **designers of behaving systems:**

That standpoint raises the question:

How could such and such **information processing capabilities** have evolved?

And a host of subsidiary questions including how behaviours can be programmed through DNA, or some combination of DNA and the development environment.

- People who have tried to design behaving systems can more easily identify the problems evolution must have solved than people who merely observe, measure, and experiment on living systems.
- In particular, trying to design a working system teaches us a lot about what sorts of information the system will need, what forms of representation will not work, etc.
- However, at present, most of the competences of animals are more sophisticated than anything we can get our robots to do – because AI has only been going for about 50-60 years, whereas evolution had many millions of years, building many layers of competence that we are nowhere near replicating, except in very special cases.

Example: Evolution of language and language learning

Common assumptions about language and its evolution

It is commonly assumed that:

- The primary or sole function of language is communication between individuals;
- Language initially evolved through primitive forms of communication
  - Vocal communication according to some theories
  - Gestural communication according to other theories;
    (E.g. see paper by Fadiga L., Craighero L. 2007)
- Only after a rich external language had evolved did internal uses of language evolve;
  E.g. evolution produced shortcuts between brain mechanisms so that people could talk to themselves silently.

By adopting **the design stance** we can challenge the most popular views of evolution of language, and replace them with something more subtle and complex.

This involves looking at the information-processing requirements of tasks performed by pre-verbal children and by animals that do not use external languages like human languages.

This work was done jointly with Jackie Chappell, in Biosciences, Birmingham.
It is still at an early stage.

Our theory summarised:

Languages evolved for **internal** purposes first, though not in the form in which we now know language, but with key features (described later) in common.

- As a result of this internal use, complex actions, especially actions based on complex intentions and plans, naturally became means of communication
- Humans evolved a sophisticated sign language capability
- Then spoken language took over
- But the evolutionary heritage of gestural language remains with us.

  (E.g. Some Down Syndrome children have difficulty learning to talk, but they learn a sign language much more easily.)

Note: The above claims do not deny that once language for communication developed, that helped to accelerate the development of human competences, partly through cultural and social evolution and partly through development of brain mechanisms well suited to learning from other humans.

What are the main features of a language that give it its power?

Three features will be presented later.

Some requirements for human and animal competences

A few reminders:

- Humans and other animals can take in information about objects and situations of varying complexity.
- We can notice and reason about portions of a situation that can change or move in relation to others, producing new situations.
- We can think about some these things even when we have never seen them happen and what we are thinking about is not happening.
- We can use these abilities in coping with dangers and opportunities in the environment, and in planning and controlling actions so as to use opportunities and avoid or solve problems in the environment.
- All those competences involve abilities to acquire, manipulate and use information about things that exist or could exist.
- That means we need mechanisms for creating, manipulating, storing, using, deriving new internal structures that encode information.
  I.e. there are mechanisms for using internal languages.
- How to do that is a major topic in research in AI – there has been some progress, but there are still many unsolved problems.

Some important features of a language

What are the main features of a language (external or internal) that give it its power?

- Structural variability in what is expressed and in the form of expression.

  So that different sorts of things, of varying complexity can be described or represented, e.g.

  - a dot on a blank surface
  - a collection of dots
  - a collection of dots and lines
  - a plan of the furniture in a room
  - a plan of a house
  - a generalisation about houses

- Compositional semantics – allowing more complex meanings to be built up from simpler meanings.

  Linguistic expressions are "combinatorial rather than wholistic".

  Meaningful components can be combined in different ways to express different things, including totally new things, e.g.: I ate a yellow crocodile with red spots for breakfast yesterday.

- Use for reasoning (predicting, explaining, hypothesising, planning) by manipulating and recombining parts of complex representing structures.

- Use for expressing motives, preferences, goals, values, etc.

  So that you can derive new predictions, plans, summaries, generalisations, explanations, hypotheses, designs, maps, computer programs, etc. from old information.

Illustrate with SHRDLU demo.

For a non-interactive video of the program running see demo 12 here
`http://www.cs.bham.ac.uk/research/projects/poplog/figs/simagent`

Generalising features of language

We can generalise the three features commonly thought to be core features of human language, as follows:

A language with structural variability, compositional semantics and means of making inferences

(a) need not be composed of things we would recognise as words:

e.g. think of musical notations, circuit diagrams, maps, charts, pictures, models of chemical molecules, computer programs, and interactive graphical design tools

(b) need not be used for communication:

e.g. it may be used entirely inside a perceiver, thinker, planner problem-solver, including uses for formulating goals, questions, hypotheses, plans, and percepts etc.

Let's use the label "Generalised Language" (GL) to refer to a form of expression or representation that has

– structural variability,
– compositional semantics
– means of making inferences,

which is capable of being used for any information-processing purpose at all, communicative or non-communicative.

Reasoning with spatial structures

Will the pen hit the rim of the mug if moved downwards?

In the scenes depicted, you can probably distinguish cases where the answer is clear from cases where the answer cannot be determined.

Where the answer is clear you can find the answer by imagining the pen moving down between the rectangular sheets, and working out whether it will hit the rim or not.

This is a simple illustration of a general point: we often reason spatially by visualising a view of some configuration and imagining parts moving around and seeing what happens.

Where the answer is uncertain, because of some ambiguity in what you see, you can probably imagine a way of moving left or right, or up or down, so as to remove, or reduce the uncertainty.

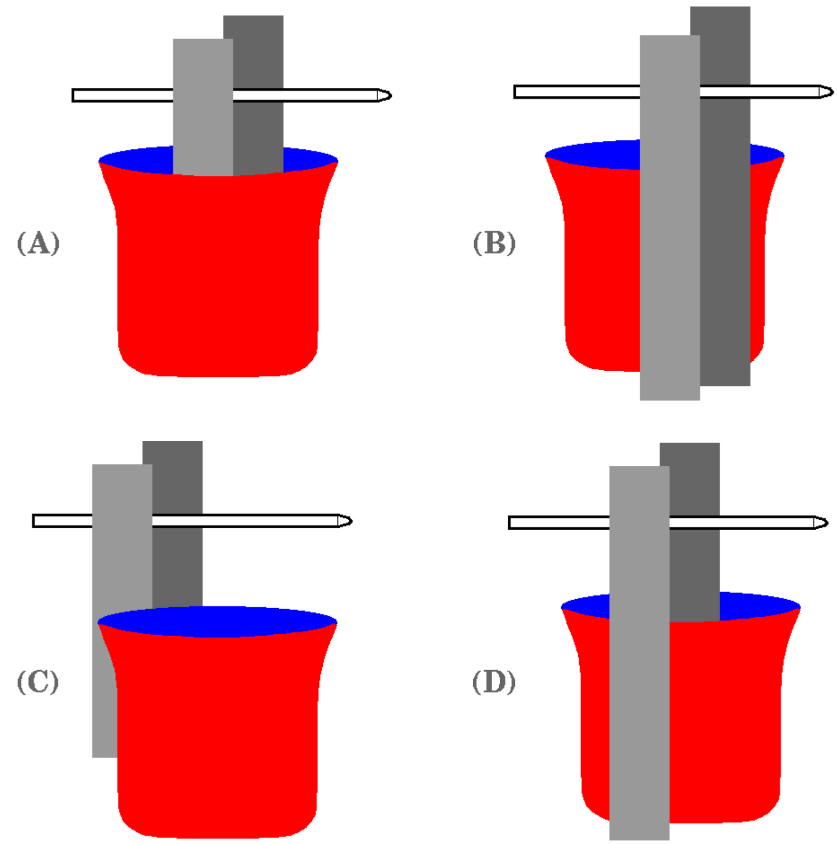

I argued in Sloman 1971 that visualisation can provide valid inferences, just as logical reasoning can, and that AI researchers need to investigate such modes of inference.

Main Theses

(a) GLs evolved in biological organisms for internal uses, before human languages developed for external communication

where internal uses included perception of complex situations and formation and execution of complex plans for action,

(b) GLs develop in pre-verbal human infants before they learn to use a human language for communication.

For examples of infant competences see
E. Gibson & A. Pike
An Ecological Approach to Perceptual Learning and Development, OUP, 2000

The main evidence for (a) is the fact that many non-human animals that do not communicate in anything recognisable as a human language, nevertheless have competences, which, from an AI standpoint, seem to require the use of internal GLs.

SHOW SOME VIDEOS, OF CHILDREN AND ANIMALS.
(See list at end of these slides.)

Conjecture: Gestural languages came first

If one of the uses of GL's was formulation of executable plans for action, then observing someone's action could provide a basis for inferring intentions: so actions could communicate meanings that had previously been expressed in internal GLs.

- In that case involuntary communication of plans by executing actions came first.
- The usefulness of such communication could have led to voluntary gestural communication, e.g. during performance of cooperative tasks.
- Since there was already a rich internal GL used for perceiving, thinking, planning, acting, etc. there could be both motive and opportunity to enrich actions to extend their voluntary communicative functions.

  The fact that there are already rich and complex meanings (including plan-structures) to be communicated, and benefits to be gained by communicating them (e.g. better cooperation) makes the evolution of rich forms of communication more likely.

- There are many explanations of the pressure to switch from gestural language (sign language) to spoken language, but that required complex evolution of the physiology of breathing, swallowing, and control of vocalisations.
- Empirical evidence of the primacy of sign languages:

  The example of Nicaraguan deaf children, and Down's syndrome children.
  (explained below)

Rival views about evolution of human language

1. First there were expressive noises which gradually became more differentiated and elaborate and then were “internalised”.

Only after that did thinking, planning, reasoning, hypothesising, goal formation, become possible.

Rival views about evolution of human language

1. First there were expressive noises which gradually became more differentiated and elaborate and then were “internalised”.

2. First there were expressive gestures, then noises, then as in 1.

Only after that did thinking, planning, reasoning, hypothesising, goal formation, become possible.

Rival views about evolution of human language

1. First there were expressive noises which gradually became more differentiated and elaborate and then were "internalised".

2. First there were expressive gestures, then noises, then as in 1.

3. First there were internal representations used for perceiving, thinking, forming goals, forming questions, planning, controlling actions; later, external forms developed for communicating meanings.

   Two options

   - externalisation was first gestural
   - externalisation was first vocal

NB: Do not assume such internal representations must be like Fodor's LOT (Language of Thought).

(For reasons explained later.)

Rival views about evolution of human language

1. First there were expressive noises which gradually became more differentiated and elaborate and then were "internalised".

2. First there were expressive gestures, then noises, then as in 1.

3. First there were internal representations used for perceiving, thinking, forming goals, forming questions, planning, controlling actions; later, external forms developed for communicating meanings.

   Two options

   - externalisation was first gestural
   - externalisation was first vocal

All the above options allow for the possibility that the existence of external languages and cultures produced evolutionary and developmental pressures that caused internal languages to acquire new functions and more complex forms.

Rival views about evolution of human language

1. First there were expressive noises which gradually became more differentiated and elaborate and then were “internalised”.

2. First there were expressive gestures, then noises, then as in 1.

3. First there were internal representations used for perceiving, thinking, forming goals, forming questions, planning, controlling actions; later, external forms developed for communicating meanings.

   Two options

   - externalisation was first gestural
   - externalisation was first vocal

All the above options allow for the possibility that the existence of external languages and cultures produced evolutionary and developmental pressures that caused internal languages to acquire new functions and more complex forms.

Our question is: what came first:

- external human languages?
- internal languages with core properties of human language? (Which core properties?)

Rival views about evolution of human language

1. First there were expressive noises which gradually became more differentiated and elaborate and then were "internalised".

2. First there were expressive gestures, then noises, then as in 1.

3. First there were internal representations used for perceiving, thinking, forming goals, forming questions, planning, controlling actions; later, external forms developed for communicating meanings.

   Two options

   - externalisation was first gestural
   - externalisation was first vocal

All the above options allow for the possibility that the existence of external languages and cultures produced evolutionary and developmental pressures that caused internal languages to acquire new functions and more complex forms.

Our question is: what came first:

- external human languages?
- internal languages with core properties of human language?

A similar question about what comes first can be asked about individual development.

Rival views about evolution of human language

1. First there were expressive noises which gradually became more differentiated and elaborate and then were "internalised".

2. First there were expressive gestures, then noises, then as in 1.

3. First there were internal representations used for perceiving, thinking, forming goals, forming questions, planning, controlling actions; later, external forms developed for communicating meanings.

   Two options

   - externalisation was first gestural
   - externalisation was first vocal

All the above options allow for the possibility that the existence of external languages and cultures produced evolutionary and developmental pressures that caused internal languages to acquire new functions and more complex forms.

Our question is: what came first:

- external human languages?
- internal languages with core properties of human language?

WHAT CORE PROPERTIES?

MORE ON CORE PROPERTIES OF GLs

Human languages (including sign languages) use many formats and have many features.

Earlier, I described three core properties required for using language in relation to novel situations for multiple uses,all found in both external human languages and internal GLs.

- **Structural variability:**

  Linguistic utterances can include varying numbers of distinct components and are not restricted to flat vectors but can have deeply nested substructures, with pronouns, other forms of anaphora and repeated elements providing cross-links.

  Familiar labels for this property include: 'generative' and 'productive'.

  An implication is that not everything that can be communicated has to be learnt, or previously agreed.

- **Compositional semantics:**

  Novel structures can be given a meaning in a systematic way on the basis of the meanings of the components and the mode of composition (i.e. structural or syntactic relationships between the components).

- **Manipulability:** (a consequence of the previous two)

  Meaningful structures can be extended, modified or combined for various purposes, discussed later.

I now explain in more detail how the idea of "compositional semantics" need to be generalised to meet all the requirements for internal GLs.

Standard compositional semantics

Conventional compositional semantics:

New combinations of words, phrases, pictures, and other components of meaningful structures, are understood because the meaning of a whole is determined by two things:

- the meanings of the parts
- the way the parts are assembled to form the whole.

However, the context in which information structures are combined also often helps to determine what the combination means for the user.

- Often it is not just the internal structure: the components of the representation and their relationships do not suffice to determine the semantics of a complex whole.
- Aspects of the context both inside the user of the representation and in the external environment are often important in determining what is expressed.

  (This is obvious with indexicals such as “now”, “here”, “this”, “you”.)
- So the standard notion of compositional semantics does not account for all uses of representational complexity.

Examples of generalised compositional semantics

Generalised (context-sensitive, situated) compositional semantics:

Meanings of complex wholes are determined by three things:

(a) meanings of parts,
(b) the way the parts are assembled to form the whole, and
(c) the linguistic and non-linguistic context (obviously true for indexicals, e.g. “this”, “here”, etc.)

Examples:

- “Put it roughly there.”

  You don’t have to be told exactly where, and there is no semantic rule determining the location.

  You have to use your judgement of the situation in selecting a location.

- “If you can’t see over the wall, find a big box to stand on.”

  You don’t have to be told exactly how big – use your understanding of what the box is wanted for.

- “The wind will blow the tarpaulin away so let’s put a pile of stones on each corner.”

  You don’t have to be told how many stones make a pile, and no semantic rule says how many: but you can work out that there must be enough at each corner to keep the tarpaulin down when the wind blows, and that will depend on how strong the wind is.

The role of context in compositional semantics generalises H.P. Grice’s “Cooperative principle” and his “Maxims of communication”, to include internal languages.

For more on this see:

Spatial prepositions as higher order functions: and implications of Grice’s theory for evolution of language.
`http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0605`

Illustrating compositional semantics

The notion of ‘compositional semantics’ was proposed by Frege and Peirce, and is normally summarised something like this

> The meaning of a complex expression is determined by the meanings of its parts, and the way in which those parts are combined.

So, a semantic function (**S**) which derives semantic content from this syntactic structure:

F(X, Y, Z)

Could be expressed as

S(F(X, Y, Z)) = **S**(F)(**S**(X), **S**(Y), **S**(Z))

For example, to evaluate the arithmetic expression: sum(33, 66, 99)
apply the procedure denoted by ‘sum’, i.e.

S(‘sum’),

to the numbers denoted by the symbols

‘33’, ‘66’ and ‘99’.

In other words: apply

S(‘sum’)

to

S(‘33’), **S**(‘66’) and **S**(‘99’)

All this will be familiar to every programmer, and even more familiar to compiler writers.

Generalising the formulation of compositional semantics

We generalise the familiar notion of compositional semantics by taking account of context and current goals (C, G) in the interpretation

We replace this equation

S(F (X, Y, Z)) =

S(F)(**S**(X), **S**(Y), **S**(Z))

with an equation in which the interpretation of every component, at every level, is (potentially) influenced by context and goals:

S(F (X, Y, Z)) =

S(F,C,G)(**S**(X,C,G), **S**(Y,C,G), **S**(Z,C,G))

Neither C nor G is a component of the representation: rather they are parts of the "environment" accessed by the processes interpreting and using the representation.

Compare C with the interpretation of non-local variables in computer programs.

In some contexts one or both of C and G may not be needed.

For instance in many mathematical and programming contexts, such as evaluation of arithmetical expressions, C and G will not be needed.

However in much human communication C and G (Context and Goals) are both needed.

For more on this, see the discussion paper mentioned above:
`http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0605`

A diagrammatic version

Generalised (context-sensitive, situated) compositional semantics can be explained diagrammatically:

New combinations of words, phrases and other components are understood because the meaning of a whole is determined by **three things**

- the meanings of the parts
- the way the parts are assembled to form the whole
- linguistic and non-linguistic aspects of the context, including
  - the physical environment
  - the goals of the speaker and hearer
  - current tasks in progress ... and other things

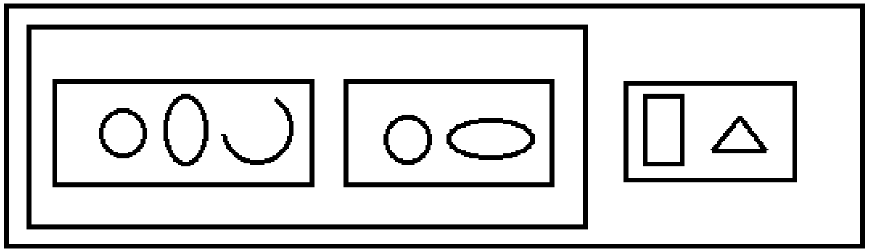

Context free semantic composition

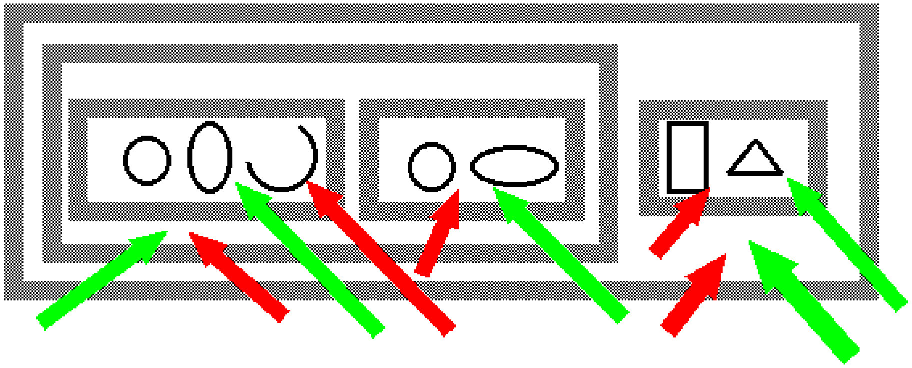

Context sensitive semantic composition

Formally, we can think of every syntactic construct (every box) as having extra arguments that enrich the interpretation: the current context and current goals (which may or may not all be shared between speaker and hearer). The contexts that enrich the semantics (green and red arrows) may come from inside the symbol user, or from the external physical or social environment.

Manipulation requires mechanisms

The mere fact that a form of representation supports manipulability as explained above does not in itself explain how actual manipulation occurs in any machine or animal.

That requires **mechanisms** to be available that can construct, modify, combine, store, compare, and derive new instances of representations.

E.g. new phrases, new sentences, new stories, new plans, new diagrams, new working models

If an animal or machine has a large repertoire of information and mechanisms, selecting the appropriate ones to use can itself require additional mechanisms and additional information about how to use the resources.

AI systems typically have powerful abilities, but current systems don't know that they have them; nor can they choose which ones would be best to use: except by following simple pre-programmed rules, which they don't learn, and don't modify.

That will need to be changed.

At present we still have a lot to learn about how to build mechanisms that grow themselves in a machine with human-like competences.

What does “internalising language” mean?

What does the blue part of this common assumption mean:

External human language evolved from primitive to complex communication, and was later internalised. (NB: I am not defending this claim.)

The reference to being internalised could mean something like this:

- Evolution several times extended brain functions so that mechanisms that originally evolved for peripheral modules become available for purely internal uses
  e.g. visual mechanisms later used for imagining?
- Modules evolved for linguistic communication were later modified for internal use, in something like this sequence of steps (e.g. proposed in Dennett 1969?):
  - After external languages evolved for communication, humans discovered that it could sometimes be useful to talk to themselves, e.g. when making plans, solving problems, formulating questions ...
  - Subsequent evolutionary changes enabled **talking silently**: i.e. brain mechanisms became able to provide inputs directly to the speech input portions of the brain, instead of having to route them externally.
  - This made it possible to construct internal meaningful, manipulable linguistic structures that could be used to think, plan, reason, invent stories, solve problems, construct explanations, remember what has happened, etc.
    (Daniel Dennett, *Content and Consciousness*, 1969.)

However, such theories of “internalisation” ignore the internal representational (GL) mechanisms required for external language use in the first place. (Sloman 1979)

Biological relevance

THESIS: Some animal competences and some competences of pre-linguistic children need richly structured internal, manipulable forms of representation with context-sensitive compositional semantics, which are constructed and used for perception, reasoning, planning and generation and achievement of goals related to complex features of the environment.

- I have tried to bring out some of the possible uses of GLs with the three core properties: structural variability, compositional semantics, manipulability.

  (Later generalised to include spatial – e.g. diagrammatic – forms of representation).

- We can point to many competences displayed by prelinguistic children and some other species that are hard to explain without the use of GLs

  Examples include nest-building, hunting, dismembering a carcass in order to eat it, playing with toys, using tools, making tools, fighting with others, collaborating with others.
  In particular both Humean and Kantian causal reasoning require use of GLs, though in different ways.

- An important point I shall not have time to go into is the need for specific forms of GL that provide meta-semantic competences, e.g. the ability to represent and reason about one's own or others' goals, beliefs, thought processes, preferences, planning strategies, etc. (So-called "mentalistic" vs "mechanistic" cognition).

  For more on that see
  `http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0604`
  Requirements for "fully deliberative" architectures.

Direction of fit of GL structures to the world

Many information structures (but not all!) are used to refer to some portion of the world and represent that portion as having certain features, possibly quite complex features:

> in principle such things can be true or false, or in some cases more or less accurate or inaccurate, more or less close to being true, etc. all depending on how the world is.

Various philosophers (e.g. Anscombe, Austin, Searle) have pointed out that two major kinds of use of such structures can be distinguished:

- where the information-user tends to construct or modify **the representation** so as to make it true or keep it true (belief-like uses)
- where the user tends to monitor and alter **the world** so as to make or keep the information structure true (desire-like uses).

Sometimes referred to as a difference in "direction of fit" between beliefs and desires.

The distinction also has a clear role from the standpoint of designers of robots or other intelligent systems, though, as I've shown elsewhere, there are more intermediate cases to consider in complex, multi-functional machines (e.g. animals).

These ideas about belief-like and desire-like states of an organism or machine are developed further in:

A. Sloman, R.L. Chrisley and M. Scheutz,

The architectural basis of affective states and processes, in
Who Needs Emotions?: The Brain Meets the Robot, Eds. M. Arbib & J-M. Fellous, OUP, 2005, pp. 203–244,
`http://www.cs.bham.ac.uk/research/cogaff/03.html#200305`

Desires, beliefs and direction of fit

Content vs function of mental states

Both beliefs and desires can be checked against current perceptual input, but the consequences of mismatches are different.

What makes something a desire, or belief, or fear, or idle thought depends not on the form of the information structure, nor its medium, but on its causal role in the whole architecture.

Simple architectures allow for only simple causal roles, whereas more sophisticated architectures allow information structures to have very varied causal roles.

To understand fully the variety of functions served by GLs in a particular type of animal (or machine) we would need to have a detailed specification of the information-processing architecture.

We are not ready for that yet!

See the presentations on architectures here
http://www.cs.bham.ac.uk/research/projects/cogaff/talks/

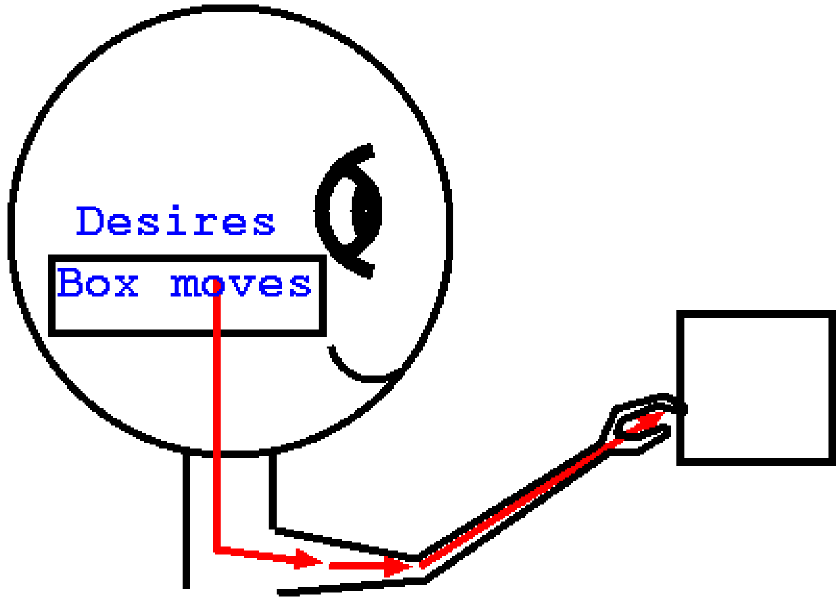

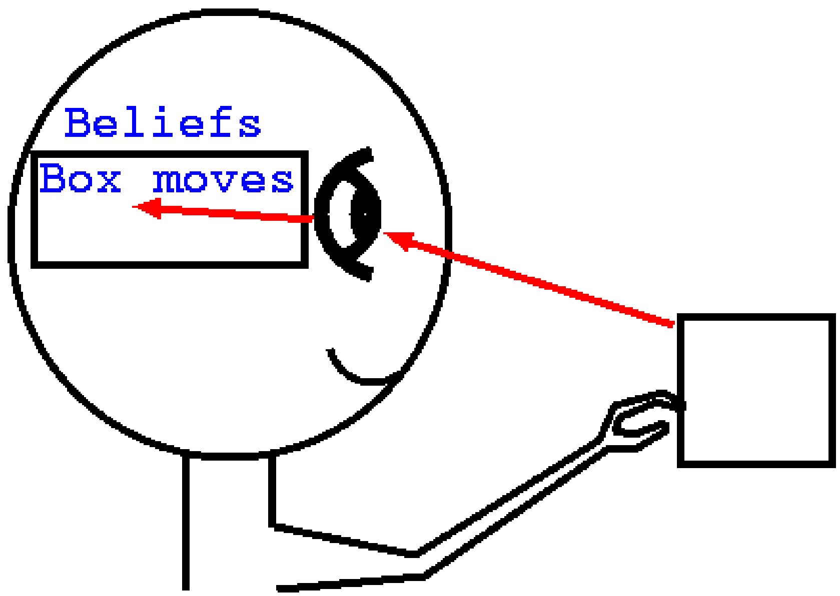

Varieties of uses of internal GLs

Within an organism or robot, a GL structure may have many different kinds of use: depending on the conditions under which it is created, how it is used, what sorts of things modify it and when, and what effects it has and what sorts of things can affect it. For example,

- The use of representations in perceptual subsystems is related to one direction of fit (produce information structures that represent how things are)
- Their role in motivational subsystems is clearly related to the other direction of fit (change the world so that an information structure represents how things are.)
- An organism's or robot's ability to have very diverse beliefs, desires and competences is connected with the structural variability and compositional semantics of its GLs.
- GLs can be substantially extended during development: they are not innately given.
- Some representations need to endure and be usable in different contexts (e.g. facts, values, competences), whereas others are needed only transiently (e.g. feedback).
- The conditions for a GL to be used for planning several steps ahead are different from the conditions for using information for online control of continuous actions.

  The former requires more complex virtual machines that evolved much later and in relatively few animals, and benefits from an animal's ability to represent states of affairs and processes independently of the sensory and motor signals involved in perceiving or producing them, using an amodal, exosomatic ontology.

  I suspect confusion about so-called mirror neurones can arise from a failure to understand that point. (Should they have been called 'abstraction neurones'?)

Other uses of GL structures in humans

Besides expressing semantic contents for desire-like and belief-like states, GL structures can have a wide variety of causal roles, depending not only on their location in the architecture, but also on their form and the mechanisms available for manipulating them. E.g.

- Comparing and evaluating things, states of affairs, possible actions, goals, policies, ...
- creating more or less complex plans for future actions
- using a plan to control actions (either continuously, as in visual servoing, or step-by-step)
- synchronising concurrent processes, or modulating ongoing processes
- expressing a question,

  i.e. constructing a GL structure that directs a search to determine whether it is true or false, or how it needs to be modified or expanded to make it true.

  `http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0502`
- considering unobserved possibilities to explain what has been observed,
- predicting things that have not yet happened

  (e.g. Humean or Kantian causal reasoning),
- fantasising, e.g. wondering what would have happened if,
- inventing stories
- day-dreaming
- meta-management functions (making use of meta-semantic competences).

Most animals, and current robots, have much simpler information processing competences.

A consequence of the core features

A consequence of the core features is that it is possible to produce well-formed linguistic expressions for which the compositional semantics will produce an impossible (internally inconsistent) interpretation.

E.g. Consider this conjunction

> Tom is taller than Mary
> and Mary is taller than Jane
> and Jane is taller than Dick
> and Dick is taller than Tom

If

> (a) 'Taller than' has its normal meaning
> (b) Each repeated occurrence of the same name refers to only one individual

then

> That conjunction is inconsistent: not all conjuncts can be true simultaneously.

We'll see a similar kind of inconsistency in non-verbal forms, later.

Inconsistency of an information structure implies that

- if it is adopted as a belief it will be a necessarily false belief,
- if adopted as a goal it will be a necessarily unachievable goal, and
- if constructed as a percept it will be a perception of an impossible state of the world.
  (illustrated later)

(Compare G. Frege on failure of reference.)

Another generalisation: Non-verbal forms

The three core features of human languages **structural variability**, **generalised compositional semantics** and **manipulability** are also features of many non-verbal forms of representation.

Given a map, a flow-chart for an algorithm, a circuit diagram, or a picture of an object to be constructed, more components can go on being added, sometimes indefinitely.

If we use paper, or some other markable surface, it is possible to

- expand a picture or diagram outwards,
- add more internal details (e.g. extra lines),
  but eventually there is 'clutter limit' because the structure is not stretchable.
  (Other kinds of limits relate to short-term memory constraints.)

Structural variability of such spatial forms of representation has recently been enhanced by the use of film or computing techniques that allow zooming in and out to reveal more or less of the 'nested' detail.

It is possible that virtual machines evolved in brains allow such 'zooming' in and out, though precise requirements for such a facility to be useful still need to be specified.

The retinoid model of Arnold Trehub's *The Cognitive Brain* (MIT Press, 1991) may be an example.
`http://www.people.umass.edu/trehub/`

Sloman 1971 (ref. at end) describes more precisely a distinction between "Fregean" and "analogical" forms of representation, claiming that both can be used for reasoning, planning, and proofs.
This was a criticism of the "Logicist" AI approach expounded by McCarthy and Hayes, in 1969.

Compositional semantics and structural variability in vision

Your familiarity with the role of low level pictorial cues in representing features like edges, orientation, curvature of surfaces, joins between two objects or surfaces, etc., allows you to use compositional semantics to see the 3-D structure, and some causal and functional relationships, in pictures you have never previously seen.

No AI vision system comes close to being able to do that – yet.

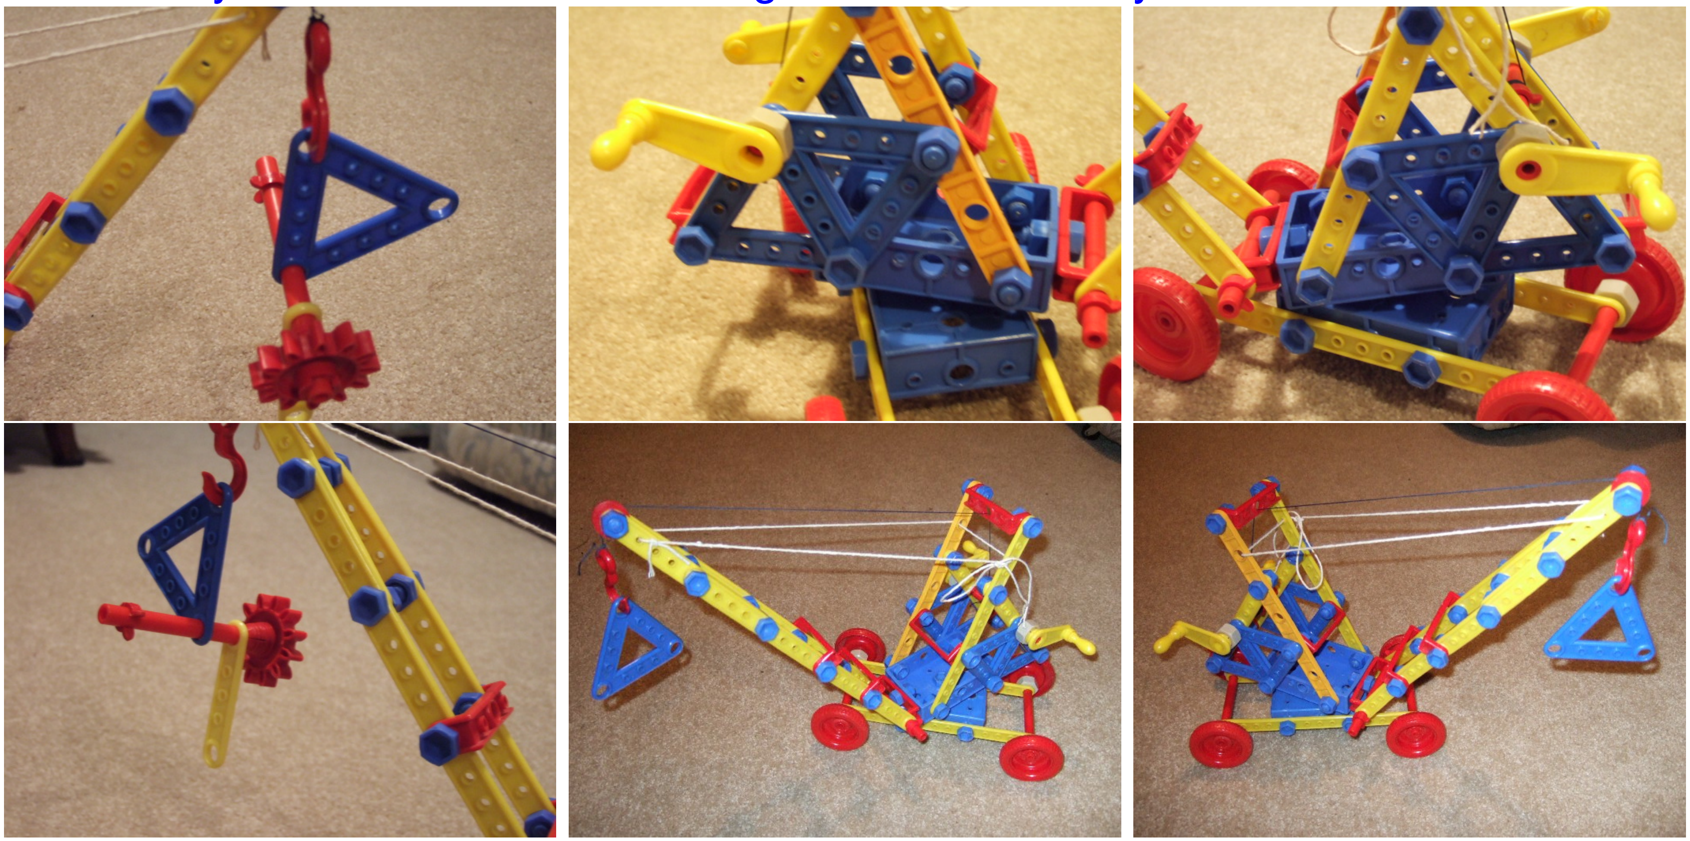

http://www.cs.bham.ac.uk/research/projects/cosy/photos/crane/

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

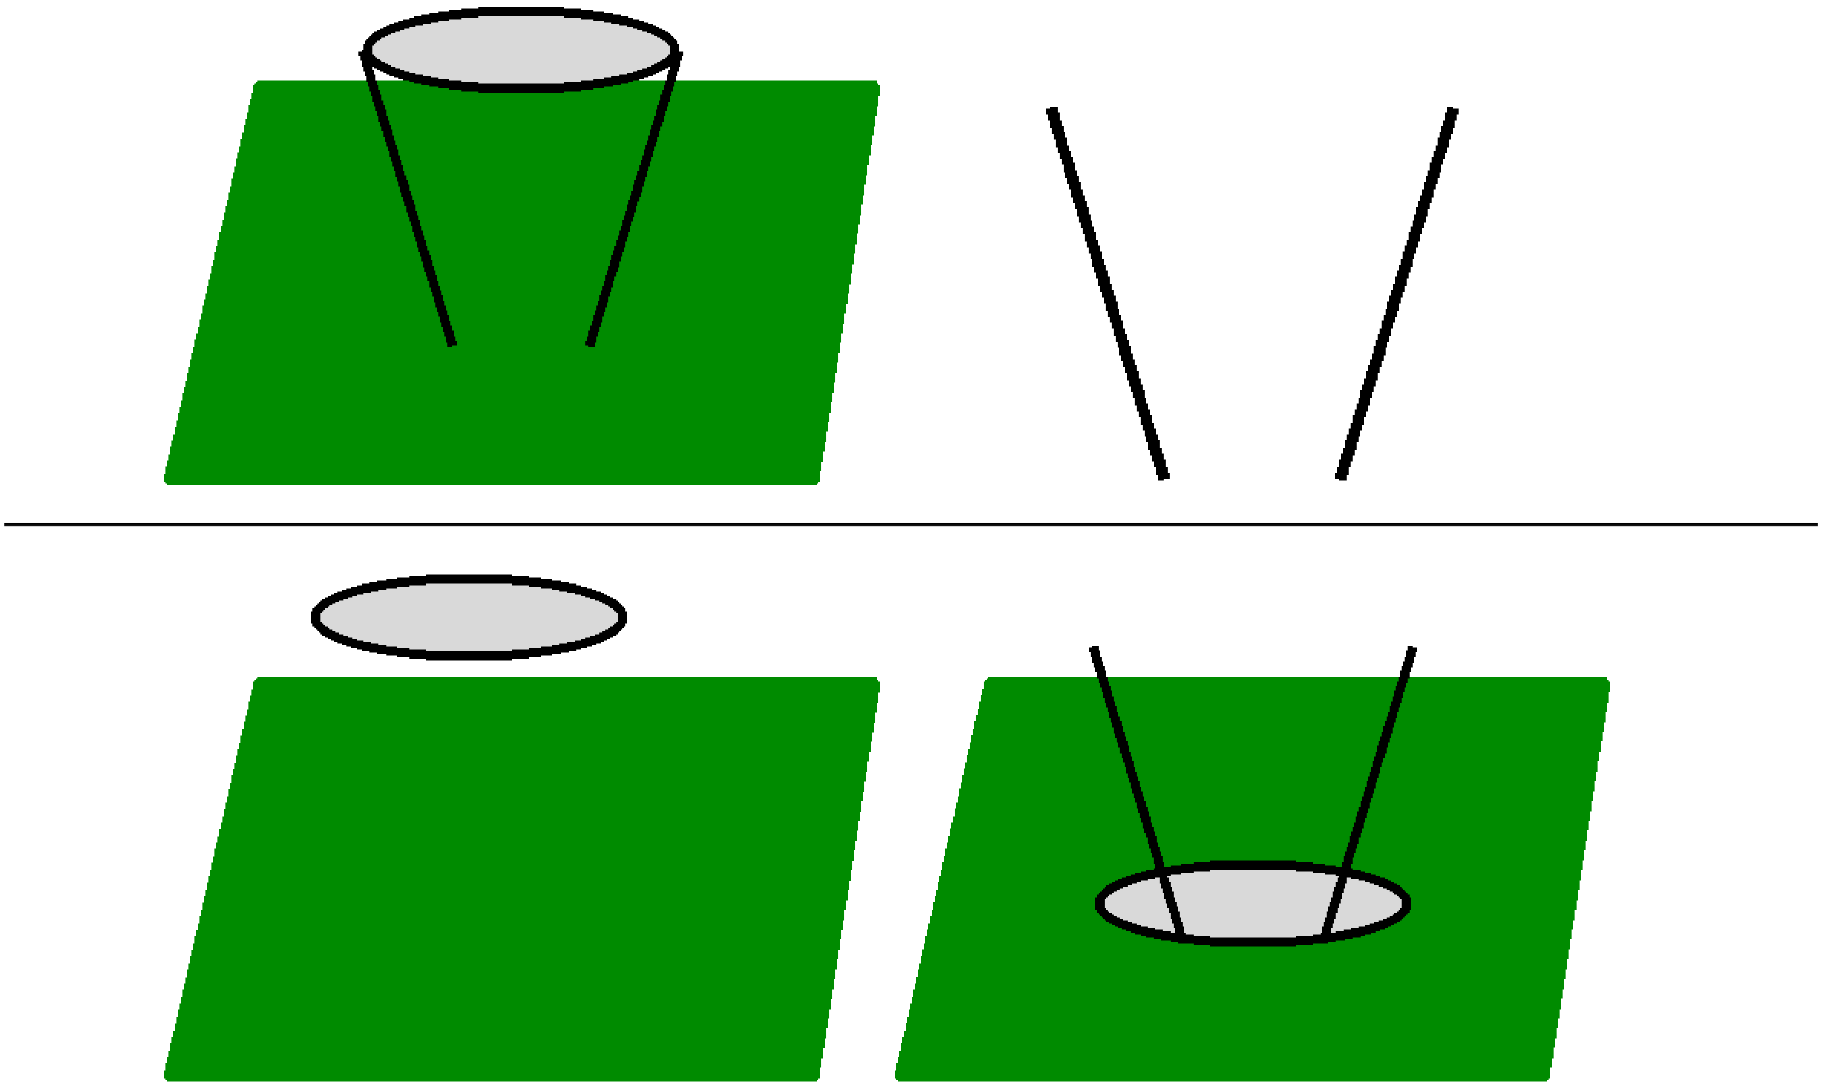

A droodle: Can you tell what this is?

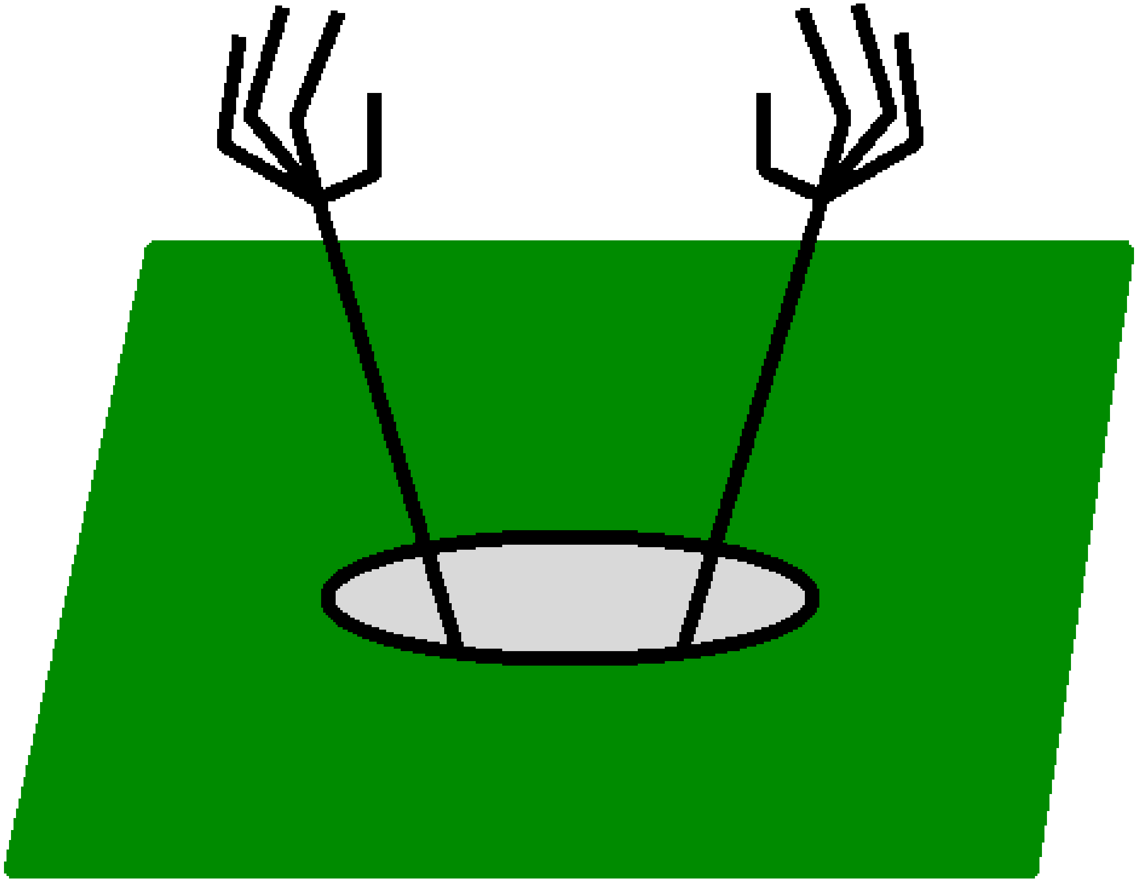

Droodles depend heavily on the fact that interpretation of visual GL instances can be can be partly driven by sensory data and partly by verbal hints ("top down").

Possible answers to droodle question

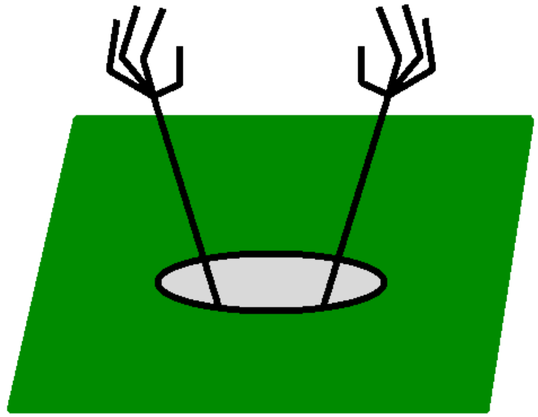

"Early bird catches very strong worm?"

"Sewer worker attacked by a shark?"

Interpretation of visual scenes can include perception of causal relationships, as in both the above droodle interpretations.

There is much to be said about droodles, but no time today.

Perceptual combination of spatial and causal relationships is also needed in use or construction of tools: e.g. shape of a spanner's head.

When objects share a region of space, indefinitely many different kinds of structural and causal relationships can be perceived and interpreted: in contrast with the constrained, rule-based, use of syntactic relations in human formal and informal languages.

Show broom video, available here (with others)
`http://www.cs.bham.ac.uk/research/projects/cosy/conferences/mofm-paris-07/sloman/vid/`

Long before children can talk, they can take in and make use of structural relationships in the environment in order to produce and control actions.

That's in addition to their ability to manipulate continuously changing dynamical systems, e.g. maintaining balance while walking, reaching, etc.

Likewise many other animals.

Perceiving spatial structure vs creating images

Information structures in a spatial GL should not be confused with images

An image is a very large collection of small image features,

which may include colour, brightness, texture, edge-features, optical flow, and various kinds of gradient, and various metrical and qualitative 'low-level' relationships such as brighter, same colour, coarser textured, so many degrees apart, etc.

For pictorial or spatial GLs to be useful in the ways described, they must be composed of larger structures with more global relationships not restricted to simple metrical comparisons.

This topic was much discussed by AI vision researchers in the 1960s. See
S. Kaneff (ed) *Picture Language Machines*, Academic Press 1970.
and `http://hopl.murdoch.edu.au/showlanguage.prx?exp=7352&language=Clowes`

The larger structures

- may be image components like lines, regions, polygons, with relationships like touching, enclosing overlapping, being collinear, approaching, etc., OR
- they may be representations of 3-D or other objects and processes represented by the 2-D structures, e.g. fingers, pools, planks, rocks, rivers, trees, trains going into tunnels, etc., with static and changing 3-D and causal relationships, e.g. supporting, penetrating, grasping, pushing, going behind, etc.

For the user of the GL to be able to perform manipulations and transformations that are useful for tasks like predicting, planning, explaining, formulating questions, it is necessary to do something like parsing of the representations, i.e. segmenting them into relatively large components with relationships, so that either components or relationships can be changed.

This is quite unlike what is called "image processing", e.g. enhancing images or applying global transformations to them, such as convolution.

Making an "H"

Making a capital "H" using an elastic band and pins

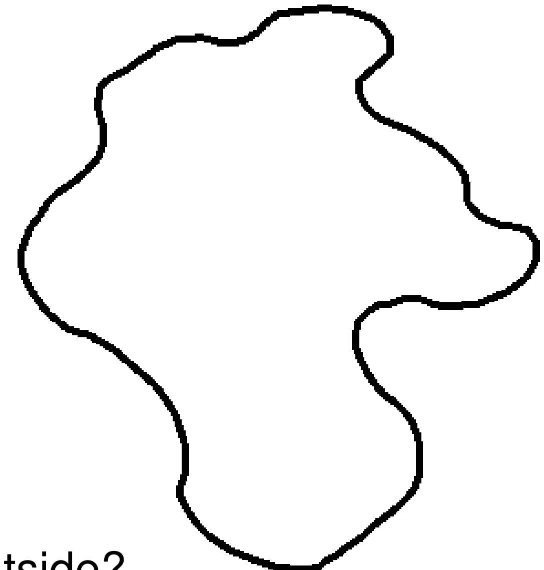

Suppose you had an elastic band and a pile of pins:

could you use the pins to hold the stretched rubber band in the form of a capital "H"?

What sort of GL is needed to make it possible to answer such a question?

- How many pins would you need?
- Could you do it using only one hand?
- In what order would you insert the pins?
- How many pins would be inside the band and how many outside?
- Could you do it if the pins were replaced with marbles?

You can probably answer the questions in two ways: by trying physically and examining what happens, and by merely thinking about it and examining what happens.

- A very young child will not be able either to construct the H physically, or to answer the questions.
- You are probably able to answer the questions just by thinking about the construction processes and the result.
- What is your brain doing while you visualise the process of creating the final configuration?
- Do you first visualise the final configuration, and then make a plan for constructing it, or do you get to the final configuration by making a plan, or visualising the construction process?
- What is your brain doing while you count the imagined pins, inside or outside the band?

Major problems for vision researchers

Relationships between **static** complex 3-D objects involve many relationships between parts, some metrical, some topological, and some causal/functional. I.e. relationships between complex, structured, objects

are multi-strand relationships.

When **processes** occur involving changing or moving 3-D objects, many relationships can change at the same time:

they are multi-strand processes.

- The changes are not just geometrical.
  They can include changing causal and functional relationships
  (e.g. supporting, compressing, obstructing, etc.).
- Perception of processes can include perception of changing affordances.
- I.e. perceived changes can involve several ontological layers.

We can perceive multi-strand processes in which complex 3-D objects change many relationships at the same time. What forms of representation and what mechanisms make that possible? As far as I know, neuroscientists have no explanations and AI vision researchers have no working models.

For more on that see

`http://www.cs.bham.ac.uk/research/projects/cosy/papers/#pr0505`
A (Possibly) New Theory of Vision (October 2005)

`http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#compmod07`
Architectural and representational requirements for seeing processes and affordances.
(31 May 2007, BBSRC Workshop)

Partial summary so far

Many familiar kinds of competence involving

- perception of 3-D structures and processes,
- planning and control of actions in a 3-D environment,
- predicting and explaining processes in the environment

require the use of structured, manipulable internal forms of representation with context-sensitive compositional semantics.

Those forms of representation, GLs, have some (but not all) features of human language, but use additional mechanisms and are used internally for information processing.

Some of the manipulations that are possible are discrete (e.g. adding or removing an object, or a contact or containment relation), others continuous e.g. sliding something, distorting a shape.

In some forms of GL, the structural and functional relationships in the interpretation arise from spatial embedding of different parts of the same information structure: rather than use of arbitrary or totally general syntactic conventions (as in language and logic).

Nevertheless the spatial form of representation is not a structure that is isomorphic with what it represents.

This can be demonstrated using pictures of impossible objects.

Some of these points were made in Sloman 1971 and in Sloman 1979

Building a configuration of blocks - 1

Given a collection of cubes and rectangular blocks could you arrange them to look like this?

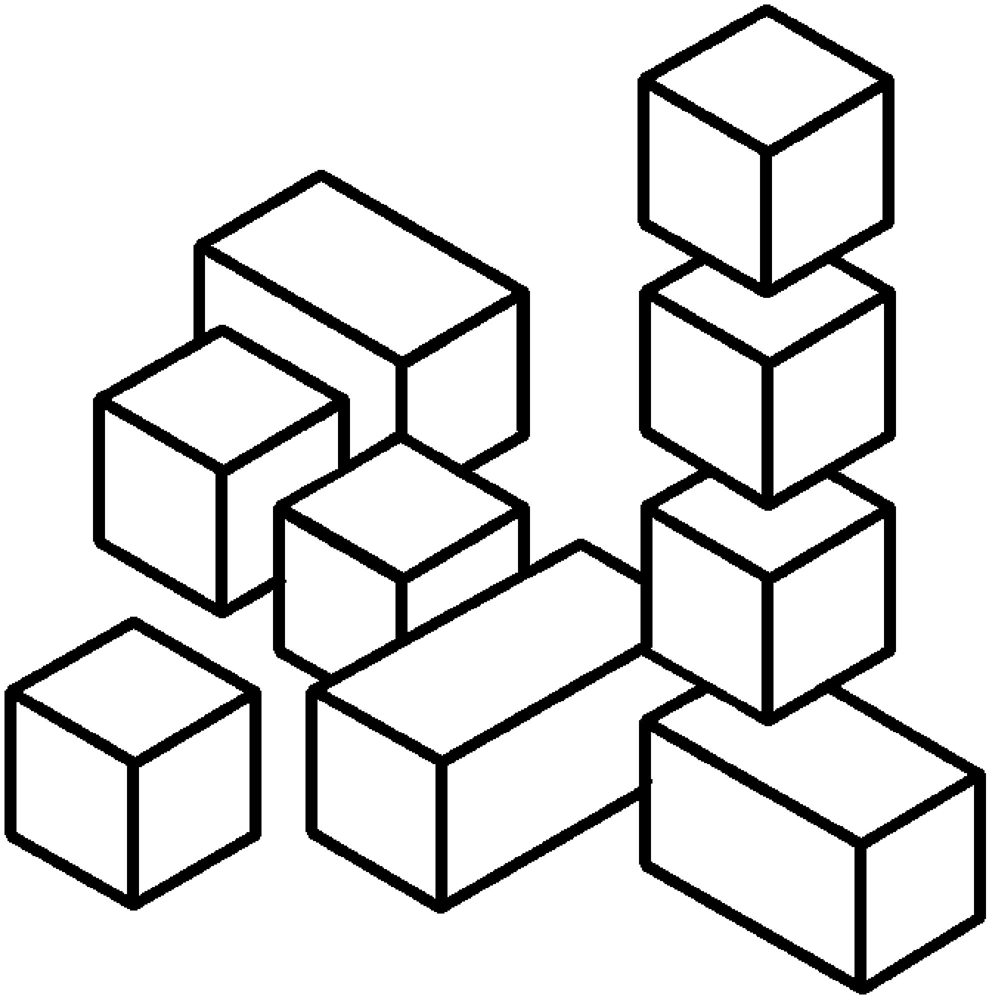

Think of locations to which you could you move the ‘loose’ cube on the left.

Building a configuration of blocks - 2

Moving one cube, could you re-arrange them to look like this?

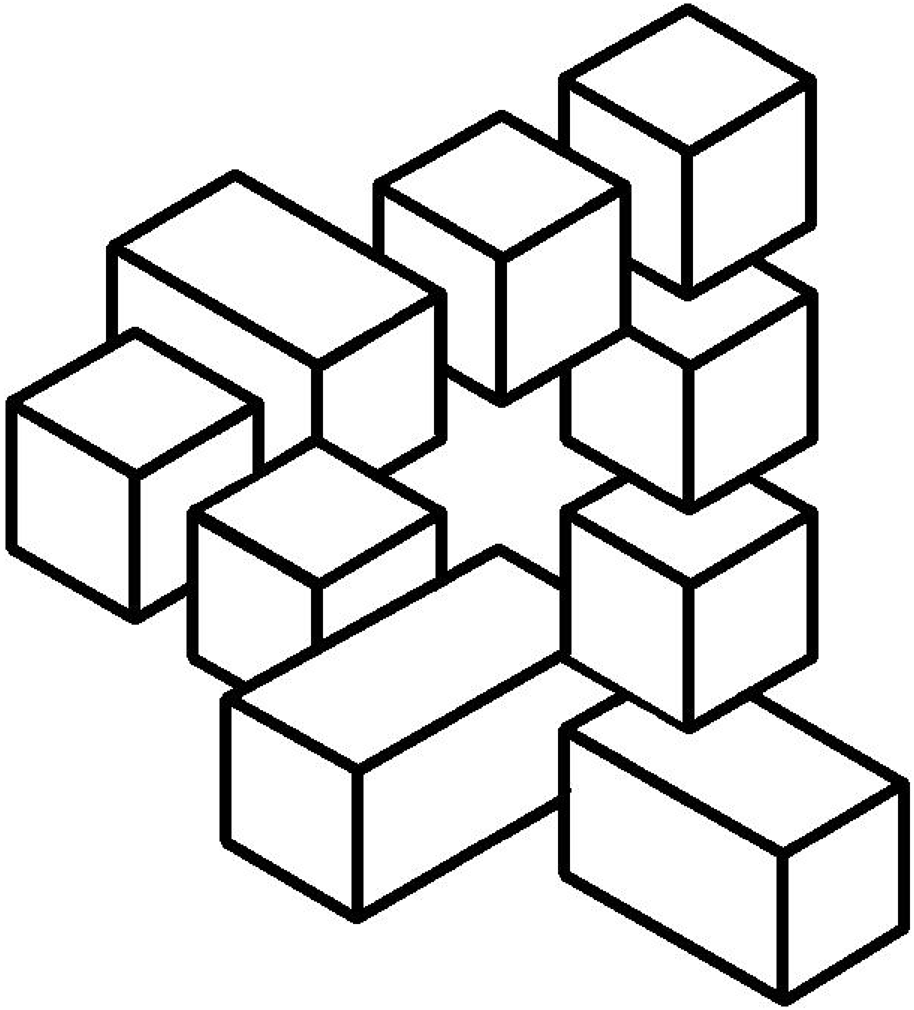

Some young children will say ‘yes’.
What has to change for them to be able to detect the impossibility?

Implications of pictures of impossible objects

The impossible pictures rule out the assumption that seeing involves building a structure that is isomorphic with what is seen: for it is impossible to build a structure that is isomorphic with an impossible structure.

What we (and other animals?) do must be much more subtle, general and powerful, and connected with manipulability, structural variation, and compositional semantics, all of which are important in seeing affordances.

The example of logic shows that it is possible to assemble coherent fragments of information into an incoherent whole: this seems also to be what happens when we see pictures of impossible objects, though in that case we do not seem to be using a logical formalism.

Exactly what sort of GL suffices for the purpose requires further research,

- We need to analyse requirements for GLs, including both being usable for representing what exists and being usable for representing and reasoning about changes that are possible
- We seem to use those features of GLs in understanding many examples of causation.

Fortunately we don't normally need to check for consistency because the 3-D environment cannot be inconsistent.

See also `http://www.cs.bham.ac.uk/research/projects/cogaff/challenge-penrose.pdf`

Examples: To be expanded

Show Felix Warneken movies showing prelinguistic children and chimps apparently spontaneously determining and responding to goals of an adult human.

This requires them not only to use GLs without being able to talk but also possessing some meta-semantic competence.

`http://email.eva.mpg.de/~warneken/`

Warneken was mainly concerned with evidence for altruism.

I am mainly concerned with the cognitive mechanisms presupposed by the performances, whatever the motives.

Nest building birds, e.g. corvids.

Could you build a rigid nest using only one hand (or hand and mouth), bringing one twig at a time?

Betty making hooks in different ways and using them for a common task.

Search using google for

betty crow hook

Humans can solve many problems about spatial structures and processes in their heads, illustrated in previous slides.

Implications of the examples

GLs are needed for many capabilities shown by other animals and capabilities shown by pre-linguistic children.

So they cannot be a by-product of evolution of human language.

Since GLs can express plans that can be used to control actions, and since actions can reveal intentions, they are already well suited as a basis for generating communicative language

Implication: sign-languages evolved first, but previous theories about how that happened must be wrong

E.g. theories claiming that simple external gestures arose first, then increasing complexity, then vocalisation and finally internalisation must be back to front.

Not Fodor's Language Of Thought

There is no implication in any of this that a human, or nest-building bird, or intelligent language user, must start with an 'innate' (or genetically determined) GL that suffices to express everything it will ever need to express, so that all future meanings are definable in terms of the initial set.

Papers with Chappell investigate ways in which boot-strapping processes can substantially extend innate competences through exploration and play in the environment along with the ability to construct new explanatory theories to account for surprises.

This can include substantial ontology extension: introducing concepts that are not definable in terms of previous ones, e.g. using model-based semantics and symbol/theory-tethering.

For more on that see
http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#models
http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#grounding
http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#grounding
http://www.cs.bham.ac.uk/research/projects/cosy/papers/#pr0702

That option was not open to Fodor because he used a model of understanding based on compiled programming languages, where all programming constructs are translated into the machine language before the programs run.

He apparently forgot about interpreted programming languages and perhaps did not know about logical programming languages (e.g. prolog).

He should have known about model-theoretic semantics, but failed to see its relevance, as described in the presentations listed above.

Unanswered questions

Despite the evolutionary continuities between humans and some other species it is clear that there are many spectacular discontinuities

(e.g. only humans make tools to make tools to make tools .... to build things, and it seems to be the case that only humans prove mathematical theorems, enjoy thinking about infinite sets, tell stories to one another, etc.).

What explains these discontinuities?

We need to consider various possibilities:

- Was there some change in degree that went past a threshold whose effects were then amplified? (E.g. some memory structure increased in size?)
- Was there a particular crucial discontinuous change in architecture, or some mechanism, or some form of representation, after which effects cascaded?
- Were there several different changes, with independent causes, which combined to produce new spectacularly different effects?
- other possibilities???

We don't know enough to answer, but I suspect the first answer (a quantitative change passed a threshold) is unlikely.

I suspect there were a few key discrete architectural changes, that modified the forms of learning and development in humans and other altricial species (see below).

Language learning vs language development

If the previous observations and speculations are correct, previous theories about language learning must be wrong!

- Previous theories imply that children do not acquire a way of representing information that supports structural variability, compositional semantics and useful manipulability until they have learnt an external human language, which they do by some sort of data-mining of perceived uses of language by others.
- If our speculations are correct, the process of language learning is primarily one of creative and collaborative problem solving in which new ways of expressing pre-existing meanings are devised collaboratively.
- This is a process of development of internal GLs along with their extension to an external mode of expression.
- The fact that learners are normally in a minority and can have little influence on the outcome makes it look as if they are absorbing rather than creating.

  But the Nicaraguan case shows that must be wrong. Nicaraguan deaf children rapidly developed a new sophisticated sign language which they used very fluently and which their teacher was unable to learn.

Once humans had acquired the ability to communicate rich and complex meanings, cultural evolution, including development of new linguistic forms and functions, could enormously speed up transmission of information from one generation to another and that might produce evolutionary opportunities to extend the internal GL-engines.

Implications for Chomskian theories

Does all the above imply that humans have anything like the kind of innate (genetically determined) Language Acquisition Device (LAD) postulated by Chomsky

(E.g. in Aspects of the Theory of Syntax, 1965

or is the learning of human language completely explained by general purpose learning mechanisms at the basis of all human intelligence?

Our theories imply that the answer is somewhere in between and back to front.

The discussion of the need for GLs in humans and other animals implies that evolution produced something used **internally** with the three core properties,

thereby supporting intelligent perception and manipulation of objects in the environment.

The use of GLs also supports the development of communicative language:

a pre-verbal child has things to communicate about

and has motives that can be served by such communication.

A different view of language development

The GL structures were not overtly communicated and did not use the grammars of later human languages.

Insofar as internal GLs are partly acquired through interaction with the environment, instead of being wholly innate, it follows that the genome of some species provides one or more GL acquisition devices (GLADS), though they are better viewed not as completely innate devices, but as self-extending mechanisms, whose self-extending capabilities are themselves extended by things derived from the environment.

When communicative uses of GLs began they would have built most naturally on the role of GLs in controlling behaviour (e.g. executing a plan), since what you do often communicates your intentions.

That probably involved many evolutionary steps that will be hard to find evidence for.

Only later would new pressures cause vocal information structures to take over.

The additional constraints of that impoverished medium (compared with the earlier gestural GL) may have driven both human languages and the brain mechanisms down narrow channels, further constraining the permitted structural variability and modes of composition.

But that's a topic for another time.

Routes from genome to behaviours

Cognitive epigenesis

The diagram shows different stages at which the environment influences processes, e.g.:

- during development of seed, egg, or embryo, and subsequent growth (i.e. it is not all controlled by DNA)
- triggering meta-competences to produce new competences or new meta-competences (e.g. after previous competences have produced exploratory and learning processes)
- during the triggering and deployment of the competences to produce behaviours

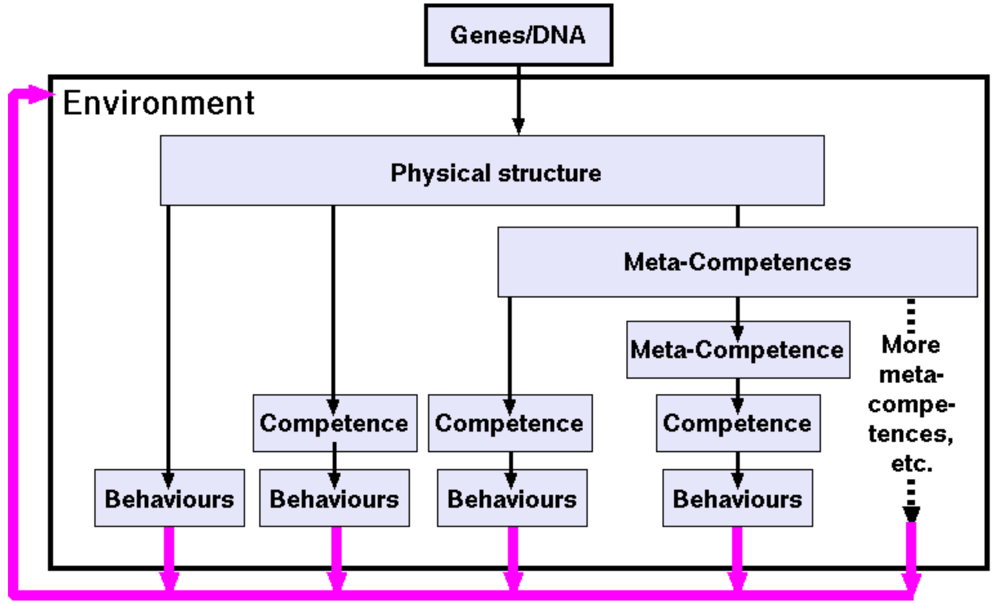

Insofar as the behaviours influence the environment there can be complex developmental feedback loops. Competences and behaviours further to the right may use several 'simpler' competences and behaviours developed on the left. Diagram is from the IJUC paper with Jackie Chappell. Chris Miall helped with the diagram.

The construction of some competences should be construed as an ongoing process, with repeated activation of the meta-competence over time.

These schematic specifications may have different sorts of instantiations in different parts of a multi-functional architecture, e.g. in reactive and deliberative components.
In reactive components many (but not all) of the processes will involve continuous control.
In deliberative and some meta-management components much will be discrete.

Cascaded development and learning

If learning has to go through partially-ordered competences, where each competence builds on what has been built in previous stages, and that involves building new layers of brain mechanism, then that might explain why each new GL extension can only happen at a certain stage of development.

A particular GL cannot be added too early because it needs prior resources to provide

- the representing structures,
- the ability to manipulate them, and
- the contents that they represent.

It can't happen too late because lots of other things are normally delayed until the appropriate GL has got going, and if that doesn't happen they may develop anyway, but in inferior forms and they cannot be disassembled and reassembled later.

There may also be facts related to the sequence in which portions of brains develop. (e.g. myelinization??)

But the stages may be only partially ordered – allowing different learners to traverse different developmental trajectories in a network of possible trajectories.
(Compare Waddington's epigenetic landscape.)

All this still needs to be made a lot more precise – preferably in working models.

The evolutionary heritage of gestural GLs

It has often been remarked that at least three remarkable facts about humans suggest that we still retain brain mechanisms that originally evolved in connection with external gestural GLs.

- It is hard for people to talk without gesturing, often highly creatively, even when they are talking on the phone to people who cannot possibly see the gestures and who do not need them – as shown by the usefulness of telephones.
- Some children with Down's syndrome find it easier to learn a sign language than to learn to talk normally.
- Nicaraguan deaf children rapidly developed a new sophisticated sign language which they used very fluently and which their teacher was unable to learn.

Moreover, if all human children develop and use rich internal GLs before they learn to talk (orally or by signing), then what we used to think of language learning should be thought of as language extension since they already have rich linguistic (GL) capabilities which they extend for communicative purposes.

Nicaraguan children showed that that should be thought of as a collaborative, creative process of developing a means to communicate, rather than a process of doing data-mining or induction on information collected from the environment.

In most cases the child learners are a small minority, and politically weak, so language creation looks deceptively like language learning.

Methodological warning

- Much of what I say is still largely descriptive and hypothetical:

  I cannot (yet) demonstrate working computer models of the mechanisms discussed.

  Although I have a lot of experience of building computer models and can see how parts of what I am talking about could be implemented, there are still many gaps in the current state of the art (especially robot perception of 3-D structures and processes.)
- Until a theory can be expressed with sufficient precision to guide the construction of a working system it should always be regarded as suspect: for example, you cannot easily tell what assumptions you are making and whether they are seriously flawed.
- Implementation on computers is a "proof of concept", but still leaves open the question whether the mechanisms proposed can be implemented on biological information processing machinery.
- Unfortunately I don't think we understand more than a subset of types of biological information processing mechanisms, and we cannot yet build convincing platforms that can be used as a basis for implementing theories of the kind discussed here.
- Therefore, much of this is still speculative and not subject to immediate testing.
- This could be the start of a progressive or a degenerating research programme.

  Deciding which it is can take many years of development of a research programme:

See: Imre Lakatos,

The methodology of scientific research programmes, in *Philosophical papers, volume I*,
Eds. J. Worrall & G. Currie, CUP, 1980, (`http://www.answers.com/topic/imre-lakatos`)

More on the implementability requirement

Remember the warning about unimplemented models.

- All three of the core properties (structural variability, compositional semantics and manipulability) have implications for mechanisms, and architectures in which they can be combined.
- Some mechanisms cannot support structural variability,
  e.g. many of those that deal only with vectors of numerical values.
- Some mechanisms have no use for compositional semantics because they do not do any significant interpretation of the structures they operate on.
- The three core properties should be regarded as properties of virtual machines implemented in brains not as properties of physical mechanisms:

  E.g. your brain does not get rewired when you see a new scene, make an inference, create and compare a set of plans, compose a poem in your head, ..., but a virtual network might be rewired.
  For a short introduction to virtual machines and supervenience see
  `http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#bielefeld`
- Current computer-based models support only a small subset of the types of manipulability discussed here.

  Current biologically-inspired mechanisms (e.g. existing neural models) are so far inadequate for these purposes.
- Perhaps animal brains run virtual machines no modellers have thought of yet?

Many unsolved problems

These slides scratch the surface of many very deep and difficult problems.

In particular, I have ignored the fact that very little is understood about what the varied functions of visual perception are, how they work, and what forms of representation (GLs) they use.

It does not seem to me that anyone in psychology, neuroscience, or AI/Robotics is near answering the questions.

Some of the points were made at the BBSRC-funded Conference in Birmingham May-June 2007 on

Closing the gap between neurophysiology and behaviour: A computational modelling approach

See this presentation:

`http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#compmod07`
Architectural and representational requirements for seeing processes and affordances.

Background to this presentation

The slides are partly based on this BBS paper (in Press), which introduced the term 'g-language', now 'GL'.

Aaron Sloman and Jackie Chappell (2007).

'Computational Cognitive Epigenetics', in *Behavioral and Brain Sciences Journal*, 30(4).
http://www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0703
Commentary on: Eva Jablonka, Marion J. Lamb,

Evolution in Four Dimensions:
Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life (MIT Press, 2005)

Precis of book: http://www.bbsonline.org/Preprints/Jablonka-10132006/Referees/

There are several other closely related joint papers by Chappell and Sloman (2005 to 2007) on the CoSy project web site:
http://www.cs.bham.ac.uk/research/projects/cosy/papers/

We also have some slide presentations on kinds of causal reasoning in animals and robots prepared for WONAC
(Workshop on Natural and Artificial Cognition), Oxford 2007, here:
http://www.cs.bham.ac.uk/research/projects/cogaff/talks/wonac

---

Much earlier, less developed, versions of some of the ideas were in these two papers, both now online.

Sloman71
http://www.cs.bham.ac.uk/research/cogaff/04.html#200407
Interactions between philosophy and AI: The role of intuition and non-logical reasoning in intelligence,
Proc 2nd IJCAI London, pp. 209–226, Reprinted in *Artificial Intelligence Journal* 1971.

Describes a distinction between “Fregean” and “analogical” forms of representation, claiming that both can be used for reasoning, planning, and proofs.

Sloman79
http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#43
The primacy of non-communicative language,
in *The analysis of Meaning: Informatics 5 Proceedings ASLIB/BCS Conference*, Oxford, March 1979, Eds. M. MacCafferty and K. Gray, Aslib, London, pp. 1–15,

Videos and demos shown in the lecture

Videos of Betty, the New Caledonian Crow making hooks are available online, e.g.

`http://users.ox.ac.uk/~kgroup/tools/introduction.shtml`

`http://news.bbc.co.uk/1/hi/sci/tech/2178920.stm`

The simulated sheepdog, uses a library in the Pop11 Simagent toolkit

`http://www.cs.bham.ac.uk/research/projects/poplog/figs/simagent`

See demo 5

The SHRDLU "language understanding" program is in Demo 12

`http://www.cs.bham.ac.uk/research/projects/poplog/packages/simagent.html`

The SimAgent toolkit

The child pushing a broom is one of the videos here

`http://www.cs.bham.ac.uk/research/projects/cosy/conferences/mofm-paris-07/sloman/`

Follow the link to videos.

The 'SHRDLU' video showing the 2-D simulated robot analysing sentences and moving blocks around is part of the Poplog RCLIB package.

In Ved or XVed you can do

```
    uses rclib
    uses rc_blocks
  then
    ENTER blocks
  At the prompt, type
    help
  to get instructions
```

More references (still incomplete)

For people who are not familiar with the story of the Nicaraguan deaf children, there are various summaries on the web including

- Brief history and some links `http://www.signwriting.org/nicaragua/nicaragua.html`
- PBS documentary including video `http://www.pbs.org/wgbh/evolution/library/07/2/l_072_04.html`
- BBC summary `http://news.bbc.co.uk/2/hi/science/nature/3662928.stm`
- Wikipedia summary `http://en.wikipedia.org/wiki/Nicaraguan_Sign_Language`

Bruce Bridgeman has a theory that overlaps significantly with the one presented here:

On the Evolution of Consciousness and Language. Psycoloquy: 3(15) Consciousness (1) (1992) `http://www.cogsci.ecs.soton.ac.uk/cgi/psyc/newpsy?consciousness.1`

References involving mirror neurons and the gestural theory of evolution of language

Fadiga L., Craighero L. Cues on the origin of language. From electrophysiological data on mirror neurons and motor representations, in In S. Breten (Ed.), *On Being Moved: From mirror neurons to empathy.* Amsterdam, John Benjamins. 2007

`http://web.unife.it/progetti/neurolab/pdf/2007_1_Fadiga-Craighero_chapter6.pdf`