

# INCOMPLETE DRAFT NOTES

(Last revised February 25, 2005)

## Towards an ontology for factual information and factual questions for a playful robot

Aaron Sloman

School of Computer Science, University of Birmingham

Birmingham, B15 2TT, UK

[http://www.cs.bham.ac.uk/~axs\\*](http://www.cs.bham.ac.uk/~axs*)

February 25, 2005

### Contents

<b>1</b>	<b>– Introduction</b>	<b>2</b>
1.1	– Questions for thinkers as well as communicators . . . . .	3
1.2	– Varieties of non-information-seeking types of questions . . . . .	4
<b>2</b>	<b>– Definitions:</b>	<b>5</b>
2.0.1	– Key ideas: . . . . .	5
2.1	– Non-factual questions . . . . .	6
2.2	– Varieties of answers . . . . .	6
2.3	– Questions and propositions in non-linguistic (pre-linguistic) information users	7
2.4	– Types of question structures and answer structures . . . . .	8
<b>3</b>	<b>– QUESTION FORMS</b>	<b>9</b>
3.0.1	– Yes-no questions: Proposition and its negation . . . . .	9
3.0.2	– Derived questions: operations on propositions with gaps . . . . .	9
3.0.3	– Some common question forms . . . . .	10
3.0.4	– Further ways of deriving questions from propositions after creating gaps . . . . .	10
<b>4</b>	<b>– Information gaps in AI systems</b>	<b>11</b>

---

\*With much help from the CoSy Project Team (See URL <http://www.cognitivesystems.org/>)  
in Birmingham (See URL <http://www.cs.bham.ac.uk/research/projects/cosy/>) and elsewhere.

<b>5</b>	<b>– PROPOSITIONAL COMPONENTS FOR A PHYSICALLY EMBEDDED INFORMATION USER</b>	<b>12</b>
5.1	– Types of entities that can be referred to	12
5.1.1	– Physical object types	12
5.1.2	– ‘Stuff’ types	14
5.1.3	– Location types	15
5.1.4	– Attributes (of objects, locations, events, etc.)	16
5.1.5	– Object relations	18
5.1.6	– Intrinsic relations: Relations based on attribute values	19
5.1.7	– Relations based on shape	20
5.1.8	– Extrinsic relations: Spatio-temporal relations	21
<b>6</b>	<b>– mess to be fixed follows</b>	<b>24</b>
<b>7</b>	<b>– What about a non-linguistic (pre-linguistic) agent?</b>	<b>25</b>
<b>8</b>	<b>– Nature Nurture Issues</b>	<b>25</b>
<b>9</b>	<b>– Some References</b>	<b>26</b>

---

## 1 – Introduction

This paper started as a presentation of some ideas about the variety of types of questions that can elicit information or at least express a need for information, or specify a gap in information available to some information-user. This was not a question about communication between agents (e.g. asking questions) but about processing *within* an agent. An agent deciding what to do might discover that it needs some information, e.g. about the current environment, about whether preconditions for some action are satisfied, about what the consequences of the action could be, etc. After some exploration of these ideas I realised that something had to be said about the space of possible facts within which questions (information gaps) could arise. So the paper changed direction, and now has two strands: how to form questions expressing information gaps, and what sorts of information could have gaps. That’s a huge topic, but it is constrained here by the needs of the PlayMate robot scenario<sup>1</sup> within the CoSy project<sup>2</sup>. This scenario is based on a robot able to perceive, manipulate and talk about 3-D objects on a table top, including eventually assembling complex objects from components (hopefully going beyond Freddy<sup>3</sup> the Edinburgh robot which was able to assemble a few toy objects about 30 years ago, e.g. making a toy wooden car from a body with holes, two axles and four wheels).

Although the PlayMate scenario ultimately has to include linguistic competence, it is likely that such competence will build on more fundamental and general forms of competence shared by pre-verbal human children and other animals that manipulate objects such as New Caledonian crows.<sup>4</sup> So this document is first draft incomplete investigation into some of the kinds of information and information gaps that a pre-linguistic robot with perceptual and manipulative skills may need to cope with, especially one which, unlike Freddy, knows what it is doing and why.

<sup>1</sup>See URL <http://www.cs.bham.ac.uk/research/projects/cosy/PlayMate-start.html>

<sup>2</sup>See URL <http://www.cognitivesystems.org>

<sup>3</sup>See URL <http://www.aiai.ed.ac.uk/project/freddy/>

<sup>4</sup>See URL <http://users.ox.ac.uk/.kgroup/tools/tools.main.shtml>

Concerning questions, the key idea was that every question content (every specification of a need for factual information) can be expressed by starting with a true/false propositional form then creating zero or more gaps in that form by removing components and then applying various question-forming operators to the resulting gappy structure, the simplest such operator being ‘Is it true that...’, which requires no gaps. This point is independent of exactly how the logical forms are represented, as long as they are ‘Fregean’ representations, in the sense of [Sloman 1971]<sup>5</sup> (also available in HTML<sup>6</sup>). There are many different syntactic forms with the same semantics that can be used in intelligent systems. The comments about question formation apply to all of them. In systems using hybrid representations, e.g. a mixture of logical forms and map-like structures (labelled ‘analogical’ in [Sloman 1971]), some generalisation of the method for deriving question forms from propositional forms would be required, but that is a topic for another occasion.

In general, several questions can be derived from any proposition, by creating gaps in different places, and by applying different question-forming operators to the gaps.

Obviously, developing the ontology for questions (information-gaps) presupposes a prior ontology for the kinds of propositions from which the questions could be derived. So a substantial part of the ontology presented here is common to both factual propositions and factual questions.

What follows is not a totally general ontology. The survey is restricted to types of factual information that may be discovered, or needed, or used, by an animal or machine with a physical body, located in space, enduring over time, and perceiving things, performing actions and thinking about objects, events, and processes (including actions) that might be relevant to an animal or robot manipulating and thinking or talking about 3-D objects on a table top, as in the PlayMate<sup>7</sup> scenario of the CoSy project.

The survey is concerned mainly with factual information and gaps in factual information, though some of the ideas might be applicable to other forms of information (e.g. control information, as expressed in imperatives, discussed briefly below).

## 1.1 – Questions for thinkers as well as communicators

Some suggestions are made about possible forms of propositions and forms of derived questions that can arise both in dialogues and in thinking. I believe that the ideas are widely applicable in the design not only of dialogues between agents (as in PlayMate) but also agents that have dialogues with themselves, for instance when planning or reasoning, or wondering what to do, or why something happened. In the latter cases the forms of question and answer need not necessarily be translated, or even translatable, into external linguistic forms. Some of them could be used by pre-verbal children or other animals. Exactly which subsets can be used by which types of individuals is an empirical question which this document does not address, though it ends with some brief comments on how this relates to pre-verbal cognitive processes in a predominantly altricial species. The altricial-precocial spectrum in organisms, and its possible application to robots, is discussed in another draft incomplete paper.<sup>8</sup>

I don’t claim that any of this is original: the idea about questions arising out of gaps is based on something dimly remembered from my youth as a philosophy student in Oxford 1957-1962, but I cannot be more specific as to where the ideas came from, except that it builds heavily on the

---

<sup>5</sup>See URL <http://www.cs.bham.ac.uk/research/cogaff/sloman-analogical-1971.pdf>

<sup>6</sup>See URL <http://www.cs.bham.ac.uk/research/cogaff/sloman-analogical-1971>

<sup>7</sup>See URL <http://www.cs.bham.ac.uk/research/projects/cosy/PlayMate-start.html>

<sup>8</sup>See URL <http://www.cs.bham.ac.uk/research/cogaff/altricial-precocial.pdf>

logical investigations of Gottlob Frege (see this<sup>9</sup> and this<sup>10</sup> among many other overviews available) about a hundred years ago. The powerful idea that gaps in propositions can be filled in many ways, including the use of quantifiers was, I believe, independently discovered by G.Frege and C.S.Peirce, breaking away from the previously held (Aristotelian?) view that in a sentence like ‘All As are Bs’ the subject of the sentence is some entity referred to as ‘All As’. On Frege’s analysis the sentence is to be understood as asserting that any gap-filler that makes ‘... is an A’ true will also make ‘... is a B’ true. (There are people, e.g. the philosopher Fred Sommers, who believe that Aristotle was right and Frege wrong, but I shall ignore that debate.)

Logicians will recognize that much of what is said about gaps here can be expressed using Lambda Calculus, the formal notation Alonzo Church developed inspired by the rather more clumsy formalisms invented by Frege. I have decided to keep this presentation informal, instead of using lambda expressions or a similar notation.

Looking for related work on the internet, I found a recent PhD thesis which develops ideas about question types very close to the ideas about questions presented here:

Debra Thomas Burhans (2002)

*A Question Answering Interpretation of Resolution Refutation*<sup>11</sup>

PhD Thesis, State University of New York at Buffalo Department of Computer Science and Engineering.

Before moving on to the task of specifying types of information relevant to our project I’ll make a few remarks about restrictions on the types of questions that are relevant at least to this document.

## 1.2 – Varieties of non-information-seeking types of questions

It is clear (as Erica Melis reminded me when I first mentioned the ideas sketched here) that questions can have many different pragmatic functions, apart from eliciting information, for example:

- testing someone’s knowledge
- getting someone to think about a problem (Socratic questions)
- teasing someone
- challenging someone
- being ironic or sarcastic
- drawing attention to something
- being rhetorical (in various ways)
- making a request (i.e. trying to get something done)
- asking for a definition (I constantly have to say to people: ‘what do you mean by "emotion"?’ in response to some comment or question about emotions.)

---

<sup>9</sup>See URL <http://plato.stanford.edu/entries/frege>

<sup>10</sup>See URL <http://www-groups.dcs.st-and.ac.uk/history/Mathematicians/Frege.html>

<sup>11</sup>See URL <http://www.cse.buffalo.edu/.burhans/debra.ps>

For more examples of types of question functions see this taxonomy<sup>12</sup> of question types, and the taxonomy often referred to as "Bloom's Taxonomy", e.g. here<sup>13</sup> and here<sup>14</sup>.

Most of those varieties of question will be ignored here. We are concerned only with the 'core' case: use of questions to specify missing but required factual information. I consider *only* the semantic content of questions and answers, ignoring all pragmatic and performative issues, and leaving open the syntax to be used in a working system. Any syntax used here is illustrative only.

## 2 – Definitions:

- A *question* as discussed here, is a specification of a factual information gap in an information user.
- An *answer to a question* is an information structure that fills, or reduces, the gap (or in some cases, specifies how the gap could be filled).

This leaves completely open what kind of medium is used to express either questions or answers, and what syntactic forms are used within any particular medium. It also leaves open what kinds of use are made of the information, how gaps are detected, how they are filled, how genetic information, learning or development can produce any of the competence required, how brains operate, etc.

### 2.0.1 – Key ideas:

The key ideas are:

1. There is a basic type of factual question, a Yes/No question, which refers to a proposition that is capable of being true or false and requests the information whether the proposition is true or false,
2. There are ways of generating other types of factual question by creating one or more gaps in the proposition and specifying requests for information about ways of filling the gaps so as to make the proposition true (an example would be 'Which individuals satisfy: X is P?')
3. Many, if not all, factual information gaps in information users involve needs that can be expressed as factual questions of the above kinds. (Note that two kinds of gaps are involved: information gaps in intelligent systems and gaps in propositions produced by removing components of those propositions.)

E.g. if a proposition asserts 'A did X' then the question form 'Who did X?' is a natural language expression of the logical question form requesting information that can fill the gap in '... did X?' so as to produce a true proposition. Some questions are related to a single gap, others to two or more gaps, as illustrated below. Some questions refer to a gap on the assumption that there is a unique gap-filler that would make the proposition true ('Who is the murderer?') whereas others allow the answer to specify any number of fillers that could produce true propositions, or to specify the number without listing them, e.g. 'Who saw the event?' (answer: Tom, Trudy, Tim, ....) or 'How many people saw the event?' (answer: none, or five, or...).

---

<sup>12</sup>See URL <http://www.usingenglish.com/glossary/question-types.html>

<sup>13</sup>See URL <http://honolulu.hawaii.edu/intranet/committees/FacDevCom/guidebk/teachtip/questype.htm>

<sup>14</sup>See URL <http://www.officeport.com/edu/bloomq.htm>

Some questions specify two or more gaps and ask about individuals standing in some relationship (e.g. ‘Who is taller than whom?’).

Since propositions can take many forms, with unbounded complexity, the variety of types of gaps and gap-fillers is essentially unbounded.

Sometimes the specification of information needed requires two gaps to be filled by the same thing. E.g. ‘Who shot himself?’ involves two linked gaps in the propositional form: ‘... shot ...’. Such linking is normally expressed in formal languages by the use of variables e.g. ‘x shot x’ or ‘Shot(x, x)’. In natural languages other devices are used e.g. use of pronouns, words like ‘himself’, ‘itself’, ‘themselves’, etc. (Failure to understand what’s going on here leads many people to believe that we all have spooky entities called ‘selves’ inside us.) The natural language devices are irrelevant to the aims of this paper, which is concerned only with cognitive structures and processes that may be common to users of many different languages and to some intelligent systems that do not use what we would call an external language.

## 2.1 – Non-factual questions

A more general definition of ‘question’ would include *control* information gaps, as expressed in questions like the following, whose answers could be imperatives rather than factual statements:

- What should I do?
- Is it better to do X or Y?

For now these are not considered, though some of what is said here would also be applicable to such cases. E.g. some of the ways in which different question forms can be generated from a single factual proposition by inserting gaps and applying appropriate operators to the proposition with gaps, would also be applicable to ways in which different control questions can be generated by inserting various kinds of gaps in an instruction, or imperative, of the form ‘A, do X?’ (e.g. talking to myself: ‘Me, do X’). Questions like ‘What should I do?’ could be interpreted (at least sometimes) as requesting the gap to be filled in the self-directed imperative ‘Me, do ...’. The difference is that instead of depending on the notion of a true proposition, control questions depend on the notion of an imperative or instruction being accepted or acceptable. There are many complications that arise from the fact that two imperatives that are independently acceptable can be inconsistent in the sense that they cannot both be acted on whereas two true propositions cannot be inconsistent. Another complication is that often a control question is not about what to do but which of two options is better<sup>15</sup>, a topic with many complications that we shall ignore.

## 2.2 – Varieties of answers

Here we discuss only answers that provide information to fill gaps specified in factual questions. However, as indicated above, questions can have many different functions, and so can answers. In dialogues between information users there are things that we would call answers that perform other functions besides giving answers, such as apologising for not knowing the answer or being too busy to respond, or refusing to give the information, or challenging some aspect of the question – e.g. its relevance to some shared goal context, its assumptions (e.g. that there is a unique gap filler), its politeness, the right of the asker to hear the answer, the availability of information that provides an

---

<sup>15</sup>See URL <http://www.cs.bham.ac.uk/research/cogaff/sloman.better.html>

answer, etc. None of those ‘pragmatic’ answers will be discussed here. The scope of this document is very limited.

I believe that every question of the types surveyed below can be expressed in English, though some of the questions have very complex structures and therefore expressing them accurately in English may be difficult, or at least clumsy. E.g. consider questions derived from various ways of creating gaps in this proposition:

*A1, B1 and C1 are in the same order on line L1 as A2, B2 and C2 are in on line L2.*

(some English speakers would omit the second ‘in’.) A possible question derived from that would be:

*Which three things in the room are on the same order on line L2 as A1, B1, C1 are in on which line?*

Linguists may be interested in debating whether this really is English or not, and if not why not, whereas for my purposes that is irrelevant since, like many other non-English sentences, the meaning is clear, as in ‘Me go home after me eat’, and many of the things young children and non-native speakers say.

The answer to the above question may specify a complex set of alternative ways of producing true propositions by referring to objects in the room to fill the gaps formed by removing A2, B2, C2 and L1 from the original proposition.

Not only generating, but also understanding an accurate expression of a complex question may also be difficult. Often there is a formal mathematical way of expressing the question that mathematicians would find clearer or more succinct, though non-mathematicians may find it incomprehensible. One of the problems of learning to do mathematics is learning ever more sophisticated question-forming and answer-forming techniques, though I have no idea how many mathematics teachers understand this. It is very likely that some of the more complex forms of propositions and questions discussed here could never be understood by young children or other animals. Likewise some of them may be beyond the grasp of our robot. Nevertheless the theory subsumes them.

### **2.3 – Questions and propositions in non-linguistic (pre-linguistic) information users**

As far as I know every human language is capable of expressing all the forms of propositions and factual information-seeking questions mentioned here, but it is theoretically possible that some natural languages lack the syntactic expressiveness required for some forms of questions, just as young children do. Whether some natural language lack the expressive power to do everything described here is an empirical issue, that is not relevant here. The study of the precise forms of syntax used in a particular language to express one of these question types is part of empirical linguistics, and I have nothing to say about that. However, I conjecture that every information-seeking question that is expressible in any language is an example of the sort of schema for generating questions presented below, and every proposition that is expressible in any language is an example of one of the propositional forms we are discussing. (Though the actual list in this document is incomplete: the paper is unfinished.)

But there are more things to be thought and questions to be asked that we can express linguistically. It may be the case that some of the contents of factual questions are sometimes referred to by non-linguistic means, e.g. using kinds of attention-focusing mechanisms (‘virtual fingers’

(Pylyshyn's FINSTs (1989))<sup>16</sup> ) to specify entities or other things referred to for oneself (even if one is a robot). The work of Trehub(1991) is also relevant.

Public, external, languages used for communication between agents constitute only a subset of the forms of representation required by intelligent agents. Many kinds of information structures are used in various parts of a complex agent architecture, to store information of different kinds, referring to external or internal entities, states or processes, at various levels of abstraction, stored for different time scales, used for different purposes. There is no reason to believe that all such semantic contents are expressible in external languages, such as human languages. For some examples, see <http://www.cs.bham.ac.uk/research/cogaff/challenge.pdf><sup>17</sup>

So there may be types of questions and types of answers that are not accurately expressible in any language but are important for the cognitive functioning of some non-linguistic animals and pre-linguistic children. It is possible that this kind of pre-linguistic mental function provides part of the infrastructure for natural language as used by humans and therefore needs to be understood in order to develop an accurate theory of language understanding by humans.

Pre-linguistic question formation and answer-seeking or answer-providing processes (within an individual) may also be important for some robots, in particular robots that are able to perceive and act intelligently while lacking the ability to communicate in any human-like language.

## 2.4 – Types of question structures and answer structures

We turn now to a first draft informal elaboration of the points made above about ways of deriving information-seeking questions from gaps in propositions. This requires a specification of forms of propositions, on which more will be said in a later section. For now I am going to assume that all the propositions required are of a type that could be expressed in first or higher order predicate calculus, with certain modal operators added (e.g. to express causes, purposes, etc.) Instead of a general formalisation, I present only examples, which should suffice to explain the structure of the ontology of questions being presented here. Remember that arguments of predicates and functions need not be denoted by words or linguistic expression: they may be objects of current attention identified by their relationship to what is currently seen or thought about.

A question specifies a request for information or an information gap. The basic form of answer is "yes" or "no" answering a question that asks what the truth-value of a proposition is. More complex answers correspond to questions that are derived from a proposition in various ways, described below. A question has a structure, which may be more or less complex, as explained below.

Answers will be information structures (or structured information items) capable of filling such information gaps, or specifications for sets of information structures that can fill gaps. In many cases the answer to a question can be given in alternative forms: e.g. if the question requests an example of something, different examples may be given in answers. However, in general the form of a question constrains the forms of appropriate answers. For example, if the question is "Is it raining?", the answer "thirty seven" is inappropriate, and if the question is "How many people are in the room?" the answer "yes" is inappropriate. If the question is "What happened next?" only a complete proposition is an appropriate answer (though the proposition may be identified by a referring device, e.g. "What I just saw happened next").

This is not a document about natural language syntax, but about kinds of factual information that can be needed by an information processing system.

---

<sup>16</sup>See URL <http://rucss.rutgers.edu/ftp/pub/papers/cognit89.pdf>

<sup>17</sup>See URL <http://www.cs.bham.ac.uk/research/cogaff/challenge.pdf>

### 3 – QUESTION FORMS

This section describes, rather informally, and incompletely, a principled collection of transformations of propositions in order to specify information gaps. A question, from this standpoint is just an expression of an information gap, and an answer is an expression of the information that fills the gap.

#### 3.0.1 – Yes-no questions: Proposition and its negation

Those are the fundamental sorts of questions, merely asking about the truth or falsity of a proposition.

#### 3.0.2 – Derived questions: operations on propositions with gaps

There are many question forms that can be generated by producing one or more gaps in a proposition. The gaps can be of many kinds: object, attribute-type, attribute value, relation, action, purpose, manner, instrument, cause, location (different sorts), time (different sorts),

Some of these occur so often that we have special words to identify the type of gap to be filled:

who, what, which, where, when, how, why, ....

Sometimes there is more than one gap

‘Who broke what?’

‘Who murdered whom when?’

Different question forms relate to whether a unique gap filler (or a unique tuple of fillers) is requested, or a the set of all, or the number of the set of all, or some statistical property of the set (are most X’s P?), whether referents are identical, or whether an identified entity fills some gap, etc. etc. For example:

‘Is fred the person who....?’

‘Did fred solve it in the same way as Mary?’

‘How many people solved it in the same way as Mary?’

‘How many things did Fred do in the same way as Mary?’

‘Is Fred cleverer, bigger than Joe, ...’

It is worth noting that from a logical point of view many things are clearly interpretable whose most ‘logical’ expression in natural language would be regarded as ill-formed or bad style, such as

‘How many people did fred solve it in the same way as?’ ‘Fred solved it in what as Mary?’

Possible answers to the latter include: ‘in the same way’, ‘in the same room’, ‘in the same time’, etc.

From the present point of view the best way to express these questions in natural language, or even whether they can be expressed in a particular natural language, is of no consequence.

### 3.0.3 – Some common question forms

- **WH-singular-unary:** Proposition with one gap, requests unique filler:

Who married Mary?

- **WH-singular-binary:** Proposition with two gaps, requests unique filler-pair

Who married whom? (also ternary, etc. ‘Who gave what to whom?’)

- **WH-plural-unary:** allows plural answer:

Who was at the party?

Answer: Fred, Mary, the man in the moon, and my uncle.

- **WH-plural-binary:**

Who talked to whom at the party?

Answer: Fred talked to Mary, Sue talked to Joe and to Tom, etc.

Note that for every type of proposition P more complex propositions can be formed by embedding P in a larger context, and then creating a gap in the context from which questions can be derived. Examples include:

- P because Q
- Q because P
- P in order to Q  
where P describes an action that has been done, is being done, or will be done.
- Q in order to bring it about that P  
where P describes a state of affairs

Other examples specify means by which something is done, route taken, style or manner, or other qualifiers of P.

For each question form and each situation determining an answer there may be varied verbal answer forms possible (e.g. different ways of identifying an object) with no context-independent means of saying which answer or answer form is best.

Sometimes selecting between them depends on facts about the dialogue context other than whatever is referred to in the question (e.g. the answerer may be impatient with the questioner, or may want to draw attention to a mistaken presupposition of the questioner, or .....)

### 3.0.4 – Further ways of deriving questions from propositions after creating gaps

To be continued .... there is a lot of existing work on this, including work in linguistics, NLP, Burhans (2002) and others. Some sample ways of deriving questions by filling gaps

- How many ...?
- Are there more than N ...?
- Is the number of ... the same as the number ...?

- If P will Q ??
- Why P ? — possible answer: P because Q
- Why Q ? — possible answer: P caused Q
- Why P ? — possible answer: P in order to Q
- How P ? — possible answer: P using X
- etc.

The above set of points about forming questions by applying various kinds of operators to gaps in more or less complex propositional forms generates a VERY large variety of types of questions (And answers).

Since for any proposition different gaps and different numbers of gaps can be created in it, and for each gap or each combination of gaps different questions can be formed using different question-generating operators, the variety of forms of questions (or more precisely information requirement specifications) must be greater than the variety of types of propositions.

## 4 – Information gaps in AI systems

A more detailed theory will need to specify features of cognitive architectures that support the discovery by information users that they need some information that they do not have, and the mechanisms that allow the missing information to be identified with sufficient precision to generate action to remedy the deficiency, which may involve

- re-focusing attention
- re-processing current sensory information in some new way
- accessing and manipulating previously acquired information
- performing some external action to get more information (e.g. looking in a different direction)
- communicating with another agent in some way, whether non-verbally or verbally (e.g. a child may say ‘What’s that?’, and turn an adult’s head to look in the right direction.)

One very obvious way in which an AI system can encounter an information gap occurs if it is attempting to form a plan to achieve some goal by combining action operators for which preconditions and consequences are known. The system may discover that some potentially relevant action has a precondition, and then on attempting to decide whether the precondition is satisfied or not it may discover that it does not have the information. (This is unlike a typical implementation of prolog, which assumes that everything it cannot prove is false.) Another example would be a visual system attempting to identify some object or to perceive the relationship between two objects (e.g. whether they are touching), which then discovers that it cannot decide because some or all of what it is looking at is hidden by a large object. Having identified missing information an intelligent agent with sufficient meta-level knowledge may be able to take steps to fill the information gap. In more primitive animals or robots a purely reactive mechanism might generate information-acquiring actions without any self-knowledge about what is going on.

SOAR is an example of a problem solver that can discover that in order to solve one problem it may have to solve another which includes discovering something. Many such mechanisms are theoretically possible and we shall not attempt a survey here.

We can expect different subsets of the theoretically possible variety of question forms and information-obtaining processes to occur in different species, in different robots, and perhaps in the same individual at different stages of development. Where linguistic forms are used, different languages will encode them in different ways.

For now we leave the topic of question formation and the nature and uses of information gaps, and instead turn to the varieties of types of information that a robot might use.

## **5 – PROPOSITIONAL COMPONENTS FOR A PHYSICALLY EMBEDDED INFORMATION USER**

There may not be any point trying to cover all possible propositional forms and question forms in the CoSy project. E.g. the ones a two or three year old child can understand and generate may be plenty, at least for the next few years. It is probably wise to start with mechanisms and forms of representation that do not presuppose the ability to use an external language, since those evolved first and are present first in young children, so it is very likely that they are used by the language using mechanisms. In any case it seems that the variety of pre-linguistically comprehensible propositions is very large, and worth understanding in order to investigate how robots might work, independently of what is true of humans and other animals.

An animal, child, or robot may have a complex architecture with many different components using information of different sorts in performing different tasks, often concurrently, e.g. controlling eye vergence, controlling posture, controlling breathing, controlling direction of gaze, causing linguistic input to be parsed and interpreted, controlling the digestion of food, the insertion of hormones into the blood, the pumping of blood, etc. Different subsystems will use various specialised representations for their tasks. Even information from the same sensory source, e.g. visual information, may be transmitted to different subsystems that use it in different ways to derive different information, represented in different formats, for instance in controlling current actions and in planning future ones, or predicting what will happen next in the environment. For now we ignore all those differences and focus only on subsystems that can use factual information that might usefully be expressed, at least partly, in a propositional form (including non-linguistic internal indexical referring devices, as mentioned previously).

We consider, as an example, kinds of propositional structures that might be relevant to a robot perceiving and manipulating objects on a table top. (Simpler versions of the robot will cope with only a small subset of ontology presented here.)

### **5.1 – Types of entities that can be referred to**

#### **5.1.1 – Physical object types**

For now ‘object’ will not be defined, though most of the things called objects will be enduring, spatially-bounded, possibly moving, entities, which may have fixed or changing attributes. A more general notion of object is definable as ‘anything that can be referred to’, which amounts to the same as Quine’s notion of whatever can be a value of a variable. This would include such things as times, numbers, colours, shapes, strategies, styles, theories, proofs, explanations, problems, functions of

objects, etc. For now we focus on an ill-defined subset of spatio-temporally located objects (what many philosophers would call ‘particulars’, as opposed to universals, like shapes, colours, numbers and proofs).

Any object will be an instance of one or more (usually a whole hierarchy) of object types. Some of the object types considered here are fairly abstract, but they all have instances with a spatial location and the possibility of relationships to other objects. Examples would be types of physical object for which we have names:

- Object types categorised on the basis of function and shape: box, tray, cup, mug, saucer, lid, bowl, ball (?)
- Object types categorised only on the basis of shape: cube, sphere, cylinder, ball (?)
- Biological object types: lemon, apple, dog, ...
- Body parts: arm, hand, finger, fingernail, fingertip
- Parts of objects defined by shape, function, ....: Handle, keyhole, rim, base, ...
- Generic object parts: surface, face, edge, corner (2-D or 3-D), hole, crack,  
**NOTE:** We could also have typical parts for other things, e.g. plants, pieces of fruit, various animals, various utensils. various kinds of rooms, various kinds of buildings, various kinds of towns, etc.
- Abstract entities that exist by virtue of relations between other entities, gap, opening, mouth, passage, enclosure, corridor, ...
- various entities that can occur on or in a surface:
  - 2-D surface features: marks, texture boundaries, colour boundaries, shadows, edges of those, etc...
  - 3-D (shape-determining) surface features: furrows, indentations, protrusions, curvature extremes, saddle points, edges (where two surfaces meet in a well defined curved or straight line)
- Entities that exist relative to a viewpoint, e.g. visible portion of a surface or object, occluding edge of a surface (e.g. of a sphere viewed from a certain direction – the spherical surface itself contains no surface edges)
- thing (catch all ? Or meta-level concept, like a variable)

[The above is to be extended. Some more relevant ideas about types of spatial entities can be found in CYC<sup>18</sup>.

Some half-baked thoughts about requirements are here.<sup>19</sup> ]

In a different environment there could be different object types including various plant and animal types, types of furniture, parts of buildings (walls, doors, windows), and various types of out-door objects, such as rocks, trees, clouds, lawns, roads, etc. We ignore those complications for now.

---

<sup>18</sup>See URL <http://www.cyc.com/cycdoc/vocab/spatial-vocab.html>

<sup>19</sup>See URL <http://www.cs.bham.ac.uk/.axs/misc/ontology-for-a-manipulator.txt>

There is no requirement that an animal or robot should be able to see *only* things for which it has names, or which it can recognize. It is clear that we are perfectly capable of seeing, and thinking about a complex structure that we have never seen before, but on parts of which we are capable of performing many actions (grasping, prodding, pressing, pushing, pulling, twisting, bending etc.) The ability to see named, or nameable things almost certainly rests on this more basic ability shared with many other animals that act in an environment about which they cannot talk.

A pre-linguistic animal or robot may, however, use some object types that correspond to recognition states of some internal pattern recogniser without being able to use any external labels for those types even if there are internal labels of some kind. (A draft document discussing, among other things, the evolution of internal labels in biological organisms is available at <http://www.cs.bham.ac.uk/research/cogaff/vis-affordances.pdf><sup>20</sup>)

### 5.1.2 – ‘Stuff’ types

The object types have well defined instances with boundaries between individuals so that they can be counted. The table-top environment might also include kinds of stuff (labelled in natural languages by ‘mass nouns’). If X is a kind of stuff (kind of material), such as water, sand, mud, bread, dough, cotton, etc., then one cannot have two Xs though one can have two pieces, lumps, pools, spoonfuls, (in some cases) squirts, or stretches of X, where the additional (chunker) noun, ‘piece’, ‘lump’, ‘stretch’. etc, is used to refer to a bounded portion of space filled with X.

For a stuff type X it is not possible to ask how many Xs there are without adding a chunker noun, but it is possible to ask how much X there is, since adding more X to a portion of X, or removing some X from it does not stop what is there being X, it merely alters the amount of X (which is the size, volume, weight, etc. of the portion of X).

In a table-top scenario some subset of the following stuff types might occur:

- Rigid Types of stuff wood, metal, plastic, glass, (stiff card?)
- Non-rigid Types of stuff
  - Dry cohesive: paper, string, cotton wool, wire, foil, plastic (e.g. film), plasticine, various plant materials (non-dry twigs, leaves),
  - Dry particulate (pourable, stirrable, spillable): sand, sugar, salt, pepper, .... (Compare piles of smaller and smaller marbles, or cubes: type boundaries may not be sharp, or may be task dependent.)
  - Liquid (e.g. pourable, stirrable, spillable): water, paint (could be more or less viscous, sticky, etc. type boundaries may not be sharp, or may be task dependent.)
  - Gaseous: Steam, smoke, wind, clouds.... maybe in another project

Note that many attributes that an object has, as described below (e.g. rigidity, hardness), will be inherited from properties of the stuff of which it is composed. Understanding this relationship may or may not be required for a particular level of child or animal competence. For particular purposes more sub-divisions may be required, e.g. between breakable and non-breakable rigid objects. But such things could come later.

Insofar as surfaces are important for the PlayMate robot, it should be noted that the perceptual and manipulable qualities of surfaces will be heavily dependent on the kind of stuff involved.

---

<sup>20</sup>See URL <http://www.cs.bham.ac.uk/research/cogaff/vis-affordances.pdf>

Picking up a portion of liquid, a piece of thread, a lump of plasticine, a wooden cube all involve different percepts and skills.

### 5.1.3 – Location types

Physical objects (including chunks of stuff) have locations: they occupy space. Locations of objects are important for many purposes, including grasping them, putting things in them, avoiding them, throwing things at them, identifying them when communicating with others, working out where to go in order to see them, etc.

In addition to objects, events and processes can also occupy or occur at a location. Some locations are not filled by the object (or event) but moved through. For instance a pea in a constantly shaken small box moves around in the volume that is enclosed by the box, and that volume moves around in a larger space as the box is shaken. If a ball is thrown the location through which it moves can be variously identified as the room through which it moves the volume above the table over which it is thrown, the ‘tube’ which forms the space enclosing all the volumes it occupies during its motion, and no doubt many more.

Besides having spatial locations events and processes can also have temporal locations. As with spatial locations there are different sorts of temporal locations.

Not all types of locations are relevant to all contexts. For example, spatial locations on a large farm may include fields, boundaries, paths, ponds, hills, hillsides, hilltops, valleys, passes, etc. none of which occur on a table-top (although sometimes toy versions do).

A significant part of the history of mathematics has been concerned with space and its properties. Euclidean geometry and more recently topology can be seen as attempts to generalise many of the sorts of notions referred to above, while both abstracting away from features of particular physical objects and processes, and also exploring limiting cases, e.g. as locations get smaller indefinitely, or routes get thinner indefinitely, or lines get longer indefinitely, or structures get more and more symmetrical indefinitely. (Straightness, circularity, perpendicularity, are all examples of symmetry in the limit.) It may be that the human ability to understand such mathematics depends on our evolutionarily older ability to understand, perceive, act on or in, and reason about the less abstract, usually physically instantiated, spatial structures, even though it is possible in principle to produce axiomatic specifications of the abstract limiting cases and reason about them using only general logical capabilities.

There has been an enormous amount of research on ways of referring to or describing spatial locations and regions of various kinds (e.g. by Tony Cohn and collaborators at Leeds University), although much of this research has been done without regard to what the information about locations is to be used for. This paper will not attempt to summarise or reinvent that work, but merely gives some simple examples of what might be useful for a robot that manipulates and converses about objects on a table top, or perhaps a young child playing with toys, or even a nest-building bird.

Spatial locations that an object can occupy or move through, or at which an event or process can occur, that might be relevant to our ‘play-and-manipulate’ context include at least the following, where the word ‘generalised’ is used to indicate that we are not restricted to the corresponding mathematical notions.

- generalised points  
(the location at which an object is, or to which it is being moved, or which it has come from in a movement: where the sort of location that is relevant will depend on the sort of object)

- generalised regions of surfaces  
(surfaces can be carved up into regions in many different ways for many different purposes, including possible areas where something could be placed, where something could move, where an individual can reach, or see, or in some contexts regions owned or controlled by someone or something).
- generalised volumes of space  
(Like generalised regions, only 3-dimensional).
- routes of various kinds  
(including 2-D and 3-D routes, both relative to fixed objects or relative to parts of moving objects – e.g. the route of a teardrop down one’s face).
- locations on or in other objects, or locations.  
These include regions of surfaces, e.g. ‘on his left cheek’, or portions of volumes, e.g. ‘in his mouth’, ‘around his left kidney’, and also parts of locations and routes, e.g. near the centre or boundary of R, where R is a region, at the beginning of R, where R is a route.

A more complete discussion would also specify different kinds of temporal locations that can be referred to in propositions, including absolute locations and relative locations,

It is worth noting that the forms of representation for locations, surfaces, routes, etc. that a mathematician might first think of are not necessarily the most appropriate for a robot or animal. In particular mathematical specifications will normally be very precise so that when what is scene is unclear or ambiguous a large variety of mathematical expressions may be needed to express all the possibilities, and subsequent reasoning or planning may have to handle a large disjunction, which can lead to combinatorial search. It may be possible to simplify thinking, reasoning, planning and control of actions, if, instead of a precise specification a less precise form of description can be used, which inherently covers a range of different possibilities (like specifying a time interval rather than a time point, or a spatial region rather than a precise location, or referring to a polyhedron rather than a cube, or tetrahedron, etc.)

CONJECTURE: the move towards less precise concepts makes it possible to avoid complex mechanisms for dealing with uncertainty, as discussed here.<sup>21</sup>

#### 5.1.4 – Attributes (of objects, locations, events, etc.)

All sorts of things can have attributes (and also relationships, discussed later). E.g. physical objects can have size, shape, colour, location (and many other attributes); locations can have extent (length, area, volume, width), shape, or location in a larger location; events can have locations, durations, start or end times, participants, causes, effects, etc.

In the sort of domain we are considering, each sort of entity has a collection of attributes which belong to different types. Each attribute type has a set of possible values, e.g. the type colour can have values red, blue, green, and possibly many others, the attribute height may have either chunked qualitative values, such as small, large, huge, or continuous measure values, e.g. 5cm, 5.1cm. 5.11 cm, etc. In some cases the values are linearly ordered (e.g. height) in others not (e.g. shape, discussed further below). When the values are ordered there may be some direction in the ordering that is naturally construed as *increasing* the value whereas the opposite direction is a *decreasing* direction.

---

<sup>21</sup>See URL <http://www.cs.bham.ac.uk/.axs/misc/ontology-for-a-manipulator.txt>

Note: insofar as objects grow, develop, learn, become ill, get damaged, etc. their attributes are not fixed. It is not just that the values can change (e.g. size getting larger, colour getting greyer), it is also possible to acquire attributes in a space that was previously totally absent in the individual. E.g. when a child has learnt to do mental arithmetic it makes sense to ask about the speed at which he multiplies two numbers, whereas at an earlier stage there is no speed because he cannot multiply. Likewise if someone has no arms you cannot ask about the extent of his reach or the strength of his grip. In our initial domain we can ignore growth, learning, development, etc., though we may need to consider objects that get damaged.

Some attributes will have possible values that form ordered sets (e.g. height) whereas others will not (e.g. shape, material, uses, or species in the case of living things). Some may be ordered in several dimensions, especially functional attributes such as usefulness for a task, which may have dimensions like ease of use, quality of result of use, difficulty in learning to use, cost of use, speed of operation, etc. (The multi-dimensionality of 'better' is discussed further here.<sup>22</sup>)

Here are examples of types of attributes that objects in a simple robot-manipulation domain might have

- colour
- height
- width
- overall size
- weight
- shape
- volume (e.g. of a container)
- hardness
- rigidity/flexibility
- smoothness
- graspability
- stability
- wetness
- stickiness
- warmth
- material
- function

---

<sup>22</sup>See URL <http://www.cs.bham.ac.uk/research/cogaff/sloman.better.html>

That is not supposed to be an exhaustive list. In fact for an agent that can act on objects it is a seriously incomplete list, for there will also be a range of important object attribute types that are not definable in terms of intrinsic features of the object (e.g. as size, weight, shape and material are) but depend on what the agent can and cannot do to the object and the consequences of doing, or trying to do those things.

In fact some of the attributes listed can be interpreted that way. For instance, instead of hardness being an objective (or intrinsic) attribute of objects (or surfaces of objects), it can be interpreted as relative to a particular agent, depending on how much resistance the surface offers to pressure that the agent can apply. Thus a surface that is hard for a weak, small agent might be soft for a much larger stronger agent. Similarly, smoothness instead of being objective may be relative to the amount and kind of resistance to motion when an agent attempts to slide one of its body parts along the surface. What smoothness amounts to for such an agent will depend on what sorts of sensors are activated by such sliding actions.

In many cases each of the possible values of an attribute can be applied as a predicate, producing a proposition that is in some sense an abbreviation of a proposition specifying both the attribute and the value. E.g. ‘The block is red and square’ could be regarded as an abbreviation for ‘The block has colour red and shape square’. The abbreviations lose no information because the name of the attribute value unambiguously (at least in some contexts) identifies the relevant attribute.

However when forming questions if we simply start from the unexpanded abbreviation and create a gap by removing the predicate, the result could be too uninformative as regards what sort of gap filler is required. Thus going from the proposition ‘the block is red’ to the question ‘what is the block?’ (or what gap-filler makes the proposition ‘the block is ...’ true?) does not specify precisely enough the intended question better expressed as ‘what colour is the block?’, which imposes much stronger limits on what counts as an appropriate answer, insofar as it is equivalent to asking what gap-filler makes true ‘The colour of the block is ...’.

Note that the ability to refer explicitly to the attribute by some sort of label (internal or external) need not occur first in development or in evolution. There could be simpler stages where the values of particular attributes are implicitly identified as being of a certain type because of how they are produced by sensory mechanisms and how they are used. Thus the output of a temperature detector fed into some sort of neural control system need not be labelled as being a value of temperature, because the output is connected only to parts of the system that make use of temperature information. So in that context, simple predication without any possibility of expansion to the attribute value form is all that can occur and it is all that is needed. Only when more sophisticated architectures evolve, or develop, performing more complex functions, is it necessary or possible to use explicit attribute labels.

(**NOTE:** Explicitness here has nothing to do with consciousness: it is merely a matter of whether some structure or process exists which has a certain sort of function in the system.)

### **5.1.5 – Object relations**

Physical objects, locations, routes and other things can have relations to all sorts of things, of the same type or different type. The relations differ in ‘arity’, depending on how many things are involved. E.g. they may be binary, ternary, or n-ary for any integer n.

A major distinction can be made between relations based on attribute values (e.g. X is taller than Y, means the value of X’s height attribute is larger than the value of Y’s height attribute), and likewise for other relations of ordering, e.g. X is between Y and Z in height, relations involving measures, e.g. X is 14 cm taller than Y, etc. The types of attribute-based relations involving attribute

A that are possible will depend on the structure of the set of possible values for attribute A, e.g. whether it is totally or partially ordered, whether it is continuous, dense, discrete, etc., whether the values are themselves structured (e.g. vectors or trees or nets), and so on. For example, one object can be the mirror image of another if the shape of the first and the shape of the second are essentially the same except for a single reflection, so that together the objects can form a symmetric structure. Another kind of relationship that can hold between X and Y depending on shape is X being isomorphic with a part of Y, e.g. if X is a sphere and Y consists of two spheres joined by a cylindrical rod.

There are other relations that are more complex and are not derivable from or definable in terms of the intrinsic properties of the objects but depend on where they are located in space and or time. We'll call those extrinsic relationships, and return to them later.

### 5.1.6 – – Intrinsic relations: Relations based on attribute values

Some of the relations involve comparisons of attribute values. There are two main sub-cases, namely equality and inequality of values, and the inequality sub-cases can be further sub-divided, depending on the structure of the set of possible values, e.g. whether it is totally ordered, whether the total ordering has a kind of asymmetry supporting a notion of more or less, whether it is a continuum of values or not, whether the values vary in different dimensions so that they form some sort of vector space, etc.

For example, if attribute A has possible values  $Av_1, Av_2, \dots$  then propositions can be expressed regarding two objects X and Y saying

- 'X is the same A as Y' (e.g. X is the same colour/height/size as Y), meaning something like: the value that fills the gap in 'The A of X is ...' and the value that fills the gap in 'the A of Y is ...' are the same thing.

Variants of this use a relation weaker than having identical values, and instead specify values that are close in the ordering, which could be expressed as 'X is like Y in respect of A', or using phrases like 'similar to' 'close to', etc.

- 'X is more A than Y', or 'X's A exceeds Y's A' (e.g. X's size, height, hardness, smoothness, exceeds Y's) meaning something like: the value that fills the gap in 'The A of X is ...' is 'bigger' (or 'higher') than the value that fills the gap in 'the A of Y is ...'.

In many natural languages there are alternative constructs for expressing this sort of idea, some of which use the attribute name with a modifier, e.g. 'more height', 'less weight', some of which use a comparative form of an attribute value, e.g. bigger, weightier, smaller. As usual we are not concerned with the precise syntax but with what might be expressed.

Although most of the examples given involve relations between physical objects there are also relations between events and processes that are intrinsic. For instance one of the attributes of a process is its duration. So one process can have a longer or shorter duration than another, or a duration that differs by 2 seconds or 2 days, etc. One process may involve something happening faster than another, e.g. rotation, or colour changes, or speed of motion.

The examples given so far of relations based on attribute values are all binary. However insofar as attribute values are ordered, and many different relations can exist between items in an ordered set, there will be derived non-binary relations between objects with those attributes. For instance

- Ternary relations

For example: X is between Y and Z in A (e.g. in height, in colour, in size) means something like the attribute of type A of X is between the attributes of type A of Y and Z (in the ordering of attribute values of type A). This could be true if A is height, and the attribute values of X, Y and Z are respectively 6cm, 3cm and 22cm, or, 6cm, 22cm and 3cm (since the proposition says nothing about the relative position of Y and Z). There are many other Ternary relations, such as that X and Y differ from Z in attribute A by the same amount.

- Quaternary relations

As with ternary relations many quaternary relations can be formulated on the basis of attribute values, for example the proposition that W differs from X in A more than/the same amount as/less than Y differs from Z in A.

Exactly which sorts of n-ary relations can exist between objects with attributes of type A will depend on the structure of the space of values of A. If there is a small discrete set of values, e.g. small, medium and large, then far fewer relationships will be possible than if there are many values, or if the values form a continuum (or dense set).

If the set of attribute values has some sort of distance metric, then comparisons of distances can be used to construct yet more attribute-based relationships, which are not necessarily elegantly expressible in English. Are some languages more suited to saying this sort of thing than others? Examples are:

- X is more tall than Y by a larger amount than Z is more tall than Y.
- W is more red than X by the same amount as Y is more red than Z.

In the case of some qualitative attributes, such as colour, taste and smell, we find it useful to describe things in terms of closeness to particular important or interesting attributes, for instance describing X as more red than Y (redder) or more sweet than Y (sweeter). In contrast we don't say of X's height that it is more 10cm than Y, as a way of saying that X's height is closer to 10cm than Y's height. However in principle there is no reason why a particular language should not have such syntax added to it, if that would prove useful, e.g. because some heights have particular social, religious or economic importance, for certain objects.

All the types of propositions expressing relations derived from ordering or other relations between attribute values allow the removal of components leaving gaps on which question-forming transformations can operate. For example, introducing different gaps in the proposition '*X is between Y and Z*' in A, could produce such questions as

- What is between Y and Z in colour? (one gap)
- In what respect is X between Y and Z? (one gap)
- What is X between in colour? (two gaps)

### 5.1.7 – Relations based on shape

The examples so far have used attributes whose values are essentially points in a space of possible values, where the points themselves (e.g. heights, widths, colours(?)) have no structure, though the space may be wholly or partially ordered, discrete or not, finite or not, and so on. One particular set of attributes has elements which instead of being mere points in a space (or even vectors in a

fixed dimensional space) are structures of varying complexity, namely shapes. Two objects can have shapes that in turn can have attributes of varying complexity, for instance, number of edges, number of holes, relative lengths of holes and edges, being convex or concave all over or in parts of the surface, having grooves, having dents and many more.

The precise set of shapes that any individual can think about and ask questions about will vary from species to species as well as from individual to individual within a species. Of particular interest in robotics and (human or animal) psychology are the relationships between shapes of objects and the actions that individuals can perform. Whatever set of names we have for shapes will never be enough because shapes can be made more and more complex indefinitely by adding components, replacing components, joining two or more shapes in various ways to make a more complex shape, adding grooves, hollows, bumps, holes, and so on. The fact that such transformations produce shapes that cannot be recognised does not imply that they cannot be seen. On the contrary, seeing shapes needs to come before recognising them when they are re-encountered.

The point of all that is that it is a mistake to think of perceptual systems as producing only propositions about features and relationships of recognised (or named) whole objects. Rather there will be many object fragments and surface fragments that are seen, and relations between them that are seen, including not only relations within images, which don't concern us for the moment, but also relations within the 3-D environment.

Thus if there are propositions expressing what is seen they may be in part like parse trees summarising a network of relationships between simple or complex fragments, which may or may not be recognised fragments.

The process of question formation through gap manipulation then will have to operate on these structures, and the results may not be easily expressible in familiar human language. For example it may be possible to wonder whether a partly visible object or fragment has a shape for which there is no label but is (internally) identified by the perceiver as 'that shape', referring to the shape of a wholly visible component of the scene.

Similar comments apply to perceived events and processes, such as the motion of an object. For instance a perceiver may wonder what kind of action that it can produce could cause X to move as Y is seen to move, where there is no name or prior label for the latter type of motion. (E.g. think of a dancing student wondering how to replicate a teacher's actions.)

### **5.1.8 – Extrinsic relations: Spatio-temporal relations**

So far we have considered intrinsic relations based on attribute values, including relations like being more or less A, where A is an attribute, or having an attribute value between other attribute values. Among these are relations like being the same shape or having a shape with fewer concavities, that depend on the shapes of the objects standing in the relations.

There are also *extrinsic* relations that depend not only on the attributes of the objects but also where they are, and how they are oriented in space, and in the case of events and processes some of the relations depend on temporal location, i.e. when they occur, when they start, when they end, etc.

There are also very many relations that depend on spatial location, orientation, and shapes of objects. Some of these involve metrical relationships, such as being similar in shape, or being a certain distance apart, or being adjacent, whereas others can be thought of as purely topological, such as being inside, or being linked (as two rings can be). Others are a mixture.

For instance when a hook and a ring are joined up as in a toy train where one truck pulls another, it is possible to think of the relation between hook and ring that makes pulling happen as metrical

rather than topological because bending the hook so that it is lifted out of the ring and then makes pulling impossible will change only metrical properties, not topological properties. However if we think of the hook as part of a ring with a missing part then we can think of the hook and its missing part as linked to the ring, a topological relationship. The need to move between more or less metrical and topological relationships in solving problems is discussed in a little more detail here

A. Sloman, *Diagrams in the mind*,  
Ed. M. Anderson and B. Meyer and P. Olivier,  
*Diagrammatic Representation and Reasoning*,  
Springer-Verlag, Berlin, 2001,

A.G. Cohn, S.M. Hazarika, *Qualitative Spatial Representation and Reasoning: An Overview*, (2001).  
<http://www.comp.leeds.ac.uk/qsr/pub/funinfreview.ps.gz>

**ADD Distinction between relations in some ‘global’ space vs relations within an object or surface. E.g. fingers have locations relative to the hand, the arm, and the whole animal, as well as locations in the room, etc.**

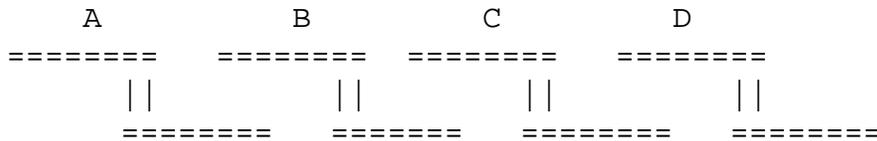
A complete survey of spatial relations would require a book at least. However a few points are worth making about some of the kinds of variation a perceiver may need to take account of.

Some relations are relative to an object’s intrinsic frame of reference, e.g. in front of, on the left side of, behind, when applied to something like a car, a horse, or a person. For objects that don’t have an intrinsic front and back, such as a ball, a box or a banana, the very same words can be used to specify a relationship based on the perceiver’s viewpoint, so that ‘in front of’ is interpreted as equivalent to something like ‘nearer to me’.

Many of the relationships for which we have names, like the ones mentioned above also have a kind of indeterminacy as to what regions are involved: e.g. if you look at a box in the middle of the table, and consider under what conditions you would describe a ball placed anywhere on the table as being to the left of the box, there may be considerable regions of indeterminacy. It is sometimes thought that such cases reflect statistical concepts with a probability distribution. An alternative view is that they are actually higher order concepts with an implicit reference to the *purpose* or *task* for which the relationship is being described. If there is no specific purpose or task the question about where the boundaries between the regions (left, right, front, back, etc.) are is essentially pointless, like asking whether a specific rock is or is not big without having any criterion for things being big enough. If forced to answer such pointless questions (e.g. in a psychology experiment asking which locations are to the left of that box) people may then use a default rule that covers many of the normal reasons for mentioning the ‘left of’ relation.

Instead of postulating a probability distribution it may be more accurate to regard this as a case where there is no answer, and therefore the statement ‘X is to the left of Y’ in such a case has no truth value, like the statement ‘My lawnmower is better than my car’ taken out of any context specifying the relevant respects of comparison.

Many spatial relations involve an order which depends on location or a combination of location and other attributes. For instance things may be ordered according to their distance from some reference location (e.g. the viewer) or according to what partially occludes what. ‘X occludes Y’ often implies ‘Y is further away than X’ but need not where complex shapes are involved. For



\O/

[Better diagram needed (see text below)]

instance, when the figure is viewed from somewhere below D, A occludes B, B occludes C, C occludes D, but A may be furthest from the viewer and D nearest.

For a perceiver that is capable of manipulating objects, many spatial relations are concerned with what motions are possible or constrained. For example saying that there is a gap between X and Y could in some contexts refer to the possibility of some object currently under consideration being able to move between X and Y. There will also be many relations for which we don't have verbal names but which an expert manipulator learns to perceive and think about as relations that facilitate or obstruct actions, or which change during actions. For example, when a child lifts a cut-out picture of a car from its 'home' in a flat sheet of wood, and then attempts to replace it, she may be unaware of the requirement to line up the boundary of the car shape with the boundary of the hole it left behind, and merely think of putting it back in the same general location. Such a child can be seen to try to get the piece back into its hole merely by pressing hard on it.

At a later stage the child has learnt about the significance of the additional spatial relationship, and appreciates the need to do some sliding and rotating as well as pressing, in order to get the shape back in place. At that stage the child need not have any *name* for either the relationship of alignment or coincidence between corresponding parts of the outline of the car and of the hole, nor for the changing relationships that occur while it is being moved around in order to get it into position to drop into place. Yet the child may understand the relationships, may be able to think about what sort of misalignment exists (e.g. where there is a boundary mismatch) may be aware of information gaps, may change its viewpoint, or move its hand out of the line of sight, in order to obtain missing information and then use the newly acquired information in order to complete the task. Thus there is an intermediate stage in which a question is considered, for which no linguistic expression is available (at least not to the child), and an answer found, for which there is also no linguistic expression available.

Similar remarks can be made about tool manipulation and manipulation of items of food or nest-building materials in other animals, e.g. squirrels, monkeys, chimpanzees, crows, etc., none of which can use anything like a human language.

**NB:** Nothing said here about needing, acquiring and using information of a particular type im-

plies that the child or the animal *knows* what it is doing. That kind of self knowledge requires additional architectural resources and meta-semantic capabilities supporting reflective self-description. It is possible that many animals that are very good at thinking about and solving problems are totally incapable of thinking about or even being aware of their own problem-solving. This is probably true of much of what happens in young children: the meta-management architectural layer required for self-understanding does not seem to develop until some time after many other capabilities have developed. In general, knowing X, and knowing that you know X require quite different mechanisms, including different forms of representation whose semantic contents presuppose very different ontologies. The latter requires a meta-semantic ontology which includes things that manipulate information, whereas the former does not.

---

## 6 – mess to be fixed follows

further moreleft moreright moreinfront morebehind closerin (colour, height, size, distance, .....)

Ternary relations can be generated from features that allow ‘more’ or less’, e.g more red, more high, more fast

- – **QUATERNARY**

- W is as far from (close to, higher than) X, as Y is (to) Z

- – **GEOMETRICAL RELATIONS**

These can involve arbitrarily many objects. I.e. the relations can be

- unary (is a triangle),
- binary (contains)
- ??
- ternary
  - form an equilateral triangle
  - form an isosceles triangle
  - are collinear
  - form a acute/obtuse/right angle (X Y Z)

etc.

### PROPOSITIONAL FORMS

x is

x is at

x is a

x’s is

morethan (X’s Y’s )

Existentially quantified Universally quantified Most (Xs are Ys) More (Xs are P than Ys are)  
 Numberof Xs is N Numberof Xs has some numerical property.

## 7 – What about a non-linguistic (pre-linguistic) agent?

A significant subset (what subset?) of the above proposition types and question types may be capable of being considered by animals without human language and pre-linguistic children. The questions will not be posed externally, but internally, and will determine goals for cognitive and other processes.

So that raises the question of what sort of formalism could do these things \*within\* a cognitive agent, as opposed to \*between\* speakers?

It may be that answering that question will give us deeper insights into what goes on in speakers, for they will presumably, to some extent, build on the pre-verbal capabilities that evolved earlier.

Later on the internal processes were certainly expanded by the availability of linguistic constructs – e.g. things like ‘By how much is the distance from A to B bigger than the distance from C to D?’

Could a pre-linguistic animal (or child) wonder about such a question?

Some of these issues are posed in relation to images of a cup, saucer and spoon in this short presentation: <http://www.cs.bham.ac.uk/research/cogaff/challenge.pdf><sup>23</sup>

## 8 – Nature Nurture Issues

Suppose it is true that at least some natural language users are capable of understanding the full variety of types of propositions and questions sketched here, in at least some sense of "understanding" involving knowing what sort of thing counts as a correct answer, without necessarily knowing the answer or even known how to find or verify the answer – because that might require technology that has not yet been invented, or in some cases might require use of a theory that has not yet been discovered). This raises the question how such understanding arises?

One possible answer is that it arises out of the use of the sorts of mechanisms sketched in this draft paper (co-authored with bio-scientist Jackie Chappell) <http://www.cs.bham.ac.uk/research/cogaff/altricial-precocial.pdf><sup>24</sup> discussing the precocial/altricial spectrum, where altricial species (or altricial skills) instead of being restricted to (a) learning by the use of positive and negative reinforcement mechanisms that gradually transform weights in some sort of statistical mechanism (e.g. a neural net) can also be learnt by (b) a combination of more ‘symbolic’ mechanisms:

- exploring (randomly or otherwise) effects of many kinds of distinct sensory patterns and action types (starting with innate ones)
- noticing, categorising, and inventing (internal) labels for ‘interesting’ cases
- using some (innate? learnt?) ‘syntactic’ mechanism for internally recombining the labels in various ways to generate new, larger perceptual and action structures, which are also explored and if found interesting stored, labelled, and made available for re-use

Such mechanisms can support discrete, creative learning and discovery steps producing quite large changes in competence.

(The idea goes back at least to work by Oliver Selfridge, as reported in this toy demonstration

---

<sup>23</sup>See URL <http://www.cs.bham.ac.uk/research/cogaff/challenge.pdf>

<sup>24</sup>See URL <http://www.cs.bham.ac.uk/research/cogaff/altricial-precocial.pdf>.

<http://www.cs.bham.ac.uk/research/poplog/teach/finger><sup>25</sup> which can be run using the poplog system.<sup>26</sup>

Defining that hypothesis about altricial mechanisms properly will have to wait for another occasion. For now it is important only that these mechanisms have implications both for the innate ontology of the robot or animal and for how the ontology develops over time.

---

## 9 – Some References

These references were found with the help of Google, after the first draft of these notes, concerned only with question formation, was written. The list is neither complete nor authoritative nor representative. It's just what I found in a fairly short time spent searching. The Burhans thesis mentioned at the beginning seems to be closest in spirit to what I have been proposing regarding question formation, though I have not yet read it all.

List of question types.<sup>27</sup>

IR-323: (2003) Metzler, D. and Croft, W.B., " Analysis of Statistical Question Classification for fact-based Questions " *Information Retrieval*. PDF<sup>28</sup>

Jean-Pierre Ko (2000)

Any questions left? Review of Ginzburg & Sag's *Interrogative Investigations*. *J. Linguistics* 40 (2004), 131-148. f 2004 Cambridge University Press DOI : 10.1017/S0022226703002354, UK. PDF<sup>29</sup>

Jonathan Ginzburg and Ivan Sag, (2000) *Interrogative Investigations*, CSLI Lecture Notes 123. Stanford, California: CSLI publications,, Pp. xii + 449.

Pylyshyn, Z.W., 'The role of location indexes in spatial perception: A sketch of the FINST spatial-index model,' *Cognition* 32, 65-97, 1989.

<http://ruccs.rutgers.edu/ftp/pub/papers/cognit89.pdf><sup>30</sup>

See OpenCyc on 'Spatial Relations'

<http://www.cyc.com/cycdoc/vocab/spatial-vocab.html><sup>31</sup>

This document describes collections, predicates and other Cyc constants that are used to represent spatial objects and relations. See also documents for Groups, Quantities, Movement, Paths & Trajectories, Parts Of Objects, and Geography.

A. Trehub, 1991, *The Cognitive Brain*, MIT Press, Cambridge, MA,

---

<sup>25</sup>See URL <http://www.cs.bham.ac.uk/research/poplog/teach/finger>

<sup>26</sup>See URL <http://www.cs.bham.ac.uk/research/poplog/freepoplog.html>

<sup>27</sup>See URL <http://www.usingenglish.com/glossary/question-types.html>

<sup>28</sup>See URL <http://ciir.cs.umass.edu/pubfiles/ir-323.pdf>

<sup>29</sup>See URL <http://lingo.stanford.edu/sag/revs/koenig-rev.pdf>

<sup>30</sup>See URL <http://ruccs.rutgers.edu/ftp/pub/papers/cognit89.pdf>

<sup>31</sup>See URL <http://www.cyc.com/cycdoc/vocab/spatial-vocab.html>