euCognition
The European Network for the Advancement of Artificial Cognitive Systems

Research Roadmap Presentation: euCognition Meeting, Munich 12 Jan 2007

http://www.eucognition.org/six_monthly_meeting_2.htm

What's a Research Roadmap For?
Why do we need one?
How can we produce one?

Aaron Sloman

(http://www.cs.bham.ac.uk/~axs)

University of Birmingham, United Kingdom

In collaboration with Bill Sharpe and members of the EU CoSy Robotic Project

These slides are available online here:

http://www.cs.bham.ac.uk/research/projects/cosy/papers/#pr0701

See also the euCognition Research Roadmap project:

http://www.eucognition.org/wiki/index.php?title=Research_Roadmap

Caveat

My aim is not to produce intelligent machines

(except as experiments testing the theories and models).

This is a presentation from the viewpoint of a researcher trying primarily to understand and model aspects of natural intelligence

not only human intelligence, but also insects, birds, deer, hunting mammals, primates, landmark-using ants, portia spiders, Betty the hook-making crow, ...

However if we ever do understand the products of biological evolution well enough to produce convincing working models, that may enable us to be far more effective at producing intelligent robots to meet many kinds of practical needs.

Such as looking after me in my helpless old age?
Or maybe looking after your great, great... grandchildren in their old age!

Extravagant claims and predictions have been made repeatedly since the early days of computers, compared with which progress has been very slow, though real. (Often referred to as 'hype'!)

So we need to get better at predicting what we can and cannot do when, and we need to replace warring factions with collaborating researchers combining different, complementary, approaches.
That requires everyone to become more broad-minded.

A note about this slide presentation

This is one of many (partially overlapping) slide presentations located at

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

and

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

My slides are primarily intended for online reading.
As a result they usually contain far too much text for use in presentations.
But I don't have time to make two versions of each.

This is one of many papers and presentations written since 1971 aiming to clarify requirements for AI systems that are intended to incorporate human competences.

A draft incomplete list of the papers on requirements can be found here

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

An unusual kind of research: into REQUIREMENTS

People normally assume requirements are given before research or development starts.

However, for AI/Robotics/Cognitive Systems research understanding requirements is itself a major research activity.

This is part of a long term study of the space of requirements
to be satisfied by theories, explanations, and working systems –
a space that needs to be related to the space of possible designs.
and mappings between those spaces (not just univalued fitness functions).

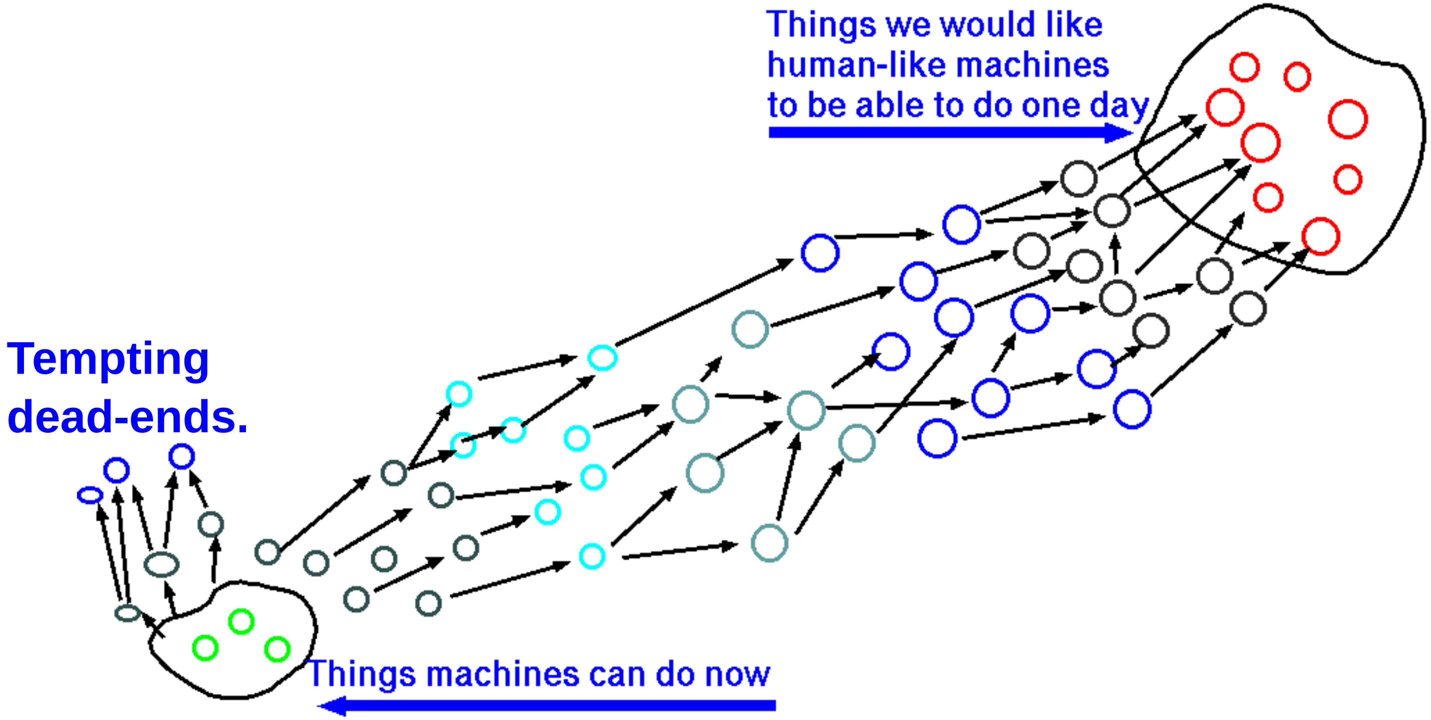

What? Why? How? What?

- **What's a Research Roadmap For?**
  - **An agreed specification of what the problems are: what we are trying to do.**
  - **When the problem is very complex, a roadmap can break the problem down into significant sub-problems, helping with research planning.**
  - **It can be used to specify milestones and routes through them.**
- **Why do we need one?**
  - **Because there have been so many past optimistic predictions that failed!**
  - **Because even people who disagree on mechanisms, architectures, representations, etc. may be able to agree on requirements.**
- **How can we produce one?**

What? Why? How? What?

- **What's a Research Roadmap For?**
  - **An agreed specification of what the problems are: what we are trying to do.**
  - **When the problem is very complex, a roadmap can break the problem down into significant sub-problems, helping with research planning.**
  - **It can be used to specify milestones and routes through them.**
- **Why do we need one?**
  - **Because there have been so many past optimistic predictions that failed!**
  - **Because even people who disagree on mechanisms, architectures, representations, etc. may be able to agree on requirements.**
- **How can we produce one?**
  - **NOT BY ASKING YET MORE 'EXPERTS' TO MAKE YET MORE PREDICTIONS!**

What? Why? How? What?

- **What's a Research Roadmap For?**
  - **An agreed specification of what the problems are: what we are trying to do.**
  - **When the problem is very complex, a roadmap can break the problem down into significant sub-problems, helping with research planning.**
  - **It can be used to specify milestones and routes through them.**
- **Why do we need one?**
  - **Because there have been so many past optimistic predictions that failed!**
  - **Because even people who disagree on mechanisms, architectures, representations, etc. may be able to agree on requirements.**
- **How can we produce one?**
  - **NOT BY ASKING YET MORE 'EXPERTS' TO MAKE YET MORE PREDICTIONS!**
  - **Instead, collect very many possible future scenarios described in great detail based on human or animal existence proofs.**
  - **Analyse in depth requirements for achieving those details.**
  - **Order the scenarios, and requirements (by difficulty and by dependence).**

.... to be continued, after some examples ....

Example: vision is much, much, more than recognition

What competences are required in a visual system to enable a child (or a robot) to get from the first configuration to the second?

- **in many different ways,**
- **with different variations of the first configuration,**
- **with different variations of the second configuration,**
- **using the right hand,**
- **using the left hand,**
- **using both hands,**
- **using no hands, only mouth...?**

Can you visualise such processes – including interacting curved surfaces?

For more on this see

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

The multiple functions of vision

There is a vast amount of research on visual classification and recognition.

However, most of that research ignores the fact that people, and other animals can see, and manipulate things they do not recognise, e.g. strange animals or machines.

Moreover much of what we see involves processes not just static objects and configurations of objects. Partly that's because we use vision to work out what actions are possible, prior to deciding what to do, and then to help control actions that result. This implies that vision includes seeing what is possible, as opposed to what exists or is happening. That includes seeing empty spaces.

That's a part of what is involved in perceiving affordances.

For more on the non-recognition functions of vision see

- `http://www.cs.bham.ac.uk/research/projects/cogaff/06.html#0604`
  Image interpretation: The way ahead? (1982)
- `http://www.cs.bham.ac.uk/research/projects/cogaff/81-95.html#7`
  On designing a visual system: Towards a Gibsonian computational model of vision (1989)
- `http://www.cs.bham.ac.uk/research/projects/cogaff/challenge.pdf`
  A challenge for vision researchers
- `http://www.cs.bham.ac.uk/research/projects/cosy/papers/#pr0505`
  A (Possibly) New Theory of Vision
- `http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk7`
  When is seeing (possibly in your mind's eye) better than deducing, for reasoning?
- `http://www.cs.bham.ac.uk/research/projects/cogaff/misc/nature-nurture-cube.html`
  On seeing changing binary pixels as a rotating wire-frame cube.

Another example: Ontologies for getting at something

Understanding varieties of causation involved in learning how to get hold of a toy that is out of reach, resting on a blanket, or beyond it.

Some things to learn through play and exploration

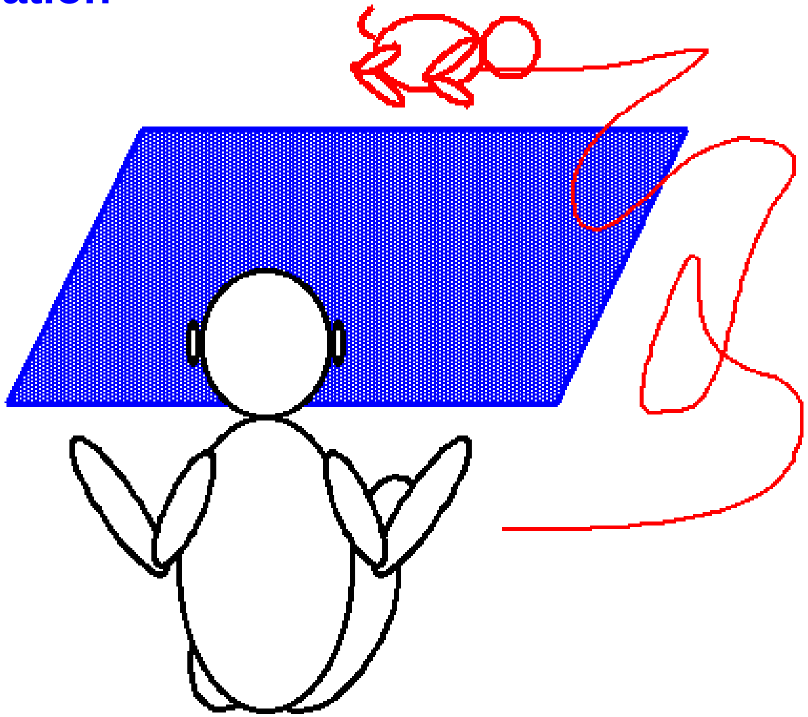

Toy on short blanket Grab edge and pull

Toy on long blanket Repeatedly scrunch and pull

Toy on towel Like blanket

Toy on sheet of plywood
Pull if short(!!), otherwise crawl over or round it

Toy on sheet of paper Roll up?
(But not thin tissue paper!)

Toy on slab of concrete Crawl over or round

Toy at end of taut string Pull

Toy at end of string with slack Pull repeatedly

String round chair-leg Depends

Elastic string
?????

See this discussion of learning orthogonal recombinable competences
http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0601

It takes a lot of learning to develop all the visual and reasoning competences required for seeing and understanding these affordances – including visualising what would have happened if you had done something different, or if someone else were to move something.
Our spatial and visual competence goes far beyond actually doing.

Two notions of causation: (Humean & Kantian)

Understanding causation is one of the requirements for competence.

Consider two gear wheels attached to a box with hidden contents.

Can you tell simply by looking, without actually turning anything, what will happen to one wheel if you rotate the other about its central axis?

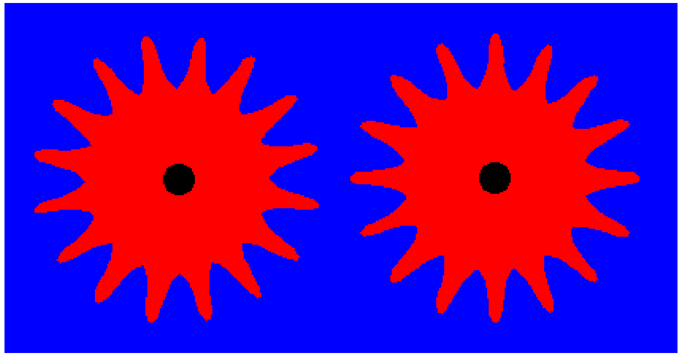

Only in the lower case. (How do you tell?)

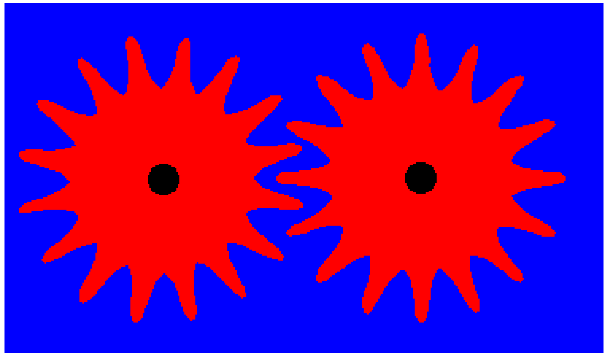

Seeing what must happen if....

A child, and an intelligent robot able to find out how things work in the environment, will learn the difference between causation that is merely correlational (Humean) and causation that is based on intelligible structure and is therefore also deterministic (Kantian).

Invent more examples to fit both cases.

Which other animals learn about Kantian causation?

Watch 'clever or funny animals' TV shows!

Why will robots need this?

For more on this see

`http://www.cs.bham.ac.uk/research/projects/cosy/papers/#pr0506`

Two views of child as scientist: Humean and Kantian

Some more abstract competences

Besides being able to see, hear, feel, manipulate, use, and react to processes involving configurations of objects in the environment, humans can also learn about, think about, reason about, be puzzled about, communicate about, prove things about, many kinds of abstract entities, including:

- **Games – including their rules, and many entities that can be involved in games, e.g. goals, threats, illegal moves, etc.**
- **Numbers**
- **Calculations**
- **Proofs**
- **Stories**
- **Plans**
- **Strategies**
- **Values**
- **Theories and explanations**
- **Their own and other people's goals, beliefs, desires, hopes fears, puzzles, confusions, strategies, etc.**
- **Social groups, social relations, social processes (e.g. revenge, punishment)**

These require significant representational and architectural competences beyond those commonly required for perceiving and acting in the environment as all other animals do, including microbes and insects. E.g. meta-semantic competences and deliberative competences are needed.

See requirements for 'fully deliberative' architectures:

`http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0604`

Making high level competences fast and fluent

Humans often learn a new competence that involves deliberately following procedures, thinking about options, either remembering rules, using charts or maps, or working out new plans.
After much practice using that competence in many different contexts, the deployment of the competence because much more fast and fluent. How does that happen?

Examples:

- **Learning to feed oneself, to dress and undress, tie shoelaces, wash dishes,**
- **Learning ones way around a building or a town**
- **Learning to play physical games like tennis, football, boxing, hockey,**
- **Learning to ride a bicycle or drive a car.**
- **Learning a board game like chess or othello.**
- **Learning to play a musical instrument.**
- **Learning to talk.**
- **Learning to read text, music, computer programs, flow-charts...**
- **Learning to count, to answer questions about numbers, to do calculations, to prove mathematical theorems, to think about infinite sets.**
- **Learning to see moods, feelings, intentions, etc. in other people.**
- **Learning to interact socially.**

Conjecture: brain mechanisms that originally evolved to support genetically determined competences (e.g. deer running soon after birth) later evolved to start with large amounts of spare capacity that could be used to acquire new competences under the control of deliberative and meta-semantic competences.

.....continued What? Why? How? What?

- **What's a Research Roadmap For?**
  - An agreed specification of what the problems are: what we are trying to do.
  - When the problem is very complex, a roadmap can break the problem down into significant sub-problems, helping with research planning.
  - It can be used to specify milestones and routes through them.
- **Why do we need one?**
  - Because there have been so many past optimistic predictions that failed!
  - Because even people who disagree on mechanisms, architectures, representations, etc. may be able to agree on requirements.
- **How can we produce one?**
  - NOT BY ASKING YET MORE 'EXPERTS' TO MAKE YET MORE PREDICTIONS!
  - Instead, collect very many possible future scenarios described in great detail

    based on human or animal existence proofs.
  - Analyse in depth requirements for achieving those details.
  - Order the scenarios, and requirements (by difficulty and by dependence).
- **What will we gain from doing it?**
  - More collaboration between currently warring factions.

    See next slide
  - More progress, avoiding dead ends
  - Better ways of evaluating progress.

Here are some preconditions for progress

- **Recognition that we all share a problem: how little our current systems can do, and how wrong most predictions have been.**

  Many expert systems, theorem provers, planners, trainable classifiers, evolutionary problem solvers, robot learning mechanisms have proved interesting and useful. But they are ALL narrowly restricted, and PATHETIC compared with a squirrel, a raven, a human 3 year old, even leaf-cutting ants, in their ability to cope with a structured physical environment

- **Willingness to try to develop a shared ontology for talking about:**

  behaviours, requirements, kinds of competence, kinds of information, kinds of mechanism, kinds of representation, kinds of architecture.... (Not just our favourite ones.)

- **Willingness to try to agree on some diagrammatic and notational conventions for presenting types of requirements, types of architectures and other designs.**

- **Avoidance of questions like 'is it really X?' (X=intelligent, conscious, cognitive, emotional): postpone the grand philosophical questions till we have FAR more interesting working systems.**

- **Willingness to examine theories and data from many disciplines.**

- **Willingness to teach our students to develop competences using several different approaches, instead of telling them 'the others have failed'.**

- **Willingness to reclassify assumed established truths as controversial claims.**

See the 'controversies' section of the euCognition wiki

`http://www.eucognition.org/wiki/index.php?title=Controversies_in_Cognitive_Systems_Research`

A partially ordered network of stages

The process of extending competence is not continuous (like growing taller):

- **A child has to learn about**
  - **distinct new kinds of objects, properties, relations, process structures, e.g. for rigid objects, flexible objects, stretchable objects, liquids, sand, treacle, plasticine, pieces of string, sheets of paper, construction kit components in Lego, Meccano, Tinkertoy, electronic kits...**
  - **new forms of representation, new kinds of transformations, new constraints on transformations, new applications of previously acquired information.**

  It is easy to come up with hundreds of significantly different examples of things to be learnt.
- **There are not fixed stages: there is no order in which things have to be learnt.**
- **There are many dependencies but not enough to generate a total ordering – each learner finds routes through several partially ordered graphs.**
- **What can be learnt varies both from one generation to another and from one location to another.**
- **Provision of new kinds of playthings based on scientific and technological advances is a major form of communication across generations.**

  Likewise games, stories, poems, languages, pictures, theories, ...

A collection of research milestones may also have many discontinuities, of different kinds, and will also form a partial ordering.

All this is not required for production of highly specialised robots (like robots in car factories.) Standard engineering approaches are OK for narrow objectives.

How to develop a research roadmap

Observe feats of humans (e.g. young children, playing, exploring, communicating, solving problems) and other animals (e.g. nest-building birds, tool makers and users, berry-pickers and hunters). These provide many existence proofs, not of specific mechanisms, but of a wide variety of possible behaviours for intelligent embodied individuals i.e. requirements.

Previous slides have given a tiny, but varied, subset of examples.

How to develop a research roadmap

Observe feats of humans (e.g. young children, playing, exploring, communicating, solving problems) and other animals (e.g. nest-building birds, tool makers and users, berry-pickers and hunters). These provide many existence proofs, not of specific mechanisms, but of a wide variety of possible behaviours for intelligent embodied individuals i.e. requirements.

Use those observed behaviours to develop and document a partially ordered network of more or less challenging scenarios – ordered by dependency, complexity, difficulty and variety of competences each scenario requires.

How to develop a research roadmap

Observe feats of humans (e.g. young children, playing, exploring, communicating, solving problems) and other animals (e.g. nest-building birds, tool makers and users, berry-pickers and hunters). These provide many existence proofs, not of specific mechanisms, but of a wide variety of possible behaviours for intelligent embodied individuals i.e. requirements.

Use those observed behaviours to develop and document a partially ordered network of more or less challenging scenarios – ordered by dependency, complexity, difficulty and variety of competences each scenario requires.

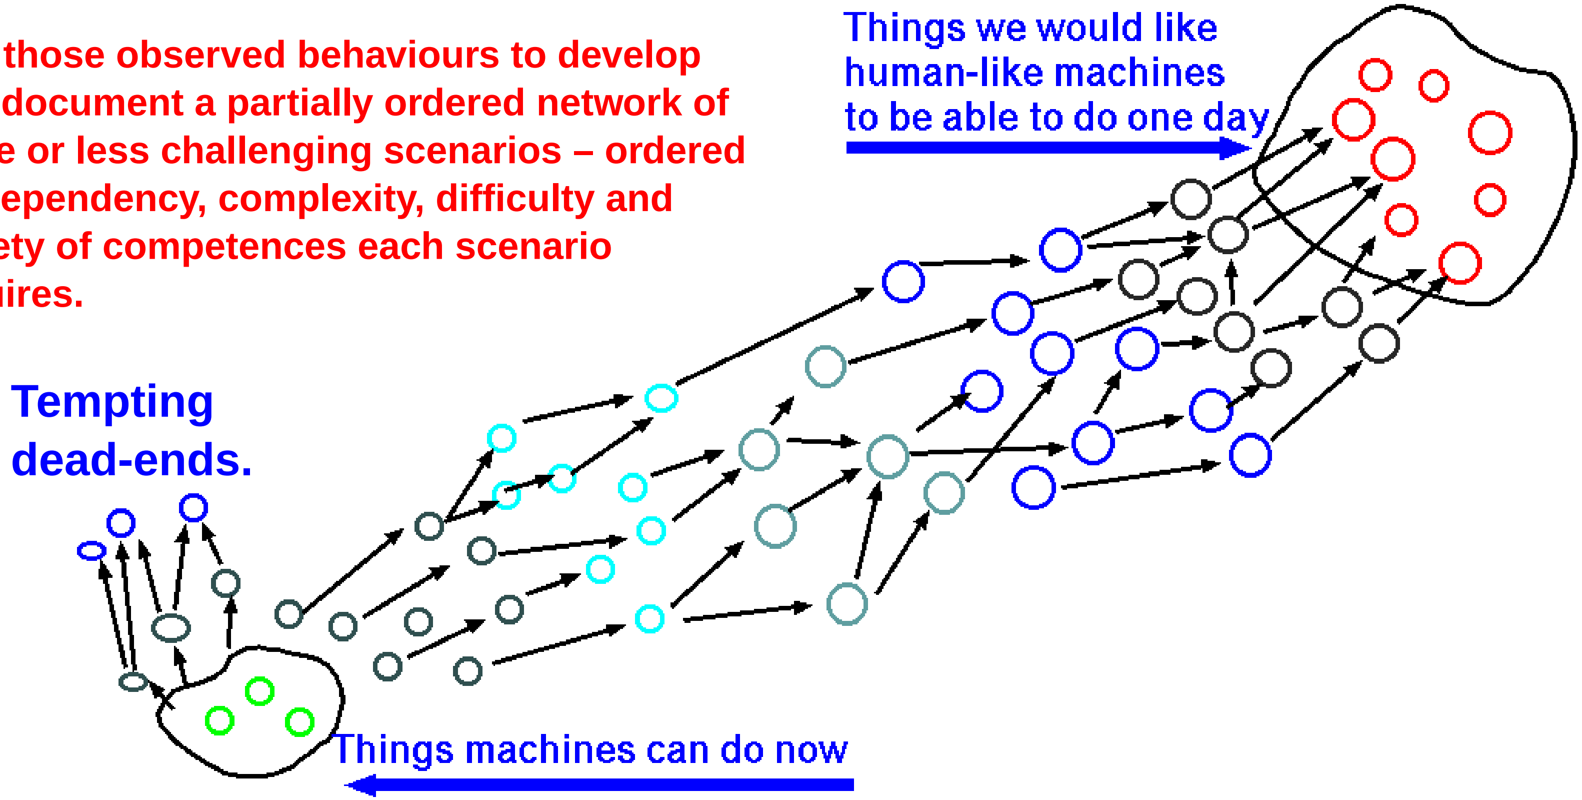

How to develop a research roadmap

Observe feats of humans (e.g. young children, playing, exploring, communicating, solving problems) and other animals (e.g. nest-building birds, tool makers and users, berry-pickers and hunters). These provide many existence proofs, not of specific mechanisms, but of a wide variety of possible behaviours for intelligent embodied individuals i.e. requirements.

Use those observed behaviours to develop and document a partially ordered network of more or less challenging scenarios – ordered by dependency, complexity, difficulty and variety of competences each scenario requires.

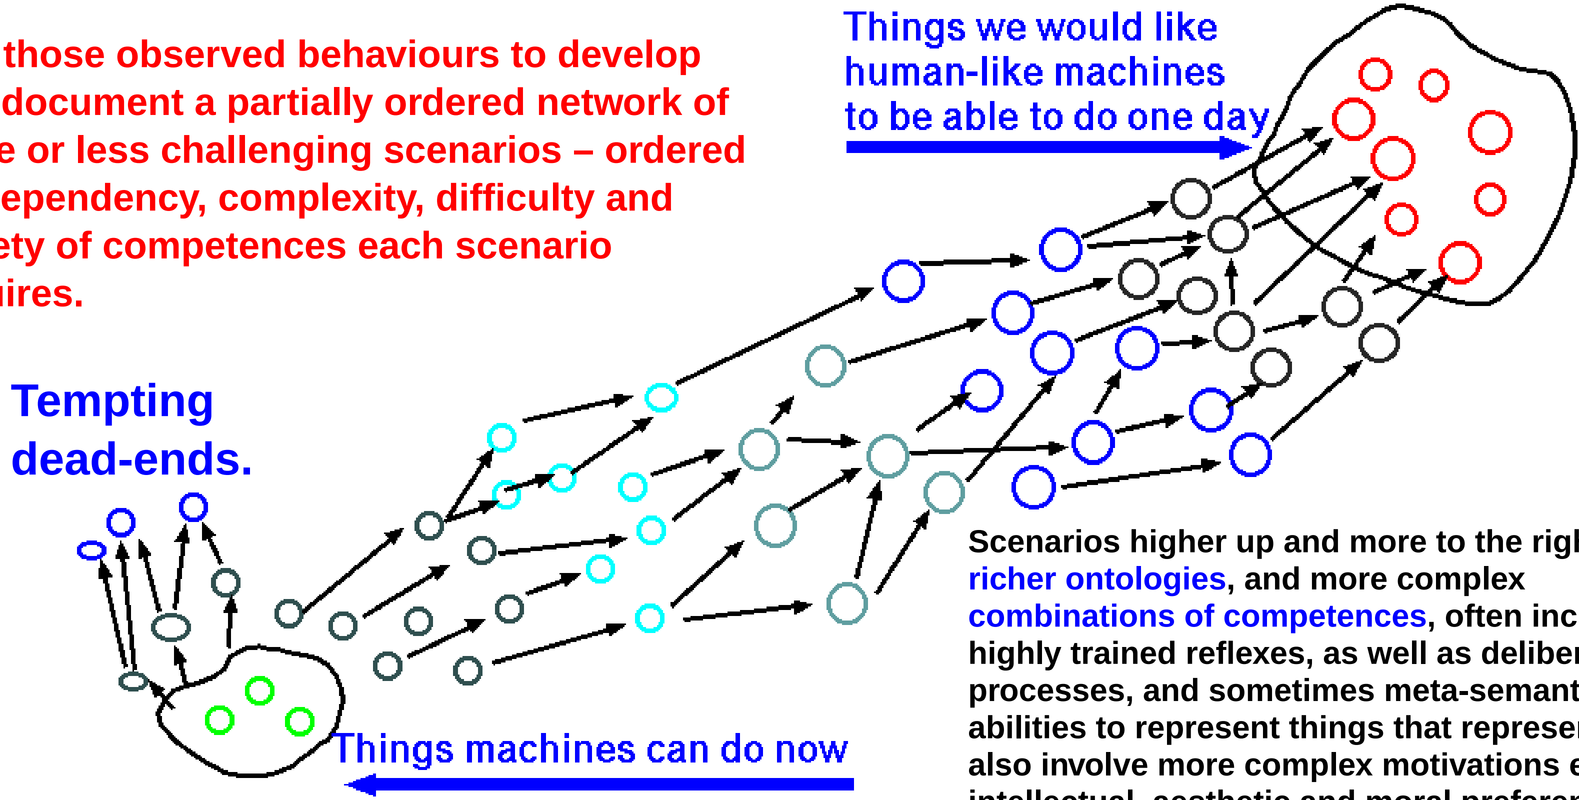

Scenarios higher up and more to the right use richer ontologies, and more complex combinations of competences, often including highly trained reflexes, as well as deliberative processes, and sometimes meta-semantic abilities to represent things that represent. They also involve more complex motivations e.g. intellectual, aesthetic and moral preferences.

How to develop a research roadmap

Observe feats of humans (e.g. young children, playing, exploring, communicating, solving problems) and other animals (e.g. nest-building birds, tool makers and users, berry-pickers and hunters). These provide many existence proofs, not of specific mechanisms, but of a wide variety of possible behaviours for intelligent embodied individuals i.e. requirements.

Use those observed behaviours to develop and document a partially ordered network of more or less challenging scenarios – ordered by dependency, complexity, difficulty and variety of competences each scenario requires.

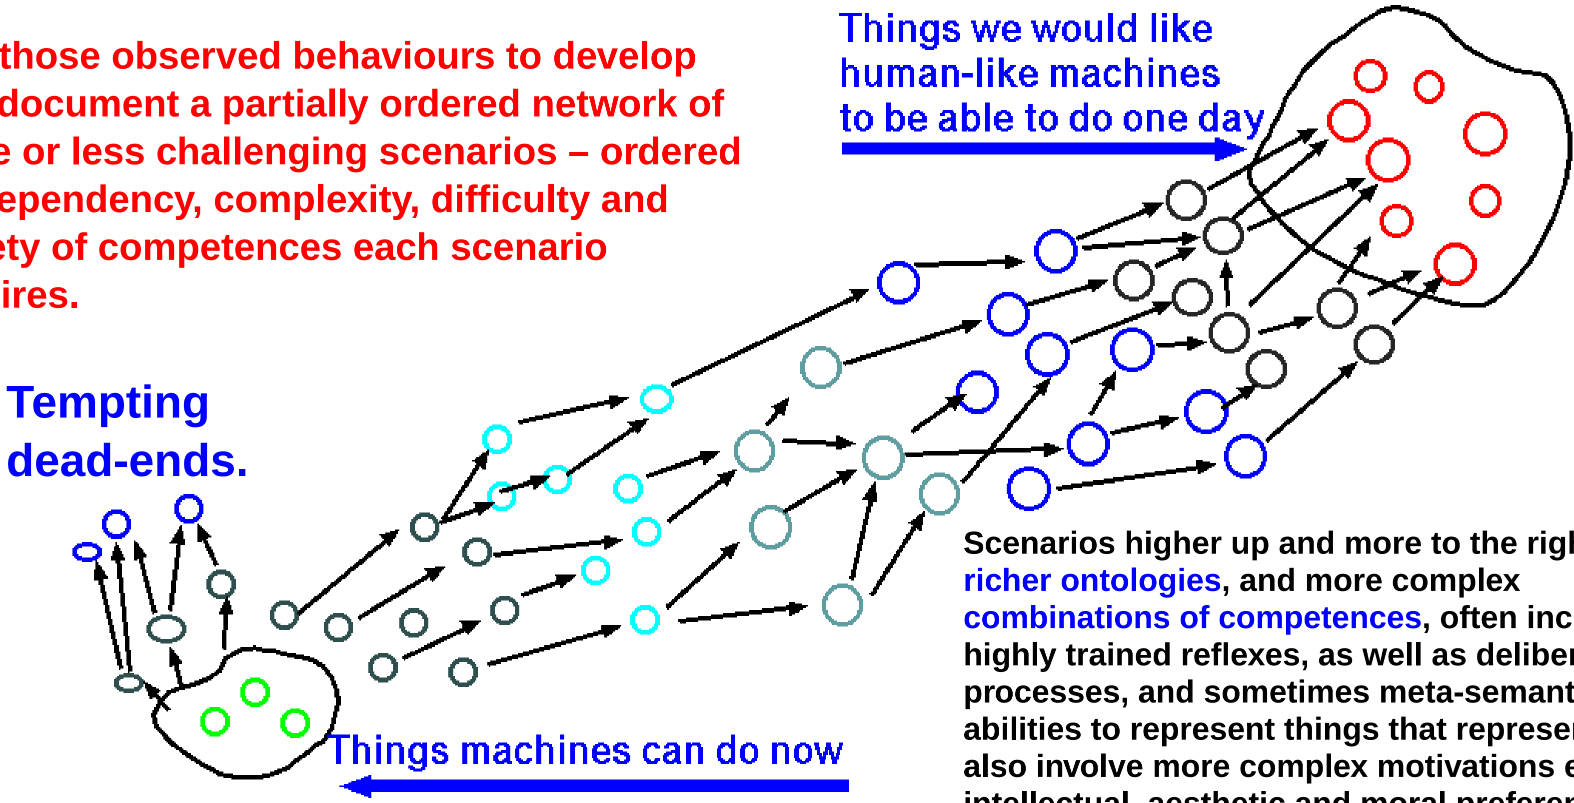

Scenarios higher up and more to the right use richer ontologies, and more complex combinations of competences, often including highly trained reflexes, as well as deliberative processes, and sometimes meta-semantic abilities to represent things that represent. They also involve more complex motivations e.g. intellectual, aesthetic and moral preferences.

We should plan more of our research by identifying long term requirements in great detail and working back through less demanding requirements.

More complex behaviours need to be specified in terms of the competences they require.

Behaviour type XXX has features requiring competence of type YYY

NB: such claims — e.g. ‘These linguistic and planning behaviours require mechanisms with recursive capabilities and compositional semantics’ — must be open to testing and refutation.

Competences lead to ideas about architectures. There are too many different, arbitrary, terminologies, diagrammatic conventions, preferred architectures. We need a common vocabulary for talking about architectures. A first draft example is the CogAff schema.

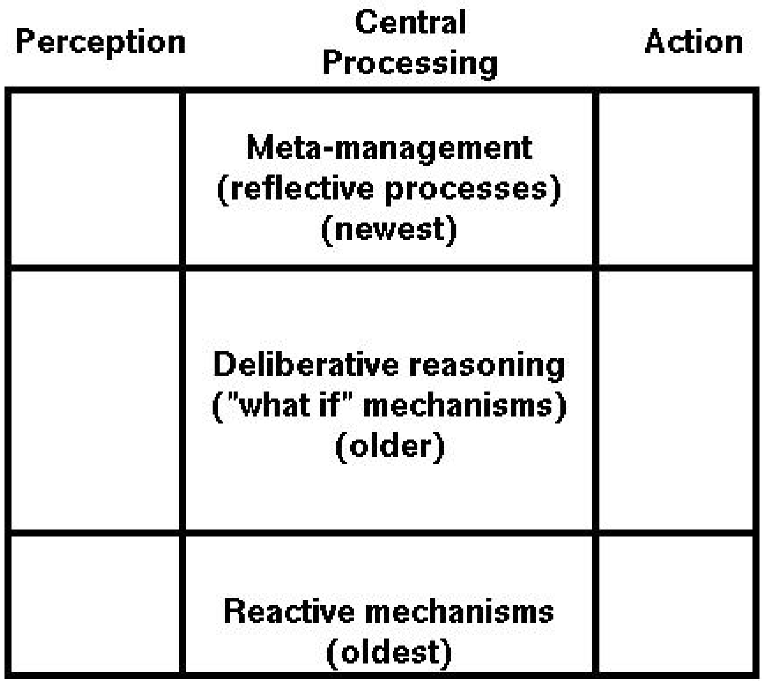

Requirements for subsystems can refer to

- Types of information used (ontology used)
- Forms of representation (continuous, discrete, Fregean, diagrammatic, distributed, dynamical...)
- Uses of information (controlling, describing, planning, teaching, questioning, instructing...)
- Types of mechanism (many examples have already been explored – there may be lots more ...).
- Ways of putting things together in an architecture or sub-architecture

Architectures vary according to which of the boxes contain mechanisms, what those mechanisms are required to do, which mechanisms are connected to which others, what sorts of connections there are, what sorts of learning can occur, whether the architecture grows itself....

Two conjectures

A: The most general capabilities of humans, which are those provided by evolution, and which support all others, develop during the first few years of infancy and childhood. We need to understand those in order to understand and replicate the more ‘sophisticated’ and specialised adults that develop out of them.

Attempting to model the adult competences directly will often produce highly specialised, unextendable, and probably very fragile systems – because they lack the child’s general ability to accommodate, adjust, and creatively re-combine old competences.

B: There are many aspects of human cognition that evolved originally to meet requirements for 3-D vision and action — including intricate manipulations of 3-D structures — using exosomatic ontologies. The mechanisms and forms of representation are now used for many other purposes, and can be used by people who have been blind from birth or who were born without limbs.

Embodiment of your ancestors is more important for development of your mental competences than your own embodiment.

Tempting traps to avoid: DON'T:

- **DON'T:** Worry about scaling up

  Humans don't scale up – they scale out. (I.e. combing old competences in new ways.)
  Machines can beat humans on almost any specific task. But humans can use any competence in combination with others in creative ways.
  Examples ....
- **DON'T:** Focus on benchmarks.

  If you focus all energy on improving performance on fixed tasks, you'll (probably) produce something that does not 'scale out'.
- **DON'T:** Focus on forms of representation whose semantics is purely somatic

  i.e. concerned with sensor signals, motor signals and internal states. Much greater power, and animal competence can come from exosomatic ontologies.

  Including the ability to teach, to imitate, to help someone else...
- **DON'T:** Equate perceiving with recognising

  Vision enables controlling actions, perceiving structure, understanding how something works, communicating, noticing possibilities and constraints....
- **DON'T:** Ignore the details of your own environment and what you do with it.
- **DON'T:** Search for the BEST representation.

  For many problems, perhaps all, different forms of representation are needed for different sub-tasks – e.g. planning vs control of fast actions, vs understanding what went wrong, vs answering questions about what you did.

Varieties of learning and development

Summary – explained in more detail later.

An observation: **Individual learning and development, at least in humans, and to some extent other species, involves various combinations of learning and development, based on various combinations of genetically provided competences, including meta-competences that provide new competences, at different levels of abstraction. Humans also have meta-competences that lead to even more sophisticated meta-competences, e.g. when studying at university level.**

Multiple routes from genome to behaviours

(Environment affects all embedded processes)

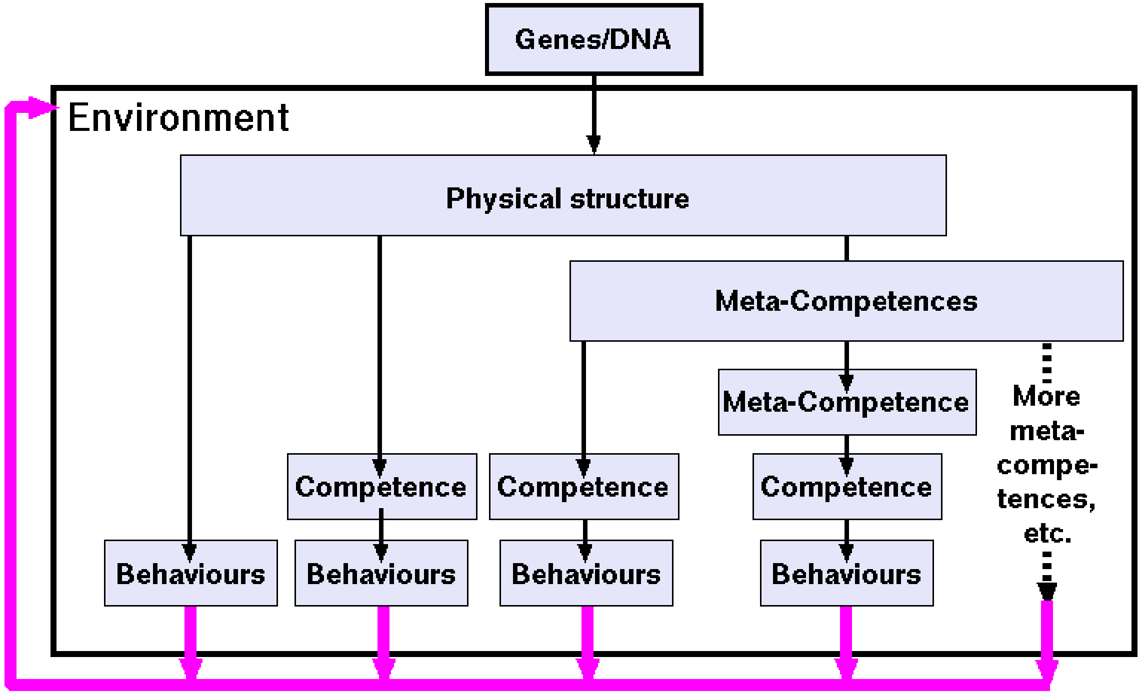

NB: All boundaries are somewhat fuzzy

Based on work with Jackie Chappell. (To appear in IJUC)

Thanks to Chris Miall for suggesting this format.

Sometimes development of a new competence or meta-competence requires ontology extension: different sorts – how are they achieved?

Methods and tools to help build roadmaps

Many people find it very difficult to think up a systematic and comprehensive collection of future scenarios of the kind required. We have been working on a methodology to help with development of this network of roadmaps, using a 3-D 'Grid of Competences'

Columns represent types of *entity* to which competences can be applied (e.g. 2-D and 3-D spatial locations, regions, routes, inert objects, mobile objects, objects that have perception, goals and actions, and more abstract entities such as beliefs, proofs, numbers, plans, concepts, questions, problems).

Rows represent types of *competence* that can be applied to instances of some or all of the types of entities; e.g. competences like perceiving, manipulating, referring to in thought, referring to in language, constructing, dismantling, ...

Entity-types / Competences	E1	E2	E3	E4	E5	E6	E7	E8
C1								
C2								
C3								
C4								
C5								
C6								
C7								

The third dimension is depth of items in the boxes representing difficulty of the competence.

The degree and kind of difficulty will affect time required to produce working systems.

NOTE:

a more complex topology than a rectangular grid is required: refinements and elaborations of the grid are topics for future research. (For more detail see the introduction to GC5 symposium in proceedings or website http://www.cs.bham.ac.uk/research/cogaff/gc/aisb06/).

For a first draft sample grid see http://www.cs.bham.ac.uk/research/projects/cosy/matrix

The Grid is Over Simple

The grid generates various types of competence applied to various types of entity.

E.g. consider the many kinds of things, of different sizes, shapes, weights, kinds of fragility, that you can grasp in different ways, using two or more fingers, two hands, your mouth, using tweezers, etc., and the subtle and complex requirements for vision in these tasks

- **Combining different subsets of the grid, at different depths, produces scenarios of varying complexity, creating milestones on the long term roadmaps/graph, defining scientific challenges that everyone will agree are hard.**
- **Progress can be measured by which portions of the graph have been achieved.**
- **Benchmarks requiring integration of different combinations of competences can be defined by people who disagree on methods and mechanisms.**
- **The grid is an oversimplification: some boxes need subdivisions, and other boxes will be empty.**

  You can refer to anything, concrete or abstract, but many things cannot be acted on physically, pointed at, disassembled, etc.

  The ability to refer to some things, e.g. macroscopic physical objects, requires simpler conceptual apparatus than the ability to refer to other things, e.g. to transfinite ordinals or to the referring capabilities of symbols.

  So finding the grid's topology is a research goal.

The space of sets of requirements:

'niche space' for biological and non-biological machines

Things researchers and designers need to think about:

- Types of entity
- Types of competence
- Types of combined competence-typePLUSentity-type
- Somatic ontologies: entities and abstractions internal to the individual

  E.g. multimodal sensorimotor relationships
- Exosomatic ontologies:

  Entities, processes, relations in the environment, or in other individuals.

  E.g. inferred properties of materials, hypothesised causal relations.
- Abstraction ontologies:

  Beliefs, goals, numbers, proofs, plans, theories, ...
- Types of complex competence based on combinations of simpler competences.

  E.g. seeing or imagining or describing a hippo swallowing a fly.

Recent concerns about embodiment, sensorimotor contingencies, symbol grounding, dynamical systems, situatedness, mainly arise from a consideration of only a subset of the requirements for a human- (or chimp- or crow-) like information-processing machine, namely the subset shared with microbes, insects, fishes, reptiles, etc. using only (or mostly) somatic ontologies.

The space of designs

We need a meta-theory of types of

- **mechanisms**
- **forms of representation**
- **types of sub-functions**
- **architectures in which the above can be combined**

Evolution produced a wide variety of which we still understand only a tiny subset.

For example there are ill-informed debates about whether things do or do not use representations, which need to be replaced with investigations into the variety of types of information acquired, manipulated, stored, combined, transformed, derived, used ....,

That requires investigating types of ways in which information structure can differ and can change – i.e. types of 'syntax'.

Individuals exist in ecosystems

In an ecosystem (e.g. our planet) there are myriad widely diverse sets of requirements (niches) which change over time.

There are myriad widely diverse types of designs, also changing, with many instances that develop over time.

Changes in the design of a type of organism can affect the niches for that sort of organism and other species.

Changes in niches often (not always) lead to changes in designs.

As designs or niches change there are different sorts of trajectories, evolutionary, individual, social/cultural, etc.

These trajectories can involve complex feedback loops (on different time scales) between designs and niches, with multiple complex, structured, 'fitness' relationships, at different levels.

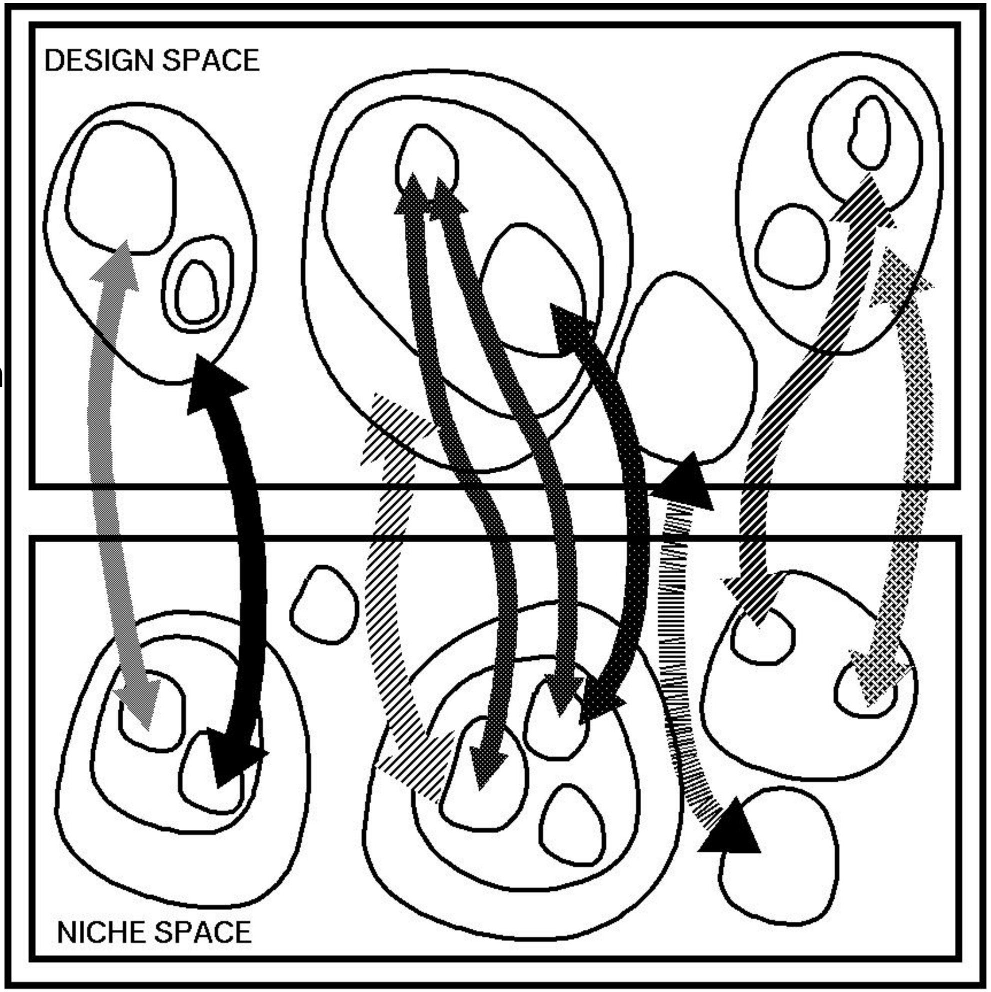

This can also apply to parts or aspects of organisms.

See http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#talk6

Don't ask for numeric fitness functions to assign a total ordering: not even if you are a UK government: There are structured relations between designs and niches.

Trajectories in design space and in niche space

There are different sorts of trajectories in both spaces:

- **i-trajectories:**
  Individual learning and development
- **e-trajectories:**
  Evolutionary development, across generations, of a species.
- **r-trajectories:**
  Repair trajectories: an external agent replaces, repairs or adds some new feature. The process may temporarily disable the thing being repaired or modified. It may then jump to a new part of design space and niche space.
- **s-trajectories:**
  Trajectories of social systems.

Some e-trajectories may be influenced by cognitive processes (e.g. mate-selection). We can call them c-trajectories (not shown separately).

All except r-trajectories involve continuously viable fully functioning working systems at every stage.

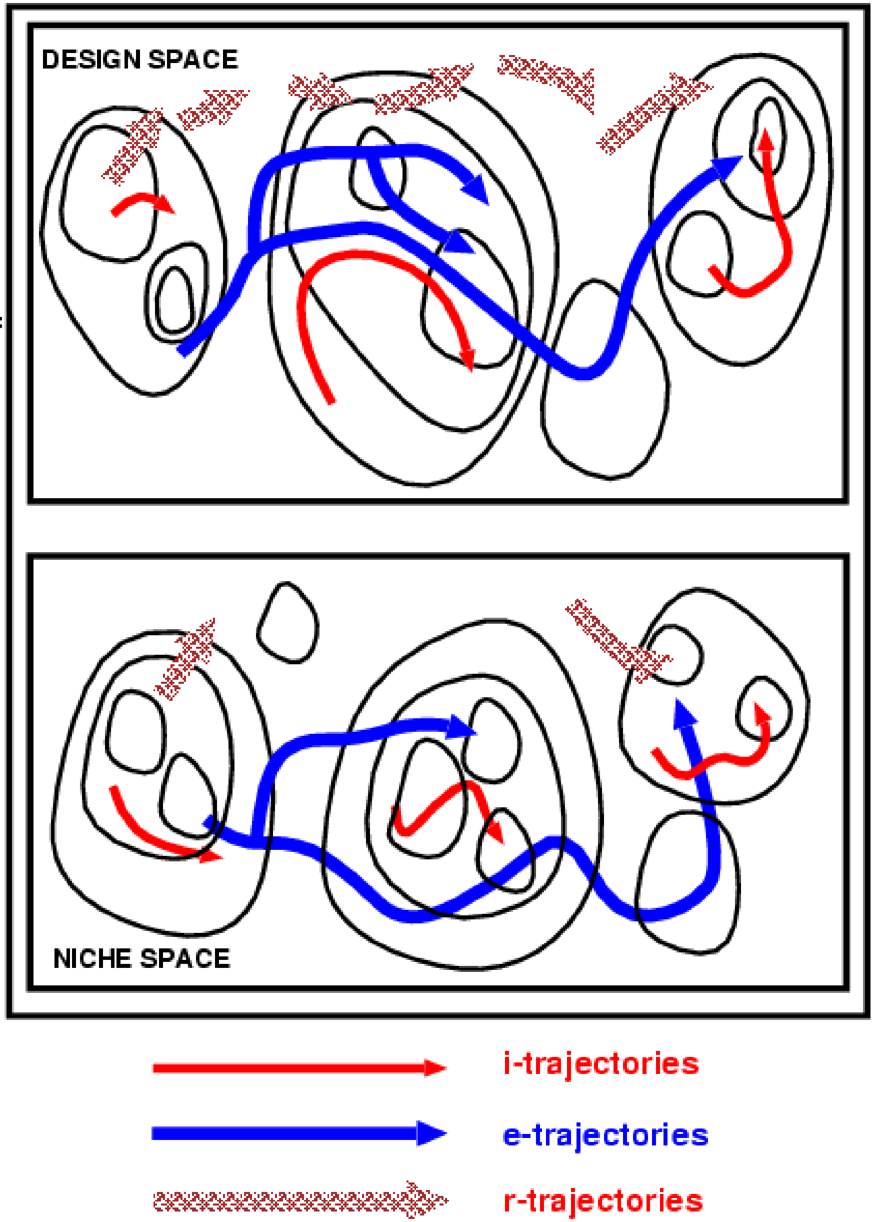

Meta-level requirements

Often requirements for advanced intelligent systems are specified using labels like ‘robustness’, ‘flexibility’, ‘creativity’, ‘autonomy’.

What do they mean?

In some sense we know what they mean – we can recognize instances and non-instances – but specifying explicitly what we are looking for is hard, partly because what we understand by each term is quite complex.

These are names for meta-requirements — analogous to higher order functions. They need to be given additional information to provide actual requirements.

E.g. a robust lawnmower is different from a robust operating system, or a robust planner.

Similar remarks apply to ‘efficient’, ‘flexible’, ‘creative’, ‘agile’, ‘autonomous’.

These meta-requirements are discussed at much greater length in this discussion paper:

`http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0701`

A First Draft Analysis of some Meta-Requirements for Cognitive Systems in Robots (January 2007)

Kinds of information-processing

All organisms are information processors.

What sorts of information they process, whether explicitly or implicitly, can be classified in different ways, e.g.:

- **The ontology used:**
  kinds of things, relations, processes, generalisations, etc. referred to.
- **What the information is used for:**
  e.g. online-control (servoing), selecting goals, recording specific facts, formulating generalisations, formulating questions, formulating hypotheses, making plans, executing plans, learning of many kinds, enjoying, suffering, ... etc.
- **How the information is represented,**
  e.g. only implicitly in transient patterns of activation during use, or in enduring, re-usable structures.
- **And also:**
  how it is matched, how it is combined, how it is modified, now it is used, how it is tested, whether it decays, etc., etc., etc.

The remaining slides

The remaining slides present some conjectures about natural intelligence that suggest possible future scenarios in which robots demonstrate complex mixtures of kinds of learning, some of them more sophisticated than others, including some kinds that are genetically determined (pre-configured) and some kinds that are produces as a result of interactions between pre-configured competences and the environment, physical and social.

Designing and evaluating such scenarios as part of a research roadmap involves hard methodological questions not discussed here.

Two kinds of ontology

All information-processing systems have direct access only to limited sources of information about states and processes within the system, e.g. sensor, motor and other internal signals and states. We call that 'somatic' information.

- **For some systems it suffices to detect and use a 'somatic' ontology: referring only to patterns and associations (at different levels of abstraction) found in those sources, including conditional probabilities linking input and output and other signals within the body,**
- **Sometimes it is necessary to refer beyond the available data to entities that exist independently of the information-processing system, and which have properties and relationships that are not definable in terms of patterns in sensed data: use of an 'exosomatic' ontology.**

The first sort can be compared with Plato's cave-dwellers.

Remember Plato’s Cave?

Plato likened us to inhabitants of a cave who could see only shadows on the cave wall, cast by entities out of sight – around a bend.

That’s what much current research is re-inventing, e.g. much work focused on sensorimotor contingencies.

All the information-processing being investigated in such research is concerned with patterns and associations between signals within the body whether going through the skin or between subsystems.

Such systems can never think about or refer to anything apart from the “shadows on the wall” – not what produces them:

The causes of shadows may have a very different structure from the shadows: different properties, different behaviours, different causal relationships.

Exosomatic ontologies in other animals?

Of course, it is possible that many organisms, perhaps the vast majority, are worse than Plato's cave-dwellers: they cannot represent, think about, ask questions about, find out about, what's happening outside them, for they can only handle information about their own sensorimotor signals.

Contrast the portia spider taking a complex spatial structure visually then following a route to a point above its prey,

Ants and wasps that use landmarks learnt and recognised visually.

Betty the famous hook-making crow, twiddling food with a stick, and apparently deciding that she needs to go round to the side to see more clearly what's going on.

(In one of Jackie Chappell's videos.)

Why do these capabilities require exosomatic ontologies?

Importance of exosomatic ontologies

Exosomatic ontologies can represent things independently of how they are sensed or acted on. ‘Concept empiricism’ (recently reinvented as symbol-grounding theory) deems that impossible. But Kant (in 1780) argued otherwise against Hume.

Use of exosomatic (not just amodal or multimodal) ontologies is a significant feature of human information processing.

- **This is commonplace in science: genes, neutrinos, electromagnetic fields, and many other things are postulated because of their explanatory role in theories, despite never being directly sensed or acted on.**
- **Does this also go on in learning processes in infants and hatchlings that discover how the environment works by playful exploration and experiment?**
- **Are ‘ontologies’ that refer beyond the sensor data also set up in the genome of some species whose young don’t have time to go through that process of discovery but must be highly competent at birth or hatching? (precocial species)**
- **Could the portia spider perform her amazing feats if she had only representations of her sensory and motor signals**

  See `http://www.freerepublic.com/focus/f-chat/1640513/posts`
- **Is there anything in common between the different ways ontologies get expanded in biological systems (e.g. in evolution, in development, in social processes)?**
- **This relates to questions about what a genome is, and about varieties of epigenesis.**

This is more than philosophy. There are engineering implications: use of exosomatic ontologies can reduce search spaces, allow much greater generalisation, make possible imitation and ability to help others.

Two sorts of ontology extension

There are two importantly different ways in which an ontology can be extended.

- **Definitional extension:** introducing a new label for a complex form of representation expressible using existing resources.
- **Substantive extension:** introducing a new symbol that is not definable in terms of pre-existing resources. (How? See history of science and maths for clues: new entities are assumed to be causally related to old ones.)
- According to many tempting theories of meaning substantive extension is impossible — yet it has happened throughout the history of science and mathematics.
- I claim it also happens during human development: concepts used by a professor of theoretical physics, like most other adult concepts, are not definable using resources in a newborn infant.

  So the thesis of Fodor's The language of thought is wrong. There is no innate language adequate to express everything a child learns.
- It may also happen to some other animals, though many animals (the 'precocial' species) are born with almost everything they need to know ready for use – e.g. most invertebrates and many others, including deer born ready to run with the herd long before there's time to learn.

  See Sloman & Chappell IJCAI 2005, and Chappell & Sloman IJUC (to appear)

Implications for theories of meaning

The existence of precocial species refutes 'symbol-grounding' theory

(One version of 'concept empiricism' – the theory that all meaning has to be derived by processes of abstraction from sensory experiences, which is clearly not required for precocial species that are competent at birth).

In our IJCAI paper we distinguish two sources of meaning

- the structure of a theory in which 'undefined terms' occur
  (where the structure limits the class of possible models/interpretations)
- links to sensing and acting (e.g. through tests and predictions)
  (where the links – e.g. Carnapian 'meaning postulates' further reduce the set of possible interpretations, tethering the interpretation – though there is always residual indeterminacy.)

The second picture seems to represent how terms in scientific theories get their meaning, i.e. largely from the structure of the theory, which constrains possible models. So why not concepts in toddler theories?

Compare 20th century philosophy of science after crude empiricism was shown to be wrong: Popper, Carnap, Hempel, Pap, ...

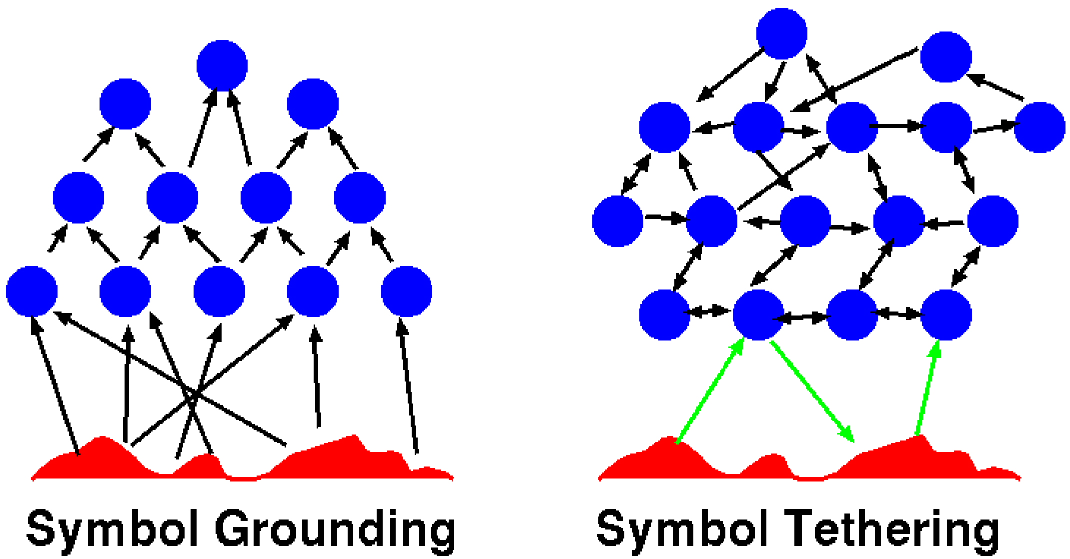

Pre- and Meta-configured competences

Every biological individual is a product of two sources of information

- **The genome:** information produced by evolution plus fertilisation processes (made available in fertilised eggs, seeds, or ‘parent’ cells, etc.)
- **Information from the environment:** acquired/produced/derived during development (starting from fertilisation/germination of eggs)

In most species behavioural competences derived from those two sources are common across all individuals, with minor differences based on physical differences and adaptation/calibration:

> Environmental variations have little impact on some genetically preconfigured competences – unlike fine details of behaviour using the competences e.g. when a bee lands on a flower.

In a few species the second source of information has much greater impact on competences, and individuals can develop very different competences in different environments, often very rapidly and without needing lengthy training (conditioning) regimes: this requires genetically determined meta-competences, i.e. competences to create competences by interacting with the environment.

We say that in the first case the competences are **preconfigured** and in the second case they are **meta-configured.**

Since the second case often involves ontology extension, we can say that for some species their ontologies are preconfigured, and for others at least parts of the ontology are meta-configured: created during development.

Explaining evolution of behaviours

As the figure indicates, the connection between DNA and structure of an organism is more direct than connections between DNA and behaviours of an organism.

Physical structures develop via sequential physical and chemical processes. Slight changes in the DNA may produce small changes in the resulting structure, but some are catastrophic.

Behaviours are produced less directly (even when preconfigured). Development produces structures (e.g. in nervous systems) encoding behaviours to be produced when needed. Even 'innate' behaviours, e.g. pecking for food, migrating, imprinting, running with the herd, are not rigidly determined collections of motor signals but high level patterns that are finely controlled by sensory input (e.g. jumping over or avoiding obstacles).

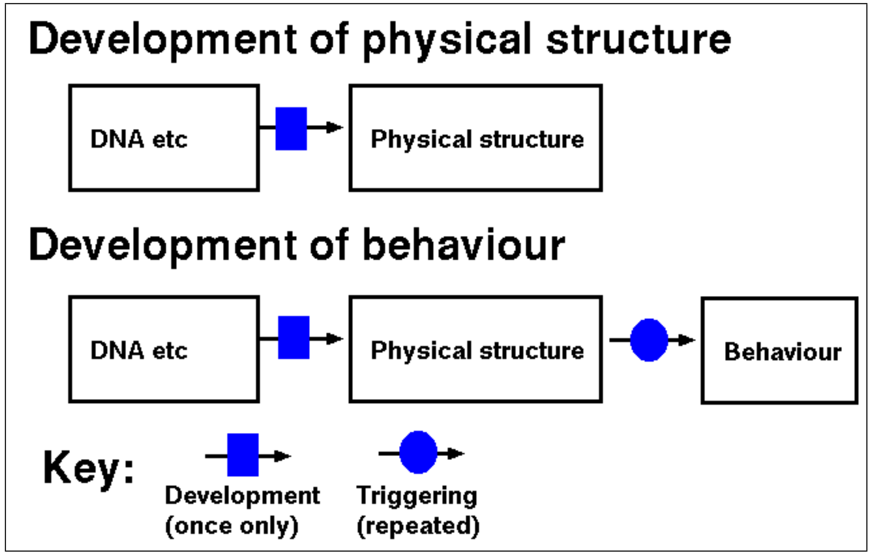

Although DNA changes can produce advantageous behaviour changes the chances seem infinitesimal, making evolution of behaviour a very slow process. Can it be sped up?

Conjecture: evolution discovered how to produce larger re-usable behaviour-generating genetic building blocks, with separate parameters. Then mutation of parameters would allow 'behaviour space' to be searched more safely.

If this were combined with mechanisms that can combine separate behaviours to produce new complex behaviours then different sub-behaviours could evolve, producing minor changes in macro behaviours. Engineers often use such modular decomposition to aid design and maintenance.

This idea is also used in Genetic Programming (GP), a form of evolutionary computation.

From preconfigured to meta-configured competences

Another development would allow some of the evolved behaviours to operate on internal (e.g. neural) rather than on external structures. This would be useful for many different purposes, e.g. changing a context involved in the control of behaviours, so that behaviours are different when food is needed, when a predator might be near, etc.

An even more useful internal behaviour could construct new competences, based on what has been found in the environment. This requires meta-competences.

Some of the competence developers might have their actions delayed until after the organism has become active and is able to explore the environment.

If previously mentioned mechanisms for combining old competences to form new ones are available, new powerful combinations of environmentally determined competences could be formed, tested, and stored for re-use if found valuable.

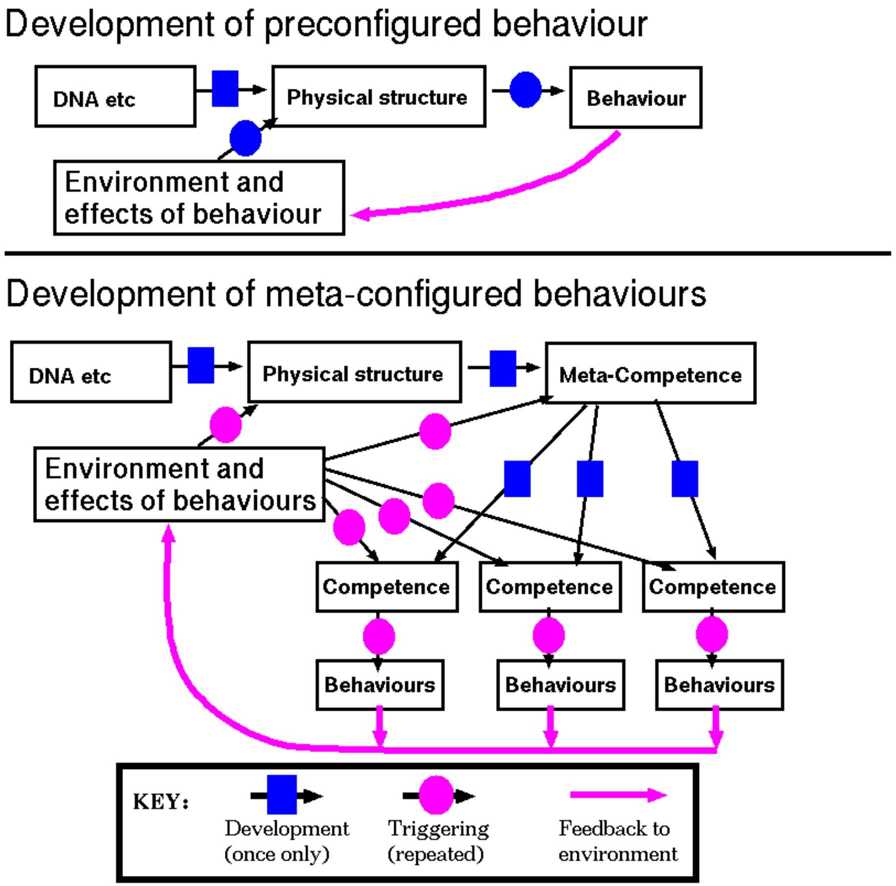

If applied internally this would allow new meta-competences to develop, e.g. enhancing the learning capability to suit the environment, producing even more varied and powerful end results.

Routes from genome to behaviours

This shows more ways in 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

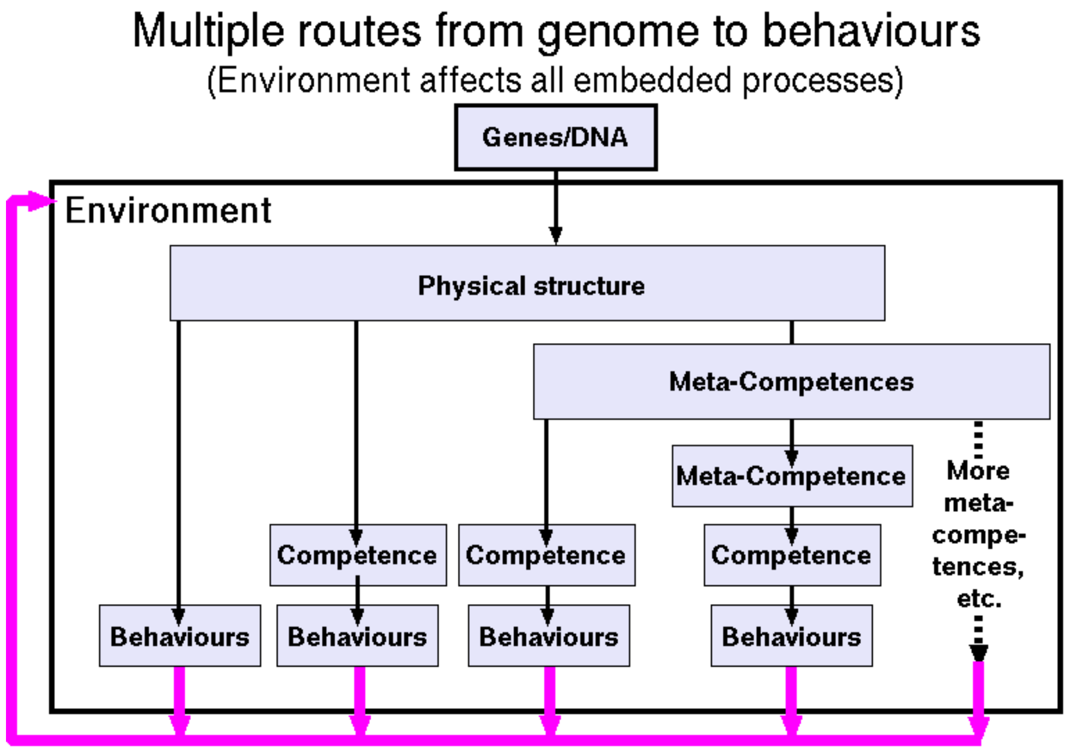

Insofar as the behaviours influence the environment there can be complex feedback loops. Competences and behaviours further to the right may use several 'simpler' competences and behaviours to the left.

Strictly, the construction of some competences should be shown 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 of the processes will involve continuous control.
In deliberative components much will be discrete.

Meta-competences in different architectural layers

The diagram on the previous page can be seen as representing a general schema (the meta-competence schema) which could be implemented in different ways in different organisms and machines, and also in different subsystems of the same organism or machine.

For example if an animal has an architecture that includes different layers such as reactive, deliberative, or meta-management layers, or the more finely distinguished layers described in Minsky's *The emotion machine*, then different meta-competences might be implemented in different layers, extending the capabilities of those layers as a result of various kinds of interactions with the environment.

- **In particular in the lower level portions of a reactive subsystem where all the mechanisms are based on continuous control loops a meta-competence might allow the construction of a new subsystem with its own control loops for performing a new task (e.g. standing on one leg, hopping, catching a moving object, etc.) (Compare the H-Mosaic model).**
- **A different sort of meta-competence in the reactive layer might allow creation of new reactive plans, for doing things by triggering sequences of discrete steps (compare Nilsson's Teleoreactive systems).**
- **A deliberative meta-competence might be able to generate new kinds of planning capabilities, e.g. planning actions that involve going to locations that are not currently visible, planning actions that require some tasks to be done in parallel, planning actions that involve getting others to cooperate, etc.**
- **Meta-management meta-competences might allow new forms of self-monitoring and self-control to be learnt.**

Eventually very complex architectures will be required

– with many different sorts of components and very rich interactions, including different layers of abstraction in perception and action, and global alarm mechanisms.

H-Cogaff:

This human-inspired collection of requirements could be relevant to future robots.

Arrows represent information flow (including control signals)

If meta-management processes have access to intermediate perceptual databases, then this can produce self-monitoring of sensory contents, leading robot philosophers with this architecture to discover "the problem(s) of Qualia?"

'Alarm' mechanisms can achieve rapid global re-organisation.

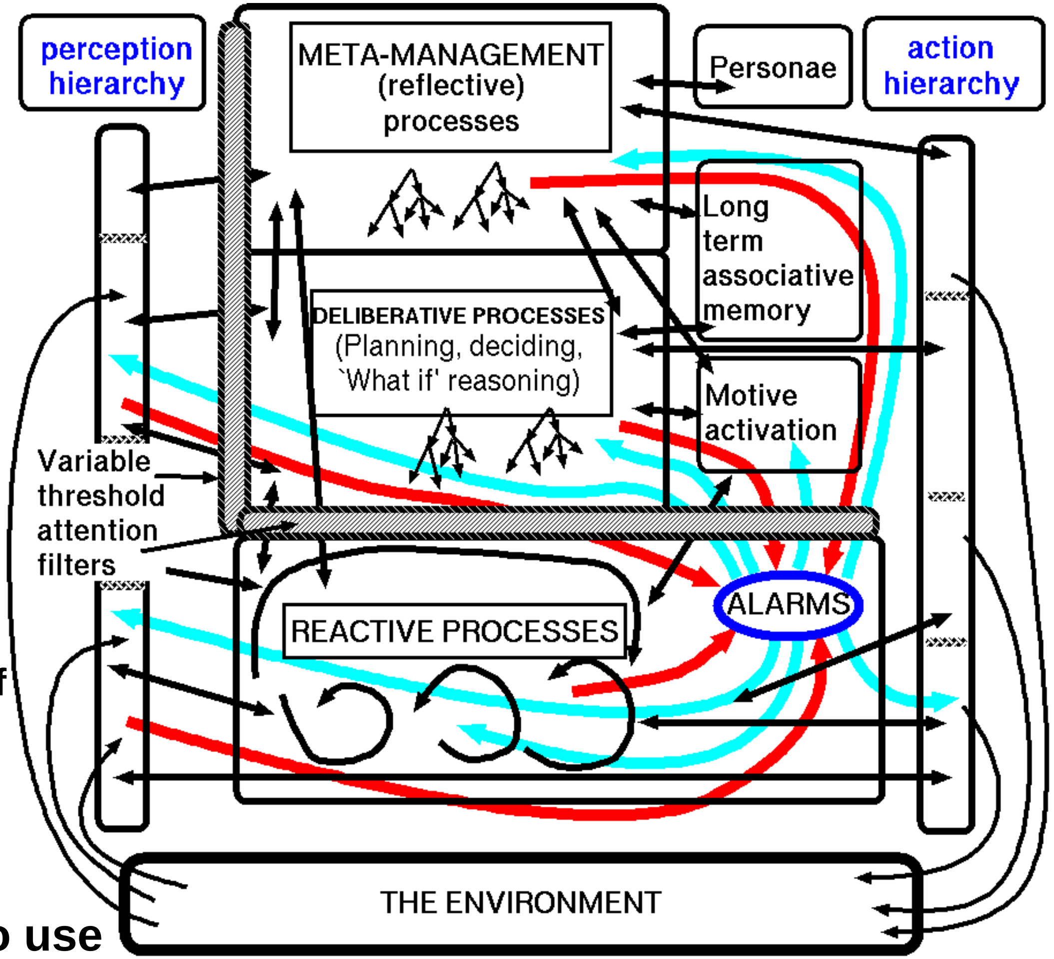

Meta-management systems need to use meta-semantic ontologies: they need the ability to refer to things that refer to things.

Some Implications

Within this framework we can explain (or predict) many phenomena, some of them part of everyday experience and some discovered by scientists:

- **Several varieties of emotions: at least three distinct types related to the three layers: primary (exclusively reactive), secondary (partly deliberative) and tertiary emotions (including disruption of meta-management) – some shared with other animals, some unique to humans. (For more on this see Cogaff Project papers)**
- **Discovery of different visual pathways, since there are many routes for visual information to be used.**
  (See talk 8 in http://www.cs.bham.ac.uk/˜axs/misc/talks/)
- **Many possible types of brain damage and their effects, e.g. frontal-lobe damage interfering with meta-management (Damasio).**
- **Blindsight (damage to some meta-management access routes prevents self-knowledge about intact (reactive?) visual processes.)**

This helps to enrich the analyses of concepts produced by philosophers sitting in their arm chairs: for it is very hard to dream up all these examples of kinds of architectures, states, processes if you merely use your own imagination.

Implications continued ....

- **Many varieties of learning and development**
  (E.g. "skill compilation" when repeated actions at deliberative levels train reactive systems to produce fast fluent actions, and action sequences. Needs spare capacity in reactive mechanisms, (e.g. the cerebellum?). We can also analyse development of the architecture in infancy, including development of personality as the architecture grows.)
- **Conjecture: mathematical development depends on development of meta-management – the ability to attend to and reflect on thought processes and their structure, e.g. noticing features of your own counting operations, or features of your visual processes.**
- **Further work may help us understand some of the evolutionary trade-offs in developing these systems.**
  (Deliberative and meta-management mechanisms can be very expensive, and require a food pyramid to support them.)
- **Discovery by philosophers of sensory 'qualia'. We can see how philosophical thoughts (and confusions) about consciousness are inevitable in intelligent systems with partial self-knowledge.**

For more see papers here: http://www.cs.bham.ac.uk/research/cogaff/

REMAINING SLIDES

The remaining slides point to examples of how un-obvious kinds of competence develop in very young children.

Often these are hard to see in real life – because things happen so quickly – but become visible in a video that is watched several times.

Many scenarios can be based on child development

Infants may not see causal relations adults experience as obvious

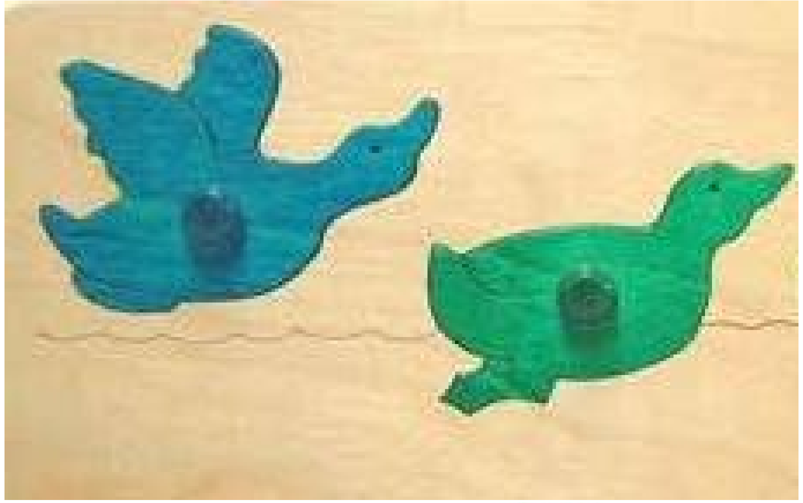

A child C learns that she can lift a piece out of its recess, and generates a goal to put it back, either because C sees the task being done by others or because of an implicit assumption of reversibility. At first, even when C has learnt which piece belongs in which recess there is no perception of unaligned boundaries, so there is only futile pressing.
Later C may succeed by chance, using nearly random movements, but the probability of success with random movements is very low. Why?

Memorising the position and orientation with great accuracy might allow toddlers to succeed: but there is no evidence that they can memorise precise orientation and location of an irregular shape. Can you?

Stacking cups simplify the cognitive task, partly through use of symmetry, partly through sloping sides — so they are much easier.

Eventually C's (still pre-linguistic) ontology includes something like 'boundary' and 'alignment'. Only then can she learn that if the boundaries are not aligned the puzzle piece cannot be inserted — probably some time after learning how to cope with symmetric stacking cups.

Conjecture: many changes in perception and action competence require the child to extend its ontology for representing objects, states and processes in the environment. The enriched ontology is used by the child's pre-linguistic perception and problem-solving mechanisms. HOW?

Example: watch toddlers and children, and ask: how could we design something that does that?

Yogurt can be food for both mind and body in an 11 month baby.

Video available at `http://www.cs.bham.ac.uk/~axs/fig/yog.mpg`

J discovered he could transfer yogurt to his leg, and picture 1 shows him trying to transfer more. His ontology seems not yet to include the orientation of the bowl, and its effects. Picture 2 shows J trying to place a piece of yogurt picked up from the carpet into the spoon, prior to transferring it into his mouth. Picture 3 shows him trying, and failing, to put another piece of yogurt on the carpet, still apparently not experiencing the orientation of the bowl. Later J manages to transfer his grasp of the spoon handle from one hand to another. What mechanisms would allow a robot to learn like this?

J seems to experiment with his hands, legs, spoon, yogurt and the carpet. He sees opportunities and tries them out, notices things and tries to recreate them (often unsuccessfully). His ontology is quite rich but some gaps are evident.

He probably doesn't know he is doing all this! That would require a sophisticated self-monitoring architecture that is probably still being constructed. A baby is not just a tiny adult!

Not understanding hooks and rings at 19 months

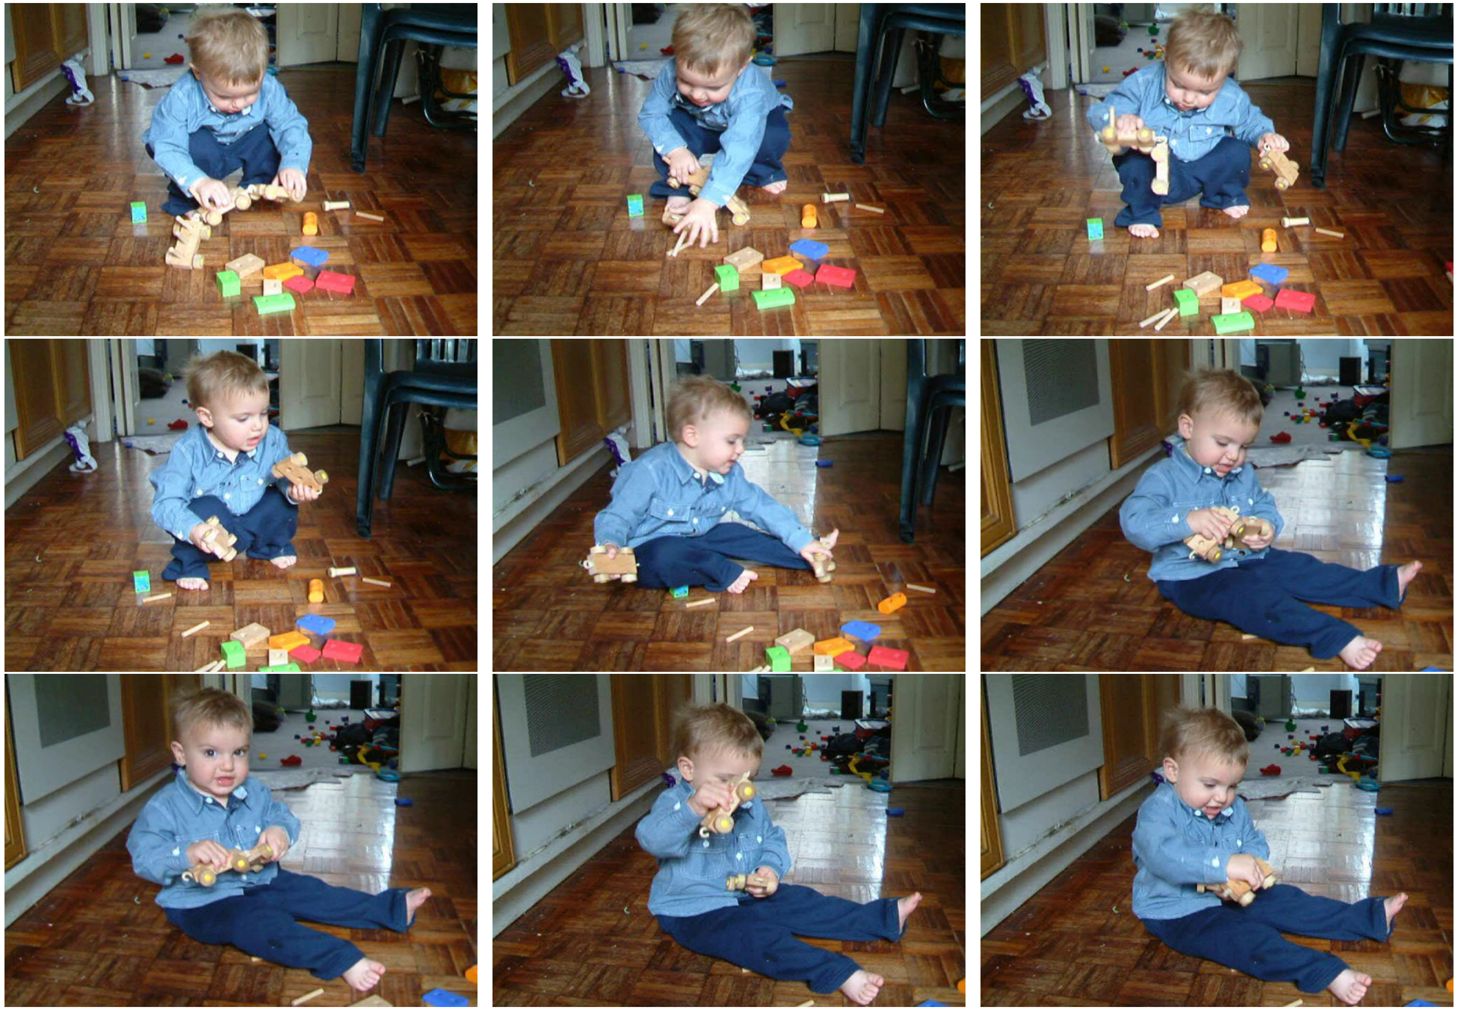

1: Lifting two trucks makes the third disengage. 2-3: He picks it up with his left hand & shakes off the hanging truck with his right. 4: He notices the blank end & puts the truck down, rotating it. 5: He makes a complex backward move from crouching to sitting – while leaning forward to pick up the rotated truck. 6: He sees two rings. 7-9: He tries to join the rings, ignoring the hook, fails and gets frustrated, bashing trucks together and making an angry sound.

See the video `http://www.jonathans.me.uk/josh/movies/josh34_0096.mpg`

Within a few weeks, he had learnt to see and use the hook-affordances. How? (Nobody saw how.)

More on Blanket and String

If a toy is beyond a blanket, but a string attached to the toy is close at hand, a very young child, whose understanding of causation involving blanket-pulling is still Humean, may try pulling the blanket to get the toy.

Later the child may either have extended the ontology used in its conditional probabilities, or have learnt to simulate the process of moving X when X supports Y. (Kantian causation.)

As a result of either, he doesn't try pulling the blanket to get the toy lying just beyond it, but uses the string.

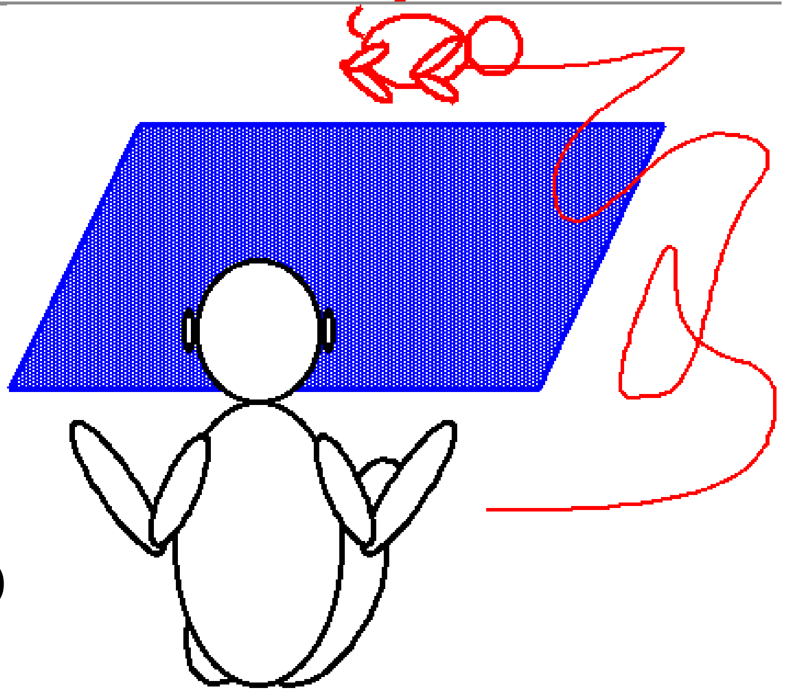

If he has a Kantian understanding he can solve a wider range of problems.

However the ontology of strings is a bag of worms, even before knots turn up.

Pulling the end of a string connected to the toy towards you will not move the toy if the string is too long: it will merely straighten part of the string. The child needs to learn the requirement to produce a straight portion of string between the toy and the place where the string is grasped, so that the fact that string is inextensible can be used to move its far end by moving its near end (by pulling, though not by pushing). Try analysing the different strategies that the child may learn to cope with a long string, and the perceptual, ontological and representational requirements for learning them.

Perceiving diverse multi-strand processes

Ontological and representational requirements for experiencing commonality.

What is common to performing the same sort of grasping action (bringing two surfaces together to clasp part of an object) using

- **your left hand,**
- **your right hand,**
- **two hands grasping a large object,**
- **one hand and your chest,**
- **your teeth?**

How much of what you need in order to see and do those things also applies to thinking about past or future occurrences, or similar actions performed by others?

And much much more.....

See the conjecture (possibly) new theory of vision
`http://www.cs.bham.ac.uk/research/projects/cosy/papers/#pr0505`

To be continued

Although I've focused on requirements for modelling and explaining competences of humans (young and old) and other kinds of animals all of this is relevant to a vast array of potential applications.

Many other projects and proposals refer to practical applications in which intelligence, flexibility, human-like competences are deployed.

I believe that most such projects will fail if they do not start from the existence proofs – actual working systems – in order to find out in more detail what their competences really are (not just what they appear superficially to be.

That does not imply that exactly the same mechanisms are required in order to implement new versions of those competences: that is a separate (and open) question.

Maybe I'll extend this document later, if I have time, especially if you send me suggestions, criticisms, corrections. ...

Aaron Sloman

A.Sloman@cs.bham.ac.uk

http://www.cs.bham.ac.uk/˜axs/