Putting the Pieces of AI Together Again

Aaron Sloman,
http://www.cs.bham.ac.uk/˜axs/
and the CoSy team.
University of Birmingham, United Kingdom

Fragmentation of AI: (1) Growth

Growth, and discovery of new deep problems led to fragmentation of the research community.

- **During first two decades people in AI mostly knew about one another's research and went to the same conferences, covering all subfields, e.g IJCAI.**

  There were exceptions of course

- **As numbers of AI researchers grew and more and more problems were seen to have deep sub-problems, the field naturally fragmented, with people working on more narrowly-focused problems as the years progressed.**
- **Individual researchers found it hard enough to keep up with the latest developments in their own area, without also trying to keep up with what was happening in other subfields.**
- **So they formed sub-communities studying their own problems, using their own methods, going to their own conferences, publishing their own journals etc.**
- **Moreover, most people come into AI only as graduates in another discipline so they have to pick up an accelerated education to get to PhD level,**
  and often that means reading only things their supervisor read – a much narrower range than they would get from e.g. an undergraduate degree in AI taught by many people.

Fragmentation of AI: (2) Mis-diagnosis

Another cause of fragmentation:

- The failures of optimistic (extravagant) predictions for AI, were mis-diagnosed as being due to
  - the wrong tools,
  - the wrong forms of representation,
  - the wrong mechanisms, or architectures.
- So waves of fashion surged and ebbed, focusing on alternative tools, representations, mechanisms, etc., for instance, connectionism, evolutionary computation, dynamical systems, situated cognition, use of reactive systems, subsumption architectures, ...
- This led to more fragmentation and mutual ignorance between factions.
- It also led to wasteful competition for resources, when cooperation might have been more productive, because AI needs diverse forms of representation, algorithms, architectures and tools.
- As in other fields, e.g. study of consciousness, people tend to define the problems in terms of what their own theories and tools can illuminate.

Can we collaborate in finding out what the problems really are: e.g. what needs to be explained?

Can researchers start to collaborate? First agree on what the problems are.

I heard an engineer from industry at an AI robotics research planning workshop say:

Every few years I hear AI researchers say what they are going to do in the next few years, and it never changes.

Do you recognise that phenomenon?

Why does that happen? Some suggestions:

- **People think they know what humans can do, e.g. what vision is, but they grossly underestimate its complexity, and embark on projects with faulty requirements analysis.**
- **This leads to over-optimistic predictions and over-hasty selection of strategies and designs.**

  Example: frequent talk about the need to 'scale up' to human-level competence, ignoring the fact that humans scale up very badly – though they do 'scale out': combining old competences in new, productive, ways.

- **Progress will require collaboration on the task of identifying**
  - **requirements and goals**
  - **gaps in our knowledge in relation to those requirements and goals.**

Partial Remedy: Build a Roadmap Backwards

In the EU-funded CoSy robotics project some of us have been trying to devise a way of countering the fragmentation and deepening our understanding of the problems by:

- **creating a network of scenarios of varying degrees of difficulty and depth,**
- **starting from detailed descriptions of machines doing things that AI systems are nowhere near doing at present but might do in the distant future,**
- **focusing on things humans (e.g. young children) and some other animals can do now, to provide existence proofs,**
- **working back from long-term difficult scenarios through medium-term easier scenarios to near-term that are less and less complex and demanding**
- **some fragments of which look as if they might be doable in the next 3 to 5 years.**

- **I.e. we need to get very clear about requirements before producing designs**
- **analysis of scenarios can help focus attention on detailed requirements**
- **however there will be feedback loops involving scenarios, requirements, designs and (partial or inadequate) implementations**

This can produce a partially ordered collection of requirements for future systems, ordered by dependency relations and difficulty, where early items are not very far from what we can do now.

Picture of a Research Roadmap

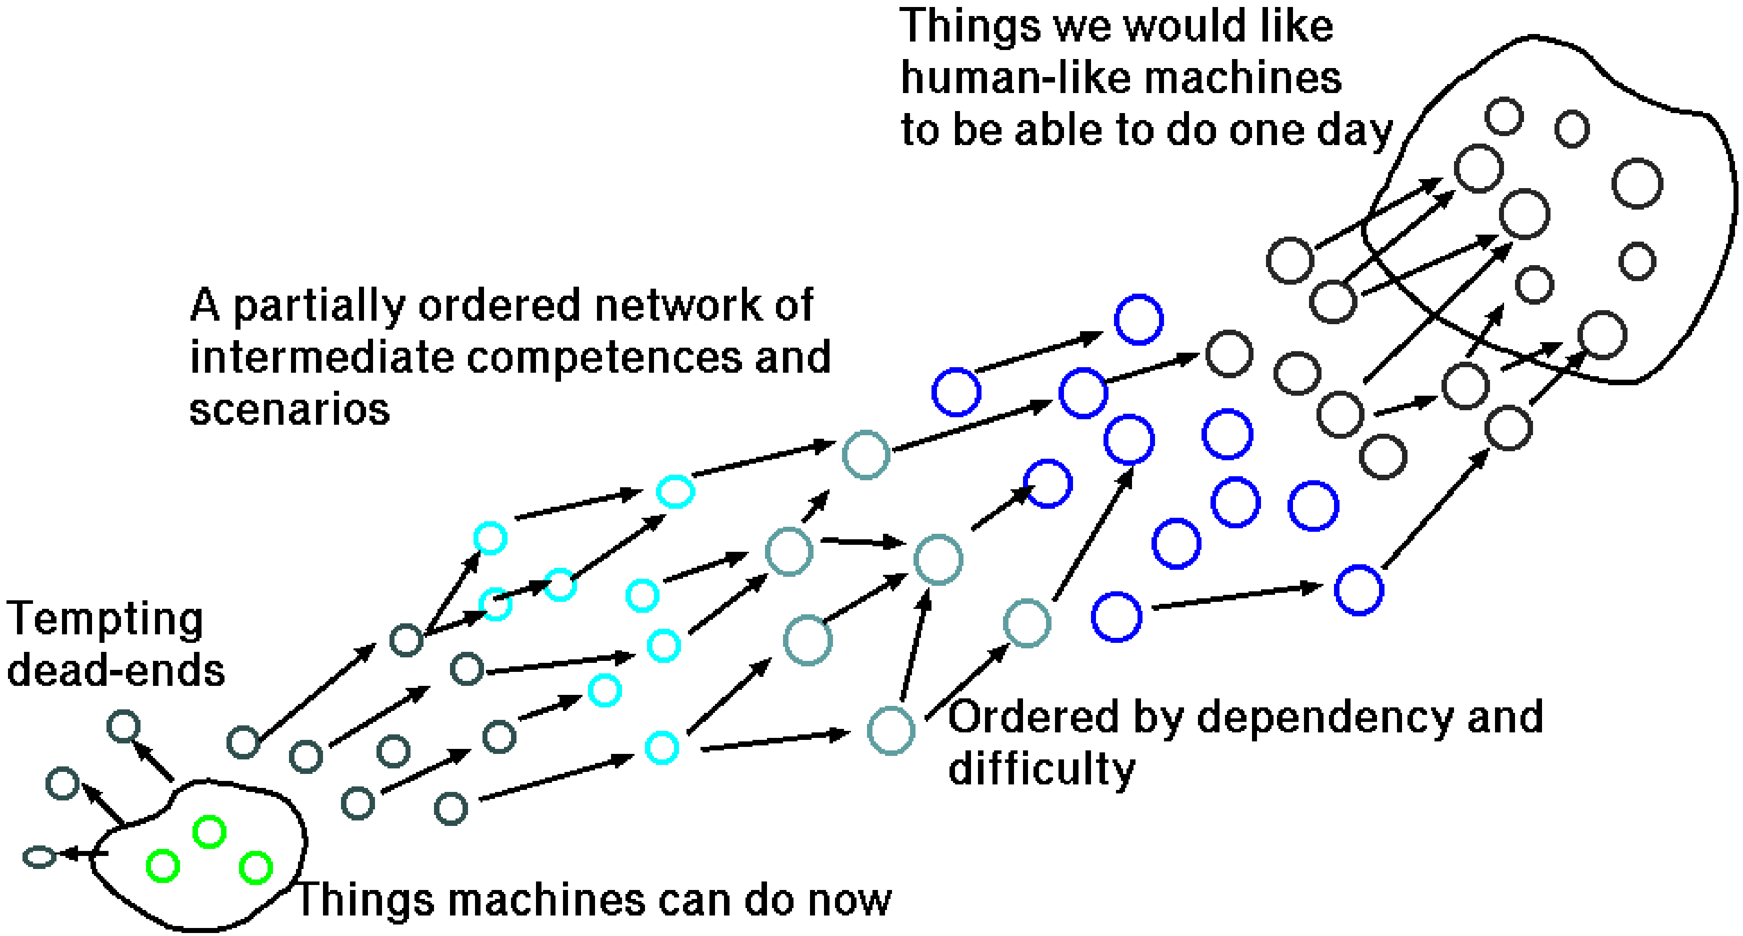

Forward chaining research asks: how can we improve what we have already done?

Backward chaining research asks: what is needed to achieve our long term goals?

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).

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 (implying time required to produce working systems).

(Actually a more complex topology than a rectangular grid is required.)

For more detail see:

http://www.cs.bham.ac.uk/research/cogaff/gc/aisb06/
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.

Example: Vision

Applying this methodology to 3-D manipulation requirements for a future domestic robot revealed new requirements for vision.

- **It is tempting to regard vision as primarily being concerned with information about the spatial structures and affordances in the environment, at many levels of abstraction.**

- **A scenario for a future robot performing domestic chores pointed to the requirement for vision to be concerned primarily with *processes* and *affordances* rather than *structures*,**

  A static scene is a process in which all velocities, etc. are zero.

- **Seeing may need to involve concurrent processes at different levels of abstraction in partial registration with the optic array.**

  Seeing one object being placed on another includes concurrently seeing:

  - **processes involving changing topology (e.g. A moves into the convex hull of B).**
  - **processes with qualitative information (e.g. A is moving towards B's left side),**
  - **processes with precise metrical information (e.g.required for grasping)**

  All in partial registration with each other and the optic array.

  See http://www.cs.bham.ac.uk/research/projects/cosy/papers/#pr0505

Example: Types of Learning & Representation

Compare:

- **kinds of learning required when a new robot is delivered**
- **varieties of learning and development in different species of animals.**

- **Most species are born or hatched with all the competences they will need (foals can run with the herd) and there is only minor adaptation in the life of an individual**
- **Humans, and a few other altricial species, go through a process of creative and playful exploration and experiment as a result of which they somehow decompose what they find into collections of (nearly) orthogonal competences that can be recombined in new ways to solve new problems or perceive unfamiliar structures and processes.**

  See: http: www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0601
  'Orthogonal recombinable competences'.

- **To reduce problems of combinatorial explosions of 'somatic' sensorimotor patterns, evolution seems to have developed (in some species) abstract 'exosomatic', or objective, ontologies and forms of representation that ignore details of sensory inputs and motor outputs.**

  This made possible planning future actions, and reasoning about past or hidden processes where precise physical details and therefore sensory and motor signals are not known.

  Example: grasping can be represented in terms of relative movements of two surfaces and something that comes between them – leaving out details of how it looks or feels, or the motor signals involved.

  (Perhaps the real function of 'mirror' neurons is abstraction.)

  Developmental psychology provides many examples to go into the roadmap.

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. 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!

Hypothesis

Alongside the innate physical sucking reflex for obtaining milk to be digested, decomposed and used all over the body for growth, repair, and energy, there is a genetically determined information-sucking reflex, which seeks out, sucks in, and decomposes information, which is later recombined in many ways, growing the information-processing architecture and many diverse recombinable competences.

HOW ???

What mechanisms explain this?
What exactly changes as a result?
How does the child's ontology grow?
What are the nature-nurture trade-offs?

No current AI system comes close.
AI vision systems are restricted to much simpler tasks,
e.g. recognition, tracking, prediction,
point-wise reconstruction of surfaces, etc..

We cannot do it all from birth

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?

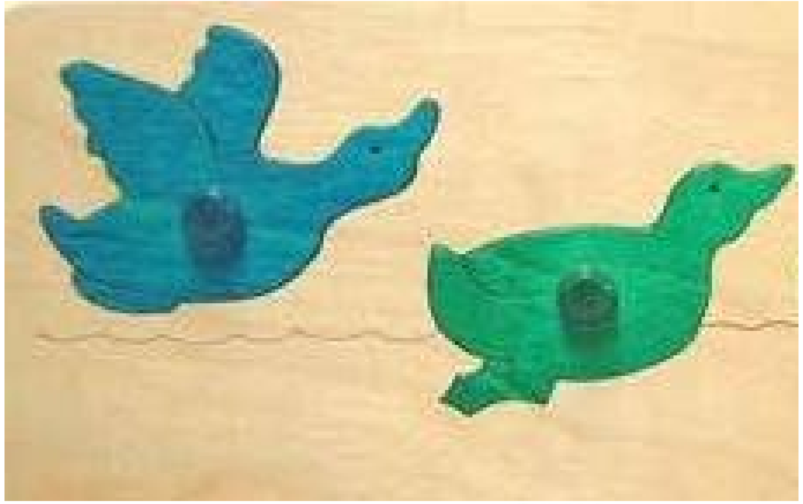

Memorising the position and orientation with great accuracy might allow toddlers to succeed: but there is no evidence that they can memorise precise orientation and location of an irregular shape. Can you?

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?

Failing to deal with hooks at 19 months

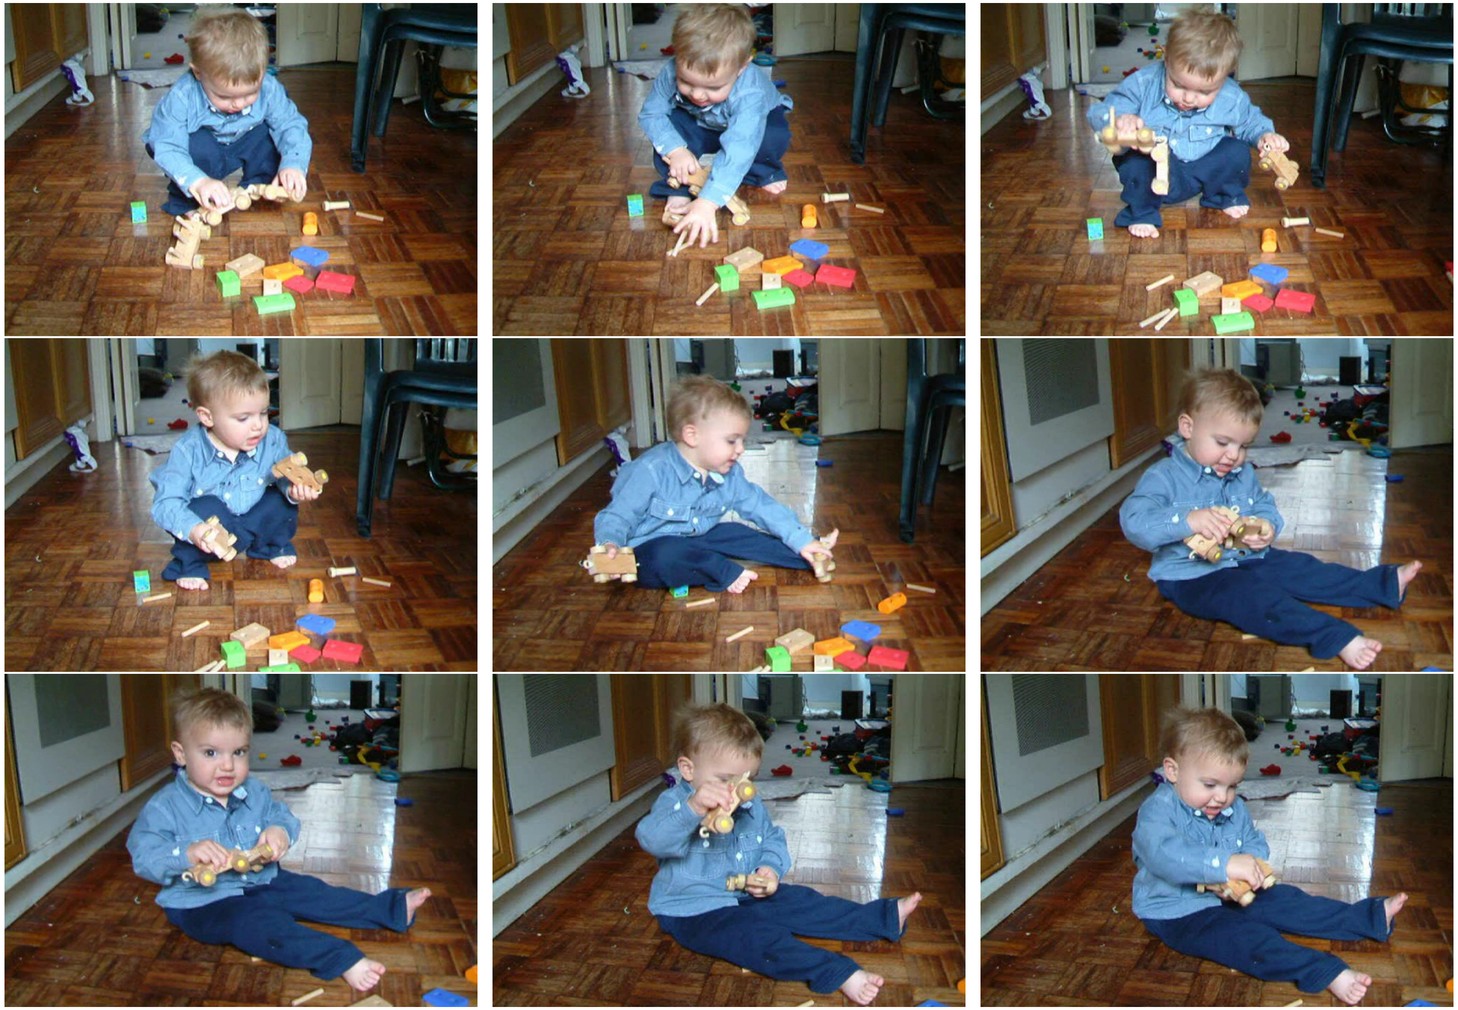

1: Lifting two trucks makes the third disengage. 2-3: He picks it up with his left hand & shakes off the hanging truck with his right. 4: He notices the blank end & puts the truck down, rotating it. 5: He makes a complex backward move from crouching to sitting – while leaning forward to pick up the rotated 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.)

A partially ordered network of stages

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

  (Compare: Waddington's epigenetic landscape)
- I don't know how many different things of this sort have to be learnt, but it is easy to come up with hundreds of significantly different examples.
- What can be learnt keeps changing from one generation to another: provision of new kinds of playthings based on scientific and technological advances is a major form of communication across generations. (Not for all species!)

  CONJECTURE: in the first five years, using a genetically determined bootstrapping mechanism, a child learns to run at least least hundreds, possibly thousands, of different sorts of simulations, using different ontologies – with different materials, objects, properties, relationships, constraints.

  This involves acquiring two kinds of causal knowledge: Humean and Kantian.

Example: Humean causation

Two gear wheels attached to a box with hidden contents.

Can you tell by looking what will happen to one wheel if you rotate the other about its central axis?

Not really!

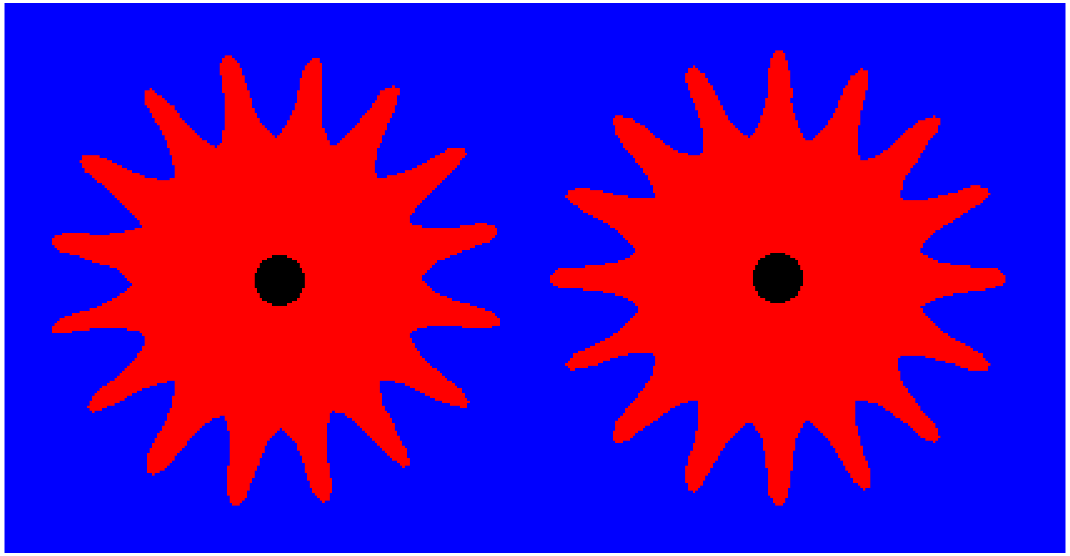

- **You can tell by experimenting: you may or may not discover a correlation – depending on what is inside the box.**
- **In more complex cases there might be a knob or lever on the box, and you might discover that which way the second wheel rotates depends on the position of the knob or lever. (Compare learning about gears in driving a car.)**
- **In still more complex cases there may be various knobs and levers, modifying one another's effects through hidden mechanisms. There could also be motors turning things in different directions, competing through friction devices, so that the fastest one wins.**
- **Humean causation involves unintelligible sets of conditional probabilities (e.g. Bayesian nets?).**

Knowledge of Humean causation amounts to having a convenient summary of raw data.

Example: Kantian Causation

Different gear wheels (still pivoted at their centres):

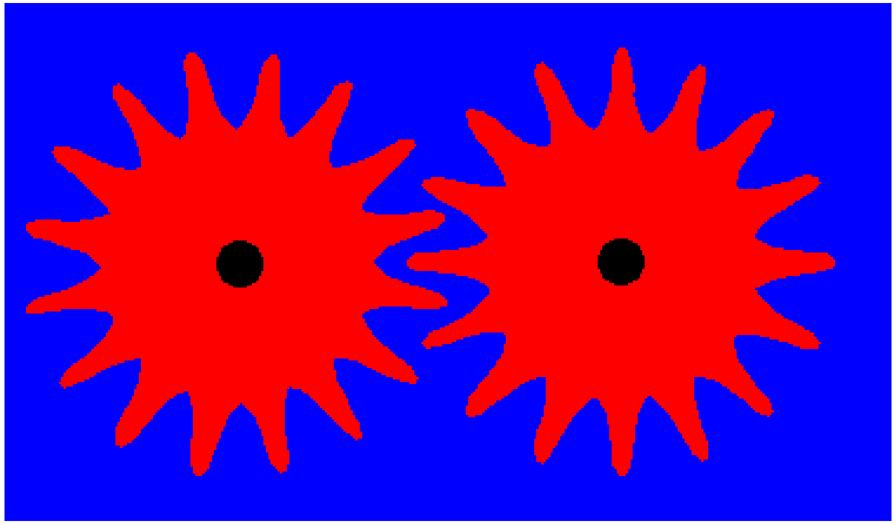

You (and some children) can tell, by looking, how rotation of one wheel will affect the other.

How? Assuming rigidity and impenetrability, you can simulate rotations and 'inspect' the consequences.

What you can see includes this:

As a tooth near the centre of the picture moves up or down it will come into contact with a tooth from the other wheel. If both are rigid and impenetrable, then if the first tooth continues moving, it will push the other in the same direction, causing its wheel to rotate in the opposite direction.

(I am not claiming that children need to reason verbally like this: consciously or unconsciously. It's a deeper, pre-linguistic ability at first. Probably some other animals grasp Kantian causality, e.g. nest-building birds?)

We seem to be able to run simulations using not only perceived shape, but also unperceived constraints: in this case rigidity and impenetrability.

Such constraints must be part of the perceiver's ontology and used in the simulations, for the simulation to be deterministic. How do 'rigidity' & 'impenetrability' enter a toddler's, or chimp's, or crow's ontology? The constraints and the processes using the constraints need not be conscious, or expressed in a logical, or verbal, form.

Reuniting factions as well as fragments

Very detailed observation of actual human and animal competences constantly informed by the design stance (posing questions about architectures, mechanisms, forms of representation, ontologies, etc.) can help us generate hundreds of target scenarios of varying complexity and difficulty, and produce a partial ordering of difficulty and dependency.

- **Requiring separate mechanisms to be combinable with others later on in integrated systems will prevent dead ends.**
- **Creating such a roadmap, and identifying milestones, can be a collaborative task between people who disagree on how to solve the problems.**
- **Several different roadmaps could be built with different sorts of long term targets.**
- **Producing such a roadmap is non trivial: it needs a lot of work, a lot of time and a lot of human intelligence.**
- **But it could also save a lot of work leading to dead ends.**
- **And there could be major benefits for the research community.**

A possible long term benefit

Combining and comparing work on different projects on different parts of the same roadmap, or on different roadmaps can lead to a new understanding of the whole field.

We can systematically study different sets of requirements (niches) and different designs, and the tradeoffs between designs in relation to niches.

Any set of requirements may be satisfied (in different ways) by different designs (see consumer reports on domestic equipment).

Any design will relate in different ways to different sets of requirements (niches)

Researchers using different designs, mechanisms, etc. can contribute to a common study of options and tradeoffs instead of wastefully competing to find the right design – when there isn't one.

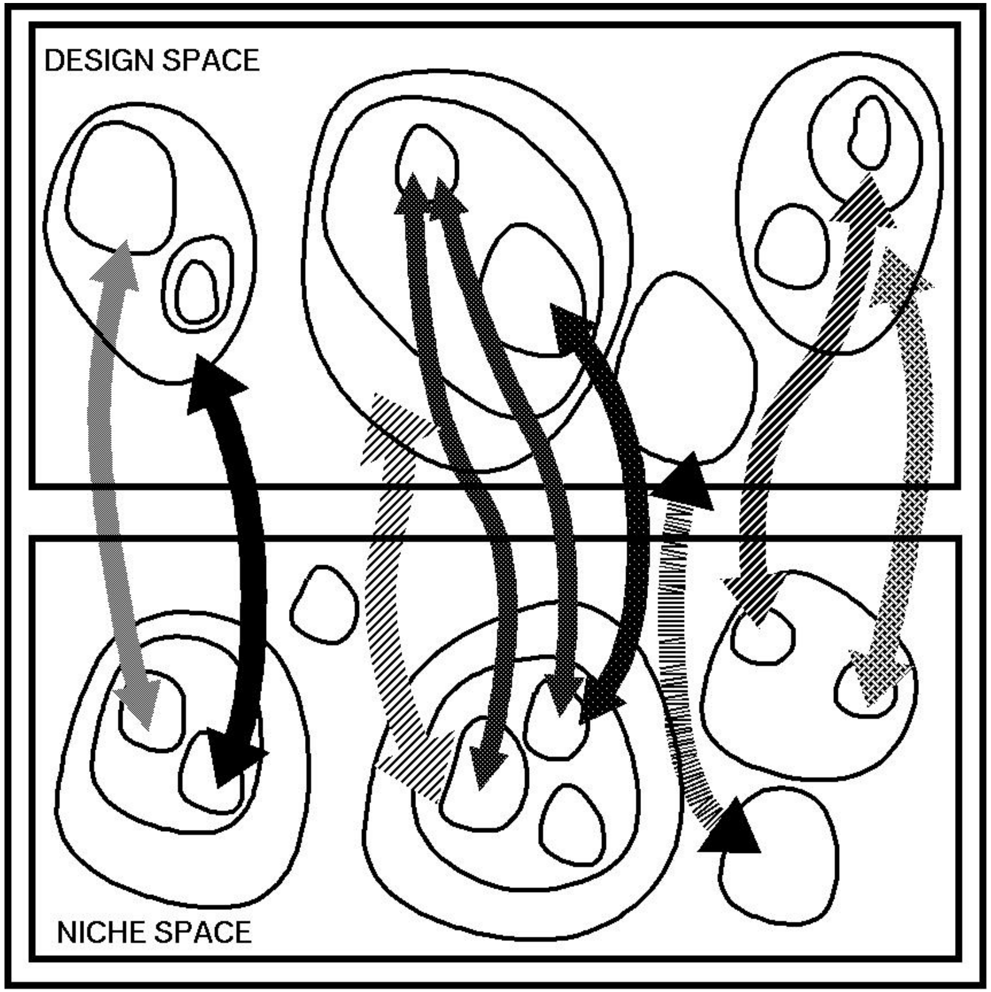

Acknowledgements

This work was done with colleagues in the EU-Funded CoSy project team at the University of Birmingham:

Jeremy Wyatt, Nick Hawes, Somboon Hongeng, Marek Kopicki

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

Jackie Chappell, in the School of Biosciences, University of Birmingham:

http://www.biosciences.bham.ac.uk/staff/staff.htm?ID=90