

How to Put the Pieces of AI Together Again

Aaron Sloman

School of Computer Science, University of Birmingham UK

A.Sloman@cs.bham.ac.uk

<http://www.cs.bham.ac.uk/~axs/>

Abstract

Since the 1970s AI as a science has progressively fragmented into many activities that are very narrowly focused. It is not clear that work done within these fragments can be combined in the design of a human-like integrated system – long held as one of the goals of AI as science. A strategy is proposed for reintegrating AI based around a backward-chaining analysis to produce a roadmap with partially ordered milestones, based on detailed scenarios, that everyone can agree are worth achieving, even when they disagree about means.

The Fragmentation of AI

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 came into AI only as graduates in another discipline so they had to pick up an accelerated education to get to PhD level, and often that meant reading only things their supervisor read.

A further source of fragmentation arose from the fact that extravagant predictions for AI had been made, which had not been fulfilled. When the optimism proved ill founded, the reason for lack of success was repeatedly mis-diagnosed as being due to the wrong tools, the wrong forms of representation, the wrong mechanisms. 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, and probably more. This led to more fragmentation and mutual ignorance between factions.

A Better Diagnosis

The main reason for the fragmentation and disappointing progress in ‘integrated’ AI has not been choice of mechanisms, or tools, or forms of representation (e.g. Symbolic

Copyright © 2006, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

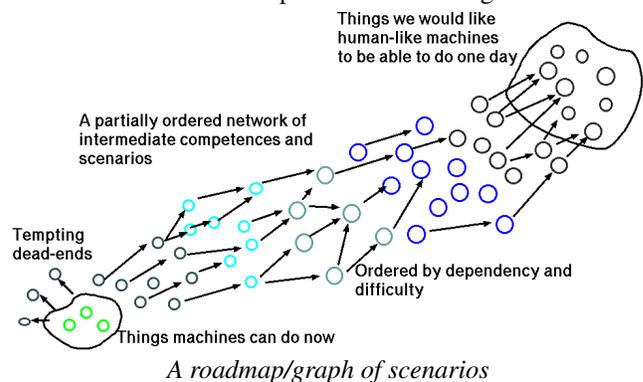
AI) but an inadequate analysis of long term *requirements* for success, leading to over-optimistic predictions and over-hasty selection of strategies and designs.

An example of this failure is the tendency to talk about the need to ‘scale up’ to human-level competence, ignoring the fact that humans scale up badly – though they do ‘scale out’: combining old competences in new ways.

All of this explains what an engineer from industry recently said at an AI robotics planning workshop: every few years he hears AI researchers say what they are going to do in the next few years, and it never changes.

Remedy: Build a Roadmap Backwards.

With colleagues in the EC-funded CoSy robotics project ¹ I have been trying to devise a way of countering the fragmentation and deepening our understanding of the problems. One idea² is to create a network of scenarios³ of varying degrees of difficulty and depth, starting from detailed descriptions of machines doing things that AI systems are nowhere near doing at present, and working back through scenarios that are less and less complex and demanding.



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

¹<http://www.cs.bham.ac.uk/research/projects/cosy>

²Reported at an interdisciplinary tutorial at IJCAI 2005 www.cs.bham.ac.uk/research/projects/cosy/papers/#tr0504 ‘AI in a New Millenium: Obstacles & Opportunities’

³Some draft scenario templates can be found here: <http://www.cs.bham.ac.uk/research/projects/cosy/scenarios/>

we can do now and more remote items correspond to distant goals. This ‘backward chaining’ research methodology contrasts with the ‘forward chaining’ that most researchers do — asking what improvements can be added to their current theories, techniques and systems, often defining ‘improvement’ in terms of performance on some arbitrary set of benchmarks. There is no guarantee that those improvements will take us nearer to achieving the long term scientific goals, even if they help to solve immediate engineering problems, such as recognising faulty parts on a production line, or steering a vehicle. In contrast, initiatives like Robocup rescue promote study of scenarios of varying difficulty, requiring multiple capabilities to be integrated. A different family of scenarios relates to a ‘general purpose’ domestic robot helper for the aged or infirm.

Scenario Design is Hard

It is difficult to develop the distant scenarios in sufficient detail. generate the backward chaining. So we are experimenting with a methodology for generating robot scenarios of various sorts by using a 3-D grid of requirement types. One dimension is concerned with *types of entities* (concrete, abstract, inert, mobile, totally dumb, intelligent, etc.) and another dimension with *things that can be done*, e.g. perceiving, pushing, assembling, referring to, thinking about, etc. Those two dimensions produce a grid of types of competence/entity pairs. The third dimension, depth, represents complexity and difficulty.

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

A 3D grid of competence+entity types.
The third dimension is difficulty.

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 then be measured by which portions of the graph have been achieved. Benchmarks requiring integration of different combinations of competences can be defined.

The grid is an oversimplification: some boxes need subdivisions, and other boxes will be empty. For example, you can refer to anything, concrete or abstract, but many things

cannot be acted on physically, pointed at, disassembled, etc. However 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. transfinite ordinals or the referring capabilities of symbols. So finding the grid’s topology is a research goal.

Examples

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* rather than *structures*, in particular concurrent processes at different levels of abstraction in partial registration with the optic array. The analysis is presented in more detail here www.cs.bham.ac.uk/research/projects/cosy/papers/#pr0505

Further analysis of the kinds of learning required when a robot arrives in a new household led to comparisons of varieties of learning and development in different species of animals (in collaboration with a biologist, Jackie Chappell, studying manipulation and cognition in birds.) It seems that humans, and a few other altricial species, unlike most animal 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. Some of the ideas are presented at this web site www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0601 on ‘Orthogonal recombinable competences’.

To avoid explosive combinations of possibilities, evolution seems to have developed (in a subset of species) abstract 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. This involves discovery not just of ‘subjective’ sensorimotor contingencies but also of ‘objective’ condition-consequence contingencies in the environment using an ontology that makes no reference to sensory or motor modalities. Developmental psychology provides many examples to go into the roadmap.

Examination of scenarios where causation is merely *observed* and those where it is *understood* (e.g. levers, gears) points to a need for Kantian, structure-based, deterministic causation in addition to Humean causation captured by Bayesian nets. This is related to the ability to learn mathematics, as explained in www.cs.bham.ac.uk/research/projects/cosy/papers/#pr0506

The requirements-based roadmap requires far more work by far more people. They need to agree on problems, not solutions. The long term payoff could be very high.

Acknowledgement

This work was done with the help of colleagues working in the EC-funded CoSy robot project.