



REQUIREMENTS FOR GOING BEYOND SENSORIMOTOR CONTINGENCIES TO REPRESENTING WHAT'S OUT THERE (Learning to see a set of moving lines as a rotating cube.)

Aaron Sloman
Last updated: 23 Jul 2008;16 May 2010;25 Sep 2010; 19 Sep 2011

This file is

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/nature-nurture-cube.html>

There is also a pdf version (which may sometimes lag behind the html version):

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/nature-nurture-cube.pdf>

I may try later to improve the format of this file.

An email list for discussion of this and related topics is available here:

<https://mailman.cs.bham.ac.uk/mailman/listinfo/nature-nurture>

(Note added four years later: the list produced no significant responses to the challenge.)

Background


This started life as a message circulated to colleagues in May 2006 about preconditions for a certain demanding kind of learning -- which does not seem to be easily expressed in terms of curve fitting, matching humans on some classification test, or any other simple criterion for success.

The question relates to the possibility of a learning system that starts with 2-D visual capabilities and an ontology that includes 2-D points, lines, junctions, and motion in a 2-D plane, and then invents for itself the idea of a 3-D space with structures and motions of various kinds which explains the complex sensed 2-D phenomena as a projection from a changing 3-D world.

Compare Plato's metaphor of people imprisoned in a cave able to see only shadows on the cave wall: what learning mechanism could cause them to come up with the idea of the moving 3-D objects generating those shadows, if they had not had concepts of 3-D structures and processes?

Many learning systems do *dimensionality reduction*. This problem requires *dimensionality expansion*.

The Necker Cube



You should be able to see this as either a 2-D collection of lines or as a 3-D wire-frame cube. In the latter case, if you look at it long enough it should 'flip' between two 3-D structures which differ only in the 3-D locations and orientations of the parts, not in their 2-D projections. So the difference does not involve any change in the visual input signal. Some examples of moving 2-D cubes are given below. The first is a rotating wire-frame cube. Because of the ambiguity displayed above, the rotating cube is ambiguous between rotating left and rotating right: it can be seen both ways, without any change in the 2-D visual signal.

Note: the existence of pictures whose interpretation can flip between alternatives without anything changing in the visual input is one of many reasons for taking qualia seriously, as discussed in <http://www.cs.bham.ac.uk/research/projects/cogaff/talks/#inside>

Online rotating cubes and other 3-D structures.

- <http://labix.org/snippets/cube>
Ambiguous rotating wire frame cube
If your browser allows you to expand the image try zooming to 200% or more (if you are using a high resolution screen).
This one sometimes flips for me, but mostly it rotates one way. I think that for some reason I have a strong tendency to experience it as if I am looking down towards the top surface rather than looking up towards the bottom surface.
- <http://www.math.ubc.ca/~morey/java/rotator/rot.html>
A rotating opaque cube (which you can control with the mouse).
- <http://www.math.union.edu/software/StageTools/StageManager/Samples/simple/cube.gif>
Rotating coloured opaque cube
- <http://www.math.union.edu/software/StageTools/StageManager/Samples/torus/Torus-1.gif>
Rotating opaque torus
- <http://www.cyberkerala.com/java/cubes.htm>
CyberKerala

You can use your mouse to control the speed and direction of rotation of each cube: just move the mouse pointer over the images, without pressing any buttons.

- <http://www.ies.co.jp/math/java/trig/kanran/kanran.html>

Rotating Ferris wheel

(As you click on the 'left' and 'right' buttons the wheel plane rotates about a vertical axis. But there's a 'necker-cube-like' ambiguity, and you can see it tilted either way while it rotates, which will affect whether you see individual buckets on the wheel moving towards you or away from you. Try this with the wheel stopped and with the wheel rotating (in its plane, about its axis).

- <http://www.math.union.edu/software/StageTools/StageManager/Samples/welcome.html>
More examples (click on the images to make them start moving.)

Message posted on May 19th 2006

From Aaron Sloman Fri May 19 12:44:32 BST 2006
To: nature-nurture
Subject: a theoretical question for theorists of various kinds

[Sorry I could not find a way to compress this question.]

As I am sure you are all aware there are debates about which concepts, knowledge, formalisms, algorithms, architectural features are innate and which products of learning and development.

In part that's an empirical question which may have different answers for different animal species -- for instance newly hatched chicks and new born deer can do things almost immediately that other species seem to have to learn to do. Some of you have heard Jackie Chappell talk about that with examples.

There are also theoretical/mathematical questions about what sorts of learning can and cannot be done by various kinds of mechanisms.

Is there a general purpose learning mechanism that can reliably and quickly discover the need for a 3-D (e.g. euclidean) ontology of lines, surfaces, opaque bodies, etc, simply on the basis of 2-D visual data and motor control signals?

There are some people in AI, philosophy, and psychology who believe that that there exist very general learning mechanisms that make it possible for something like a human being, or a future human-like robot, to function with concepts *all* of which have been derived from experience (e.g. by data-mining in observed sensorimotor contingencies, or by training neural nets, or some such thing).

Such people often talk about 'symbol grounding', claiming that all meanings have to be 'grounded' in sensory experiences, or some such thing. (I think the phrase was invented by Stevan Harnad, about 15 years ago.)

So that raises the question: assuming that all sensory inputs and all sensory outputs form collections of one or two-dimensional arrays of signals, is there some way in which a general purpose learning mechanism with NO built in knowledge of structures and processes in 3-D euclidean space (e.g. something based on trying to minimise algorithmic complexity of prediction mechanisms, or something else), can 'discover' the need for an ontology which includes changing 3-D structures, including rigid and flexible rods and strings, and surfaces in various orientations some of which are curved some not, some of which are wholly and partly invisible some of the time, but can become visible through movements in 3-D, either of the objects containing those surfaces, or other intervening surfaces or the viewer?

Can this richer ontology be discovered simply by running some general purpose learning mechanism?

In some sense the answer is obviously Yes: evolution produced animals such as us with a grasp of concepts of 3-D structures and processes, starting from a world that did not contain entities with such concepts.

And evolution is a general purpose learning mechanism.

But evolution takes evolutionary time-scales, and in a different environment this might not happen.

Is there a learning mechanism that can do it reliably, repeatedly (e.g. for all instances of a species) in a fairly short time, e.g. at most a few months or years of learning?

Here is an example:

Suppose you have a machine (Robby the robot) that gets 2-D visual input from a translucent screen. Behind the screen is a 3-D wire-frame cube and a distant light projecting shadows of the wire edges onto the screen.

light	cube	screen	robot
*	[]		o [+]

If normal adult humans see the shadows we typically see not only the 2-D configuration of dark and light points, and lines made of such points, but also a 3-D structure. Actually we may find that it is ambiguous (i.e. the Necker cube) and we switch between seeing two different 3-D structures while the 2-D structure remains unchanged.

Suppose the cube rotates about a 3-D axis and there is an array of knobs, levers or other devices whose state allows 1-D continuous change (e.g. rotation of a knob), whose effect is to alter the position, or orientation of the axis, and the speed of rotation of the cube, or perhaps just the orientation relative to the axis. (I think 4 knobs suffice to determine the exact 3-D location of the axis, additional knobs (4 more?) can determine the position of the cube relative to the axis and a the orientation of the cube around the axis.)

Now suppose Robby can turn the knobs, or move the levers that determine the 3-D position and orientation of the wire cube -- though Robby initially does not have any conception of 3-D position or orientation. However Robby can see the changing 2-D patterns produced by those movements.

Now suppose Robby has to be able to predict the consequences of any combination of knob or lever changes. Initially these consequences will be sensed only in terms of points on the screen changing from black to white or vice versa. But suppose Robby has 2-d line-finding algorithms, and can also therefore identify lines and observe how they move, which junctions between lines form or disappear as the lines move, when the lines change direction of movement, or perhaps change their size (if the projection of the shadows is not fully orthogonal).

Your task is to produce a general-purpose learning algorithm that can be given merely the 2-D visual data, the ability to change the state of the controls, and the task of continually predicting what will happen next.

On this basis it should be able to discover the power of a 3-D representation which can be projected onto a 2-D surface, where in the 3-D representation the metrical relations of the edges relative to each other do not change, even though in 2-D the relations between lines (e.g. angles, relative sizes, intersection points) are constantly changing.

This is a case of reducing dimensionality (of one kind -- variation in data) by increasing dimensionality of another kind (in the explanatory hypotheses). But there is structure as well as dimensionality.

It does this without being given at the start anything that includes a specification of 3-D space. E.g. it can't have the general notion of N dimensional euclidean space with hyperplanes, hypercubes, etc. and then an algorithm that searches for the smallest N that will capture all the data. It must merely try to learn to predict changes describable in a formalism that it has to start with (2-D geometry and topology)

and end up discovering a richer formalism (3-D geometry) as the solution to the learning problem.

Can that be done?

Has anyone done it?

I have a web-page criticising people who talk about learning sensorimotor contingencies without addressing this issue.

<http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0603>

ADDED 25 Sep 2010: A more recent discussion of ontology expansion vs dimensionality reduction:

<http://www.cs.bham.ac.uk/research/projects/cogaff/misc/simplicity-ontology.html>

This is also related to the discussion of learning orthogonal recombinable competences here:

<http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0601>

Also the distinction between a Humean (correlational, probabilistic) notion of causation and a Kantian (structure-based, deterministic) notion of causation, here

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

COSY-PR-0506: Two views of child as scientist: Humean and Kantian

Aaron

=====

Notes added in response to some comments from David Dowe 6 Sep 2006

He drew my attention to the Bayesian MML (Minimum Message Length) approach described in

> Dowe, Gardner and Oppy (2006+) Brit J Philos Sci (BJPS) paper,
> "Bayes Not Bust! Why Simplicity is no problem for Bayesians" at
> <http://www.csse.monash.edu.au/~dld/David.Dowe.publications.html#DoweGardnerOppy2006+>
> under "current version".

In one of my discussions of the rotating cube problem (I forget where: it was probably correspondence with a colleague) I pointed out that if the learning system had a meta-syntactic competence that allowed it to generate (enumerate) every kind of grammar expressible in one of the standard formalisms, and for each such grammar to enumerate every possible theory expressible in that grammar then by systematically exploring ever more complex theories (which would be enumerable), including theories with increasing numbers of undefined symbols in addition to the symbols required to describe the directly sensed 2-D structures and processes specified in the experiment) it could look for those theories that allowed the observed behaviour of the moving shadows to be explained and predicted, and would eventually come to the shortest one, which may or may not be one that we would interpret as describing a rotating 3-D structure.

(But it would also eventually come to one like that, even if it is not the simplest theory meeting the requirements.)

However I have not yet tried to formulate a theory in a logical format of 3-D spatial structure that would accommodate the rotating cube and its projection to a 2-D plane. I know it can be done -- and perhaps David Hilbert's axiomatisation of Euclidean geometry is rich enough?

I assume a bayesian version would replace the systematic search by a randomised one using some biases, but the principle is the same?

I suspect, but can't prove, that such processes involve such vast search spaces that they would take evolutionary time scales rather than individual learning time-scales, unless there was some bias in the ordering of grammars and theories, so as to favour the early production of what we would recognise as a theory whose ontology includes 3-D spatial structures and processes.

That, of course, would amount to an innate preference for theories that refer to 3-D structures. (Perhaps millions of years of evolution have produced such innate preferences in human and some other animal brains. Can we model this)

I.e. it would not be a *totally general* learning mechanism, without any innate biases.

My interests are in explaining how humans and other animals work, and why they are so diverse in their competences. Most cope with 3-d structures and processes of some sort, e.g. when feeding, mating, burrowing, making nests, escaping predators, etc.

Some seem to start with a pretty good, and therefore presumably genetically determined grasp of at least some aspects of 3-D structures and processes -- e.g. enough to let new-born deer run with the herd.

It's not clear what humans, chimps or crows have initially, but I suspect its much more specific than a preference for an ordering of the sort mentioned above. It probably has a lot to do with requirements for being able to suck and explore many movements of eyes, hands legs, and later fingers, tongue, etc., plus a meta-level competence that drives playful exploration of those devices and uses a criterion of 'interestingness' to store and generalise some results which are later used as new building blocks for the exploratory process. I've been trying to develop that idea with a biologist, Jackie Chappell, who works on animal cognition. Our papers on the topic are included here, including one presented at IJCAI 2005.

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

Our argument that some of that competence needs to be genetically determined is based on the claim that no totally general learning system could achieve the same results in the same time.

Hence the example of the rotating sphere as a challenge for someone to prove that claim wrong.

Further reading

- <http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0603>
Sensorimotor vs objective contingencies
- <http://www.cs.bham.ac.uk/research/projects/cosy/papers/#dp0601>
Orthogonal Recombinable Competences Acquired by Altricial Species (Blankets, string, and plywood)
- <http://www.cs.bham.ac.uk/research/projects/cosy/papers/#pr0506>
Two views of child as scientist: Humean and Kantian
- <http://www-formal.stanford.edu/jmc/child/>
John McCarthy, "The Well-Designed Child"
Also published in the AI Journal 2008.

Maintained by [Aaron Sloman](#)
[School of Computer Science](#)
[The University of Birmingham](#)