Re: Aaron Sloman's "The Irrelevance of Turing Machines to AI" article

Source: http://sci.tech−archive.net/Archive/sci.cognitive/2004−08/0076.html

---

From: Eray Ozkural exa (*erayo_at_bilkent.edu.tr*)
Date: 08/03/04

Date: 3 Aug 2004 03:54:19 −0700

"Sergio Navega" <snavega@intelliwise.com> wrote in message
news:<410ea576$1_1@news.athenanews.com>...
> "Neil W Rickert" <rickert+nn@cs.niu.edu> escreveu na mensagem
> news:ceir6u$rt4$1@usenet.cso.niu.edu...
> > >Acquisition of knowledge (or perhaps more appropriately, construction
> > >of knowledge) is also among my main interests, and I so far haven't
> > >found good enough reasons to think that a computational system
> > >is incapable of building it (provided that this system is left to
> > >interact with a natural environment).
> >
> > It is difficult to prove that a computational system couldn't
> > build knowledge. One of the difficulties is an inadequate
> > account of what is knowledge.
>
> That's an important concept, and perhaps we should try to come
> up with a "non philosophical" version of the word knowledge.

Yes, a scientific theory of knowledge is needed.

> > Still, machine learning has not shown anything that approaches human
> > learning. And it isn't for lack of trying. The best "computational"
> > learning systems cheat. That is, they depend on reward systems
> > (reinforcement systems) which are outside the computation. You might
> > say that my current direction is to investigate autonomous reward
> > systems.
>
> I agree that machine learning, the way it is today, is a far cry from
> what has to be done. But I believe that initial autonomous rewards
> (and with initial I mean "primary" or infant−like) can be obtained
> from the recognition of very salient statistical properties of the
> sensory signals. These properties are insufficient for the development
> of high−level cognition, but they seem to be important "first steps"
> towards it.

All thinking is statistical to some extent, in the sense that thought

must deal effectively with random sources encountered in the nature
(or in thought itself).

However, the more fundamental problem of perception is:
i) Abstraction
ii) Acquiring order in the nature

With these the agent can
i) Predict the future
ii) Manipulate the future

Which is more or less the base of intelligence.

> > I don't have any problem with that. But notice how you have to leave
> > computer systems and go to animal systems for good examples. So what
> > is it that the animal has that computers don't have?
>
> That's the million−dollar question (or better, its answer). I would
> say that animals have "embedded" the ability to adapt their behavior
> to the dynamics of the surrounding environment. This tends to make
> me think about an adaptive dynamical system. My hypothesis is that
> computers can be programmed to "simulate" the behavior of this
> dynamical system.

In a sense yet, but the nomenclature and theory of "dynamical systems"
is in my opinion insufficient to solve the 4 problems I have stated
above. Discrete theories are good enough.

On the other hand, if we take the statements above as natural (and not
idealized in terms of bifurcation, chaos, attractors, etc.), we will
need detailed models to capture the notion of adaptation and dynamics.
I believe the (postmodern) computational metaphor is sufficient for
that.

> > With some simple programming, say code that evaluate [f(x)]^2 − x,
> > a computer might be able judge the adequecy of a square root
> > computation.
> >
> > A human, using that ice−cream tastes good, might get to discovering
> > life forms on Mars.

Really? What is the causal link? What is the motivation? Making better
ice−cream? :)

> >
> > With computational standards, how do you get outside the narrow
> > ranged determined by those standards?
>
> I think I understand your question and it seems obvious that
> we have to let the organism come up with its own standard. What
> I propose is that the common idea behind all the standards of

> *all intelligent animals of the planet has something to do with*
> *the discovery of the most relevant statistical properties of*
> *stimuli.*

There are researchers that would argue against letting the organism come up with its own standard. This should be restricted by the cognitive architecture, or otherwise the organism might generate a utility function which will result in its immediate death (such as finding death a pleasureable experience).

Juergen Schmidhuber argues that the organism adapts a new utility function *only* when it can *prove* that the new utility function does not "topple down" the old one. That's a revolutionary idea, in my opinion. Nobody thought of the relation between a situated robot and theorem−proving, in this fashion. Such a constraint draws a border on the space of useless plans, and those that would not violate the primal utility functions.

BTW, Sergio and Neil, please join ai−philosophy group, your comments would be much appreciated.

Regards,

−−
Eray Ozkural