Minutes from "Intelligence Without Representation"

Daniel Buettner
Lecture by Cory Lueninghoener
CSCE 976 Advanced Artificial Intelligence

April 26, 2002

1 Introduction

In the lecture from 22 April, we covered material from a 1987 paper by Rodney Brooks entitled *Intelligence without representation*. This paper discusses Brooks ideas about how best to achieve intelligent agents.

2 Cory's lecture

Brooks begins his paper by defining the original goal of Artificial Intelligence, which was to replicate human intelligence. However, this task was very difficult and researchers split off into more specialized areas like NLP, representation, vision, *&c*. Brooks believes that there is insufficient communication between researchers in these areas. Consequently, even if they make a break through, their "piece" of the puzzle won't fit together with the work of any other researchers.

Instead, Brooks urges AI researchers to incrementally build complete systems that can be tested in the real world, instead of simply being tested on toy problems. However, he doesn't come out and say that duplicating human level intelligence is impossible, since the existence of humans proves that such intelligence can be found in the universe.

Then, Cory introduced us to Brooks's allegory of artificial flight researchers from the 1890s being presented with definitive proof of the possibility of flight in the form of a 747. Returning to their own time, they began attempting to duplicate what they had observed. In Brooks opinion, this is

a major problem with "traditional" AI efforts. They try to duplicate human intelligence while only being able to observe the external actions of others.

Then Brooks discusses a complaint of some AI researchers: AI is denied its rightful successes. He points out that people tend to classify "hard" problems as being AI problems. However, when someone finds a clever way of solving a problem, AI detractors say that the problem was simply solved by an algorithm, so it was not really an AI problem after all. At this point, Brooks points out that AI is seldom said to have failed, since researchers in the area are dealing with hard problems.

Next Brooks goes on to say that many AI researchers are guilty of performing an amazing feat of doublethink. Having complained that solved work was pigeonholed as "Not AI", these researchers define only the solved portions of their own work as "AI". For example, Brooks is very angry about the use of "abstraction" in defining the problem. He feels that

> ... abstraction is the essence of intelligence and the hard part of the problems being solved. Under the current scheme the abstraction is done by the researchers leaving little for the AI programs to do but search.

Then Brooks returns to his allegory and points out that researchers can be misguided by relying on merely being able to observe what they wish to duplicate. The artificial flight researchers saw lots of chairs on the plane and got to work building nice and comfortable seats for their plane. Unfortunately, they decided that they didn't need to worry about weight, since the giant plane was clearly heavy and yet it could fly. Additionally, having questioned some of the passengers about aerodynamics and finding no one knowledgeable in that subject area, the AF researchers decided that it must not be important. This illustrates the pitfall of assuming that human perception is capable of extracting the important information from a situation.

After explaining what's wrong, Brooks turns his attention to explaining The Way Things Ought To Be. He firmly believes that intelligence can be created by building "creatures" which are autonomous agents in the real world. They should be able to operate in a dynamic environment without breaking when the world changes. Additionally, he believes that they should have some "purpose" in life.

Brooks outlines two possible way to decompose the construction of such a creature. First he describes decomposition by function, which he dislikes. His chief criticism is that input and central reasoning are done by different groups with disjoint assumptions. For this reason, Brooks is a proponent

of decomposition by activity. He thinks that the agent should be divided into activity-producing subsystems or layers. In his world, each layer runs in parallel and is disconnected from the other layers.

By creating separate activity-producing layers, you can avoid explicitly representing the world. This allows for fast low-level reactions, and helps to allow the system to cope with changes in the world. To put his theories into practice he has built several robots using this architecture. He reports that they are very successful at operating in the real world. At the time his paper was written, he had built 4 robots. During Cory's talk, we checked the web and found that Brooks has been in charge several advanced robots that have been built since then.

At the end, Brooks points out that his ideas are not simply Connectionism, Neural Networks, Production Rules, Blackboard System, or German Philosophy. He claims that his robots were on the verge of "simple insect intelligence". However, he did not explain what exactly he meant by that. It seems doubtful that he has any robots that can propagate their own species nor collect food.

3 Student interaction

Lin Xu:	Has he done any new work?
A:	We checked the web, and he *does* have more work!
Tibor:	This paper is mighty old. Isn't it outdated?
A:	Well, we checked his web page, and it seems that he's still doing the same stuff. It must still be working!
Daniel:	The robots that he built had really bad controllers like PALs and ROMs. Maybe that was why he thought the activity layers should be so simple.
A:	Could be.
Daniel:	Also, the robot that my group built for senior design essentially used Brooks's activity decomposition. Although we'd never heard of Brooks.
A:	Maybe he didn't really change the world with this paper.
Rob:	Not all AI is robotics!
A:	Brooks seems to be preoccupied with reproducing animal intelligence.
Tibor:	He mentions a blind spot in the centre of your field of view. What's that all about?
A:	Cory looked it up on the web and found a simple way to demonstrate it. It's true! There *IS* a blind spot in the centre of your field of view!
Shabbir:	Saying that unsolved problems are AI and solved problems aren't seems questionable.
A:	Well, maybe Brooks was just trying to illustrate a point.
Praveen:	Human intelligence is extremely complex. Why is there so much emphasis on trying to duplicate it?
A:	That's a good question. It certainly isn't clear that human intelligence is necessarily the *best* form of intelligence possible. Especially because humans are mean to each other. We don't want an army of robots rising up and killing us all!
Amy:	How are his robots even close to bug intelligence?
A:	Well, they can avoid running into things! But they can't find their own food or reproduce. They can't communicate with one another. Maybe some of his newer work comes closer.