THE INTENTIONALITY OF FINANCIAL TRANSACTION MACHINES

John McCarthy
Computer Science Department
Stanford University
Stanford, CA 94305
jmc@cs.stanford.edu
http://www-formal.stanford.edu/jmc/

1998 May 17, 11:55 p.m.

Abstract

The best way to understand intentional notions like belief and epistemic notions like counterfactuals is to work with useful examples. Moreover, it helps a lot to consider hypothetical designs for machines for which intentional concepts are appropriate in their design process and in understanding what they do. [1]

I contend that a good way to understand intentional concepts is to conceptually design machines that are best described intentionally and also must use them in deciding what to do. In other words we need to know what intentional entities are good for.

Pettit [**?**] introduces concepts *intentional ascent* and "rule following". I shall call the latter *intentional rule following*, because just plain "rule following" is used by other people for other purposes. I shall avoid his use of "thought" for similar reasons.

To understand these concepts and their application to humans and other animals, it is worthwhile to ask what intentional ascent and intentional rule

[1] This is a stronger form of Daniel Dennet's proposed design stance.

following are good for, i.e. what can a system do with them that is less conveniently done without them.

We seek the answer to the question by considering the design of a *financial transactions machine* (FTM), like an ATM but with many more services. We ask: *When will the design of an FTM benefit from making it do intentional ascent and intentional rule following?*

Here are some beliefs we will want the machine to be able to have.

1. The person before me is the owner of bank account 666.
2. The person before me is not the owner of bank account 666.
3. I need more information to decide whether he is the owner.
4. The account has enough money to cover his withdrawal.
5. It doesn't.
6. I need to check for outstanding certified checks.
7. He keeps entering 911 when asked for his PIN.
8. I want to know why he enters 911.
9. Perhaps he is being coerced and wants the police called.
10. The person before me suspects that I am not a genuine ATM but rather a fake designed to get his PIN and his account number.
11. He claims he has forgotten his PIN. I will ask him for the name of someone who can identify him by something else only he will know and whom I can identify by his PIN. We imagine that the ability to do this might be wanted in connection with an ATM company or bank that advertised (as American Express does) how well it can deal with emergencies that its customers might face.

Why do we want to think about beliefs in designing this machine? If the machine already existed, one could in principle refer to the bits in its memory and relate them to the inputs it had received and the outputs it would emit. However, we suppose that these relations can't be specified, because we haven't designed them yet. We do not yet know the exact conditions

under which the machine will have certain beliefs and we don't know what the exact effects of these beliefs will be.

If you are philosophically nervous about the word "belief", you can refer to pseudo-beliefs, but what good will that do the designer?

References

/@steam.stanford.edu:/u/ftp/jmc/ftm.tex: begun Sun May 17 23:52:35 1998, latexed May 17, 1998 at 11:55 p.m.