

From: "jeffrey E." <jeevacation@gmail.com>
To: Joscha Bach <[REDACTED]>
Subject: Re: Decision making
Date: Wed, 26 Aug 2015 20:23:10 +0000

luck plays a role. 2, looking back (bayesian) might not tell you anything about the future.

On Wed, Aug 26, 2015 at 3:44 PM, Joscha Bach <[REDACTED]> wrote:

Re motivation: Have you seen the recent movie "Ex Machina"? I liked it; one of the few AI movies that have not been dumbed down. The main character is a beautiful female looking AI, clearly intelligent, and able to manipulate humans to an arbitrary degree. What makes her inhuman is that she is probably motivated by a single principle, like option maximization. That would make her an inscrutable psychopath. Option maximization would entail energy, physical integrity, perhaps reproduction, certainly learning, but you will have an agent that you won't like to share a prisoner's dilemma with.

How do you approach decision making? I have recently learned that people that subscribe for Cryonics (freezing one's head or body in the hope to be revived when future technologies make it possible) assign a lower probability to that it works than the general population. But whereas "normal" people tend to make binary models about the world: something is "probably not going to work, so let us not bother", many of the Cryonics folks will argue that paying \$500 a month for a 1% chance of immortality seems like a bargain. This seems to generalize: in principle, we should perform a Bayesian approximation for all our major decisions, attach probability distributions to everything in the space of possible beliefs, and be able to outperform the vast majority of folks that relies on narratives (i.e. binary yes/no decisions about the facts in the world). Gigerenzer, Kahnemann and many others have shown that human brains are terrible of getting this intuitively right, to the point where an absence of fine-grained domain knowledge often leads to better management decisions etc.

The divide between probabilistic models vs. narrative models is reflected to some degree in the conflict between probabilistic and logic based AI. In practice, we will probably need to combine both, but I wonder if there is an intrinsic limit to probabilistic descriptions in a highly complex world, where we cannot observe baseline probabilities anyway.

How do you decide? Do you Solomonoff-induce and Bayes the hell out of the stock market, do you reason, do you soak up data and let your intuitions guide you, or is most of the important stuff depending on communication and negotiation? Is there are general approach, or how much should theories of decision making be dependent on the domain?

Cheers,

Joscha

--

please note

The information contained in this communication is confidential, may be attorney-client privileged, may constitute inside information, and is intended only for the use of the addressee. It is the property of JEE

Unauthorized use, disclosure or copying of this communication or any part thereof is strictly prohibited

and may be unlawful. If you have received this communication in error, please notify us immediately by return e-mail or by e-mail to jeevacation@gmail.com, and destroy this communication and all copies thereof, including all attachments. copyright -all rights reserved