

From: Joscha Bach <[REDACTED]>
To: Jeffrey Epstein <jeevacation@gmail.com>
Subject: Re: Decision making
Date: Wed, 26 Aug 2015 21:03:35 +0000

Finding food depends on the species.

Some are only functioning if you throw them into it. Higher organisms cannot afford the cost of failure, so they will either hard-wire the types and probable places of food, or they rely on teaching (humans are born without a direct pre-wired connection between hunger and eating, and parents tell us what kind of food is safe to eat).

Successfully eating food creates pleasure signals, which act as rewards for reinforcement learning of a connection between (specific) hunger urges, food stuffs and food locations/acquisition strategies. Reinforcement learning is mathematically quite equivalent to probabilistic models (POMDPs etc.).

At some point we introduce meta-strategies (somebody tells us what food to try, and our response to these suggestions is a combination of hardwired social urges and learning).

The series continuation might be found in similar ways as Wolfram suggests in New Kind of Science: come up with a generator function for series generator functions, and let it generate results until you find a matching generator function. You will need to tune the meta-generator to make it produce most likely candidates first. Pick the solution with the lowest Kolmogoroff complexity you can find.

In practice, we often find these solutions by case-based reasoning: remember the same sequence from a previous occurrence and look up the solution, or find a solution for a similar sequence and mutate it until it fits.

Am 26.08.2015 um 16:48 schrieb jeffrey E. <jeevacation@gmail.com>:

not sure of course. but trying many things , finding a solution and repeating it until it no longer works. . to find the series solution to 7 19 37 88. is virtually impossible , but merely the difference of cubes, , how do you get to that answer.? or how do you find food? ::))

On Wed, Aug 26, 2015 at 4:44 PM, Joscha Bach <[REDACTED]> wrote:

So you are suggesting that probabilities are not the best way to frame decision making? (We can overcome things like independence, at the cost of sufficiently complex models, and in a world without repetition = some kind of regularity, all bets are off, but of course probabilistic models might still be impractical in practice.) What would be a better framework?

Am 26.08.2015 um 16:39 schrieb jeffrey E. <jeevacation@gmail.com>:

probabilities vs similiarities. I suggest you take care with the concept of probabillites which really requires repetition, symmetry, and independence.

On Wed, Aug 26, 2015 at 4:26 PM, Joscha Bach <[REDACTED]> wrote:

1. Does luck mean gambling, or getting the probabilities right by intuition?
2. What else tells you anything about the future? (The past also tells us how much is does not tell us about the future)

Am 26.08.2015 um 16:23 schrieb jeffrey E. <jeevacation@gmail.com>:

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

--

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

--

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