

---

**From:** Joscha Bach [REDACTED]  
**Sent:** Thursday, April 28, 2016 3:57 PM  
**To:** Jeffrey Epstein  
**Subject:** Re:  
**Attachments:** signature.asc

I remember how I played with neural networks as a student, and stumbled on a solution where I could speed up the learning dramatically by letting the network estimate how well it knew certain parts of the problem space, and focus its learning on that area. The next step was to let the net decide what training data it wanted to look at. Obviously, asking is better than observing. But for current AI systems, questions are much, much harder, of course, at least in the general case. To ask a question, I must know what I do not know, and I must know what someone else knows, and how to translate and fit the answer into my knowledge. Consequently, even people seem to find asking much harder than answering. (Different rules may apply to socially regulated contexts.)

Children ask "true questions", in which they try to form complex models for the first time. Once people have a model in place, they tend to only ask for variables ("at which time does the train leave"), and when they converse about each others life, they tend not to be interested in model revisions, but in mapping what they hear to existing variants of the established. This seems to be even true when scientists extend their theories: the natural inclination seems to be to find supporting evidence to allow me to let the model stay in place, and simply adjust it. Since a model revision would require to adjust all inferences and connections tied to the discarded model, this makes sense.

Looking at question asking systems seems to be an extremely promising venue for AI research. I suppose we have to treat knowledge classes quite differently: questions about basic linguistic structure can only be answered by other speakers. Questions about conceptual maps can only be inferred from statements of others. Questions about content can in principle be deduced from individual observations, and asking questions (to others) is simply a shortcut.

The nature of question-asking seems to differ significantly between most aspergers and most neurotypicals. The former try to find out what is true, with the null hypothesis usually being that their own idea is right when it conflicts with what others tell them. Aspergers tend to change their opinions based on how they judge the logical truth of the arguments of the other side. Neurotypicals tend to look for the "right" answer, not the true one, i.e. they try to align themselves with the sentiments and norms of people they assign social status. The primary mechanism for this seems to be "emotional synesthesia", i.e. they literally feel the emotions of other people as their own, including the sense of rightfulness that accompanies a statement. Successful speakers can project both conviction and high social status. Neurotypicals erect social boundaries to a large extent to avoid "infection" with "bad" opinions. I find that my intuitions about building AI (and about forming knowledge, negotiating norms etc.) are largely influenced by being not neurotypical (and that is definitely true for the most capable people in the field). It is very visible in Marvin Minsky and Noam Chomsky, I think. I wonder what we are missing, especially since my intuitions and apparent experience tell me that non-nerds rarely stumble on the right answers to complex problems.

What do you think?

— J

> On Apr 26, 2016, at 08:45, jeffrey E. <jeevacation@gmail.com> wrote:

>

> how are questions formulated by ai systems answers are easier than formulating good question

>  
> --  
> 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