From: "Joi Ito" <[redacted]>
To: "Jeffrey Epstein" <jeevacation@gmail.com>
Subject: Fwd: Re: MDF
Date: Wed, 23 Oct 2013 13:55:56 +0000

—
Sent from Mailbox for iPhone

---------- Forwarded message ----------
From: **Sebastian Seung** <[redacted]>
Date: Wed, Oct 23, 2013 at 9:48 AM
Subject: Re: MDF
To: "Joi Ito" <[redacted]>
Cc: "Joscha Bach" <[redacted]>, "takashi ikegami" [redacted] "Ari Gesher" [redacted], "Kevin Slavin" <[redacted]>, "Martin Nowak" <[redacted]>, "Greg Borenstein" <[redacted]>

> I've described today's dominant theory of vision in Chapter 4 of my book Connectome:
> http://www.amazon.com/Connectome-How-Brains-Wiring-Makes/dp/B00BQAPTL8
> and in Appendix E of this textbook:
> http://www.amazon.com/Principles-Neural-Science-Edition-Kandel/dp/0071390111
>
> I call this theory the "hierarchical perceptron," and credit it to Fukushima's Neocognitron (1980). His work in turn was inspired by discoveries and speculations of neuroscientists Hubel and Wiesel (1962). Today's deep learning architectures for vision continue in this tradition.
>
> For 50 years, neuroscientists have failed to conduct conclusive empirical tests of the theory. This situation is about to change, due to new technologies like connectomics. Over the next 10 years, neuroscientists will finally figure out how feature selectivity and invariance are related to neural connectivity. We are already succeeding in the retina, and the cortex will be next.
>
> In the 1990s, my research was focused on machine learning, which can be seen as a rebranding of the pattern recognition camp of AI. As far as I can tell, AGI is an attempt to revive and rebrand the reasoning camp of AI. I'm sympathetic to this goal. If I were starting over in AI today, I'd study reasoning rather than join the pattern recognition bandwagon.
>
> That being said, AGI will have trouble succeeding because it is following the scruffy tradition. Perhaps the main failing of this tradition is its refusal to define objective (and preferably quantitative) measures of success. In his infamous report, Lighthill wrote that "a relationship which may be called pseudo-maternal rather than Pygmalion-like comes into play between a Robot and its Builder."

http://www.chilton-computing.org.uk/inf/literature/reports/lighthill_report/p001.htm
Lighthill meant that scruffies love their creations, and hence do not evaluate them. In contrast, the pattern recognition camp has quantitative measures of success, which is arguably the main reason that they have made recognizable progress.

I find it paradoxical that Minsky was king of the scruffies. Perhaps because he was personally so talented at mathematics, he had no respect for it. When I read his inspirational writings, I recognize his genius. At the same time, I'm reminded of the saying that "A science is any discipline in which the fool of this generation can go beyond the point reached by the genius of the last generation." Minsky is bitter because he failed to turn his field into a science.

[redacted] writes:

> I'm adding Sebastian who is working on retinal neuroscience and would have a view on how the brain does this.
>
> - Joi
>
> On Oct 21, 2013, at 23:56 , Joscha Bach <[redacted]> wrote:
>
>> Hi Takashi, hi Ari, hi all,
>>
>> finally I got around to look at Takashi's talks and his 2010 ACM article. The first thing that came to mind was the distinction between "neat" and "scruffy" AI, which might be described as the clash between folks that wanted to construct AI by adding function after function, vs. those that want to take a massively complex system and constrain it until it only does what it is supposed to do.
>>
>> The idea of starting from massive data flows is very natural and theoretically acknowledged, even it is often practically neglected. Cognition, by and large, is an organism's attempt to massively reduce complexity, by compressing, encoding, selectively ignoring, abstracting, predicting. controlling it. Thus, it seems natural to focus on the mechanisms that handle this complexity reduction, which I think is exactly what most research in computer vision, machine learning, classification, robot control etc. is doing. A lot of the work on problem solving and learning within cognitive science even works _only_ on the highest level of abstraction, i.e. grammatical language, regular concept structures, ontologies and so on.
>>
>> If I understand Takashi correctly, he points towards another perspective: (please forgive and correct me if I should oversimplify too much here)
>> 1. Cognitive systems do not only need to reduce complexity, but also build it (for instance, take simple cues or abstract input and use it to seed a rich, heterogenous, ambiguous and dynamic forest of representations).
>> 2. Cognitive processes that work directly on and with high complexity data are under-explored.
>> 3. The study of systems that are immersed in such complexity might open the door to understanding intelligence and cognition.
>>
>> There is really much more in Takashi's talk, but let me respond to these in turn:
>>
>> 1. I believe that cognition is really about handling massive data flows, by encoding it in ways that the cognitive agent can handle and use to fulfill its demands. This works mostly by identifying apparent regularities and turn them into perceptual categories, features, objects, concepts, ontologies and so on. Our nervous system offers several levels and layers of such complexity reduction, the first one of course at the transition between sensory inputs and peripheral nervous system (for physiological, tactile, proprioceptive input), or, in the case of visual perception, the compression we see between retina and optic nerve. The optic

nerve transmits massively compressed data from the retina to the thalamus, and from there to the striate cortex (the primary visual cortex, V1). V1 is the lowest level of a hierarchy of visual and eventually semantic processing regions: from here, the dorsal and ventral processing streams head off into the rest of the cortex. V1 contains filtering mechanisms, which basically look for blobs, edges, movements, directions and so on, based on local contrasts. V2 organizes these basic features into a map of the visual field, including contours, V3 detects large, coherently moving patterns, V4 encodes simple geometric shapes, V5 seems to take care of moving objects, and V6 self-motion. The detection of high-level features always projects back into the lower levels, to anticipate and predict the lower level features that should be isolated based on the higher-level perceptual hypothesis. The story is similar for auditory processing, and eventually the integration of basic visual and optical percepts into semantic content: at each level, we take extremely rich and heterogeneous patterns and reduce their complexity.
>>
>> The transformation from concepts to language also represents another, incredible level of complexity reduction.
>>
>> The highest complexity reduction, however, takes place at the interface between conscious thought and all the other processes. I believe that the prefrontal cortex basically holds a handful of pointers into the associative cortical representations, skimming off only a handful objects, relations or features at a time, and bring them into the conscious focus of attention.
>>
>> The perspective of the need for staying at a complex level is entirely warranted, though: there are many intermediate representations that allow cognitive processes only if the complexity stays high, and might even need to increase it. This includes many sensor-motor coordination processes, but also most creative, more intuitive exploration.
>>
>> This is not the same complexity as the one at the input, however! This as a level where data is already split into modalities, semantically organized and so on. On the other hand, it is much more complex as linguistic or cognitively accessible types of mental content.
>>
>> 2. Scientists tend to have a fixation on thinking with language, and it is quite natural to fall for abstract, a-modal representations, such as predicate logic systems or extensions of these when it comes to modeling cognition and problem solving. This might explain the fixation of cognitive architectures like Act-R and Soar on rule-based representations, and the similar approaches of a lot of work in classical AI.
>>
>> On the other hand, there is a lot of work on learning and classification to handle vast complexity, with the goal of reducing it. (A particular beautiful example was Andrew Ngs work on deep learning, where his group took 30 million randomly chosen frames from Youtube, and trained an unsupervised neural net to make sense of them. They ended up with spontaneously emerging detectors for many typical object categories, including cats and human faces. I could not avoid to think of that paper when Takashi mentioned his fascination with looking at TV pixels directly…) --> http://arxiv.org/pdf/1112.6209.pdf
>>
>> Thus, the typical strategies seem to encompass "abstract 2 abstract" cognition, and "complex 2 abstract" cognition. What about "abstract 2 complex" and "complex 2 complex"? Most of the existing approaches on "complex 2 complex" cognition are not really cognitive, such as Ansgar Bredenfeld's "Dual Dynamics" architecture, or Herbert Jaeger's Echo State Networks. The current proponents of such complex cognition are also often radical embodimentalists (cognition as an extension of sensor motor control, neglecting dreams, creativity, imagination, and capabilities for abstract thinking).
>>
>> 3. The idea of getting to artificial intelligence _just_ by "looking at" (blind deep learning) on complex data flows is not new. I think that there are at least two aspects to it: deriving a content structure that allows the identification and exploitation of meaningful semantic relationships (for instance, discerning space, color, texture, causal order, social structure, … for instance simply by analyzing all of Youtube, or by collecting data from a robotic body and camera in a physical world), and the integration of that structure with an architecture that is capable of thought, language, intention, goal directed action, decision making, and so on. The former is

tricky, the latter impossible. Complexity itself does not define intentional action, and the differences between individuals and species should not be reduced to differences in complexity perceived by the respective agents.
>>
>> I agree that we need to gain a much better understanding of "complex 2 complex" cognition, but that must integrate, not replace what we already know about the organization of cognitive processes. I am certain that our current models are a long way off from capturing the richness of conscious experience of our inner processes, and even more so from the much greater complexity of those processes that cannot be experienced.
>>
>>
>> Another interesting point I gathered from Takashi's talk is the idea of something we might call "hyper-complex" cognition. The complexity handled by our human minds (as well as the one of Andrew Ng's deep learning Youtube watching networks) builds on very simple stimuli. But what if the atoms themselves are abstract or highly complex, for instance because they are already semantic internet content? The cognitive agents handling those elements may essentially be operating at a level above human cognition if they are capable of operating on that complexity without reducing it. Unlike humans which are forced to translate and reduce all content into their individual frame of reference, and access it only through a single perspective at a time, artificial agents do not need to obey such restrictions. Today's Big Data moniker probably marks just the beginnings of the abilities of machines to make sense of abstract and complex input data.
>>
>>
>> Cheers,
>>
>> Joscha
>>
>>
>>
>>>>> Fascinating. Ikegami is taking a very interesting tack:
>>>>>
>>>>> http://www.youtube.com/watch?v=tOLIHhjNlBc
>>>>> http://sacral.c.u-tokyo.ac.jp/pdf/ikegami_ACM_2010.pdf
>>>>>
>>>>> For me, this is similar to the discussions that you and I and Kevin have been having about auto-didactism: starting from complexity rather than abstraction (which is generally antithetical to academic learning). It would seem to me that most artificial intelligence research has started from abstraction (and forgive my ignorance if I'm off base here) and attempted to build up to complexity. My very cursory look at the Joscha's MicroPSI work seems to show an approach moving in the direction of the what Ikegami did with the MTM from the classical abstraction-first approach. MicroPSI places its constructs in a reduced fidelity virtual environment, has lower-level abstractions, and brain structures/dynamic pre-synthesized for things like motivation, emotion, (please correct me if I'm off base - like I said: cursory). The brain structures in living systems have have evolved as low-energy means of processing brain signals (both sensory data flows and internally routed streams) once they have showed fitness - ultimately, they were sand-blasted into their shape by generations of massive data flows. We have an understanding of what purpose they serve but not a good understanding of how they work (maybe I'm behind on the state of the art in neuroscience on that point?).
>>>>>
>>>>> Ikegami is starting from the complexity and seeing what emerges - which seems to me to mirror the rise of consciousness in natural systems. Mind is the surfer that hangs on the eternal wave of the massive data flow of sensory input without wiping out. Somehow, the reality of the temporally continuous observer arose from exposure to sensory data flows and the evolution of the complexity of the brain. Ikegami is shortcutting the snail's pace of the physical evolution of natural systems by synthesizing a neural network of sufficient complexity as well as high-resolution sensors.
>>>>>
>>>>> Thinking about modern synthetic data flows (you know…. the internet!) as being as rich as sensory data leads one to imagine some interesting possibilities in a) whimsically, the spontaneous emergence of consciousness and b) practically, new techniques for dealing with that massive data flow that mimic something

like natural consciousness. There's nothing in the practical world of big data that really looks like the MTM (that anyone is talking about - who knows what lurks in the high frequency trading clusters busily humming in the carrier hotels). Everything that Google and Facebook and the like seems to be doing is much simpler than anything like this.
>>>>>
>>>>>
>>>>> On Oct 19, 2013, at 9:37 AM, Joi Ito <[redacted]> wrote:
>>>>>
>>>>>>
>>>>>> http://www.dmi.unict.it/ecal2013/workshops.php#4th-w
>>>>>>
>>>>>> - Joi
>
> --
> Please use my alternative address, [redacted] to avoid email auto responder