

From: Greg Borenstein <[REDACTED]>
To: Joscha Bach <[REDACTED]>
Cc: Sebastian Seung <[REDACTED]>, Joi Ito <[REDACTED]>, takashi ikegami
<[REDACTED]>, Ari Geshner <[REDACTED]>, Kevin Slavin
<[REDACTED]>, Martin Nowak <[REDACTED]>, Jeffrey Epstein
<jeevacation@gmail.com>
Subject: Re: MDF
Date: Thu, 24 Oct 2013 00:56:54 +0000

On Oct 23, 2013, at 11:09 AM, Joscha Bach <[REDACTED]> wrote:

The question of good benchmark tasks is haunting AI since its inception. Usually, when we identify a task that requires intelligence in humans (playing chess or soccer or Jeopardy, driving a car etc.) we end up with a kind of very smart (chess-playing, car-driving) toaster. That being said, AI was always very fruitful in the sense that it arguably was the most productive and useful field of computer science, even if it fell short of its lofty goals. Without a commitment to understanding intelligence and mind itself, AI as a discipline may be doomed, because it will lose the cohesion and direction of a common goal.

I think this issue of the changing definition of intelligence being a moving goal post is absolutely critical, Joscha. And it's one that long-predates 20th century digital computation-based AI efforts.

Recently, I've been reading the work of Jessica Riskin, a Stanford historian who studies the long history of AI and Artificial Life. Specifically, Riskin's been writing about a strange phase in the history of mechanical automatons that happened in the second half of the 18th century. Previously, automatons had always been built with their mechanism in one place (i.e. in a hidden box or platform) that then drove their figures via a series of rods or connectors. The figures, the representative part of the automaton were like the birds in a cuckoo clock with no relation to the mechanism that made them move.

Then, suddenly, in the second half of the 18th century, a series of automaton makers started to produce automatons that were built in a way that was analogous to the thing they represented. The French maker Vaucanson built a flute-playing automaton that had a flexible tube in its neck that connected its bellows to a set of flexible lips made of leather that actually played a real flute. Vaucanson ending up proving some new facts about how human lips shape multiple overtones when playing the flute -- things that actual flute players didn't know they did, but that he had to figure out since he was trying to make an automaton that played the flute the same way that they did.

He also famously built a Defecating Duck -- an automaton duck that seemed to eat food from someone's hand and then defecate onto the floor. Like its more famous cousin, the Mechanical Turk, the duck was a deception: it had to be pre-loaded with fake scat. However, like the flute-player, its gears and workings were concealed inside of it and, to the best of Vaucanson's ability, emulated the workings of a real duck.

Riskin emphasizes that whenever one of these types of simulating automatons would achieve a result that was previously considered the sole domain of life or nature, people would respond by simply re-defining the natural so that the behavior that had been simulated was somehow not the essence of what it meant to be alive or human.

This strikes me as very much like the process we go through with the defining AI tasks like chess or Jeopardy or car driving. We start off believing that these are tasks that only human intelligence can achieve. Then we build computational systems that can do them. Those systems are often inspired by the way humans achieve the tasks,

but in the end work in extremely non-human ways. Google's self-driving car uses massive satellite data and laser scanning to drive, Deep Blue doesn't play chess like a human does.

This gap between how the technical systems achieve these tasks and how human intelligence does ends up emphasizing the difference between human intelligence and computation rather than erasing it. If we evaluate the project of AI based on how closely its products reproduce human forms of problem solving then this gap points towards failure. However, if we can stop making this comparison and instead value the products of this research for their own properties, especially insofar as those properties are different from human capabilities and can supplement them (as in Doug Engelbart's Augment model, recent work on human-in-the-loop, and Ari's human-computer symbiosis framing all do) we might make a lot more headway (and we might also have a more productive relationship with the technologies that we actually produce).

Document search, face detection, and spam identification were all parts of the original AI program in so many words. Now they've been deployed as Google search, on Facebook and in every digital camera in the world, and in every email service. I'm not sure we're poorer because they spread out and diversified into these various technologies rather than congealing into strong AI.

Riskin's story has a strange ending. Suddenly, around the end of the 18th century, this trend towards simulation disappeared. Automaton makers went back to building automatons that externally resembled their subjects in form but didn't attempt to reproduce their workings. This drive towards simulation didn't really return until the digital era in the second half of the 20th century. However, a lot of the mechanical know-how that got produced did continue on into the industrial revolution.

-- Greg

Two good Riskin papers to read are:

Eighteenth-Century Wetware <http://www.stanford.edu/dept/HPS/representations1.pdf>

and The Defecating Duck or the Ambiguous Origins of Artificial Life
<http://www.stanford.edu/dept/HPS/DefecatingDuck.pdf>