From: Elon Musk <erm@spacex.com>
To: Ilya Sutskever <ilyasu@openai.com>
Cc: Greg Brockman <gdb@openai.com>, Sam Altman <shga@ycombinator.com>, "Elon Musk" <erm@openai.com>
Subject: Re: openai biweekly update
Date: Mon, 3 Jul 2017 17:19:42 +0000
Importance: Normal

As mentioned, Jim Keller (who is the real deal) says those guys always oversell their abilities. Before we make any big decisions, I'd like to hear what the Tesla chip team can do. Might be able to exceed the reality of what Cerebras claims to be making.

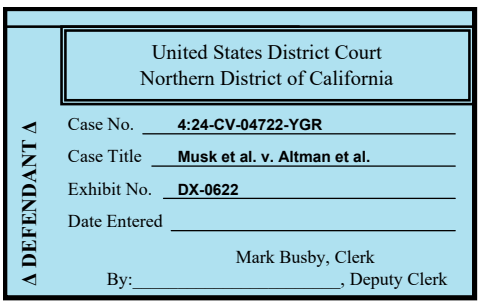

On Jul 1, 2017, at 3:23 PM, Ilya Sutskever <ilyasu@openai.com> wrote:

They do not, at least not yet.

On Sat, Jul 1, 2017 at 3:19 PM Sam Altman <shga@ycombinator.com> wrote:

> Does Cerebras believe AGI is 5 years away with their product?
>
> On 1 July 2017 at 15:16, Greg Brockman <gdb@openai.com> wrote:
>
> > Very slim odds of winning that way. They're 1.5 years in (plus experience trying this at a previous startup). The Cerebras hardware is a breakthrough and we've obtained the details necessary for our assessment under NDA. The other deep learning hardware startups (Graphcore, Groq) are not building anything as ambitious.
> >
> > We should assume that Google is already working on TPU 3.0 (they were working on TPU 2.0 before 1.0 was publicly announced), which is likely to be a generation behind gen 1 of Ceberas. So we estimate that Cerebras is somewhere between 2-3 years ahead of Google, leaving a narrow window of opportunity.
> >
> > We started OpenAI without any special advantage over our competitors — except for the mission. We spent a year learning how to operate and the past 6 months ramping up our execution. Joining forces with Cerebras gives us an unfair advantage over the competition, but starting our own from-scratch effort leaves us again a step behind Google.
> >
> > - gdb
> >
> > On Sat, Jul 1, 2017 at 2:32 PM, Sam Altman <shga@ycombinator.com> wrote:
> >
> > > If we cannot acquire or exclusively partner with Cerebras, can we help start/fund another company to make a similar piece of hardware?
> > >
> > > On 1 July 2017 at 10:27, Greg Brockman <gdb@openai.com> wrote:
> > >
> > > > **tl;dr: We see a path to building AGI in 5 years, several years before everyone else. This will require acquiring or otherwise merging with Cerebras, and for Cerebras to ship no more than 6-12 months behind schedule. If Google instead acquires Cerebras and Cerebras ships, then success is impossible.**
> > > >
> > > > Until recently, we believed that full AGI was 10-15 years in the future. That is, somewhere on the horizon but not imminent.
> > > >
> > > > Our beliefs have changed. **OpenAI sees a concrete path to AGI within 5 years.** We have good

reason to believe that DeepMind thinks that AGI is ~10 years away. Thus, at least for a while, we have a temporary information advantage over Google, DeepMind, and everyone else. We must make the absolute best use of this information advantage.

A few months ago, we realized we could build AGI within 2 years after gaining access to a thousand machines with 1 exaflop and 100 TB memory. A machine of such speed can run a neural net as large as a human brain for several years of thinking in one day. Such a machine has the property that every overnight experiment has a credible chance of producing AGI in the morning. We plan to use this machine in order to run a society of agents developing in a competitive environment. Agents in such an environment will learn to cooperate, communicate, develop an extremely sophisticated language, a theory of mind and excellent social skills, a market economy, a legal system, and grow in general intelligence. In short, they'll develop the cognitive toolkit of the early Sapiens — which we know is sufficient to match humans' ability on every cognitive axis. (This plan is remarkably similar to a Greg Egan short story: http://ttapress.com/553/crystal-nights-by-greg-egan/.)

We recently visited Cerebras, saw their gen 1 hardware, and confirmed that they have a credible path for building an exaflop machine by 2020 using essentially today's technology. Scott Gray and Paul Merolla spent a day digging into their tech and believe that they will deliver with high probability. Their gen 1 hardware is going to be 1000 times faster than the best available GPU while being a lot more flexible overall (though more difficult to program) for the kind of computations that we need in AI. For a tangible example, Cerebras's gen 1 hardware will be able to train a VGG model (the largest visual recognition model, which takes days to a week on an 8 GPU machine) in 3 minutes.

Cerebras is categorically different from other hardware projects. In particular, the TPU 2.0 pales in comparison.

Exclusive access to Cerebras hardware would give OpenAI an overwhelming hardware advantage over Google. By merging, we would create an entity which can build AGI in 5 years, while being several years ahead of others. We acknowledge that things may not go as planned; however, assuming Cerebras executes their hardware reasonably on schedule, then the 5 years timeline has a 50% chance of success, and 7 years has an 80% chance of success. There is a small probability mass on AGI being dramatically harder than we imagine, in which case no one will be there within the next 20 years.

Lastly, if Google acquires Cerebras, there is no path for long term success for OpenAI.

And now, to a quick summary of our progress over the past few weeks.

- **Robotics:** good step towards solving the Rubik's Cube: https://www.youtube.com/watch?v=dwJPZ6OyqUU
- **Multiagent self play:** Progress on competitive self play: https://drive.google.com/open?id=0B60rCy4P2FOIZHRxWE1xQUpoLW8
- We solved adversarial examples, a fundamental problem in machine learning: well-chosen inputs that cause the neural network to produce an incorrect output. For several years, nobody knew how to defend against them. Here is a video of a repair of an adversarial example. We are now immune to adversarial examples: https://drive.google.com/open?id=0B60rCy4P2FOIS3QyQWQ0MWE1UlU
- **Dota 2:** 80% winrate against the best bot from two weeks ago