From: Greg Brockman <gdb@openai.com>
To: Elon Musk <erm@spacex.com>
Cc: "ilyasu@openai.com" <ilyasu@openai.com>, Sam Altman <shga@ycombinator.com>, Sam Teller <steller@spacex.com>, Shivon Zilis <Shivon@spacex.com>
Subject: Re: Top AI institutions today
Date: Wed, 31 Jan 2018 22:56:04 -0800
Importance: Normal
Inline-Images: pastedImage.png

π PLAINTIFF π
United States District Court
Northern District of California
Case No. 4:24-CV-04722-YGR
Case Title Musk et al. v. Altman et al.
Exhibit No. 164
Date Entered
Mark Busby, Clerk
By: , Deputy Clerk

Hi Elon,

Thank you for the thoughtful note. I have always been impressed by your focus on the big picture, and agree completely we must change trajectory to achieve our goals. Let's speak tomorrow, any time 4p or later will work.

My view is that the best future will come from a major expansion of OpenAI. Our goal and mission are fundamentally correct, and that will increasingly be a superpower as AGI grows near.

Fundraising

Our fundraising conversations show that:

- Ilya and I are able to convince reputable people that AGI can really happen in the next <=10 years
- There's appetite for donations from those people
- There's *very* large appetite for investments from those people

I respect your decision on the ICO idea, which matches the evolution of our own thinking. Sam Altman has been working on a fundraising structure that does not rely on a public offering, and we will be curious to hear your feedback.

Of the people we've been talking to, the following people are currently my top suggestions for board members. Would also love suggestions for your top picks not on this list, and we can figure out how to approach them.

- Reid Hoffman
- Gabe Newell
- Adam d'Angelo
- Jed McCaleb
- Herb Allen
- Terah Lyons (she heads Partnership on AI, originally created by Demis to steal OpenAI's thunder — would bring a lot of outside credibility)

The next 3 years

Over the next 3 years, we must build 3 things:

- Custom AI hardware (such as Jim's computer)
- Massive AI data center (likely multiple revs thereof)
- Best software team, mixing between algorithm development, public demonstrations, and safety

We've talked the most about the custom AI hardware and AI data center. On the software front, we have a credible path (self-play in a competitive multiagent environment) which has been validated by Dota and

AlphaGo. We also have identified a small but finite number of limitations in today's deep learning which are barriers to learning from human levels of experience. And we believe we uniquely are on trajectory to solving safety (at least in broad strokes) in the next three years.

We would like to scale headcount in this way:

- Beginning of 2017: ~40
- End of 2018: 100
- End of 2019: 300
- End of 2020: 900

Note that many of DeepMind's 600+ people are not working on AGI, but instead DeepMind Health or DeepMind for Google. Every single one of our people will be working on AGI.

Moral high ground

Our biggest tool is the moral high ground. To retain this, we must:

- Try our best to remain a non-profit. AI is going to shake up the fabric of society, and our fiduciary duty should be to humanity.
- Put increasing effort into the safety/control problem, rather than the fig leaf you've noted in other institutions. It doesn't matter who wins if everyone dies. Related to this, we need to communicate a "better red than dead" outlook — we're trying to build safe AGI, and we're not willing to destroy the world in a down-to-the-wire race to do so.
- Engage with government to provide trusted, unbiased policy advice — we often hear that they mistrust recommendations from companies such as Intel, Google, etc..
- Be perceived as a place that provides public good to the research community, and keeps the other actors honest and open via leading by example.

The past 2 years

I would be curious to hear how you rate our execution over the past two years, relative to resources. In my view:

- Over the past five years, there have two major demonstrations of working systems: AlphaZero [DeepMind] and Dota 1v1 [OpenAI]. (There are a larger number of breakthroughs of "capabilities" popular among practitioners, the top of which I'd say are: ProgressiveGAN [NVIDIA], unsupervised translation [Facebook], WaveNet [DeepMind], Atari/DQN [DeepMind], machine translation [Ilya at Google — now at OpenAI], generative adversarial network [Ian Goodfellow at grad school — now at Google], variational autoencoder [Durk at grad school — now at OpenAI], AlexNet [Ilya at grad school — now at OpenAI].) We benchmark well on this axis.
- We grew very rapidly in 2016, and in 2017 iterated to a working management structure. We're now ready to scale massively, given the resources. We lose people on comp currently, but pretty much *only* on comp. I've been resuming the style of recruiting I did in the early days, and believe I can exceed those results.
- We have the most talent dense team in the field, and we have the reputation for it as well.
- We don't encourage paper writing, and so paper acceptance isn't a measure we optimize. For the ICLR chart Andrej sent, I'd expect our (accepted papers)/(people submitting papers) to be the highest in the field.

- gdb

On Wed, Jan 31, 2018 at 2:02 PM, Elon Musk <erm@spacex.com> wrote:

> OpenAI is on a path of certain failure relative to Google. There obviously needs to be immediate and dramatic action or everyone except for Google will be consigned to irrelevance.
>
> I have considered the ICO approach and will not support it. In my opinion, that would simply result in a massive loss of credibility for OpenAI and everyone associated with the ICO. If something seems too good to

be true, it is. This was, in my opinion, an unwise diversion.

The only paths I can think of are a major expansion of OpenAI and a major expansion of Tesla AI. Perhaps both simultaneously. The former would require a major increase in funds donated and highly credible people joining our board. The current board situation is very weak.

I will set up a time for us to talk tomorrow. To be clear, I have a lot of respect for your abilities and accomplishments, but I am not happy with how things have been managed. That is why I have had trouble engaging with OpenAI in recent months. Either we fix things and my engagement increases a lot or we don't and I will drop to near zero and publicly reduce my association. I will not be in a situation where the perception of my influence and time doesn't match the reality.

Begin forwarded message:

From: Andrej Karpathy <akarpathy@tesla.com>
Date: January 31, 2018 at 1:20:42 PM PST
To: Elon Musk <erm@tesla.com>
Cc: Shivon Zilis <shivon@tesla.com>
Subject: Top AI institutions today

The ICLR conference (which is the top deep learning - specific conference (NIPS is larger, but more diffuse)) released their decisions for accepted/rejected papers, and someone made some nice plots that show where the current deep learning / AI research happens at. It's an imperfect measure because not every company might prioritize paper publications, but it's indicative.

Here's a plot that shows the total number of papers (broken down by oral/poster/workshop/rejected) from any institution:

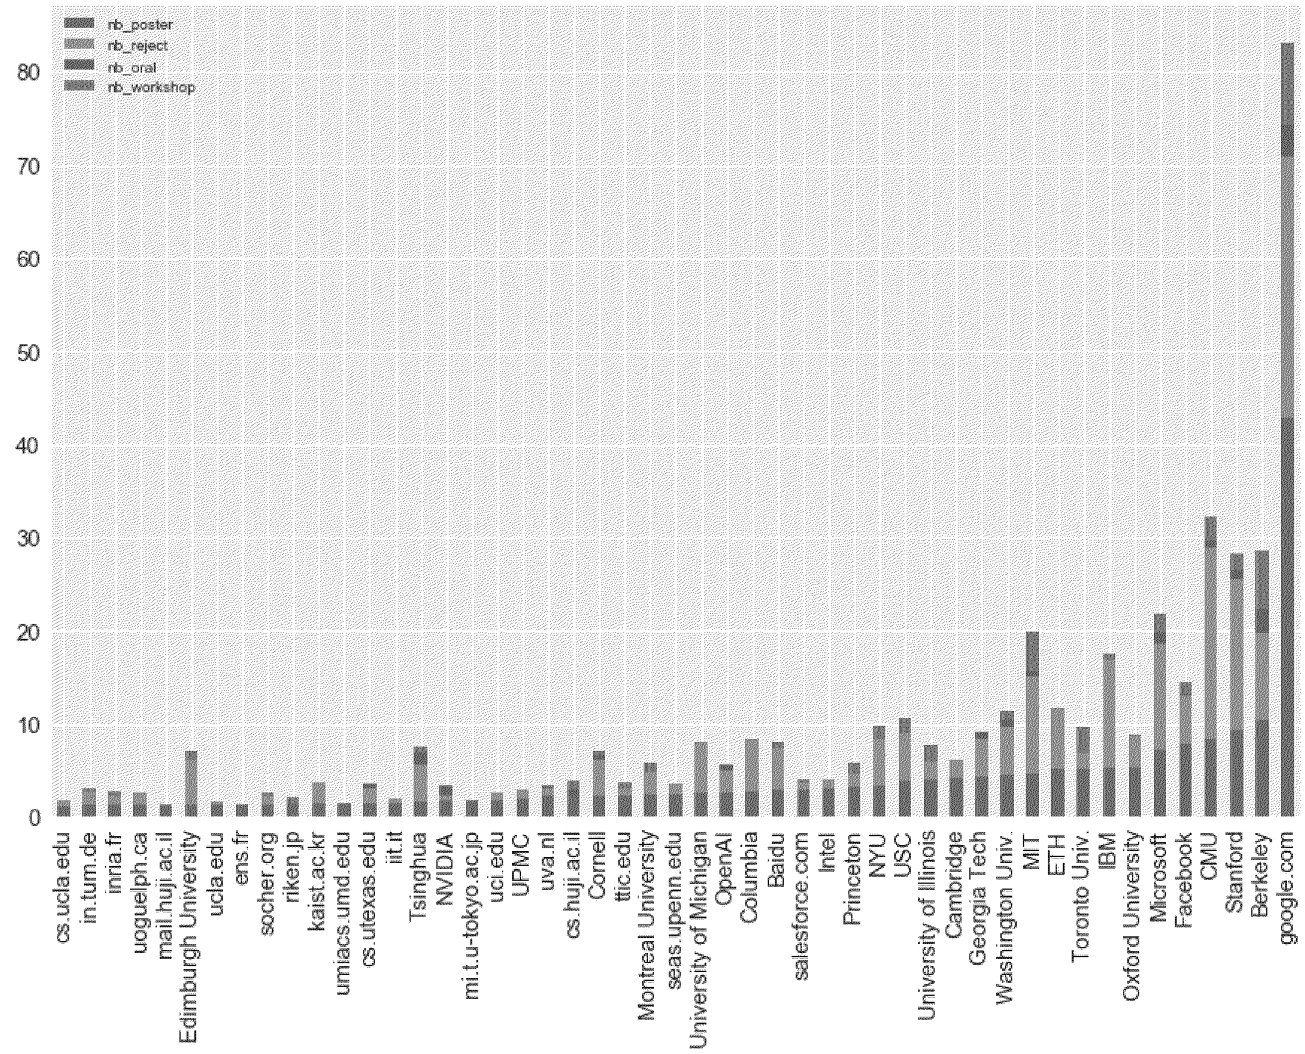

Long story short, Google is dominating with 83 paper submissions. The academic institutions (Berkeley / Stanford / CMU / MIT) are next, in 20-30 ranges each.

Just thought it was an interesting snapshot of where all the action is today. The full data is here: http://webia.lip6.fr/~pajot/dataviz.html

-Andrej