

From: "Elon Musk" <erm@spacex.com>

To: "Ilya Sutskever" <[REDACTED]>

Cc: "sam@ycombinator.com" <sam@ycombinator.com>, [REDACTED]

Greg Brockman

Subject: Re: congrats on the falcon 9

Date: Sat, 02 Jan 2016 17:11:18 -0000

Importance: Normal

United States District Court Northern District of California	
A DEFENDANT A	Case No. <u>4:24-CV-04722-YGR</u>
	Case Title <u>Musk et al. v. Altman et al.</u>
	Exhibit No. <u>DX-0516</u>
	Date Entered _____
By: _____, Deputy Clerk	

Yup

On Jan 2, 2016, at 11:06 AM, Ilya Sutskever <[REDACTED]> wrote:

The article is concerned with a hard takeoff scenario: if a hard takeoff occurs, and a safe AI is harder to build than an unsafe one, then by opensourcing everything, we make it easy for someone unscrupulous with access to overwhelming amount of hardware to build an unsafe AI, which will experience a hard takeoff.

As we get closer to building AI, it will make sense to start being less open. The Open in openAI means that everyone should benefit from the fruits of AI after its built, but it's totally OK to not share the science (even though sharing everything is definitely the right strategy in the short and possibly medium term for recruitment purposes).

On Sat, Jan 2, 2016 at 8:18 AM Elon Musk <erm@spacex.com> wrote:

Begin forwarded message:

From: Demis Hassabis <demishassabis@google.com>

Date: January 2, 2016 at 10:12:32 AM CST

To: Elon Musk <erm@spacex.com>

Subject: congrats on the falcon 9

Hi Elon

Happy new year to you, Talulah and the boys!

Congratulations on landing the Falcon 9, what an amazing achievement. Time to build out the fleet now!

I've seen you (and Sam and other OpenAI people) doing a lot of interviews recently extolling the virtues of open sourcing AI, but I presume you realise that this is not some sort of panacea that will somehow magically solve the safety problem? There are many good arguments as to why the approach you are taking is actually very dangerous and in fact may increase the risk to the world. Some of the more obvious points are well articulated in this blog post, that I'm sure you've seen, but there are also other important considerations:

<http://slatestarcodex.com/2015/12/17/should-ai-be-open/>

I'd be interested to hear your counter-arguments to these points.

Best
Demis