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. 66
Date Entered
Mark Busby, Clerk
By: , Deputy Clerk

From: Greg Brockman [redacted]
on behalf of Greg Brockman [redacted]
Sent: 12/6/2015 4:33:04 PM
To: Sam Altman [sam@ycombinator.com]
Subject: Fwd: Decision

(Good timing on this + Sergey emails.)

- gdb

---------- Forwarded message ----------
From: Ilya Sutskever [redacted]
Date: Sun, Dec 6, 2015 at 7:26 AM
Subject: Re: Decision
To: Greg Brockman [redacted]

Strong, true words.

On Sun, Dec 6, 2015 at 10:16 AM Greg Brockman [redacted] wrote:
>
>
>
> ---------- Forwarded message ----------
> From: Kingma, D.P. [redacted]
> Date: Sunday, December 6, 2015
> Subject: Decision
> To: Greg Brockman [redacted]
> Cc: Sam Altman <sam@ycombinator.com>, John Schulman [redacted]
>
>
> Hi,
>
> I accept! :-)
>
> You might wonder what made me change my mind. First, perhaps most importantly, I realised that doing AI research the way OpenAI intends to, is almost certainly the more ethical choice in the long run. Second, fundamental AI research is simply much more interesting in the long run, and I really enjoy research work. Third, contributing to fundamental AI research will have the greatest expected long-term impact. Also, from my earlier e-mail: you have assembled an extraordinary initial team (and being part of the founding team has some advantages), I think the environment is going to be really special, and Elon's involvement is awesome. It's an honour to be part of this team. And, of course, the weather is a big step up ;-)
>
> Together, these things have made me very excited to join your lab. It's hard to describe how much I'm looking forward to it.
>
> D
>
> On Sat, Dec 5, 2015 at 7:22 PM, Greg Brockman [redacted] wrote:
>>
>> Fantastic news!!!
>>
>> - I've attached a letter with an updated date of Dec 8.
>> - We're flexible on start date. Starting after graduating from your
>> PhD in Summer 2016 works just fine. (We'll of course be happy for you
>> to be as involved as you like prior to then, and could have you on the
>> email lists, etc..)
>>
>> - gdb
>>
>>
>>
>> On Sat, Dec 5, 2015 at 6:17 AM, Kingma, D.P. [redacted] wrote:
>> > Hi Greg,
>> >
>>
>> > Makes sense!
>> >
>> > Some details before I sign:

>> >
>> > - The offer letter has deadline Dec 1st. Not sure if it's legally valid if I
>> > accept that version.
>> >
>> > - Regarding my start date: I can most probably finish my PhD in summer 2016.
>> > Please let me know what your constraints are regarding the start date.
>> >
>> > Regards, Durk
>> >
>>
>> > On Thu, Dec 3, 2015 at 6:55 PM, Greg Brockman [redacted] wrote:
>> >>
>> >> Hey Durk,
>> >>
>> >> Yep, makes sense.
>> >>
>> >> To add to what John said, I strongly agree with your statements. (B)
>> >> is actually the core reason we're starting this lab. Even more
>> >> strongly, on the world's current trajectory, AGI seems most likely to
>> >> come out of a for-profit company, or maybe a government intelligence
>> >> agency as part of an arms race. We have an opportunity to set the
>> >> stage so that doesn't happen: I think it should happen as an
>> >> international coalition, more like the ISS than the space race.
>> >>
>> >> The lab's mission is ensuring that AGI is beneficial, and that those
>> >> benefits are distributed widely to the world rather than making anyone
>> >> into a quadrillion-dollar company or omnipotent surveillance state. I
>> >> don't care if we're the ones that ultimately make AGI, but I care that
>> >> we shape the direction of progress so that AGI ultimately happens in
>> >> the right way.
>> >>
>> >> Achieving this requires being a leading research institution. Our
>> >> initial team is fantastic, and I want to keep the quality bar
>> >> consistent over time: there's an unreasonable power in small,
>> >> tight-knit groups of super talented people (PARC for example did their
>> >> best work with only 24 researchers!). The most important thing for us
>> >> will be producing great research (I'd like us to get to the point
>> >> where whenever we release a paper everyone drops what they're doing to
>> >> read it). But I also think it'll be important for us to influence the
>> >> whole picture for how AGI should happen, for example:
>> >>
>> >> - Open up the dialog about safety in the community: right now, I'm
>> >> pretty sure the majority AI researchers believe safety is worth
>> >> starting to think about, but it's pretty hard to talk about openly due
>> >> to the vocal minority.
>> >> - Set the standard for how AI research should happen (not just for
>> >> AGI, but also for reproducibility, etc), and how to actionably
>> >> incorporate safety into the work.
>> >> - Figure out how to communicate to the mainstream: I think most people
>> >> see the costs of AI (a la Terminator) but don't know what the benefits
>> >> would be. Maybe this requires something crazy like getting more movies
>> >> like Her made.
>> >>
>> >> I'll note it's quite possible that this is the only time one could
>> >> start a group like this - for example, if Google/DeepMind were to
>> >> cement their talent monopoly over the next 5 years, it'd just be
>> >> impossible to get anything new started. (Of course, there are many
>> >> more people entering the field now, so perhaps their monopoly would
>> >> break naturally - it's hard to know right now.)
>> >>
>> >> So I think that now is the right time. I think you are exactly the
>> >> right person to help make this effort succeed. And I think you'd learn
>> >> a ton from people like Elon, Sam, and Alan Kay - some of the best
>> >> people in the world at getting things done.
>> >>
>> >> As a parting thought, here is a quote from Elon (from a private email,
>> >> please keep confidential):
>> >>
>> >> "On the time allocation front, I've decided to dedicate whatever
>> >> amount of my time that is materially useful, even though it will come
>> >> at some cost to SpaceX and Tesla. If I really believe that this is
>> >> potentially the biggest near-term existential threat, then action
>> >> should follow belief."
>> >>
>> >> Anyway, I get to NIPS on Sunday, and would be happy to chat whenever;
>> >> also happy to schedule an Elon call - let me know what would be most
>> >> helpful!
>> >>

>> >> - gdb
>> >>
>> >>
>> >> On Wed, Dec 2, 2015 at 10:32 PM, John Schulman
>> >> [redacted] wrote:
>> >> > Hey Durk, I'm really excited to hear that you're still considering
>> >> > joining
>> >> > us.
>> >> >
>> >> > I think you're the top person in the world in your field--certainly,
>> >> > you've
>> >> > written my favorite papers on generative models and semisupervised
>> >> > learning--so it would be sad if you weren't able to use your hard-earned
>> >> > mastery of these topics to make the next breakthroughs.
>> >> > And if you join us, you'd help us build our research program and culture
>> >> > so
>> >> > we can get the next generation of researchers up to speed, to build on
>> >> > your
>> >> > ideas in various directions.
>> >> >
>> >> > I agree with your P.S. regarding (A) vs (B).
>> >> > Regarding human-level AI, one hard part of this discussion is that we
>> >> > don't
>> >> > really know what we want or what would be a good outcome (at least
>> >> > that's
>> >> > how I feel.) But you've hit the nail on the head that the best thing we
>> >> > can
>> >> > do right now is try to make it happen in a place where people have the
>> >> > right
>> >> > intentions and values that are consistent with ours.
>> >> > This lab will be unique as a top-notch research group where we mutually
>> >> > agree it is worth taking this issue seriously and talking about it
>> >> > candidly,
>> >> > so that as we understand it better, we can try to steer things in the
>> >> > right
>> >> > direction.
>> >> >
>> >> > John
>> >> >
>> >> >
>> >> >
>> >> > On Wed, Dec 2, 2015 at 3:07 PM, Kingma, D.P. [redacted] wrote:
>> >> >>
>> >> >> To avoid any confusion: my previous email was to confirm that my
>> >> >> earlier
>> >> >> decision should indeed not be considered final. thanks again for your
>> >> >> patience.
>> >> >>
>> >> >> D
>> >> >>
>> >> >> On Dec 2, 2015 16:29, "Kingma, D.P." [redacted] wrote:
>> >> >> >
>> >> >> > Hi Greg,
>> >> >> >
>> >> >> > Good points... Let's keep in touch. Thanks for your patience.
>> >> >> >
>> >> >> > D
>> >> >> >
>> >> >> > P.S. Regarding the ethical aspect, I gave it a bit more thought this
>> >> >> > week, and my current take on it is pasted below. TL;DR: it's actually
>> >> >> > not
>> >> >> > unethical to contribute to AI research, as long as it's at the right
>> >> >> > place.
>> >> >> >
>> >> >> > AI technology will eventually have huge positive or negative impact
>> >> >> > on
>> >> >> > society. Therefore, the ethical aspects should not be neglected when
>> >> >> > making
>> >> >> > career choices (like e.g. working on nuclear technology).
>> >> >> >
>> >> >> > One big concern is the possibility of an apocalyptic AI scenario, and
>> >> >> > being partially responsible for it. (Of course, there are also great
>> >> >> > potential benefits from AI). As an individual, there are roughly two
>> >> >> > things
>> >> >> > that one can influence through AI research:
>> >> >> > (A) when AI will arrive. An individual's contribution might bring
>> >> >> > forward the arrival time of powerful AI by at most a few years.
>> >> >> > (B) where powerful AI is developed, by choosing where to research.

>> >> >> > Where
>> >> >> > powerful AI is developed will indirectly greatly affect whether it
>> >> >> > turns out
>> >> >> > OK. This will obviously have an effect over much longer timescales.
>> >> >> >
>> >> >> > In the grand scheme of things, the effect of (A) is actually
>> >> >> > negligible
>> >> >> > compared to (B). Therefore, from an ethical viewpoint, AI researchers
>> >> >> > one
>> >> >> > shouldn't be concerned by (A) (speeding up development of powerful AI
>> >> >> > that
>> >> >> > could later becomes hazardous); rather, career decisions should be
>> >> >> > guided by
>> >> >> > a much greater degree by (B): whether it's developed in the right
>> >> >> > environment: a civilised western institution with some oversight is a
>> >> >> > much
>> >> >> > better environment than some close, obscure, foreign institution
>> >> >> > without
>> >> >> > oversight.
>> >> >> >
>> >> >> >
>> >> >> > On Sun, Nov 29, 2015 at 6:45 AM, Greg Brockman [redacted]
>> >> >> > wrote:
>> >> >>
>> >> >> >> Hey Durk,
>> >> >>
>> >> >> >> I spent some time thinking about this, and was wondering: would you
>> >> >> >> be
>> >> >> >> up for chatting with Elon before you consider your decision final
>> >> >> >> (and
>> >> >> >> presumably after your voice returns :))?
>> >> >>
>> >> >> >> I think it'd be interesting for you to hear why he's decided to get
>> >> >> >> more involved, and what he thinks is necessary for the world to have
>> >> >> >> a
>> >> >> >> good outcome from AI. I think he's probably the person with the best
>> >> >> >> perspective in the world on how to actually achieve great outcomes
>> >> >> >> on
>> >> >> >> some of humanity's hardest problems (and how to compare it to a
>> >> >> >> scientific startup with a large market potential!).
>> >> >>
>> >> >> >> Let me know. I definitely understand if you want to consider your
>> >> >> >> mind
>> >> >> >> made up. At the same time, I've been *really* excited about the
>> >> >> >> chance
>> >> >> >> to work with you, and would hate for you to make up your mind in the
>> >> >> >> absence of important data. You've done extremely impressive work to
>> >> >> >> date, and I want to help you be even more successful.
>> >> >>
>> >> >> >> I think with the right environment, you could achieve truly
>> >> >> >> world-altering things. Don't settle for anything less. (And I could
>> >> >> >> certainly buy that includes the startup as well - I'm not in a
>> >> >> >> position to know.)
>> >> >>
>> >> >> >> I honestly don't know what Elon would advise, but I think it's worth
>> >> >> >> taking seriously whatever he says. I'll of course be more than happy
>> >> >> >> to chat at any time (and hope to hang out at NIPS!).
>> >> >>
>> >> >> >> In any case, regardless of which path you take, I mean it when I say
>> >> >> >> that I'll be as helpful to you as I can.
>> >> >>
>> >> >> >> - gdb
>> >> >>
>> >> >>
>> >> >> >> On Sat, Nov 28, 2015 at 2:02 PM, Kingma, D.P. [redacted]
>> >> >> >> wrote:
>> >> >> >> > Hi Greg,
>> >> >> >> >
>> >> >> >> > Thanks. I must say that part of me is sad as well.
>> >> >> >> >
>> >> >> >> > Regarding Skype: yesterday I completely lost my voice due to a
>> >> >> >> > cold,
>> >> >> >> > and
>> >> >> >> > Skype is infeasible. I just read that it usually takes at least a
>> >> >> >> > couple of
>> >> >> >> > days to heal. I'll let you know as soon as my voice recovers!
>> >> >> >> >
>> >> >> >> > Regards

>> >> >> >> >
>> >> >> >> >
>> >> >> >> >
>> >> >> >> > On Sat, Nov 28, 2015 at 7:37 PM, Greg Brockman
>> >> >> >> > [redacted]
>> >> >> >> > wrote:
>> >> >> >> >>
>> >> >> >> >> Aww, definitely sad to hear it, though I'm excited for you. I'm
>> >> >> >> >> certainly happy to be as useful as I can.
>> >> >> >> >>
>> >> >> >> >> And sure, I'd love to chat on Skype - I'm around the rest of
>> >> >> >> >> today,
>> >> >> >> >> on
>> >> >> >> >> a plane tomorrow, and should be free Monday. When's best?
>> >> >> >> >>
>> >> >> >> >> - gdb
>> >> >> >> >>
>> >> >> >> >>
>> >> >> >> >> On Sat, Nov 28, 2015 at 6:02 AM, Kingma, D.P.
>> >> >> >> >> [redacted]
>> >> >> >> >> wrote:
>> >> >> >> >> > Hi Greg, Sam,
>> >> >> >> >> >
>> >> >> >> >> > This was not an easy decision, but after much deliberation, I
>> >> >> >> >> > decided to
>> >> >> >> >> > go
>> >> >> >> >> > with the automated radiology startup.
>> >> >> >> >> >
>> >> >> >> >> > You have assembled an extraordinary initial team, I think the
>> >> >> >> >> > environment is
>> >> >> >> >> > going to be really special. Of course, working with Elon would
>> >> >> >> >> > be
>> >> >> >> >> > awesome.
>> >> >> >> >> > And of course being part of the founding team comes with a big
>> >> >> >> >> > set
>> >> >> >> >> > of
>> >> >> >> >> > advantages. I really enjoy research work. I'm also honored to
>> >> >> >> >> > be
>> >> >> >> >> > asked
>> >> >> >> >> > for
>> >> >> >> >> > the position. All this together made it a tough decision. Sorry
>> >> >> >> >> > for it
>> >> >> >> >> > taking so long. In the end, the lure and potential of the
>> >> >> >> >> > startup
>> >> >> >> >> > persisted.
>> >> >> >> >> >
>> >> >> >> >> > As I told you, I do plan to keep publishing papers. Also I want
>> >> >> >> >> > to
>> >> >> >> >> > return to
>> >> >> >> >> > go back to do pure AI research eventually. In that case, I'll
>> >> >> >> >> > of
>> >> >> >> >> > course
>> >> >> >> >> > come
>> >> >> >> >> > back begging for a position in your lab :)
>> >> >> >> >> >
>> >> >> >> >> > If you want to do Skype call, please let me know.
>> >> >> >> >> >
>> >> >> >> >> > Regards, Durk
>> >> >> >> >> >
>> >> >> >> >> > P.S. We plan to incorporate in SF next year, and to do a
>> >> >> >> >> > financing
>> >> >> >> >> > round
>> >> >> >> >> > in
>> >> >> >> >> > Q1/Q2 of 2016. However our funding network isn't great in SF.
>> >> >> >> >> > You
>> >> >> >> >> > or Sam
>> >> >> >> >> > might be able to open some doors for us?
>> >> >> >> >> >
>> >> >> >> >> > Some info: one of the co-founders is another top ML researcher,
>> >> >> >> >> > Tim
>> >> >> >> >> > Salimans, who multiple prizes at Kaggle. He is also a well
>> >> >> >> >> > known
>> >> >> >> >> > figure
>> >> >> >> >> > in
>> >> >> >> >> > the Bayesian statistics community, and winner of the
>> >> >> >> >> > prestigious
>> >> >> >> >> > Lindley
>> >> >> >> >> > prize. The others are a biz dev and CS guy, both excellent. Tim

>> >> >> >> >> > just
>> >> >> >> >> > finished a prototype that outperforms radiologists on
>> >> >> >> >> > classifying
>> >> >> >> >> > arthritis
>> >> >> >> >> > in the knee. On the business side is going well too. We would
>> >> >> >> >> > gladly
>> >> >> >> >> > give a
>> >> >> >> >> > talk/presentation when we're in SF to anyone that is
>> >> >> >> >> > interested.
>> >> >> >> >> >
>> >> >> >> >> >
>> >> >> >> >
>> >> >> >> >
>> >> >> >
>> >> >> >
>> >> >
>> >> >
>> >
>> >
>
>
>
>
> --
> Sent from mobile