EXHIBIT E

Message

From: Stella Biderman [stella@eleuther.ai]
Sent: 5/3/2022 6:27:28 PM
To: Stephen Roller [redacted]@fb.com]
Subject: Re: Access to OPT

Stephen,

Wow, this log book is phenomenal. It alone is a massive research contribution. I've wanted to do something like it, but we didn't have a system set up for collecting and managing this data and as every new day went by the task of retroactively doing it became more and more infeasible.

I have more or less infinite storage… if the main limitation to sharing the checkpoints is storage space can I buy you a drive to put them on and come physically pick it up? Unfortunately intermediary check-points is essential to much of the the research I would like to do. I don't have a fixed number in mind, but the more the merrier. If you can only release some of them, it would be preferable to me for them to be logarithmically spaces throughout training instead of linearly spaced.

What's next for you? I would love to find some find ti talk about research priorities and potential future collaborations if that's something you're able to do.

If you ever find yourself in the DC area, definitely let me know. I would love to take you out for a drink!

Stella

On Tue, May 3, 2022 at 12:06 AM Stephen Roller [redacted]@fb.com> wrote:

> Yes a committee of RAI experts has the final say.
>
> We kept checkpoints at every 250 steps and may not be able to share at that granularity for space reasons, but I'll move to offer you intermediate steps.
>
> Thanks for the social media support!
> Stephen
>
> > On May 2, 2022, at 23:46, Stella Biderman <stella@eleuther.ai> wrote:
> >
> > Ah, I see. Am I correct in surmising that is going to be something decided by a committee and probably involving signing paperwork then?
> >
> > I appreciate your vote of confidence, and I hope we'll have opportunities to collaborate in the future. I mentioned partially trained checkpoints in my email but rereading the paper I don't see any mention of that… do you think that that's potentially available? Partially trained checkpoints + training data (in order) is pretty important to the research I want to do.
> >
> > Stella
> >
> > On Mon, May 2, 2022 at 10:57 PM Stephen Roller [redacted]@fb.com> wrote:
> >
> > > Off the record: Follow up regarding your request for weights: I'm not in a position to make that decision alone, but I will personally (speaking only as Stephen) advocate that EleutherAI

is a critical party for receiving the weights. Susan and Naman will take the same position I'm certain.

> On May 2, 2022, at 22:50, Stephen Roller [redacted]@fb.com> wrote:
>
> Thank you Stella!
>
> Thanks for pointing out Figure 3. I wrote that sentence when WIC was included in the average and then I cut it last second. It changes the figure a lot. We will update the caption tomorrow.
>
> I agree duplication is an important research question. I don't pretend to know the answer as to which is best, just report accurately our decisions. If you would like changes to the specific wording of our draft, we want to oblige, but please be extremely explicit in your ask (and please be patient with iterations if necessary). I fully recognize this is an important, open research question.
>
> Very proud to join y'all in the "this should be open" community!
>
> Stephen
>
> > On May 2, 2022, at 22:19, Stella Biderman <stella@eleuther.ai> wrote:
> >
> > Hello,
> >
> > On behalf of EleutherAI, congratulations on announcing the release of OPT! We are very excited for the changes in research that access to such models can bring.
> >
> > I would also like to request access to the suite of models, any intermediary checkpoints you have, and the underlying training data for us to use for research purposes.
> >
> > Two notes on the paper so far:
> > 1. I believe the caption of Figure 3 is incorrect
> > 2. You write "We found the Pile was particularly full of duplicate documents, and ad- vise future researchers using the Pile to perform additional de-duplication processing." The Pile was *deliberately duplicated* as that's what was considered appropriate at the time the work was done (GPT-3 is trained on deliberately duplicated data too). We also released the individual components so that users who wish to use a different mix of the data (either deduplicated or with different duplication weights) can do so.
> >
> > Again, congratulations and I look forward to doing research with these models,
> > Stella Biderman