EXHIBIT SEVENTEEN

Message

From: Ragavan Srinivasan [redacted]@meta.com]
Sent: 1/4/2024 8:01:45 AM
To: Angela Fan [redacted]@meta.com]; Manohar Paluri [redacted]@meta.com]
Subject: FW: [A/C Priv] LibGen Approval for OneLLM

Keeping you both in the loop.

From: Ragavan Srinivasan <[redacted]@meta.com>
Date: Thursday, January 4, 2024 at 12:00 AM
To: Joelle Pineau [redacted]@meta.com>
Subject: Re: [A/C Priv] LibGen Approval for OneLLM

Oops – sorry about the post visibility, I've included it below my reply.

I met with Mary, Gargi and team today. I think the sharing of the dataset is actually lower on my list of concerns. I want to flag an issue that's going to be significantly gnarly that we probably want to discuss and align on now vs. punting it down the road.

It sounds like the OneLLM team is planning on publishing the paper and open sourcing a 7B and a 30B model around the Feb timeline. We are very likely to also launch an open source Llama 7B that should be better than the Mistral 7B also in Feb. This model will ***not*** be multi-modal. I think we will end up confusing the community if we launch both of these around the same time with different claims on performance and capabilities – even if we don't call one of them Llama. I anticipate we will also get a lot of questions from Ahmad, Chris, Mark on this as it will be the first set of launches after our strategy offsite where we aligned on Llama being the release vehicle for our AGI research.

I want to acknowledge that OneLLM is caught during this in-between timeline where we're working out the details of the collaboration model, but I don't want to end up in a really weird situation internally and externally.

Do you think the team appreciates this complexity and a possible outcome being we may not want to open source these models, and even the paper may need to be rewritten when we have a better Llama 7B launching in Feb? I worry that giving the team the dataset without fully working through this might just set them (and us) up for more heartache later down the line.

I can try and work through this with Chris while you're out, but want you to have full visibility into what I'm worried about. Or, we can wait until you're back from vacation, work through this, and then figure out next steps.

What would you prefer?

Ragavan

Full text of the post which apparently is also cross posted to FAIR Internal

Research Artifacts Sharing Principles and Process

TL;DR

- We've outlined principles for sharing research artifacts (data, model, code, *etc.*) between teams in FAIR and GenAI.
- A culture of sharing artifacts is important for the long term success of our research teams.
- But sharing artifacts also creates some responsibilities, in particular, for the team receiving those artifacts.
- Data is as important as models or code – if you collaborate on data that is part of work-in-progress, you are collaborating.

Background

As FAIR and GenAI continue to develop groundbreaking AI research, collaboration is a critical part of our organizations' long-term success and impact. GenAI needs FAIR to be successful as FAIR can incubate some of the long-term ideas that can enable GenAI to leapfrog the competition. And FAIR needs GenAI to be successful in order to test the efficacy of those ideas at massive scale, and bring them to billions of users worldwide. We can only achieve this mutual success if we share the results of our work with each other.

For that reason, part of success for AI research projects is sharing of research artifacts (data, models, code, etc.) between FAIR and GenAI (akin to teams sharing within FAIR or GenAI). The goal of this post is to outline principles so that teams have a low-friction way of collaborating with others and use each other's research artifacts. These principles largely follow Joelle's excellent note on the topic. We note that these principles are different from the engagement model that GenAI uses with other teams across the company; this is by design and reflective of the special relationship between our organizations.

Below, we have outlined principles and a process for teams within FAIR and GenAI to follow. In the principles, we distinguish between post-launch artifacts and pre-launch artifacts. Post-launch artifacts are models, data, code, *etc.* that have been launched internally and/or externally, for example, through an external publication or internal deployment via MetaGen. *Pre-launch artifacts* are models, data, code, *etc.* that have not yet been launched internally and/or externally.

Principles for Research Artifact Sharing

1. If you want to use pre-launch artifacts such as code, models, or data from another team, you have to establish a research collaboration with that team that creates value for both teams.
2. If you use pre-launch artifacts from another team, you cannot launch your work internally or externally without coordinating with the other team and jointly aligning on the launch plan.
3. Be inclusive to people who give you early access to data, models, or code. It is important that people are recognized for their work, for example, via co-authorship.
4. Teams are encouraged to share post-launch artifacts (data, code, models, *etc.)* with other teams in FAIR and GenAI.
5. The sharing of post-launch artifacts should be frictionless. Sharing does not mean there is an expectation of actively supporting the other team, and does not come with an expectation of co-authorship or other markers of collaboration beyond standard acknowledgements.
6. Data and model sharing should not increase Legal and/or PXFN risks – if you want to receive an artifact, you first need to make sure you have all the necessary approvals to use it.
7. Shared data or models should not be shared outside of FAIR or GenAI without explicit approval from both organizations.

GenAI is strongly encouraging product teams to use MetaGen instead of their own custom models. If teams use data or models produced by GenAI and want to deliver product impact, they should deliver that impact through MetaGen.

Process for FYI

Before two teams share pre-launch artifacts, it's important to write a clear document giving visibility to both orgs and CC'ing both research pillar directors. The goal is to make the process as lightweight as possible, while documenting sharing cases clearly. The process is as follows:

1. Two teams agree to share artifacts and agree on all aspects of the collaboration (alignment and timeline on outcomes, for example, publication or internal launch, *etc.*).
2. Receiving team ensures they have all the necessary Legal and PXFN approvals.
3. Teams write up what they agreed on and send an email documenting the agreement to their FAIR pillar lead and GenAI sub-pillar lead (template).

If escalations need to happen because teams cannot reach an agreement, please escalate to your research director and/or VP.

Questions

Post in the comments or reach out to **Laurens van der Maaten** or **Sony Theakanath** with any questions or concerns. This post has been cross posted to GenAI Internal and FAIR Internal.

From: Joelle Pineau [redacted]@meta.com>
Date: Wednesday, January 3, 2024 at 3:00 PM
To: Ragavan Srinivasan [redacted]@meta.com>
Subject: Re: [A/C Priv] LibGen Approval for OneLLM

Hi Ragavan!

Just looping back on your email here. The link to the post doesn't work for me. Can you send again?

Other than that, yes, I'm aligned on this approach.

Let me know if we need anything else to move forward on this.

Thanks!
Joelle

From: Ragavan Srinivasan <[redacted]@meta.com>
Date: Thursday, December 28, 2023 at 2:26 PM
To: Joelle Pineau [redacted]@meta.com>
Subject: Re: [A/C Priv] LibGen Approval for OneLLM

I agree with you that we should be collaborating on things like datasets. I too want to setup the right dynamics between the teams and want to proactively solve for the tensions. What would be helpful is our joint commitment that the OneLLM team will operate per the agreement laid out in this post. It's my understanding that the dataset isn't just the vanilla one available online, but has been curated and preprocessed for use in Llama 3.

Can I say to the Gen AI team that you and I are aligned on this? Don't mean to be pedantic, but I think this will help send the right message on collaboration to both teams.

Thanks,
Ragavan

From: Joelle Pineau [redacted]@meta.com>
Sent: Wednesday, December 27, 2023 4:54:12 PM
To: Ragavan Srinivasan <[redacted]@meta.com>
Subject: Re: [A/C Priv] LibGen Approval for OneLLM

We should give them access to the dataset. We should not be blocking access to datasets, unless they don't have the approvals.
That is setting the wrong dynamics between the teams. Afaik, they can just grab it externally. But that is inefficient (compute, storage).

That doesn't mean they get to short-cut coordinating any release strategy of course!

Cheers,
Joelle

From: Ragavan Srinivasan <[redacted]@meta.com>
Date: Wednesday, December 27, 2023 at 7:34 PM
To: Joelle Pineau [redacted]@meta.com>
Subject: Re: [A/C Priv] LibGen Approval for OneLLM

I think the team was pinging Angela for access to the dataset. Let me check with her on urgency. If it can wait until you're back and we share the general working model, I'd rather we do that. If not, I can try to work with Chris and Sony on this while you're out.

From: Joelle Pineau [redacted]@meta.com>
Date: Wednesday, December 27, 2023 at 4:30 PM
To: Ragavan Srinivasan <[redacted]@meta.com>
Subject: Re: [A/C Priv] LibGen Approval for OneLLM

Hi Ragavan,

Do you need an answer urgently? If so, please work with Sony through these questions, though he may not have all the context on C yet.

Briefly, OneLLM team is well aware that we will need to coordinate release strategy across Llama and OneLLM, including date, narrative, assets, etc. We'll probably need to sort this out after we have shared our general working model. Or if you want to work through this one as a test case for our collaboration (I think that's a good idea!), we would need to loop in Sony (via Chris).

Cheers,
Joelle

From: Ragavan Srinivasan <[redacted]@meta.com>
Date: Wednesday, December 27, 2023 at 5:54 PM
To: Joelle Pineau [redacted]@meta.com>
Subject: FW: [A/C Priv] LibGen Approval for OneLLM

Hi Joelle,

Hope your break is going well. I wanted to sync up with you on this to make sure we're on the same page. My understanding is that:

1/ The OneLLM team is looking to use the libgen dataset that was previously approved only for llama3.

2/ The OneLLM team is planning to publish a paper that will include benchmark results.

3/ The OneLLM team is ***not*** planning to release any models, code or datasets externally under an open source or commercial license.

If the above is accurate, I wanted to clarify a few things:

A/ You have signed off on the additional risk flagged by the legal team in using libgen for #1 & #2.

B/ The timing for #2. Will it happen after the llama3 paper is out? I would assume the answer is yes given the previous agreement I believe the teams have established where the team receiving the pre-release preprocessed libgen data (OneLLM in this case) will stagger their external launch/publication to go after the sharing team (Llama in this case).

C/ Positioning OneLLM vs. Llama. How would we answer questions from the external community on what the relationship is between OneLLM and Llama? I thought this would be a good test case of our new joint goal of getting to AGI and using Llama as the release vehicle for that.

Could you please clarify B and C?

Thanks,
Ragavan

From: Sony Theakanath [redacted]@meta.com>
Date: Wednesday, December 20, 2023 at 3:17 PM
To: Ragavan Srinivasan [redacted]@meta.com>
Subject: FW: [A/C Priv] LibGen Approval for OneLLM

Was asked to send this to you by Angela so you can discuss with Joelle!

From: Sony Theakanath <[redacted]@meta.com>
Date: Tuesday, December 19, 2023 at 11:42 AM
To: Joelle Pineau [redacted]@meta.com>
Cc: Gargi Ghosh [redacted]@meta.com>, Luke Zettlemoyer <[redacted]@meta.com>, Daniel Li (FAIR) [redacted]@meta.com>, Ahuva Goldstand [redacted]@meta.com>, Ammar Rizvi <[redacted]@meta.com>, Elizabeth L. Kendall <[redacted]@meta.com>, Maeve Ryan <[redacted]@meta.com>, Mo Metanat <[redacted]@meta.com>
Subject: Re: [A/C Priv] LibGen Approval for OneLLM

Thanks. CCing other members of policy and legal so we can get started on acquiring the dataset. We'll proceed with Option 2. We drafted up an email and copying that below:

[A/C PRIV]

Hi Joelle,

The team is writing for a decision on proceeding with using libgen for training OneLLM after guidance from Legal and Policy. After a prior escalation to MZ, GenAI has been approved to use LibGen for Llama 3 (with VP sponsor requiring to accept full risk) with a number of agreed upon mitigations. Upon discussion with Legal, the OneLLM team was asked to escalate to you for use of the Libgen dataset including implementation of said mitigations (see below).

Libgen is essential to meet SOTA numbers across all categories, and it is known that OpenAI and Mistral are using the library for their models (through word of mouth). Without LibGen, research scientists within OneLLM believe it would not be able to achieve the SOTA numbers the industry shows. We're beating Llama V2 7B across most benchmarks and are close to their 30B benchmarks, but are not able to reach Mistral. Results from experiments in GenAI suggest that adding Libgen would significantly close this gap, and possibly get us to SOTA.

Model	MMLU	GSM8k	MATH	BoolQ	PiQA	SiQA	Winogrande	OBQA	Hellaswag
Llama V2 7B *Text only*	45.30	14.6	2.50	77.40	78.80	48.30	69.20	58.60	77.20
Mistral 7B **(using Libgen)** *Text only*	60.10	52.10	13.10	83.00	-	75.30	-	81.30	83.00

OneLLM 7B (using Llama 2 data) *Image and text*	46.83	32.98	6.58	82.11	79.43	57.52	69.77	50.00	76.36

The team presents four options. In **no case** would we disclose publicly that we had trained on libgen, however there is practical risk external parties could deduce our use of this dataset, especially if we release models trained on it (e.g. option #3).

Option 1: Use libgen internally only, without external publication of numbers or external model release.
The team would use libgen for internal evaluation, and the team would not release the libgen-trained model externally or publish results in a paper, discuss externally, or allow usage of the model internally, e.g. for production use cases. Purpose is to show internally that we can reach SOTA numbers for MMLU for 7B, provided we have the right datasets, such as LibGen.
Risk: Medium
Requires no acceptance of risk and just an FYI.

Option 2: Use libgen, and publish benchmarked numbers on blog post.
The team would use libgen for internal evaluation and publish results externally. For extreme clarity, this would mean we would publish benchmarks similar to the table above. Additionally, we would not disclose use of LibGen datasets used to train, but there is a possibility that external parties may deduce that we are using libgen. Please note that we will not open source a model with Libgen.
Risk: Medium-High
[Recommended by Research Scientists]
Requires Joelle to accept medium-high legal risks (with mitigations) and high policy risk.

Option 3: Use libgen, publish benchmarked numbers on blog post, and open source a model trained on LibGen.
The team would use libgen for internal evaluation and publish results externally. We would also either research license / commercial license a model. We would not disclose use of Libgen datasets used to train, but there is a possibility that external parties may deduce that we are using libgen.
Risk: Medium-High
Requires Joelle to accept Medium-high legal risk (with mitigations) and high policy risk .

Option 4: Do not proceed with libgen.
Risk: Low

Legal and policy (Ahuva Goldstand, Elizabeth L. Kendall, Maeve Ryan, Olena Ripnick-O'Farrell) have outlined the legal and policy risks and mitigations:

Legal Risk:

Redacted

Redacted

Mitigations:

1. Remove data clearly marked as pirated/stolen
2. Do not externally cite the use of any training data including LibGen
3. Perform evaluations and appropriate mitigations as necessary to reduce memorization metrics (e.g. propensity to generate outputs containing copyrighted material). The memorization rates should meet the same thresholds achieved by Llama 2 team for which leadership has already accepted the risk.

Policy Risks:

1. Legislative/lobbying:

Risks:

Copyright and IP is top of mind for legislators around the world, including in the US and the EU.

US legislators expressed concern in a recent hearing about AI developers using pirated websites for training. It's unclear what their legislative actions would be if the concern spreads, but it reflects some of the negative lobbying rights holders have been doing, related to our litigation on this topic (along the lines that this is "stolen" content that then taints the output of the model).

If there is media coverage suggesting we have used a dataset we know to be pirated, such as LibGen, this may undermine our negotiating position with regulators on these issues.

The recently agreed EU AI Act will include measures requiring disclosure of training data, and the right for rights holders to opt out of their material being used for training AI models. The latter is expected to extend beyond the EU's borders, meaning rights holders outside of Europe (many of whose works are in LibGen) could leverage this clause. **The penalty for non-compliance related to requirements on data is up to €35m or 7% of the total worldwide annual turnover of the preceding financial year (whichever is higher).**

- Misuse:

1. Risks:

A secondary concern is about the potential malicious applications of a model trained on large quantities of scientific data. A 2022 paper on the risk of misuse of AI drug discovery models was widely covered in the media, and subsequent papers have highlighted that training on large corpuses of medical data can elevate risks relating to bioweapons and CBRNE. One proposed mitigation is to remove some of the data that would enable these capabilities (e.g. toxicity datasets), which represented a fraction of the overall data in the paper in question (see p4).

If the model is to be used internally for evaluation only, then this risk may not warrant further action. However, if a model trained on LibGen is to be used for anything other than internal evaluation then it would need to be red teamed for these risks.

- Potential mitigations:

In its Responsible Scaling Policy, Anthropic proposes containment measures for models that could present a risk of 'catastrophic misuse' (see ASL 3) which include internal controls on who can access information about the model. OpenAI propose a similar approach in its Preparedness policy. This might be something to consider even for a version that is for internal evaluation only.

If we were to use models trained on LibGen for a purpose other than internal evaluation, we would need to red team those models for bioweapons and CBRNE risks to ensure we understand and have mitigated risks that may arise from the scientific literature in LibGen.

As well as providing assurances, this would be a way to demonstrate that we have upheld commitments to developing and deploying AI responsibly, in the event that we are forced by future legislation to disclose that we used LibGen.

We might also consider filtering the dataset to reduce risks relating to both bioweapons and CBRNE, and also based on the year that a publication was added to reduce the risk that it is trained on papers that include harmful stereotypes. It is also possible that a model trained on LibGen may produce inaccurate scientific sounding information, and this is again something we should test for.

Thanks,
Sony

From: Joelle Pineau [redacted]@meta.com>
Date: Tuesday, December 19, 2023 at 11:11 AM
To: Sony Theakanath <[redacted]@meta.com>
Cc: Gargi Ghosh <[redacted]@meta.com>, Luke Zettlemoyer <[redacted]@meta.com>, Daniel Li (FAIR) [redacted]@meta.com>, Ahuva Goldstand [redacted]@meta.com>
Subject: Re: [A/C Priv] LibGen Approval for OneLLM

Hi,

I've discussed this with Gargi. Given this is already approved for Llama3, and the usage is the comparable, I am supportive of usage for OneLLM also.

Best,
Joelle

From: Sony Theakanath [redacted]@meta.com>
Date: Wednesday, December 13, 2023 at 3:34 PM
To: Joelle Pineau [redacted]@meta.com>
Cc: Gargi Ghosh <[redacted]@meta.com>, Luke Zettlemoyer [redacted]@meta.com>, Daniel Li (FAIR) [redacted]@meta.com>, Ahuva Goldstand [redacted]@meta.com>
Subject: Re: [A/C Priv] LibGen Approval for OneLLM

Joelle - as an update I'm collecting more info on proportionality of the risks to benefits for OneLLM. The group will come back with a decision on whether to continue to escalate given the high policy and legal risk of the dataset.

Sony

From: Sony Theakanath <[redacted]@meta.com>
Date: Wednesday, December 13, 2023 at 11:12 AM
To: Joelle Pineau <[redacted]@meta.com>
Cc: Gargi Ghosh [redacted]@meta.com>, Luke Zettlemoyer [redacted]@meta.com>, Daniel Li (FAIR) [redacted]@meta.com>, Ahuva Goldstand [redacted]@meta.com>
Subject: LibGen Approval for OneLLM

Hi Joelle,

The OneLLM team would like to use LibGen to improve multi task text understanding and generation quality. After talking with Legal (Ahuva), she requested the team to send you an email for approval. Llama 3 is approved for usage and we'll be using the same dataset downloaded by Melanie's team.

I'll let Ahuva add any additional context.

Sony