Meeting of the LF AI & Data Technical Advisory Council (TAC)

June 16, 2022

⊡LF AI & DATA

Antitrust Policy

› Linux Foundation meetings involve participation by industry competitors, and it is the intention of the Linux Foundation to conduct all of its activities in accordance with applicable antitrust and competition laws. It is therefore extremely important that attendees adhere to meeting agendas, and be aware of, and not participate in, any activities that are prohibited under applicable US state, federal or foreign antitrust and competition laws.

› Examples of types of actions that are prohibited at Linux Foundation meetings and in connection with Linux Foundation activities are described in the Linux Foundation Antitrust Policy available at http://www.linuxfoundation.org/antitrust-policy. If you have questions about these matters, please contact your company counsel, or if you are a member of the Linux Foundation, feel free to contact Andrew Updegrove of the firm of Gesmer Undergone LLP, which provides legal counsel to the Linux Foundation.

Recording of Calls

Reminder:

TAC calls are recorded and available for viewing on the TAC Wiki

LF AI & DATA

Reminder: LF AI & Data Useful Links

› Web site: lfaidata.foundation
› Wiki: wiki.lfaidata.foundation
› GitHub: github.com/lfaidata
› Landscape: https://landscape.lfaidata.foundation or https://l.lfaidata.foundation
› Mail Lists: https://lists.lfaidata.foundation
› Slack: https://slack.lfaidata.foundation
› Youtube: https://www.youtube.com/channel/UCfasaeqXJBCAJMNO9HcHfbA
› LF AI Logos: https://github.com/lfaidata/artwork/tree/master/lfaidata
› LF AI Presentation Template: https://drive.google.com/file/d/1eiDNJvXCqSZHT4Zk_-czASlz2GTBRZk2/view?usp=sharing

› Events Page on LF AI Website: https://lfaidata.foundation/events/
› Events Calendar on LF AI Wiki (subscribe available): https://wiki.lfaidata.foundation/pages/viewpage.action?pageId=12091544
› Event Wiki Pages: https://wiki.lfaidata.foundation/display/DL/LF+AI+Data+Foundation+Events

LF AI & DATA

Agenda

› Roll Call (2 mins)

› Approval of Minutes from previous meetings (2 mins)

› Beyond ML – new sandbox project proposal (35 minutes)

› Retiring the Acumos Project – proposal (10 minutes)

› Update from the Outreach Committee (5 minutes) (tentative)

› LF AI General Updates (2 min)

› Open Discussion (1 min)

TAC Voting Members - Please note

Please ensure that you do the following to facilitate smooth procedural quorum and voting processes:

- Change your Zoom display name to include your First/Last Name, Company/Project Represented
  - example: Nancy Rausch, SAS
- State your First/Last Name and Company/Project when submitting a motion
  - example: First motion, Nancy Rausch/SAS

LF AI & DATA

TAC Voting Members

Note: we still need a few designated backups specified on wiki

Member Representatives (9 out of 17 required for quorum)

Member Company or Graduated Project	Membership Level or Project Level	Voting Eligibility	Country	TAC Representative	Designated TAC Representative Alternates
4paradigm	Premier	Voting Member	China	Zhongyi Tan	
Baidu	Premier	Voting Member	China	Ti Zhou	Daxiang Dong, Yanjun Ma
Ericsson	Premier	Voting Member	Sweden	Rani Yadav-Ranjan	
Huawei	Premier	Voting Member	China	Howard (Huang Zhipeng)	Charlotte (Xiaoman Hu) , Leon (Hui Wang)
Nokia	Premier	Voting Member	Finland	@ Michael Rooke	@ Jonne Soininen
OPPO	Premier	Voting Member	China	Jimin Jia	
SAS	Premier	Voting Member	USA	*Nancy Rausch	JP Trawinski
ZTE	Premier	Voting Member	China	Wei Meng	Liya Yuan
Acumos Project	Graduated Technical Project	Voting Member	USA	Amit Kumar	Prasanna Kulkarni
Adversarial Robustness Toolbox Project	Graduated Technical Project	Voting Member	USA	Beat Buesser	
Angel Project	Graduated Technical Project	Voting Member	China	Bruce Tao	Huaming Rao
Egeria Project	Graduated Technical Project	Voting Member	UK	Mandy Chessell	Nigel Jones, David Radley, Maryna Strelchuk, Ljupcho Palashevski, Chris Grote
Flyte Project	Graduated Technical Project	Voting Member	USA	Ketan Umare	
Horovod Project	Graduated Technical Project	Voting Member	USA	Travis Addair	
Milvus Project	Graduated Technical Project	Voting Member	China	Xiaofan Luan	Jun Gu
ONNX Project	Graduated Technical Project	Voting Member	USA	Alexandre Eichenberger	Prasanth Pulavarthi, Jim Spohrer
Pyro Project	Graduated Technical Project	Voting Member	USA	Fritz Obermeyer	

*Current TAC Chairperson

LF AI & DATA

Minutes approval

Approval of May 5, 2022 Minutes

Draft minutes from the May 5th TAC call were previously distributed to the TAC members via the mailing list

Proposed Resolution:

> › That the minutes of the May 5th meeting of the Technical Advisory Council of the LF AI & Data Foundation are hereby approved.

Approval of June 2, 2022 Minutes

Draft minutes from the June 2 TAC call were previously distributed to the TAC members via the mailing list

Proposed Resolution:

› That the minutes of the June 2 meeting of the Technical Advisory Council of the LF AI & Data Foundation are hereby approved.

LF AI & DATA

Beyond ML

New Project Proposal at the Sandbox level

Presenter: Jacob Renn – jacob.renn@squared.ai

Resources: Source Code - https://github.com/Beyond-ML-Labs

Sandbox Project Hosting Proposal Review/Discussion/Vote: BeyondML

Description:

BeyondML is a framework for developing sparse neural networks that can perform multiple tasks across multiple data domains. This framework simplifies the development and deployment of advanced machine learning capabilities on low-end devices and in dynamic environments characteristic of the resource-constrained edge.

Presenter: Jacob Renn – jacob.renn@squared.ai

Resources: Source Code - https://github.com/Beyond-ML-Labs

Proposal Link - https://github.com/lfai/proposing-projects/blob/master/proposals/BeyondML_proposal

Why contribute BeyondML to Linux Foundation?

Neutral holding ground

- Vendor neutral, not for profit

Open governance model

- Transparent and open governance model

- Instill trust in contributors and adopters in the management of the project

- Neutral management of projects' assets by the foundation

Growing community

- Increase visibility of the project through LF ecosystem

- Increase contributors by converting new & existing users

- Opportunities to collaborate with other hosted projects

BeyondML

A framework for Multitask, Multimodal ML Built for Robust Performance at the Edge

Background: Neural Nets have a Size Problem

- ML models are vastly overparameterized, requiring ever more computational power
  - Everything is about MORE, MORE, MORE, optimization concerns minimizing loss not compute

- Implications:
  - High cloud utilization costs
  - Energy inefficiency
  - Unable to get SOTA performance without high-end hardware

Motivation for our work

- Simplify model deployment
  - Reduce the complexity of deploying multiple models by combining multiple tasks into a single, easily-deployed neural network

- Minimize the costs of cloud resource utilization
  - Reduce the hardware requirements of running a model or system of models
  - Reduce the number of necessary endpoints from n to 1

- Bring SOTA performance to edge devices
  - Empower autonomous systems with more flexible and robust decision making

- Reduce the carbon footprint of deployed neural networks
  - Fewer parameters = lower computational cost = fewer cycles on less powerful hardware needed

BeyondML: Overview

A framework to...

- ## Make neural networks more computationally efficient
  - o Utilize sparse tensors to represent multiple tasks in the same network, reducing computational resource requirements without sacrificing model performance

- ## Create Multimodal Multitask Neural Networks
  - o Allow multiple tasks to be learned by a single sparse neural network. Crucially, NNs developed by this framework are not susceptible to catastrophic interference and can learn to perform multiple disparate tasks on multiple disparate data domains
  - o We have shown that we can use BeyondML to train a single neural network capable of learning both a computer vision classification and a numerical regression task, and performing both with no performance decrement when compared with representing each task with an individual neural network

BeyondML: Multitask Representational Structure
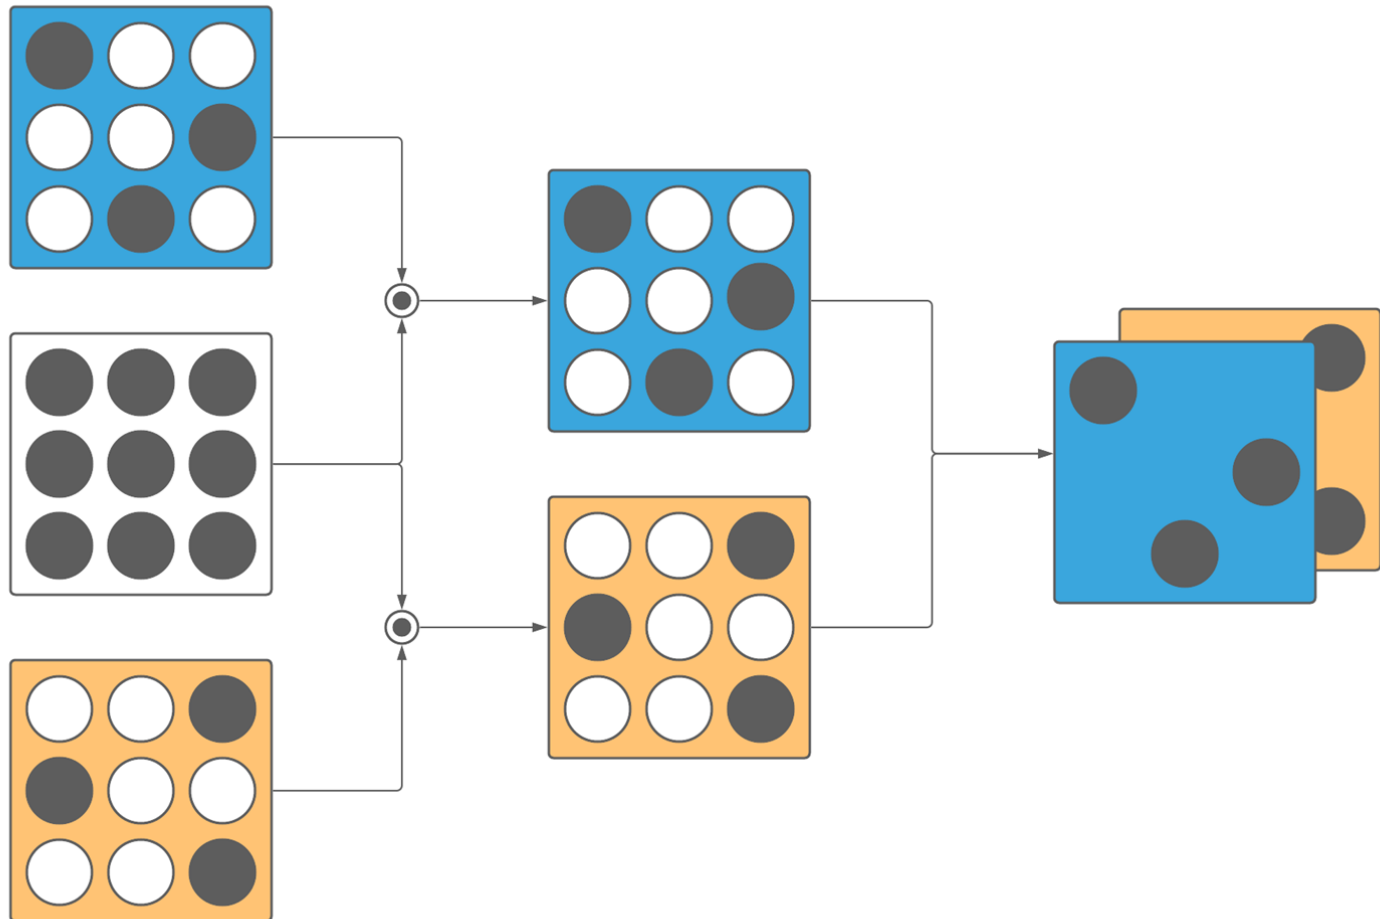
What is Currently Working?

- Support for Tensorflow layers and models
- Support for combining multiple tasks as a single neural network
- Support for several pruning methodologies to remove unnecessary parameters
- Support for model quantization
- Support for transformer models

Towards a Sparse Multitask Inference Framework

PyTorch

- Lacks edge support
- Less mature in deploying to different languages and environments (mobile, JavaScript, etc.)

TensorFlow

- Lacks support for sparse tensors as network weights
- Rigid WRT multi-layered weights & channeling of these weights

> We are leveraging the best features of industry-leading frameworks, but are working towards adding the necessary ingredients for creating sparse multitask models

16JUN2022

What is Currently Missing? - Roadmap

- Supporting PyTorch models and layers
  - Tentative ETA 07/2022

- Implement support for sparse tensors
  - Tentative ETA 11/2022

- Creation of sparse multitask inference framework
  - Tentative ETA 04/2023

- Support for between-task subnetwork communication
  - Tentative ETA 08/2023

Synergies
with other
LF AI &
Data
Projects

16JUN2022
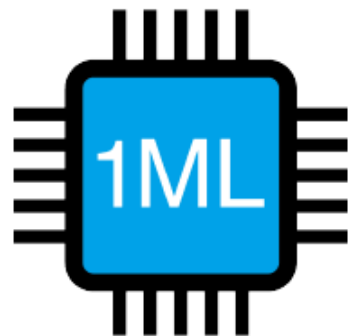
TAC Vote on Project Incubation Proposal

Proposed Resolution:

The TAC approves BeyondML as a Sandbox project of the LF AI & Data Foundation

History

- The Acumos project was the first project to join LF AI & Data in 2018 via AT&T and Tech Mahindra (both founding members of LF Deep Learning)

- Over the next 3+ years, Acumos grew and was the center of collaboration for a number of organizations that worked together, contributed to the project, and deployed it either internally or in products/services

The Linux Foundation Launches Open Source Acumos AI Project

Delivers a Framework and Marketplace to Empower Data Scientists and Developers to Build the Future of Artificial Intelligence

LOS ANGELES (Open Networking Summit), March 26, 2018

Current situation

- For the past year until today:
  - Interest in Acumos has decreased significantly
  - AT&T and TechM are no longer members of LF AI & Data
  - Development on the project came to a complete halt a few months ago
  - Acumos has contributed to the ecosystem, but there are more actively maintained projects
  - The ongoing IT costs of the project are too high (~270K/year)

Sep 2019	Dec 2019	Mar 2020	Jun 2020	Sep 2020	Dec 2020	Mar 2021	Jun 2021	Sep 2021	Dec 2021	Mar 2022	Jun 2022

LOC Added   **LOC Deleted**
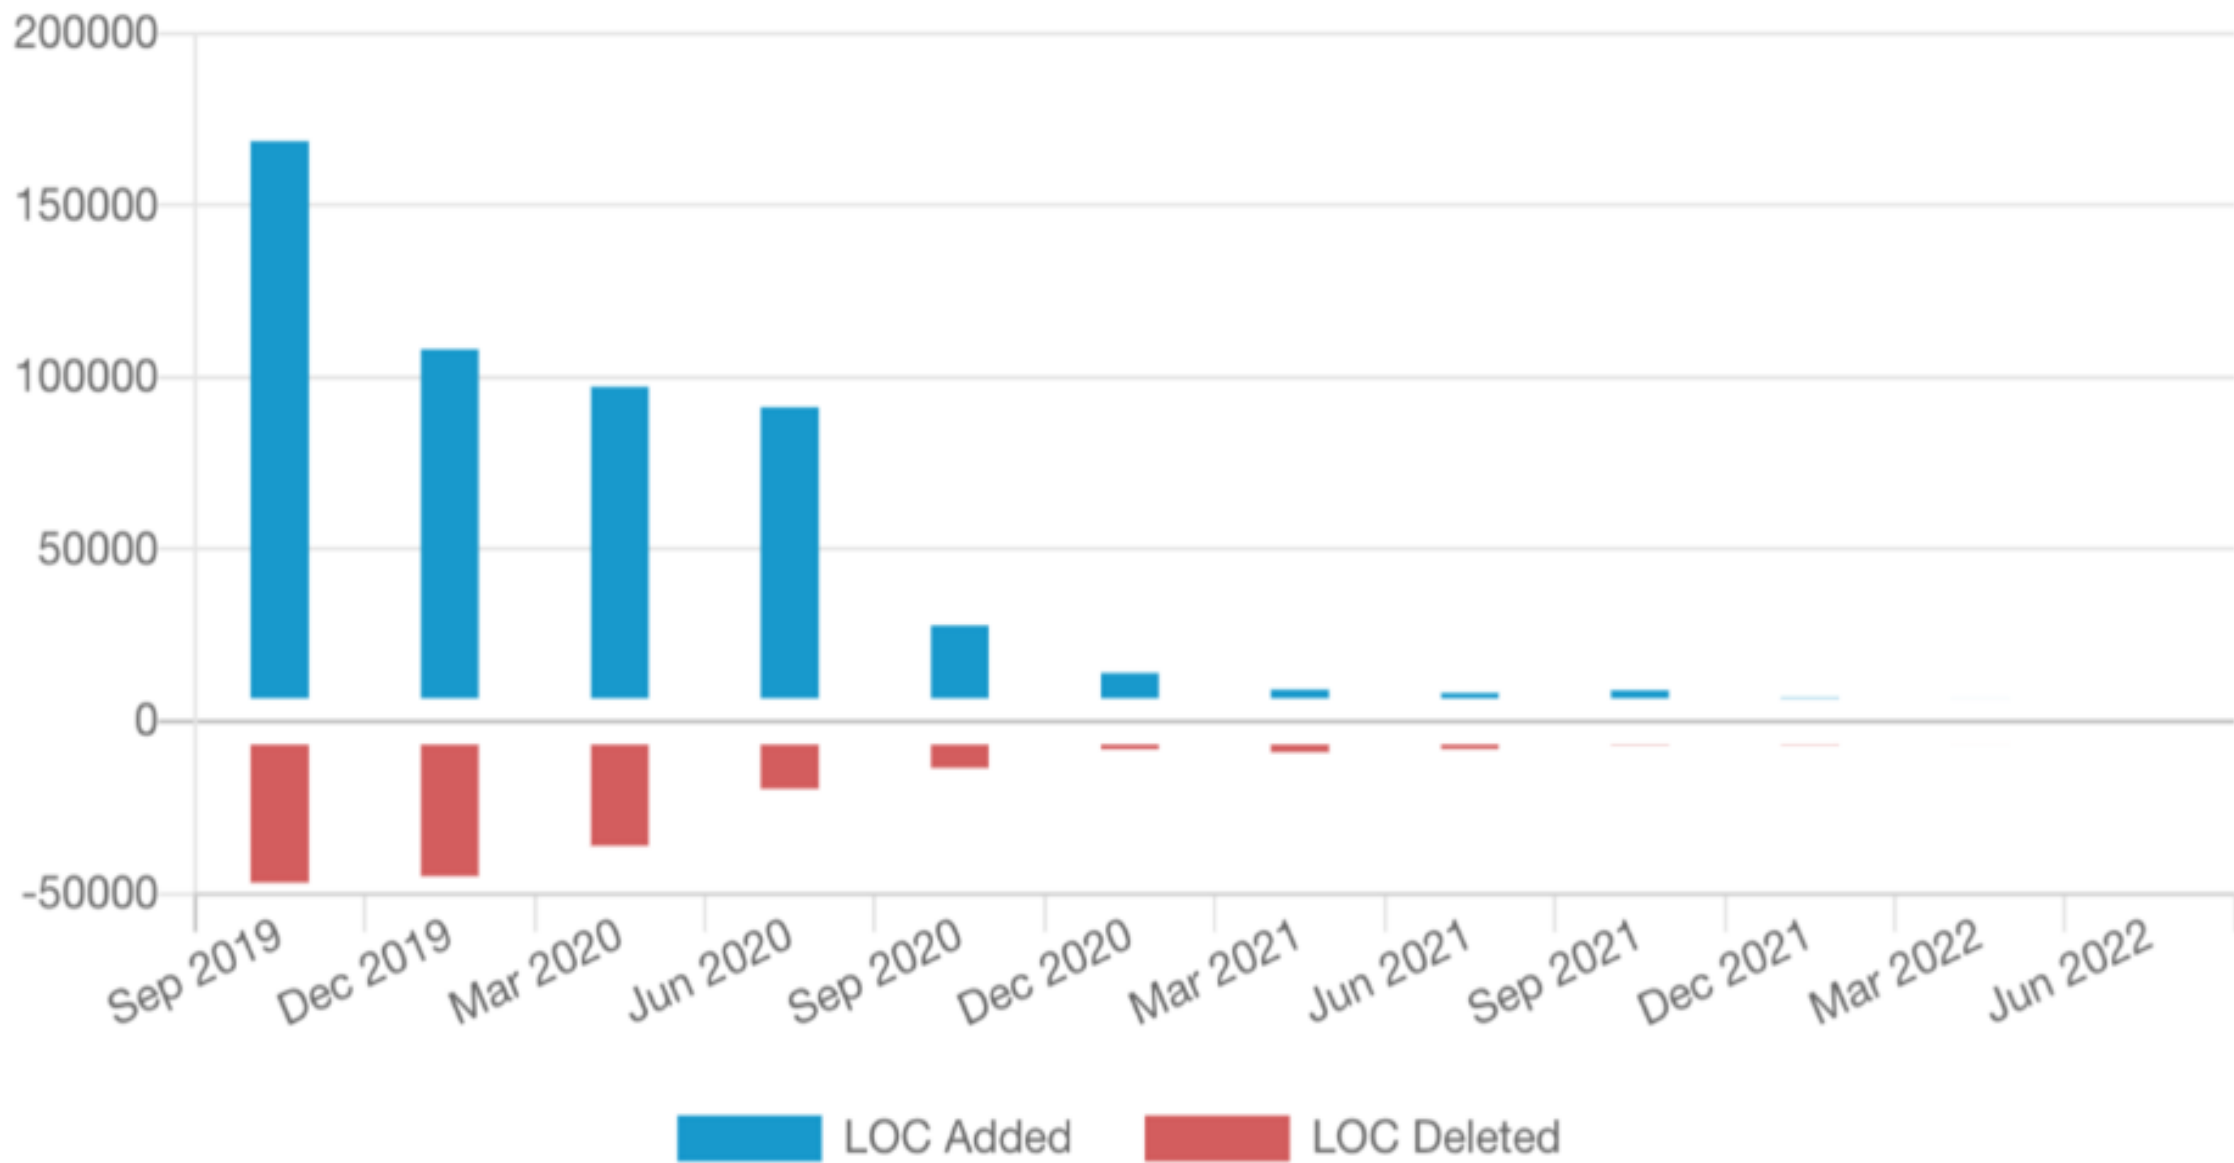
Proposal

- We would like to propose to retire the Acumos project into the Emeritus stage.
- Projects in this stage are not in active development. Their maintainers may infrequently monitor their repositories, and may only push updates to address security issues, if at all.
- The foundation will continue to hold any project assets including trademarks and domains, but the project does not draw on foundation resources.

Plan to Shutdown Acumos IT

Wednesday and Thursday, June 8-9:

- Shut down Jenkins & Sandbox
- Begin full replication from Gerrit to Github (currently only merged changes are present; we want all change information).
- Replicate all docker images from Nexus 3 to DockerHub.

Friday June 10:

- Shut down Gerrit.
- Shut down Nexus 3.
- Move maven artifacts from Nexus 2 to S3 storage.

Mon & Tue June 13-14:

- Validate Nexus 2 backups, then shut down Nexus 2.

Wednesday June 15:

- Remove Jenkins & Sandbox systems.

Friday June 17:

- Remove Gerrit and Nexus 3 systems.

Tue June 21:

- Remove Nexus 2 system.
- Backup and then shut down Jira and Confluence.

Resolution Vote

Proposed Resolution

The LF AI & Data Foundation Technical Advisory Council approves the transition of the Acumos project from the Graduation stage into the Emeritus stage.

Thank you!

The Foundation appreciates the contributions of Acumos, the organizations that supported the project and the Acumos community, and the role they have played in moving the ecosystem forward.

Upcoming TAC Meetings

Upcoming TAC Meetings

› June 30

  › FateAI – Sandbox Proposal from BAAI

  › Using LFX Insights - Jen Shelby (LF)

Please note we are requesting special topics for future meetings.

If you have a topic idea or agenda item, please send agenda topic requests to tac-general@lists.lfaidata.foundation

LF AI & DATA

Open Discussion

□LF AI & DATA

TAC Meeting Details

› To subscribe to the TAC Group Calendar, visit the wiki:
https://wiki.lfaidata.foundation/x/cQB2 _____

› Join from PC, Mac, Linux, iOS or Android: https://zoom.us/j/430697670

› Or iPhone one-tap:

  › US: +16465588656,,430697670# or +16699006833,,430697670#

› Or Telephone:

  › Dial(for higher quality, dial a number based on your current location):

  › US: +1 646 558 8656 or +1 669 900 6833 or +1 855 880 1246 (Toll Free) or +1 877 369 0926 (Toll Free)

› Meeting ID: 430 697 670

› International numbers available: https://zoom.us/u/achYtcw7uN

LF AI & DATA

Legal Notice

□LF AI & DATA

16JUN2022