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

August 10, 2023

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

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

Agenda

- Roll Call (1 mins)
- Approval of Minutes from previous meeting (2 mins)
- LF Edge Update 10 minutes
- DeepCausality proposal for Incubation 45 minutes
- Open Discussion

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

TAC Voting Members - Please note

- TAC members must attend consistently to maintain their voting status
- After 2 absences voting members will lose voting privileges
- Voting privileges will only be reinstated after attending 2 meetings in a row

TAC Voting Members

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

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	Jun Zhang	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)
IBM	Premier	Voting Member	USA	Susan Malaika	Beat Buesser, Alexandre Eichenberger
Nokia	Premier	Voting Member	Finland	@Michael Rooke	@Jonne Soininen
OPPO	Premier	Voting Member	China	Jimmy (Hongmin Xu)	
SAS	Premier	Voting Member	USA	*Nancy Rausch	Liz McIntosh
ZTE	Premier	Voting Member	China	Wei Meng	Liya Yuan
Adversarial Robustness Toolbox Project	Graduated Technical Project	Voting Member	USA	Beat Buesser	Kevin Eykholt
Angel Project	Graduated Technical Project	Voting Member	China	Jun Yao	
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	Andreas Fehlner, Prasanth Pulavarthi, Jim Spohrer
Pyro Project	Graduated Technical Project	Voting Member	USA	Fritz Obermeyer	
Open Lineage Project	Graduated Technical Project	Voting Member	USA	*Awaiting confirmation from Project Lead*	

LF AI & DATA

Minutes approval

LF AI & DATA

Approval of July 27, 2023 Minutes

Draft minutes from the July 27 TAC call were previously distributed to the TAC members via the mailing list

Proposed Resolution:

- That the minutes of the July 27 meeting of the Technical Advisory Council of the LF AI & Data Foundation are hereby approved.

AI & Edge/IOT

Sunny Cai, Joseph Pearson

LF Edge Projects

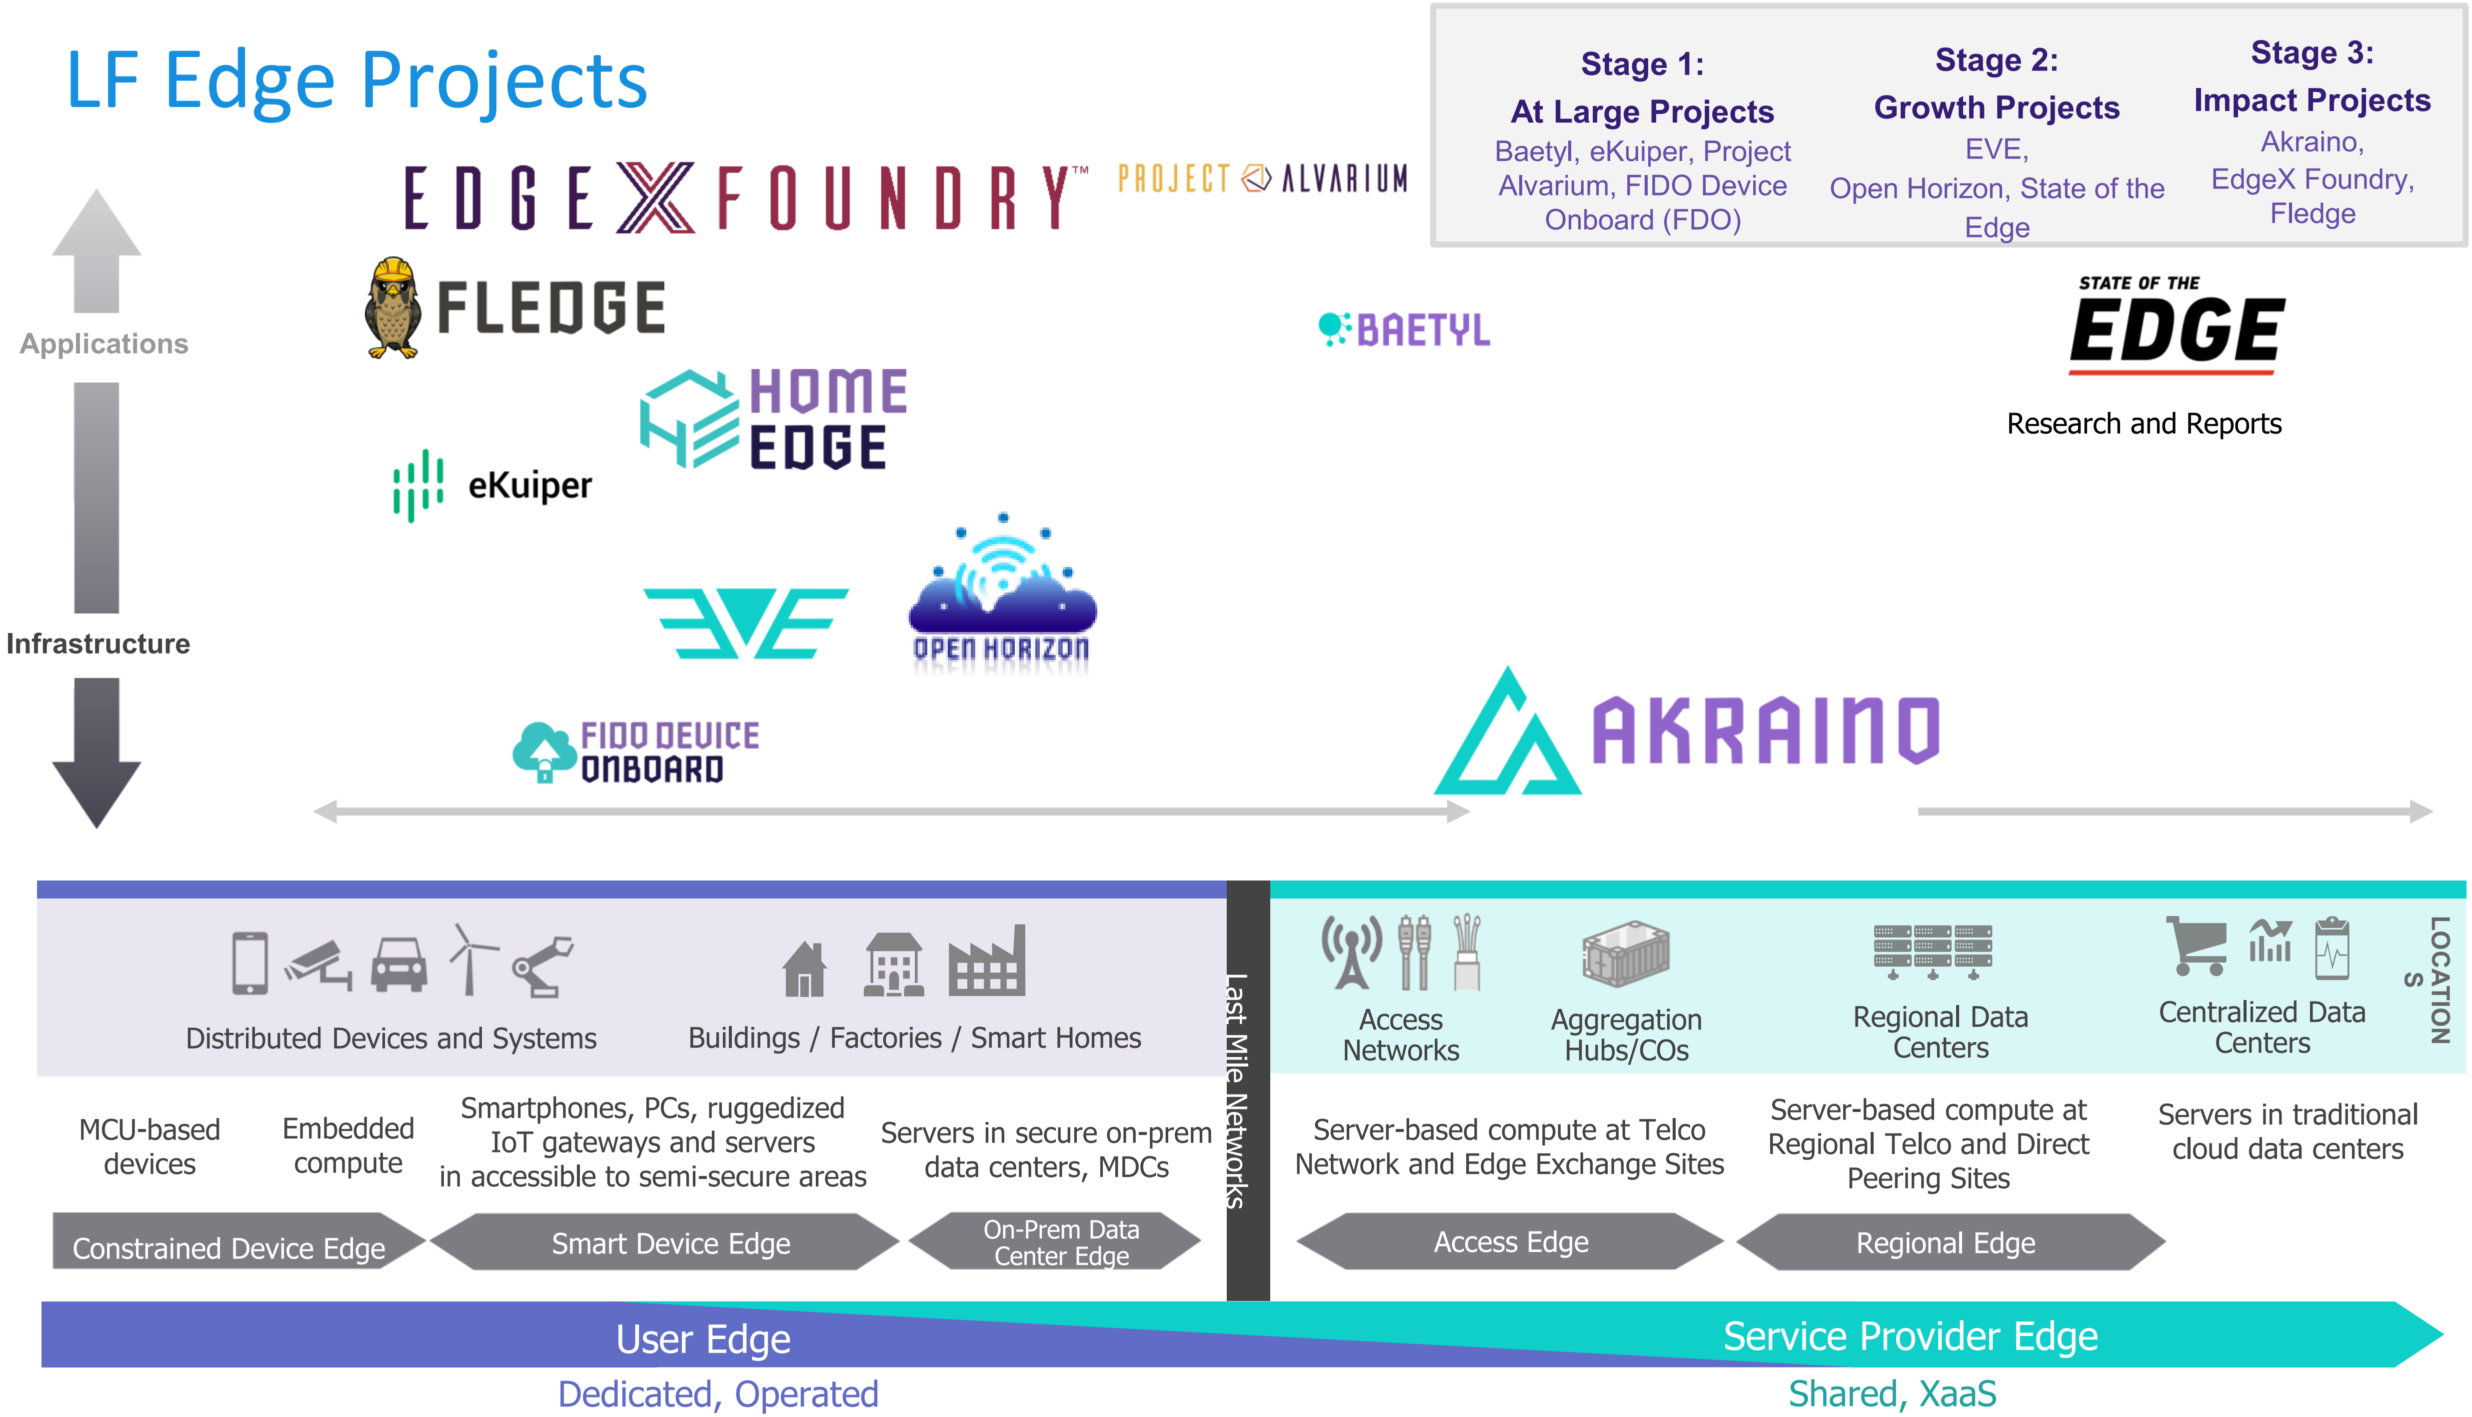

Motivation for AI focus in LF Edge

- The convergence of AI and IoT also has attracted significant investment in recent years.
    - According to a report by IDC, global spending on AI and IoT is projected to reach $1.1 trillion this year.
    - This substantial investment reflects the growing recognition of the transformative power of AI and IoT when combined.

Driving LF Edge's Thought Leadership in AI

- A few different types of usage of AI in our LF Edge projects:
  - Integrating AI into an Edge platform to assist with analytics: EdgeX Foundry, Fledge
  - Managing AI applications on the Edge - Deploying models consumed by applications: Open Horizon, Baetyl
  - Optimized Edge infrastructure for AI: Akraino
- Several of the hot topics in AI are highly aligned with the LF Edge scope:
  - Model Federation at the edge
  - Foundational Models to accelerate and optimize model development (and how it applies to the edge)

Some of the most discussed Edge AI use cases

Industrial	Predictive Maintenance using failure prediction, Quality Control with real time detection of flaws, Equipment efficiency, Yield Optimization by re-calibrating machines based on applying ML algorithms to sensor data
Healthcare	Monitoring hospital rooms, fall detection for assisted living, AI inference at the edge for radiology,cardiovascular and skeletal imaging
Smart Home	Keep data local and perform AI analytics to detect unwanted activity. Process entertainment recommendation data locally.
Video Surveillance	Perform real-time analysis at the edge while minimizing delay and network traffic
Retail	Just-walk-out-stores, smart shopping carts that combine data from camera and other sensors. Video processed at the edge for efficiency. Retail surveillance and edge processing for loss prevention
Banking and financial	Robots deployed in retail bank locations and interact with customers using AI at the edge

Source: XeonStack

Collaboration opportunities with LF AI & Data

- Webinar ideas:
  - Introduce LF Edge and LF AI & Data
  - Evolution of AI at the Edge
  - Real-world applications of AI at the Edge
- Or more

Thank you!

Questions?

EMET-LABS PROPOSAL TO HOST DEEPCAUSALITY IN LF AI & DATA

AUGUST 10, 2023

MARVIN HANSEN, DIRECTOR EMET-LABS.COM

MARVIN.HANSEN@EMET-LABS.COM

WHY CONTRIBUTE DEEPCAUSALALITY TO LINUX FOUNDATION

EMET-LABS
FinTech Research

- 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 project through LF ecosystem
  - Increase contributors by converting new & existing users
  - Opportunities to collaborate with other projects

Emet-Labs
FinTech Research

AGENDA

1) Problem
2) Challenge
3) How DeepCausality helps
4) Next steps

PROBLEM

WHEN DEEP LEARNING HITS ITS LIMITS

EMET-LABS
FinTech Research

Emet Labs specializes in modelling volatility in financial markets

- Financial markets are increasingly interrelated
- Internal context (say month low/high) collides with external context (say Twitter)
- Complex spatial-temporal patterns in interconnected contextual time-series data

Deep Learning applied to market volatility falls short:

- Context-blind
- World-model-blind
- Data-relation-blind

ROOT CAUSE

1) Context-blindness roots in the universal approximation theorem

2) World-Model-blindness means DeepLearning does not have a model of the world that generates the data it uses to learn

3) Data-relational-blindness roots in the Independent and identically distributed (IID) data assumption

DARPA ACKNOWLEDGES THE PROBLEM:

''ANSR hypothesizes that several of the limitations in ML today are a consequence of the inability to incorporate contextual and background knowledge, and treating each data set as an independent, uncorrelated input. In the real world, observations are often correlated and a product of an underlying causal mechanism, which can be modeled and understood."

Assured Neuro Symbolic Learning and Reasoning (ANSR) Program
[HTTPS://WWW.DARPA.MIL/PROGRAM/ASSURED-NEURO-SYMBOLIC-LEARNING-AND-REASONING](https://www.darpa.mil/program/assured-neuro-symbolic-learning-and-reasoning)

CHALLENGES

THREE BIG CHALLENGES

1) **Context & background knowledge**
2) **Relations between observations**
3) **Causal mechanism**

CONTEXT & BACKGROUND KNOWLEDGE

EMET-LABS
FinTech Research

- Context & background knowledge is independent of the model
- Context data:
  - Differ over time i.e. today's report is different from last week
  - Differ across time i.e. weekly reports contain different data than the annual report
  - Differ across space i.e. the Chicago office requires a different report than London
  - Differ across space *and* time: Zoom meeting across five different time zones…
- Context may need regular updates

EMET-LABS
FinTech Research

RELATIONS BETWEEN OBSERVATIONS

- Temporal relations:
  - Last month high/low serves as reference point for the current data
  - One dimensional over T (Time)
- Spatial relations:
  - Geometric patterns i.e. candle sticks relate to current data
  - Two or three dimensional, depending on the choice of representation
- Temporal-spatial relations
  - Time varying geometric patterns that inform current observation in a progressing time continuum
  - Four-dimensional: 1D Time + 3D Space

CAUSAL MECHANISM (1/2)

EMET-LABS FinTech Research

- Correlation: IF A occurs, B occurs as well
  - Correlations are omnipresent i.e. rain and umbrella
  - Can be reversed i.e. umbrella and rain
- Causation: IF A then B; AND if NOT A then NOT B
  - Causation is an underlying structure:
    - IF power on, then lightbulb on; AND If the power is not on, the light bulb is not on
    - Cannot be reversed under the time arrow assumption i.e. the lightbulb cannot be on without its cause
- All causation can be expresses as correlation i.e. power on and light on
- Correlation cannot be expressed as causation. It's impossible in absence of a causal relation
- Finding causal relations and structures is fundamentally an exceptionally hard problem

EMET-LABS
FinTech Research

CAUSAL MECHANISM (2/2)

- Single cause: These are extremely rare, but easy
  - Smoking cigarette at an open petrol tank -> Explosion
- Multi-cause, single effect: A bit more common
  - High unemployment, high interest rates, high inflation / costs of living (causes) lead to decreased disposable income (effect) resulting in lower real-estate sales (higher order effect).
- Multi-cause, multiple effects: Very common
  - High unemployment, low security standards, and poor governance may lead to vandalism and higher crime
- Multiple-stages, multiple causes, and multiple effects
  - Let's call this actual reality: Degree of complexity approximates infinity the more you dive into details
- And causality is contextual over space, time, and space-time

Emet-Labs
FinTech Research

SOLUTION SPACE

- Judea Pearl at UCLA: Foundational work & structural causal models
- Ilya Shpitser, Causal AI Lab at Johns Hopkins University: Causal and semi-parametric inference
- Miguel Hernan, Causal Lab at Harvard University: Causal models applied to Infectious diseases, mental health, and veterans' health
- Elias Bareinboim at Columbia University: Causal inference with decision-making/reinforcement learning
- Lucien Hardy at the Perimeter Institute for theoretical physics: Causality foundation for Quantum Gravity
- Causality and Machine Learning at Microsoft Research: ALICE, Causica, DiCE, DoWhy, EconML
- CausalML at Uber: Causal inference with machine learning algorithms
- Causal Inference Applications at Netflix: Causal recommendation models & subscriber retention

OBSERVATIONS

1) Advanced research available
2) Microsoft and Uber at the forefront of industry adoption
3) Very diverse projects: Foundational work, Causal RL, EconML, Healthcare…
   - Impossible to compare projects in a feature matrix
4) Projects focus more on algorithm and application
5) Python is the lingua franca; Rust not yet explored
6) Contextualized causal inference not yet explored
7) Causal (hyper) geometric structures not yet explored

EMET-LABS
FinTech Research

EMET-LABS
FinTech Research

KEY FACTORS TO SUCCESS

1) Causal structure:
   - Expresses arbitrary complex causal relations between causes, data, and context
2) Context:
   - Relates causal relations to contextual data
3) Explanation
   - Causal reasoning with explanation helps to understand the process

EMET-LABS
FinTech Research

DEEPCAUSALITY

HOW DOES DEEP CAUSALITY DIFFER FROM DEEP LEARNING?

EMET-LABS
FinTech Research

DEEP CAUSALITY

- Free of the IID assumption & explicit assumptions
- Deterministic causality
- Specializes
- Explainable
- Context aware
- Content with little data
- Small model of causal relations
- Good for reasoning, control systems & anomaly detection

DEEP LEARNING

- Subject to the IID assumption & implicit assumptions
- Non-deterministic correlation
- Generalizes
- Hard to explain; possible but needs work
- Context-free
- Requires big data
- Big model
- Good for generation, classification & language models

NOT ALL CAUSAL PROBLEMS ARE CREATED EQUAL

Simple vs. Complex	Static vs. Dynamic	Deterministic vs Probabilistic
Simple problems need simple solutions with simple structures.	Static problems need efficient representation guaranteed to be invariant.	Deterministic problems require verification and explanation.
Complex problems need complex structures and techniques to manage complexity.	Dynamic problems need an efficient update mechanism.	Probabilistic problems require a simple integration into the causal world.

FOUR LAYERS OF CAUSAL REASONING

Observations: Data we observe in the world

Assumptions: What we assume about the data we observe

Inference: What we can infer from the data with those assumptions

Causality: The causal relationships that structure inferences

ARCHITECTURE (EXCL CONTEXT)

CONTEXT ARCHITECTURE

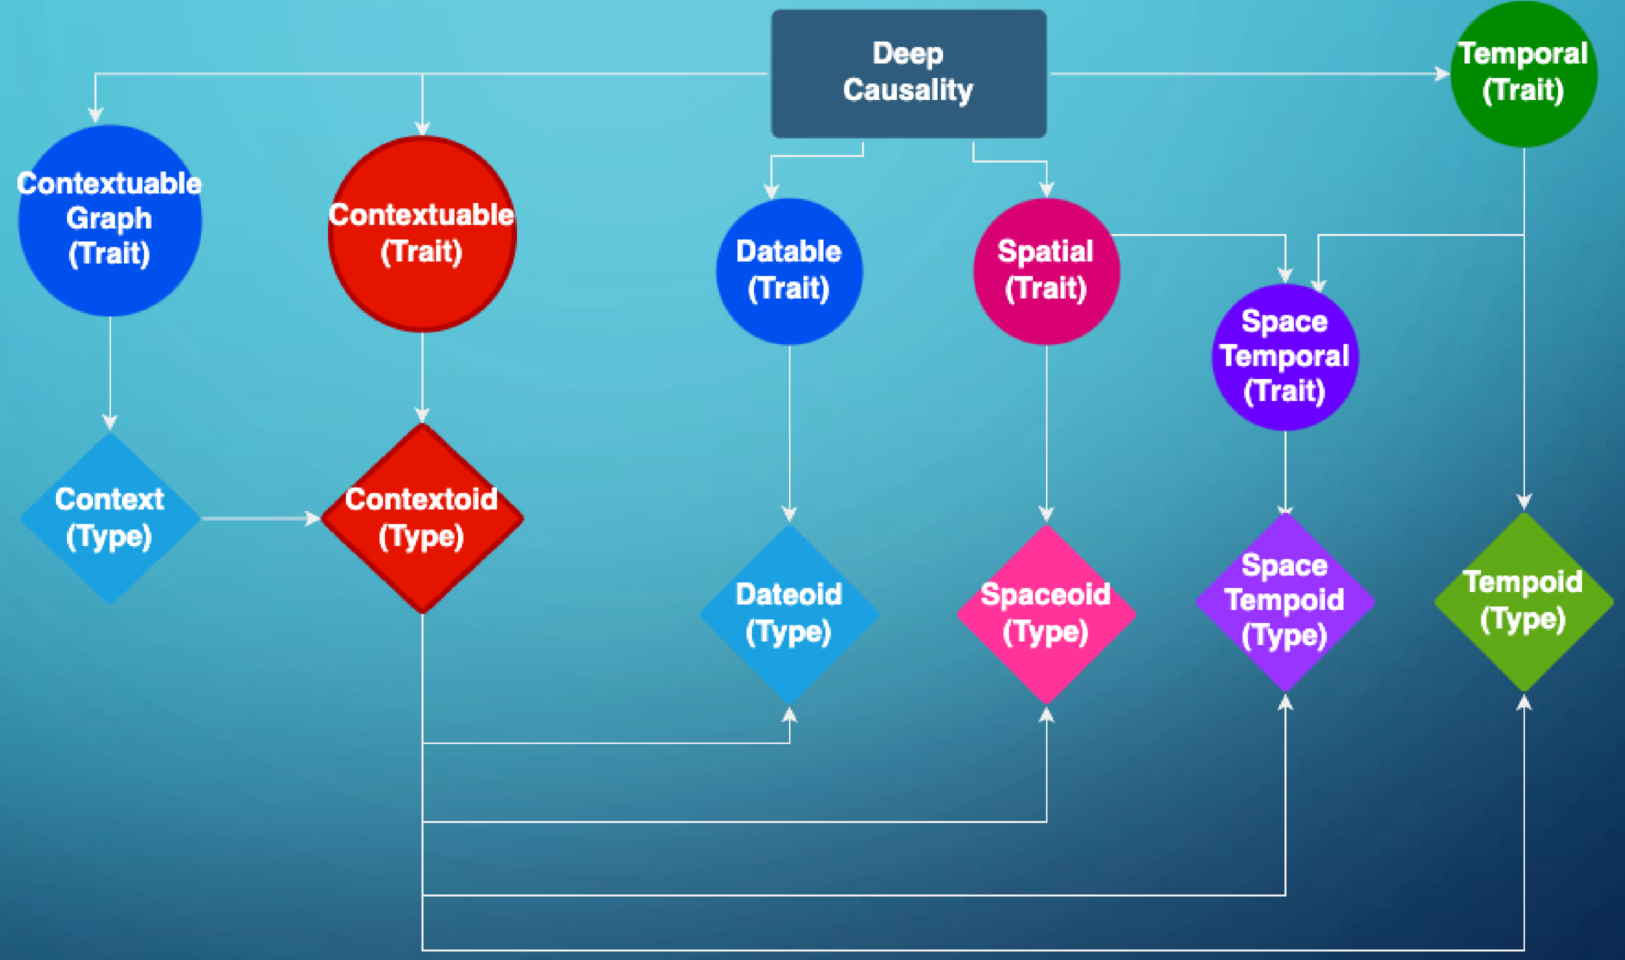

EMET-LABS
FinTech Research

EMET-LABS
FinTech Research

IMPLEMENTATION PRINCIPLES

- Protocol based causality
  - Inspired by differential programming pioneered by Google research in 2020, adopted by PassiveLogic
  - Defines traits with default implementation (Think protocols and extensions in Swift)
  - Extends standard collections in extension traits with functionality from a generic default implementation
- Recursive isomorphic data structures
  - Recursion in Rust only requires a layer of indirection i.e. a reference
  - Isomorphism requires that all data entities share the same structure: Hence Causaloid & Contextoid
- Disjoint algebraic data types a.k.a nested Enum
  - Solves the problem of missing inheritance & shared supertype
  - Nests various traits in an enum and then wraps the enum in a struct

FOUR PILLARS OF DEEP CAUSALITY

Emet-Labs
FinTech Research

1) **Hypergeometric deep causality**
2) **Recursive causal data structures**
3) **Hypergraph context**
4) **Causal state machines**

Emet-Labs
FinTech Research

HYPER-GEOMETRIC COMPUTATIONAL CAUSALITY

Emet-Labs
FinTech Research

HYPER-GEOMETRIC COMPUTATIONAL CAUSALITY

- **More than one way to represent causal structure:** Algebra vs Geometry
  - **Algebra: Complex arithmetic**
  - **Geometry: Complex structure**
- **DeepCausality choses a hypergraph geometric representation**
- **DeepCausality solves complexity with recursive causal data structures**

EMET-LABS
FinTech Research

RECURSIVE CAUSAL DATA STRUCTURES

RECURSIVE CAUSAL DATA STRUCTURES

Emet-Labs
FinTech Research

- Causaloid central structure
- Causaloid: Concept borrowed from Lucien Hardy's work on Quantum Gravity
- Causaloid defines a causal relation as a causal function
- Causal collections scale from simple to complex: Vector, Map, Hypergraph
- The causal graph is a hypergraph comprising of causaloids
  - A collection stores multiple causaloids, the collection is stored in a causaloid, and the causaloid may then be stored in a hypergraph, which itself is stored as a causaloid
  - Causal reasoning across the entire graph, selected nodes, selected path, or shortest path

RECURSIVE CAUSAL DATA STRUCTURES

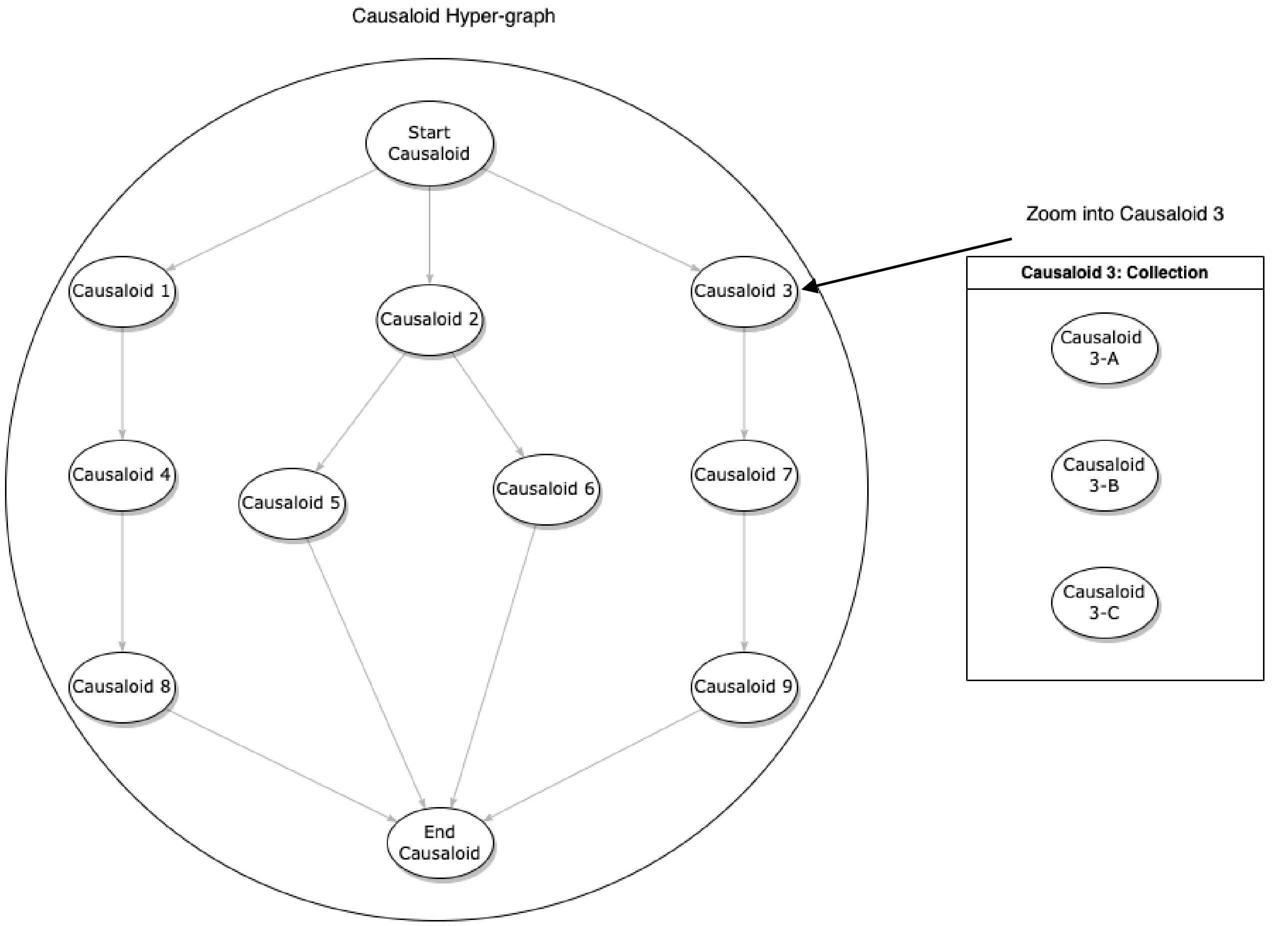

IMPLEMENTATION

- Causalgraph builds upon a sparse matrix hypergraph from the petgraph lib
- Trait defines behavior relative to the causable trait and its causaloid impl.
- Causal collection implemented as extension trait
- Causaloid uses an internal type Enum to determine if If it's a singleton, graph, or a collection
  - Depending on the internal type, different inference functions are called
  - However, the last element must be singleton to let the recursion terminate eventually
- Causaloid uses an internal bool flag to determine if it has a context
  - Depending on the flag, it either calls a causal function with context or one without context
  - Constructor sets bool flag, type Enum, and context.

HYPERGRAPH CONTEXT

HYPERGRAPH CONTEXT

- Time:
  - Conventionally seen as continuous discrete series: T0, T1, … Tn
- Space:
  - Conventionally seen as static location i.e. coordinates that are time invariant
- Space-Time:
  - Conventionally not much considered outside theoretical physics and space industry
- Data context:
  - Might be time varying, space-varying, or even space-time varying

HYPERGRAPH CONTEXT IN DEEPCAUSALITY

EMET-LABS
FinTech Research

- Time:
  - Considered as temporal hypergraph comprising of tempoids
  - Tempoid is one unit of time with a scale: Month: 12 (December)
- Space:
  - Considered as spatial hypergraph comprising of spaceoids
  - Spaceoid is one unit of space with coordinates
- Space-Time:
  - Considered as spacetime hypergraph comprising of spacetempoids
  - Spacetempoid is one unit of spacetime with coordinates at one unit of time
- Data:
  - Might be anything depending on how you define your data

HYPERGRAPH CONTEXT IN DEEPCAUSALITY

Emet-Labs
FinTech Research

- Contextoid is the central structure
- Contextoids can be data-like, time-like, space-like, or any combination
- Contextoids might be referenced by one or more causaloids
- Context hypergraph
  - Graph of contextoids that represent, data, time, space, or space-time
  - A contextoid may link to any number of other contextoids

HYPERGRAPH CONTEXT

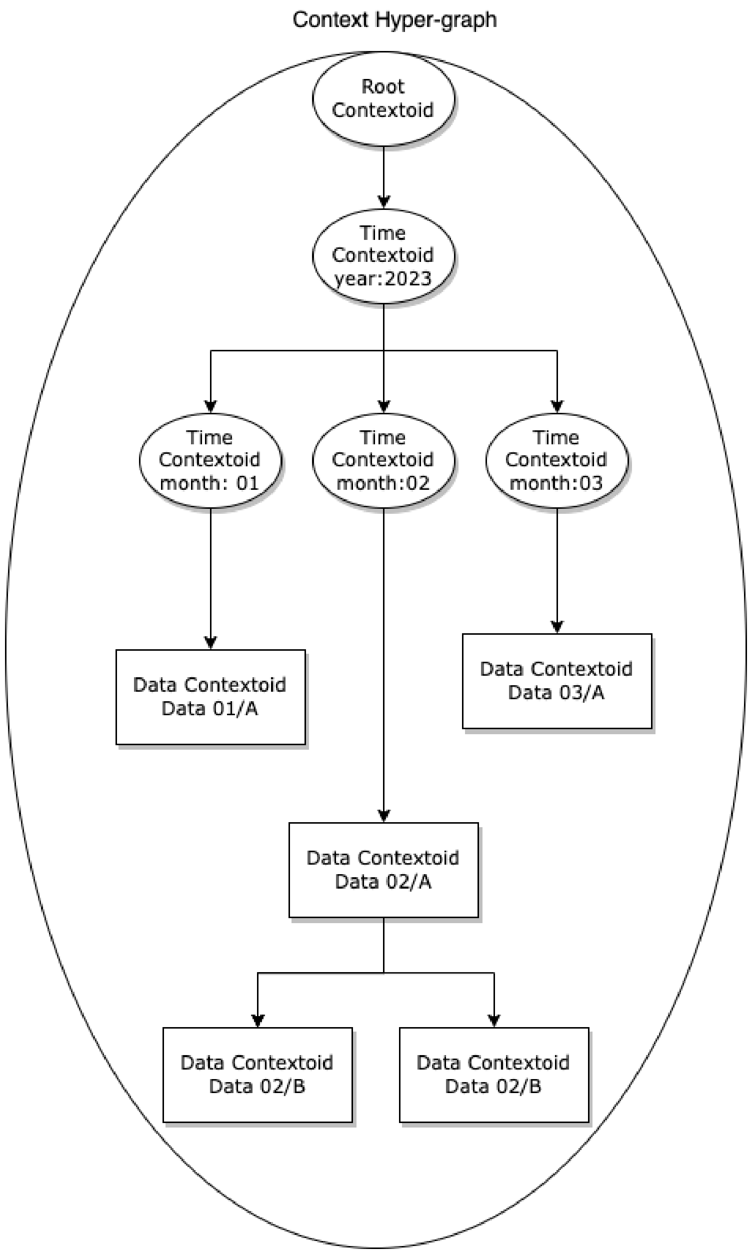

IMPLEMENTATION

Emet-Labs FinTech Research

- Like the causal structure, context builds upon a sparse matrix hypergraph
- Unlike the causal structure, **<u>context is not recursive</u>**
- Traits serve more as an interface for both, the contextoid and the context
- Contextoid wraps disjoint algebraic data structure a.k.a nested Enum
  - Each Enum embeds a type trait i.e. Temporal to allow custom types
- Custom types must extend the corresponding (super) trait i.e. temporal and implement both, the custom trait and the super trait

CAUSALOID + CONTEXT
=
CAUSAL MODEL

CONTEXTUALIZED CAUSAL MODELS

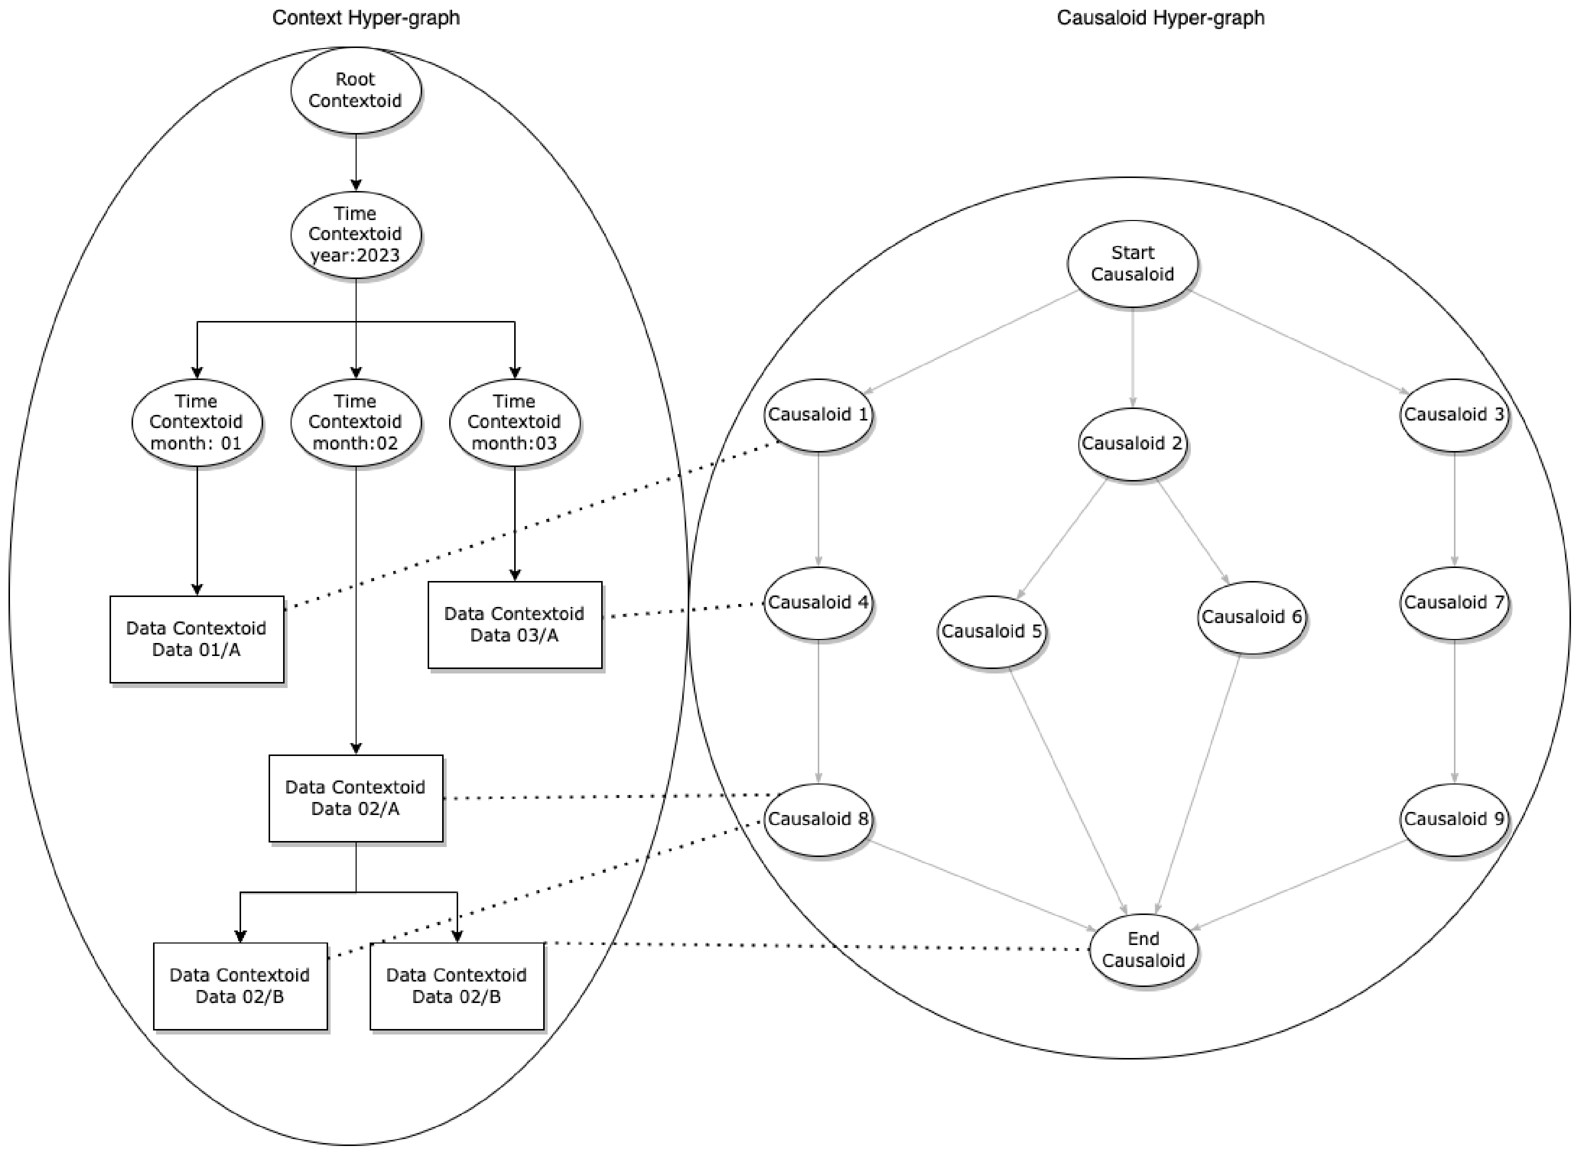

IMPLEMENTATION

- The context comes in two variants: BaseContext and Context.
  - BaseContext is an alias to the default (basic node type) that removes generic type annotations
  - Context requires to overwriting each node type with either a default or a custom type
  - In practice, only few custom types are needed, and default types fill the rest
- The causaloid must be constructed with context and with ContextCausalFunction
  - Functionality of custom types is only accessible inside the context function when the custom trait has been imported
  - The Causaloid dispatches automatically between functions with or without context.
- BTW, in DeepCausality it's all static dispatching so no dynamic overhead

CAUSAL STATE MACHINES

EMET-LABS
FinTech Research

CAUSAL STATE MACHINES (CSM)

- Problem: Finite State Machines (FSM) require a priori knowledge of all states
- Solution: Causal State machine (CSM) generalizes over FSM and allows dynamic state management
- CSM
  - Causal State = Causaloid with its causal function
  - Causal Action = The effect function triggered when cause was detected
  - Context-free = State machine can only work with defined states.
- Implemented via safe function references (pointers) in Rust
- CSM can be generated, evaluated, and executed on-demand
- Ideal for dynamic control systems

IMPLEMENTATION

- Uses the BaseCausaloid since neither customization nor context is needed
- Adds two more types:
  - CausalState: defines a target state
  - CausalAction: Defines what to do when the target state has been reached
- CasualStateMachine
  - Evaluates if the target state has been reached and then triggers the action
- Can dynamically add, remove, or evaluate states at runtime

CAUSALOID VS. CAUSAL STATE MACHINES

Causaloid

- Contextualized
- Dynamic w.r.t. model & context
- Flexible: Can be probabilistic, deterministic, or both.
- Action defined outside the model

Causal State Machine

- Context-free
- Static
- Strictly deterministic
- Defined causal action

THE VALUE OF DEEPCAUSALITY

EMET-LABS
FinTech Research

VALUE

1) Comprehensive data enrichment
2) Efficient causal representation
3) Gives reason across the chain of causality
4) Inspires new directions i.e. contextualized deep learning

DEEPCAUSALITY FEATURE

1) Context
2) Causaloid data structure
3) Explainability
4) Exploration of new ideas, data structures, and implementation techniques

NEXT STEP

NEXT STEP

Roadmap:

- Support multiple contexts
  - Complex models may require more than one context
- Explore causal learning
  - Right now, causal models are crafted by hand
  - Groundwork has been laid via assumable & inferable protocol
  - However, existing work of causal RL learning does not transfer well
  - Deep Neuroevolution might offer a novel path towards causal learning
- Expand docs & code examples

EMET-LABS
FinTech Research

INFORMATION FOR PROPOSAL

- License: MIT
- GH repo: https://github.com/deepcausality-rs
- Proposal: https://github.com/lfai/proposing-projects/blob/master/proposals/deepcausality.adoc
- Possible Collaboration in LF AI & Data:
  - 1chipML: Machine learning for microcontrollers
  - ForestFlow: Machine learning model server
  - Xtreme1: ML platform for multisensory training data
  - BentoML: The Unified AI Application Framework

EMET-LABS
FinTech Research

Thank You

Marvin Hansen

marvin.hansen@emet-labs.com

Approval on DeepCausality

Proposed Resolution:

- DeepCausality as a sandbox project of the LF AI & Data Foundation is hereby approved.

Upcoming TAC Meetings

LF AI & DATA

Upcoming TAC Meetings

- August 24 – SapientML new Sandbox Project from Fujitsu; Trusted AI Committee update
- September 7 – Update from the MLSecOps committee, <open slot>

Please note we are always open to special topics as well.

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

Upcoming Events of Interest

- 2023 AICON Middle East Summit - October 8th to 9^{th} in Riyadh https://lfaidata.foundation/blog/2023/07/18/2023-aicon-middle-east-summit-call-for-topics-from-around-the-world/
- Open Source Summit Europe in Bilbao, Spain, September 19-21 – LF AI&Data will have a booth https://events.linuxfoundation.org/open-source-summit-europe/

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

Legal Notice

- The Linux Foundation, The Linux Foundation logos, and other marks that may be used herein are owned by The Linux Foundation or its affiliated entities, and are subject to The Linux Foundation's Trademark Usage Policy at https://www.linuxfoundation.org/trademark-usage, as may be modified from time to time.
- Linux is a registered trademark of Linus Torvalds. Please see the Linux Mark Institute's trademark usage page at https://lmi.linuxfoundation.org for details regarding use of this trademark.
- Some marks that may be used herein are owned by projects operating as separately incorporated entities managed by The Linux Foundation, and have their own trademarks, policies and usage guidelines.
- TWITTER, TWEET, RETWEET and the Twitter logo are trademarks of Twitter, Inc. or its affiliates.
- Facebook and the "f" logo are trademarks of Facebook or its affiliates.
- LinkedIn, the LinkedIn logo, the IN logo and InMail are registered trademarks or trademarks of LinkedIn Corporation and its affiliates in the United States and/or other countries.
- YouTube and the YouTube icon are trademarks of YouTube or its affiliates.
- All other trademarks are the property of their respective owners. Use of such marks herein does not represent affiliation with or authorization, sponsorship or approval by such owners unless otherwise expressly specified.
- The Linux Foundation is subject to other policies, including without limitation its Privacy Policy at https://www.linuxfoundation.org/privacy and its Antitrust Policy at https://www.linuxfoundation.org/antitrust-policy. each as may be modified from time to time. More information about The Linux Foundation's policies is available at https://www.linuxfoundation.org.
- Please email legal@linuxfoundation.org with any questions about The Linux Foundation's policies or the notices set forth on this slide.

LF AI & DATA