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

March 24, 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

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 (2 mins)
- Approval of Minutes from previous meeting (2 mins)
- Datashim annual review (20 minutes)
- Interpretable Deep Learning:
  Interpretation, Interpretability, Trustworthiness, and Beyond (25 minutes)
- Outreach committee update (5 minutes)
- LF AI General Updates (2 min)
- Open Discussion (2 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

TAC Voting Members

* = still need backup specified on wiki

Member Representatives

Member Company or Graduated Project	Membership Level or Project Level	Voting Eligibility	Country	TAC Representative	Designated TAC Representative Alternates
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)
IBM	Premier	Voting Member	USA	Susan Malaika	Saishruthi Swaminathan
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
Tech Mahindra	Premier	Voting Member	India	Amit Kumar	Prasanna Kulkarni
Tencent	Premier	Voting Member	China	Bruce Tao	Huaming Rao
ZTE	Premier	Voting Member	China	Wei Meng	Liya Yuan
Acumos Project	Graduated Technical Project	Voting Member	USA	Amit Kumar	Prasanna Kulkarni
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 Chair

LF AI & DATA

Minutes approval

LF AI & DATA

Approval of March 10, 2022 Minutes

Draft minutes from the March 10th TAC call were previously distributed to the TAC members via the mailing list

Proposed Resolution:

- That the minutes of the March 10 meeting of the Technical Advisory Council of the LF AI & Data Foundation are hereby approved.

Annual Review for Datashim

03/24/2022

Srikumar Venugopal

LF AI & DATA

Datashim

Brief Description:

Datashim is enabling and accelerating data access for Kubernetes/Openshift workloads in a transparent and declarative way. Open-sourced since September of 2019 and is growing to support use-cases related to data access in AI projects. It brings benefits across different entities:

- Data scientists/engineers: Focus on workload/experiments development and not on configuring/tuning data access
- Storage Providers: Increase adoption since the framework is extensible without hindering the User Experience
- Data-oriented Frameworks: Can build capabilities (caching, scheduling) on top of Datashim using a declarative way to access/manage data sources

Contributed by:

IBM in January 2021 as an Incubation Project

Key Links:

Github: https://github.com/datashim-io/datashim

Website: https://datashim.io/

Artwork:

https://github.com/lfai/artwork/tree/master/projects/datashim

Mailing lists:

- datashim-announce
- datashim-technical-discuss
- datashim-tsc

Contributions

211.40K
Lines Of Code Changed

122
Commits

12
Contributors

1
No Of Sub Projects

1
Repositories

Top 10 Contributors By Commits View All

NAME	LINES OF CODE	COMMITS	%
Srikumar Venugo...	15.03K	43	7.11%
Yiannis Gkoufas	5.24K	24	2.48%
Panagiotis Kouts...	1.69K	22	0.80%
-MISSING-NAME	129.62K	10	61.31%
pkoutsov	264	4	0.12%
Tommy Li	42	3	0.02%
Antony Chazapis	1.23K	3	0.58%
Evangelos Maliar...	2.43K	3	1.15%
Gareth Western	1	1	0.00%

Top 10 Organizations By Commits View All

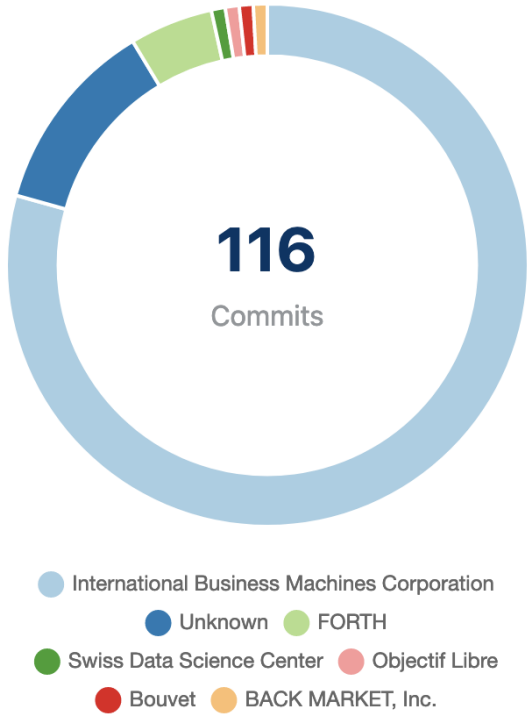

LF AI & DATA

Organizations contributing

- IBM
- Swiss Data Science Center
- European Bioinformatics Institute
- FORTH, Greece

Downstream

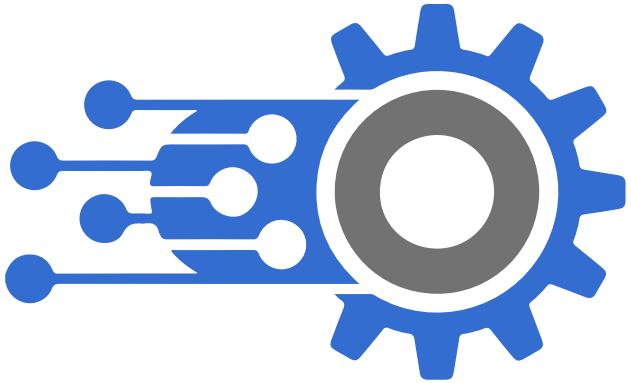

renku

Fybrik

karvdash

Key Achievements in the past year

- Technical
  - Plugin for transparent caching of object storage buckets
  - PoC of integration with Kubernetes Scheduling Framework
  - Support for NFS, HostPath, H3 storage backends
- Open Source
  - 207 stars, 44 forks, 11 contributors
  - Adoption by EBI, SDSC, FORTH, CODAIT
  - Upstream contributions to COSI, IBM Object Storage Plugin, Red Hat Noobaa
  - 4 research publications, 5 community call presentations

Roadmap

- Bi-Monthly releases
- Ephemeral volume support for S3
- Integration with COSI (when finalised)
- Auto-discovery of CSI implementation capabilities
- Support for more frameworks (Tekton, Flyte)
- Focus on observability (Design phase)

Areas the project could use help on

- Publicity
  - Venue suggestions - advance notifications
  - Joint webinars
  - Collaborations with relevant communities (CNCF, LF Edge)

Feedback on working with LF AI & Data

- Immense boost to project profile
- Outreach from linked communities (Flyte)

TAC Open Discussion

LF AI & DATA

Interpretable Deep Learning:

Interpretation, Interpretability, Trustworthiness, and Beyond

Xuhong Li
Big Data Laboratory, Baidu Research

Overview

1. Introduction: AI and XAI
2. Taxonomy and Trustworthiness
3. Open Source Toolkit: InterpretDL
4. Advances and Applications

Deep Learning and Artificial Intelligence

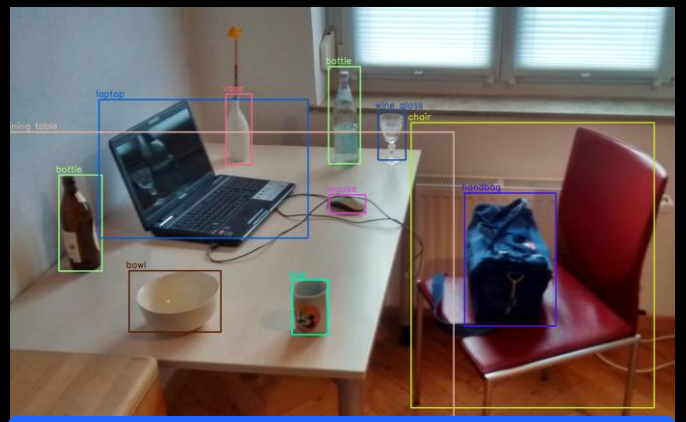

Visual Recognition

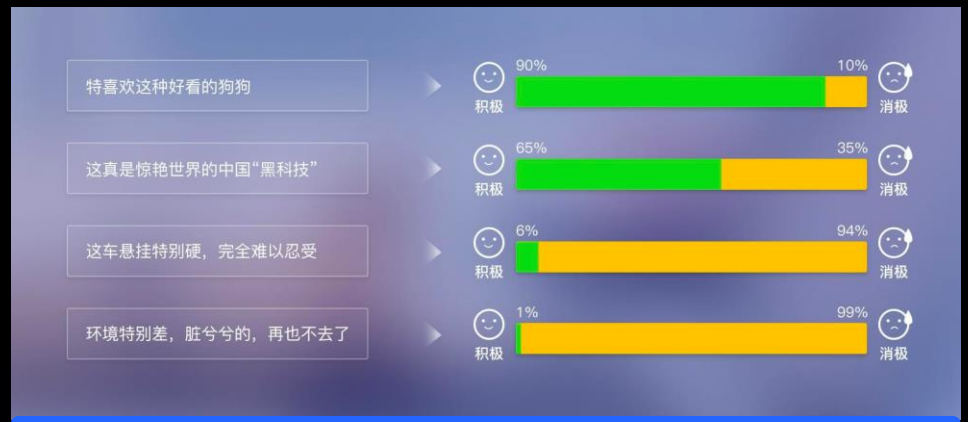

Natural Language Processing

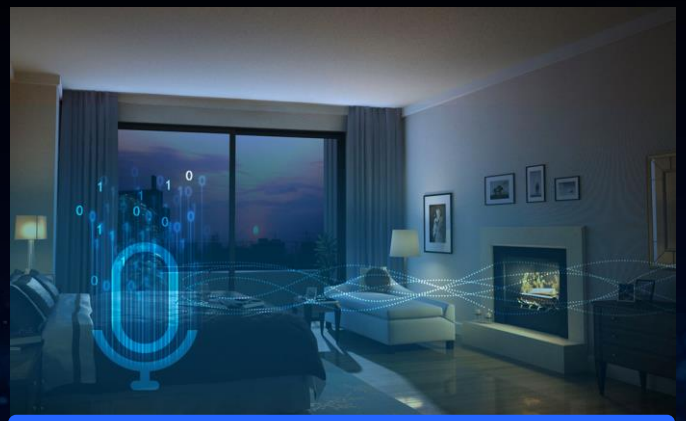

Speech Recognition

Reinforcement Learning

Deep Learning and Artificial Intelligence

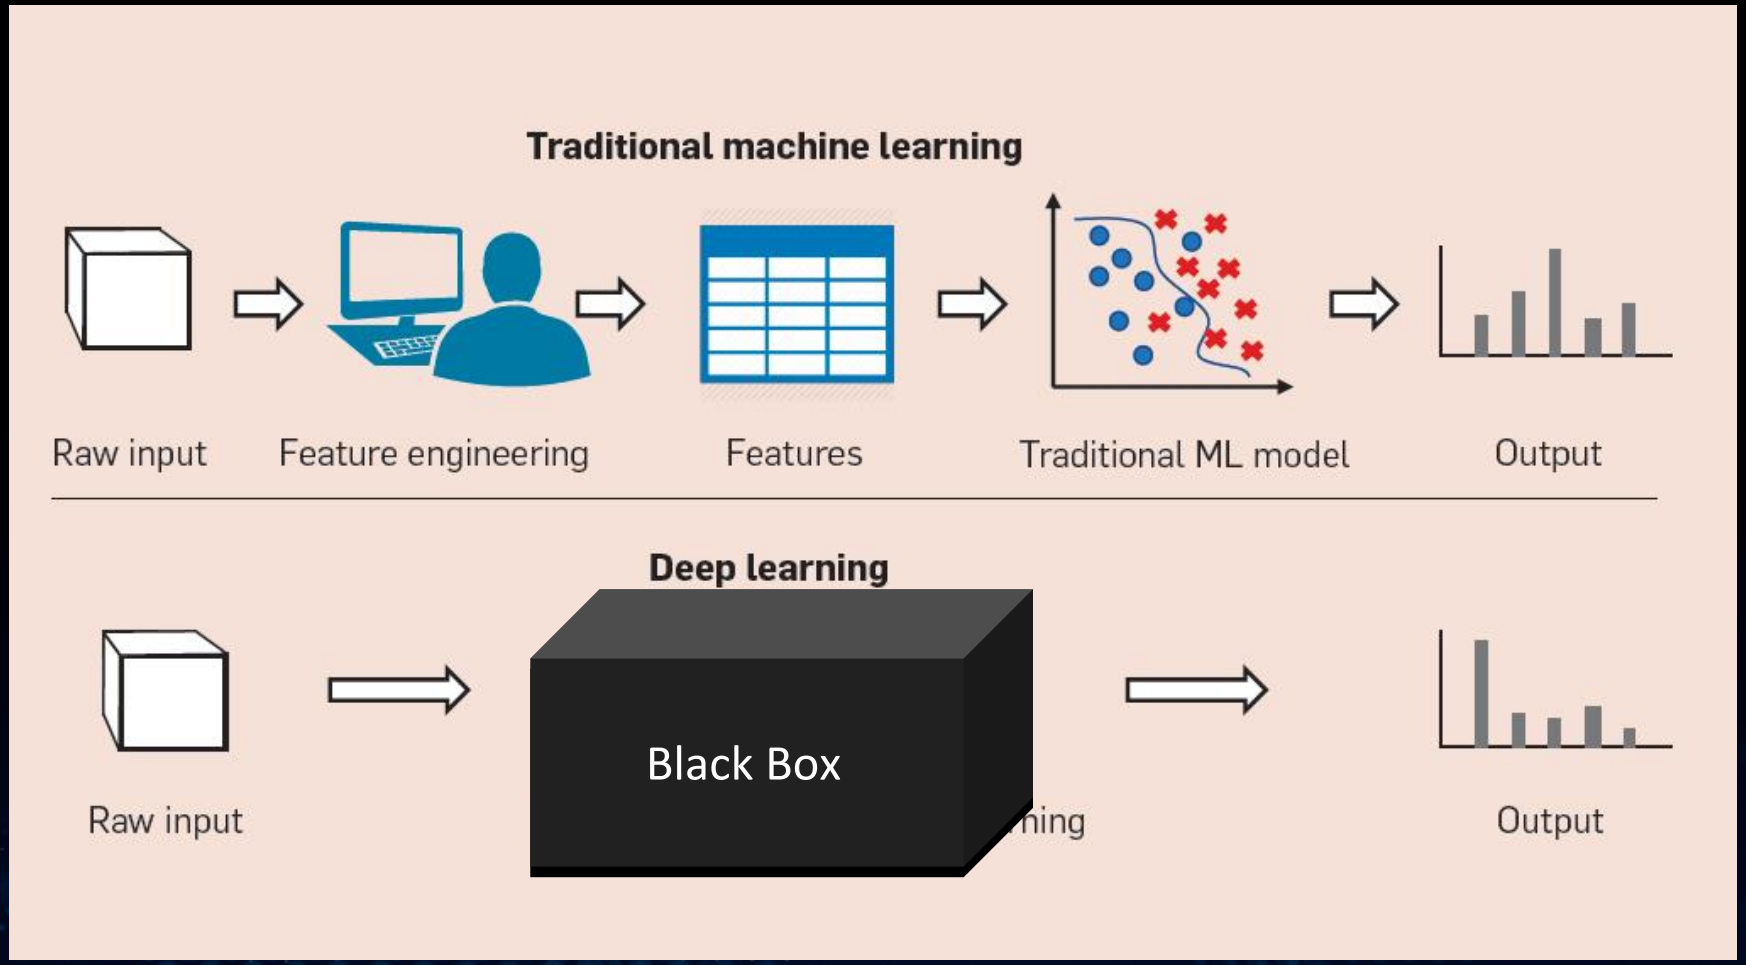

Black Box and Explainable AI

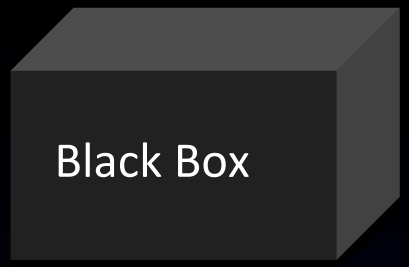

Why Interpretability?

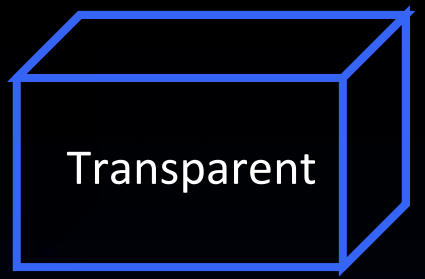

Black Box and Explainable AI

Why Interpretability?

1. Model Understanding
2. Model Debug and Improvement
3. Safety and Social Issues
4. Human Curiosity and Scientific Exploration

Taxonomy of Interpretation Algorithms

Taxonomy

Explainable Model

Self-Explainable Models, designed with Bayes, Causal Inference, or Logicism *etc*.

Interpretation on Training

Post-hoc Interpretation

Taxonomy

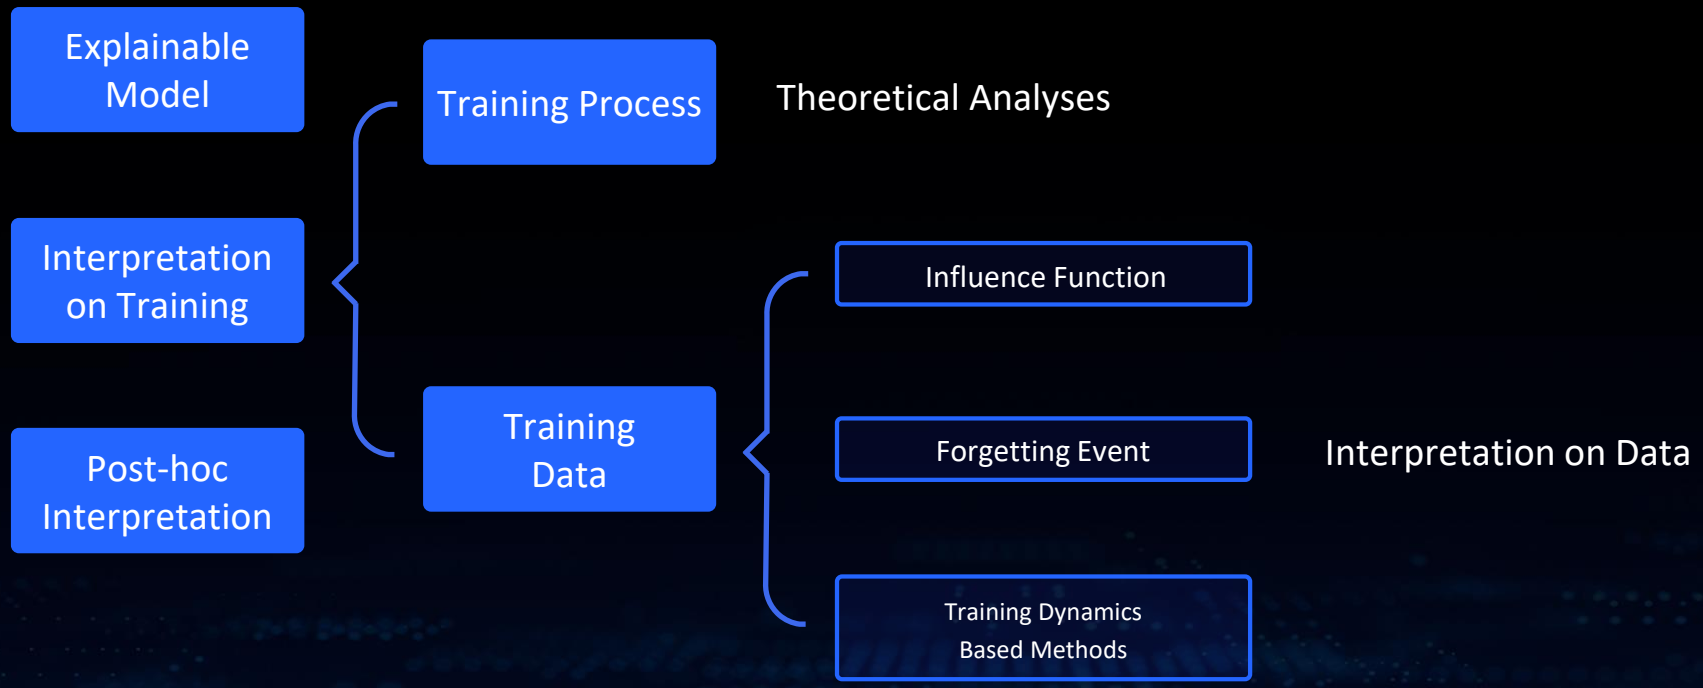
Explainable Model
Interpretation on Training
Post-hoc Interpretation
Training Process
Training Data
Theoretical Analyses
Influence Function
Forgetting Event
Training Dynamics Based Methods
Interpretation on Data

Taxonomy

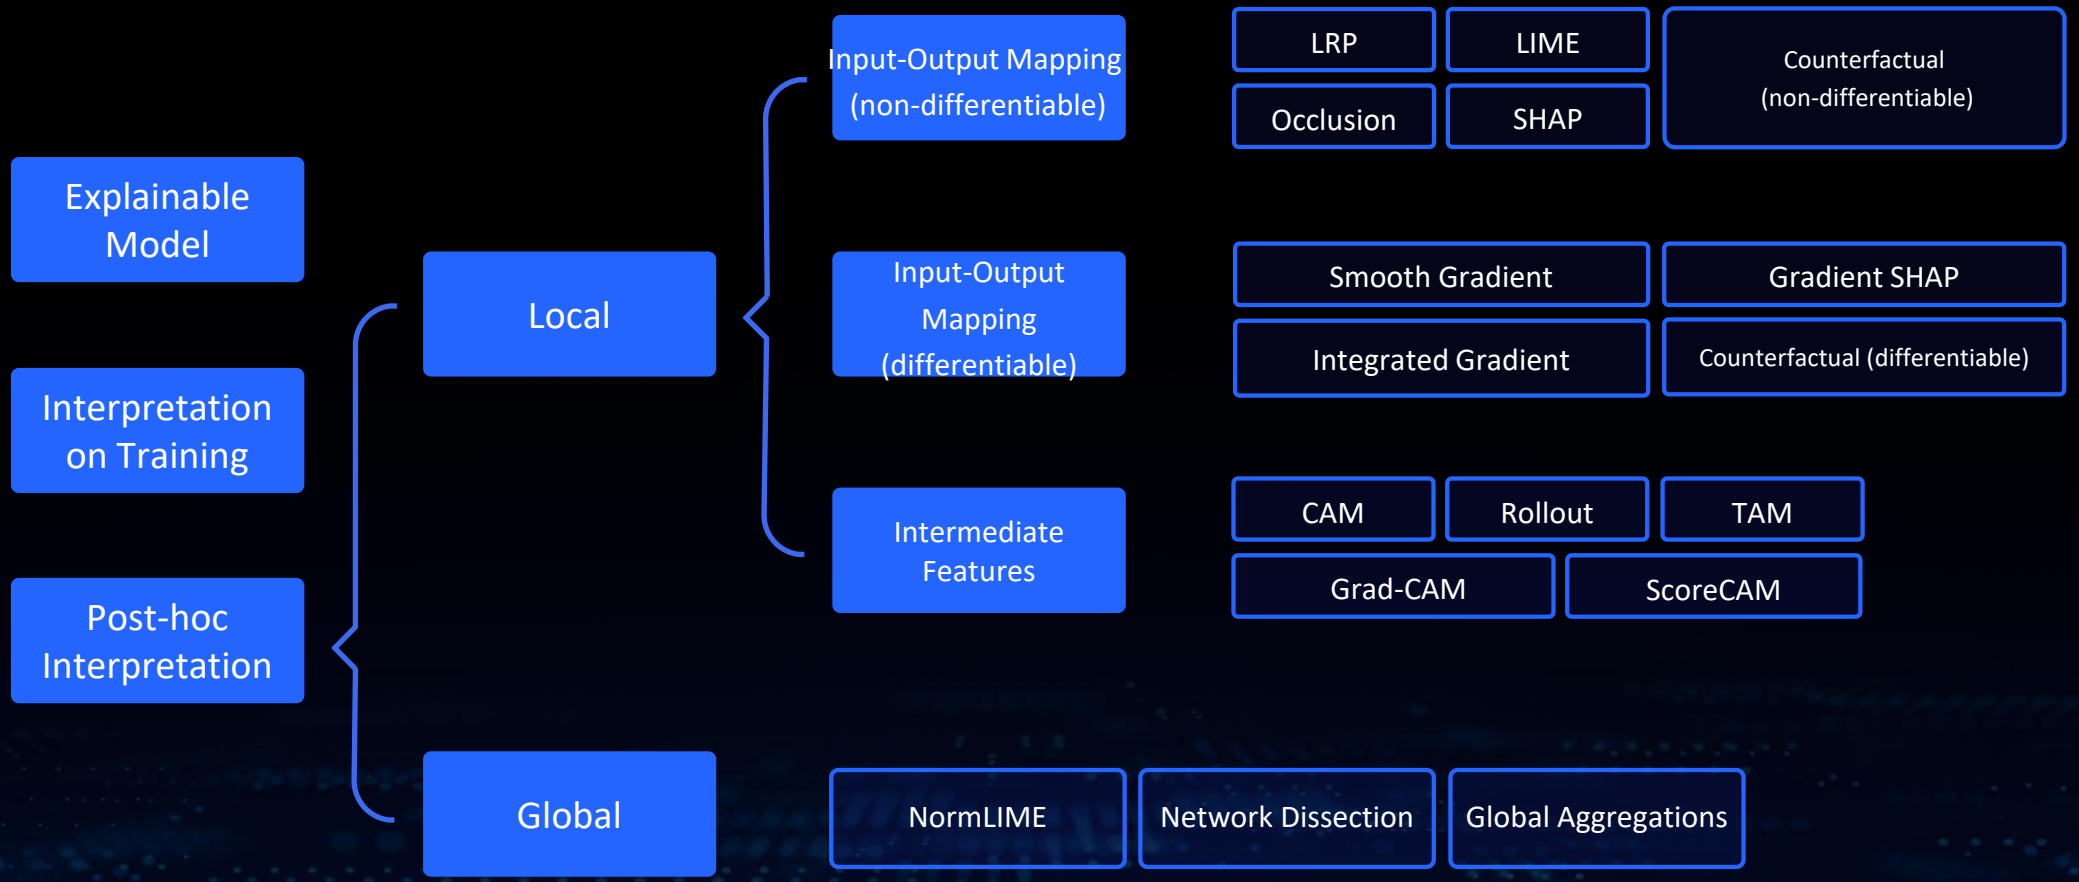
Explainable Model
Interpretation on Training
Post-hoc Interpretation
Local
Global
Input-Output Mapping (non-differentiable)
Input-Output Mapping (differentiable)
Intermediate Features
LRP
LIME
Occlusion
SHAP
Counterfactual (non-differentiable)
Smooth Gradient
Gradient SHAP
Integrated Gradient
Counterfactual (differentiable)
CAM
Rollout
TAM
Grad-CAM
ScoreCAM
NormLIME
Network Dissection
Global Aggregations

Typical Interpretations

Typical Interpretations

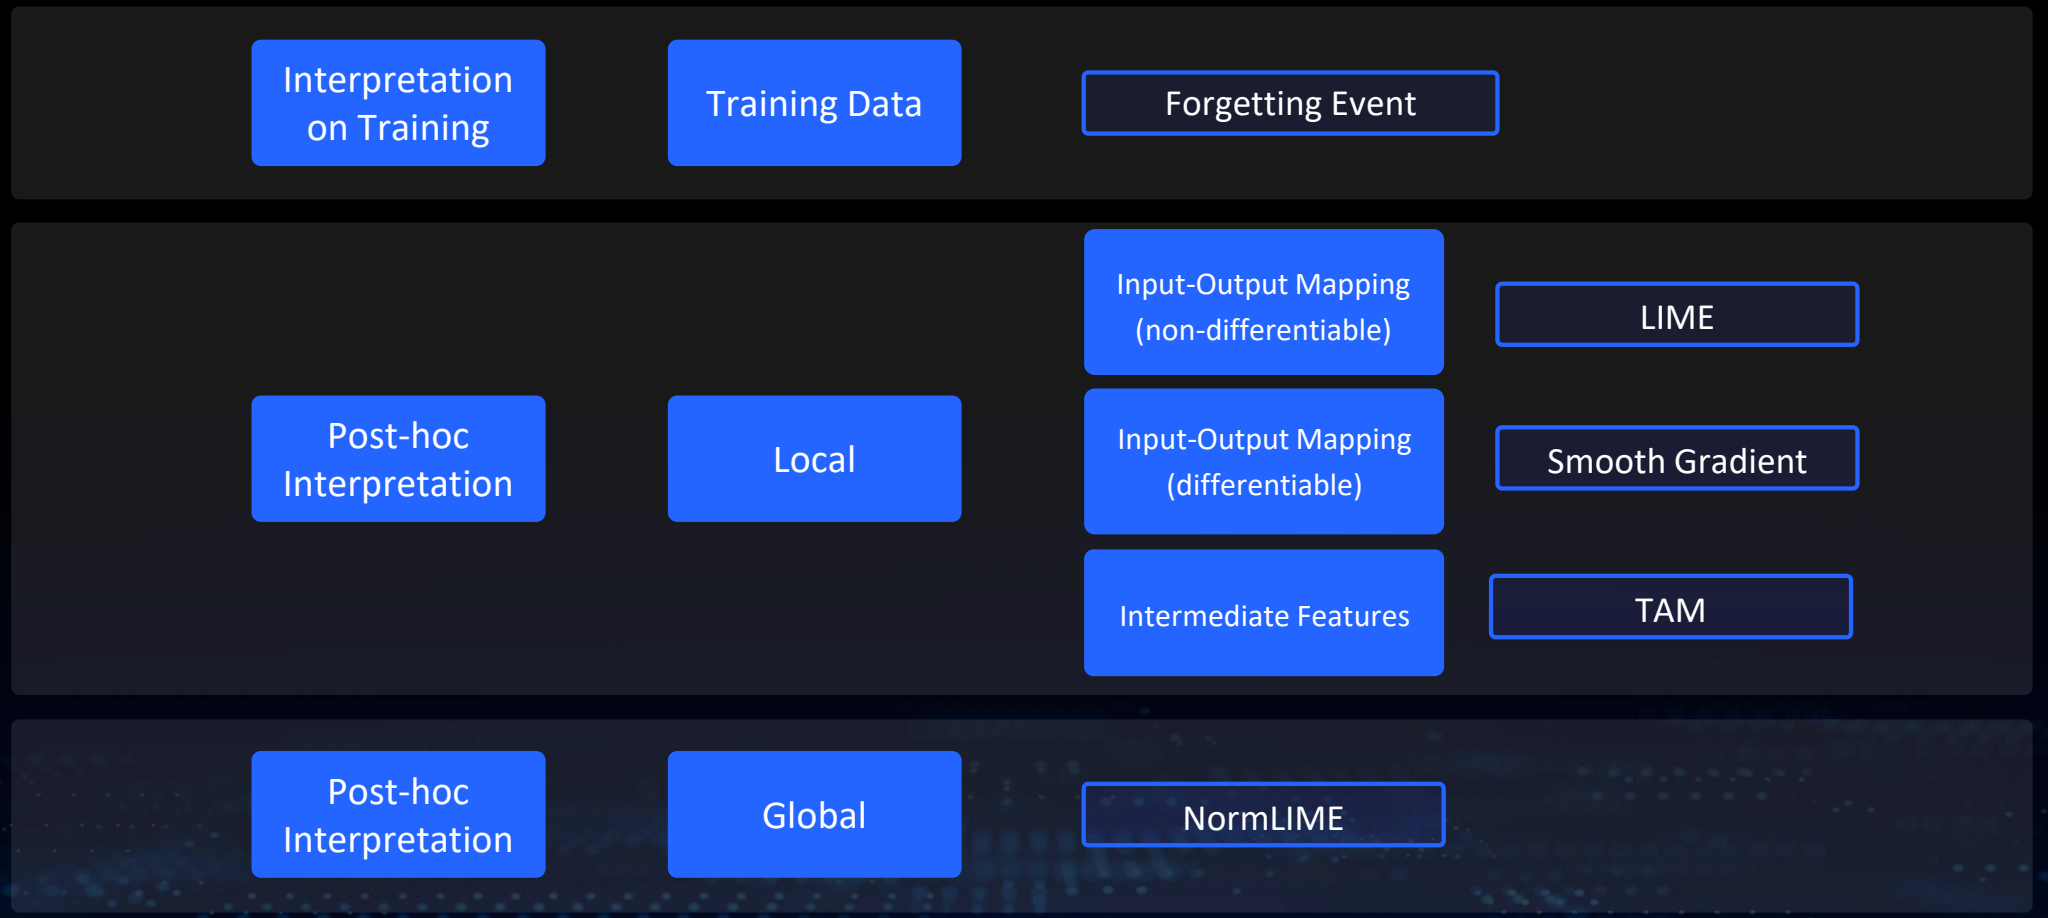

Forgetting Event

def

A forgetting event is defined as an event when a sample is correctly classified at *t* epoch but misclassified at *t+1* epoch during the training process.

way

Record the frequency of each sample happening forgetting events, and do the analyses.

ex

(left) Removing samples of no forgetting events does not affect the model performance;

(right) Mislabeled samples have higher frequencies of forgetting events.

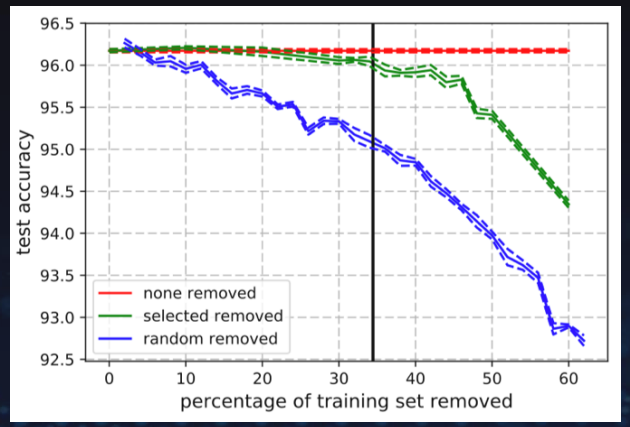

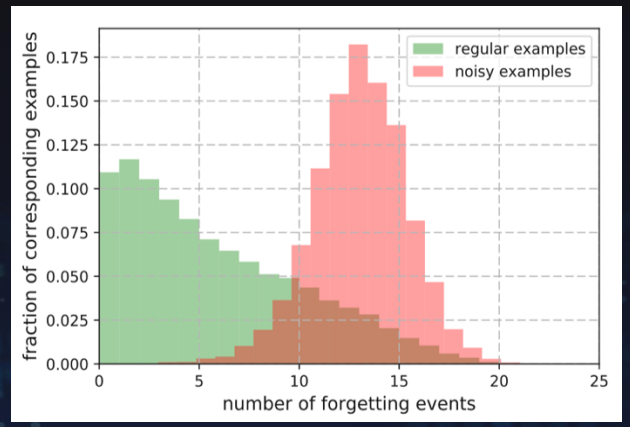

Toneva et al. “An Empirical Study of Example Forgetting During Deep Neural Network Learning.” in ICLR’19.

LIME: Local Interpretable Model-agnostic Explanation

idea

To fit a deep model at a local point using a linear model

way

(1) Generate samples at a local point and compute the responses of the model to these samples;

(2) Fit these samples and their responses using a linear model.

ex

LIME is a model-agnostic explanation algorithm and applicable on various tasks.

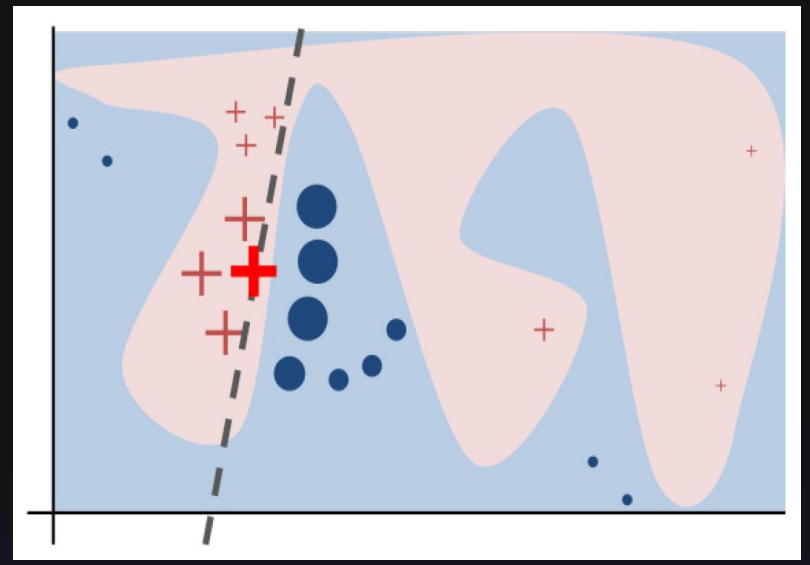

True Label	Predicted Label (Prob)	Target Label	Word Importance
1	1 (1.00)	1	it ' s a charming and often affecting journey .
1	1 (1.00)	1	the movie achieves as great an impact by keeping these thoughts hidden as . . . (quills) did by showing them .
0	0 (0.93)	0	this one is definitely one to skip , even for horror movie fanatics .
0	0 (0.97)	0	in its best moments , resembles a bad high school production of grease , without benefit of song .

Ribeiro et al. “Explaining the Predictions of Any Classifier.” SIGKDD’16.

NormLIME

intro

LIME is a local explanation but sometimes global one is needed.

way

Normalization and Average

$$\mathcal{S}(c_i) := \frac{1}{|E(c_i)|} \sum_{\boldsymbol{w}_{\boldsymbol{x}_j} \in E(c_i)} \gamma(\boldsymbol{w}_{\boldsymbol{x}_j}, i) \left| w_{x_j,i} \right|$$

$$\gamma(\boldsymbol{w}_{x_j}, i) := \frac{\left| w_{x_j,i} \right|}{\sum_k \left| w_{x_j,k} \right|} = \frac{\left| w_{x_j,i} \right|}{\| \boldsymbol{w}_{x_j} \|_1}$$

Important Words in Sentiment Analyses (Chinese)

Top words for "positive":

	word	weight	freq
0	适中	0.225913	3
1	超高	0.195469	3
2	爽	0.191485	7
3	很满意	0.187091	14
4	极好	0.177923	2
5	再来	0.177336	2
6	开心	0.175077	2
7	轻便	0.173926	3
8	五星	0.170384	3
9	真好	0.170218	4
10	非常棒	0.160273	2
11	挺好	0.158340	10
12	强悍	0.156671	2
13	放心	0.154761	3
14	合适	0.153143	4

Top words for "negative":

	word	weight	freq
0	不稳	0.319744	2
1	糟糕	0.262375	2
2	较差	0.241238	6
3	崩溃	0.234540	2
4	极差	0.234481	5
5	垃圾	0.231528	6
6	缝隙	0.229802	4
7	不佳	0.226448	2
8	太烂	0.222260	3
9	不符合	0.222123	4
10	失望	0.221986	16
11	很一般	0.216074	18
12	很失望	0.207618	4
13	上当受骗	0.203699	2
14	最差	0.190424	9

Ahern et al. “Normlime: A new feature importance metric for explaining deep neural networks”. arXiv:1909.04200, 2019.

Smooth Gradient

intro

Vanilla gradients *w.r.t.* inputs should be intuitive, but deep models consist of non-linear components, leading to gradient saturation or vanishing.

way

Add noises to remove the noise.

$$SG = \frac{1}{n}\sum^{n}\frac{\partial F(x + \mathcal{N}(0, \sigma^2))}{\partial x}$$

ex

Smooth Gradient is applicable to all differentiable models.

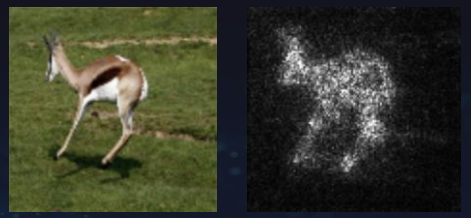

Smilkov et al. “Smoothgrad: Removing noise by adding noise”. arXiv:1706.03825, 2017.

TAM: Transition Attention Maps

idea

TAM connects the attention results to the transition matrices of Markov Chain, to explain the information flow across the layers of the Transformer.

ex

Explanations of different decisions

zebra elephant

Explanation of the wrong decision

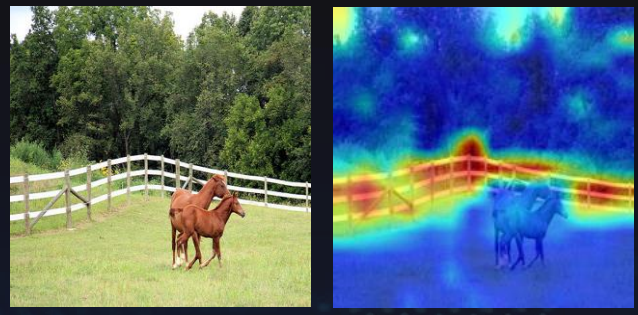

The image is labeled as *horse* but recognized as *fence*.

Yuan et al. “Explaining Information Flow Inside VisionTransformers Using Markov Chain”. XAI4Debugging@NeurIPS2021.

Trustworthiness Evaluations

Feature Importance	The model (supposedly) uses the most important features to make decisions.
Local Fidelity	Compute the fidelity between the proxy model and the deep model (only available to LIME and variants).
User Study	By experts.

Open Source Toolkit: InterpretDL

Open Source Toolkit InterpretDL

- Mainstream SOTA algorithms
- Plug-and-Play Designs
- Simple Usage: 5 lines
- Universal API: interpret
- Tutorials
- Various Visualizations
- Trustworthiness Evaluation

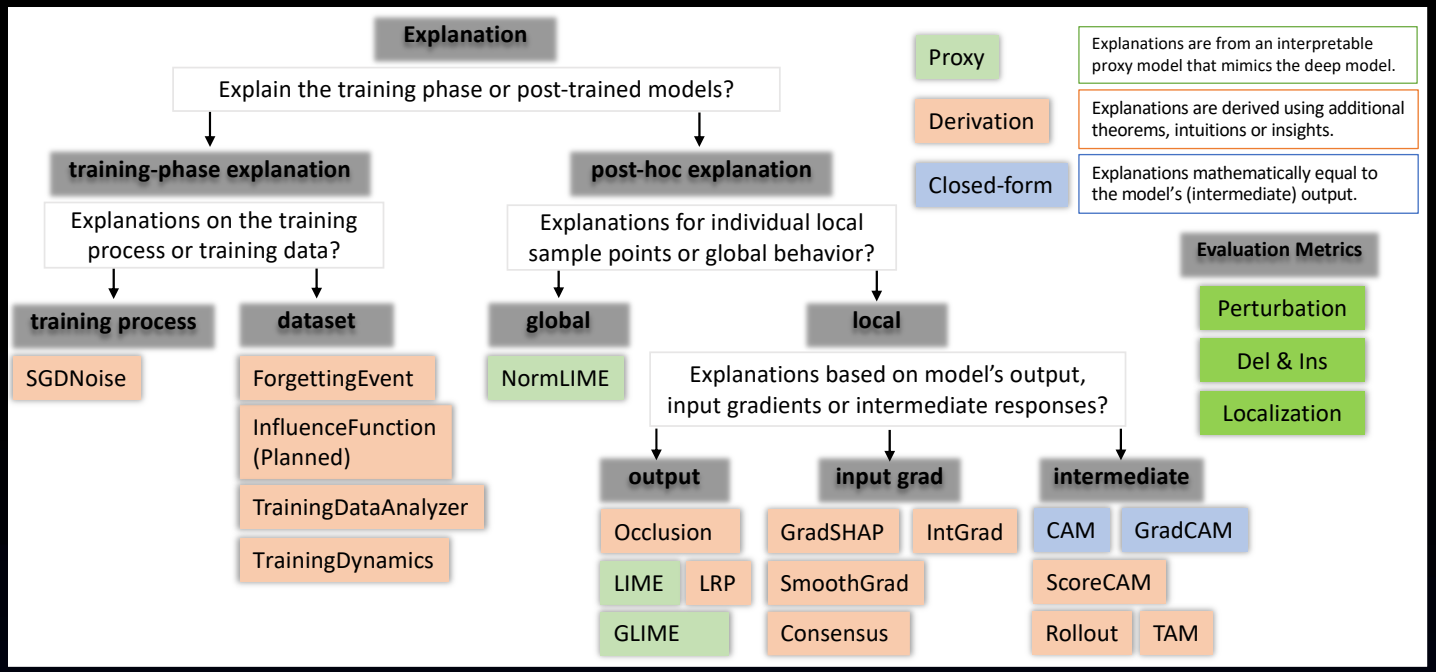

https://github.com/PaddlePaddle/InterpretDL

Open Source Toolkit InterpretDL

- Mainstream SOTA algorithms
- Plug-and-Play Designs
- Simple Usage: 5 lines
- Universal API: interpret
- Tutorials
- Various Visualizations
- Trustworthiness Evaluation

Feature-level Interpretation Algorithms

- Target at Input Features
  - [x] SmoothGrad
  - [x] IntegratedGradients
  - [x] Occlusion
  - [x] GradientSHAP
  - [x] LIME
  - [x] GLIME (LIMEPrior)
  - [x] NormLIME/FastNormLIME
  - [x] LRP
  - [ ] More ...
- Target at Intermediate Features
  - [x] CAM
  - [x] GradCAM
  - [x] ScoreCAM
  - [x] Rollout
  - [x] TAM
  - [ ] More ...

Dataset-level Interpretation Algorithms

- [x] Forgetting Event
- [x] SGDNoise
- [x] TrainIng Data analYzer (TIDY)
- [ ] Influence Function
- [ ] More ...

```
pip install interpretdl
```

```
import interpretdl as it
from paddle.vision.models import resnet50
paddle_model = resnet50(pretrained=True)
sg = it.SmoothGradInterpreter(paddle_model, use_cuda=True)
gradients = sg.interpret("test.jpg", visual=True, save_path=None)
```

The following table gives visualizations of several interpretation algorithms applied to the original image to tell us why the model predicts "bull_mastiff."

Original Image	IntGrad (demo)	SG (demo)	LIME (demo)	Grad-CAM (demo)
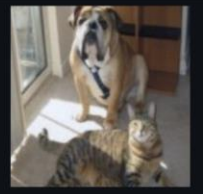				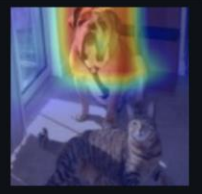

For sentiment classfication task, the reason why a model gives positive/negative predictions can be visualized as follows. A quick demo can be found here. Samples in Chinese are also available here.

True Label	Predicted Label (Prob)	Target Label	Word Importance
1	1 (1.00)	1	it ' s a charming and often affecting journey .
1	1 (1.00)	1	the movie achieves as great an impact by keeping these thoughts hidden as . . . (quills) did by showing them .
0	0 (0.93)	0	this one is definitely one to skip , even for horror movie fanatics .
0	0 (0.97)	0	in its best moments , resembles a bad high school production of grease , without benefit of song .

https://github.com/PaddlePaddle/InterpretDL

InterpretDL Application1 - Credit Scoring

Public Dataset - *GiveMeSomeCredits*

Variable Name	Description	Data Type
30-59d	number of times 30-59 days past due	int
60-89d	number of times 60-89 days past due	int
90d	number of times 90 days late	int
income	monthly income	float
debt	debt ratio	float
family	number of dependents	int
age	age	int
...	...	...

InterpretDL Application1 - Credit Scoring

Model Performance

Logistic Regression (ROC AUC : 0.784)

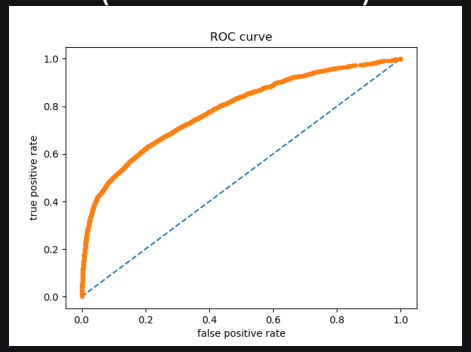

DNN (ROC AUC : 0.846)

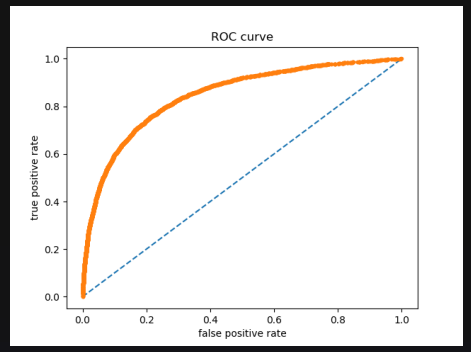

Interpretation

Logistic weights

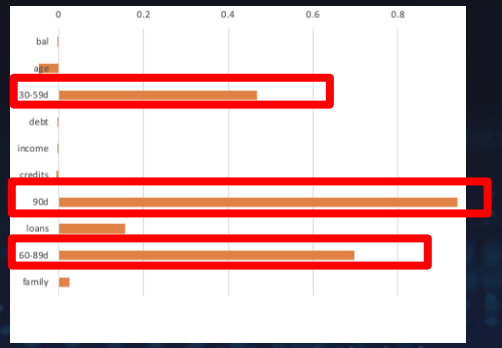

DNN + LIME weights

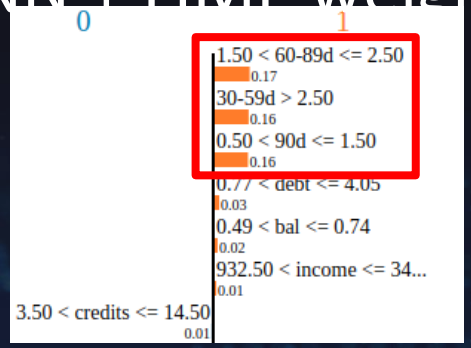

Better model with similar explanations

InterpretDL Application2 - Training Data Analyses

Synthesized mislabeled MNIST + Interpretation algorithm on training data:

To distinguish mislabeled data, samples that are easy to learn and those that are difficult.

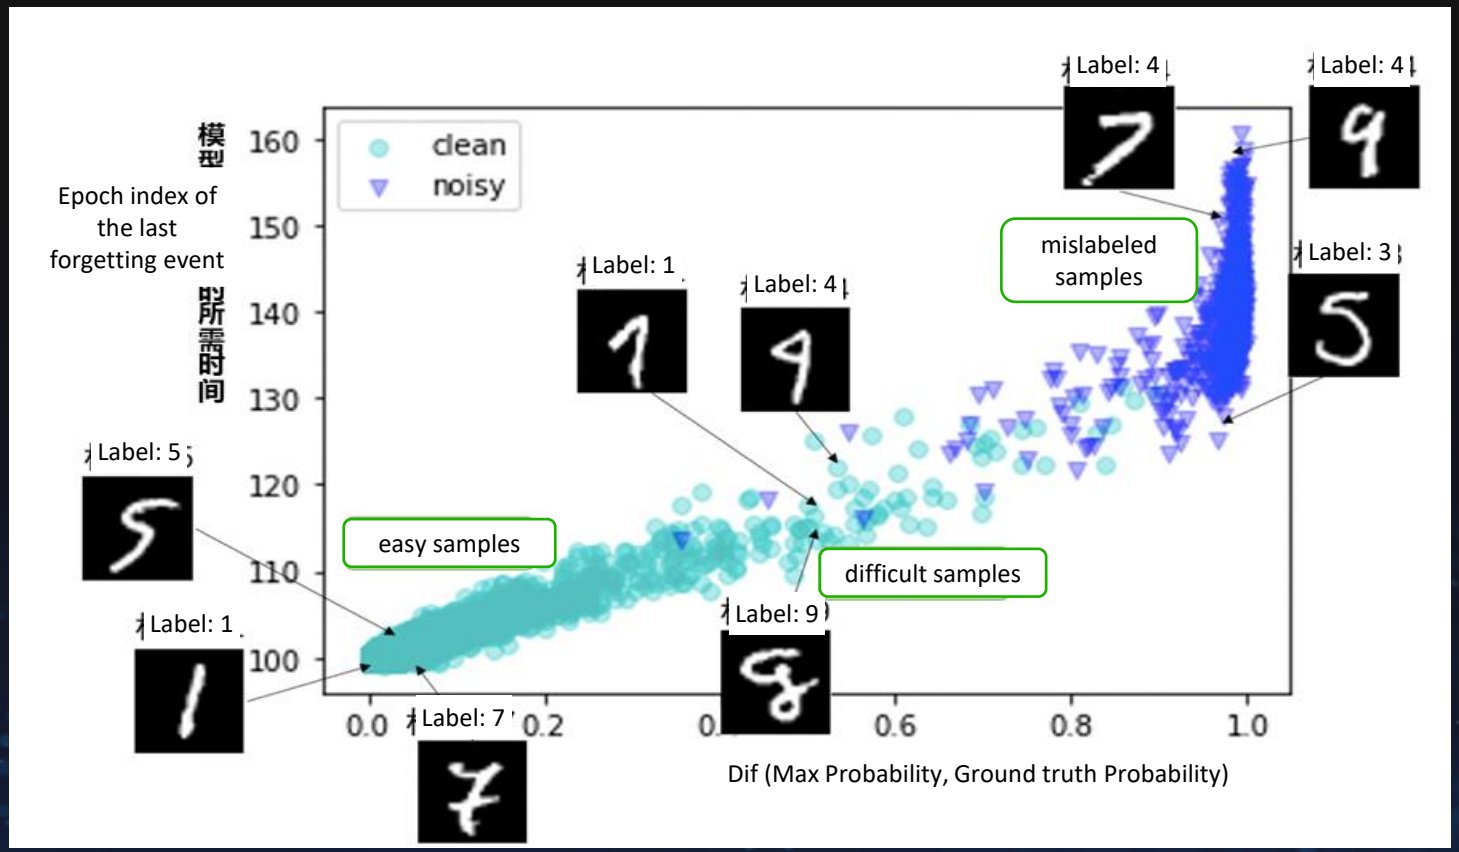

InterpretDL Application3 - Image Semantic Segmentation

InterpretDL Consensus + 100+ image classification models from PaddlePaddle

Vacuum

Cross-Model Consensus

LIME

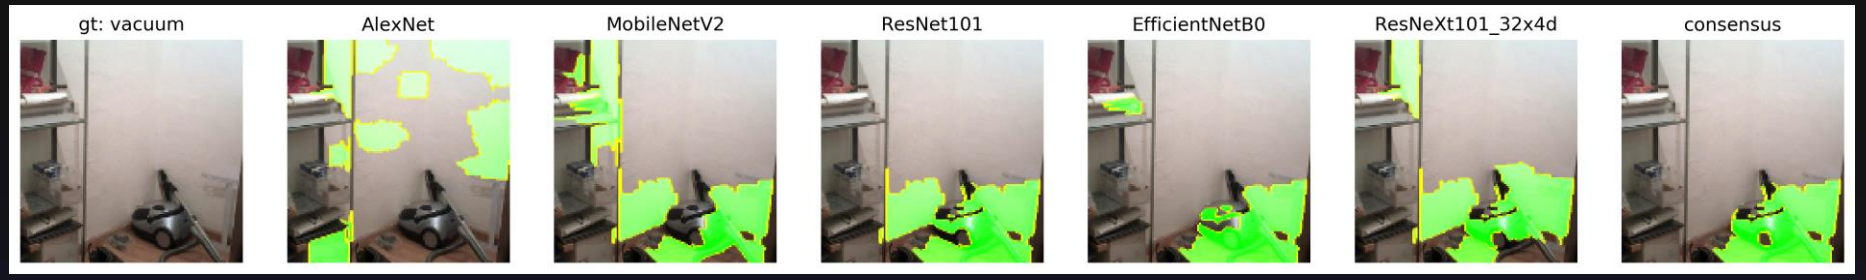

Smooth Gradient

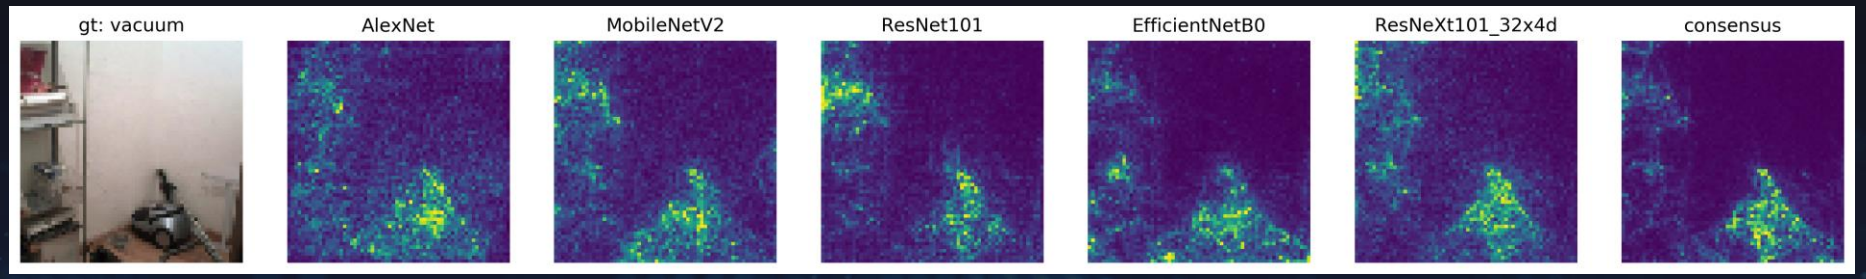

Li et al. "Cross-Model Consensus of Explanations and Beyond for Image Classification Models: An Empirical Study." arXiv:2109.00707 (2021).

InterpretDL Application3 - Image Semantic Segmentation

We have created a dataset named PSSL using Cross-Model Consensus, and pre-trained image segmentation models, to improve the performance on downstream tasks.

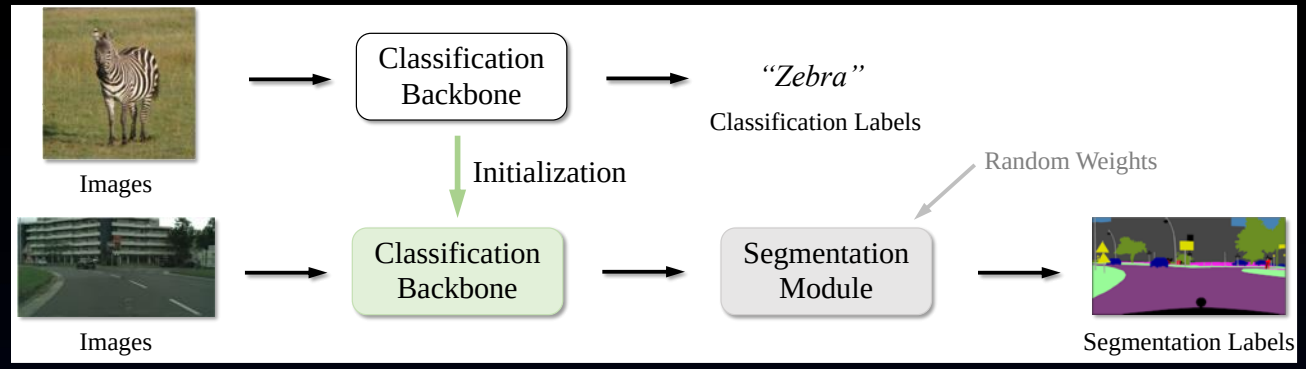

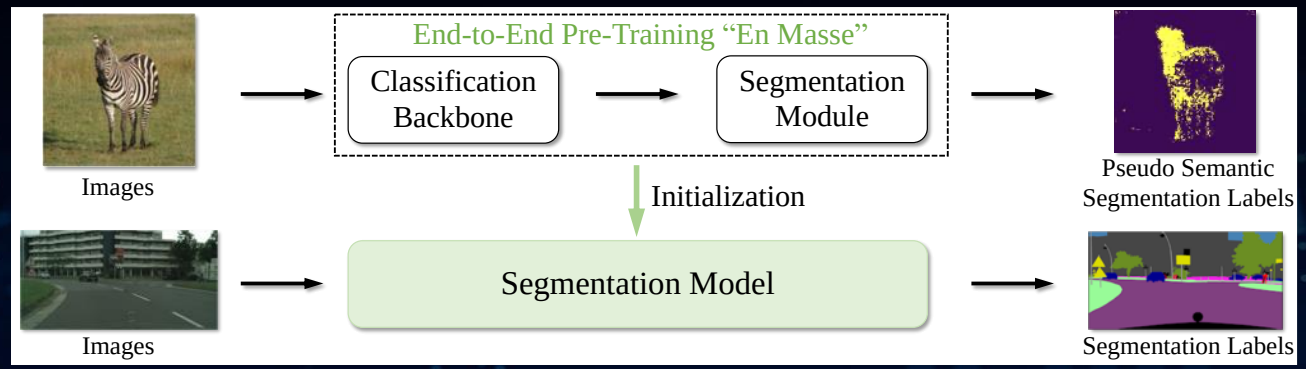

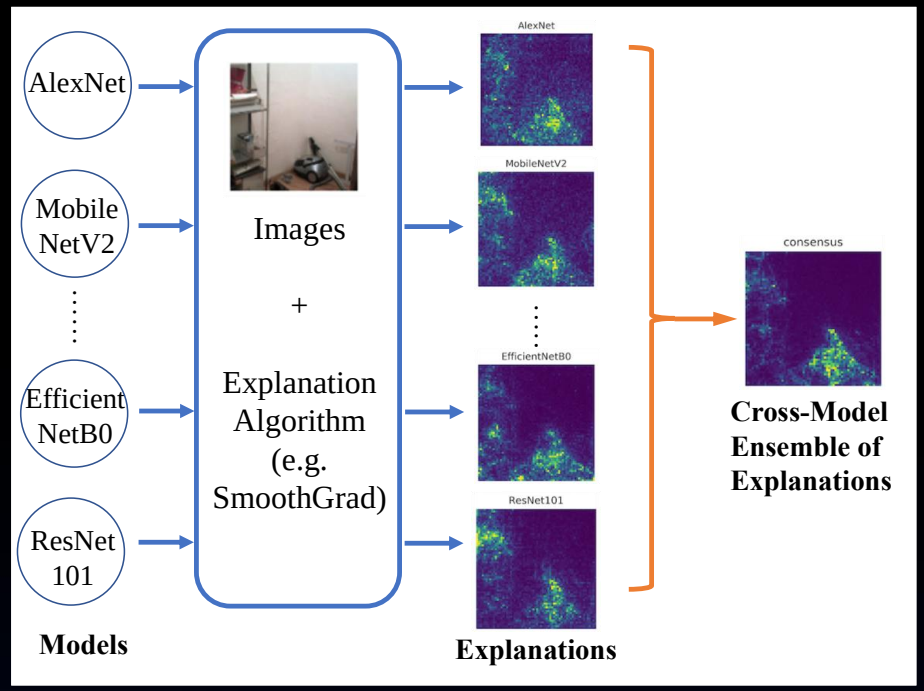

	PSPNet-ResNet50		DeepLabV3-ResNeSt50		OCRNet-HRNetW18	
	ImageNet	PSSL	ImageNet	PSSL	ImageNet	PSSL
CamVid	65.9	68.1 (+2.2)	66.6	69.1 (+2.5)	59.2	62.8 (+3.6)
VOC-A	79.4	80.3 (+0.9)	79.1	80.1 (+1.0)	76.4	77.1 (+0.7)
VOC-C	47.0	48.5 (+1.5)	48.8	49.4 (+0.6)	44.5	45.7 (+1.2)
ADE20K	42.9	43.8 (+0.9)	45.2	45.8 (+0.6)	40.0	40.9 (+0.9)
Cityscapes	78.7	78.9 (+0.2)	79.0	79.7 (+0.7)	79.6	79.8 (+0.2)

Li et al. "Distilling Ensemble of Explanations for Weakly-Supervised Pre-Training of Image Segmentation Models." Machine Learning Journal (accepted with minor revision).

References

- Isaac Ahern, Adam Noack, Luis Guzman-Nateras, Dejing Dou, Boyang Li, and Jun Huan. 2019. “NormLIME: A new feature importance metric for explaining deep neural networks.” arXiv preprint arXiv:1909.04200 (2019).
- Xuhong Li, Haoyi Xiong, Xingjian Li, Xuanyu Wu, Xiao Zhang, Ji Liu, and Dejing Dou. "Interpretable Deep Learning: Interpretation, Interpretability, Trustworthiness, and Beyond." arXiv preprint arXiv:2103.10689 (2021).
- Xuhong Li, Haoyi Xiong, Siyu Huang, Shilei Ji and Dejing Dou. "Cross-Model Consensus of Explanations and Beyond for Image Classification Models: An Empirical Study." arXiv preprint arXiv:2109.00707 (2021).
- Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 2921-2929, 2016.
- David Bau, Bolei Zhou, Aditya Khosla, Aude Oliva, and Antonio Torralba. Network dissection: Quantifying interpretability of deep visual representations. In IEEE Transactions on Pattern Analysis and Machine Intelligence, pages 6541–6549, 2017.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. “why should i trust you?” explaining the predictions of any classifier. In Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining, pages 1135–1144, 2016.
- Daniel Smilkov, Nikhil Thorat, Been Kim, Fernanda Vi´egas, and Martin Wattenberg. Smoothgrad: removing noise by adding noise. arXiv preprint arXiv:1706.03825, 2017.
- Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In International Conference on Machine Learning (ICML), 2017.
- Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In Proceedings of the IEEE international conference on computer vision, pages 618–626, 2017.
- Haofan Wang, Zifan Wang, Mengnan Du, Fan Yang, Zijian Zhang, Sirui Ding, Piotr Mardziel, and Xia Hu. Score-cam: Score-weighted visual explanations for convolutional neural networks. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, pages 24–25, 2020.
- Mariya Toneva, Alessandro Sordoni, Remi Tachet des Combes, Adam Trischler, Yoshua Bengio, and Geoffrey J Gordon. An empirical study of example forgetting during deep neural network learning. International Conference on Learning Representations (ICLR), 2019.
- Pang Wei Koh and Percy Liang. Understanding black-box predictions via influence functions. In International Conference on Machine Learning (ICML), pages 1885–1894. PMLR, 2017.
- Petsiuk, Vitali, Abir Das, and Kate Saenko. "Rise: Randomized input sampling for explanation of black-box models." arXiv preprint arXiv:1806.07421 (2018).
- Chih-Kuan Yeh, Cheng-Yu Hsieh, Arun Sai Suggala, David I Inouye, and Pradeep Ravikumar. On the (in) fidelity and sensitivity for explanations. arXiv preprint arXiv:1901.09392, 2019.
- Tingyi Yuan, Xuhong Li, Haoyi Xiong, Hui Cao, and Dejing Dou. Explaining Information Flow Inside Vision Transformers Using Markov Chain. In XAI4Debugging @Neurips 2021 Workshop.

Thanks

Outreach Committee Update

March 24, 2022

Hu Xiaoman (Charlotte)

THE LINUX FOUNDATION

LF AI & DATA

Upcoming Events

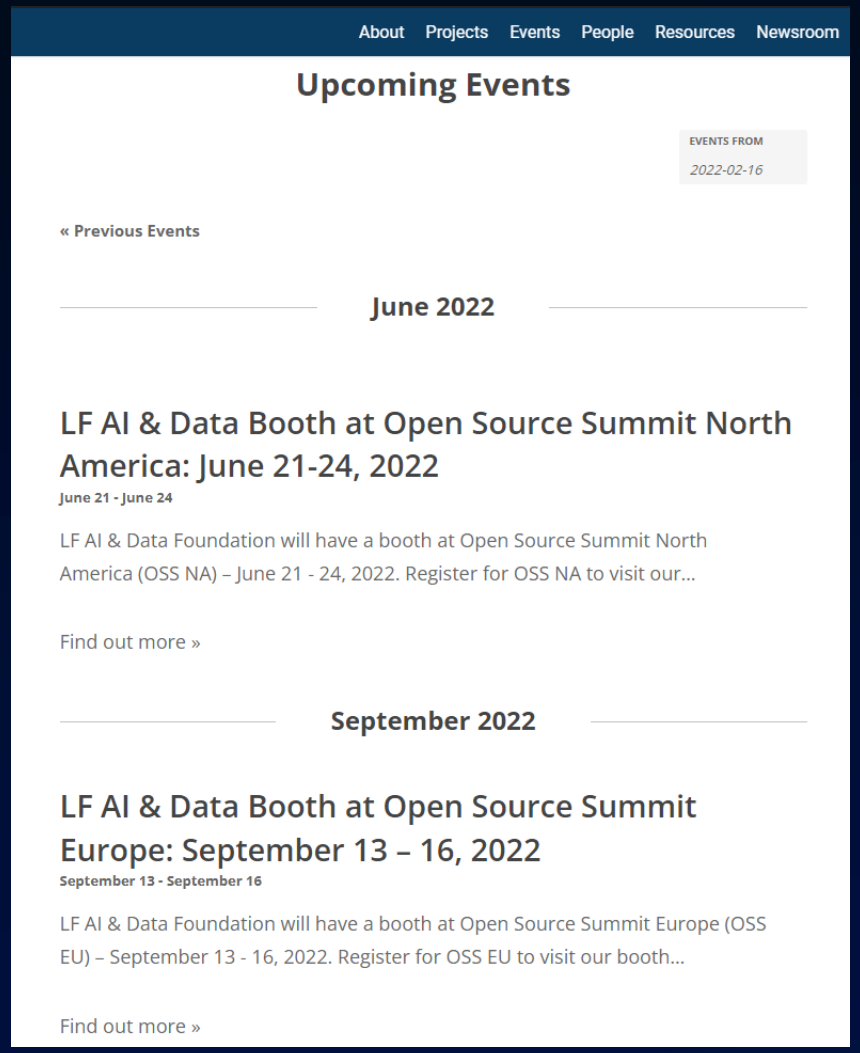

Visit the LF AI & Data Events Calendar or the LF AI & Data 2022 Events wiki for a list of all upcoming events

To discuss participation in an event or to host an event, please email events@lfaidata.foundation

Event Call For Proposals

All **Linux Foundation** events for 2022 are published here

We have an AI/Data track in KubeCons and Open Source Summits, in addition to various smaller events with an AI/Datafocus

Browse the events list and determine if you'd like to submit a proposal to present your project

Recent Announcements

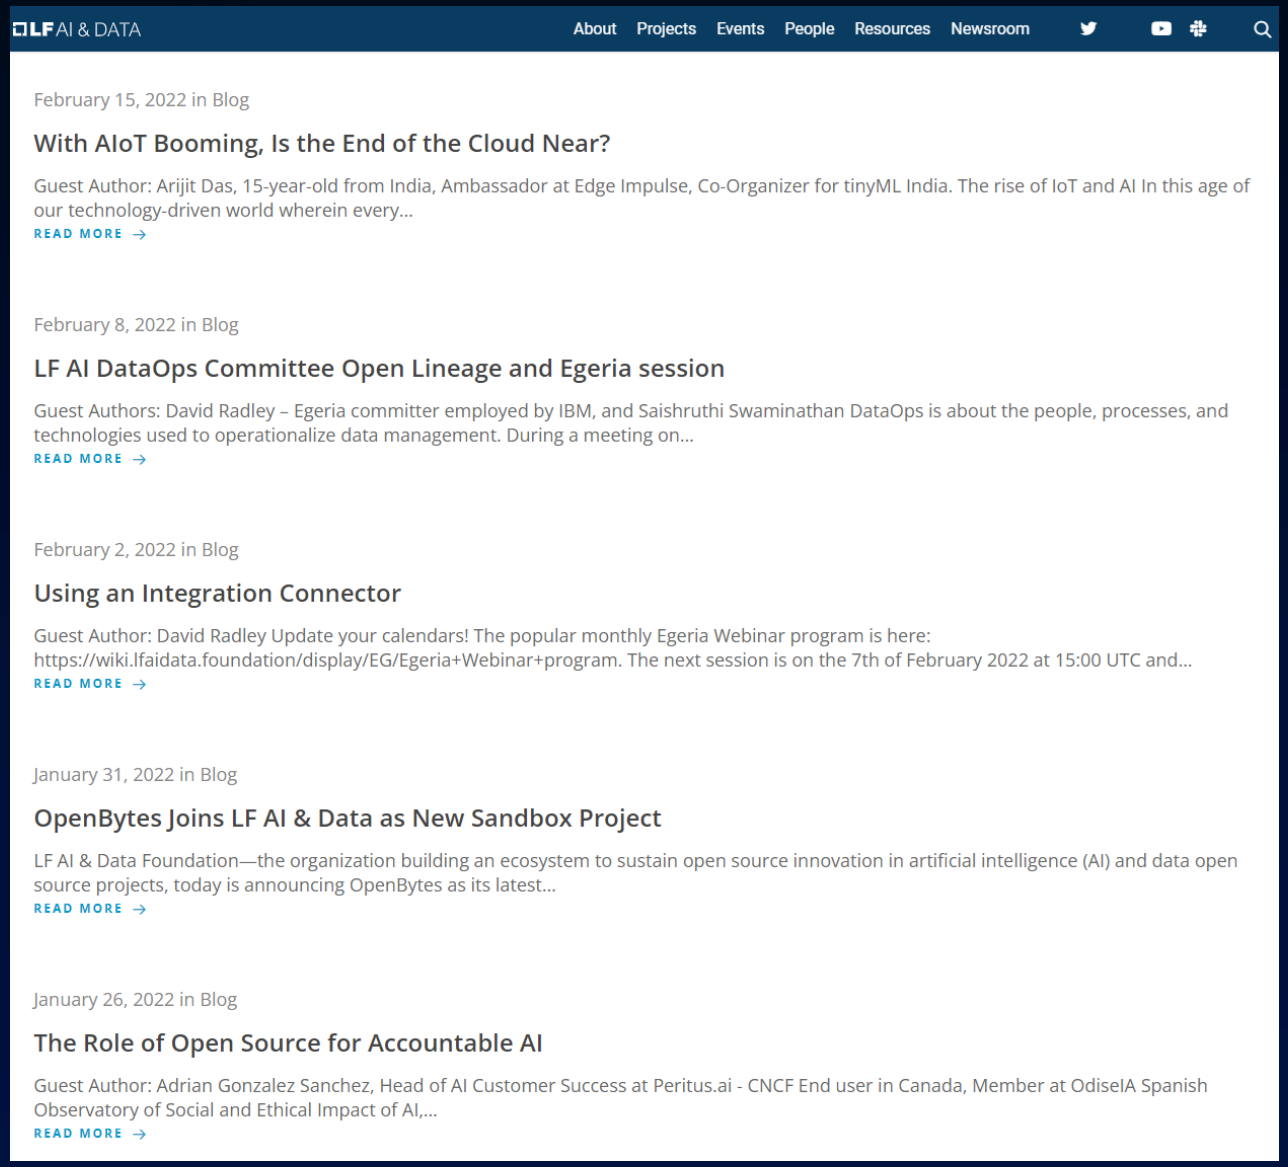

Visit the LF AI & Data Blog for announcement details

Subscribe here to tac-general mailing list to get emails with recent announcements

LF AI & DATA

Promoting Upcoming Project Releases

We promote project releases via a blog post and on LF AI & Data Twitter and/or LinkedIn social channels

If you are an LF AI & Data hosted project and would like LF AI & Data to promote your release, reach out to pr@lfaidata.foundation to coordinate in advance (min 2 wks) of your expected release date - Review guidelines here

LF AI & Data PR/Comms

Please follow LF AI & Data on Twitter (@LFAIDataFdn) & LinkedIn and help amplify news via your social networks - Please retweet and share!

Open call to publish project/committee updates or other relevant content on the LF AI & Data Blog - Review guidelines here

To discuss more details on participation or upcoming announcements, please email pr@lfaidata.foundation

Upcoming TAC Meetings

LF AI & DATA

Upcoming TAC Meetings

- April 7, 2022 – RosaeNLG Annual Review
- April 21, 2022 – Micro ML – new project (tentative)

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

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