Argonne NATIONAL LABORATORY

APPLICATION OF PRIVACY PRESERVING FEDERATED LEARNING IN BIOMEDICAL APPLICATIONS – LESSONS LEARNED FROM THE PALISADE-X PROJECT

APPFL

globus

funcX

RAVI MADDURI
Data Science and Learning Division
madduri@anl.gov

U.S. DEPARTMENT OF ENERGY
Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.

Nov 16, 2022

FUNDING ACKNOWLEDGEMENTS

THIS MATERIAL IS BASED UPON WORK SUPPORTED BY THE U.S. DEPARTMENT OF ENERGY, OFFICE OF SCIENCE, UNDER CONTRACT NUMBER DE-AC02-06CH11357.

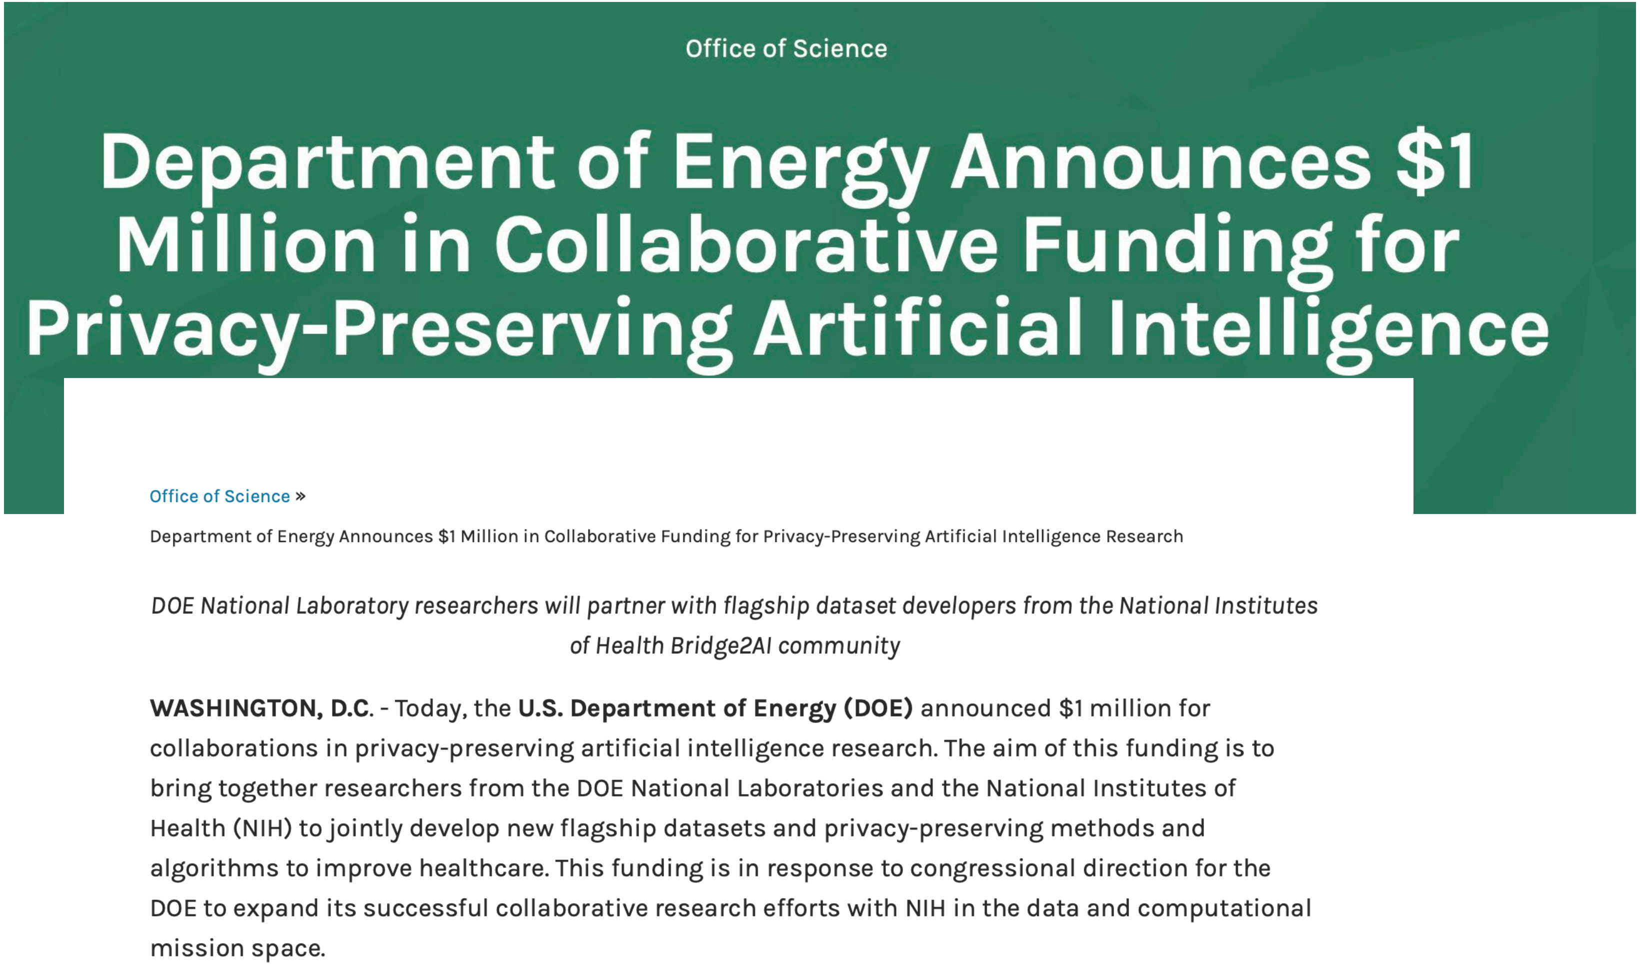

Office of Science

Department of Energy Announces $1 Million in Collaborative Funding for Privacy-Preserving Artificial Intelligence

Office of Science »

Department of Energy Announces $1 Million in Collaborative Funding for Privacy-Preserving Artificial Intelligence Research

DOE National Laboratory researchers will partner with flagship dataset developers from the National Institutes of Health Bridge2AI community

WASHINGTON, D.C. - Today, the **U.S. Department of Energy (DOE)** announced $1 million for collaborations in privacy-preserving artificial intelligence research. The aim of this funding is to bring together researchers from the DOE National Laboratories and the National Institutes of Health (NIH) to jointly develop new flagship datasets and privacy-preserving methods and algorithms to improve healthcare. This funding is in response to congressional direction for the DOE to expand its successful collaborative research efforts with NIH in the data and computational mission space.

Source: Press release on PALISADE-X

TEAM

Multi-disciplinary team from ANL, LLNL, Harvard, UChicago, UIUC

- Team ANL
  - Ravi Madduri (Computer Scientist)
  - Kibaek Kim (Computational Mathematician)
  - Minseok Ryu (Postdoctoral Appointee)
  - Miao Li (Predoctoral Appointee)
  - Hieu Hoang (Computational Researcher)
  - Eliu Huerta (Computational Scientist)
- Team UChicago (Radiology)
  - Maryellen Giger
  - Sam Armato
  - Jordan Fuhrman
- Team Harvard (Genomics and EHR)
  - Pradeep Natarajan
  - Puneet Batra
  - Marcus Klarquist
  - Sarah Urbut
- Team LLNL (Secure Infrastructure)
  - Kyle Halliday
  - Megan Utter

Ravi Madduri

Kibaek Kim

Ryu Minseok

Eliu Huerta

Pradeep Natarajan

Puneet Batra

Maryellen Giger

Sam Armato

MOTIVATION FOR PALISADE-X

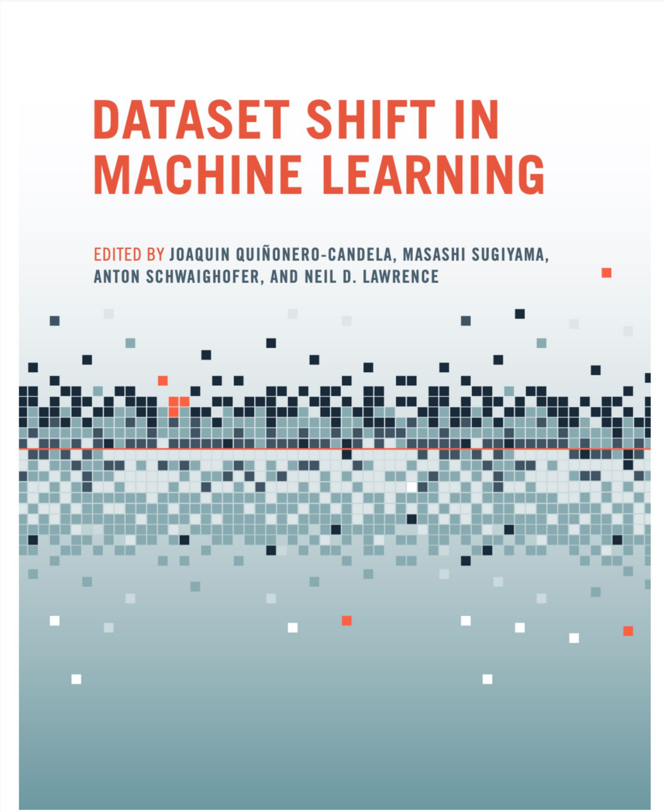

DATASET SHIFT IN MACHINE LEARNING

EDITED BY JOAQUIN QUIÑONERO-CANDELA, MASASHI SUGIYAMA, ANTON SCHWAIGHOFER, AND NEIL D. LAWRENCE

THE LANCET
Digital Health

ARTICLES | VOLUME 4, ISSUE 6, E406-E414, JUNE 01, 2022

AI recognition of patient race in medical imaging: a modelling study

The NEW ENGLAND JOURNAL of MEDICINE

CORRESPONDENCE

The Clinician and Dataset Shift in Artificial Intelligence

Source: PMID: 34260843 DOI: [10.1056/NEJMc2104626]

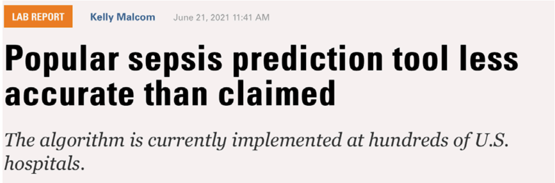

LAB REPORT Kelly Malcom June 21, 2021 11:41 AM

Popular sepsis prediction tool less accurate than claimed

The algorithm is currently implemented at hundreds of U.S. hospitals.

Source: https://labblog.uofmhealth.org/lab-report/popular-sepsis-prediction-tool-less-accurate-than-claimed

KEY CAPABILITIES OF PALISADE-X

To Build Models that are Fair and Trustworthy using PPFL

- End-to-End strong IAM
  - Enable setting up Secure Federation across organizational boundaries
- Easy to leverage HPC for training
  - Integrate heterogenous computing resources
- Secure Enclaves for compliant data storage and facilitate secure federation
  - Software to create virtual private enclaves on supercomputers
- Framework to evaluate privacy protection possible for different data modalities
  - Different levels of privacy budgets to conform with different compliance – FISMA High, Medium, Low
  - Frameworks and approaches to measure privacy protection by attacking models
- APIs and Plug-and-Play architecture
  - To integrate into existing services and add new capabilities/algorithms

CHALLENGES UNIQUE TO AI IN BIOMEDICINE

As they relate to building better AI models

- AI is a data hungry sport
- In Biomedicine, Data is
  - Often private and sensitive
  - Comes in different modalities
  - With different distributions
- In Biomedicine, you will hear about
  - FISMA High, Moderate
  - Covered Entities
  - HIPAA
  - DUAs
  - IRBs
- The non-technical/policy challenges of general data availability leads to models that are under specifed and overfitted

PALISADE-X

Privacy-preserving analysis and learning in secure and distributed enclaves and exascale systems

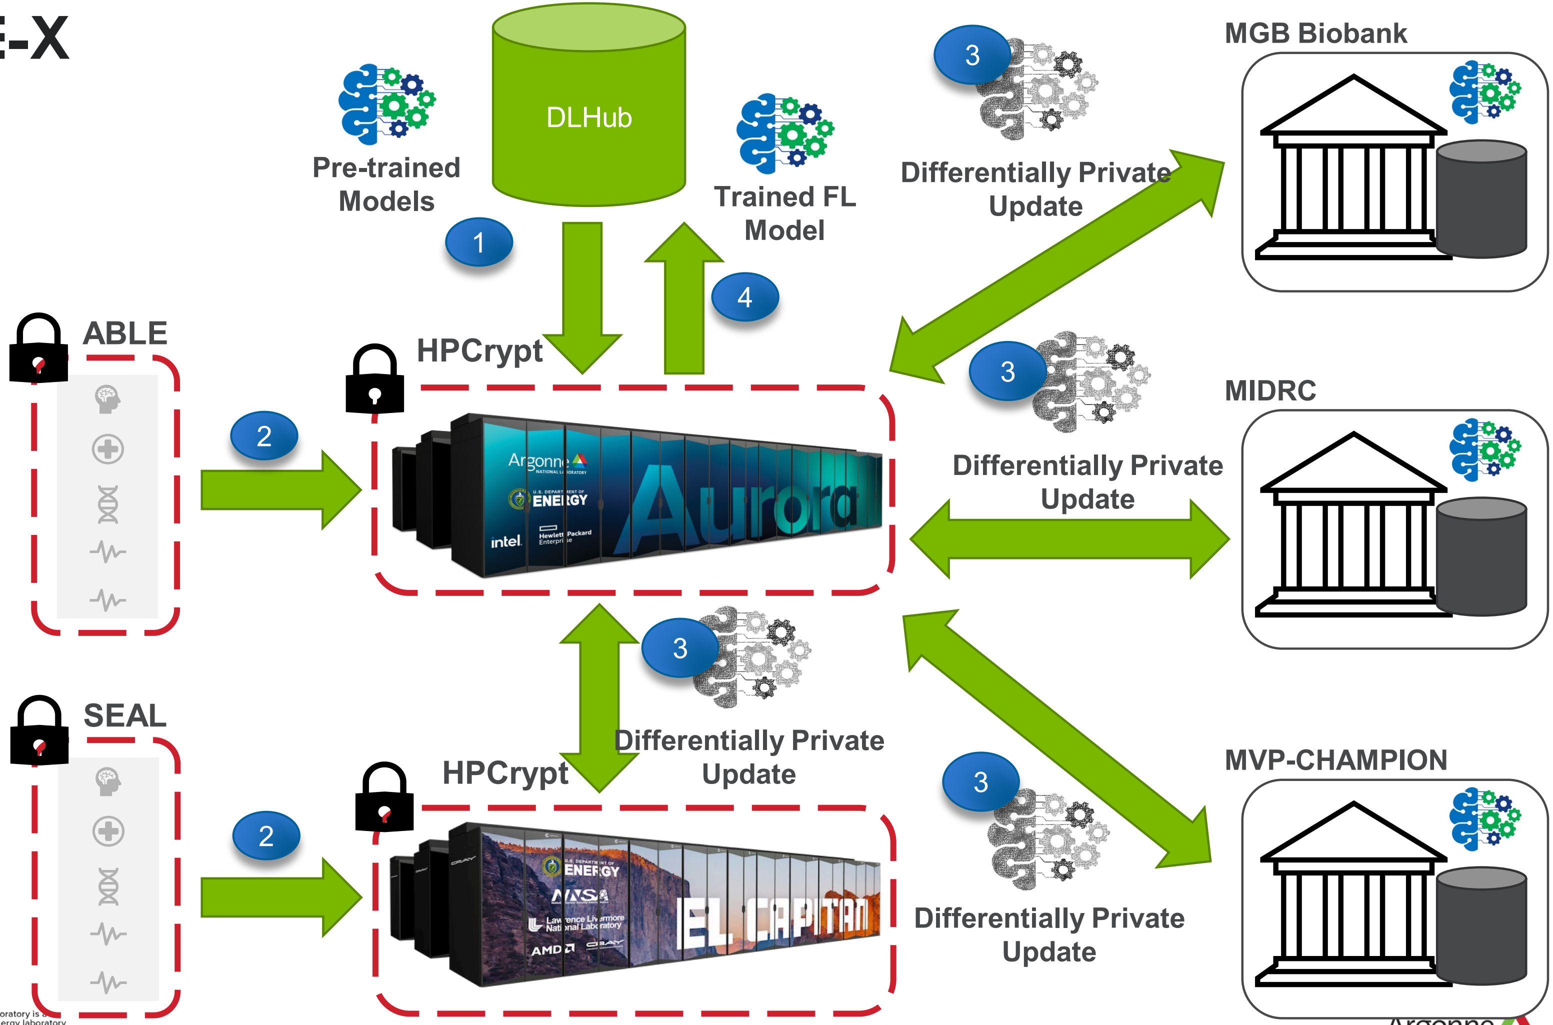
DLHub
Pre-trained Models
1
Trained FL Model
4
3
Differentially Private Update
MGB Biobank
ABLE
HPCrypt
2
Aurora
3
Differentially Private Update
MIDRC
3
Differentially Private Update
SEAL
HPCrypt
2
EL CAPITAN
3
Differentially Private Update
MVP-CHAMPION

KEY CAPABILITIES – SECURE ENCLAVES

Computing, Policies

Argonne biomedical learning enclave (ABLE)

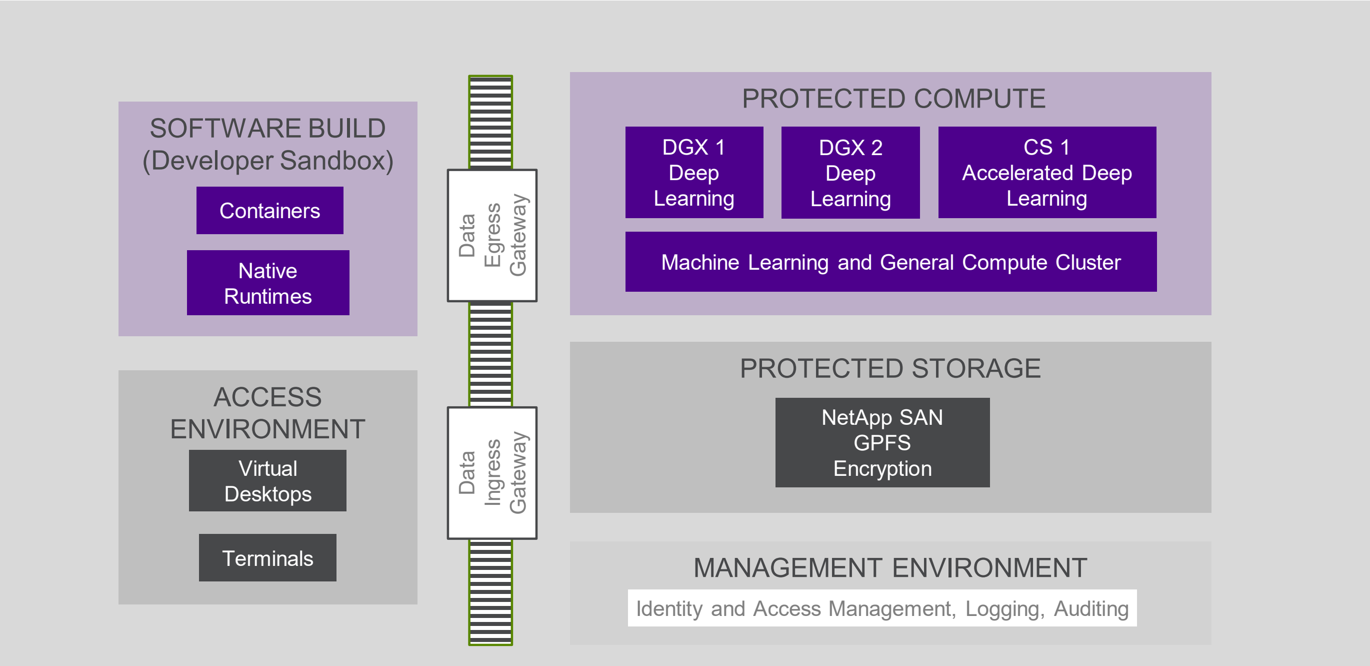

Compliance Shield
Sword & Shield
HIPAA COMPLIANT

ARGONNE BIOMEDICAL LEARNING ENCLAVE

Certified for HIPAA Compliance

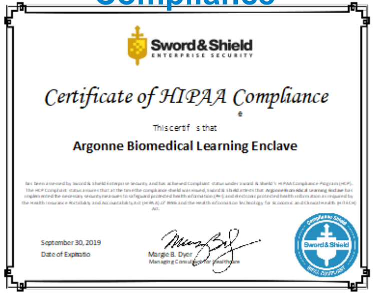

Sword & Shield
ENTERPRISE SECURITY

Certificate of HIPAA Compliance
e

This certif s that

Argonne Biomedical Learning Enclave

has been assessed by Sword & Shield Enterprise Security and has achieved Compliant status under Sword & Shield's HIPAA Compliance Program (HCP). The HCP Compliant status assures that at the time the compliance shield was issued, Sword & Shield attests that Argonne Biomedical Learning Enclave has implemented the necessary security measures to safeguard protected health information (PHI) and electronic protected health information as required by the Health Insurance Portability and Accountability Act (HIPAA) of 1996 and the Health Information Technology for Economic and Clinical Health (HITECH) Act.

September 30, 2019
Date of Expiratio

Margie B. Dyer
Managing Consultant for Healthcare

KEY CAPABILITIES HPCRYPT

Software to securely leverage supercomputers for AI and HPC

- HPCrypt enables securely leverage HPC resources
- HPCrypt creates *virtual enclaves* that provide additional security within HPC systems
- Integrated HPCrypt capabilities with ABLE

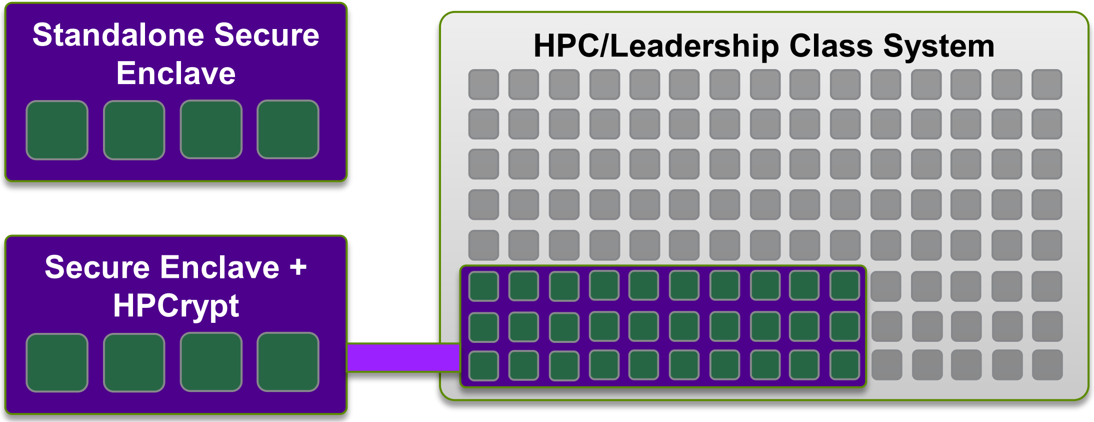

Virtual enclave extends into HPC system for duration of job execution

ARGONNE PRIVACY PRESERVING FEDERATED LEARNING FRAMEWORK

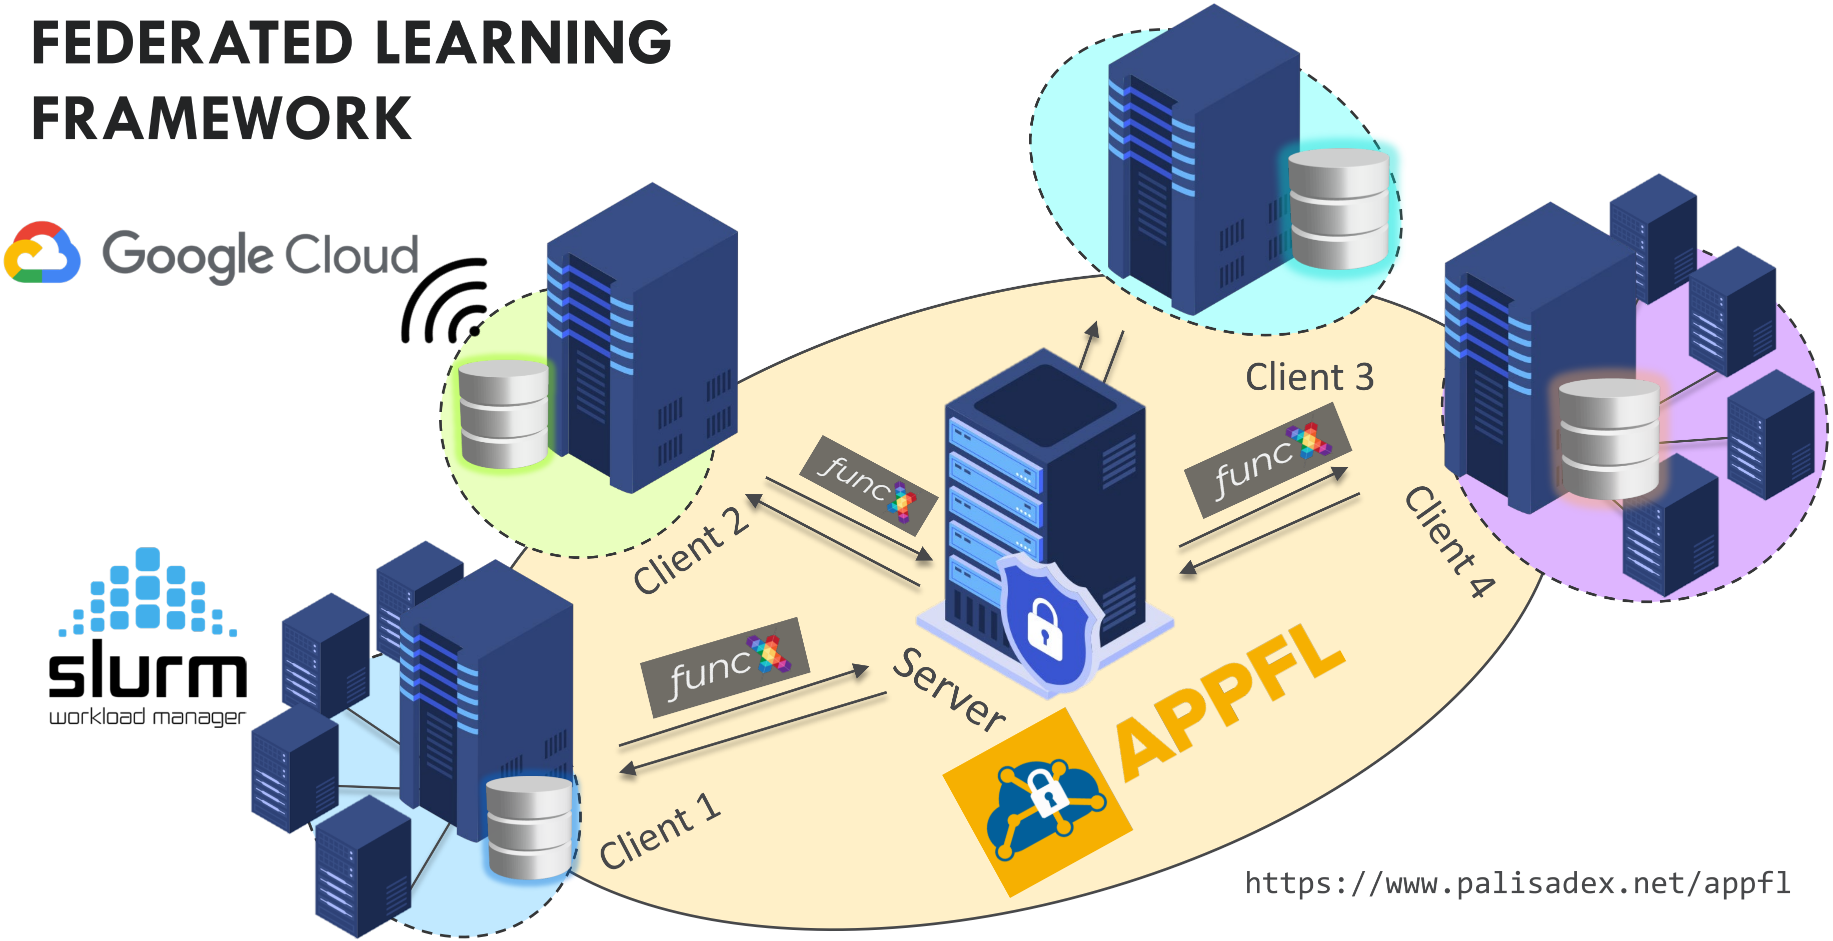

https://www.palisadex.net/appfl

IDENTITY AND ACCESS MANAGEMENT FOR EASY AND SECURE FEDERATION

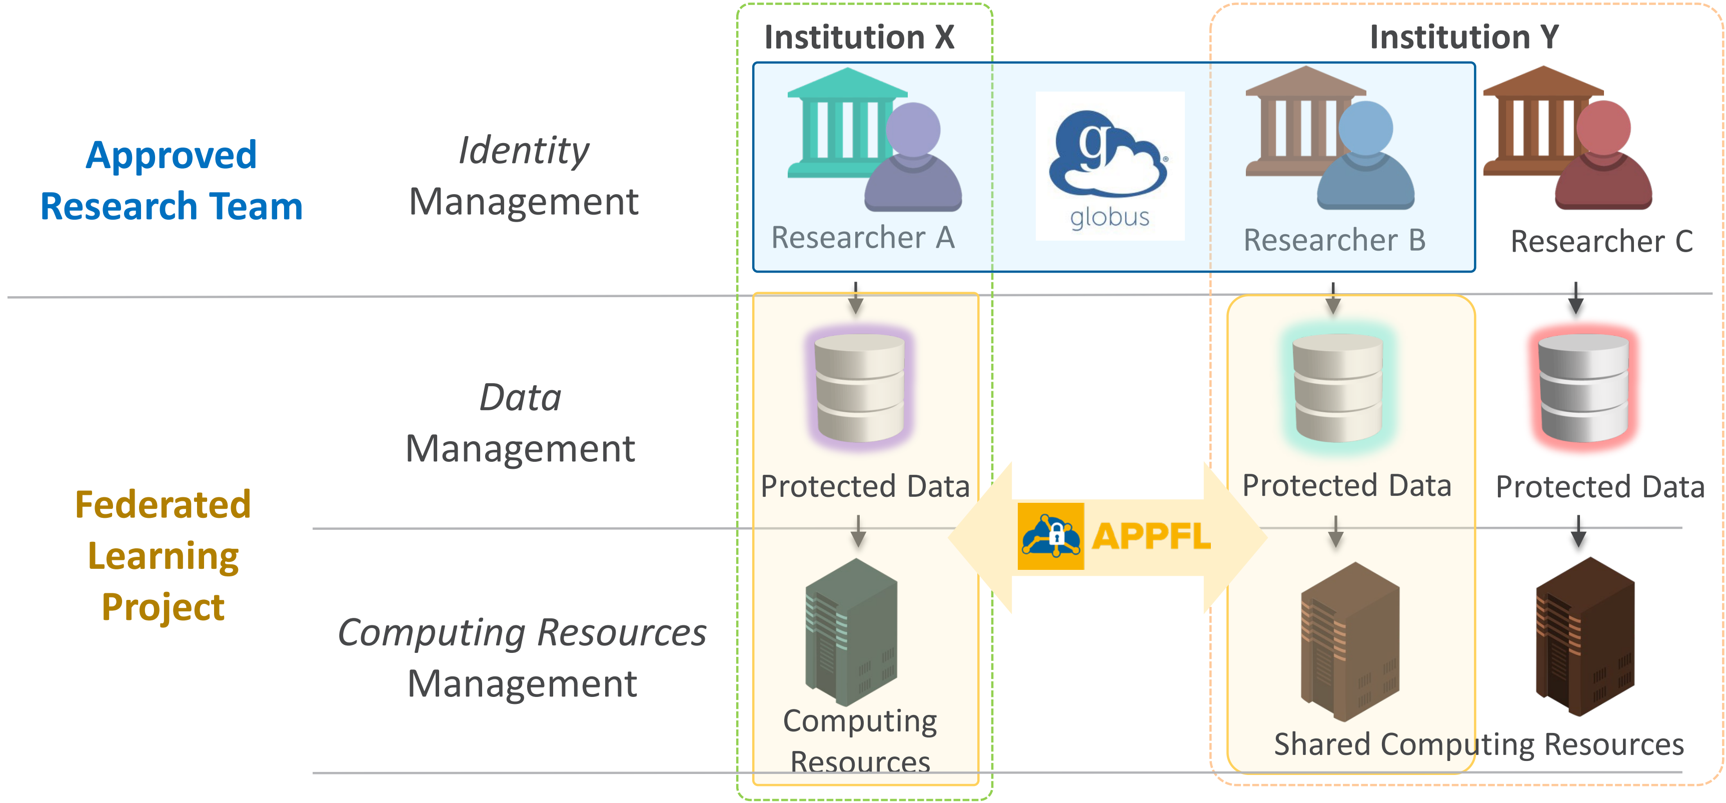

END-TO-END PPFL IN PRACTICE

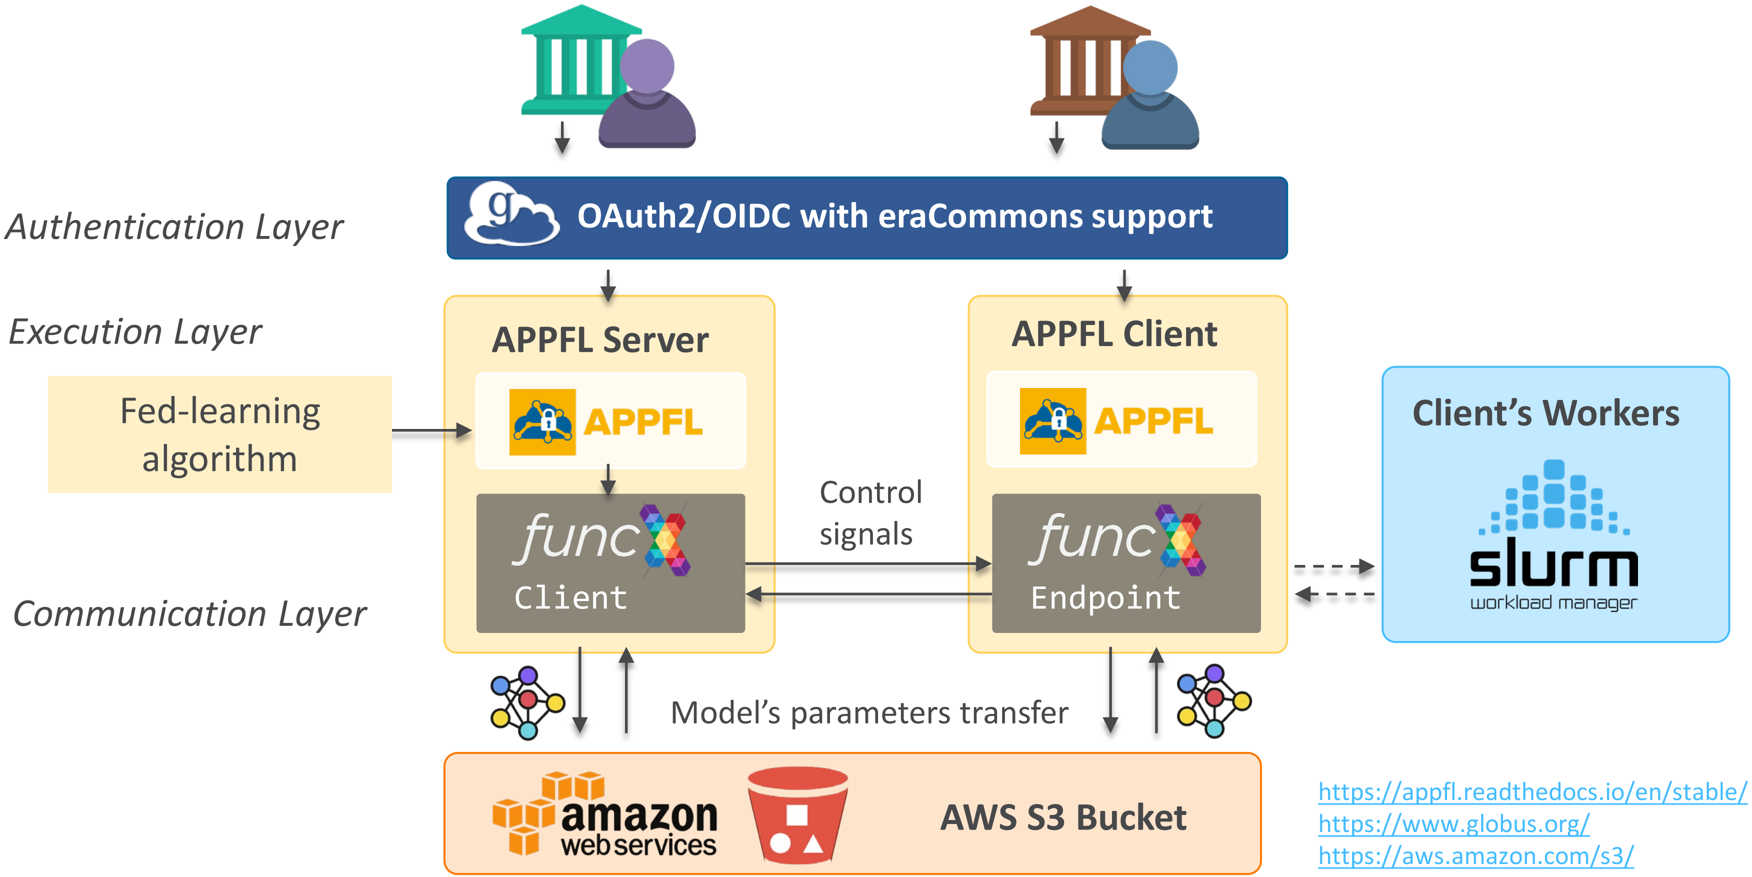

Authentication Layer

Execution Layer

Communication Layer

OAuth2/OIDC with eraCommons support

Fed-learning algorithm

APPFL Server

APPFL Client

Client's Workers

Control signals

Model's parameters transfer

AWS S3 Bucket

https://appfl.readthedocs.io/en/stable/
https://www.globus.org/
https://aws.amazon.com/s3/

ADDITIONAL ONGOING WORK

- Systematic evaluation of different attack modalities
  - Joint work with Miao Li and Mihai Anitescu
  - Attack models include inverse gradient approach, Optimization-based approach like Deep Leakage from Gradients (DLG) and Solving a sequence of linear equations in the R-Gap(Recursive Gradient Attack On Privacy)
- Develop and apply a methodology for providing tiered levels of privacy assurance for a privacy-preserving federated learning framework, while validating the security of the overall system against risks such as model poisoning/corruption, denial of service, or intentional prevention of convergence
  - Joint work with Argonne’s Cyber team (Blakely et al.)

APPLYING APPFL IN BIOMEDICINE APPLICATIONS & CHALLENGES

BIOMEDICAL APPLICATIONS

Detection of COVID-19 from Chest X-Rays

Prediction of age from ECGs to use in models predicting risk for a CVD event

DETECTION OF COVID-19 FROM CHEST X-RAYS

- **Datasets:**
  - ANL-COVID: the dataset is aggregated from multiple open-source datasets
  - Uchicago-COVID: private dataset collected by UChicago

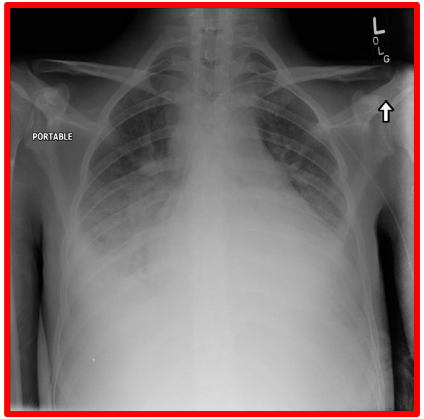

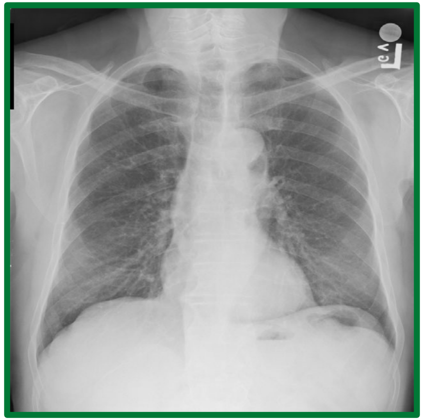

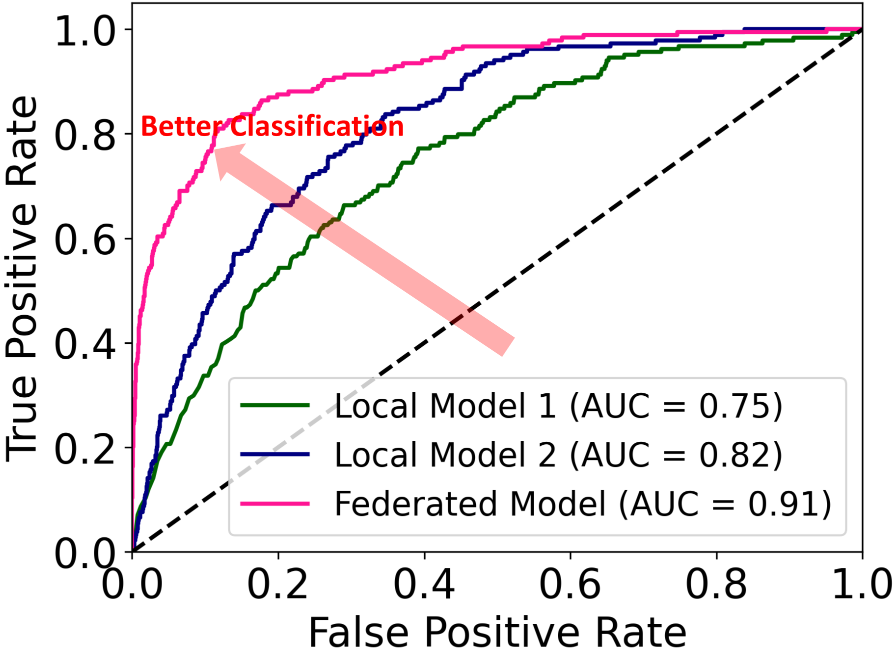
Better Classification
True Positive Rate
False Positive Rate
0.0
0.2
0.4
0.6
0.8
1.0
Local Model 1 (AUC = 0.75)
Local Model 2 (AUC = 0.82)
Federated Model (AUC = 0.91)

Results from initial proof of principle experiments show improved performance

TRAINING DATA/CLASS DISTRIBUTION

Dataset	Train	Val	Test
ANL- COVID19 +/-	13992+/16490-	200+/200-	200+/200-
Uchicago – COVID19 +/-	974+/5336-	244+/1334-	305+/1667-

Training set at ANL is roughly 5 times larger than the training set at UChicago/MIDRC

MIDRC
MEDICAL IMAGING AND DATA RESOURCE CENTER.

DETECTION OF COVID-19 FROM CHEST X-RAYS

Validation accuracy at ANL Client

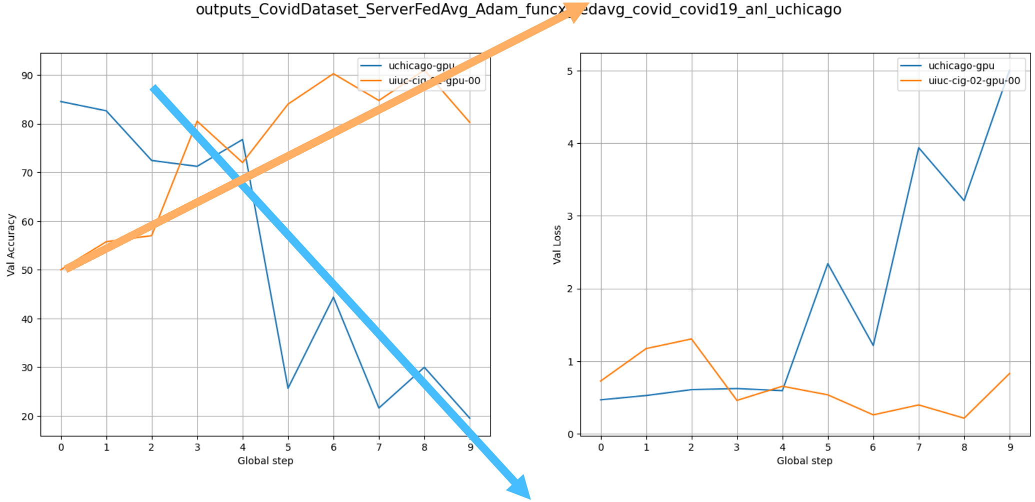

Validation accuracy at Uchicago Client

Since the ratio $\frac{n_1}{n} \gg \frac{n_2}{n}$, the client at ANL (1) has more influence!

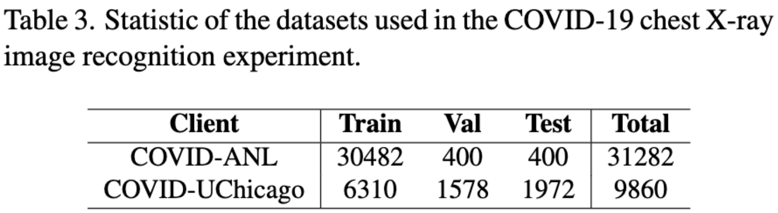

Table 3. Statistic of the datasets used in the COVID-19 chest X-ray image recognition experiment.

Client	Train	Val	Test	Total
COVID-ANL	30482	400	400	31282
COVID-UChicago	6310	1578	1972	9860

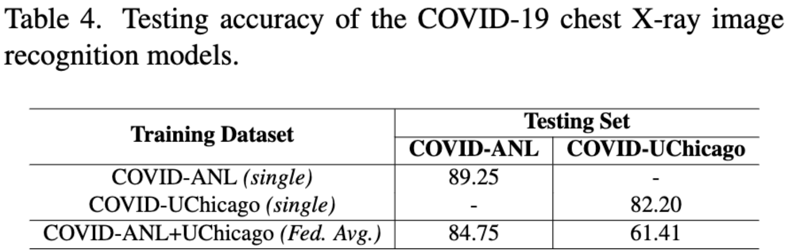

Table 4. Testing accuracy of the COVID-19 chest X-ray image recognition models.

Training Dataset	Testing Set	
	COVID-ANL	COVID-UChicago
COVID-ANL *(single)*	89.25	-
COVID-UChicago *(single)*	-	82.20
COVID-ANL+UChicago *(Fed. Avg.)*	84.75	61.41

OVERVIEW OF ECG USECASE

- Chronologic age can be a poor predictor of lifetime CVD risk, particularly among younger individuals
- Augmenting with additional variables can help to refine these estimates, but still ignore the component of variation explained by age
- Replacing with more biological proxies for age, such as from ECGs, can resolve these issues
- ECGs are typically not shared across, or sometimes even within, institutions

WORST PERFORMING ECG: 21-22 YEARS OLDER COMPARED TO CHRONOLOGICAL AGE

Predicted: 78.342468

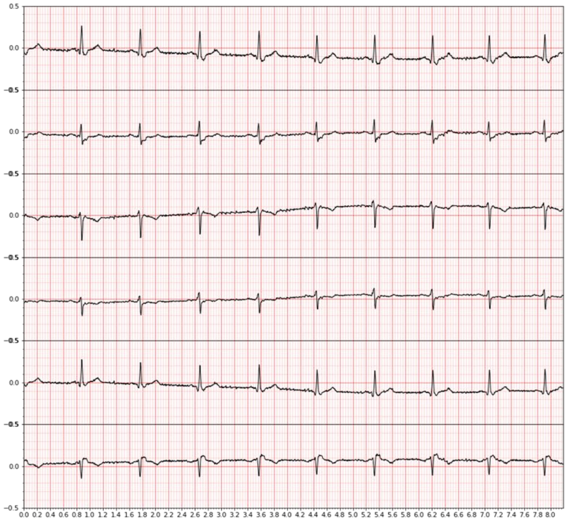

Real: 56.861546

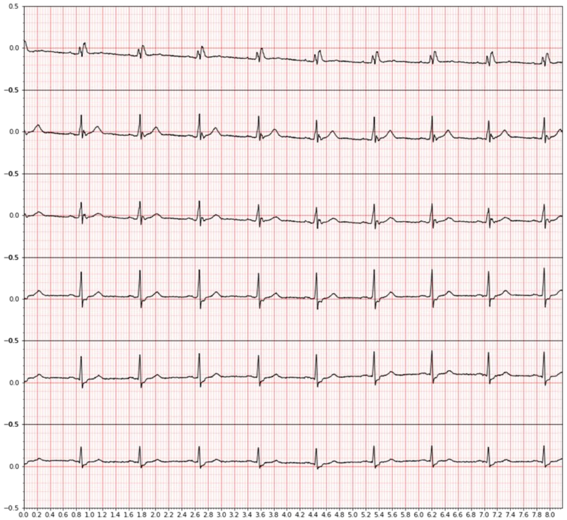

How does this age-proxy affect disease risk?

REGRESSED AGE IMPROVES PREDICTIVE POWER FOR YOUNGER AND HARDER-TO-PREDICT SUBJECTS

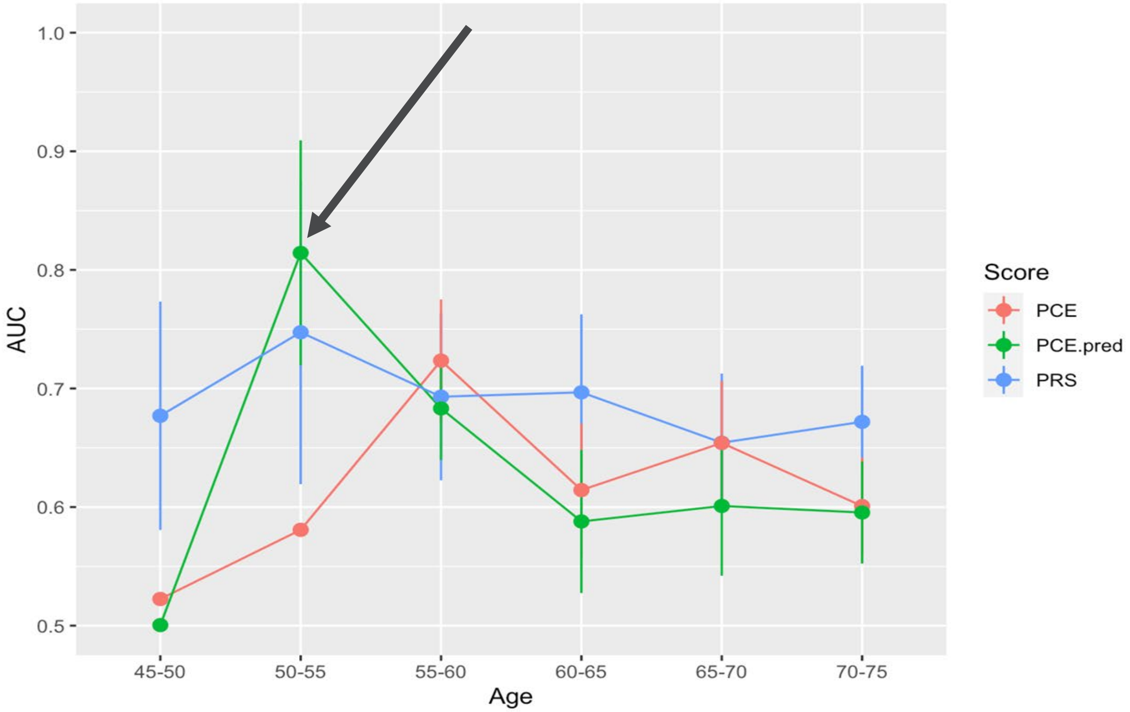

Urbut et al. (In preparation)

TRAINING DATA/CLASS DISTRIBUTION

Biological Aging Prediction from ECG Signal

- **Training Data Details**
  - ECG-ANL: 80,569 ECGs in PhysioNet dataset
  - ECG-Broad: private dataset with 37,623 ECGs collected from the UK Biobank
- Age at the time of ECG reading is computed as DOB - date of reading

BIOLOGICAL AGING PREDICTION FROM ECG SIGNAL

Training with FL on two clients

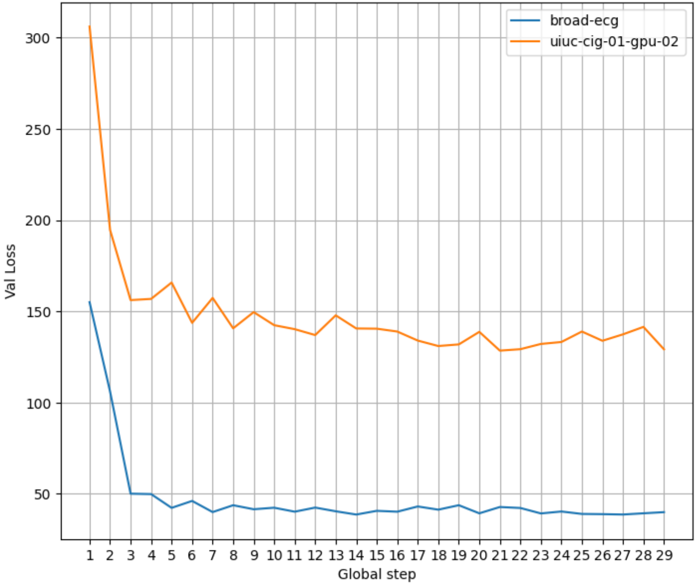

Best MSE on ECG-ANL = 125.00
Best MSE on ECG-Broad = 41.70

FL can learn a global model that performs relative well on both datasets

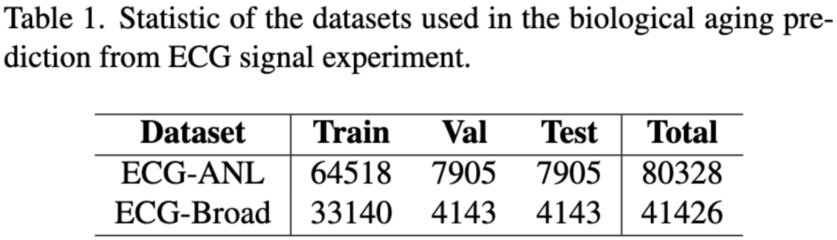

Table 1. Statistic of the datasets used in the biological aging prediction from ECG signal experiment.

Dataset	Train	Val	Test	Total
ECG-ANL	64518	7905	7905	80328
ECG-Broad	33140	4143	4143	41426

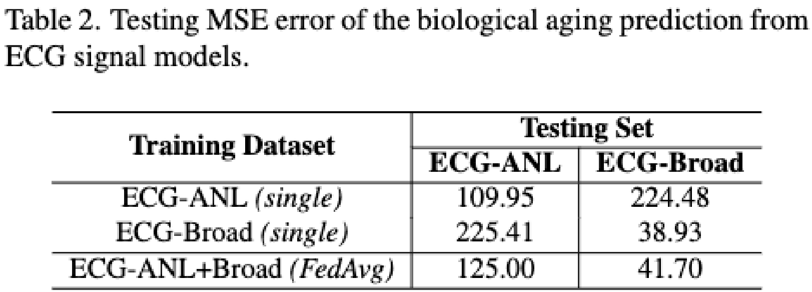

Table 2. Testing MSE error of the biological aging prediction from ECG signal models.

Training Dataset	Testing Set	
	ECG-ANL	ECG-Broad
ECG-ANL *(single)*	109.95	224.48
ECG-Broad *(single)*	225.41	38.93
ECG-ANL+Broad *(FedAvg)*	125.00	41.70

LESSONS LEARNED

- Federated learning (FL) models can be unfair
- FL models can be biased towards clients that have a larger number of samples
- Alternatives for FedAvg are being investigated to ensure fairness when using PPFL (Agnostic federated learning methods and biomedical adaptation)
- Adapting with the distribution shift among groups of clients
  - Balancing the number of training samples by adding more samples/ using data augmentation
  - Data regularization through normalization techniques

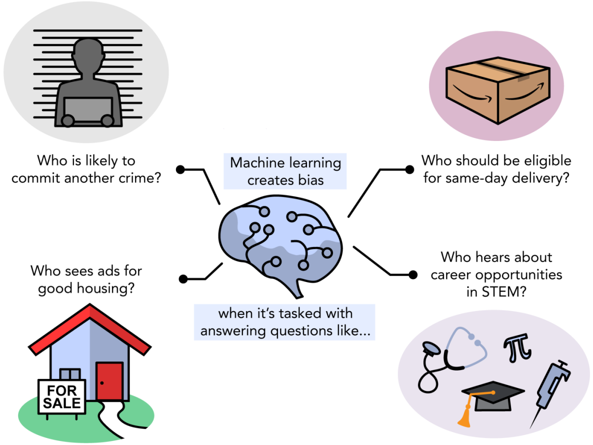
Who is likely to commit another crime?
Machine learning creates bias
Who should be eligible for same-day delivery?
Who sees ads for good housing?
when it's tasked with answering questions like...
Who hears about career opportunities in STEM?
FOR SALE
π

Examples of how bias in machine learning can affect our daily lives.

CONCLUSION

Privacy Preserving Analysis and Learning in Secure and Distributed Enclaves and Exascale systems (PALISADE-X)

- Dataset Shift challenge in AI are real
  - Models don’t do well when applying in settings different from settings and data used in training
  - Bigger challenge in Biomedicine where data is not shared because of policy issues
- We presented PALISADE-X where we
  - Developed APPFL (Argonne Privacy-Preserving Federated Learning) framework that implements end-to-end secure framework that leverages *differential privacy algorithms* along with capabilities to leverage heterogenous HPC resources easily
  - We discussed how we integrated APPFL framework with our existing computing and data infrastructure (i.e., ABLE, SEAL, HPCrypt, funcX, and DLHub) with focus on validating and evaluating APPFL framework by using the multi-institutional biomedical datasets
- We presented results and lessons learned when applying APPFL to Biomedical datasets

Q&A

https://github.com/APPFL/APPFL/tree/thoang/funcx