Private Federated Learning

Jose-Luis Ambite, PhD
Research Team Leader, Information Sciences Institute
Associate Research Professor, Computer Science
University of Southern California

Can we learn from distributed data respecting privacy?

Learn a model over all the data without sharing data?

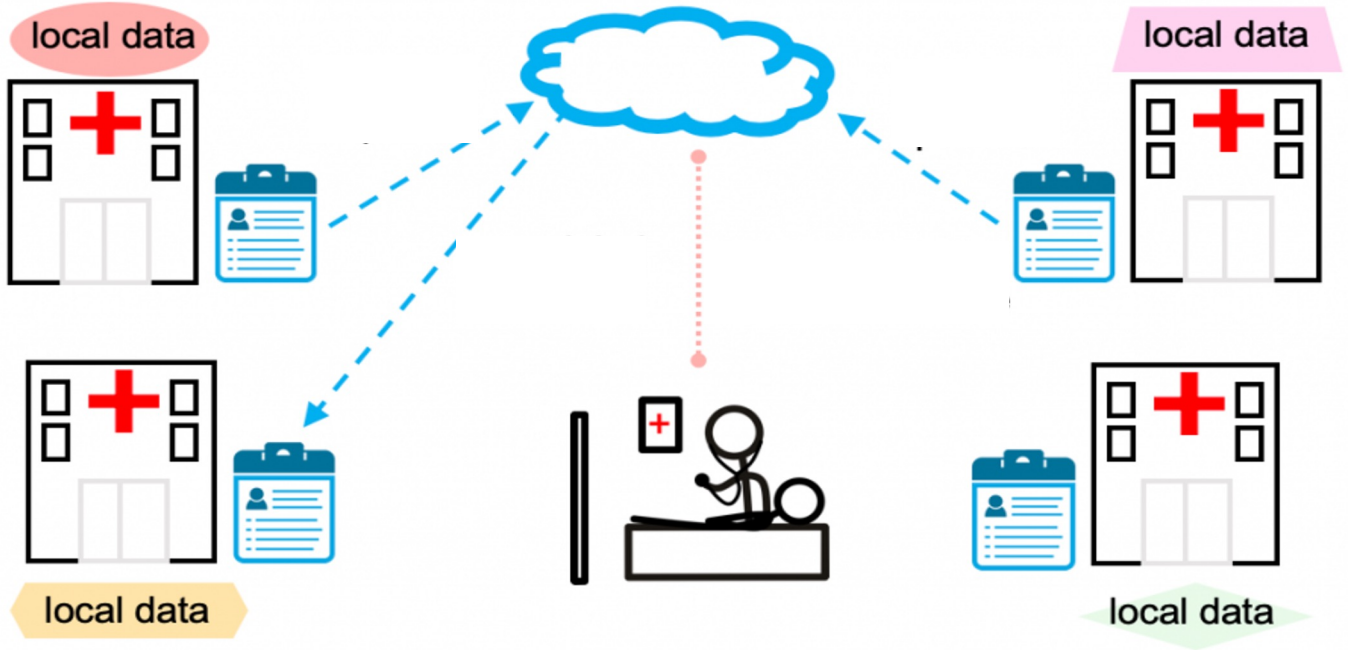

image courtesy: https://blog.ml.cmu.edu/2019/11/12/federated-learning-challenges-methods-and-future-directions/

Federated Learning
Learning without sharing data!

- No data shared among data sources
- Training is pushed down to data sources
- Sources share parameters (e.g., gradients)

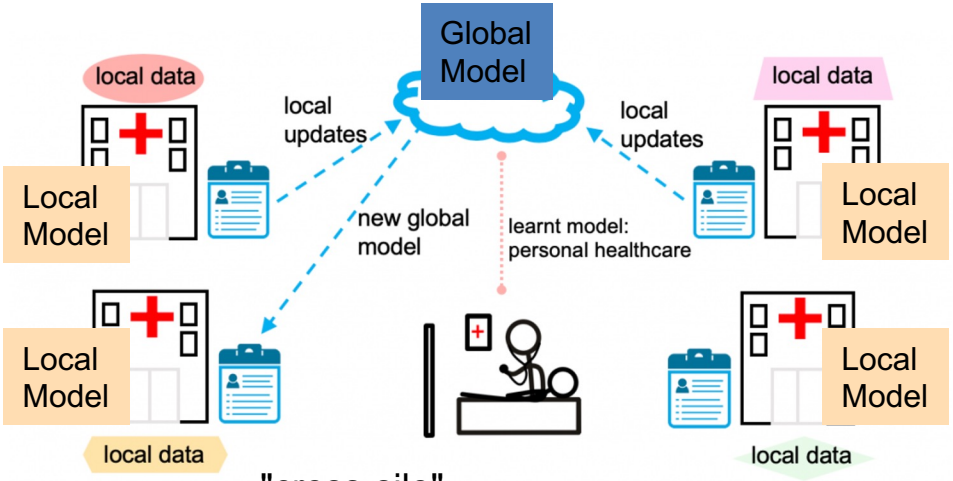

"cross-silo"

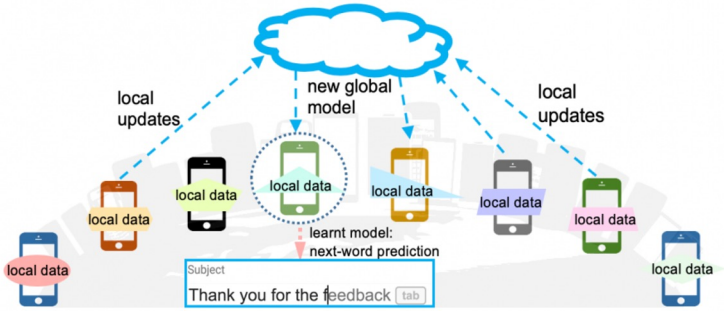

"cross-device"

Challenges

- Efficient learning
- Privacy/Security
- Heterogenous environments in
  - Computational power
  - Data distributions (size, target classes)

Private Federated Learning

Learn a Neural Network without sharing data

Encrypted Computation

Encrypted Transmission (of model parameters)

Data never leaves sites

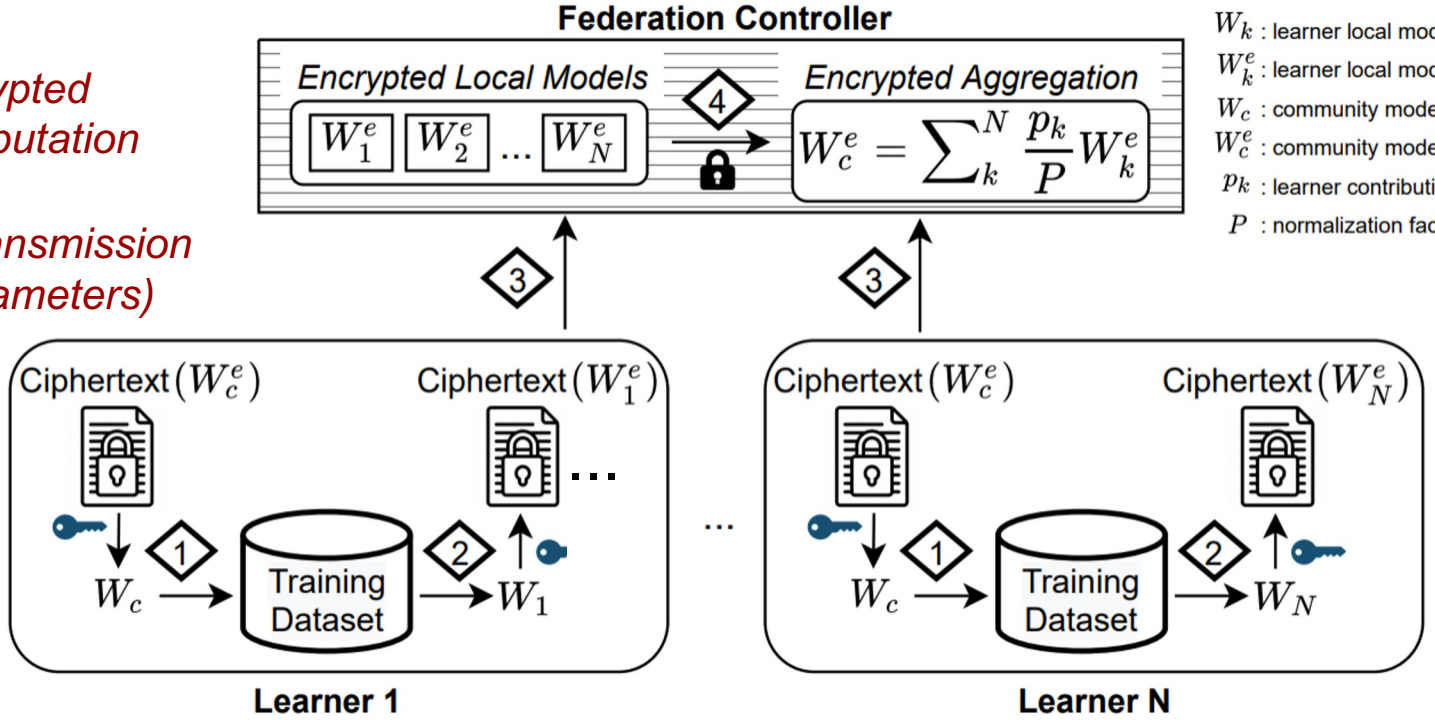

Federated Model Outperforms any Local Model: Uniform Data Amount, IID

All sites incentivized to join!

Federated Learning Environment

- **8 sites**
- **Domain:** BrainAge, predict age from MRIs
- **Data Amount:** Uniform, equal number of samples per site
- **Data Distribution:** IID
- **Model:** 5-CNN

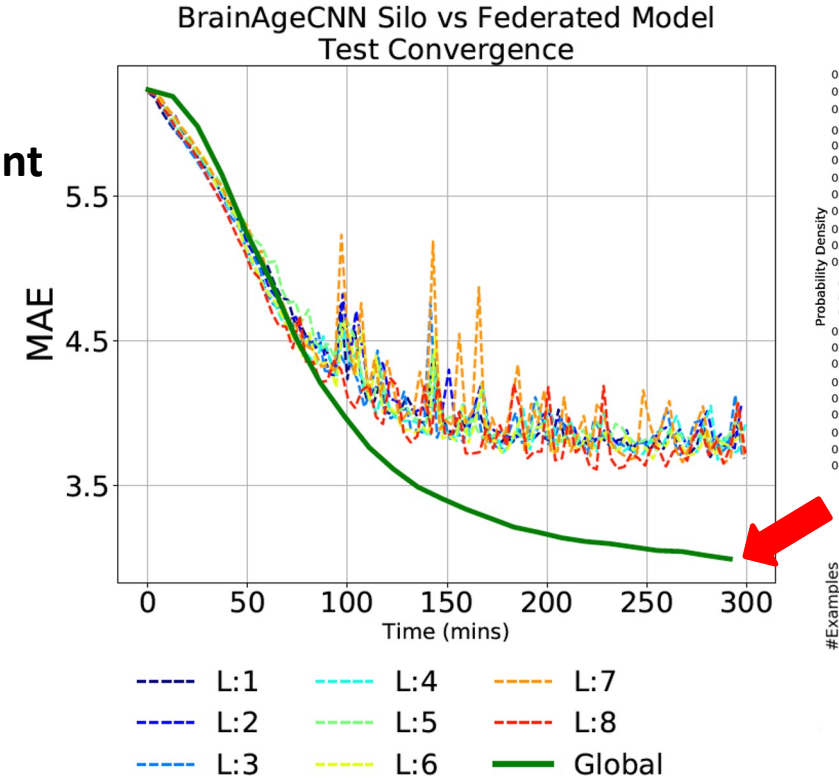

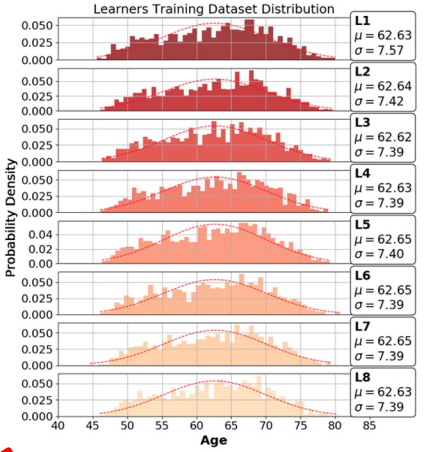

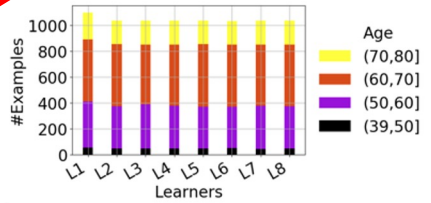

Federated Model Outperforms any Local Model: Skewed Data Amount, Non-IID

All sites incentivized to join!

Federated Learning Environment

- **8 sites**
- **Domain:** BrainAge, predict age from MRIs
- **Data Amount:** Skewed, different number of samples per site
- **Data Distribution:** Non-IID
- **Model:** 5-CNN

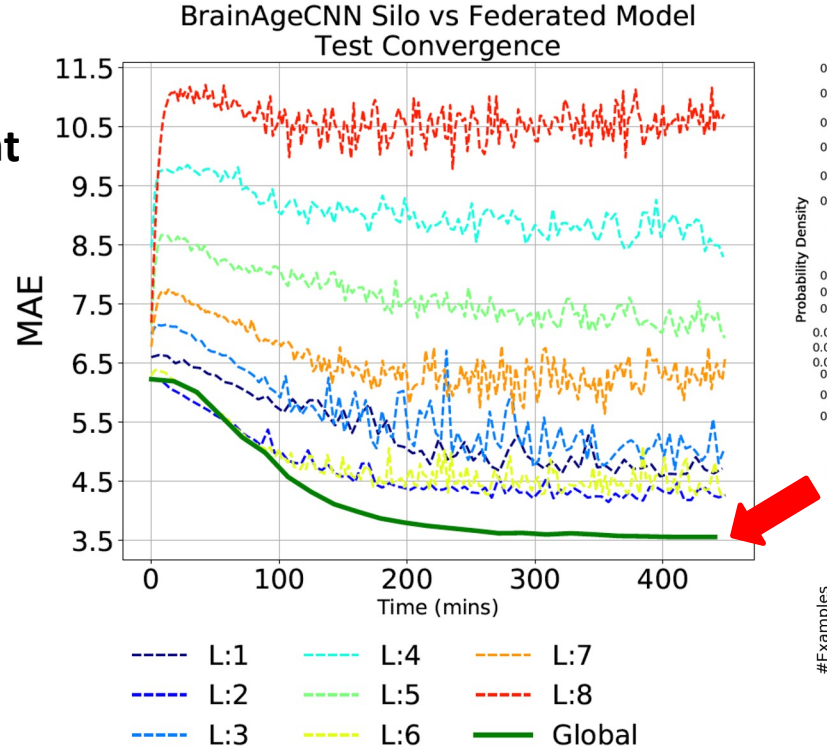

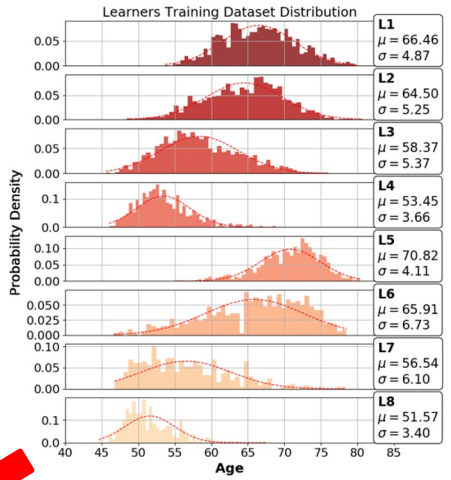

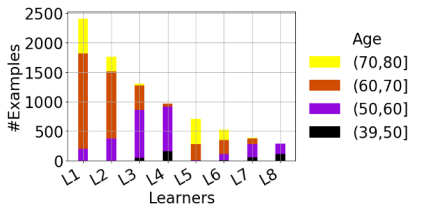

Federated Training Policies on Benchmark domains

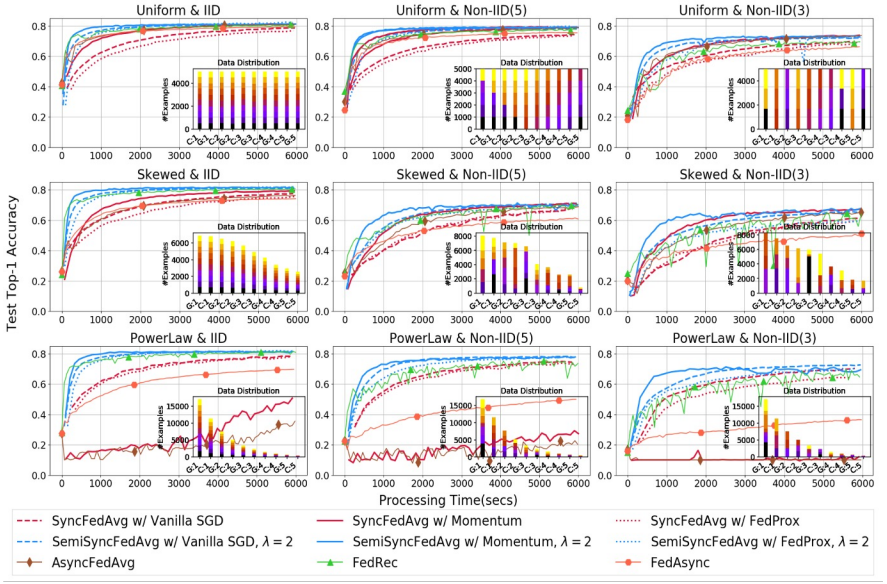

CIFAR-10

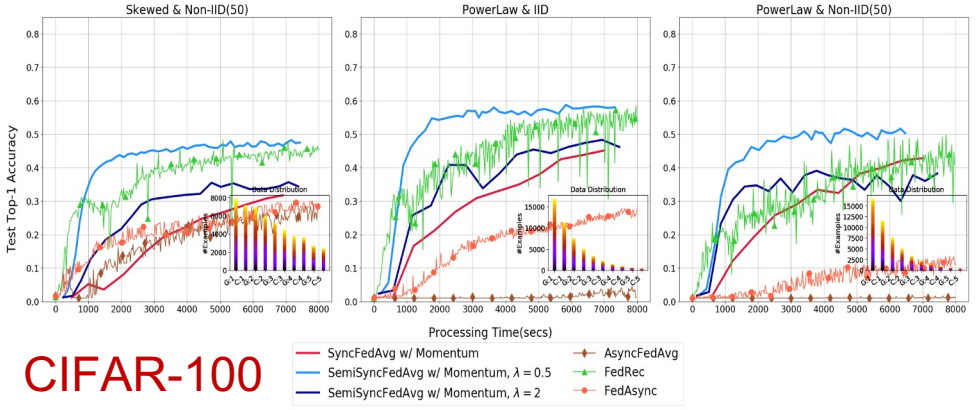

CIFAR-100

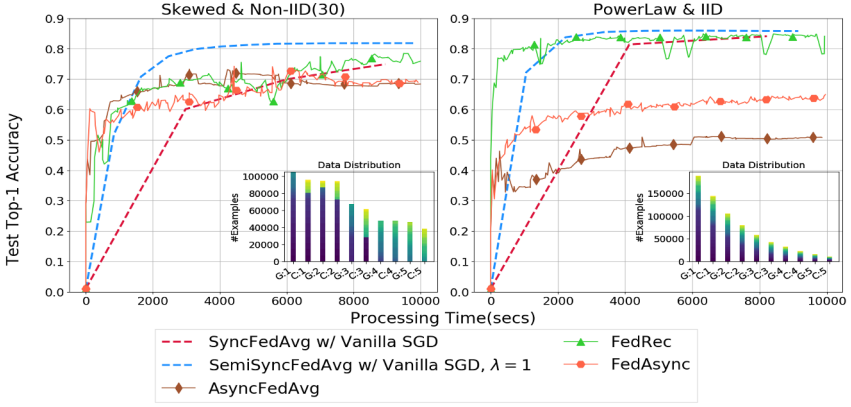

Extended MNIST (62 classes)

Predicting **Brain Age** from MRI Scans

- **Dataset:** UK Biobank
- **Data Distribution:** Non-IID (diverse age ranges across learners)
- **Computational Environment:** 8xGPU
- **Model:** 5-CNN

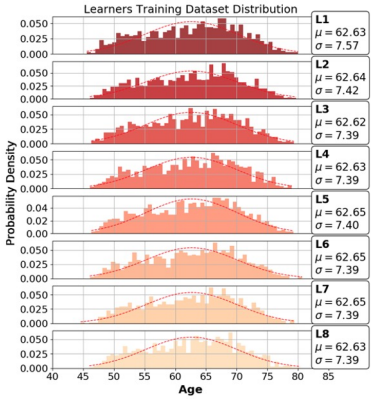

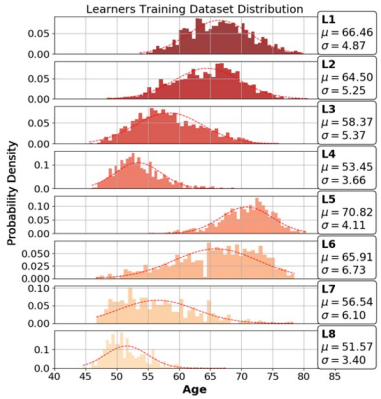

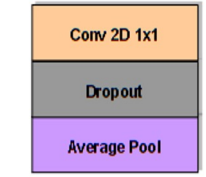

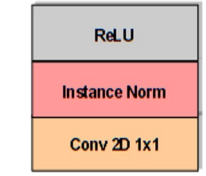

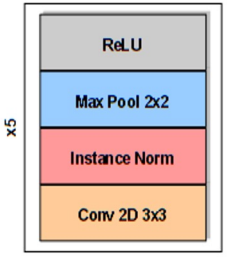

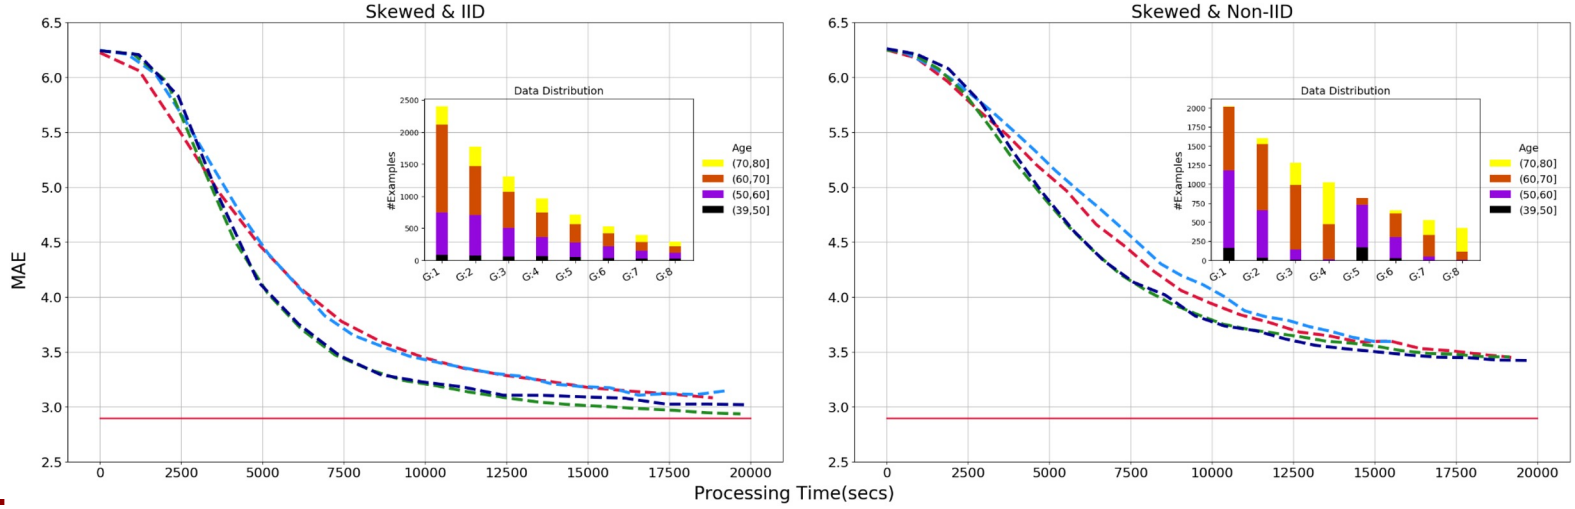

Predicting **Alzheimer's Disease** from MRIs

Cohort	Training Set		Test Set	
	Alzheimer's	Controls	Alzheimer's	Controls
ADNI1	1,313	1,832	324	458
ADNI2	384	457	95	115
ADNI3	51	232	13	59
OASIS	36	150	10	38
AIBL	112	641	29	161
Total	*1,896*	*3,312*	*471*	*831*

TABLE I: AD: Train/test splits per cohort and target label.

- **Federated similar performance to Centralized**
- 3D-CNN Neural Model
- Pretraining helps
  (sex prediction from MRI on UK Biobank)

Model	Accuracy	Precision	Recall	F1	AUC PR	AUC ROC
ADNI only	0.8570 ± 0.0090	0.7940 ± 0.0311	0.8270 ± 0.0288	0.8095 ± 0.0080	0.8639 ± 0.0052	0.8954 ± 0.0057
OASIS only	0.4428 ± 0.0194	0.3686 ± 0.0036	0.7447 ± 0.0518	0.4927 ± 0.0091	0.3396 ± 0.0020	0.4631 ± 0.0047
AIBL only	0.8050 ± 0.0044	0.7246 ± 0.0153	0.7577 ± 0.0172	0.7405 ± 0.0022	0.7990 ± 0.0005	0.8526 ± 0.0017
Centralized 5AOB	0.8612 ± 0.0106	0.7977 ± 0.0350	0.8287 ± 0.0271	0.8122 ± 0.0091	0.8683 ± 0.0130	0.8986 ± 0.0051
Federated 3A	0.8462 ± 0.0043	0.8148 ± 0.0189	0.8048 ± 0.0194	0.8095 ± 0.0038	0.8791 ± 0.0039	0.8967 ± 0.0012
Federated 4AO	0.8474 ± 0.0073	0.7955 ± 0.0126	0.8296 ± 0.0026	0.8121 ± 0.0077	0.8766 ± 0.0007	0.8920 ± 0.0014
Federated 5AOB	0.8633 ± 0.0013	0.8098 ± 0.0043	0.8132 ± 0.0097	0.8114 ± 0.0031	0.8682 ± 0.0009	0.8971 ± 0.0006

TABLE II: Alzheimer's Disease Prediction. Test results on a global stratified test dataset (5 sites), for each dataset by itself; 3 sites, ADNI1,2,3 (3A); 4 sites, ADNI1,2,3 + OASIS (4AO), and 5 sites, ADNI1,2,3 + OASIS + AIBL. In federated environments each dataset is at a different learner. Centralized environments are trained over all the corresponding datasets.

Same learning performance with and without encryption

CKKS Homomorphic Encryption (Optimized), low time overhead (~7%)

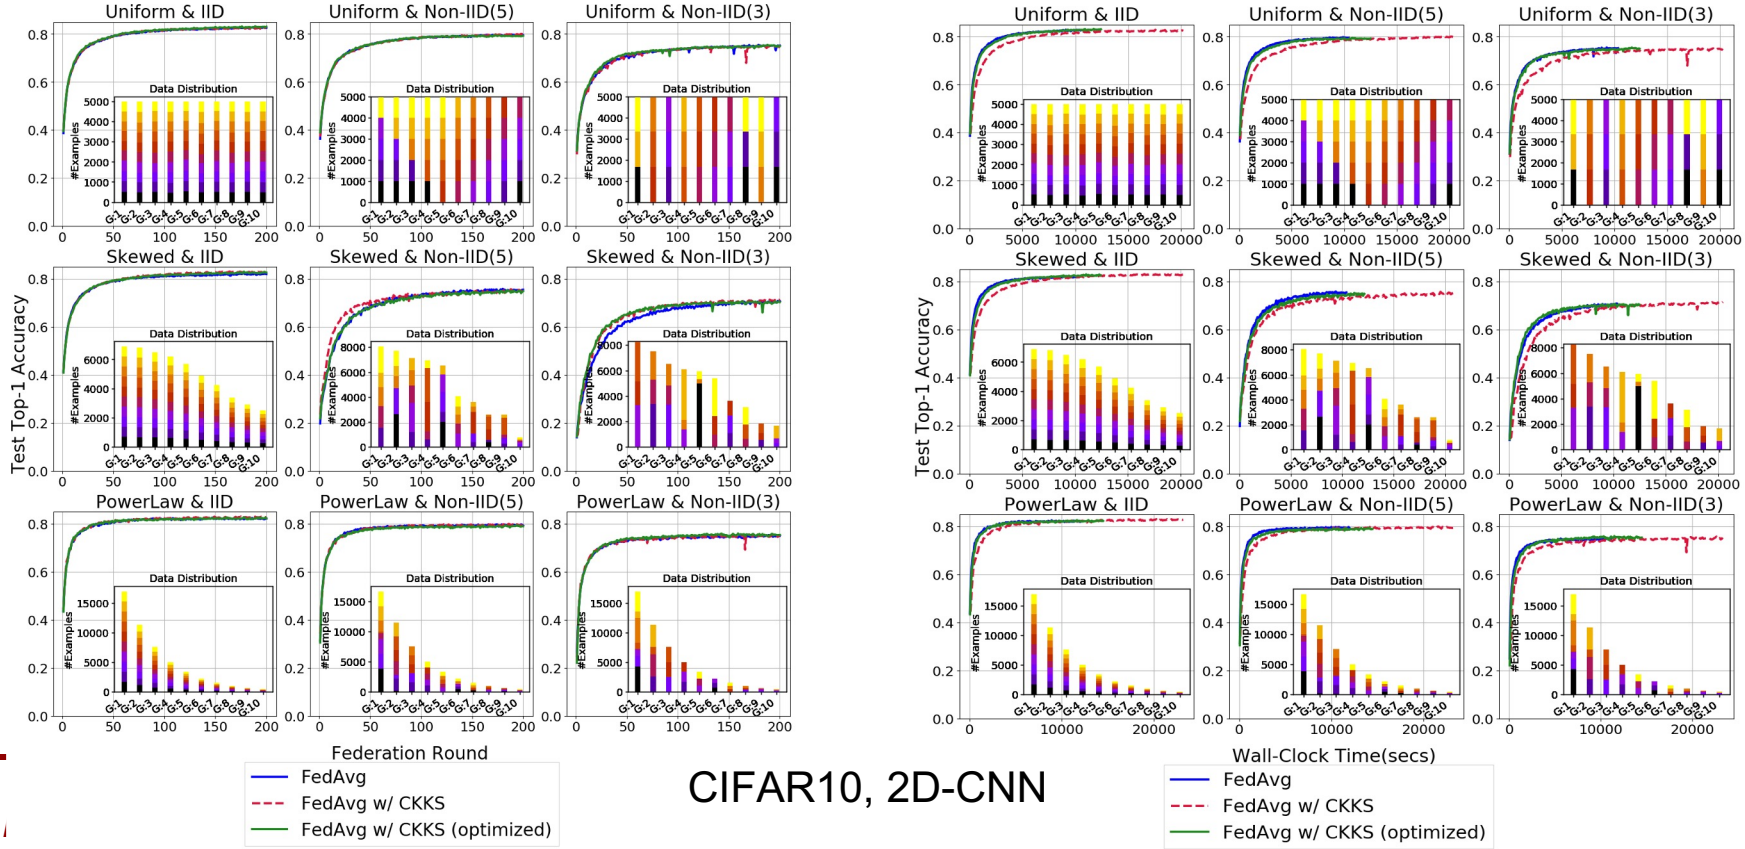

CIFAR10, 2D-CNN

Same learning performance with and without encryption

BrainAge: 3D-CNN Model

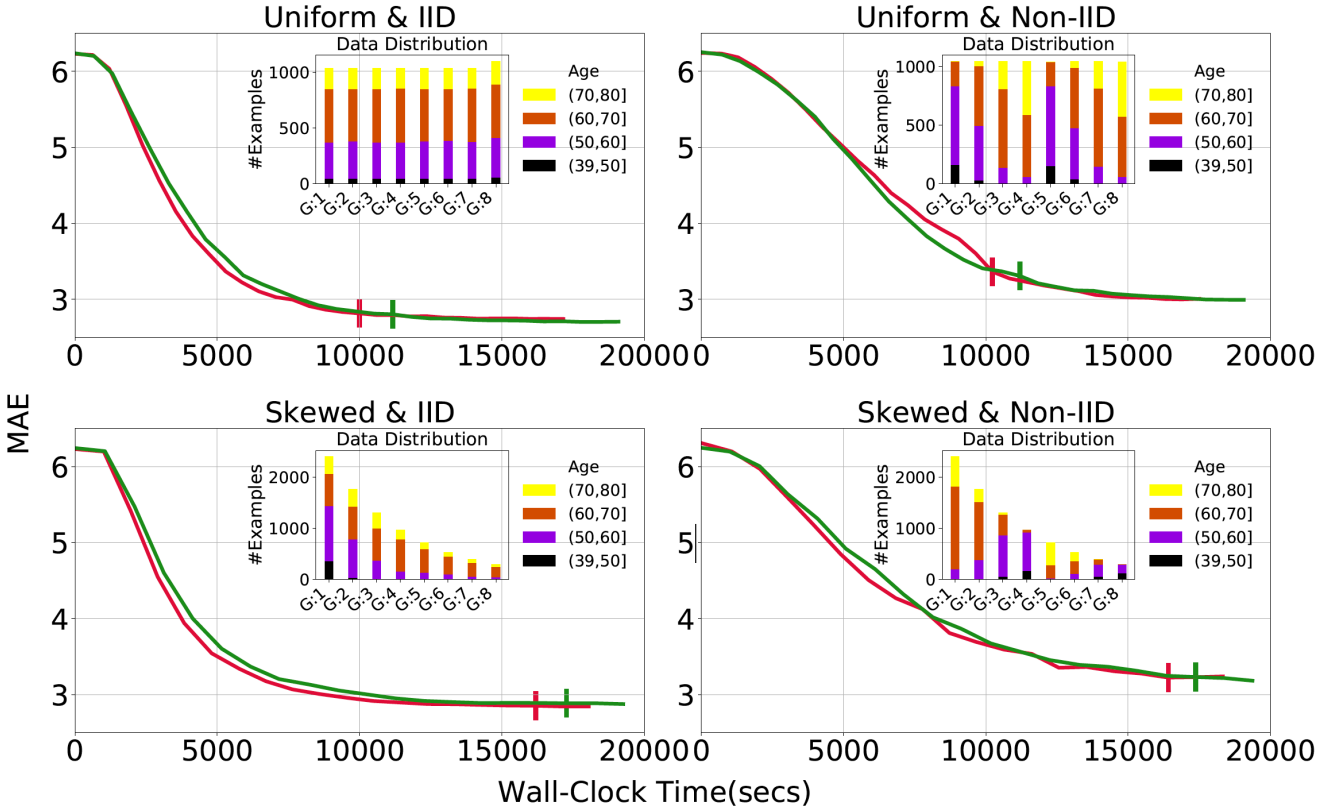

Private Federated Learning

Learn a Neural Network without sharing data

Encrypted Computation

Encrypted Transmission (of model parameters)

Data never leaves sites

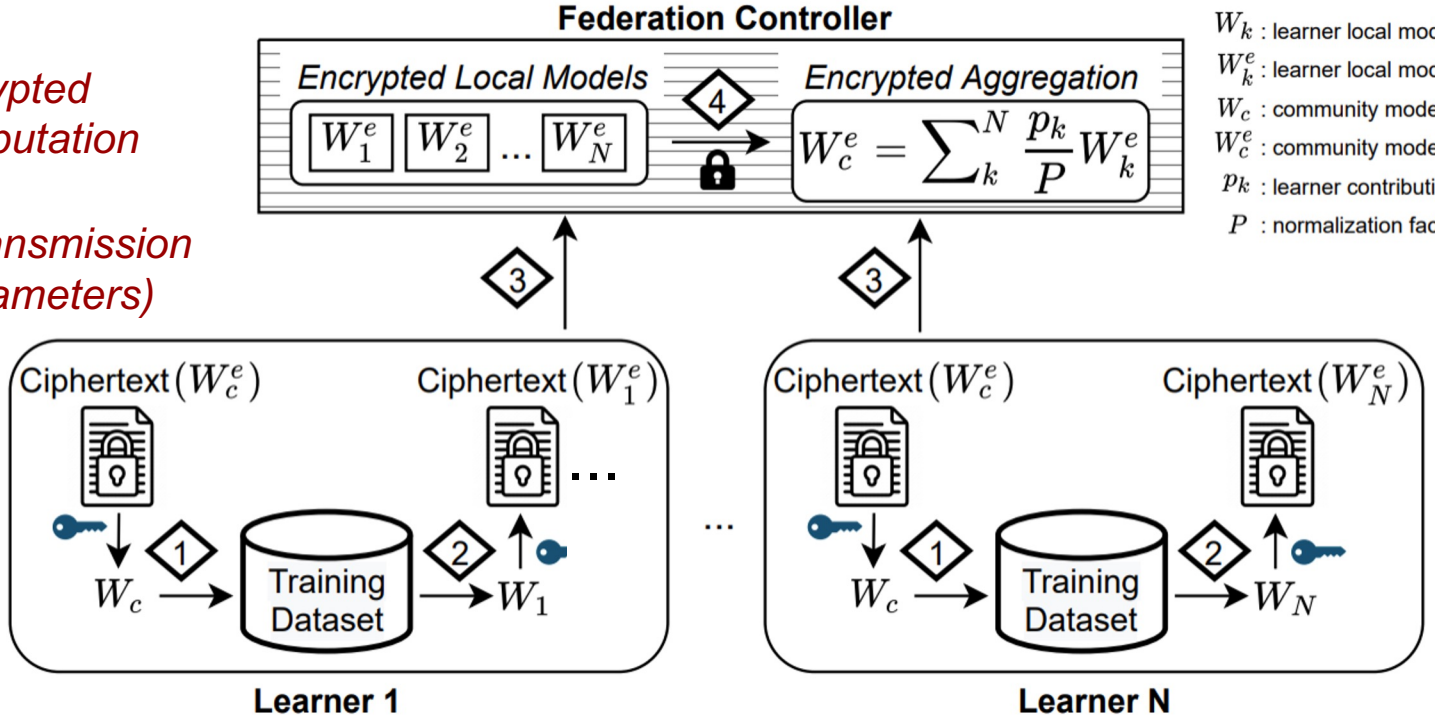

Membership Inference Attacks

- Since the federated controller aggregates models using homomorphic encryption and data is transmitted through secure channels, data is kept private at the controller and outside the federation.
- However, an *honest, but curious* site may attack the neural model:
  - *Model inversion*: data memorized in the model, usually not very successful since global model aggregates many local models
  - *Membership Inference Attack*: A site participating in a federation studying a sensitive topic (e.g., schizophrenia, HIV in the brain, …) may be curious if one of the subjects they have data on (e.g., an MRI), but not provided to the federation, participated in training that federated model at other sites. Unacceptable privacy violation.

Membership Inference Attacks on Unprotected Federated Models are very Successful

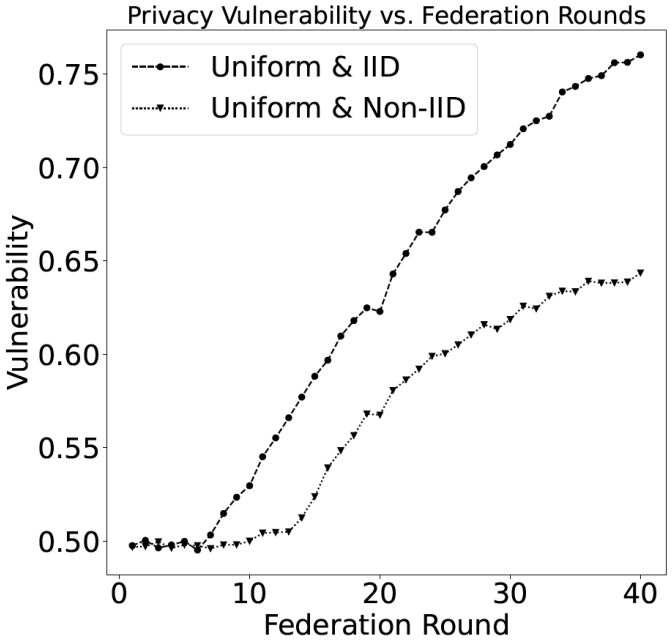

Fig. 6: Privacy vulnerability increases with federation rounds. Vulnerability is measured as the average accuracy of distinguishing train samples vs unseen samples across learners.

Defeating Membership Attacks

Membership attack success % (lower is better)

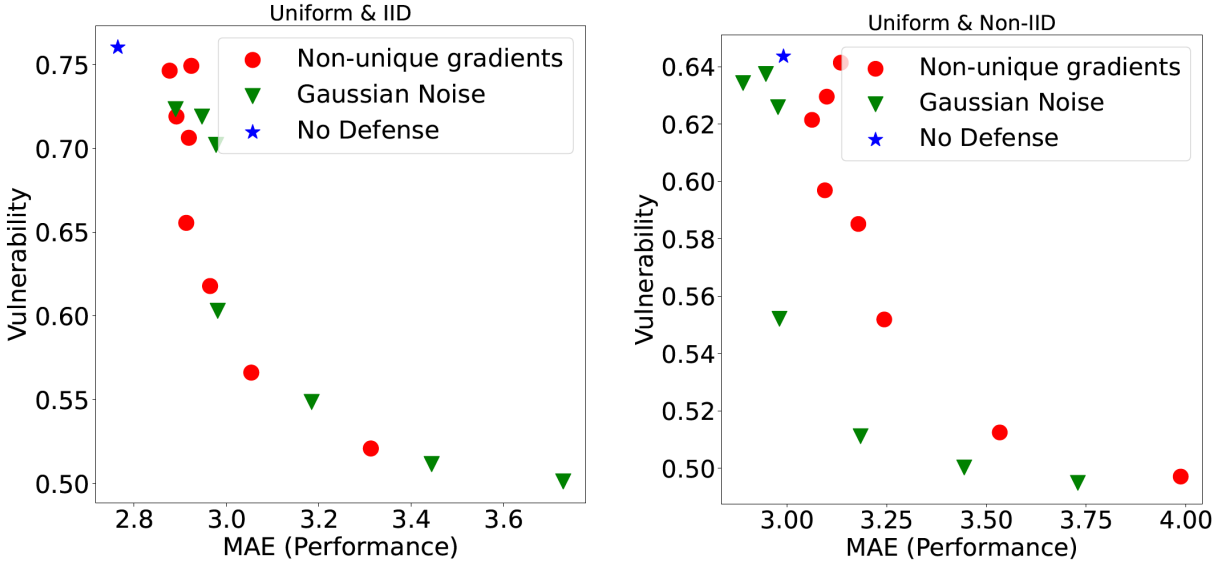

(a) Uniform & IID (b) Uniform & Non-IID

Fig. 7: Vulnerability vs Performance trade-off when training learners with Differential Privacy (Gaussian Noise) and Non-unique gradients approach. Lower vulnerability and lower MAE is desired, i.e., points towards the bottom left are better.

Private Federated Learning

Learn a Neural Network without sharing data

Encrypted Computation

Encrypted Transmission (of model parameters)

Data never leaves sites

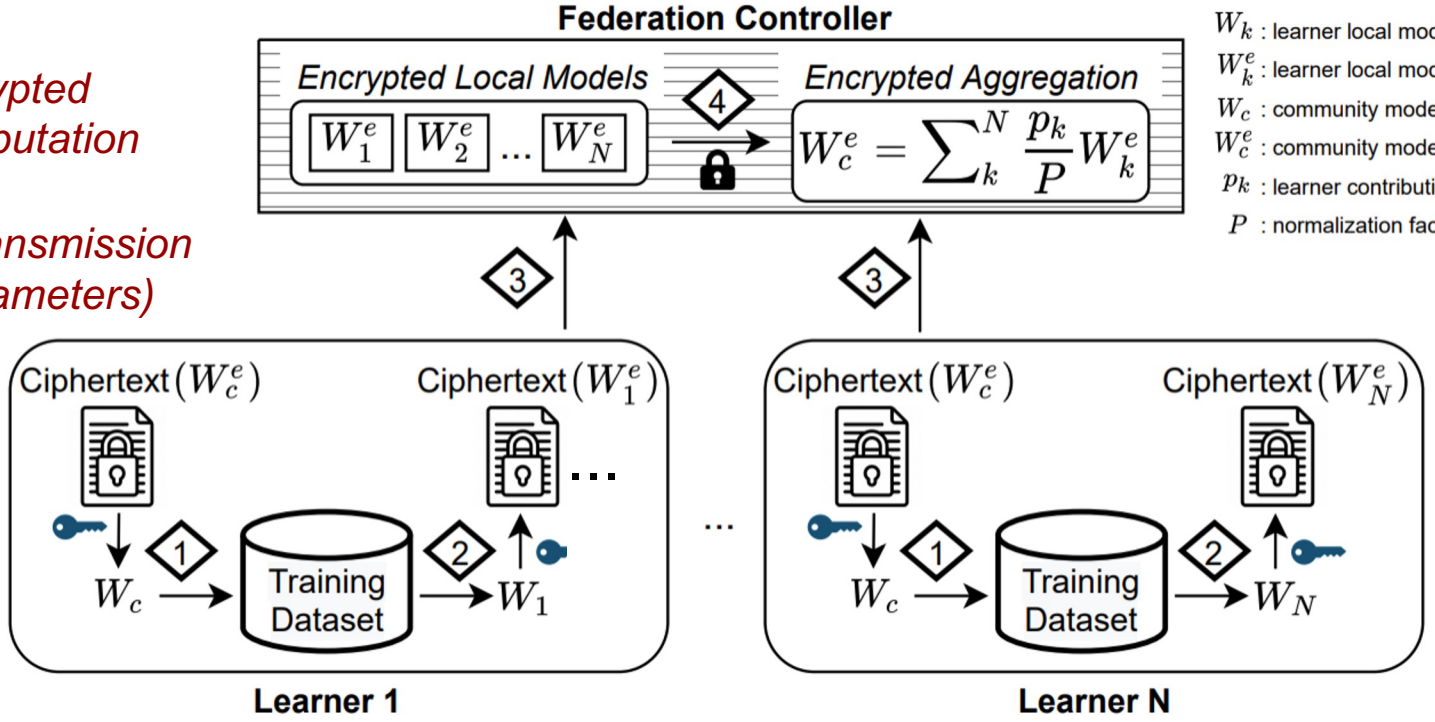

Private Federated Learning

- Easily deployable, containerized Private Federated Learning architecture
- Fast convergence, energy-efficient, heterogeneous domains
- Same learning performance with and without encryption, low overhead
- Enforcing privacy while maintaining learning performance
  - Limit information leakage via gradient noise and "unique gradient info"
- Biomedical domains:
  - Critical need for federated learning: private, distributed datasets
  - Great potential: clinical diagnosis, prognosis, drug development
  - Good predictive power for Alzheimer's Disease, working on Parkinson's Disease
    - disease classification, predicting cognitive scores, cognitive decline, prognosis from MRI+
- Ongoing: model compression/sparsification

Publications

- Dimitris Stripelis, Paul Thompson, and Jose Luis Ambite. **Semi-Synchronous Federated Learning for Energy-Efficient Training and Accelerated Convergence in Cross-Silo Settings**. *ACM Transactions on Intelligent Systems and Technology, Special Issue in Federated Learning*. Forthcoming 2022.
- Dimitris Stripelis, Hamza Saleem, Tanmay Ghai, Nikhil Dhinagar, Umang Gupta, Chrysovalantis Anastasiou, Greg Ver Steeg, Srivatsan Ravi, Muhammad Naveed, Paul M.Thompson, and Jose Luis Ambite. **Secure neuroimaging analysis using federated learning with homomorphic encryption**. *In 17th International Symposium on Medical Information Processing and Analysis (SIPAIM),* Campinas, Brazil, 2021.
- Nikhil J. Dhinagar, Sophia I. Thomopoulos, Conor Owens-Walton, Dimitris Stripelis, Jose Luis Ambite, Greg Ver Steeg, Daniel Weintraub, Philip Cook, Corey McMillan, and Paul M. Thompson. **3D Convolutional Neural Networks for Classification of Alzheimer's and Parkinson's Disease with T1-weighted brain MRI**. In *17th International Symposium on Medical Information Processing and Analysis (SIPAIM)*, Campinas, Brazil, 2021
- Umang Gupta, Dimitris Stripelis, Pradeep K. Lam, Paul M. Thompson, Jose Luis Ambite, and Greg Ver Steeg. **Membership inference attacks on deep regression models for neuroimaging**. In *Medical Imaging with Deep Learning (MIDL)*, Zürich, Switzerland, 2021.
- Dimitris Stripelis, Jose Luis Ambite, Pradeep Lam, and Paul Thompson. **Scaling neuroscience research using federated learning**. In *IEEE International Symposium on Biomedical Imaging,* Nice, France, 2021.

Team

Jose-Luis Ambite, PhD
Research Team Leader, ISI
Associate Res. Prof., CS

Muhammad Naveed, PhD
Assistant Prof., CS

Srivatsan Ravi, PhD
Research Scientist, ISI
Assistant Res. Prof., CS

Greg Ver Steeg, PhD
Research Team Leader, ISI
Associate Res. Prof., CS

Paul Thompson, PhD
Professor, Neurology, Psychiatry, Radiology, Ophthalmology, and Engineering

Dimitris Stripelis, MS
PhD Student,
Computer Science

Rafa Sanchez, BS
MS Student,
Computer Science

Hamza Saleem, BS
PhD Student,
Computer Science

Tanmay Ghai, BS
MS Student,
Computer Science

Umang Gupta, MS
PhD Student,
Computer Science

Nikhil Dhinagar, PhD
Analyst,
Imaging Genetics Center

\+ **Armaghan Asghar,** MS CS ; **Oneeb Khan**, MS CS ; **Chrysovalantis Anastasiou**, PhD CS

Pediatric Research using Integrated Sensor Monitoring Systems (PRISMS) Data and Software Coordination and Integration Center (DSCIC)

- Big data integration & analytics
- Novel statistical/ML methods
- prisms-study.org

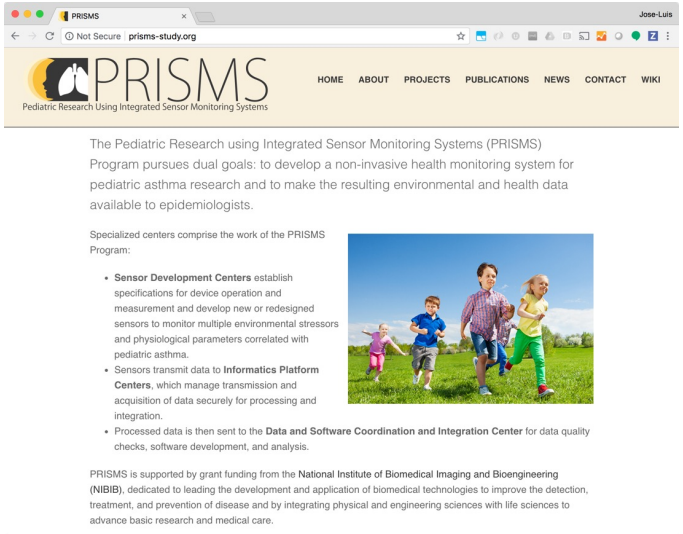

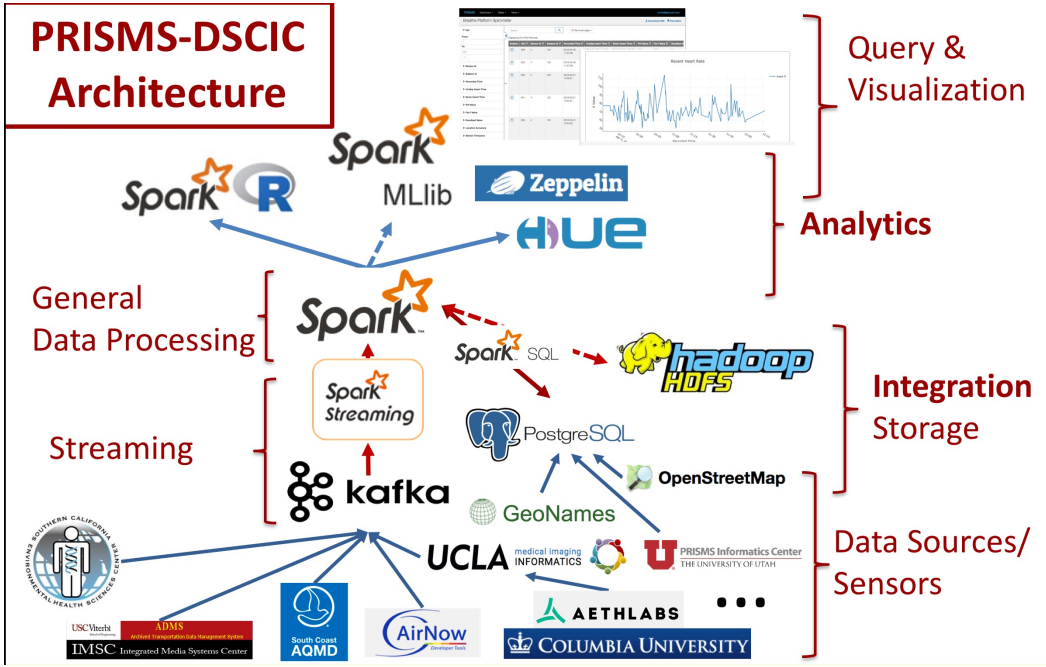

NIMH REPOSITORY & GENOMICS RESOURCE

Sign In | Request Access | Help | Search

EXPLORE OUR DATA | DOWNLOAD DATA | ORDER BIOSAMPLES | SUBMIT YOUR DATA | LEARN ABOUT US | RESOURCES

Comprehensive Data + Tools to Support Your Research

Explore our phenotypic & genetic data for a particular mental disorder or explore our stem cell data.

Choose Collection by Disorder | Stem Cell Data

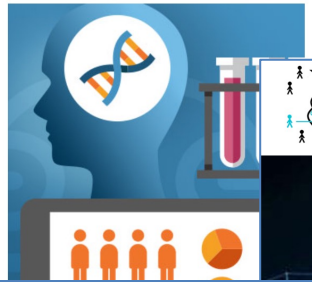

nature

Explore content | About the journal | Publish with us

nature > letters > article

Letter | Published: 19 June 2019

Genetic analyses of diverse populations improves discovery for complex traits

Genevieve L. Wojcik, Mariaelisa Graff, ... Christopher S. Carlson + Show authors

Nature **570**, 514–518 (2019) | Cite this article

Home | Studies | Multi-ethnic Genotyping Array | PAGE Investigators | Publications | Contact

The PAGE Study

Population Architecture using Genomics and Epidemiology

nature communications

Explore content | About the journal | Publish with us

nature > nature communications > articles > article

Article | Open Access | Published: 10 June 2021

Rapid detection of identity-by-descent tracts for mega-scale datasets

Ruhollah Shemirani, Gillian M. Belbin, Christy L. Avery, Eimear E. Kenny, Christopher R. Gignoux & José Luis Ambite

Nature Communications **12**, Article number: 3546 (2021) | Cite this article

Although many genome-wide association studies (GWAS), whole-genome sequencing (WGS), omics, and polygenic risk score (PRS) efforts are underway, there is still a notable gap in leveraging diversity to empower discovery and improve our understanding of genotypic and phenotypic architecture across all populations. The Population Architecture through Genomics and Environment (PAGE) Study aims to help fill this gap. Our recent flagship paper (Wojcik et al., 2019) shows how much more can be learned by including people from multipe ancestries.