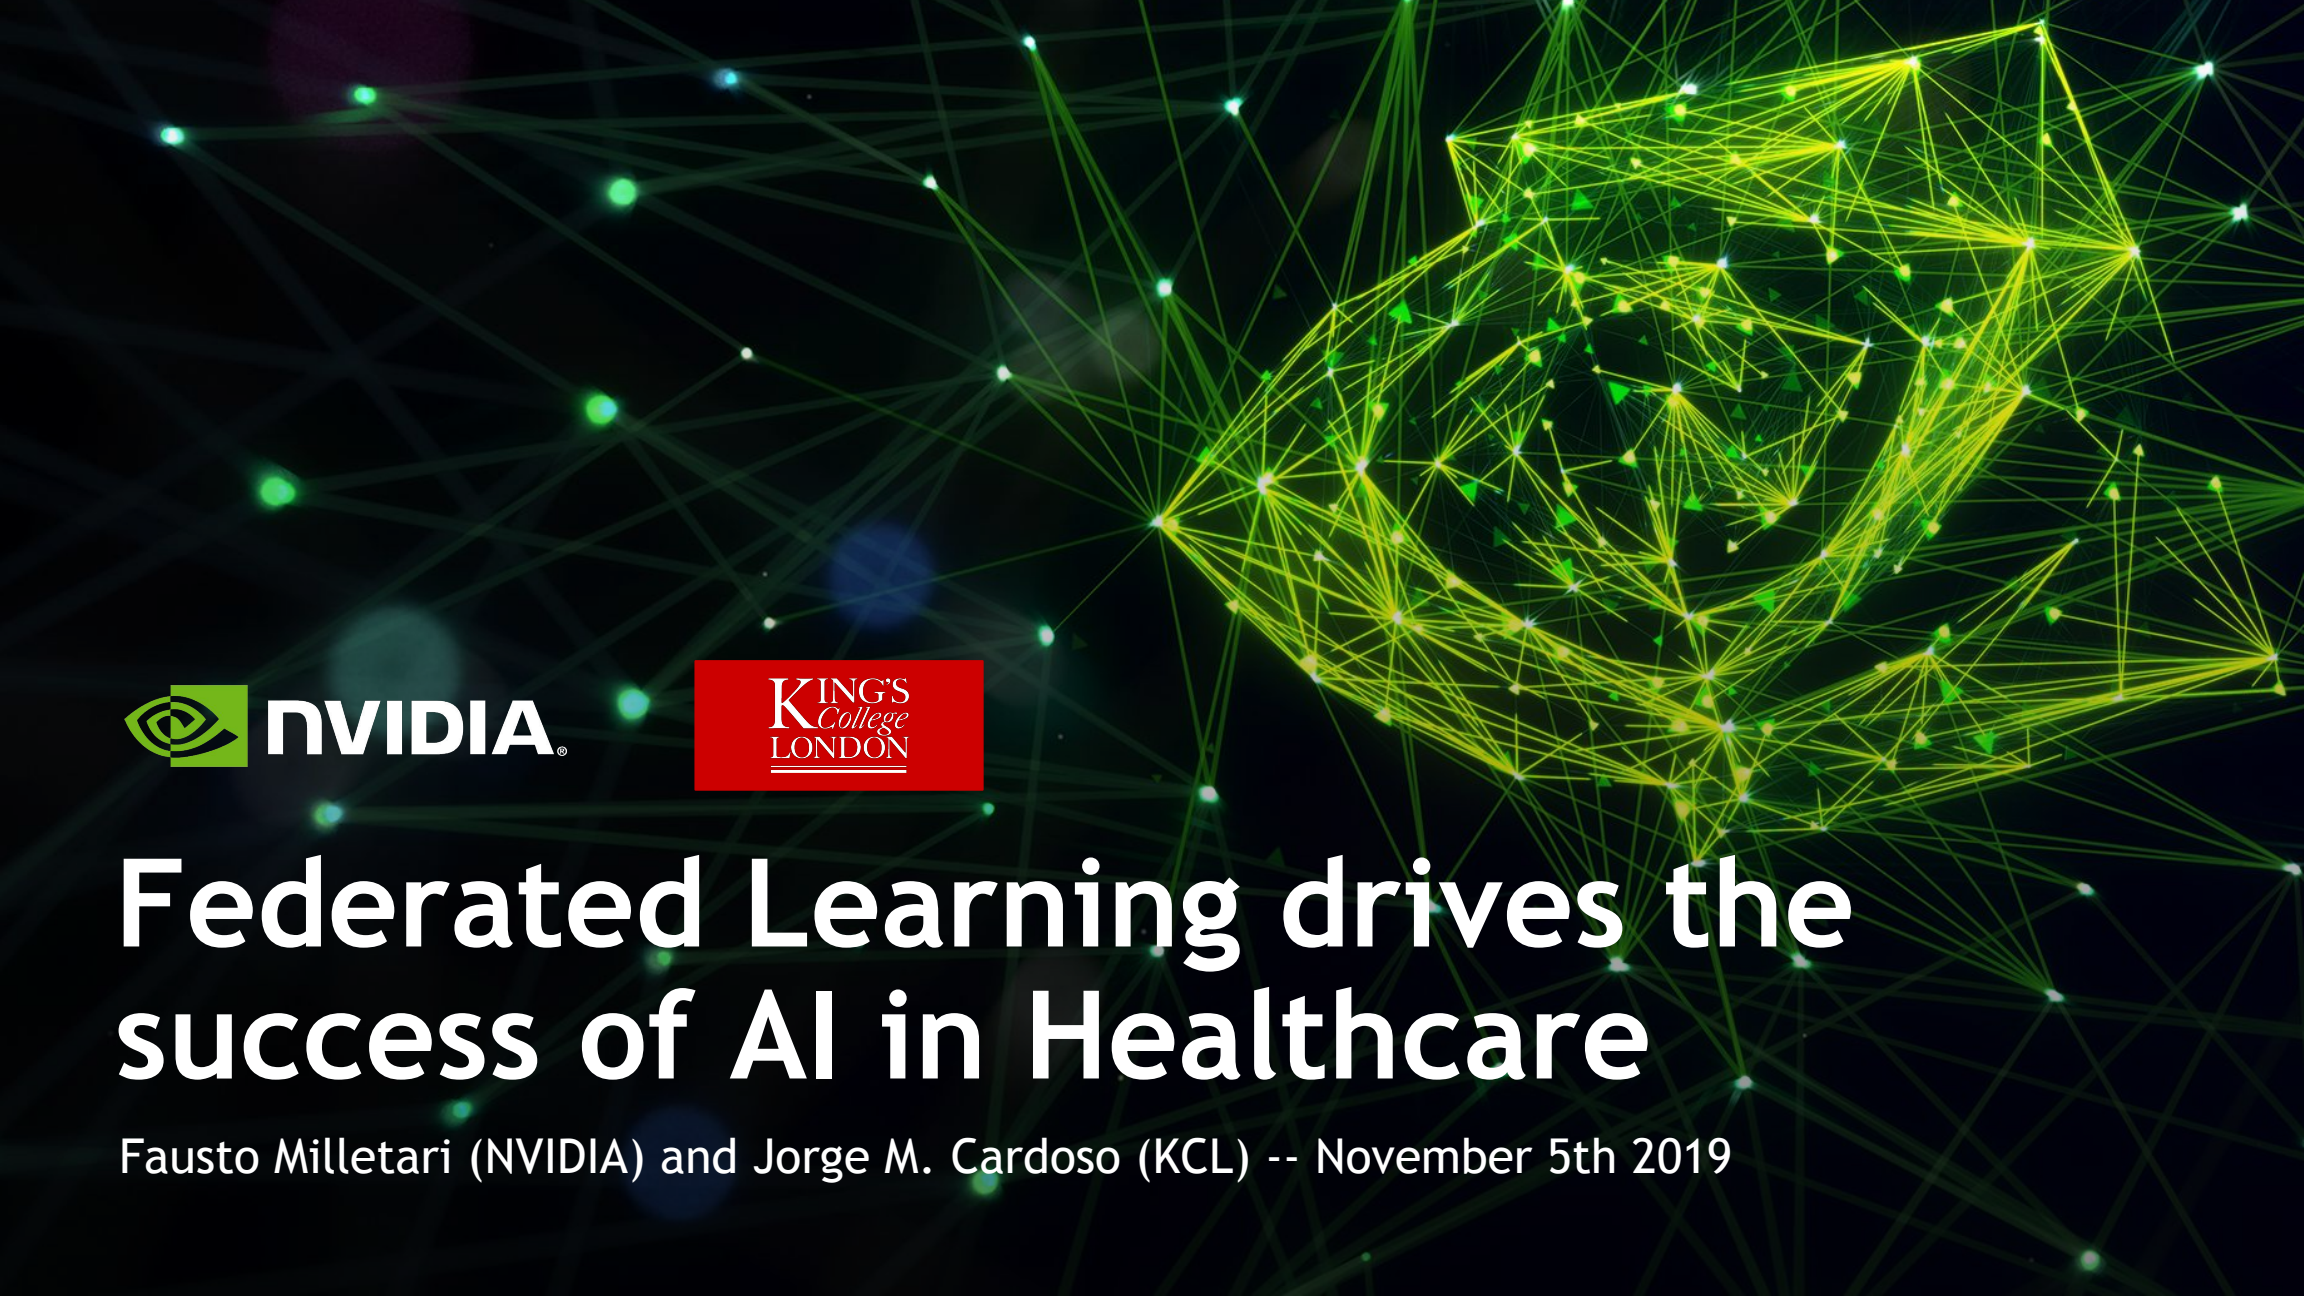

Federated Learning drives the success of AI in Healthcare

Fausto Milletari (NVIDIA) and Jorge M. Cardoso (KCL) -- November 5th 2019

“Up to three-to-five years ago all that [medical] data was just sitting there. Now it is being analyzed and interpreted. It is the most radical change happening in healthcare.”

-- Eric Topol

AI IN MEDICINE

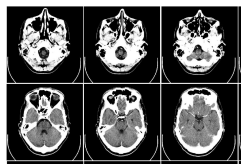
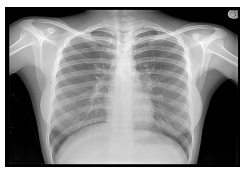
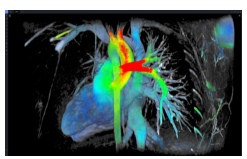


RADIOLOGY
CT, MR, US, X-RAY

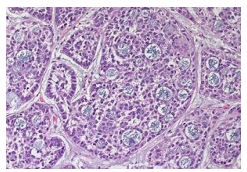
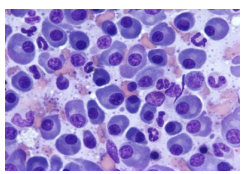

PATHOLOGY
TISSUE & CELL

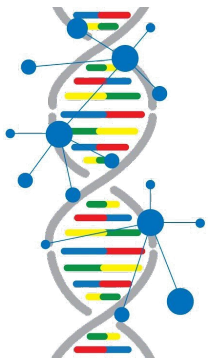

GENETICS

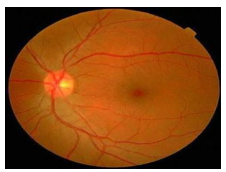


DERMATOLOGY
OPHTHALMOLOGY

• • •

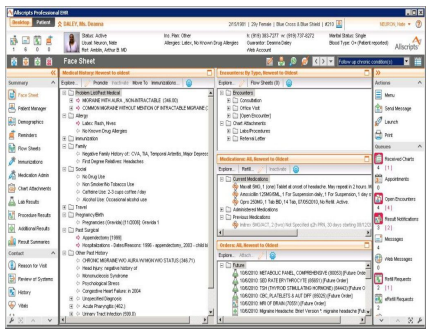

ELECTRONIC HEALTH
RECORDS

27K Medical AI papers - ~30 FDA Approved products
~7 Billion USD investment by 2021

FROM RESEARCH...

Improving state of the art performance in controlled settings

Convolutional Networks

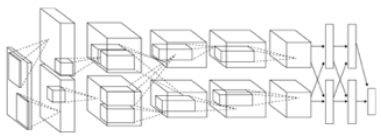

Encoder/Decoder

ReLu

BatchNorm

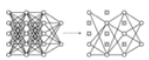

Concat Dropout Pooling

Recurrent Networks

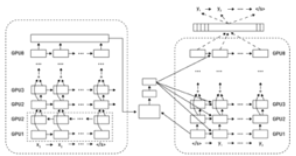

LSTM

GRU

Beam Search

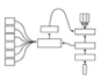

WaveNet CTC Attention

GAN

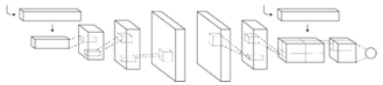

3D-GAN

MedGAN

Conditional GAN

Coupled GAN Speech Enhancement GAN

Reinforcement Learning

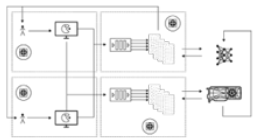

DQN

Simulation

DDPG

New Species

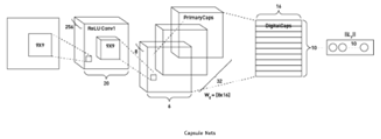

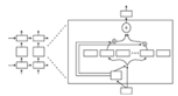

Mixture of Experts

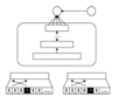

Neural Collaborative Filtering

Block Sparse LSTM

...TO APPLICATIONS

Achieving human level performance on large data and clinical settings

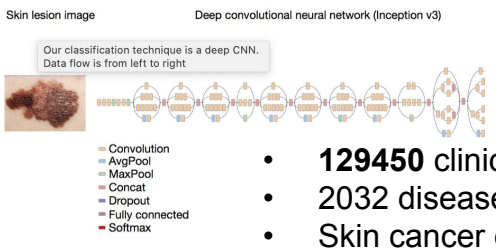

- **129450** clinical images
- 2032 diseases
- Skin cancer detection
- comparable to dermatologists

Esteva et al, Dermatologist-level classification of skin cancer with deep neural networks, Nature 2017

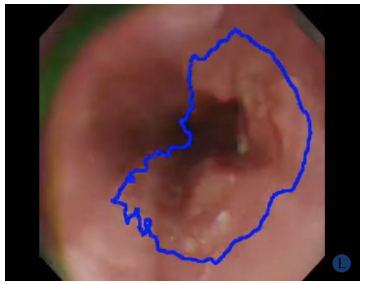

- 6 hospitals in China
- 84424 individuals
- **1036496** endoscopy images
- Gastrointestinal cancer detection
- Perf. similar to the expert endoscopist

Luo et al., Real-time artificial intelligence for detection of upper gastrointestinal cancer by endoscopy: a multicentre, case-control, diagnostic study, Lancet Oncology 2019

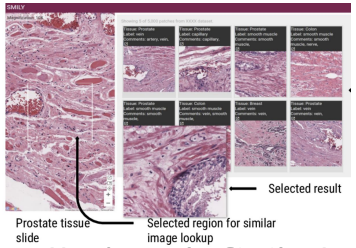

- **127000** image patches
- 128,000,000 8 × 8 μm regions
- Histopathology image search

Hegde et al., Similar Image Search for Histopathology: SMILY, Nature digital medicine 2019

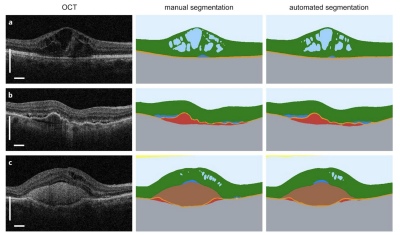

- "Only" **14884** OCT 3D scans
- Resolution ~ 5 μm
- Volumetric multi-region segmentation
- Performance comparable to humans

De Fauw et al., Clinically applicable deep learning for diagnosis and referral in retinal disease, Nature Medicine 2018

HOW FAR WOULD WE BE ABLE TO GO IF WE HAD DATA FROM 10 MILLION PATIENTS?

No more human errors in medicine?

Best outcome for each patient?

Affordable care?

No bias due to small demographics?

Automated diagnosis for most diseases?

Instantaneous diagnoses and wide screenings?

Ability to deal with rare diseases?

“Deep learning solutions can result in annual savings worth over USD 35 billion in the (medical) diagnostics segment alone by 2035”

-- Roots Analysis Business Research and Consulting

LARGE SCALE TRAINING

Creating large, centralized medical dataset is challenging

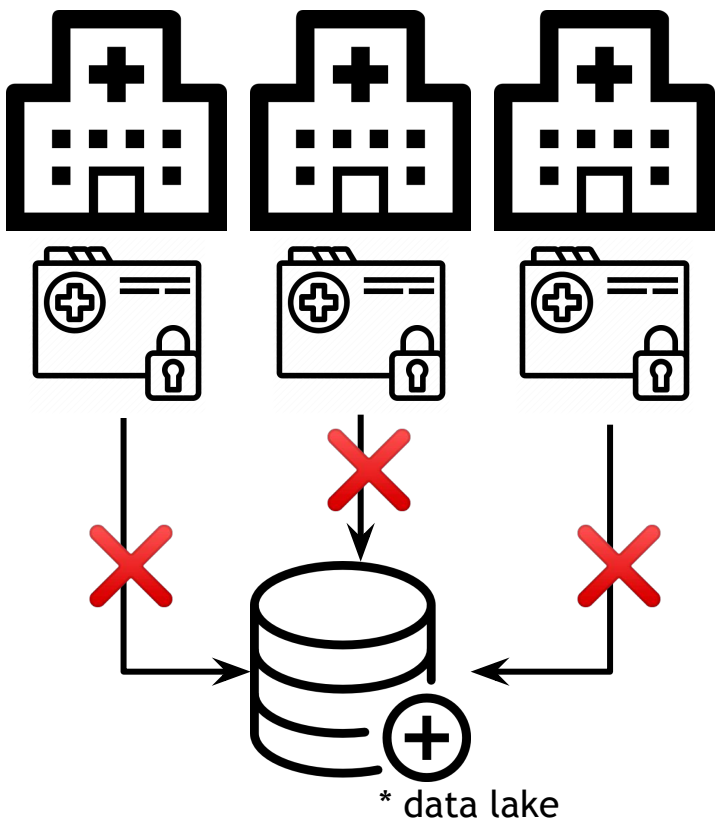

- Private data can’t be shared
- Anonymization is not truly effective
- Data annotation is costly. Data is an asset
- Bureaucracy of data sharing is complex

Wachinger, Christian, et al. "BrainPrint: A discriminative characterization of brain morphology." *NeuroImage* 109 (2015): 232-248.

LARGE SCALE TRAINING

Collaborative learning solves crucial issues in healthcare

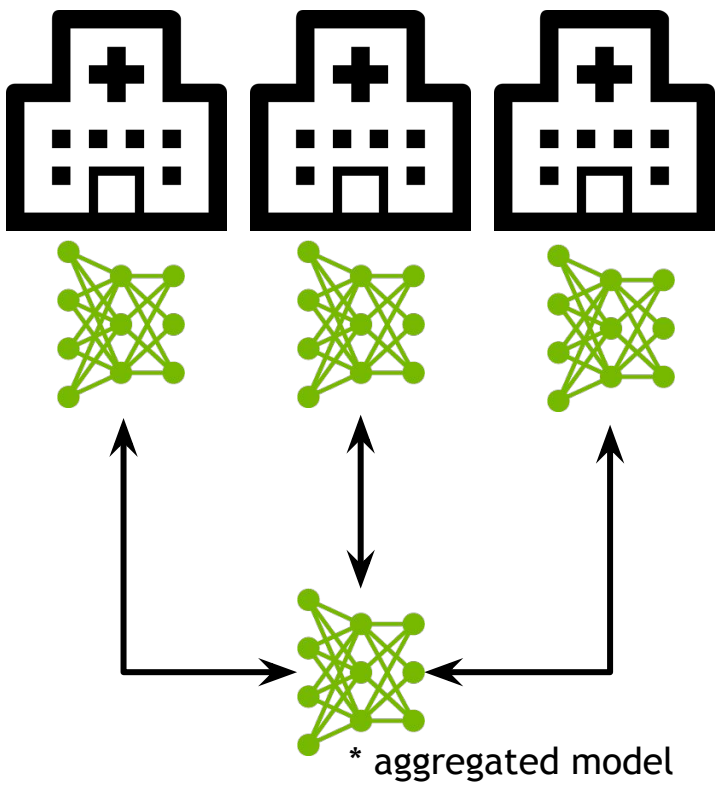

- Private data can’t be shared
- Anonymization is not truly effective
- Data annotation is costly. Data is an asset
- Bureaucracy of data sharing is complex
- Share models, not data!

McMahan, Brendan, and Daniel Ramage. "Federated learning: Collaborative machine learning without centralized training data." *Google Research Blog* 3 (2017).

COLLABORATIVE TRAINING METHODS

Obtaining strong models without pooling data

CYCLICAL LEARNING

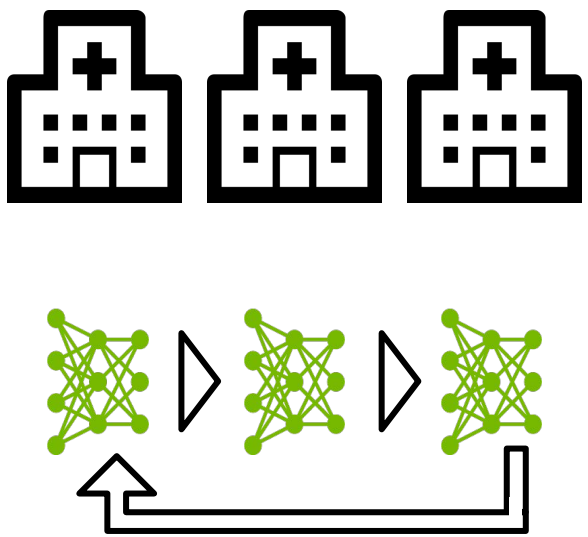

FEDERATED LEARNING

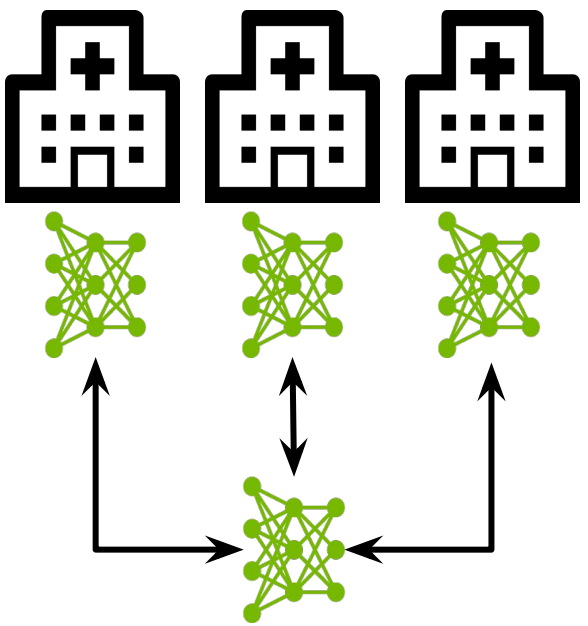

SERVER-CLIENT FEDERATED LEARNING

Changing the way AI algorithms are trained

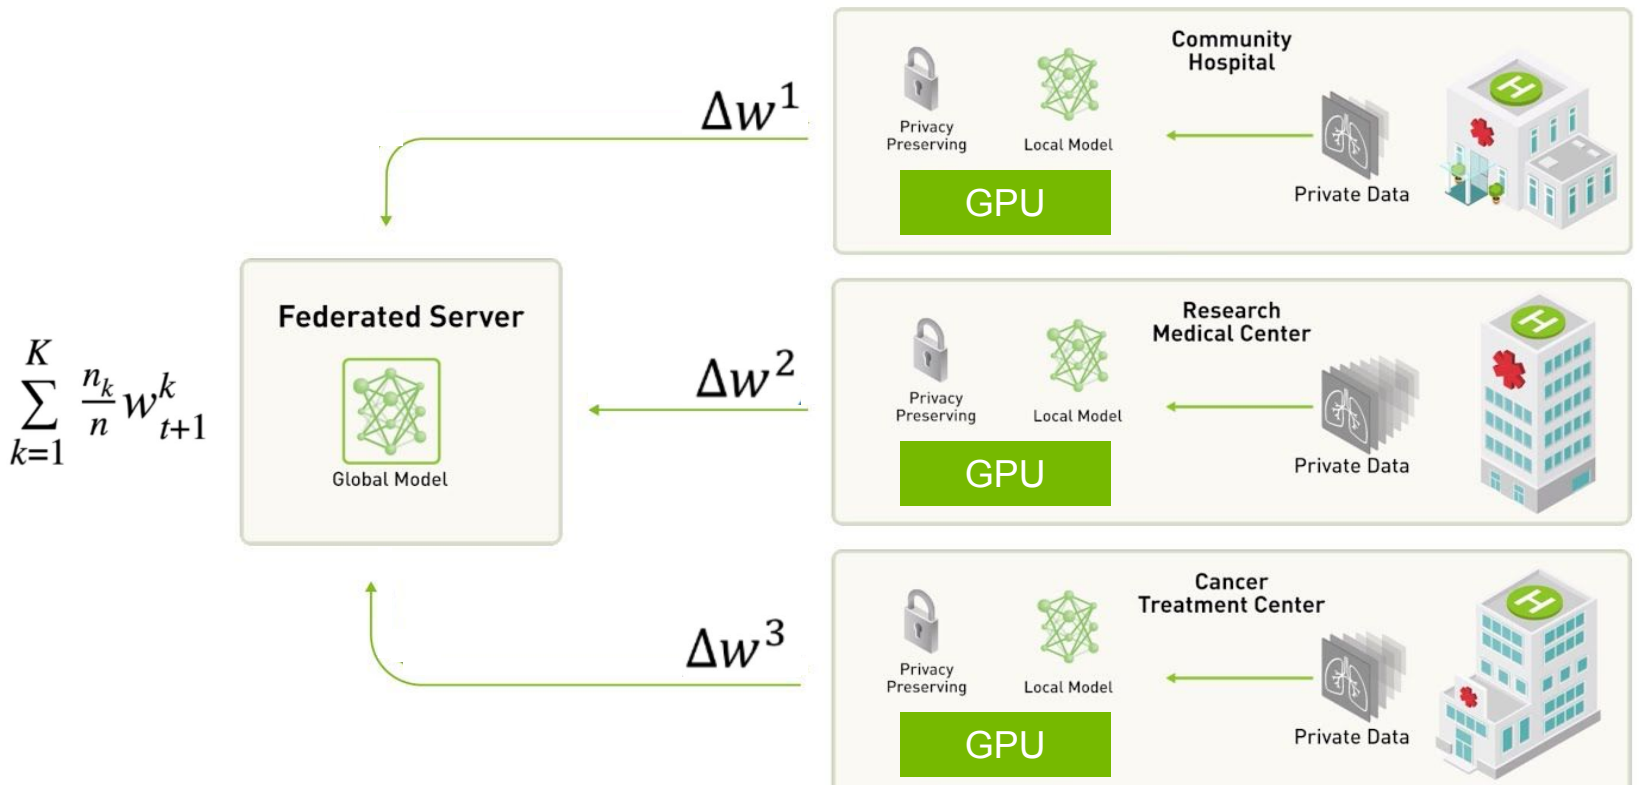

SERVER-CLIENT FEDERATED LEARNING

Changing the way AI algorithms are trained

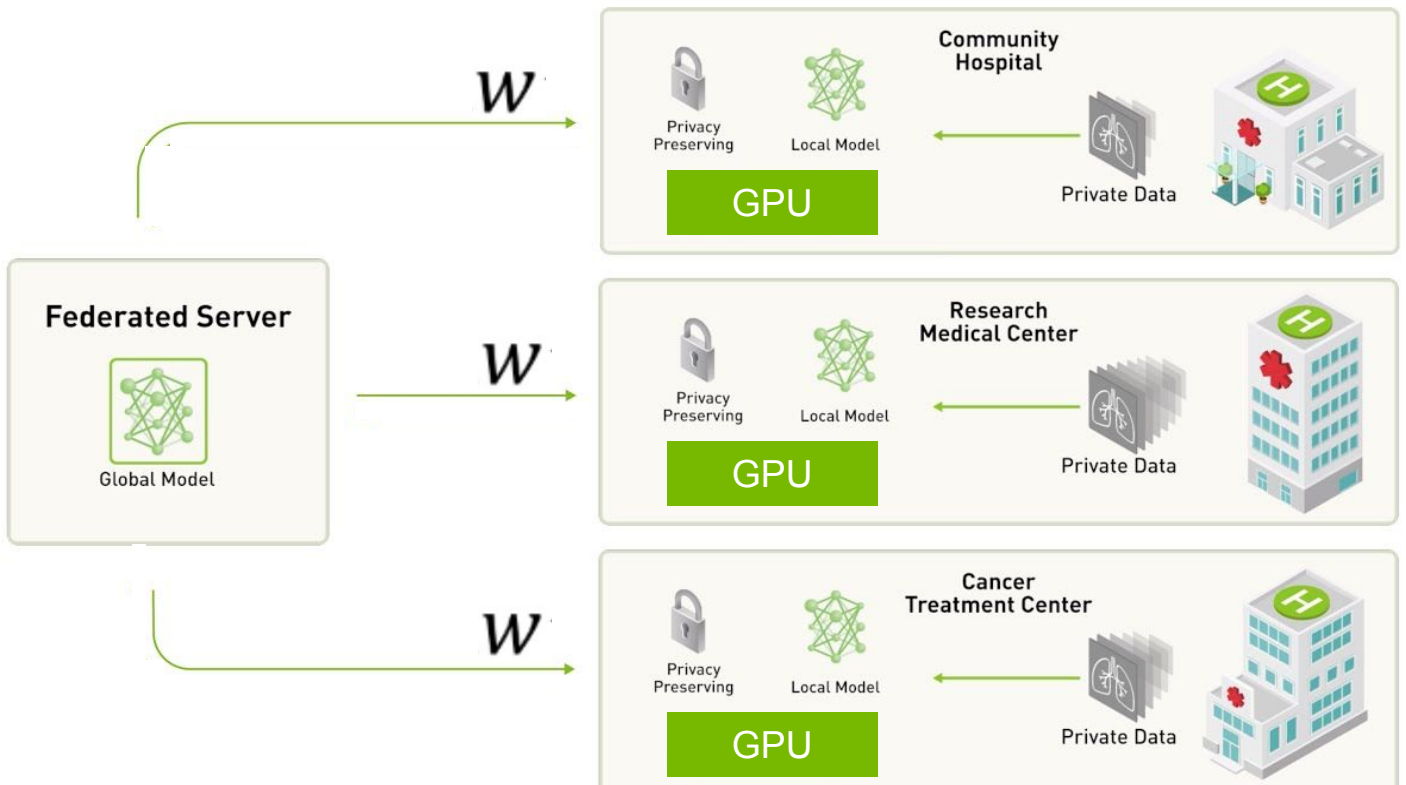

CURRENT CHALLENGES

SYSTEM ARCHITECTURE

Local training in each institution requires computational infrastructure available on-site. Computational requirements for training shift from centralized data-centers attached to data lakes to single institutions or even edge devices.

TRACEABILITY & ACCOUNTABILITY

Data is the new oil! Due to the value and cost of data, it is necessary to trace how different participant to FL training contribute to the final model. A system to link back contributions to the model to participants is also needed.

INITIATIVES & CONSORTIA

Defining collaborations and initiatives leveraging federated learning requires rethinking current agreements between institution. Even though data governance issues are solved, institutions share models that might correspond to IP.

PRIVACY & SECURITY

Deep learning models have a large number of parameters, this may cause parts of the training set to be memorized in their parameters. In federated learning this mechanism determines a possible leakage of private information.

FEDERATED LEARNING AT NVIDIA

Research, development and collaborations

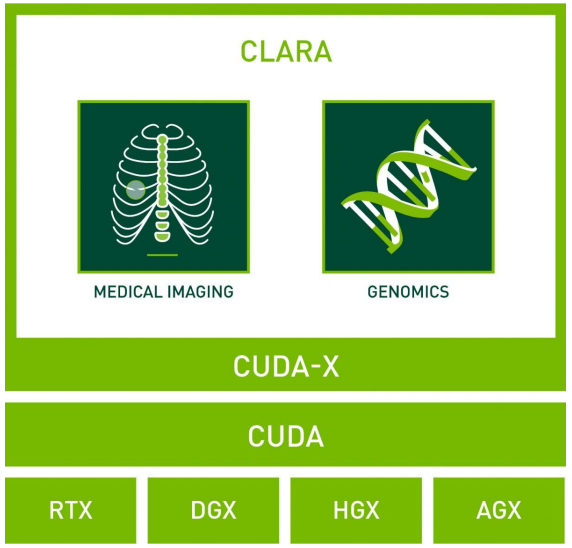

Infrastructure

Developing tools for medical AI on GPU

Collaboration

Partnering up with KCL and Owkin in UK

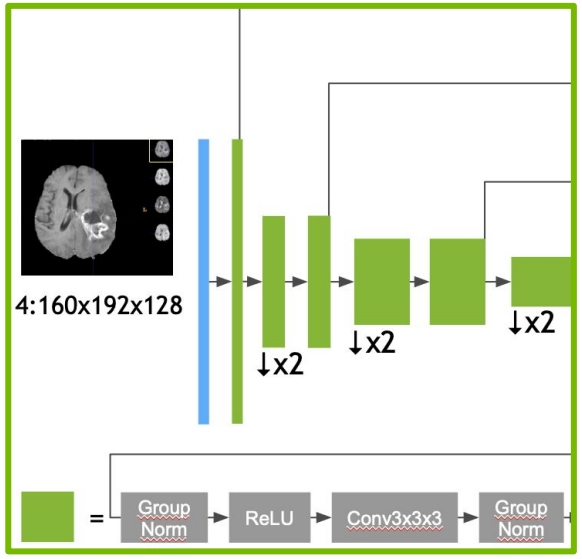

Research

Research of novel approaches for secure FL

CLARA TECHNOLOGY STACK

https://developer.nvidia.com/clara-medical-imaging

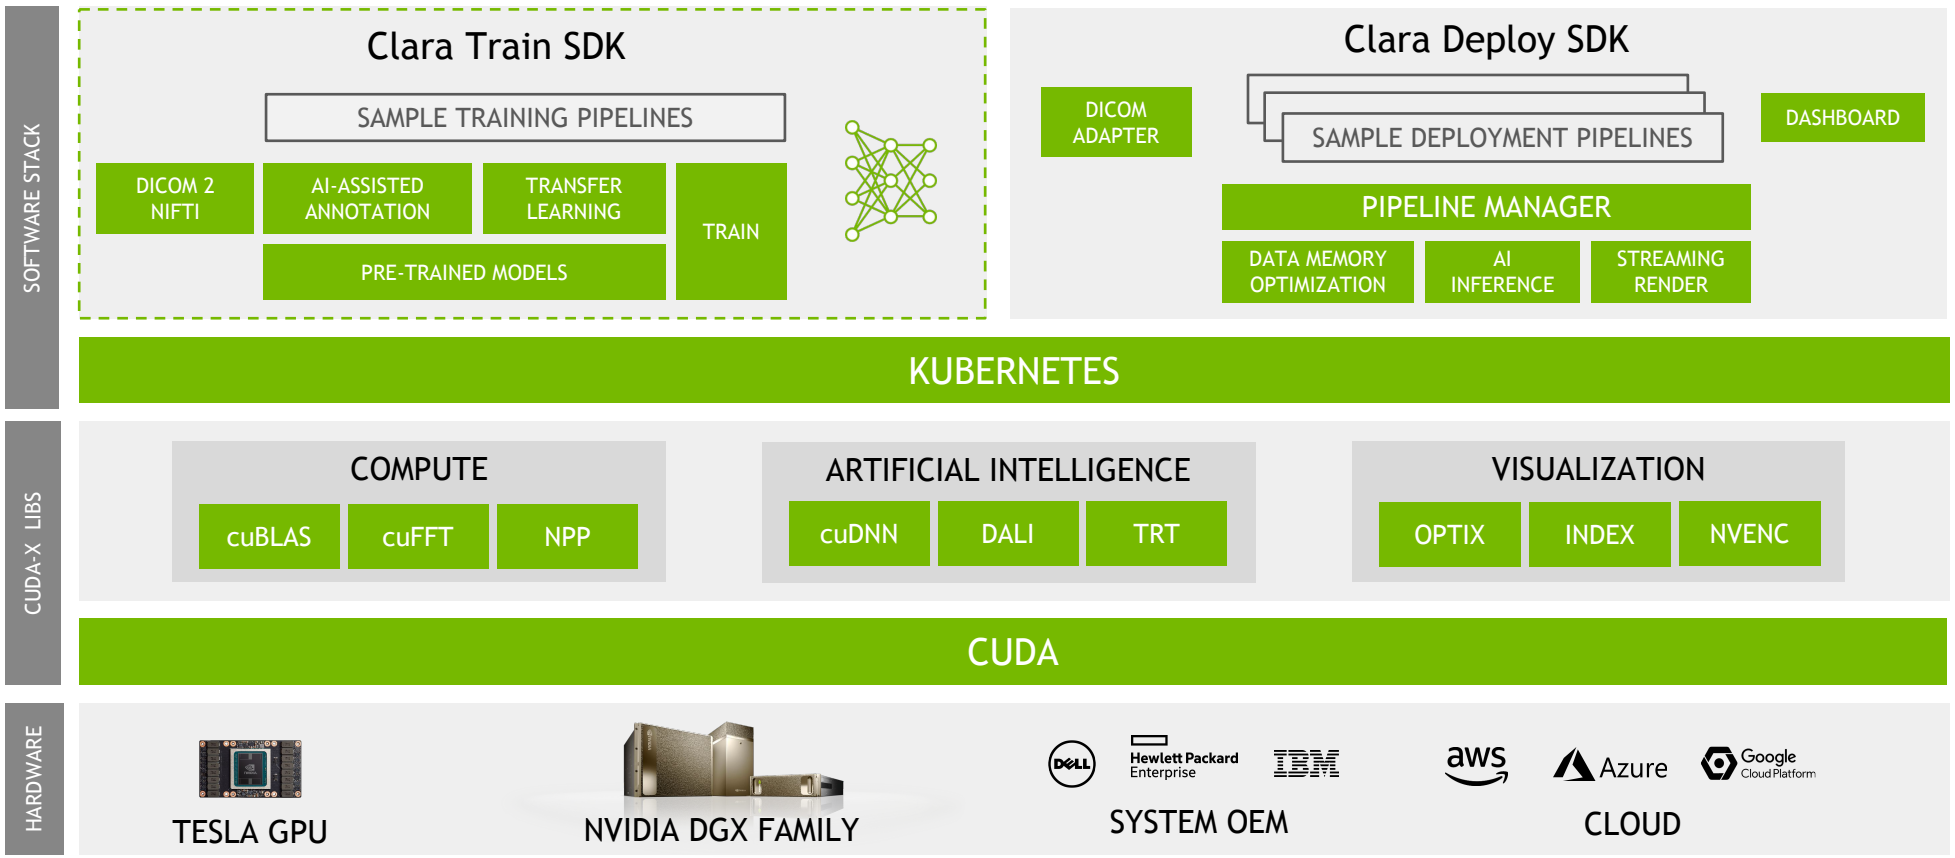

MODULAR TRAINING FRAMEWORK

Feature that Allows Users to Bring Your Own Model

Sample model application framework that standardizes model development and training workflow

Clara Train SDK provides a layer of tools for abstraction

Flexible to make use of any underlying components

Supports, bring your own model, transforms, losses or any component

Users include a library in their code and extend the model class

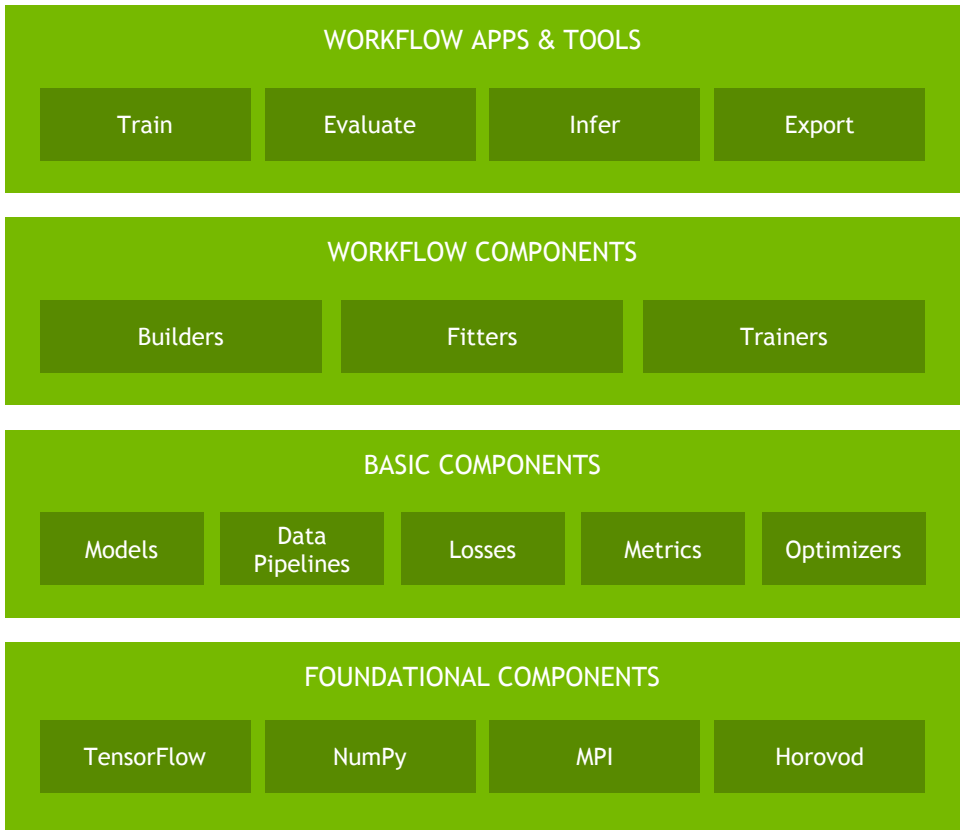

CLARA FEDERATED TRAINING

Building federated learning capabilities into Clara Train SDK

Basic FL capabilities

- Planned for next Clara Train SDK release
- Mainly focused on research needs

Support for centralized model

- “Classic” topology for FL: master + multiple workers

Private deep learning

- Partial model sharing and differential privacy
- Avoid data leakage through model parameters

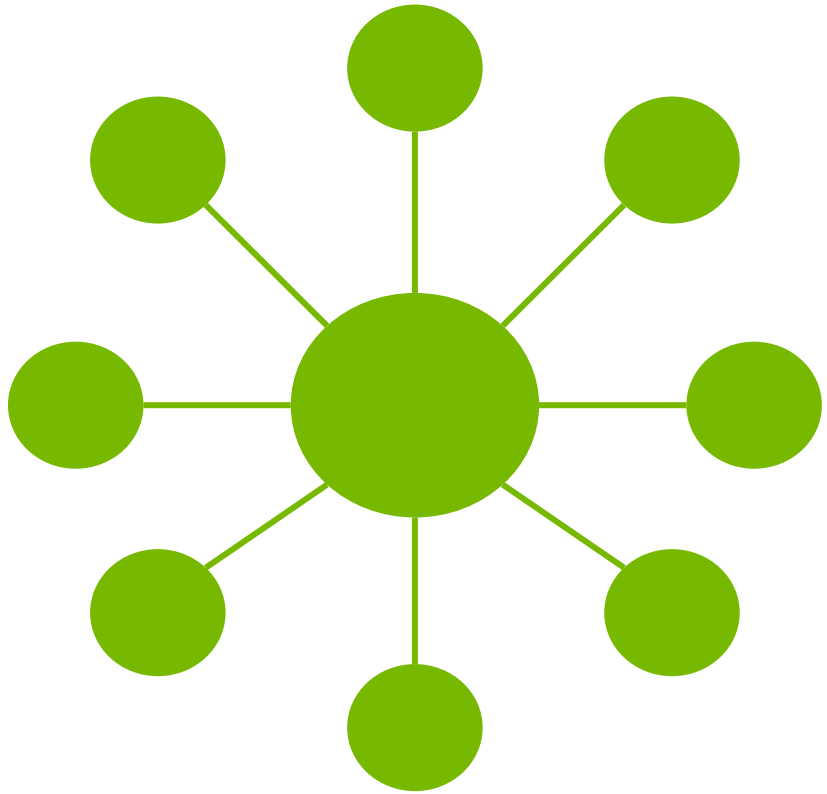

PRIVACY PRESERVING FL

Building differential privacy for collaborative training

Federated privacy baseline

- No direct data sharing
- Authentication, authorisation, secure connection

Private information can still leak out...

- Network may memorize data in its parameters
- Additional privacy mechanisms need to be enforced
- Literature proposes various strategies for FL privacy

https://xkcd.com/2169/

WHEN YOU TRAIN PREDICTIVE MODELS ON INPUT FROM YOUR USERS, IT CAN LEAK INFORMATION IN UNEXPECTED WAYS.

NVIDIA + KCL FL RESEARCH

Privacy-preserving federated brain tumor segmentation

Going beyond the baseline

- Exploring the concept of differential privacy
- Investigate robustness against model inversion

Our work in Federated Learning privacy

- Use data from multiple institutions
- Find trade-off between privacy and performance
- Research improved optimization and sharing schemes
- We show results of different techniques and methods

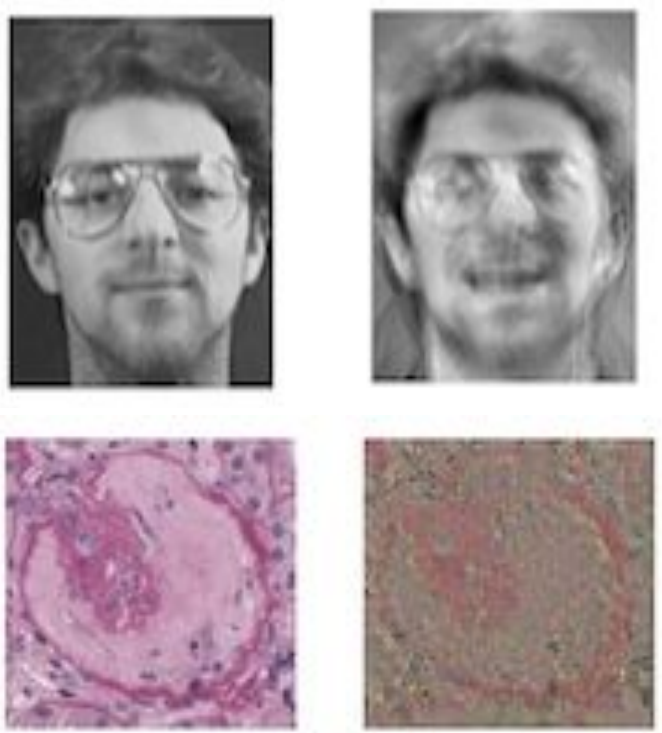

Fredrikson et al. "Model Inversion Attacks that Exploit Confidence Information and Basic Countermeasures" CCS 2015
Wu et al., "P3SGD: Patient Privacy Preserving SGD for Regularizing Deep CNNs in Pathological Image Classification" CVPR 2019

STATE OF THE ART TRAINING

Using best ranking models for FL within Clara Train SDK

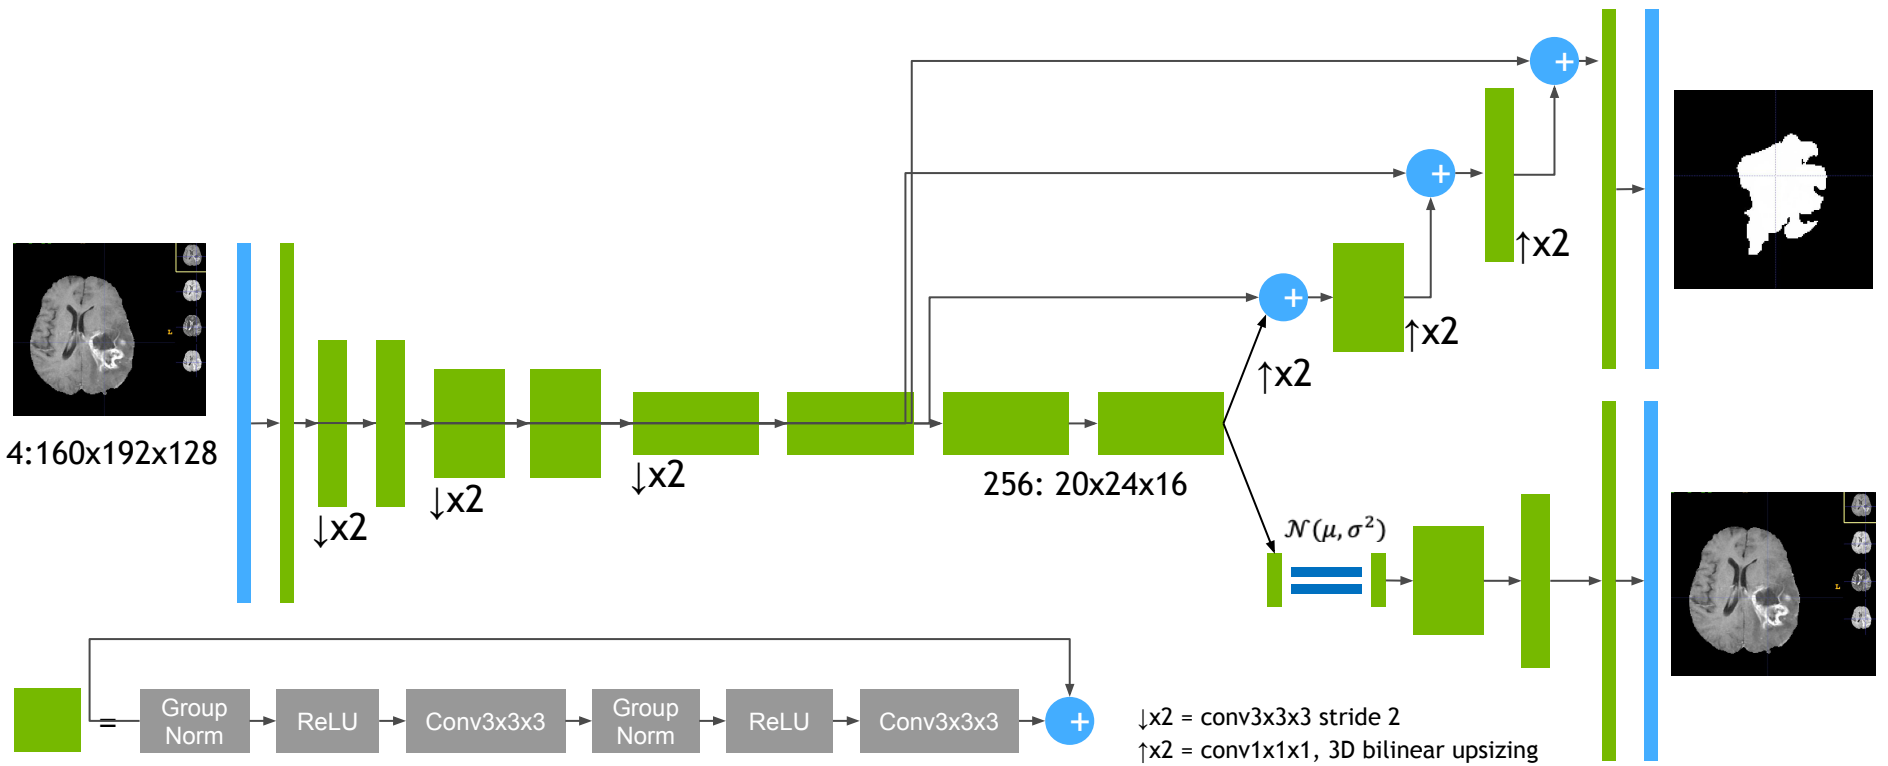

Myronenko A.: "3D MRI brain tumor segmentation using autoencoder regularization", BrainLes, MICCAI 2018

MULTI INSTITUTION DATA

Training a common model without leaking private information

- Each institution holds variable amounts of data
- Protect privacy of model updates from each institution
- Different techniques and best practices for private FL
  - Truncating values of model updates
  - Selecting partial models based on a noisy threshold on updates
  - Adding noise to selected model before sharing

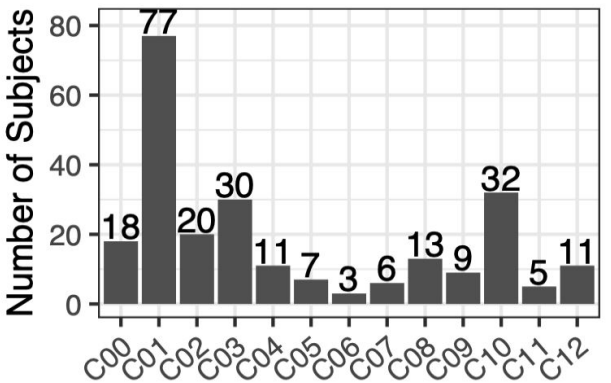

Data distribution
across institutions

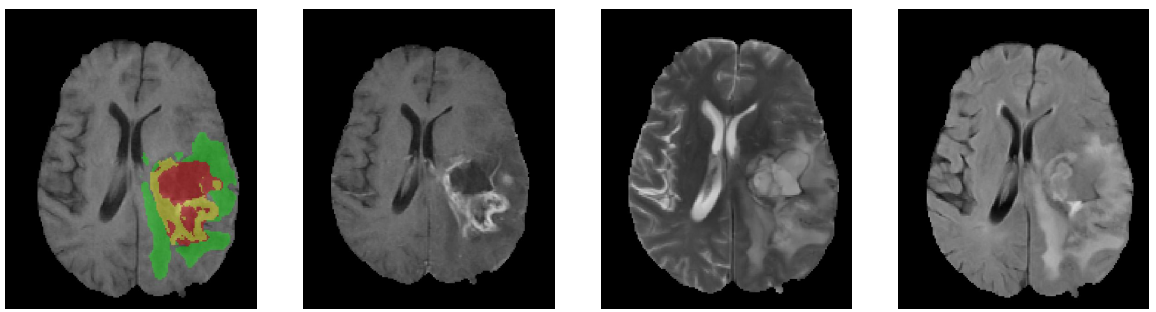

Multimodal Brain Tumor Segmentation Challenge 2018

OUR FINDINGS

Protecting privacy while ensuring best performance

PARTIAL MODEL SHARING

We reduced the amount of parameters shared in each FL round by each client.

These parameters have been selected at random in each round.

Results show great performance when sharing only 40% of model.

MOMENTUM RESTARTING

We used an optimizer with momentum.

The variables holding momentum should not be shared with FL server and kept in sync.

We got best results when re-initializing these variables in clients.

DIFFERENTIAL PRIVACY

Legacy differential privacy can be achieved by adding Laplacian noise to model parameter updates.

We observed that sharing less parameters (10% or 40%) improves the performance in presence of noisy updates.

Li, et al. "Privacy-preserving Federated Brain Tumour Segmentation." MICCAI MLMI 2019

OUR FINDINGS

Further insights on privacy and model inversion attacks

- There's a trade-off between privacy protection and quality of model
- The system can provide good privacy protection with a reasonably small cost in model performance
- Privacy and accuracy are a trade-off

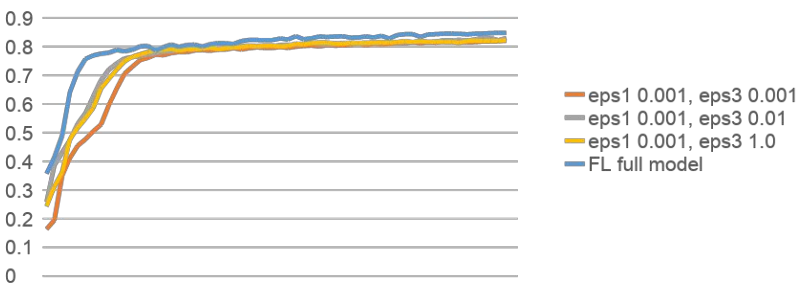

Slices from the training volumes

Reconstruction by inverting a standard model after training

Reconstruction by inverting model with privacy-preservation

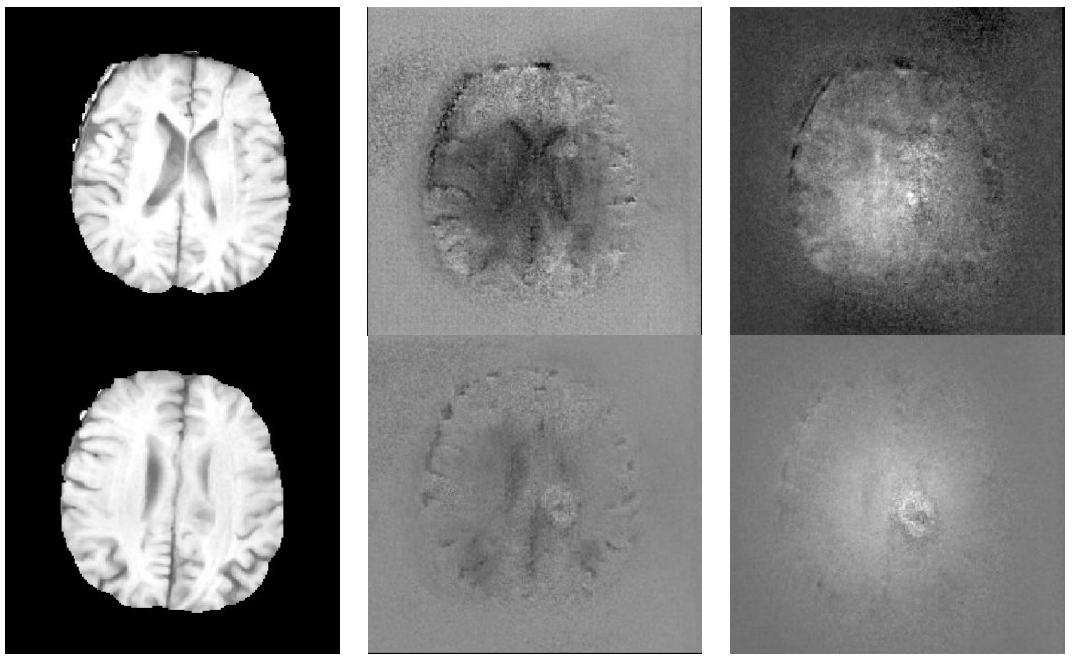

Li, et al. "Privacy-preserving Federated Brain Tumour Segmentation." MICCAI MLMI 2019

FEDERATED LEARNING PARTNERSHIP

Towards real-life applications of FL

King’s Health Partners

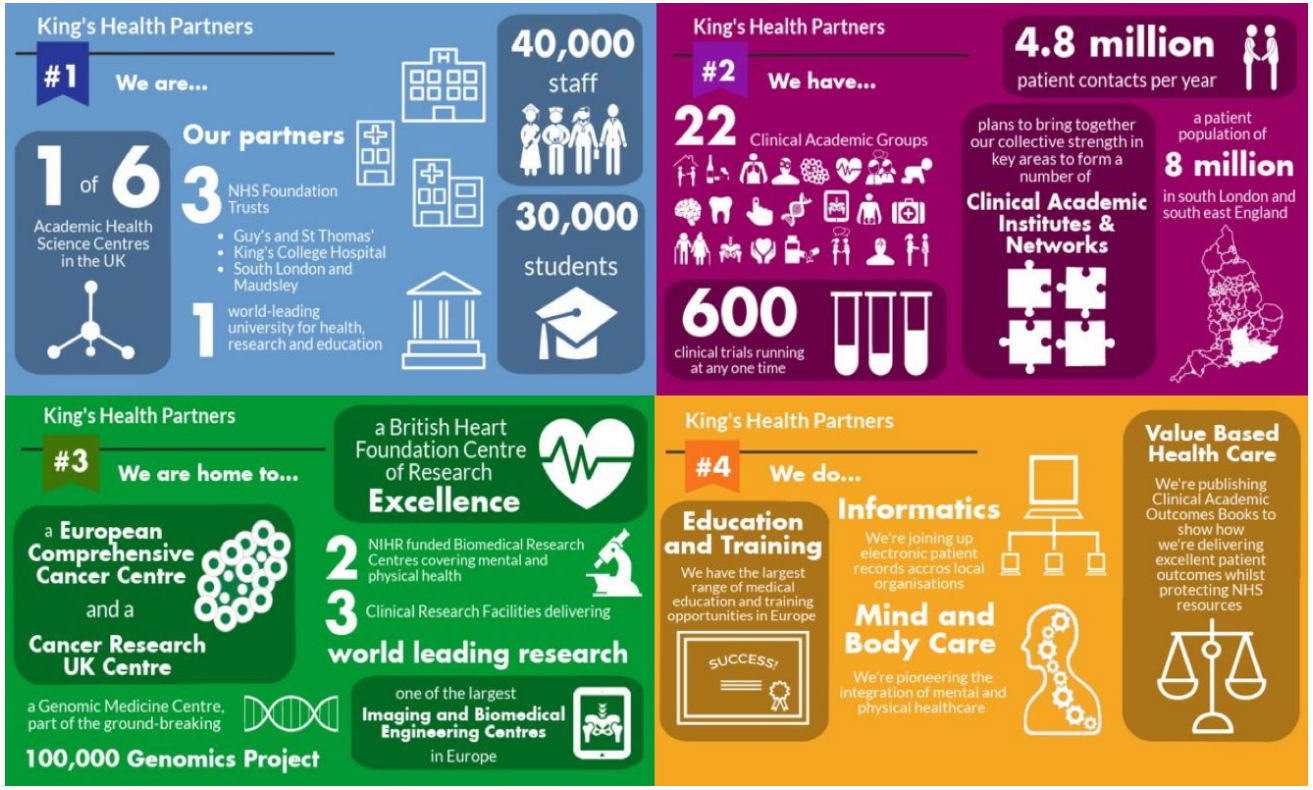

KING'S College LONDON

NHS King’s College Hospital NHS Foundation Trust

NHS South London and Maudsley NHS Foundation Trust

NHS Guy’s and St Thomas’ NHS Foundation Trust

AI4VBH: Full spectrum of applications

Centre funded as part of the UK industrial strategy

3 universities (KCL, ICL, QM), 4 hospitals (GSTT, KCH, SLaM, Barts)

4 multinationals (Siemens, Nvidia, IBM, GSK), 12 startups

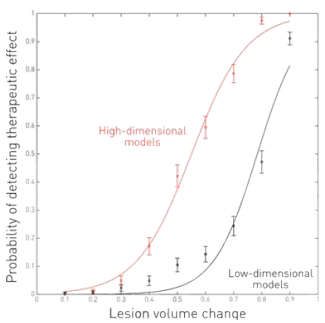

Patient Care

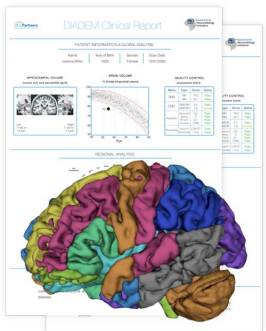

Clinical Efficiency

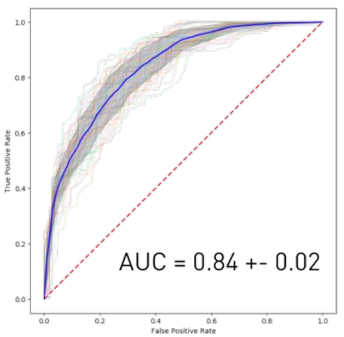

Live Auditing

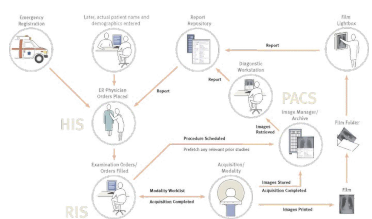

Operational Efficiency

Value Based Heathcare

From snapshots to real-time analytics

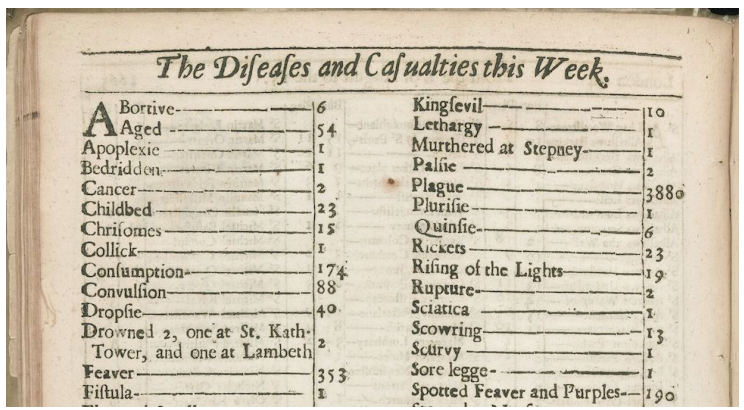

The Diſeaſes and Caſualties this Week.

Abortive	6	Kingſevil	10
Aged	54	Lethargy	1
Apoplexie	1	Murthered at Stepney	1
Bedridden	1	Paliſie	2
Cancer	2	Plague	3880
Childbed	23	Pluriſie	1
Chriſomes	15	Quinſie	6
Collick	1	Rickets	23
Conſumption	174	Riſing of the Lights	19
Convulſion	88	Rupture	2
Dropſie	40	Sciatica	1
Drowned 2, one at St. Kath. Tower, and one at Lambeth	2	Scowring	13
		Scurvy	1
Feaver	353	Sore legge	1
Fiſtula	1	Spotted Feaver and Purples	190

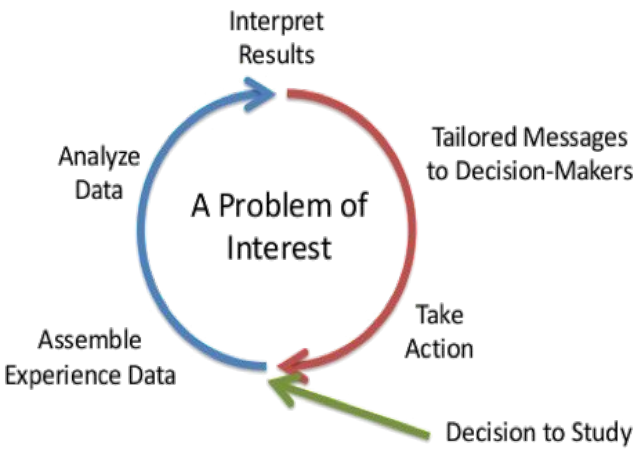

End-to-end Actionable Analytics

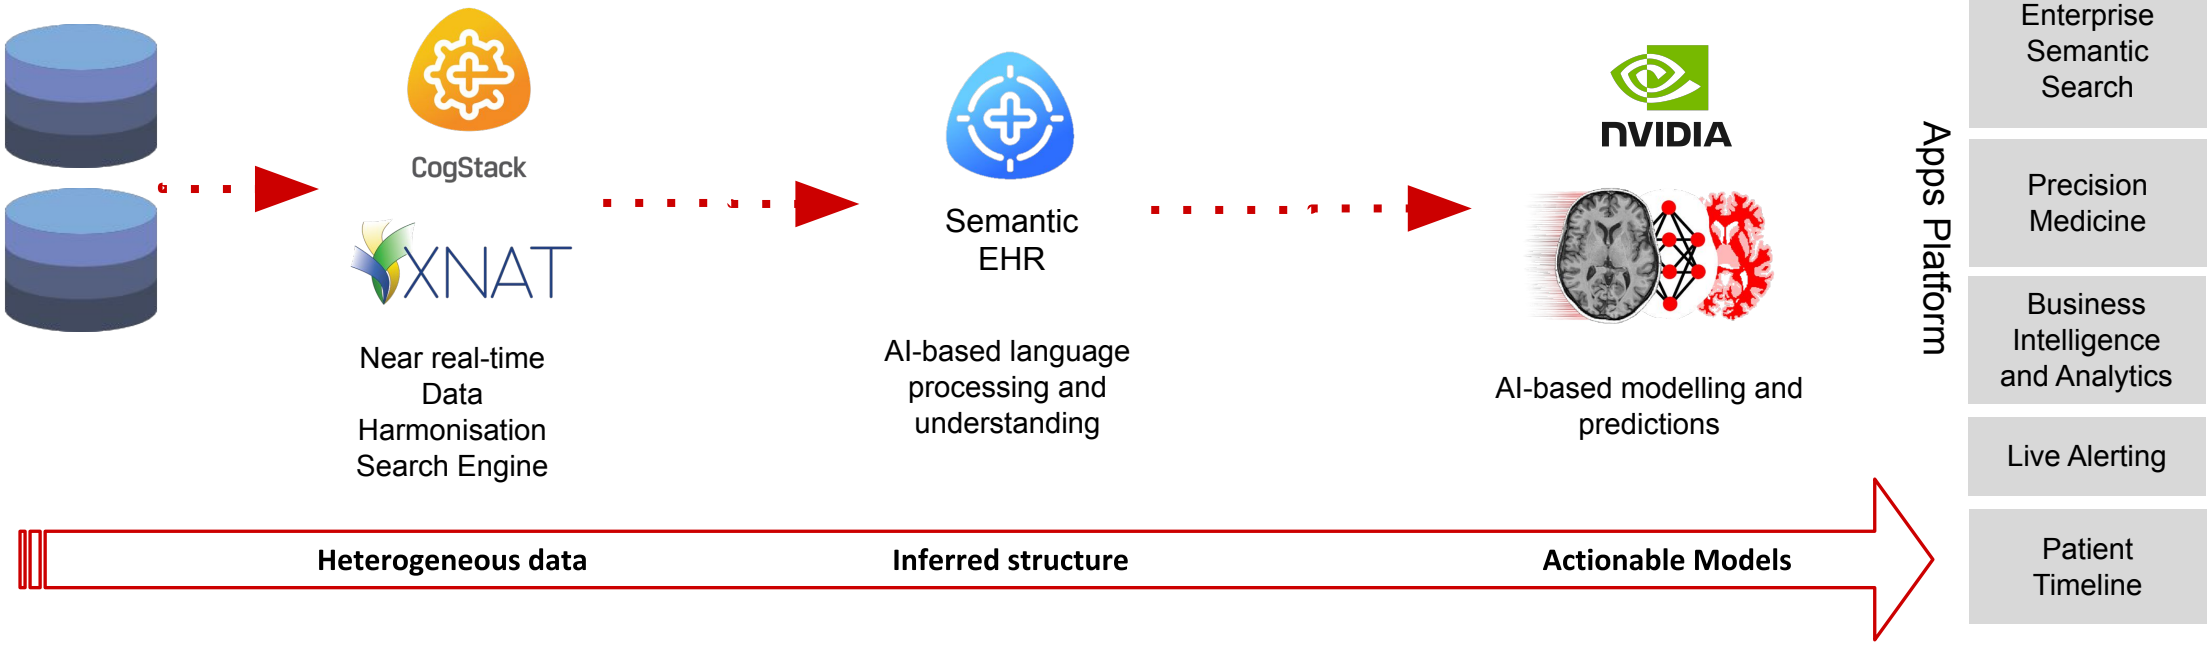

Federated Learning @ KHP

- Safe
- Open Source
- Privacy-preserving
- Traceability
- Scalable
- Governance compliance

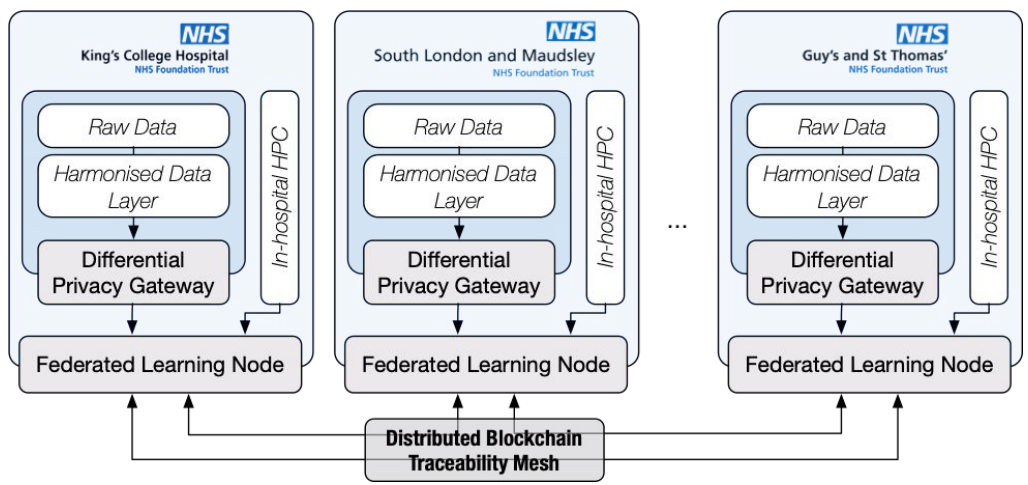

IMPACT OF FEDERATED LEARNING

Increasing the value of AI for all healthcare stakeholders

Doctors
Accurate assistance tools, unbiased decisions

Patients
Accurate and unbiased AI, cost reductions

Researchers
Safe collaboration, access to large datasets, impact

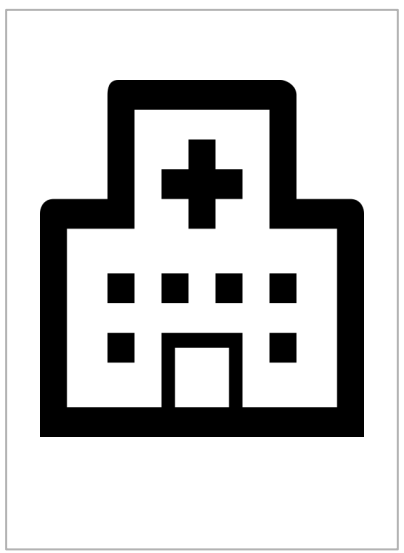

Healthcare Providers
Access accurate AI while monetizing data via FL

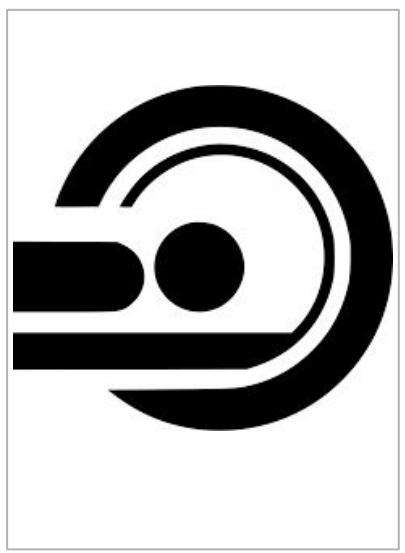

Manufacturers
Learning from users for ever-improving products

Built on partnerships

Semantic EHR

XNAT

CogStack

Software

SCAN

NVIDIA

Mellanox TECHNOLOGIES

NetApp

Hardware

omni·sci

IBM

SIEMENS Healthineers

OWKIN

Platforms

Create a startup ecosystem

- Create ecosystem for tech transfer
  - Spin-offs and spin-ins
- Bring together:
  - Academics
  - Clinicians
  - Industry

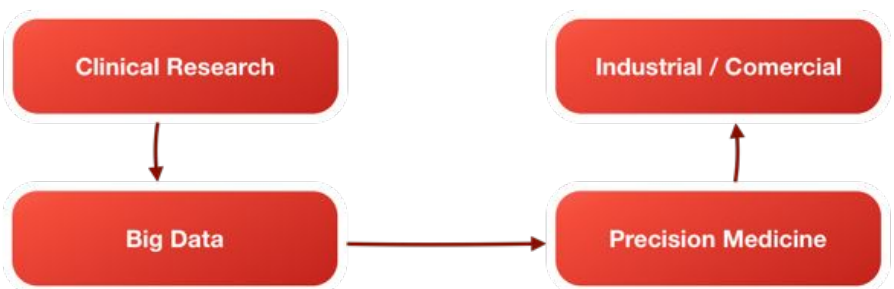

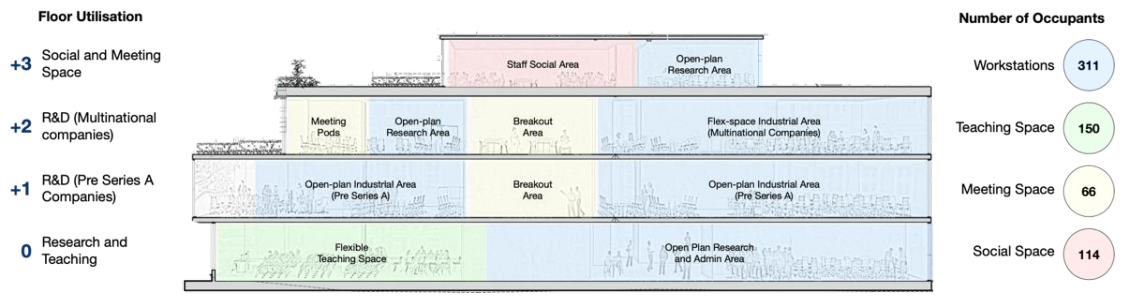

Conclusion

Factors for the success of AI in healthcare

- Large scale data access is necessary for accurate, safe and ethical AI
- Governance and privacy currently limits the potential of healthcare AI
- Federated learning solves governance and large scale data access
- Differential privacy protects user information
- Federated learning + differential privacy will drive the success of AI in healthcare


THANK YOU!
QUESTIONS?
NVIDIA
KING'S College LONDON