Imperial College London

Klinikum rechts der Isar
Technische Universität München

Human-centered AI for Medical Imaging

Daniel Rueckert, FREng, FMedSci, FIEEE

Alexander von Humboldt Professor of AI for Medicine and Healthcare
School of Medicine and Informatics, TU Munich, Germany

and

Biomedical Image Analysis Group
Department of Computing, Imperial College London, UK

AI/ML in Medicine

23,216 views | Apr 30, 2017, 12:10pm

AI In Medicine: Rise Of The Machines

Paul Hsieh Contributor

I cover health care and economics from a free-market perspective.

THE NEW YORKER

APRIL 3, 2017 ISSUE

A.I. VERSUS M.D.

What happens when diagnosis is automated?

By Siddhartha Mukherjee

AI/ML in Medicine: There is a lot of hype

MIT Technology Review

Artificial intelligence / Machine learning

Hundreds of AI tools have been built to catch covid. None of them helped.

Some have been used in hospitals, despite not being properly tested. But the pandemic could help make medical AI better.

by **Will Douglas Heaven**

July 30, 2021

AI/ML in Medical Imaging

- Out of 64 AI/ML based, FDA approved medical devices and algorithms, 30 (46.9%) for focus on radiology

npj | digital medicine

Explore content ⌄ Journal information ⌄ Publish with us ⌄

nature > npj digital medicine > articles > article

Article | Open Access | Published: 11 September 2020

The state of artificial intelligence-based FDA-approved medical devices and algorithms: an online database

Stan Benjamens, Pranavsingh Dhunnoo & Bertalan Meskó

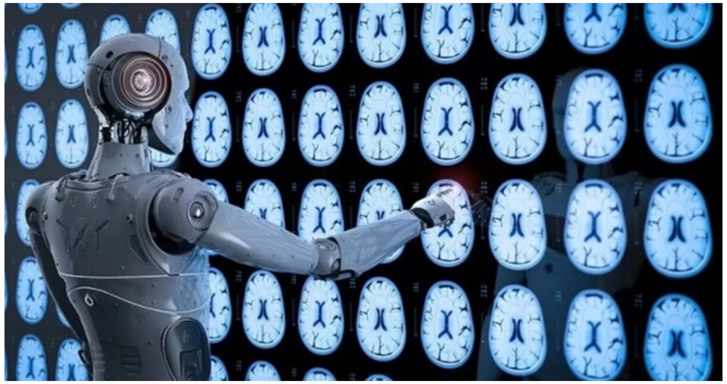

MIT Technology Review

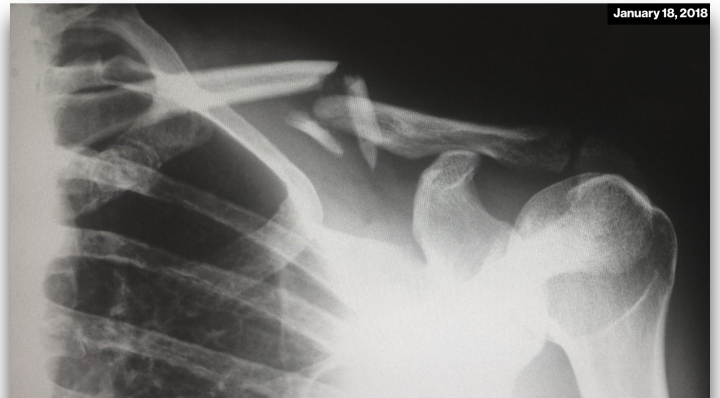

AI Is Continuing Its Assault on Radiologists

A new model can detect abnormalities in x-rays better than radiologists—in some parts of the body, anyway.

AI in Medical Imaging: Opportunities

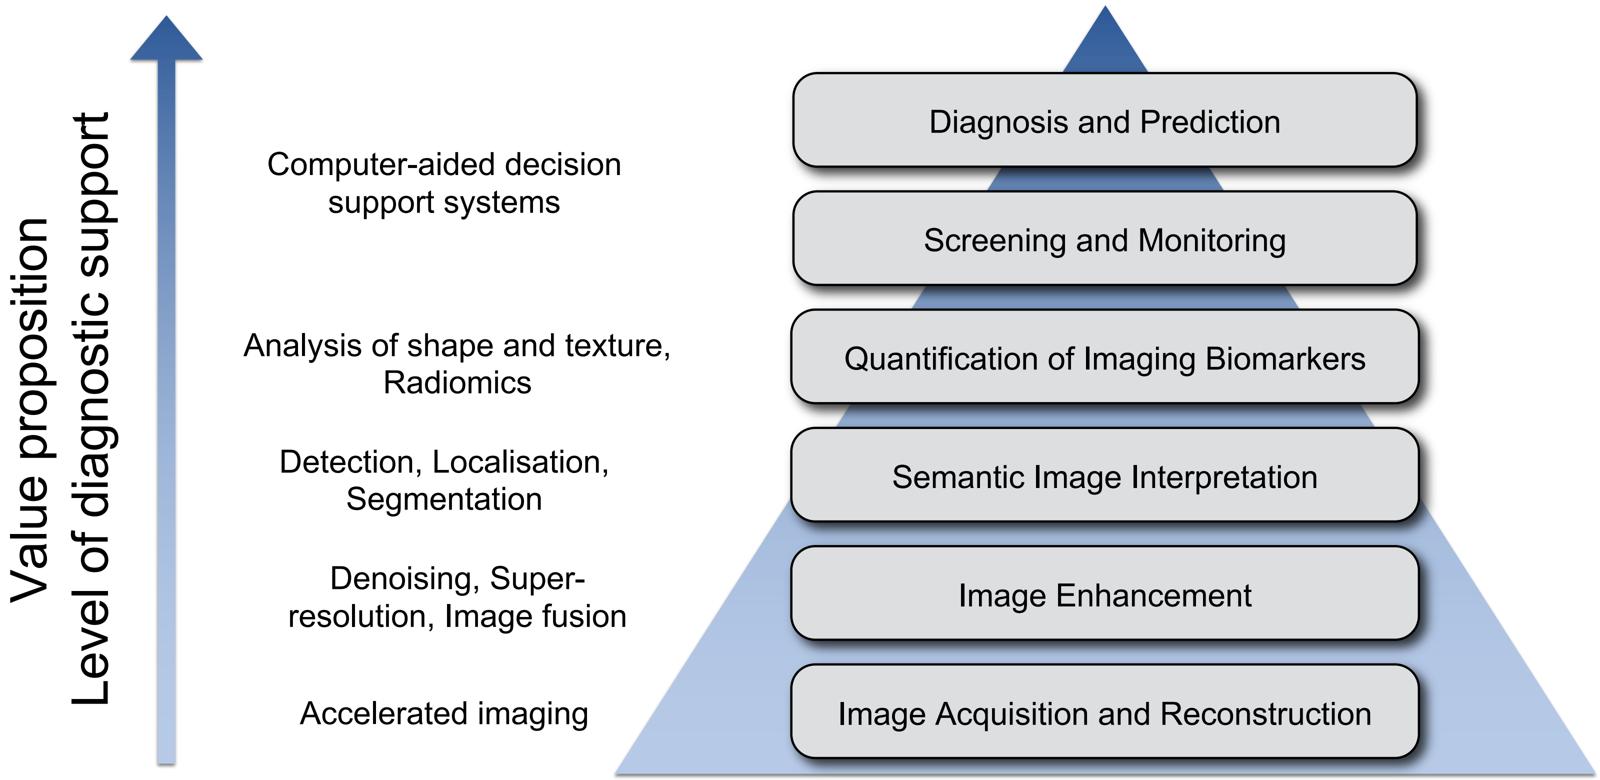

Learning to reconstruct cardiac MRI

K-space | **Signal space**

Full sampling (slow)

$\mathcal{F}^{-1}\{.\}$

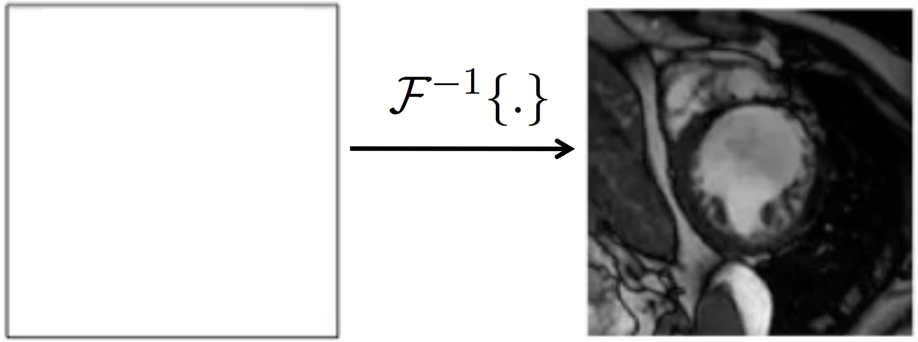

$\mathbf{x}_f$

learning-based reconstruction

e.g. compressed sensing

Deep learning for image reconstruction

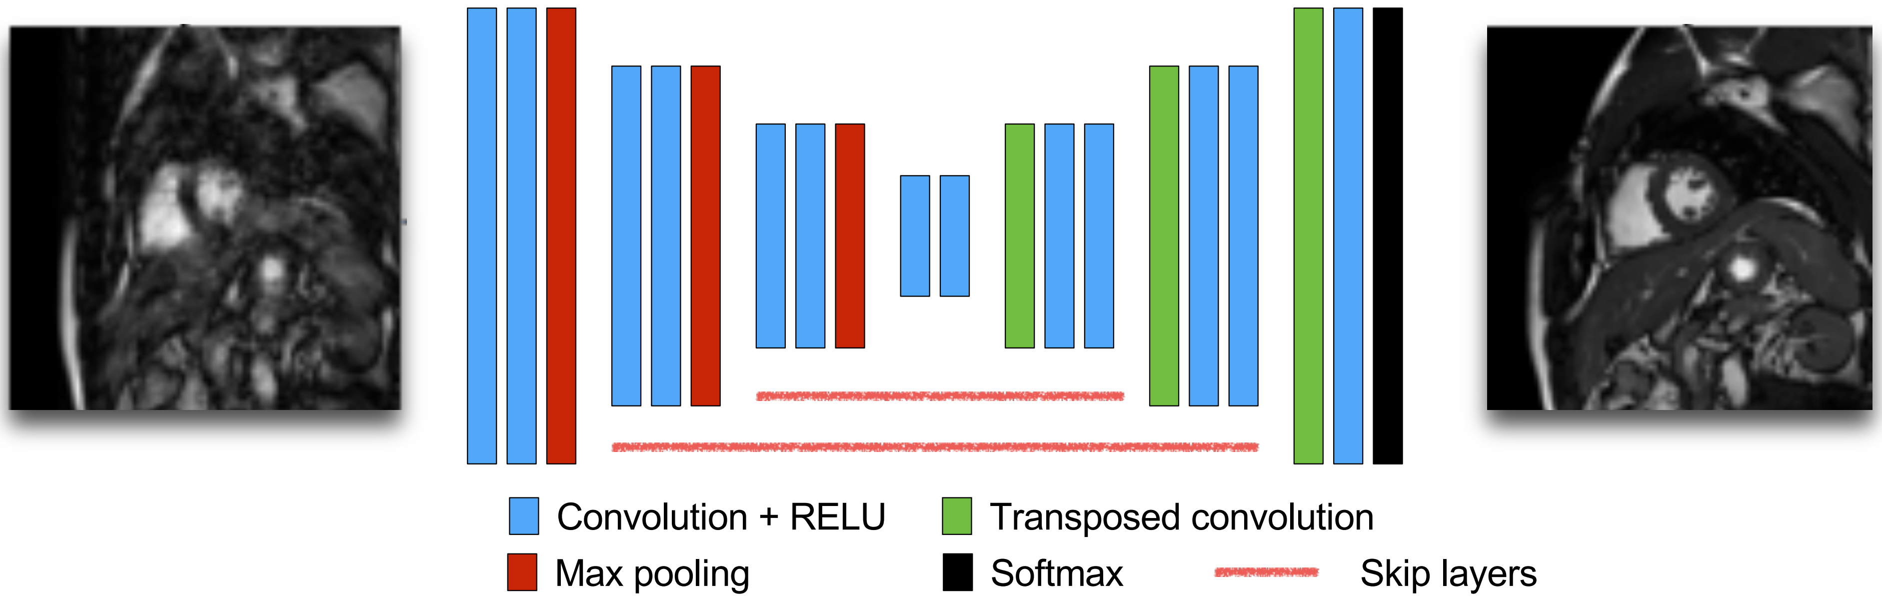

Deep learning for image reconstruction

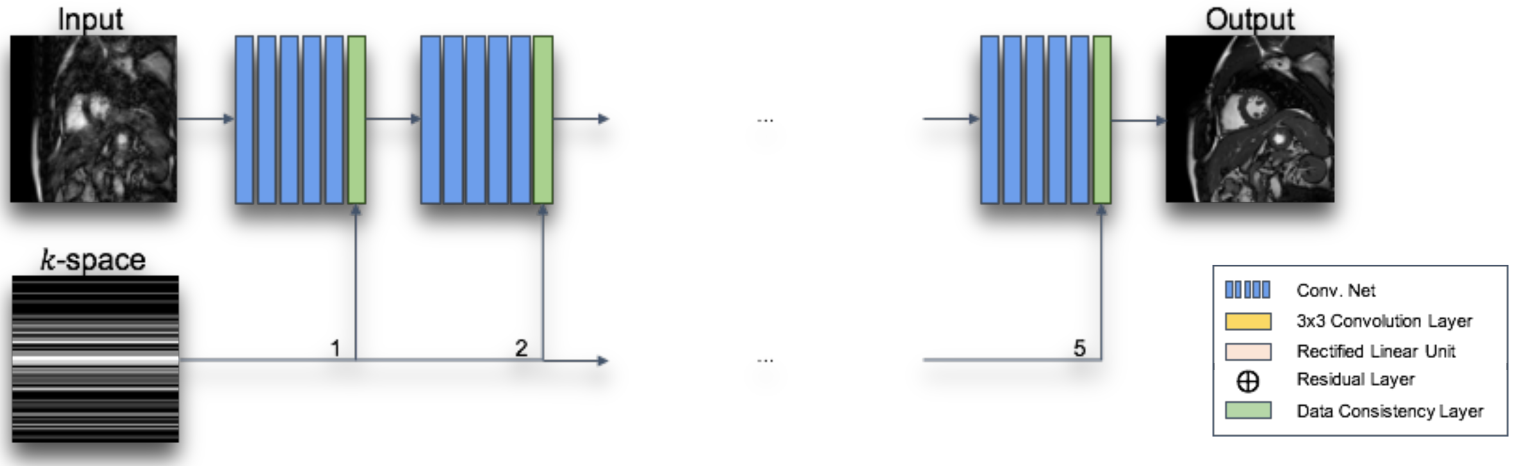

Schlemper et al. IEEE TMI 2017

Deep learning for image reconstruction

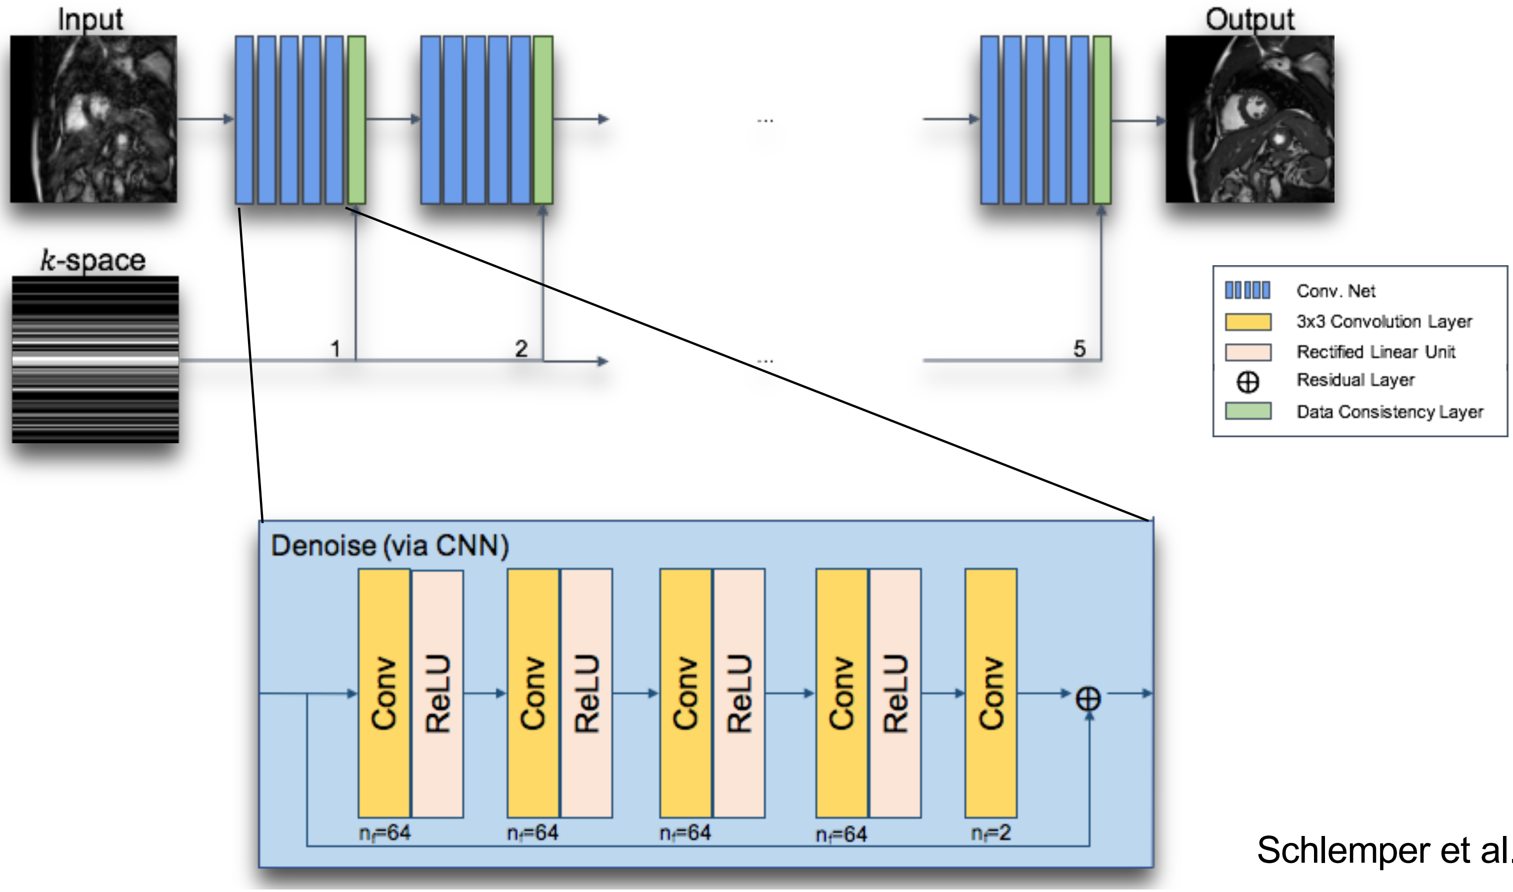

Schlemper et al. IEEE TMI 2017

Deep learning for image reconstruction

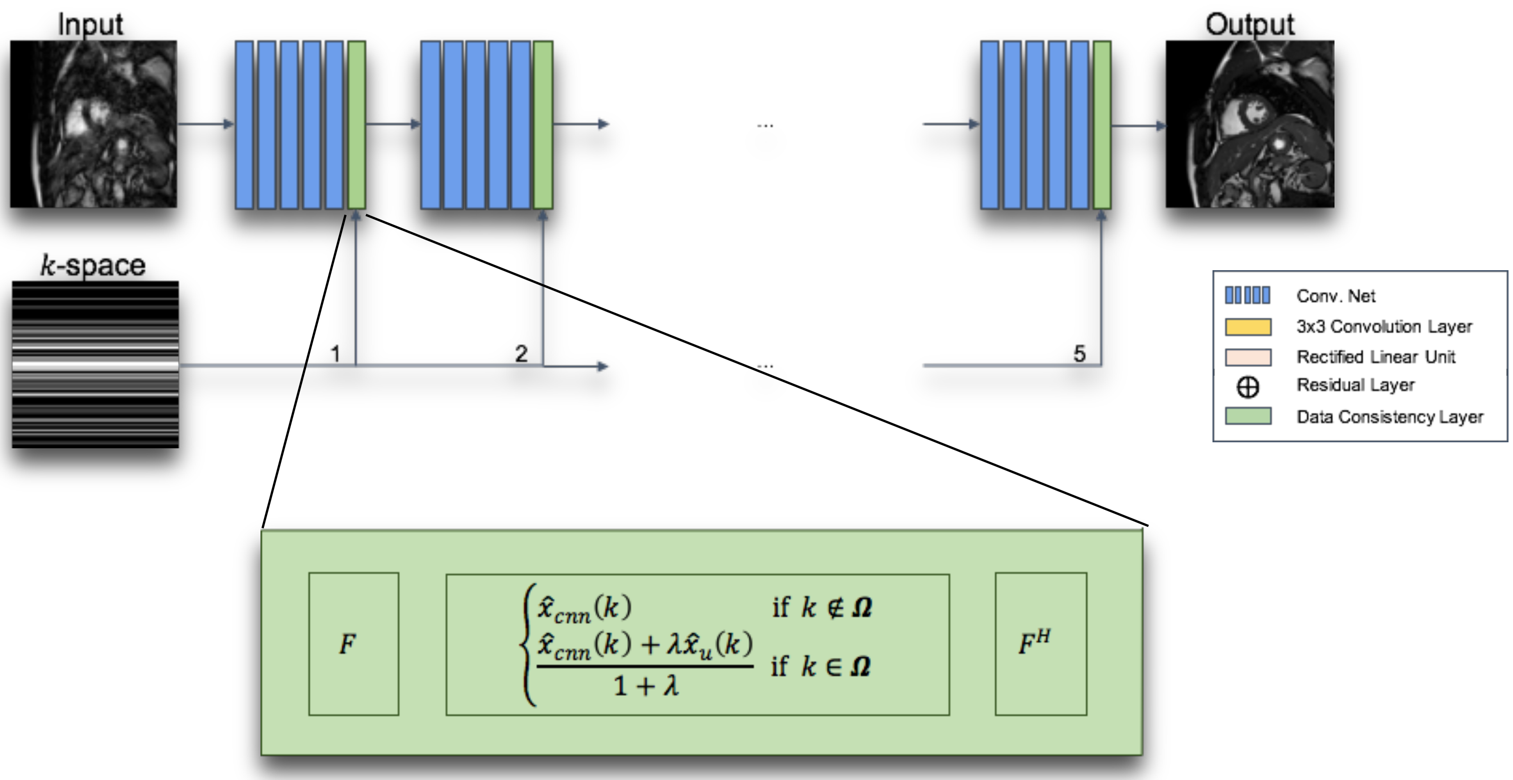

Schlemper et al. IEEE TMI 2017

Magnitude reconstruction (6-fold)

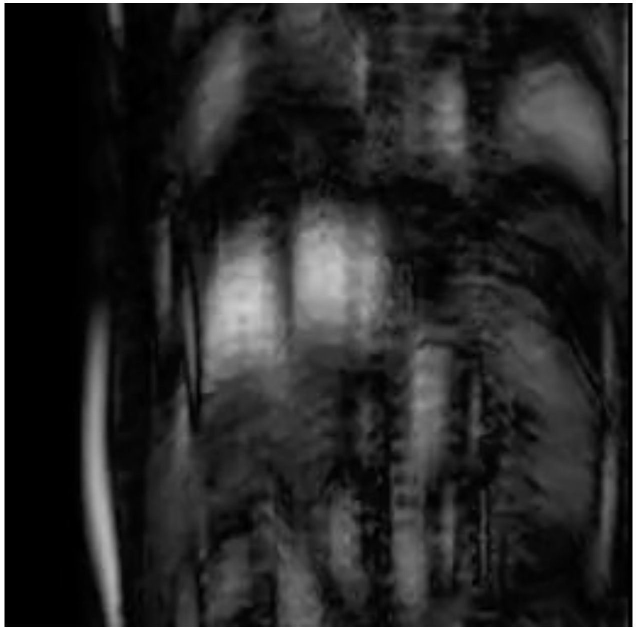

(a) 6x Undersampled

Schlemper et al. IEEE TMI 2017

Magnitude reconstruction (11-fold)

(a) 11x Undersampled

Schlemper et al. IEEE TMI 2017

AI-enabled image super-resolution

- Acquisition of cardiac MRI typically consists of 2D multi-slice data due to
  - constraints on SNR
  - breath-hold time
  - total acquisition time
- This leads to thick slice data (thickness 8-10 mm per slice)

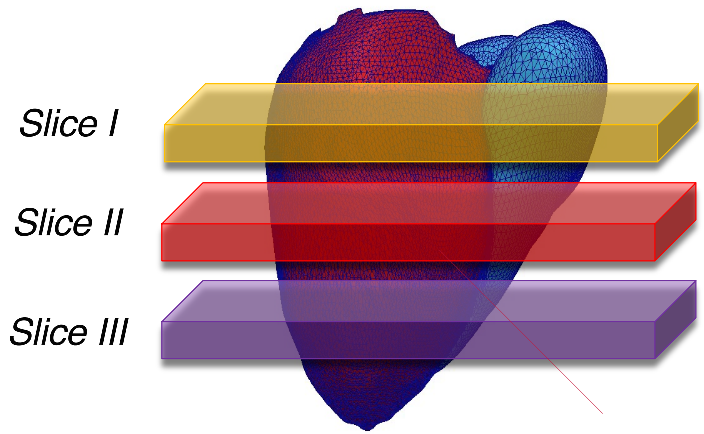

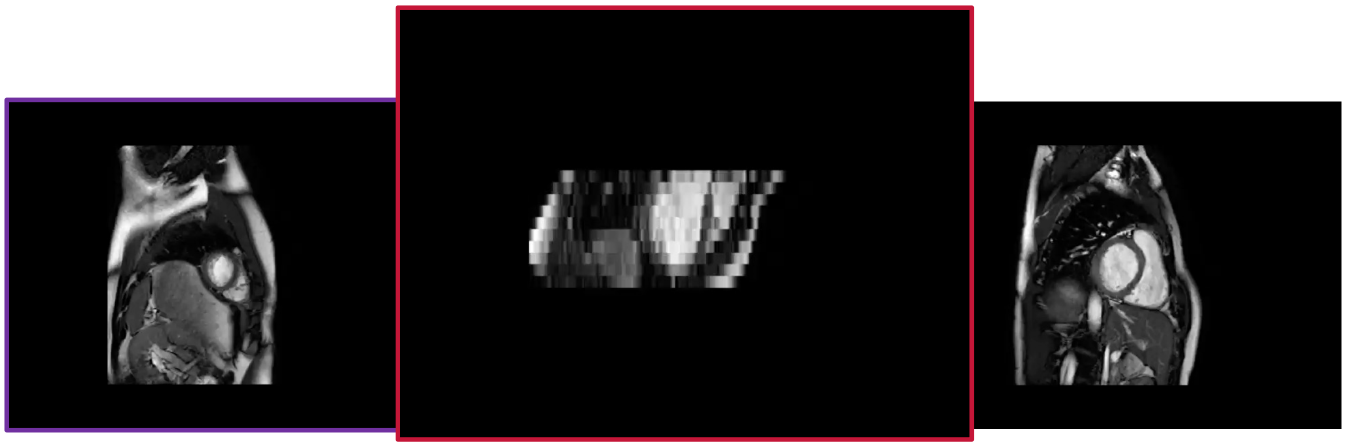

AI-enabled image super-resolution

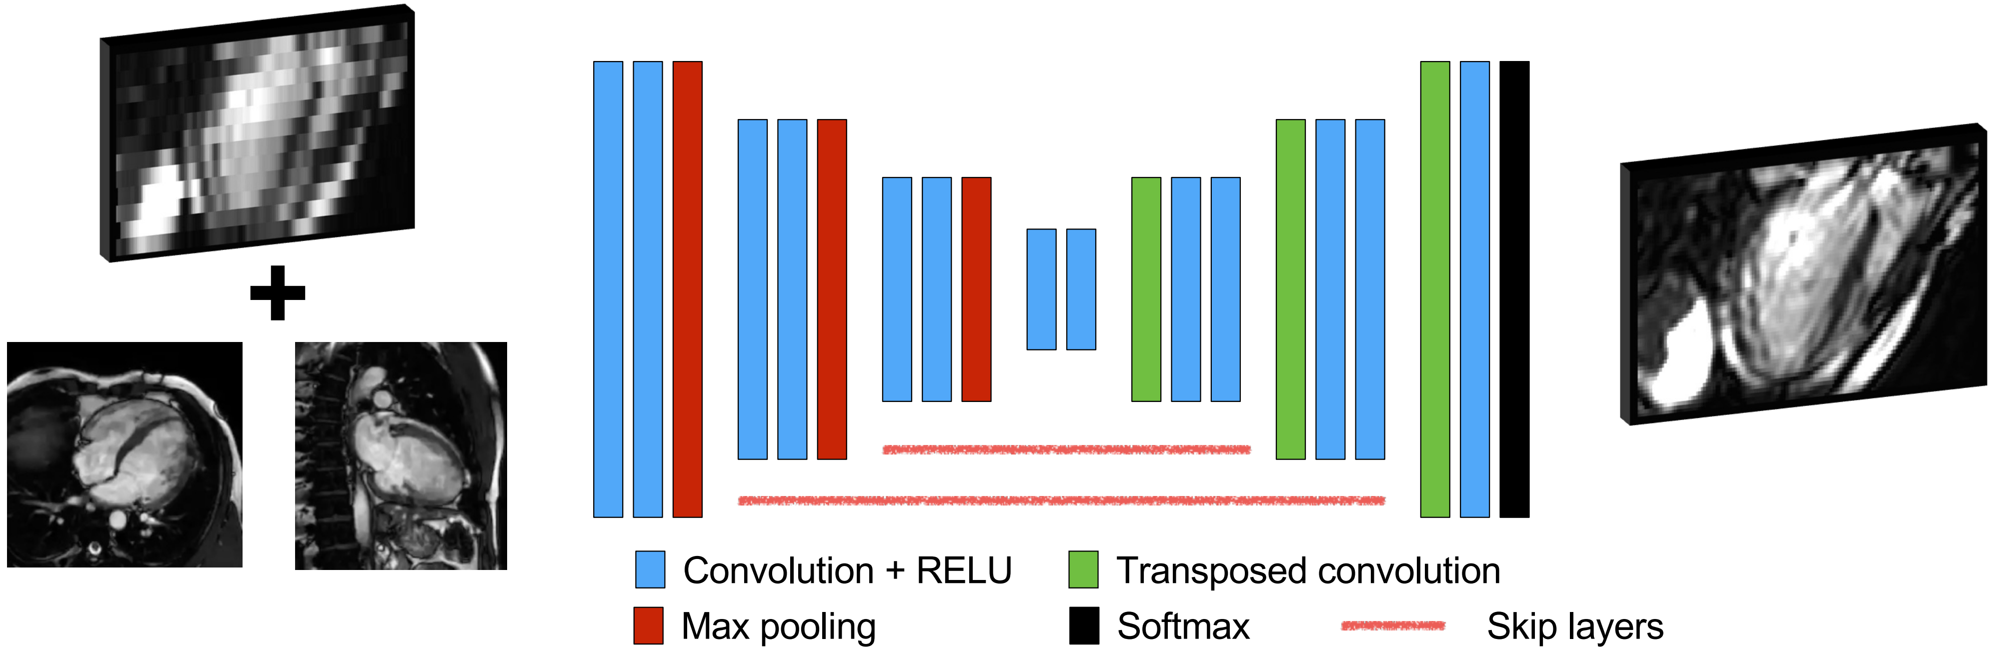

AI-enabled image super-resolution

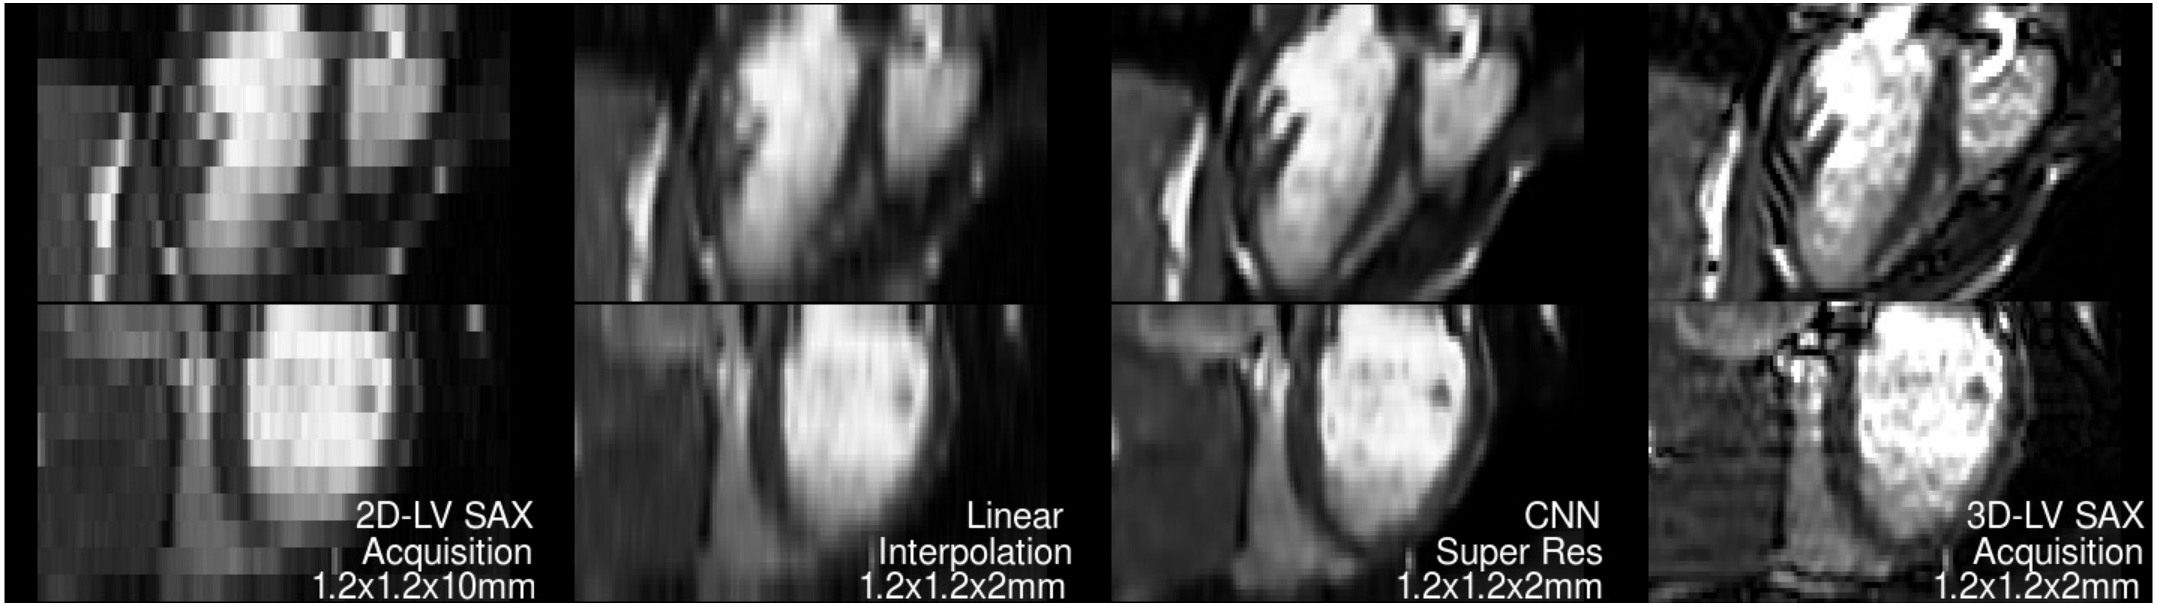

O. Oktay et al. IEEE TMI 2018

Application to fetal MR imaging

Fetal example:

1. Long acquisition times
2. Fetal motion and maternal breathing

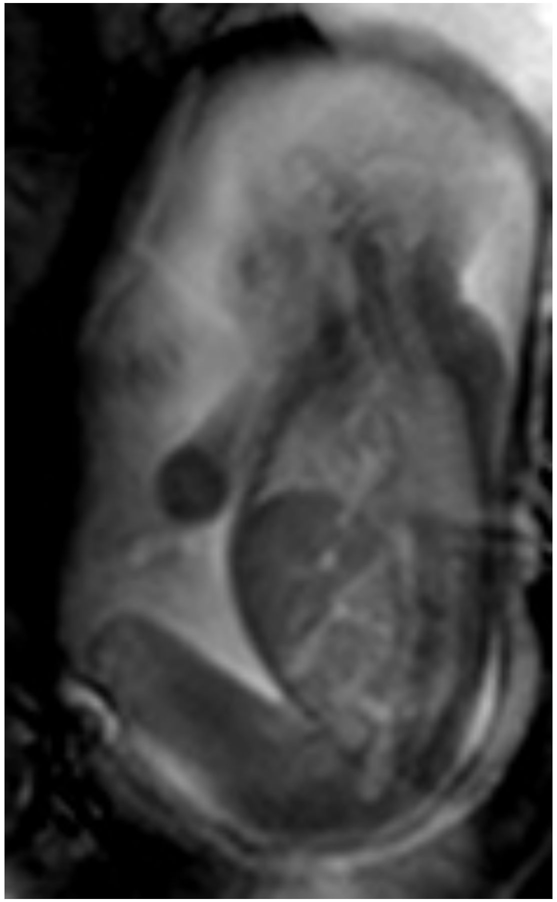

fast single-shot techniques are 2D acquisitions that freeze the motion in time
but ...

Application to fetal MR imaging

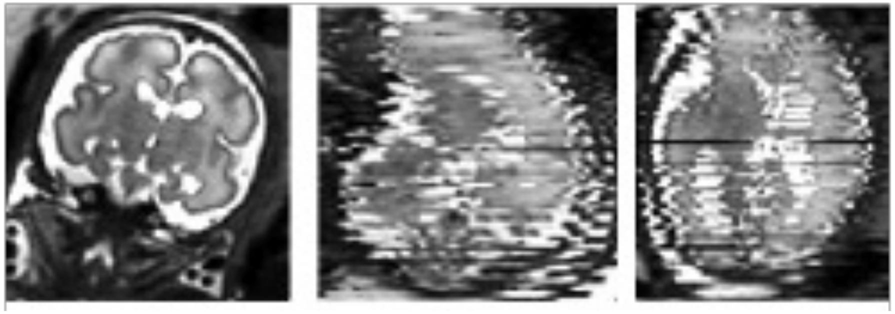

Application to fetal MR imaging

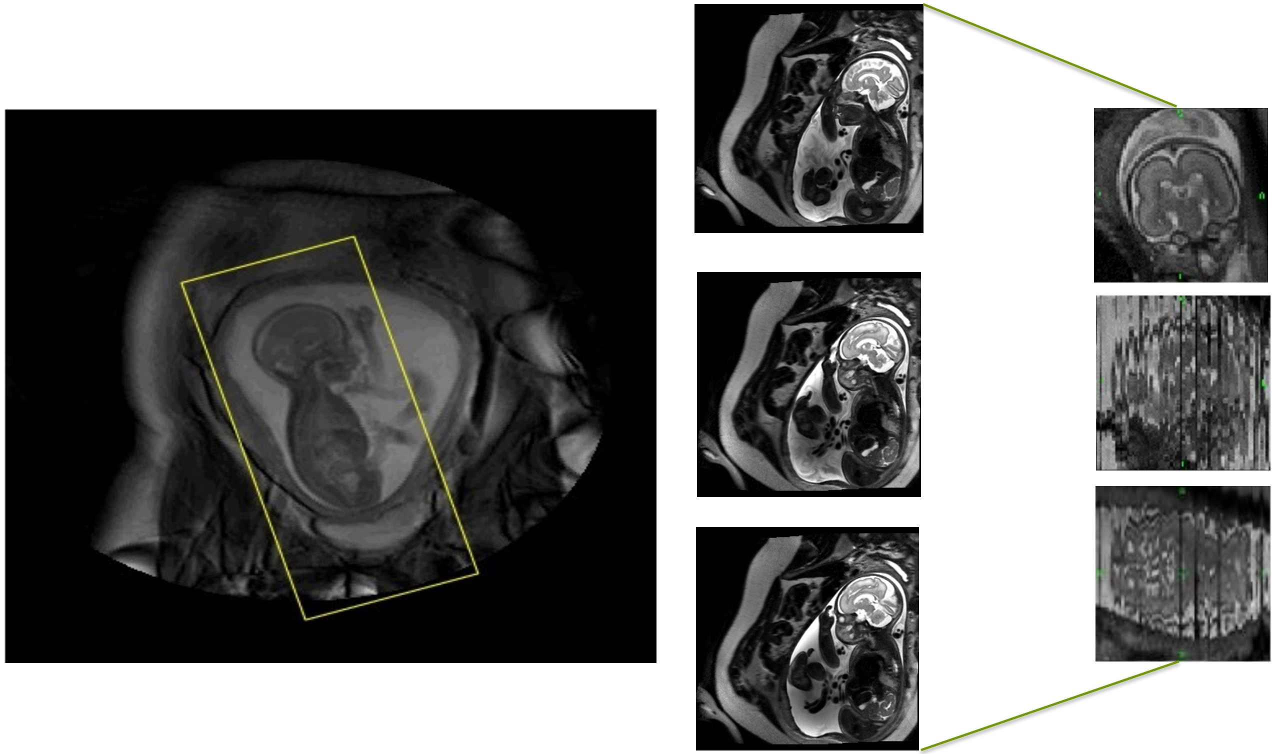

Application to fetal MR imaging

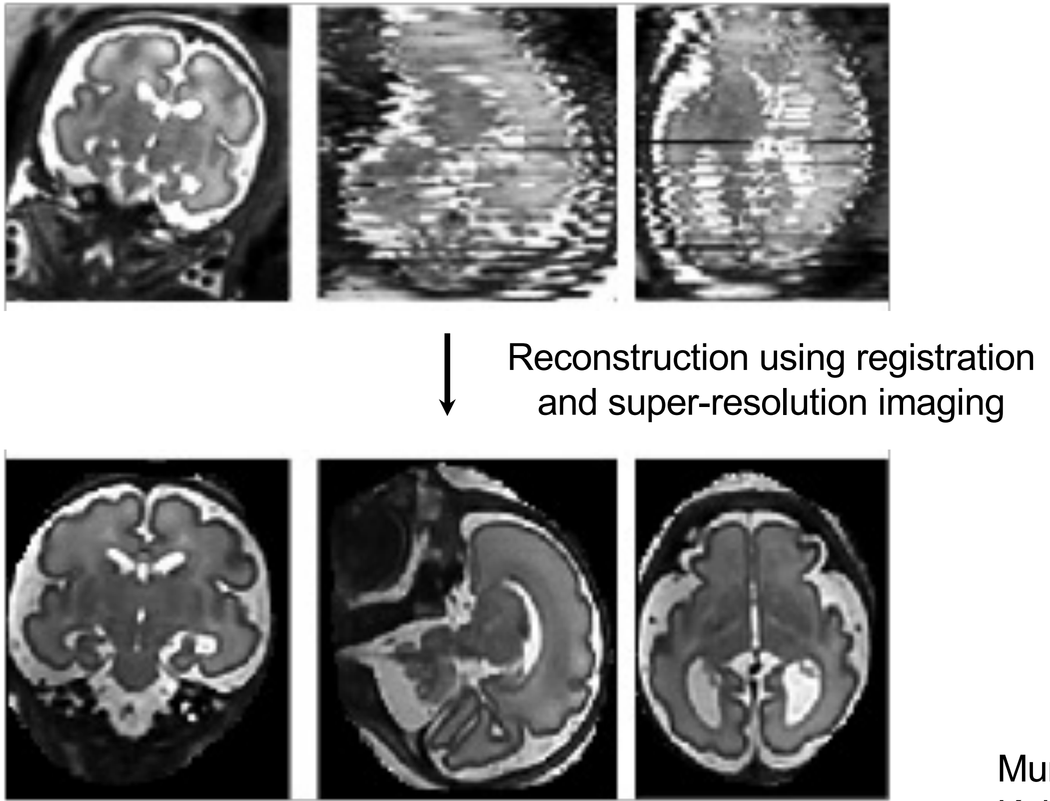

Murgasova et al., MEDIA, 2012
Kainz et al., IEEE TMI 2015
Alansary et al., IEEE TMI 2017

AI-enabled image recognition

- Potential applications:
  - Guidance: Assist inexperienced sonographers
  - Convenience: Automatically make a check list of visited planes
  - Reproducibility: Reduce variability between operators

Fetal brain standard planes
a: Transventricular plane
b: Transthalamic plane
c: Transcerebellar plane

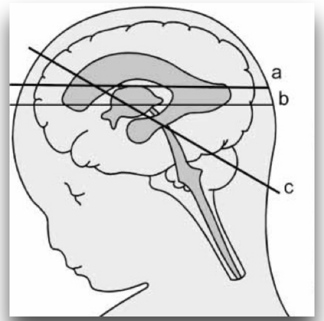

Image from *Ultrasound Obstet Gynecol*, 29: 109-116

AI-enabled image recognition: Automatic Scan Plane Detection

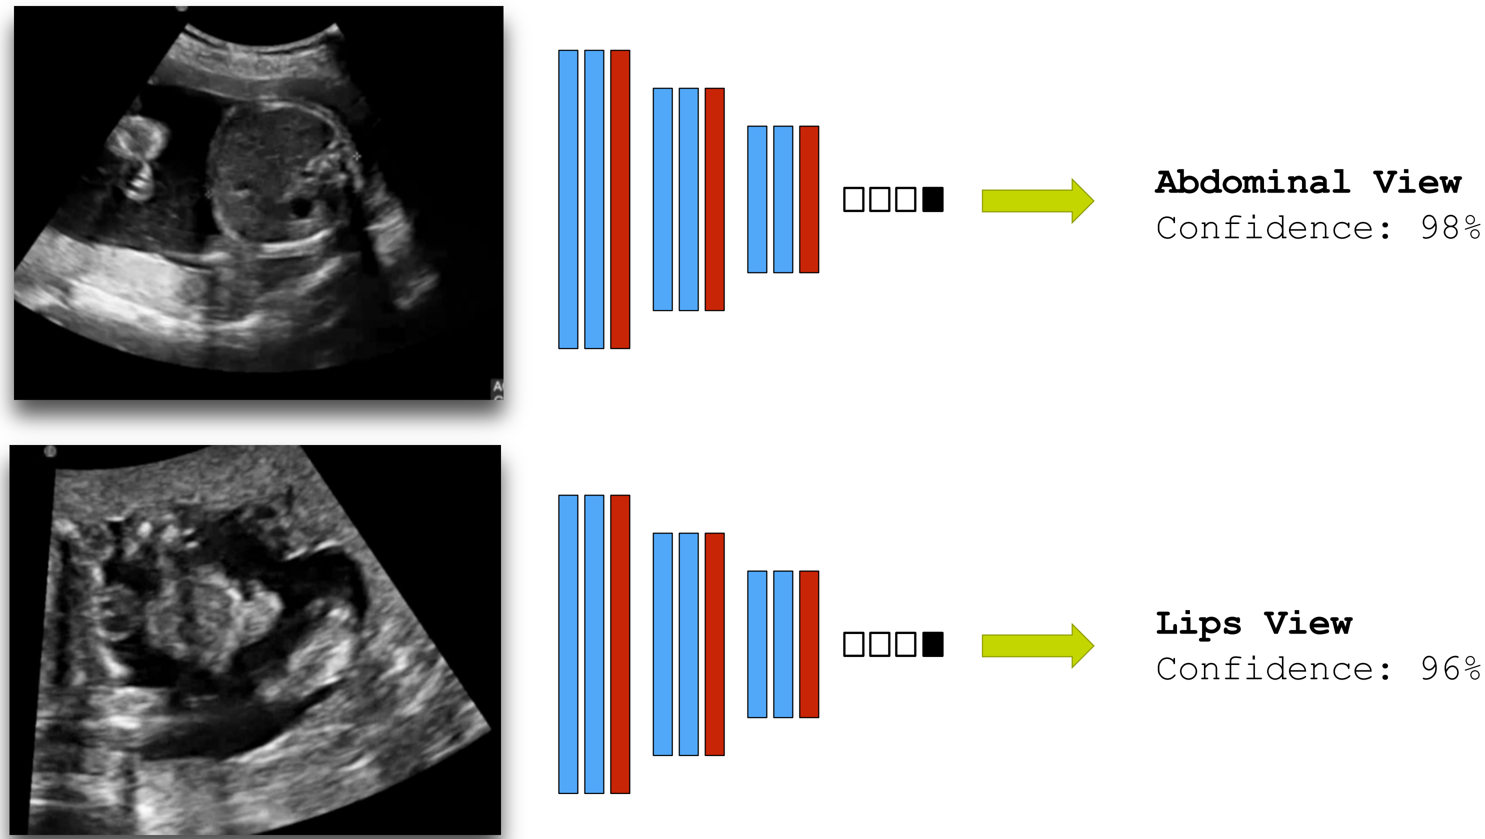

Goal: Do this in real-time on images straight from US machine

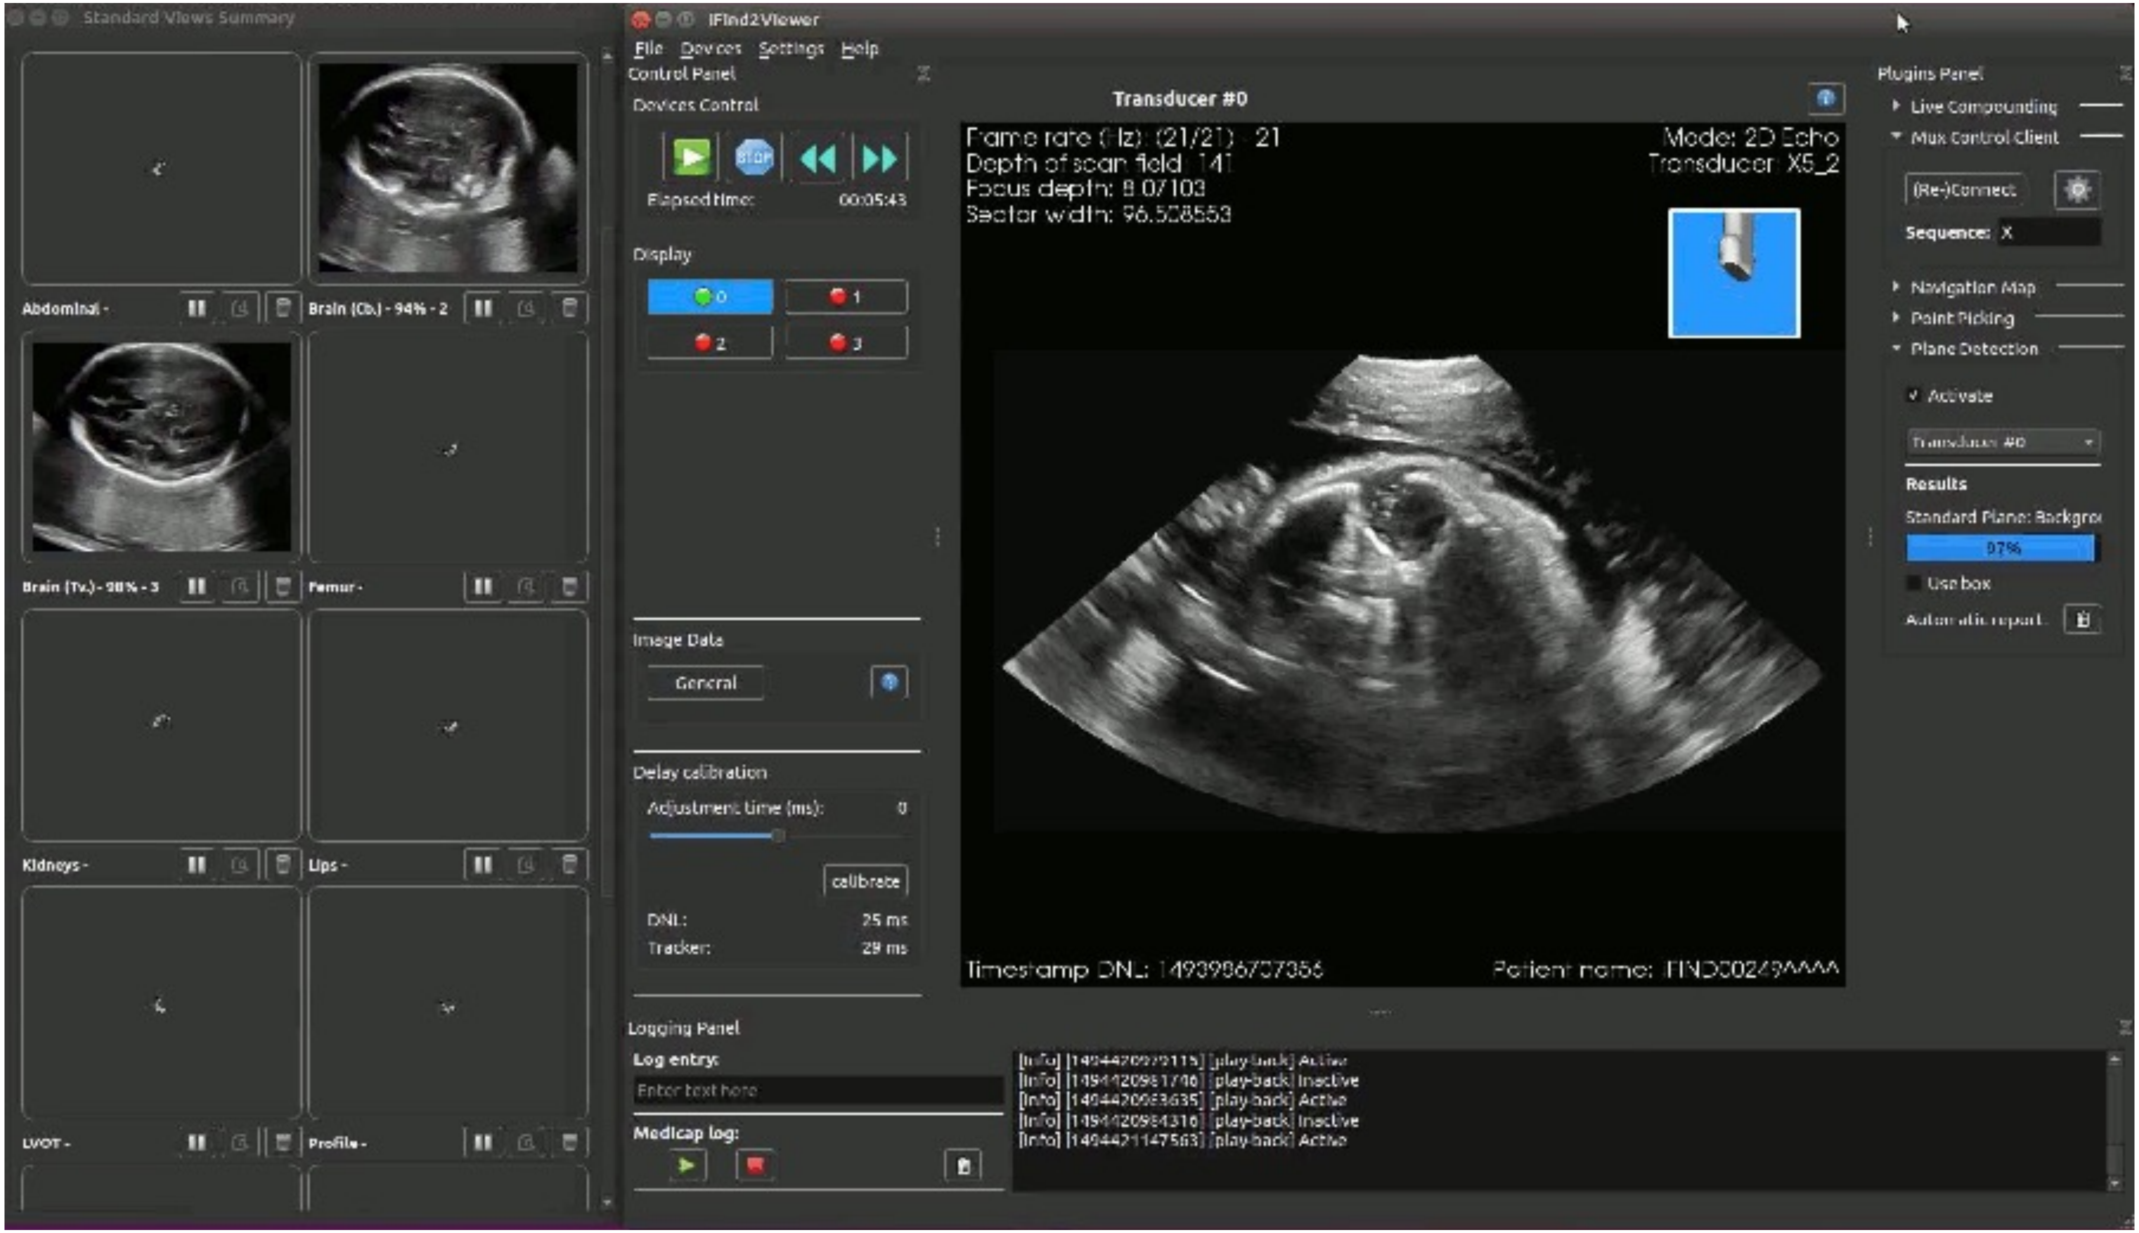
Standard Views Summary
IFind2Viewer
File Devices Settings Help
Control Panel
Devices Control
Elapsed time: 00:05:43
Display
Image Data
General
Delay calibration
Adjustment time (ms): 0
calibrate
DNL: 25 ms
Tracker: 29 ms
Transducer #0
Frame rate (Hz): (21/21) - 21
Depth of scan field: 141
Focus depth: 8.07103
Sector width: 96.508553
Mode: 2D Echo
Transducer: X5_2
Plugins Panel
Live Compounding
Mux Control Client
(Re-)Connect
Sequence: X
Navigation Map
Point Picking
Plane Detection
Activate
Results
97%
Use box
Abdominal -
Brain (Cb.) - 94% - 2
Brain (Tv.) - 98% - 3
Femur -
Kidneys -
Lips -
LVOT -
Profile -
Logging Panel
Log entry:
Enter text here
Medicap log:

Automatic Standard Scan
Plane Detection: Attention models

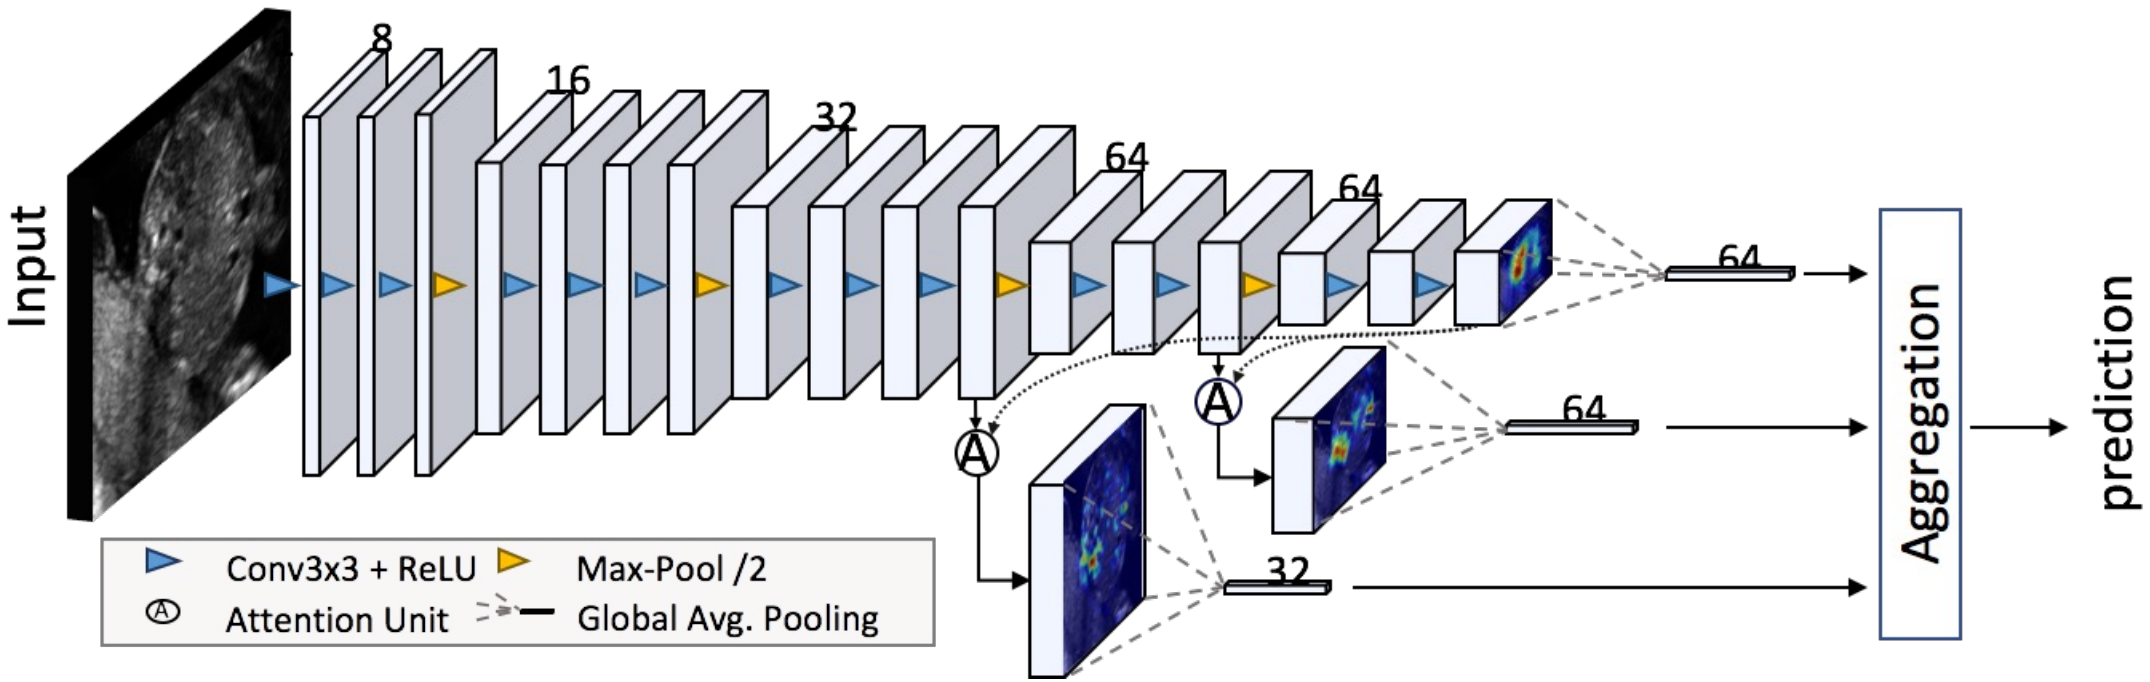

Schlemper et al. MedIA 2019

Automatic Standard Scan Plane Detection: Attention models

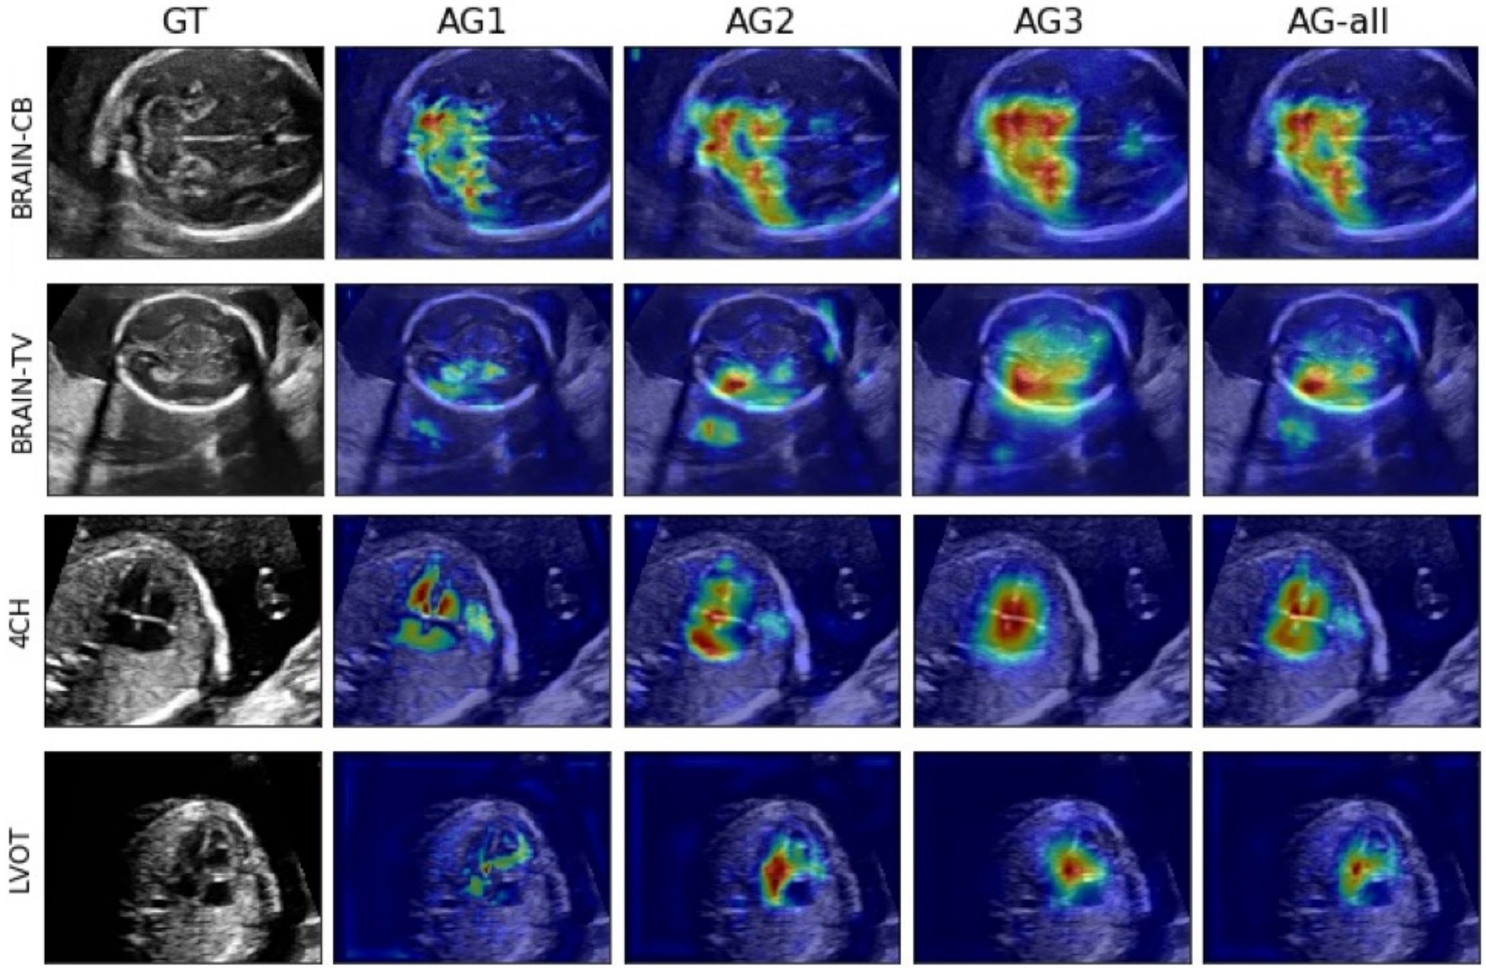

Schlemper et al. MedIA 2019

Automatic Standard Scan Plane Detection in 3D

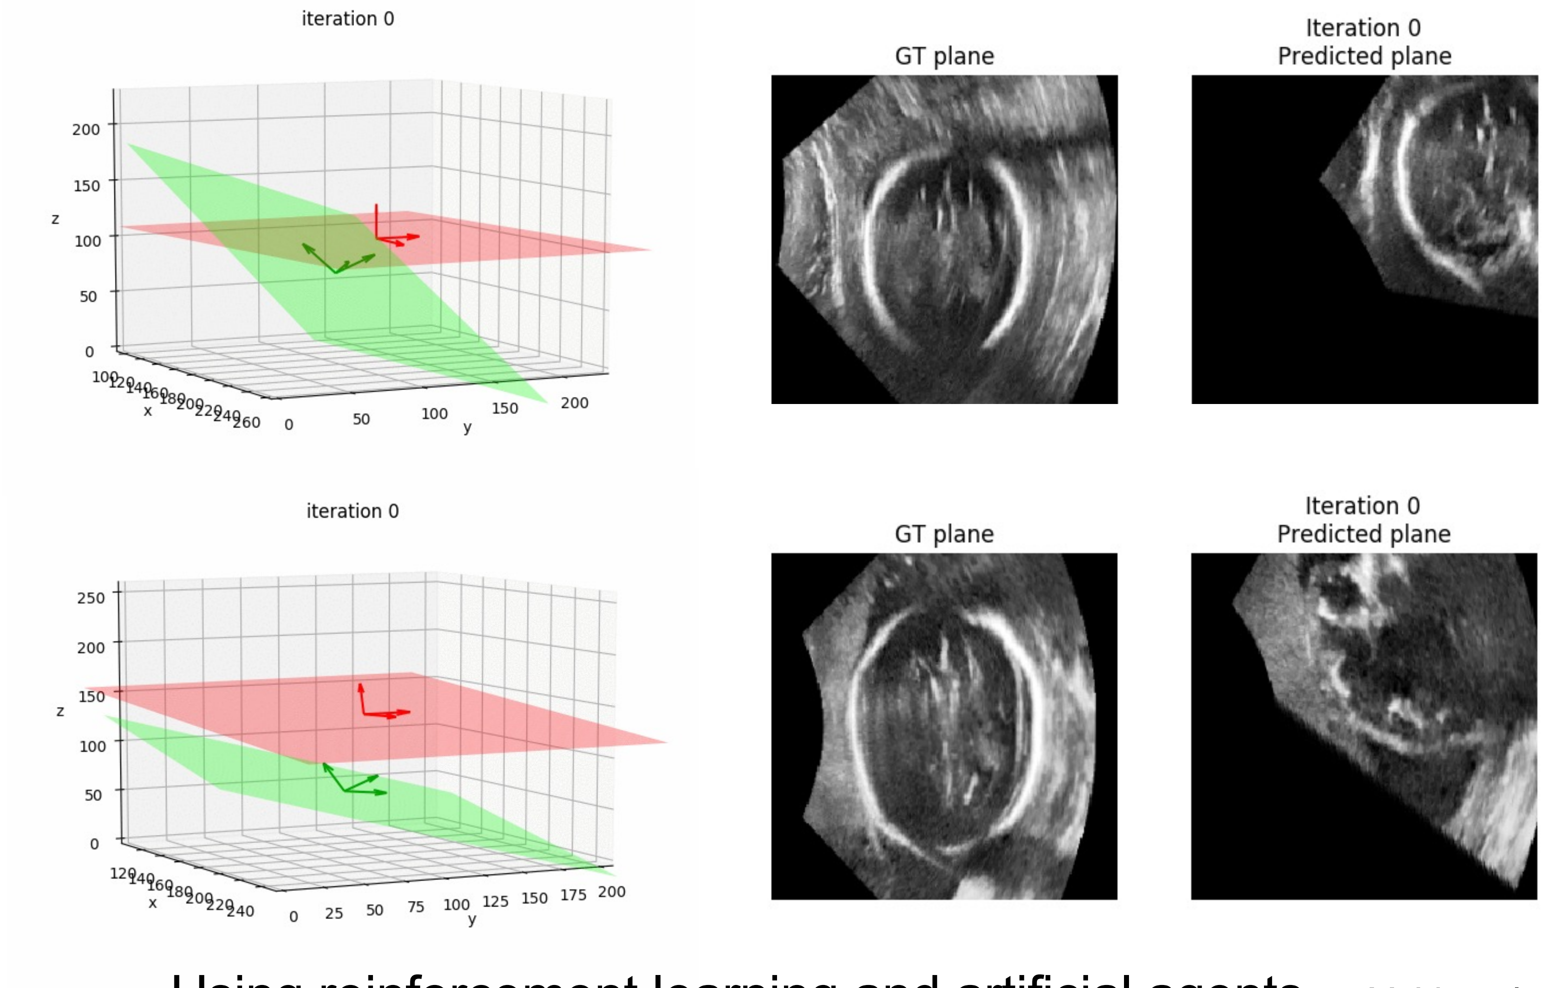

Using reinforcement learning and artificial agents

Y. Li et al. MICCAI 2018

AI-enabled image segmentation

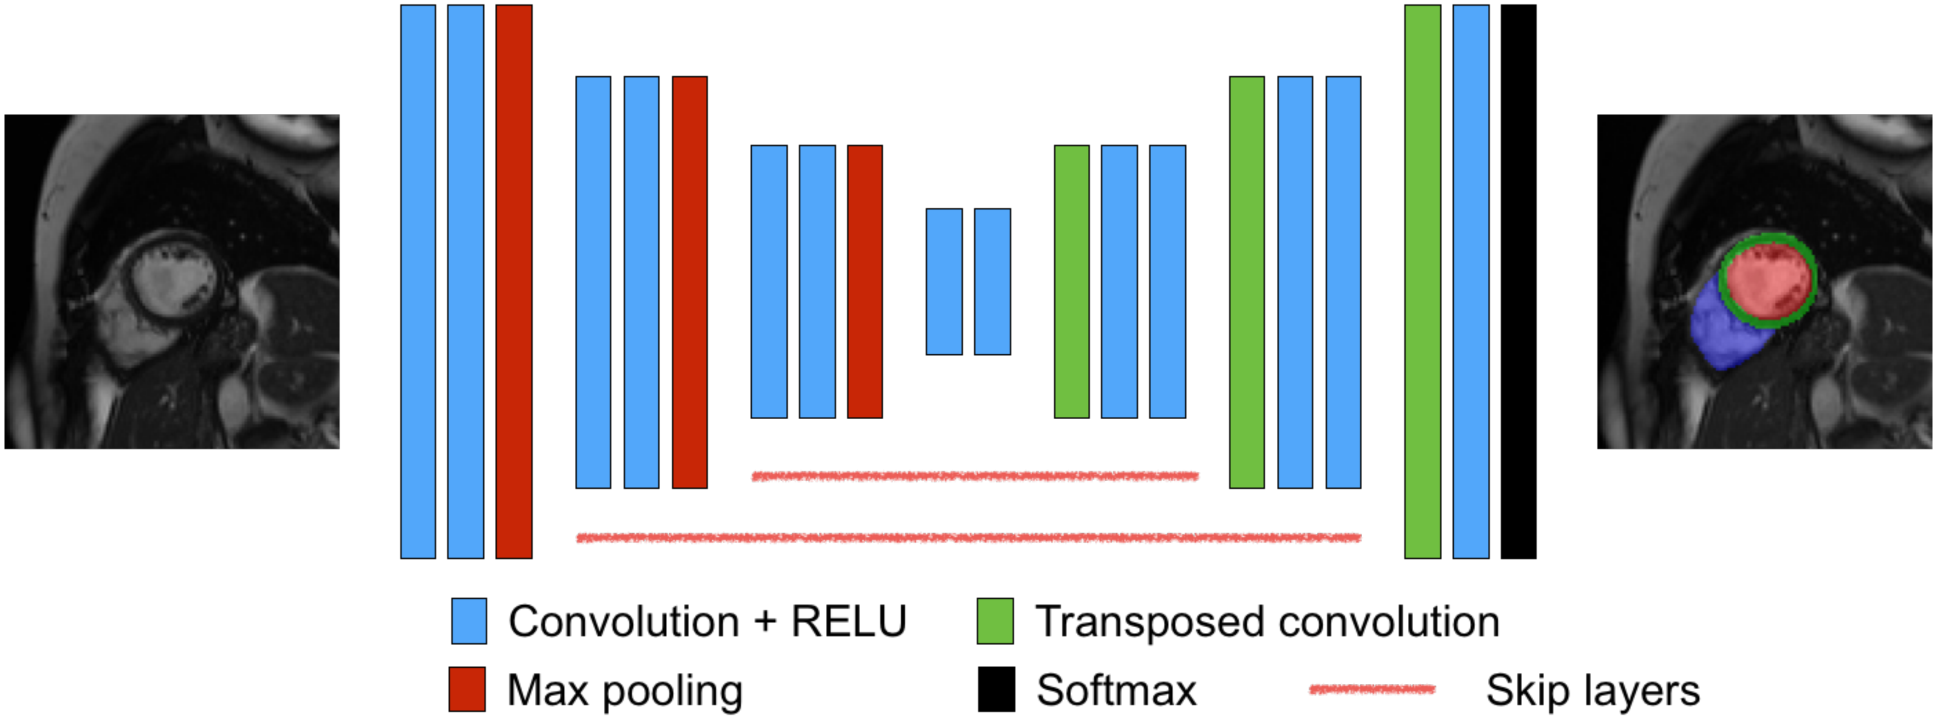

Bai et al., JCMR 2018

AI-enabled image segmentation

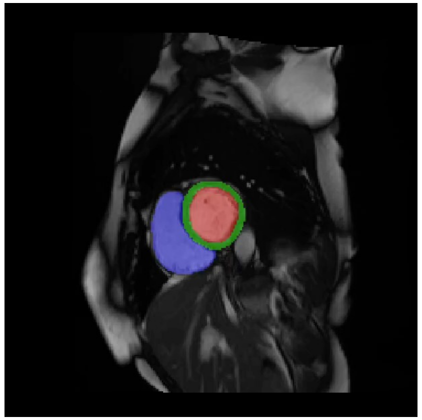

SA, basal

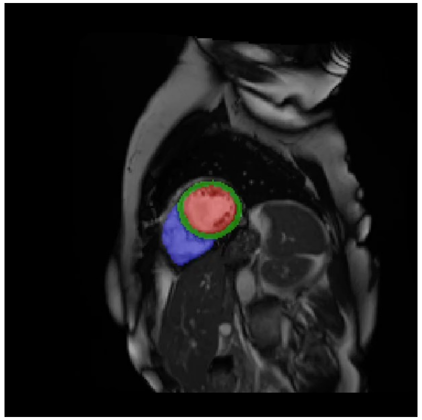

SA, mid-ventricular

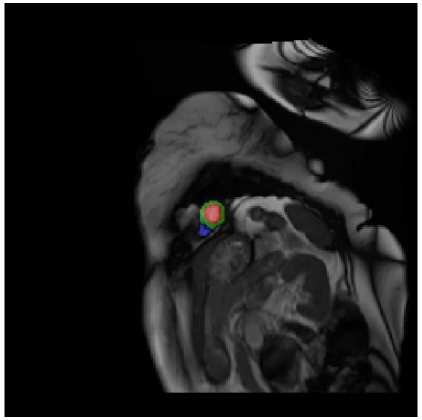

SA, apical

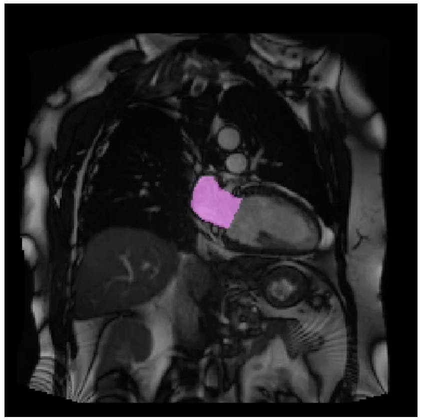

LA, 2 chamber

LA, 4 chamber

Bai et al., JCMR 2018

Large-scale population analysis

- In 2014, UK Biobank began the process of inviting back 100,000 of the original volunteers for brain, heart and body imaging.
- Imaging is done across several dedicated centres in the UK

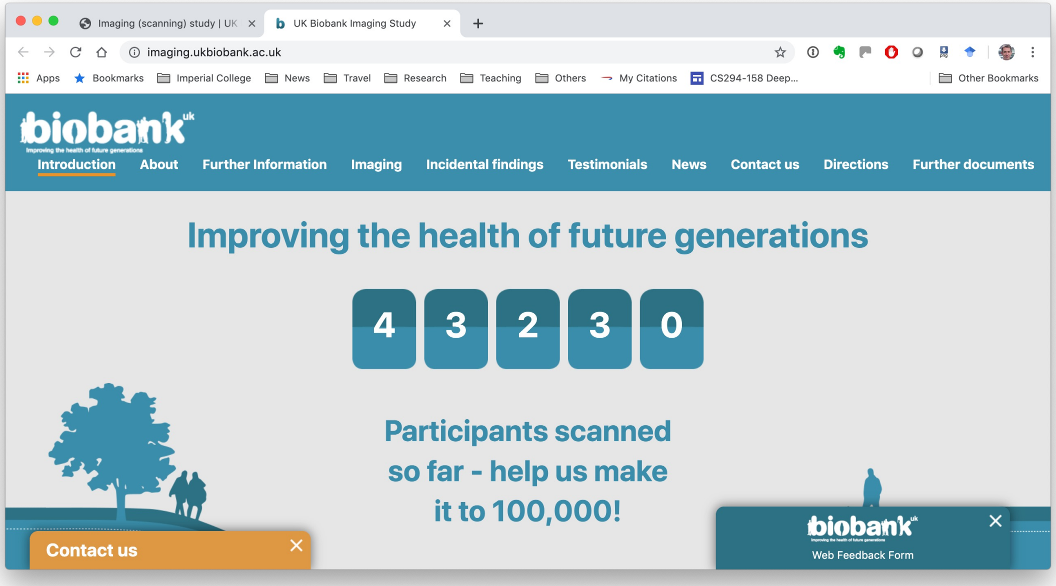

UK Biobank: Imaging

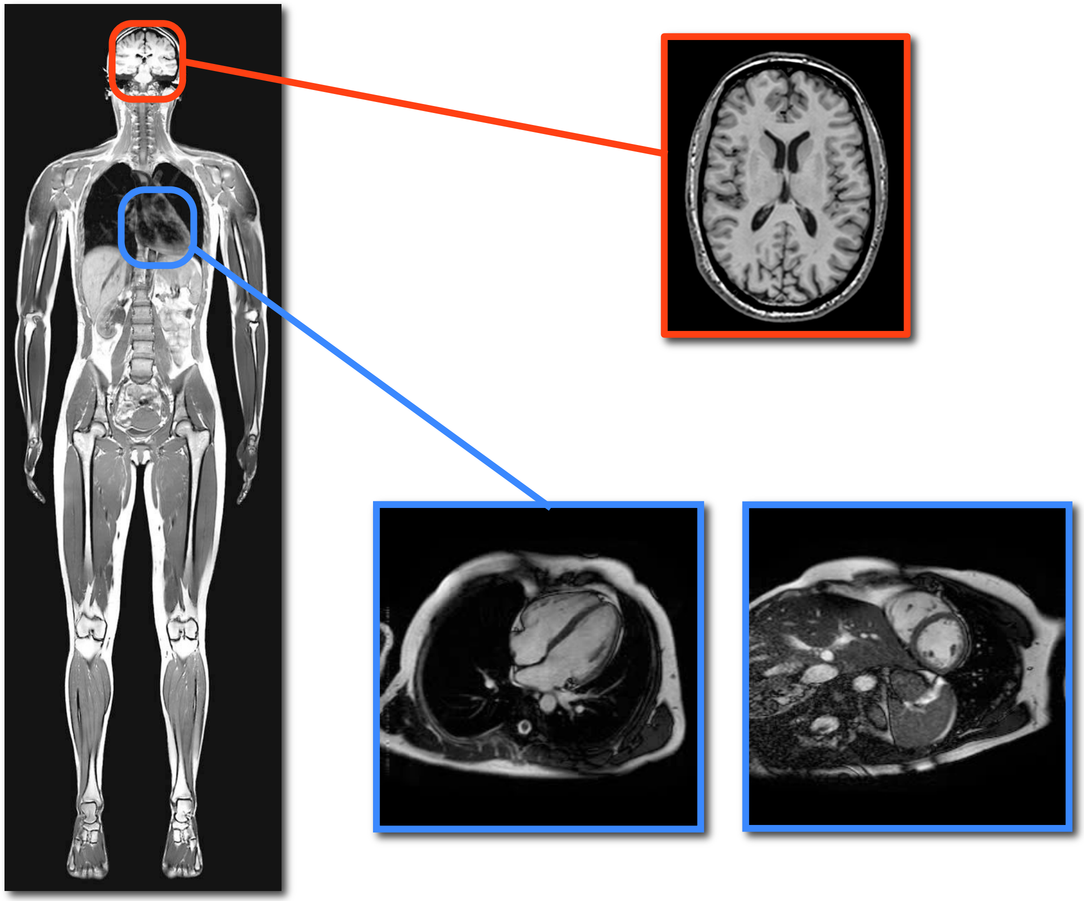

+

Lifestyle

Genetics

Clinical records

Large-scale population analysis

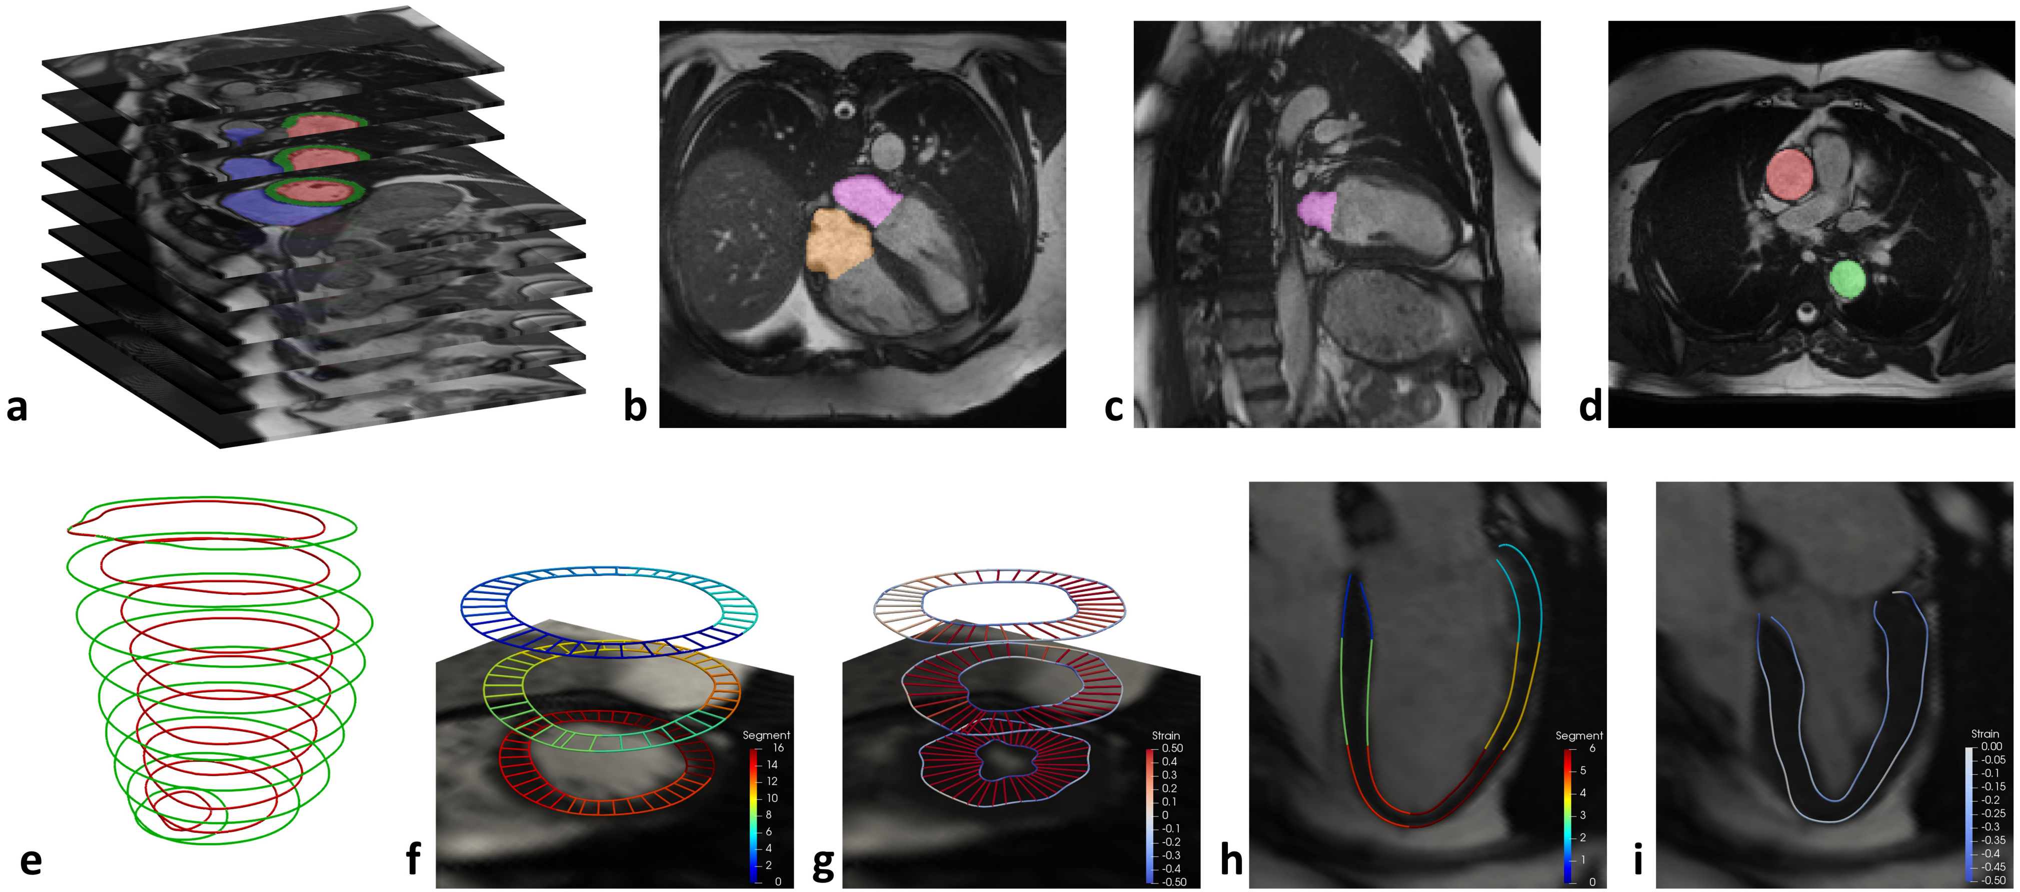

W. Bai et al., Nature Medicine, 2020

Cardiac IDPs from 26,893 subjects
Associations with sex and age

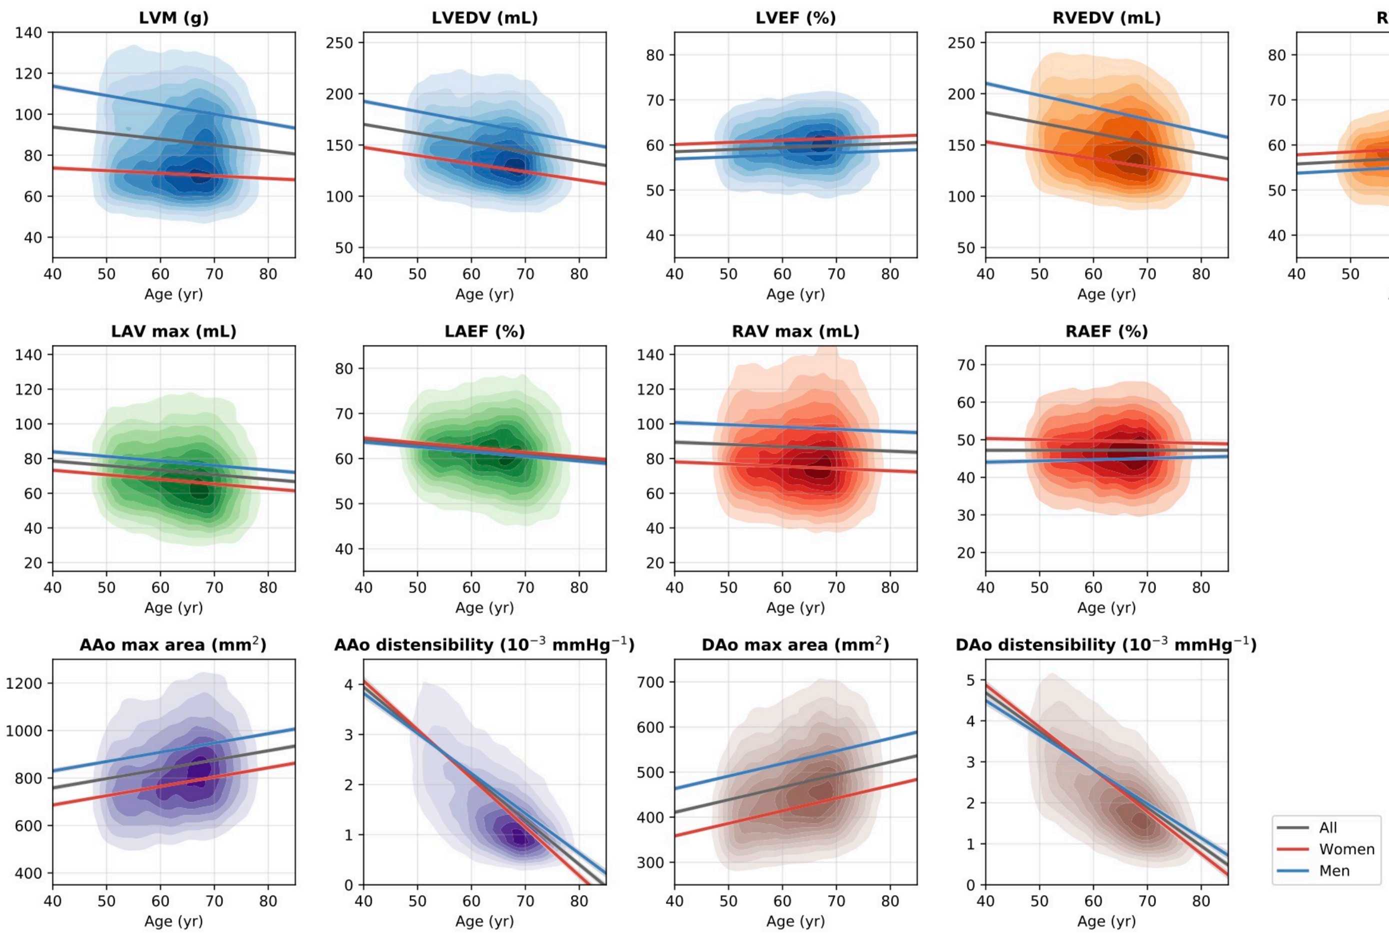

W. Bai et al., Nature Medicine, 2020

Phenome-wide association study

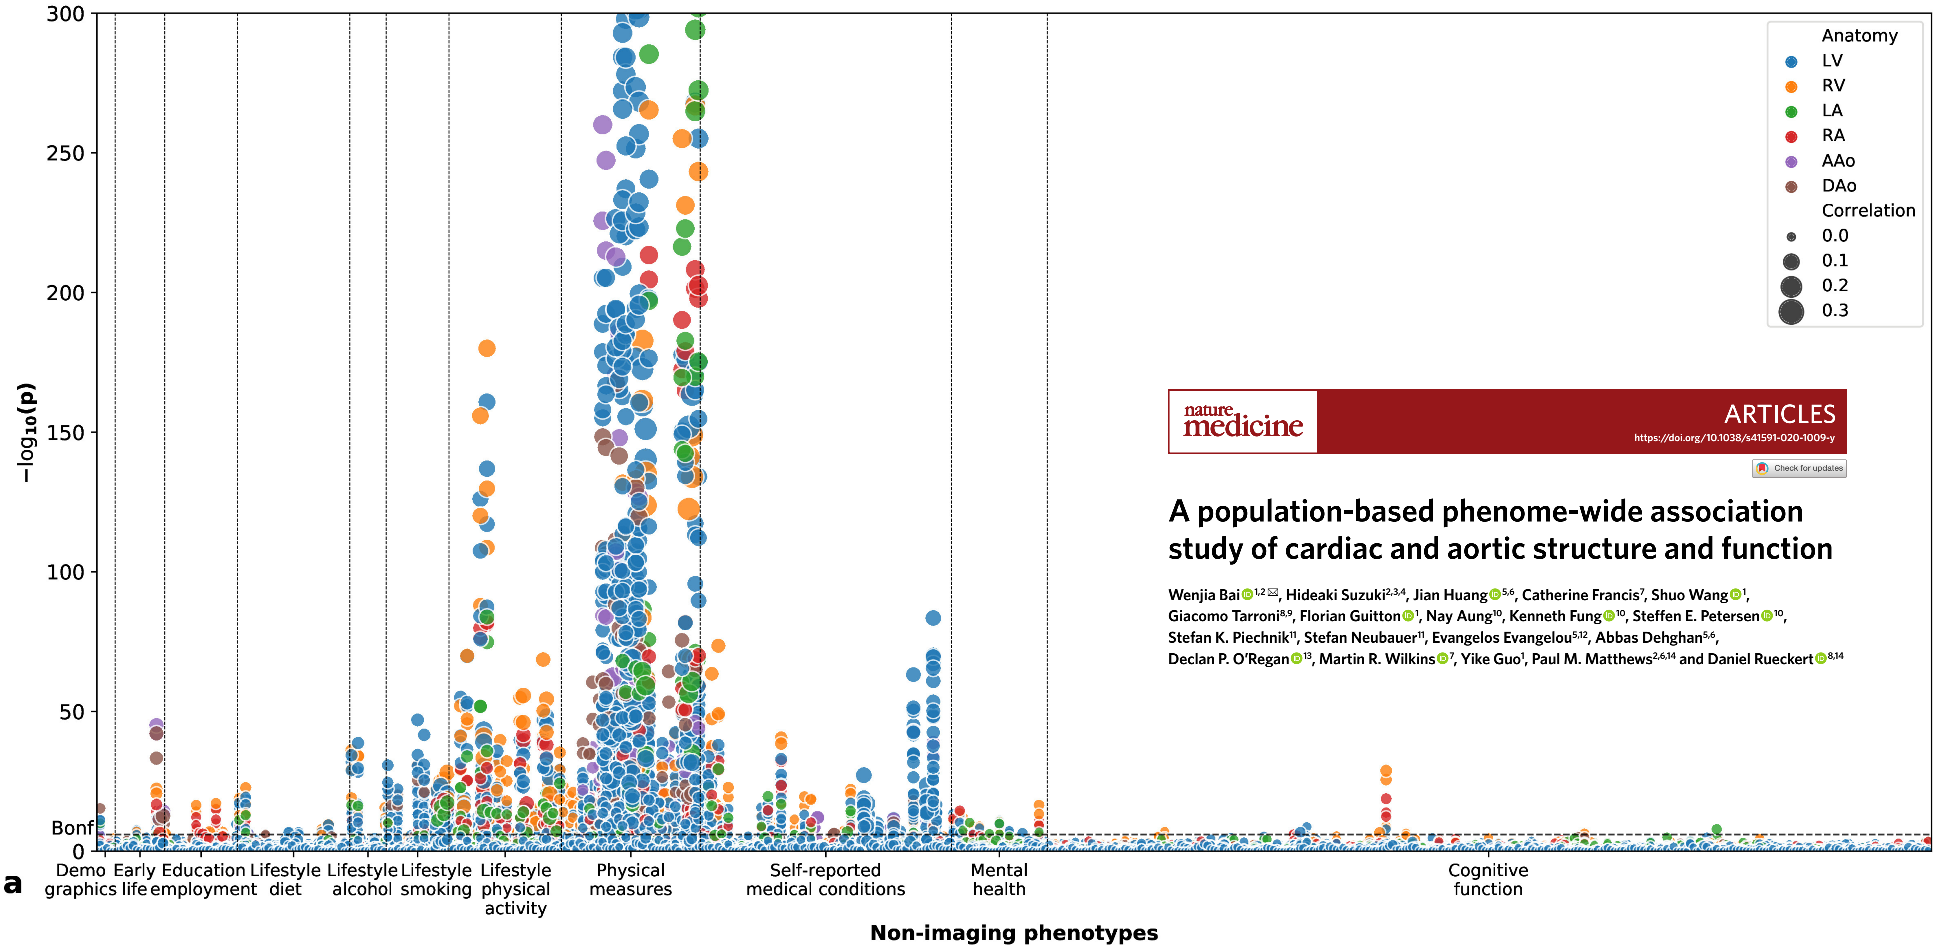

nature medicine

ARTICLES

https://doi.org/10.1038/s41591-020-1009-y

Check for updates

A population-based phenome-wide association study of cardiac and aortic structure and function

Wenjia Bai[1,2], Hideaki Suzuki[2,3,4], Jian Huang[5,6], Catherine Francis[7], Shuo Wang[1], Giacomo Tarroni[8,9], Florian Guitton[1], Nay Aung[10], Kenneth Fung[10], Steffen E. Petersen[10], Stefan K. Piechnik[11], Stefan Neubauer[11], Evangelos Evangelou[5,12], Abbas Dehghan[5,6], Declan P. O'Regan[13], Martin R. Wilkins[7], Yike Guo[1], Paul M. Matthews[2,6,14] and Daniel Rueckert[8,14]

W. Bai et al., Nature Medicine, 2020

AI for decision support: Survival prediction

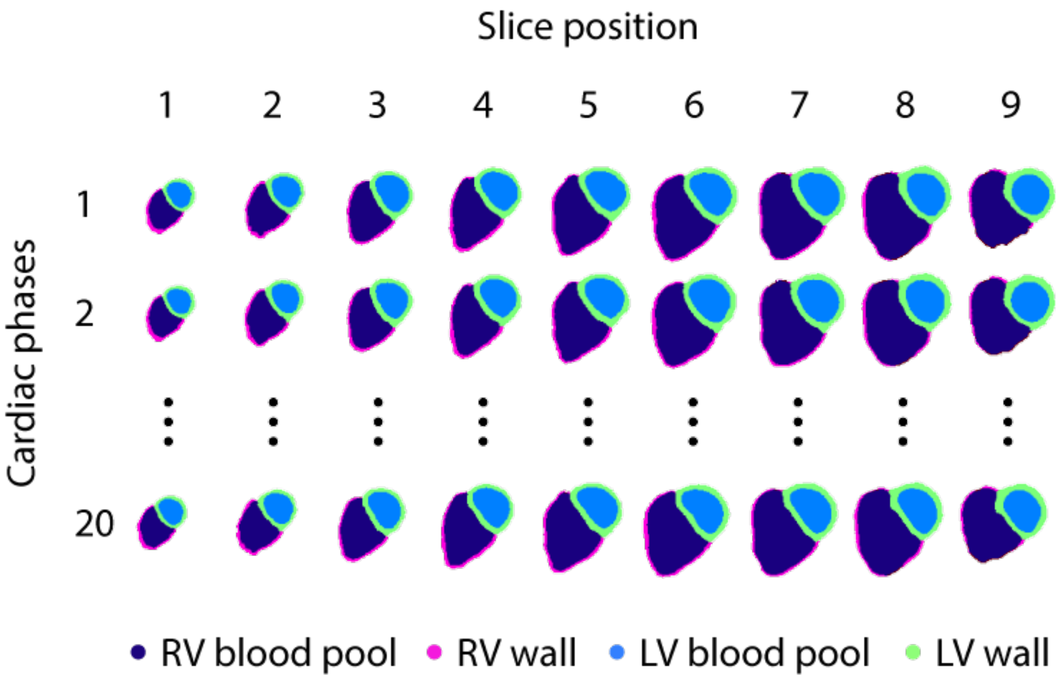

Bello et al. Nature Machine Intelligence 2019

AI for decision support: Survival prediction

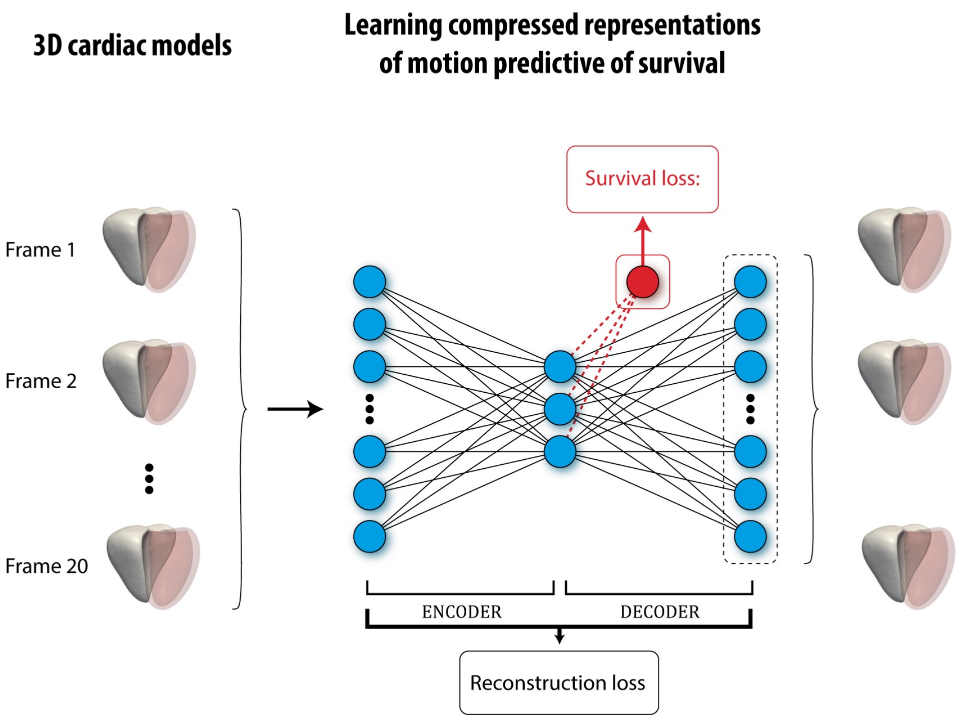

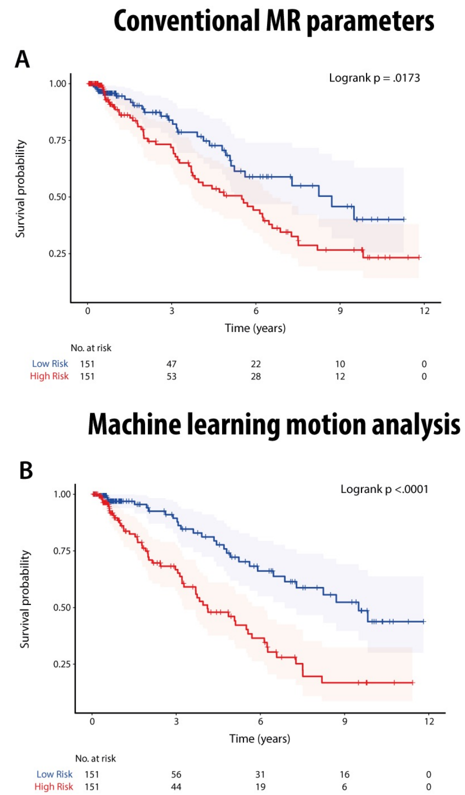

Imperial College London

MRC London Institute of Medical Sciences

Bello et al. Nature Machine Intelligence 2019

AI has the potential to revolutionize medicine and healthcare

But what are the challenges?

Lack of sufficient data: Bias and fairness

RESEARCH ARTICLE Obermeyer et al., Science 2019

ECONOMICS

Dissecting racial bias in an algorithm used to manage the health of populations

Ziad Obermeyer[1,2]*, Brian Powers[3], Christine Vogeli[4], Sendhil Mullainathan[5]*†

Lack of sufficient data: Variability

- How to deal with variability?
  - Population variability (normal vs pathologies)
  - Image acquisition variability (e.g. due to scanner differences)

Data during training

Data during deployment

Lack of sufficient data: Variability

- How to deal with variability?
  - Population variability (normal vs pathologies)
  - Image acquisition variability (e.g. due to scanner differences)

Different hardware

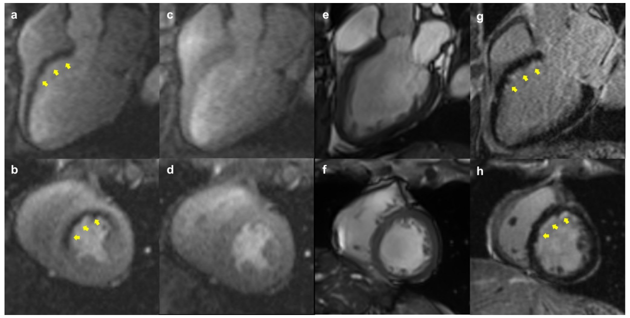

Stress Rest Cine LGE

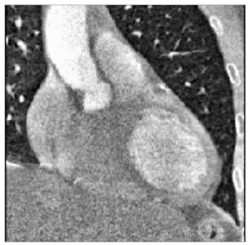

CT

MR

Lack of sufficient data: Domain shift

Source (S)	
Domain:	$D_S = \{\mathcal{X}_S, P(X_S)\}$
Task:	$T_S = \{\mathcal{Y}_S, f'_S : \mathcal{X}_S \mapsto \mathcal{Y}_S\}$
Given:	(X_S, Y_S)
	$X_S = \{x_{S1}, ..., x_{Sn}\}, x_{Si} \in \mathcal{X}_S$
	$Y_S = \{y_{S1}, ..., y_{Sn}\}, y_{Si} \in \mathcal{Y}_S$
Learn:	$f_S \approx f'_S$
	$f_S(x) \approx P_S(y\|x)$

Target (T)	
Domain:	$D_T = \{\mathcal{X}_T, P(X_T)\}$
Task:	$T_T = \{\mathcal{Y}_T, f'$
<u>Here:</u>	$\mathcal{Y}_T = \mathcal{Y}$
Domain Shift:	
	$P(X$
	$P_T(y\|x$ $\quad$ $y\|x)$
	$f'_T \neq f'_S$

Domain Shift !

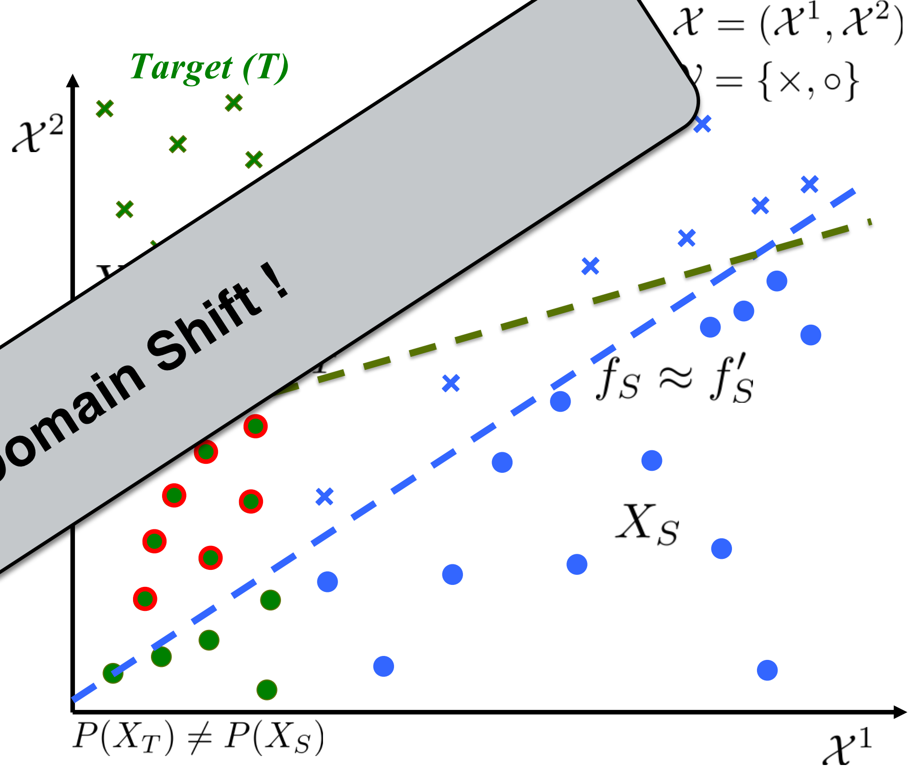

Lack of sufficient data: How to address?

Learning domain invariant features

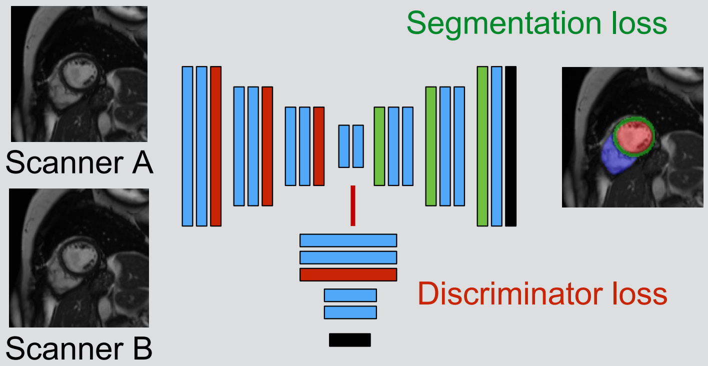

K. Kamnitsas et al. IPMI 2017

Data augmentation

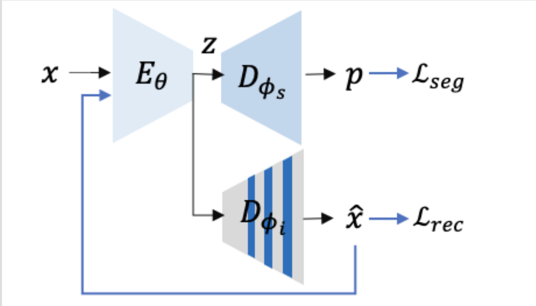

C. Chen et al. MICCAI 2020, 2021, 2022

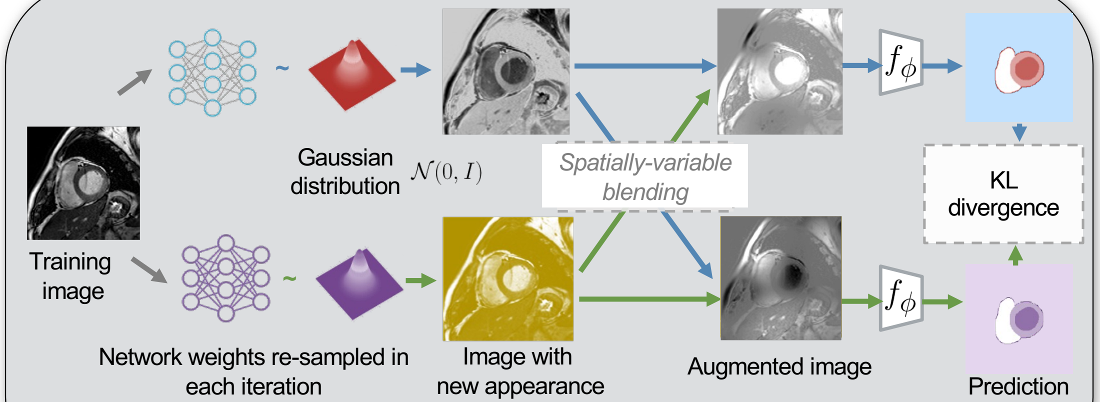

C. Ouyang et al., IEEE TMI in press

Privacy-preserving AI/ML

Access to large datasets during training is critical …

… but how do we ensure privacy?


nature machine intelligence

PERSPECTIVE

https://doi.org/10.1038/s42256-020-0186-1

Check for updates

Secure, privacy-preserving and federated machine learning in medical imaging

Georgios A. Kaissis[1,2,3], Marcus R. Makowski[1], Daniel Rückert[2] and Rickmer F. Braren[1]

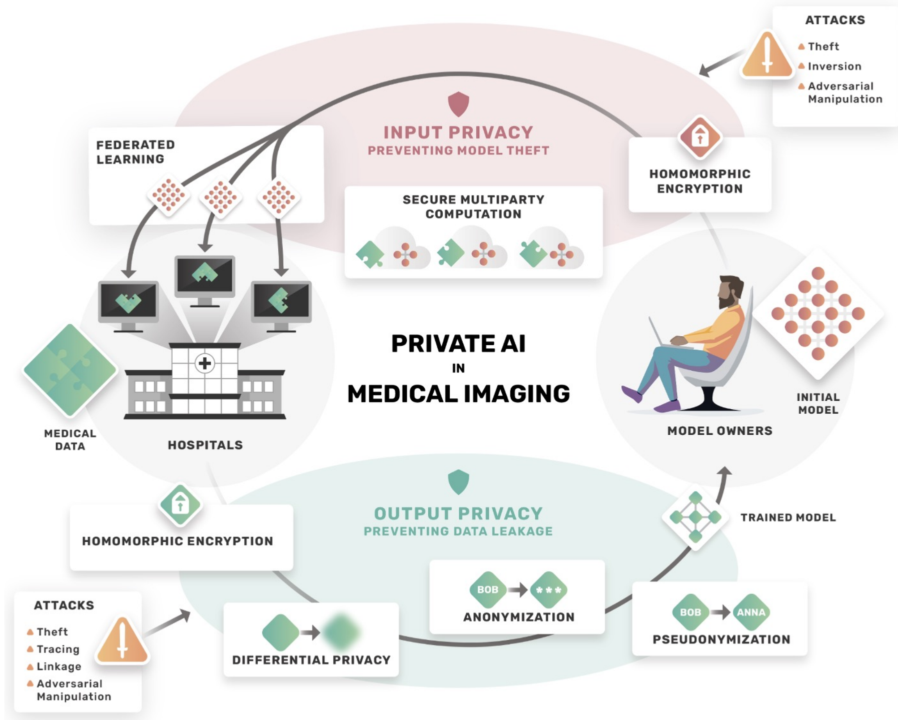

Federated learning

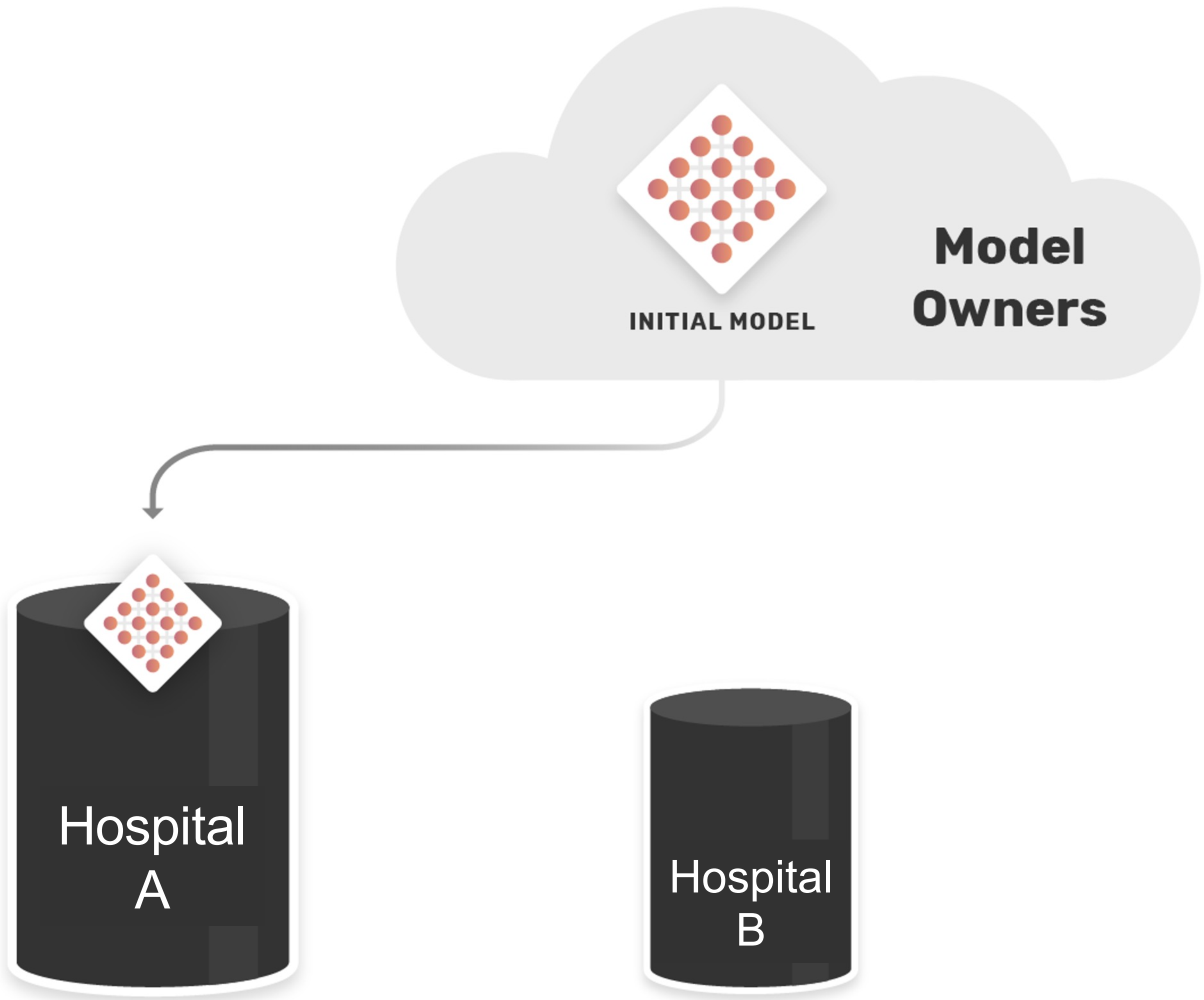

Federated learning

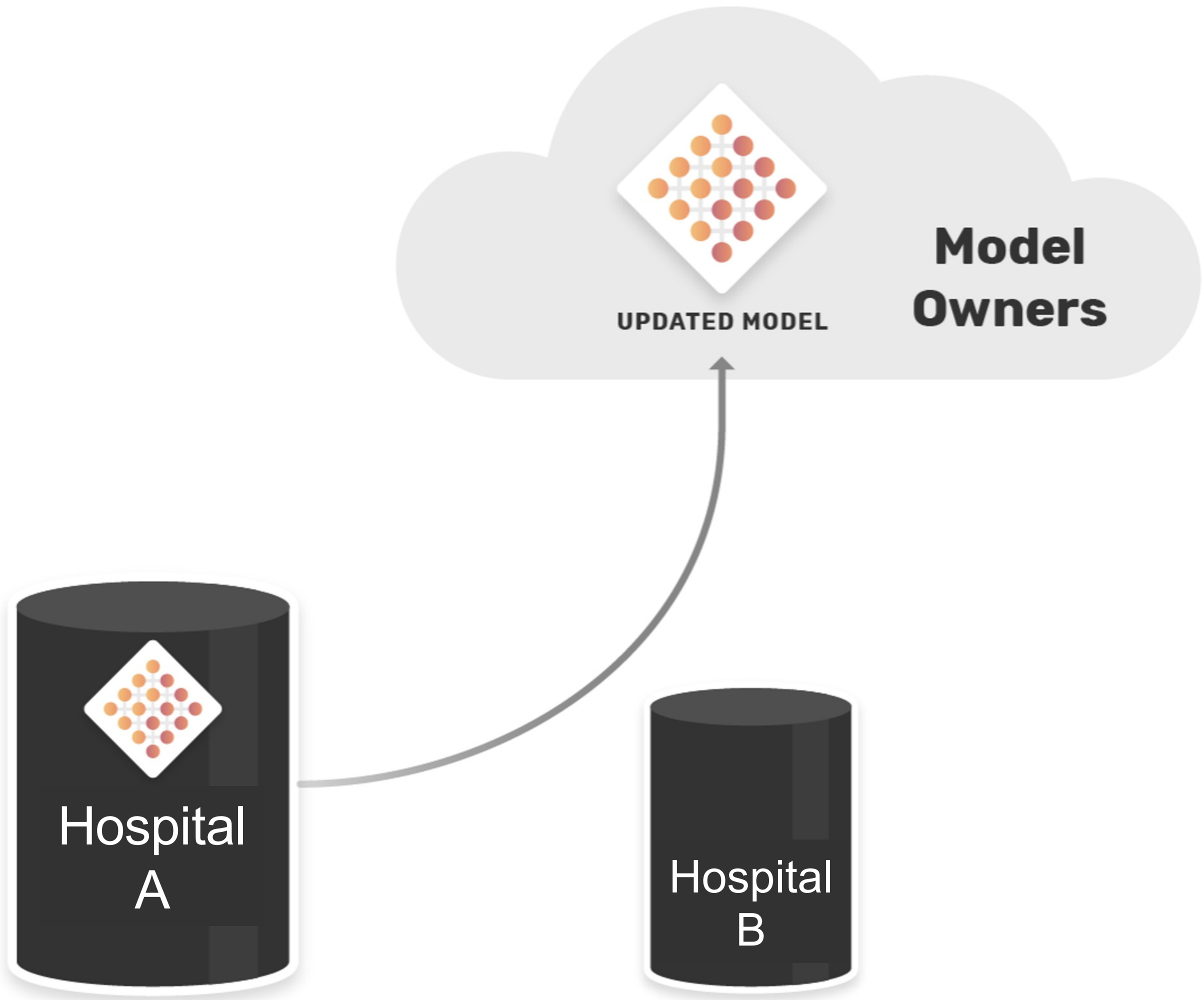

Federated learning

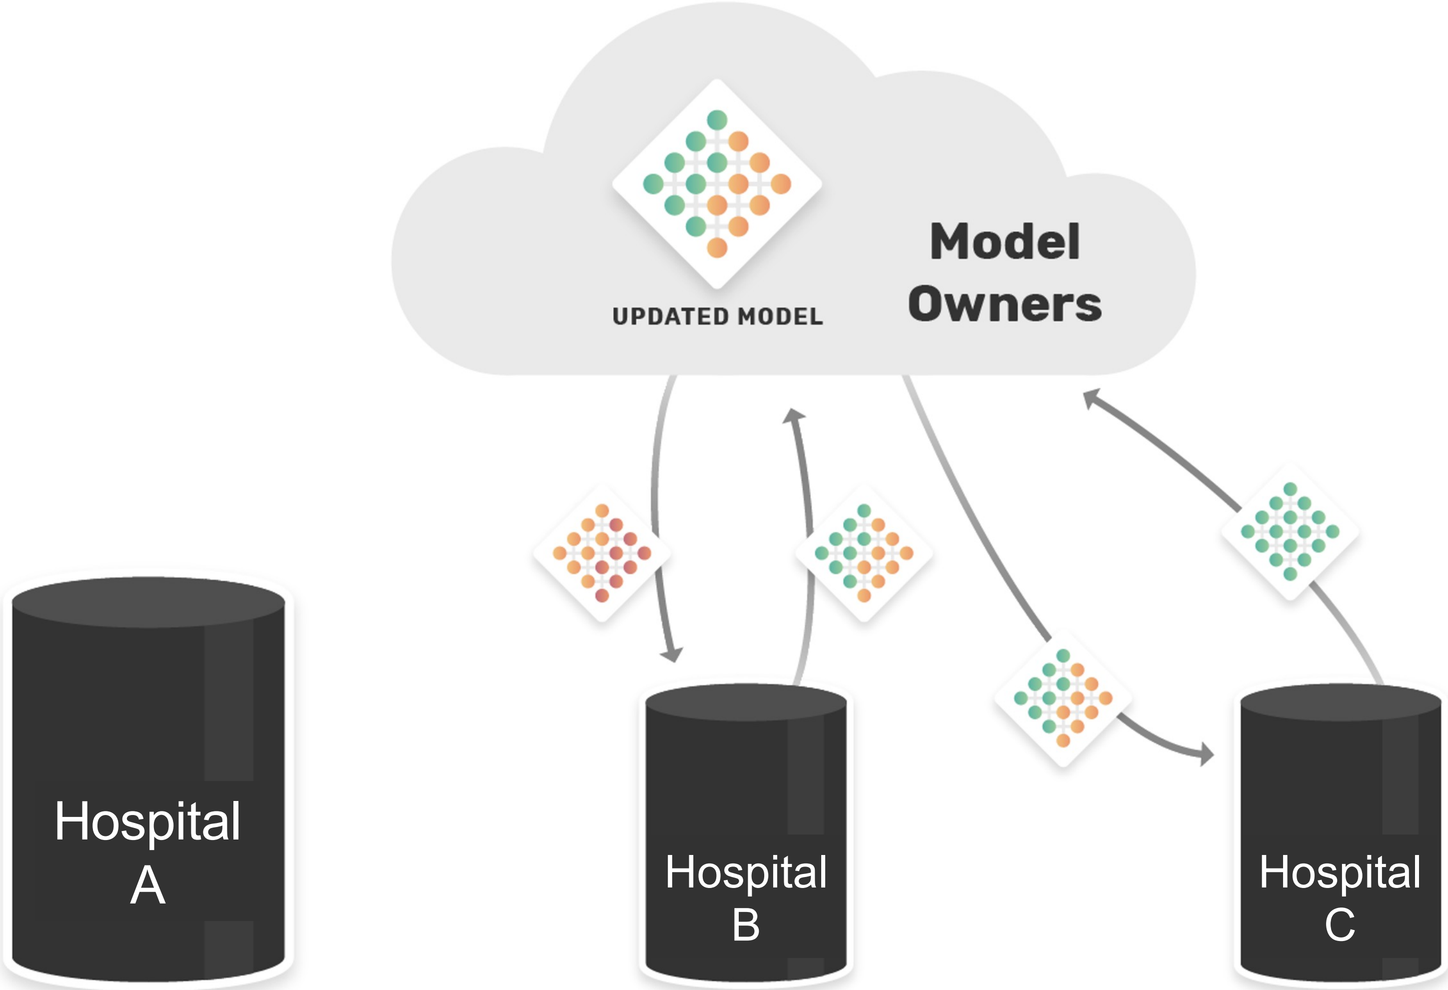

But federated learning is not enough!

Privacy-Centred attacks:

- Attempt to disclose information participants did not consent to disclosing
- Examples include:
  - Membership
  - Sensitive attributes
  - Training records
  - Reconstruction
  - etc.

Utility-Centred attacks:

- Attempt to subvert the protocol and alter the utility of the model
- Examples include:
  - Crafting malicious data or updates
  - Hidden collateral tasks
  - etc.

But federated learning is not enough!

Kaissis et al. Nature Machine Intelligence, 2021

Bring privacy-preserving machine learning to clinical routine

nature machine intelligence

ARTICLES

https://doi.org/10.1038/s42256-021-00337-8

Check for updates

End-to-end privacy preserving deep learning on multi-institutional medical imaging

Georgios Kaissis[1,2,3,4,13], Alexander Ziller[1,2,4,13], Jonathan Passerat-Palmbach[3,4,5], Théo Ryffel[4,6,7], Dmitrii Usynin[1,2,3,4], Andrew Trask[4,8], Ionésio Lima Jr[4,9], Jason Mancuso[4,10], Friederike Jungmann[1], Marc-Matthias Steinborn[11], Andreas Saleh[11], Marcus Makowski[1], Daniel Rueckert[2,3] and Rickmer Braren[1,12] ✉

Privacy-preserving machine learning: Differential privacy

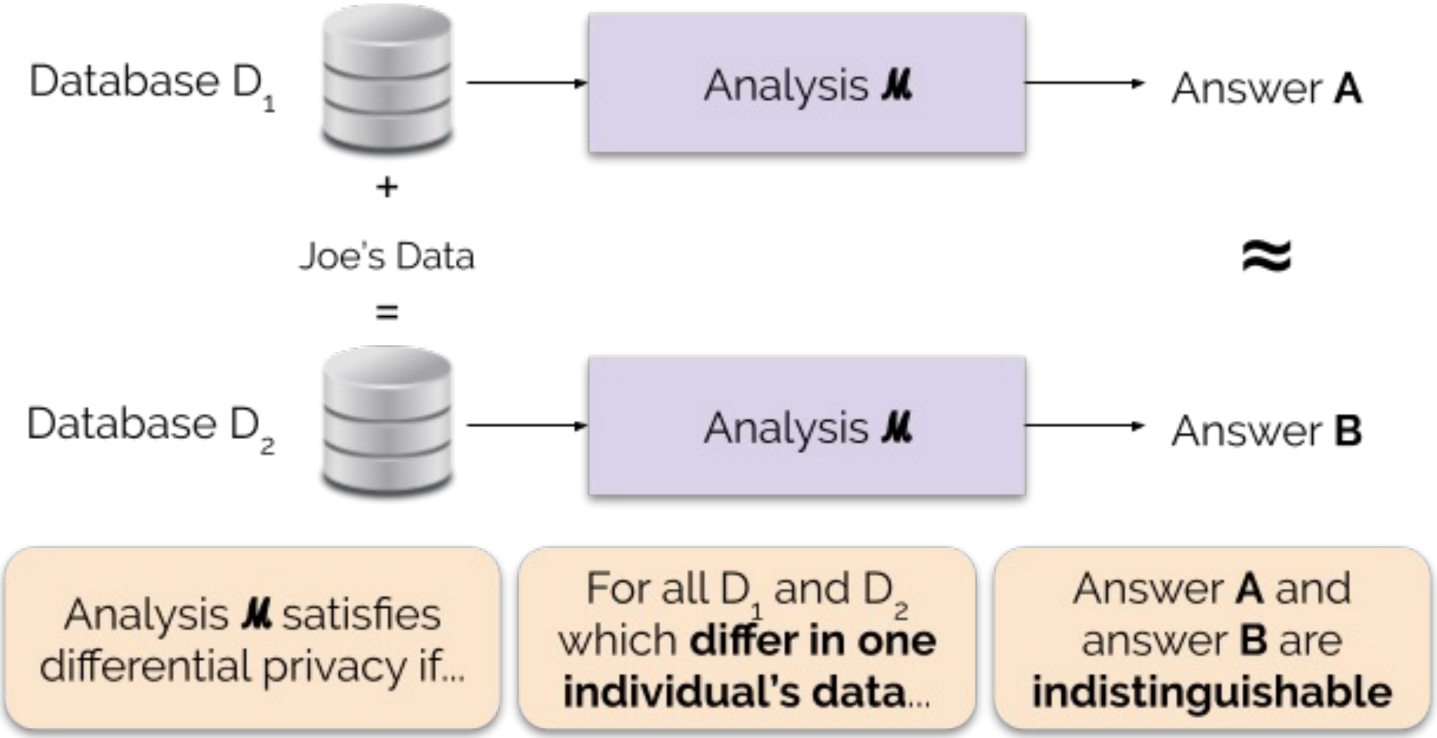

Figure from https://www.nist.gov/blogs/cybersecurity-insights/differential-privacy-privacy-preserving-data-analysis-introduction-our

Privacy-preserving machine learning: Differential privacy

Probability of seeing output O on input D_1

Probability of seeing output O on input D_2

$$\frac{\mathbf{Pr}[\mathcal{M}(D_1) \in O]}{\mathbf{Pr}[\mathcal{M}(D_2) \in O]} \leq e^{\varepsilon}$$

Indistinguishability: bounded ratio of probabilities

Figure from https://www.nist.gov/blogs/cybersecurity-insights/differential-privacy-privacy-preserving-data-analysis-introduction-our

Privacy-preserving machine learning: Differentially private stochastic gradient descent

- Algorithm:
  1. Compute gradients for each individual sample (they represent independent clients)
  2. Clip the calculated gradients to obtain a known sensitivity
  3. Add the noise scaled by the sensitivity from step 2
  4. Perform the gradient descent step

Algorithm 1 Differentially private SGD (Outline)

Input: Examples $\{x_1, \ldots, x_N\}$, loss function $\mathcal{L}(\theta) = \frac{1}{N}\sum_i \mathcal{L}(\theta, x_i)$. Parameters: learning rate η_t, noise scale σ, group size L, gradient norm bound C.

Initialize θ_0 randomly

for $t \in [T]$ **do**

Take a random sample L_t with sampling probability L/N

Compute gradient

For each $i \in L_t$, compute $\mathbf{g}_t(x_i) \leftarrow \nabla_{\theta_t} \mathcal{L}(\theta_t, x_i)$

Clip gradient

$\bar{\mathbf{g}}_t(x_i) \leftarrow \mathbf{g}_t(x_i) / \max\left(1, \frac{\|\mathbf{g}_t(x_i)\|_2}{C}\right)$

Add noise

$\tilde{\mathbf{g}}_t \leftarrow \frac{1}{L}\sum_i \left(\bar{\mathbf{g}}_t(x_i) + \mathcal{N}(0, \sigma^2 C^2 \mathbf{I})\right)$

Descent

$\theta_{t+1} \leftarrow \theta_t - \eta_t \tilde{\mathbf{g}}_t$

Output θ_T and compute the overall privacy cost (ε, δ) using a privacy accounting method.

Abadi, Martin, et al. "Deep learning with differential privacy." *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*. 2016.

Privacy-preserving machine learning: Adding differential privacy

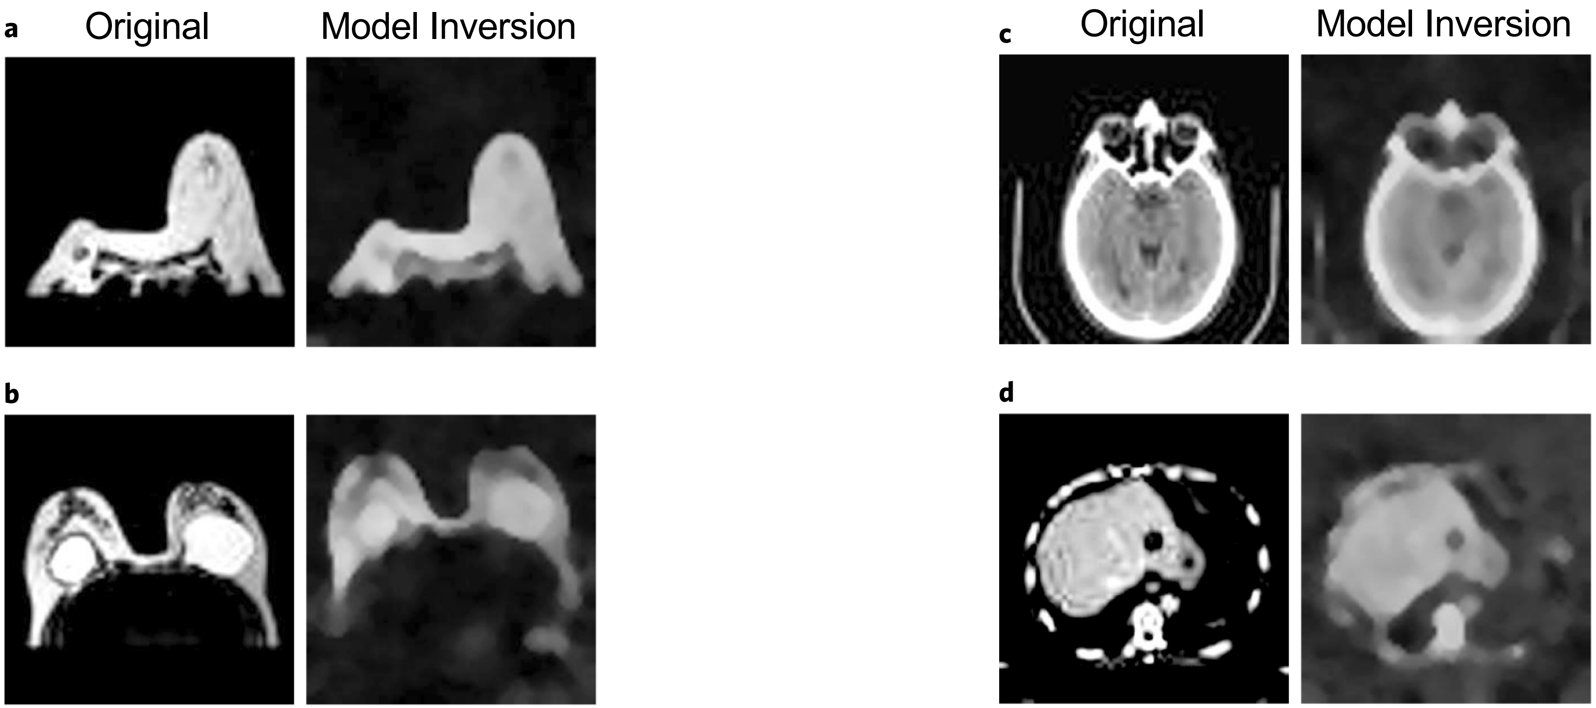

Kaissis et al. Nature Machine Intelligence, 2021

Privacy-preserving machine learning: Adding differential privacy

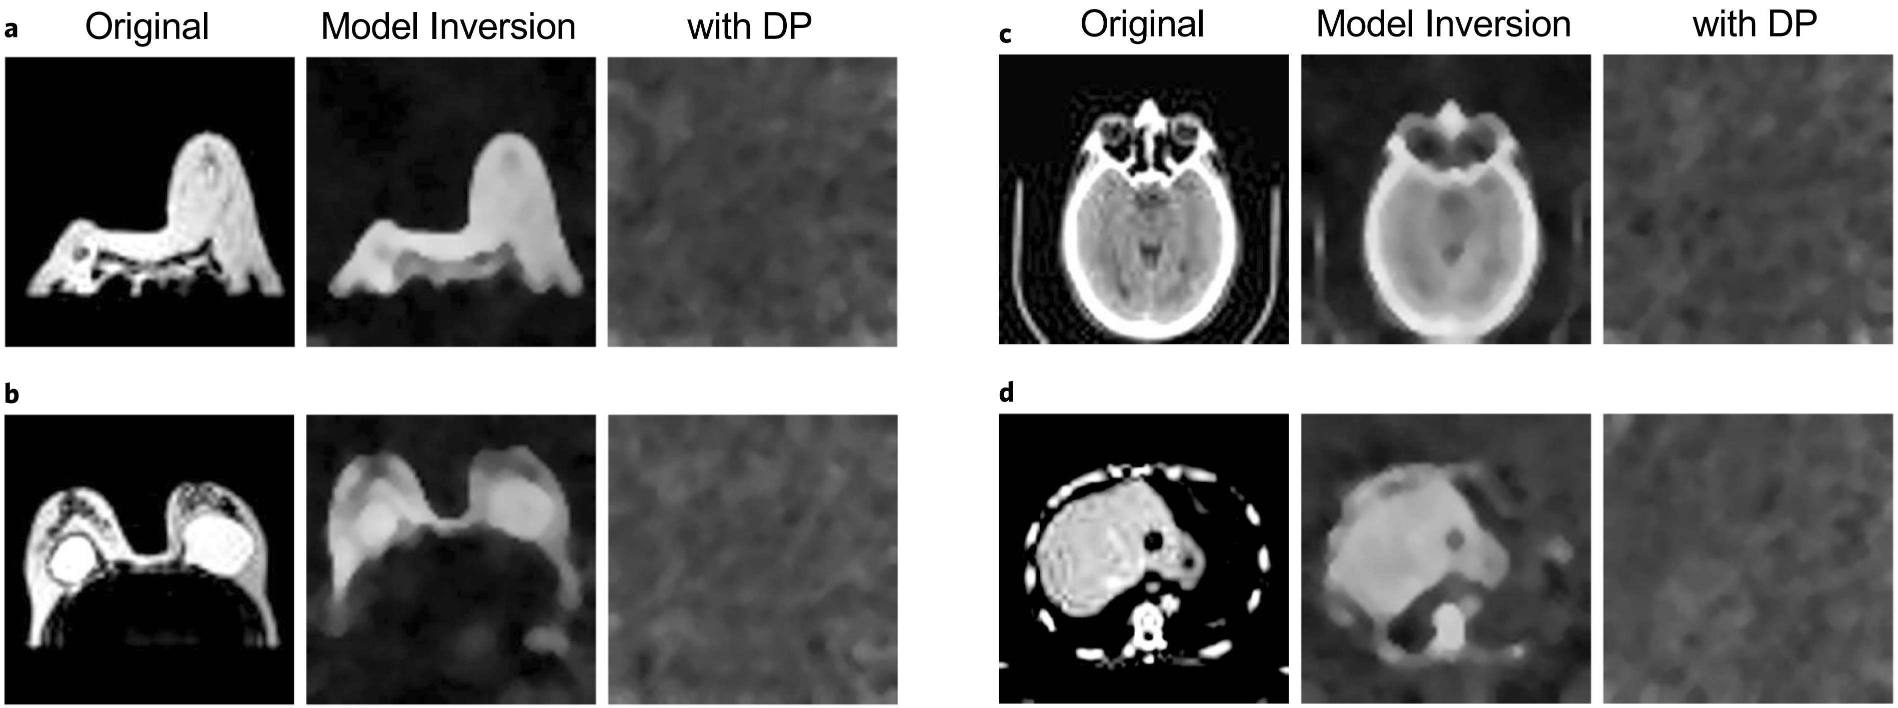

Kaissis et al. Nature Machine Intelligence, 2021

What’s next?

Traditional medical imaging

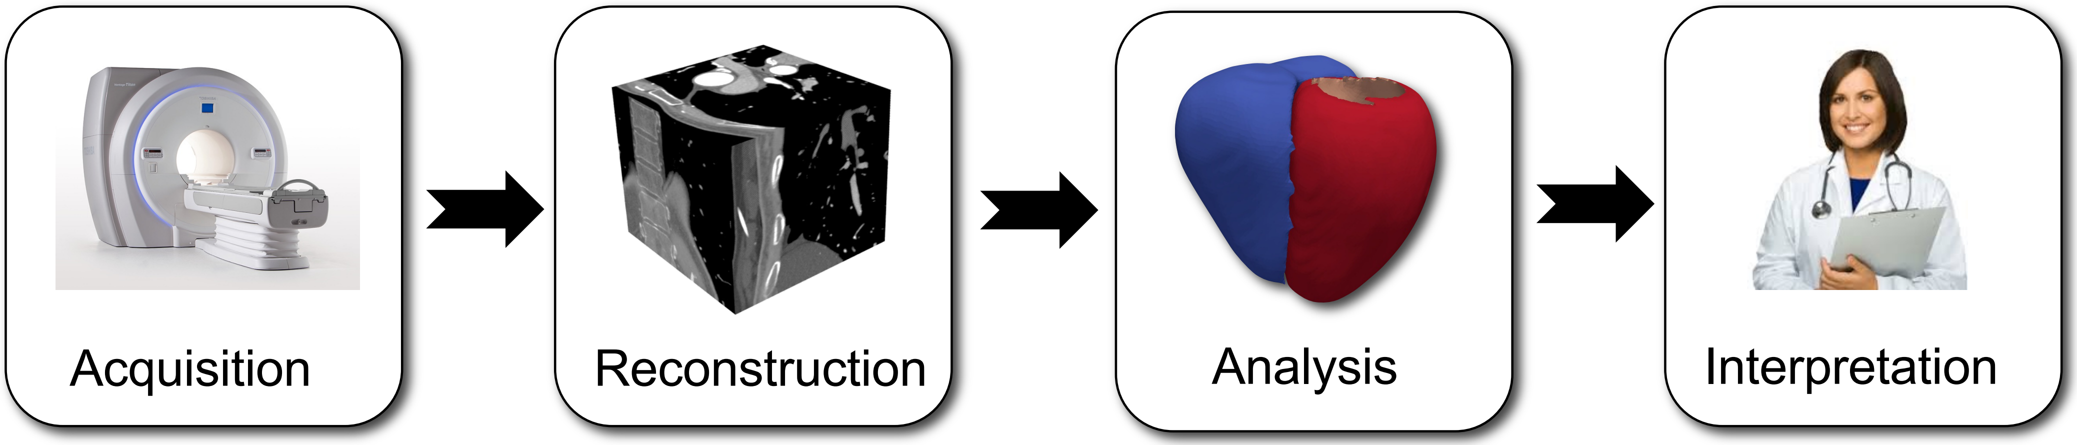

- ✗ Serial process with no interaction between different components of imaging pipeline
- ✗ Limited ability for adjustment of upstream imaging pipeline based on downstream requirements
- ✗ Stages of imaging pipeline not optimized for clinical endpoint

AI-enabled medical imaging

Acquisition

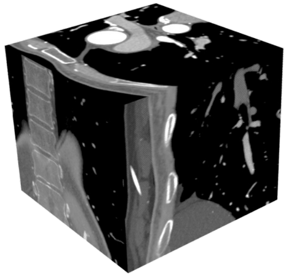

Reconstruction

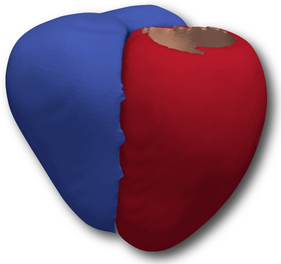

Analysis

Interpretation

- ✓ Close coupling of acquisition, reconstruction, analysis and interpretation
- ✓ Feedback and interaction between components of imaging pipeline
- ✓ Optimization of whole imaging pipeline with respect to clinical endpoint

AI-enabled medical imaging

“They should stop training radiologists now.”
Geoffrey Hinton (godfather of deep learning) in 2017

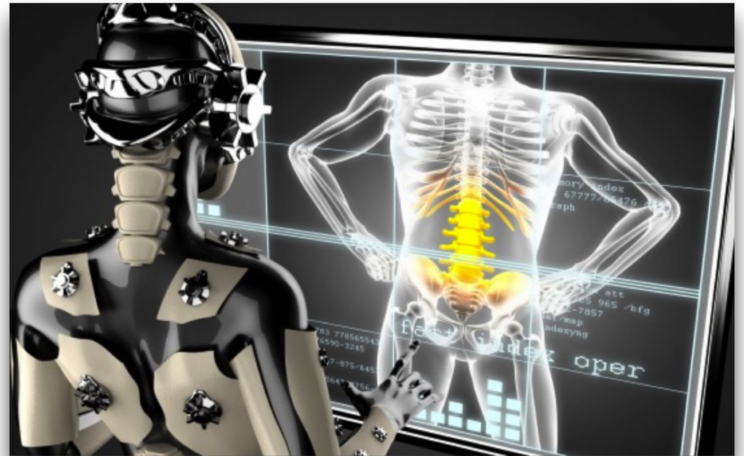

"To the question, will AI replace radiologists, I say the answer is no…”

“… but radiologists who do AI will replace radiologists who don’t."
Curtis Langlotz in 2017


RSNA News

Machine Learning Plays Central Role at RSNA 2017

BY MIKE BASSETT

November 1, 2017

Share Tweet Share Email

Machine Learning (ML) and the role it will play in the future of radiology will be central to a broad scope of programming at RSNA 2017.

Langlotz

AI-enabled medical imaging

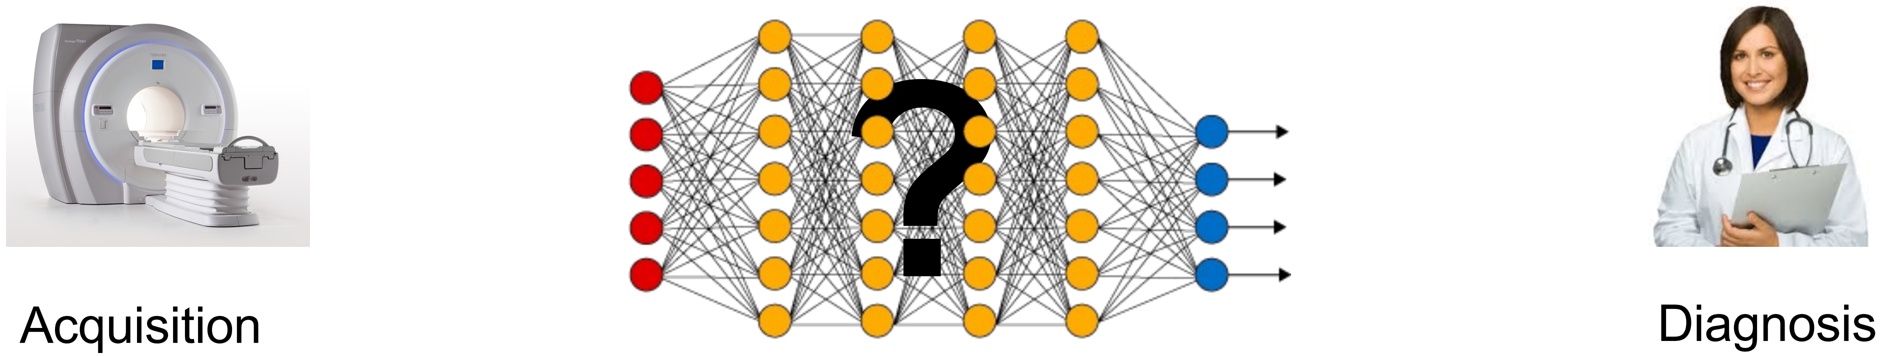

Do we need images at all?

Acknowledgements

EPSRC
Engineering and Physical Sciences Research Council

British Heart Foundation

erc

wellcome trust

biobank uk
Improving the health of future generations

Alexander von Humboldt
Stiftung/Foundation

Acknowledgements

Lab for AI in Medicine @

https://aim-lab.io/

BioMedIA @ Imperial College London

https://biomedia.doc.ic.ac.uk/