Imperial College London

# Deep learning

## ~~Learning~~ for Extracting Clinically Useful Information from Medical Images

Professor Daniel Rueckert, FREng, FIEEE

Biomedical Image Analysis Group

Department of Computing

Imperial College London, UK

# Machine Learning in Medical Imaging: Opportunities

- Machine learning techniques are starting to reach levels of human performance in challenging visual tasks
- Big data is slowly arriving in medical imaging

UK Biobank will provide large-scale imaging data from 100,000 subjects

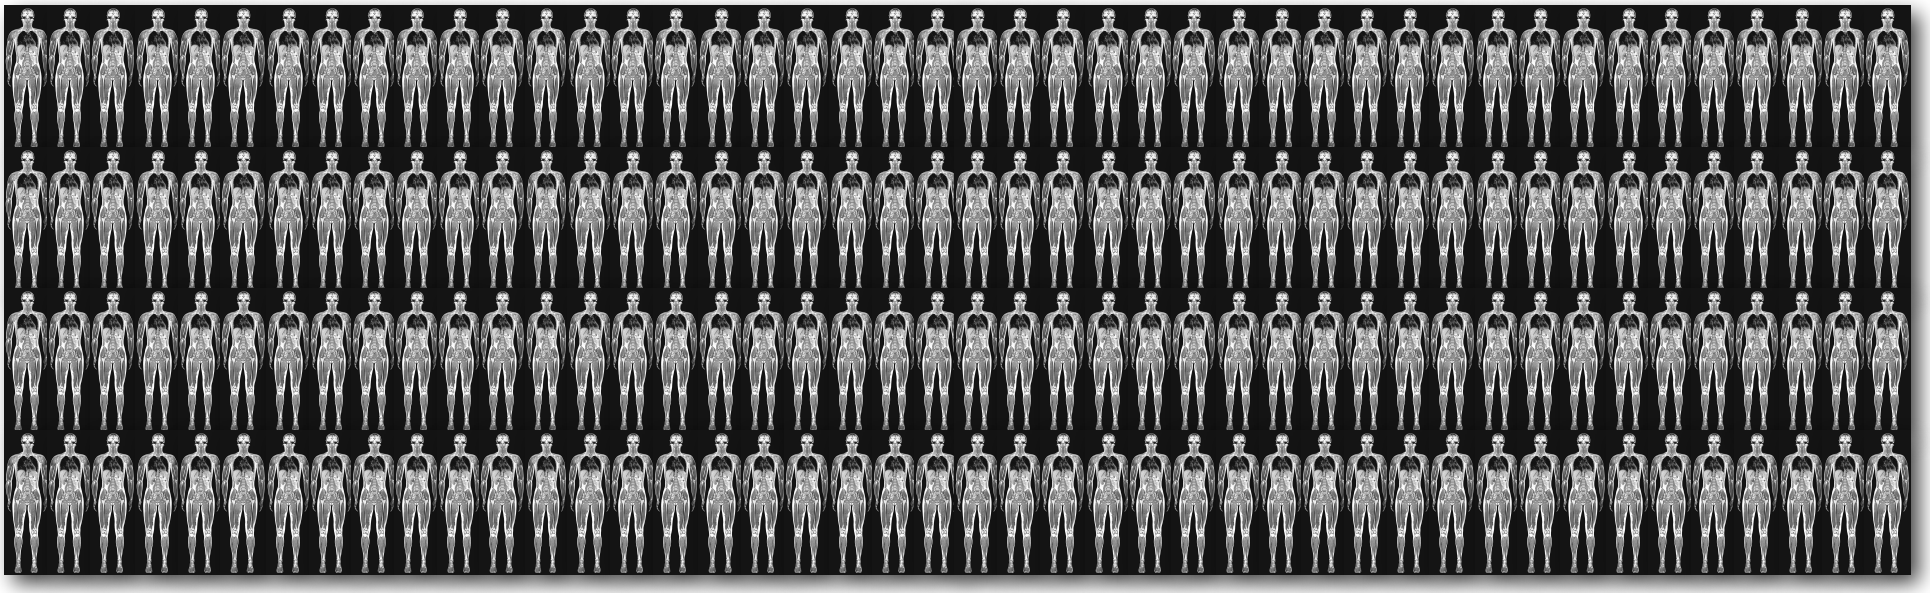

# Machine Learning in Medical Imaging: Opportunities

- Machine learning techniques are starting to reach levels of human performance in challenging visual tasks
- Big data is slowly arriving in medical imaging

UK Biobank will provide large-scale imaging data from 100,000 subjects

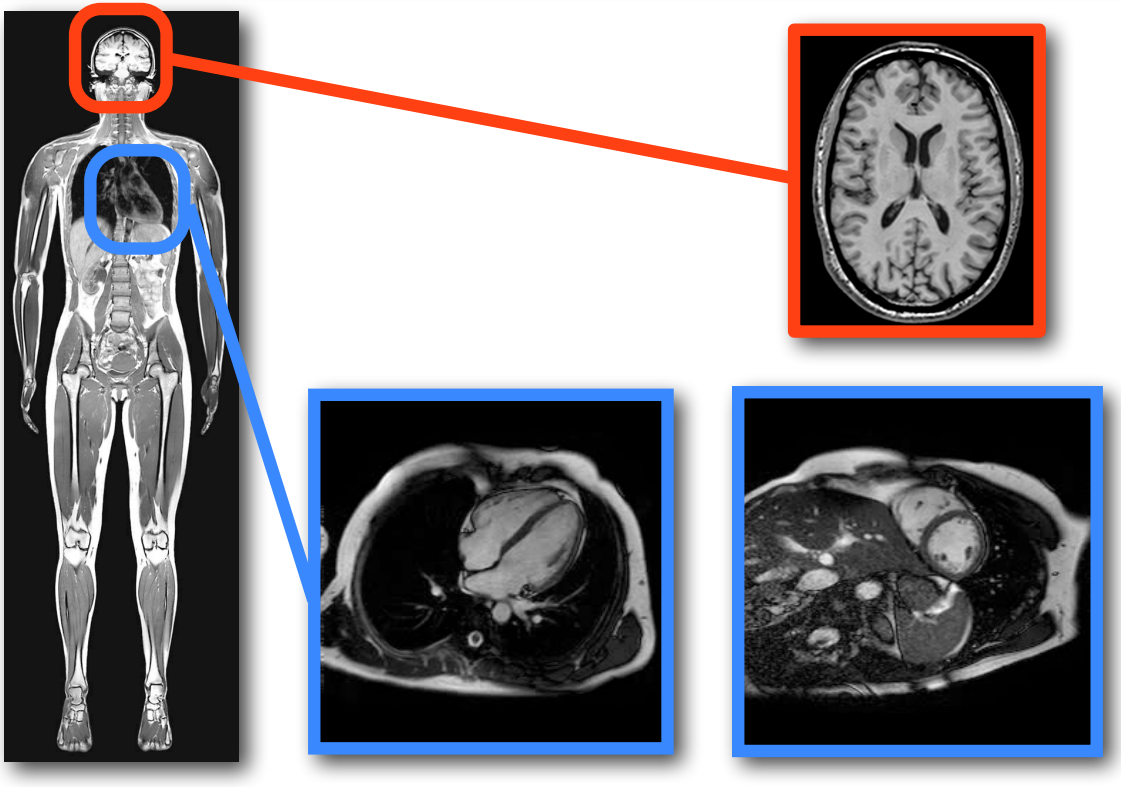

+

Lifestyle

Genetics

Clinical records

# Machine Learning in Medical Imaging: Opportunities

- Machine learning techniques are starting to reach levels of human performance in challenging visual tasks
- Big data is slowly arriving in medical imaging

UK Biobank will provide large-scale imaging data from 100,000 subjects

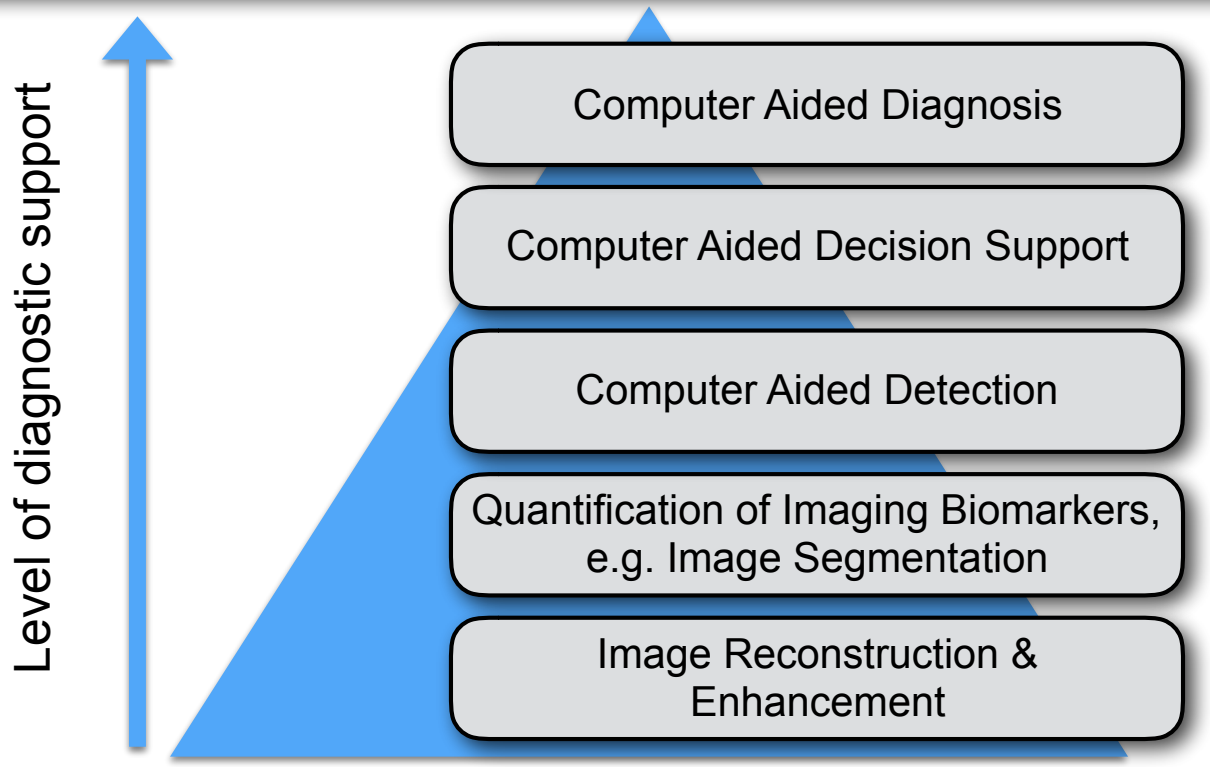

# Machine Learning in Medical Imaging: Challenges

- Images are often 3D or 4D:
  - # of voxels and # of extracted features is very large
- Number of images is limited:
  - large data set means typically 100 to 1000 images
  - “small sample size problem”
- Training data is expensive:
  - annotation of images is resource intensive (manpower, cost, time)
  - crowdsourcing not possible
- Training data is imperfect:
  - training data may be wrongly labelled, e.g. for diseases such as Alzheimer’s confirmation requires pathology (difficult and costly to obtain)

# Machine Learning in Medical Imaging: Challenges

- Images are highly variable
  - Different scanner hardware
  - Different imaging sequences
  - Images have artefacts

Siemens

Philips

GE

- Presence or absence of pathologies

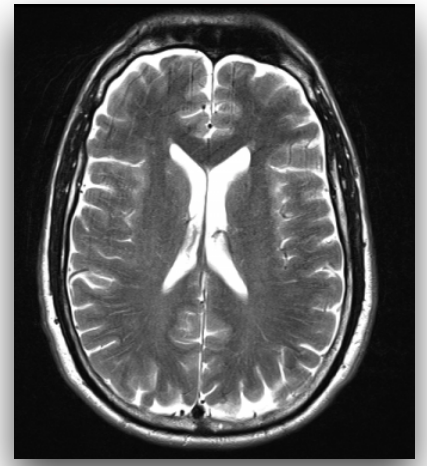

Normal

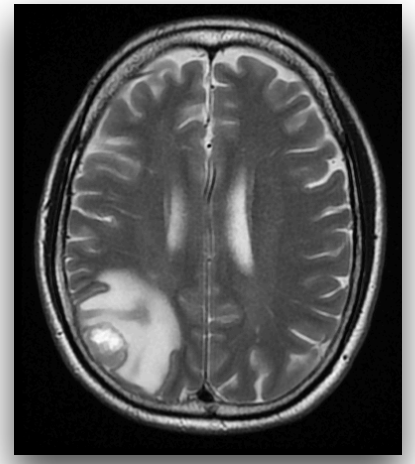

Abnormal

# Overview

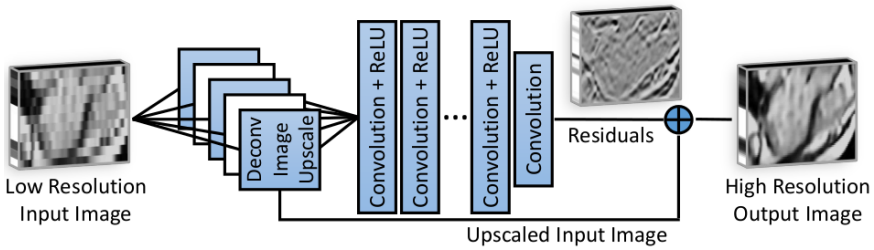


Image reconstruction &
super-resolution

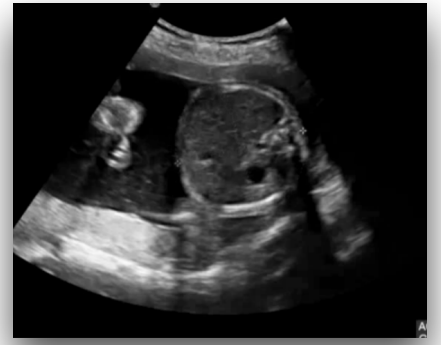

**Abdominal View**
Confidence: 98%

Image classification

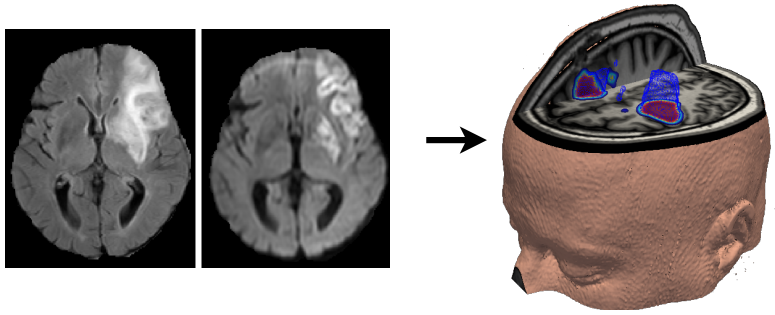

Image segmentation

# MR image acquisition: Challenges

- Magnetic Resonance Imaging (MRI)
  - MRI acquisition is inherently a slow process
  - Slow acquisition is
    - ok for static objects (e.g. brain, bones, etc)
    - problematic for moving objects (e.g. heart, liver, fetus)
  - Options for MRI acquisition:
    - real-time MRI: fast, but 2D and relatively poor image quality
    - gated MRI: fine for period motion, e.g. respiration or cardiac motion but requires gating (ECG or navigators) leading to long acquisition times (30-90 min).

# Example: Cardiac imaging

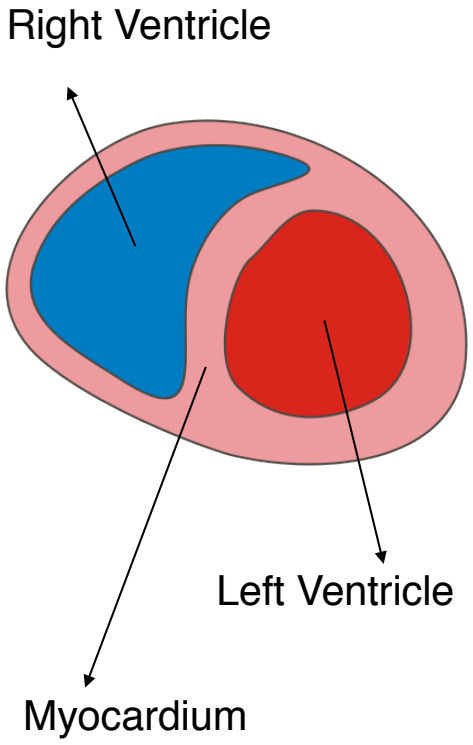


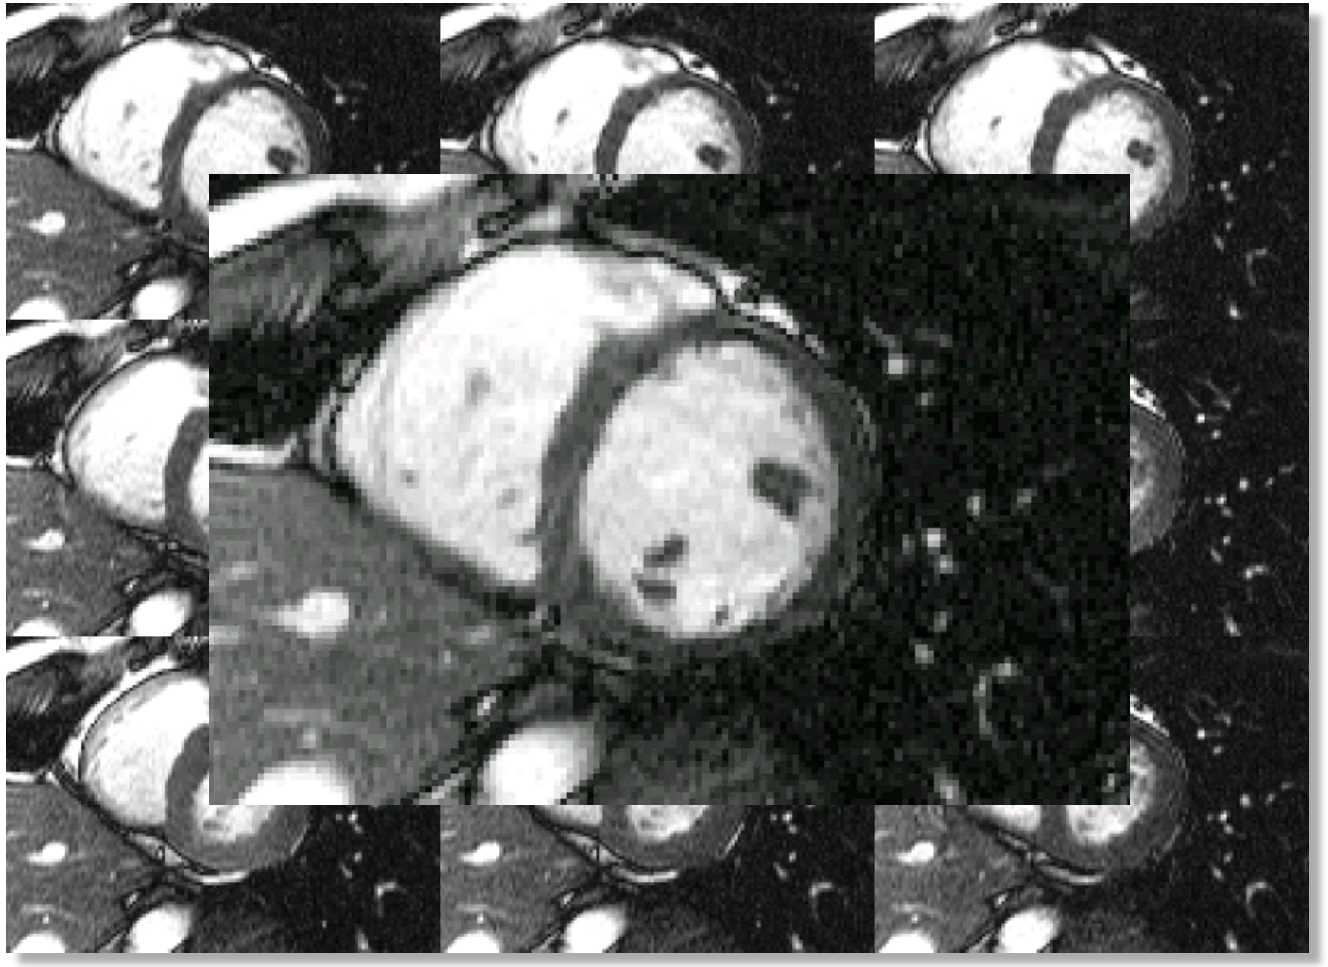

# MR full acquisition

- MRI acquisition is performed in k-space by sequenti
  traversing sampling trajectories.

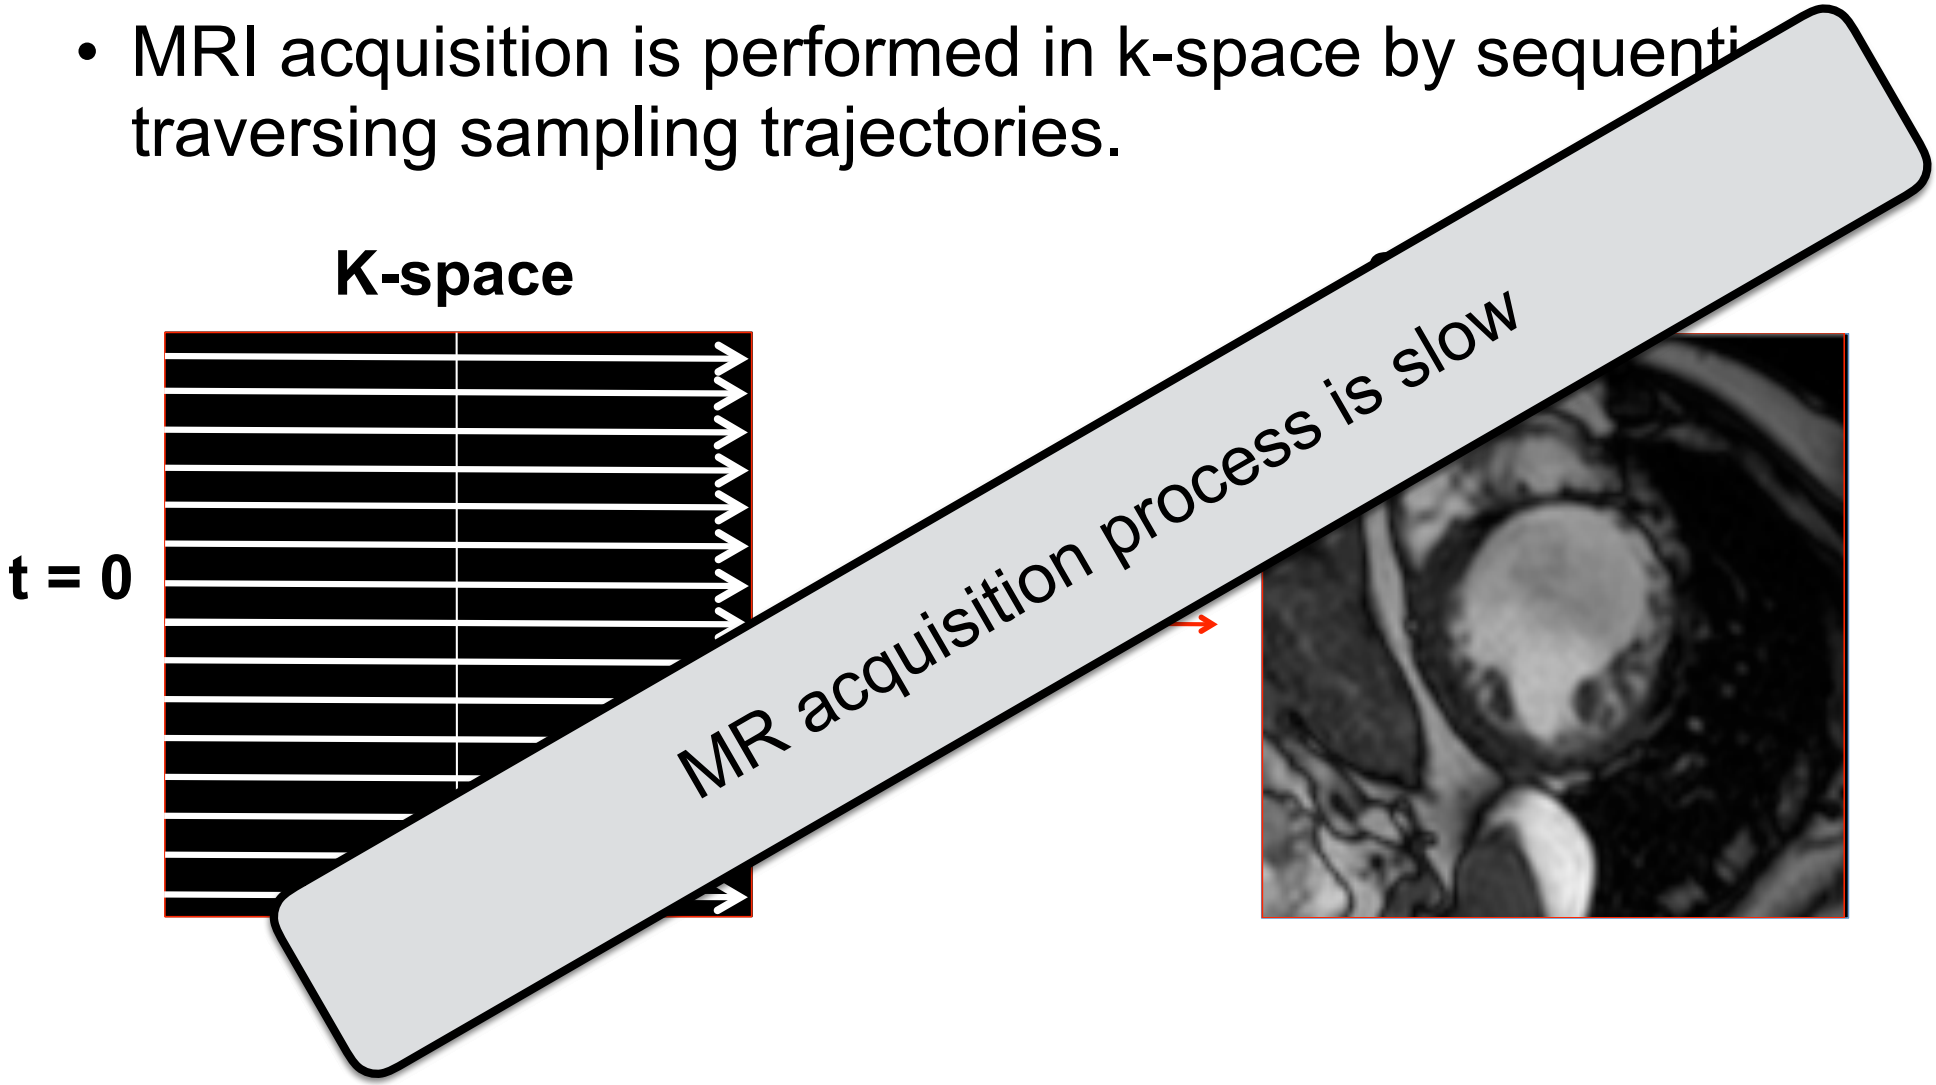

# MR full acquisition

- MRI acquisition is performed in k-space by sequentially traversing sampling trajectories.

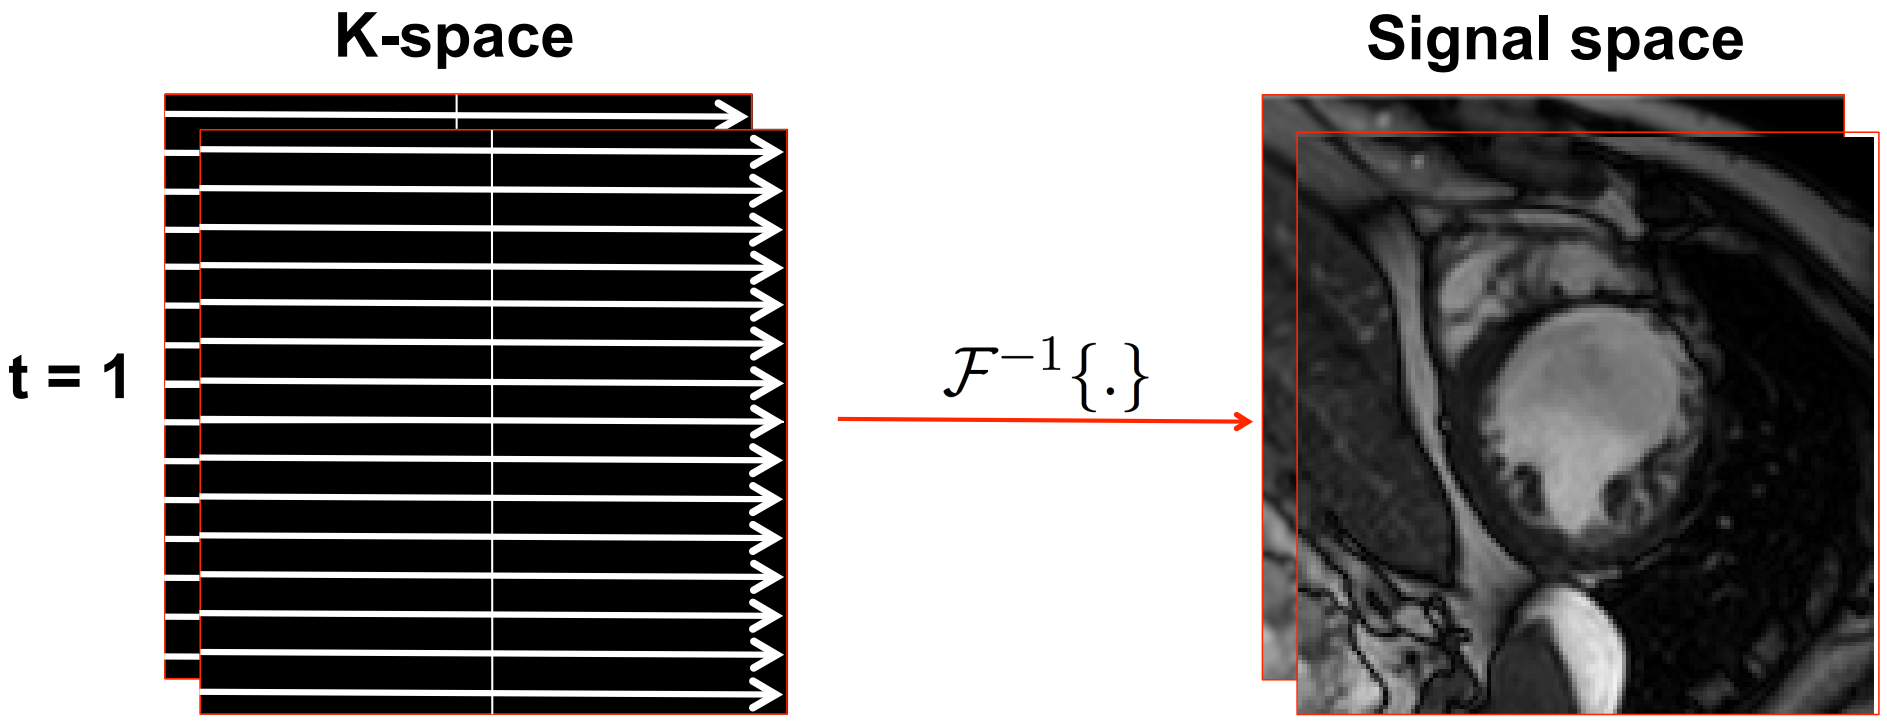

# MR full acquisition

- MRI acquisition is performed in k-space by sequentially traversing sampling trajectories.

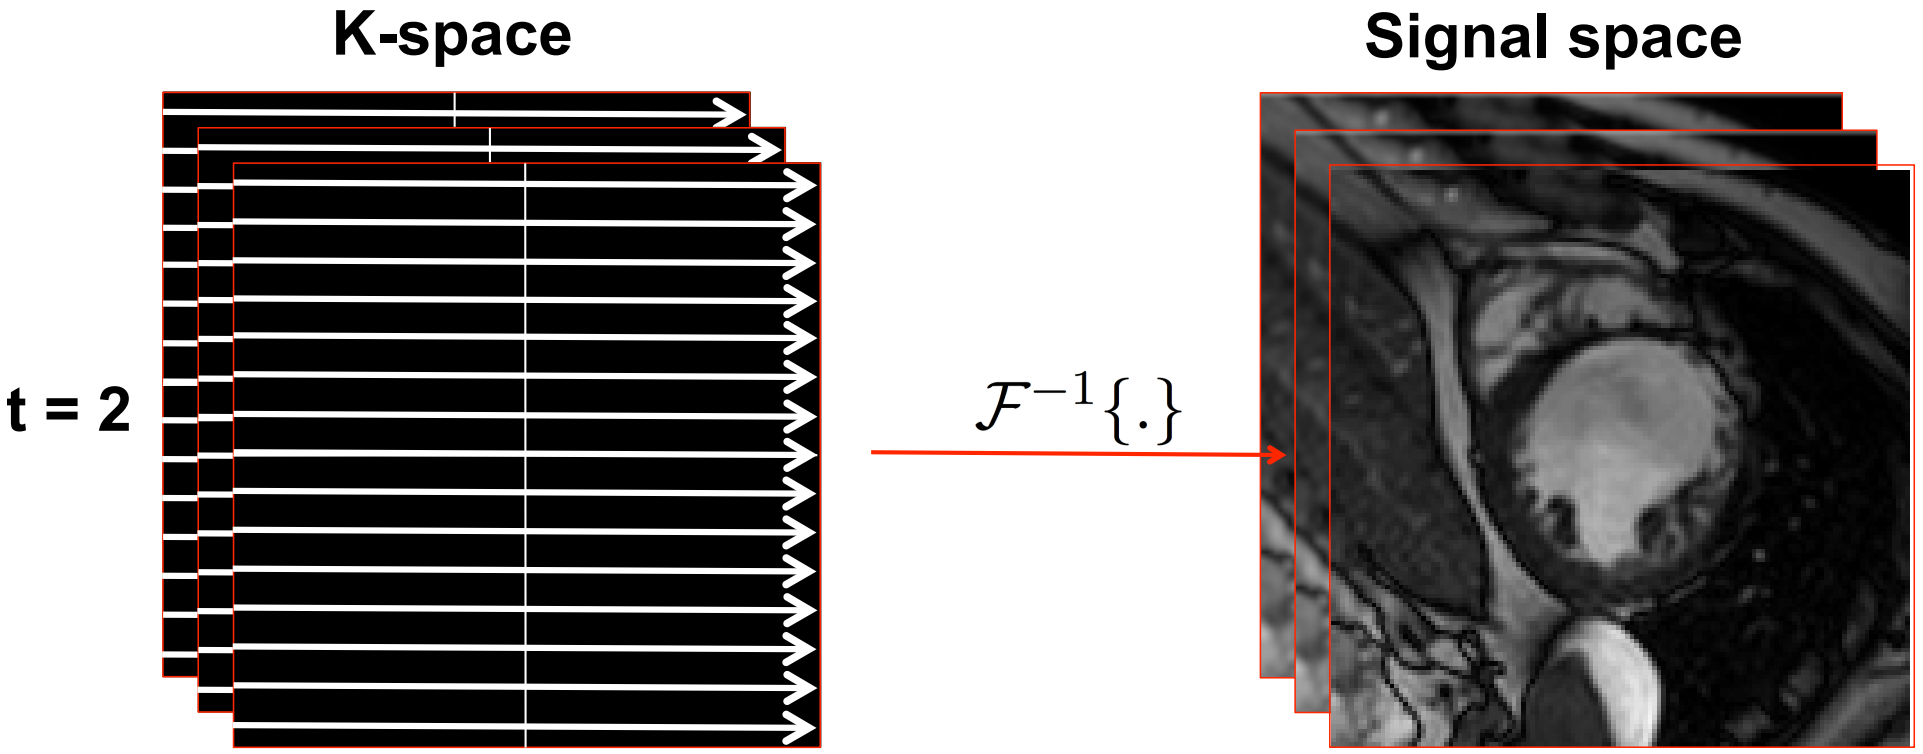

# MR full acquisition

- MRI acquisition is performed in k-space by sequenti
  traversing sampling trajectories.

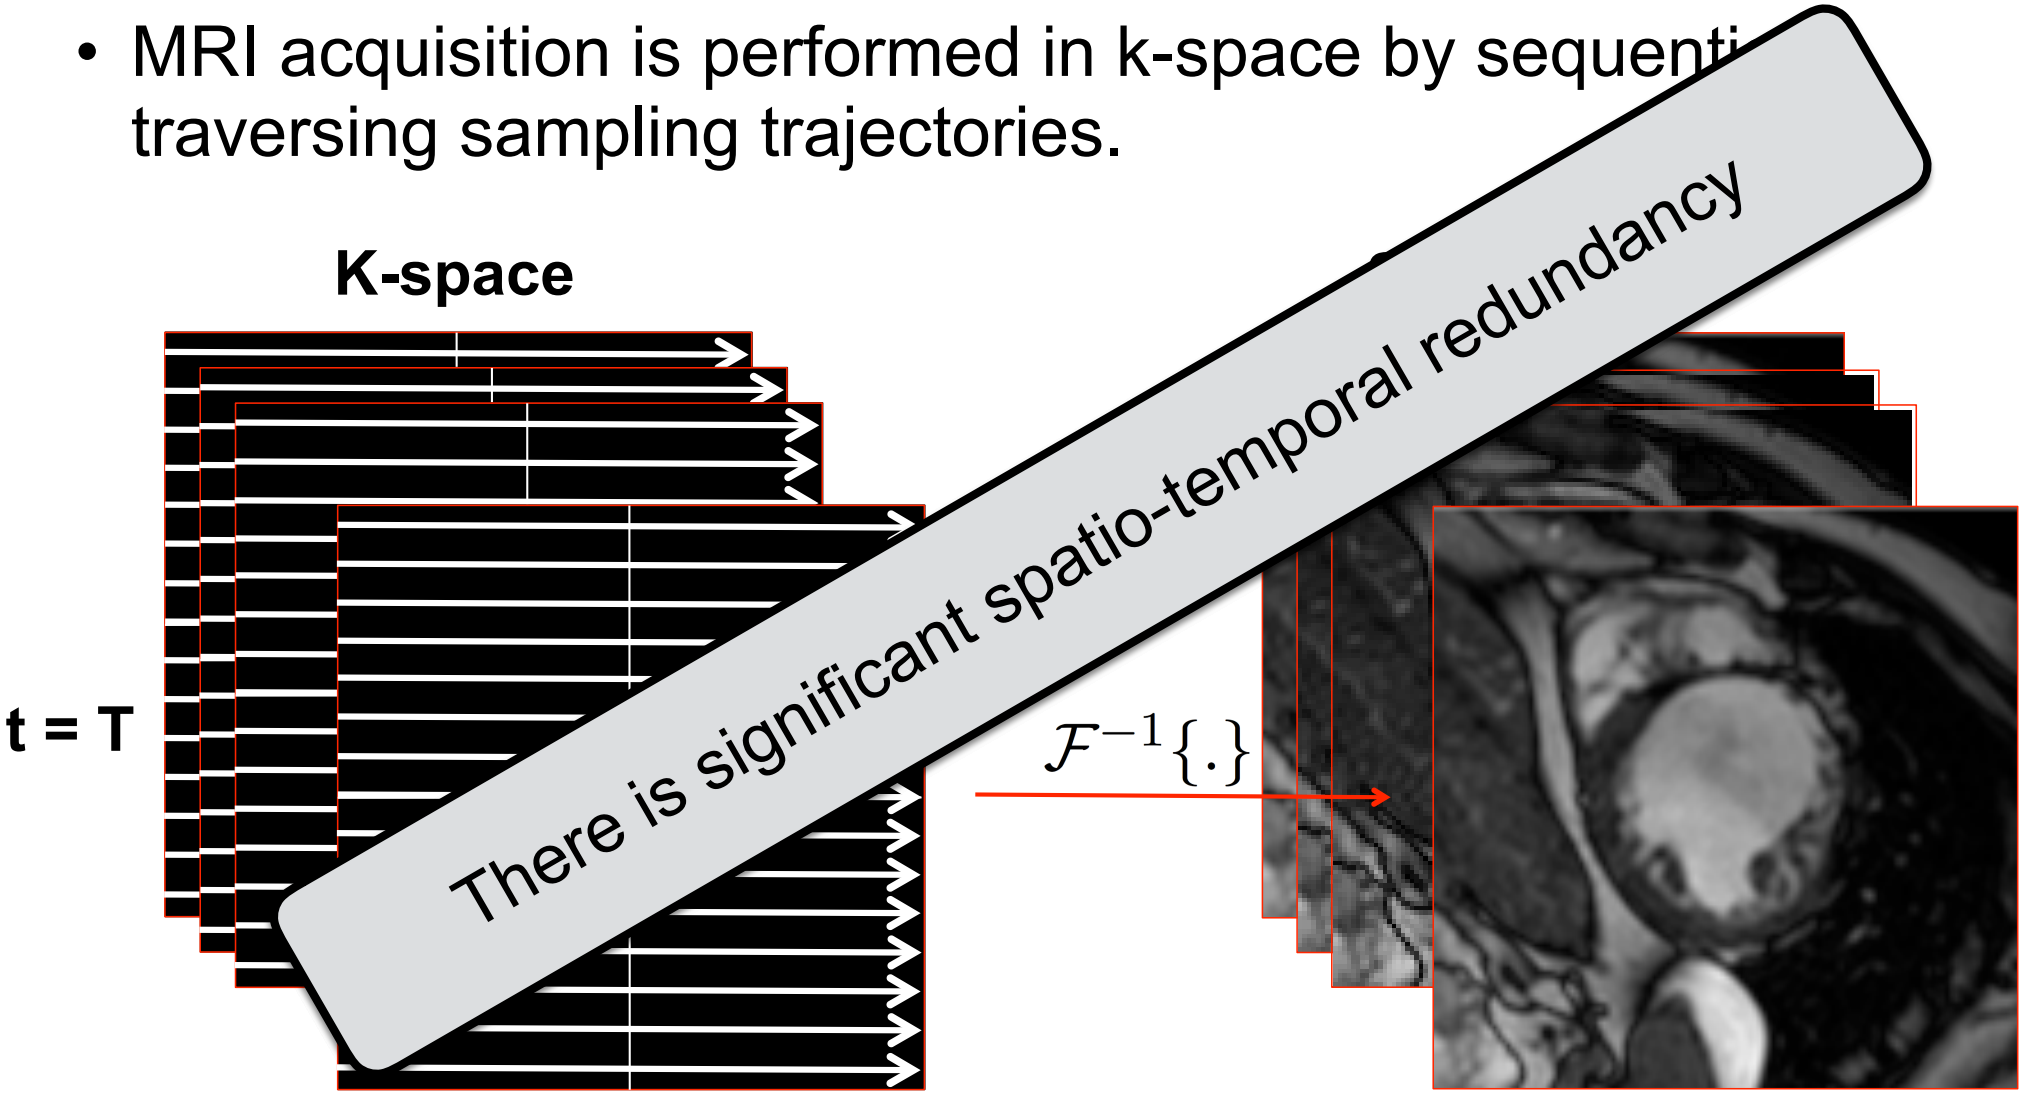

# K-space undersampling

- Acquiring a fraction of k-space **accelerates** the process but introduces **aliasing** in signal space.

# K-space undersampling

- Acquiring a fraction of k-space **accelerates** the process but introduces **aliasing** in signal space.

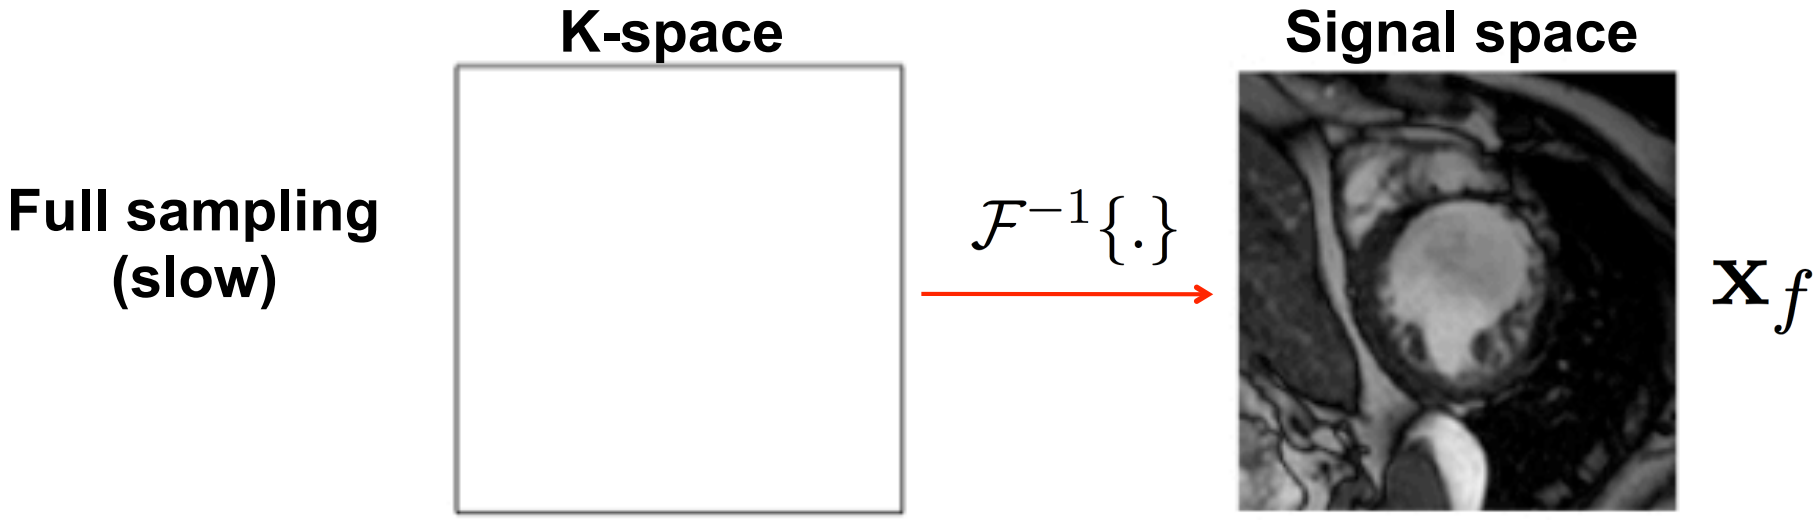

# Deep Cascade of CNNs for MRI Reconstruction

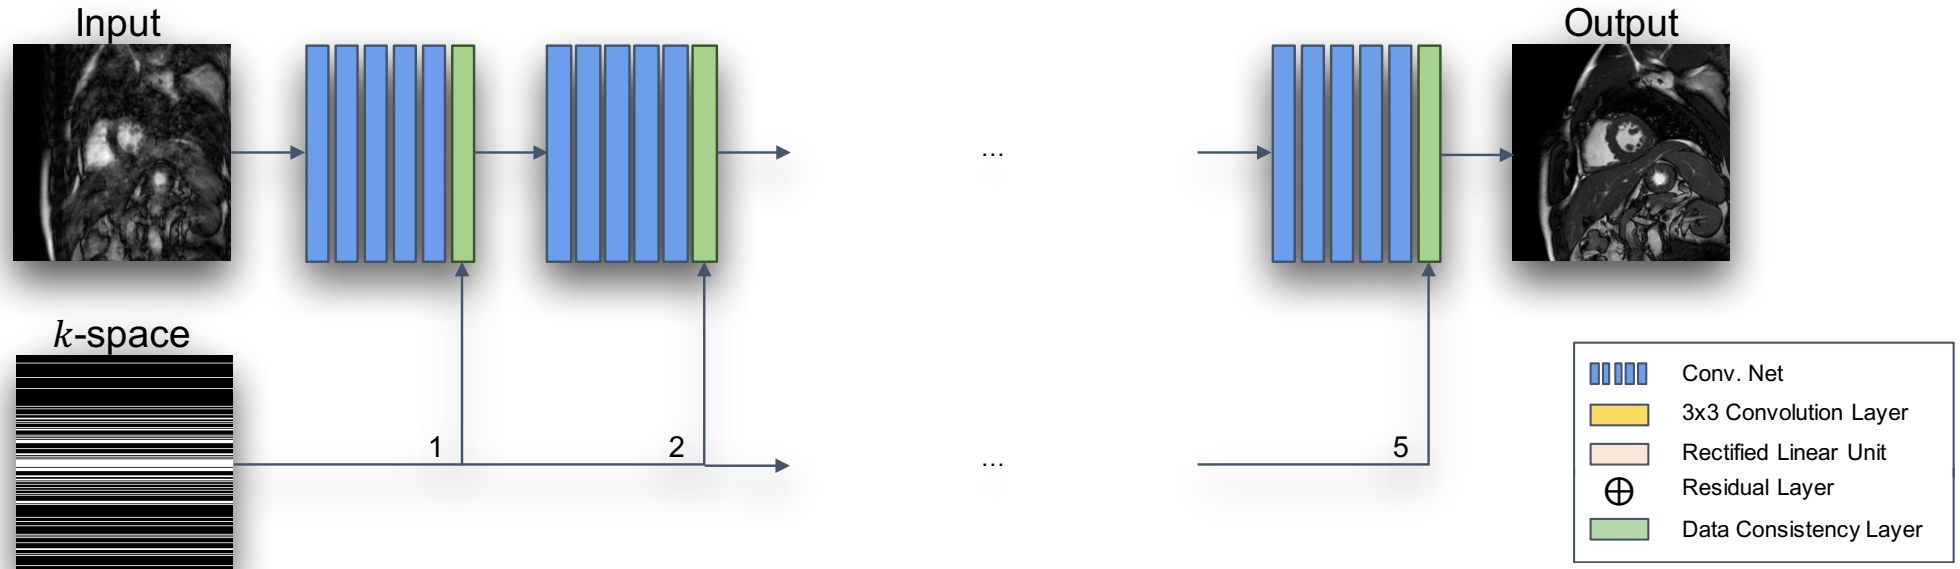


Schlemper et al. IEEE TMI 2017

# Deep Cascade of CNNs for MRI Reconstruction

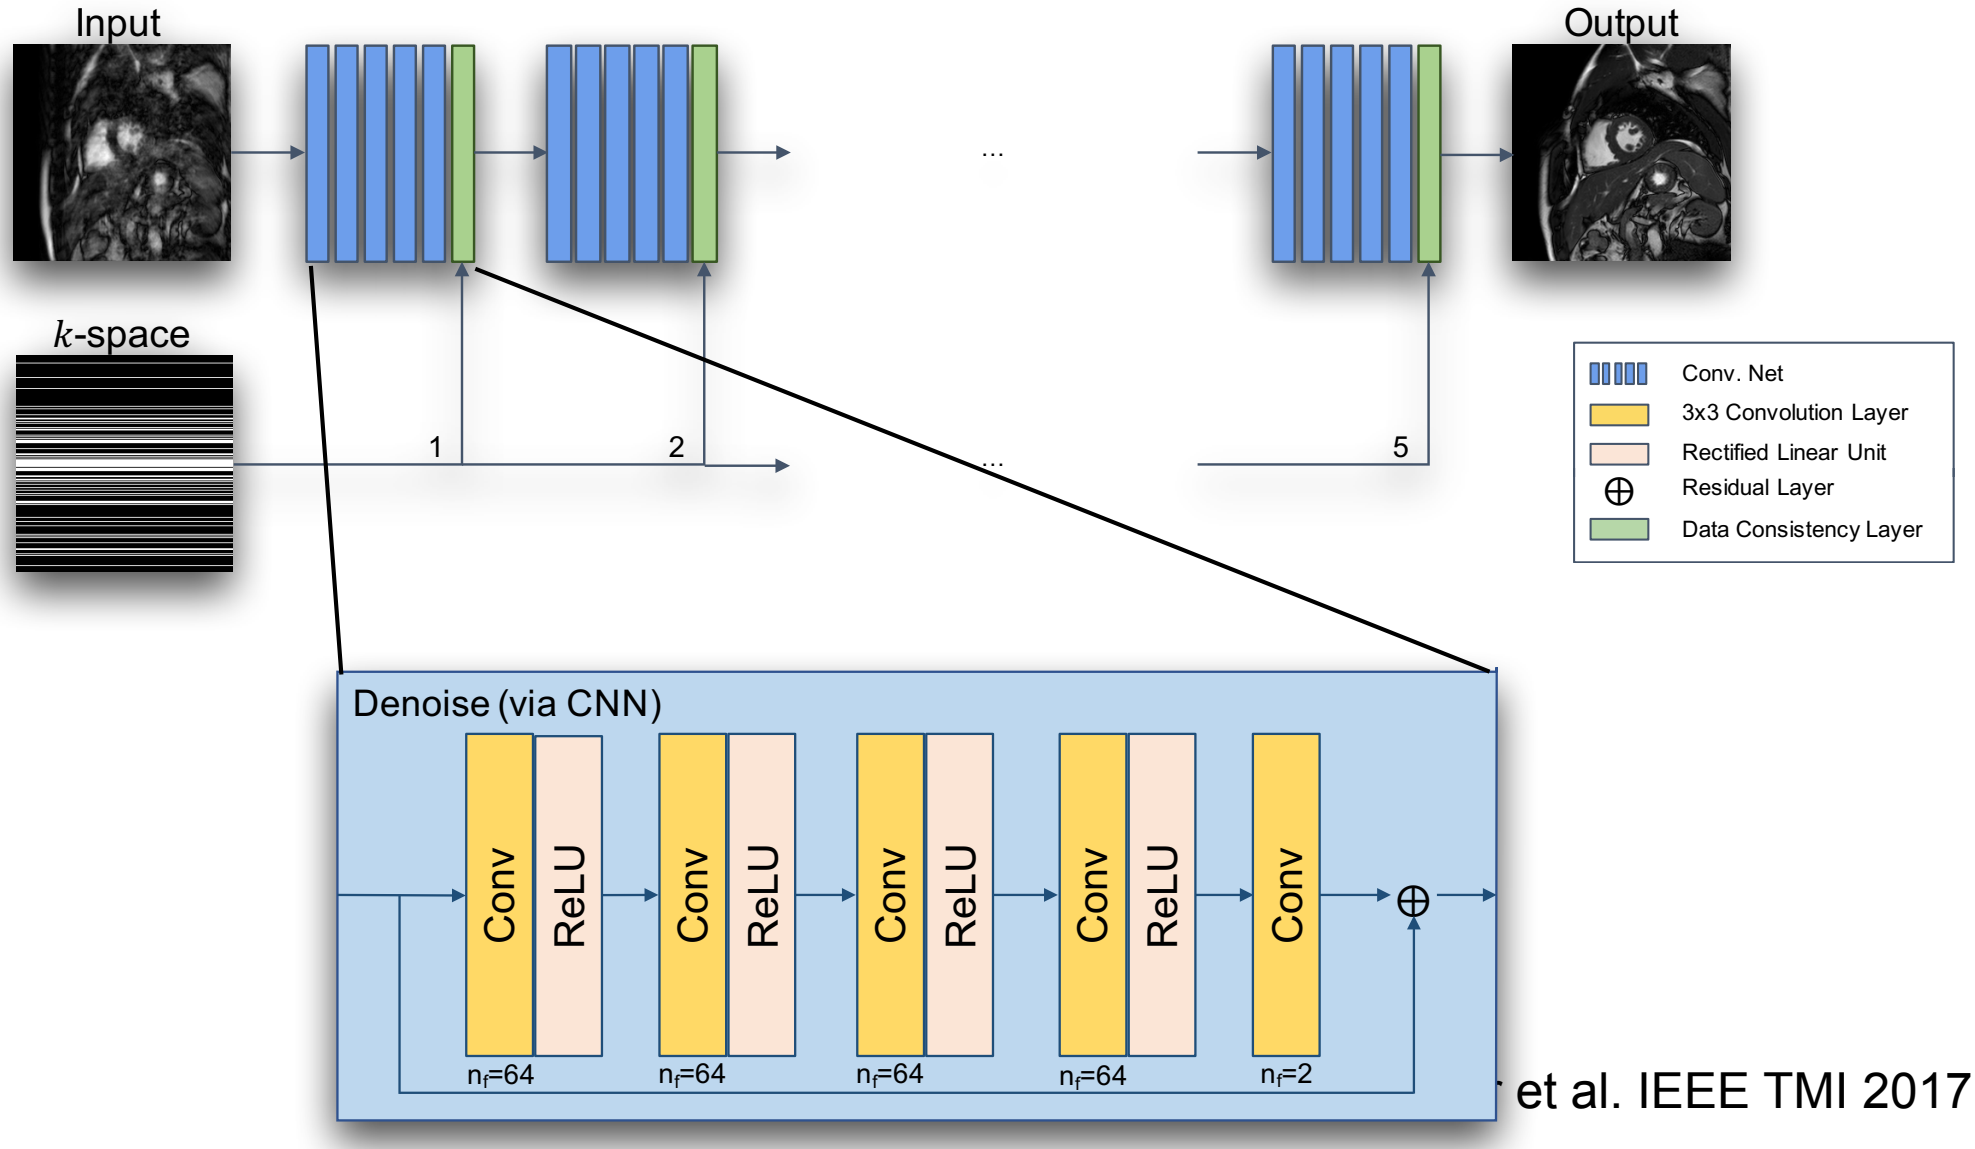


et al. IEEE TMI 2017

# Deep Cascade of CNNs for MRI Reconstruction

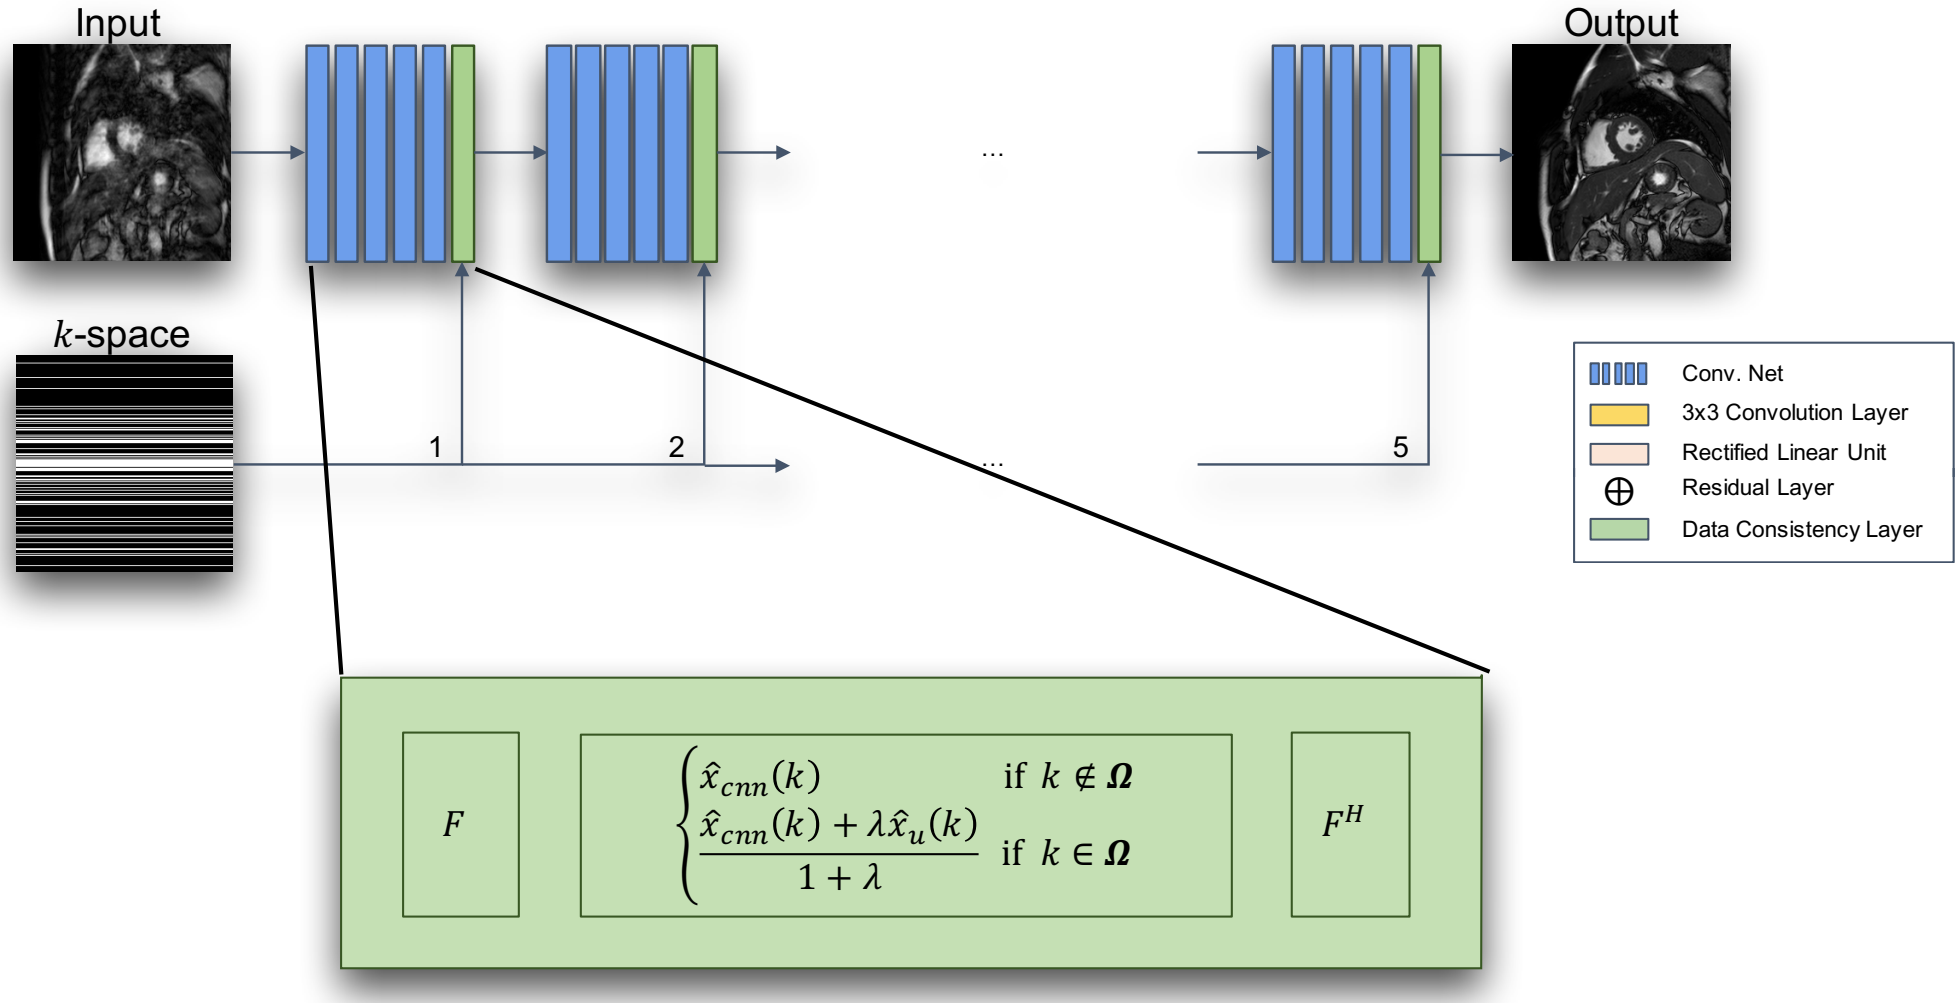


Schlemper et al. IEEE TMI 2017

# Magnitude reconstruction (6-fold)

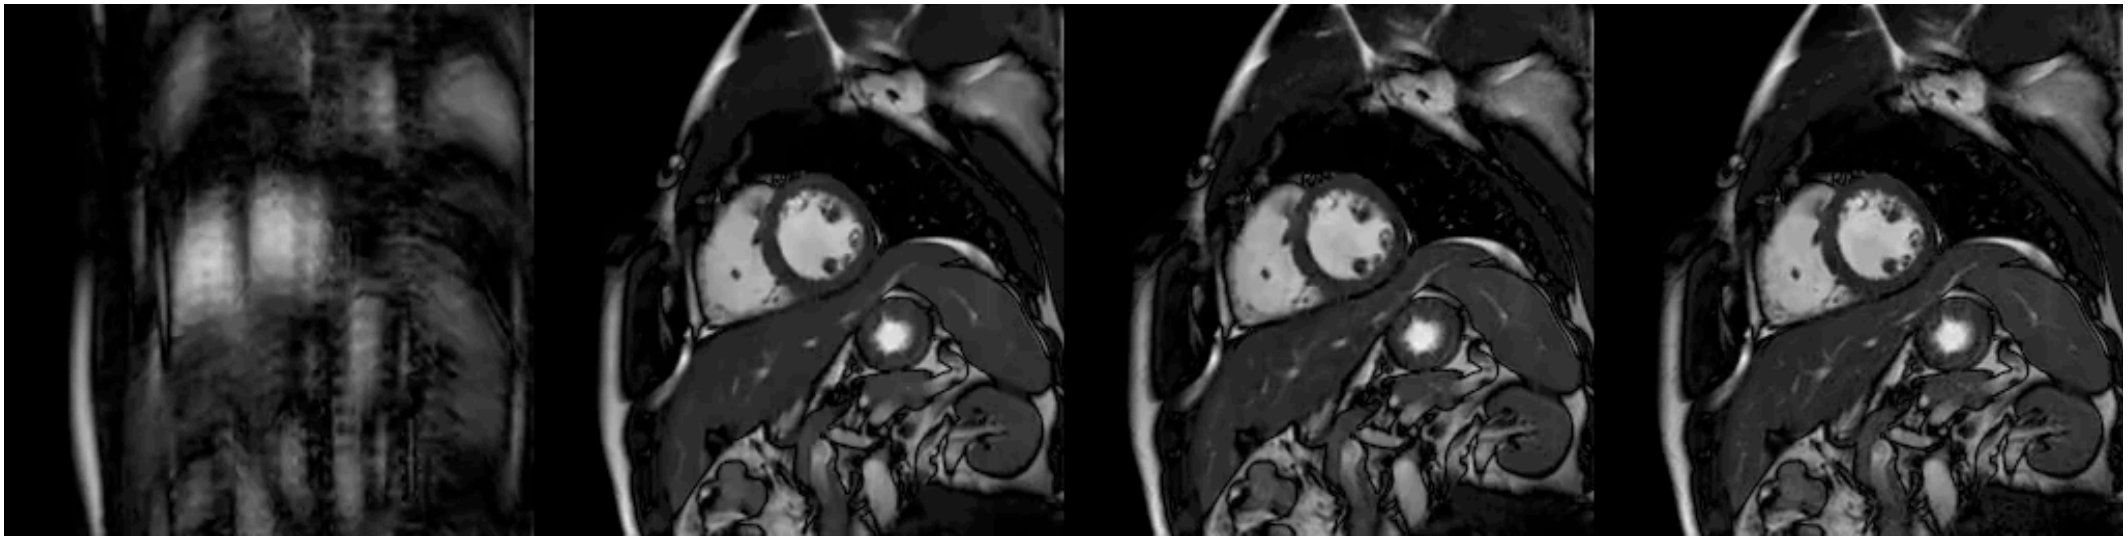

**(a) 6x Undersampled** **(b) DLTG** **(c) CNN** **(d) Ground Truth**

Schlemper et al. IEEE TMI 2017

# Magnitude reconstruction (11-fold)

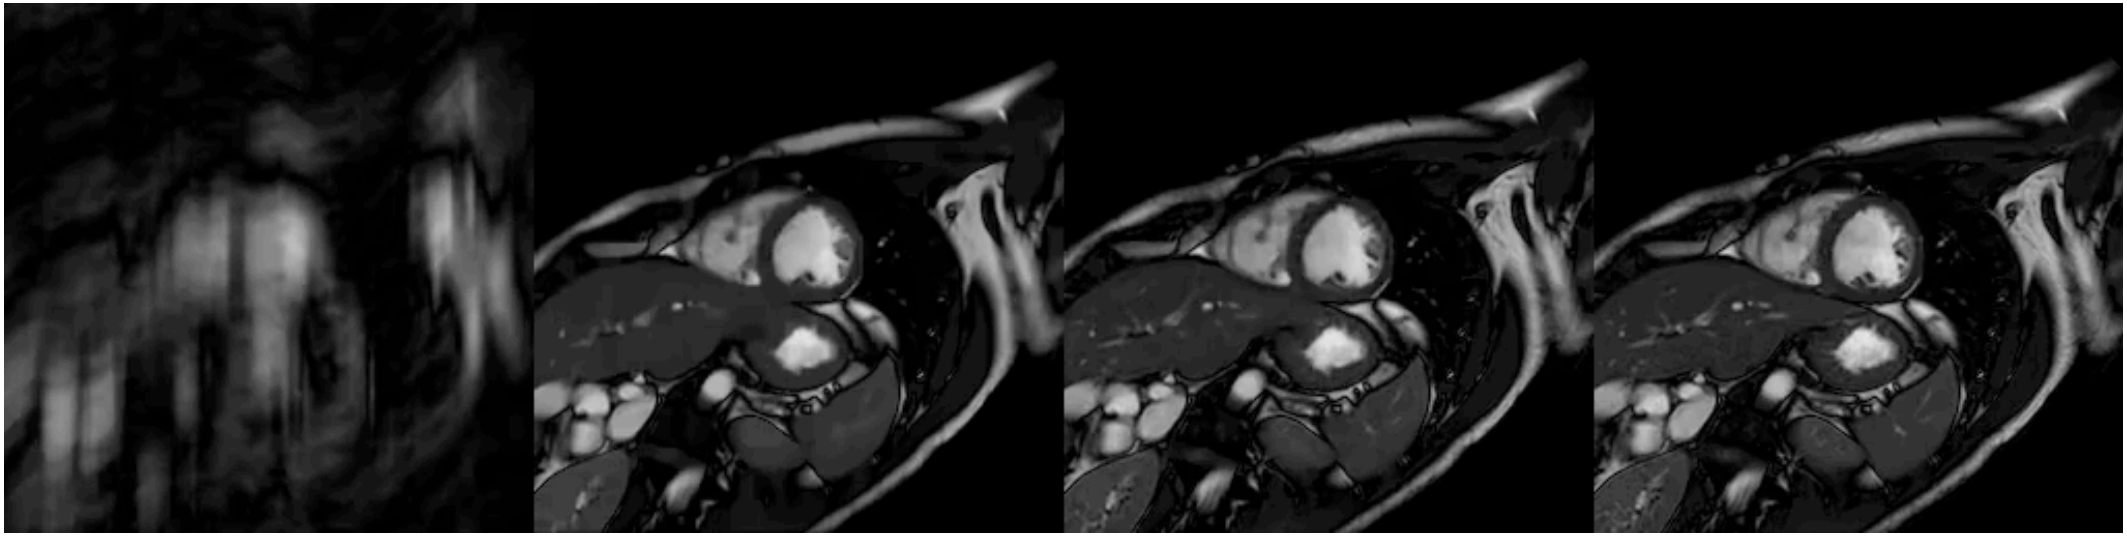

**(a) 11x Undersampled** **(b) DLTG** **(c) CNN** **(d) Ground Truth**

Schlemper et al. IEEE TMI 2017

# Deep Cascade of CNNs for MRI Reconstruction: Results

- Test error across 10 subjects:

2D, PSNR: mean (sd)

| Model | R=4 (dB) | R=8 (dB) |
|---|---|---|
| DLTG | 27.5 (1.31) | 22.6 (0.95) |
| CNN | **31.0 (1.08)** | **25.2 (1.00)** |

2D+t (vs. DLTG)

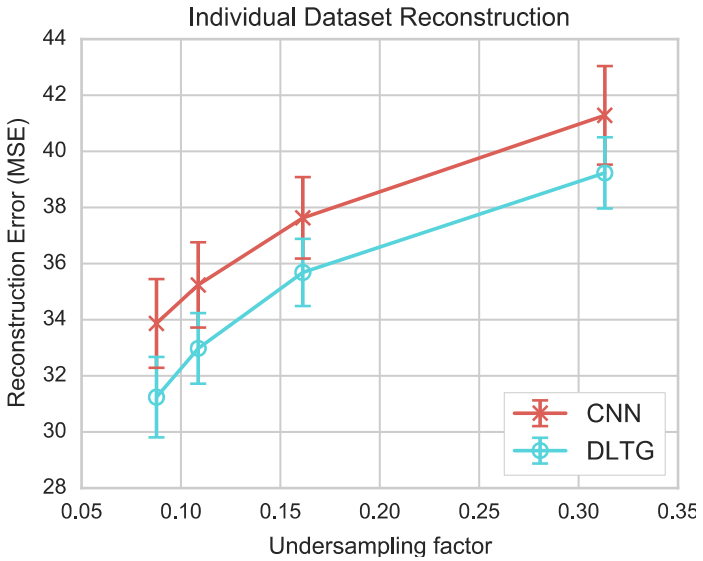


Schlemper et al. IEEE TMI 2017

# Deep Cascade of CNNs for MRI Reconstruction: Results

- Test error across 10 subjects:

2D, PSNR: mean (sd)

| Model | R=4 (dB) | R=8 (dB) |
|---|---|---|
| DLTG | 27.5 (1.31) | 22.6 (0.95) |
| CNN | **31.0 (1.08)** | **25.2 (1.00)** |

Reconstruction speed

| Model | Time |
|---|---|
| DLMRI/DLTG | ~6 hr (CPU) |
| CNN (2D) | 0.69 s (GPU) |
| CNN (2D+t) | 10 s (GPU) |

2D+t (vs. DLTG)

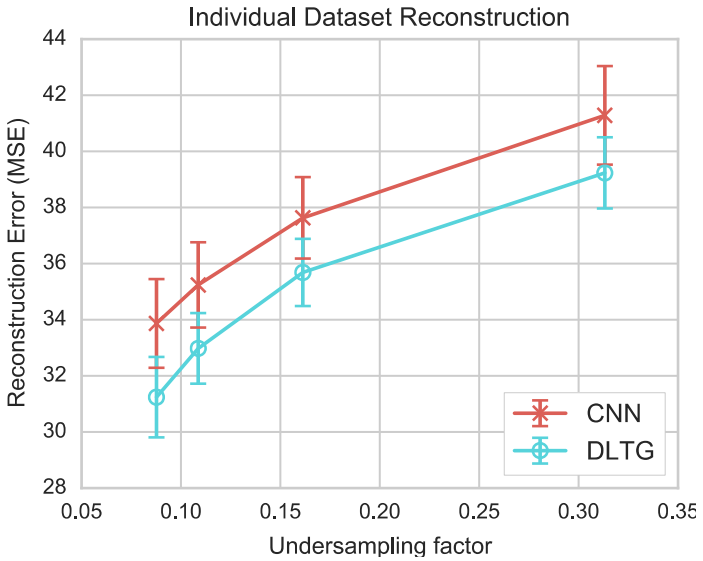


Schlemper et al. IEEE TMI 2017

# Still, cardiac imaging is challenging

- 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)

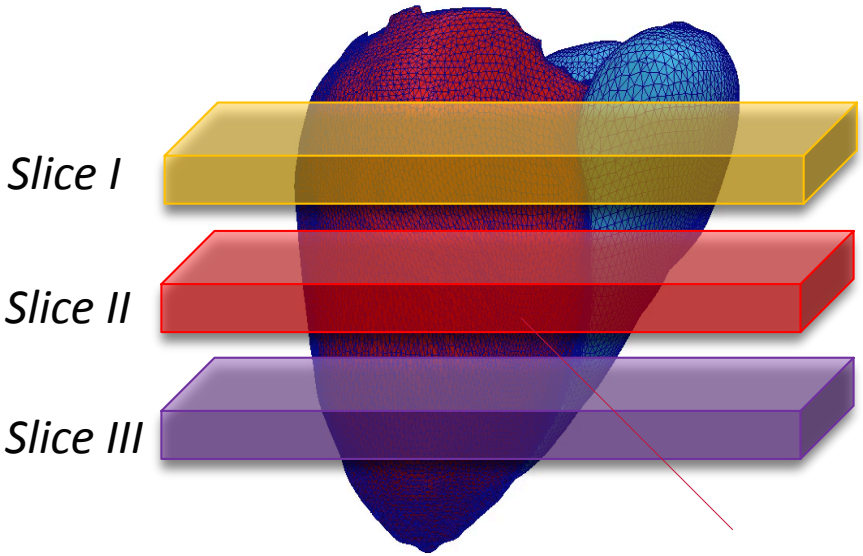


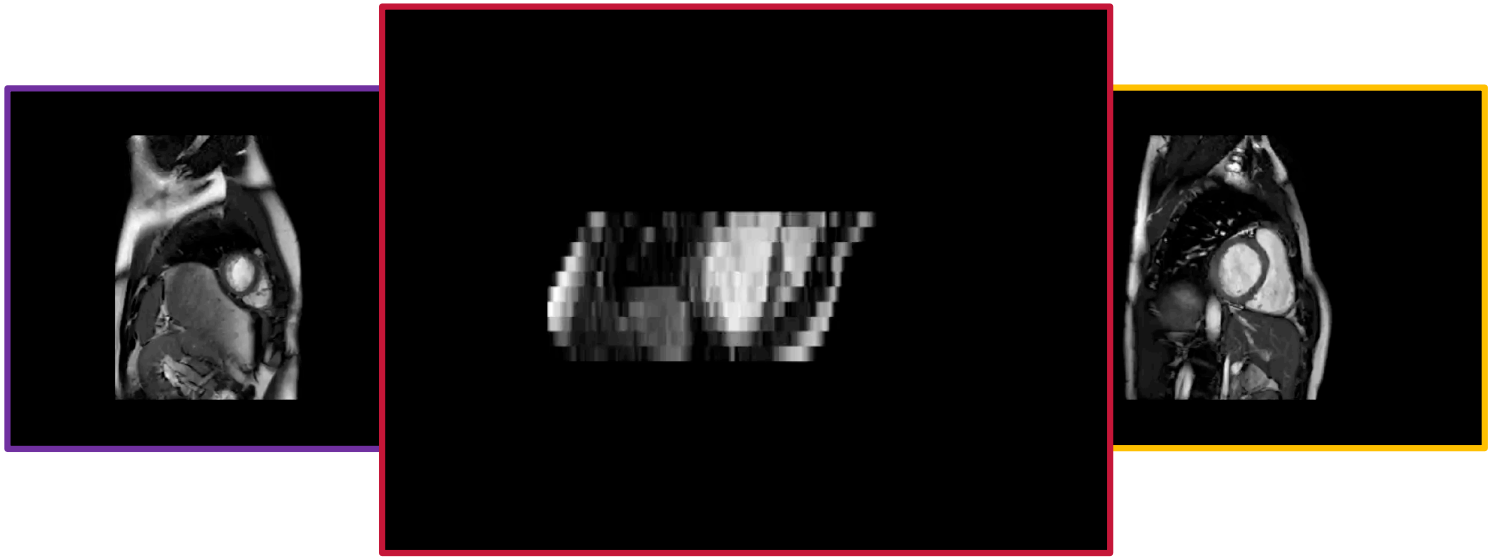

# Still, cardiac imaging is challenging

- 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)
- Images are acquired in different orientations:
  - short-axis
  - long-axis, i.e.
    - 2CH or 4CH views

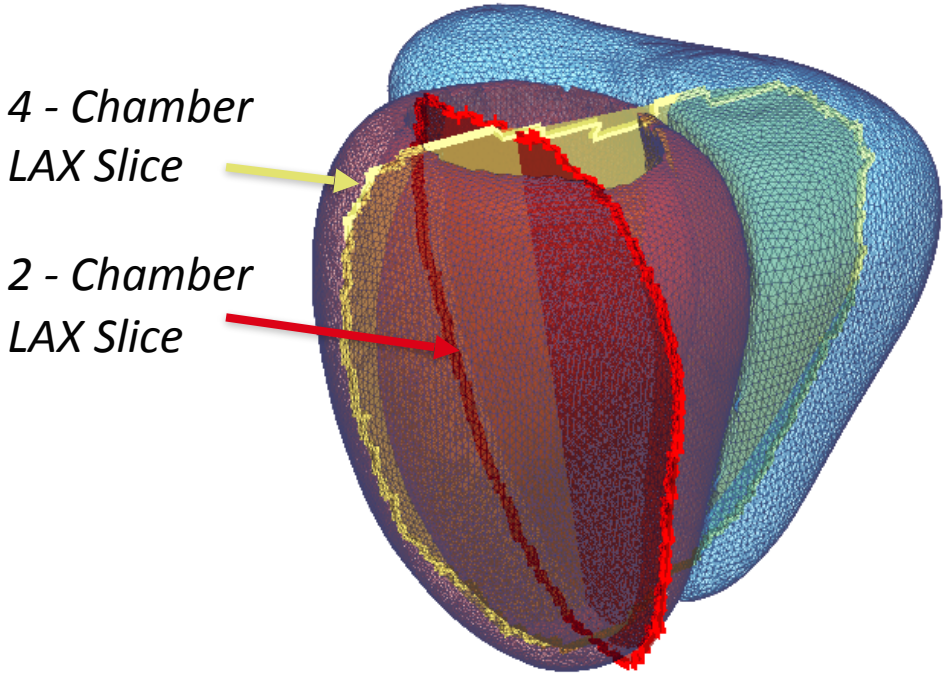


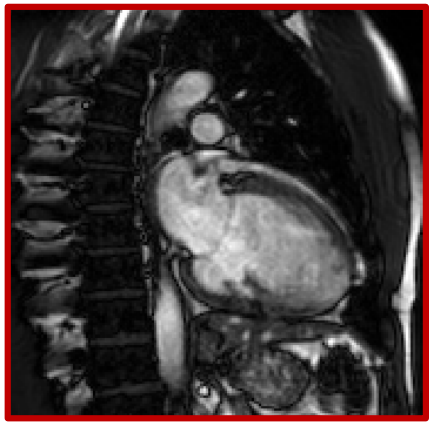

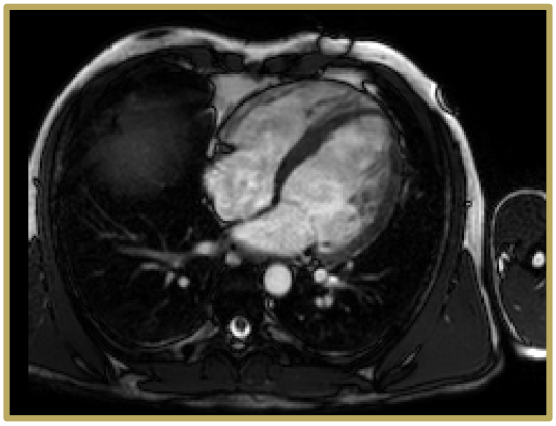

# Still, cardiac imaging is challenging: Low and High Resolution Images

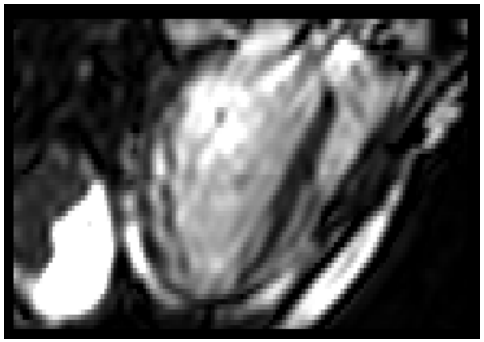

3D HR Image

PSF kernel and patient motion

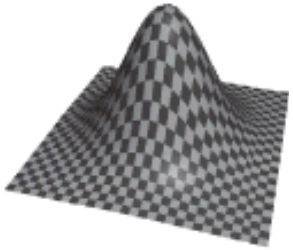

Sinc Filter

Down-sampling

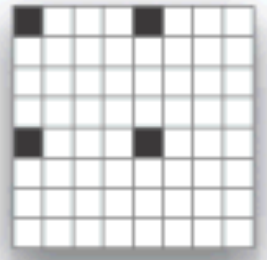

Sub-sampling grid

????

Output

Image Super Resolution Model

Input

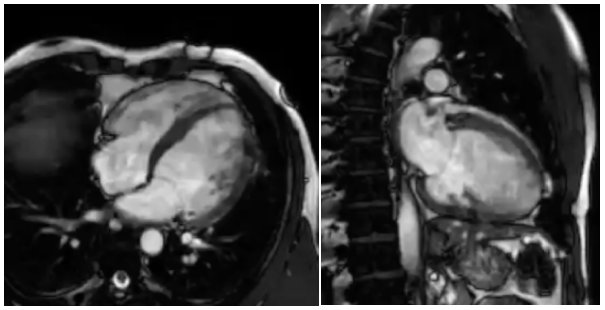

2D LAX Images

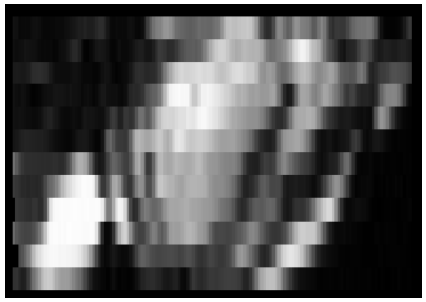

2D SAX Images

# Proposed 3D-SR Model (Single-Image)

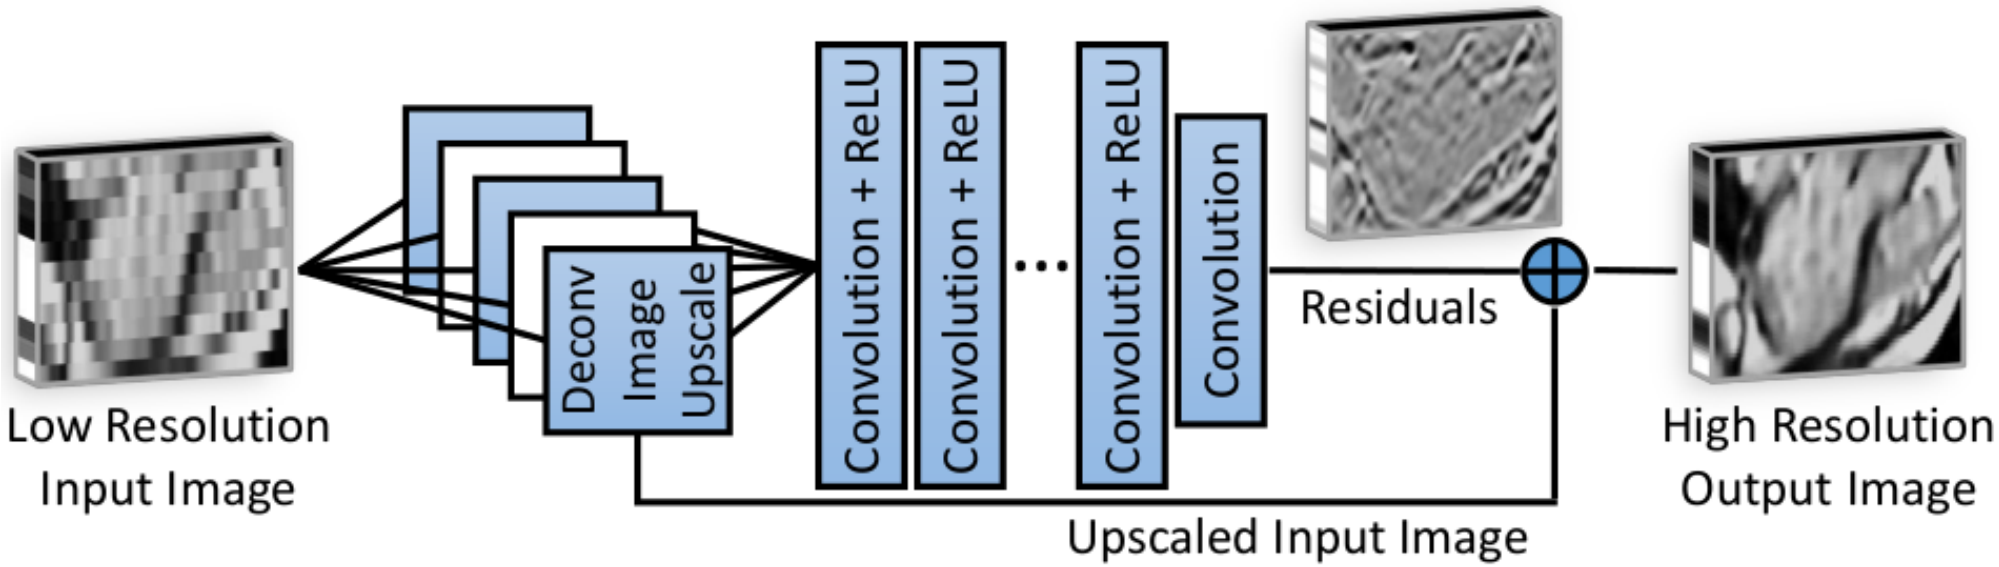


O. Oktay et al. MICCAI 2016

# Proposed 3D-SR Model (Multi-Image)

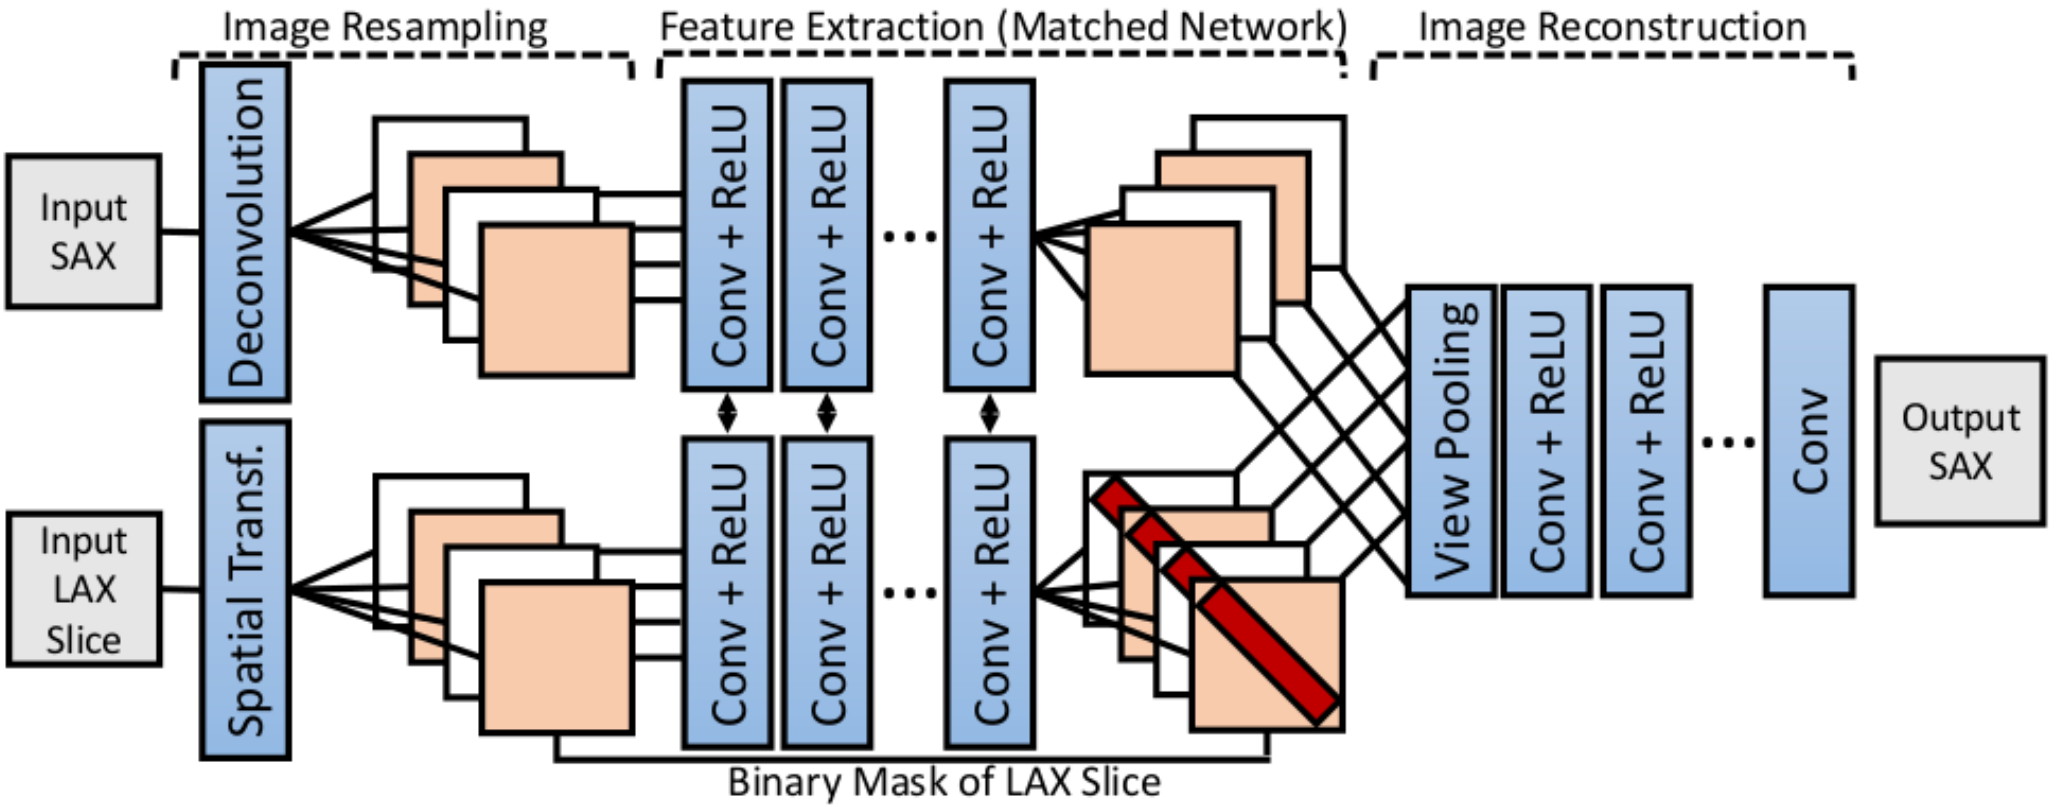


– Siamese model is used to combine information from multiple stacks

– The learned kernels can be easily integrated in this multi-model

O. Oktay et al. MICCAI 2016

# Image Quality Assessment

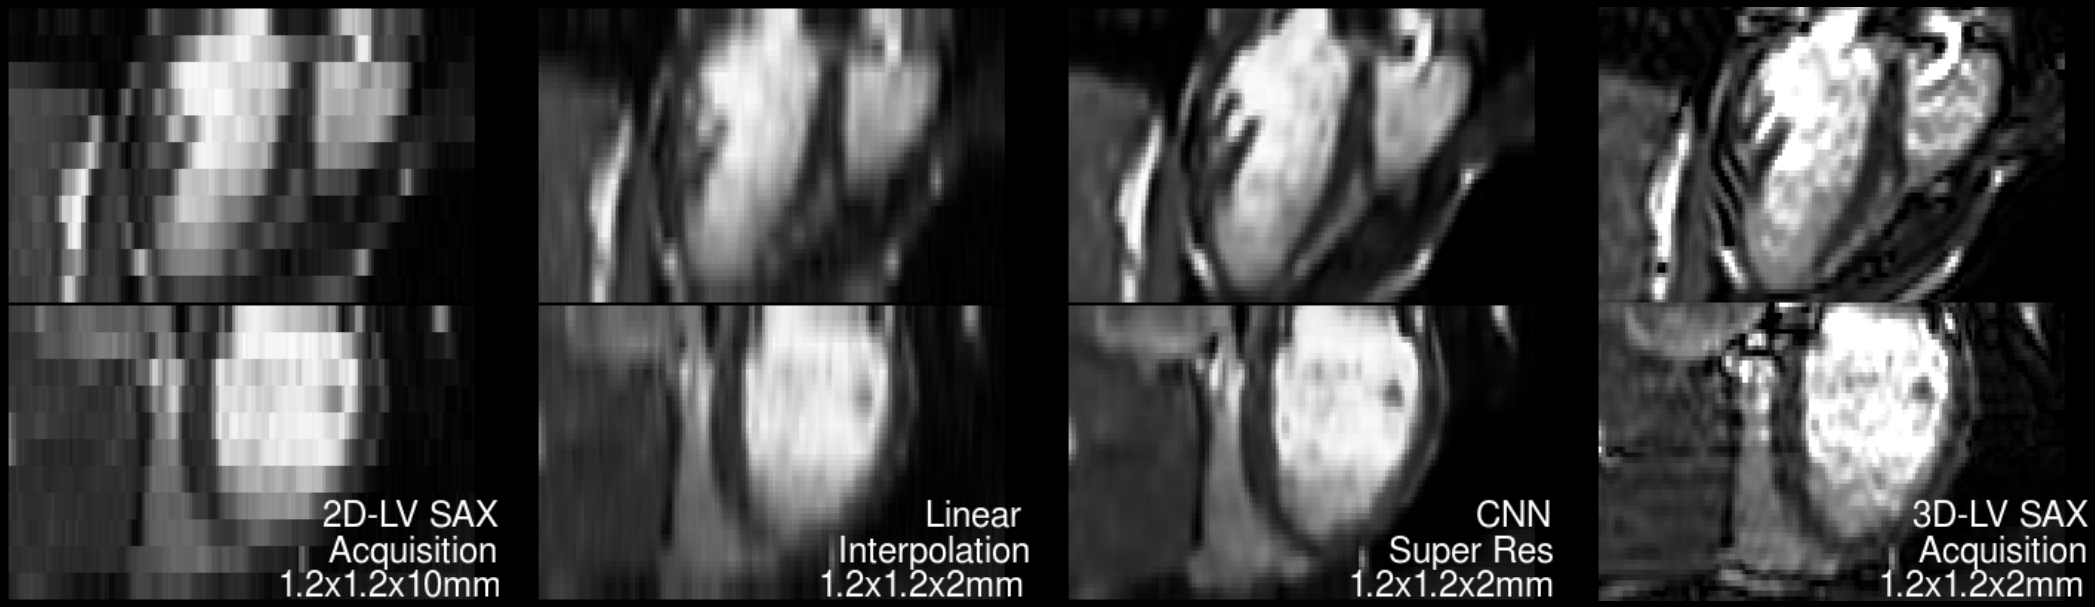


Low Resolution Input Image

Linear Interpolation

Proposed Method

High Resolution Ground-truth

Upsampling x5

Inference Time: 6-8 Seconds for image size (140x140x10)

O. Oktay et al. MICCAI 2016

# Motion Tracking Experiments
# (SR is used as a preprocessing method)

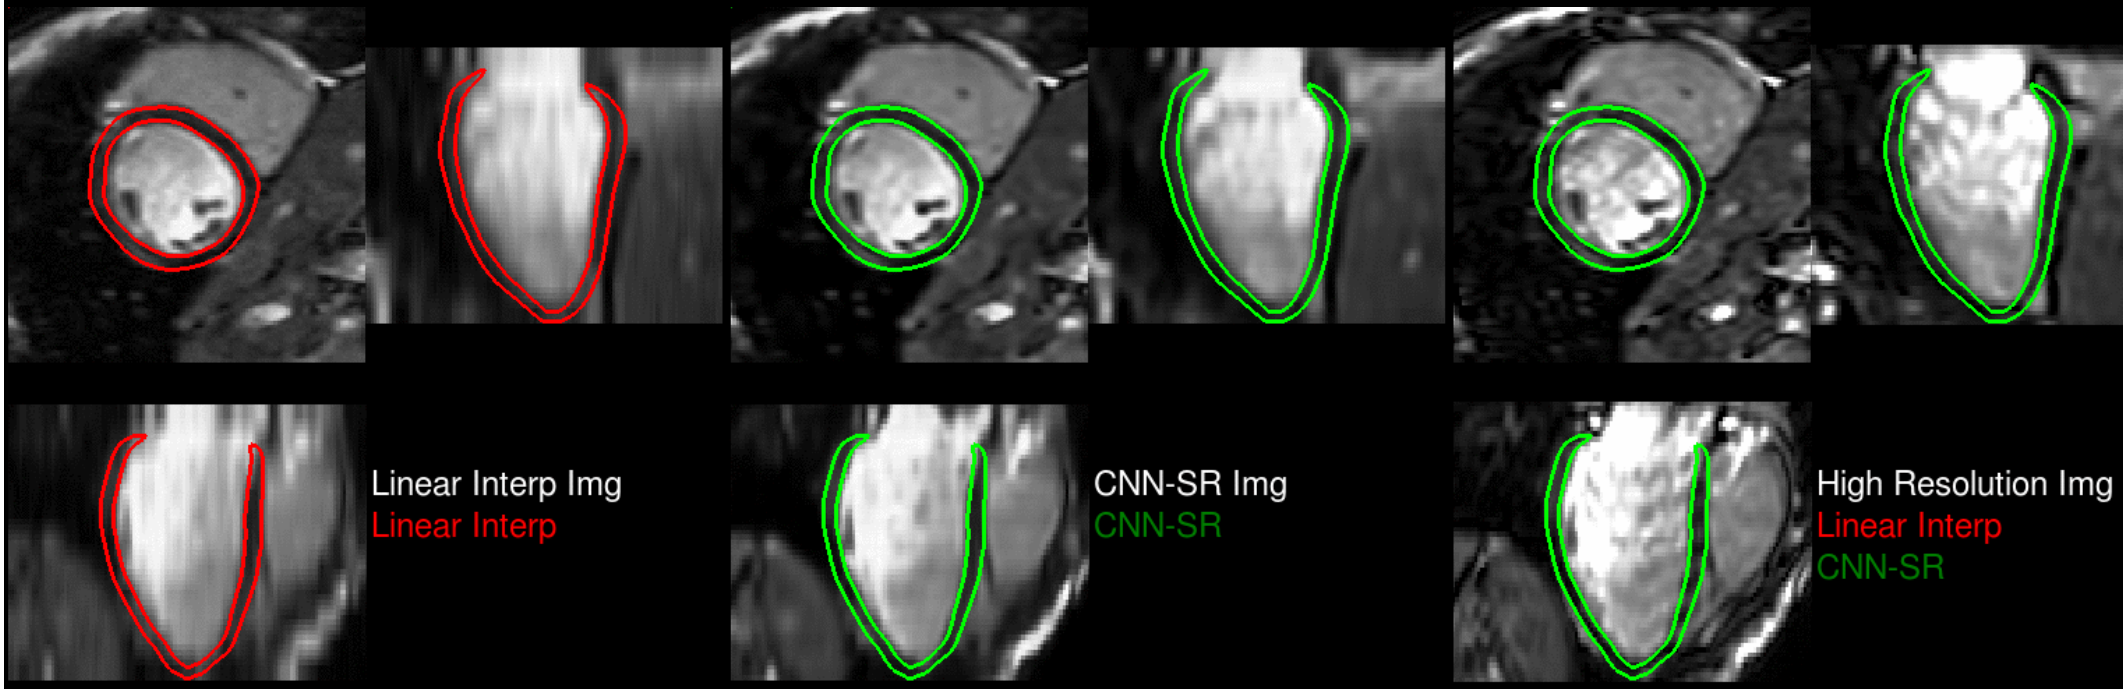


Surface to Surface Distance
(Linear vs HR) 5.50 mm

Surface to Surface Distance
(Proposed vs HR) 4.73 mm

# Overview

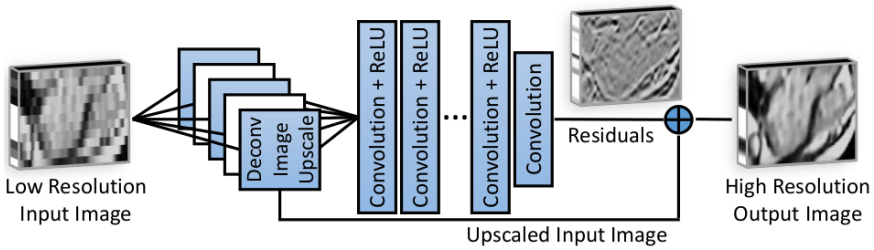


Image reconstruction &
super-resolution

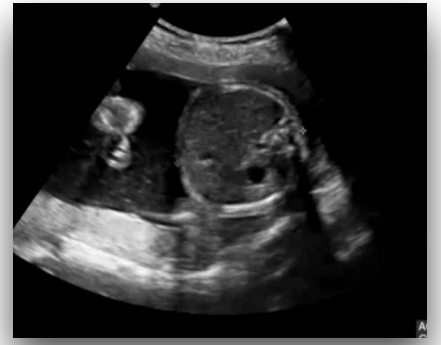

**Abdominal View**
Confidence: 98%

Image classification

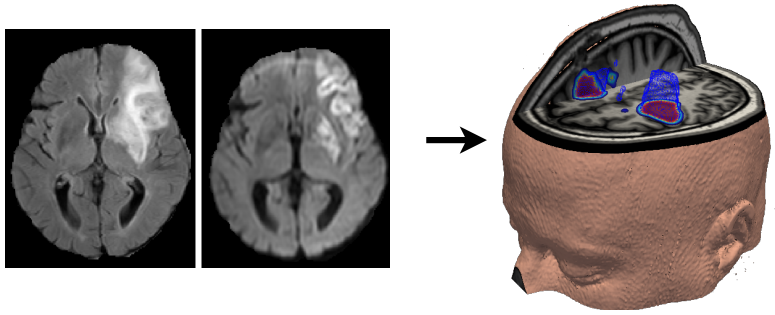

Image segmentation

# iFIND

intelligent Fetal Imaging and Diagnosis

- Aim: Better fetal screening with US
  - Use multiple US probes to acquire more comprehensive imaging data
  - Use robotic control of US probes to ensure wider field of view
  - Use machine learning for automated US image acquisition and interpretation

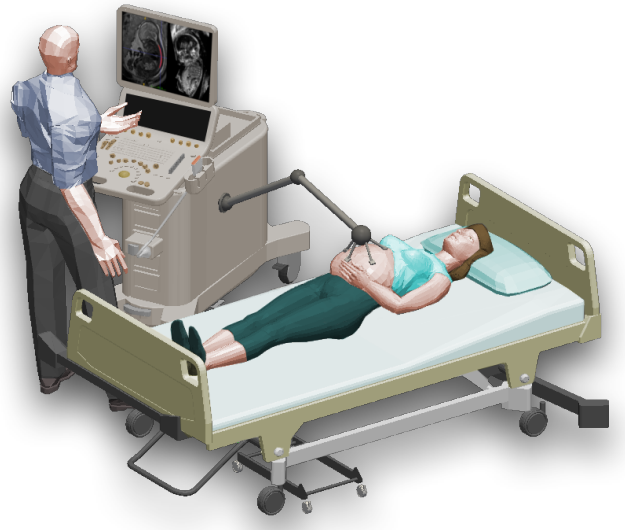

- Partners:

wellcome trust PHILIPS EPSRC

KING'S College LONDON Imperial College London UNIVERSITÀ DEGLI STUDI FIRENZE SickKids

# Automatic Standard Scan Plane Detection

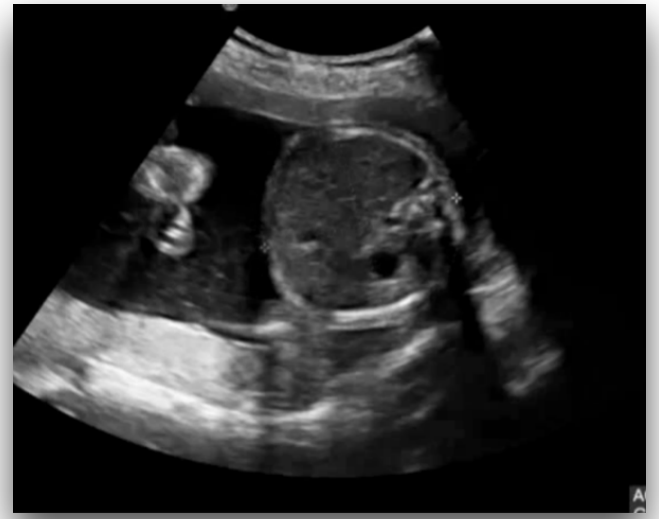

**Abdominal View**
Confidence: 98%

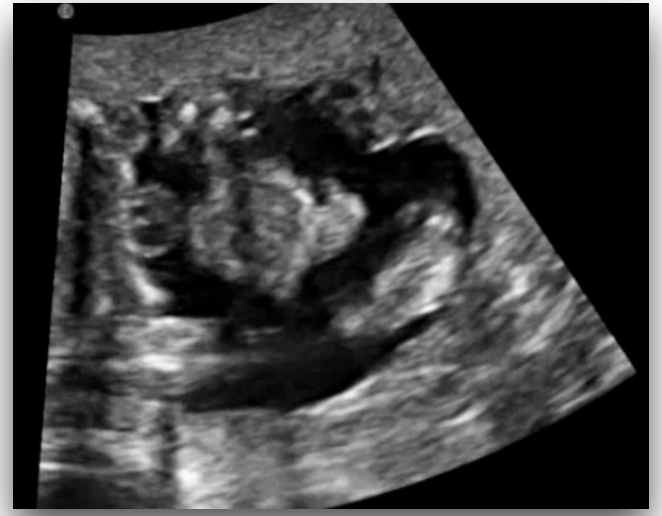

**Lips View**
Confidence: 96%

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

# Automatic Standard Scan Plane Detection

- Potential applications:
  - Guidance: Assist inexperienced sonographers

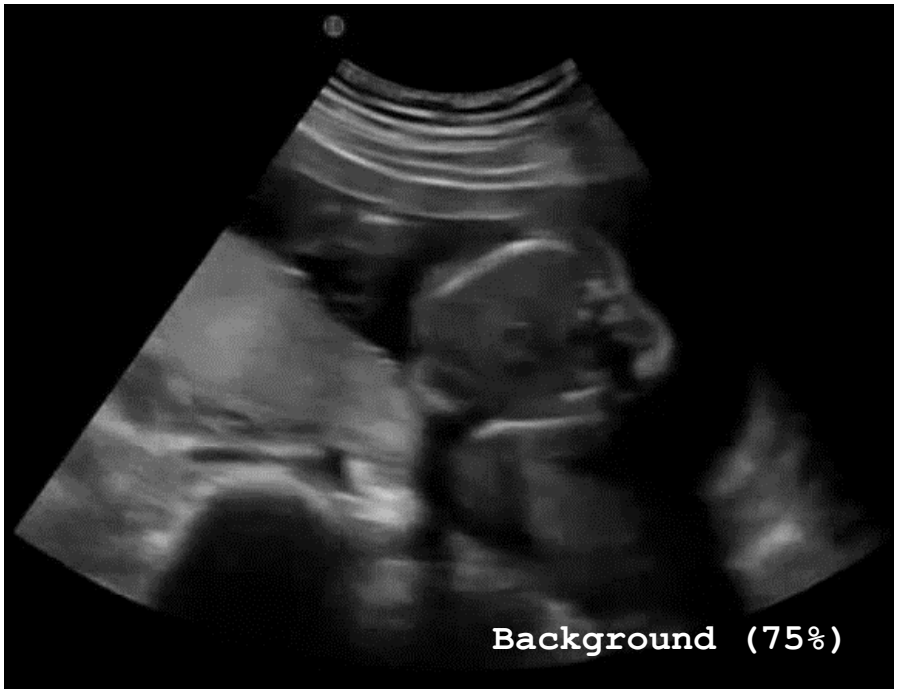

# Automatic Standard Scan Plane Detection

- Potential applications:
  - Guidance: Assist inexperienced sonographers
  - Convenience: Automatically make a check list of visited planes

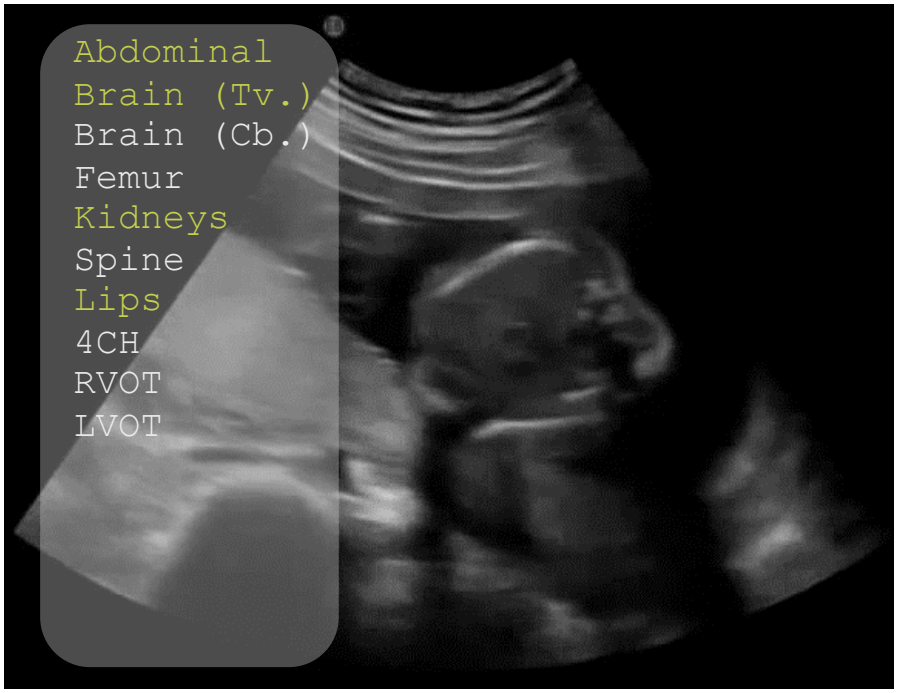

# Automatic Standard Scan Plane Detection

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

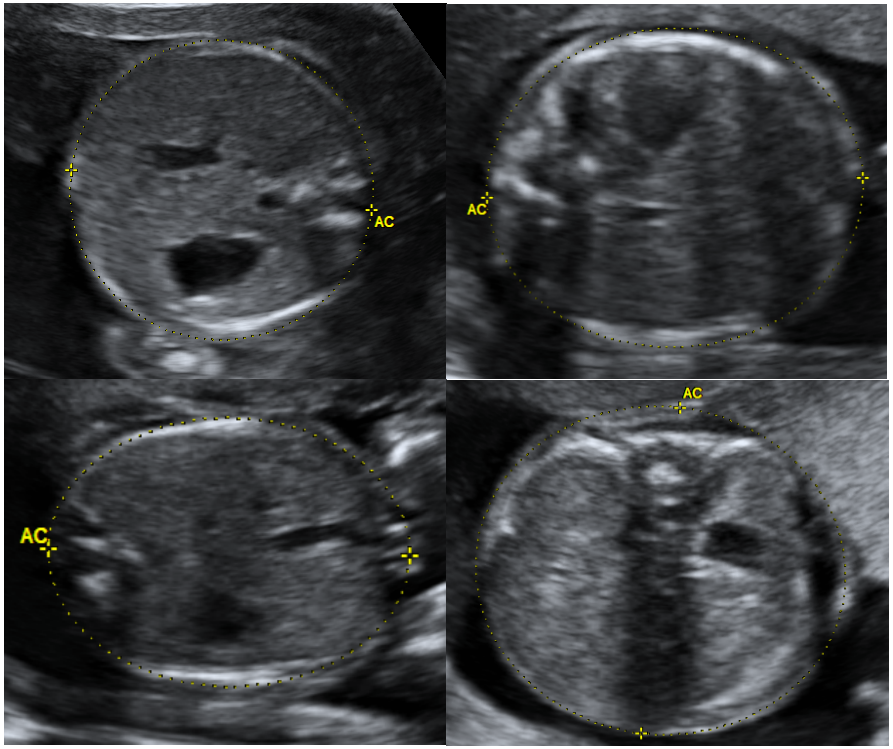

# Automatic Standard Scan Plane Detection: Method

- Fully convolutional neural network:

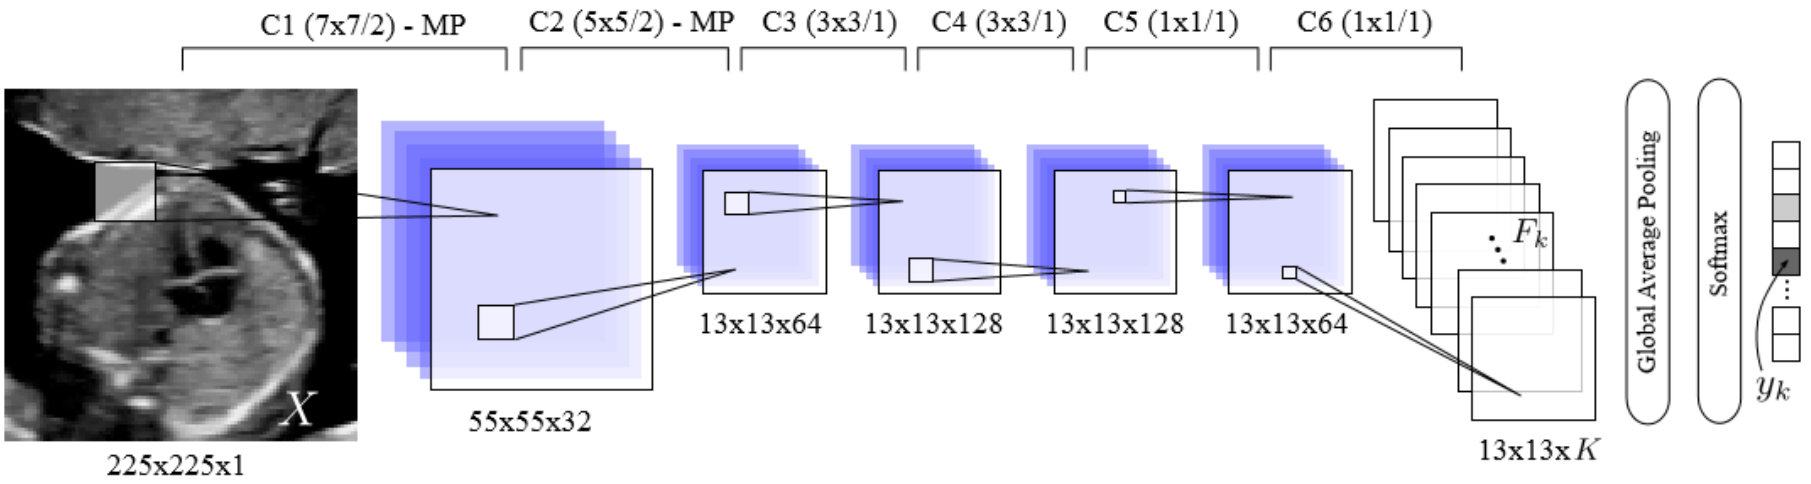


  - Very fast
  - Very accurate

C. Baumgartner et al. MICCAI 2016, IEEE-TMI 2017

# Automatic Standard Scan Plane Detection: Data

- We use very large 2D ultrasound dataset consisting of ***images*** of standard views and ***videos***

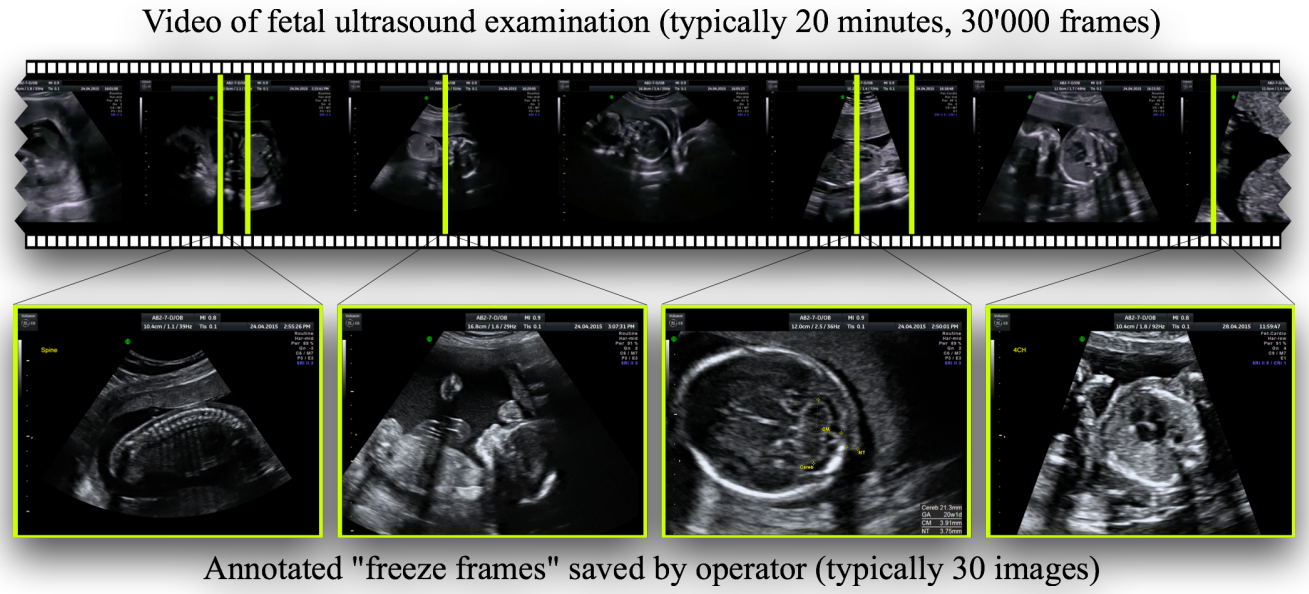


- Data from
  - 2700 patients
  - Between 1200 and 4800 images for each standard plane

# Demo

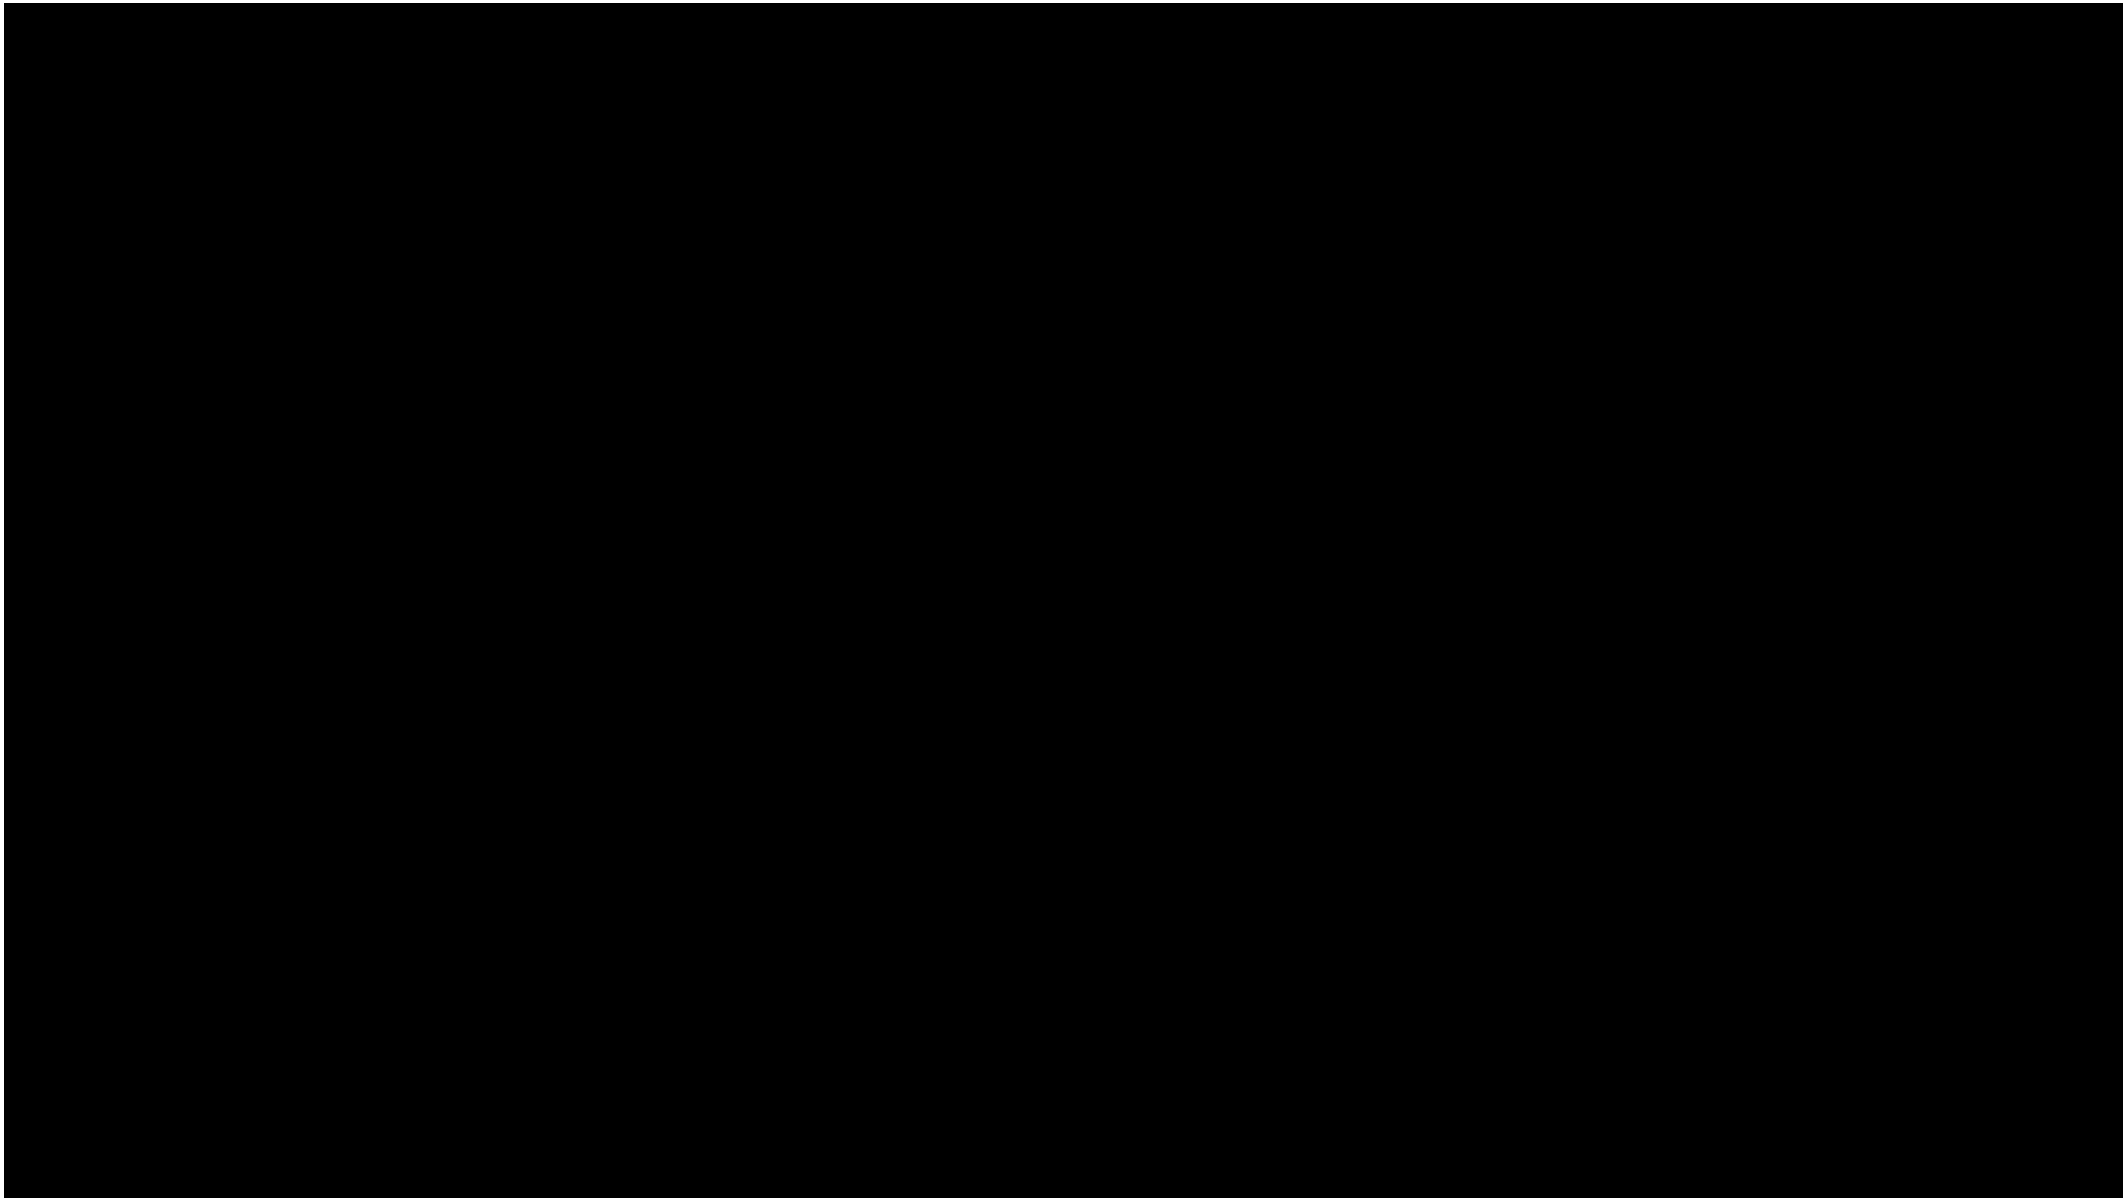

# Automatic Standard Scan Plane Detection: Localisation

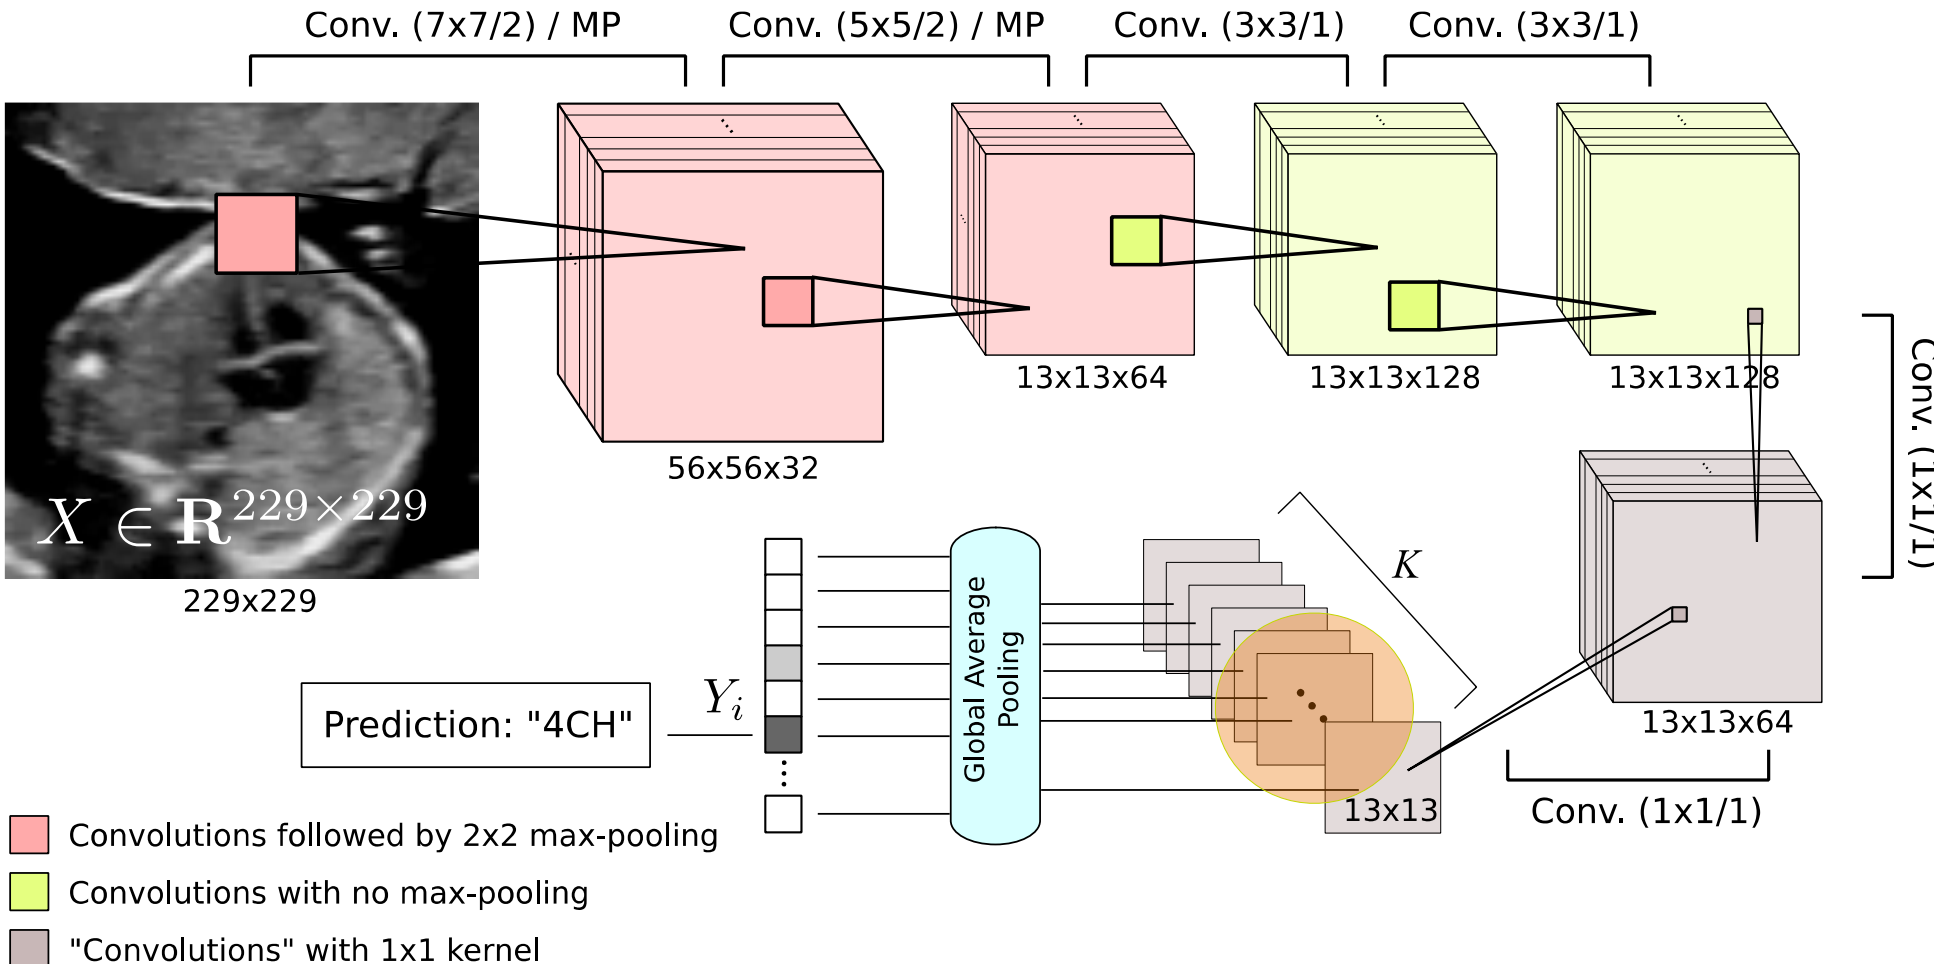


Localisation is (almost) for free in this framework!

C. Baumgartner et al. MICCAI 2016, IEEE-TMI 2017

# Automatic Standard Scan Plane Detection: Localisation

- Can also identify which regions of a frame caused it to make a particular prediction

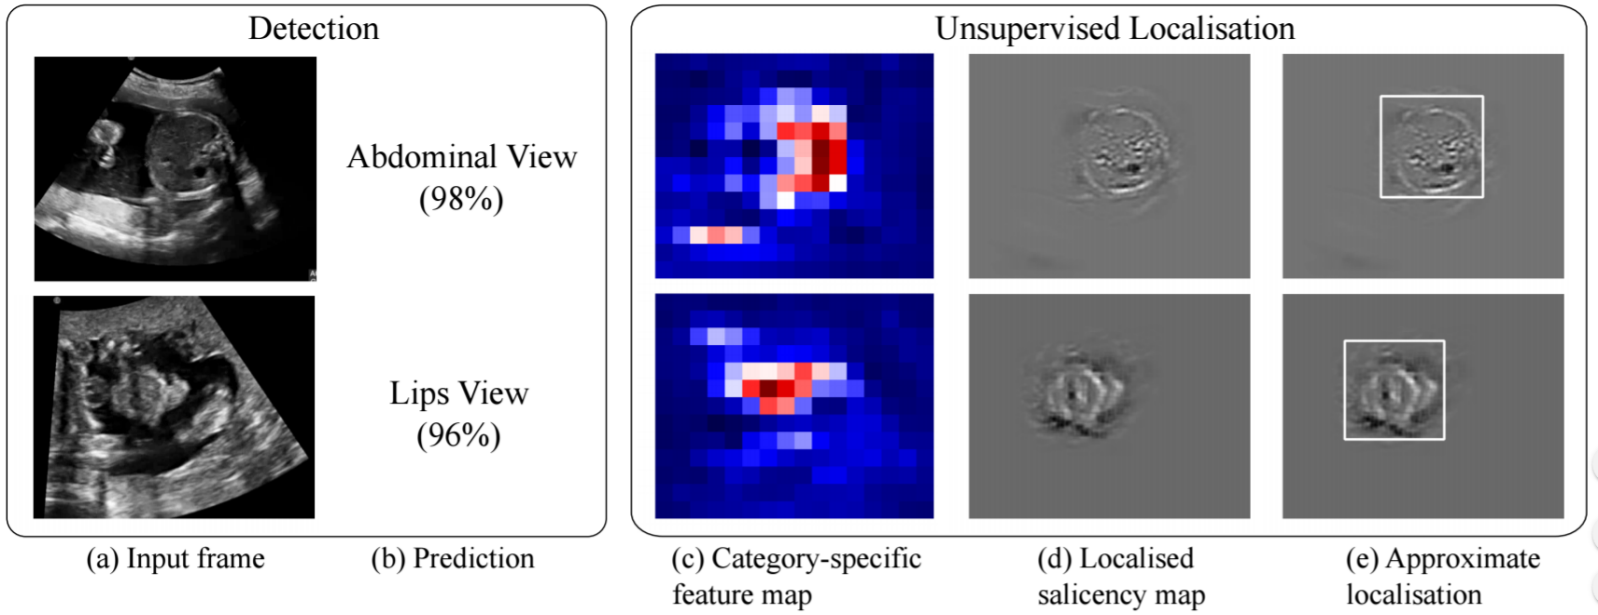


- This can be used for localisation of the fetal anatomy without having bounding boxes for training

C. Baumgartner et al. MICCAI 2016, IEEE-TMI 2017

# Demo

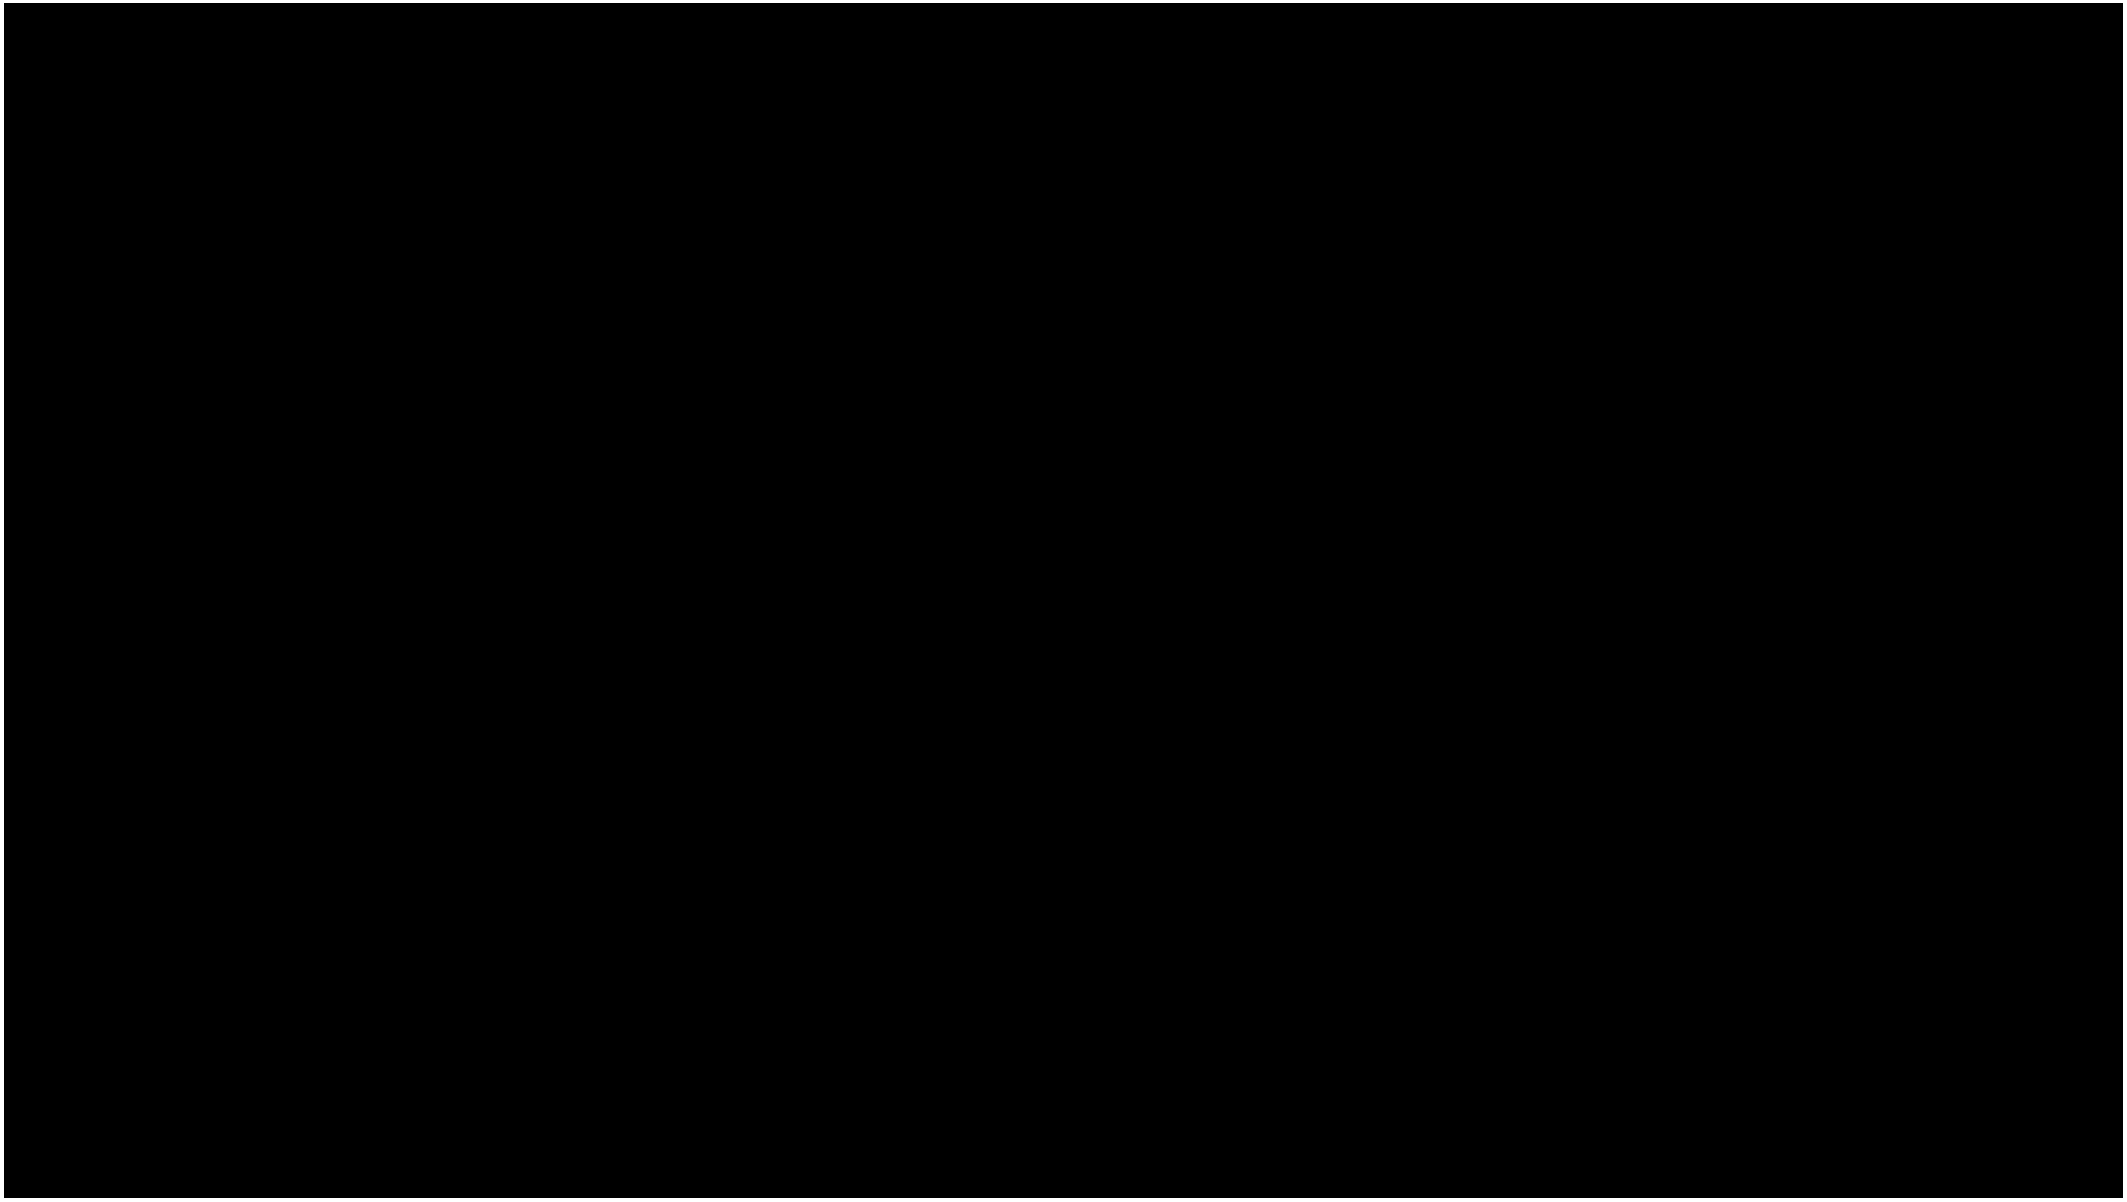

# Standard Plane Detection from 3D Ultrasound

- Motivation
  - 3D ultrasound data is hard to interpret directly on the US scanner
  - We aim for a system that can automatically extract standard 2D views from a 3D view at any probe position
- Eventually, we would like this to work in real-time

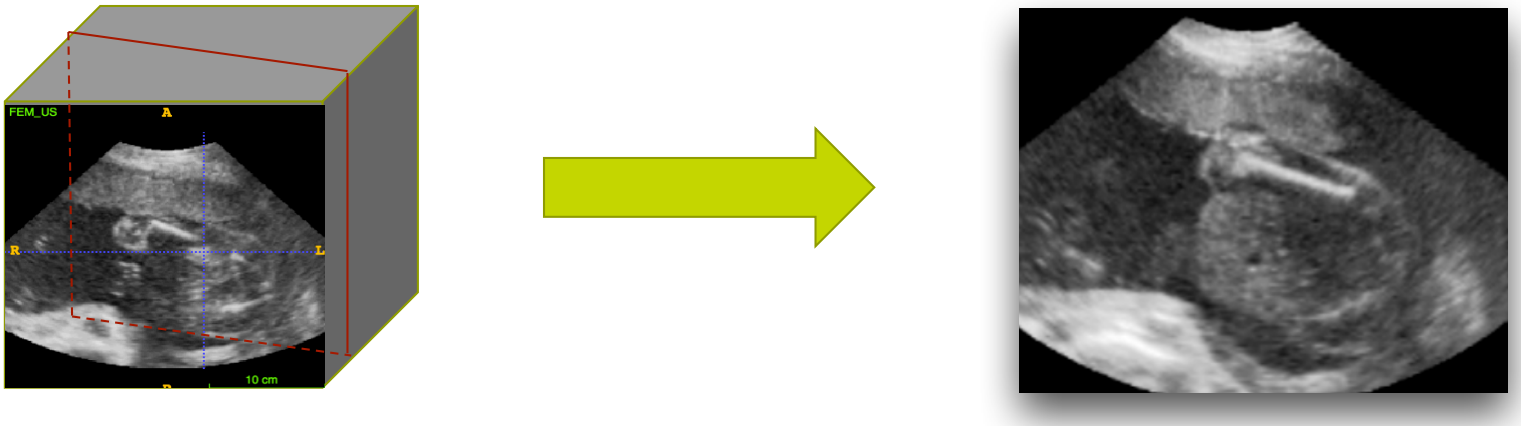

# Overview

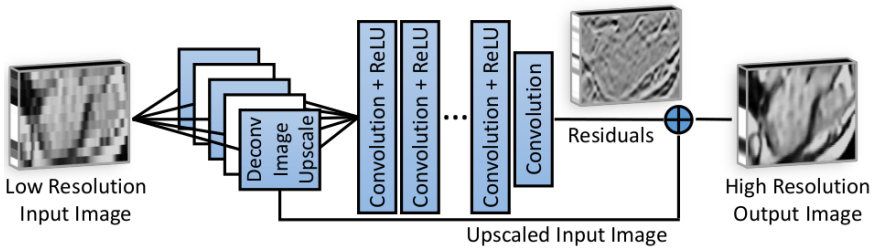


Image reconstruction &
super-resolution

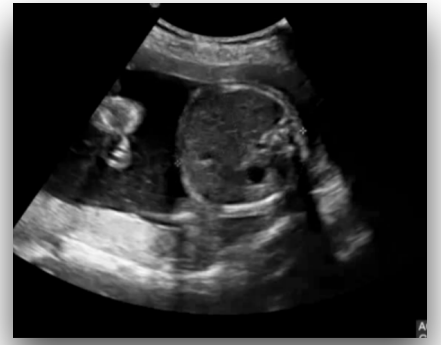

**Abdominal View**
Confidence: 98%

Image classification

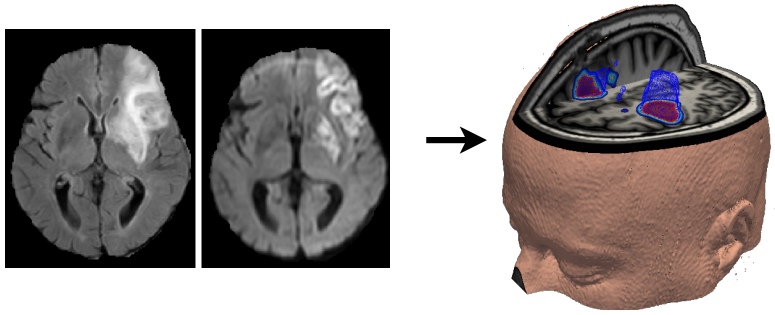

Image segmentation

# Convolutional Neural Networks

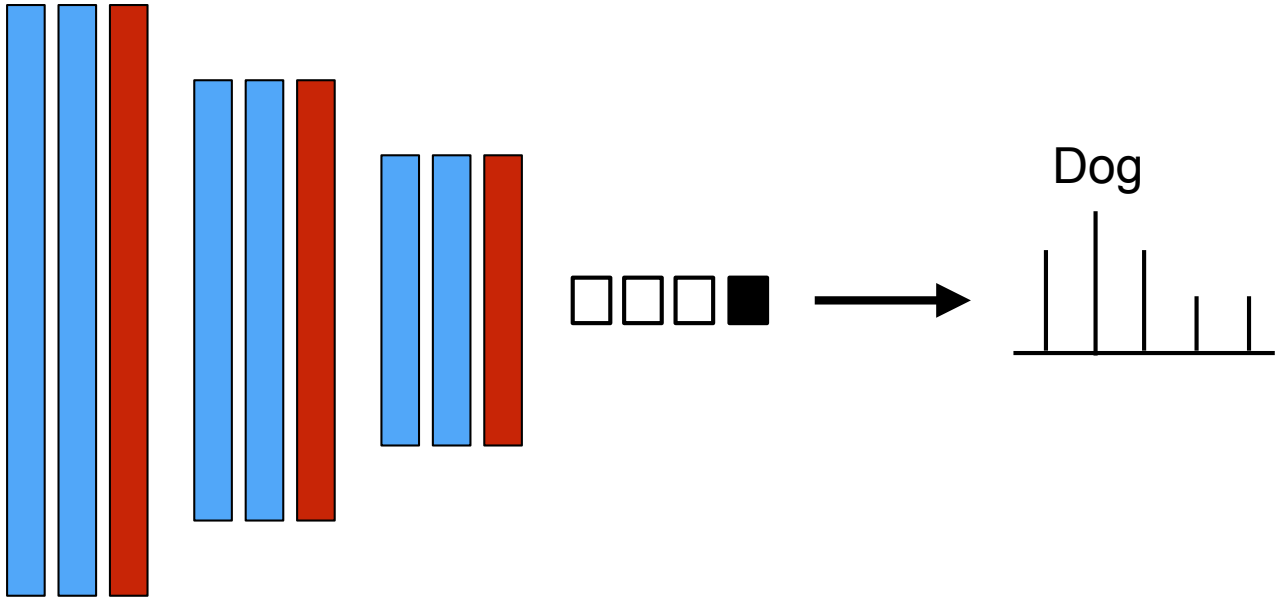


Convolution + RELU

Fully connected layer

Max pooling

Softmax

# Convolutional Neural Networks

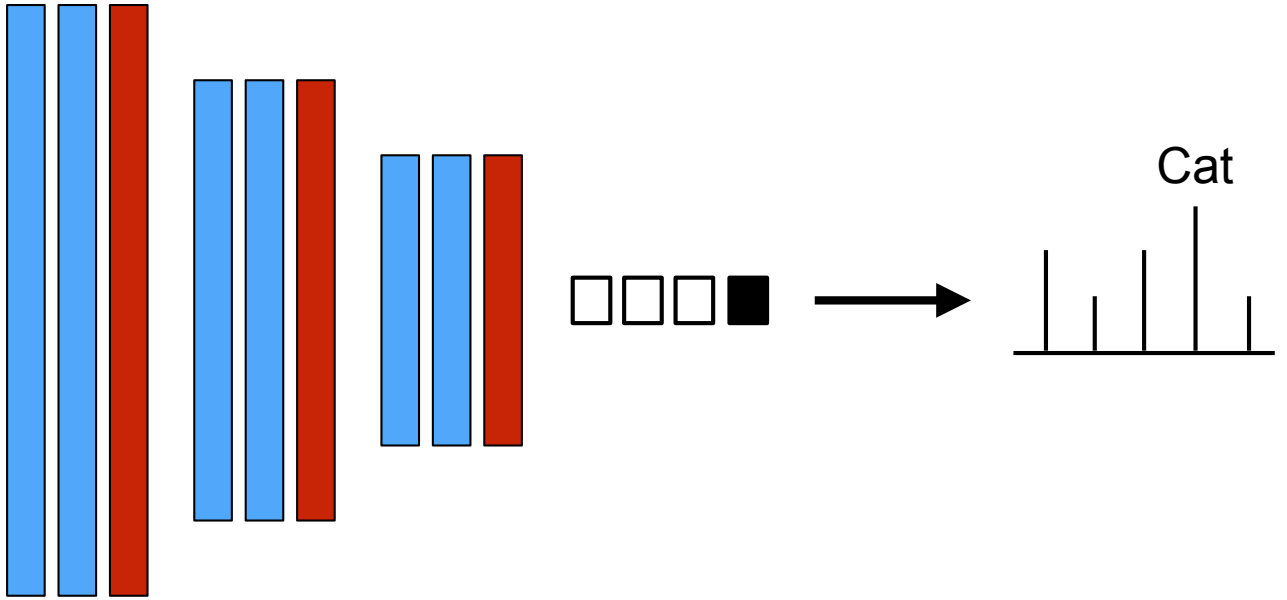


Convolution + RELU

Fully connected layer

Max pooling

Softmax

# Convolutional Neural Networks for Image Segmentation

- Different architectures:
  - Fully convolutional networks (Long et al., 2015)
  - U-Net (Ronneberger et al., 2015)
  - DeepMedic (Kamnitsas et al., 2016)

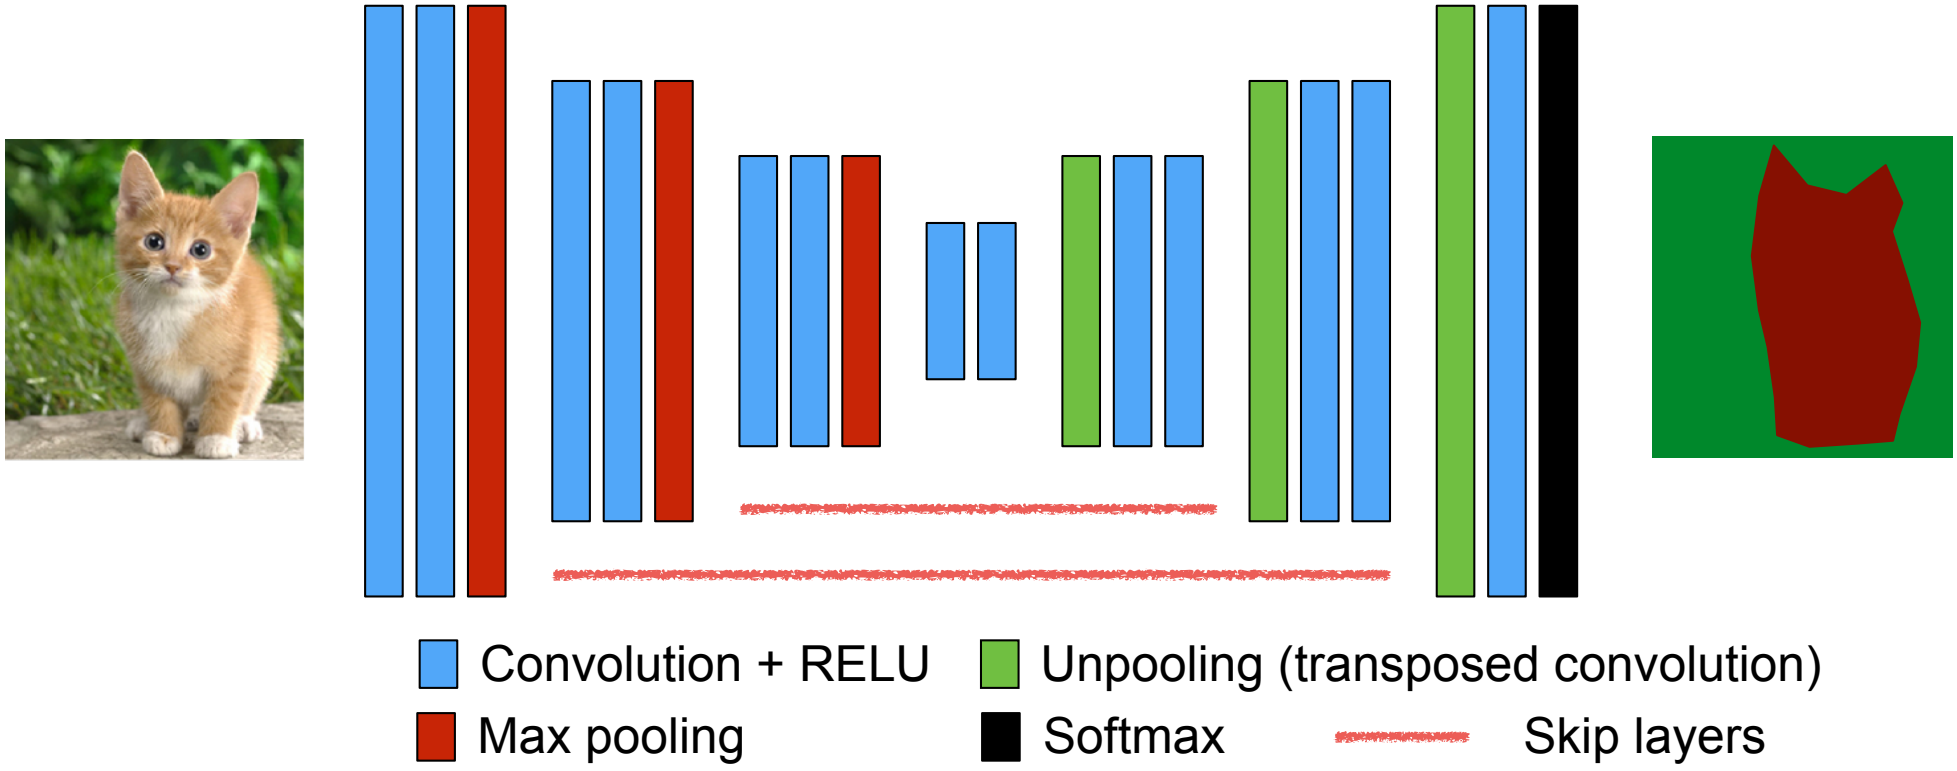

# Convolutional Neural Networks for Image Segmentation

- Different architectures:
  - Fully convolutional networks (Long et al., 2015)
  - U-Net (Ronneberger et al., 2015)
  - DeepMedic (Kamnitsas et al., 2016)

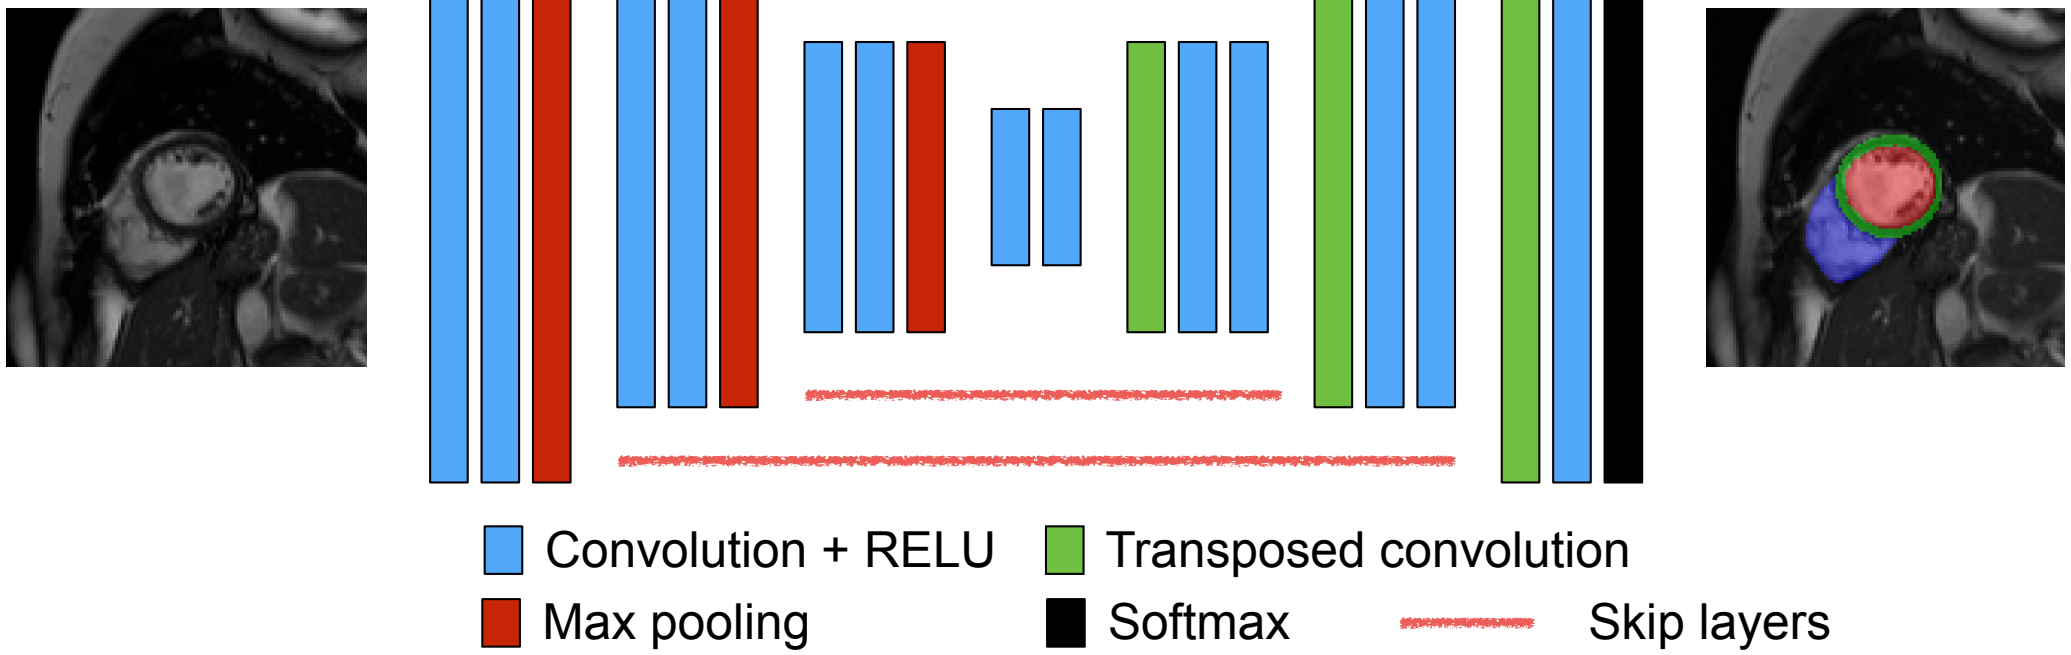

# Image segmentation as a machine learning problem

- Fully connected networks (Long et al., 2015)
- Manual annotations of **4,872 subjects** (QMUL/Oxford) with **93,128 pixelwise annotated 2D images** slices
- Divided into training/validation/test: 3,972/300/600

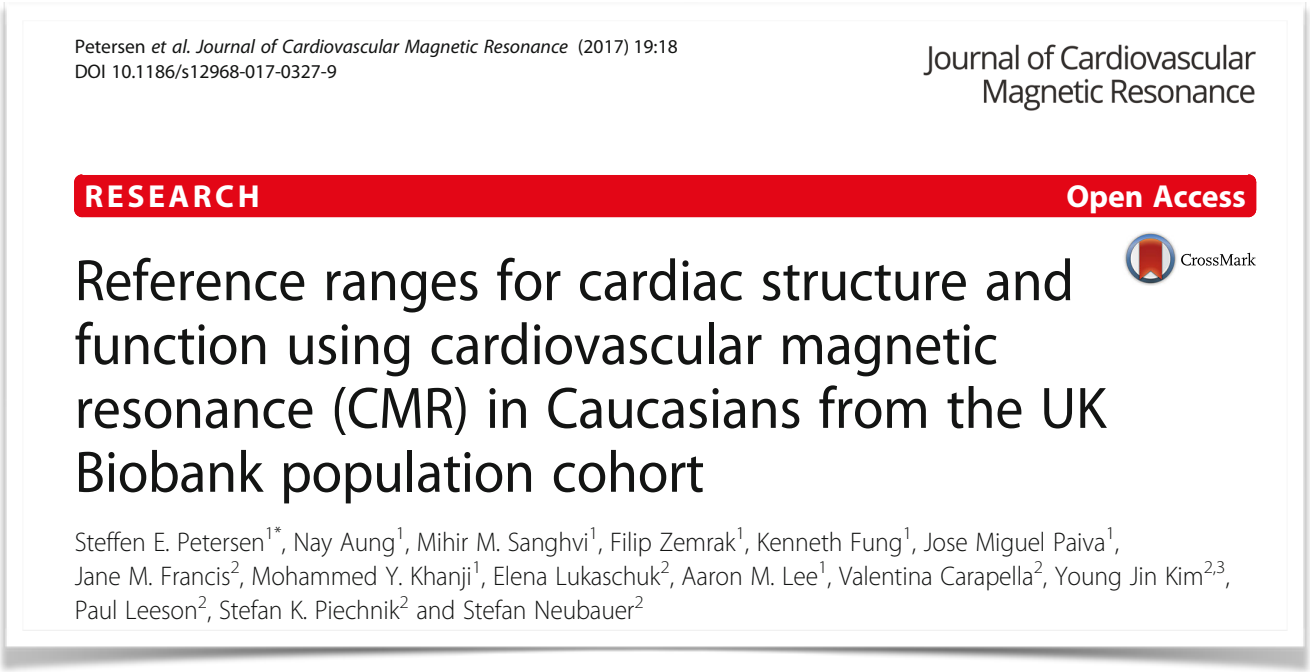

Petersen *et al. Journal of Cardiovascular Magnetic Resonance* (2017) 19:18
DOI 10.1186/s12968-017-0327-9

Journal of Cardiovascular Magnetic Resonance

RESEARCH Open Access

Reference ranges for cardiac structure and function using cardiovascular magnetic resonance (CMR) in Caucasians from the UK Biobank population cohort

Steffen E. Petersen[1*], Nay Aung[1], Mihir M. Sanghvi[1], Filip Zemrak[1], Kenneth Fung[1], Jose Miguel Paiva[1], Jane M. Francis[2], Mohammed Y. Khanji[1], Elena Lukaschuk[2], Aaron M. Lee[1], Valentina Carapella[2], Young Jin Kim[2,3], Paul Leeson[2], Stefan K. Piechnik[2] and Stefan Neubauer[2]

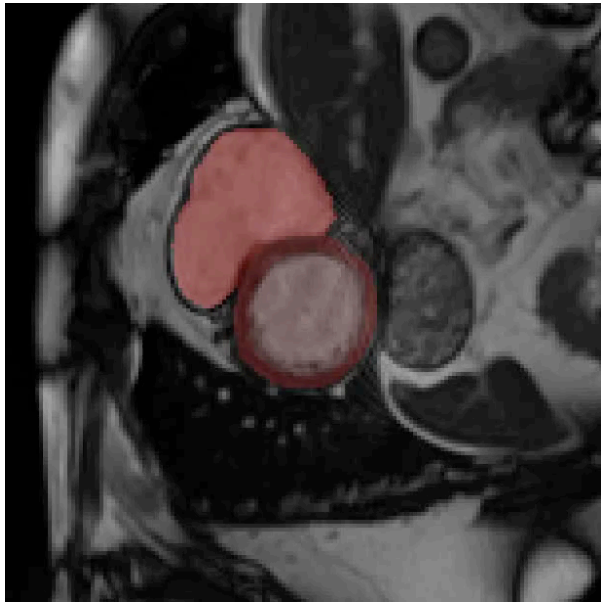

SA, basal

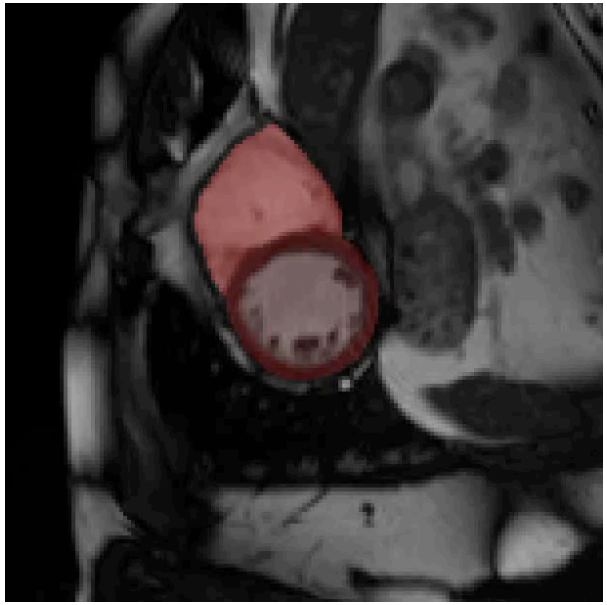

SA, mid-ventricular

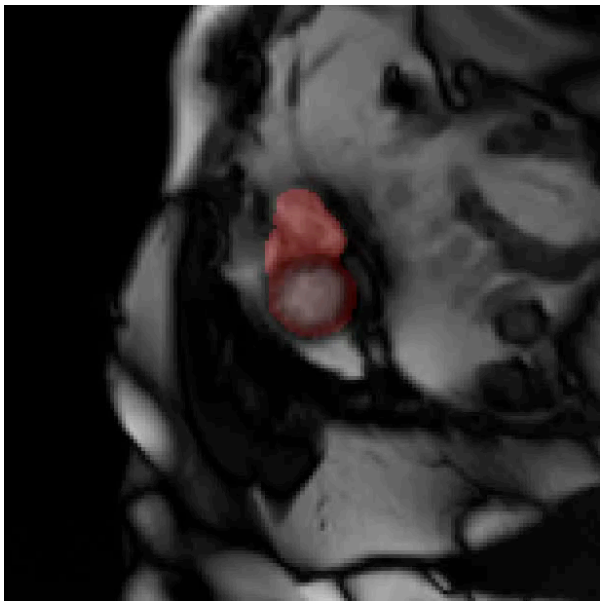

SA, apical

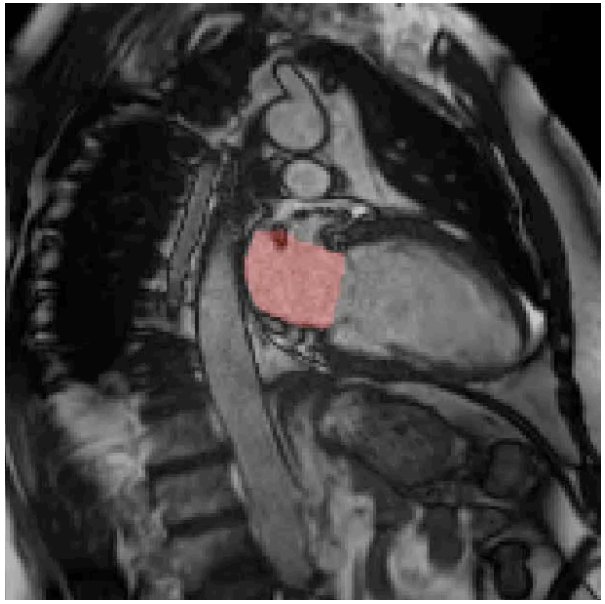

LA, 2 chamber

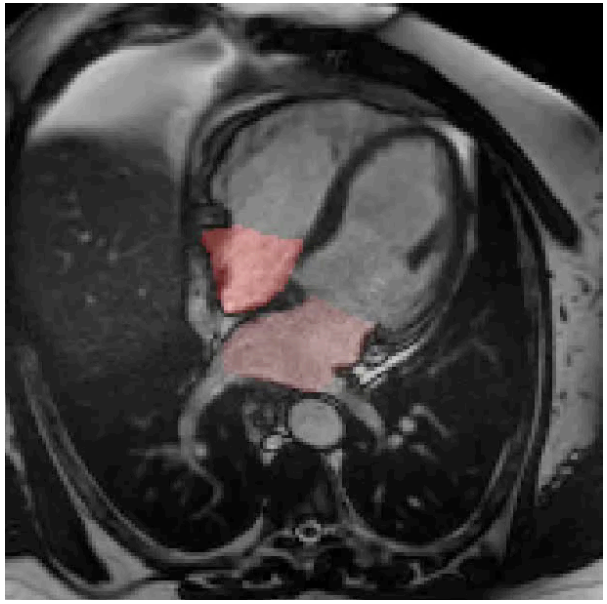

LA, 4 chamber

# Evaluation of segmentation accuracy
# Comparison to expert observers

(a) Absolute difference

| | Auto vs Man (n = 600) | O1 vs O2 (n = 50) | O2 vs O3 (n = 50) | O3 vs O1 (n = 50) |
|---|---|---|---|---|
| LVEDV (mL) | $6.1_{\pm 5.3}$ | $6.1_{\pm 4.4}$ | $8.8_{\pm 4.8}$ | $4.8_{\pm 3.1}$ |
| LVESV (mL) | $5.3_{\pm 4.9}$ | $4.1_{\pm 4.2}$ | $6.7_{\pm 4.2}$ | $7.1_{\pm 3.8}$ |
| LVM (gram) | $6.9_{\pm 5.5}$ | $4.2_{\pm 3.2}$ | $6.6_{\pm 4.9}$ | $6.5_{\pm 4.8}$ |
| RVEDV (mL) | $8.5_{\pm 7.1}$ | $11.1_{\pm 7.2}$ | $6.2_{\pm 4.6}$ | $8.7_{\pm 5.8}$ |
| RVESV (mL) | $7.2_{\pm 6.8}$ | $15.6_{\pm 7.8}$ | $6.6_{\pm 5.5}$ | $11.7_{\pm 6.9}$ |

(b) Relative difference

| | Auto vs Man (n = 600) | O1 vs O2 (n = 50) | O2 vs O3 (n = 50) | O3 vs O1 (n = 50) |
|---|---|---|---|---|
| LVEDV (%) | $4.1_{\pm 3.5}$ | $4.2_{\pm 3.1}$ | $6.3_{\pm 3.3}$ | $3.4_{\pm 2.2}$ |
| LVESV (%) | $9.5_{\pm 9.5}$ | $6.8_{\pm 7.5}$ | $12.5_{\pm 8.5}$ | $11.7_{\pm 5.1}$ |
| LVM (%) | $8.3_{\pm 7.6}$ | $4.4_{\pm 3.3}$ | $6.0_{\pm 3.7}$ | $6.7_{\pm 4.6}$ |
| RVEDV (%) | $5.6_{\pm 4.6}$ | $8.0_{\pm 5.0}$ | $4.2_{\pm 3.1}$ | $5.7_{\pm 3.6}$ |
| RVESV (%) | $11.8_{\pm 12.2}$ | $30.6_{\pm 15.5}$ | $10.9_{\pm 8.3}$ | $16.9_{\pm 9.2}$ |

Automated

Manual

# Challenges for image segmentation: Pathologies

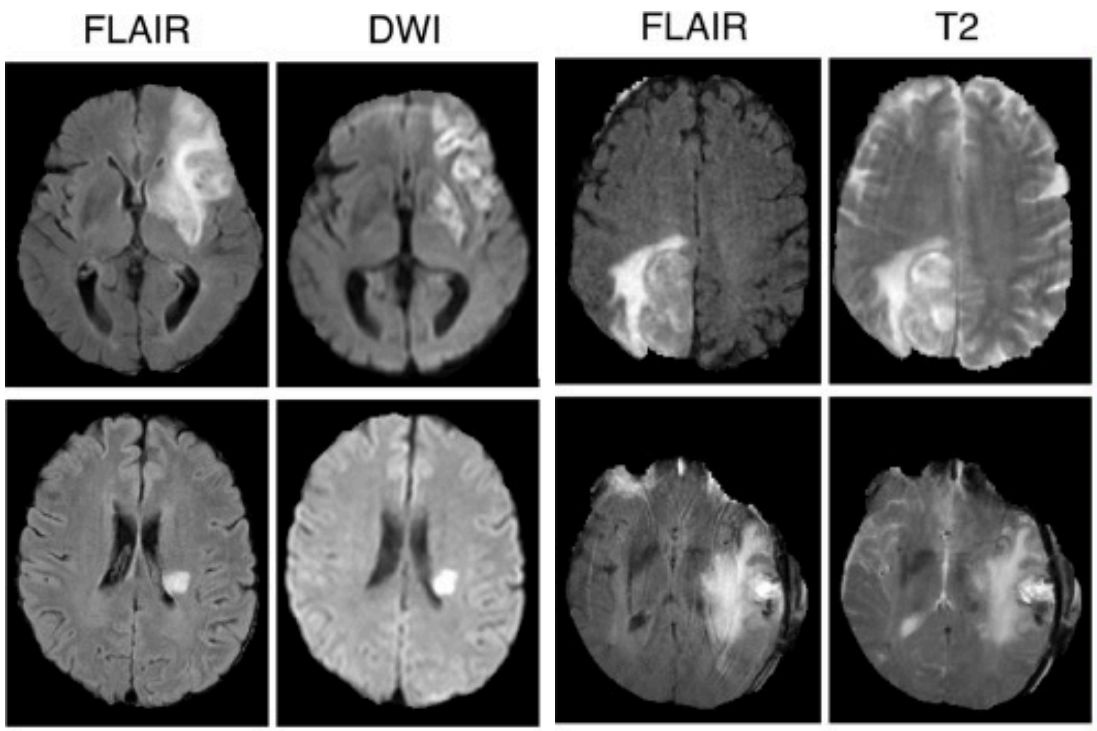


May be present or absent

Location & size of pathologies varies significantly

Appearance can local or diffuse

# DeepMedic: Overview

- Baseline CNN architecture:
  - Four layers with $5^3$ kernels for feature extraction, leading to a receptive field of size $17^3$.
  - The classification layer is implemented as convolutional with $1^3$ kernels, which enables efficient dense inference.
  - Cross-entropy as loss function

$$\mathcal{L}_{seg} = -\frac{1}{m}\sum_{i=1}^{m}[f(x_i) = y_i]log(f(x_i))$$

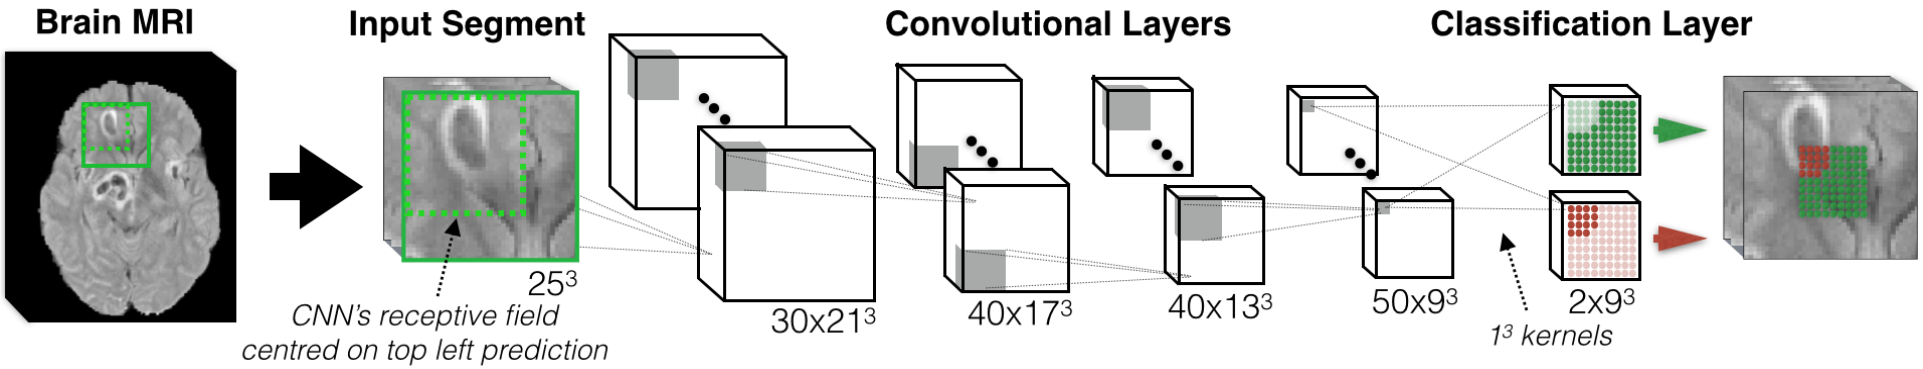


K. Kamnitsas et al. Medical Image Analysis, 2016

# DeepMedic: Overview

- The size of receptive field of CNNs is important:
  - Large receptive increases computation and memory requirements
  - Pooling leads to loss of the spatial information
- Solution: Use multi-scale approach

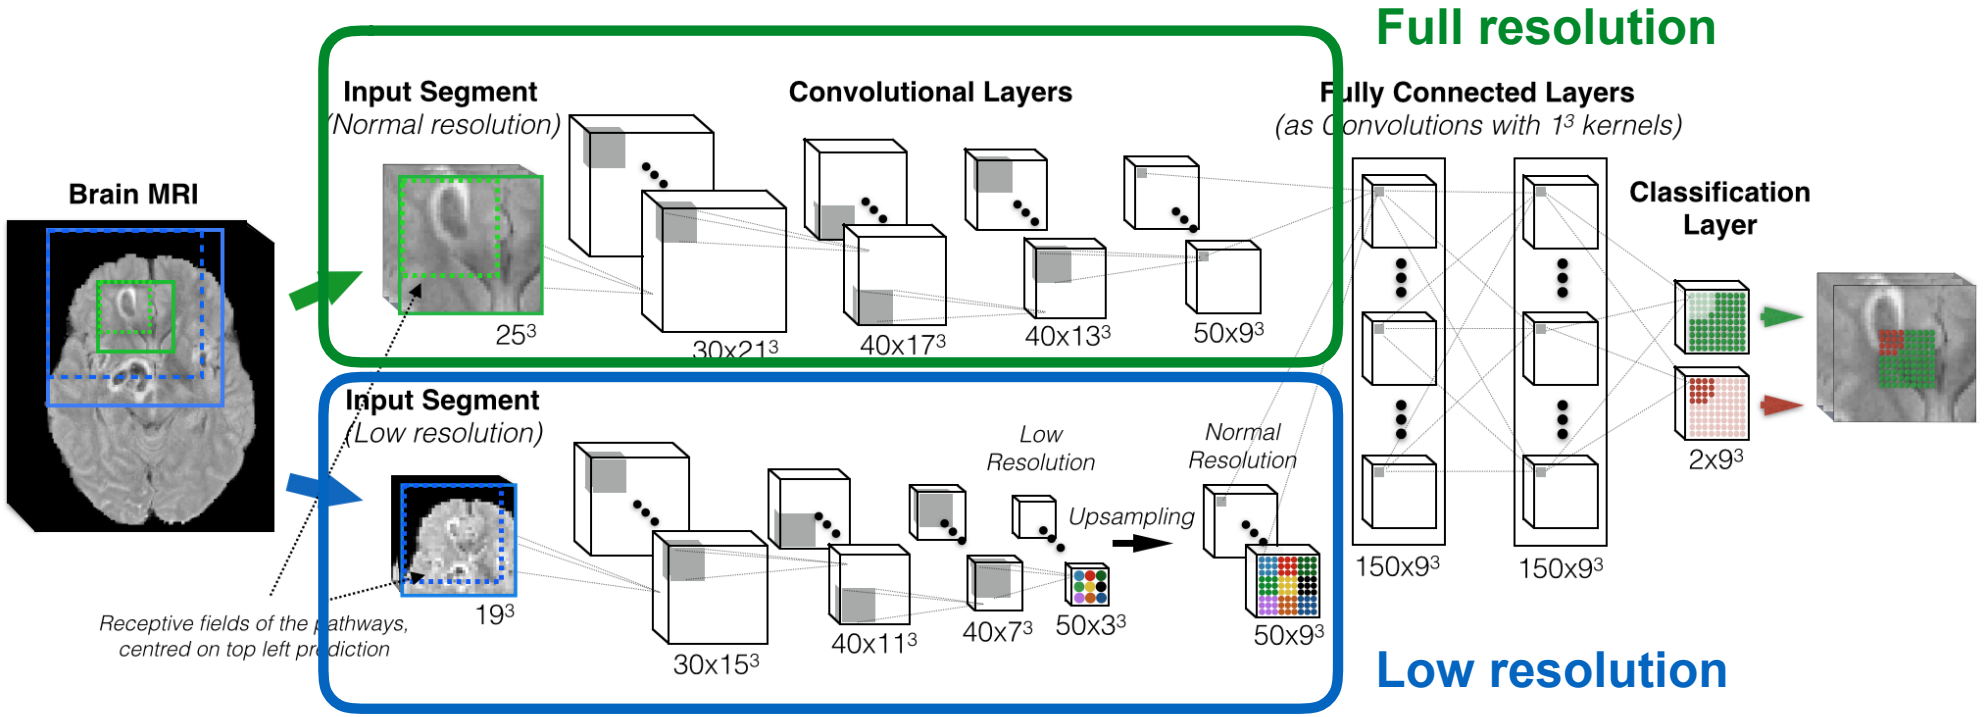


K. Kamnitsas et al. Medical Image Analysis, 2016

# DeepMedic: Results

Patients with severe TBI

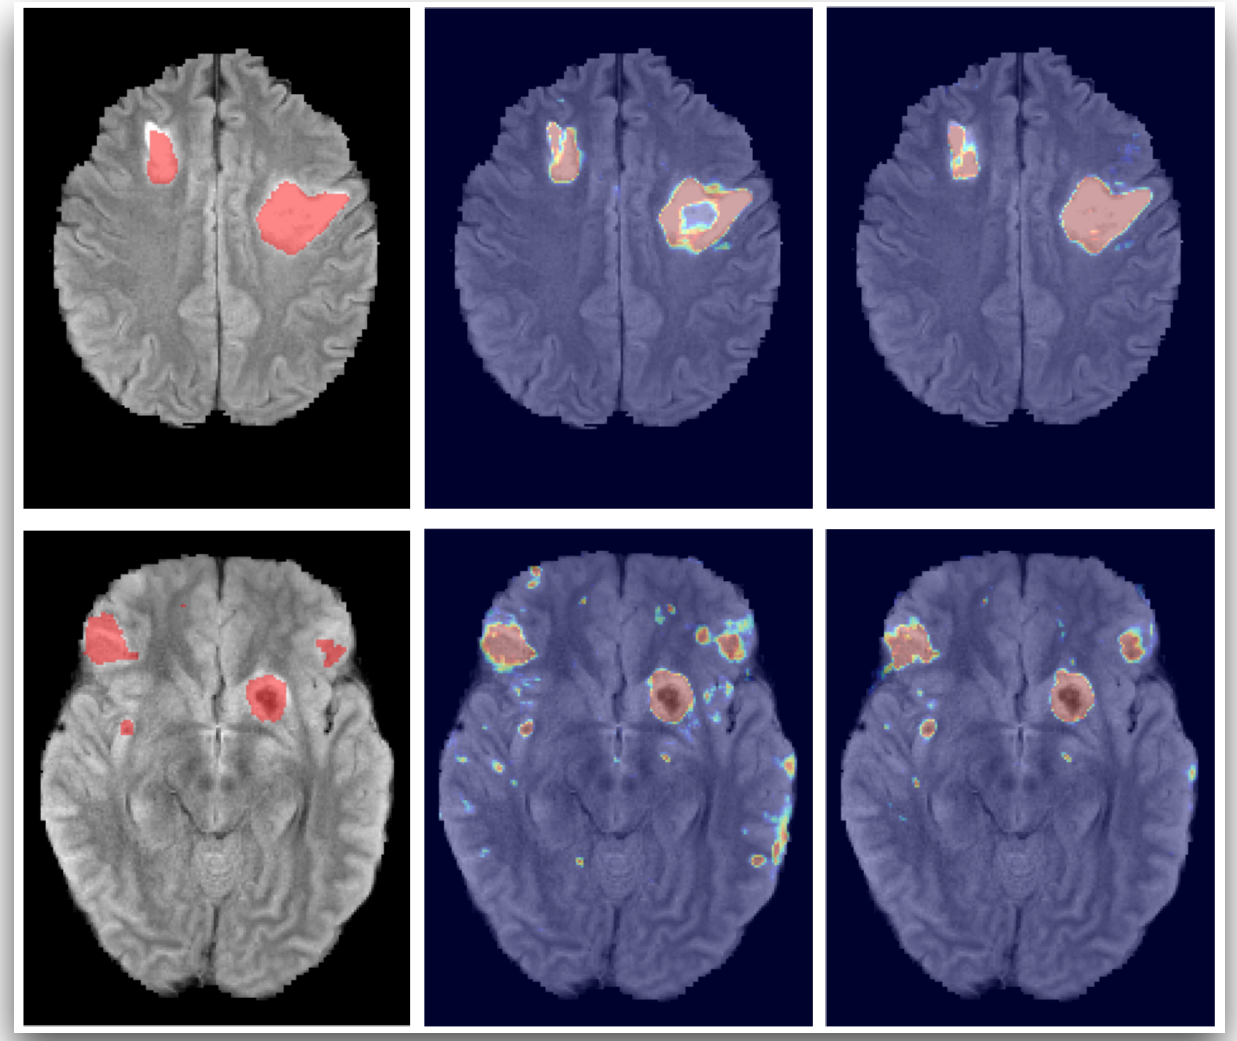

Manual

DeepMedic (single scale)

DeepMedic (multi-scale)

K. Kamnitsas et al. Medical Image Analysis, 2016

# DeepMedic in Action

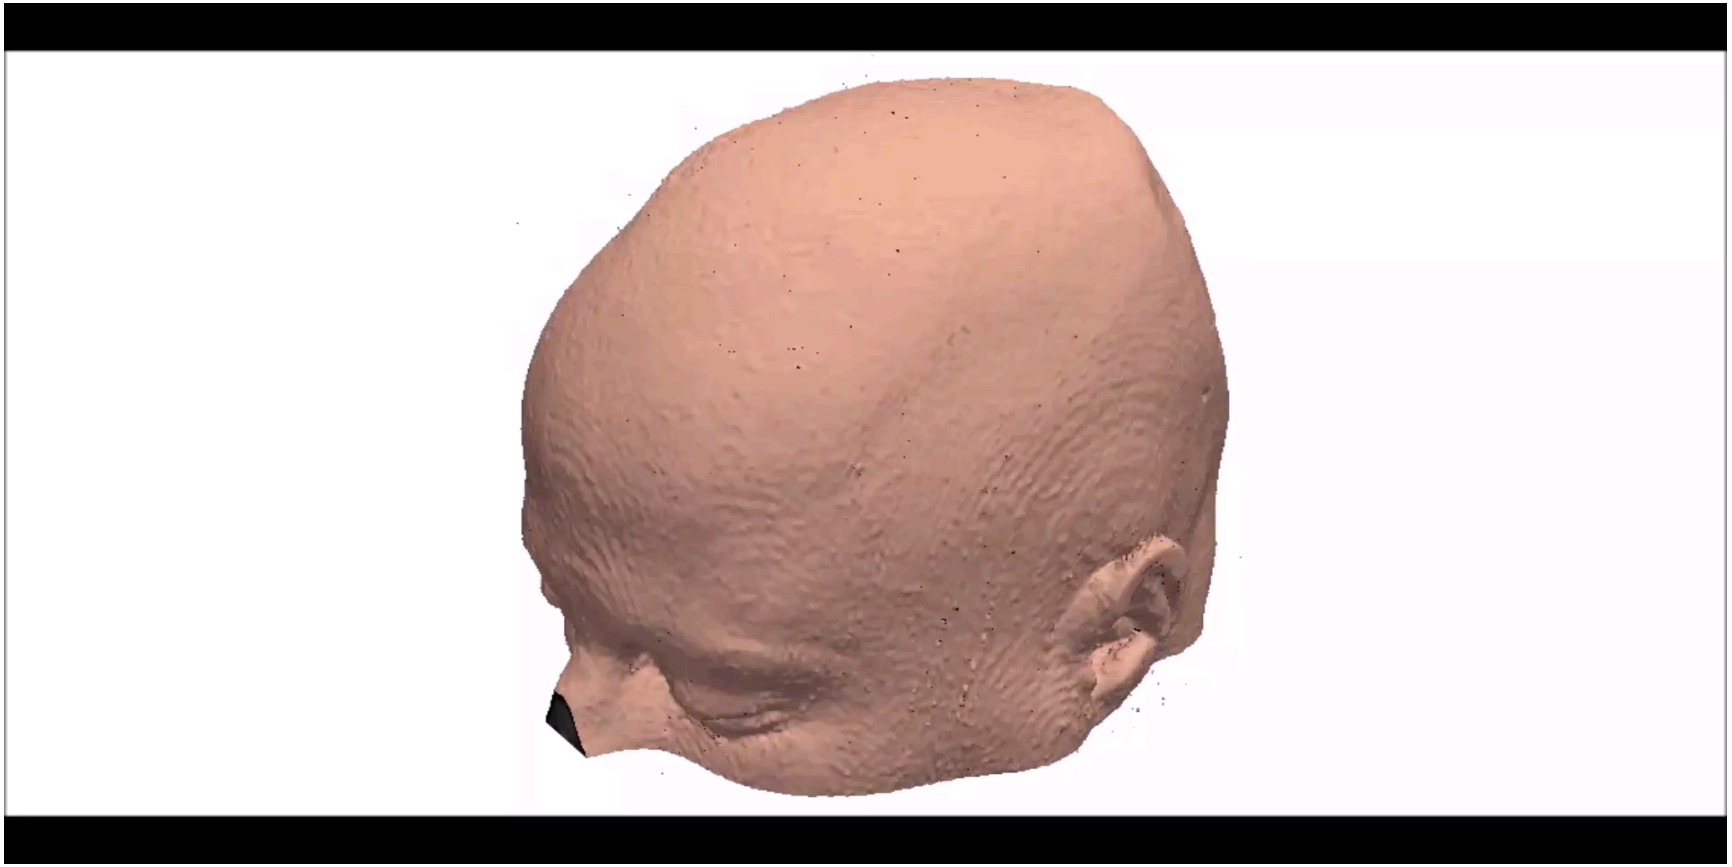

K. Kamnitsas et al. Medical Image Analysis, 2016

# DeepMedic: Results

- 66 patients with moderate-to-severe Traumatic Brain Injury (TBI)
- Imaging at Addenbrooke's Hospital, Cambridge, with 3T Siemens Trio within the first week of injury.
- MRI sequences include
  - MPRAGE
  - FLAIR, T2, Proton Density (PD) and Gradient-Echo (GE)

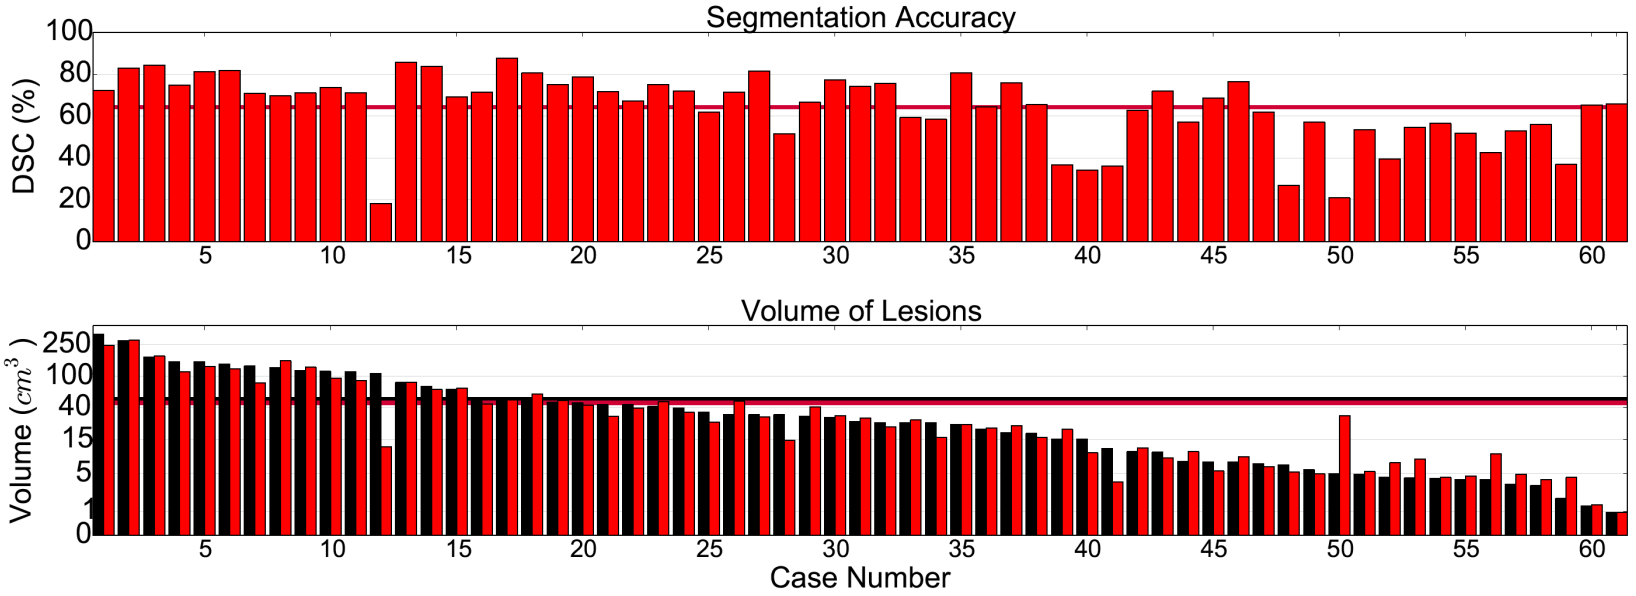


K. Kamnitsas et al. Medical Image Analysis, 2016

# Challenges for image segmentation: Deployment in the clinic

- ML-based segmentation often degrades when deployed in clinical

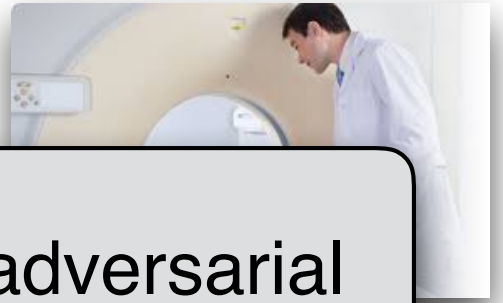

Unsupervised domain adaptation using adversarial neural networks can be used to train a CNN-based segmentation

  - which is more invariant to differences in the input data
  - which does not require any annotations on the test domain

- Manually annotating new data for each test domain is not a feasible solution

# Deploying machine learning into clinical practice: What is the problem?

**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: $D_T = \{\mathcal{X}_T, P(X$

Task: $T_T = \{\mathcal{Y}_T, f'$

Here: $\mathcal{Y}_T = \mathcal{Y}$

Domain Shift:

$P(X$

$P_T(y|x$ $|x)$

$f'_T \neq f'_S$

Domain Shift !

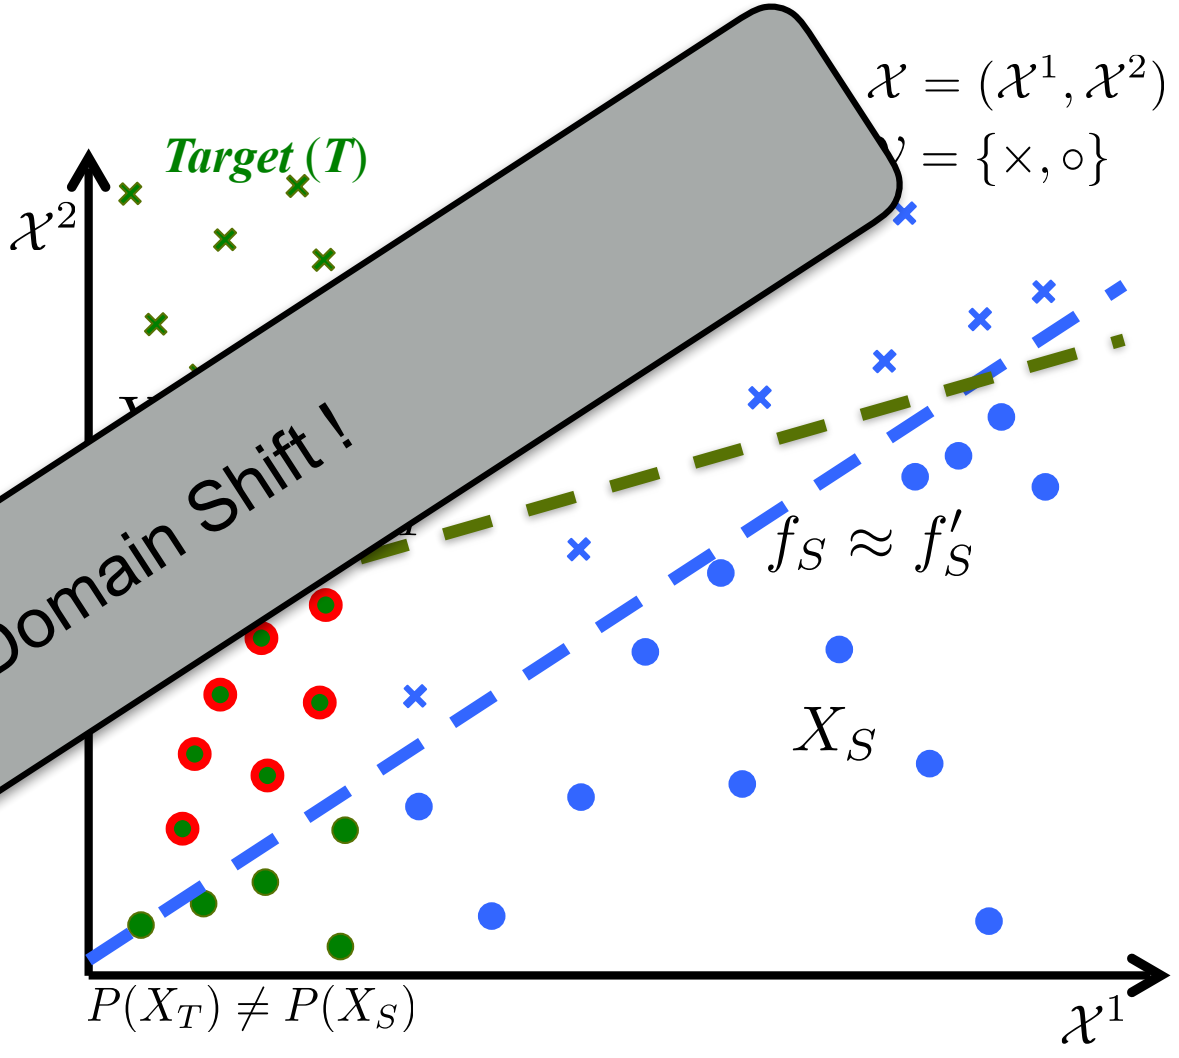

# Solution: Unsupervised domain adaptation with adversarial networks

- Learn a domain classifier $f_D$

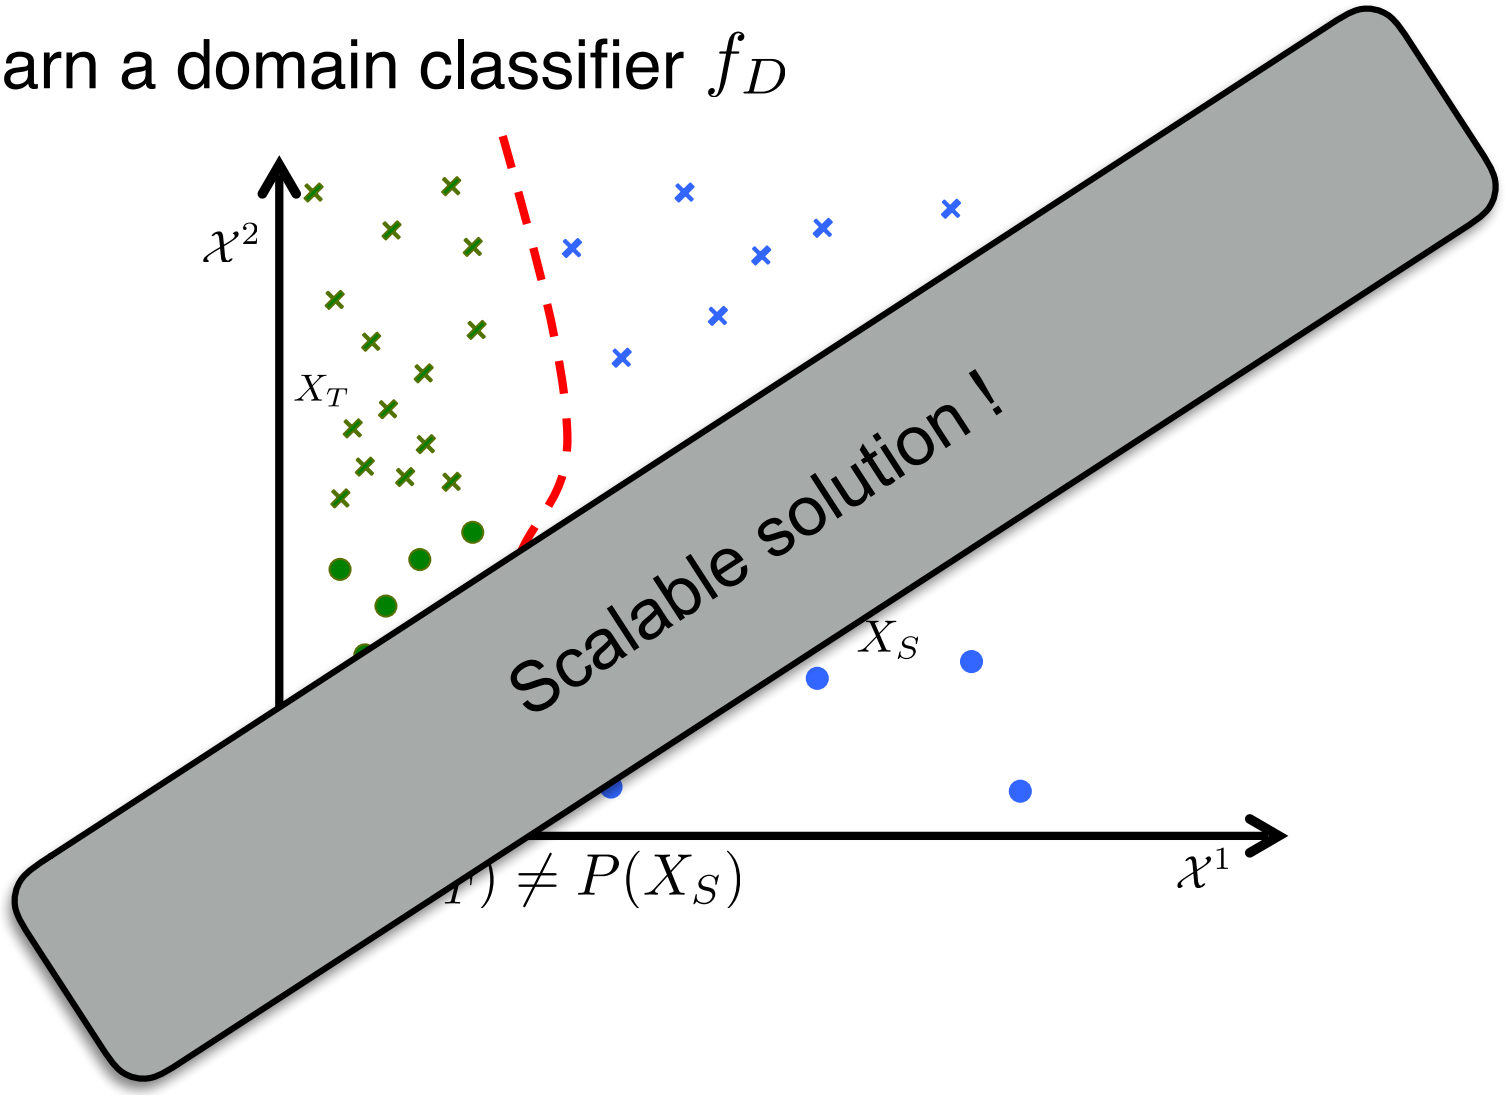

# DeepMedic: Unsupervised domain adaptation with adversarial networks

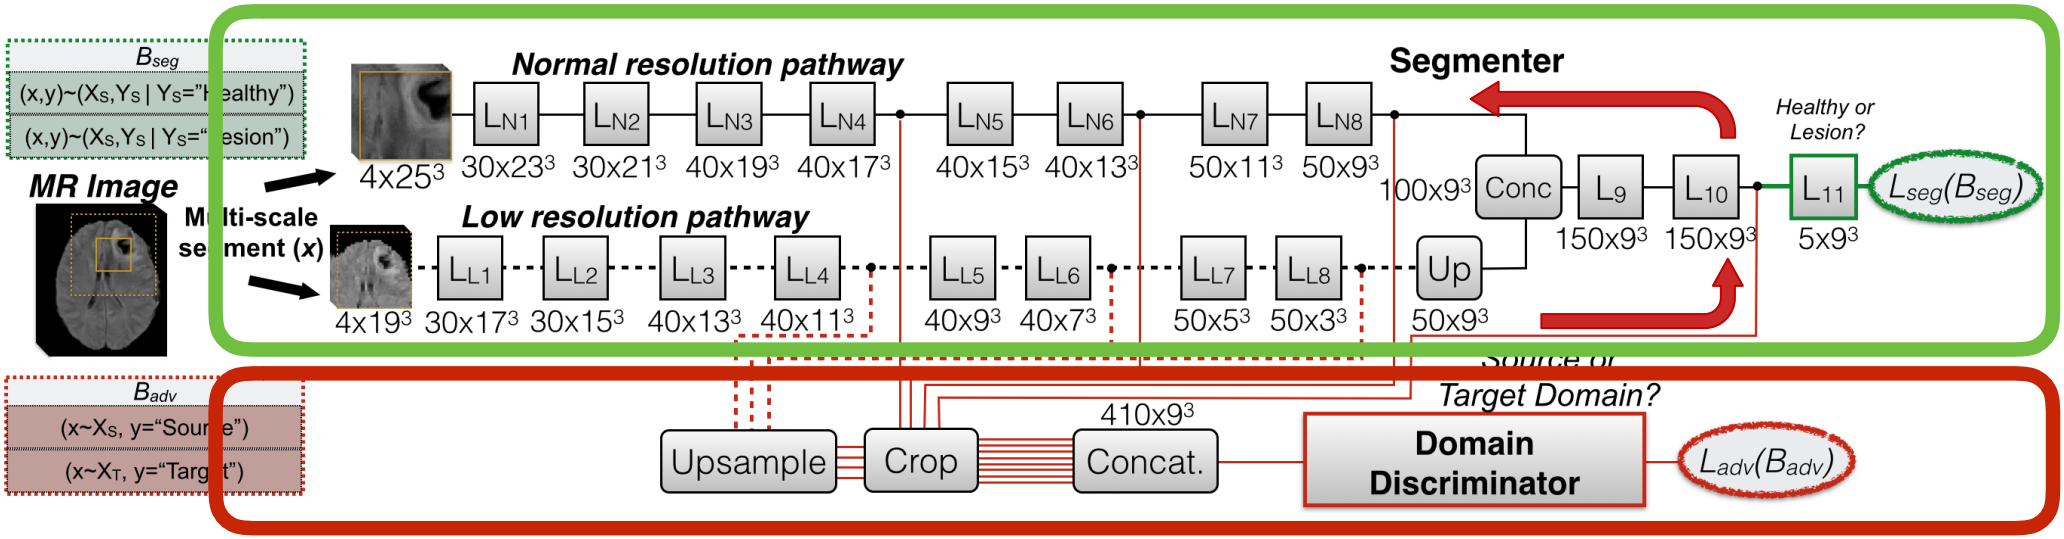


K. Kamnitsas et al. IPMI 2017, arXiv:1612.08894

# DeepMedic: Unsupervised domain adaptation with adversarial networks

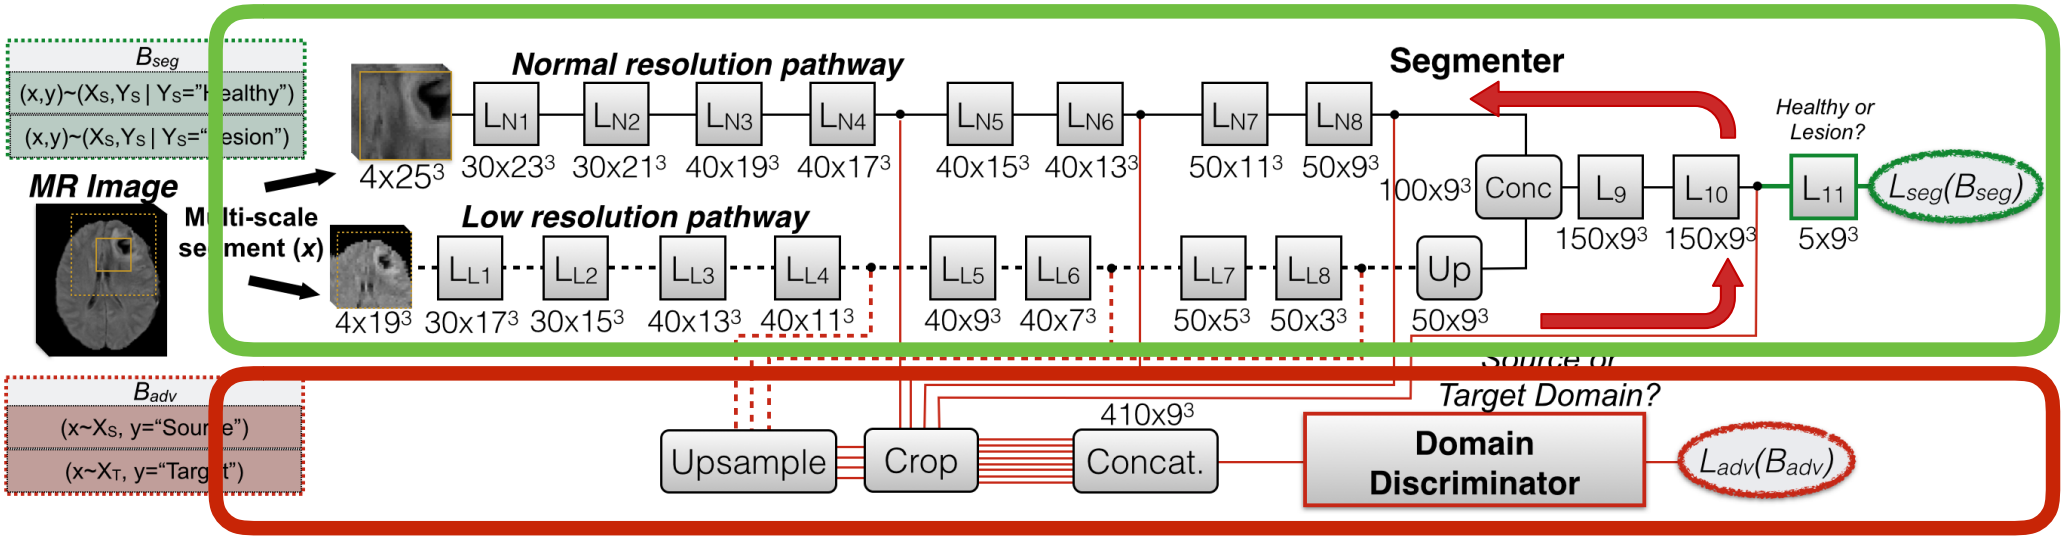


Segmenter: $\mathcal{L}_{seg} = -\frac{1}{m}\sum_{i=1}^{m}[f(x_i) = y_i]log(f(x_i)) \quad , (x_i, y_i) \sim (X_S, Y_S)$

Domain Discr.: $\mathcal{L}_{adv} = -\frac{1}{m}\sum_{i=1}^{m} log(f_D(h(x_i))) - \frac{1}{m}\sum_{j=1}^{m} log(1 - f_D(h(x_j))) \, , x_i \sim X_S \; x_j \sim X_T$

Adapted Segmenter: $\mathcal{L}_{advSeg} = \mathcal{L}_{seg} - \alpha\mathcal{L}_{adv}$

K. Kamnitsas et al. IPMI 2017, arXiv:1612.08894

# DeepMedic: Unsupervised domain adaptation with adversarial networks

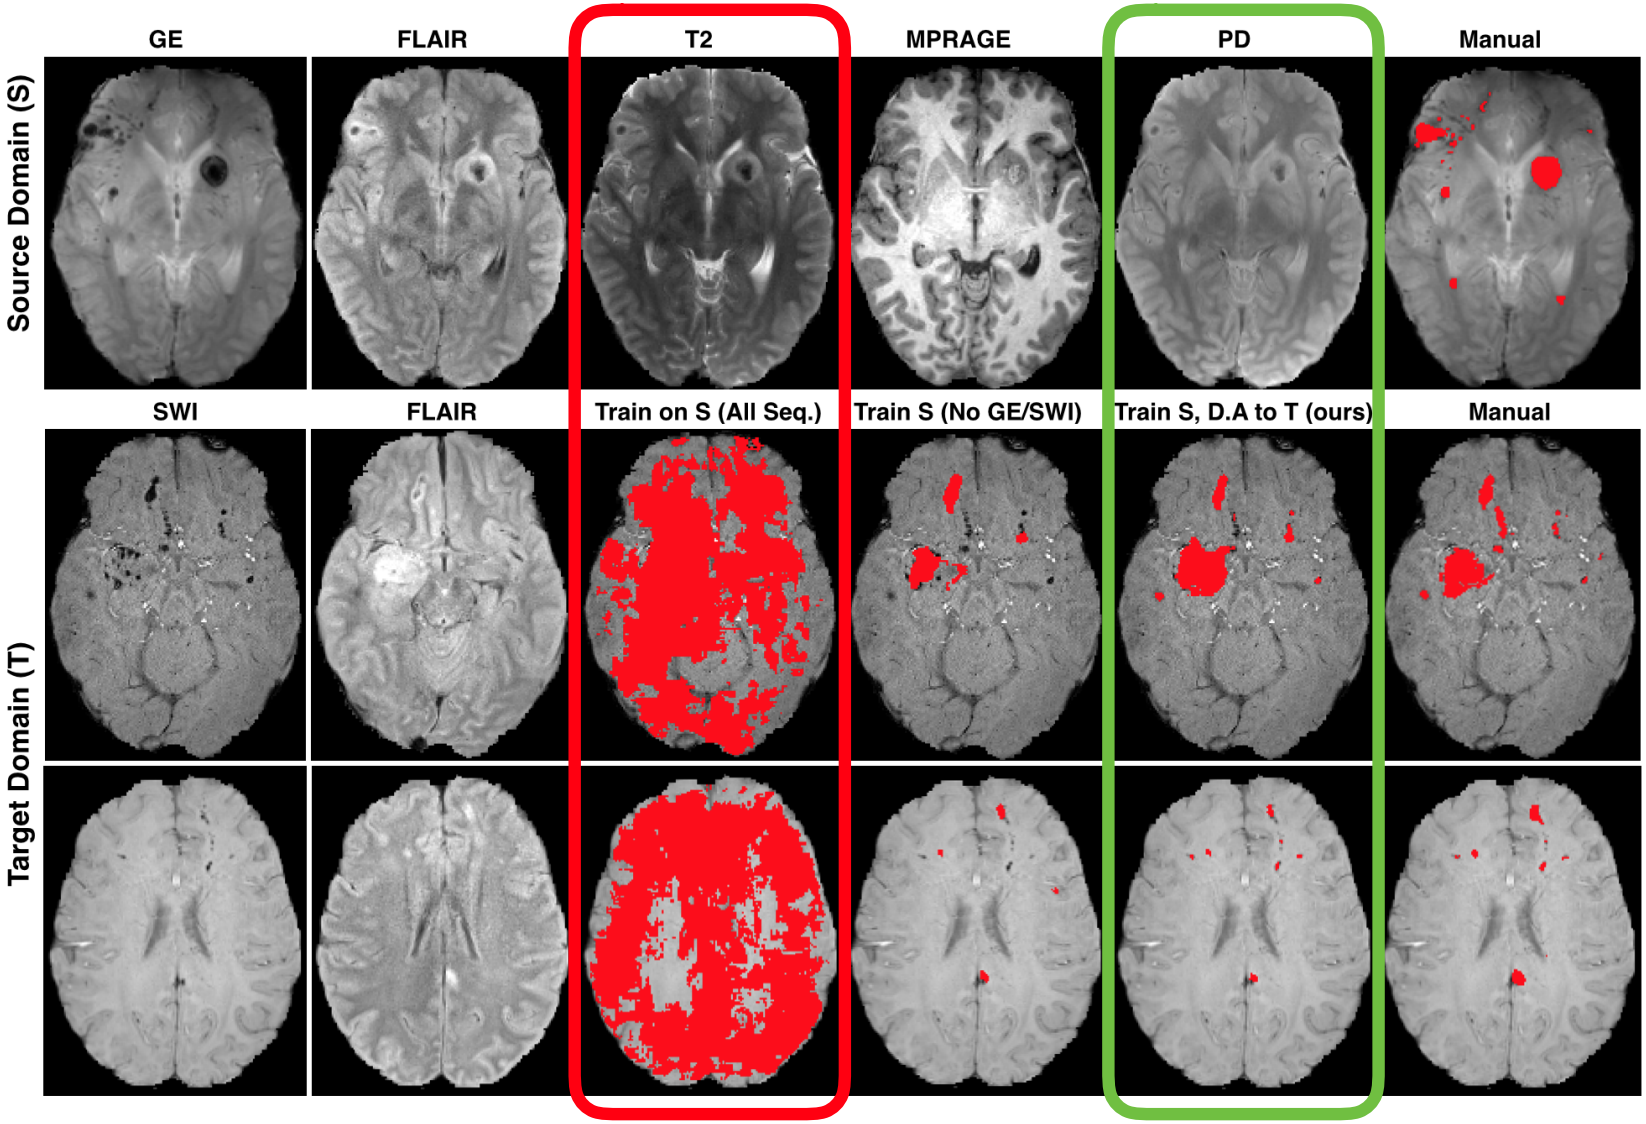


K. Kamnitsas et al. IPMI 2017, arXiv:1612.08894

# DeepMedic, FCN & U-Net

- **The good:**
  - There are some good/promising CNN-based segmentation

Ensemble of Multiple Models & Architectures (EMMA)

Performance **insensitive** to suboptimal configuration

Behaviour **unbiased** by architecture & configuration

- **The ugly:**
  - Chosen model & config may be suboptimal for other data/task
  - Results and conclusions of analysis are strongly biased

# Ensemble of Multiple Models and Architectures (EMMA)

Need to learn: $P(Y|X)$

Approximate it by model: $P(Y|X;\theta_m,m)$

with learnt parameters $\theta_m = \min_{\theta_m} d\left(P\left(Y|X;\theta_m,m\right), P(Y|X)\right)$, *d* the loss.

Model is defined by chosen meta-parameters ***m.***

**Commonly *m* is neglected, but it biases the results!**

We define stochastic variable ***M***, over configurations of interest.

Need to marginalise out influence of ***M***:

$$P(Y|X) = \sum_{\forall m \in M} P(Y, M=m|X) = \sum_{\forall m \in M} P(Y|X, M=m)\, P(M=m)$$

**EMMA approximate the joint by ensembling individual models:**

$$\sim P_{EMMA}(Y|X) = \sum_{\forall m \in M} P(Y|X;\theta_m,m)\, \frac{1}{|M|}$$

# ***M:*** Network architectures

**DeepMedic** [Kamnitsas 2015, 2016, 2017]: https://github.com/Kamnitsas

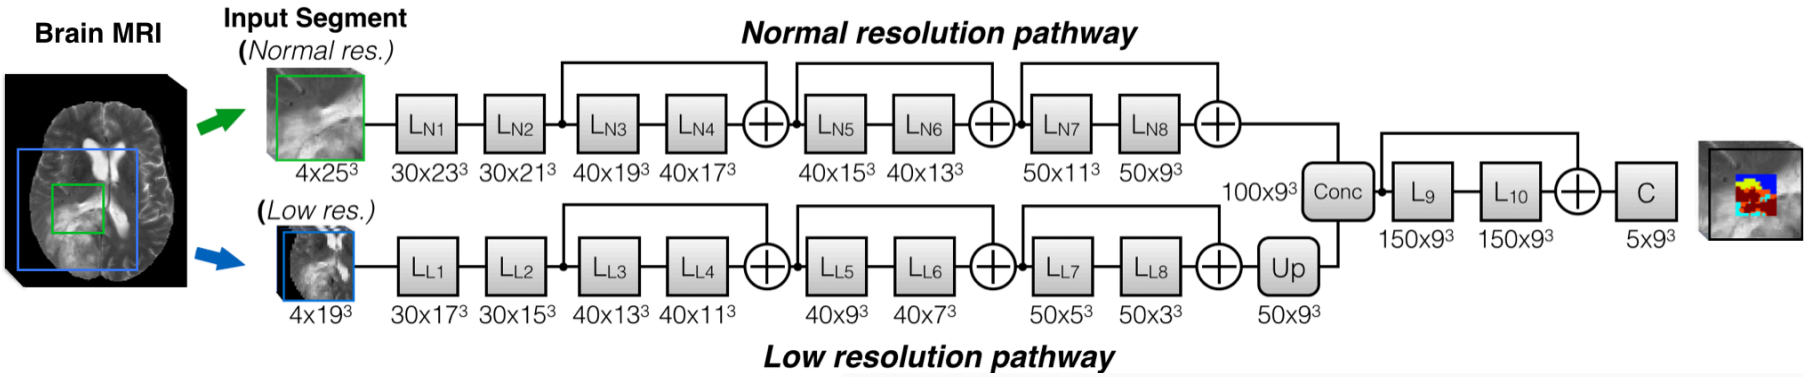


**FCN** [Long 2015]:
https://github.com/DLTK

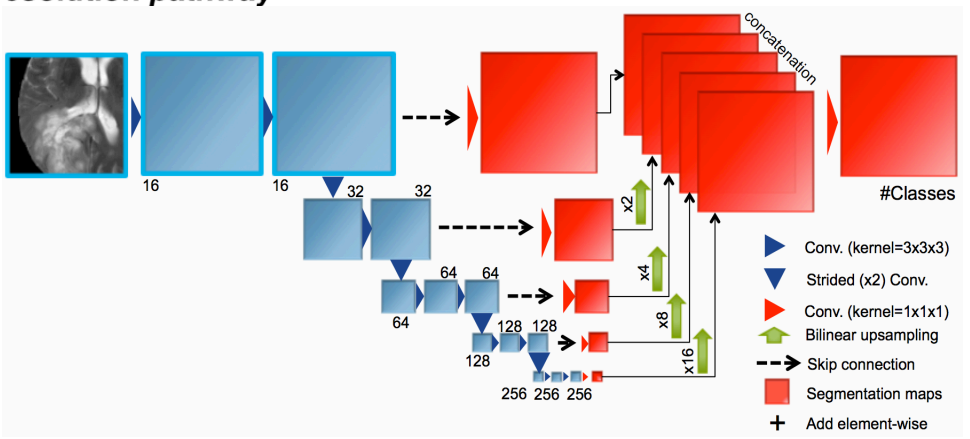


**U-Net** [Ronneberger 2015]:
https://gitlab.com/eferrante

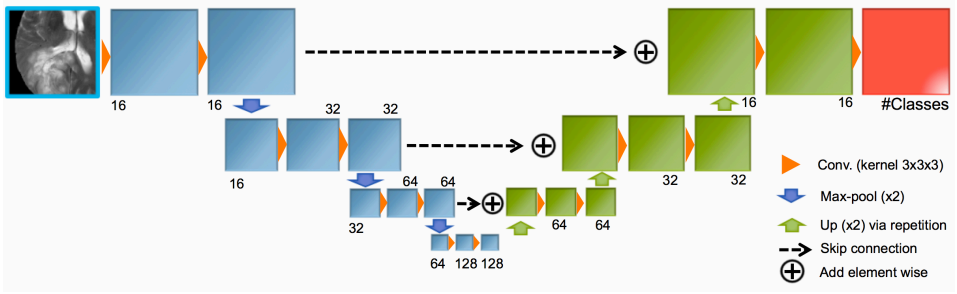

# ***M:*** Network configurations

- Architecture configuration:
  - depth, width, scales, residuals, etc.
- Training Loss:
  - Cross-Entropy, IoU, DSC, etc.
- Sampling strategy:
  - equally per class, foreground/background, etc.
- Optimisation:
  - optimizer, learning rate, momentum, regulariser…
- Data normalisation:
  - z-score, bias field correction, histogram matching


*Led by CBICA*

# 1st Place

**2017 MICCAI BraTS Challenge**
**(Segmentation Task)**
**K. Kamnitsas, et al. "Ensembles of Multiple Models and Architectures for Robust Brain Tumour Segmentation"**

# BRATS'17 Challenge: Quantitative validation

- **EMMA: 2 x DeepMedic, 3 x FCNs, 1 x U-Net**
  - Different training losses, sampling strategies, widths, depths, configurations
  - No config was heavily optimised for the task (3/6 nets were quite suboptimal)

| | DSC | | | Sensitivity | | | Specificity | | | Hausdorff_95 | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | Enh. | Whole | Core | Enh. | Whole | Core | Enh | Whole | Core | Enh. | Whole | Core |
| **EMMA** | **75.7** | **90.2** | 82.0 | **79.0** | 90.9 | 78.3 | 99.8 | 99.5 | **99.9** | **4.22** | 4.56 | **6.11** |
| UCL-TIG | 75.2 | 89.7 | **82.5** | 77.1 | **91.2** | **83.9** | 99.8 | 99.4 | 99.7 | 4.78 | **3.97** | 7.60 |
| MIC_DKFZ | 73.2 | 89.6 | 79.7 | **79.0** | 89.6 | 78.1 | 99.8 | **99.6** | **99.9** | 4.55 | 6.97 | 9.48 |

- **Robustness:**
  - EMMA of all 6 was better than individuals.
  - Ensemble of 3 best nets was only marginally better than EMMA of all 6 nets.

# Conventional CNNs: Do not use prior knowledge explicitly

- Analysis of Neural Networks
  - Model parameterization
  - Model capacity / receptive field
  - Loss function / objective
- Standard Loss Functions
  - X-Entropy loss function

$$L_x = -\sum_{i \in \mathcal{S}} \sum_{c=1}^{C} \log \left( \frac{e^{f_{(c,i)}}}{\sum_j e^{f_{(j,i)}}} \right)$$

  - L2 or Smooth L1 loss function

$$\sum_{i \in \mathcal{S}} \| \Phi(\boldsymbol{x}_i, \theta_r) - \boldsymbol{y}_i \|^2$$

# Conventional CNNs: Problem

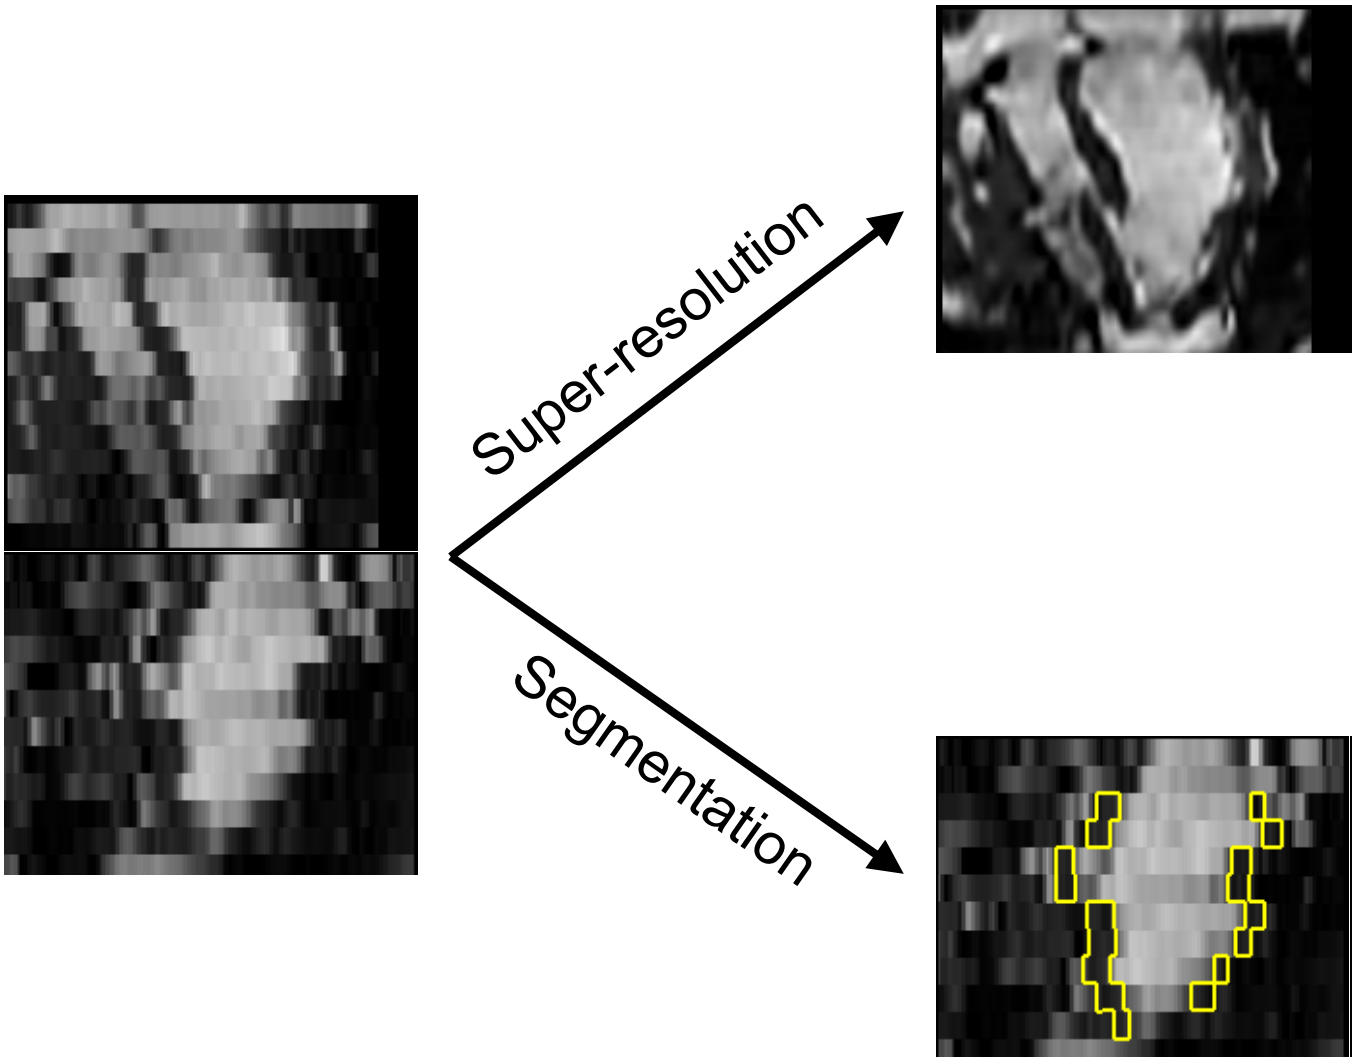

# Conventional CNNs: What we want

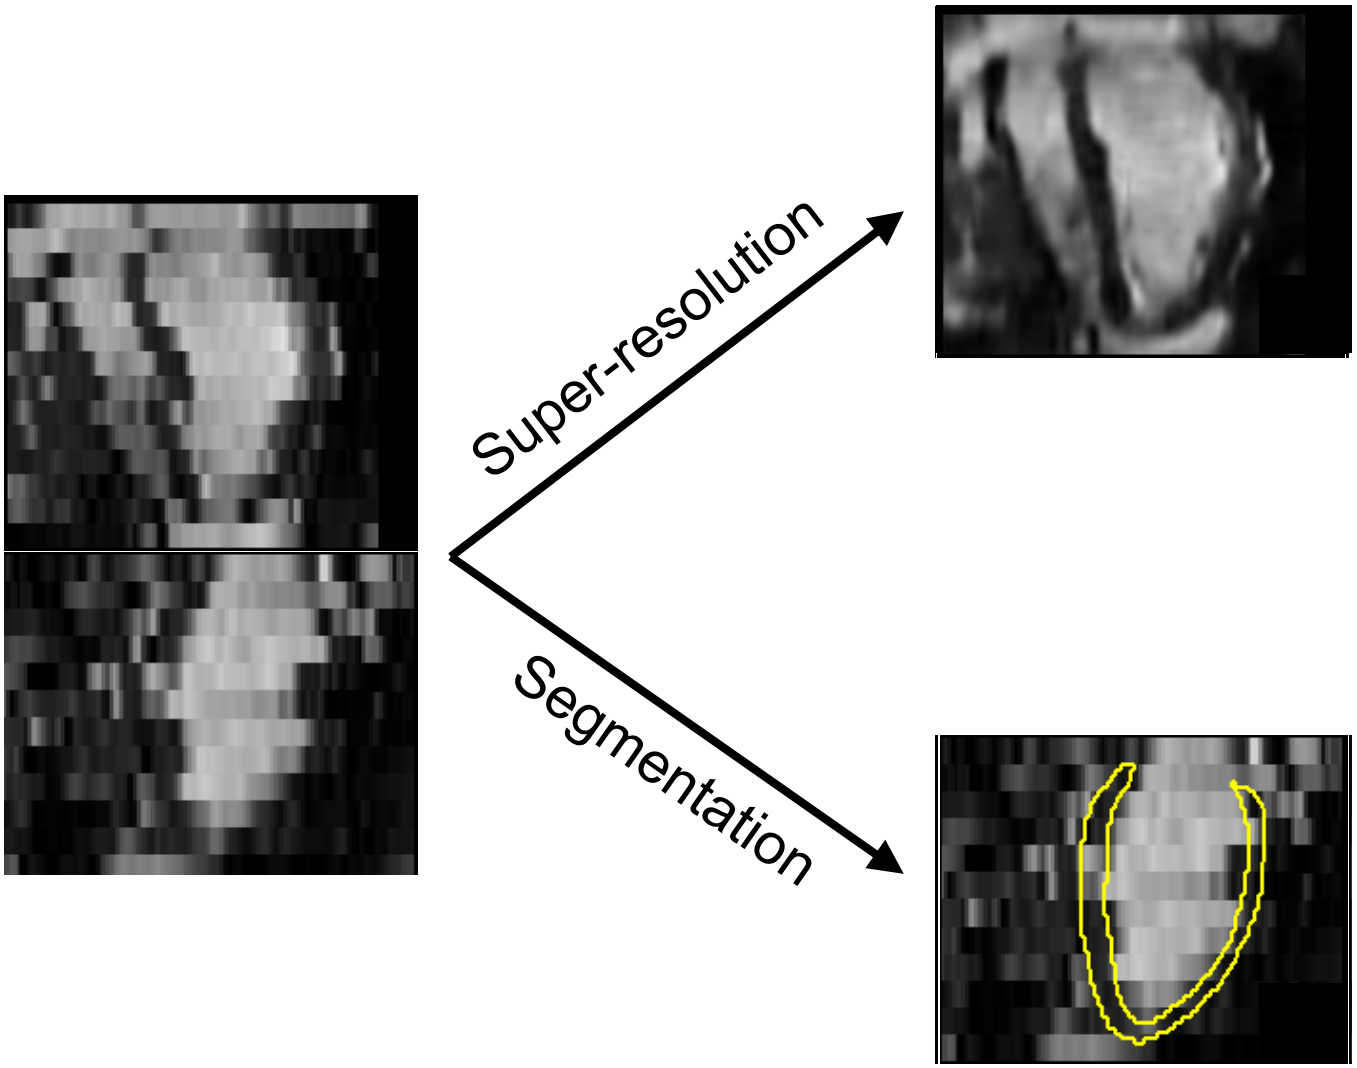

# Learn anatomical priors using a stacked convolutional autoencoder

- Provides a non-linear compact representation of the underlying anatomy

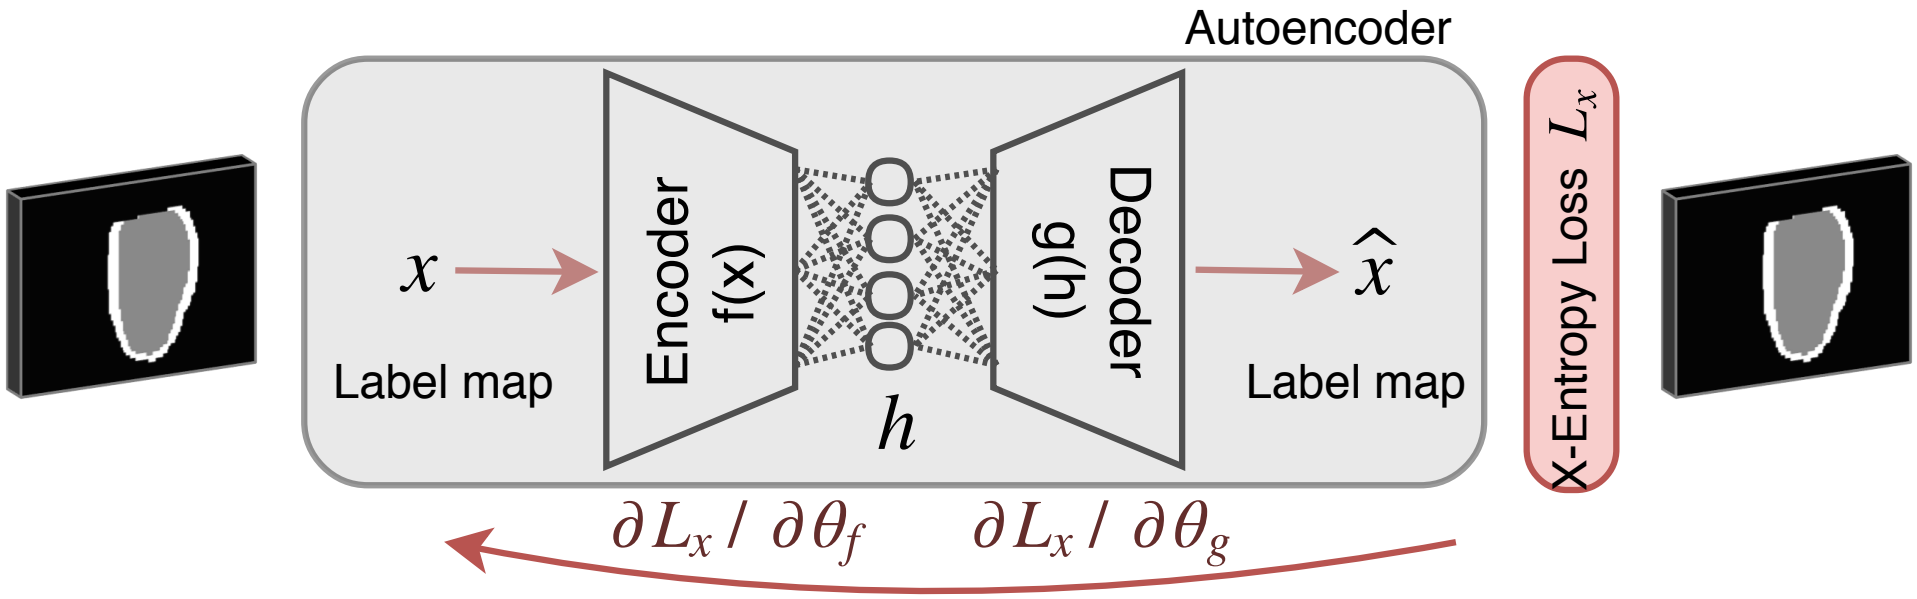


O. Oktay, in press, IEEE TMI 2017

# Anatomically constrained CNN: T-L networks for representing priors

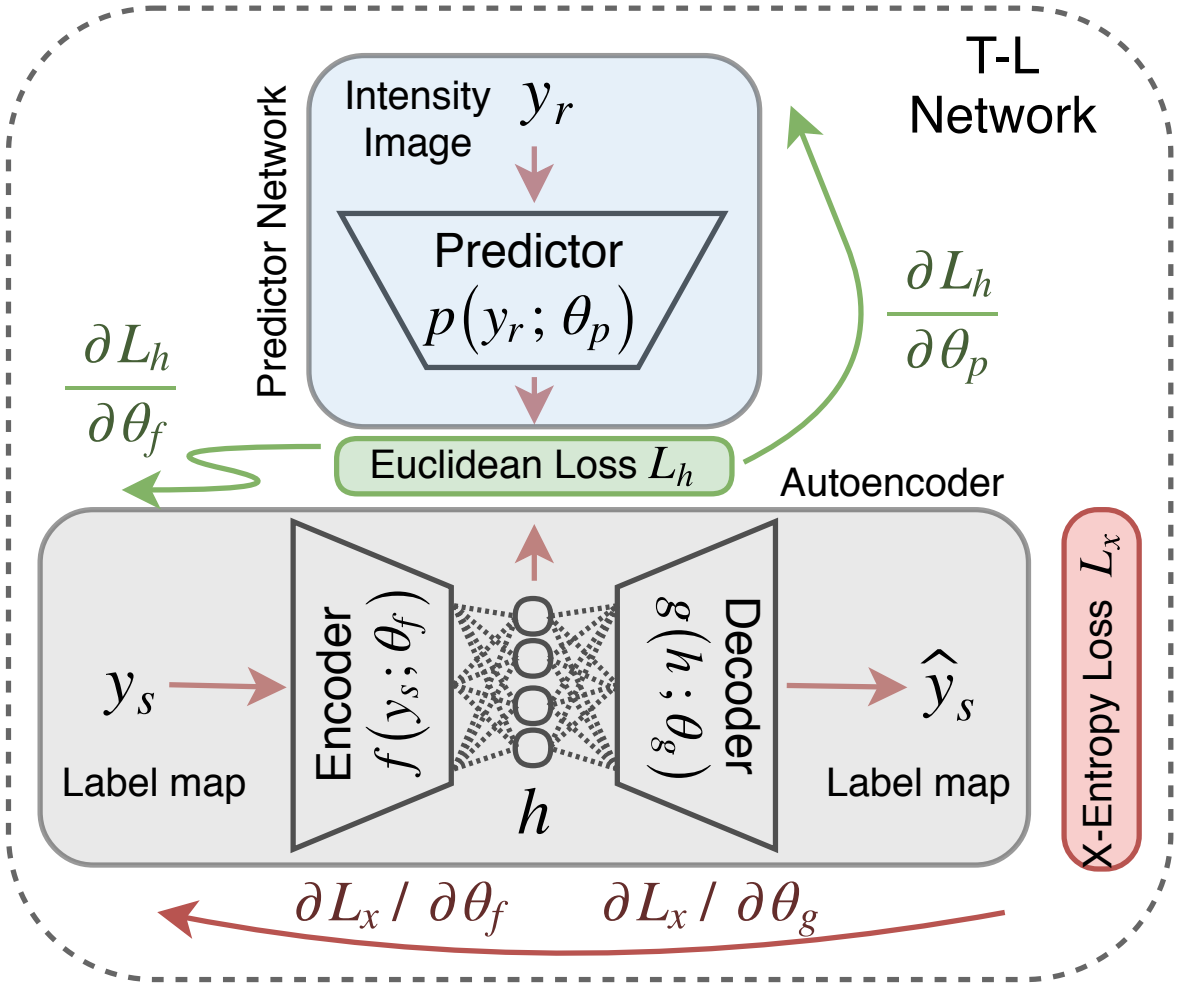


O. Oktay, in press, IEEE TMI 2017

# Anatomically constrained CNN: Segmentation framework

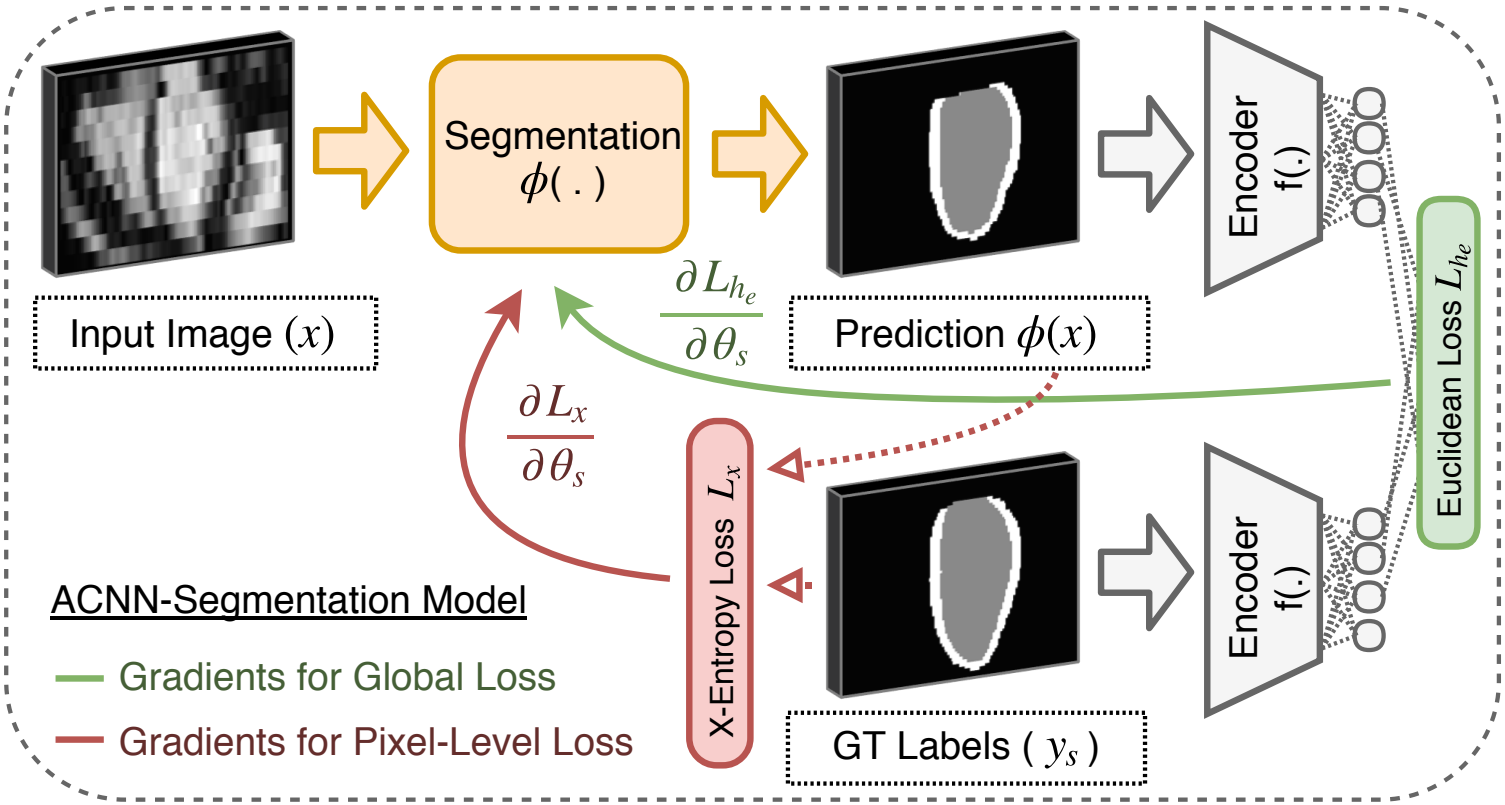


O. Oktay, in press, IEEE TMI 2017

# Anatomically constrained CNN: Segmentation results

| | Method | Mean Dist. (mm) | Hausdorff Dist. (mm) | Dice Score (%) | # Trainable Parameters |
|---|---|---|---|---|---|
| *Endo-cardium* | 2D-FCN [264] | 2.07±0.61 | 11.37±7.15 | .908±.021 | $1.39 \times 10^6$ |
| | 3D-Seg | 1.77±0.84 | 10.28±8.25 | .923±.019 | $1.60 \times 10^6$ |
| | 3D-UNet [50] | 1.66±0.74 | 9.94±9.22 | .923±.019 | $1.64 \times 10^6$ |
| | AE-Seg [217] | 1.75±0.58 | 8.42±3.64 | .926±.019 | $1.68 \times 10^6$ |
| | 3D-Seg-MAug | 1.59±0.74 | 8.52±8.13 | .928±.019 | $1.60 \times 10^6$ |
| | AE-Seg-M | 1.59±0.48 | **7.52±3.78** | .927±.017 | $1.91 \times 10^6$ |
| | **ACNN-Seg** | **1.37±0.42** | 7.89±3.83 | **.939±.017** | $1.60 \times 10^6$ |
| | p-values | $p \ll 0.001$ | $p \approx 0.890$ | $p \ll 0.001$ | - |
| *Myo-cardium* | 2D-FCN [264] | 1.58±0.44 | 9.19±7.22 | .727±.046 | $1.39 \times 10^6$ |
| | 3D-Seg | 1.48±0.51 | 10.15±10.58 | .773±.038 | $1.60 \times 10^6$ |
| | 3D-UNet [50] | 1.45±0.47 | 9.81±11.77 | .764±.045 | $1.64 \times 10^6$ |
| | AE-Seg [217] | 1.51±0.29 | 8.52±2.72 | .779±.033 | $1.68 \times 10^6$ |
| | 3D-Seg-MAug | 1.37±0.41 | 9.41±9.17 | .785±.041 | $1.60 \times 10^6$ |
| | AE-Seg-M | 1.32±0.26 | **7.12±2.79** | .791±.036 | $1.91 \times 10^6$ |
| | **ACNN-Seg** | **1.14±0.22** | 7.31±3.59 | **.811±.027** | $1.60 \times 10^6$ |
| | p-values | $p \ll 0.001$ | $p \approx 0.071$ | $p \ll 0.001$ | - |

# Anatomically constrained CNN: Segmentation results

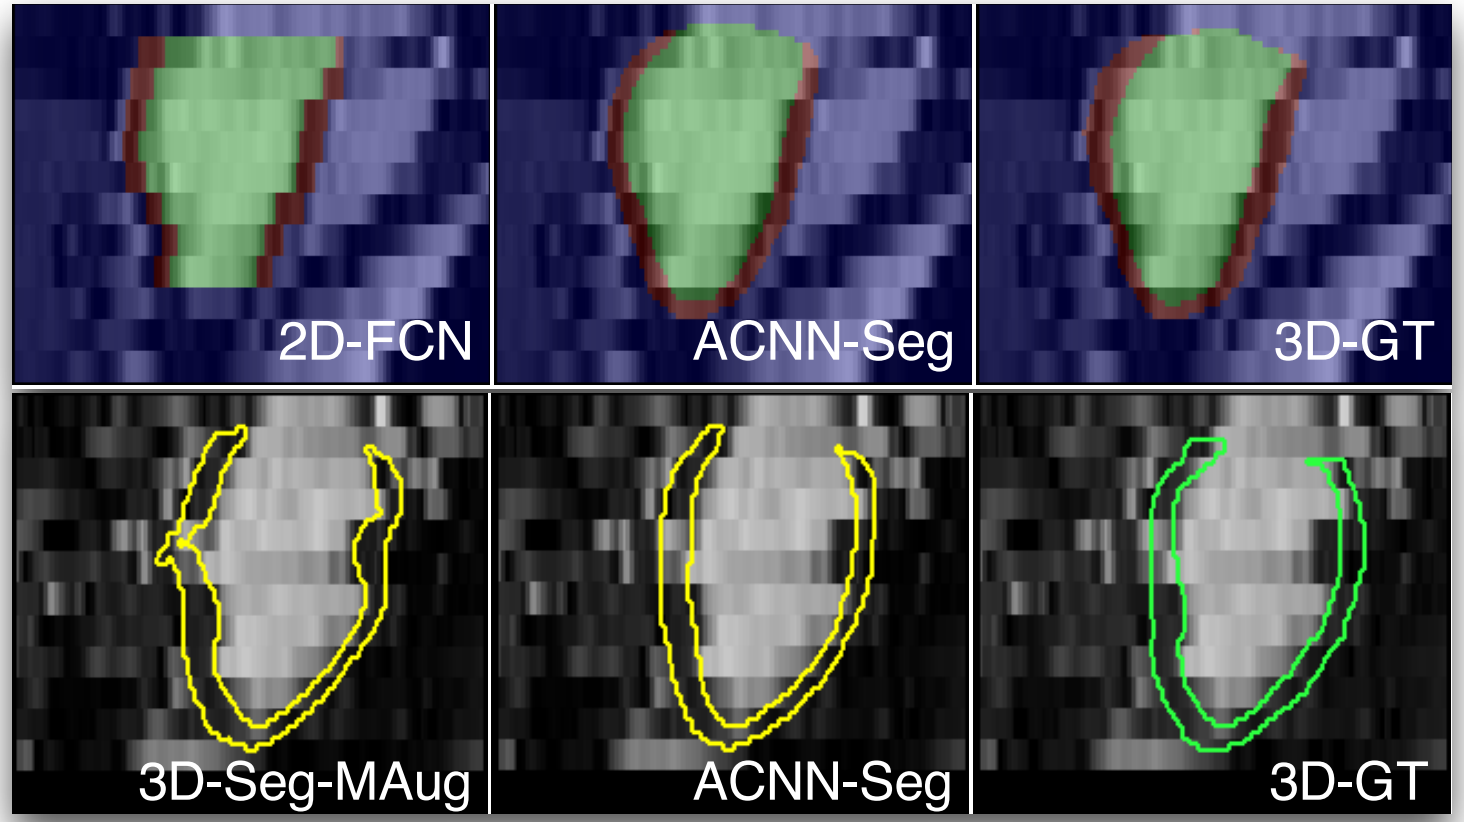


O. Oktay, in press, IEEE TMI 2017

# Anatomically constrained CNN: Super-resolution framework

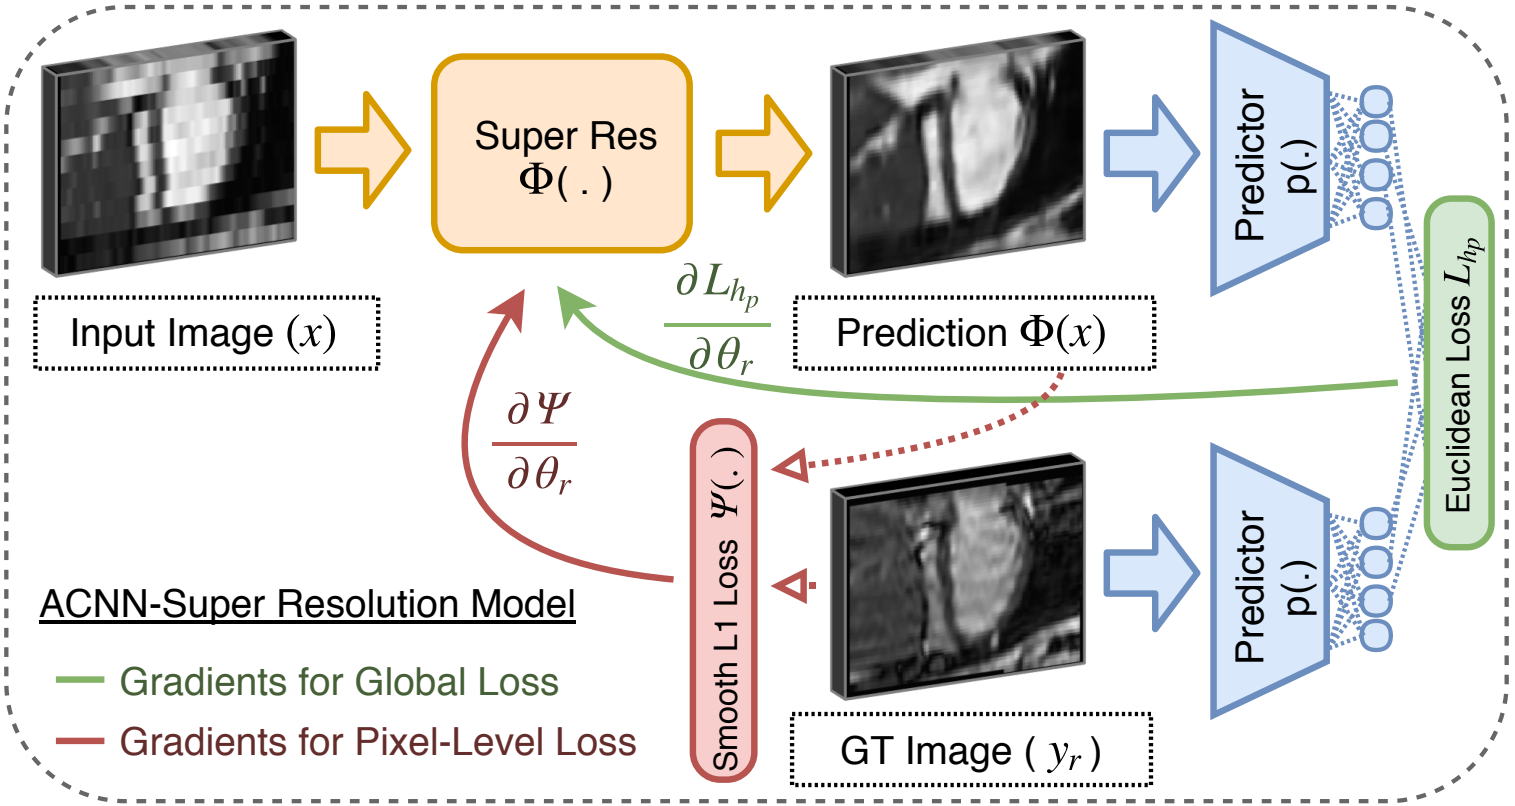


O. Oktay, in press, IEEE TMI 2017

# Anatomically constrained CNN: Super-resolution results

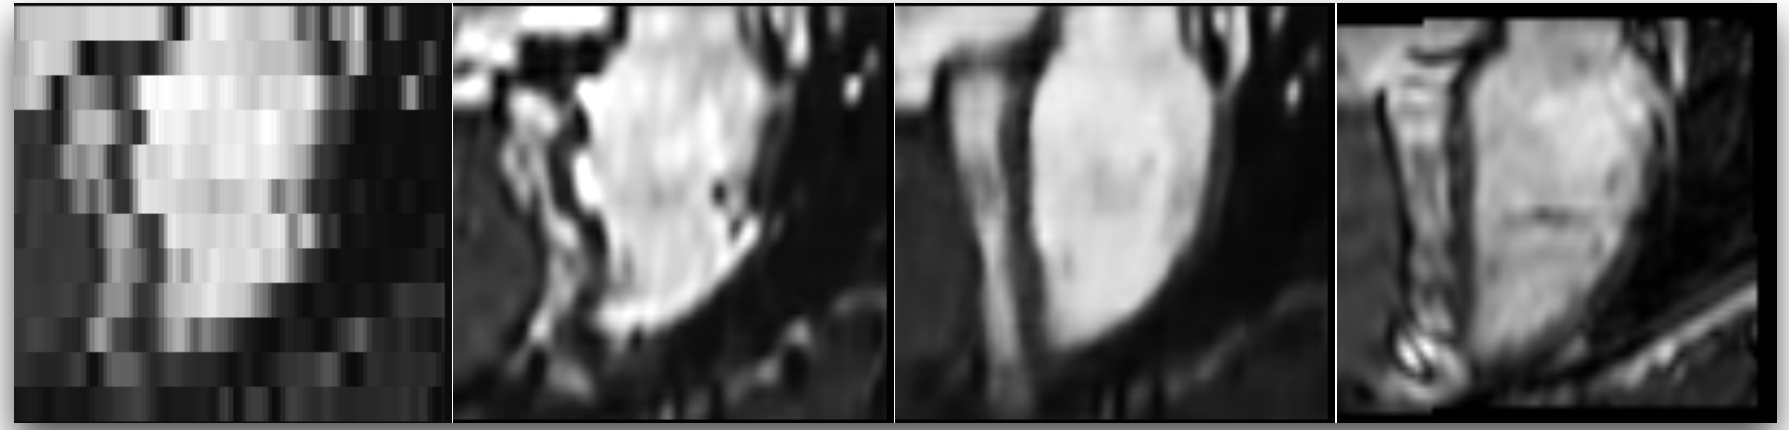

Original LR image | Baseline SR approach | Anatomically constrained SR model | Ground-truth HR image

O. Oktay, in press, IEEE TMI 2017

# Anatomically constrained CNN: Learnt Hidden Representations

- Histogram of the learnt low-dimensional latent representations (randomly selected 16 components are shown).

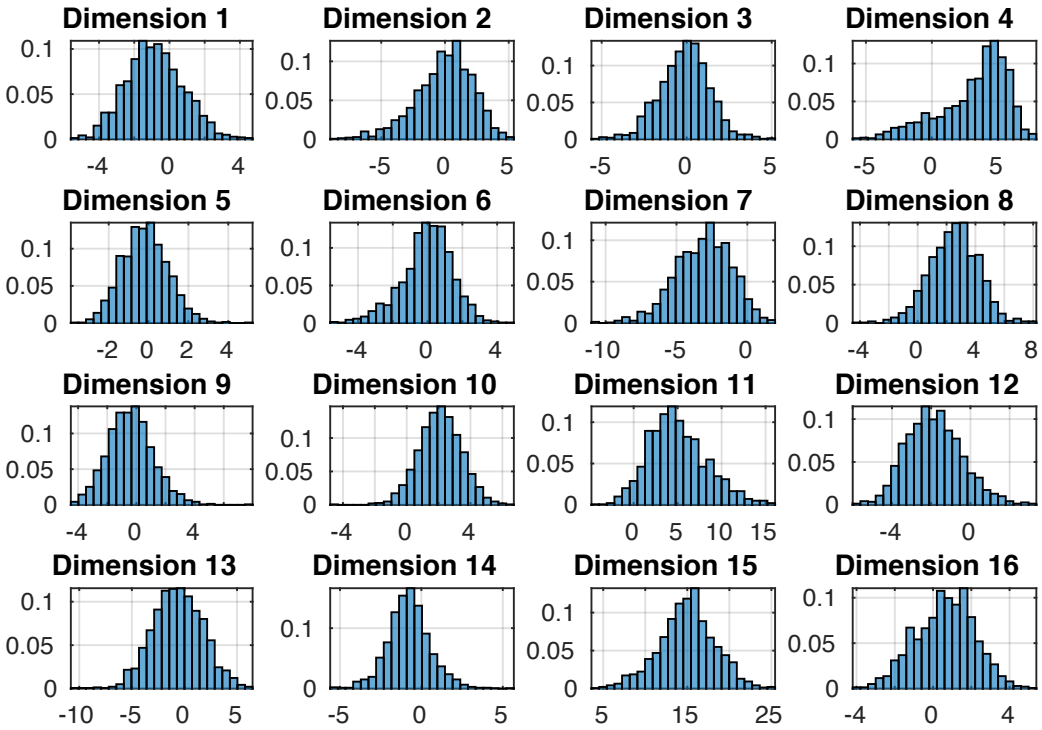


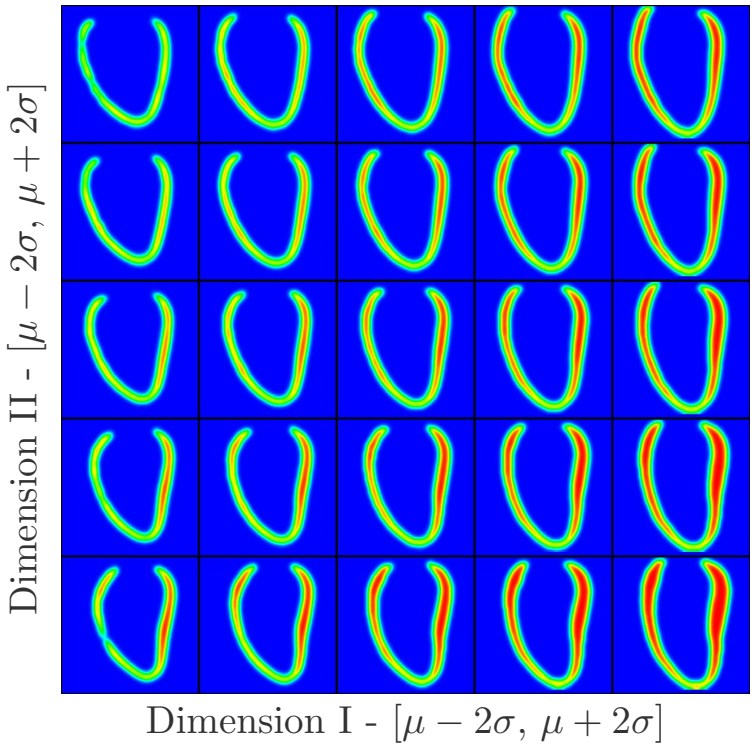


O. Oktay, in press, IEEE TMI 2017

# Summary and Conclusions

- Machine learning plays a crucial role in
  - Image reconstruction
  - Computer aided detection and decision support



Validation is challenging

- Truly intel
  diagnosis requires
  - Learning from unlabelled, large-scale population data
  - Integration of imagin
    information, e.g. clinical records and genetics

Requires collaboration between computer scientists, engineers and clinicians

Optimisation of imaging pipeline with respect to clinically useful information

# Current state-of-the-art

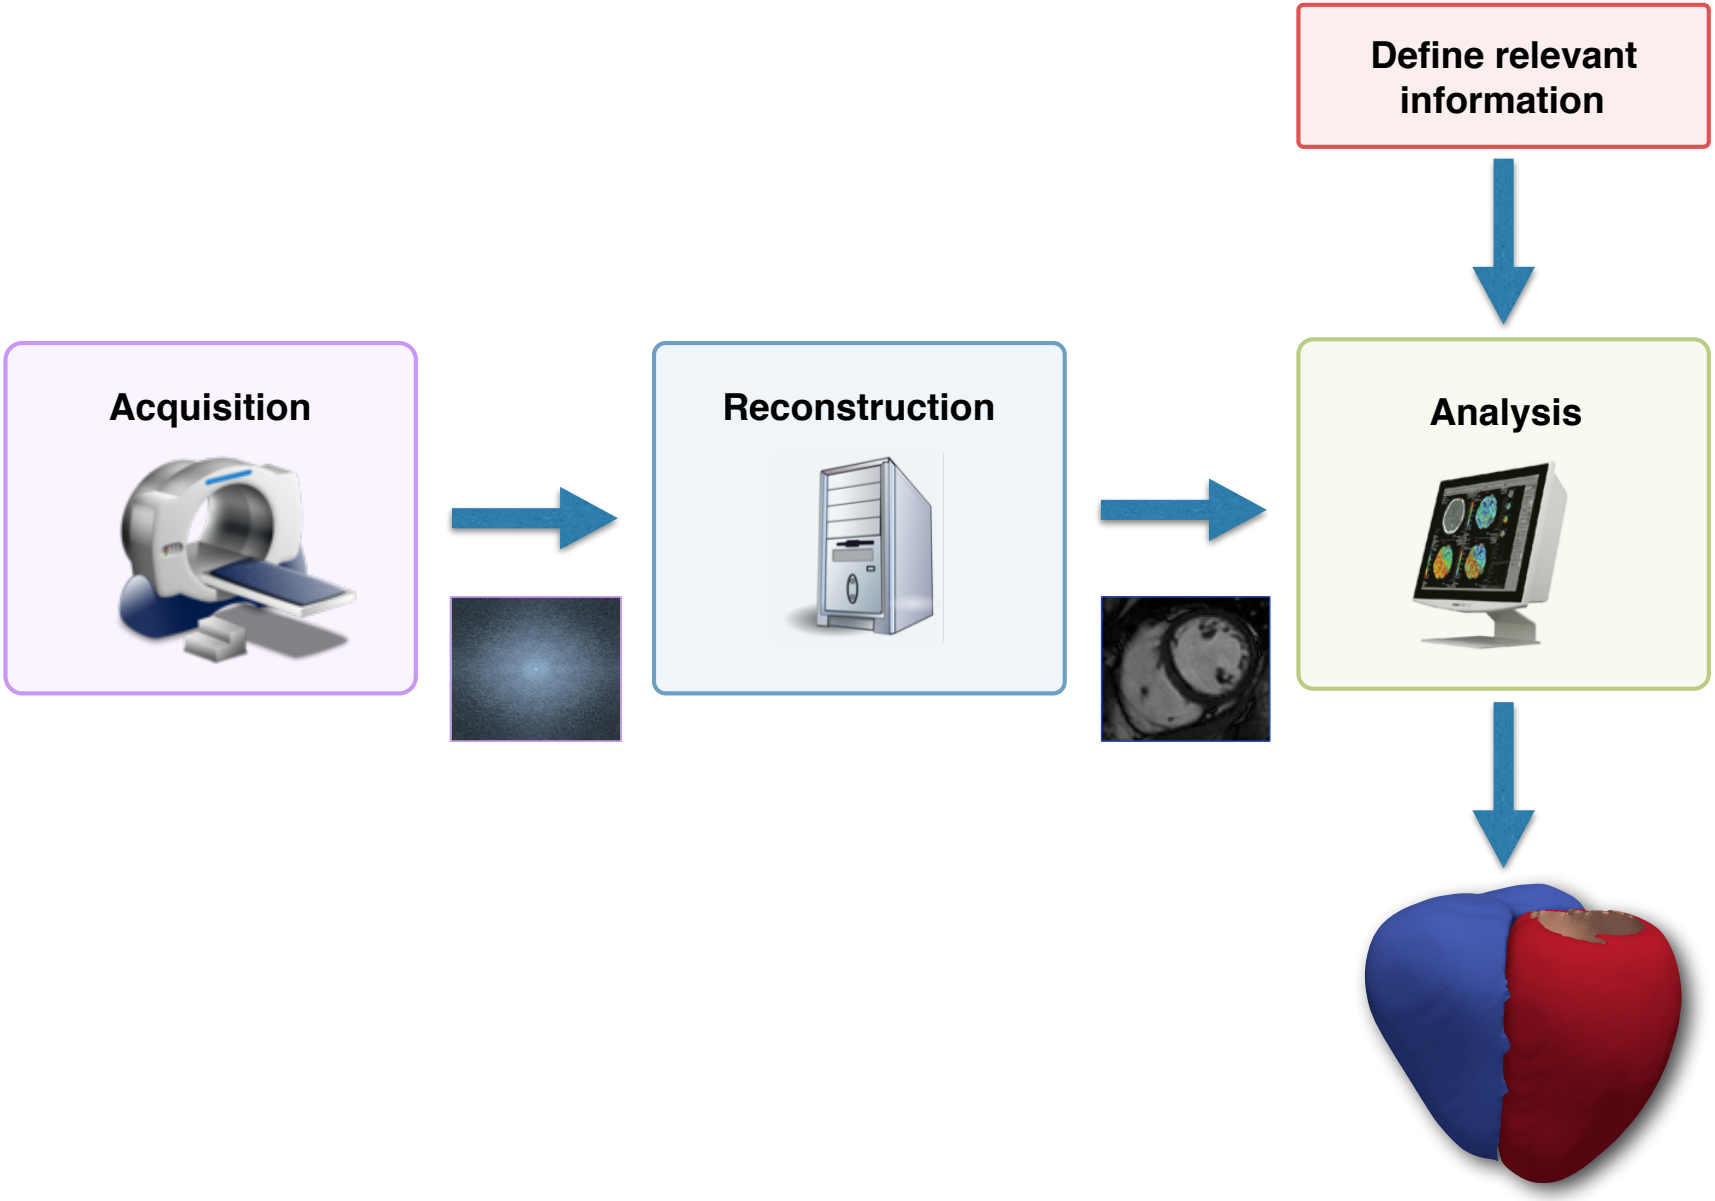

# Future: End-to-end optimisation of entire imaging pipeline via deep learning

Define relevant information

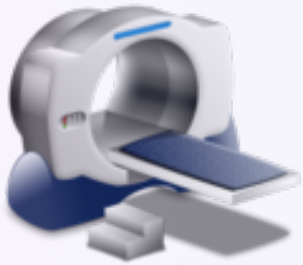

**End-to-end optimisation of acquisition, reconstruction, analysis & interpretation via deep learning**

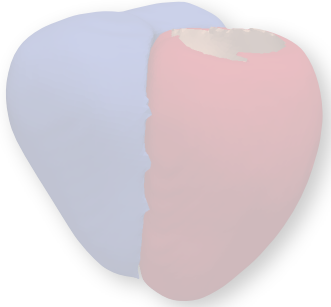

# Future: End-to-end optimisation of entire imaging pipeline via deep learning

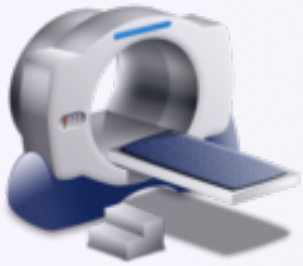


**End-to-end optimisation of acquisition, reconstruction, analysis & interpretation via deep learning**

biobank uk

Dementias Platform UK
Medical Research Council

**Big data (population data)**

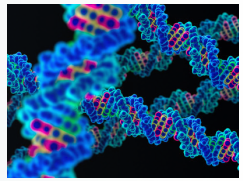

**Multi-modal data**

# Acknowledgements


Wenjia Bai
Matthew Sinclair
Giacomo Tarroni
Ozan Oktay
Martin Rajchl
Ghislain Vaillant
Aaron M. Lee
Nay Aung
Elena Lukaschuk
Mihir M. Sanghvi
Filip Zemrak
Steven McDonagh
Kenneth Fung
Jose Miguel Paiva
Valentina Carapella
Young Jin Kim
Hideaki Suzuki
Bernhard Kainz
Paul M. Matthews
Steffen E. Petersen
Stefan K. Piechnik
Stefan Neubauer
Ben Glocker
Enzo Ferrante
Jo Hajnal
Jose Caballero
Christian Ledig
Christian Baumgartner
Kostas Kamnitsas
Tong Tong
Wenzhe Shi
Martin Rajchl
Jo Schlemper
Carlo Biffi
Nick Pawlowski
Matthew Lee

This research has been conducted mainly using the UKBB Resource under Application Number 2946. The initial stage of the research was conducted using the UKBBResource under Application Number 18545.

# Acknowledgements




European Research Council











**Interested in join us?** We have several opportunities in deep learning research: https://biomedic.doc.ic.ac.uk