Imperial College London

Deep learning for medical image reconstruction, segmentation and analysis

Daniel Rueckert

Biomedical Image Analysis Group

Department of Computing, Imperial College London, UK

Disclosures

- Co-founder – IXICO
- Adviser – HeartFlow, Circle Cardiovascular Imaging
- Grant funding from:

EPSRC Engineering and Physical Sciences Research Council

British Heart Foundation

Innovate UK Technology Strategy Board

HeartFlow®

NVIDIA

wellcome trust

erc European Research Council

Deep learning for medical imaging: Opportunities

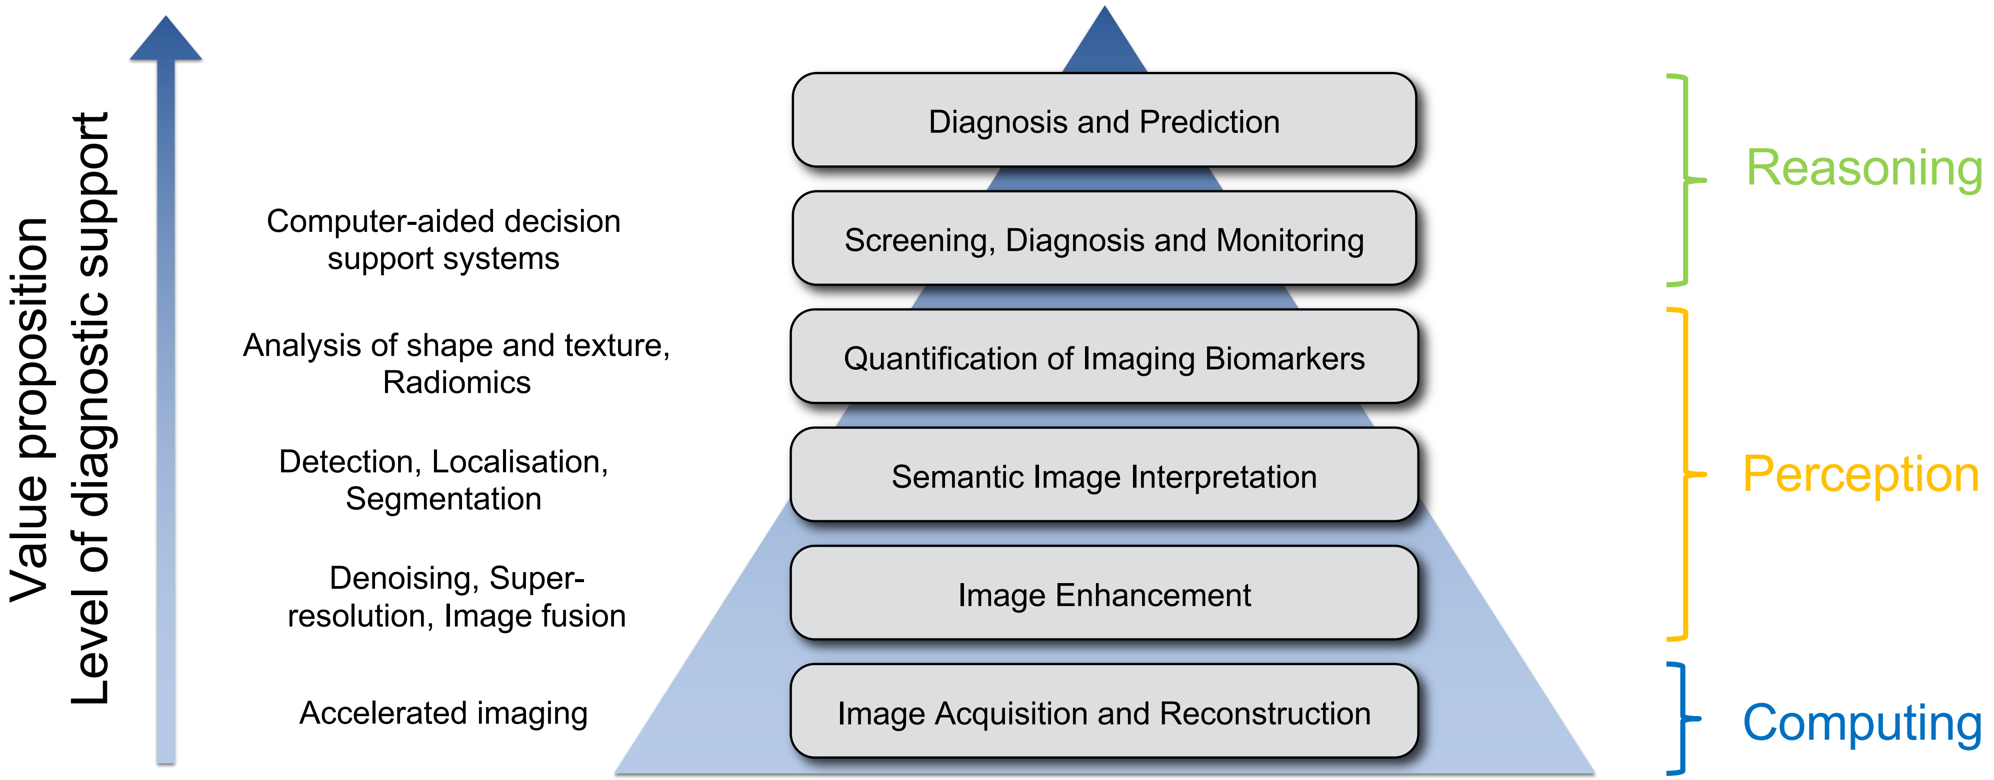

MR image acquisition: Challenges

- 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

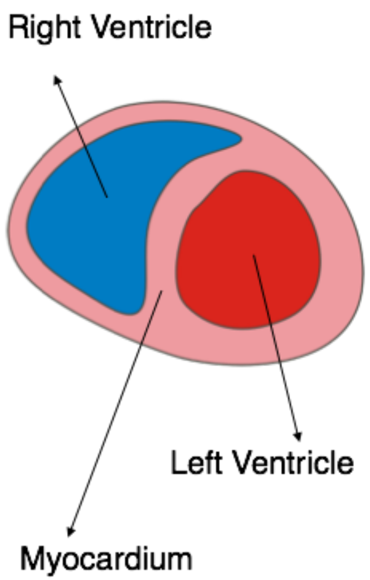

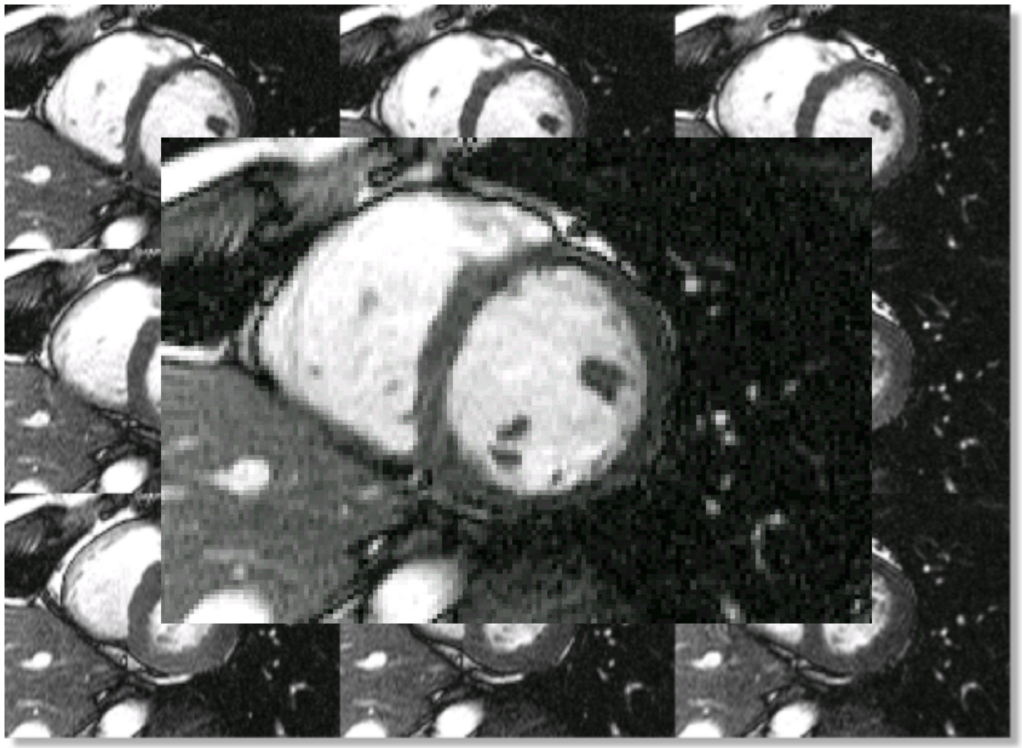

Cardiac MRI: Full acquisition is slow

- MRI acquisition is performed in k-space by sequentially t[illegible]rsing sampling trajectories.

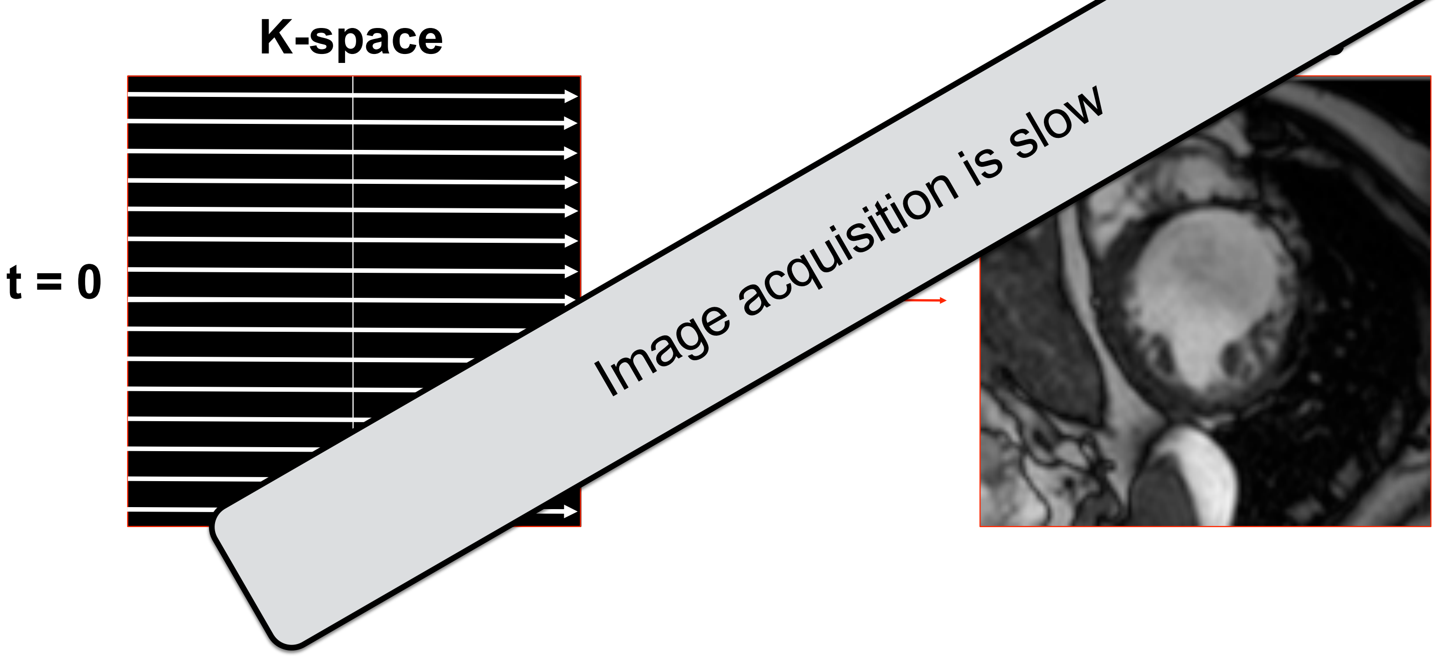

Cardiac MRI: Full acquisition is slow

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

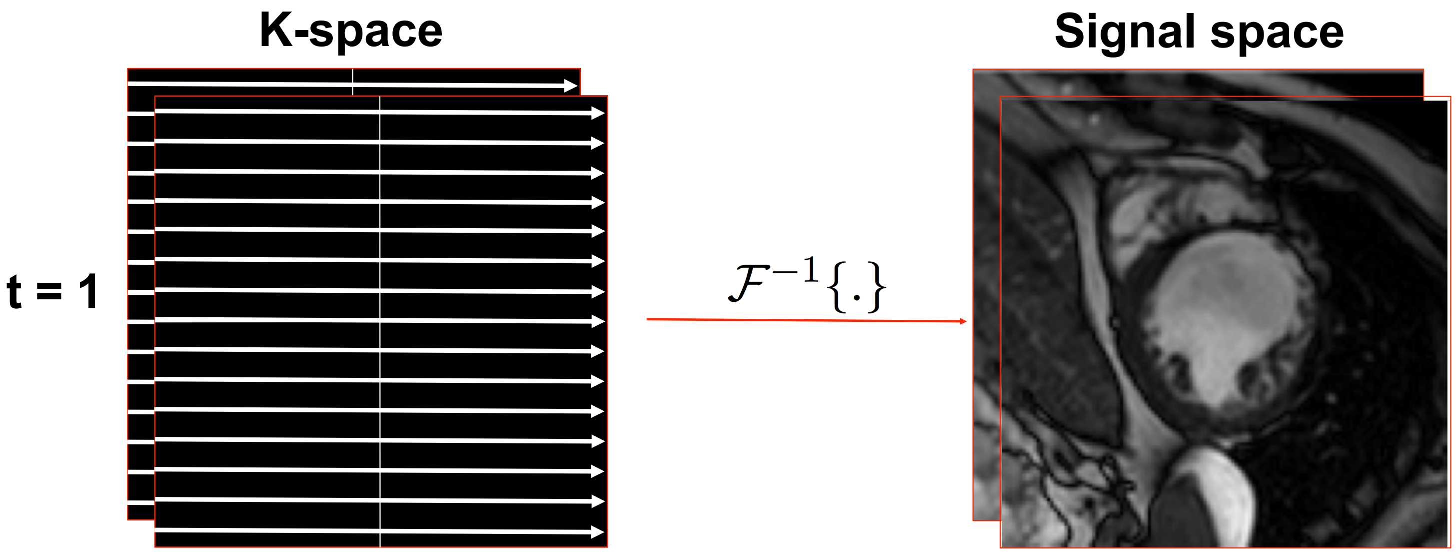

Cardiac MRI: Full acquisition is slow

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

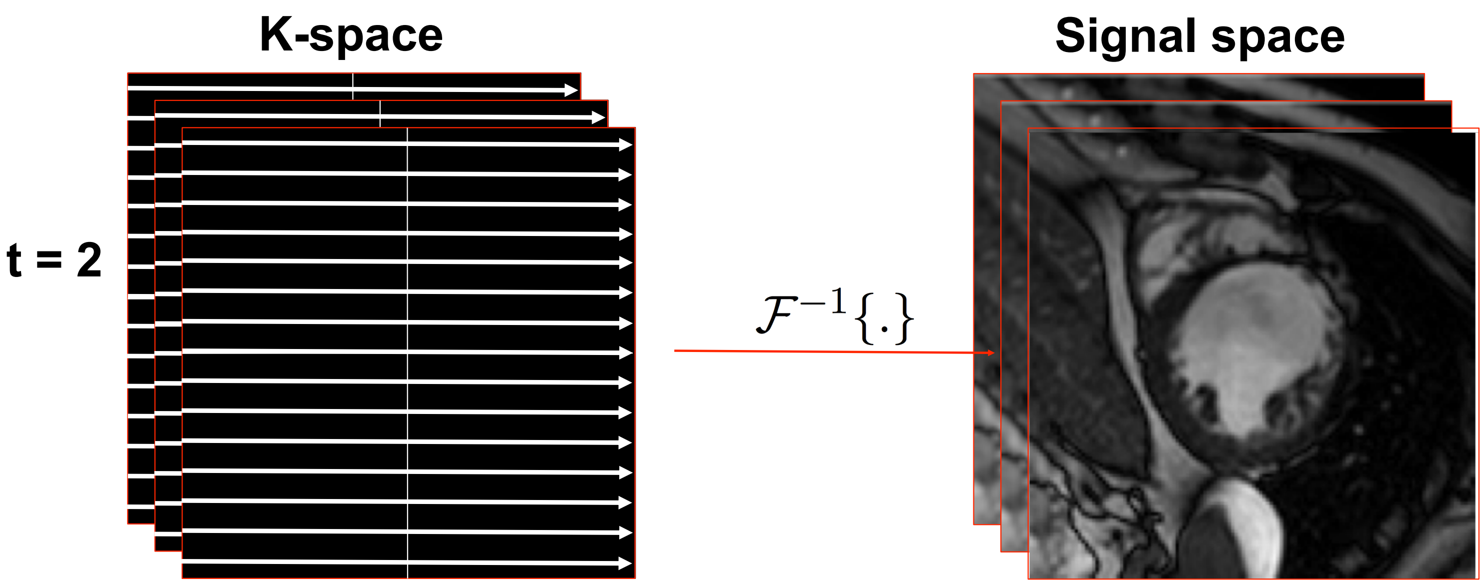

Cardiac MRI: Full acquisition is slow

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

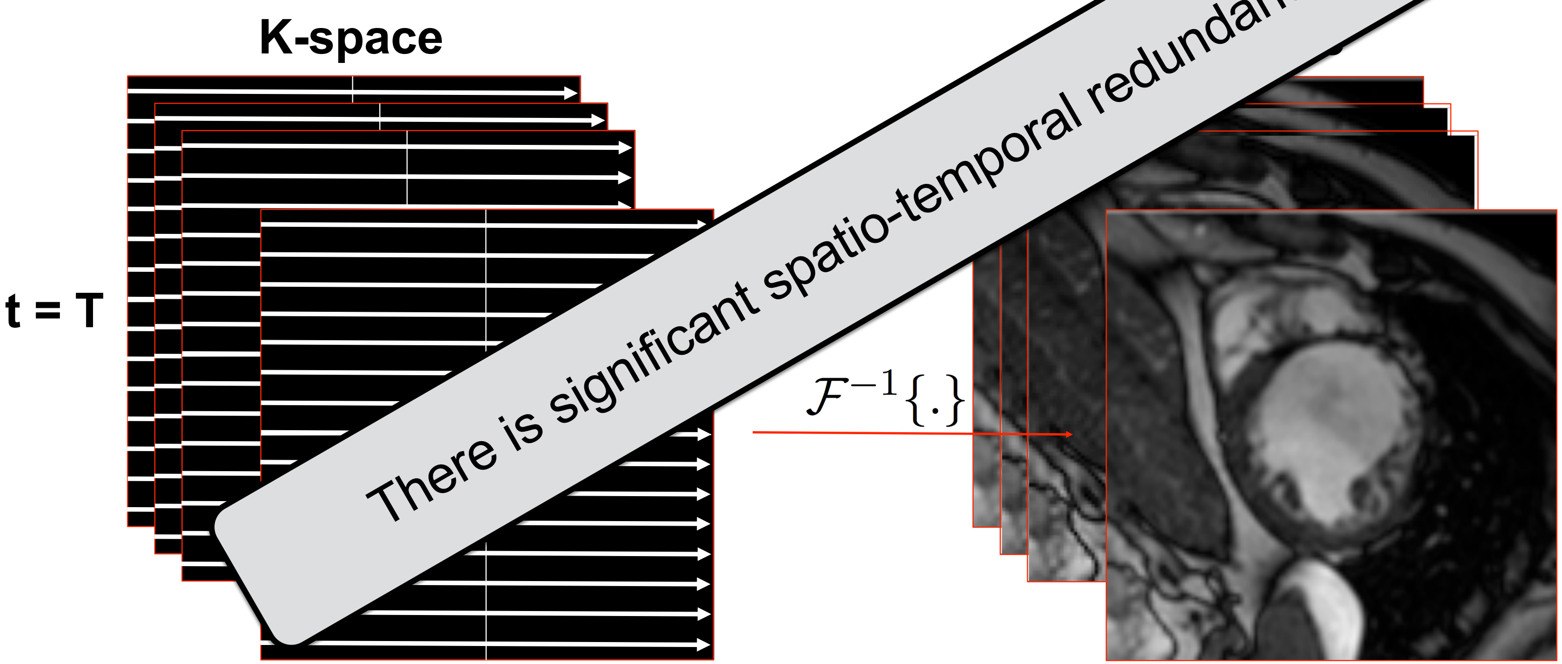

There is significant spatio-temporal redundancy

Deep learning for image reconstruction

K-space

Signal space

Full sampling (slow)

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

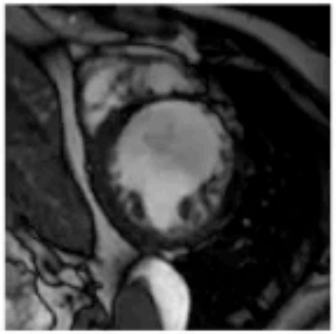

$\mathbf{x}_f$

learning-based reconstruction

e.g. compressed sensing

Deep learning for image reconstruction

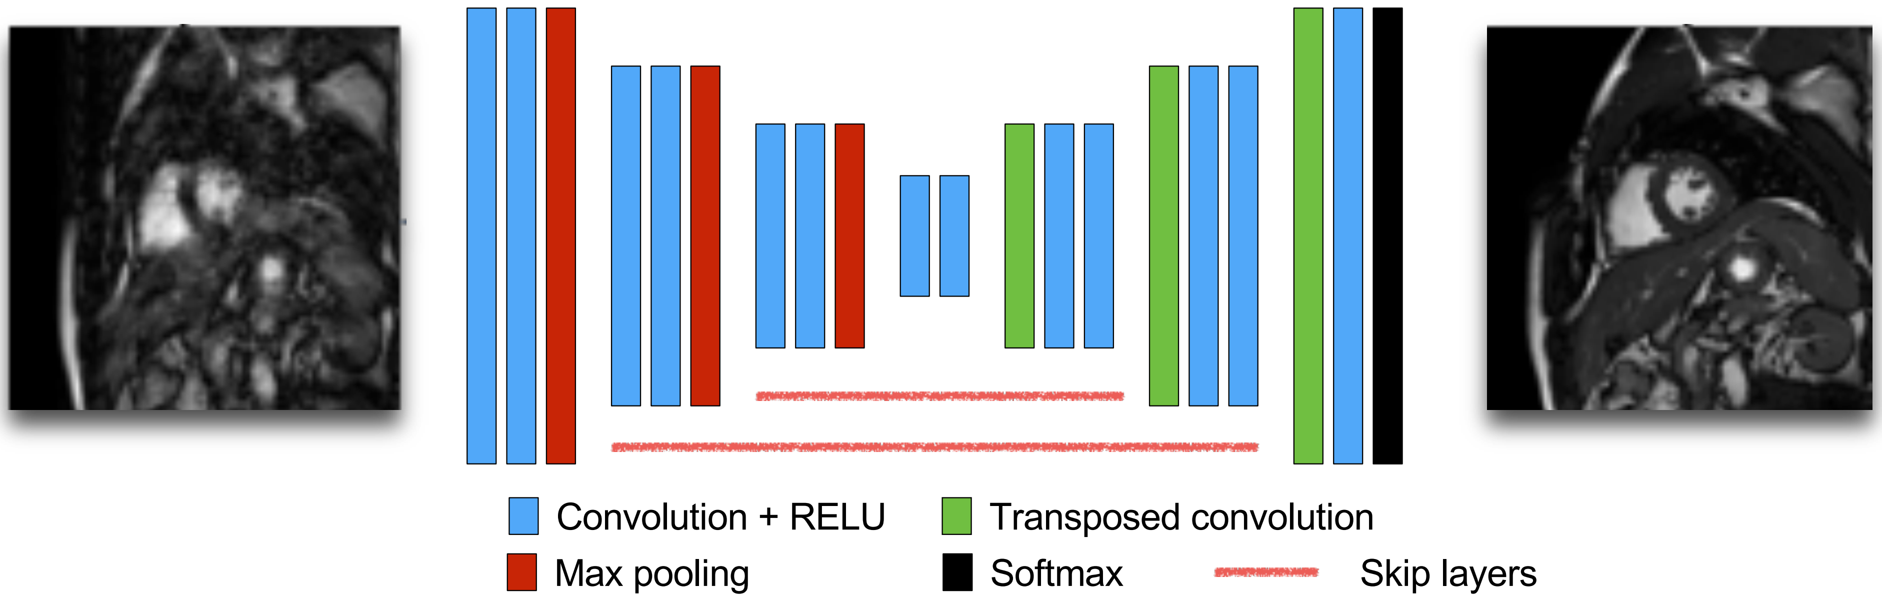

Deep learning for image reconstruction: How?

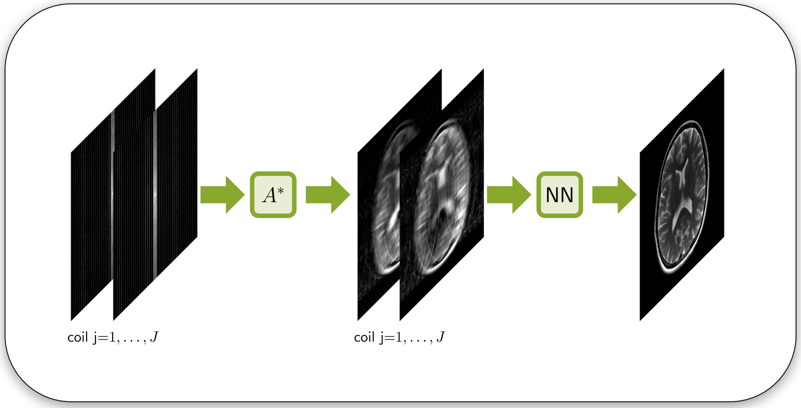

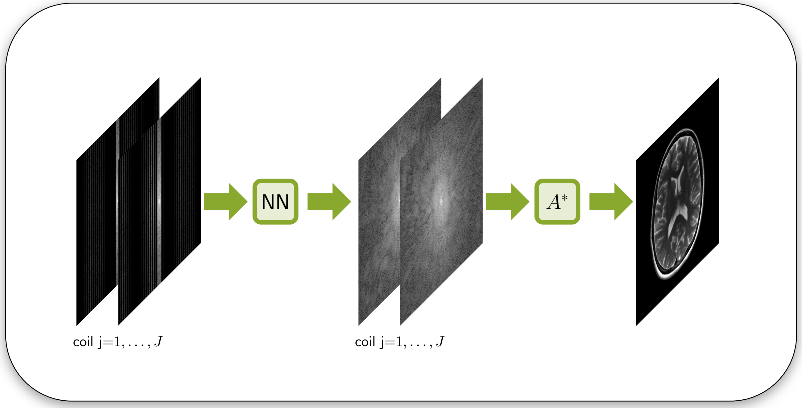

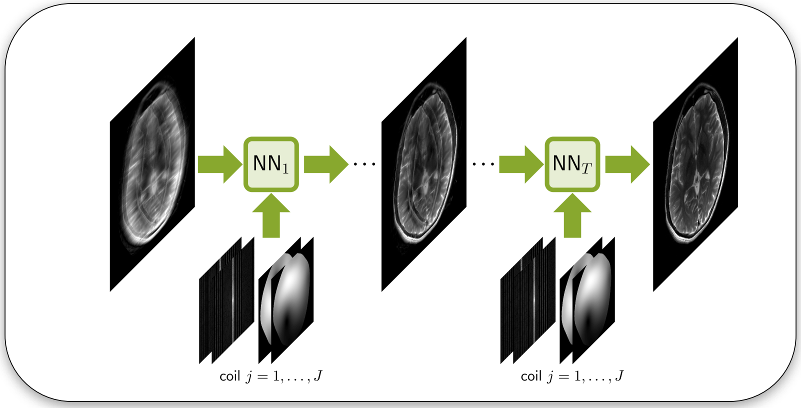

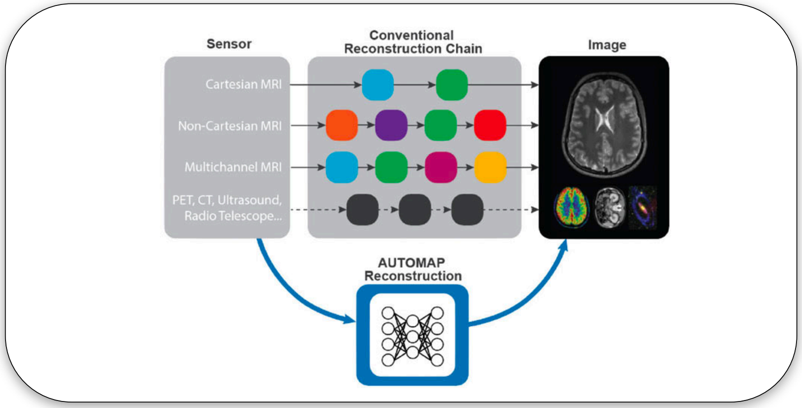

Deep learning for image reconstruction

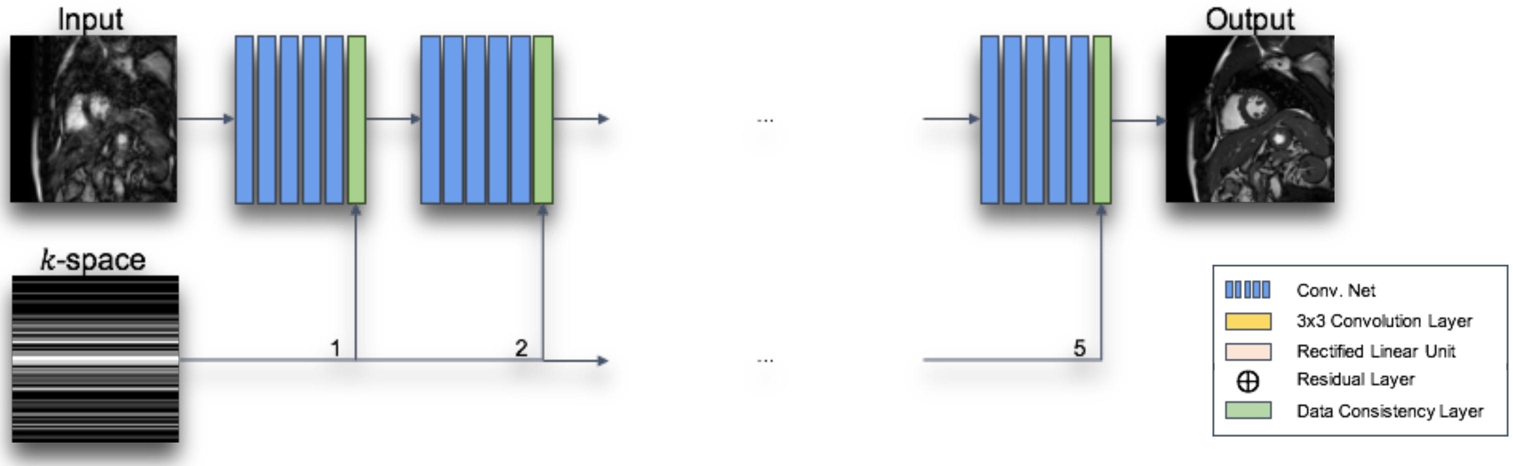

Schlemper et al. IEEE TMI 2018

Deep learning for image reconstruction

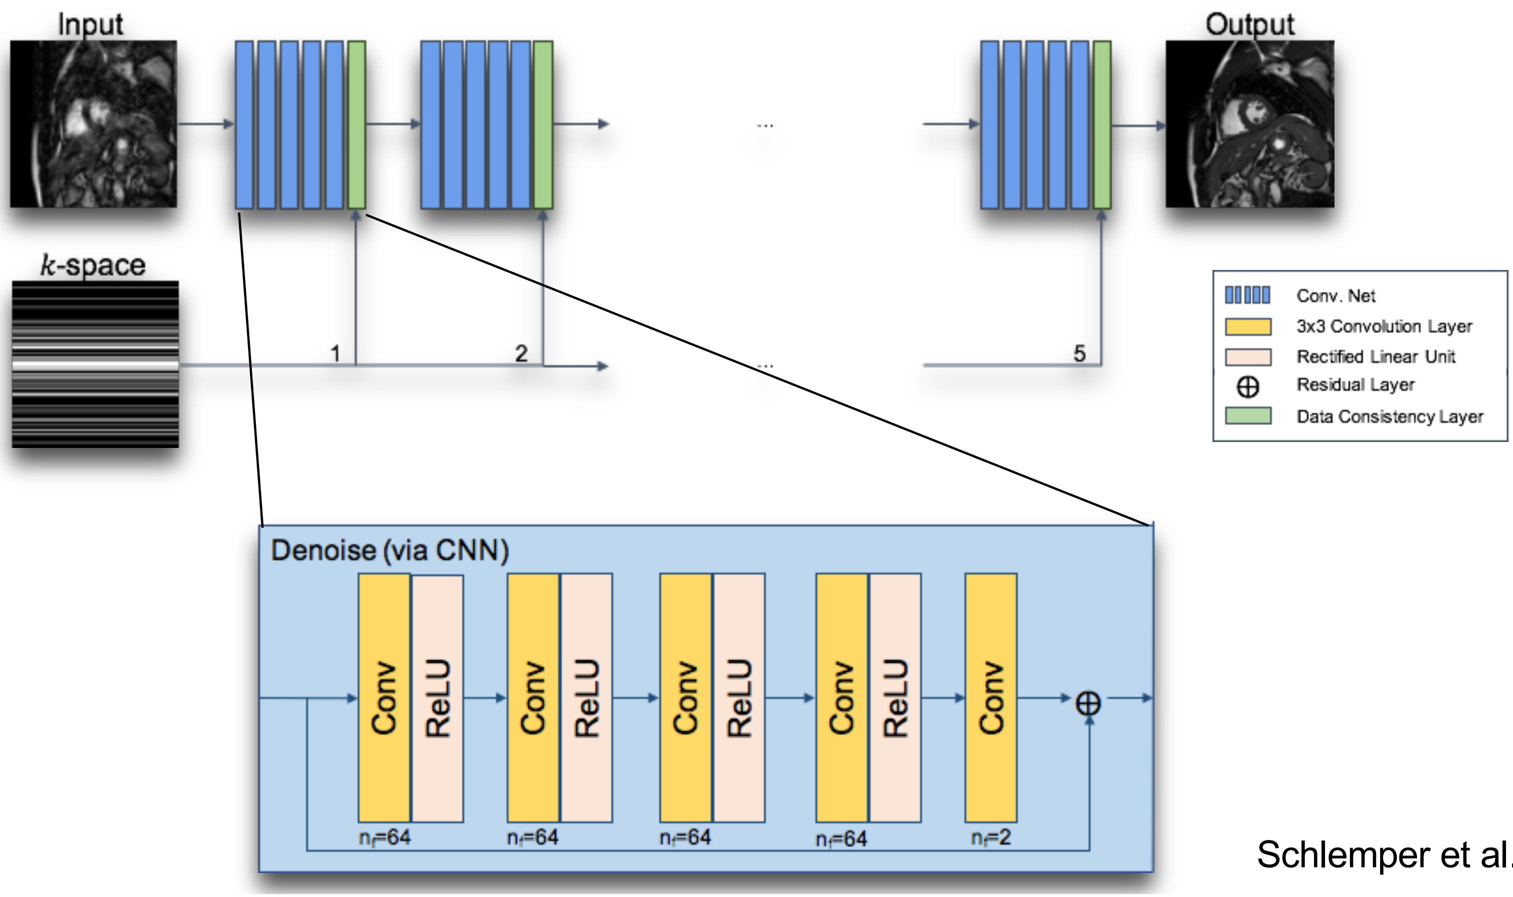

Schlemper et al. IEEE TMI 2018

Deep learning for image reconstruction

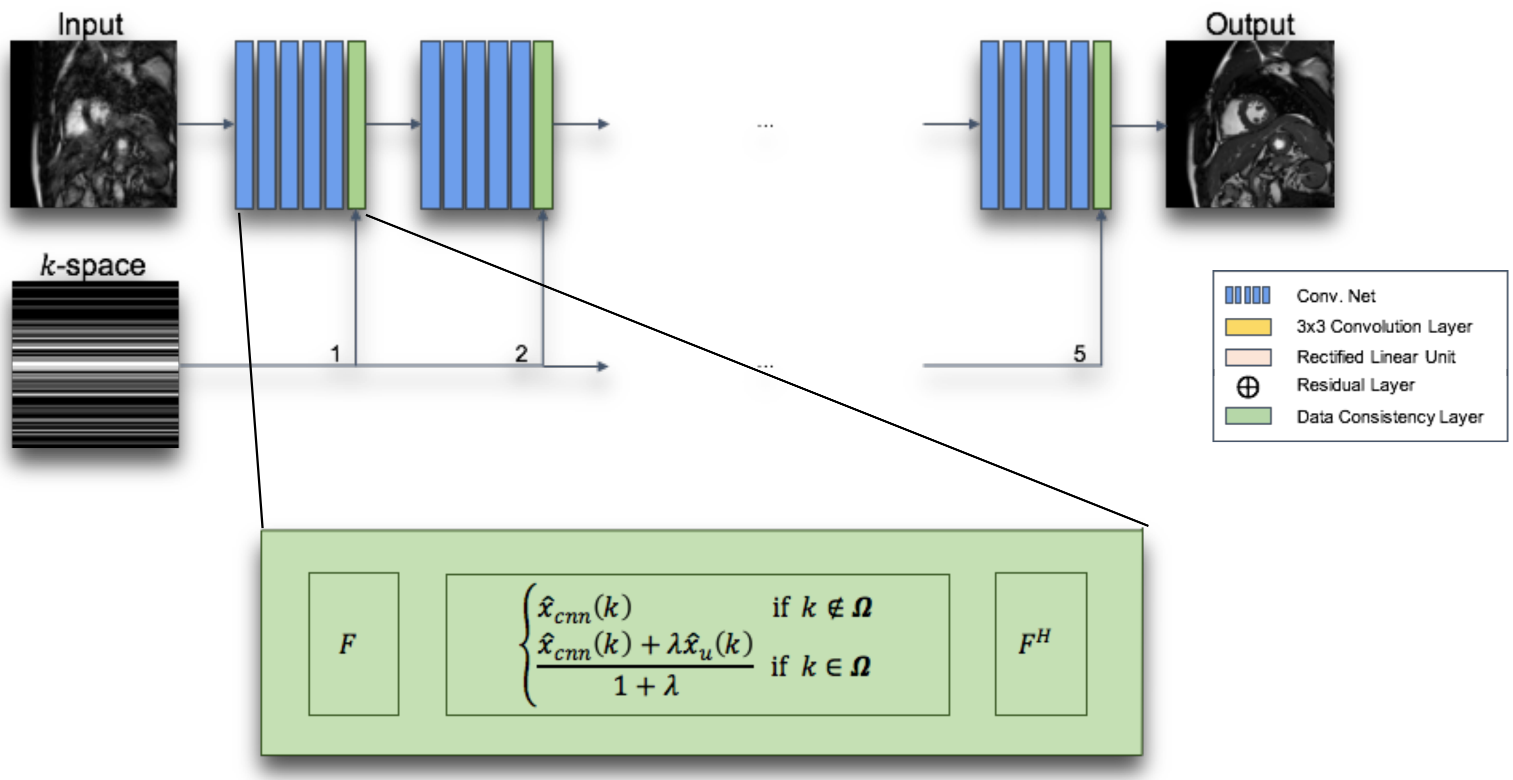

Schlemper et al. IEEE TMI 2018

Magnitude reconstruction (6-fold)

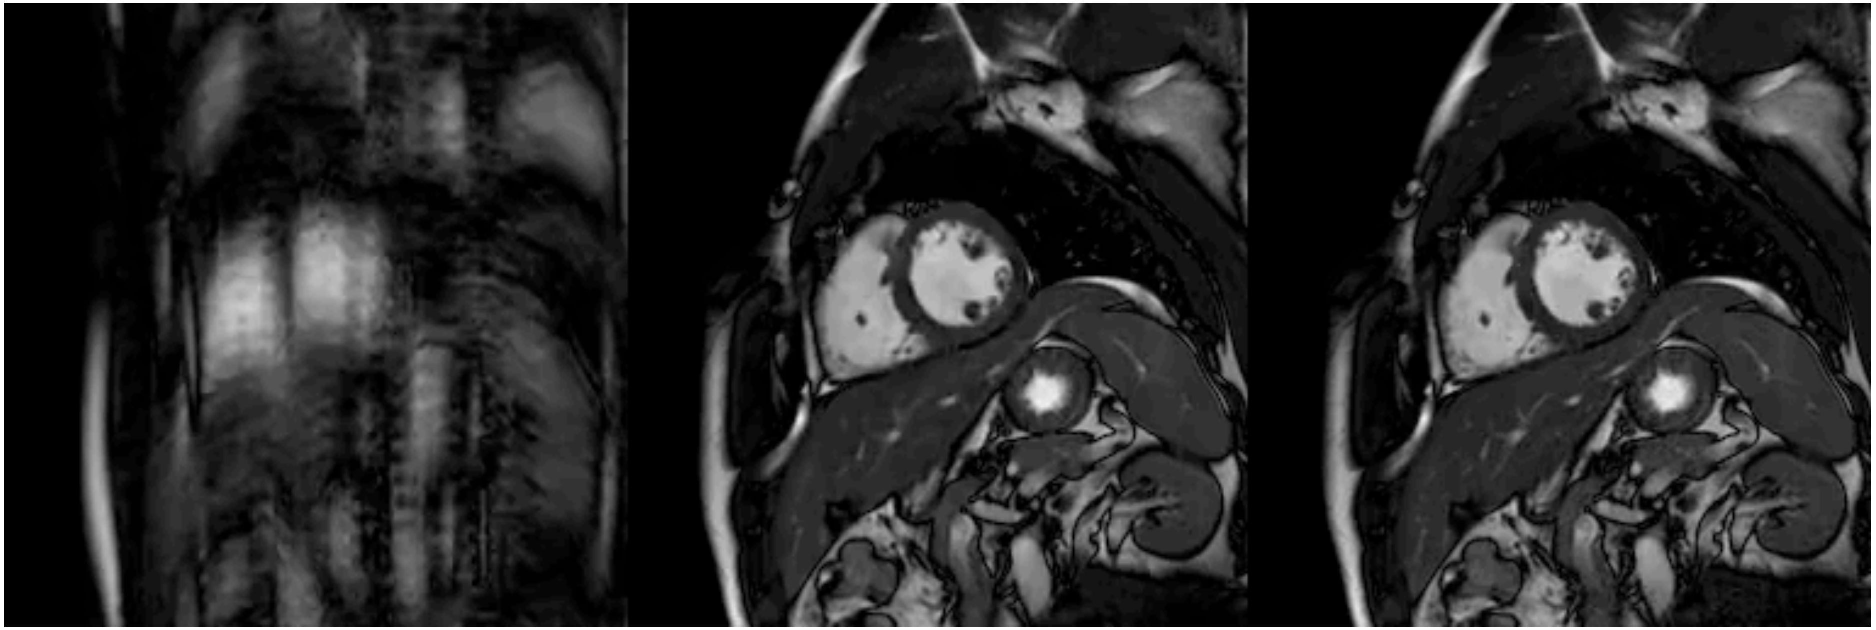

(a) 6x Undersampled

(b) CNN reconstruction

(c) Ground Truth

Schlemper et al. IEEE TMI 2018

Magnitude reconstruction (11-fold)

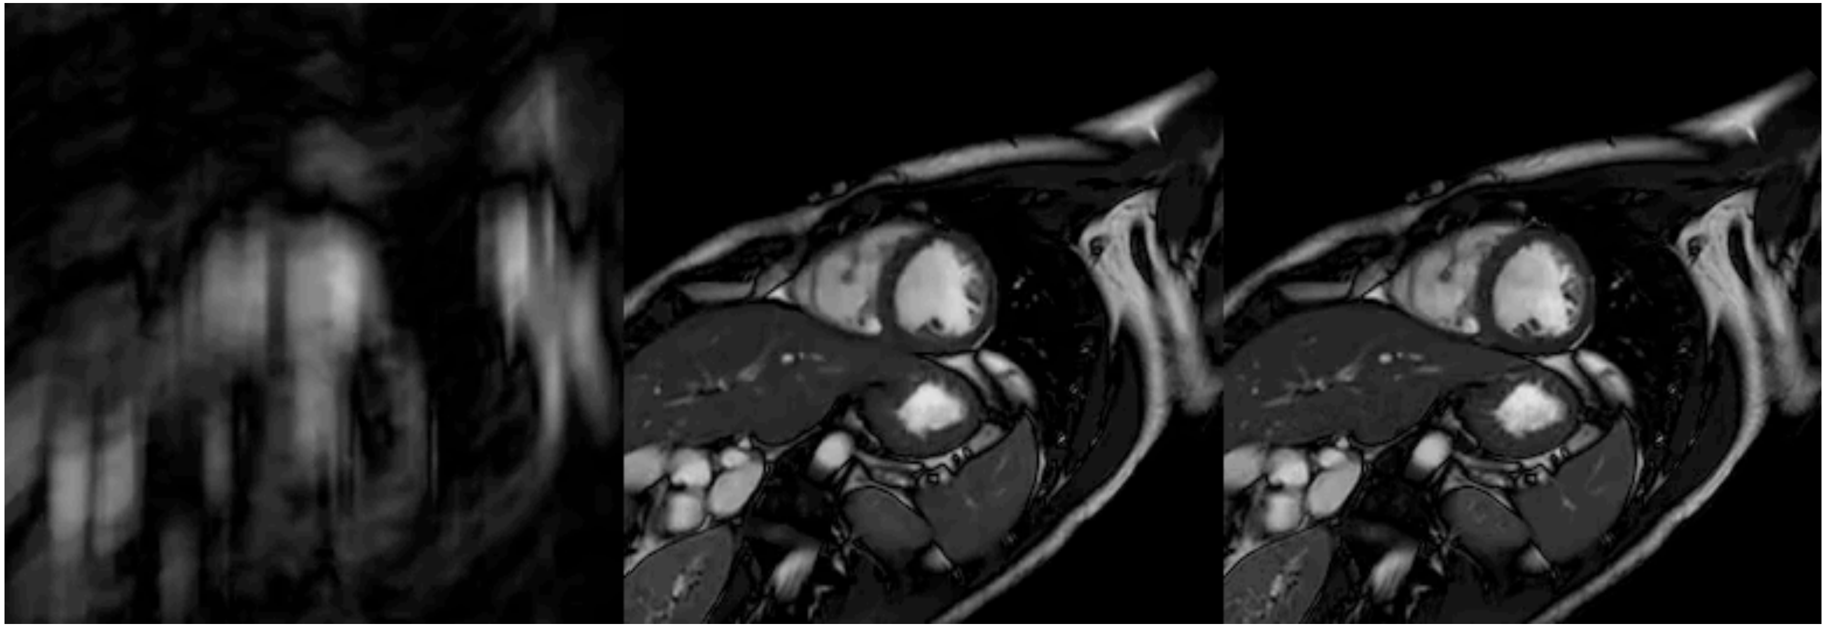

(a) 11x Undersampled (b) CNN reconstruction (c) Ground Truth

Schlemper et al. IEEE TMI 2018

Replacing Deep Cascade of CNNs with Recurrent Neural Networks

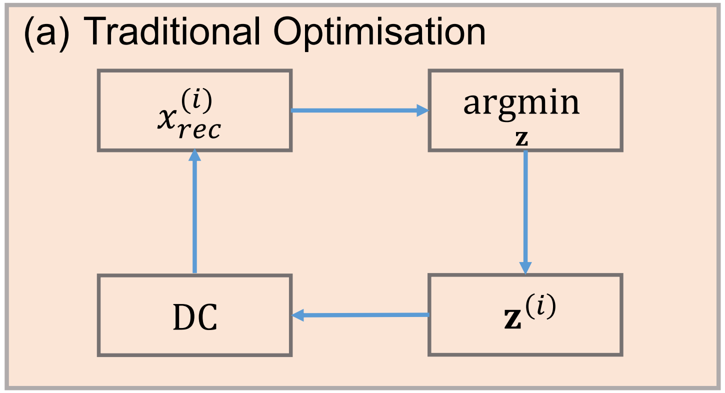

C. Qin et al. arXiv:1712.01751, IEEE TMI 2018

Replacing Deep Cascade of CNNs with Recurrent Neural Networks

RNNs over both iterations and time:

- Embed the ***iterative optimisation*** process in a learning setting
- Exploit ***spatio-temporal*** redundancies

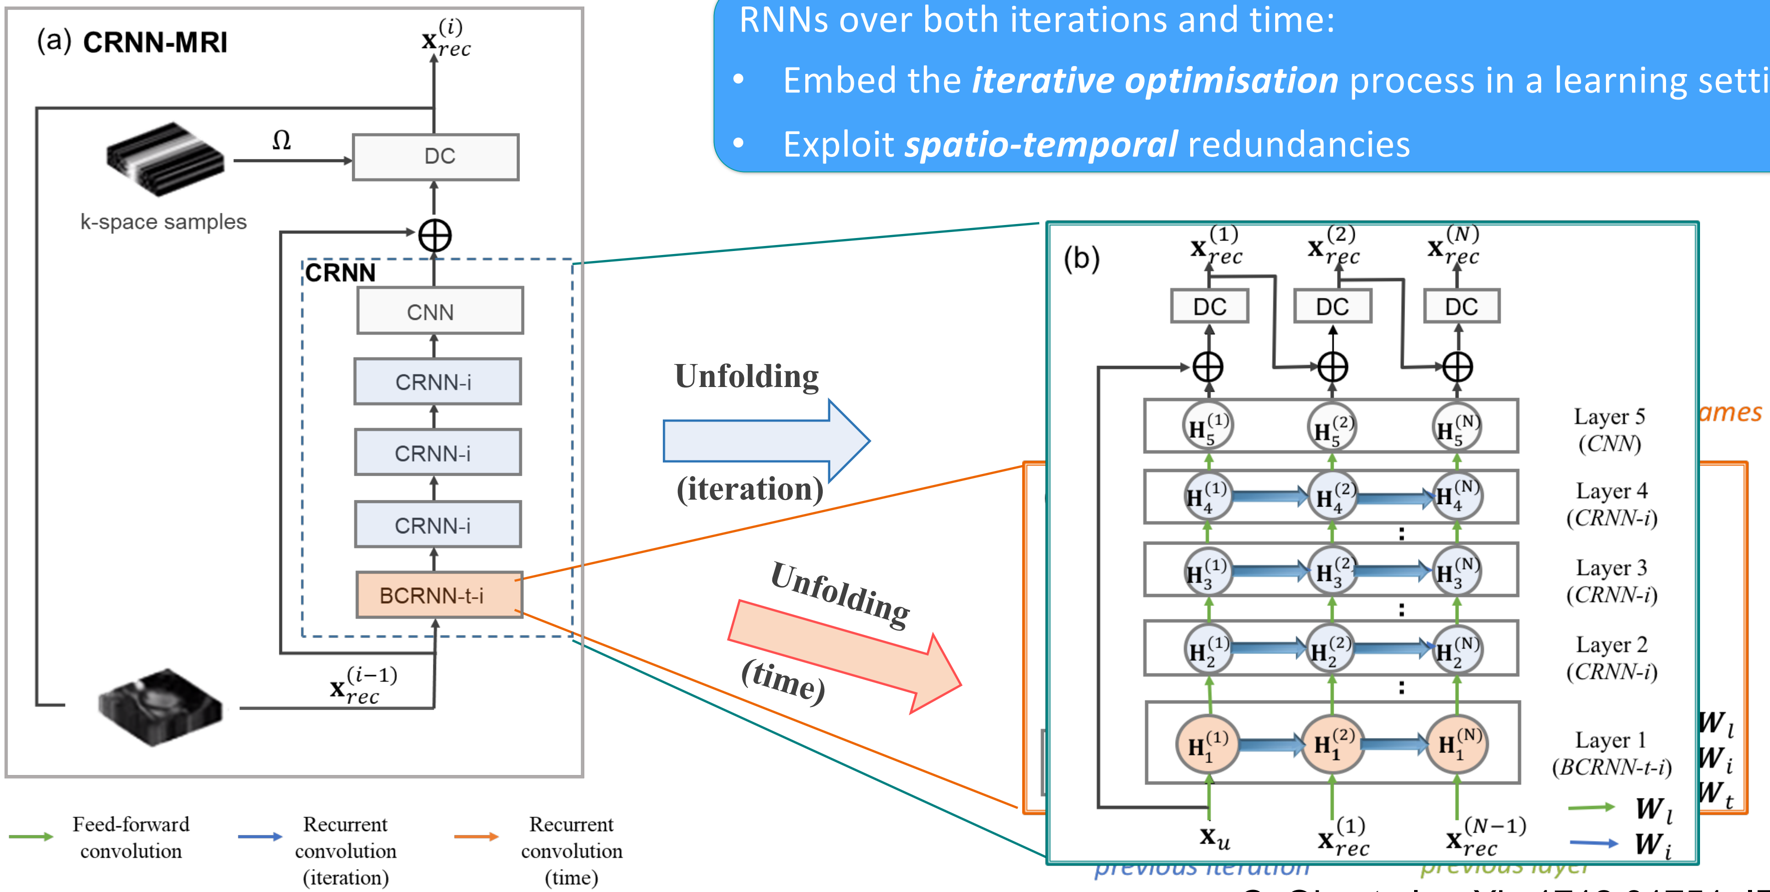

C. Qin et al. arXiv:1712.01751, IEEE TMI 2018

Replacing Deep Cascade of CNNs with Recurrent Neural Networks

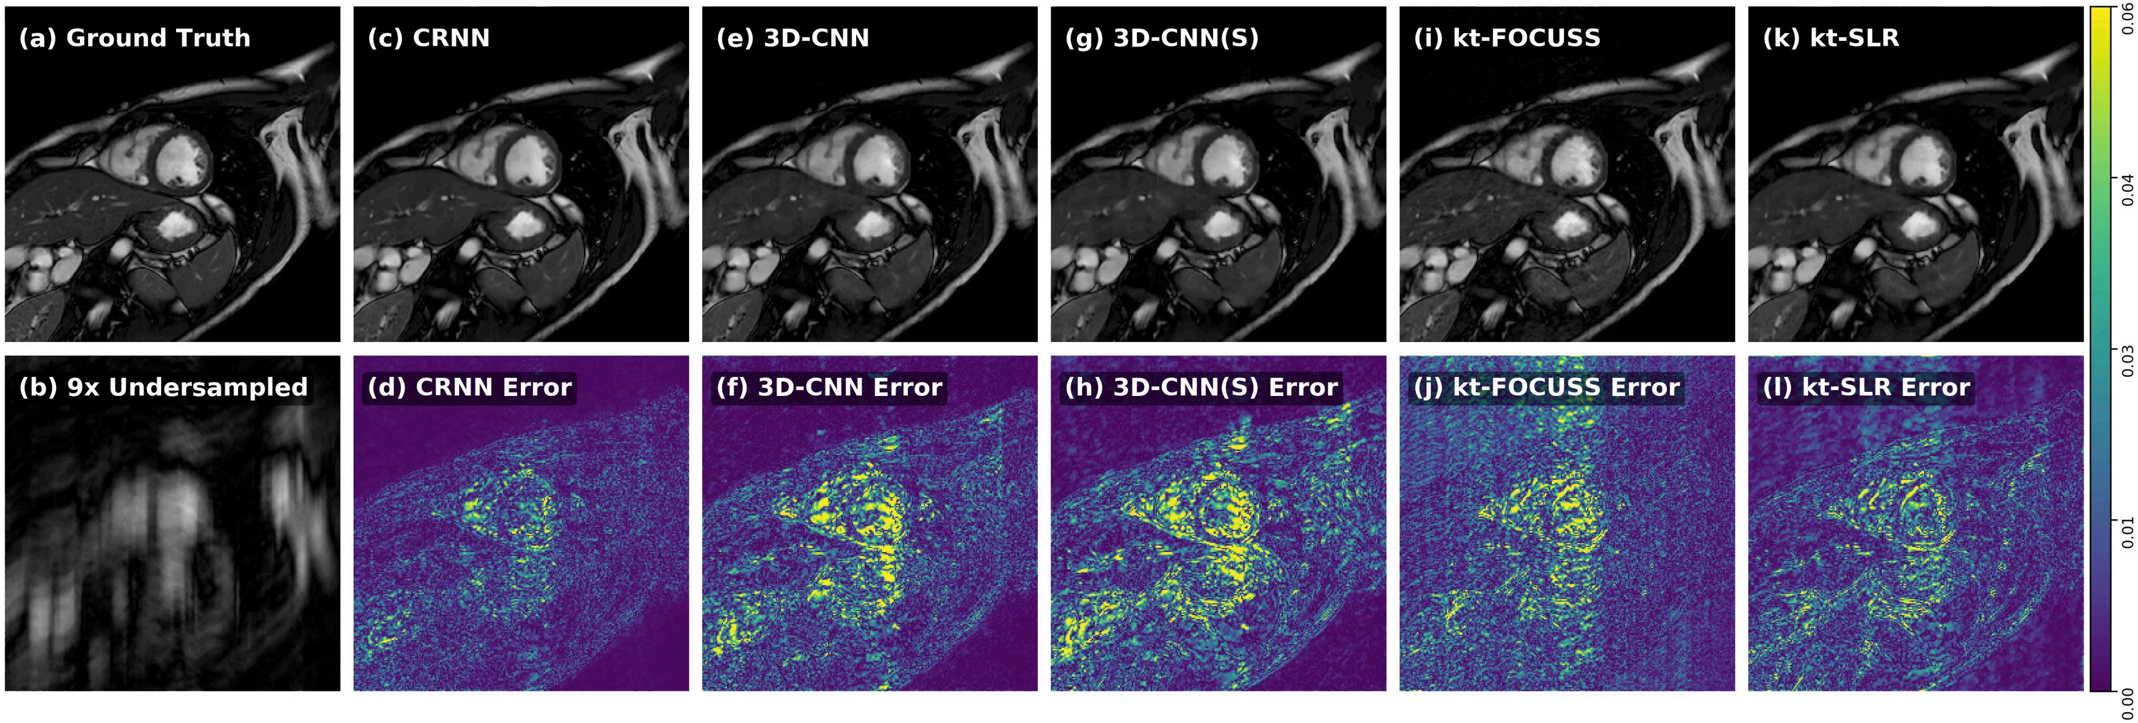

C. Qin et al. arXiv:1712.01751, IEEE TMI 2018

Exploiting *k-t* Correlations

k-t NEXT: Exploiting spatio-temporal redundancies in ***complementary*** domains

k-t NEtwork with X-f Transform:

- Exploit *k-t* correlations in ***x-f domain*** with CNNs
- Alternate between both ***x-f*** and ***image domains*** learning complementary information

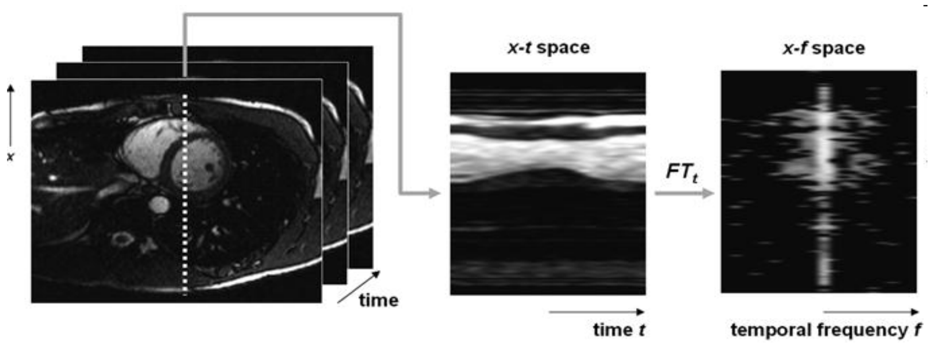

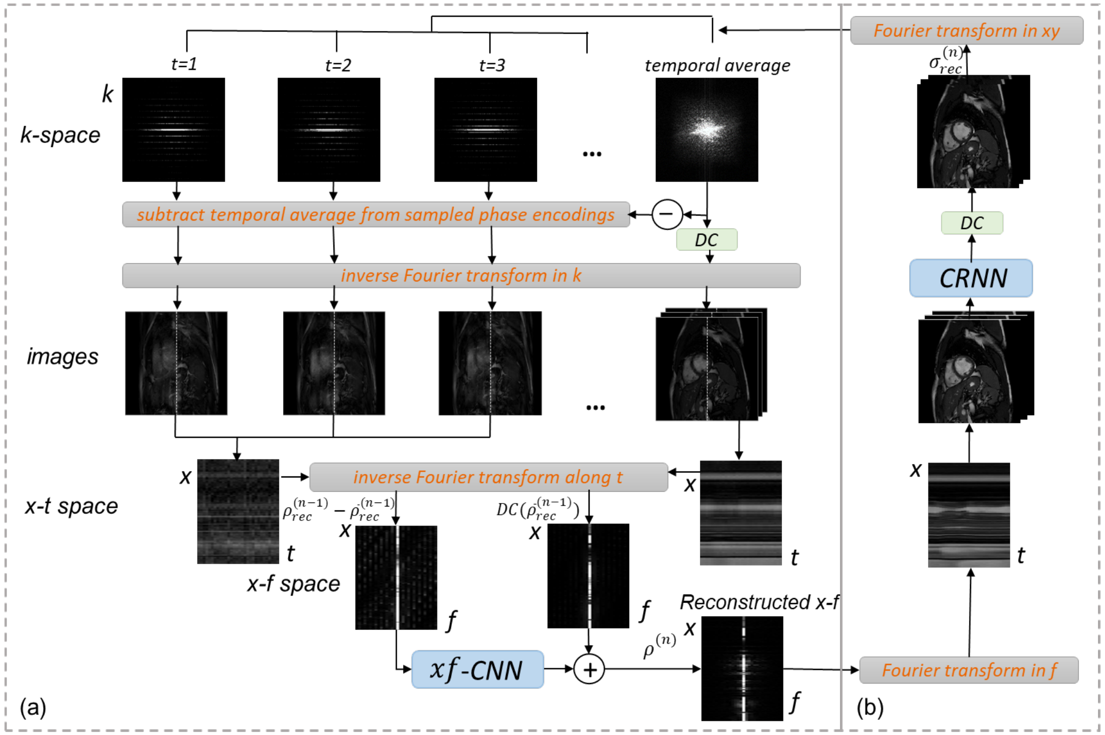

Qin et al. MICCAI, 2019

Exploiting *k-t* Correlations

k-t NEXT: Exploiting spatio-temporal redundancies in ***complementary*** domains

k-t NEtwork with X-f Transform:

- Exploit *k-t* correlations in ***x-f domain*** with CNNs
- Alternate between both ***x-f*** and ***image domains*** learning complementary information

x-f domain:

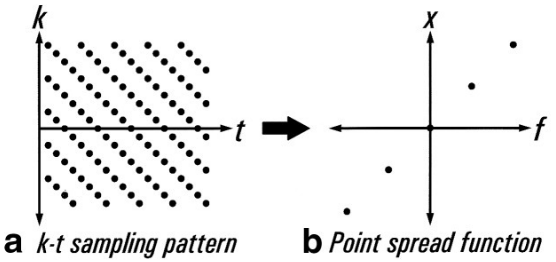

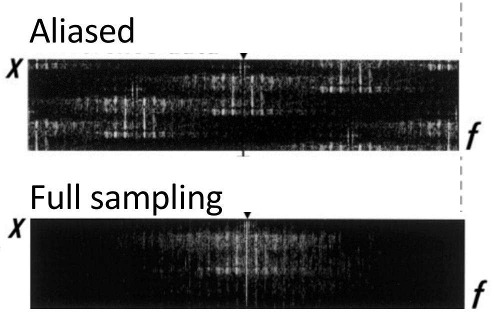

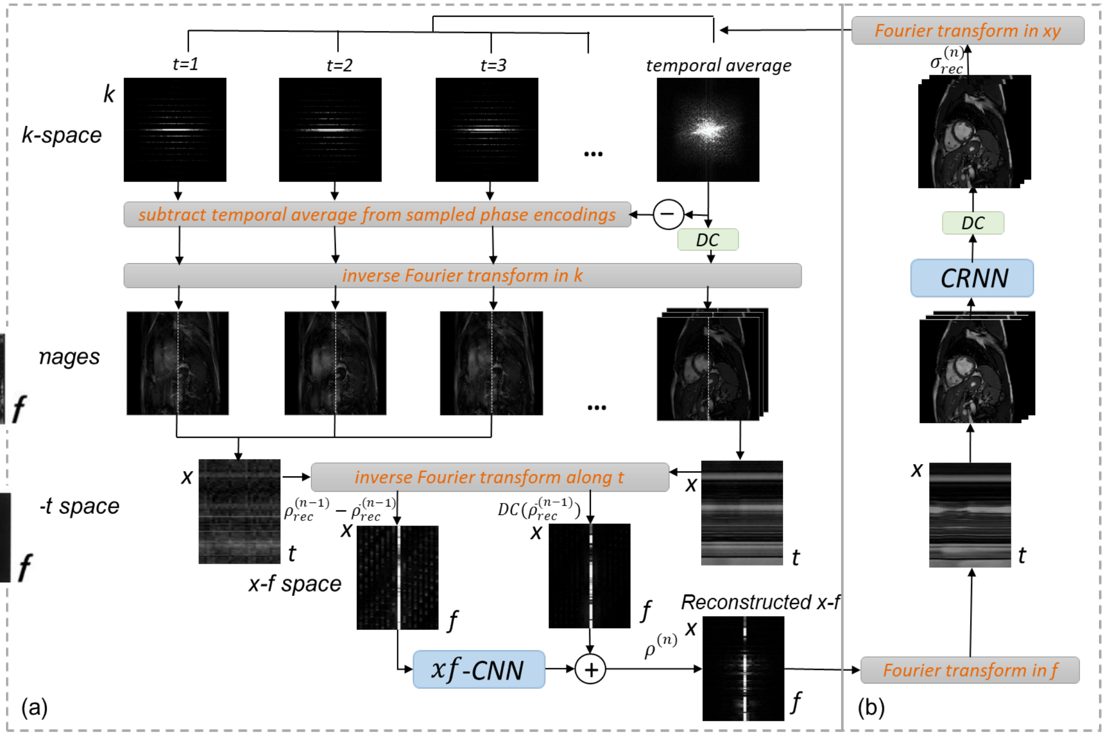

Qin et al. MICCAI, 2019

Exploiting *k-t* Correlations

k-t NEXT: Exploiting spatio-temporal redundancies in ***complementary*** domains

k-t NEtwork with X-f Transform:

- Exploit *k-t* correlations in ***x-f domain*** with CNNs
- Alternate between both ***x-f*** and ***image domains*** learning complementary information

x-f space

$$\rho^{(n)} = DC(\bar{\rho}_{rec}^{(n-1)}) + xf - \text{CNN}\left(\rho_{rec}^{(n-1)} - \bar{\rho}_{rec}^{(n-1)}\right)$$

$$\sigma_{rec}^{(n)} = \text{CRNN}\left(\mathcal{F}_f \rho^{(n)}; \mathbf{v}^{(0)}\right), \quad \rho_{rec}^{(n)} = \mathcal{F}_f^H \sigma_{rec}^{(n)}$$

image space

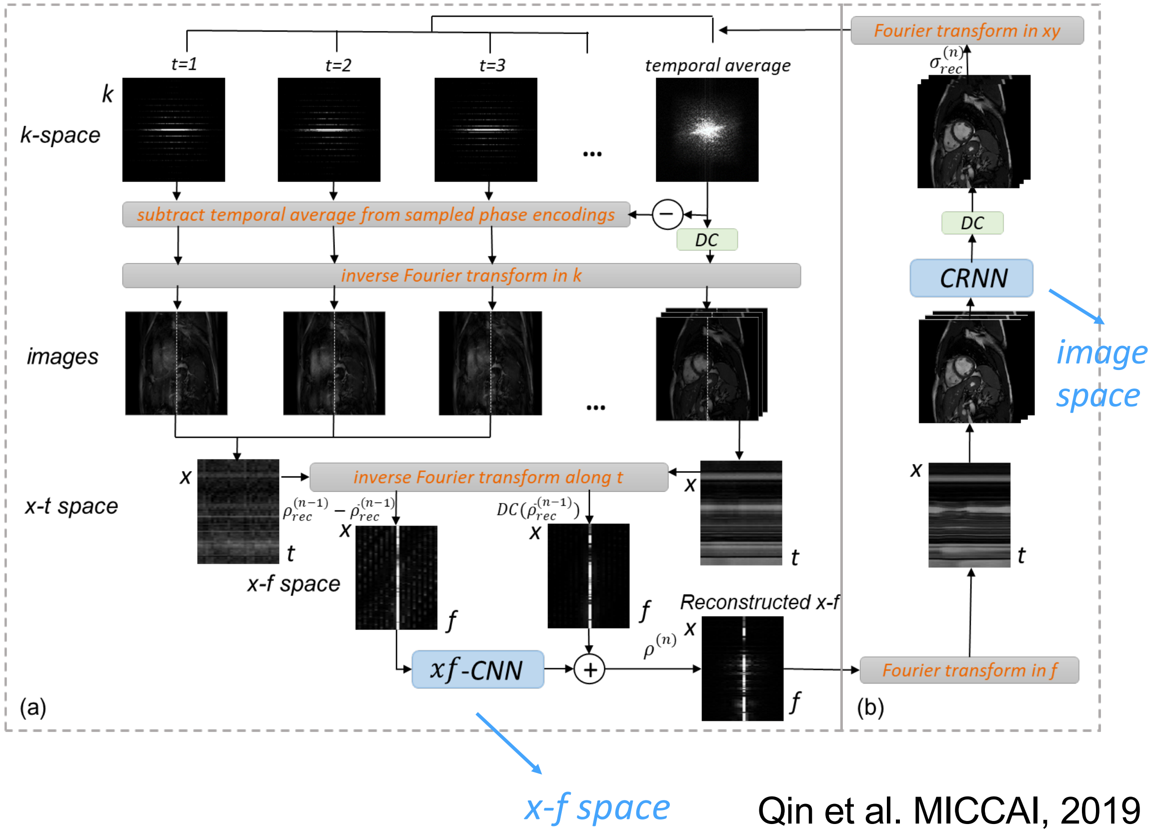

Qin et al. MICCAI, 2019

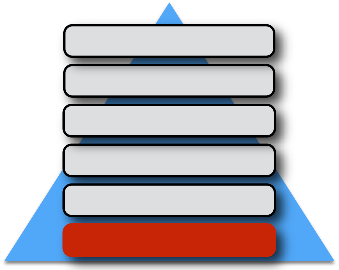

Exploiting *k-t* Correlations

PSNR (dB)/SSIM

	k-t FOCUSS	CRNN-MRI	k-t NEXT
9x	29.52/0.951	32.45/0.969	34.23/0.979
12x	28.14/0.937	31.30/0.962	33.28/0.975

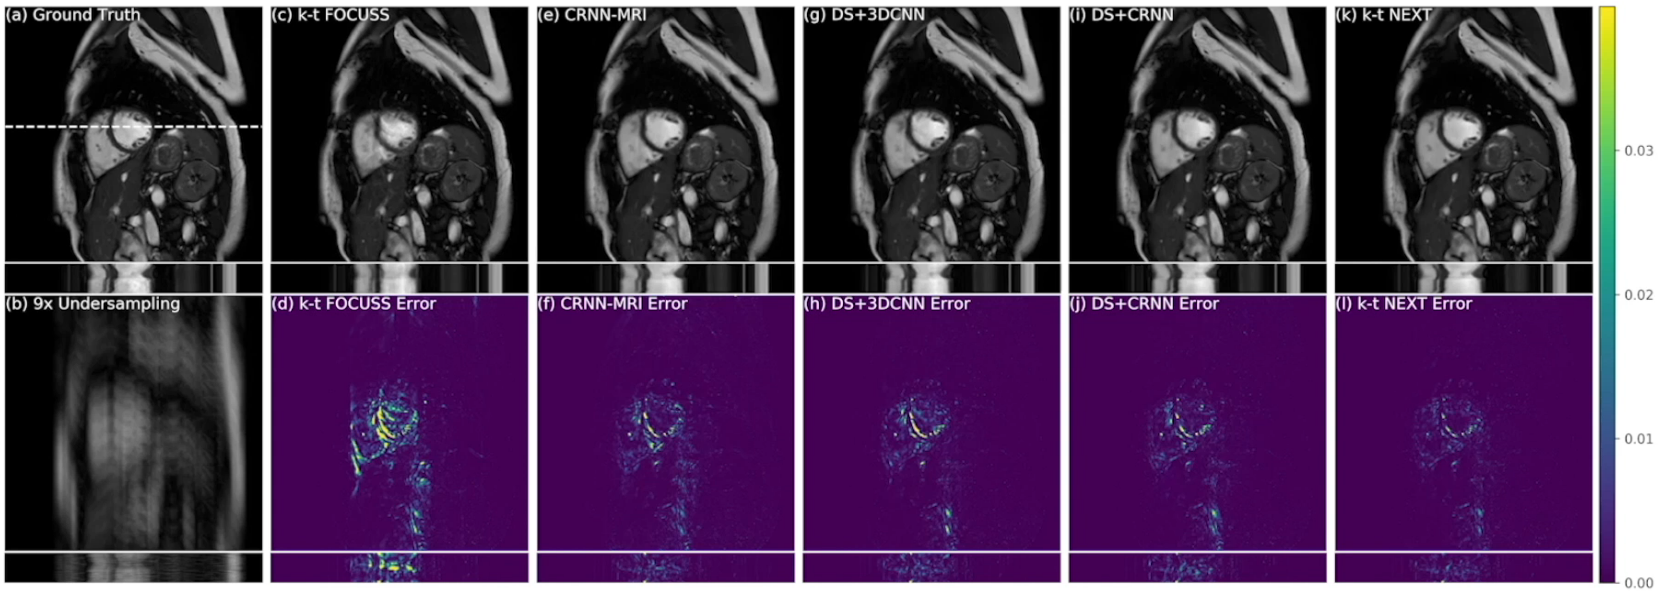

Qin et al. MICCAI, 2019

Exploiting motion for extremely undersampled dynamic MRI reconstruction

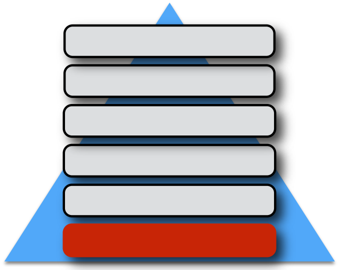

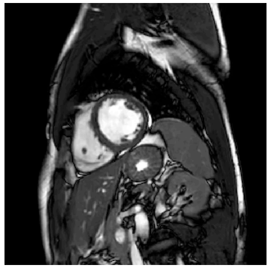

Fully sampled 'x1'
✗ High time-cost
✓ Accurate reconstruction

Exploiting motion for extremely undersampled dynamic MRI reconstruction

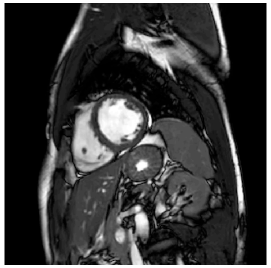

Fully sampled 'x1'
✗ High time-cost
✓ Accurate reconstruction

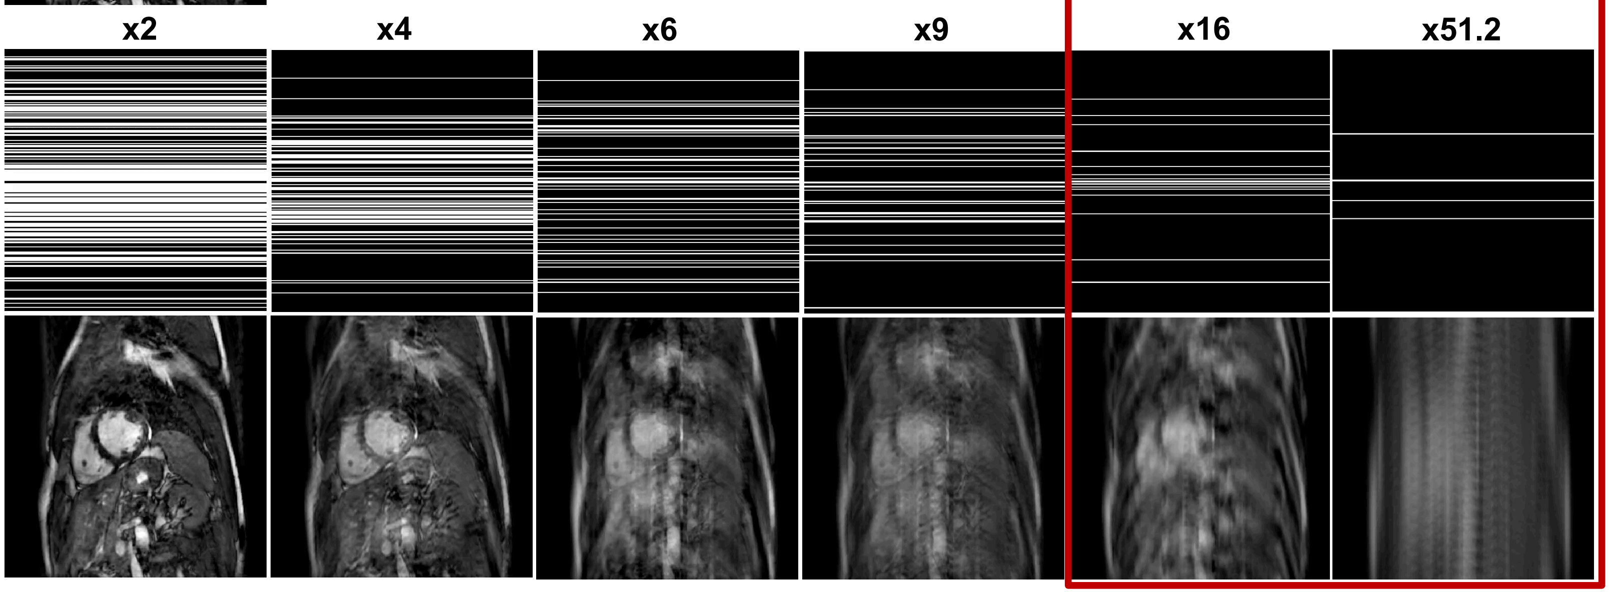

DC-CNN: Data-Consistent CNN

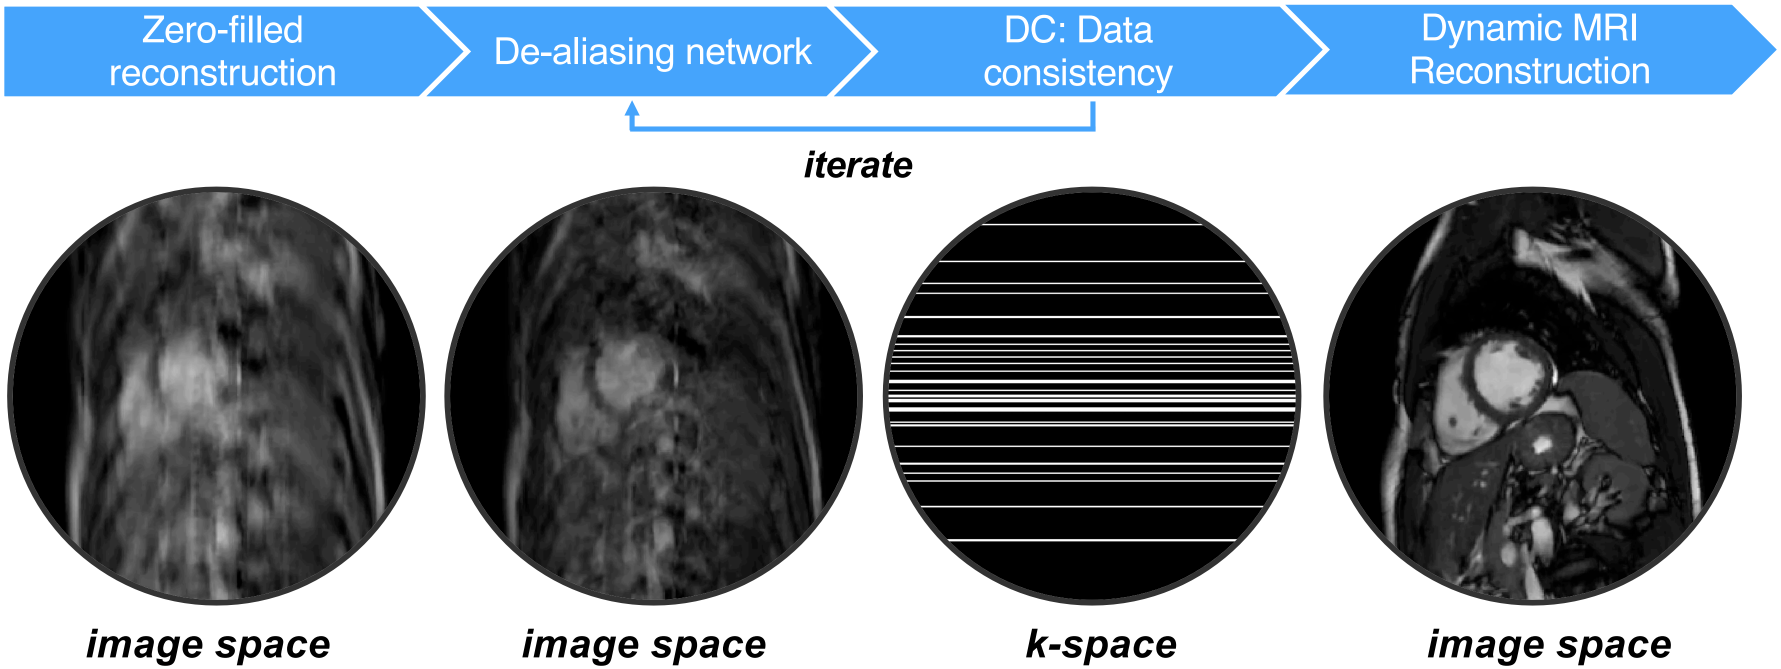

Schlemper *et. al* (2018)

Data-consistent motion-augmented cine (DC-MAC)

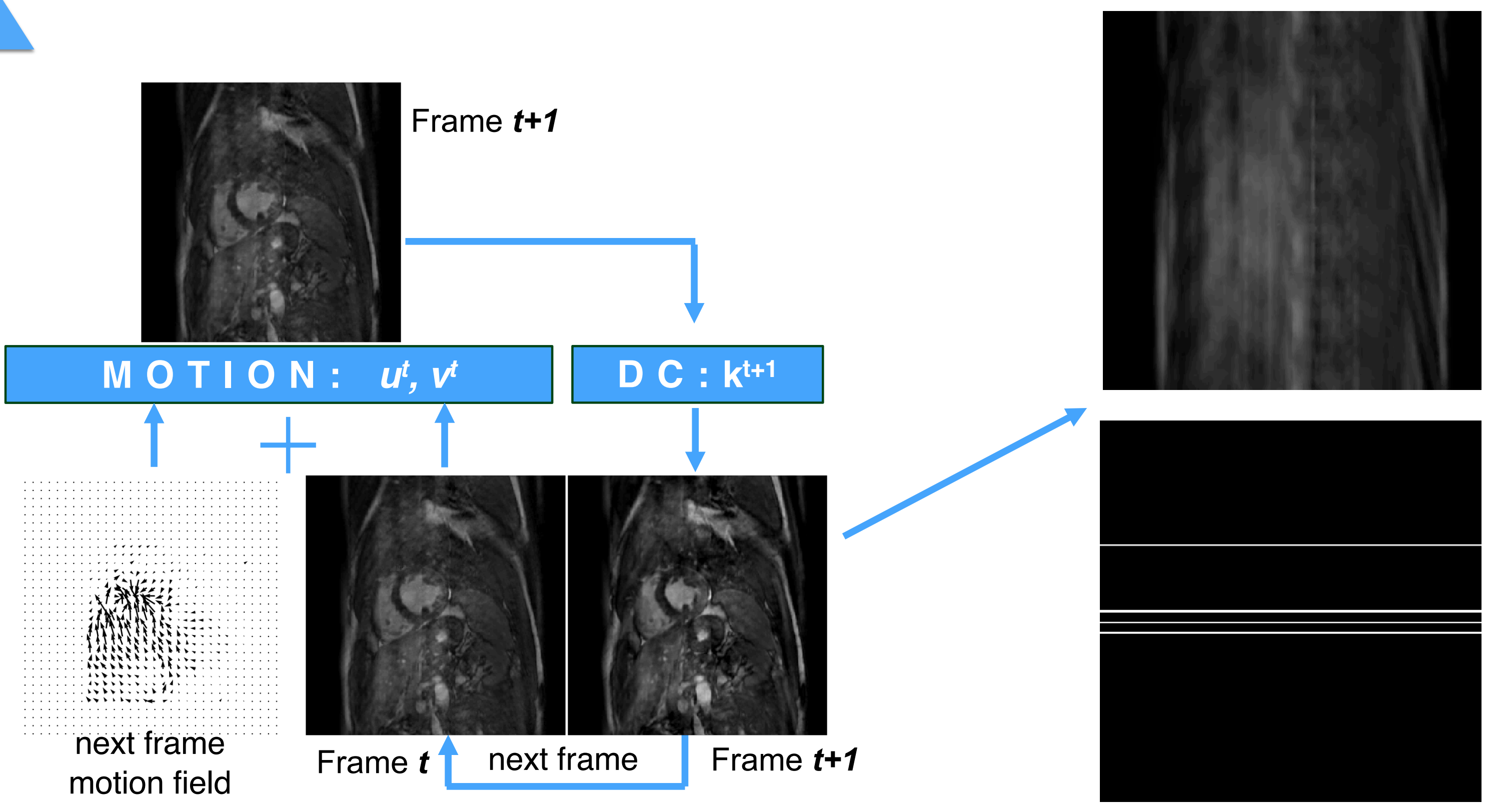

Exploiting motion for extremely undersampled dynamic MRI reconstruction
G. Seegoolam et al. MICCAI 2019, ISMRM 2020

ME-CNN: Motion-Exploiting CNN

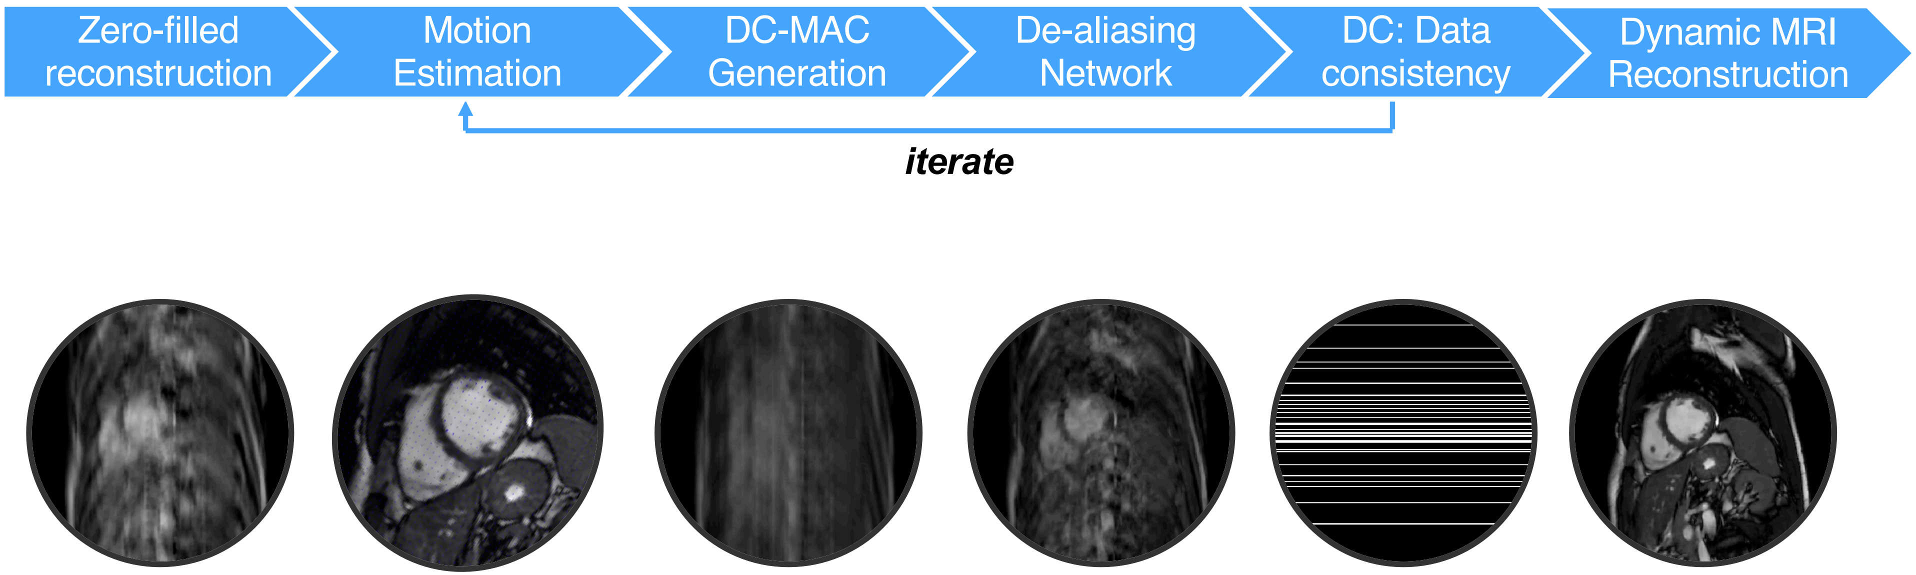

Exploiting motion for extremely undersampled dynamic MRI reconstruction
G. Seegoolam et al. MICCAI 2019, ISMRM 2020

Data-consistent motion-augmented-cine (DC-MAC) example (x16 acceleration rate)

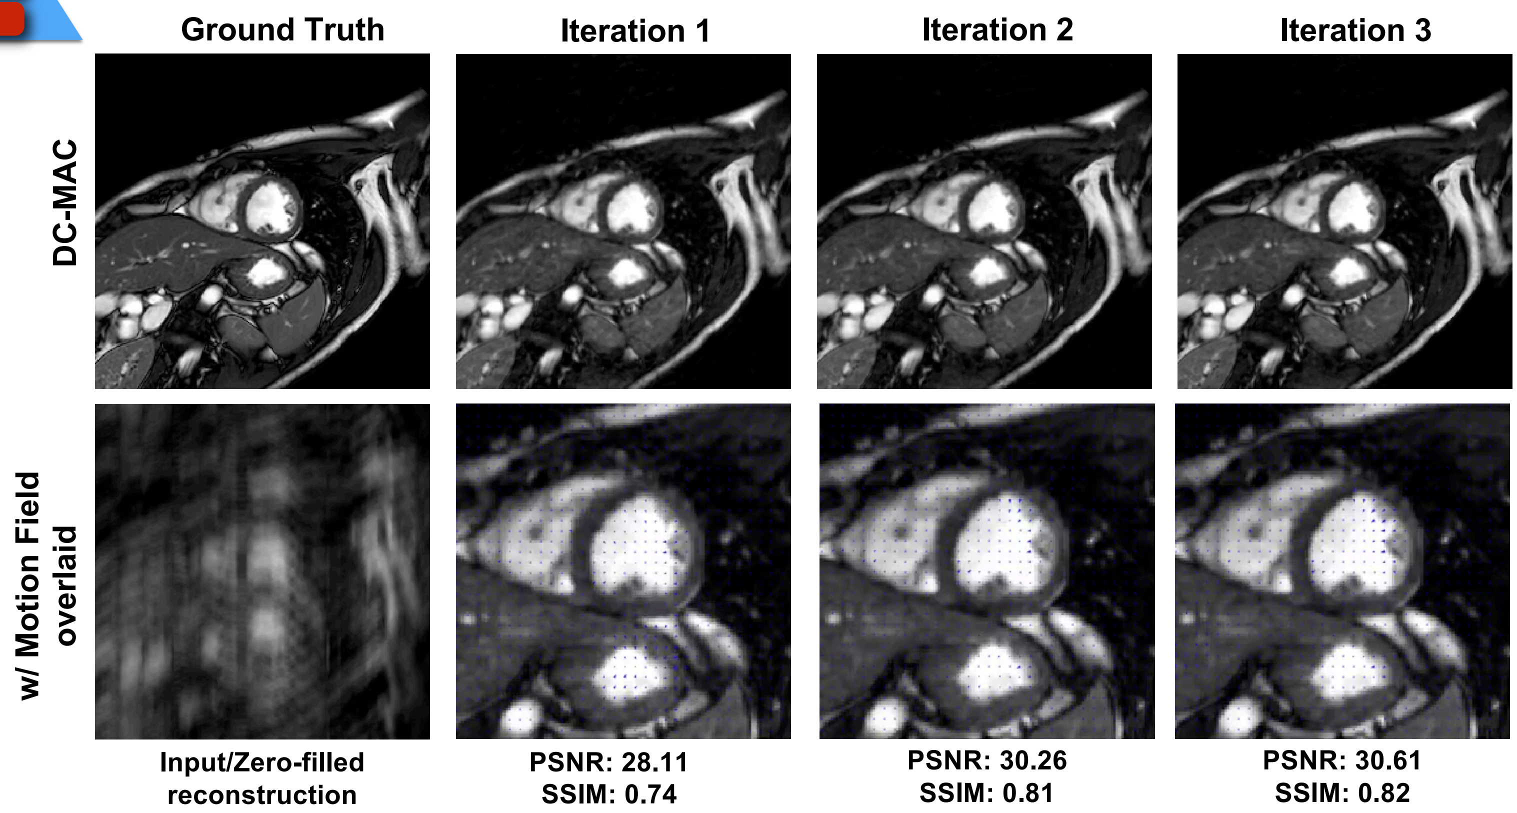

Exploiting motion for extremely undersampled dynamic MRI reconstruction
G. Seegoolam et al. MICCAI 2019, ISMRM 2020

ME-CNN vs DC-CNN example
(x16 acceleration rate)

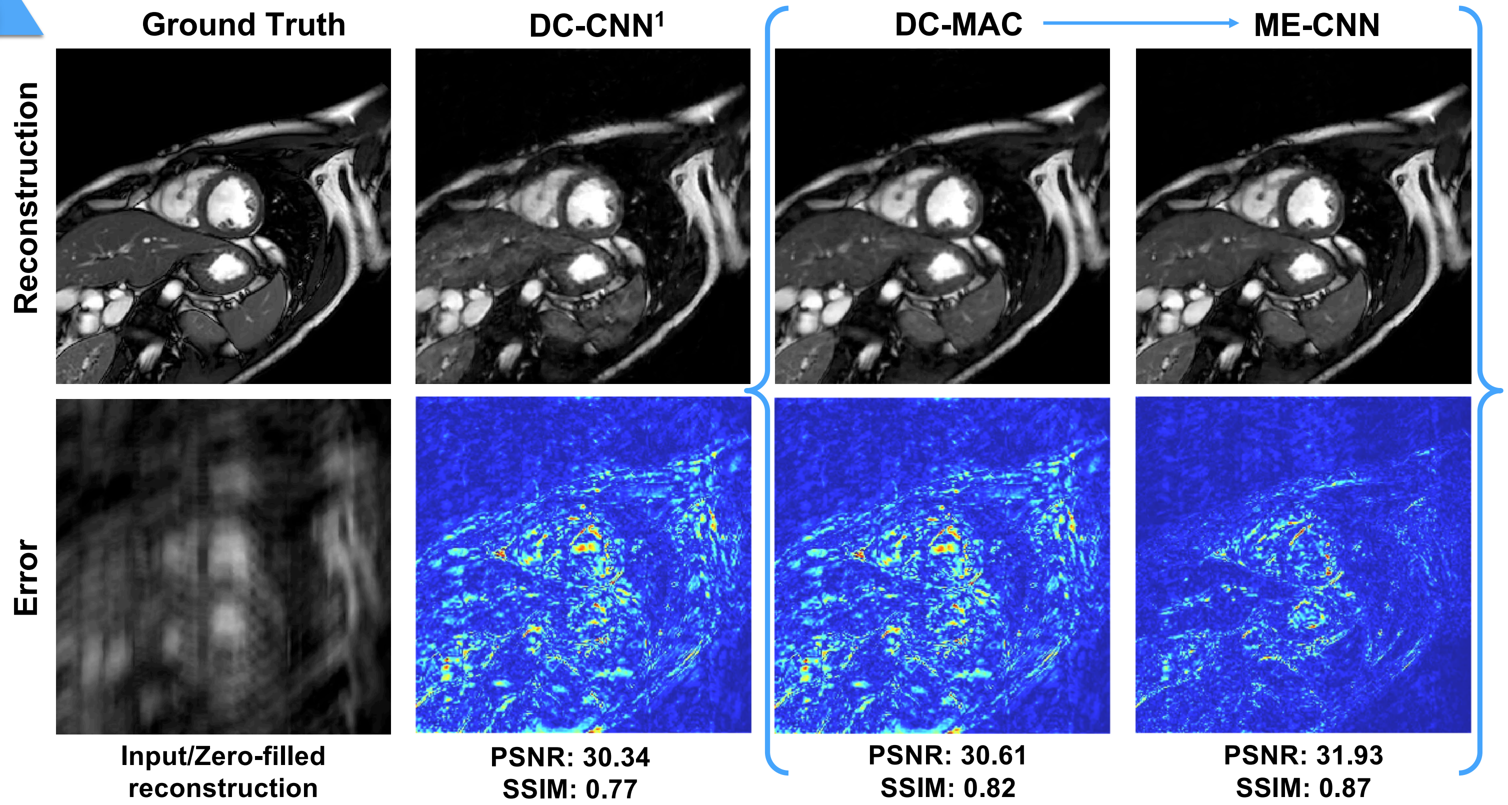

[1] Schlemper *et. al* (2018)

Exploiting motion for extremely undersampled dynamic MRI reconstruction
G. Seegoolam et al. MICCAI 2019, ISMRM 2020

ME-CNN vs DC-CNN example (x51 acceleration rate)

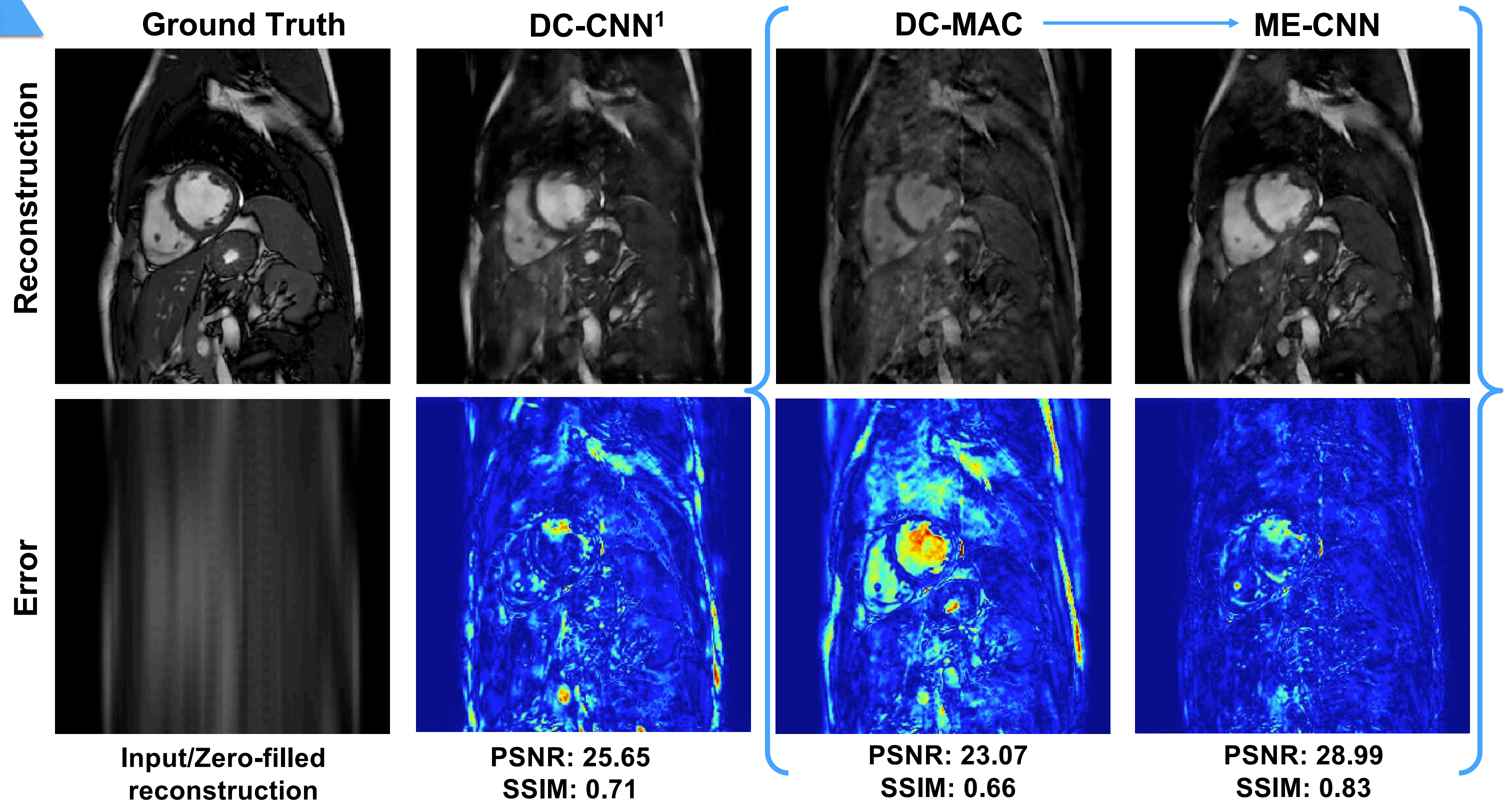

[1] Schlemper *et. al* (2018)

Exploiting motion for extremely undersampled dynamic MRI reconstruction
G. Seegoolam et al. MICCAI 2019, ISMRM 2020

Σ-net: Systematic Evaluation of Iterative Deep Neural Networks for Fast Parallel MR Image Reconstruction

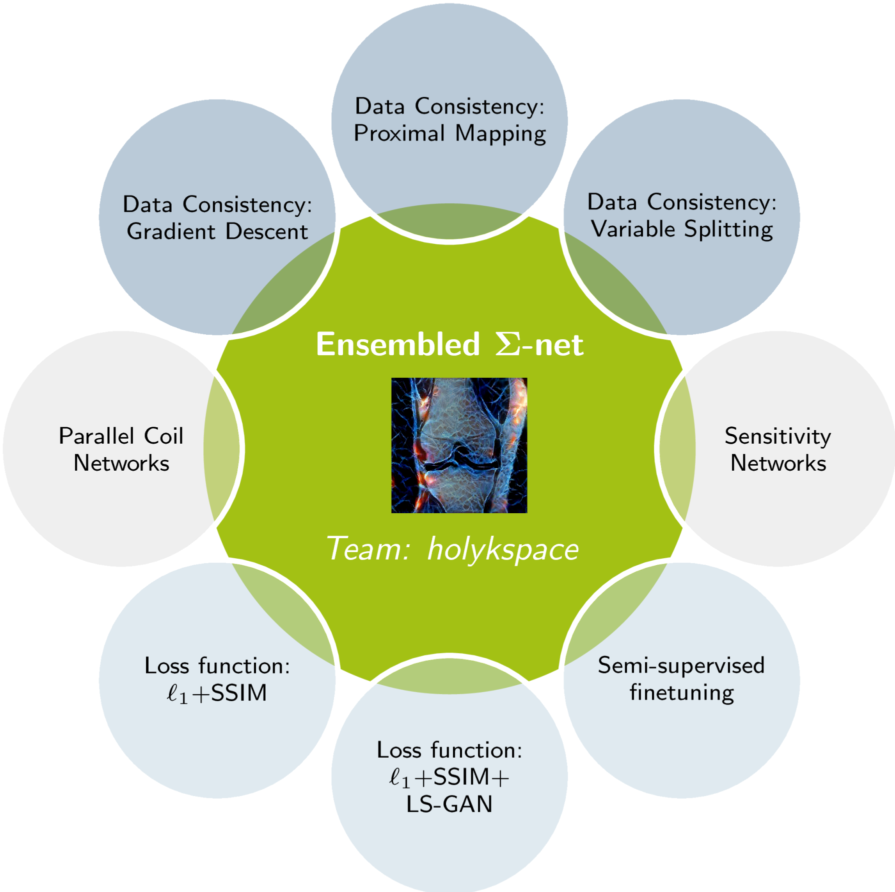

Wide-range comparison of SOTA approaches including

- various data consistency layers
- implicit and explicit coil combination
- loss functions
- model ensembling

Paper:
Hammernik et al. https://arxiv.org/abs/1912.09278
Schlemper et al. https://arxiv.org/abs/1912.05480

Code:
https://github.com/khammernik/sigmanet

Σ-net: Systematic Evaluation of Iterative Deep Neural Networks for Fast Parallel MR Image Reconstruction

Architectures

- *Sensitivity Networks*: Require explicit coil information
- *Parallel Coil Networks*: Learn implicit coil combination

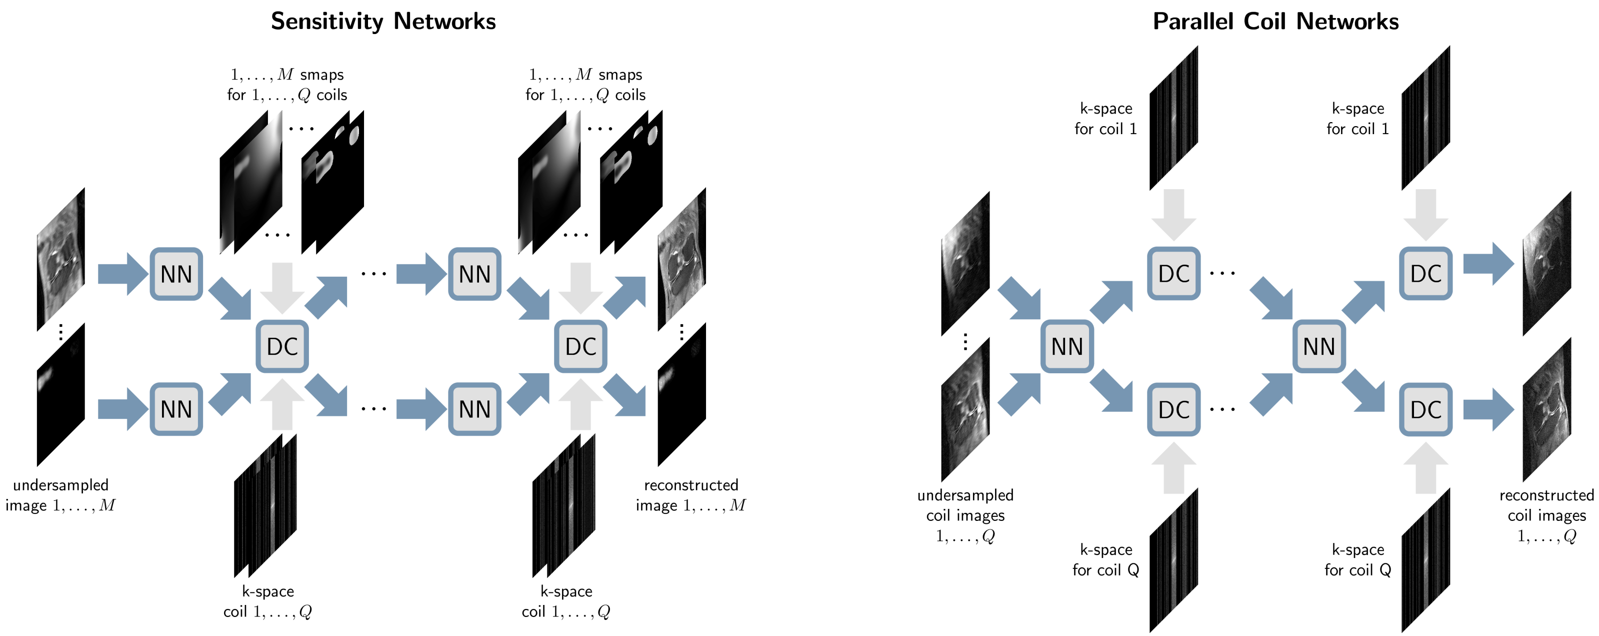

Σ-net: fastMRI challenge submission

Sensitivity networks required a *style-transfer layer* to match the intensities of the target root-sum-of-squares (RSS) reconstruction. This has NO practical relevance!

Example reconstruction for a fat-saturated PDw image of the fastMRI validation set at R=8.

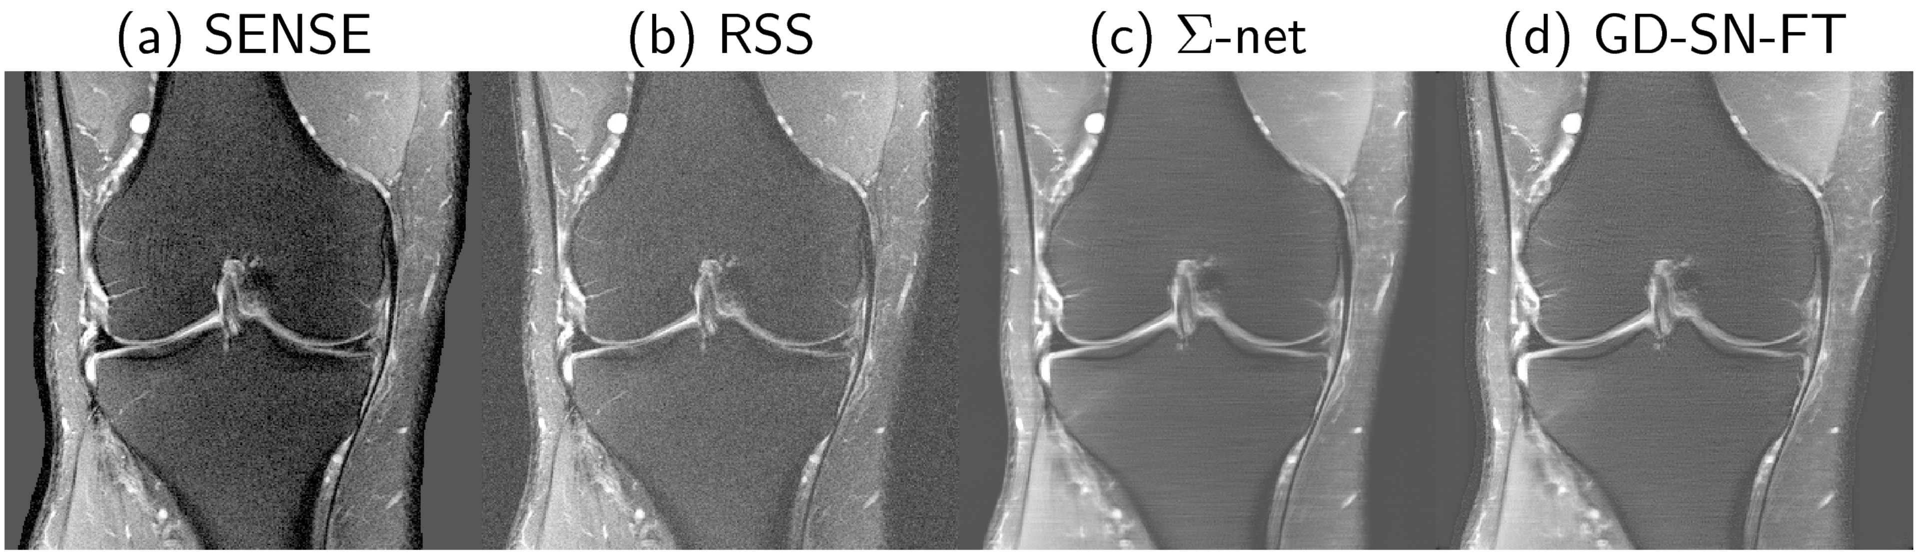

Image Enhancement vs Σ-net Reconstruction

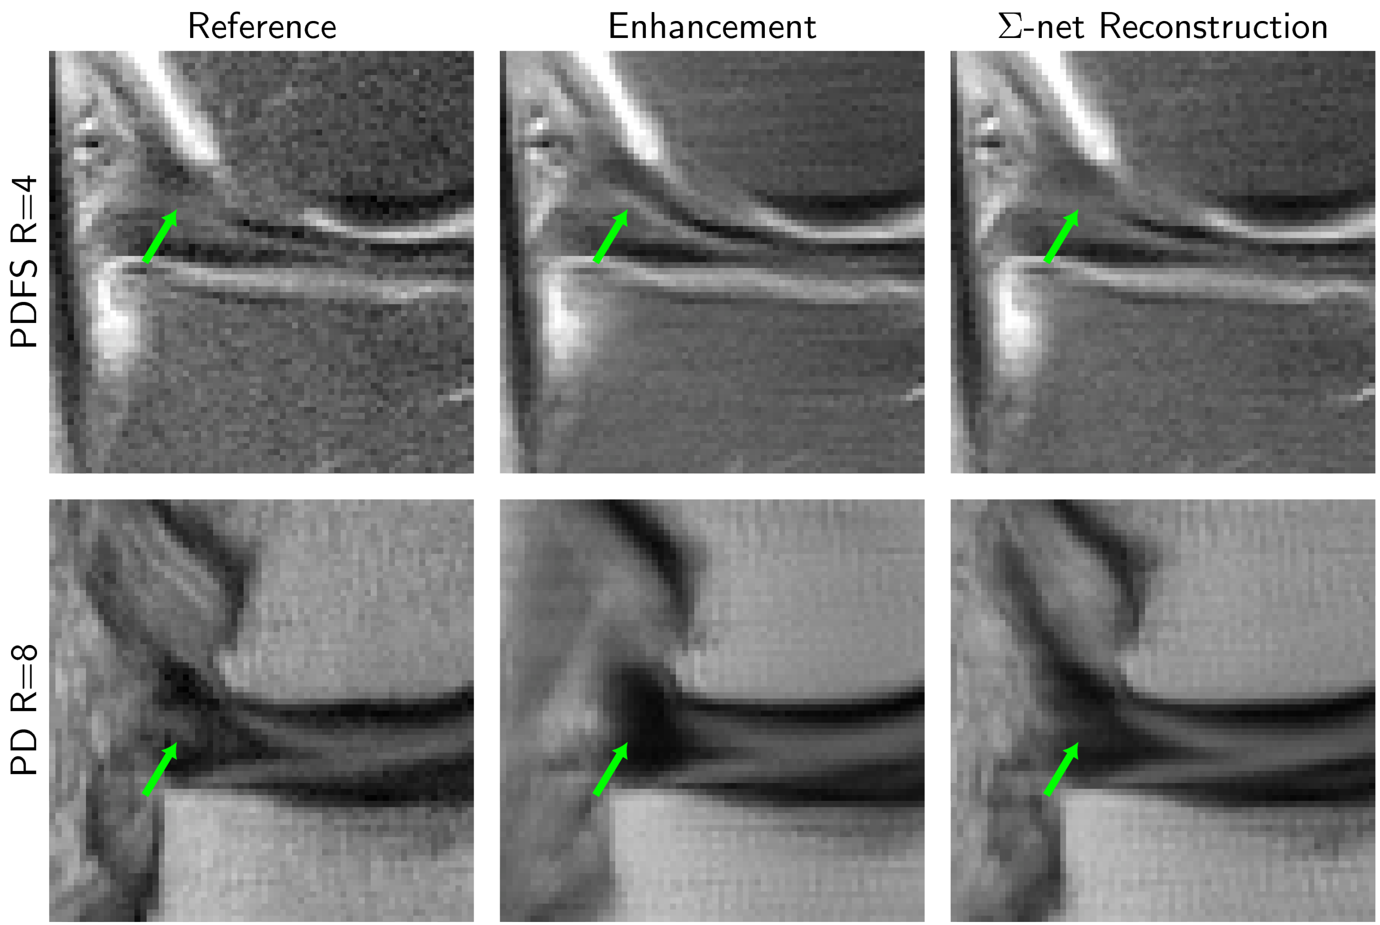

Deep learning for image super-resolution

- 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)

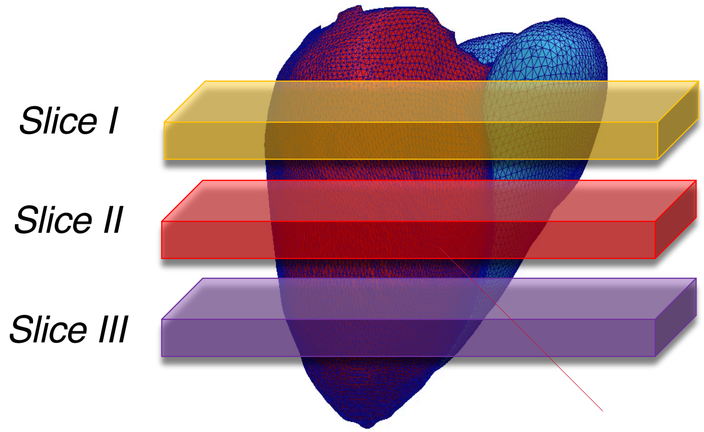

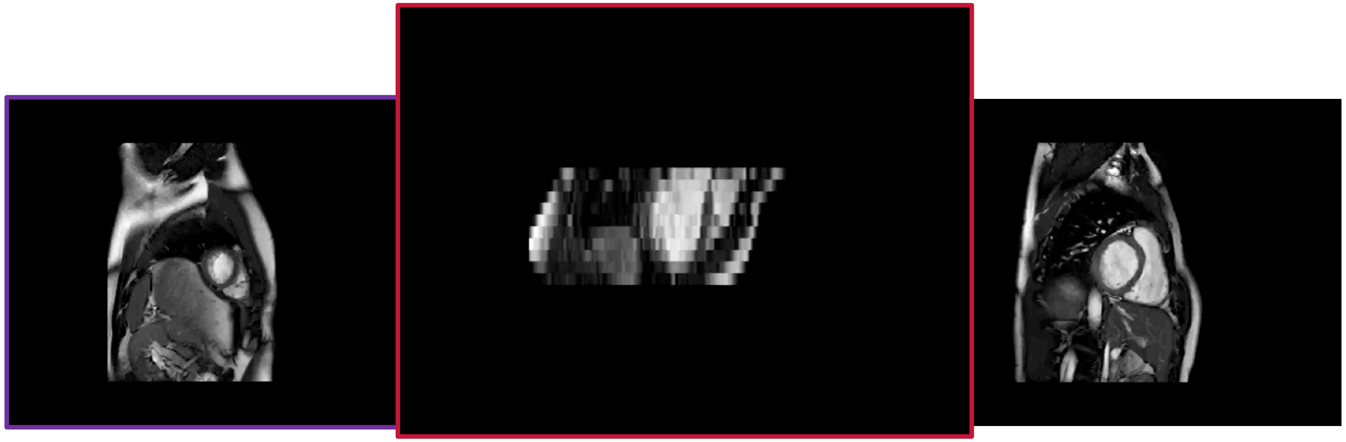

Super-resolution via image-to-image translation

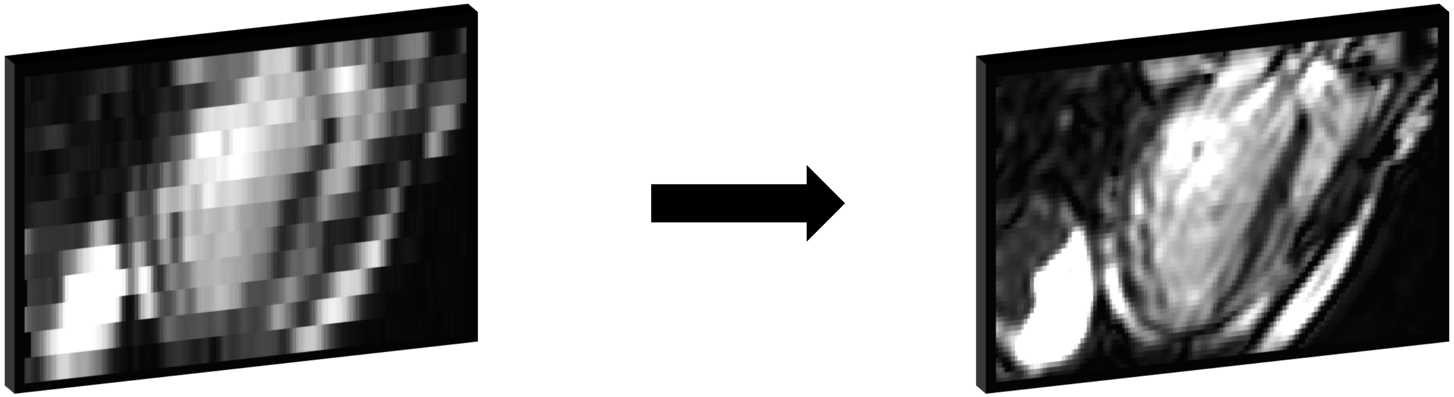

Super-resolution via image-to-image translation

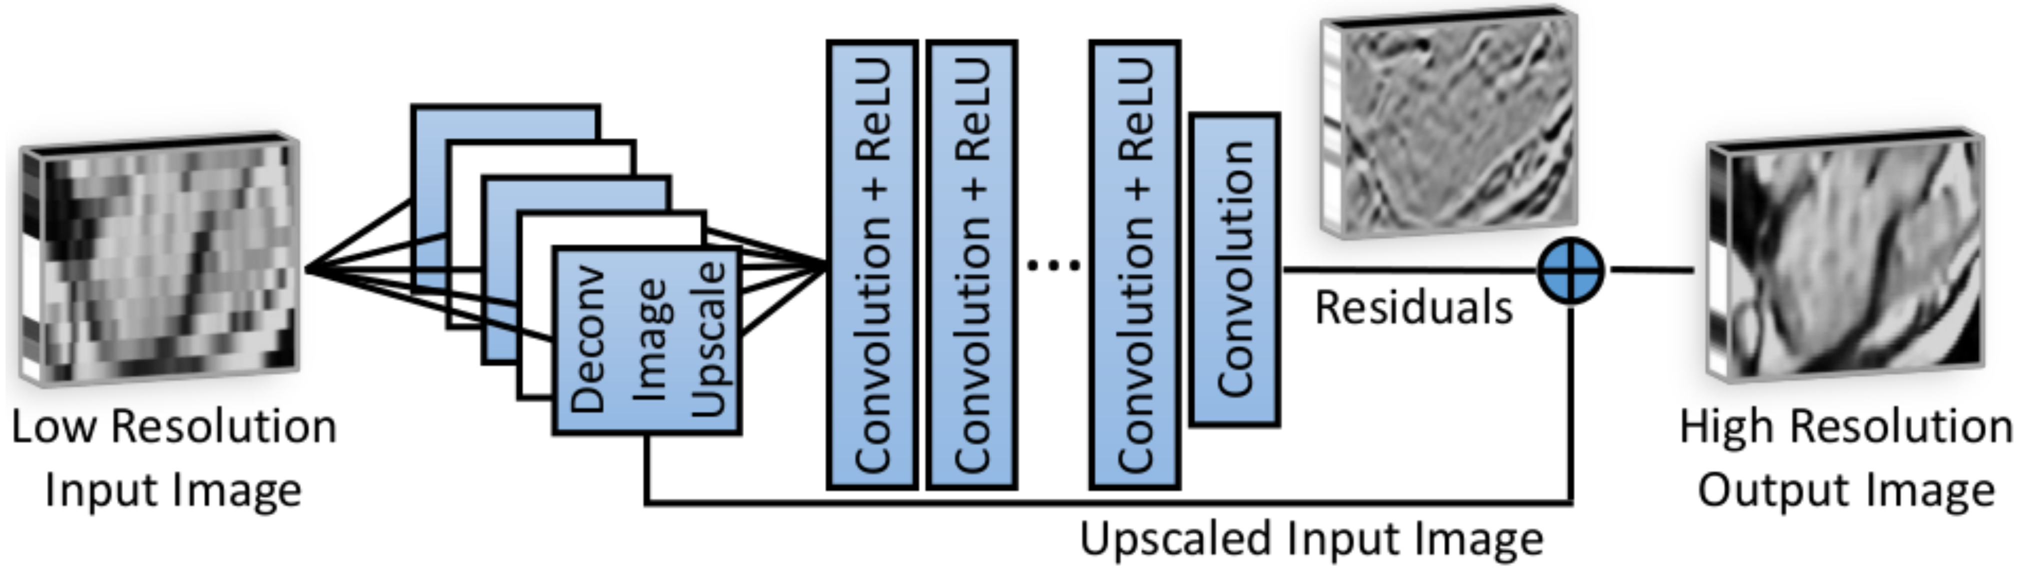

O. Oktay et al. MICCAI 2016

Super-resolution via image-to-image translation

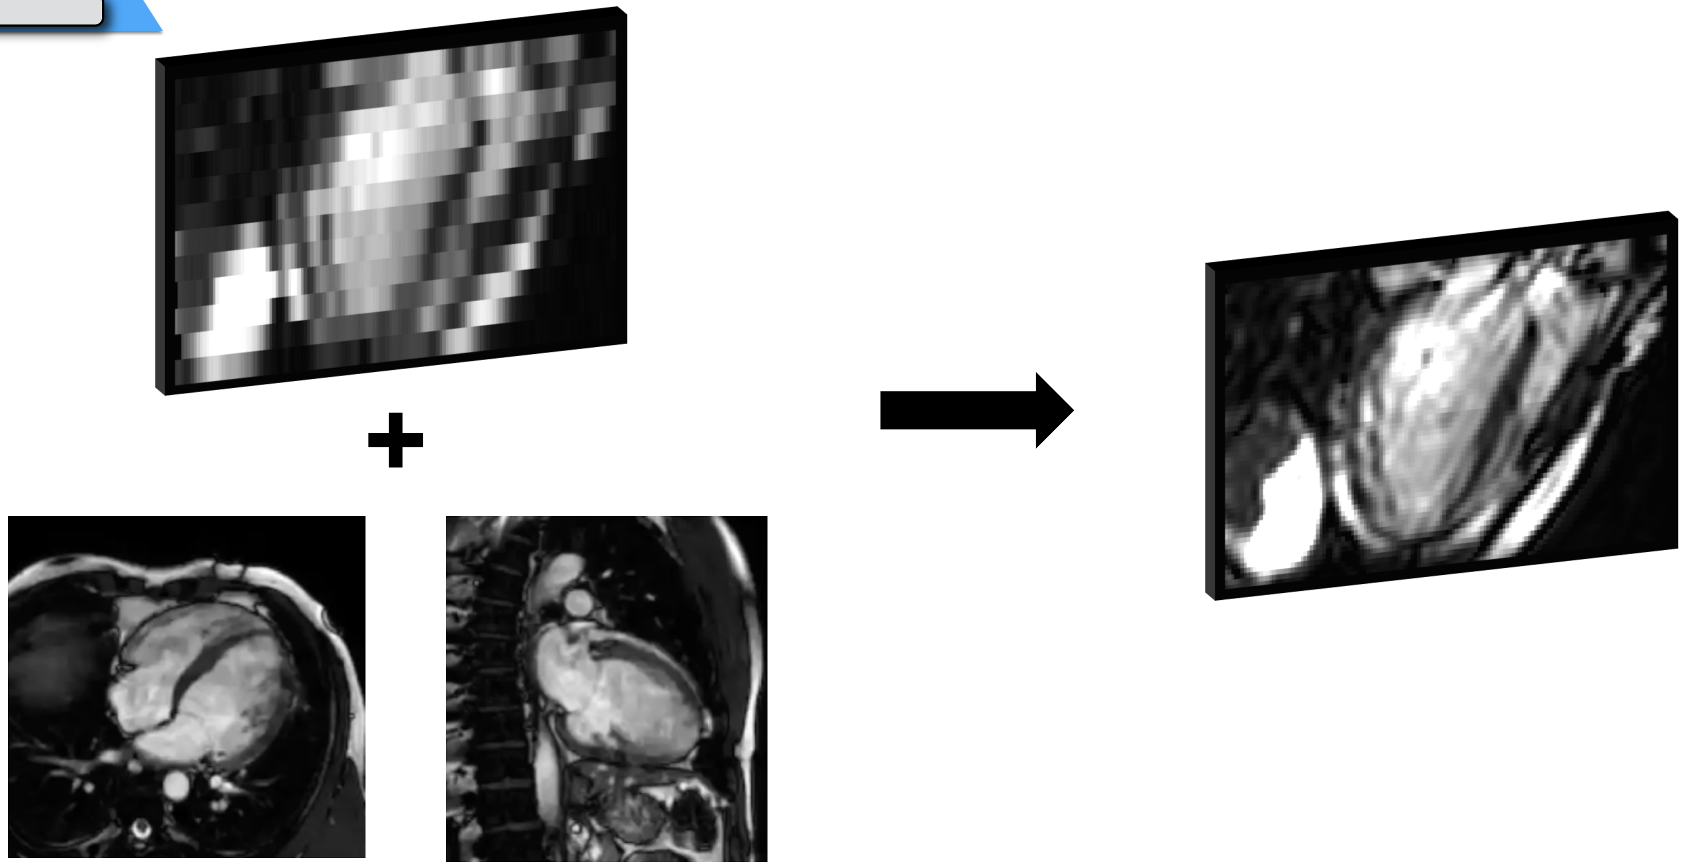

O. Oktay et al. MICCAI 2016

Deep learning for image super-resolution

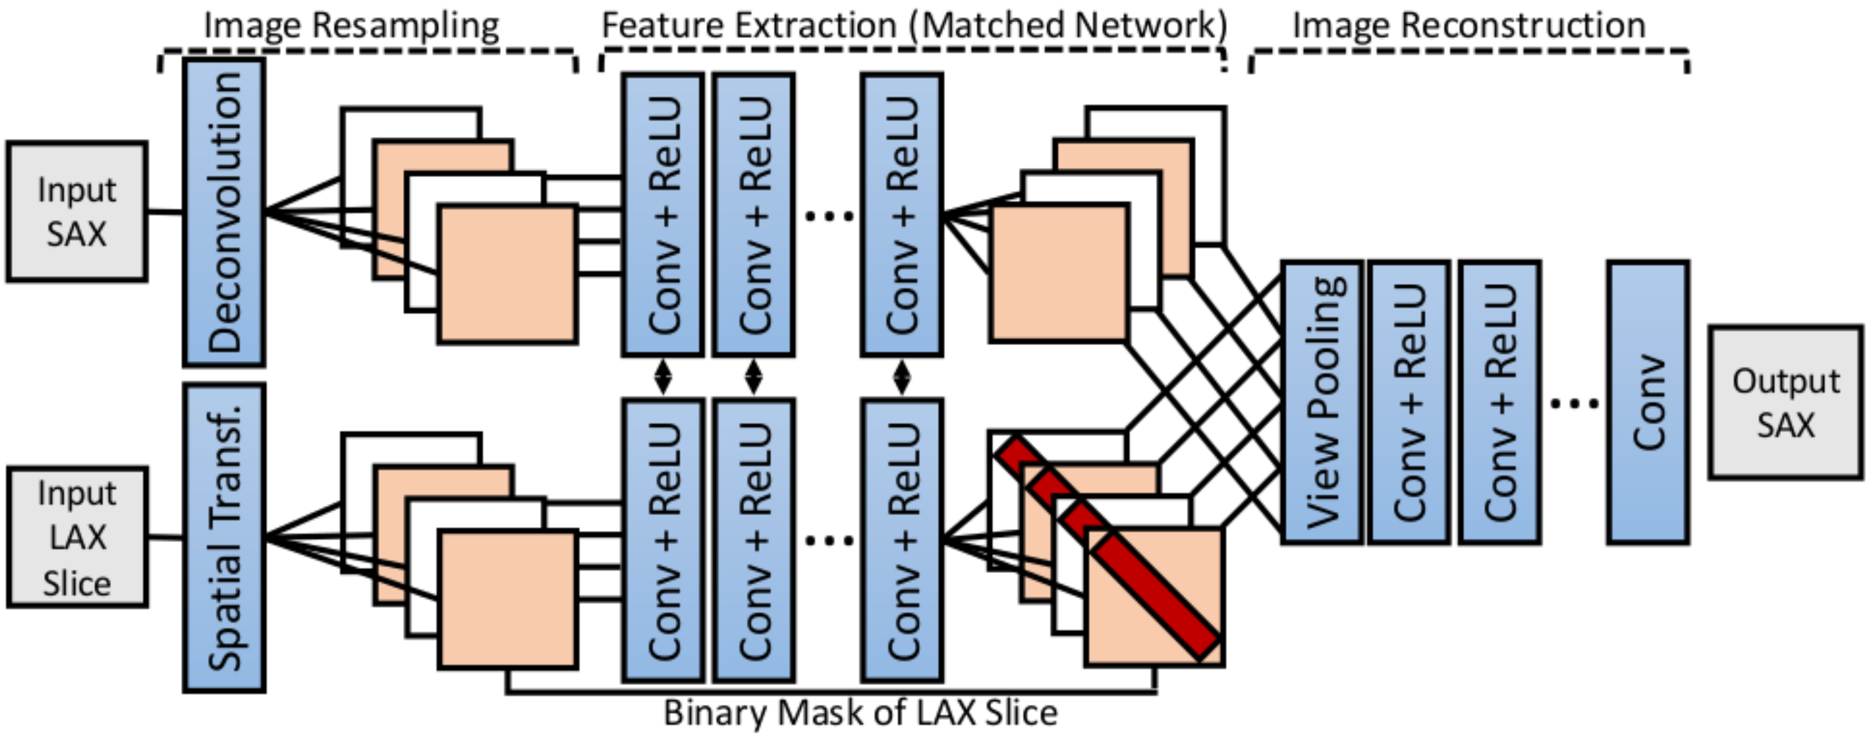

O. Oktay et al. MICCAI 2016

Deep learning for image super-resolution

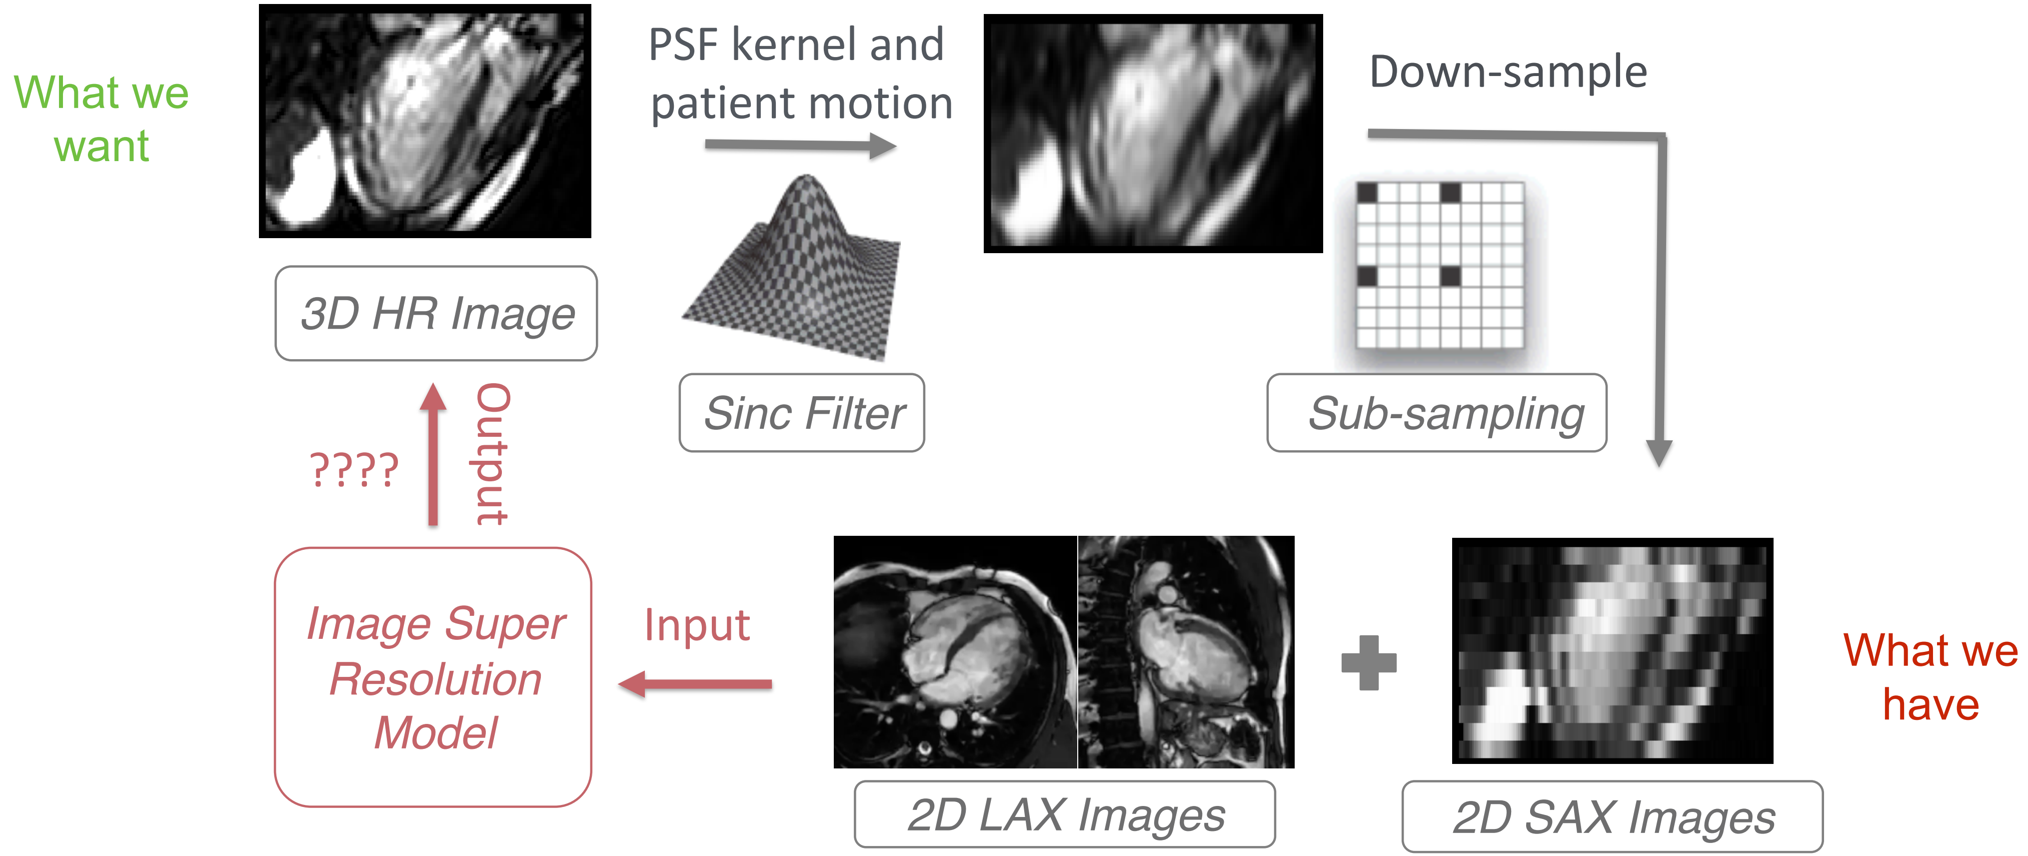

Deep learning for image super-resolution

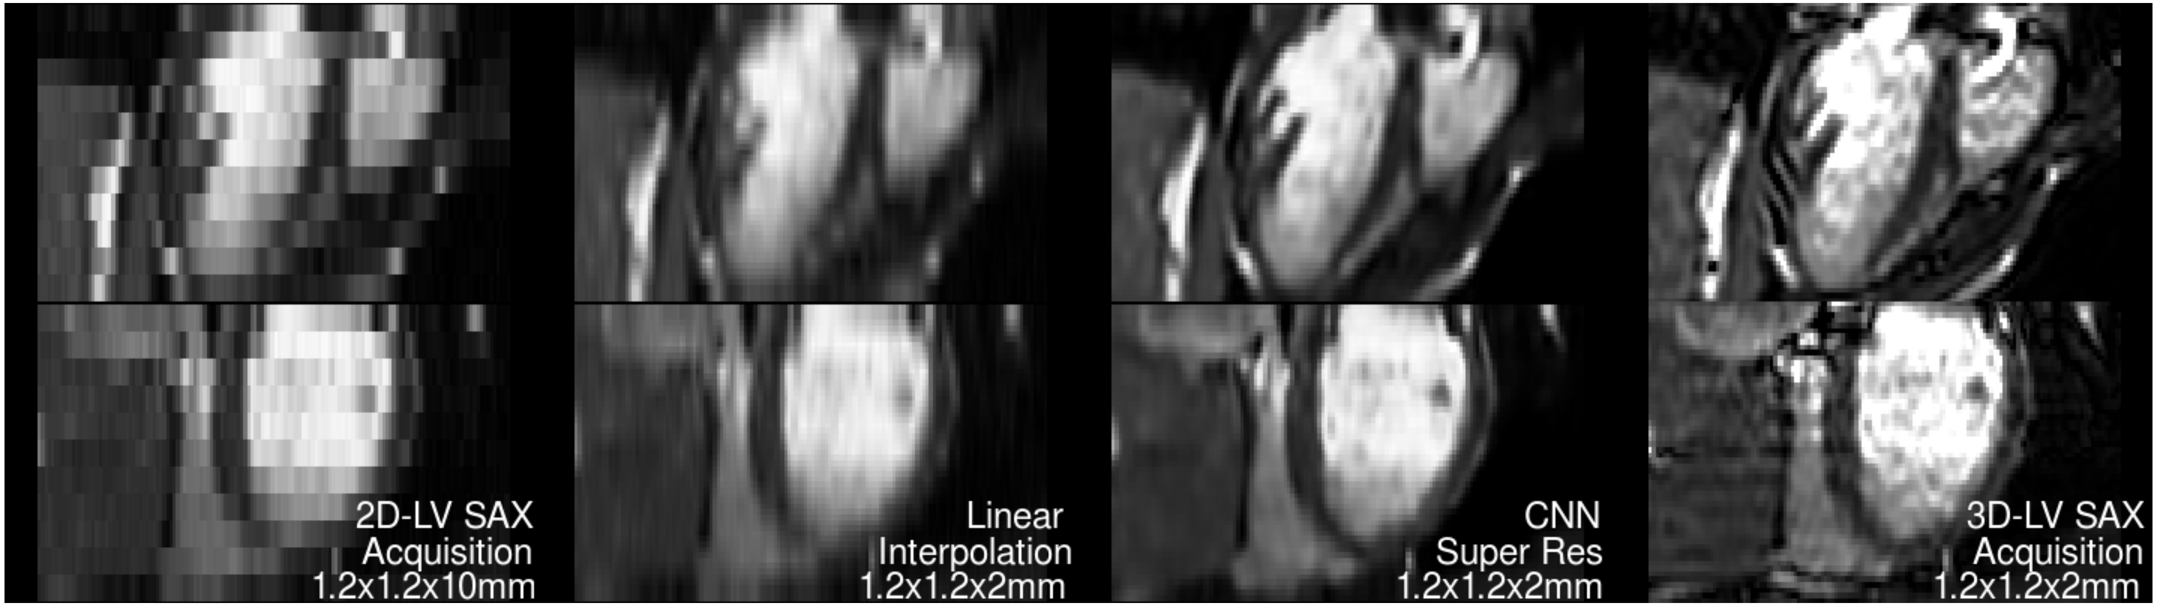

O. Oktay et al. MICCAI 2016

Deep learning for image segmentation

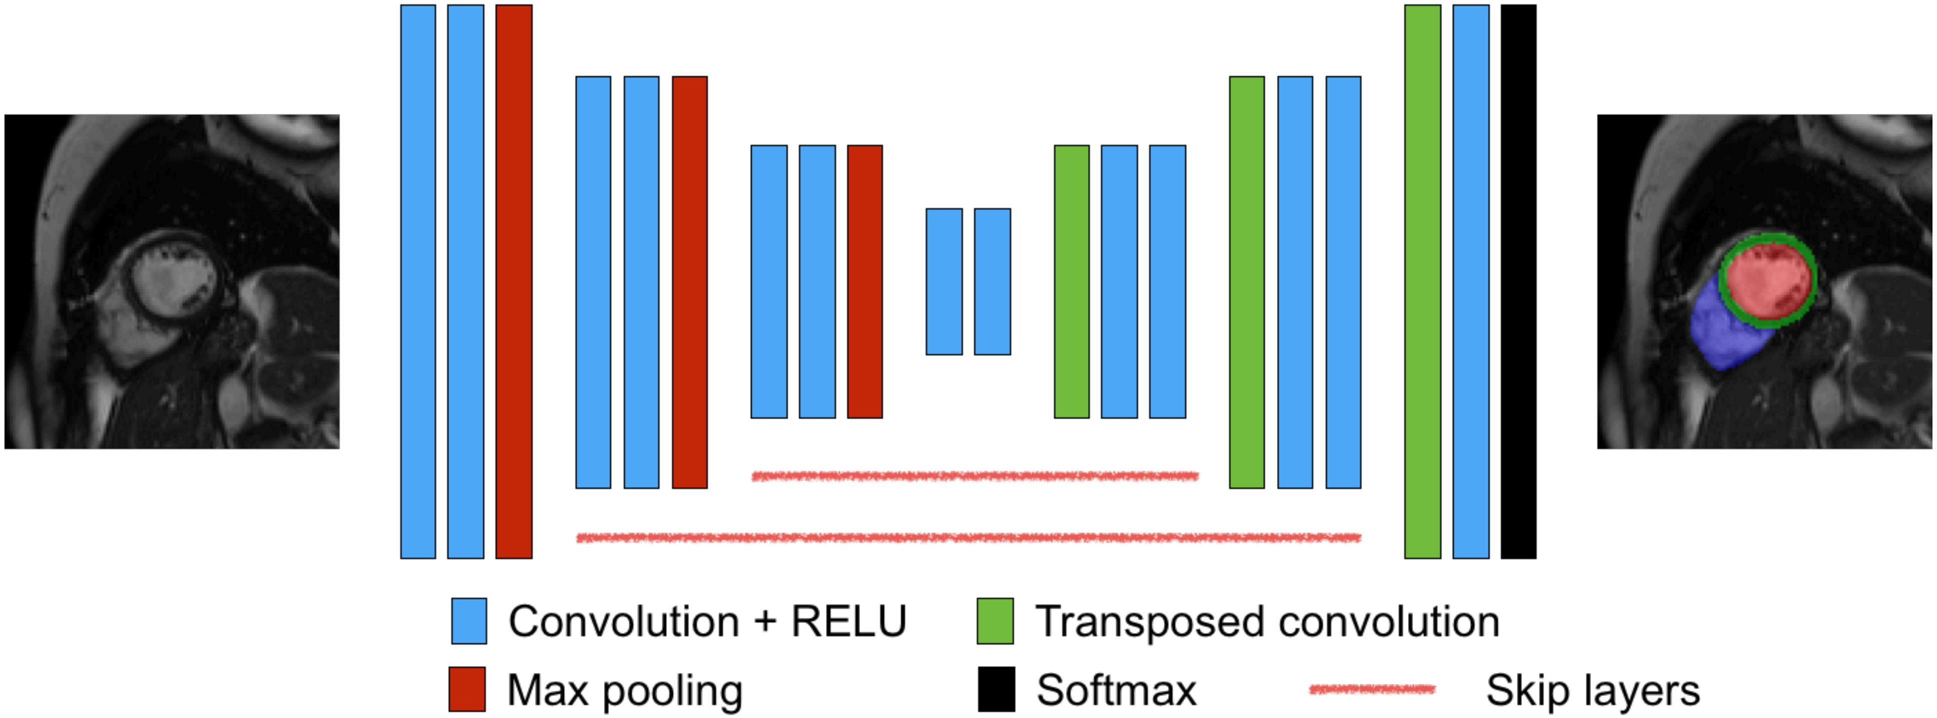

Bai et al., JCMR 2018

Deep learning for image segmentation

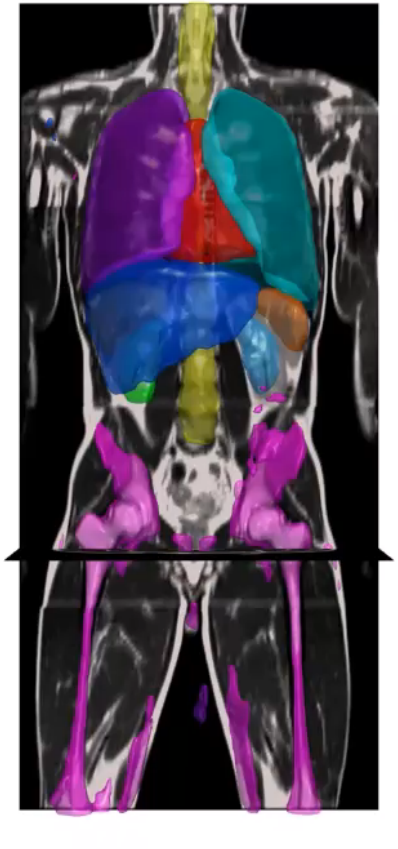

Lavdas et al. 2017,
Medical Physics

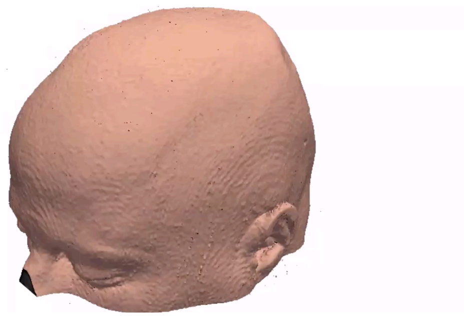

DeepMedic: K. Kamnitsas et al. Medical Image Analysis, 2017

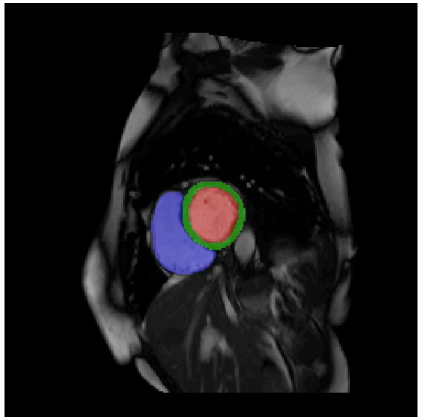

Bai et al., JCMR 2018

Deep learning for image segmentation

- 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

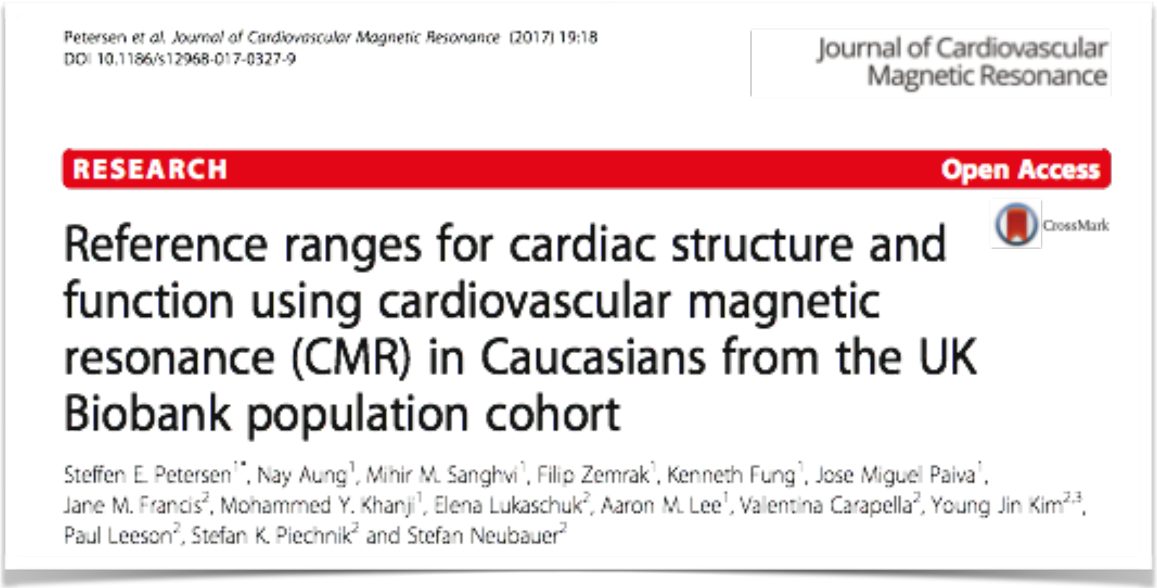

Petersen *et al. Journal of Cardiovascular Magnetic Resonance* (2017) 19:18
DOI 10.1186/s12968-017-0327-9

Journal of Cardiovascular Magnetic Resonance

RESEARCH Open Access

CrossMark

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]

Deep learning for image segmentation

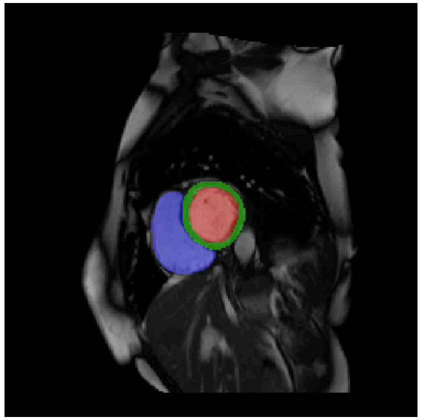

SA, basal

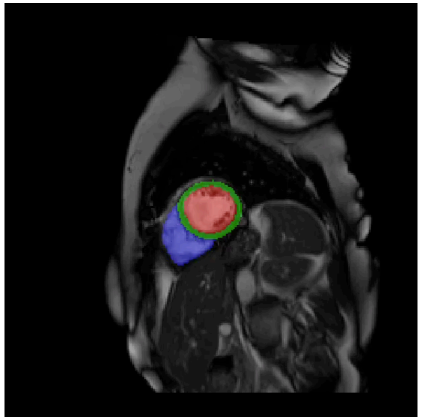

SA, mid-ventricular

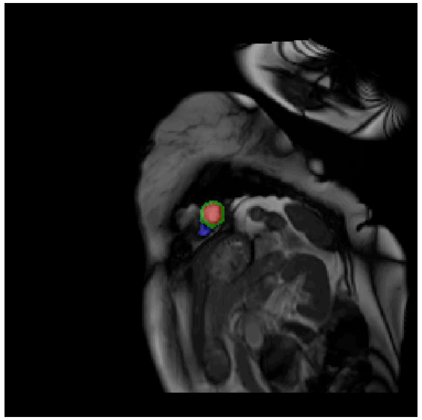

SA, apical

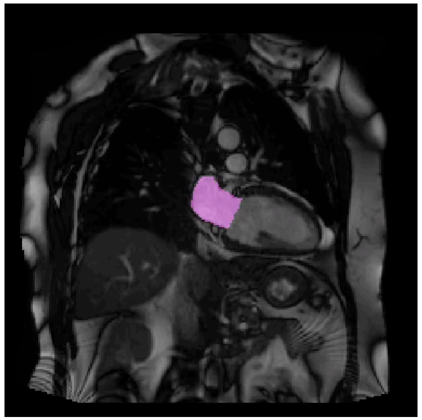

LA, 2 chamber

LA, 4 chamber

Bai et al., JCMR 2018

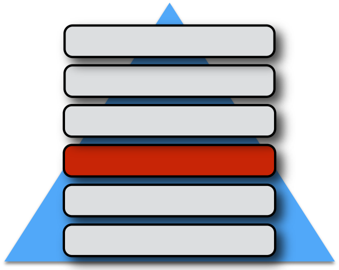

Evaluation of segmentation accuracy
Comparison to expert observers

(a) Absolute difference

	Auto vs Man (n = 600)	O1 vs O2 (n = 50)	O2 vs O3 (n = 50)	O3 vs O1 [illegible]
LVEDV (mL)	$6.1_{\pm 5.3}$	$6.1_{\pm 4.4}$	$8.8_{\pm 4.8}$	[illegible]
LVESV (mL)	$5.3_{\pm 4.9}$	$4.1_{\pm 4.2}$	[illegible]	[illegible]
LVM (gram)	$6.9_{\pm 5.5}$	$4.2_{\pm 3.2}$	[illegible]	[illegible]
RVEDV (mL)	$8.5_{\pm 7.1}$	[illegible]	[illegible]	$8.7_{\pm 5.8}$
RVESV (mL)	$7.2_{\pm 6.8}$	[illegible]	[illegible]	$11.7_{\pm 6.9}$

[illegible]	[illegible]	[illegible] vs O2 (n = 50)	O2 vs O3 (n = 50)	O3 vs O1 (n = 50)
LV[illegible]	[illegible]	$4.2_{\pm 3.1}$	$6.3_{\pm 3.3}$	$3.4_{\pm 2.2}$
[illegible]	$9.5_{\pm 9.5}$	$6.8_{\pm 7.5}$	$12.5_{\pm 8.5}$	$11.7_{\pm 5.1}$
[illegible]	$8.3_{\pm 7.6}$	$4.4_{\pm 3.3}$	$6.0_{\pm 3.7}$	$6.7_{\pm 4.6}$
[illegible]V (%)	$5.6_{\pm 4.6}$	$8.0_{\pm 5.0}$	$4.2_{\pm 3.1}$	$5.7_{\pm 3.6}$
RVESV (%)	$11.8_{\pm 12.2}$	$30.6_{\pm 15.5}$	$10.9_{\pm 8.3}$	$16.9_{\pm 9.2}$

Automated | Manual

Computer performs as well as different expert observers

W. Bai et al., JCMR 2018

Deep learning for joint image segmentation and motion tracking: Motion-Seg Net

Segmentation

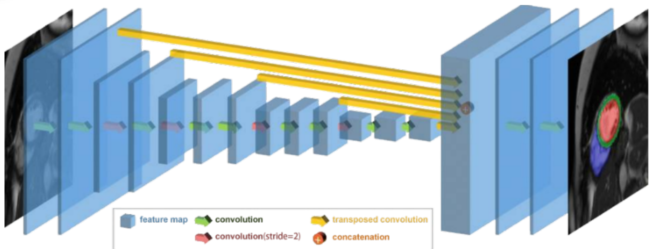

[W. Bai et al, 2018]

Complementary

Motion Estimation

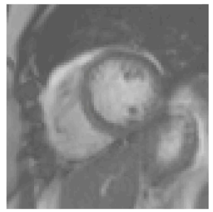

Challenges:

- Segmentation: only ED/ES frames are annotated
- Motion field: no ground truth are available

Beyond supervised learning:
Self-supervised motion estimation +
weakly-supervised segmentation

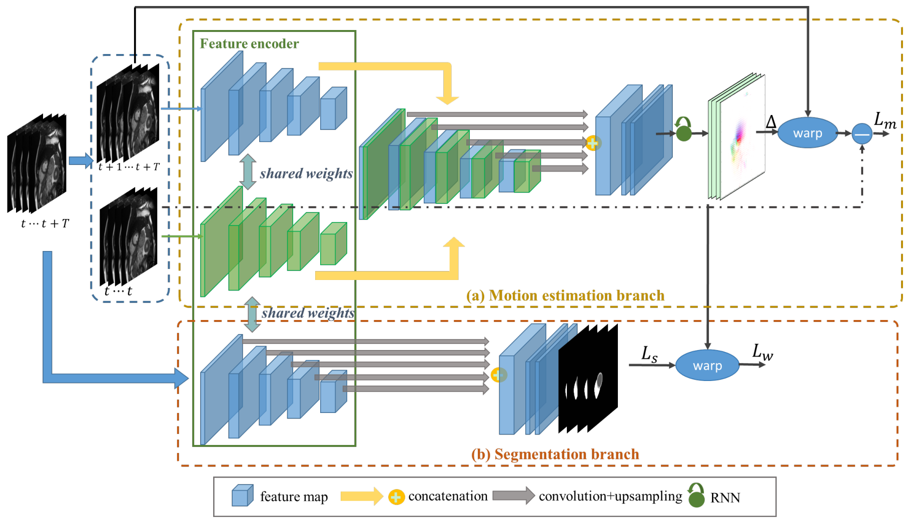

Qin et al. MICCAI, 2018

Deep learning for joint image segmentation and motion tracking: Motion-Seg Net

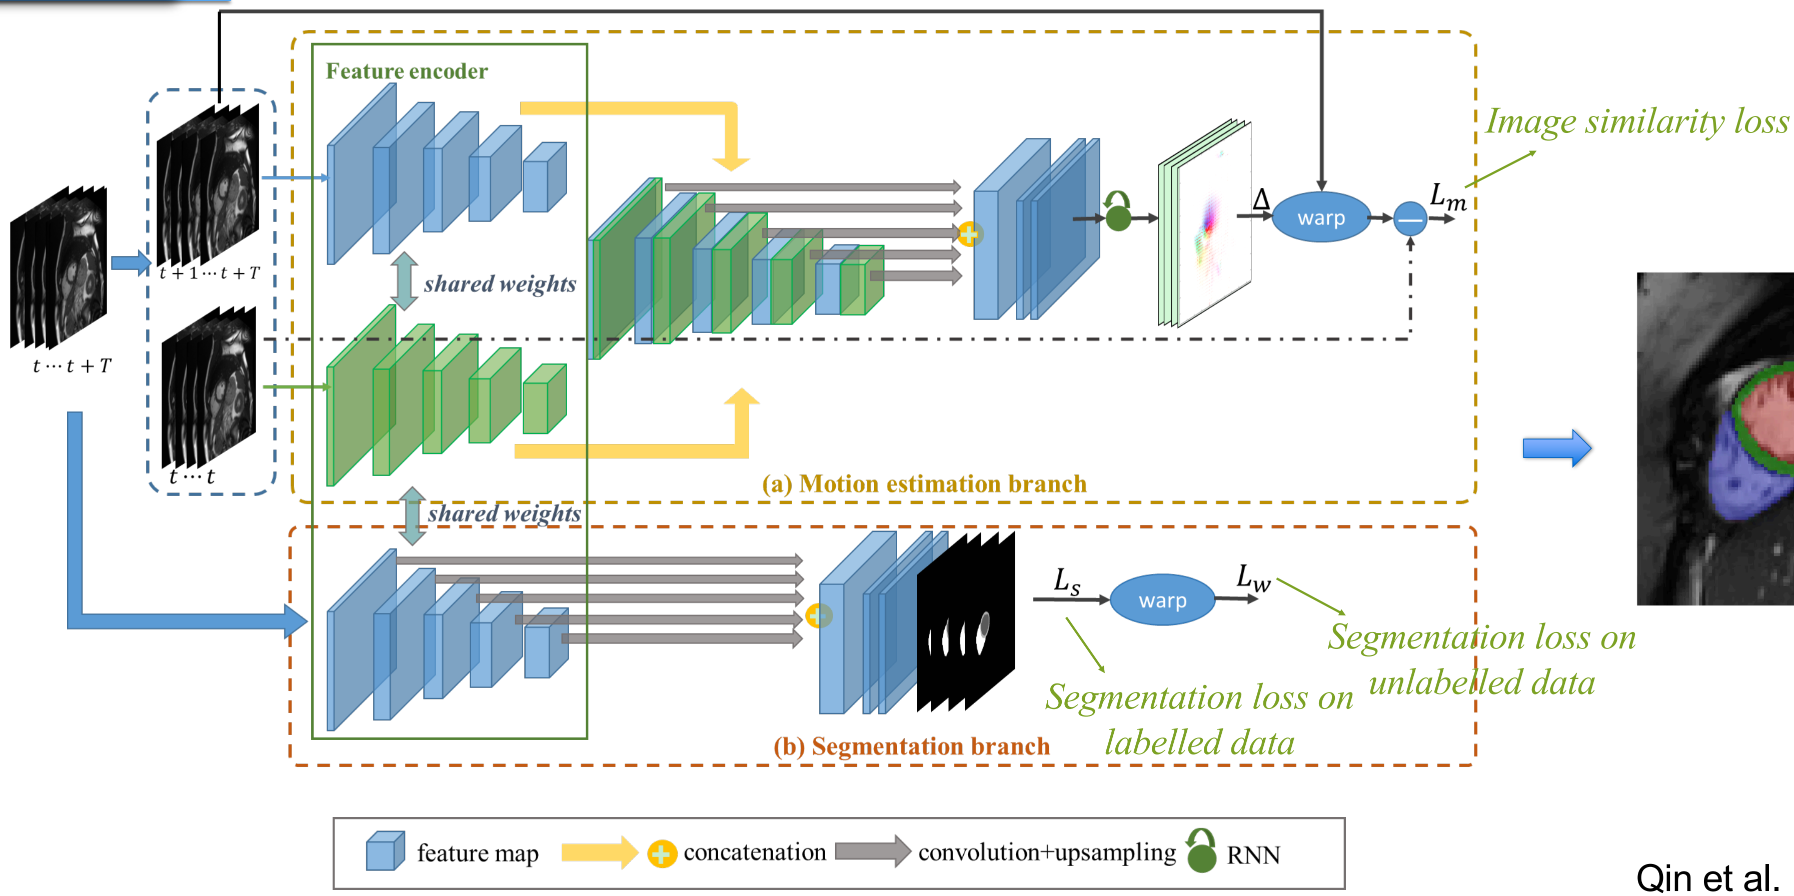

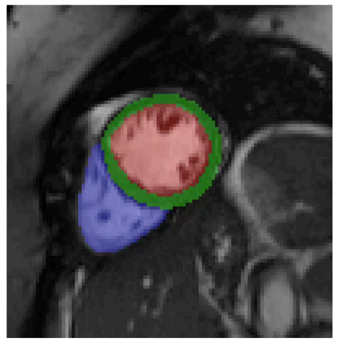

Qin et al. MICCAI, 2018

Deep learning for joint image segmentation and motion tracking: Motion-Seg Net

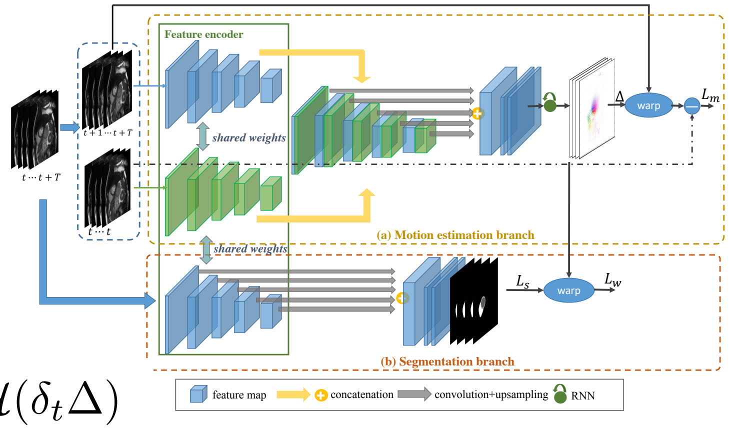

- Loss function:

$$\mathcal{L} = \mathcal{L}_m + \lambda_1 \mathcal{L}_s + \lambda_2 \mathcal{L}_w,$$

- Motion estimation loss:

$$\mathcal{L}_m = \frac{1}{T} \sum_{k=1}^{T} [\| I_t - I'_{t+k} \|^2 + \alpha \mathcal{H}(\delta_{x,y} \Delta_{t+k})] + \beta \mathcal{H}(\delta_t \Delta)$$

- Segmentation loss on labeled data:

$$\mathcal{L}_s = -\sum_{l \in L} y_l \log(f(x_l; \Theta))$$

- Warped segmentation loss on unlabeled data:

$$\mathcal{L}_w = -\sum_{n \in U} y_l \log(f_w(x_n; \Theta))$$

Qin et al., MICCAI 2018

Cardiac MR image analysis: Is the problem solved?

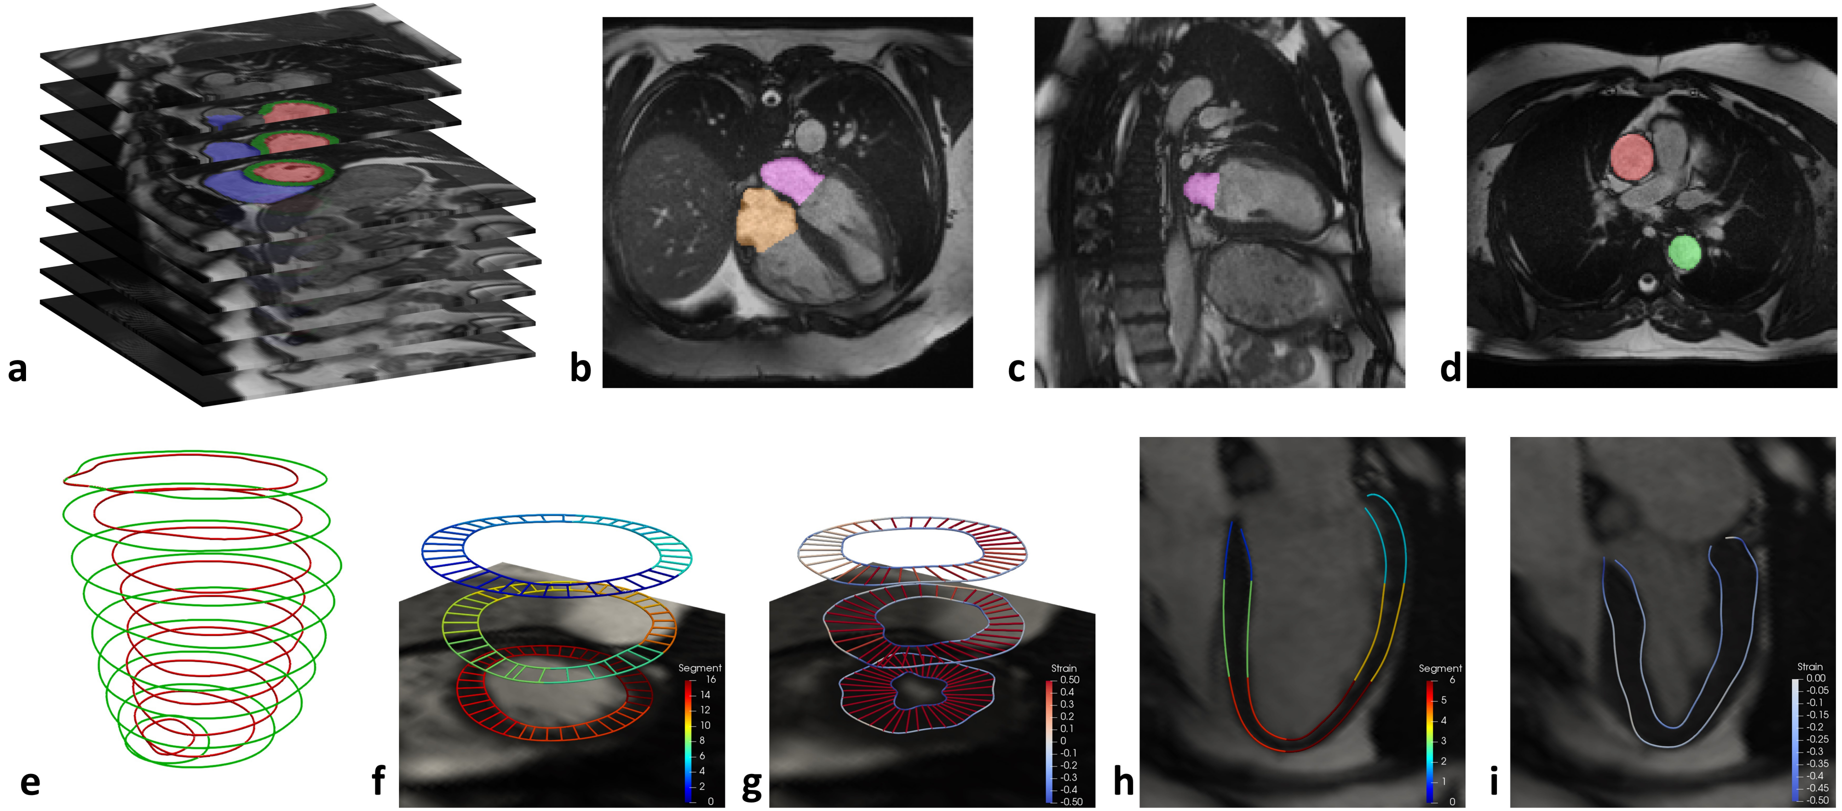

Cardiac MR image analysis: Is the problem solved?

- 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)
- Motion between slices can lead to artefacts

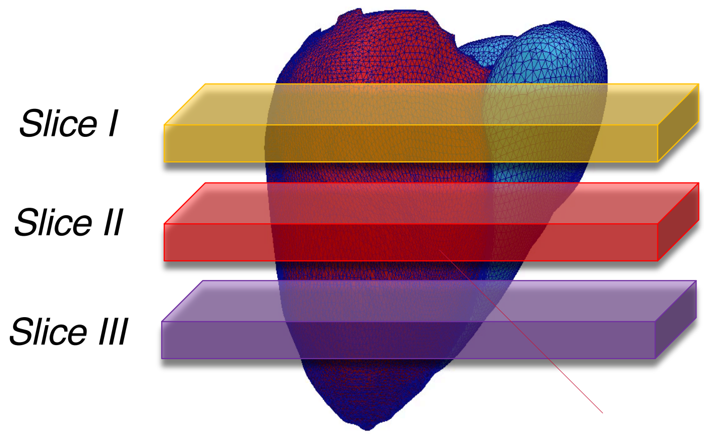

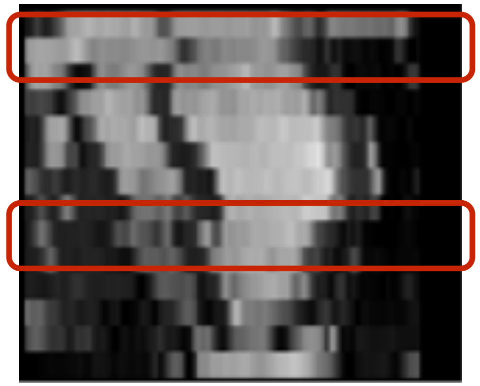

Cardiac MR image analysis:
Is the problem solved?

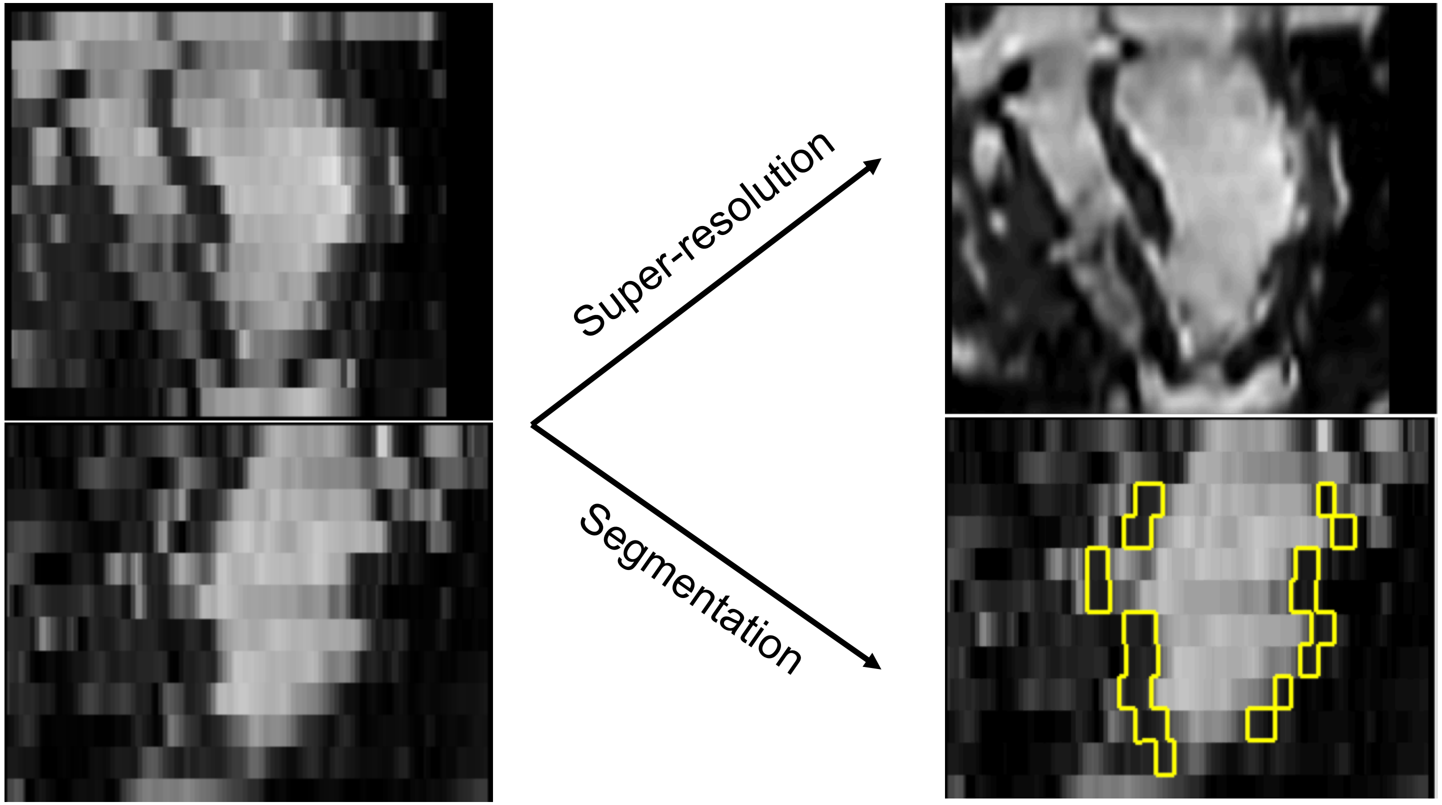

Conventional CNNs: What we want

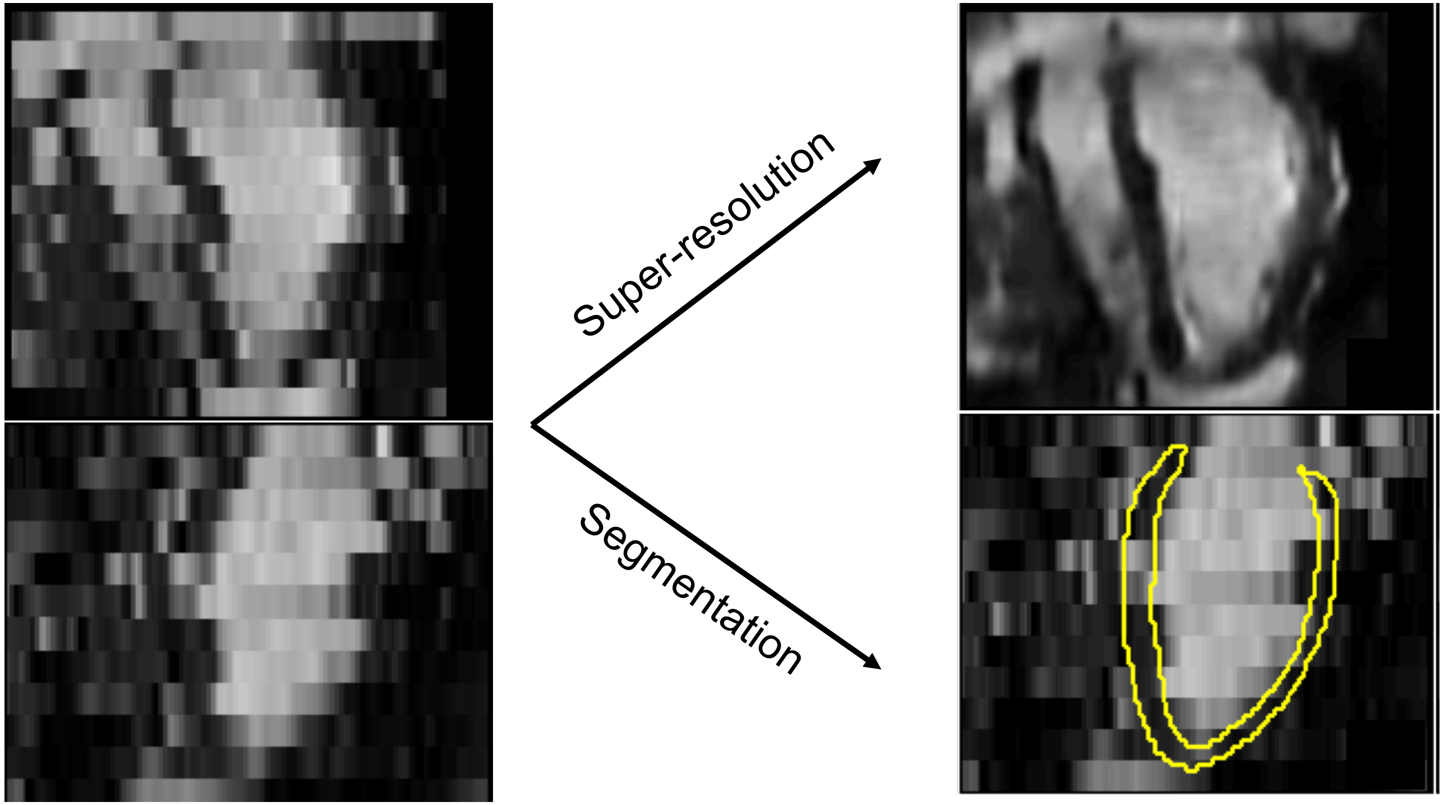

CNNs: No explicit use of prior knowledge

- Standard Loss for **segmentation**: Cross-Entropy loss

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

- Standard loss for **super-resolution**: L2 or L1 loss

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

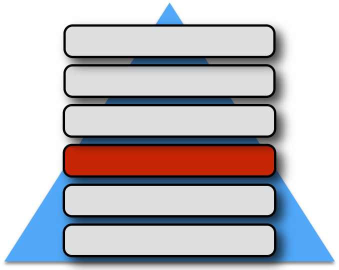

Anatomically constrained CNNs

Low-resolution input

High-resolution output

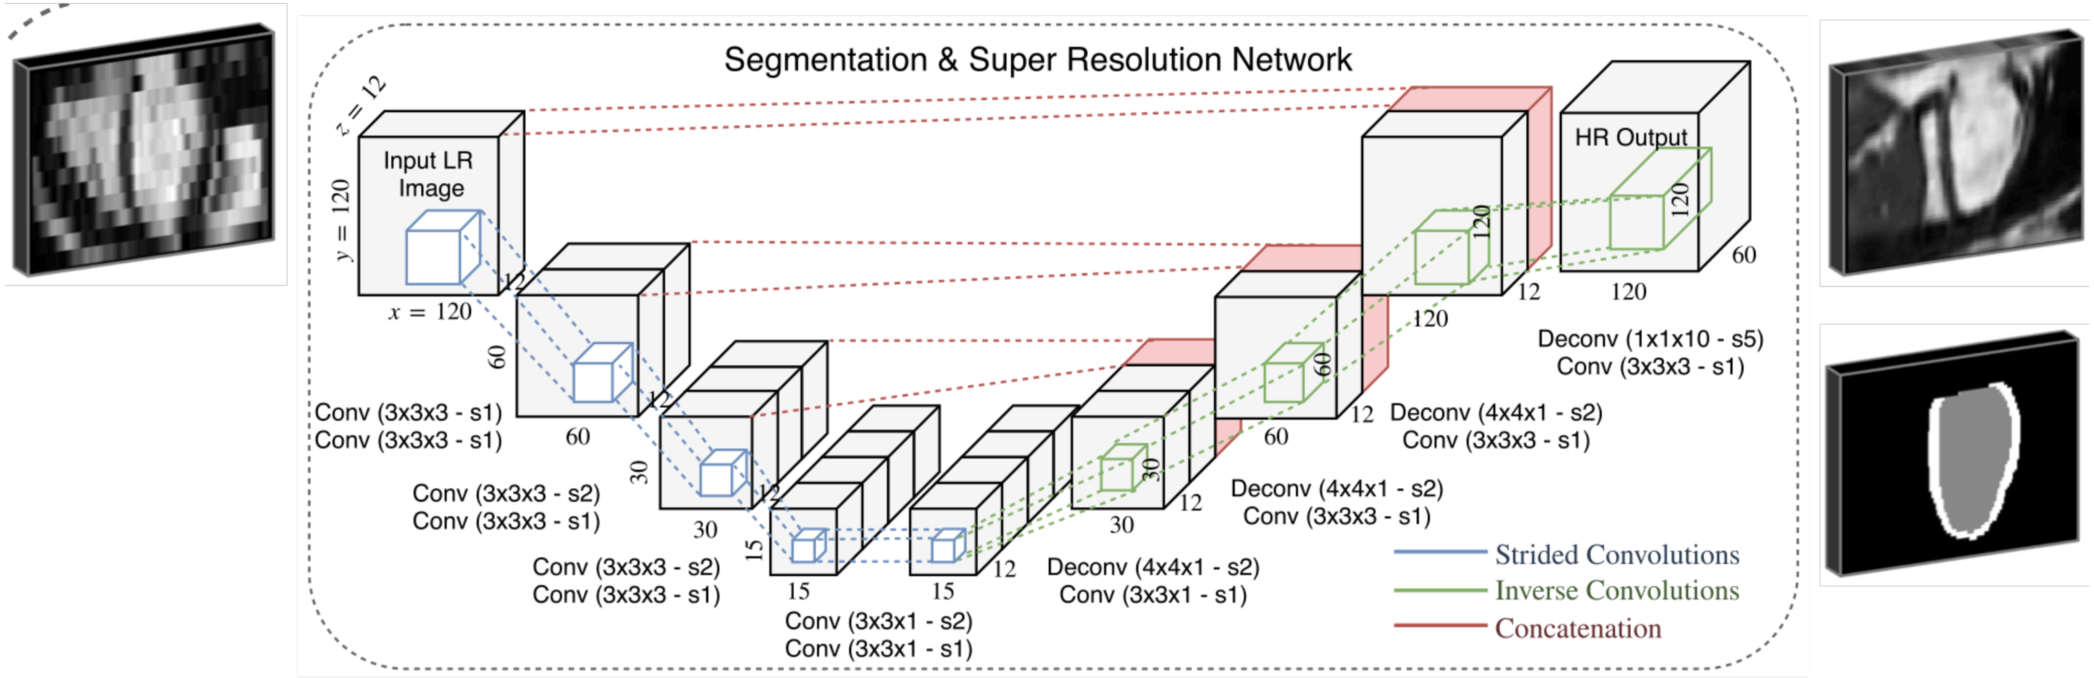

O. Oktay et al. IEEE TMI 2017

Anatomically constrained CNN: T-L networks for representing priors

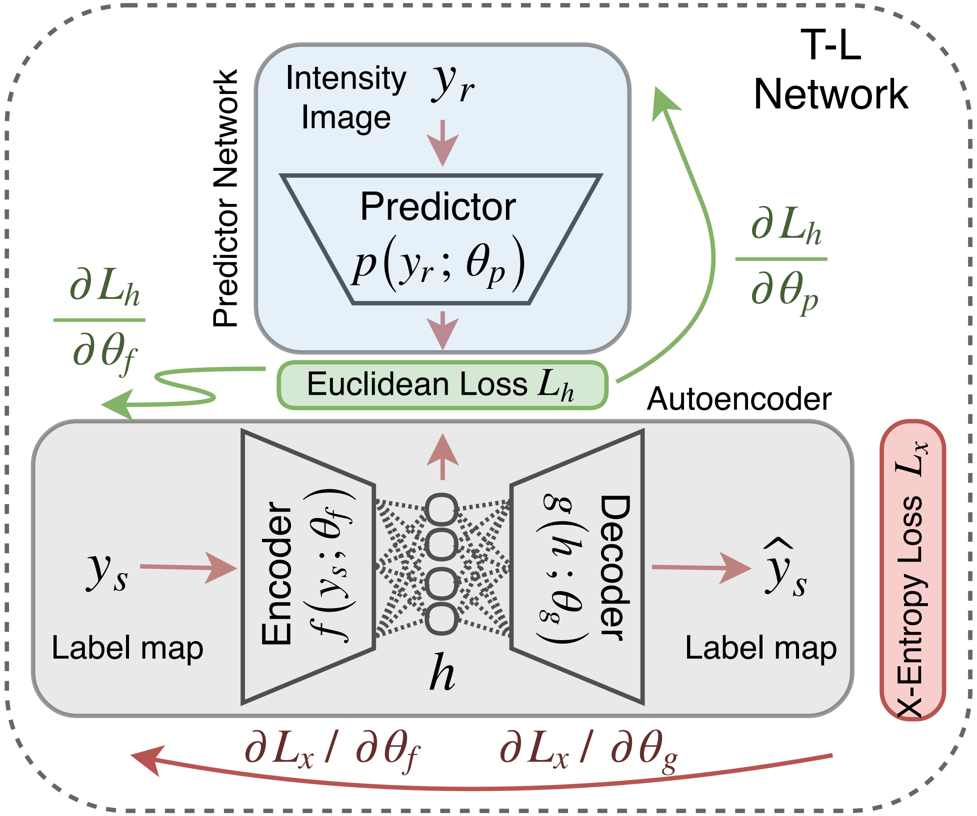

O. Oktay et al. IEEE TMI 2017

Anatomically constrained CNN: Segmentation framework

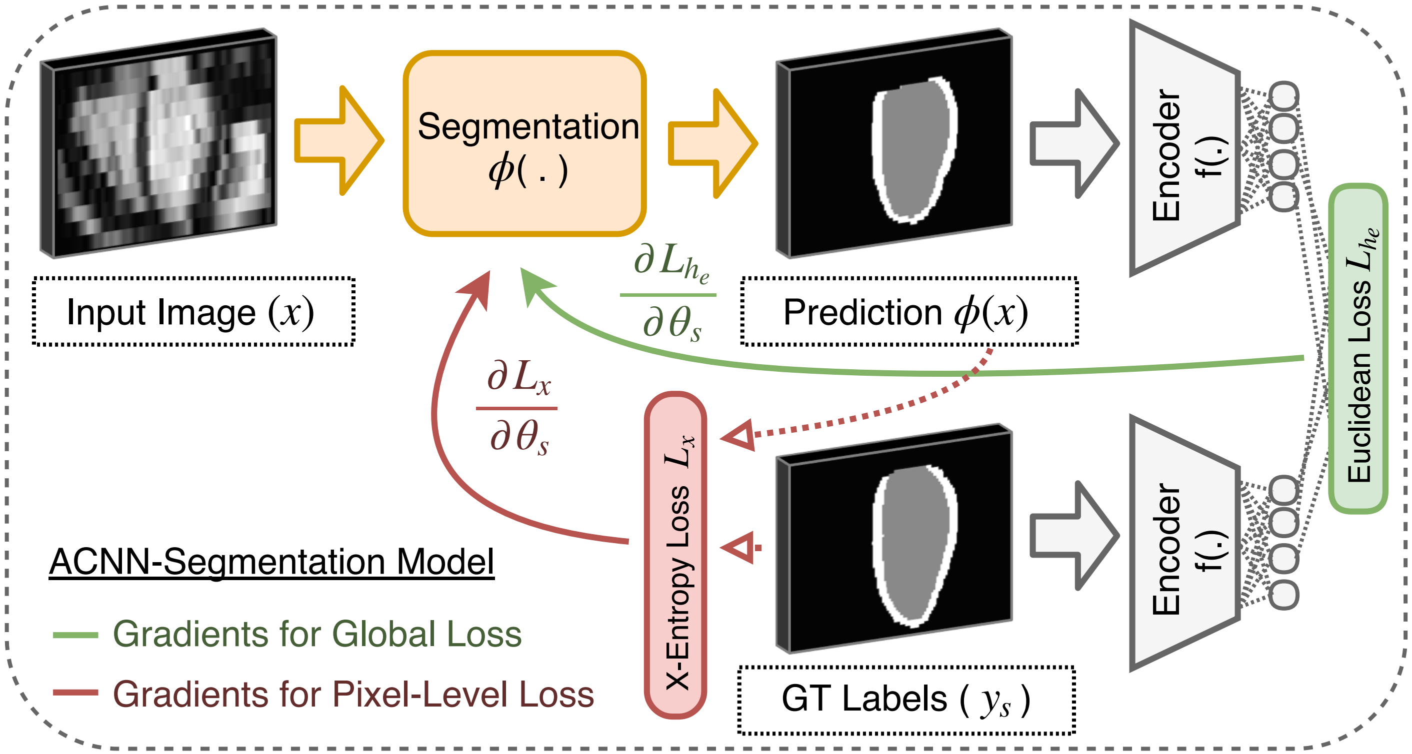

O. Oktay et al. IEEE TMI 2017

Anatomically constrained CNN: Segmentation results

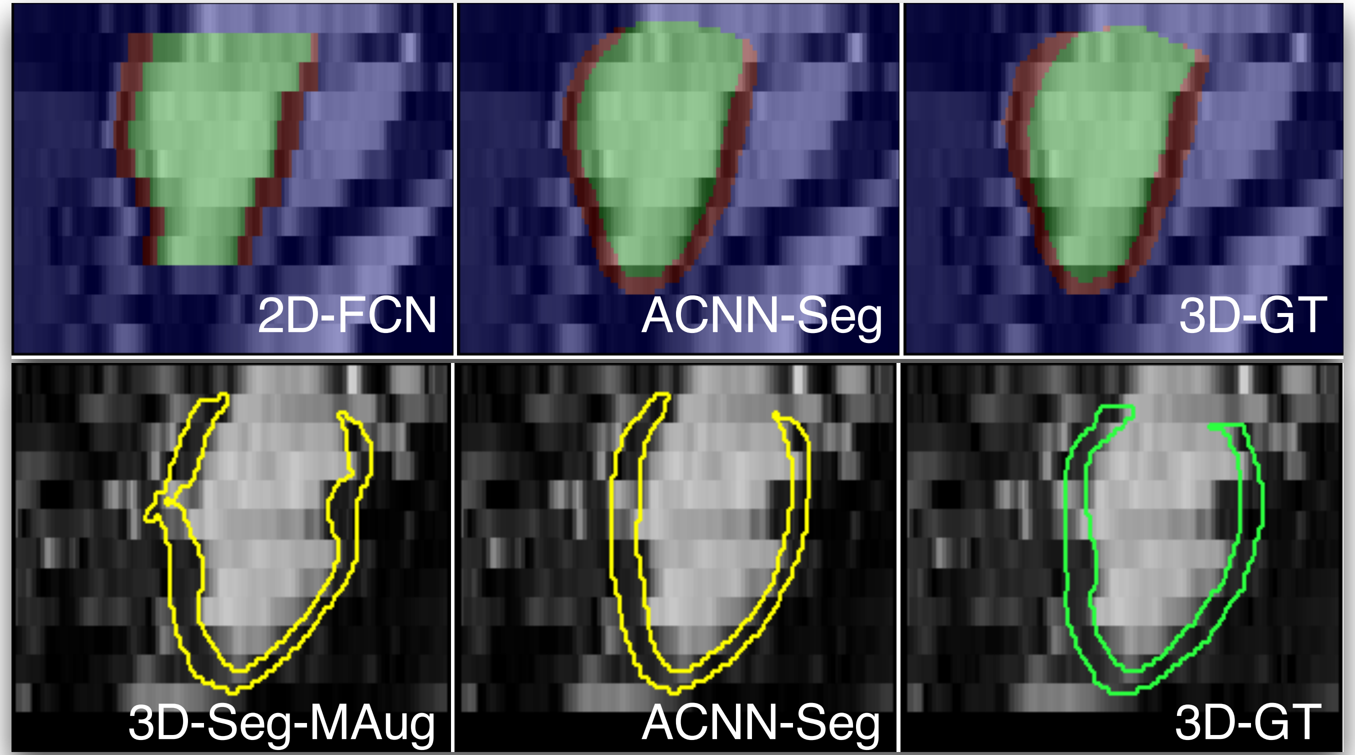

O. Oktay et al. IEEE TMI 2017

Anatomically constrained CNN: Super-resolution framework

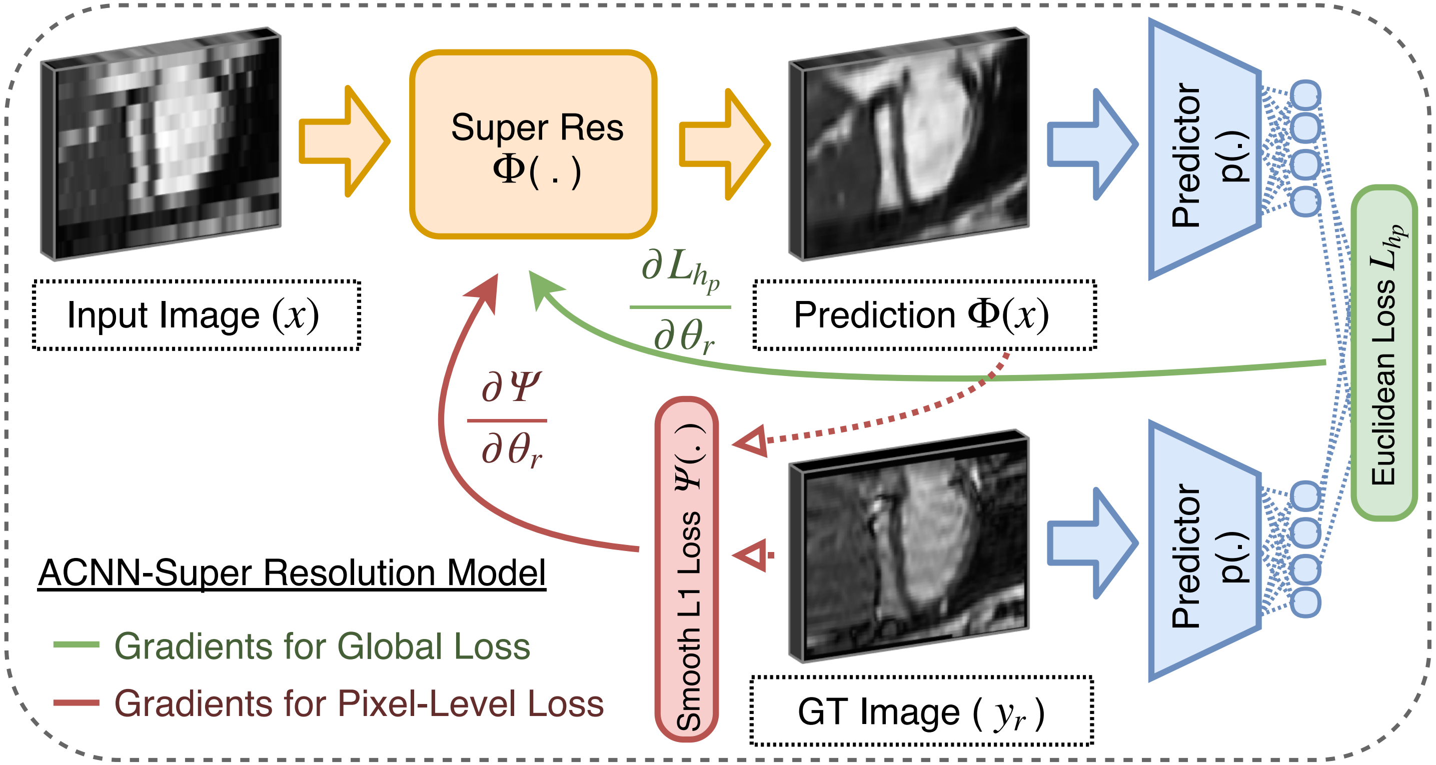

O. Oktay et al. IEEE TMI 2017

Anatomically constrained CNN: Super-resolution results

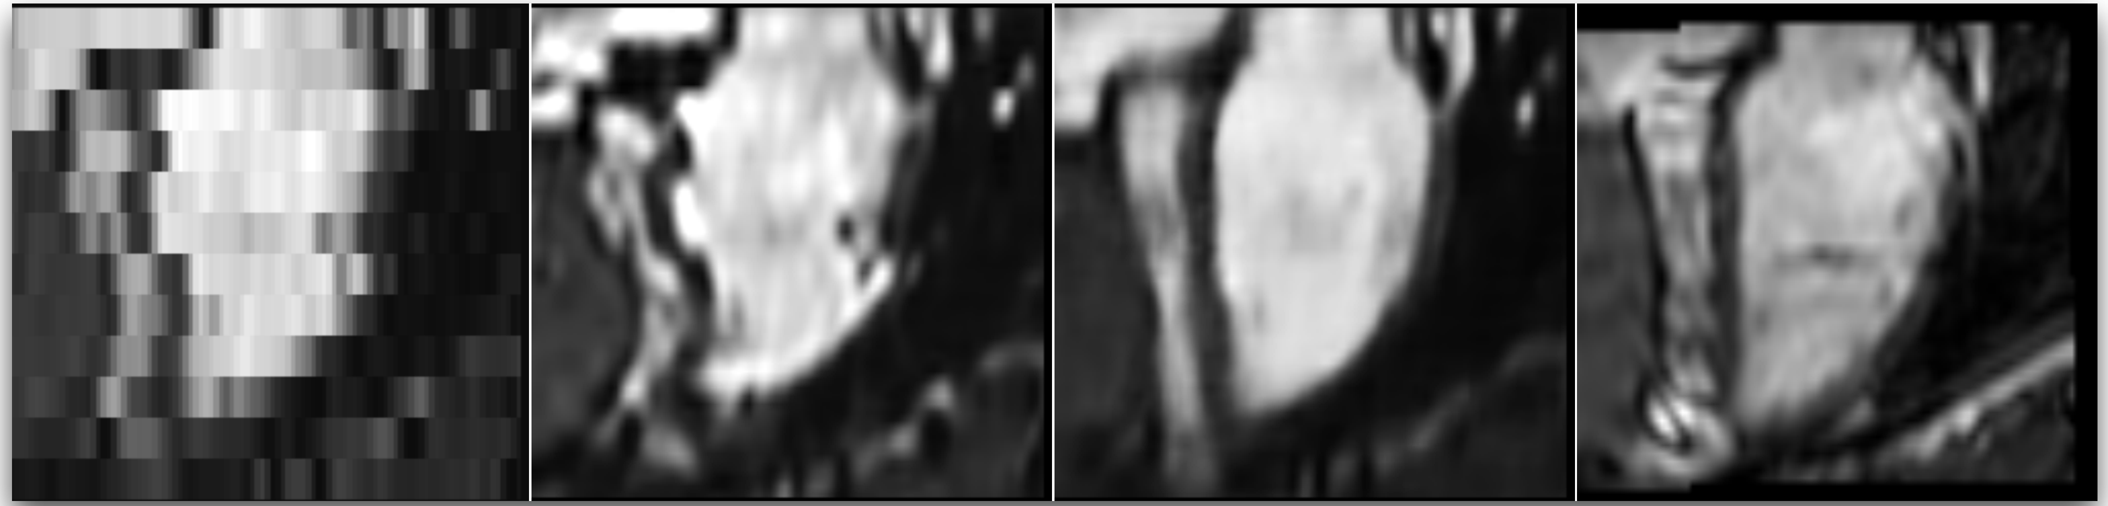

Original LR image

Baseline SR approach

Anatomically constrained SR model

Ground-truth HR image

O. Oktay et al. IEEE TMI 2017

Traditional medical imaging

Acquisition

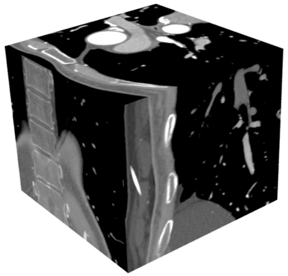

Reconstruction

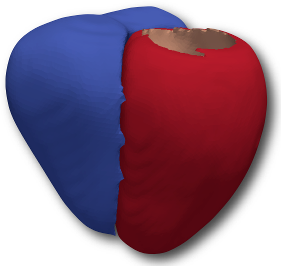

Analysis

Interpretation

- ✗ 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

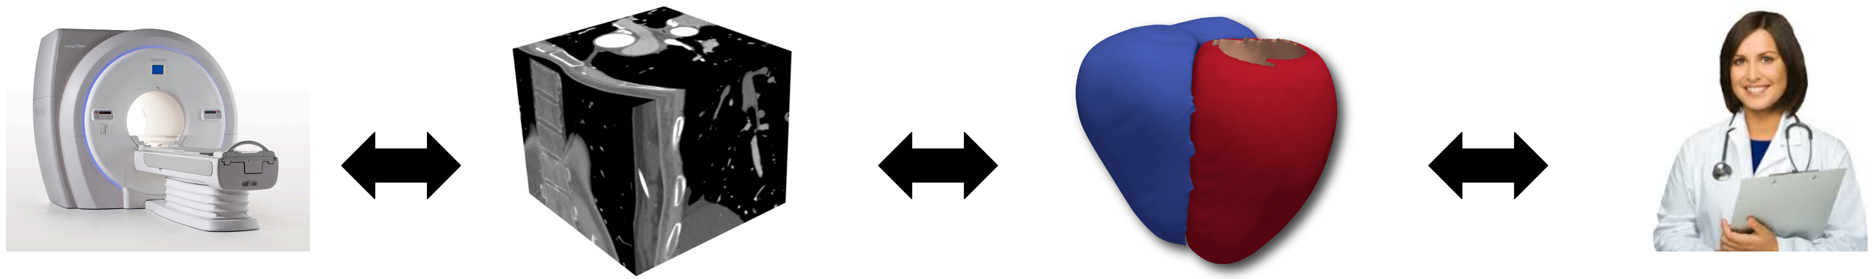

Acquisition | Reconstruction | Analysis | Interpretation

- ✓ 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

Acquisition

?

Diagnosis

Do we need images at all?

AI-enabled medical imaging: Example

Ground truth

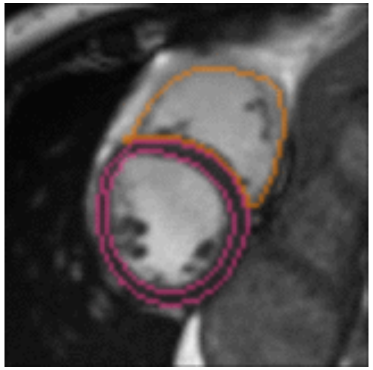

J. Schlemper et al.
MICCAI 2018

Acknowledgements

Ghalib Bello
Tim Dawes
Antonio de Marvao
Declan O'Regan
Stuart Cook
Hideaki Suzuki
Paul M. Matthews
Aaron M. Lee
Nay Aung
Elena Lukaschuk
Mihir M. Sanghvi
Filip Zemrak
Kenneth Fung
Jose Miguel Paiva
Valentina Carapella
Young Jin Kim
Steffen E. Petersen

Stefan K. Piechnik
Stefan Neubauer
Aurelien Bustin
Giulia Ginami
Gastao Cruz
Teresa Correia
Tevfik F. Ismail
Imran Rashid
Radhouene Neji
Claudia Prieto
Rene M. Botnar
Jo Hajnal
Alberto Gomez
Veronika Zimmer
Jacqueline Matthew
David Lloyd
Reza Razavi

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