AI for CMR standardization - Technique Talk II:

Deep Learning Image Reconstruction: Hope or Hype?

Martin Uecker

1. Institute of Biomedical Imaging, Graz University of Technology
2. Institute for Diagnostic and Interventional Radiology
University Medical Center Göttingen
3. German Center for Cardiovascular Research (DZHK), Partner Site Göttingen

Declaration of Interest

Co-inventor of patent for Real-time MRI.

Deep Learning in CMR

1. Segmentation ✓
2. Automatic diagnosis ?
3. Image reconstruction !

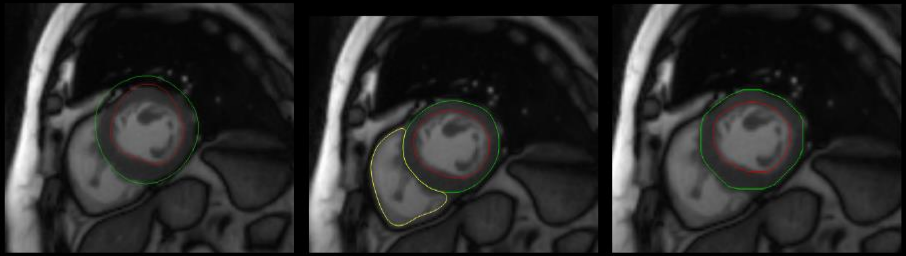

Segmentation of real-time CMR[1] images: conventional, commercial AI, own AI[2]

We should stop training radiologists now. It's just completely obvious that within five years, deep learning is going to do better than radiologists.

– Geoffrey Hinton, godfather of AI, 2016[3]

1. Uecker M et al. NMR Biomed 23:986-994 (2010)
2. Schilling M et al. ESMRMB 2020, Magma 33:69-233 (2020)
3. Machine Learning and Market for Intelligence Conference in Toronto, 2016
`https://www.youtube.com/watch?v=2HMPRXstSvQ`

Deep Learning in CMR

1. Segmentation ✓
2. Automatic diagnosis ?
3. Image reconstruction !

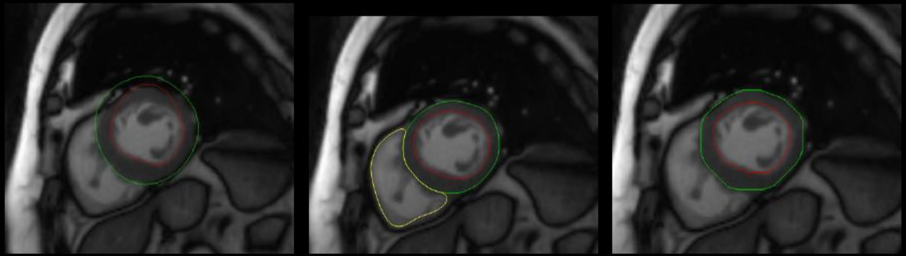

Segmentation of real-time CMR[1] images: conventional, commercial AI, own AI[2]

We should stop training radiologists now. It's just completely obvious that within five years, deep learning is going to do better than radiologists.

– Geoffrey Hinton, godfather of AI, 2016[3]

1. Uecker M et al. NMR Biomed 23:986-994 (2010)
2. Schilling M et al. ESMRMB 2020, Magma 33:69-233 (2020)
3. Machine Learning and Market for Intelligence Conference in Toronto, 2016
`https://www.youtube.com/watch?v=2HMPRXstSvQ`

Artificial Neural Networks

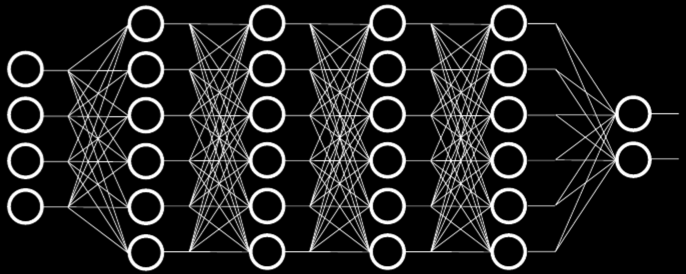

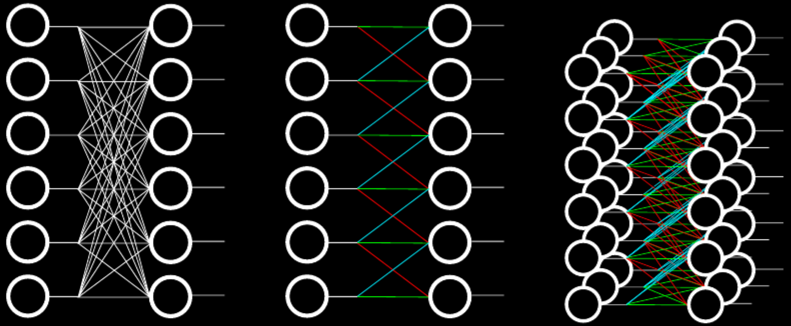

A deep neural network is a chain of multi-variate vector-valued functions $f^l(w^l, p)$ which depend on weights w^l and input p

$$h_w(\vec{y}^n) = f^N(w^N, f^{N-1}(w^{N-1}, \cdots f^1(w^1, y)))$$

Deep Learning

Now: Training of large and deep artificial neural networks

Training data sets: $(\vec{x}^n, \vec{y}^n), n = 1, \cdots, N$
Feature vector $\vec{x}$ and continuous labels $\vec{y}$

Loss function (e.g. least-squares): $E(w) = \sum_n \|\vec{y}^n - h_w(\vec{x}^n)\|_2^2$

Inference: Application of h to new data $x = h(y)$

All the impressive achievements of deep learning amount to just curve fitting.
– Judea Pearl, also a godfather of AI, 2018[1]

1. https://www.quantamagazine.org/to-build-truly-intelligent-machines-teach-them-cause-and-effect-20180515/

Learning the Reconstruction

- Learn the complete reconstruction[1]
- Training data sets (x^n, y^n) of k-space and images
- For example: least-squares loss

$$E(w) = \sum_n \|\vec{x}^n - h_w(\vec{y}^n)\|_2^2$$

- Image reconstruction: $x = h_w(y)$

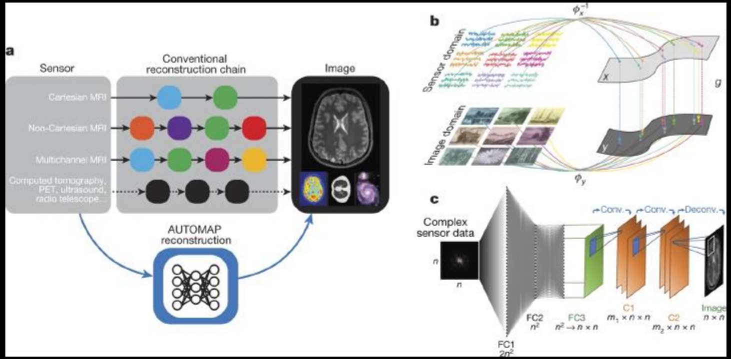

AUTOMAP (Automated Transform by Manifold Approximation)[1]

No, I'm very impressed, because we did not expect that so many problems could be solved by pure curve fitting. It turns out they can. – Judea Pearl

1. Zhu et al. Nature 555:487-492 (2018)

Image Reconstruction

State-of-the-art reconstruction based variational methods:[1]

$$x^\star = \operatorname{argmin}_x \mathcal{D}(Fx, y) + \mathcal{R}(x)$$

Data fidelity $\mathcal{D}$, MRI physics: F, data y, regularization $\mathcal{R}$

- Data fidelity term $\mathcal{D}$ ensures consistency with the acquired data.
- Regularization term stabilizes reconstruction using prior knowledge.

- Recent deep-learning methods also based on this approach![2]

$\Rightarrow$ Learned regularization term $\mathcal{R}$

1. Fessler JA. arXiv:1903.03510 (2019)
2. Knoll F et al. IEEE SPI 37;128–140 (2020)

Image Reconstruction

Is the hype justified?

Misleading Results: Metrics

- Receive-coil arrays $\Rightarrow$ No real ground truth
- FastMRI challenge:[1] RSS images as ground truth

$\Rightarrow$ Background noise, no phase[1]

$\Rightarrow$ Comparison to conventional methods that try to approximate the MVUE is then misleading:[2]

	MoDL-MVUE		MoDL-RSS		PICS (L1-Wavelet)		Zero-Filled	
	R = 4	*R* = 8	*R* = 4	*R* = 8	*R* = 4	*R* = 8	*R* = 4	*R* = 8
Test on MVUE	0.950 (38.3)	0.891 (31.3)	0.775 (33.1)	0.716 (29.5)	0.929 (37.6)	0.757 (26.8)	0.780 (27.0)	0.631 (22.5)
Test on RSS	0.782 (33.7)	0.723 (30.0)	0.945 (37.4)	0.895 (31.2)	0.751 (33.2)	0.668 (26.7)	0.793 (33.2)	0.662 (26.7)

Table 1. Average test SSIM (PSNR in parentheses) for the considered methods on fastMRI T2 brain scans.

1. Muckley MJ et al. IEEE TMI 37:2306–2317 (2020)
2. Arvinte M and Tamir J. ESMRMB MRITogether Workshop (2021)

Misleading Results: Data

- ▶ Public image data based may be preprocessed
- ⇒ Misleadingly optimistic results[1]

Data crimes:[1]

- ▶ Data was zero-padded during reconstruction
- ▶ Images were compressed using a lossly format (JPEG)

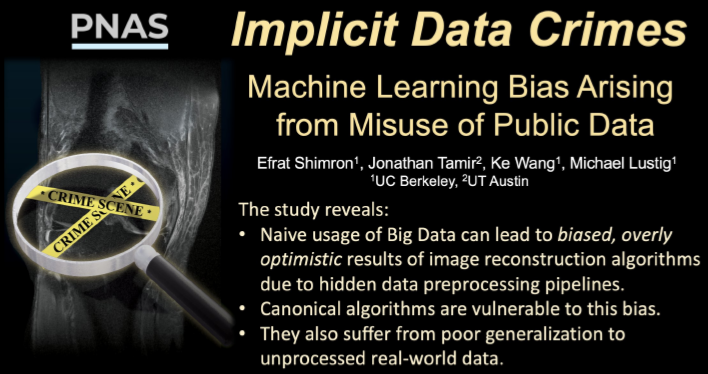

1. Shimron E et al. PNAS 119:e2117203119 (2019)

Comparison to Optimized Compressed Sensing Parallel Imaging

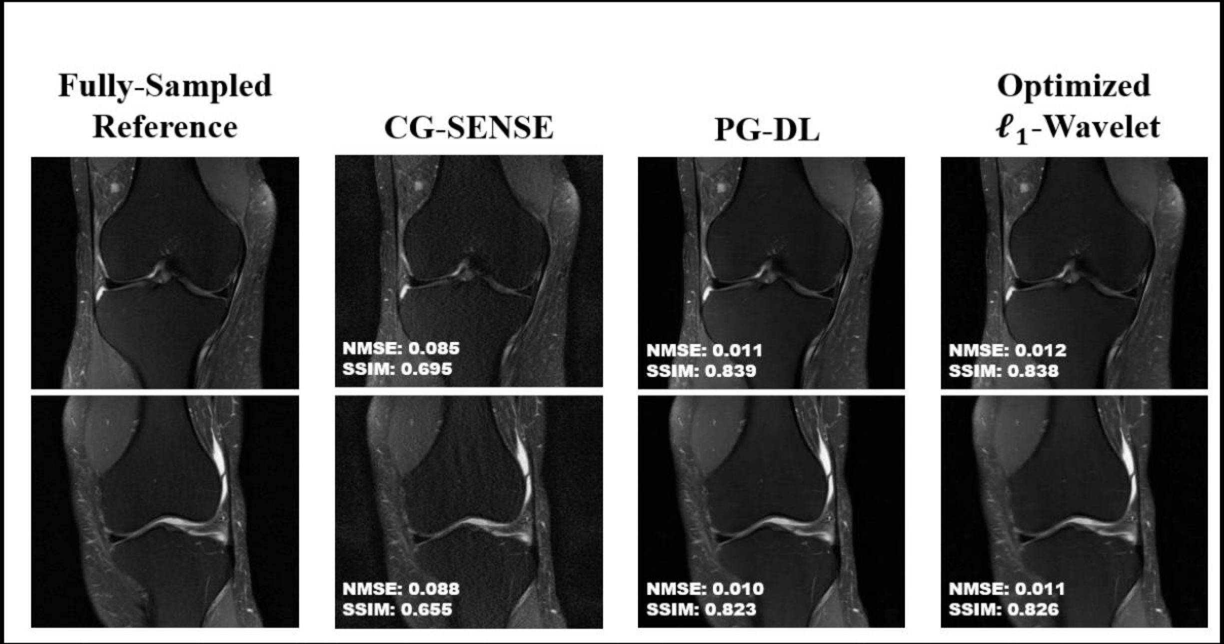

Comparison of deep-learning method to conventional methods and optimized compressed sensing parallel imaging[1]

1. Gu U et el. ISMRM 20:274 (2021)

Hallucinations

- Problem: Artificial image features learned during training appear in the reconstruction.
- No surprise: Generative models can produce realistic images from nothing.
- Data fidelity should help if acquired data are sufficient.
- Perfectly looking images do not reveal if available data are insufficient.

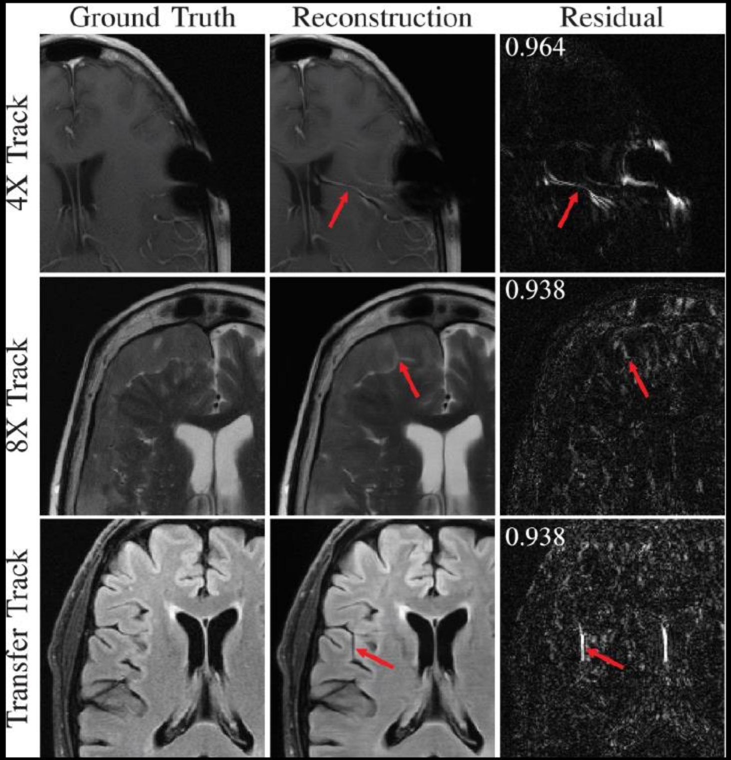

Hallucinations observed in the FastMRI challenge.[1]

1. Muckley MJ et al. IEEE TMI 37:2306–2317 (2020)

Image Reconstruction

Is there hope?

Reconstruction as Bayesian Inference

Likelihood: $p(y|x) = det(\pi\Gamma)^{-\frac{1}{2}} e^{-\|\Gamma^{-\frac{1}{2}}(y-Fx)\|_2^2}$

(non-linear) physics-based forward model F, data y, noise covariance matrix Γ

Prior: $p(x) = \cdots$ (e.g. sparsity, learned priors)

Posterior: $p(x|y) = \frac{p(y|x)p(x)}{p(y)}$ via Bayes' theorem

Data fidelity: $\mathcal{D}(Fx, y) = -\log p(y|x)$ (neg. log-likelihood)
Regularization: $\mathcal{R}(x) = -\log p(x)$ (neg. prior)

$$\text{argmin}_x \, \mathcal{D}(Fx, y) + \mathcal{R}(x)$$

Regularized reconstruction $\Leftrightarrow$ Maximum a posteriori probability (MAP)

Uncertainty Quantification

Sampling the posterior using MCMC $\Rightarrow$ MMSE, uncertainty[1–4]

Euler-Maruyama: $x_{n+1} = x_n + \frac{\gamma}{2}\nabla \log p(x|y) + \gamma z \quad (z \text{ Gaussian noise})$

Variance map quantifies uncertainty.[2,3]

1. Jalal A et al. NeurIPS 2021
2. Luo G et al. arXiv:2202.01479 (2022)
3. Luo G et al. ISMRM 30:0298 (2022)
4. others...

CMR Imaging

The future of CMR:

- Free-breathing and self-gated dynamic imaging
- High-dimensional + multi-parametric

$\Rightarrow$ Ground truth images may be impossible to acquire.

Self-Supervised Learning

- Self-supervised learning from k-space[1]
- NonLinear INVersion (NLINV): autocalibration, arbitrary sampling patterns[2]
- Application to cardiac real-time with radial acquisition[3]

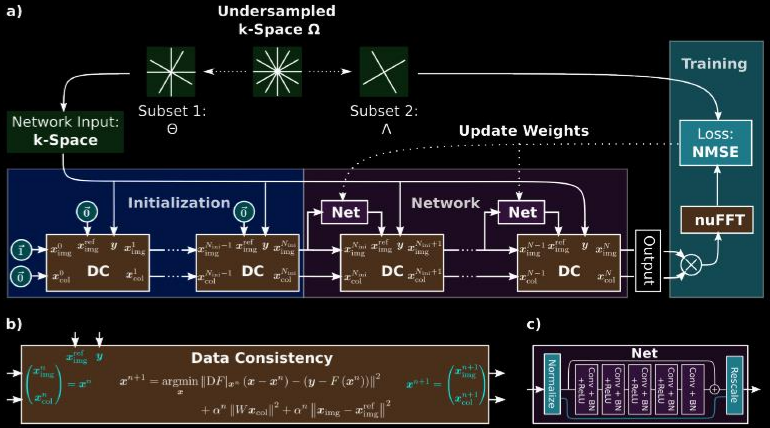

network architecture

NLINV

NLINV-NET

39 radial spokes per frame

1. Yaman B et al MRM 84:3172-3191 (2020)
2. Uecker M et al. MRM 60:674-682 (2008)
3. Blumenthal M et al. ISMRM 2022; 30:499.

Conclusion

- Published results are sometimes misleading.
- Perfectly looking images might be wrong.
- Progress is incremental but real.

Needed:

- Methods based on sound principles
- More focus on scientific understanding
- Open data sets and benchmarks
- Reproducible research based on free and open software

BART

BART Toolbox for Computational MRI
`https://mrirecon.github.io/bart/`