Deep learning for medical image reconstruction

Models, priors and data

Ruud van Sloun

Image reconstruction

Magnetic Resonance Imaging
(faster, low-field etc.)

Computed Tomography
(low-dose, sparse-view)

Ultrasound
(high-quality, reduced data-rates)

TU/e EINDHOVEN UNIVERSITY OF TECHNOLOGY

Ultrasound imaging basics

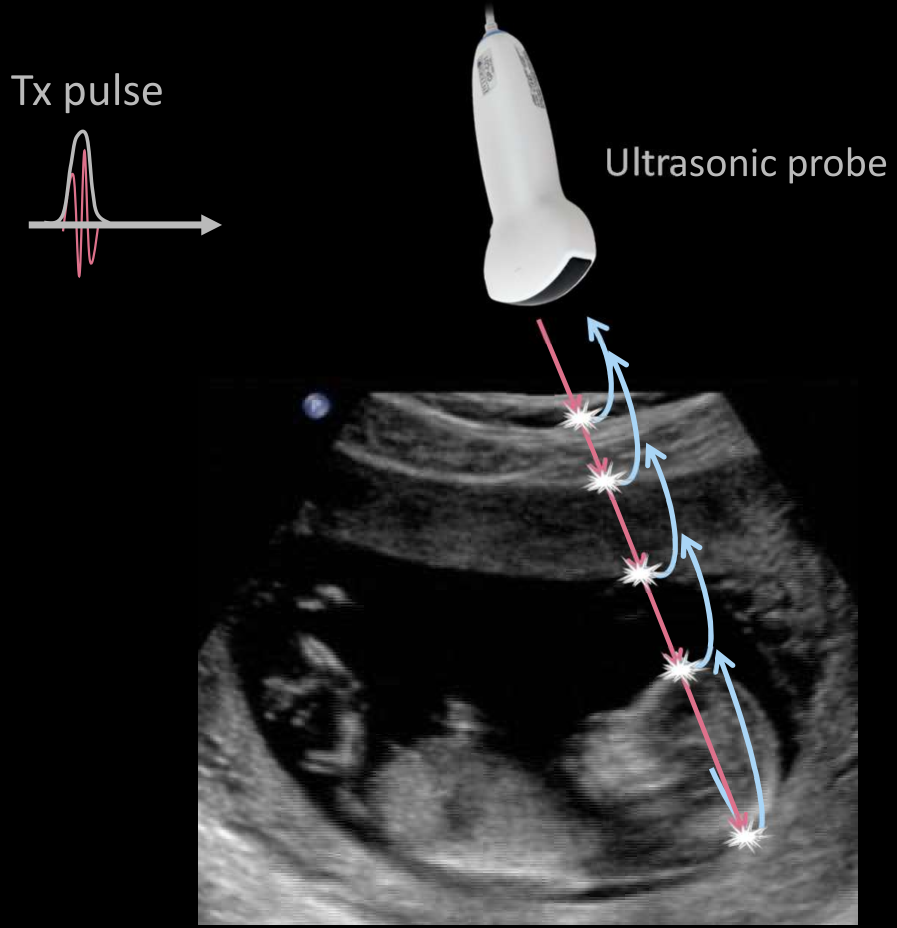

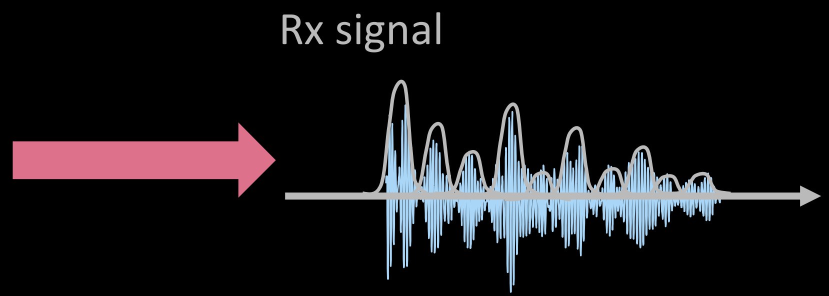

- Echoes result from scattering in the tissue
  (Coherent & incoherent - speckle)

- The image is formed by estimating (reflected) signal amplitude from a set of spatial locations (i.e. pixels) using an array of sensors = **Beamforming**

Ultrasound imaging basics

Transmission

Channel Data

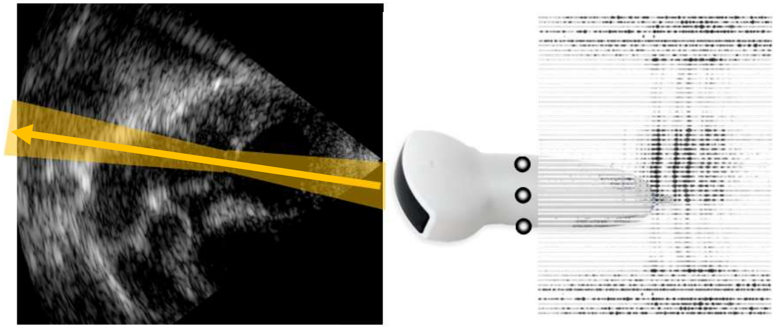

Hundreds to thousands
of channels/firing

High rate ADC
(~20-50MHz / element)

High-data rates

Digital Beamforming

Spatial filtering

Reconstructed RF scanline

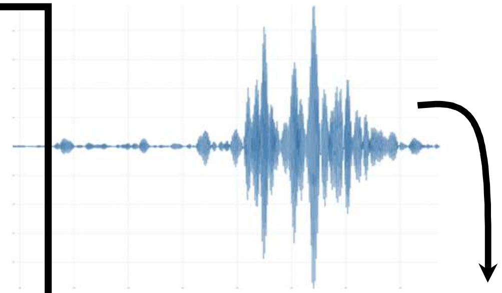

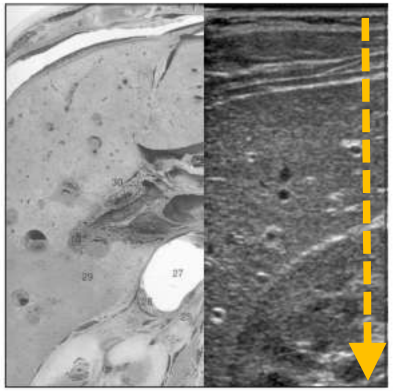

Envelope-detected
(& compressed) scanline

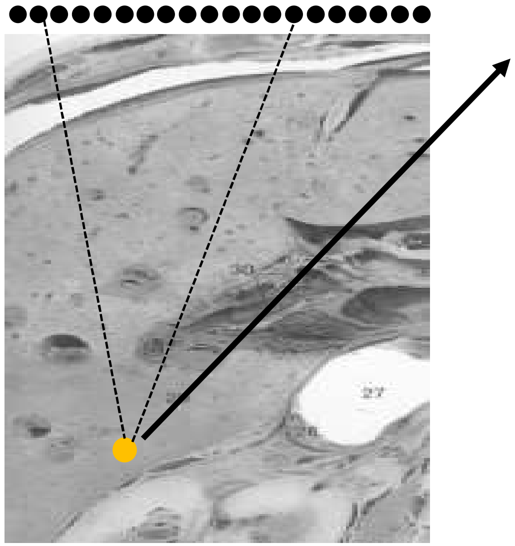

How much signal is reflected from this particular pixel?

Spatial filtering = beamforming

- Sidelobes & grating lobes
- Mainlobe width (resolution)

The better the spatial selectivity,
The better the tissue contrast and resolution

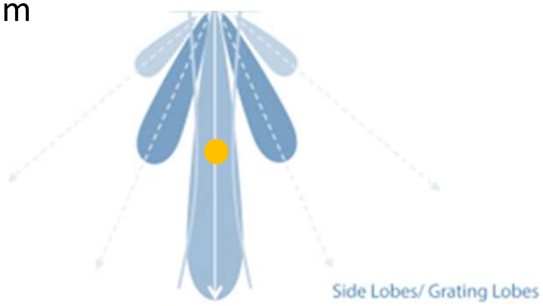

Image from “Speckle reduction imaging” by Milkowski et al.

TU/e EINDHOVEN UNIVERSITY OF TECHNOLOGY

Ultrasound imaging basics

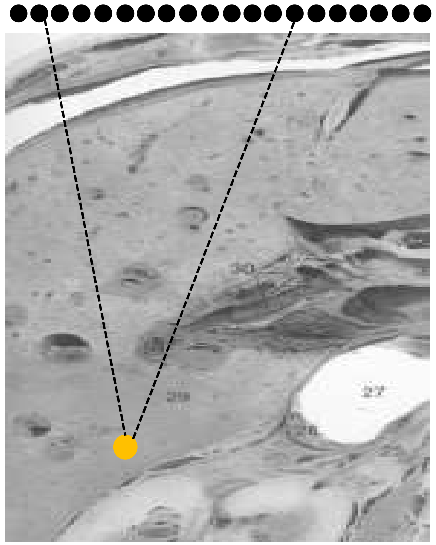

Better spatial selectivity,
Better tissue contrast and resolution

Image quality is a function of:

1. Physics: array geometry and probe bandwidth

Increased array aperture – increased angular resolution

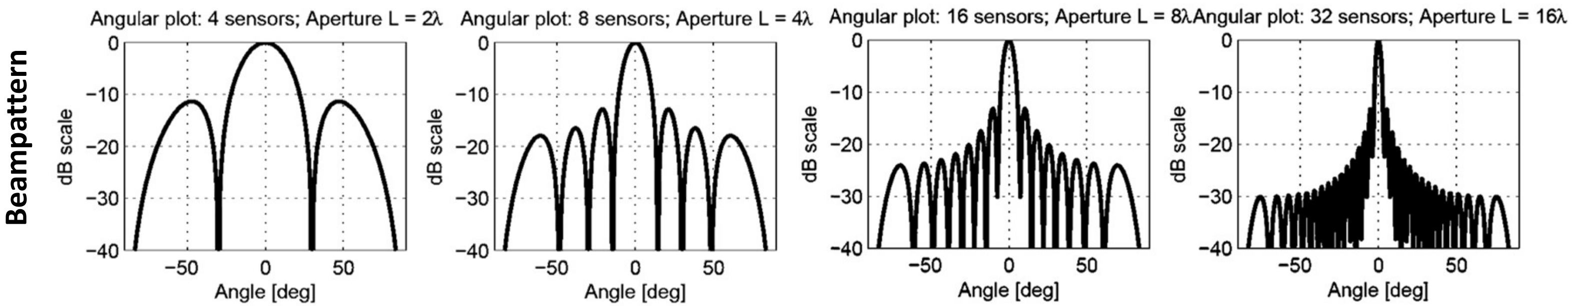

2. Algorithms: powerful digital signal processing and beamforming on RF channel data

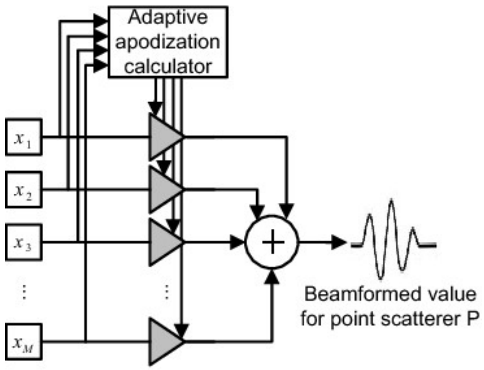

AI opportunity

Ultrasound imaging AI opportunities

Focused transmit

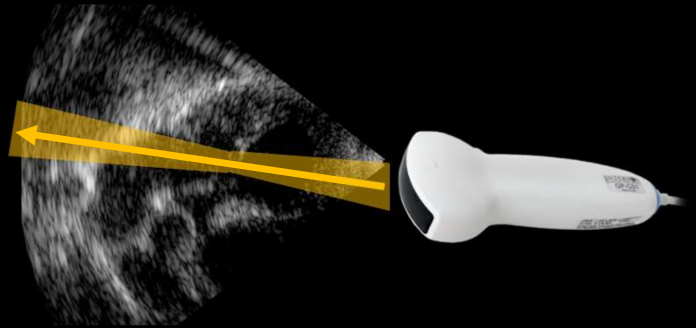

Parallel ultrafast transmit

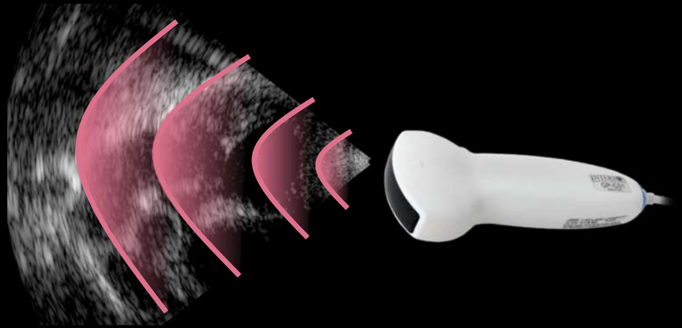

Parallel ultrafast acquisition

- High time-resolution
- Compromising spatial resolution and contrast
- Relies more heavily on receive spatial filtering/beamforming

High image quality under minimal data rates

- Improve tissue contrast (*accurate* contrast)
- Resolution depends on array aperture -> high-res with small/sparse aperture?
- Compressed sensing to reduce data rates at the probe

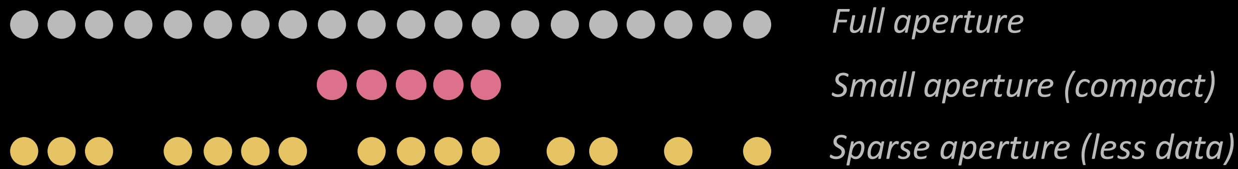

TU/e EINDHOVEN UNIVERSITY OF TECHNOLOGY

Image basics: MRI

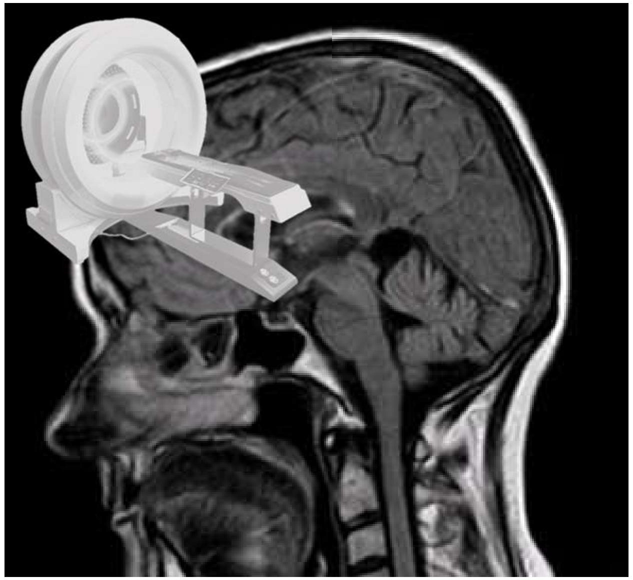

Magnetic Resonance Imaging
(faster, low-field etc.)

Inverse
FT

K-space

Pulse sequence

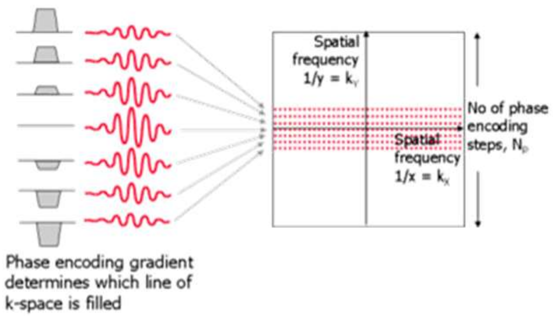

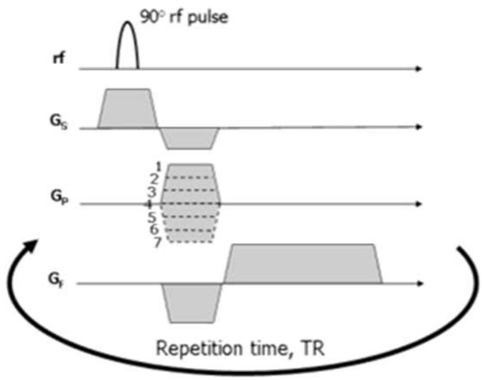

Total time *~ TR×amount of repetitions/lines*

Can we go faster?

TU/e EINDHOVEN UNIVERSITY OF TECHNOLOGY

Image basics: MRI

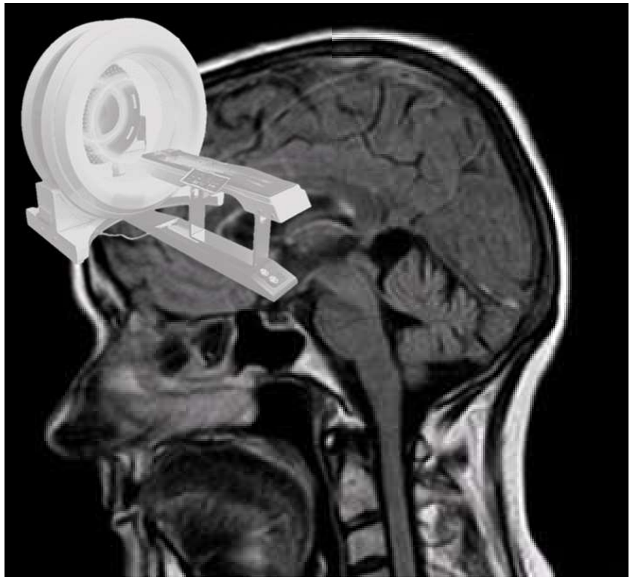

Magnetic Resonance Imaging (faster, low-field etc.)

$p(x|y)$?

Reconstruction algorithm

Priors

$p(x)$?

Undersampled K-space

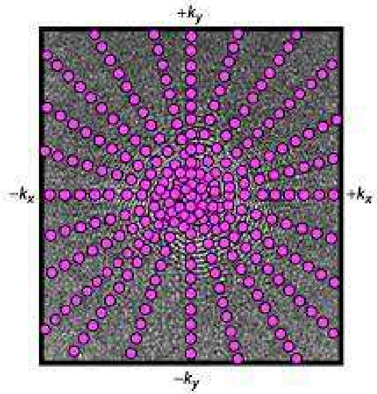

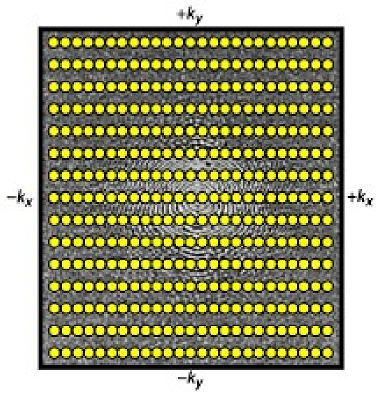

= less acquisitions = faster

$p(y|x)$?

Sampling/ acquisition

How?

Full K-space

TU/e EINDHOVEN UNIVERSITY OF TECHNOLOGY

Image reconstruction

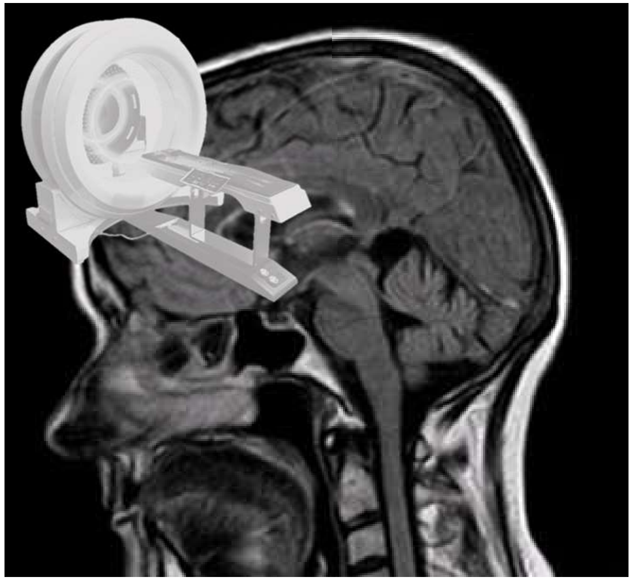

Magnetic Resonance Imaging
(faster, low-field etc.)

Fast MRI

- Compressed sensing/fewer acquisitions
- Acceleration
- Aliasing artefacts

Low-field MRI

- Lower field strengths: compact MRI machines
- Low SNR k-space -> less clear high-frequencies -> low resolution

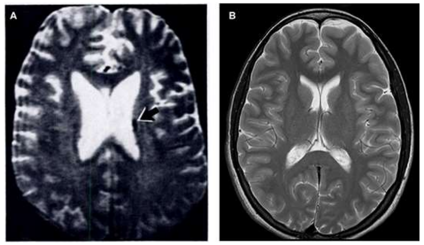

Image reconstruction

Computed Tomography
(low-dose, sparse-view)

Low-dose CT
- Safer – less ionizing radiation
- Low SNR -> limited fidelity

Low **High**

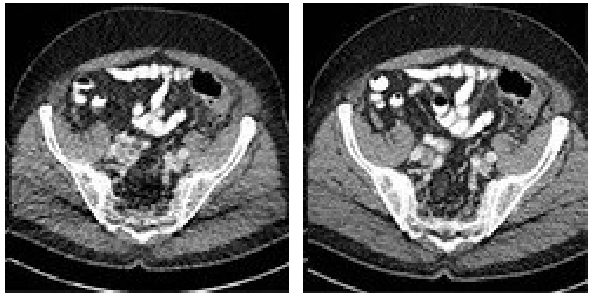

Sparse CT
- reduced scan time and improved time resolution
- Lower dose
- Undersampling artefacts

Full **Sparse**

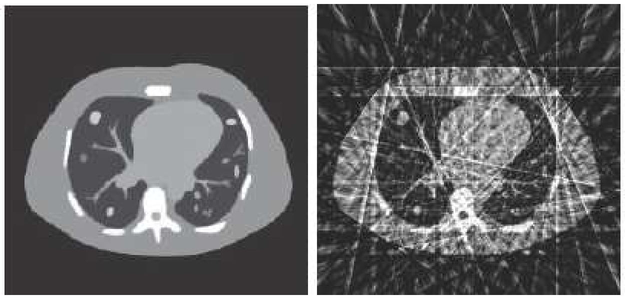

TU/e EINDHOVEN UNIVERSITY OF TECHNOLOGY

Image reconstruction

Magnetic Resonance Imaging
(faster, low-field etc.)

Computed Tomography
(low-dose, sparse-view)

Ultrasound
(high-quality, reduced data-rates)

TU/e EINDHOVEN UNIVERSITY OF TECHNOLOGY

What gaps can we fill by learning?

Challenges in classical image recon based on models

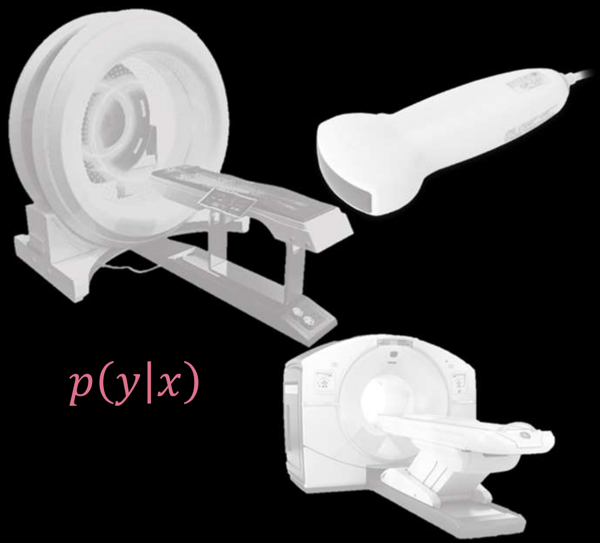

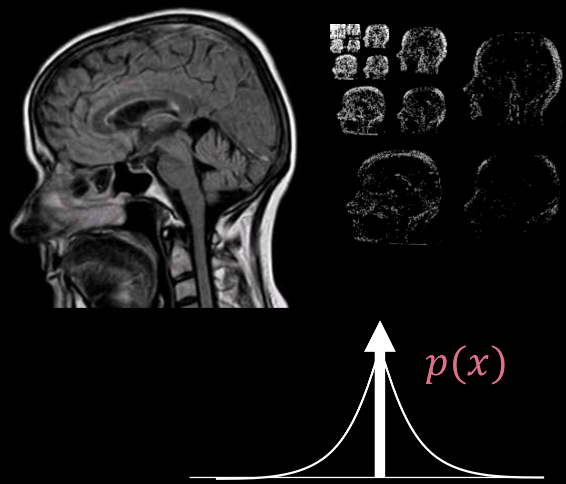

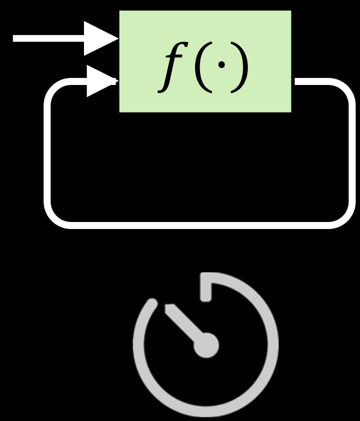

Acquisition model assumptions incorrect/imprecise

Statistical image priors not sufficiently expressive/accurate

Slow image reconstruction with high complexity (e.g. iterative, matrix inversions etc)

TU/e EINDHOVEN UNIVERSITY OF TECHNOLOGY

Models and Priors: ultrasound

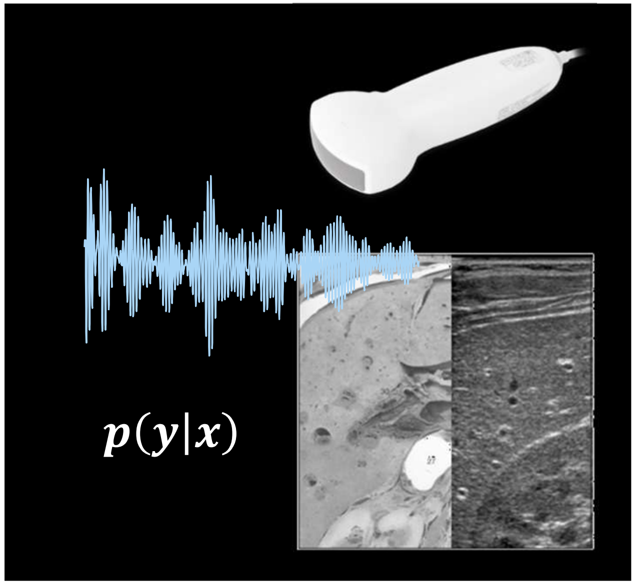

Acquisition model assumptions incorrect/imprecise (multiple scattering, aberration etc.)

$p(\mathbf{y}|\mathbf{x})$: Likelihood of measurements $\mathbf{y}$ given object $\mathbf{x}$

Across our imaging pipeline and set of applications, $\boldsymbol{y}$ can be:

(1) RF channel data (array response)

(2) Beamformed RF data

(3) Image data (beamformed + envelope detected)

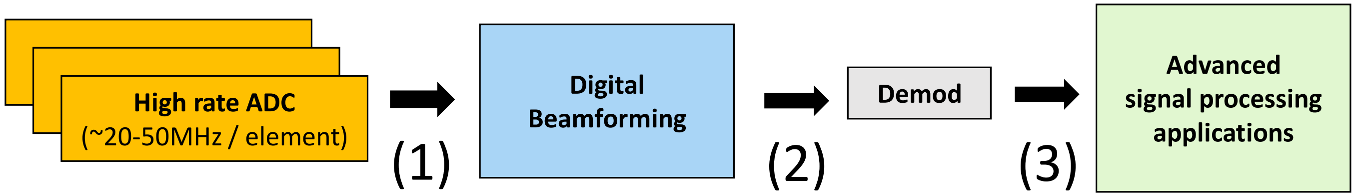

EINDHOVEN UNIVERSITY OF TECHNOLOGY

Models and Priors: ultrasound

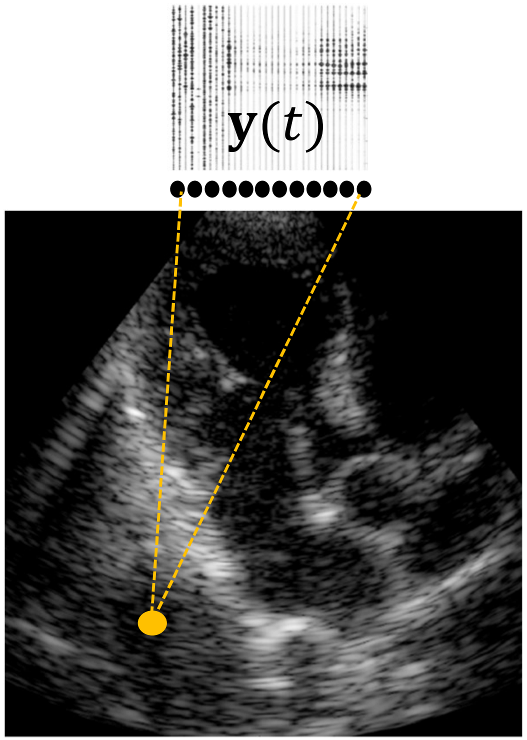

Acquisition model assumptions incorrect/imprecise (multiple scattering, aberration etc.)

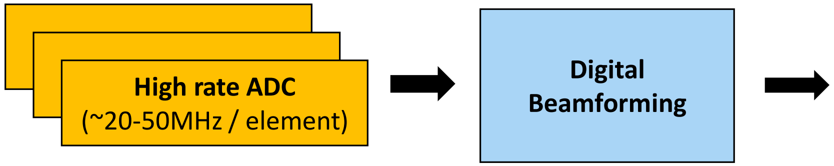

Array response (narrowband) single target:

$$\mathbf{y}(t) = \mathbf{a}(\theta)x(t) + \mathbf{n}(t)$$

Array response vector: (assumed constant SoS, no aberration)

<u>Noise vector:</u>

- Sensor noise
- Off-axis scattering/interference
- Reverberation

➡ In practice complex statistics

Typical assumption: $\mathbf{n}(t) \sim \mathcal{N}(0, \sigma_n^2 I)$

➡ *MSE-optimal* matched filter: $\hat{x}(t) = \mathbf{w}^H \mathbf{y}(t)$ **($\mathbf{w}$ = $\mathbf{a}$)**

= delay-and-sum (DAS) for wideband

Models and Priors: ultrasound

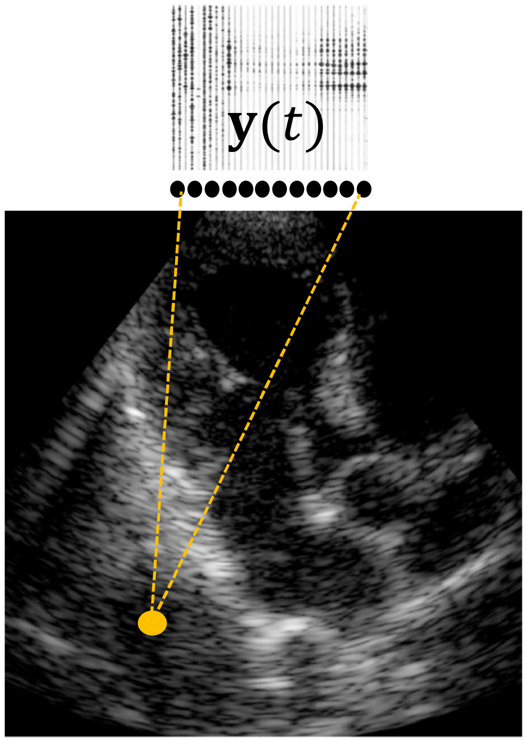

Acquisition model assumptions incorrect/imprecise (multiple scattering, aberration etc.)

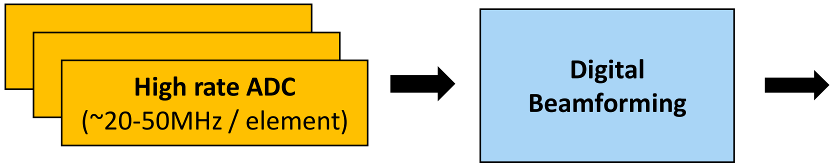

Array response (narrowband) single target:

$$\mathbf{y}(t) = \mathbf{a}(\theta)x(t) + \mathbf{n}(t)$$

Array response vector: (assumed constant SoS, no aberration)

Noise vector:

- Sensor noise
- Off-axis scattering/interference
- Reverberation

➡ In practice complex statistics

Model-based alternatives to DAS:

- **MVDR** (minimize total (noise) power but retain unity gain)
- **iMAP** (assume $x(t) \sim \mathcal{N}(0, \sigma_x^2)$; iteratively estimate σ_x, σ_n per pixel)
- **ADMIRE** (aperture domain model of signal and noise/interference/clutter; separate using optimization methods)

Jensen *et al.*
Eldar *et al.*
Byram *et al.*

Models and Priors: ultrasound

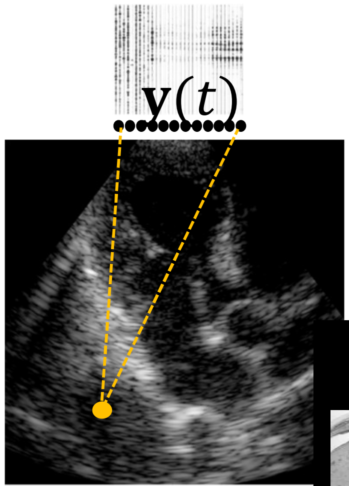

Acquisition model assumptions incorrect/imprecise
(multiple scattering, aberration etc.)

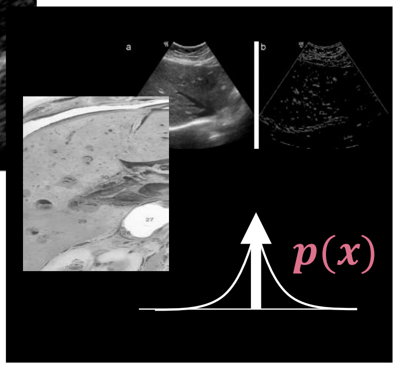

Statistical image priors not sufficiently expressive/accurate

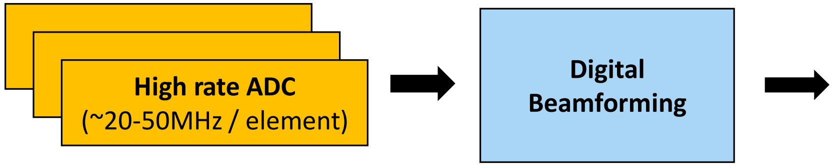

$$\mathbf{y}(t) = \mathbf{a}(\theta)x(t) + \mathbf{n}(t)$$

- **DAS** (matched filter)
- **MVDR** (minimize total (noise) power but retain unity gain)
- **iMAP** (assume $x(t) \sim \mathcal{N}(0, \sigma_x^2)$; iteratively estimate σ_x, σ_n per pixel)
- **ADMIRE** (aperture domain model of signal and noise/interference/clutter; separate using optimization methods)

General problem: model $p(\boldsymbol{y}|\boldsymbol{x})$ too simple / incomplete / or hard to solve due to suboptimal prior $p(\mathbf{x})$

TU/e EINDHOVEN UNIVERSITY OF TECHNOLOGY

Models and Priors: ultrasound

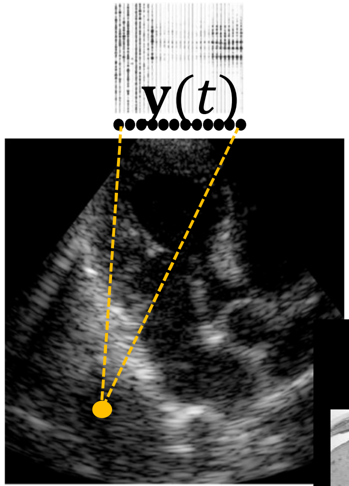

Acquisition model assumptions incorrect/imprecise (multiple scattering, aberration etc.)

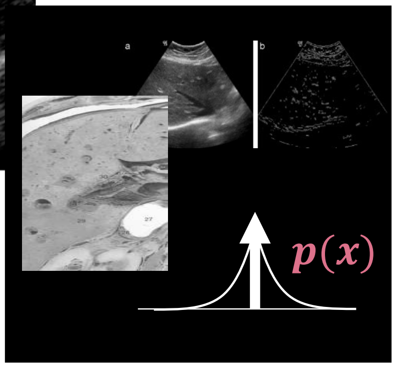

Statistical image priors not sufficiently expressive/accurate

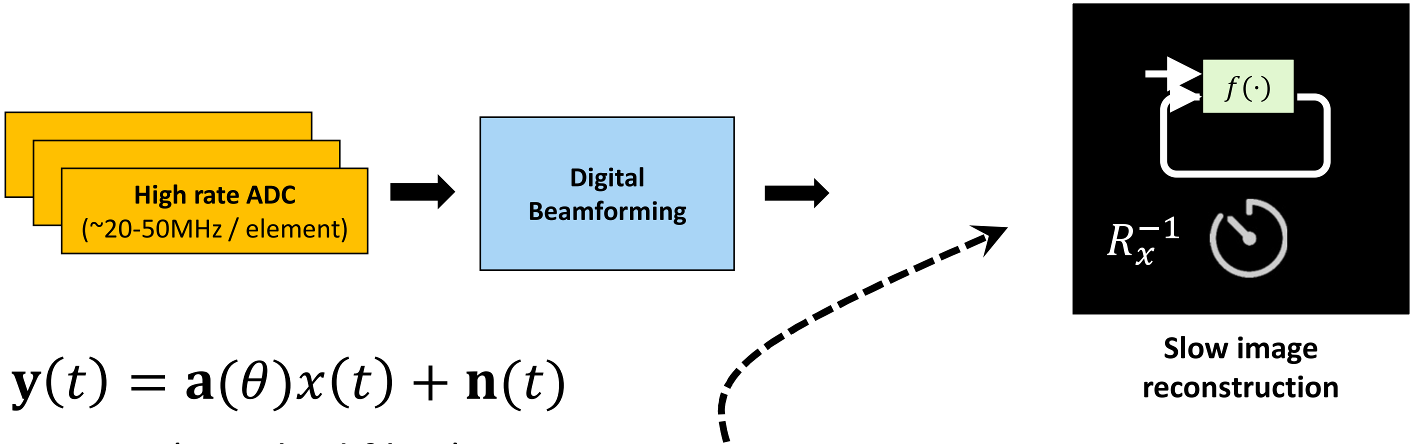

Slow image reconstruction

$$\mathbf{y}(t) = \mathbf{a}(\theta)x(t) + \mathbf{n}(t)$$

- **DAS** (matched filter)
- **MVDR** (minimize total (noise) power but retain unity gain)
- **iMAP** (assume $x(t) \sim \mathcal{N}(0, \sigma_x^2)$; iteratively estimate σ_x, σ_n per pixel)
- **ADMIRE** (aperture domain model of signal and noise/interference/clutter; separate using optimization methods)

General problem: model $p(\boldsymbol{y}|\boldsymbol{x})$ too simple / incomplete / or hard to solve due to suboptimal prior $p(\mathbf{x})$

Models and Priors: MRI

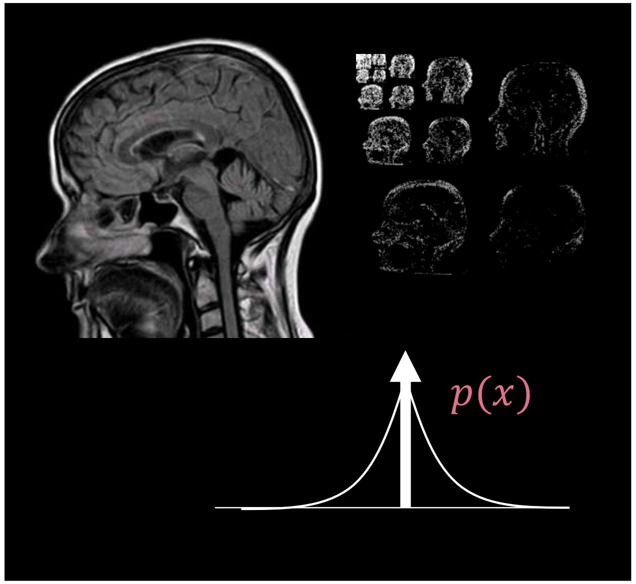

Statistical image priors not sufficiently expressive/accurate

$p(\mathbf{y}|\mathbf{x})$: Likelihood of measurements $\mathbf{y}$ given object $\mathbf{x}$

K-space measurement:

$\mathbf{y} = \mathbf{AFx} + \mathbf{n}$ → $\mathbf{A}$: Subsampling matrix
$\mathbf{F}$: 2D Fourier transform

For low-field: $\mathbf{n}$ is dominating higher frequencies
For fast-MRI: $\mathbf{A}$ is sampling below Nyquist (compressed sensing)

Classical reconstruction: $\mathbf{x} = \mathbf{F}^{-1}\mathbf{z}$ ($\mathbf{z}$ = zero-filled k-space $\mathbf{y}$)

Inverse (MAP) problem:

$$\hat{\mathbf{x}} = \text{argmax}_{\mathbf{x}}\ p(\mathbf{y}|\mathbf{x})p(\mathbf{x})$$

$$\hat{\mathbf{x}} = \text{argmin}_{\mathbf{x}}\ \frac{1}{2}||\mathbf{y} - \mathbf{AFx}||_2^2 - \log p(\mathbf{x})$$

<u>General problem</u>: prior $p(\boldsymbol{x})$ not sufficiently expressive

What if we could improve our models using deep learning without ending up just trowing away what we do know about the problem?

Model-based deep learning

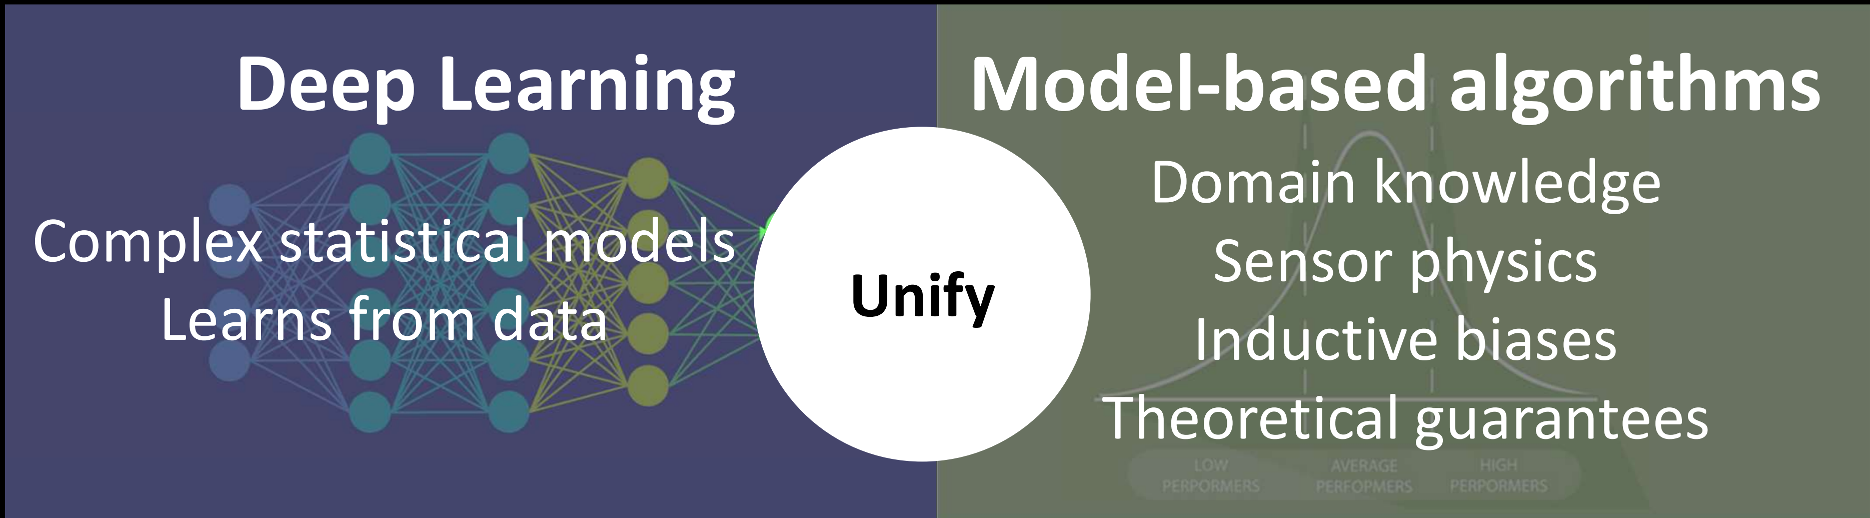

How to bring together to get the best of both?

Van Sloun *et al. Proceedings of the IEEE*, 2020
Luijten, …, Van Sloun. *IEEE trans. med. im.*, 2020
Solomon, … van Sloun, Eldar. *IEEE trans. med. im.*, 2019

One step back: Model-based optimization

General recipe: construct a model-based optimization algorithm based on a-priori knowledge of the measurement process and statistical priors

Define measurement model with image of interest $\mathbf{x}$:

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n}$$

Formulate MAP optimization problem:

$$\hat{\mathbf{x}} = \text{argmax}_{\mathbf{x}}\ p(\mathbf{y}|\mathbf{x})\, p(\mathbf{x})$$
$$\hat{\mathbf{x}} = \text{argmin}_{\mathbf{x}}\ -\log\left(p(\mathbf{y}|\mathbf{x})\, p(\mathbf{x})\right)$$

Likelihood model under Normal distribution:

$$p(\mathbf{y}|\mathbf{x}) = ce^{-\frac{1}{2}(\mathbf{y}-\mathbf{A}\mathbf{x})^T\boldsymbol{\Sigma}^{-1}(\mathbf{y}-\mathbf{A}\mathbf{x})}$$

$$\hat{\mathbf{x}} = \text{argmin}_{\mathbf{x}}\ \frac{1}{2}(\mathbf{y}-\mathbf{A}\mathbf{x})^T\boldsymbol{\Sigma}^{-1}(\mathbf{y}-\mathbf{A}\mathbf{x}) \quad - \log p(\mathbf{x})$$

(assume uncorrelated noise)

$$\hat{\mathbf{x}} = \text{argmin}_{\mathbf{x}}\ \frac{1}{2}\left|\left|\mathbf{y}-\mathbf{A}\mathbf{x}\right|\right|_2^2 - \log p(\mathbf{x})$$

Many iterative solvers (for particular choices of $p(\mathbf{x})$)

ISTA, ADMM, etc.

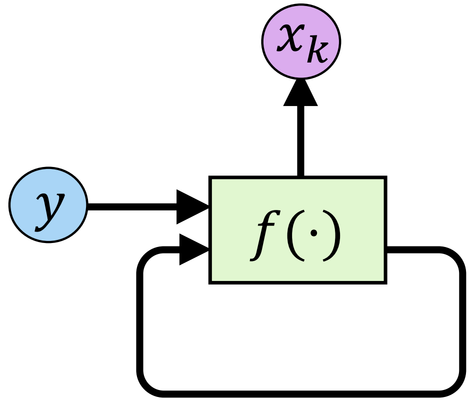

Iterative model-based algorithm with input y and output x

One step back: Model-based optimization

General recipe: <u>construct a model-based optimization algorithm</u> based on a-priori knowledge of the measurement process and statistical priors

Define measurement model with image of interest $\mathbf{x}$:

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n}$$

$$\hat{\mathbf{x}} = \mathrm{argmin}_{\mathbf{x}} \frac{1}{2}||\mathbf{y} - \mathbf{A}\mathbf{x}||_2^2 - \log p(\mathbf{x})$$

$$\hat{\mathbf{x}} = \mathrm{argmin}_{\mathbf{x}} \frac{1}{2}||\mathbf{y} - \mathbf{A}\mathbf{x}||_2^2 + g_\theta(\mathbf{x})$$ (regularizer g_θ parameterized by θ)

Iterative proximal gradient solvers:

Prior → Proximity function

$$\mathbf{z} = \mathbf{x}^k - \mu\left(\nabla_{\mathbf{x}}||\mathbf{y} - \mathbf{A}\mathbf{x}||_2^2\right)\Big|_{\mathbf{x}=\mathbf{x}^k} = f_1(\mathbf{x}^k)$$

$$\mathbf{x}^{k+1} = \mathrm{argmin}_{\mathbf{x}} \frac{1}{2}||\mathbf{z} - \mathbf{x}||_2^2 + g_\theta(\mathbf{x}) = \mathrm{Prox}(\mathbf{z})$$

→ *2-step factorized optimization*

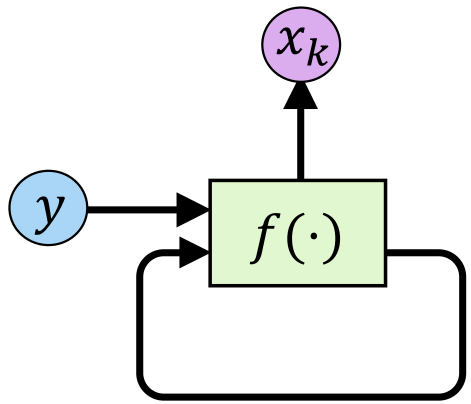

Iterative model-based algorithm with input y and output x

One step back: Model-based optimization

General recipe: construct a model-based optimization algorithm based on a-priori knowledge of the measurement process and statistical priors

Define measurement model with image of interest $\mathbf{x}$:

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n}$$

$$\hat{\mathbf{x}} = \text{argmin}_{\mathbf{x}} \; \frac{1}{2}||\mathbf{y} - \mathbf{A}\mathbf{x}||_2^2 - \log p(\mathbf{x})$$

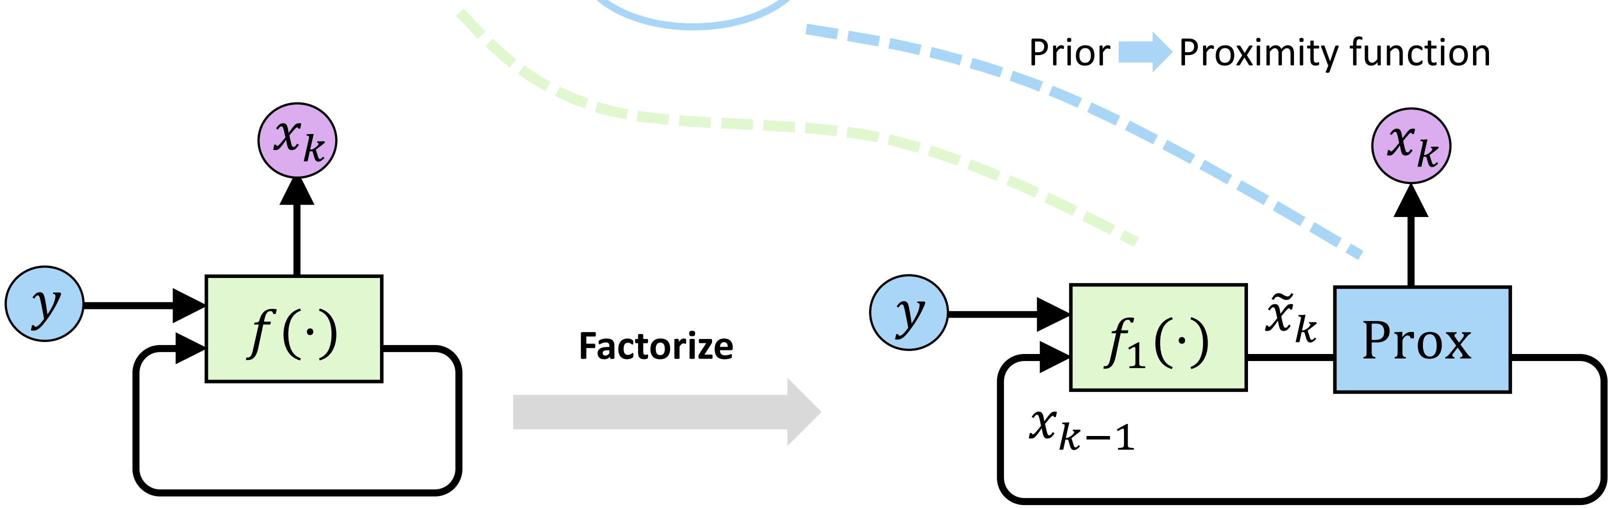

Factorized into **prior step** and **data consistency step** (=data likelihood)

Typical structure optimization algorithm

One step back: Model-based optimization

Example: sparse coding

Many applications: denoising, compressed sensing, image reconstruction, super-resolution, …

Some intuition:

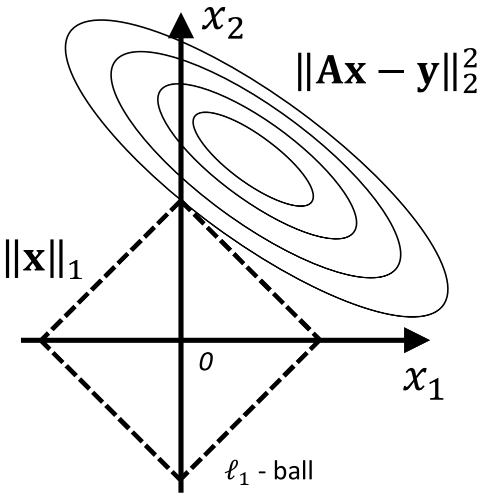

Sparse coding problem

$\mathbf{y} = \mathbf{Ax} + \mathbf{n}$ with $\mathbf{x}$ being sparse

MAP problem for $\mathbf{x}$:

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\text{minimize}} \ \|\mathbf{Ax}-\mathbf{y}\|_2^2 + \lambda\|\mathbf{x}\|_1$$

$(\mathbf{x} \sim \text{Laplace})$

⬇

$-\log p(\mathbf{x}) \sim \|\mathbf{x}\|_1$

One step back: Model-based optimization

Example: sparse coding

Many applications: denoising, compressed sensing, image reconstruction, super-resolution, …

Some intuition:

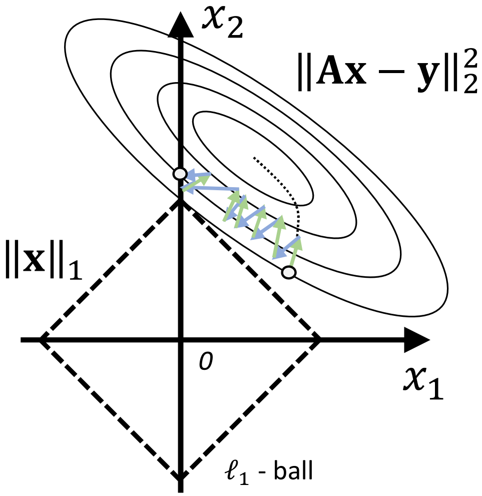

Sparse coding problem

$\mathbf{y} = \mathbf{Ax} + \mathbf{n}$ with $\mathbf{x}$ being sparse

MAP problem for $\mathbf{x}$:

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\text{minimize}} \ \|\mathbf{Ax}-\mathbf{y}\|_2^2 + \lambda\|\mathbf{x}\|_1$$

$(\mathbf{x} \sim \text{Laplace})$

⬇

$-\log p(\mathbf{x}) \sim \|\mathbf{x}\|_1$

Iterative shrinkage and thresholding (ISTA)

1. Take a gradient step towards $\min_{\mathbf{x}} \|\mathbf{Ax}-\mathbf{y}\|_2^2$ ➡ $\mathbf{x} - \mu \nabla_{\mathbf{x}} \|\mathbf{Ax}-\mathbf{y}\|_2^2$
2. Move intermediate solution towards prior

➡ *Solve*: $\min_{\mathbf{x}} \left\|\mathbf{x} - \hat{\mathbf{x}}^{(i)}\right\|_2^2 + \lambda\|\mathbf{x}\|_1$

$S_\lambda(\hat{\mathbf{x}}^{(i)})$

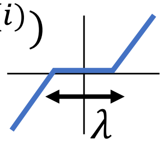

➡ = soft thresholding function

One step back: Model-based optimization

Example: sparse coding

Many applications: denoising, compressed sensing, image reconstruction, super-resolution, …

Some intuition:

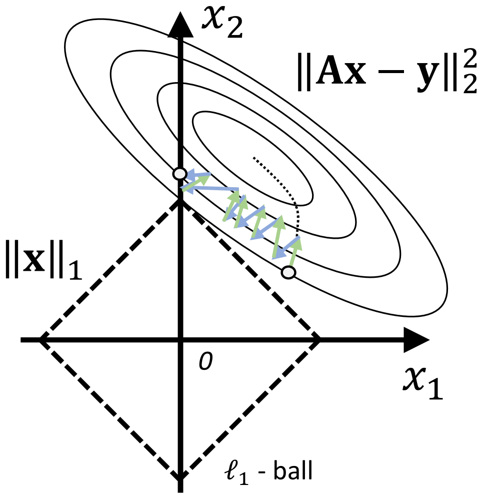

Sparse coding problem

$\mathbf{y} = \mathbf{Ax} + \mathbf{n}$ with $\mathbf{x}$ being sparse

MAP problem for $\mathbf{x}$:

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\text{minimize}} \ \|\mathbf{Ax}-\mathbf{y}\|_2^2 + \lambda\|\mathbf{x}\|_1$$

($\mathbf{x}$~Laplace)

⬇

$-\log p(\mathbf{x}) \sim \|\mathbf{x}\|_1$

Iterative shrinkage and thresholding (ISTA)

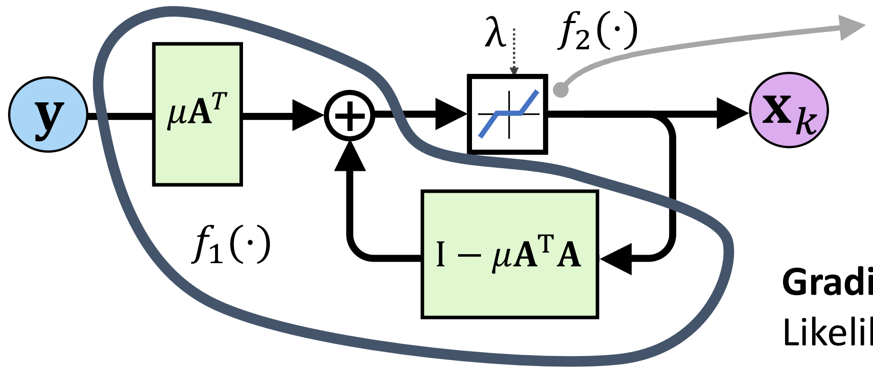

Proximity operator
Soft thresholding

$$\min_{\mathbf{x}} \left\|\mathbf{x}-\hat{\mathbf{x}}^{(i)}\right\|_2^2 + \lambda\|\mathbf{x}\|_1 = S_\lambda(\hat{\mathbf{x}}^{(i)})$$

Gradient update step w.r.t. Likelihood term

$$\mathbf{x} - \mu\nabla_{\mathbf{x}}\|\mathbf{Ax}-\mathbf{y}\|_2^2$$

Model-based DL: Deep unfolding

General recipe: take a model-based optimization algorithm, unfold it as a fixed-complexity graph, learn parameters.

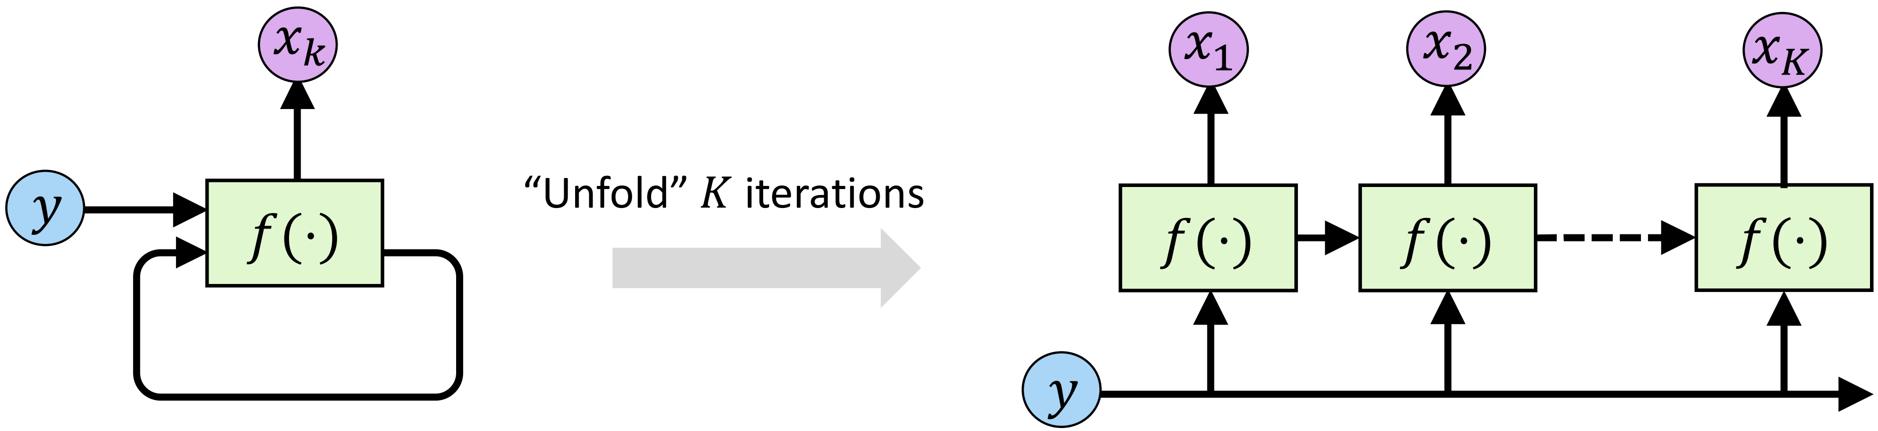

Iterative model-based algorithm with input y and output x

unfolded model-based algorithm with input y and output x

Iterative algorithm

Learned iterative algorithm

Deep unfolding for sparse coding (LISTA)

General recipe: take a model-based optimization algorithm, unfold it as a fixed-complexity graph, learn parameters.

"Unfold" K iterations

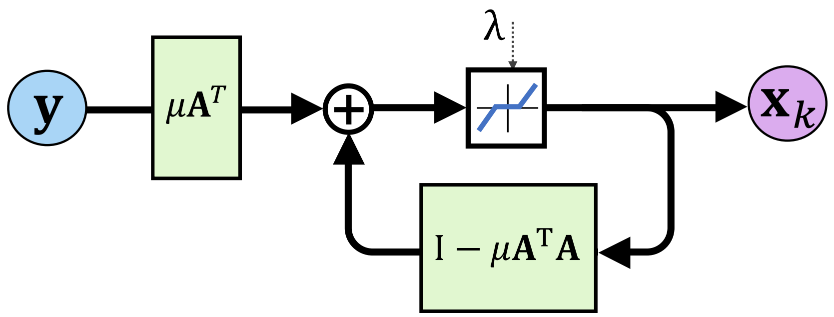

ISTA

Learned(L)ISTA

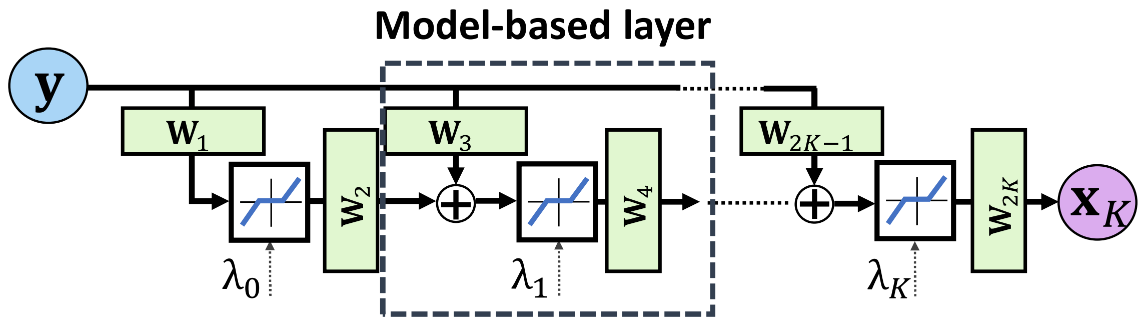

Deep learning with a model-based signal prior

- Learn the weight matrices/convolutions W_i
- Learn the thresholding parameters λ_i
- Use the prior (sparsity) and optimization structure

Note: Network nonlinearity (activation function) follows directly from the prior!

Gregor & LeCun ICML 2010

Optimization with “plug and play priors”

General recipe: take a model-based optimization algorithm, …. what if we don’t know the prior, or it is complex to describe/model?

Define measurement model with image of interest $\mathbf{x}$:

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n}$$

$$\hat{\mathbf{x}} = \mathrm{argmin}_{\mathbf{x}} \ \frac{1}{2}||\mathbf{y} - \mathbf{A}\mathbf{x}||_2^2 - \log p(\mathbf{x})$$

$$\hat{\mathbf{x}} = \mathrm{argmin}_{\mathbf{x}} \frac{1}{2}||\mathbf{y} - \mathbf{A}\mathbf{x}||_2^2 + f_\theta(\mathbf{x})$$

(regularizer f_θ parameterized by θ)

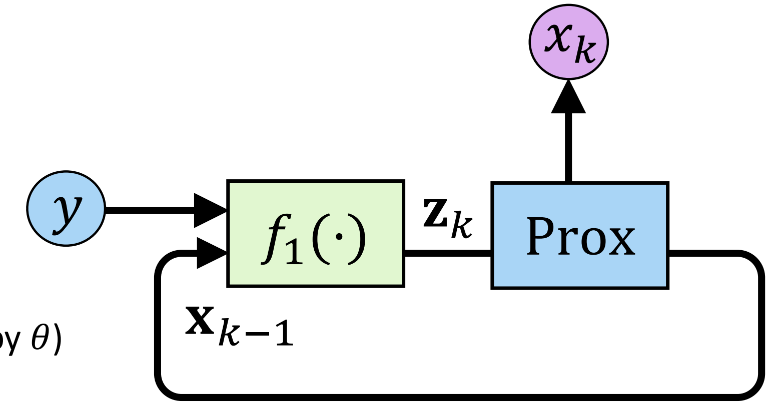

Typical structure optimization algorithm

Iterative proximal gradient methods:

Prior ⇒ Proximity function

$$\mathbf{z} = \mathbf{x}^k - \mu\left(\nabla_{\mathbf{x}}||\mathbf{y} - \mathbf{A}\mathbf{x}||_2^2\right)\Big|_{\mathbf{x}=\mathbf{x}^k} = f_1(\mathbf{x}^k)$$

$$\mathbf{x}^{k+1} = \mathrm{argmin}_{\mathbf{x}} \frac{1}{2}||\mathbf{z} - \mathbf{x}||_2^2 + f_\theta(\mathbf{x}) = \mathrm{Prox}(\mathbf{z})$$

⇒ **Question**: what is a good general choice for $\mathrm{Prox}(\mathbf{z})$ in structured signals if you don’t know $p(\mathbf{x})$?

A high-performant denoiser

⇒ *2-step factorized optimization*

Meinhardt et al. ICCV 2017

Optimization with "plug and play priors"

General recipe: take a model-based optimization algorithm, "plug in" a high-performant denoiser (e.g. trained deep neural network) as the Prox

Define measurement model with image of interest $\mathbf{x}$:

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n}$$

$$\hat{\mathbf{x}} = \text{argmin}_{\mathbf{x}} \ \frac{1}{2}||\mathbf{y} - \mathbf{A}\mathbf{x}||_2^2 - \log p(\mathbf{x})$$

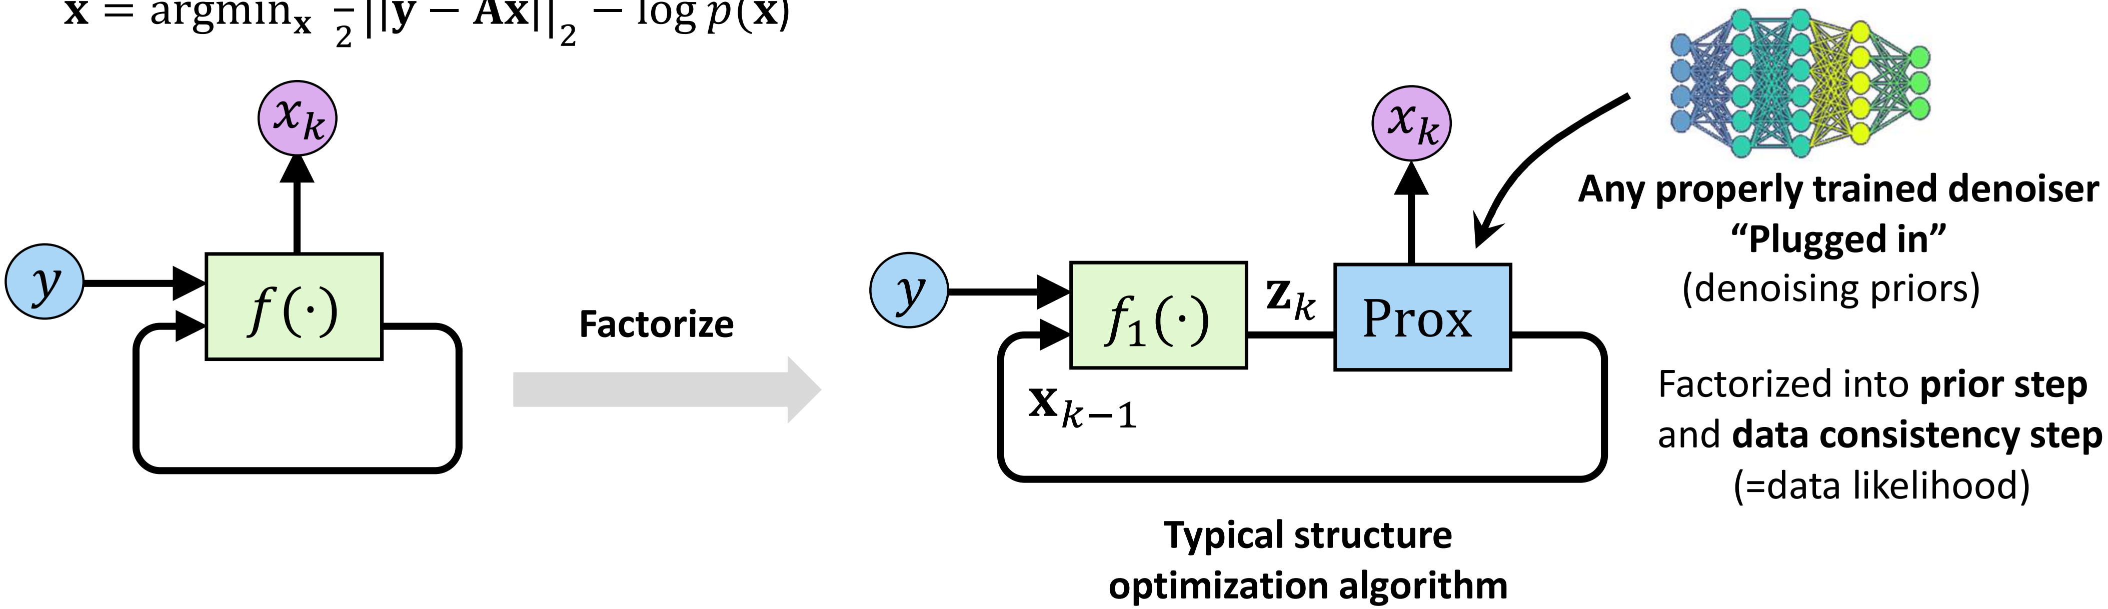

Meinhardt et al. ICCV 2017

* Properly trained: under a certain Lipschitz condition (Ryu et al., 2019)

Deep unfolding & end-to-end training

General recipe: take a model-based optimization algorithm, <u>unfold it as a fixed-complexity graph, learn proximal parameters.</u>

Define measurement model with image of interest $\mathbf{x}$:

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n}$$

$$\hat{\mathbf{x}} = \text{argmin}_{\mathbf{x}} \ \frac{1}{2}||\mathbf{y} - \mathbf{A}\mathbf{x}||_2^2 - \log p(\mathbf{x})$$

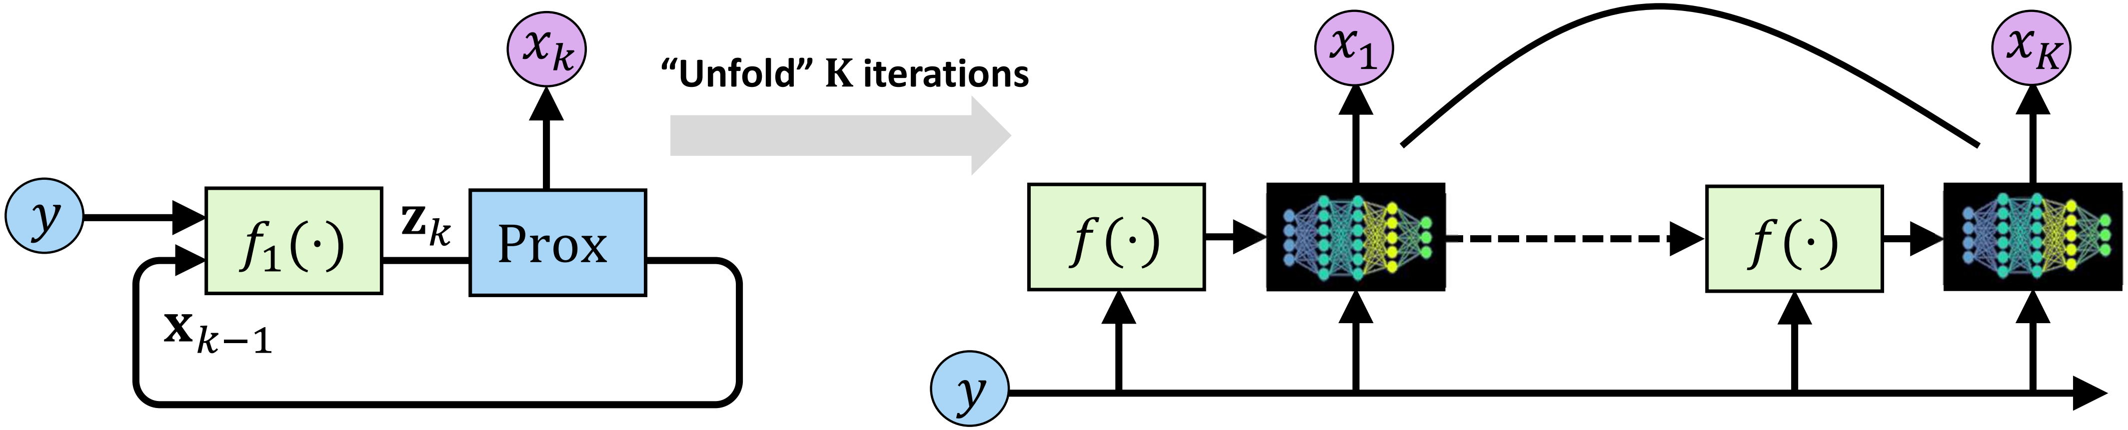

Typical structure optimization algorithm

Mardani et al. NeurIPS, 2018

Deep unfolding & end-to-end training

General recipe: take a model-based optimization algorithm, unfold it as a fixed-complexity graph, learn all parameters.

Define measurement model with image of interest $\mathbf{x}$:

$$\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n}$$

$$\hat{\mathbf{x}} = \text{argmin}_{\mathbf{x}} \frac{1}{2}||\mathbf{y} - \mathbf{A}\mathbf{x}||_2^2 - \log p(\mathbf{x})$$

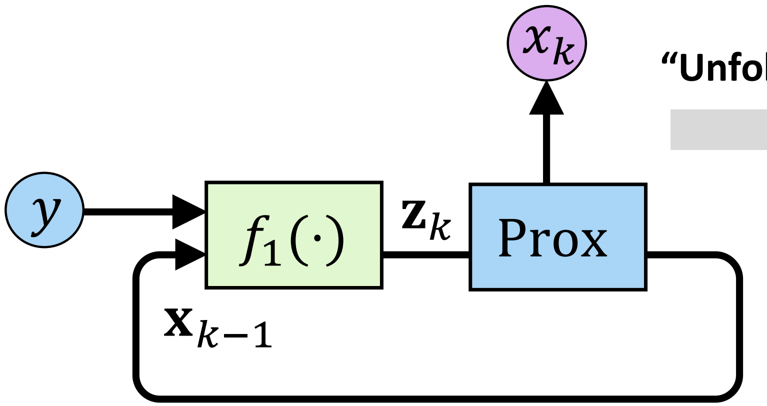

Typical structure optimization algorithm

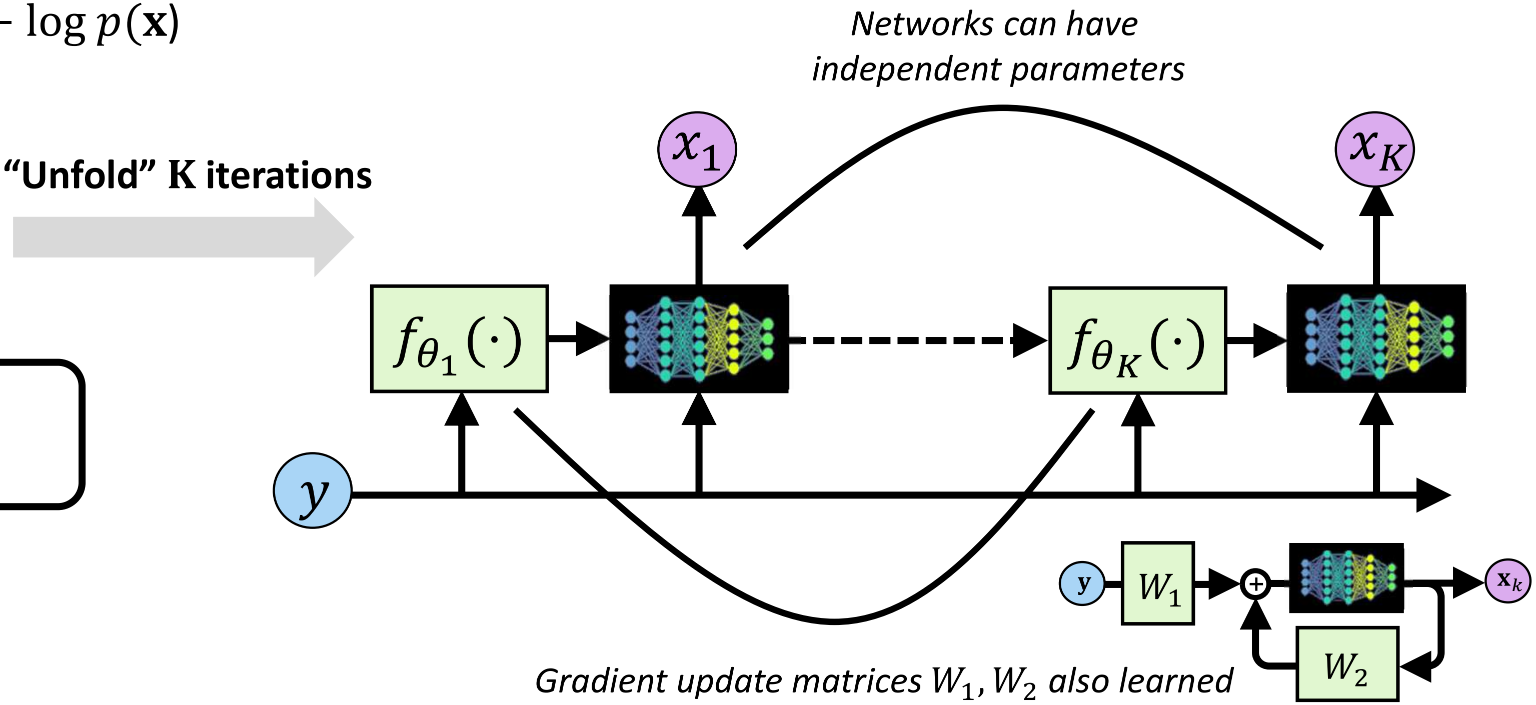

Mardani et al. NeurIPS, 2018

Combining models, priors and deep learning for image reconstruction: use-cases

Magnetic Resonance Imaging
(faster, low-field etc.)

Computed Tomography
(low-dose, sparse-view)

Ultrasound
(high-quality, reduced data-rates)

TU/e EINDHOVEN UNIVERSITY OF TECHNOLOGY

Combining models, priors and deep learning for image reconstruction: use-cases

Ultrasound
(high-quality, reduced data-rates)

TU/e EINDHOVEN UNIVERSITY OF TECHNOLOGY

Opportunities for ultrasound

Transmission

Channel Data

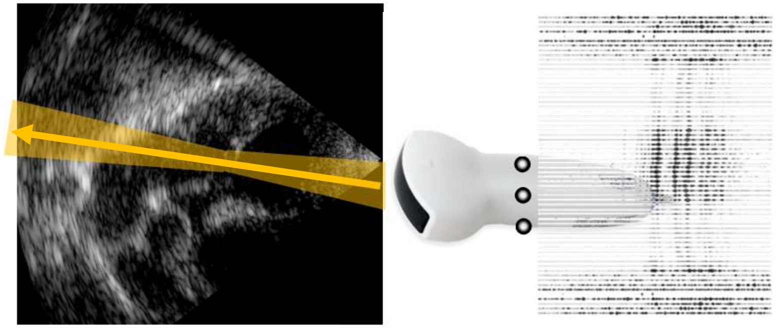

Hundreds to thousands of channels/firing

High rate ADC
(~20-50MHz / element)

High-data rates

Digital Beamforming

Spatial filtering

Advanced signal processing applications

Ultrasound localization microscopy
Shear wave imaging
Doppler/strain
Contrast agents
....

AI opportunities across the entire imaging chain

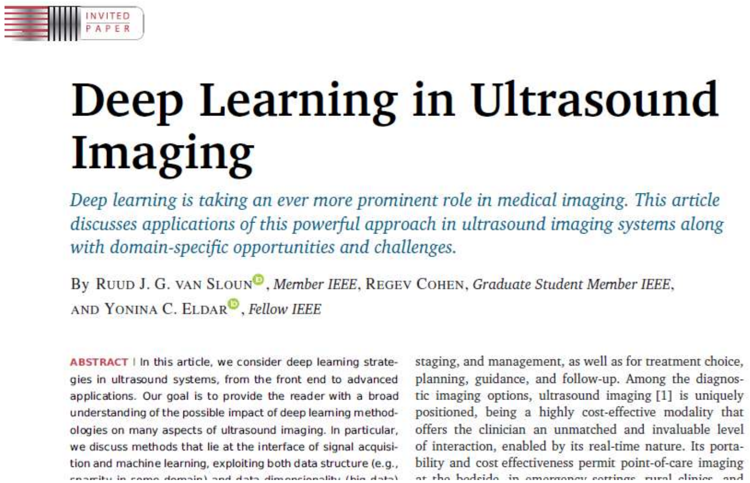

INVITED PAPER

Deep Learning in Ultrasound Imaging

Deep learning is taking an ever more prominent role in medical imaging. This article discusses applications of this powerful approach in ultrasound imaging systems along with domain-specific opportunities and challenges.

By Ruud J. G. van Sloun, *Member IEEE*, Regev Cohen, *Graduate Student Member IEEE*, and Yonina C. Eldar, *Fellow IEEE*

ABSTRACT | In this article, we consider deep learning strategies in ultrasound systems, from the front end to advanced applications. Our goal is to provide the reader with a broad understanding of the possible impact of deep learning methodologies on many aspects of ultrasound imaging. In particular, we discuss methods that lie at the interface of signal acquisition and machine learning, exploiting both data structure (e.g.,

staging, and management, as well as for treatment choice, planning, guidance, and follow-up. Among the diagnostic imaging options, ultrasound imaging [1] is uniquely positioned, being a highly cost-effective modality that offers the clinician an unmatched and invaluable level of interaction, enabled by its real-time nature. Its portability and cost effectiveness permit point-of-care imaging

Van Sloun, Cohen, Eldar, Proceedings of the IEEE, 2019

TU/e EINDHOVEN UNIVERSITY OF TECHNOLOGY

Jensen *et al.*

Adaptive beamforming by MVDR

Geometry-based time-space migraton + model-based adaptive apodization

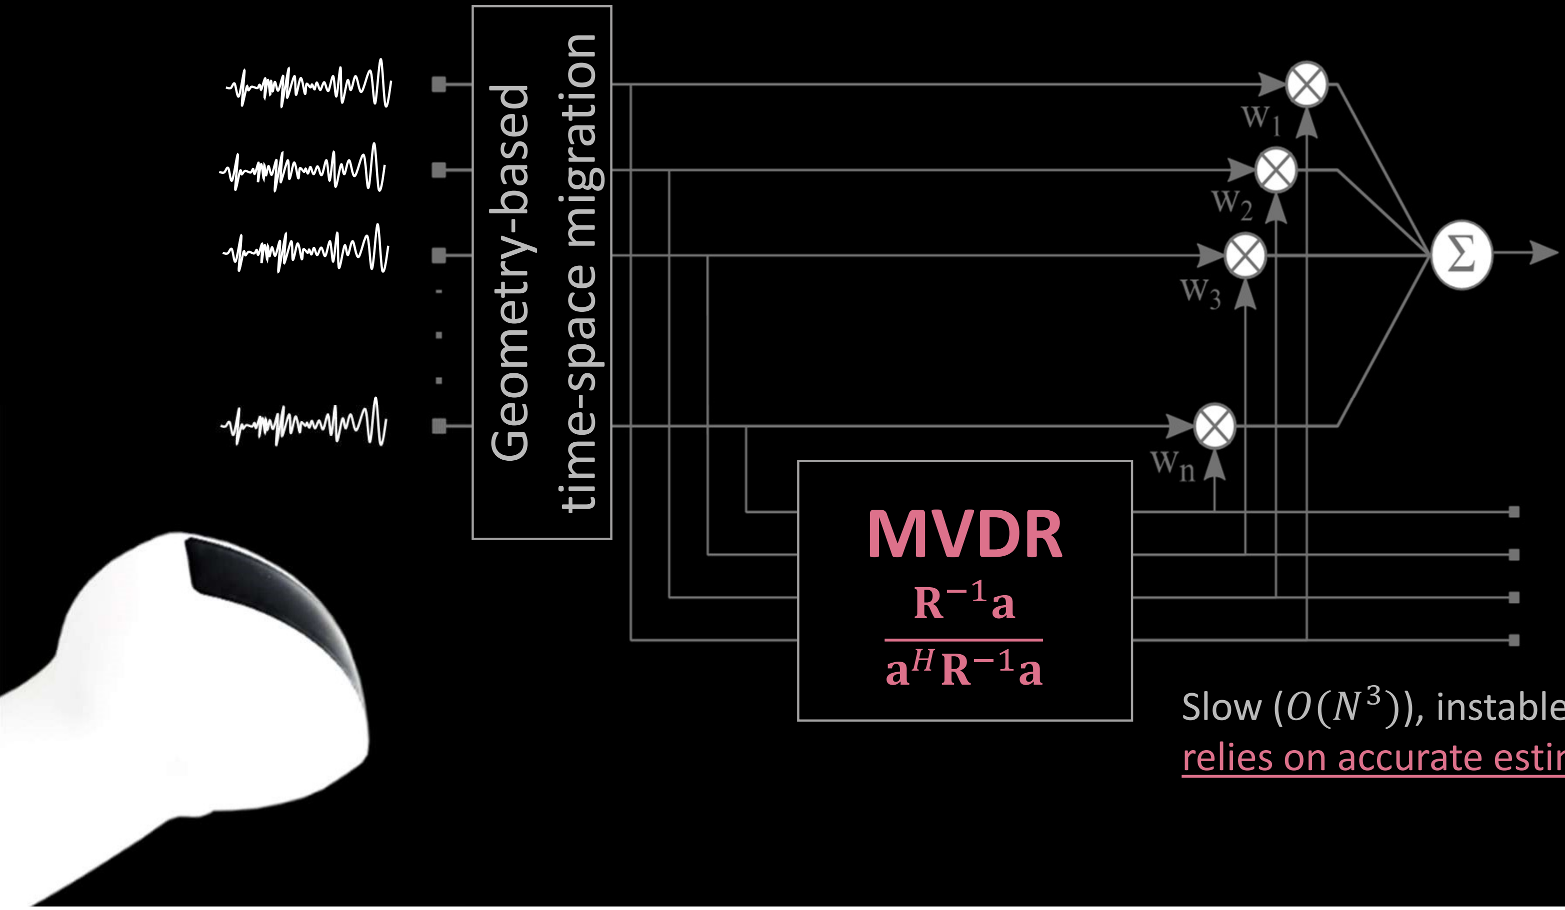

Slow ($O(N^3)$), instable matrix inversions,
relies on accurate estimates of statistics (model knowledge)

Luijten *et al.*, *IEEE ICASSP*, 2019 Luijten *et al.*, *IEEE trans. Med. Imag.*, 2020

Adaptive beamforming by deep learning (ABLE)

Hybrid inference: Model-based computational graph with integrated NN

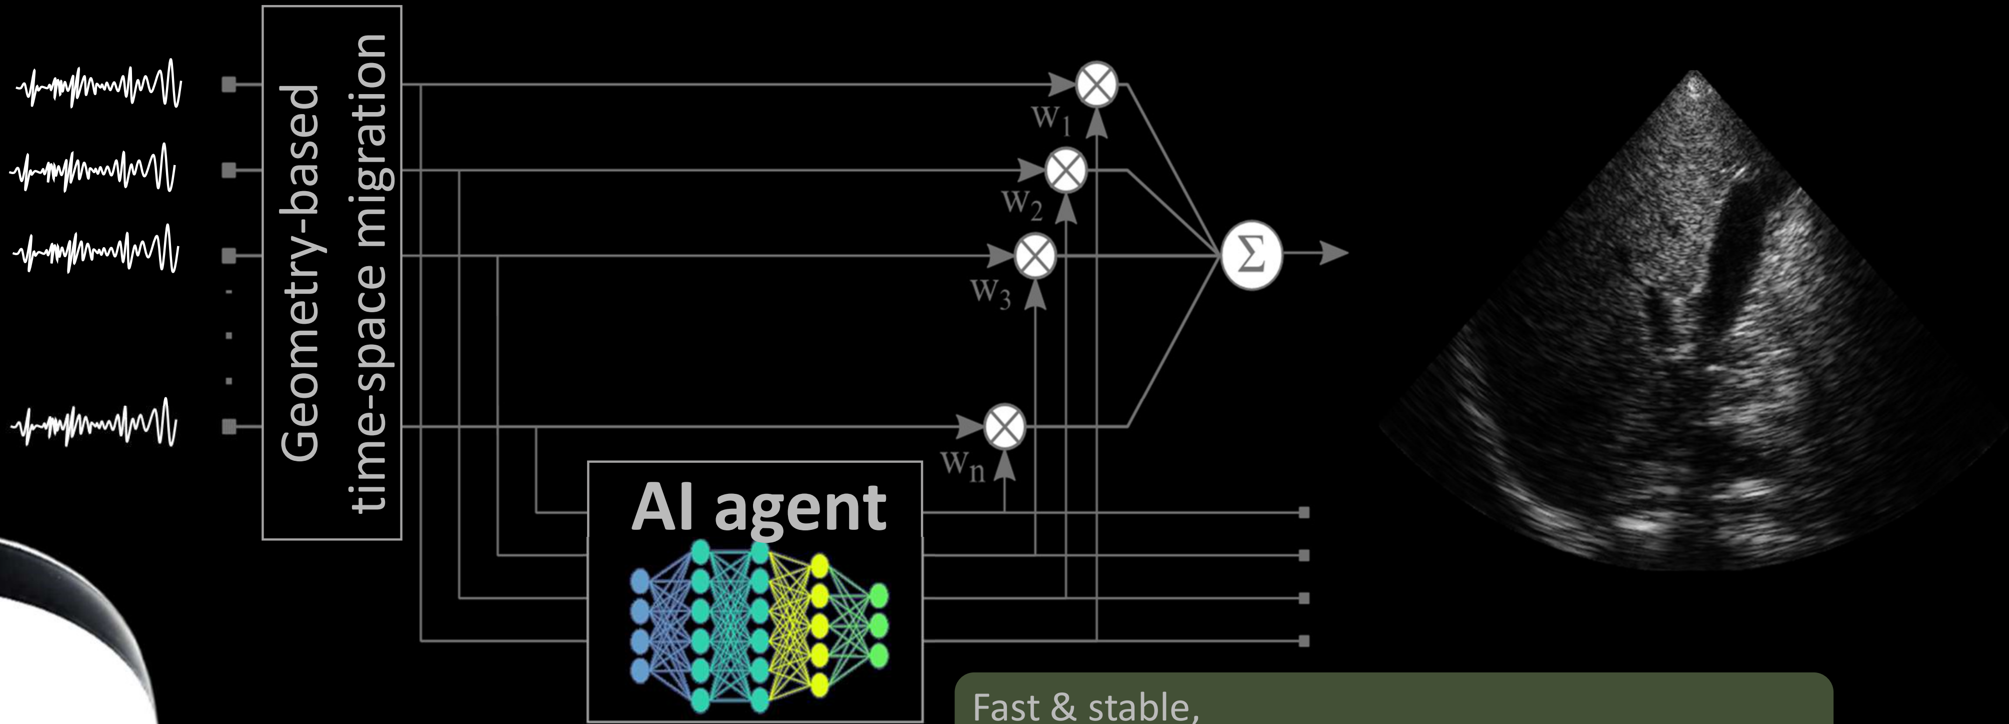

Fast & stable,
learns actionable statistics & function from data

PHILIPS TU/e EINDHOVEN UNIVERSITY OF TECHNOLOGY מכון ויצמן למדע WEIZMANN INSTITUTE OF SCIENCE

Adaptive beamforming by deep learning (ABLE)

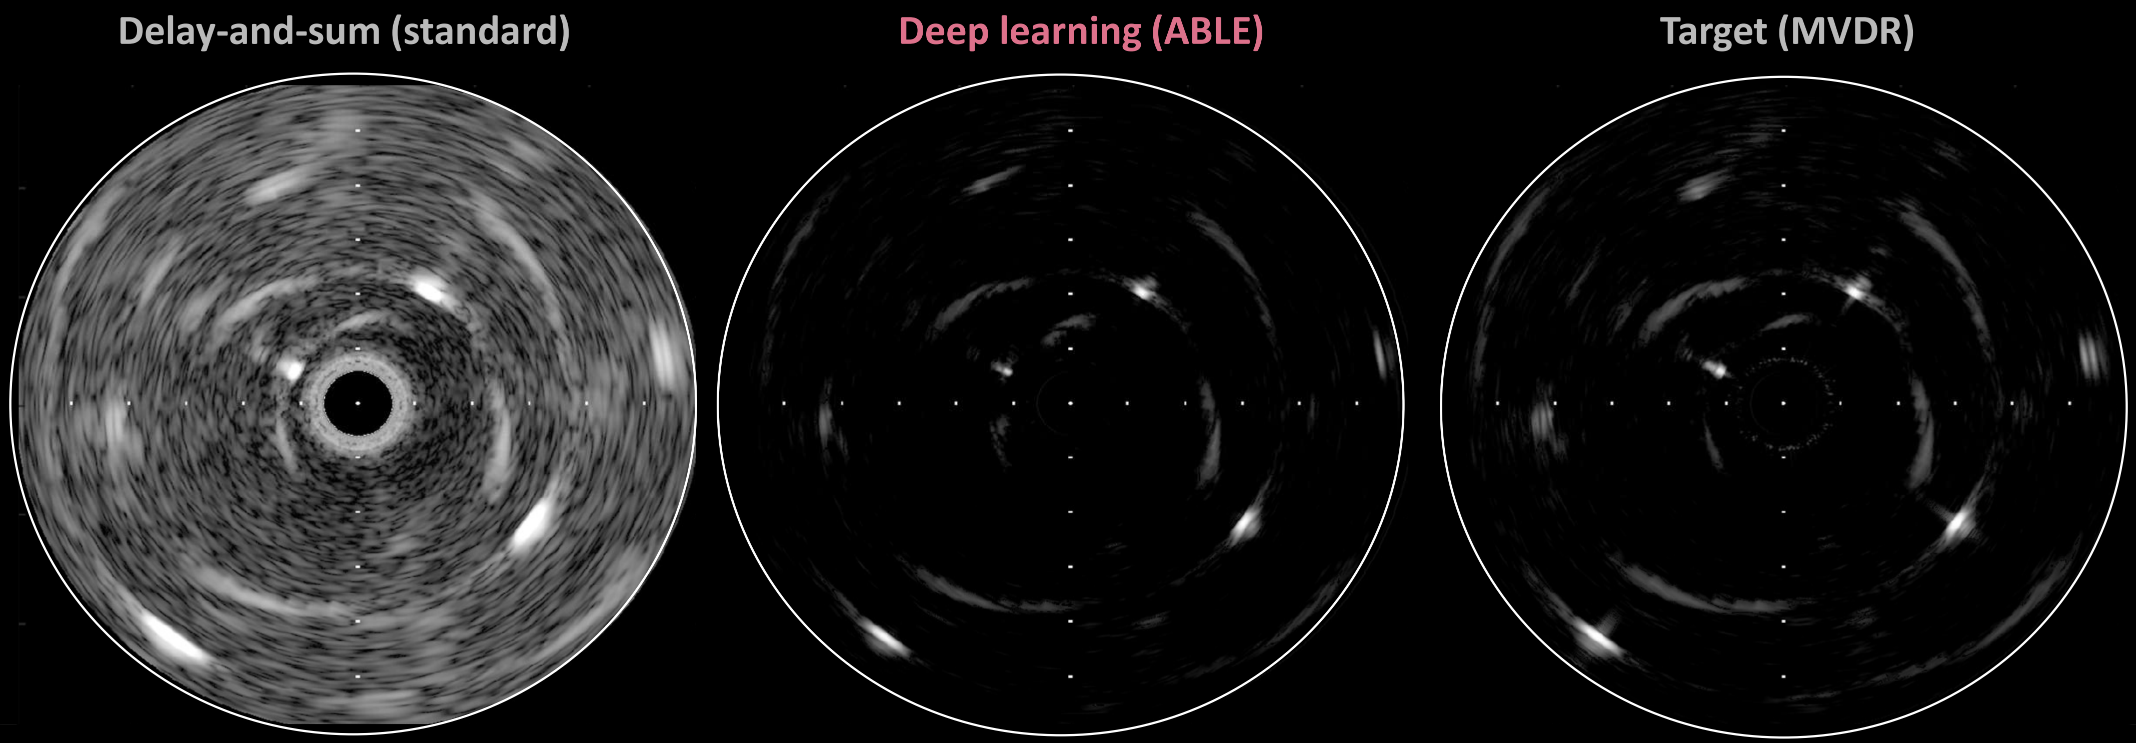

Note: without post-processing and s-curve
60 dB, log-scale

PHILIPS TU/e EINDHOVEN UNIVERSITY OF TECHNOLOGY מכון ויצמן למדע WEIZMANN INSTITUTE OF SCIENCE

Adaptive beamforming by deep learning (ABLE)

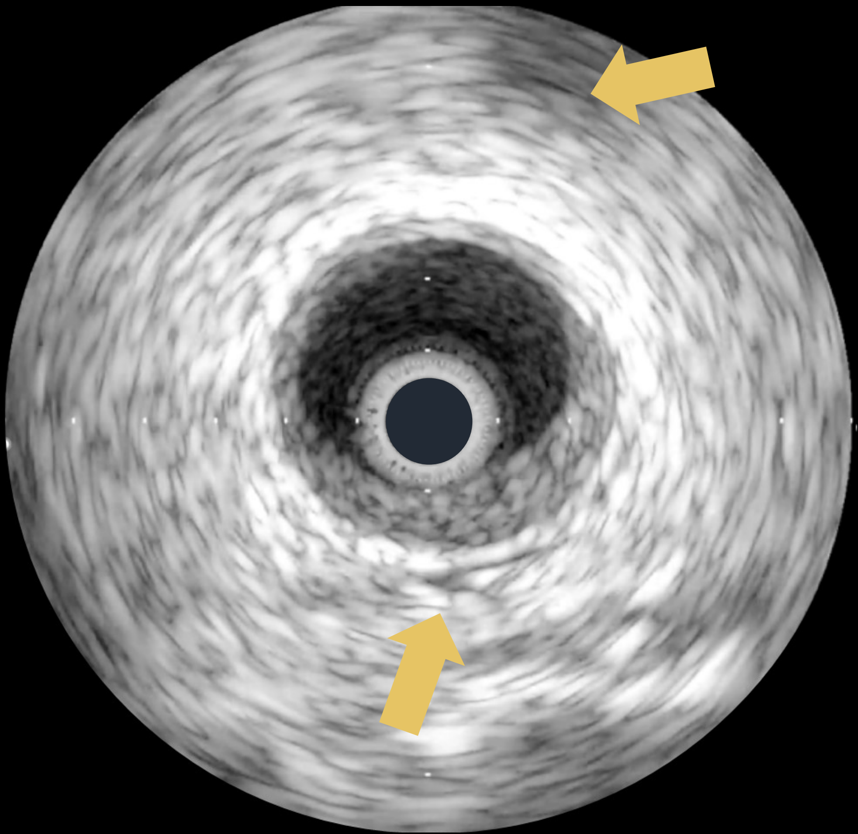

Standard

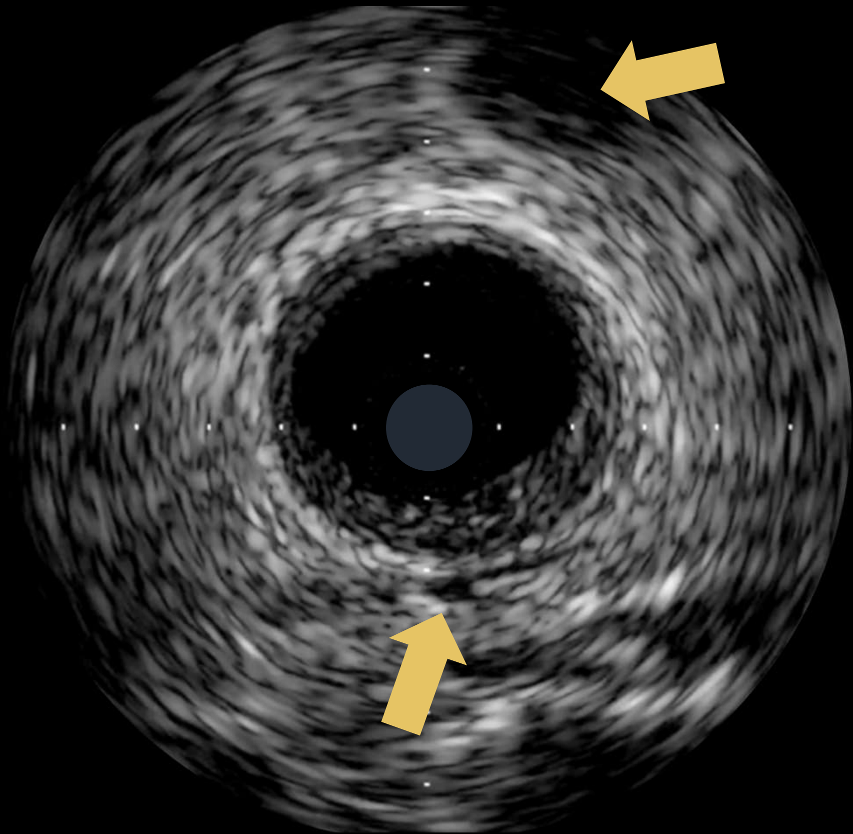

Deep learning (ABLE)

Less clutter
Higher resolution

High processing rates,
and robustness

Note: without post-processing and s-curve
60 dB, log-scale

PHILIPS TU/e EINDHOVEN UNIVERSITY OF TECHNOLOGY מכון ויצמן למדע WEIZMANN INSTITUTE OF SCIENCE

Chennakeshava *et al.*, *IEEE IUS*, 2020

High-resolution plane-wave compounding

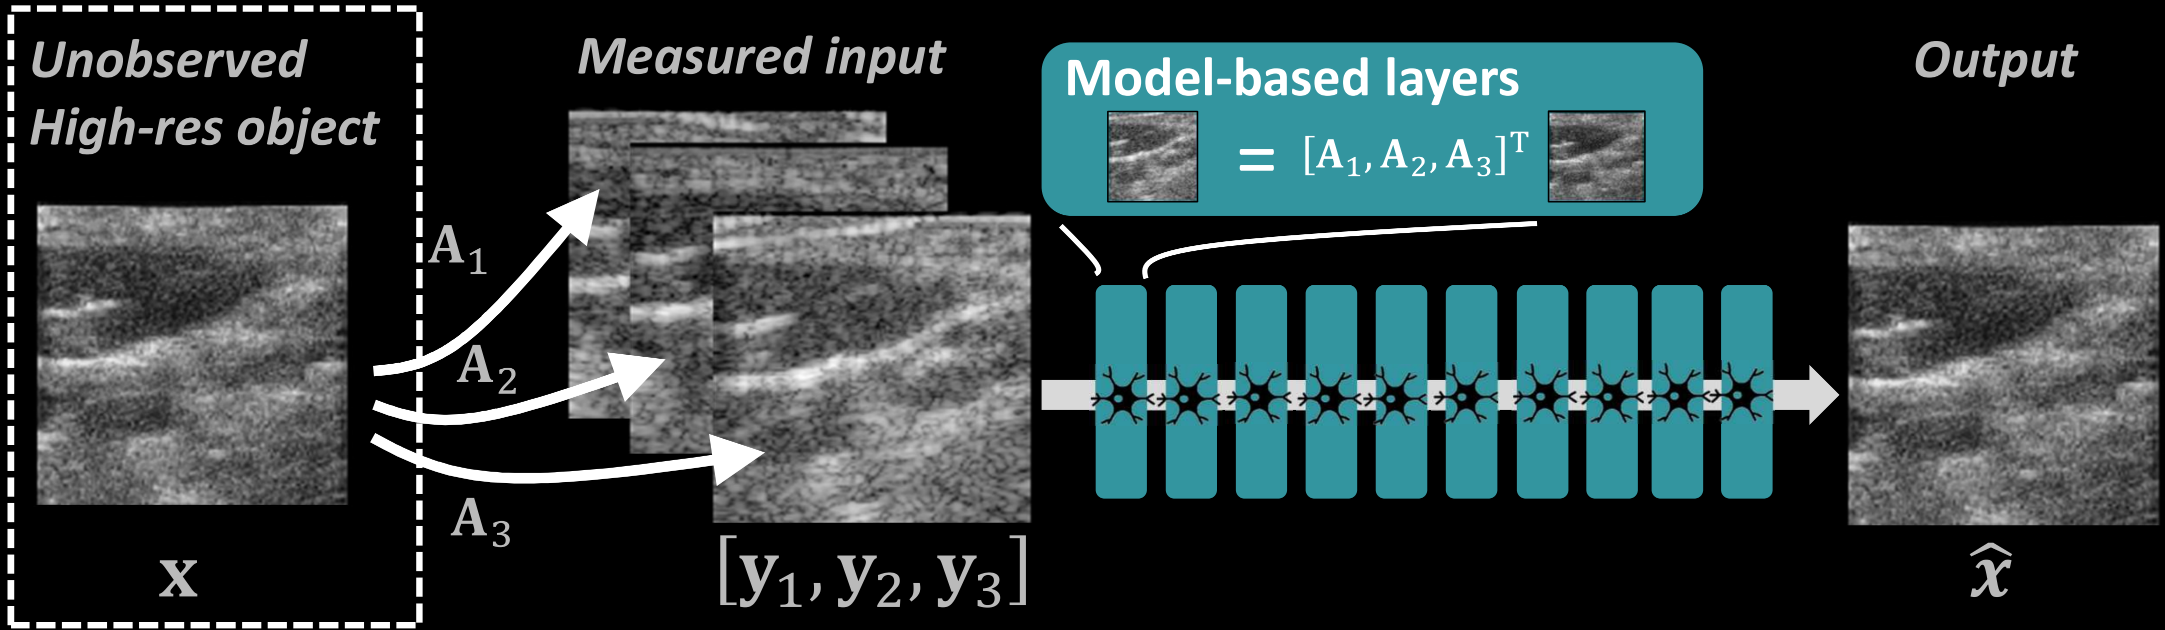

Acquisition model

Solve inverse problem:

Deep unfolding with end-to-end training

Chennakeshava *et al.*, *IEEE IUS*, 2020

High-resolution plane-wave compounding

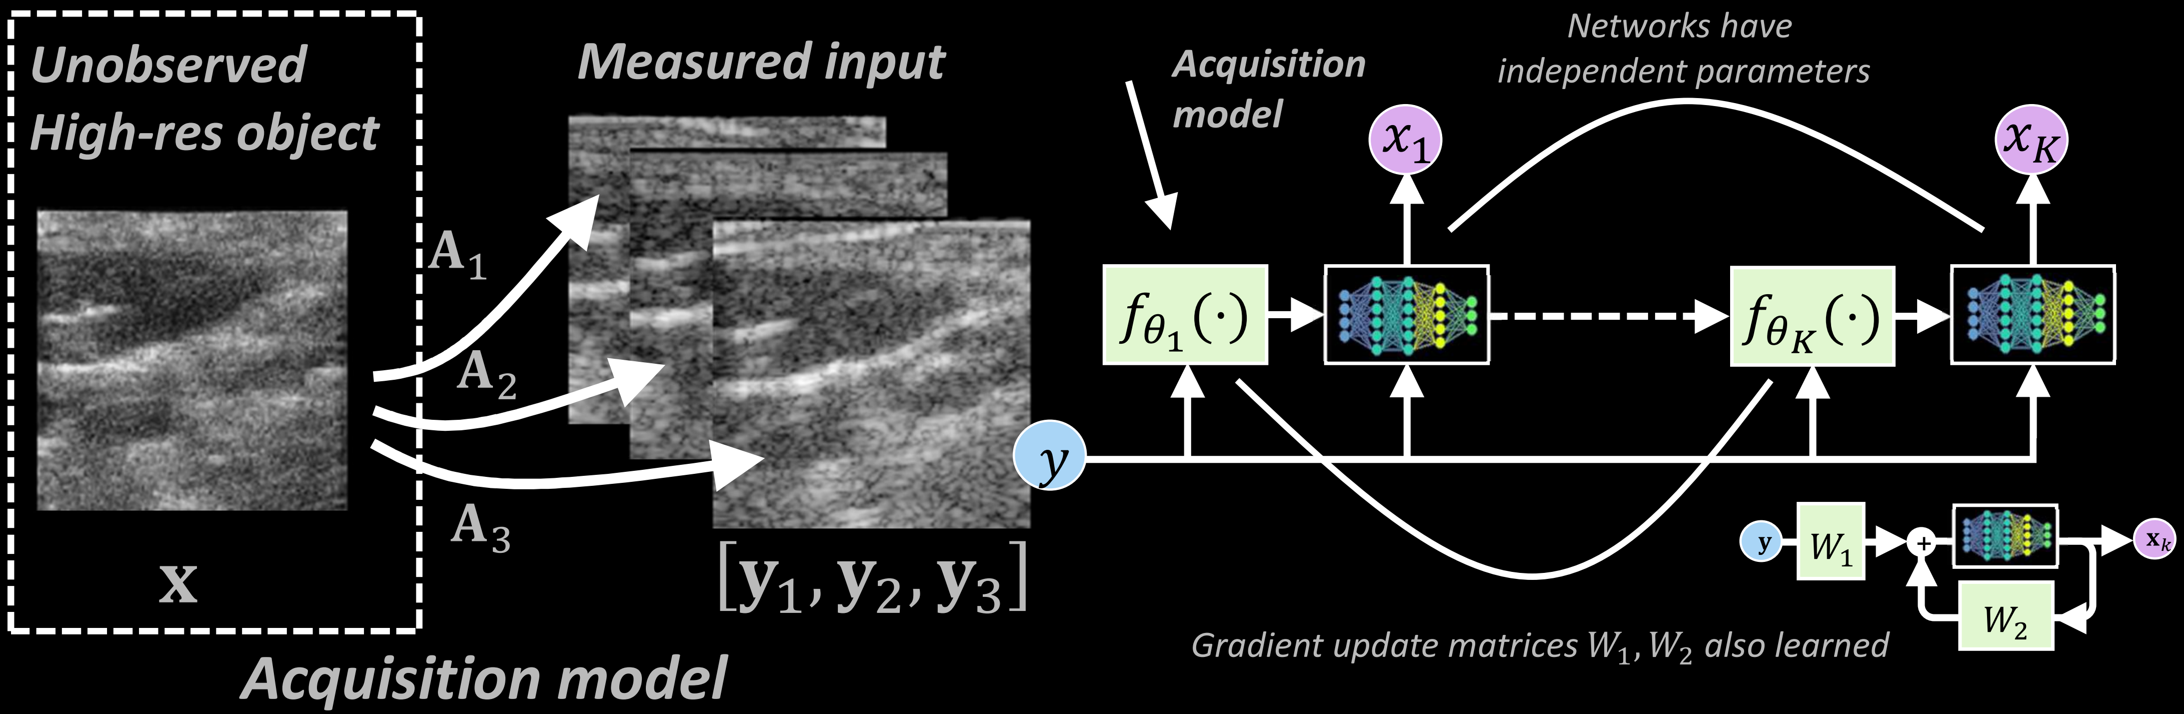

Chennakeshava *et al.*, *IEEE IUS*, 2020

High-resolution plane-wave compounding

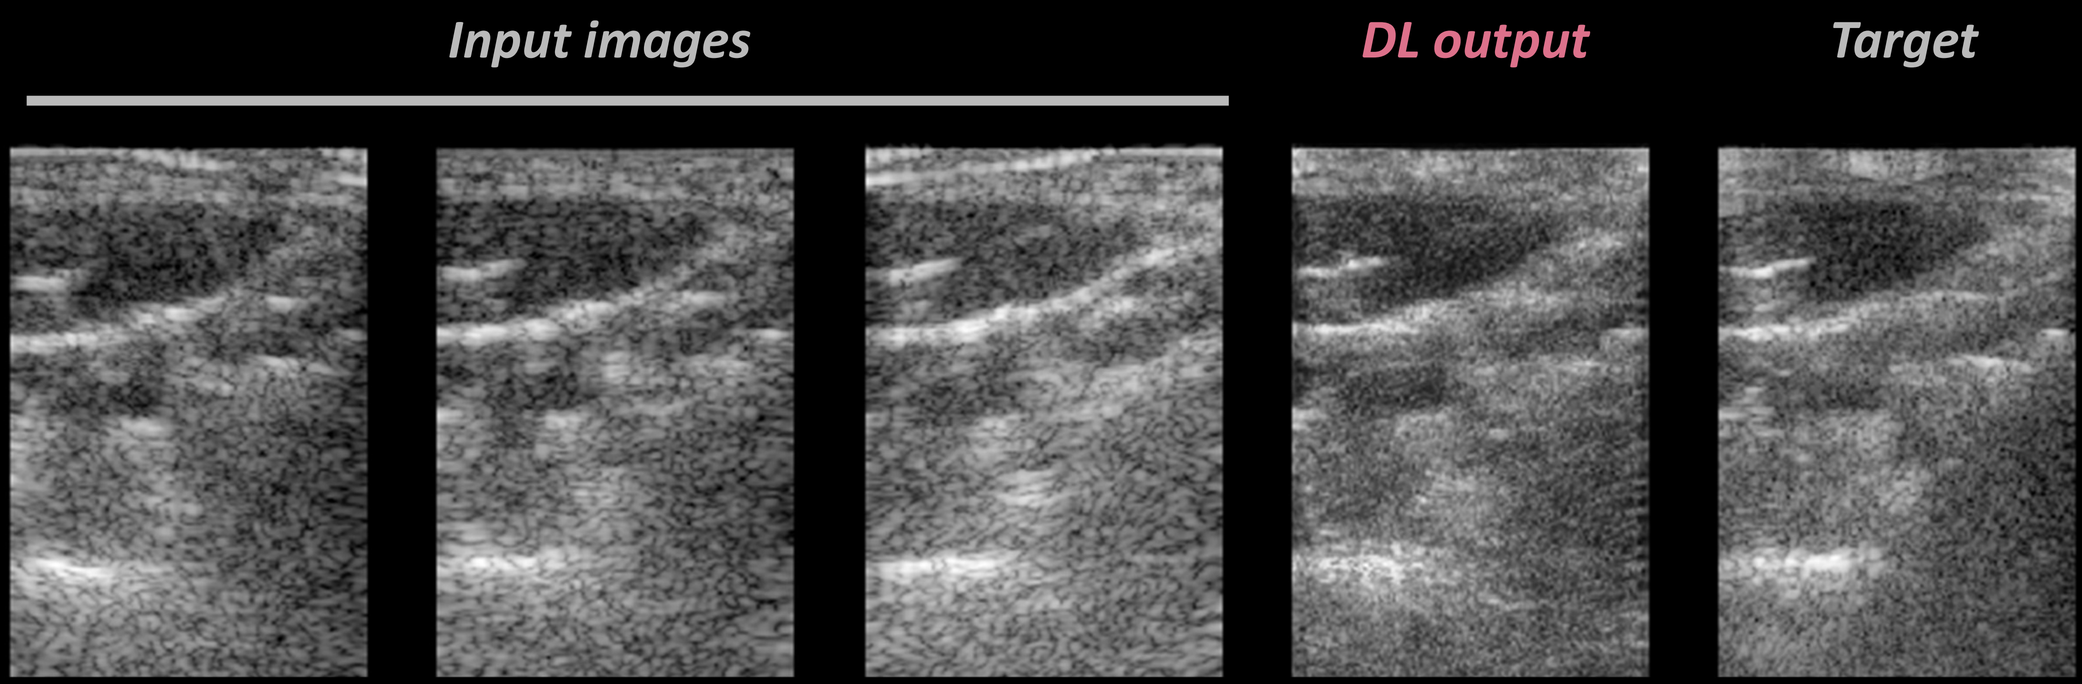

Solomon *et al. IEEE trans. Med. Imag.*, 2019
van Sloun *et al. Proceedings of the IEEE,* 2020

Spatiotemporal source-extraction/dehazing

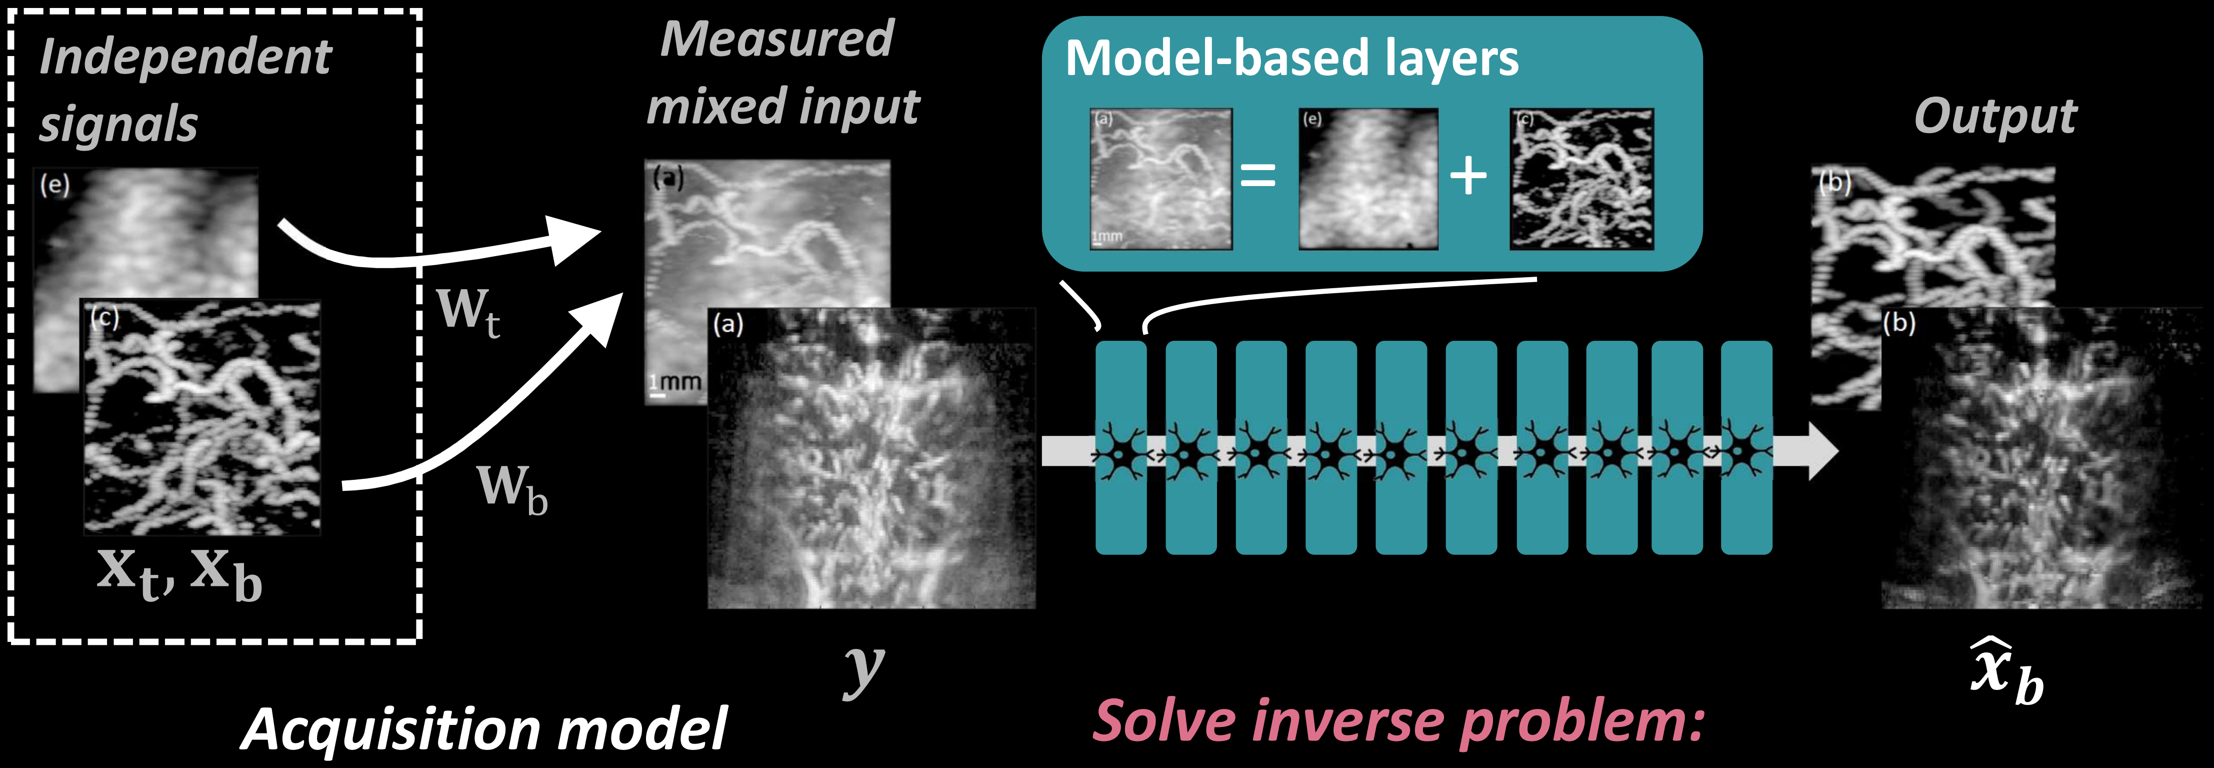

Solve inverse problem:
Deep unfolding with end-to-end training

Solomon *et al. IEEE trans. Med. Imag.*, 2019
van Sloun *et al. Proceedings of the IEEE*, 2020

Spatiotemporal source-extraction/dehazing

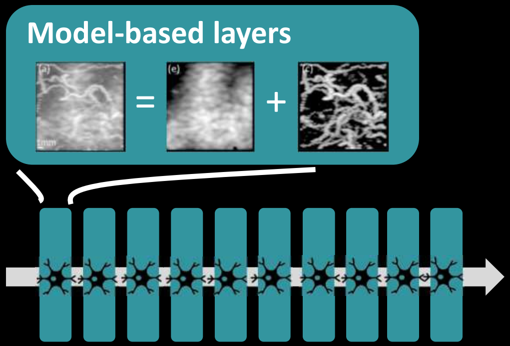

Optimization problem:

$$\min_{\mathbf{L},\mathbf{S}} \frac{1}{2}||\mathbf{D}-\mathbf{L}-\mathbf{S}||_F^2 + \lambda_1||\mathbf{L}||_* + \lambda_2||\mathbf{S}||_{1,2}$$

Prox-grad solution: (iterative)

$$\mathbf{L}^{k+1} = \mathcal{SVT}_{\lambda_1/2}\left(\frac{1}{2}\mathbf{L}^k - \mathbf{S}^k + \mathbf{D}\right)$$

$$\mathbf{S}^{k+1} = \mathcal{T}_{\lambda_2/2}\left(\frac{1}{2}\mathbf{S}^k - \mathbf{L}^k + \mathbf{D}\right)$$

Solomon *et al. IEEE trans. Med. Imag.*, 2019

van Sloun *et al. Proceedings of the IEEE*, 2020

Spatiotemporal source-extraction/dehazing

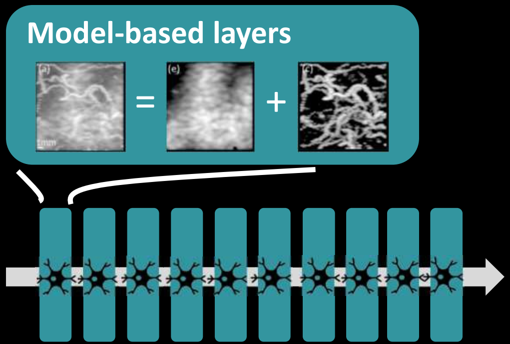

Optimization problem:

$$\min_{\mathbf{L},\mathbf{S}} \frac{1}{2}||\mathbf{D}-\mathbf{L}-\mathbf{S}||_F^2 + \lambda_1||\mathbf{L}||_* + \lambda_2||\mathbf{S}||_{1,2}$$

Prox-grad solution: (iterative)

$$\mathbf{L}^{k+1} = \mathcal{SVT}_{\lambda_1/2}\left(\frac{1}{2}\mathbf{L}^k - \mathbf{S}^k + \mathbf{D}\right)$$

$$\mathbf{S}^{k+1} = \mathcal{T}_{\lambda_2/2}\left(\frac{1}{2}\mathbf{S}^k - \mathbf{L}^k + \mathbf{D}\right)$$

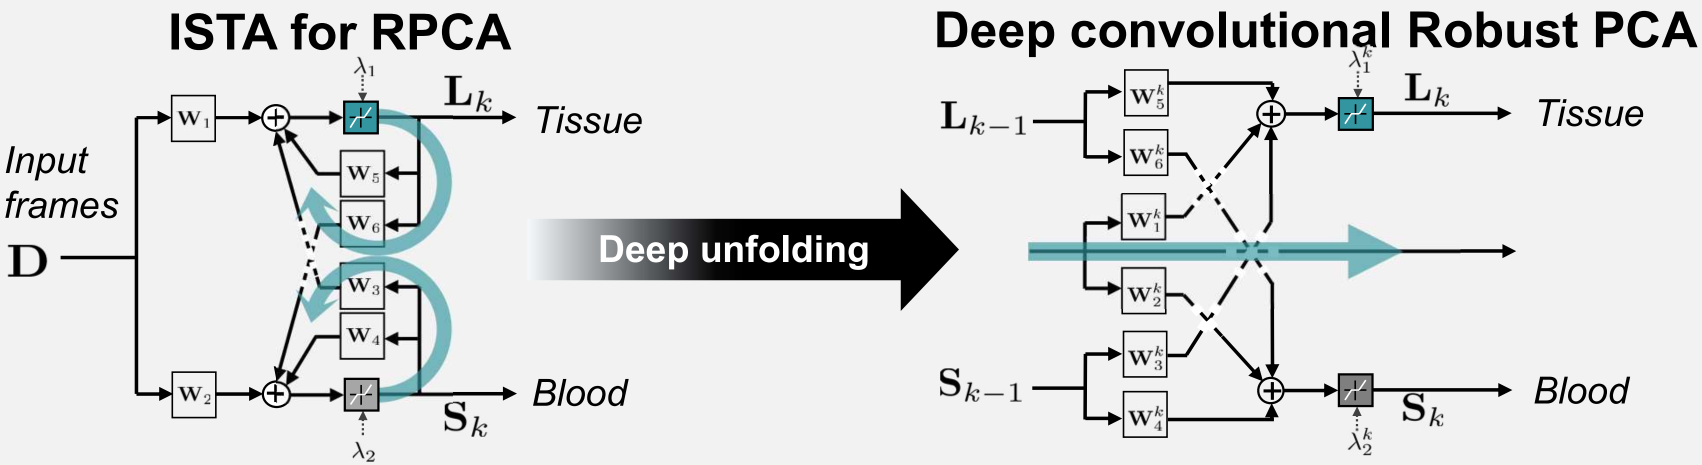

Standard contrast-ultrasound

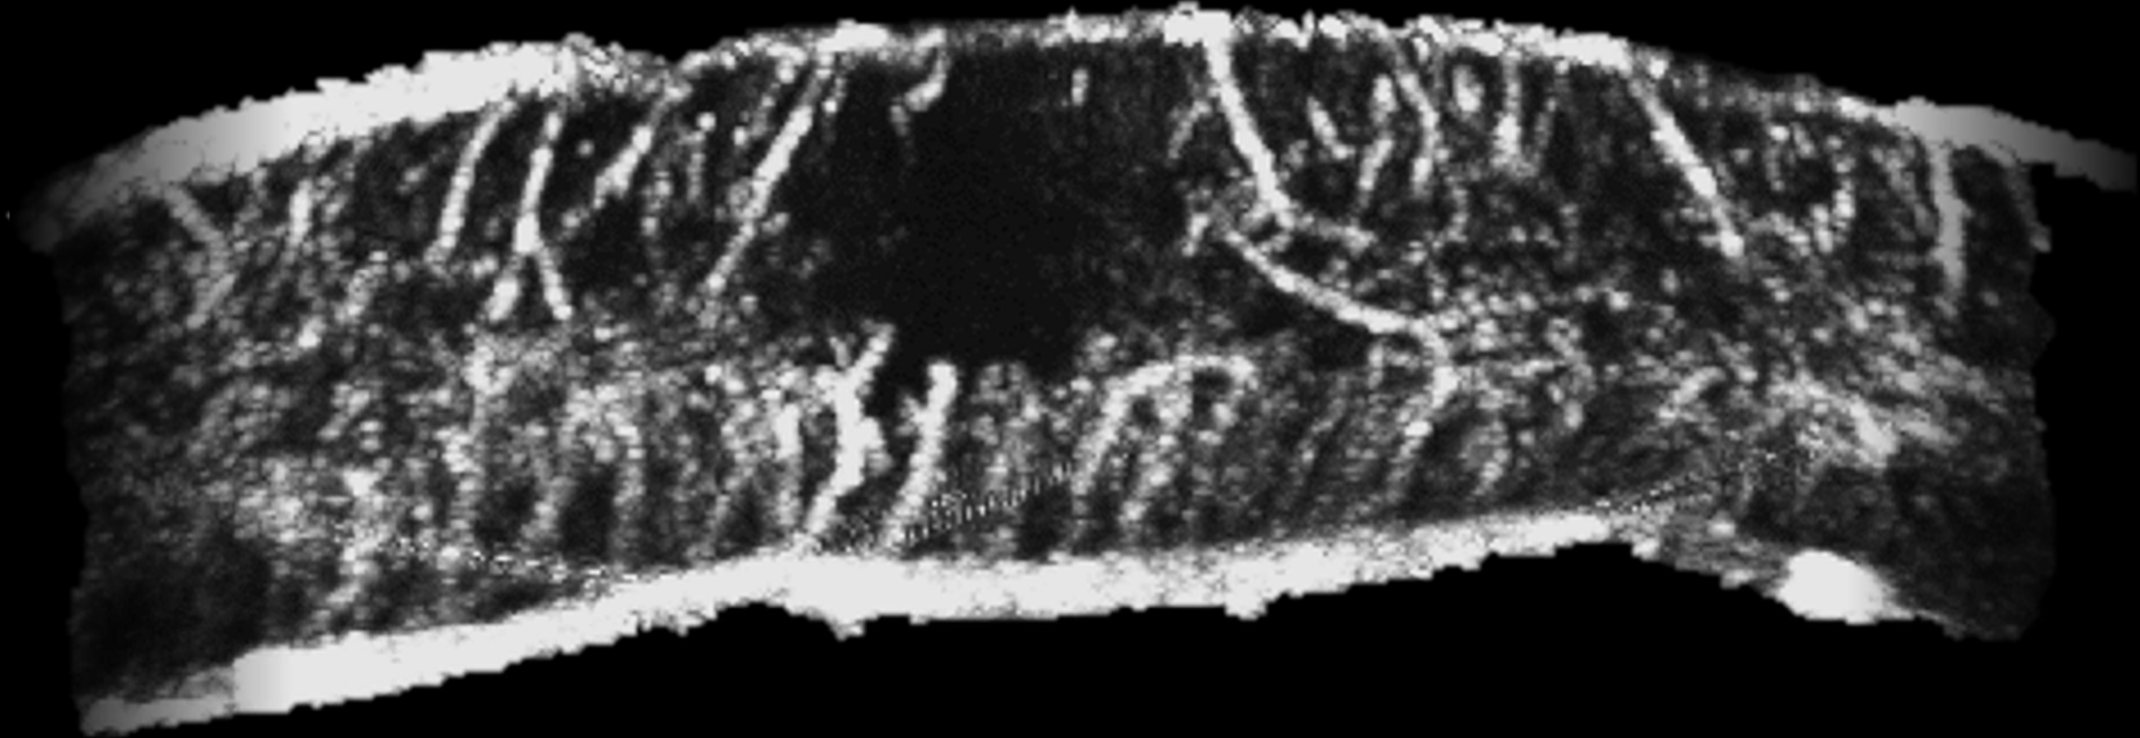

UNIVERSITY *of* WASHINGTON

Super-resolution contrast ultrsound by LISTA

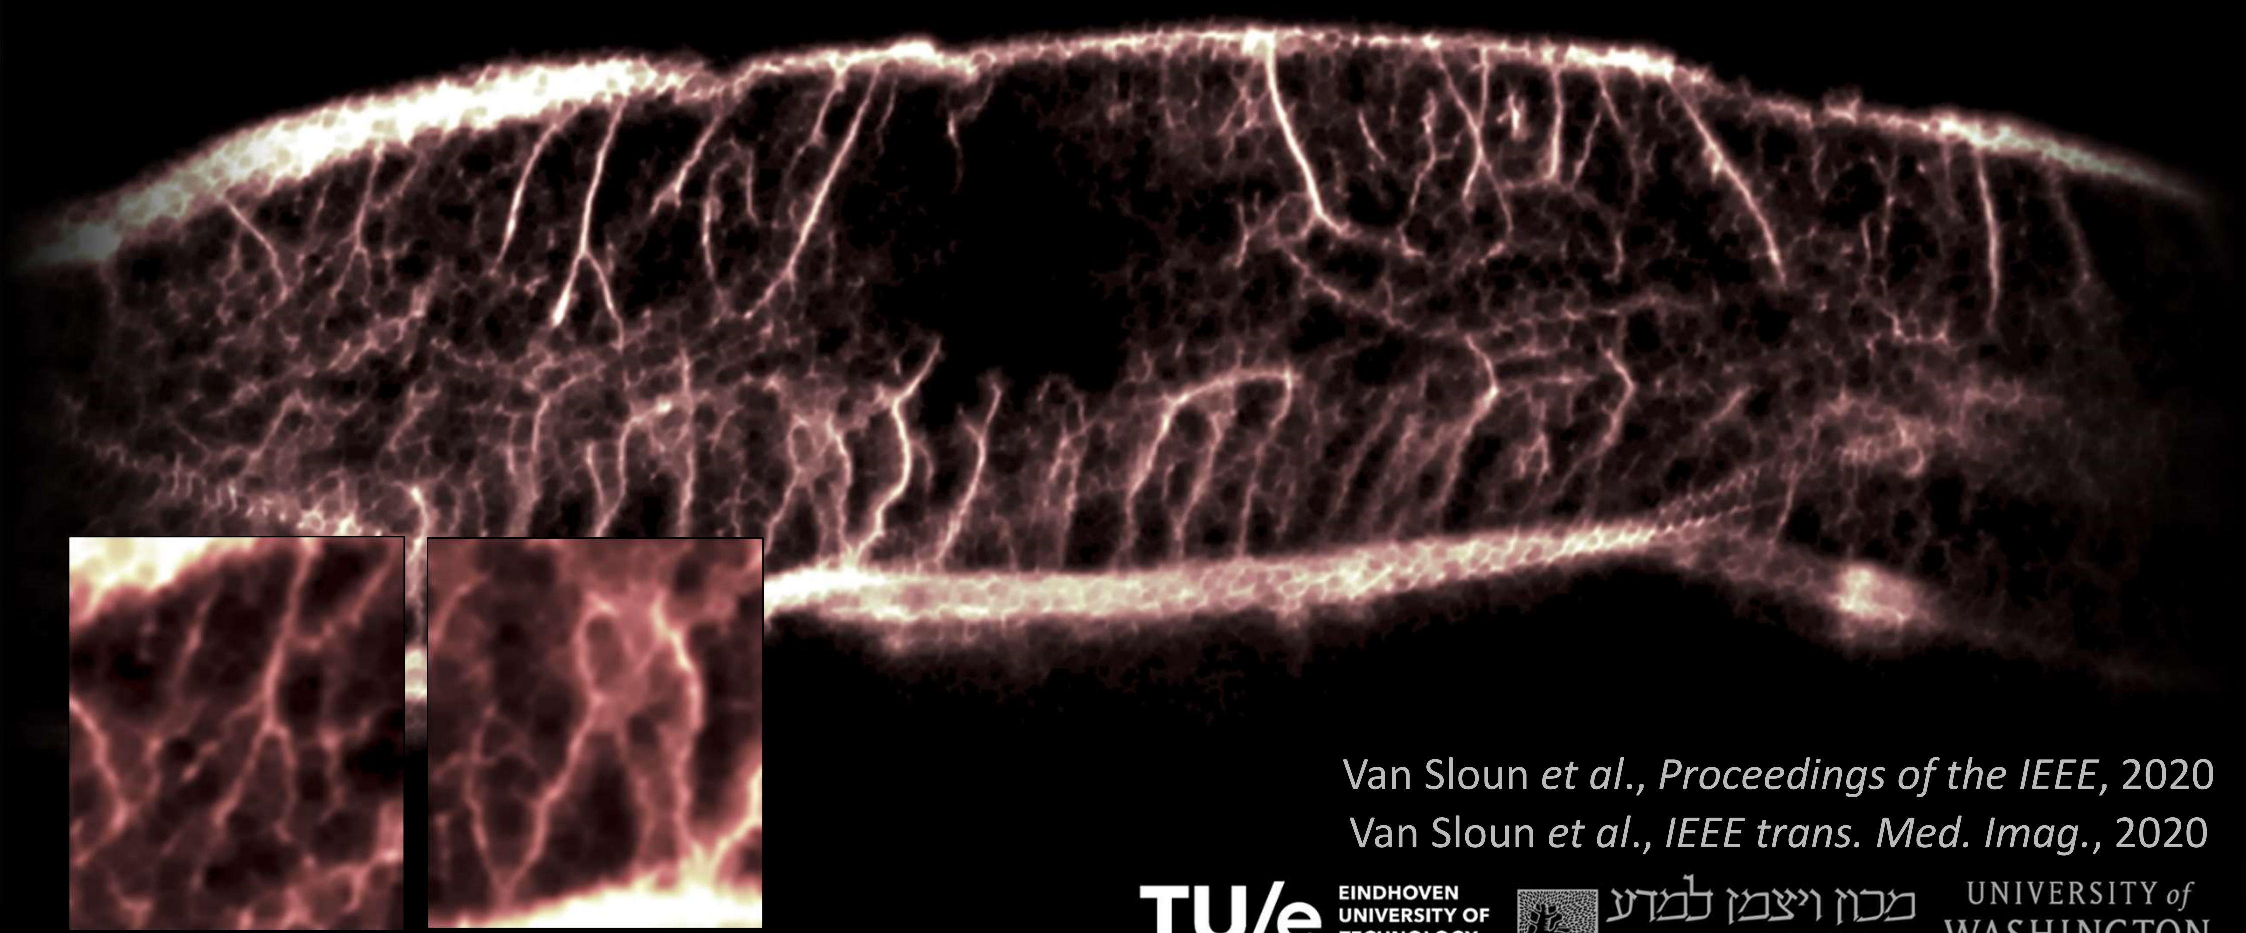

Van Sloun *et al.*, *Proceedings of the IEEE*, 2020

Van Sloun *et al.*, *IEEE trans. Med. Imag.*, 2020

TU/e EINDHOVEN UNIVERSITY OF TECHNOLOGY
WEIZMANN INSTITUTE OF SCIENCE
UNIVERSITY of WASHINGTON

Combining models, priors and deep learning for image reconstruction: use-cases

Magnetic Resonance Imaging (faster, low-field etc.)

TU/e EINDHOVEN UNIVERSITY OF TECHNOLOGY

Van Gorp *et al.*, in review Huijben *et al.*, ICASSP 2020 Huijben *et al.*, ICLR 2020

MRI: learning acquisition

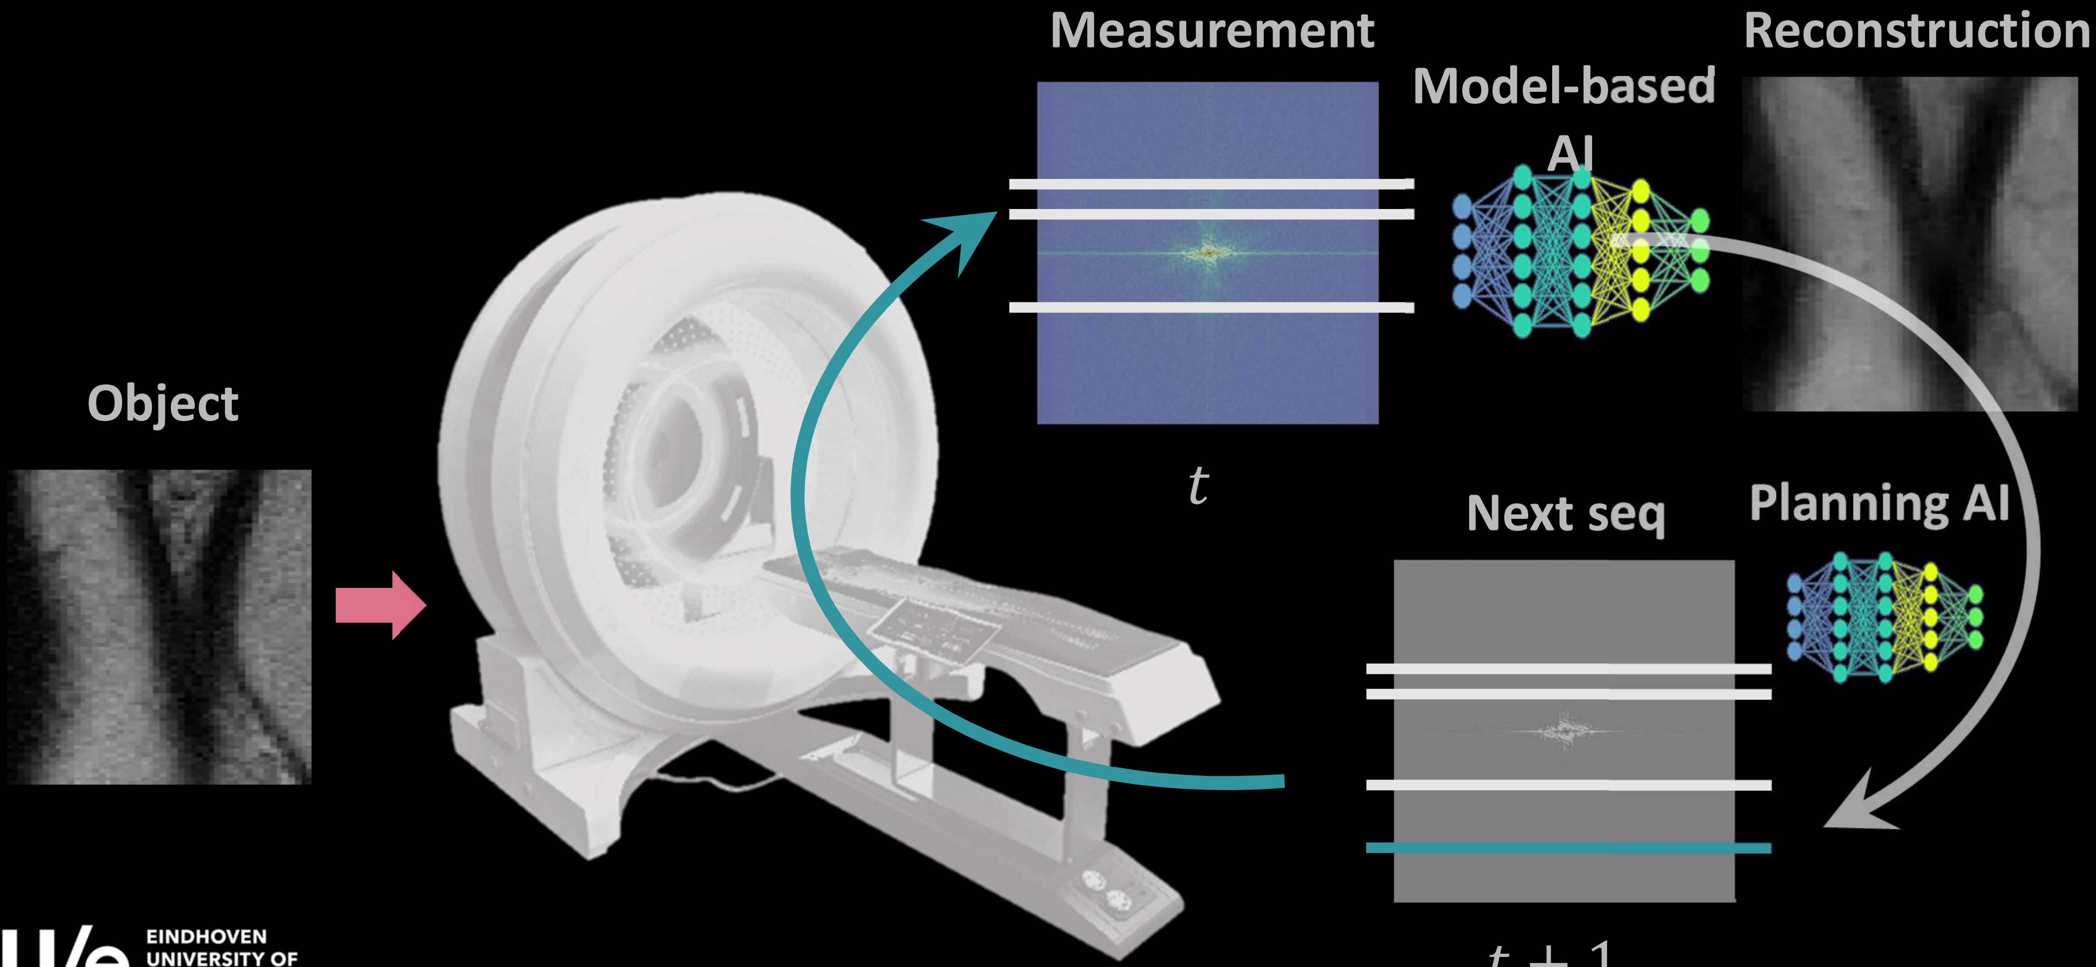

TU/e EINDHOVEN UNIVERSITY OF TECHNOLOGY

Van Gorp *et al.*, in review   Huijben *et al.*, ICASSP 2020   Huijben *et al.*, ICLR 2020

MRI: learning acquisition

Importance of model-based DL for recon:

- Factorizes knowledge
- Sampling and recon update directions 'decoupled' during learning

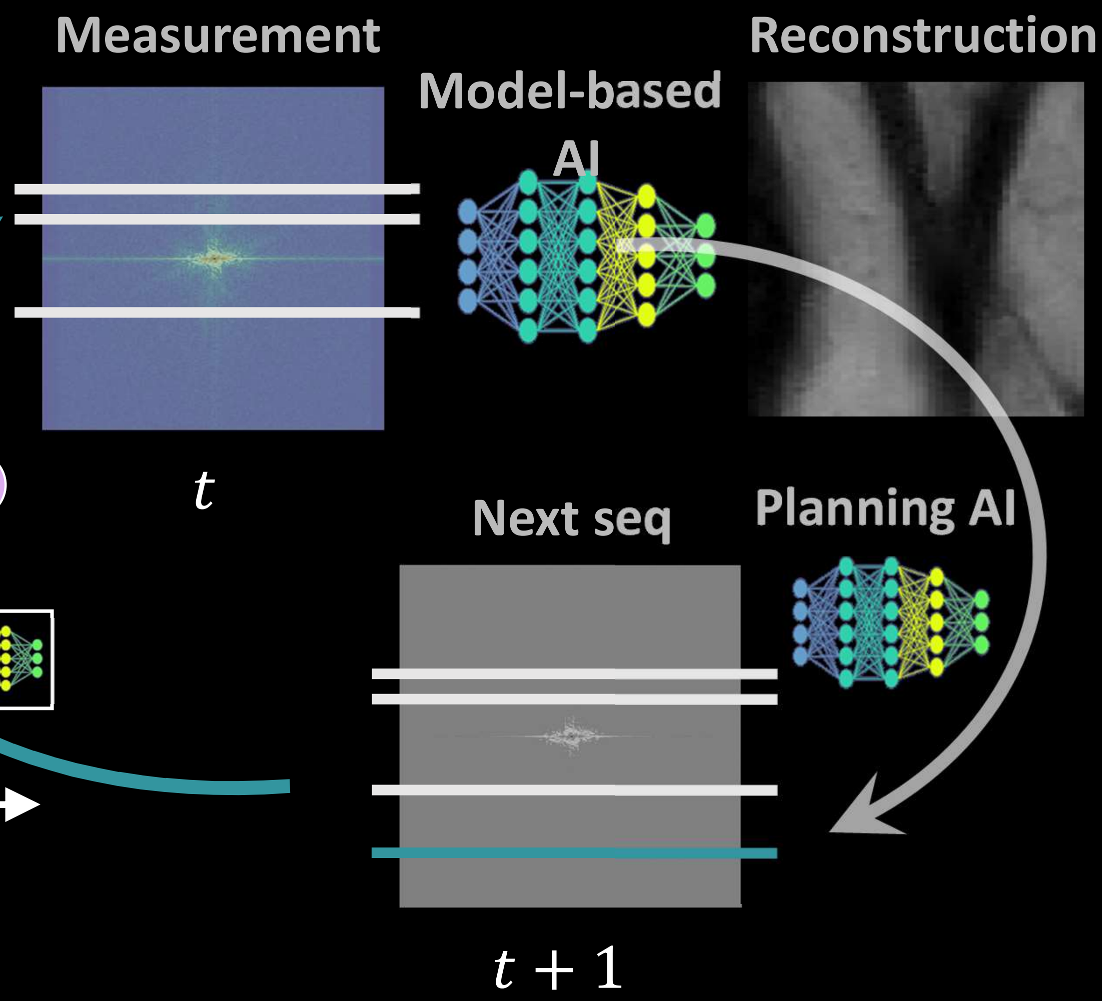

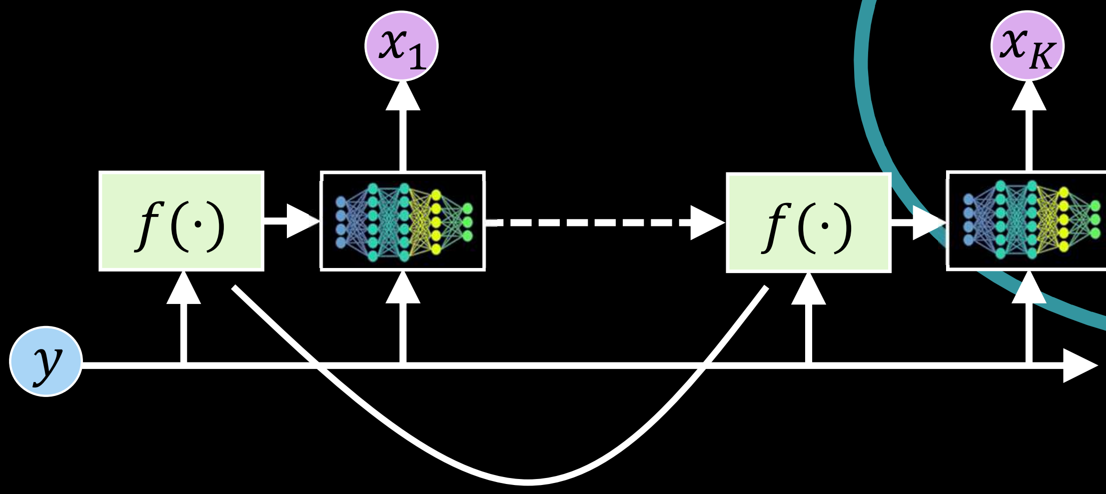

Acquisition/sampling changes directly update f()

Huijben *et al.*, ICASSP 2020   Huijben *et al.*, ICLR 2020

MRI: learning sampling density mask

Full k-space

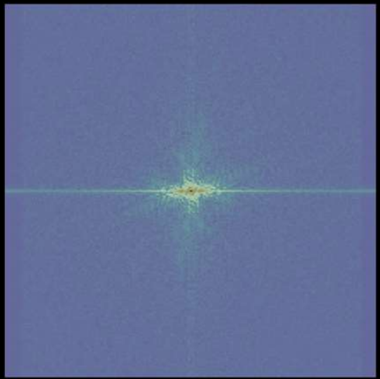

Learned mask

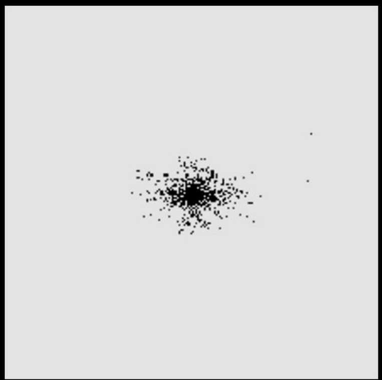

Full sampling

Sampling Low frequencies

Learned sampling

PSNR: 35.8

PSNR: 36.2

Van Gorp *et al.*, in review

MRI: learning active acquisition

Factor 8 undersampling

K-space

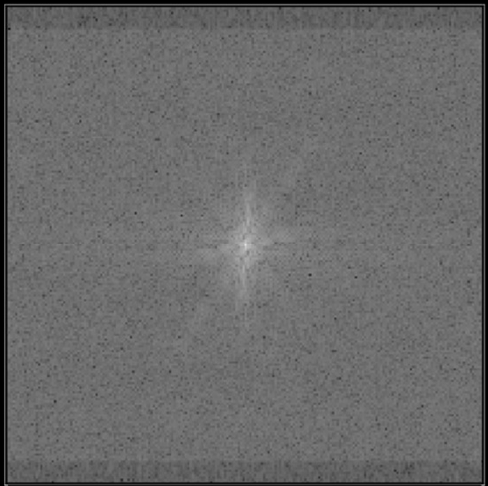

Active line sampling (cumulative)

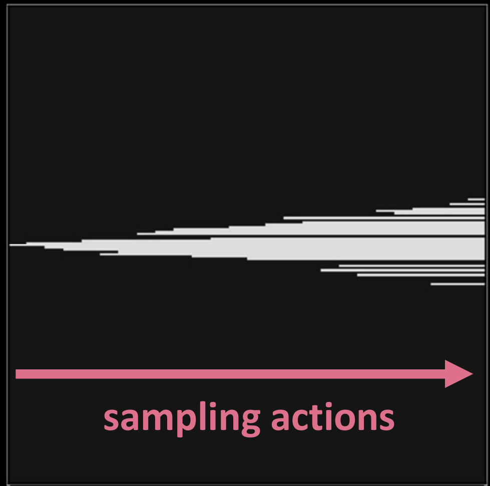

Reconstruction

Target (full acquisition)

Method	NMSE	PSNR [dB]	SSIM
Zhang et al., 2019 (active)	0.0398	28.8	0.610
Pineda et al., 2020 (active)	0.0371	29.2	0.623
Fixed learned sampling (ours)	0.0360	30.1	0.650
Active acquisition (ours)	**0.0342**	**30.2**	**0.654**

Conclusions

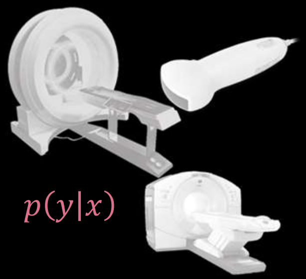

Acquisition model assumptions incorrect/imprecise

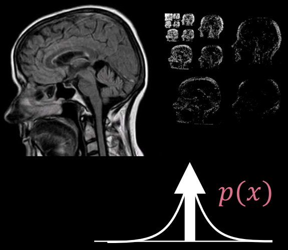

Statistical image priors not sufficiently expressive/accurate

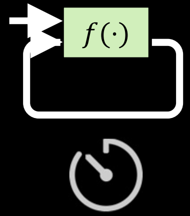

Slow reconstruction with high complexity

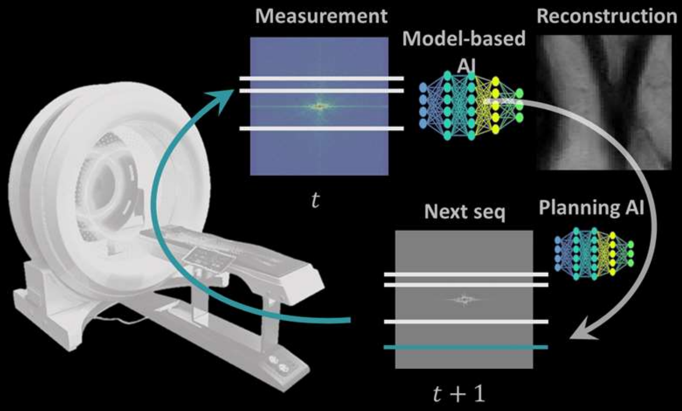

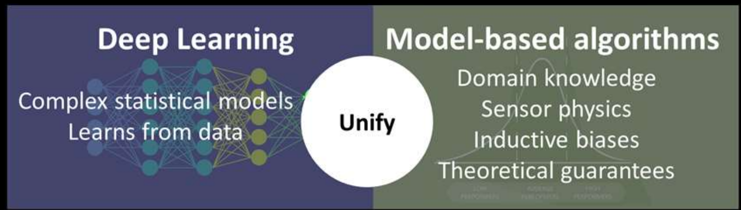

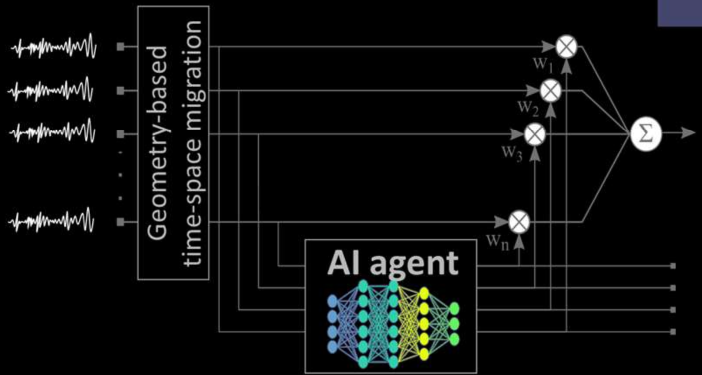

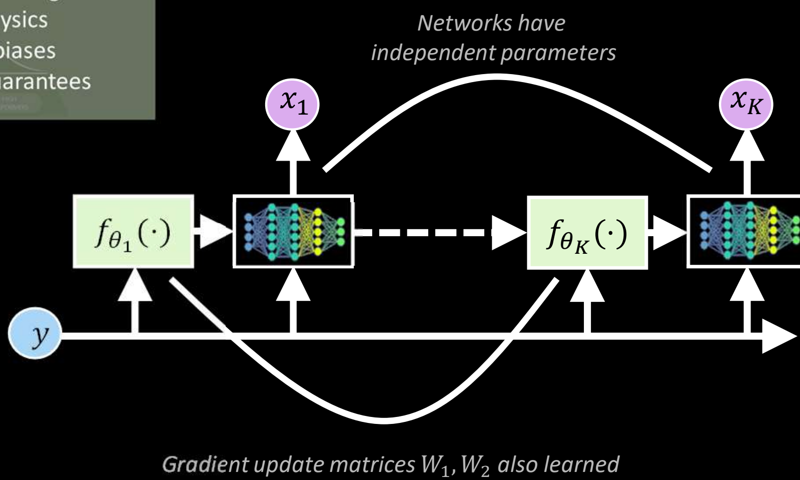

Thanks!