Near-Exact Recovery for Tomographic Inverse Problems via Deep Learning

Martin Genzel

Helmholtz-Zentrum Berlin

(work done while at Utrecht University)

Ingo Gühring

TU Berlin

(Machine Learning Group)

Jan Macdonald

TU Berlin

(Institute of Mathematics)

Maximilian März

TU Berlin

(Institute of Mathematics)

39th International Conference on Machine Learning

Baltimore, Jul 17 – Jul 23, 2022

Ill-Posed Inverse Problems

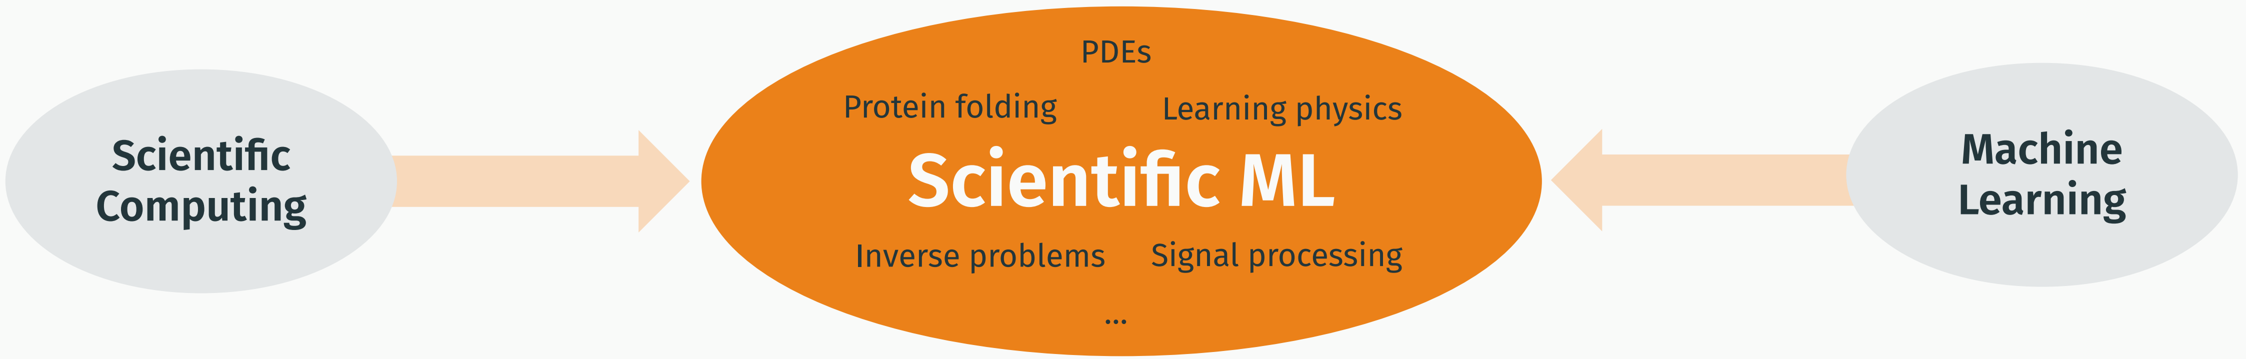

Ill-Posed Inverse Problems

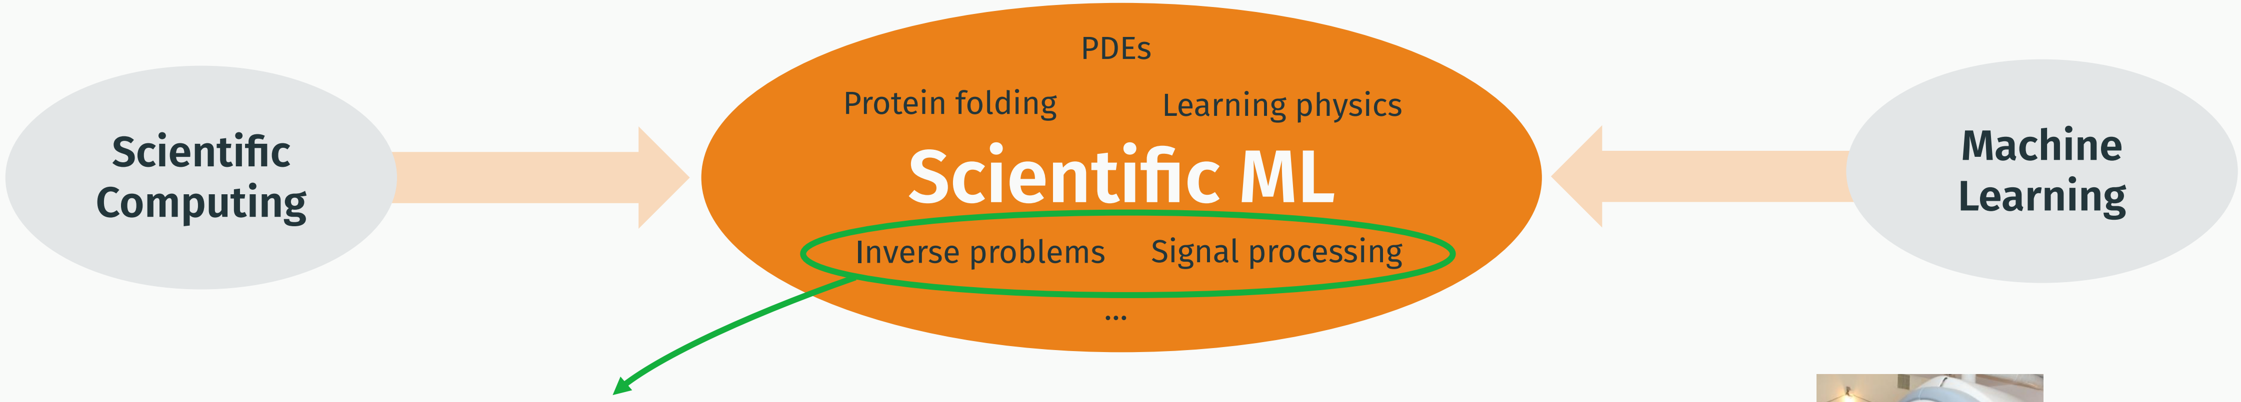

Given a forward operator $A \in \mathbb{R}^{m \times d}$ and corrupted measurements $y = Ax_0 + e$ with $\|e\|_2 \leq \eta$, reconstruct the signal x_0.

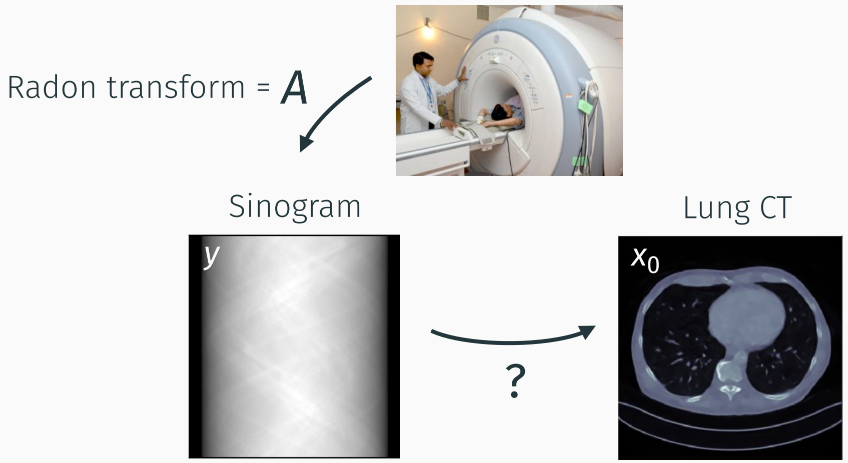

Ill-Posed Inverse Problems

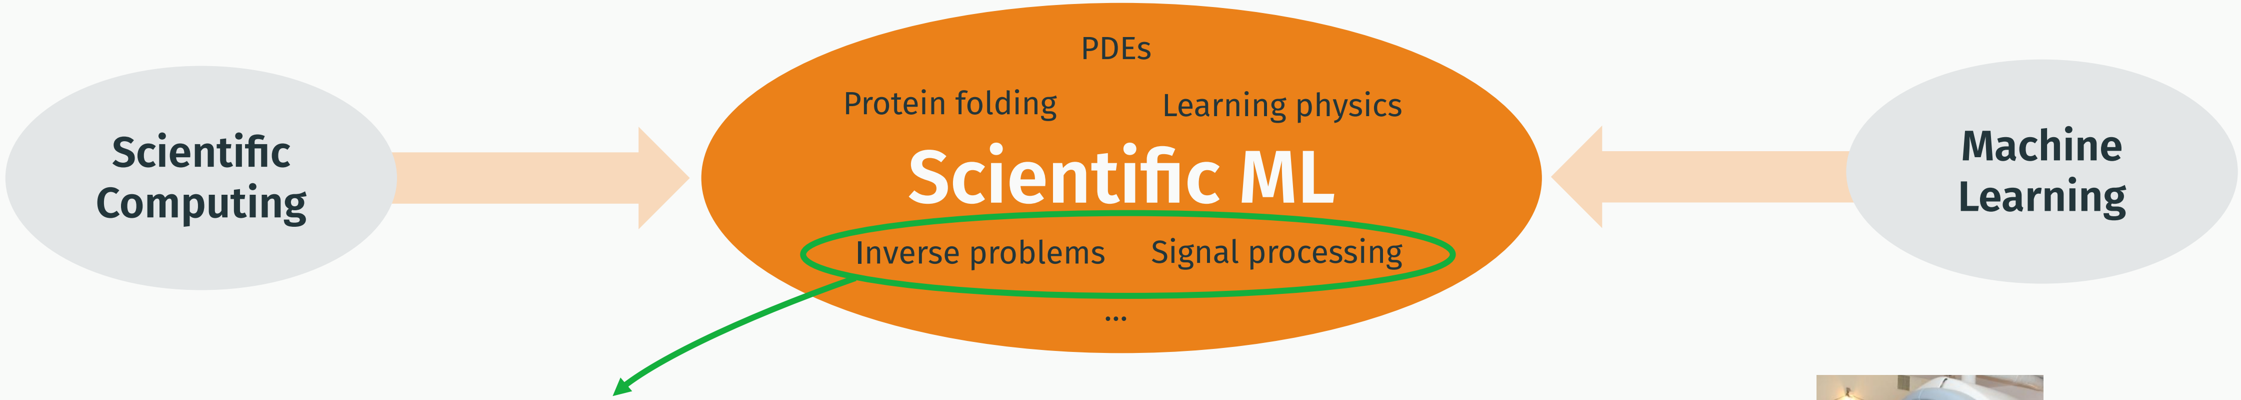

Given a forward operator $A \in \mathbb{R}^{m \times d}$ and corrupted measurements $y = Ax_0 + e$ with $\|e\|_2 \leq \eta$, reconstruct the signal x_0.

↪ Often ill-posed due to undersampling ($m \ll d$) or ill-conditioned forward operator A

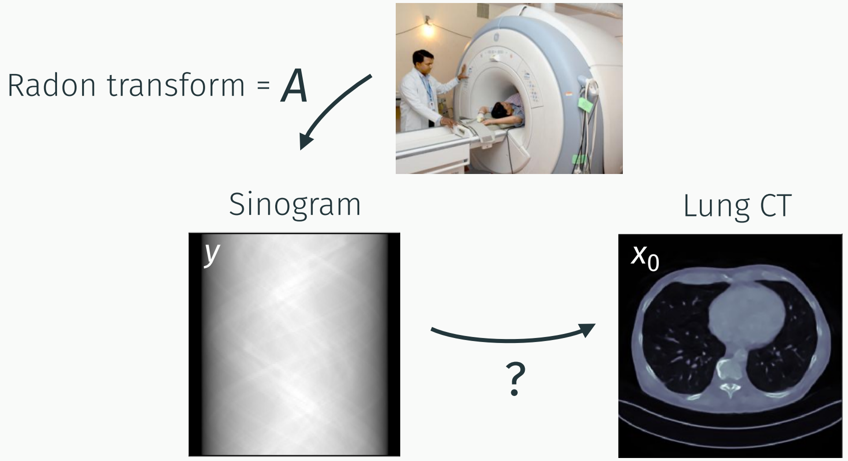

Ill-Posed Inverse Problems

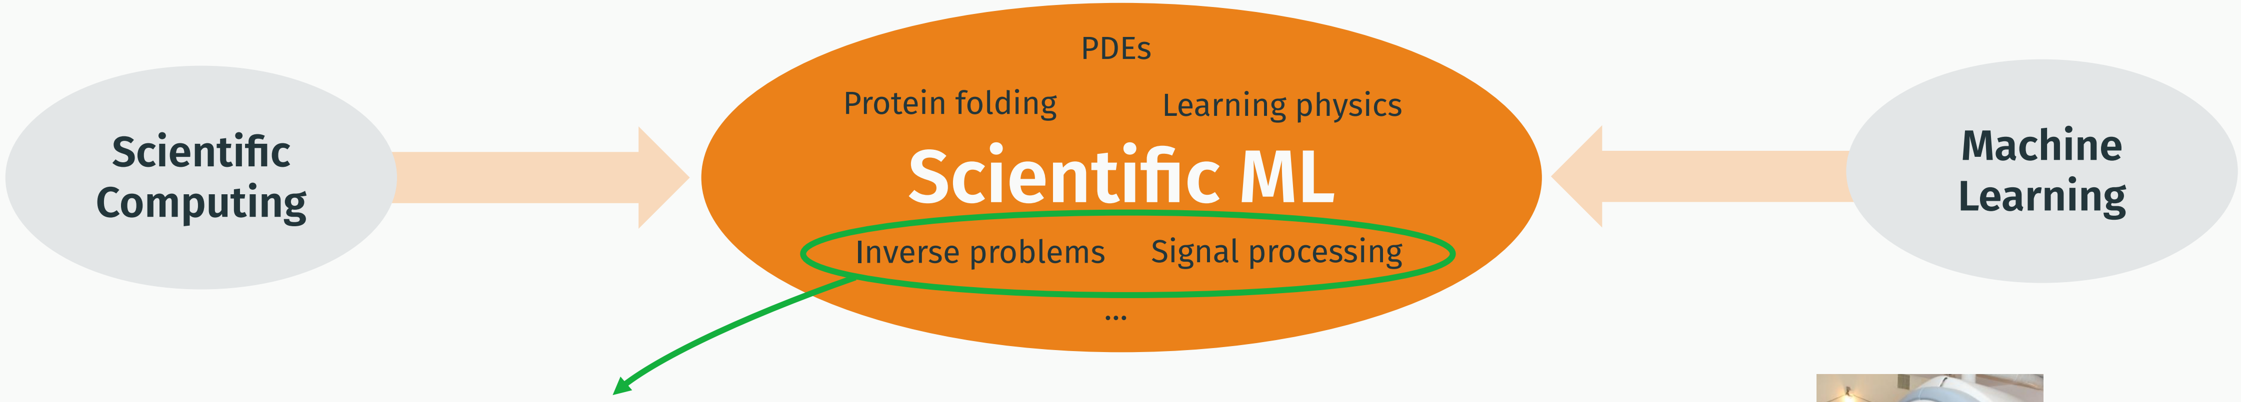

Given a forward operator $A \in \mathbb{R}^{m \times d}$ and corrupted measurements $y = Ax_0 + e$ with $\|e\|_2 \leq \eta$, reconstruct the signal x_0.

↪ Often ill-posed due to undersampling ($m \ll d$) or ill-conditioned forward operator A

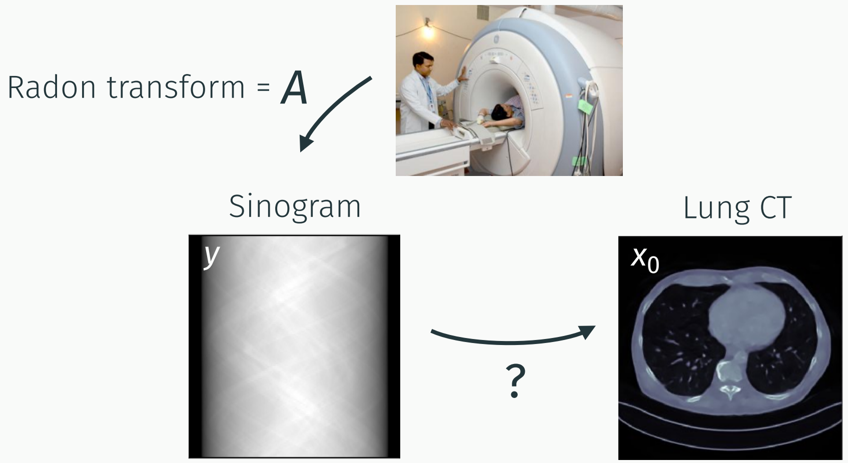

Classical variational methods

$$x_0 \approx \arg\min_x \|y - Ax\|_2^2 + \lambda \cdot R(x)$$

Ill-Posed Inverse Problems

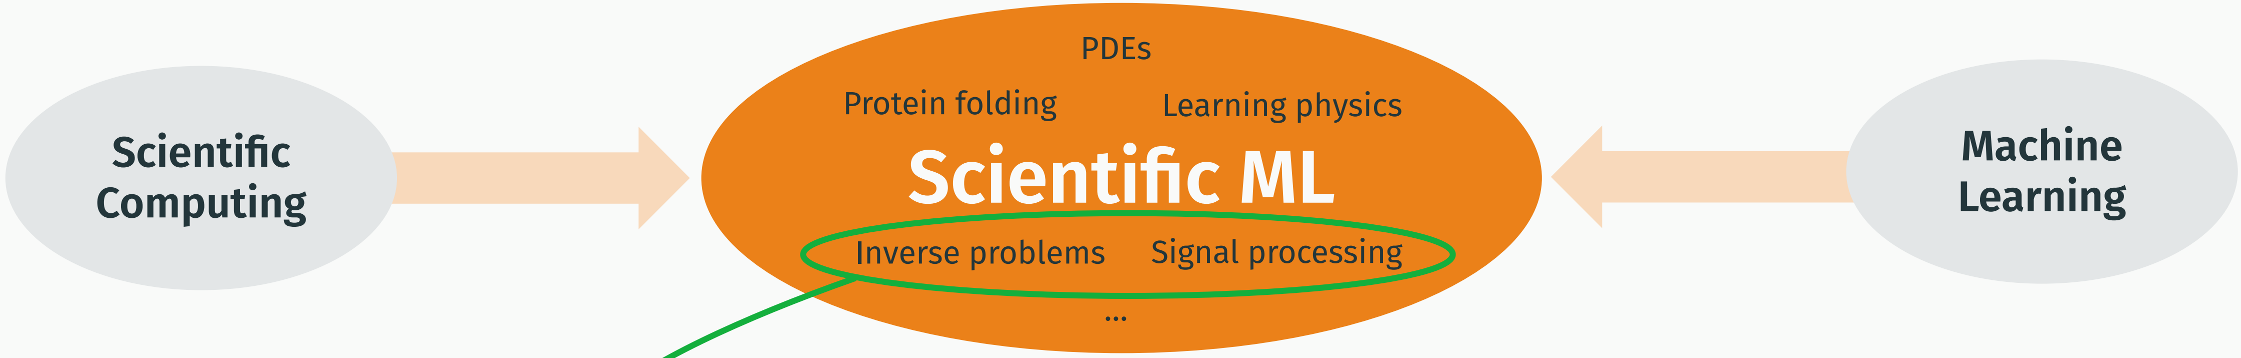

Given a forward operator $A \in \mathbb{R}^{m \times d}$ and corrupted measurements $y = Ax_0 + e$ with $\|e\|_2 \leq \eta$, reconstruct the signal x_0.

↪ Often ill-posed due to undersampling ($m \ll d$) or ill-conditioned forward operator A

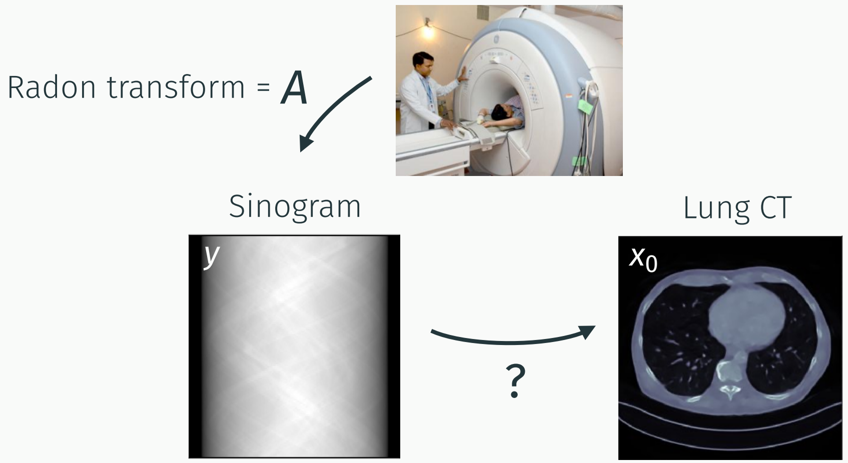

Classical variational methods

$$x_0 \approx \arg\min_x \|y - Ax\|_2^2 + \lambda \cdot R(x)$$

Modern deep learning methods

$$x_0 \approx \mathrm{Net}[\hat{\theta}](y)$$

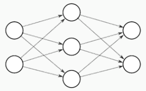

$$\min_\theta \frac{1}{M} \sum_{i=1}^{M} \|\mathrm{Net}[\theta](y^i) - x_0^i\|_2^2$$

"Infer knowledge directly from data $\{(y^i, x_0^i)\}_{i=1}^M$."

Is Deep Learning for Inverse Problems Reliable?

- **Since 2016:** Paradigm shift from sparsity-based regularization to deep learning [Arridge et al. 2019; Ongie et al. 2020]

  (post-processing, unrolling, gen. models, PnP, learned reg., DIP, ...)

$y = Ax_0 + e$ with $\|e\|_2 \leq \eta$

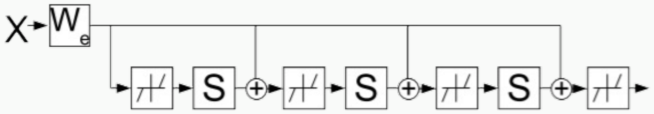

[Gregor & LeCun 2010]

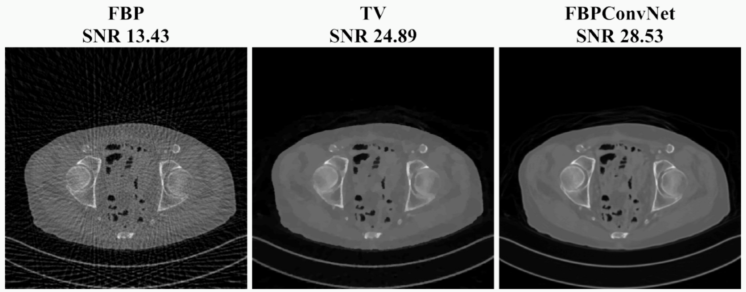

[Jin et al. 2017]

Is Deep Learning for Inverse Problems Reliable?

- **Since 2016:** Paradigm shift from sparsity-based regularization to deep learning [Arridge et al. 2019; Ongie et al. 2020]

  (post-processing, unrolling, gen. models, PnP, learned reg., DIP, ...)

- A lack of theoretical foundation has sparked a controversial debate about reliability [Antun et al. 2020]

$$\|x_0 - \mathrm{Net}(y)\|_2 \leq \underbrace{\|x_0 - \mathrm{Net}(Ax_0)\|_2}_{\text{Accuracy}} + \underbrace{\|\mathrm{Net}(Ax_0) - \mathrm{Net}(y)\|_2}_{\text{Robustness}}$$

$y = Ax_0 + e$ with $\|e\|_2 \leq \eta$

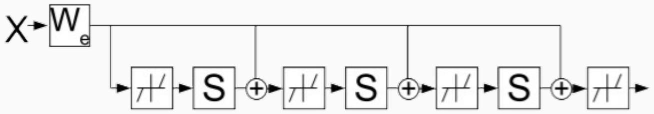

[Gregor & LeCun 2010]

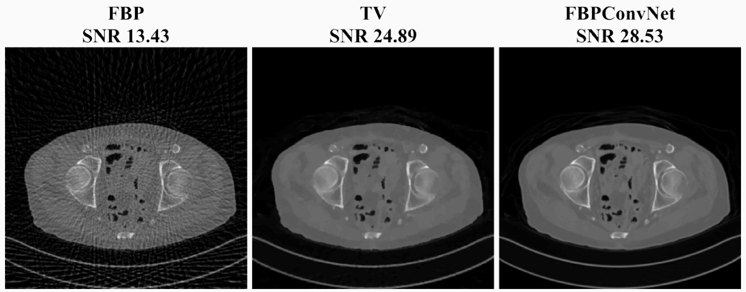

[Jin et al. 2017]

Is Deep Learning for Inverse Problems Reliable?

- **Since 2016:** Paradigm shift from sparsity-based regularization to deep learning [Arridge et al. 2019; Ongie et al. 2020]

  (post-processing, unrolling, gen. models, PnP, learned reg., DIP, ...)

- A lack of theoretical foundation has sparked a controversial debate about reliability [Antun et al. 2020]

$$\|x_0 - \mathrm{Net}(y)\|_2 \leq \underbrace{\|x_0 - \mathrm{Net}(Ax_0)\|_2}_{\text{Accuracy}} + \underbrace{\|\mathrm{Net}(Ax_0) - \mathrm{Net}(y)\|_2}_{\text{Robustness}}$$

$y = Ax_0 + e$ with $\|e\|_2 \leq \eta$

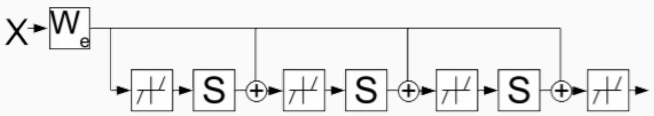

[Gregor & LeCun 2010]

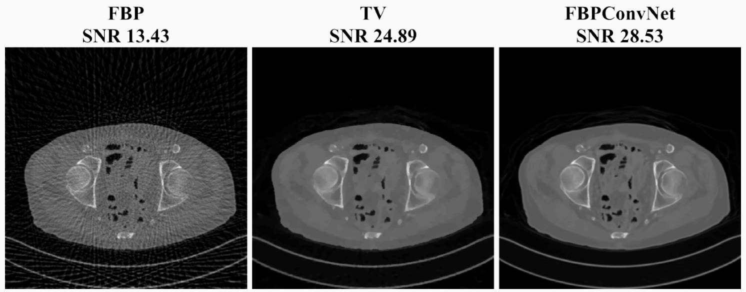

[Jin et al. 2017]

Optimistic results in

Solving Inverse Problems With Deep Neural Networks – Robustness Included?
Genzel, Macdonald, März [IEEE TPAMI 2022]

Is Deep Learning for Inverse Problems Reliable?

- **Since 2016:** Paradigm shift from sparsity-based regularization to deep learning [Arridge et al. 2019; Ongie et al. 2020]

  (post-processing, unrolling, gen. models, PnP, learned reg., DIP, ...)
- A lack of theoretical foundation has sparked a controversial debate about reliability [Antun et al. 2020]

$y = Ax_0 + e$ with $\|e\|_2 \leq \eta$

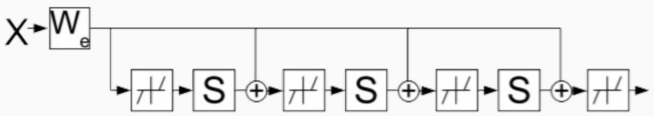

[Gregor & LeCun 2010]

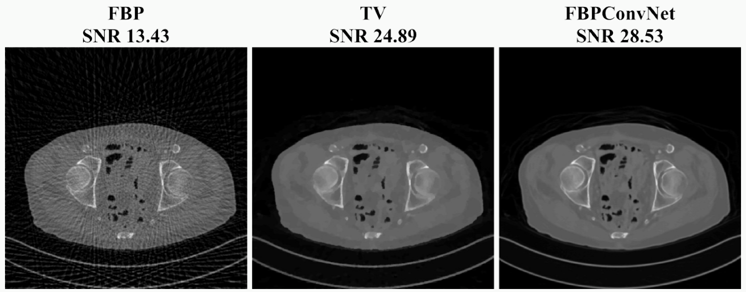

[Jin et al. 2017]

$$\|x_0 - \mathrm{Net}(y)\|_2 \leq \underbrace{\|x_0 - \mathrm{Net}(Ax_0)\|_2}_{\text{Accuracy}} + \underbrace{\|\mathrm{Net}(Ax_0) - \mathrm{Net}(y)\|_2}_{\text{Robustness}}$$

What happens at the noisefree limit $\eta = 0$?
Can NNs produce near-perfect solutions, i.e.,
$\|x_0 - \mathrm{Net}(Ax_0)\|_2 \approx 0$?

Optimistic results in

Solving Inverse Problems With Deep Neural Networks – Robustness Included?
Genzel, Macdonald, März [IEEE TPAMI 2022]

“Do CNNs Solve the CT Inverse Problem?” [Sidky et al. 2021]

- Based on research by Emil Sidky et al. on sparse-view breast CT
- **Answer for post-processing by U-net: No!**
- **Goal of AAPM Challenge:** *“The challenge seeks the data-driven methodology that provides the most accurate reconstruction of sparse-view CT data.”*

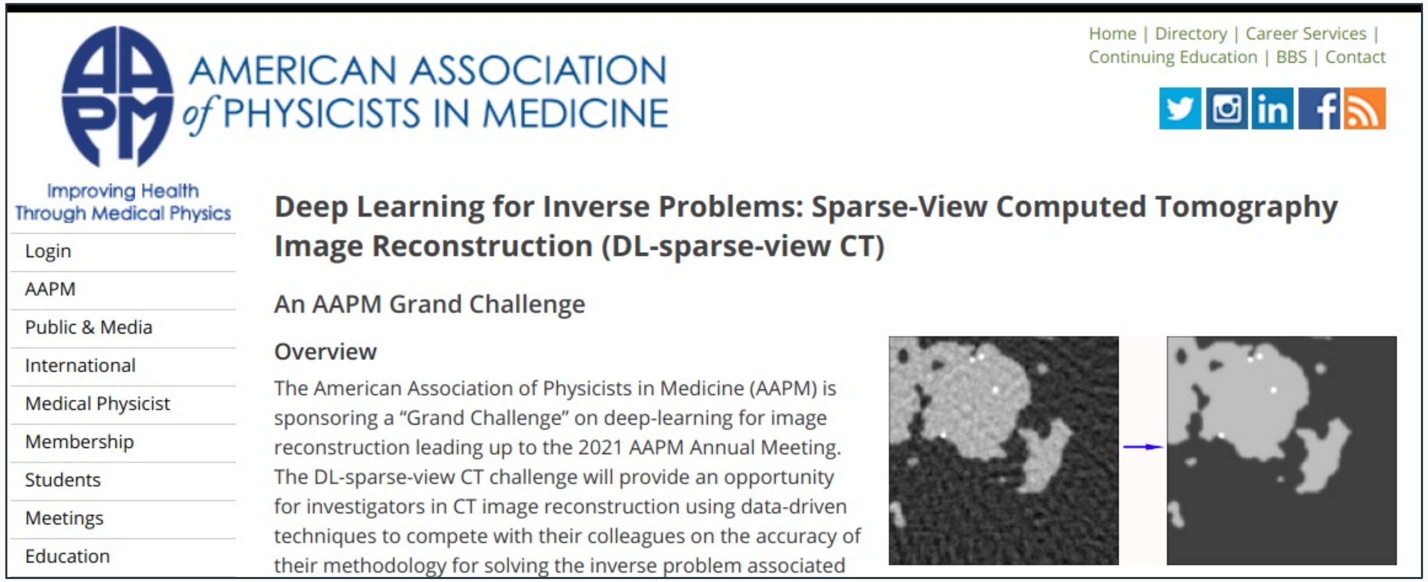
AMERICAN ASSOCIATION of PHYSICISTS IN MEDICINE

Home | Directory | Career Services | Continuing Education | BBS | Contact

Improving Health Through Medical Physics

Login
AAPM
Public & Media
International
Medical Physicist
Membership
Students
Meetings
Education

Deep Learning for Inverse Problems: Sparse-View Computed Tomography Image Reconstruction (DL-sparse-view CT)

An AAPM Grand Challenge

Overview

The American Association of Physicists in Medicine (AAPM) is sponsoring a “Grand Challenge” on deep-learning for image reconstruction leading up to the 2021 AAPM Annual Meeting. The DL-sparse-view CT challenge will provide an opportunity for investigators in CT image reconstruction using data-driven techniques to compete with their colleagues on the accuracy of their methodology for solving the inverse problem associated

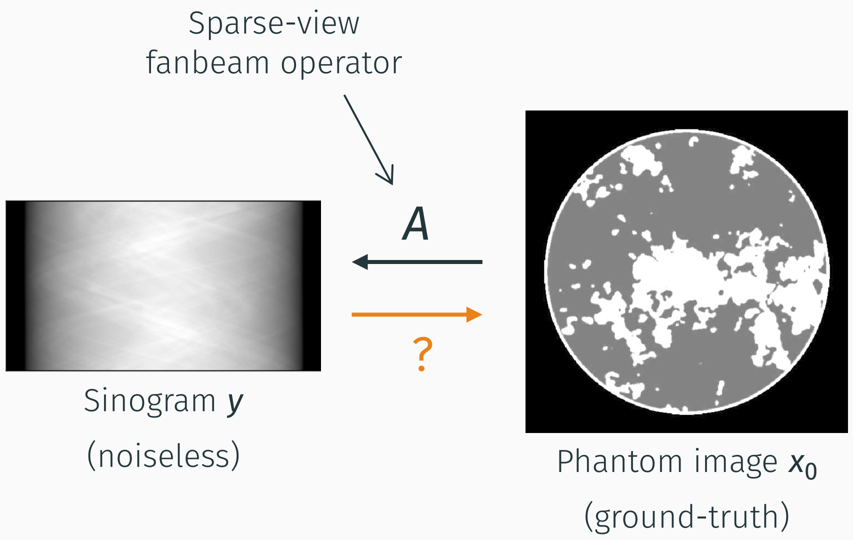

"Do CNNs Solve the CT Inverse Problem?" [Sidky et al. 2021]

- Based on research by Emil Sidky et al. on sparse-view breast CT
- **Answer for post-processing by U-net: No!**
- **Goal of AAPM Challenge:** *"The challenge seeks the data-driven methodology that provides the most accurate reconstruction of sparse-view CT data."*

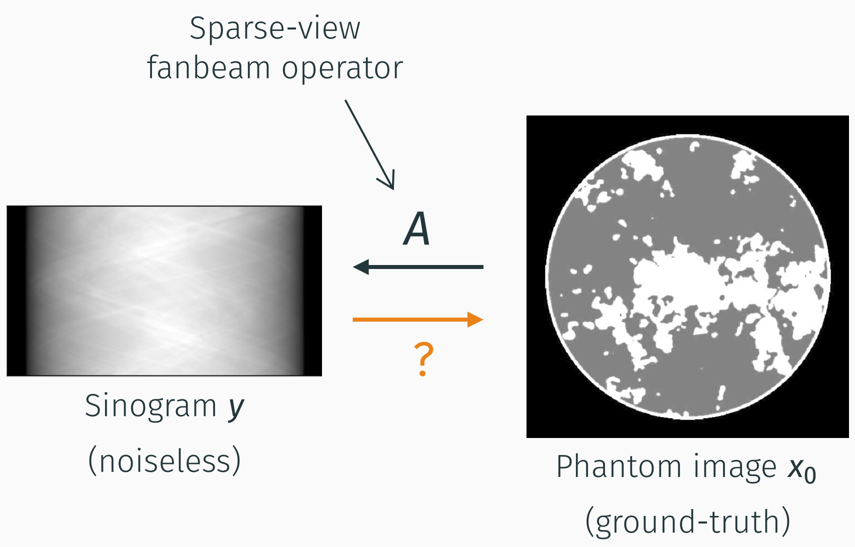

- Challenge dataset: M = 4000 pairs of
  - breast-phantom images (d=512x512)
  - noiseless sparse-view sinograms (128 views)
  - (sparse-view FBP images)
- Evaluation by RMSE ($= \frac{1}{M}\sum_{i=1}^{M}\sqrt{\|x_0^i - \hat{x}_0^i\|_2^2/d}$)

  ↪ TV minimization can solve the problem (RMSE ≈ 1e-6)
- Timeline: Mar 17 – Jun 1, 2021

  ↪ Approx. 60 groups participated (25 in final phase)

"Do CNNs Solve the CT Inverse Problem?" [Sidky et al. 2021]

- Based on research by Emil Sidky et al. on sparse-view breast CT
- **Answer for post-processing by U-net: No!**
- **Goal of AAPM Challenge:** *"The challenge seeks the data-driven methodology that provides the most accurate reconstruction of sparse-view CT data."*
- Challenge dataset: M = 4000 pairs of
  - breast-phantom images (d=512x512)
  - noiseless sparse-view sinograms (128 views)
  - (sparse-view FBP images)
- Evaluation by RMSE ($= \frac{1}{M} \sum_{i=1}^{M} \sqrt{\|x_0^i - \hat{x}_0^i\|_2^2 / d}$)

  ↪ TV minimization can solve the problem (RMSE ≈ 1e-6)
- Timeline: Mar 17 – Jun 1, 2021

  ↪ Approx. 60 groups participated (25 in final phase)

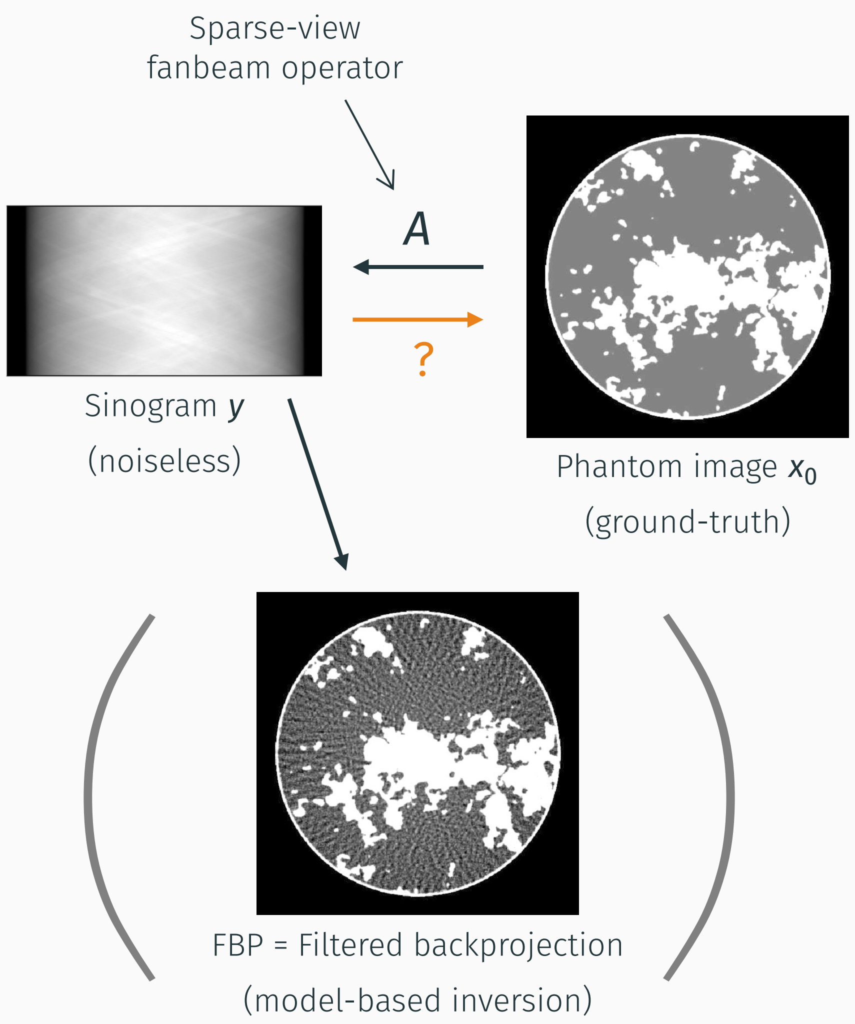

"Do CNNs Solve the CT Inverse Problem?" [Sidky et al. 2021]

- Based on research by Emil Sidky et al. on sparse-view breast CT
- **Answer for post-processing by U-net: No!**
- **Goal of AAPM Challenge:** *"The challenge seeks the data-driven methodology that provides the most accurate reconstruction of sparse-view CT data."*
- Challenge dataset: M = 4000 pairs of
  - breast-phantom images (d=512x512)
  - noiseless sparse-view sinograms (128 views)
  - (sparse-view FBP images)
- Evaluation by RMSE ($= \frac{1}{M} \sum_{i=1}^{M} \sqrt{\|x_0^i - \hat{x}_0^i\|_2^2 / d}$)

  ↪ TV minimization can solve the problem (RMSE ≈ 1e-6)
- Timeline: Mar 17 – Jun 1, 2021

  ↪ Approx. 60 groups participated (25 in final phase)

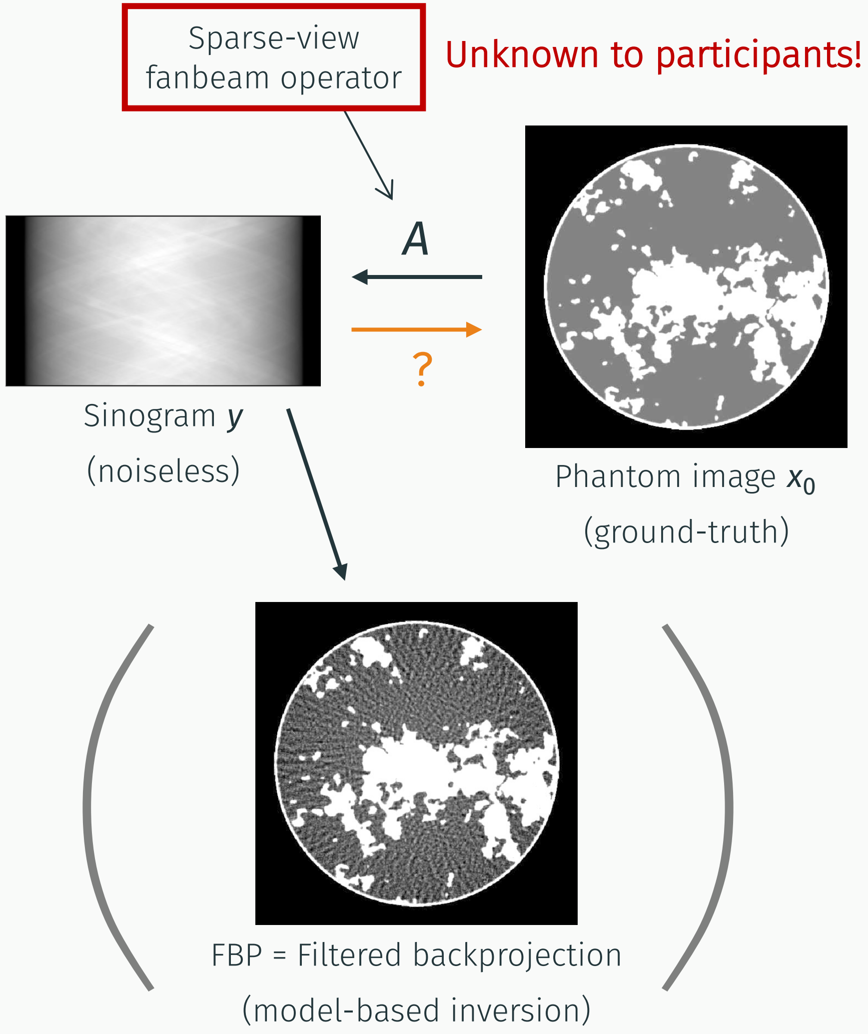

Our Approach to the AAPM Challenge

- **Step 1:** Fully data-driven operator identification based on a parameterized fwd. model $A[\theta]$

$$\min_\theta \sum_i \left\| A[\theta]\left(x_0^i\right) - y^i \right\|_2^2$$

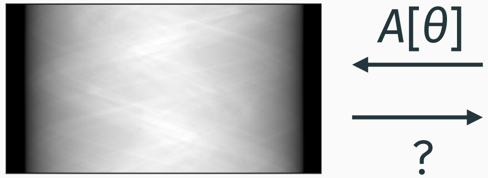

Sinogram y

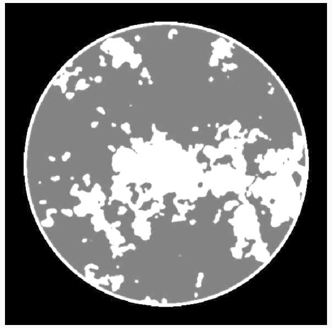
Phantom image x_0

Our Approach to the AAPM Challenge

- **Step 1:** Fully data-driven operator identification based on a parameterized fwd. model $A[\theta]$

$$\min_\theta \sum_i \left\| A[\theta]\left(x_0^i\right) - y^i \right\|_2^2$$

Deep-learning-style optimization (backprop/autodiff)

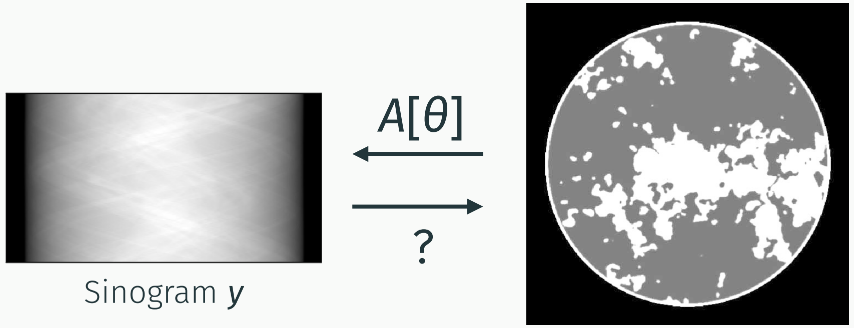

Sinogram y

Phantom image x_0

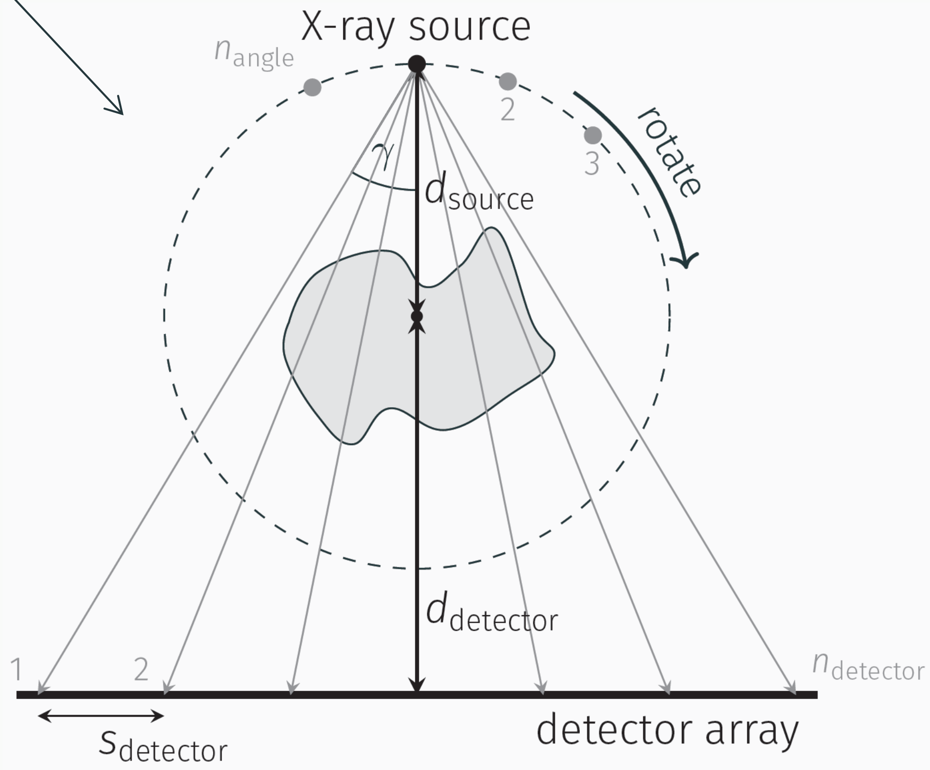

Our Approach to the AAPM Challenge

- **Step 1:** Fully data-driven operator identification based on a parameterized fwd. model $A[\theta]$

$$\min_\theta \sum_i \left\| A[\theta]\left(x_0^i\right) - y^i \right\|_2^2$$

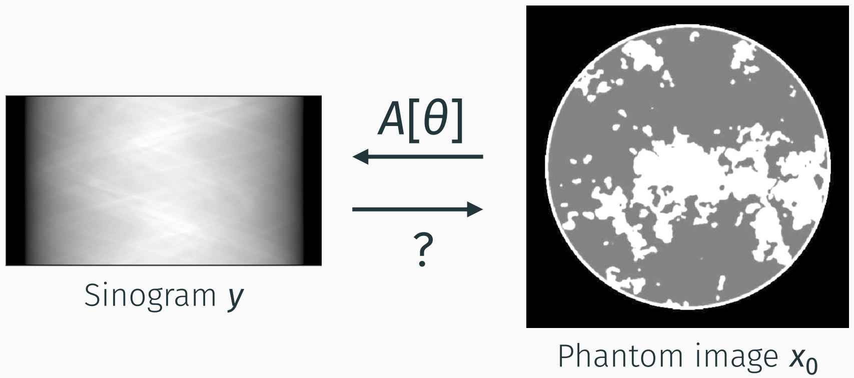

- **Step 2:** Pre-train a UNet as FBP-post-processor

[Jin et al. 2017; Kang et al. 2017; ...]

$$\min_\theta \sum_i \left\| x_0^i - [\text{UNet}[\theta] \circ \text{FBP}](y^i) \right\|_2^2$$

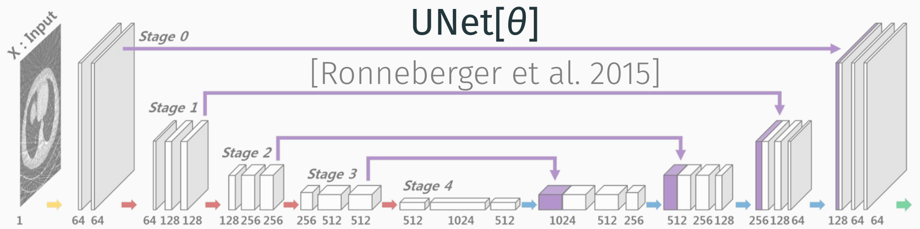

Our Approach to the AAPM Challenge

- **Step 1:** Fully data-driven operator identification based on a parameterized fwd. model $A[\theta]$

$$\min_\theta \sum_i \left\| A[\theta]\left(x_0^i\right) - y^i \right\|_2^2$$

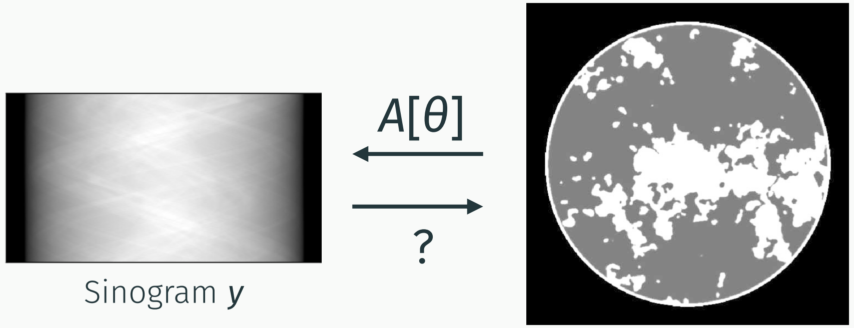

Sinogram y

Phantom image x_0

- **Step 2:** Pre-train a UNet as FBP-post-processor

[Jin et al. 2017; Kang et al. 2017; ...]

$$\min_\theta \sum_i \left\| x_0^i - [\mathrm{UNet}[\theta] \circ \mathrm{FBP}](y^i) \right\|_2^2$$

computational backbone (→ UNet[θ])

"learned" (depends on $A[\theta]$) (→ FBP)

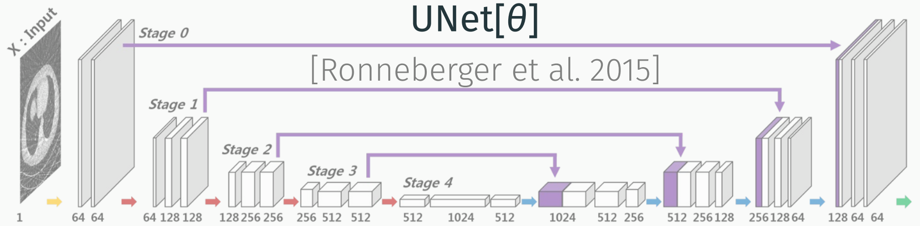

y x

Input =

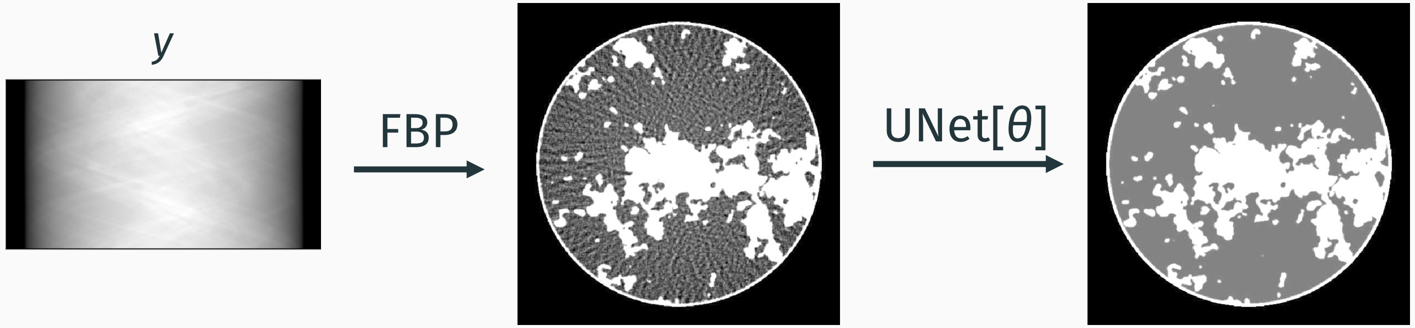

= Reconstruction

Our Approach to the AAPM Challenge

- **Step 1:** Fully data-driven operator identification based on a parameterized fwd. model $A[\theta]$

$$\min_\theta \sum_i \left\| A[\theta]\left(x_0^i\right) - y^i \right\|_2^2$$

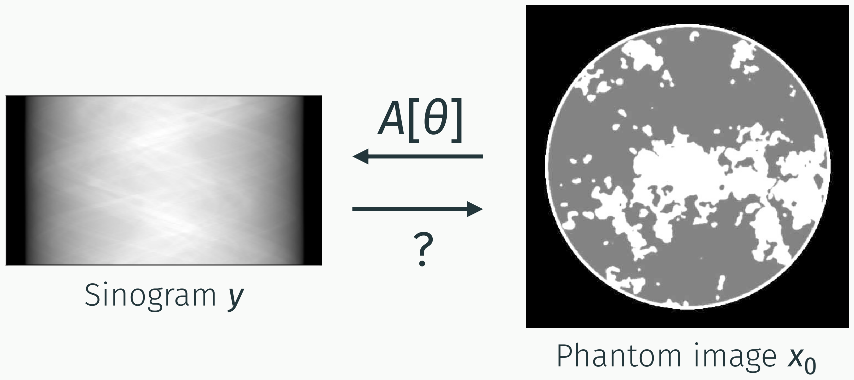

- **Step 2:** Pre-train a UNet as FBP-post-processor [Jin et al. 2017; Kang et al. 2017; ...]

- **Step 3:** Construct an iterative scheme (= ItNet) [Aggarwal et al. 2018; Schlemper et al. 2019; Hammernik, Schlemper, et al. 2021; ...]

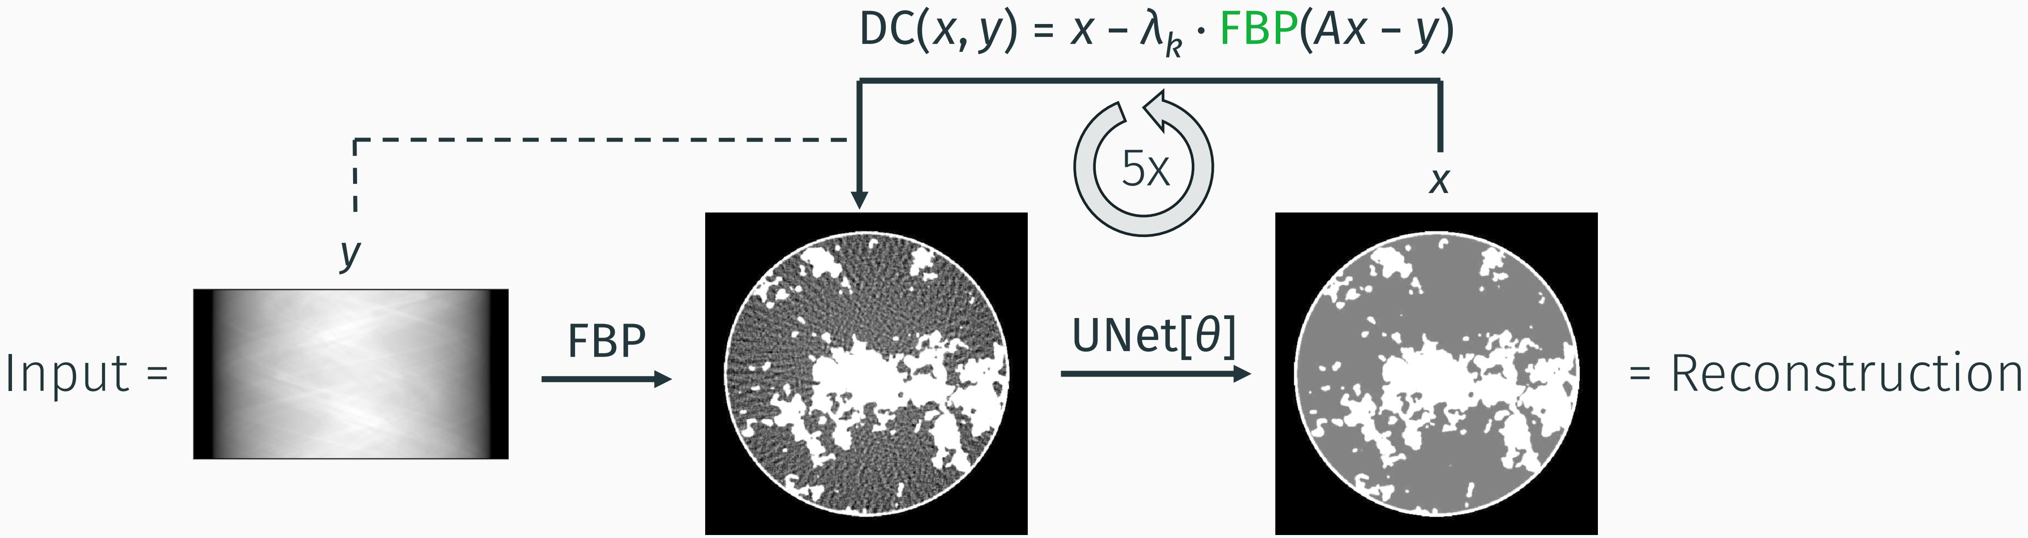

Our Approach to the AAPM Challenge

- **Step 1:** Fully data-driven operator identification based on a parameterized fwd. model $A[\theta]$

$$\min_\theta \sum_i \left\| A[\theta]\left(x_0^i\right) - y^i \right\|_2^2$$

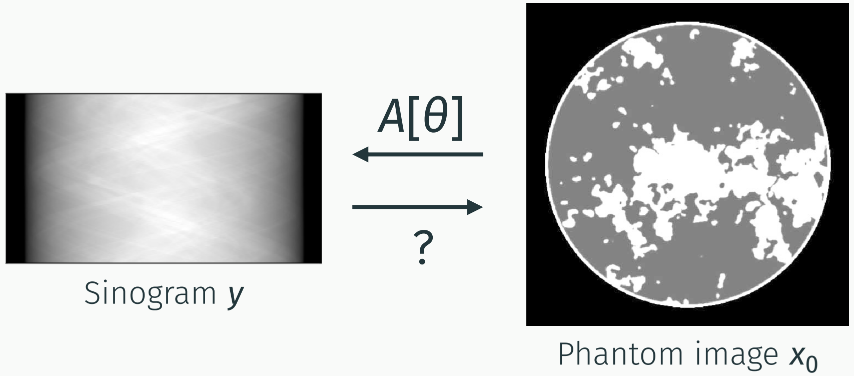

- **Step 2:** Pre-train a UNet as FBP-post-processor [Jin et al. 2017; Kang et al. 2017; ...]

- **Step 3:** Construct an iterative scheme (= ItNet) [Aggarwal et al. 2018; Schlemper et al. 2019; Hammernik, Schlemper, et al. 2021; ...]

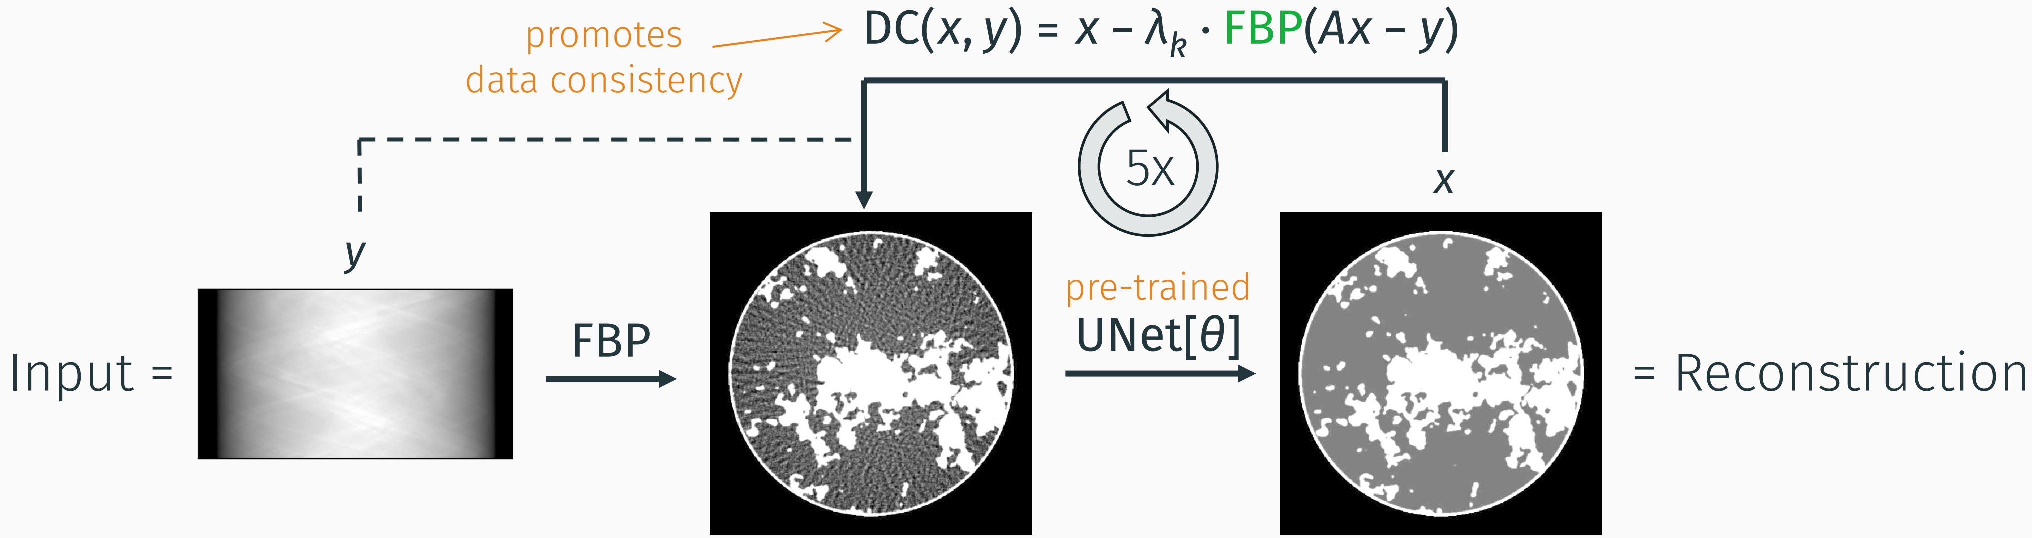

Challenge Results – Team: `robust-and-stable`

	Baselines		Our Network Variants		Comparison Networks	
	Chall. **FBP**	Our **FBP**	**U-net**	**ItNet**	**Tiramisu**	**LPD**
RMSE	5.72e-3	3.40e-3	3.50e-4	**6.37e-6**	2.24e-4	1.24e-4

[Jégou et al. 2017; Adler & Öktem 2018]

Results

#	User	Entries	Date of Last Entry	Team Name	RMSE ▲
1	**Max**	3	05/31/21	robust-and-stable	0.00000637 (1)
2	**TUM**	4	05/31/21	YM & RH	0.00003989 (2)
3	**cebel67**	4	05/31/21	DEEP_UL	0.00012923 (3)
4	deepx	3	05/31/21		0.00015935 (4)
5	Haimiao	4	05/29/21	HBB	0.00018119 (5)
6	HKim	2	05/31/21	MIR	0.00026678 (6)
7	luke199629	5	05/31/21		0.00028064 (7)
8	yume	3	05/26/21	list	0.00029180 (8)

ItNet (RMSE = 6.23e-06)

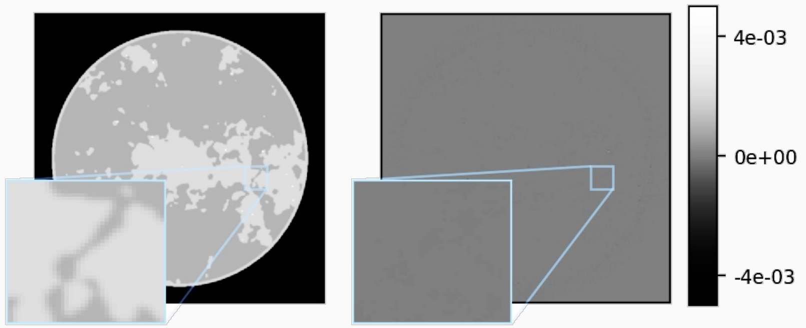

Tiramisu (RMSE = 2.75e-04)

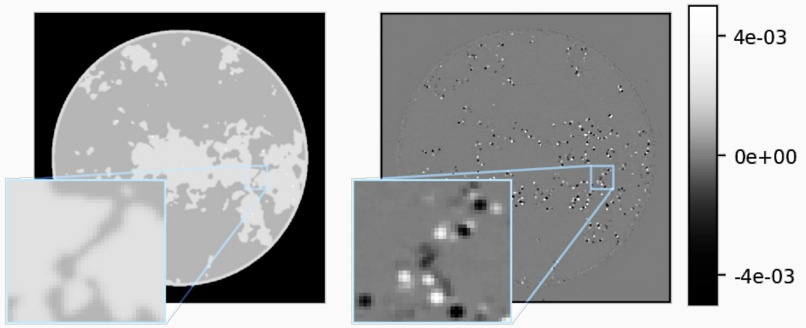

Near-exact image recovery via end-to-end NNs is possible
↪ “practical solution” to the inverse problem

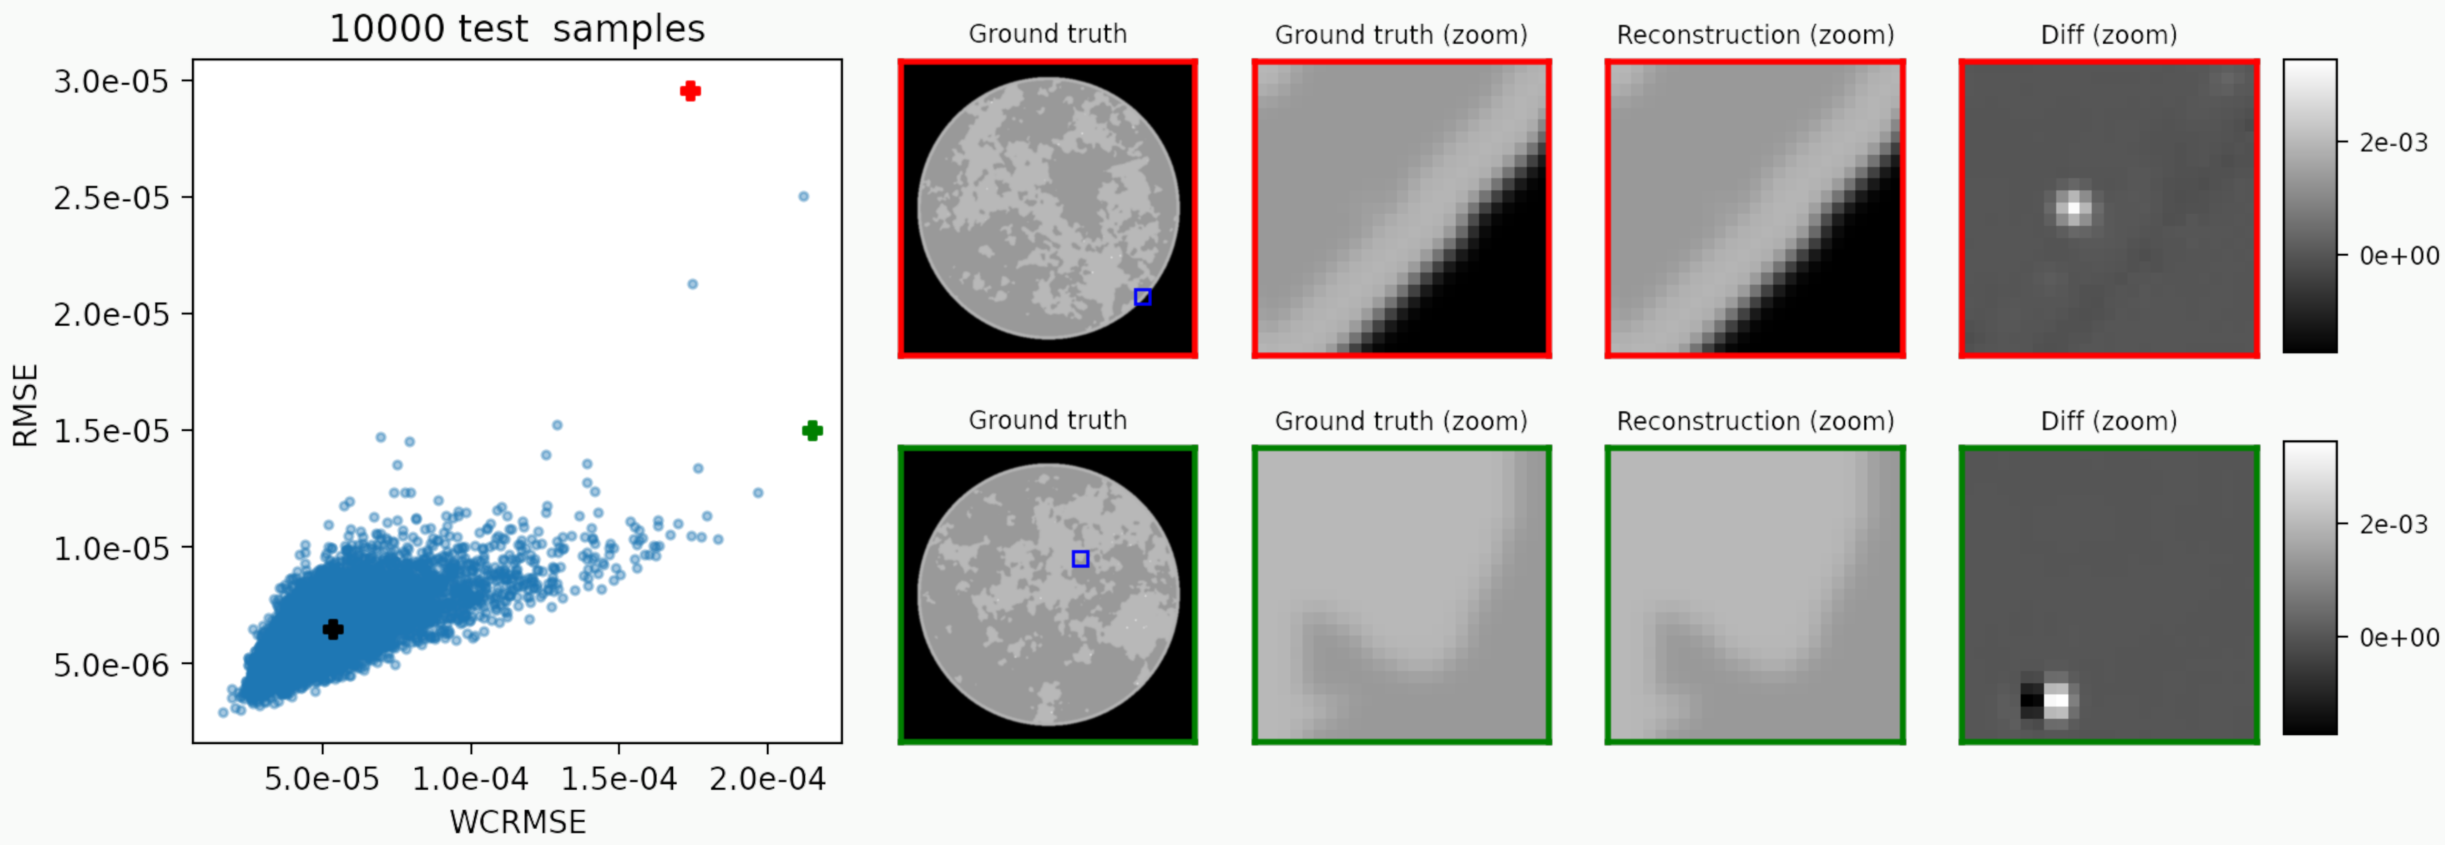

Only very few iterations of fwd. operator required

↪ very different from classical solvers

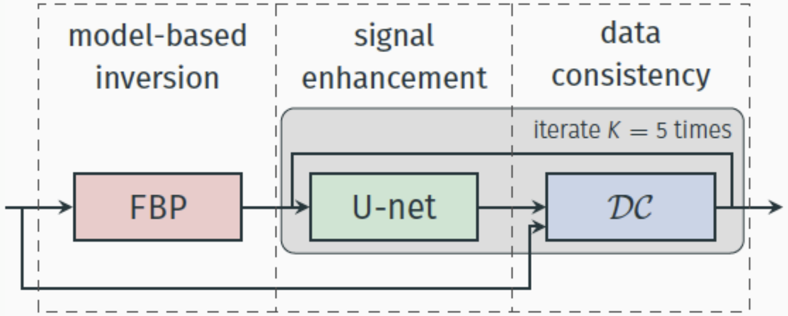

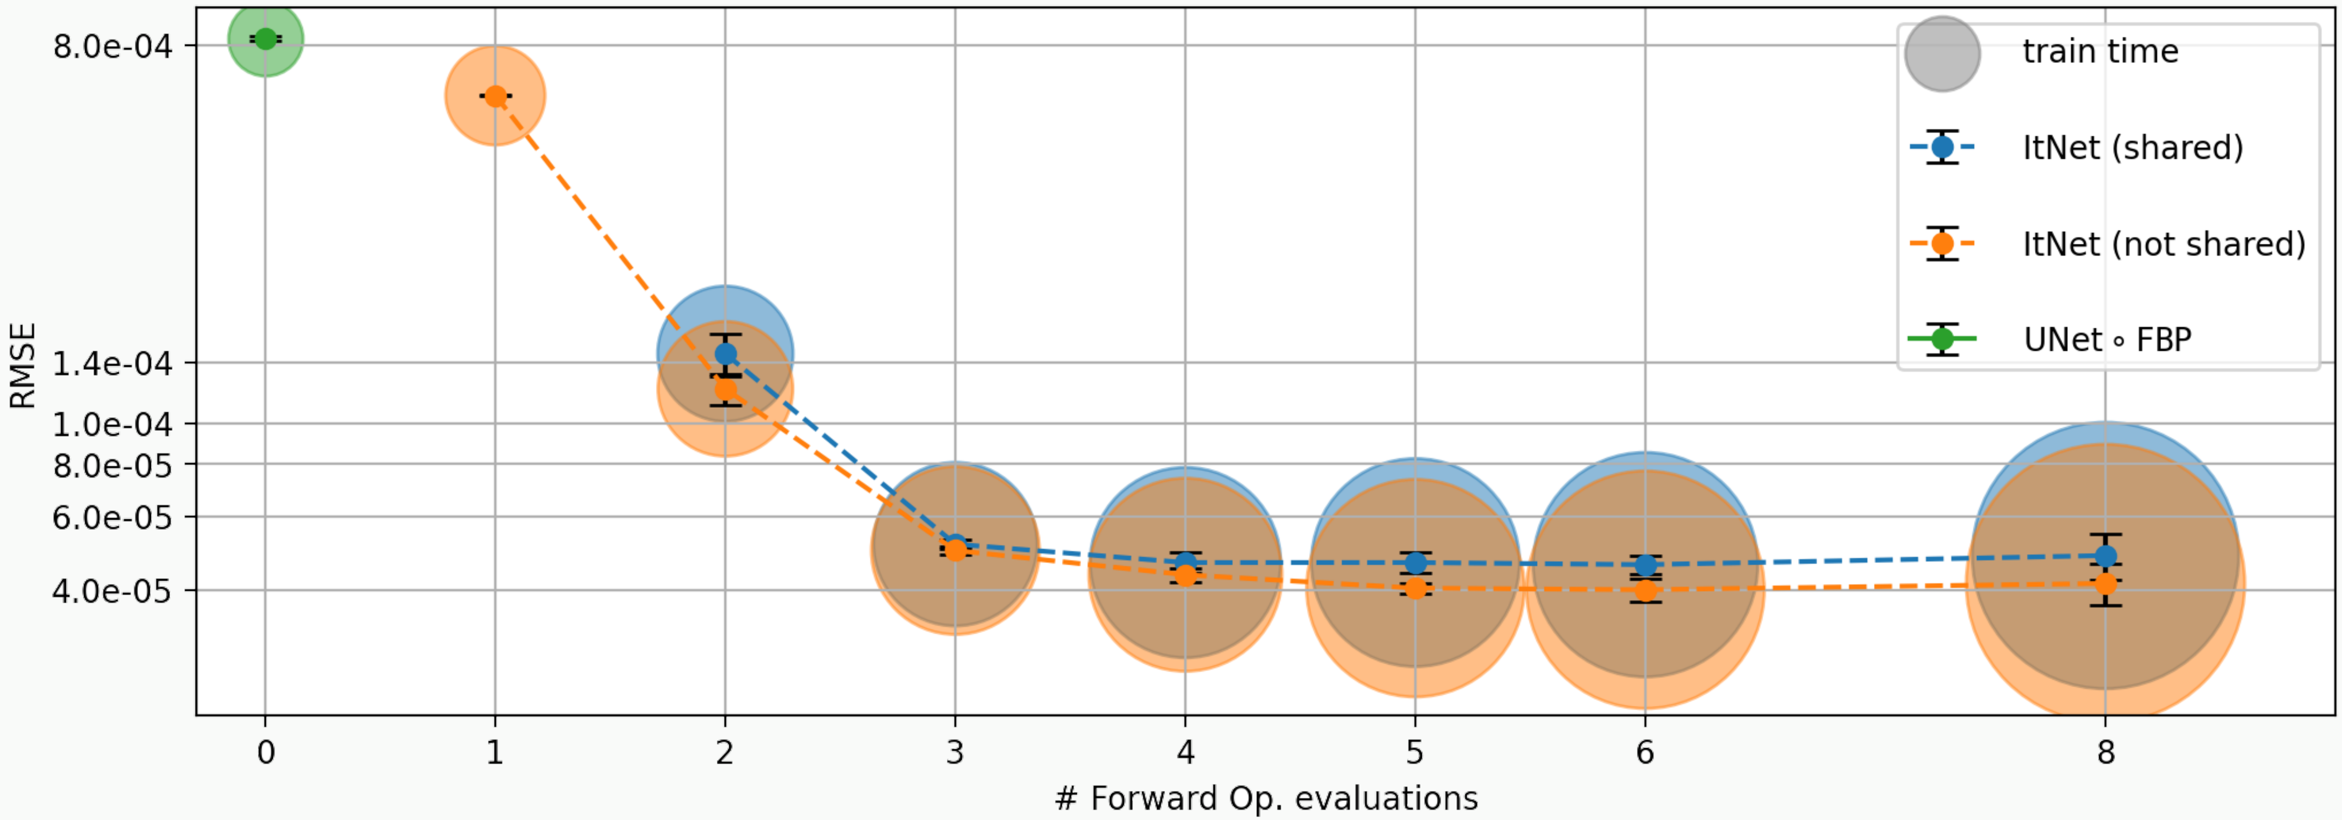

Further Analysis & Take-Aways (3/4)

Model-based knowledge and pre-training is key

↪ “Simple models, but trained well!”

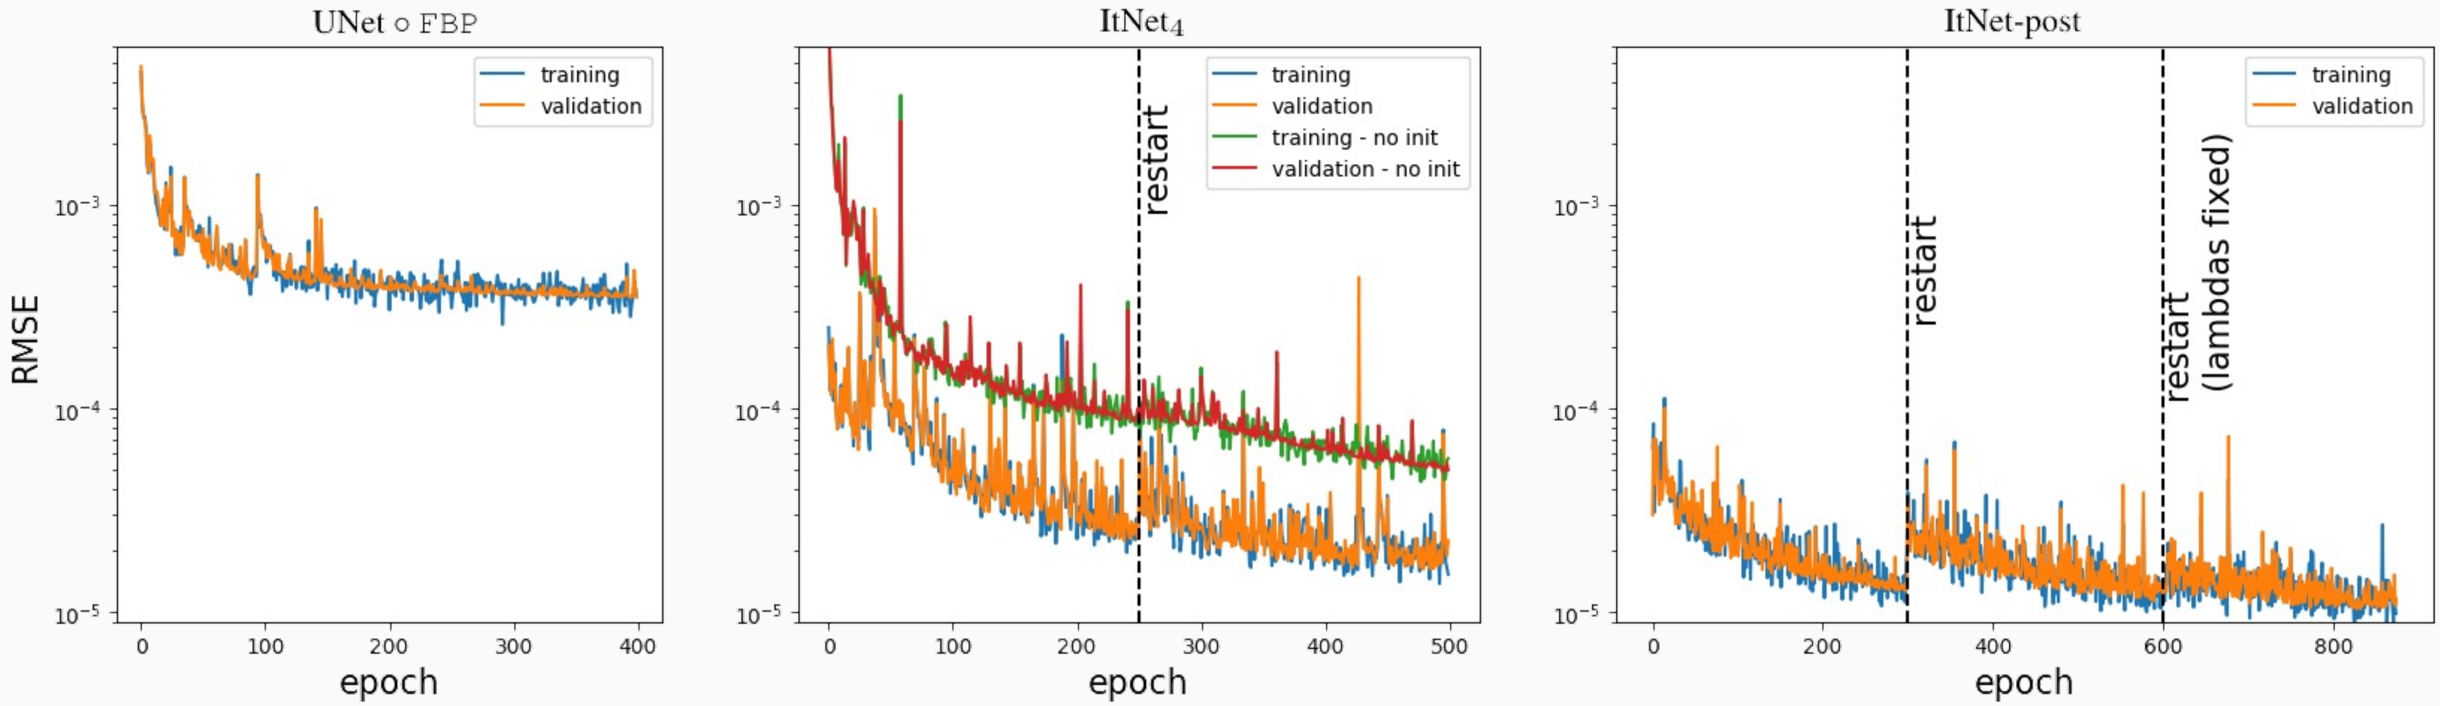

Further Analysis & Take-Aways (4/4)

ItNet is also SOTA for real-world CT image data

#	User (Team)	Created	Mean Position
1st	RobustAndStable	4 Nov. 2021	1.3
2nd	RobustAndStable	1 May 2021	2.8
2nd	iRIMforCT	19 Aug. 2021	2.8
4th	RobustAndStable	28 April 2021	5.5
5th	iRIMforCT	31 July 2021	6.8

LoDoPaB-CT – Leuschner et al. 2021

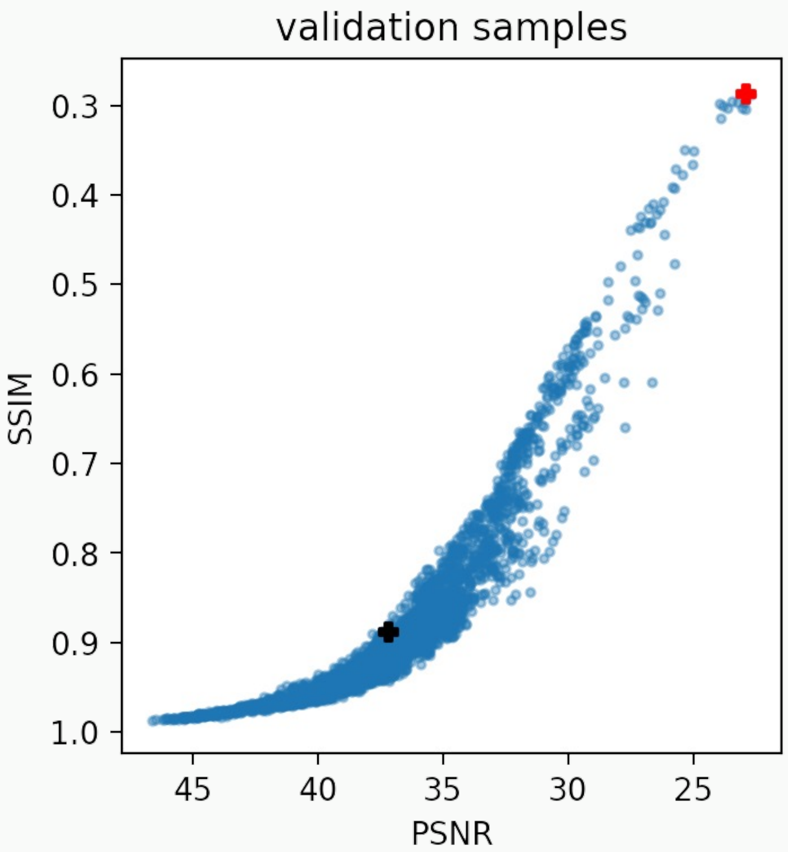

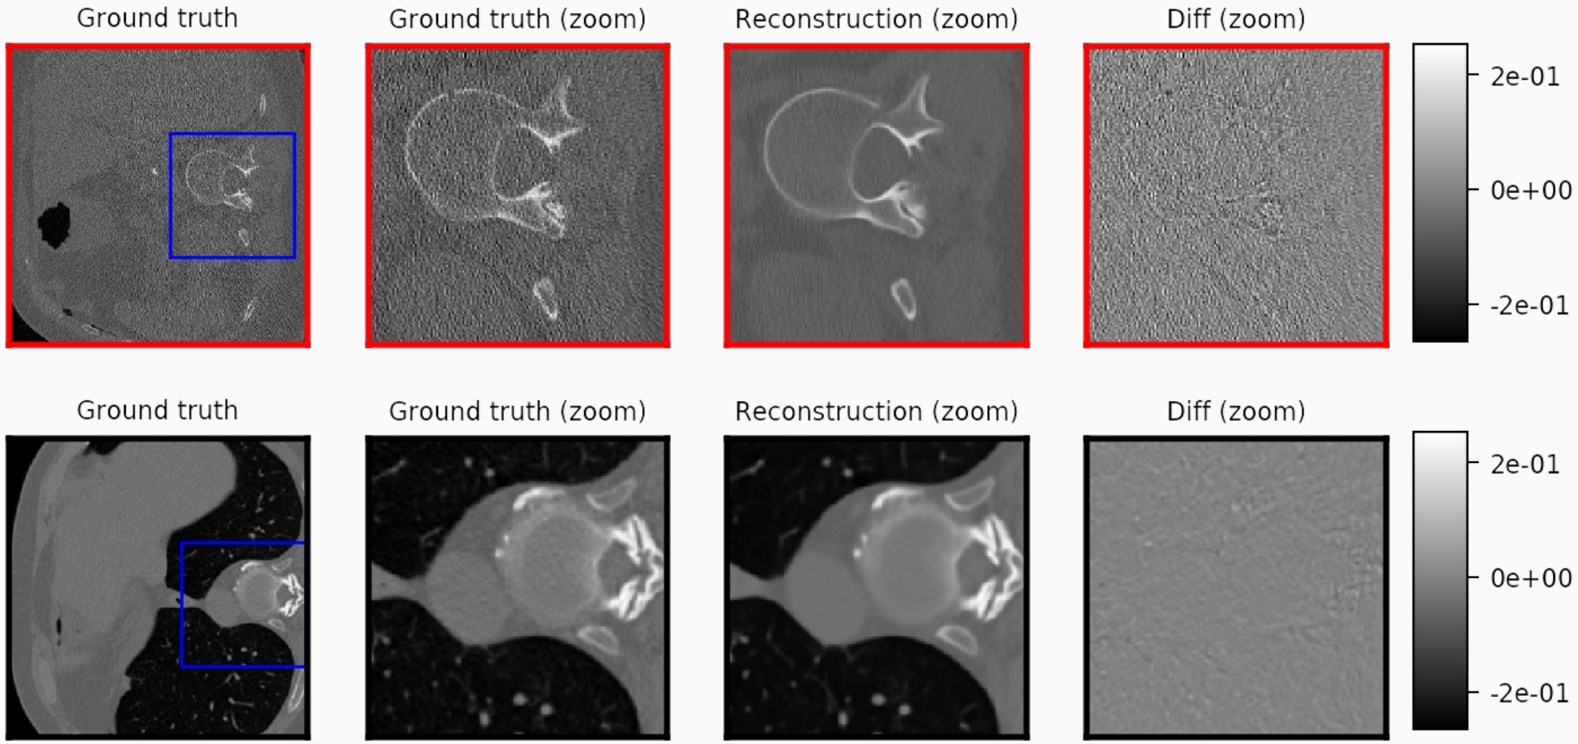

THANK YOU!

Official Challenge Report:

E. Sidky & X. Pan. Report on the AAPM deep-learning sparse-view CT (DL-sparse-view CT) Grand Challenge. *Med. Phys.* (2022), arXiv:2109.09640

Our code:

`https://github.com/jmaces/aapm-ct-challenge`

Find us on

`@MartinGenzel` `@Iguhring` `@jan_maces` `@MaximilianMarz`