Recent advances in using machine learning for image reconstruction

Ozan Öktem

Department of Mathematics
KTH - Royal Institute of Technology, Stockholm

December 6, 2017
Mathematics of Imaging and Vision
Centre for Mathematical Sciences, Cambridge

Learned iterative reconstruction

Learned iterative reconstruction
References & implementation

References Joint work with Jonas Adler (KTH/Elekta)

- J. Adler and O. Öktem, *Solving ill-posed inverse problems using iterative deep neural networks*. Inverse Problems, vol. 33, no. 12, p. 124007, 2017.
  ArXiv version at `http://arxiv.org/abs/1704.04058`
- J. Adler and O. Öktem, *Learned Primal-Dual Reconstruction*. Accepted for publication in IEEE Transactions in Medical Imaging, 2017.
  ArXiv version at `http://arxiv.org/abs/1707.06474`

Software source code

- ODL (`http://github.com/odlgroup/odl`)
- `http://github.com/adler-j`: ODL-Tensorflow-ASTRA implementation of learned primal-dual scheme for tomography

Learned iterative reconstruction
Motivation

- Inverse problem: Recover model parameters $x_{\text{true}} \in X$ given data $y \in Y$ were $y = \mathcal{A}(x_{\text{true}}) + e$.
  - $\mathcal{A}\colon X \to Y$ forward operator
  - Noise term $e \in Y$ sample of Y-valued random variable $\mathbf{E}$.
- Variational approach: Solve an optimisation problem:

$$\hat{x} := \arg\min_{x} \left[\left\| \mathcal{A}(x) - y \right\|_Y^2 + \mathcal{S}_\theta(x) \right]$$

Learned iterative reconstruction
Motivation

- Inverse problem: Recover model parameters $x_{\text{true}} \in X$ given data $y \in Y$ were $y = \mathcal{A}(x_{\text{true}}) + e$.
  - $\mathcal{A}\colon X \to Y$ forward operator
  - Noise term $e \in Y$ sample of Y-valued random variable $\mathbf{E}$.
- Variational approach: Solve an optimisation problem:

$$\hat{x} := \arg\min_x \left[\left\| \mathcal{A}(x) - y \right\|_Y^2 + \mathcal{S}_\theta(x) \right]$$

Learned iterative reconstruction
Motivation

- Inverse problem: Recover model parameters $x_{\text{true}} \in X$ given data $y \in Y$ were $y = \mathcal{A}(x_{\text{true}}) + e$.
  - $\mathcal{A}\colon X \to Y$ forward operator
  - Noise term $e \in Y$ sample of Y-valued random variable $\mathbf{E}$.
- Variational approach: Solve an optimisation problem:

$$\hat{x} := \arg\min_x \left[\left\| \mathcal{A}(x) - y \right\|_Y^2 + \mathcal{S}_\theta(x) \right]$$

Advantages

+ Generic, yet highly adaptable, with a plug-and-play structure.
  - $\mathcal{S}_\theta$ captures a priori information, like sparsity.
  - $\mathcal{A}$ (forward operator) encodes the physics.
  - Choice of $\| \cdot \|_Y^2$ reflects statistical properties of data.
  - Choice of θ governed by noise level in data.
+ Until recently, yields best reconstruction 'quality' when properly adapted.
+ Interpretation as maximum a posteriori (MAP) estimator.

Learned iterative reconstruction
Motivation

- Inverse problem: Recover model parameters $x_{\text{true}} \in X$ given data $y \in Y$ were $y = \mathcal{A}(x_{\text{true}}) + e$.
  - $\mathcal{A}\colon X \to Y$ forward operator
  - Noise term $e \in Y$ sample of Y-valued random variable $\mathbf{E}$.
- Variational approach: Solve an optimisation problem:

$$\hat{x} := \arg\min_x \left[\left\| \mathcal{A}(x) - y \right\|_Y^2 + \mathcal{S}_\theta(x) \right]$$

Drawbacks

- Lack of flexibility: $\mathcal{S}_\theta$ can only capture simplistic a priori information, parameter(s) θ need to be selected explicitly
- Computational complexity: Large computational burden

Lack of flexibility

Capturing complex a priori information: Why do we need machine learning?

A priori information: The image contains a rabbit.

Lack of flexibility

Capturing complex a priori information: Why do we need machine learning?

A priori information: The image contains a rabbit.

Designing $\mathcal{S}_\theta \Rightarrow$ parametrise the notion of a rabbit

Connected structure consisting of fused ellipsoids, two specific elongated structures near each other (ears), fibre-like texture (fur) with specific range of colours, ... then it is a rabbit!

Lack of flexibility

Capturing complex a priori information: Why do we need machine learning?

A priori information: The image contains a rabbit.

Designing $\mathcal{S}_\theta \Rightarrow$ parametrise the notion of a rabbit

Connected structure consisting of fused ellipsoids, two specific elongated structures near each other (ears), fibre-like texture (fur) with specific range of colours, . . . then it is a rabbit!

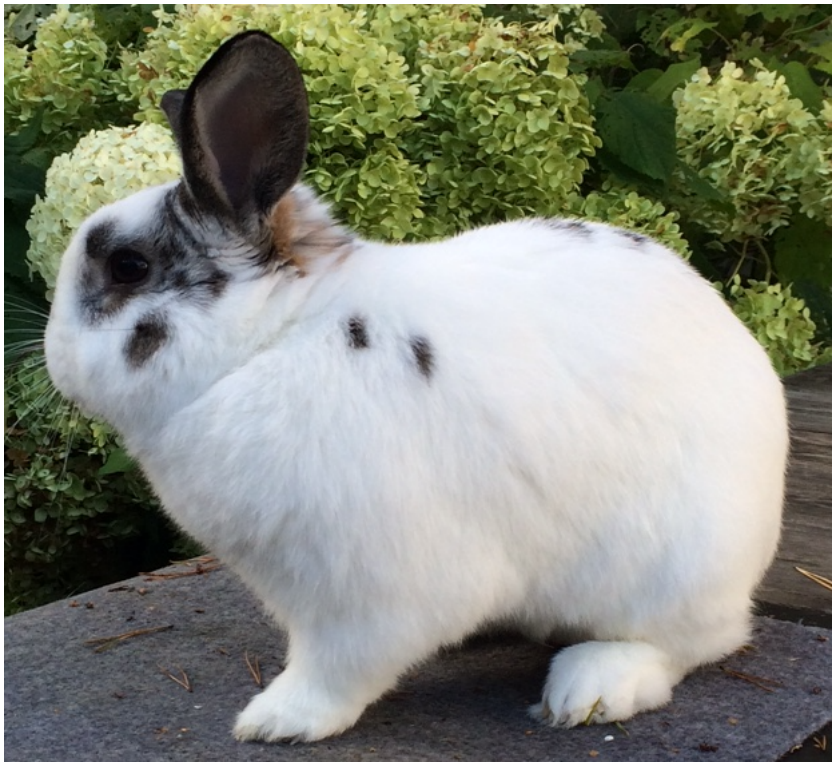

Lack of flexibility

Capturing complex a priori information: Why do we need machine learning?

A priori information: The image contains a rabbit.

Designing $\mathcal{S}_\theta \Rightarrow$ parametrise the notion of a rabbit

Connected structure consisting of fused ellipsoids, two specific elongated structures near each other (ears), fibre-like texture (fur) with specific range of colours, . . . then it is a rabbit!

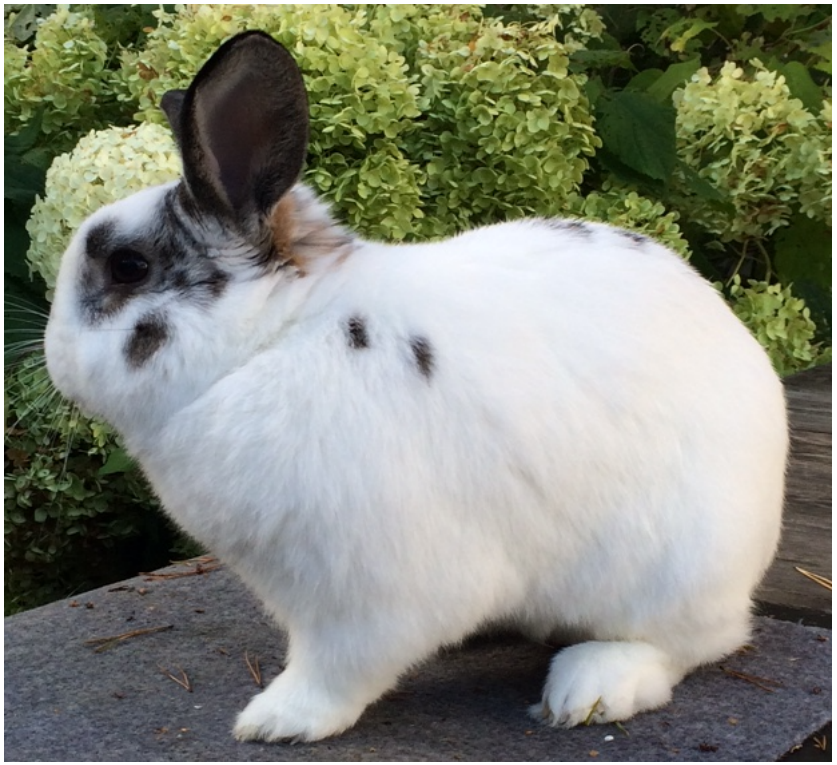

Learned iterative reconstruction
Approaches based on learning

Reconstruction methods based on learning

- Learn prior and/or regularisation parameter in variational methods $\Rightarrow$ Improves flexibility, not computational complexity

J. C. De Los Reyes et. al., *Bilevel parameter learning for higher-order total variation regularisation models*, Journal of Mathematical Imaging and Vision, 57(1):1–25, 2017

Q. Xu et. al., *Low-Dose X-ray CT Reconstruction via Dictionary Learning*, IEEE Transaction in Medical Imaging, 31(9):1682–1697, 2012

C. Rusu and J. Thompson, *Learning Fast Sparsifying Transforms*, IEEE Transactions on Signal Processing, 65(16): 4376–4378, 2017

T. Meinhardt et. al., *Learning Proximal Operators: Using Denoising Networks for Regularizing Inverse Imaging Problems*, ICCV, 2017

I. Dokmanic et. al., *Inverse problems with invariant multiscale statistics*, CoRR, 2016

Y. Romano et. al., *The Little Engine That Could: Regularization by Denoising (RED)*, SIAM Journal of Imaging Sciences, 10(4):1804–1844, 2017

Learned iterative reconstruction
Approaches based on learning

Reconstruction methods based on learning

- Learn prior and/or regularisation parameter in variational methods $\Rightarrow$ Improves flexibility, not computational complexity
- Learn how to reconstruct, i.e. learn an optimisation solver $\Rightarrow$ Improves computational complexity, not flexibility

K. Gregor and Y. LeCun, *Learning Fast Approximations of Sparse Coding*, ICML, 2010.

S. Oymak et. al., *Sharp time-data tradeoffs for linear inverse problems*, ArXiv:1507.04793v2, 2016.

R. Giryeset. al., *Tradeoffs between convergence speed and reconstruction accuracy in inverse problems*, ArXiv:1605.09232v2, 2017.

Learned iterative reconstruction
Approaches based on learning

Reconstruction methods based on learning

- Learn prior and/or regularisation parameter in variational methods $\Rightarrow$ Improves flexibility, not computational complexity
- Learn how to reconstruct, i.e. learn an optimisation solver $\Rightarrow$ Improves computational complexity, not flexibility
- Learned reconstruction: Learn both the prior and how to reconstruct

Learned iterative reconstruction
Approaches based on learning

Reconstruction methods based on learning

- Learn prior and/or regularisation parameter in variational methods $\Rightarrow$ Improves flexibility, not computational complexity
- Learn how to reconstruct, i.e. learn an optimisation solver $\Rightarrow$ Improves computational complexity, not flexibility
- Learned reconstruction: Learn both the prior and how to reconstruct
  - Fully learned reconstruction
  - Learned post-processing

M. Argyrou et. al., *Tomographic image reconstruction based on artificial neural network (ANN) techniques*, NSS/MIC, 2012

P. Paschalis et. al., *Tomographic image reconstruction using artificial neural networks*, Nuclear Instruments and Methods in Physics Research Section A, 527(1-2):211–215, 2004

E. Kang et. al., *A deep convolutional neural network using directional wavelets for low-dose X-ray CT reconstruction*, Medical Physics, 44(10):360–375, 2017

K. Hwan et. al., *Deep Convolutional Neural Network for Inverse Problems in Imaging*, IEEE Transactions on Image Processing, 26(9):4509–4522, 2017.

Learned iterative reconstruction
Approaches based on learning

Reconstruction methods based on learning

- Learn prior and/or regularisation parameter in variational methods $\Rightarrow$ Improves flexibility, not computational complexity
- Learn how to reconstruct, i.e. learn an optimisation solver $\Rightarrow$ Improves computational complexity, not flexibility
- Learned reconstruction: Learn both the prior and how to reconstruct
  - Fully learned reconstruction
  - Learned post-processing
  - Learned iterative reconstruction

Y. Yang et. al., *ADMM-Net: A Deep Learning Approach for Compressive Sensing MRI*, NIPS, 2016

P. Putzky and M. Welling, *Recurrent inference machines for solving inverse problems*, arXiv:1706.04008, 2017

J. Adler and O. Öktem, *Solving ill-posed inverse problems using iterative deep neural networks*, arXiv:1704.04058. Accepted for publication in Inverse Problems, 2017.

J. Adler and O. Öktem, *Learned Primal-Dual Reconstruction*, arXiv:1707.06474. Submitted to IEEE Transactions in Medical Imaging, 2017.

K. Hammernick et. al., *Learning a Variational Network for Reconstruction of Accelerated MRI Data*, arXiv:1704.00447, 2017.

Learned iterative reconstruction
Learned gradient descent

Algorithm Gradient descent

1: **for** $i = 1, \dots$ **do**
2: $\quad x_{i+1} \leftarrow x_i - \alpha \left[\partial \mathcal{A}(x_i)\right]^* \left(\mathcal{A}(x) - y\right)$

Solves variational problems with the following structure:

$$\min_x \left\| \mathcal{A}(x) - y \right\|_Y^2$$

Learned iterative reconstruction

Learned gradient descent

Algorithm Gradient descent

1: **for** $i = 1, \dots$ **do**
2: $\quad x_{i+1} \leftarrow x_i - \alpha \left[\partial \mathcal{A}(x_i)\right]^* \left(\mathcal{A}(x) - y\right)$

Solves variational problems with the following structure:

$$\min_x \left\| \mathcal{A}(x) - y \right\|_Y^2$$

Algorithm Learned gradient descent

1: **for** $i = 1, \dots, N$ **do**
2: $\quad x_{i+1} \leftarrow \Lambda_\theta \Big(x_i, \left[\partial \mathcal{A}(x_i)\right]^* \left(\mathcal{A}(x_i) - y\right) \Big)$
3: $\mathcal{R}_\theta(y) \leftarrow x_N$

- Finite number of iterates N.
- Λ_θ given by a convolutional neural network (CNN), learn parameter θ from training data in $X \times Y$.

Learned iterative reconstruction

Learned primal-dual scheme

Algorithm Non-linear PDHG

1: Initialise: $\sigma, \tau > 0$ s.t. $\sigma\tau\|\mathcal{A}\|^2 < 1$, $\zeta \in [0,1]$ and $x_0 \in X$, $y_0 \in Y$.
2: **for** $i = 1, \dots$ **do**
3: $\quad y_{i+1} \leftarrow \mathrm{prox}_{\sigma F^*}\big(y_i + \sigma\mathcal{A}(\bar{x}_i)\big)$
4: $\quad x_{i+1} \leftarrow \mathrm{prox}_{\tau G}\big(x_i - \tau[\partial\mathcal{A}(x_i)]^*(y_{i+1})\big)$
5: $\quad \bar{x}_{i+1} \leftarrow x_{i+1} + \zeta(x_{i+1} - x_i)$

Solves variational problems with the following structure:

$$\min_{x\in X}\Big[F\big(\mathcal{A}(x)\big) + G(x)\Big]$$

Algorithm Learned PDHG

1: Initialise: $x_0 \in X, y_0 \in Y$
2: **for** $i = 1, \dots, N$ **do**
3: $\quad y_{i+1} \leftarrow \Gamma_{\theta^d_{i+1}}\big(y_i, \mathcal{A}(x_i), y\big)$
4: $\quad x_{i+1} \leftarrow \Lambda_{\theta^p_{i+1}}\big(x_i, [\partial\mathcal{A}(x_i)]^*(y_{i+1})\big)$
5: $\mathcal{R}_\Theta(y) \leftarrow x_N$ with $\Theta = (\theta^d_1, \dots, \theta^p_N)$

- Unrolled primal-dual scheme inspired by the PDHG method.
- $\Gamma_{\theta^d_i}$ and $\Lambda_{\theta^p_i}$ given by CNNs, learn θ^d_i and θ^p_i from training data in $X \times Y$ and # iterates N.
- Primal and dual updates may have memory $\Longrightarrow$ learned primal dual method.

Learned iterative reconstruction

Learned primal-dual scheme

Algorithm Learned primal dual

1: Initialise: $x_0 \in X \times \ldots \times X$ (N_p times), $y_0 \in Y \times \ldots \times Y$ (N_d times)
2: **for** $i = 1, \ldots, N$ **do**
3: $\quad y_{i+1} \leftarrow \Gamma_{\theta^d_{i+1}}\left(y_i, \mathcal{A}(x_i^{(2)}), y\right)$
4: $\quad x_{i+1} \leftarrow \Lambda_{\theta^p_{i+1}}\left(x_i, [\partial\mathcal{A}(x_i^{(1)})]^*(y_{i+1}^{(1)})\right)$
5: $\mathcal{R}_\Theta(y) \leftarrow x_N^{(1)}$ with $\Theta = (\theta_1^d, \theta_1^p, \ldots, \theta_N^d, \theta_N^p)$

$$x_i = (x_i^{(1)}, \ldots, x_i^{(N_p)})$$

$$y_i = (y_i^{(1)}, \ldots, y_i^{(N_d)})$$

Learned iterative reconstruction

2D tomography of human phantom

Inverse problem: Recover attenuation coefficient from tomographic data (sinogram)

$$y = \mathcal{A}(x) + e$$

- Forward operator: 2D ray transform
- Geometry: Fan beam, 1000 lines/angle, 1000 angles
- Noise: Poisson noise (low dose CT)
- Image: 512×512 pixel
- Training data: 2 000 pairs (x_i, y_i)

Reconstruction methods:

- Filtered backprojection (FBP)
- Total variation (TV)
- FBP reconstruction de-noised by U-Net (FBP + U-Net)
- Learned primal-dual reconstruction

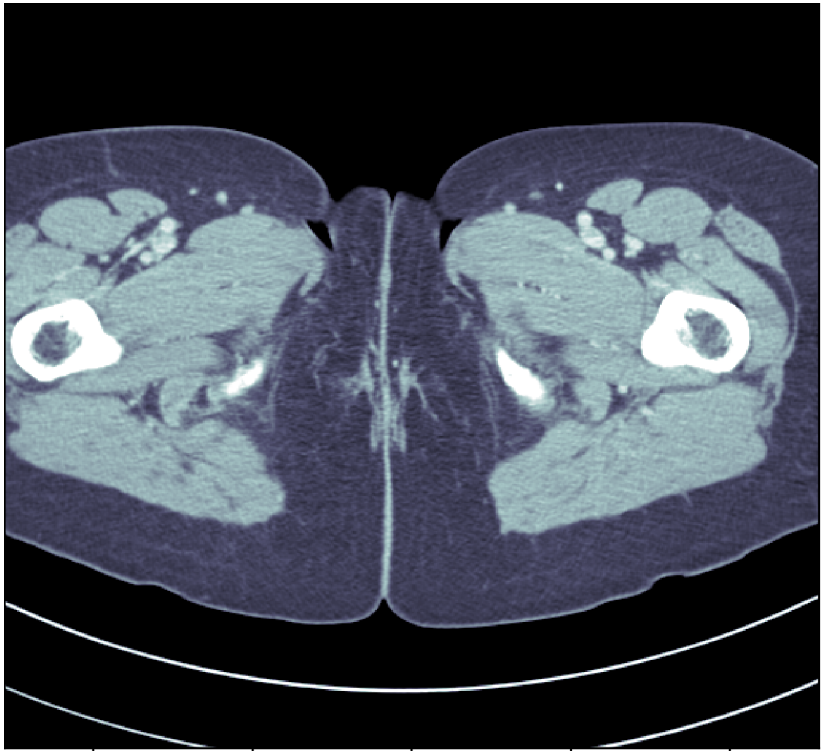

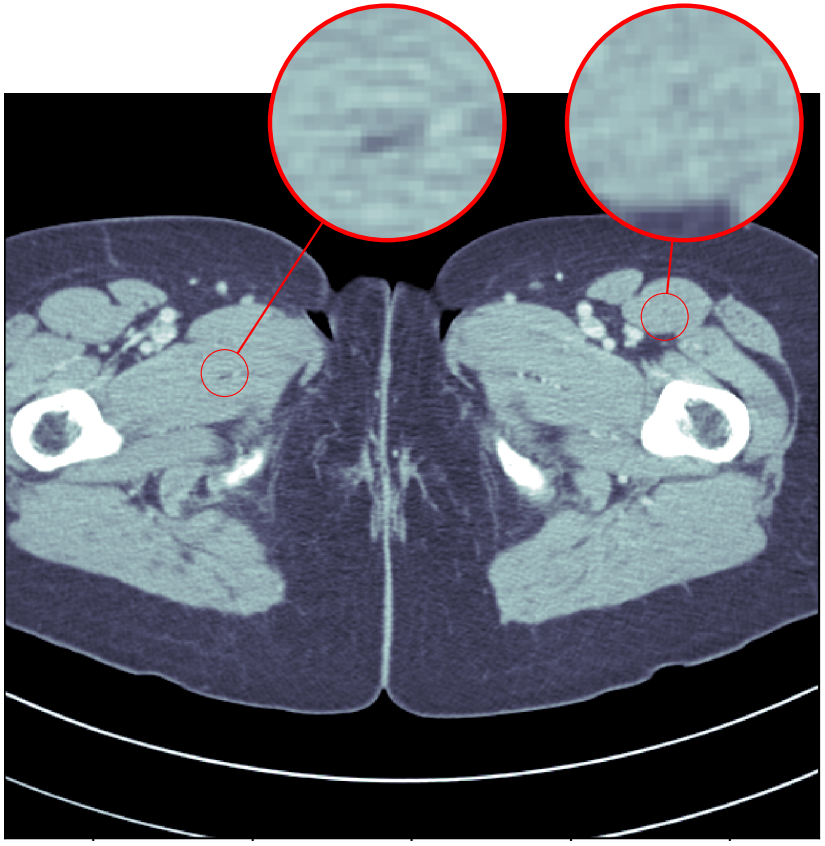

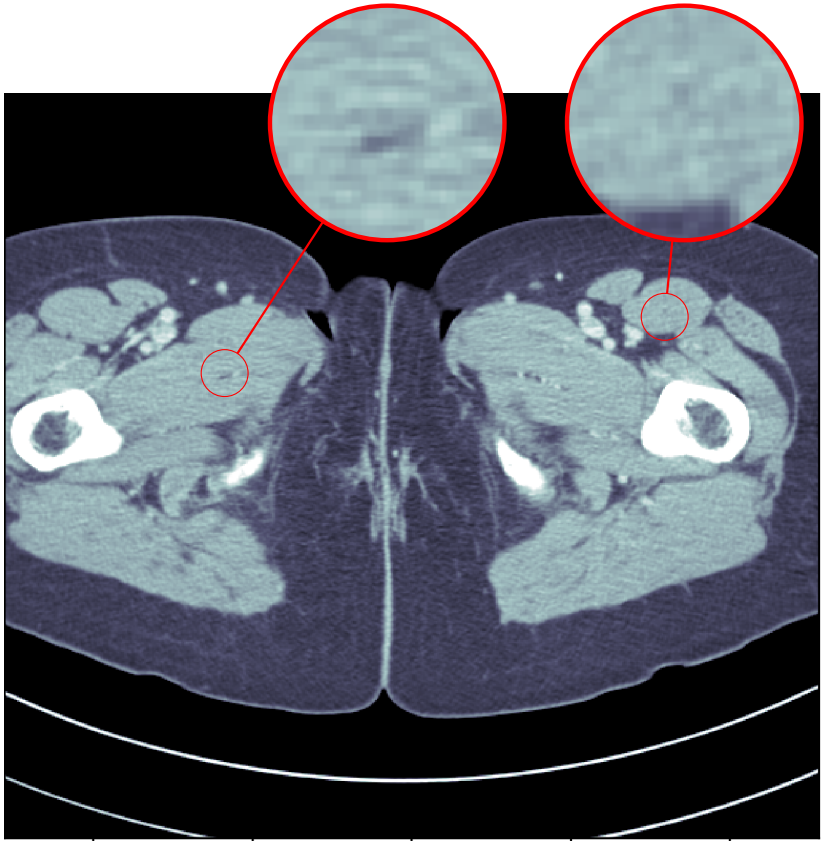

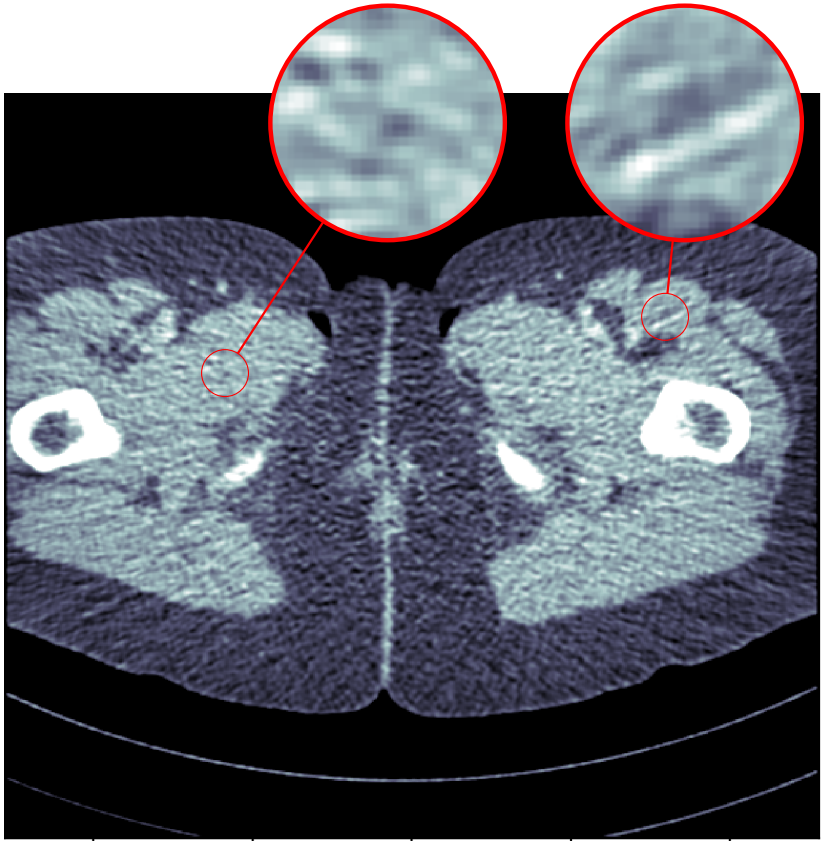

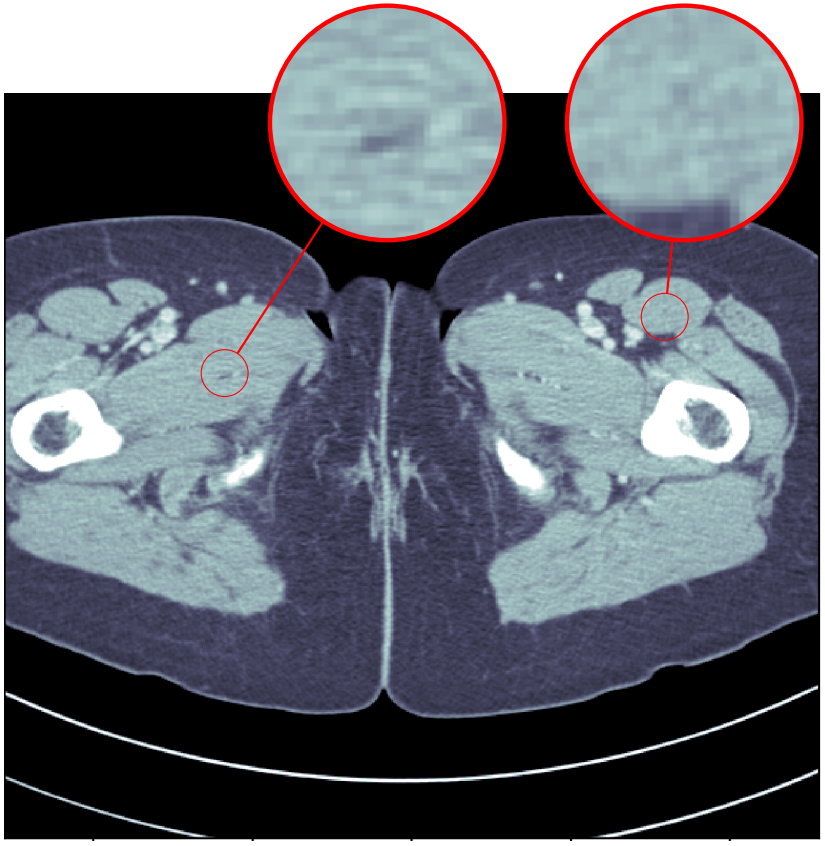

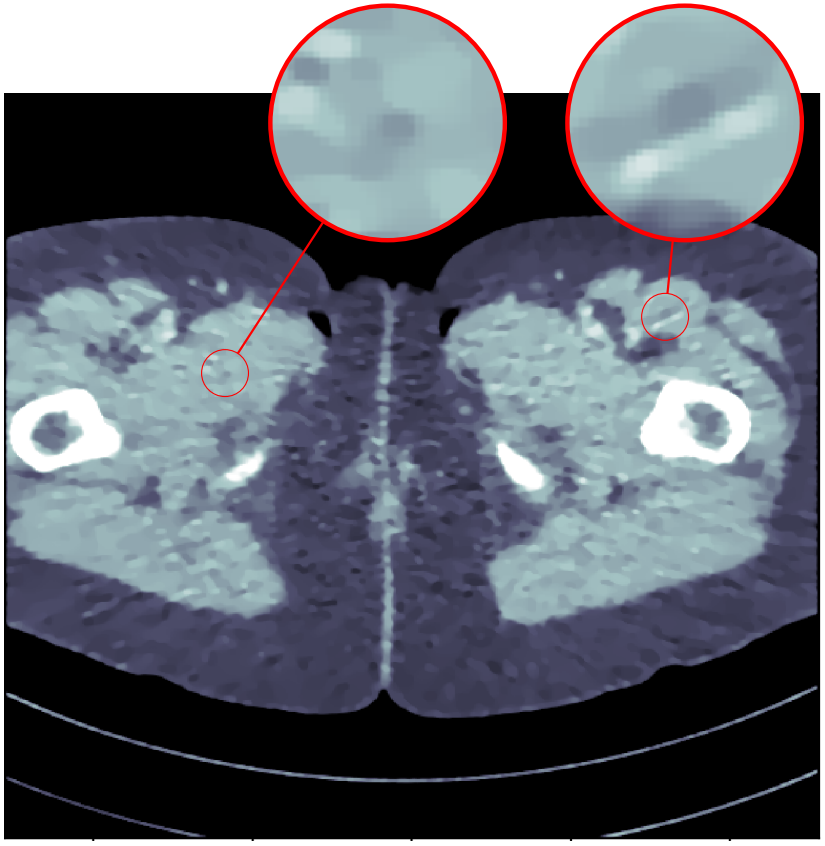

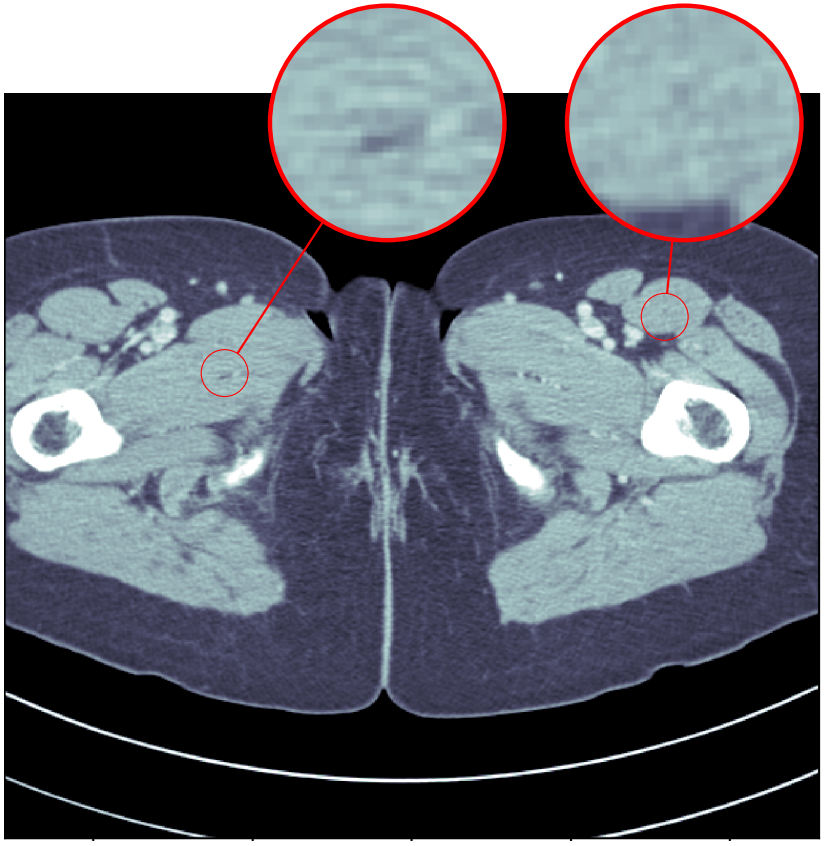

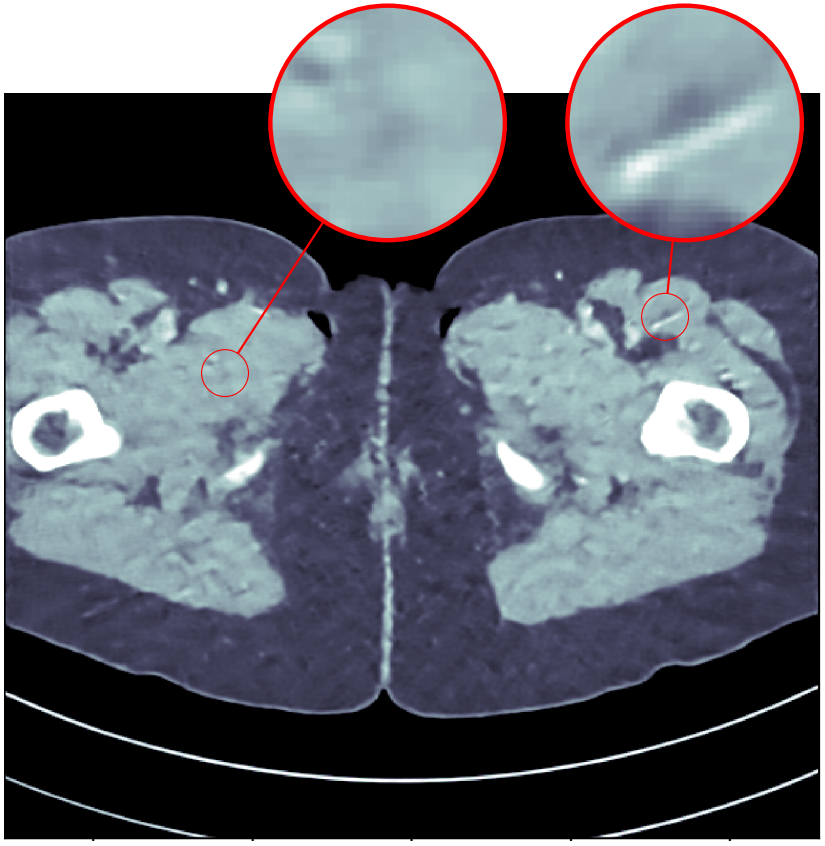

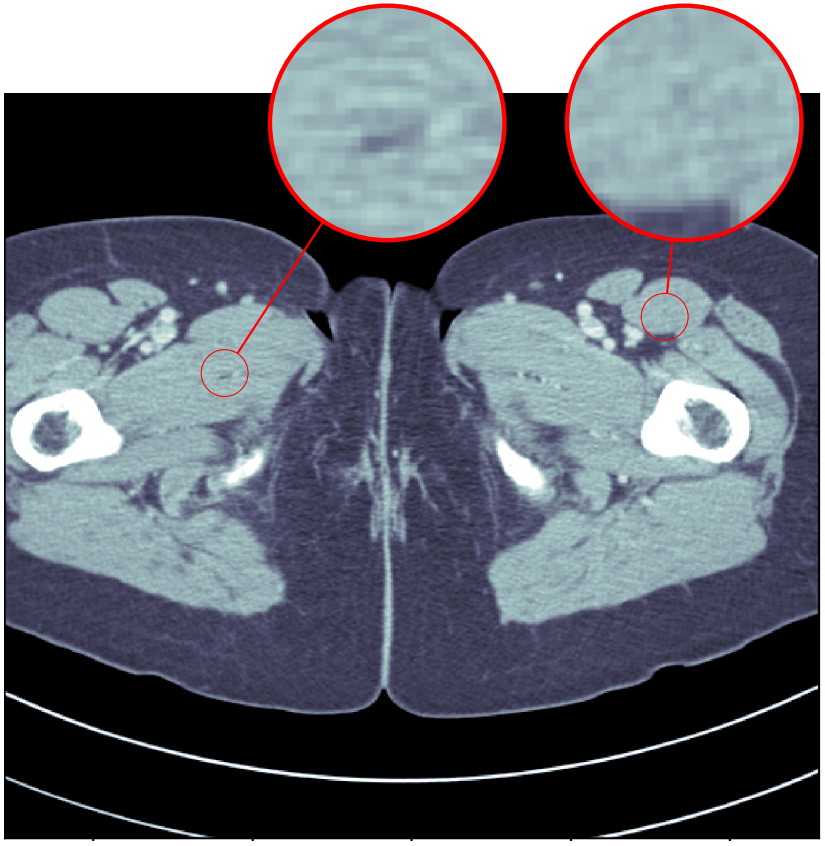

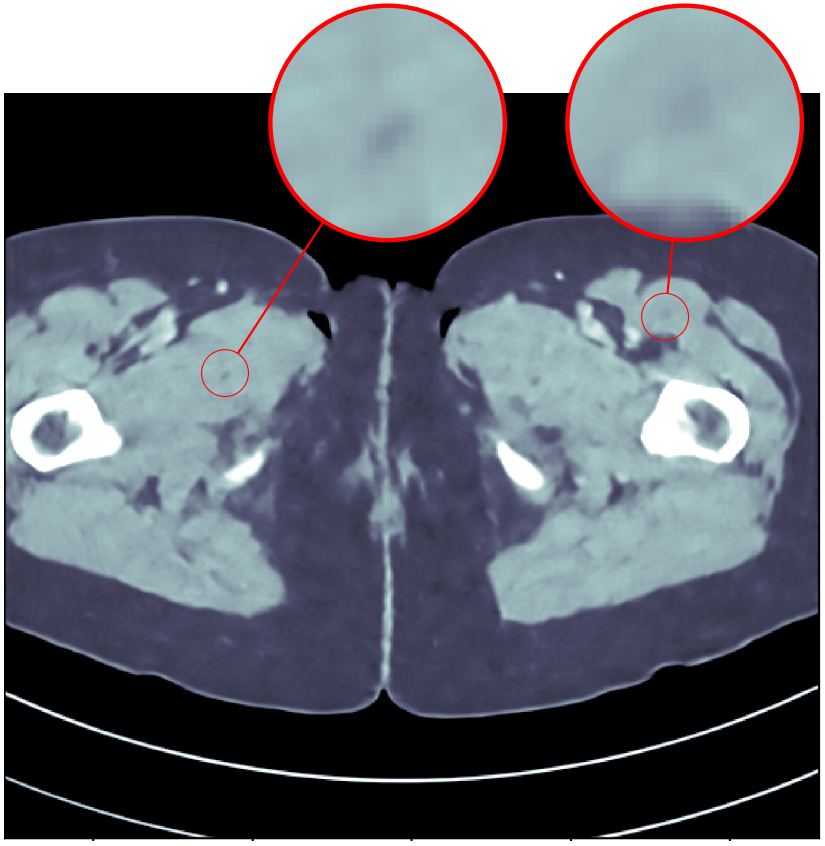

Learned iterative reconstruction
2D tomography of human phantom

Quantitative comparison

Method	PSNR (dB)	SSIM	Runtime (ms)	Parameters
FBP	33.65	0.829	423	1
TV	37.48	0.946	64 371	1
FBP + U-Net	41.92	0.941	463	10^7
Learned primal-dual	44.11	0.969	620	$2.4 \cdot 10^5$

SSIM = structural similarity index, (1 perfect match).

Comments

- Improved reconstruction quality against state-of-the-art
  - Very large improvement in PSNR
  - Significant clinically relevant improvement
- No need to manually set 'obscure' parameters
- Speedup enables clinical implementation

Learned iterative reconstruction
2D tomography of human phantom

Quantitative comparison

Method	PSNR (dB)	SSIM	Runtime (ms)	Parameters
FBP	33.65	0.829	423	1
TV	37.48	0.946	64 371	1
FBP + U-Net	41.92	0.941	463	10^7
Learned primal-dual	44.11	0.969	620	$2.4 \cdot 10^5$

SSIM = structural similarity index, (1 perfect match).

Comments

- Improved reconstruction quality against state-of-the-art
  - Very large improvement in PSNR
  - Significant clinically relevant improvement
- No need to manually set 'obscure' parameters
- Speedup enables clinical implementation

Task based reconstruction

Acknowledgements

Joint work with

- Jonas Adler (KTH/Elekta)
- Carola Schönlieb (University of Cambridge)
- Sebastian Lunz (University of Cambridge)

Task based reconstruction

- Reconstruction: Determine model parameters such that model predictions match measured data to sufficient accuracy.
- Decision making related to task: Reconstruction often only one step in pipeline for decision making related to a task.
  - Model parameters are typically summarised, either by an expert or by using specific descriptors, in an analysis step.
  - Summaries used as input for decision making.
- Natural question: Can one adapt the reconstruction method for the specific task at hand?
- Task based reconstruction: Methods for reconstruction that aim to integrate (parts of) decision making related to a task. Frequently also referred to as 'end-to-end'.

Task based reconstruction

- Reconstruction: Determine model parameters such that model predictions match measured data to sufficient accuracy.
- Decision making related to task: Reconstruction often only one step in pipeline for decision making related to a task.
  - Model parameters are typically summarised, either by an expert or by using specific descriptors, in an analysis step.
  - Summaries used as input for decision making.
- Natural question: Can one adapt the reconstruction method for the specific task at hand?
- Task based reconstruction: Methods for reconstruction that aim to integrate (parts of) decision making related to a task. Frequently also referred to as 'end-to-end'.

Task based reconstruction

- Reconstruction: Determine model parameters such that model predictions match measured data to sufficient accuracy.
- Decision making related to task: Reconstruction often only one step in pipeline for decision making related to a task.
  - Model parameters are typically summarised, either by an expert or by using specific descriptors, in an analysis step.
  - Summaries used as input for decision making.
- Natural question: Can one adapt the reconstruction method for the specific task at hand?
- Task based reconstruction: Methods for reconstruction that aim to integrate (parts of) decision making related to a task. Frequently also referred to as 'end-to-end'.

Task based reconstruction

- Reconstruction: Determine model parameters such that model predictions match measured data to sufficient accuracy.
- Decision making related to task: Reconstruction often only one step in pipeline for decision making related to a task.
  - Model parameters are typically summarised, either by an expert or by using specific descriptors, in an analysis step.
  - Summaries used as input for decision making.
- Natural question: Can one adapt the reconstruction method for the specific task at hand?
- Task based reconstruction: Methods for reconstruction that aim to integrate (parts of) decision making related to a task. Frequently also referred to as 'end-to-end'.

Task based reconstruction
Pipeline for radiomics

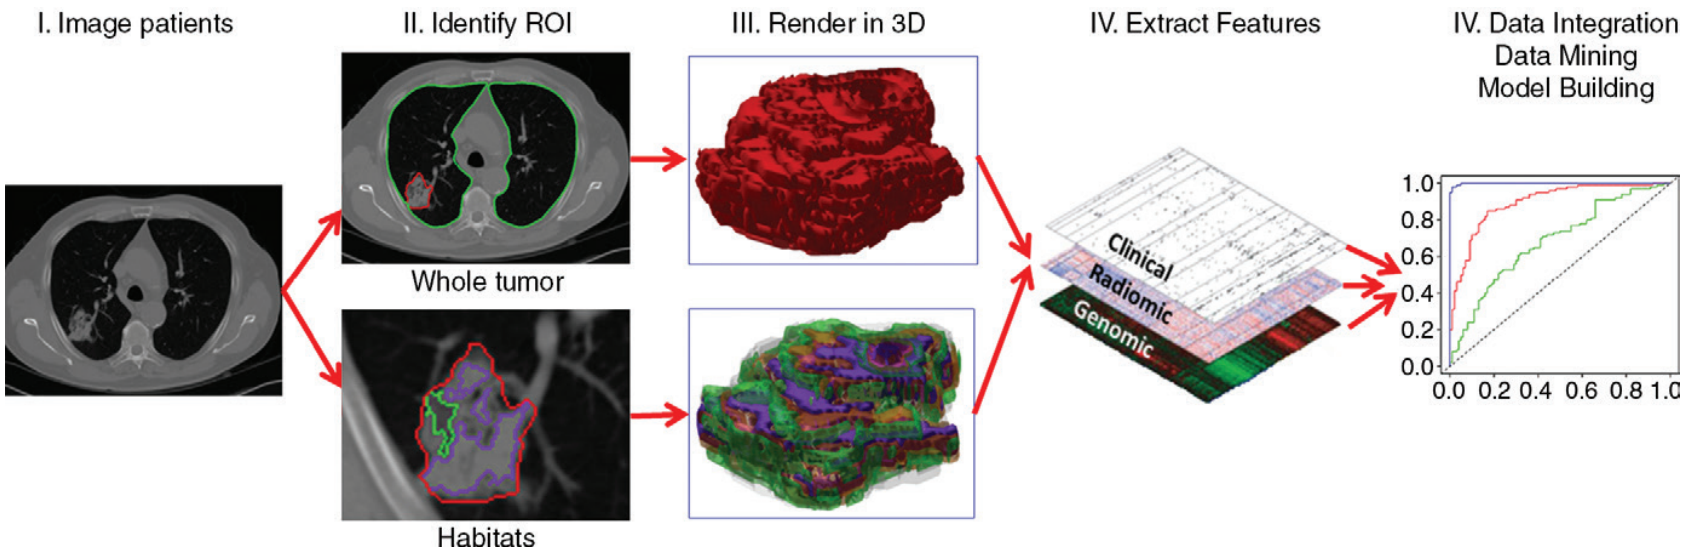

The process of radiomics and the use of radiomics in decision support.

R. J. Gillies et. al., *Radiomics: Images Are More than Pictures, They Are Data*, Radiology, 278(2):563–577, 2016.

Task based reconstruction

Pipeline for image guided patient management

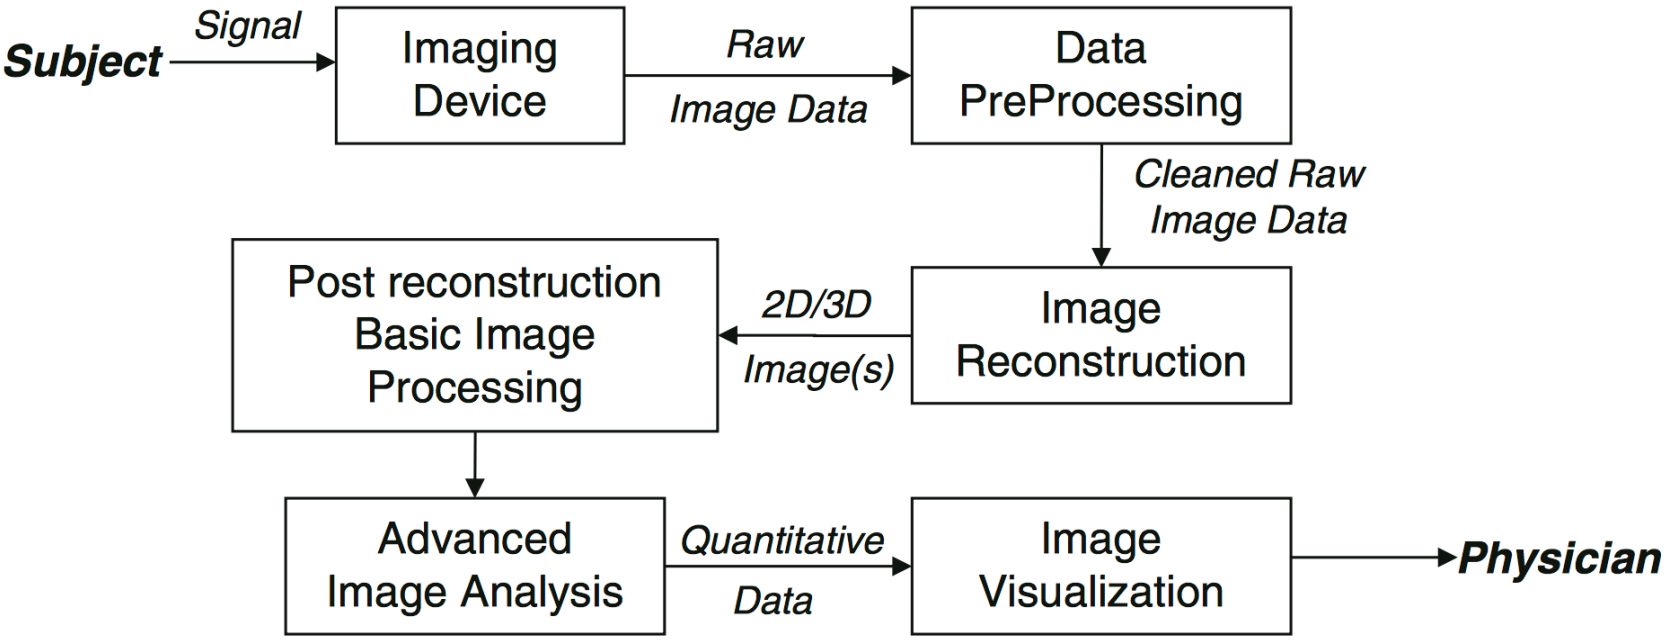

Typical workflow in medical imaging

Task based reconstruction

Pipeline for image guided patient management

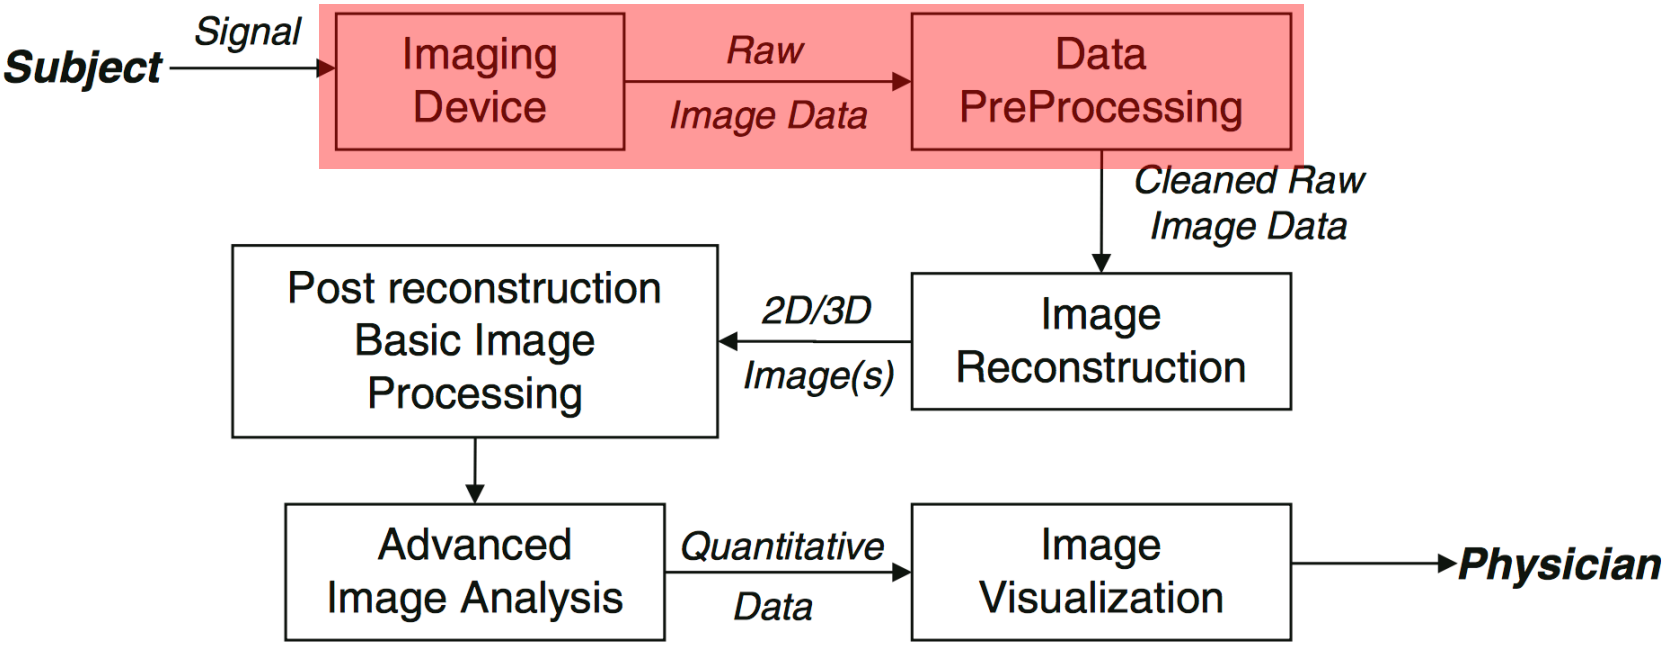

Acquisition protocol: Aquire raw data, pre-process to obtain cleaned data.

Task based reconstruction

Pipeline for image guided patient management

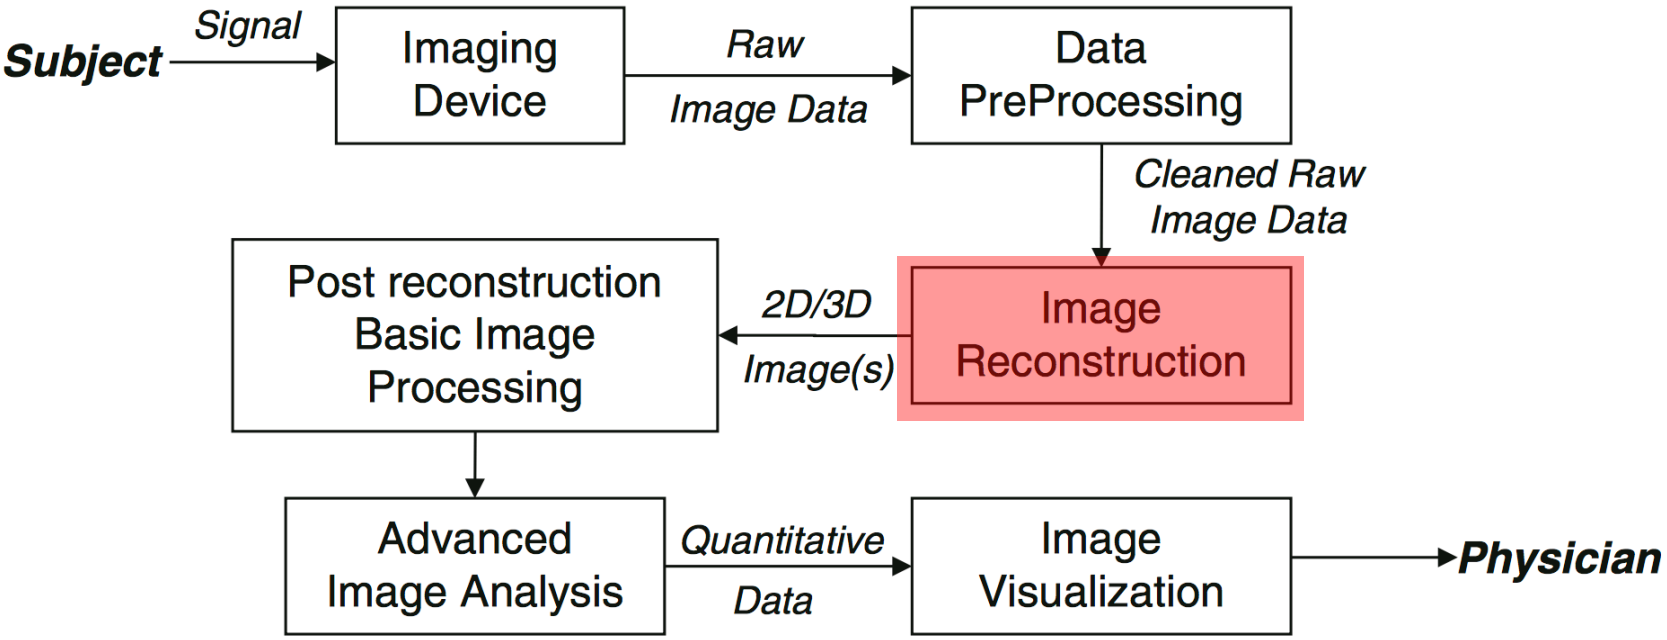

Reconstruction: Recover model parameters (image) from data

Task based reconstruction

Pipeline for image guided patient management

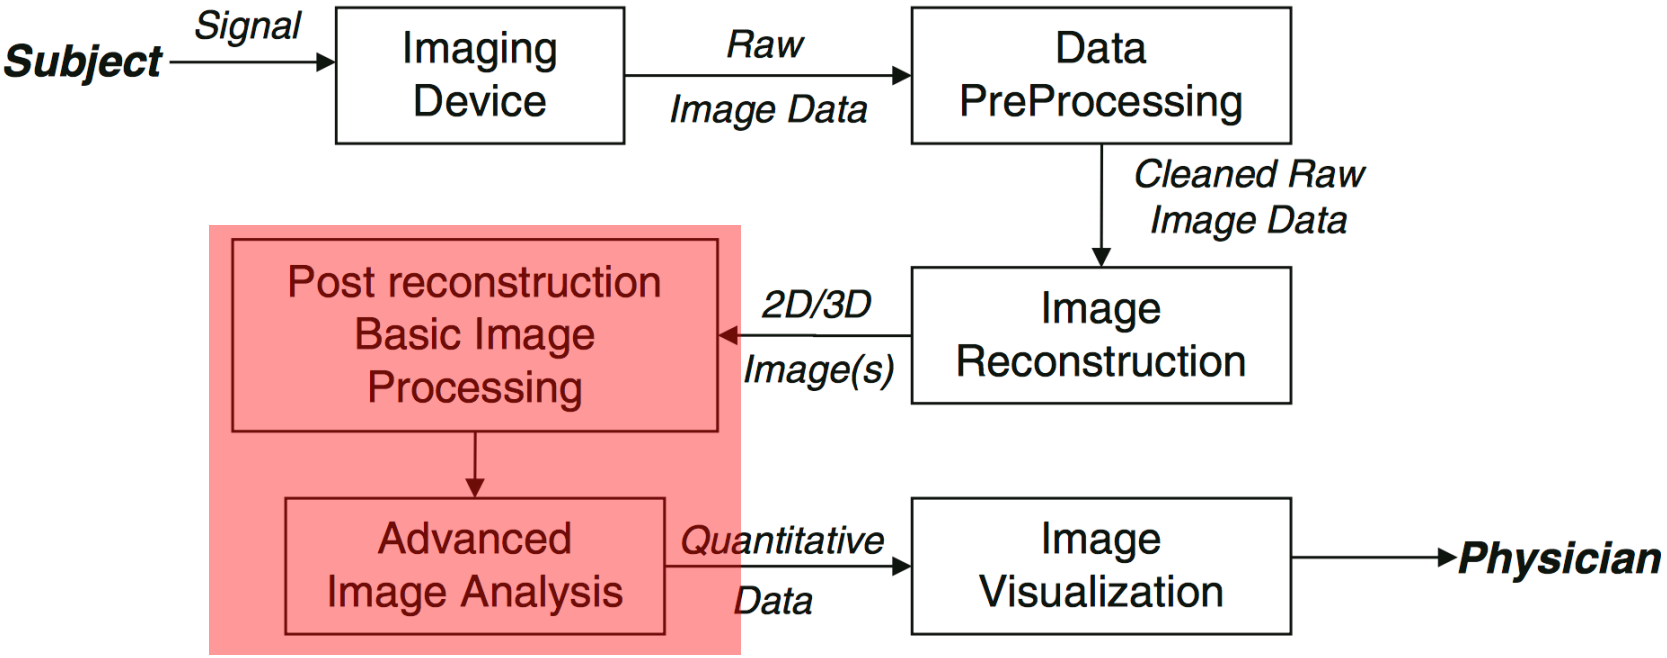

Post-processing: Extract features from model parameters

Task based reconstruction

Pipeline for image guided patient management

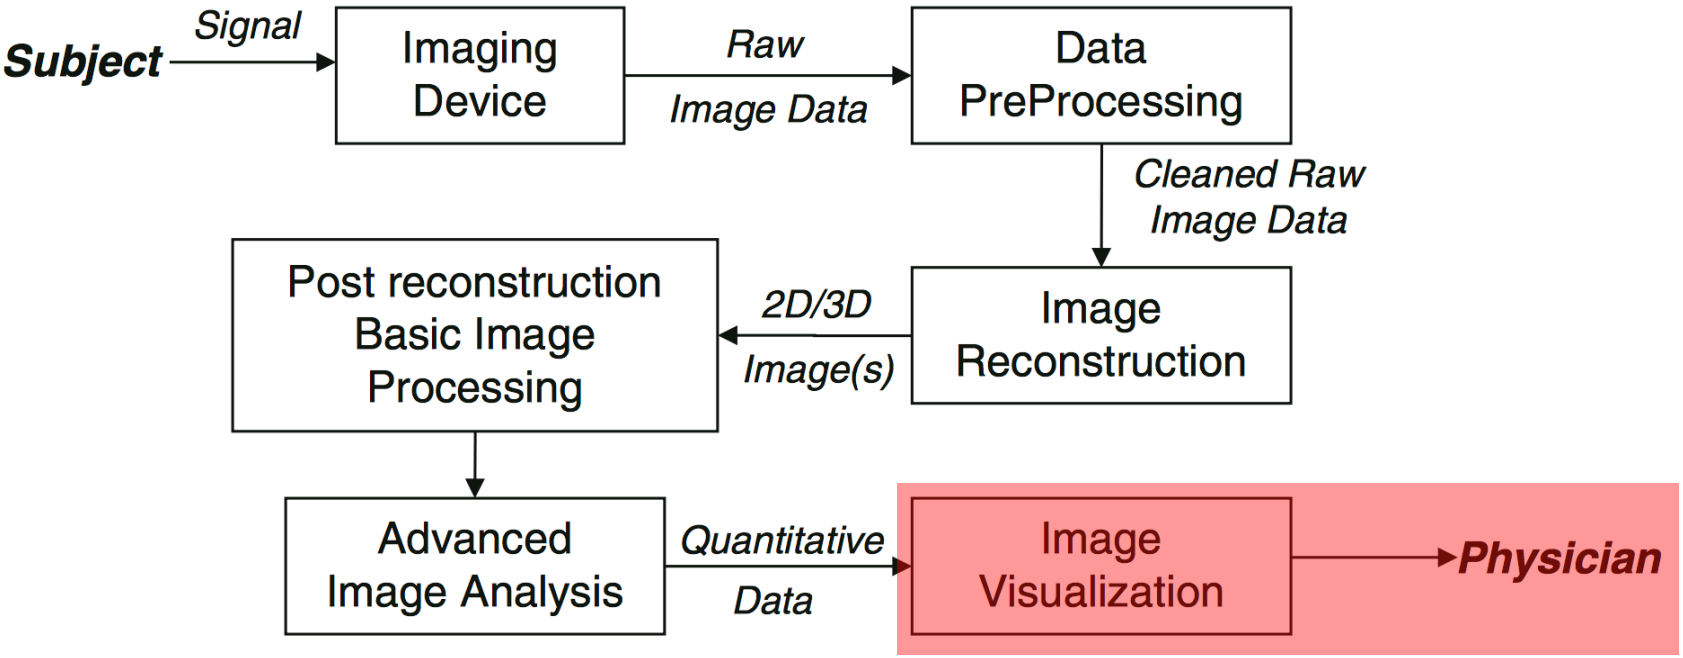

Decision making: Use extracted features for decision making.

Task based reconstruction

Pipeline for image guided patient management

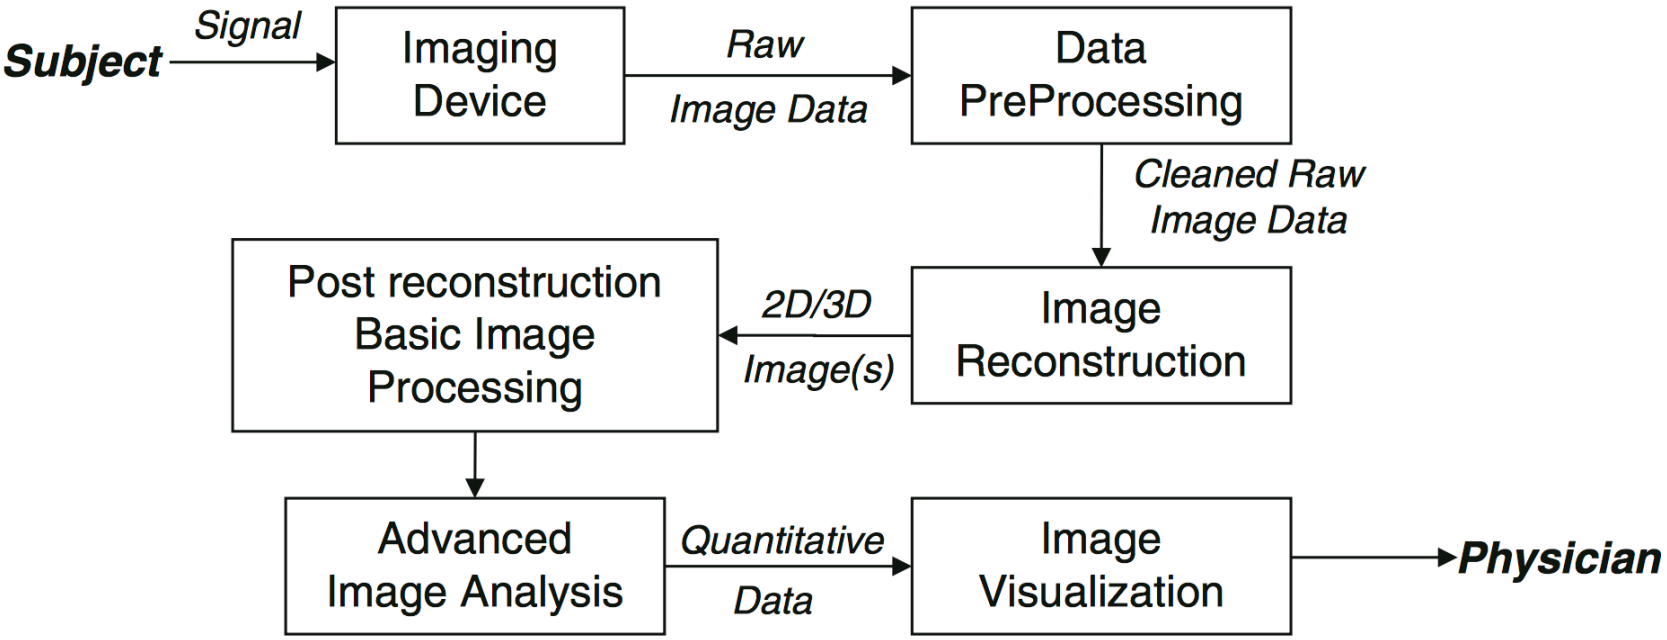

Performing these parts sequentially has several disadvantages:

- Each single step prone to introduce biases that are not accounted for by subsequent steps.
- Reconstruction does not consider end task.
- Feature extraction does not consider measured data.
- Task often only accounted for at the very final step.

Task based reconstruction

Pipeline for image guided patient management

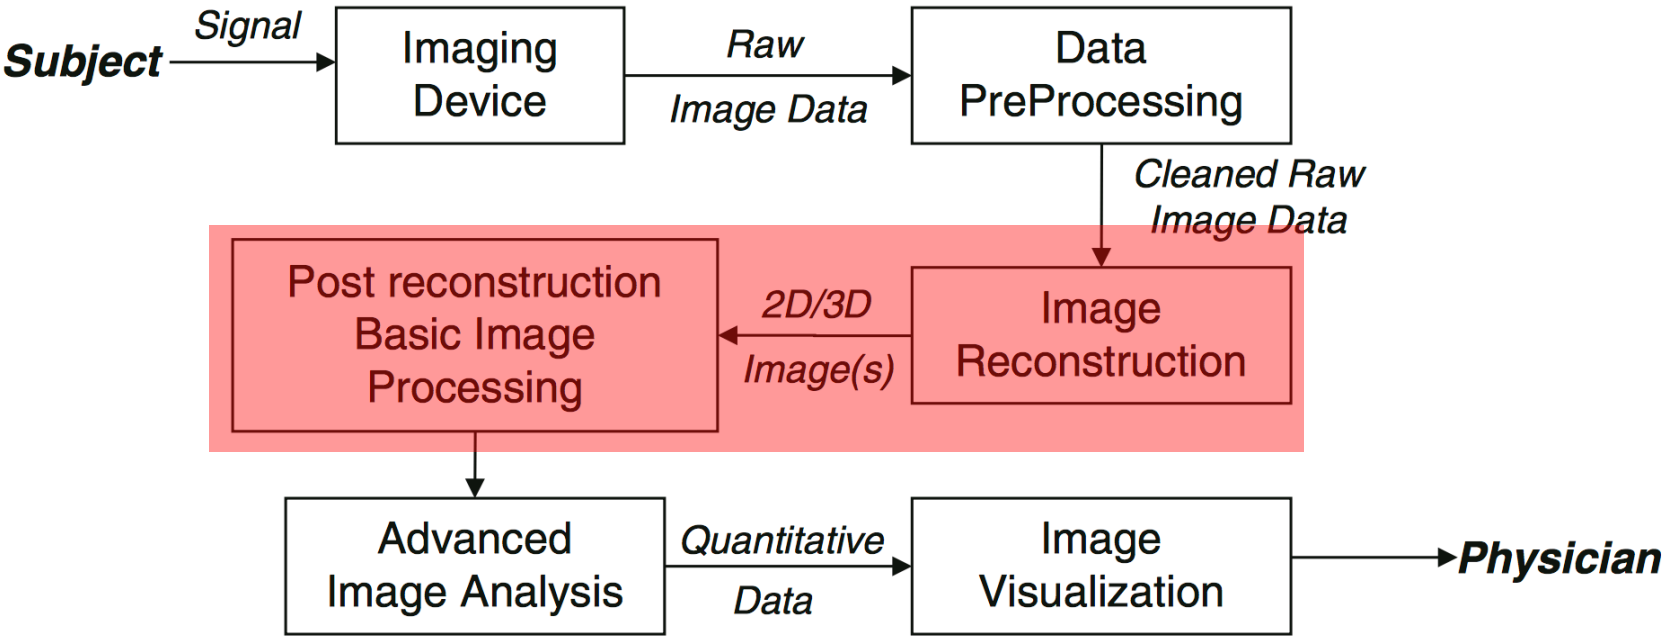

Task based reconstruction: Integrate parts into a common work-flow

- Cleaned data $\longrightarrow$ extracted simple features

Task based reconstruction

Pipeline for image guided patient management

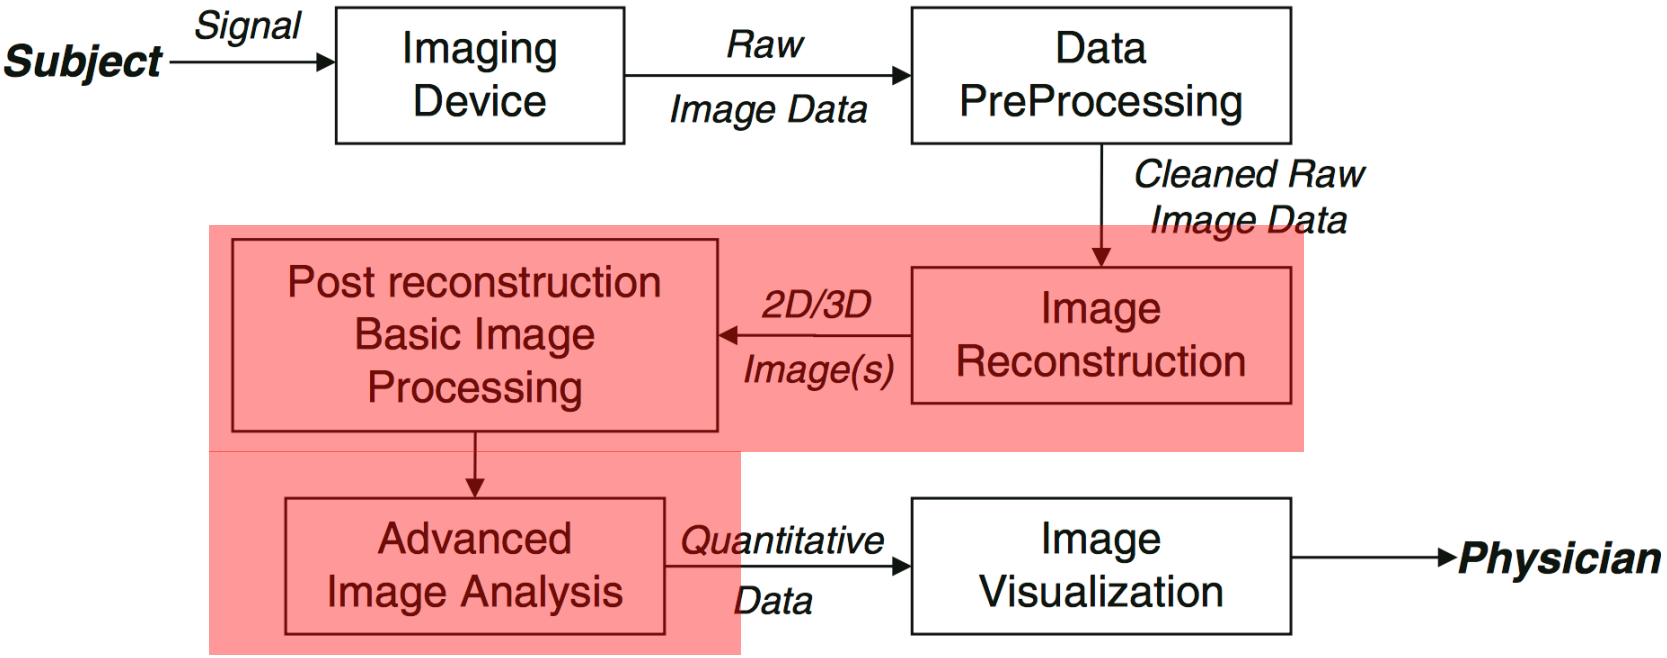

Task based reconstruction: Integrate parts into a common work-flow

- Cleaned data $\longrightarrow$ extracted simple features
- Cleaned data $\longrightarrow$ extracted advanced features

Task based reconstruction

Pipeline for image guided patient management

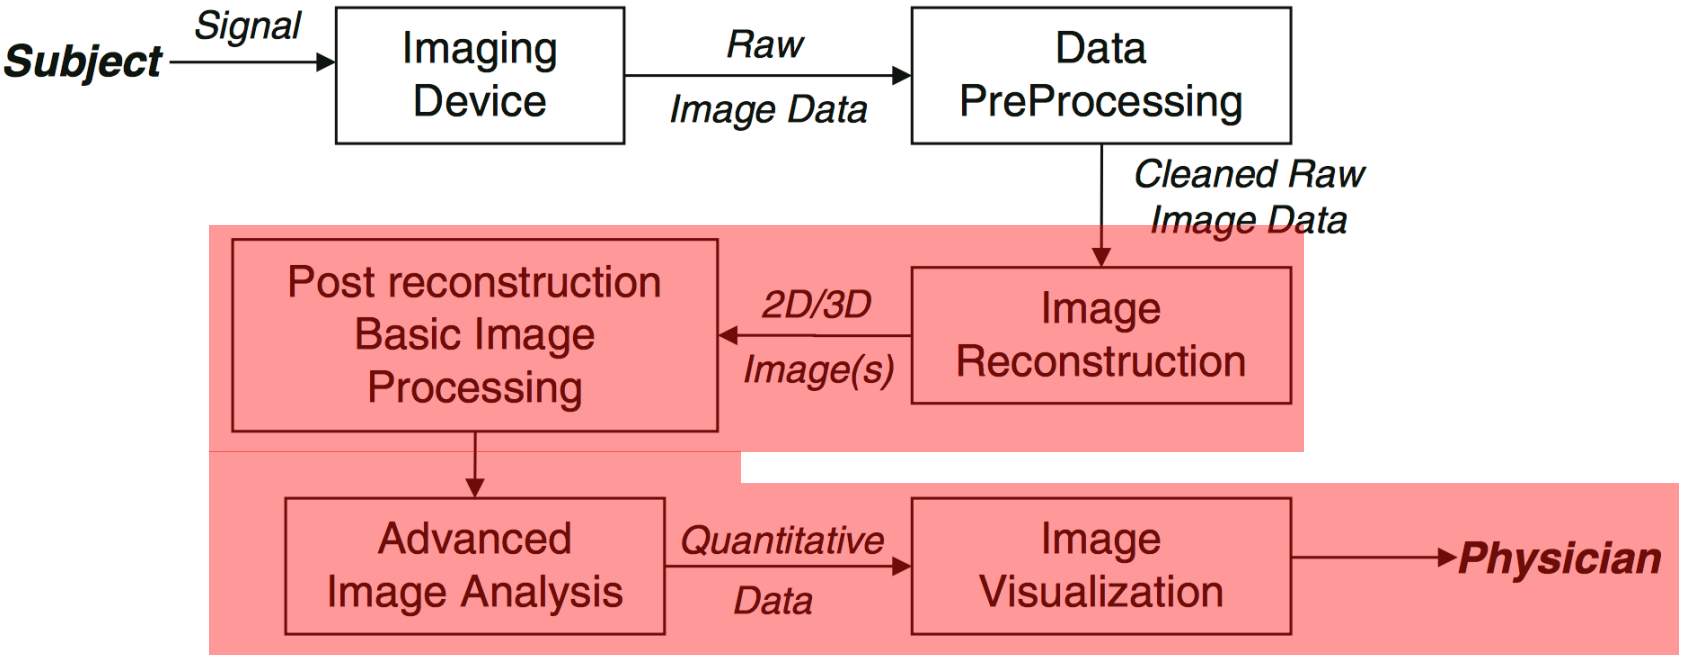

Task based reconstruction: Integrate parts into a common work-flow

- Cleaned data $\longrightarrow$ extracted simple features
- Cleaned data $\longrightarrow$ extracted advanced features
- Cleaned data $\longrightarrow$ decision making

Task based reconstruction

Pipeline for image guided patient management

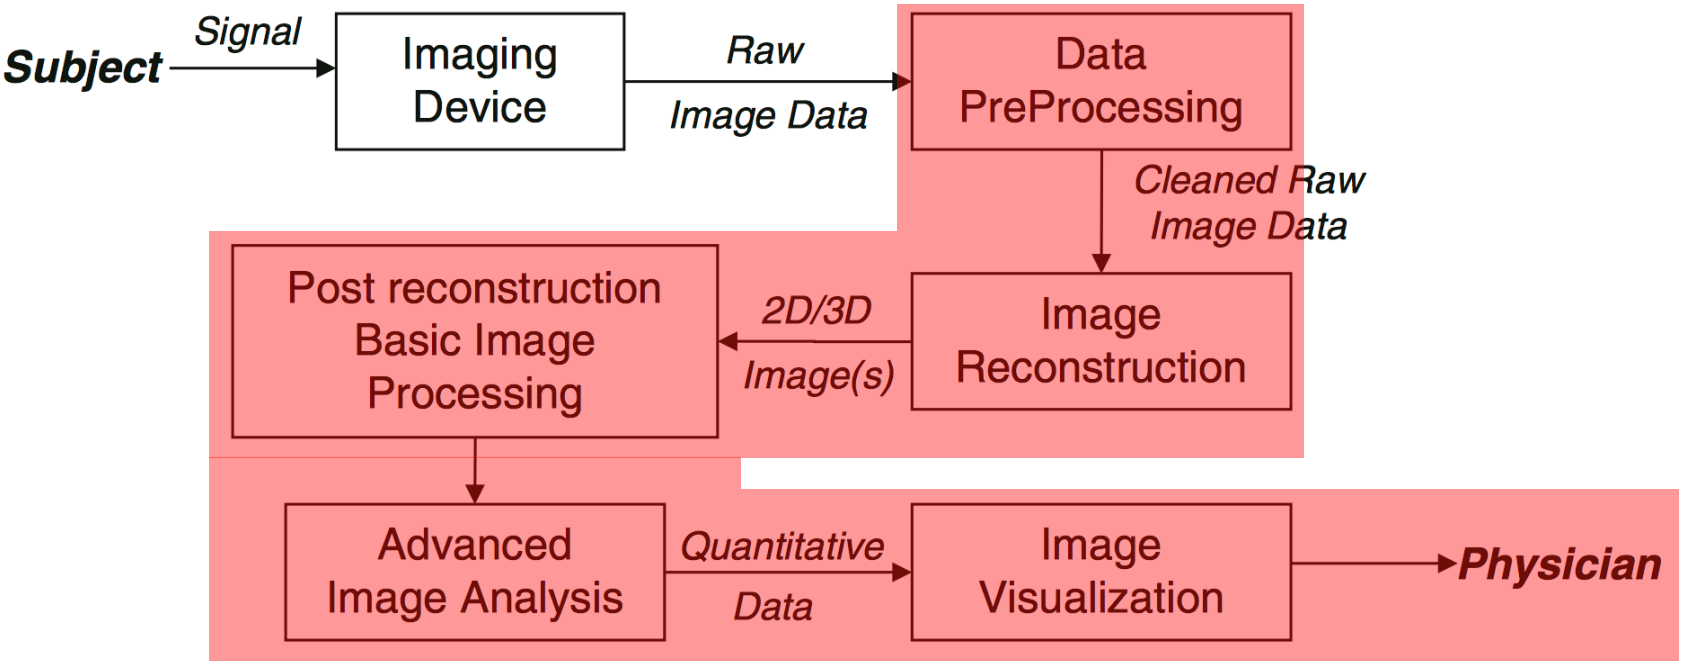

Task based reconstruction: Integrate parts into a common work-flow

- Cleaned data $\longrightarrow$ extracted simple features
- Cleaned data $\longrightarrow$ extracted advanced features
- Cleaned data $\longrightarrow$ decision making
- Raw data $\longrightarrow$ decision making

Task based reconstruction

Pipeline for image guided patient management

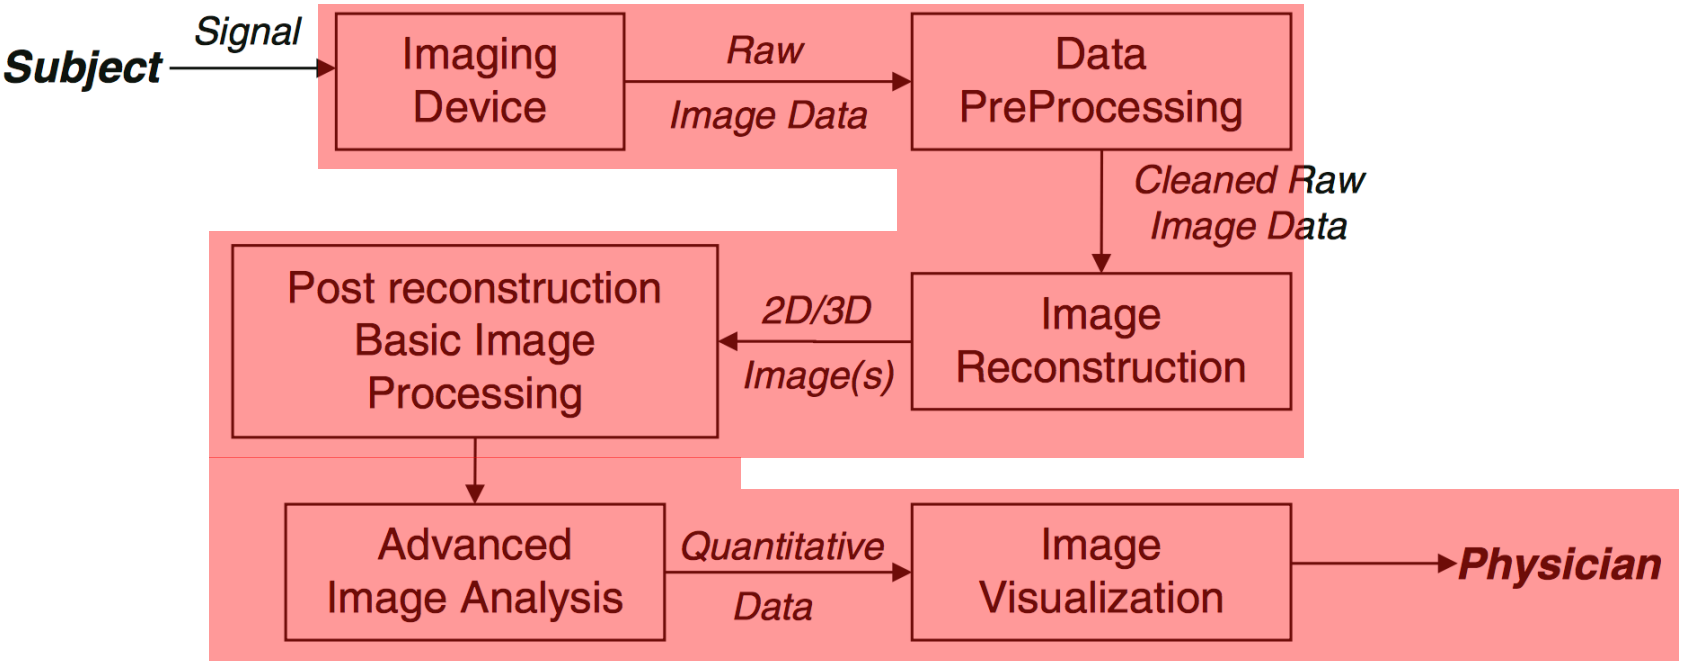

Task based reconstruction: Integrate parts into a common work-flow

- Cleaned data $\longrightarrow$ extracted simple features
- Cleaned data $\longrightarrow$ extracted advanced features
- Cleaned data $\longrightarrow$ decision making
- Raw data $\longrightarrow$ decision making
- Acquisition protocol $\longrightarrow$ decision making

Task based reconstruction

Pipeline for image guided patient management

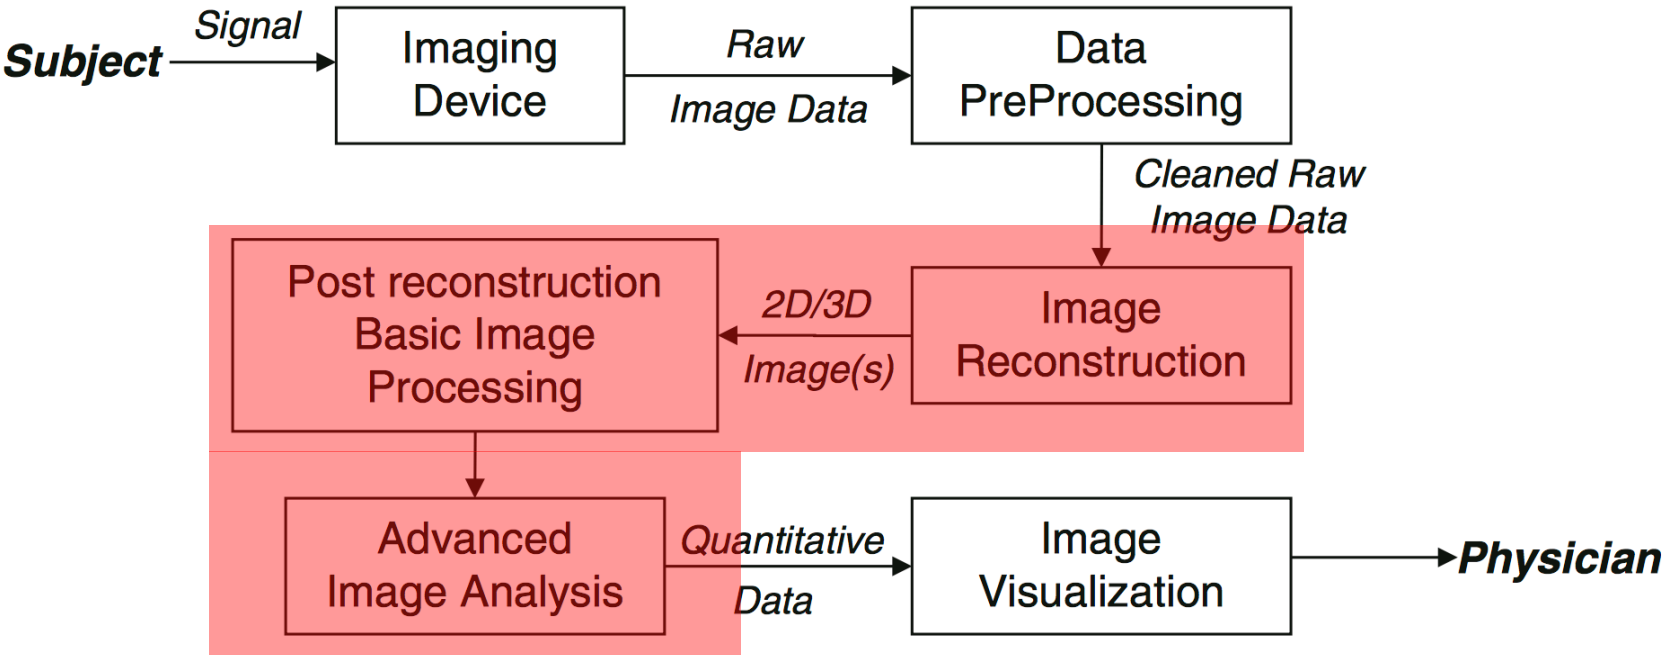

Task based reconstruction: Integrate parts into a common work-flow

- Cleaned data $\longrightarrow$ extracted simple features
- Cleaned data $\longrightarrow$ extracted advanced features
- Cleaned data $\longrightarrow$ decision making
- Raw data $\longrightarrow$ decision making
- Acquisition protocol $\longrightarrow$ decision making

Task based reconstruction

Abstract functional analytic framework

- Inverse problem: Recover model parameters $x_{\text{true}} \in X$ given data $y \in Y$ were $y = \mathcal{A}(x_{\text{true}}) + e$.
  - $\mathcal{A} \colon X \to Y$ forward operator
  - $e \in Y$ sample of Y-valued random variable $\mathbf{E}$.

Task based reconstruction

Abstract functional analytic framework

- Inverse problem: Recover model parameters $x_{\text{true}} \in X$ given data $y \in Y$ were $y = \mathcal{A}(x_{\text{true}}) + e$.
  - $\mathcal{A}\colon X \to Y$ forward operator
  - $e \in Y$ sample of Y-valued random variable $\mathbf{E}$.
- Feature reconstruction: As above, but goal is to recover features $z \in Z$ given data $y \in Y$ and $\mathcal{F}\colon X \to Z$ (feature extraction operator) such that:

Need to find the dotted operator.

$$\begin{tikzcd} X \arrow[r, "\mathcal{A}"] \arrow[d, "\mathcal{F}"'] & Y \arrow[dl, dotted] \\ Z & \end{tikzcd}$$

Task based reconstruction

Abstract functional analytic framework

- Inverse problem: Recover model parameters $x_{\text{true}} \in X$ given data $y \in Y$ were $y = \mathcal{A}(x_{\text{true}}) + e$.
  - $\mathcal{A} \colon X \to Y$ forward operator
  - $e \in Y$ sample of Y-valued random variable $\mathbf{E}$.
- Feature reconstruction: As above, but goal is to recover features $z \in Z$ given data $y \in Y$ and $\mathcal{F} \colon X \to Z$ (feature extraction operator) such that:

  Need to find the dotted operator.

$$\begin{tikzcd} X \arrow[r, "\mathcal{A}"] \arrow[d, "\mathcal{F}"'] & Y \arrow[dl, dotted] \\ Z & \end{tikzcd}$$

- $\mathcal{F}$ often highly non-injective, makes no sense to consider $\mathcal{A} \circ \mathcal{F}^{-1} \colon Z \to Y$ as 'new forward operator'.

Task based reconstruction

Abstract functional analytic framework: Examples from imaging

- Edge reconstruction
  - Lambda tomography
  - Variational regularisation using TV-type of regularisers
  - Sparse regularisation w.r.t. learned or analytic (curvelets, shearlets, beamlets, bandlets, . . .) directional dictionaries

V. P. Krishnan and E. T. Quinto, *Microlocal Analysis in Tomography* in Handbook of Mathematical Methods in Imaging, pp. 847–902, Springer, 2015.

M. Burger and S. Osher, *A Guide to the TV Zoo* in Level Set and PDE Based Reconstruction Methods in Imaging, pp. 1–70, Springer, 2013.

R. Rubinstein et. al., *Dictionaries for sparse representation modeling*, Proceedings of the IEEE, 98(6):1045?1057, 2010.

G. Kutyniok and D. Labate (eds), *Shearlets: Multiscale Analysis for Multivariate Data*, Springer, 2012.

Task based reconstruction

Abstract functional analytic framework: Examples from imaging

- Edge reconstruction
  - Lambda tomography
  - Variational regularisation using TV-type of regularisers
  - Sparse regularisation w.r.t. learned or analytic (curvelets, shearlets, beamlets, bandlets, ...) directional dictionaries
- Joint reconstruction and segmentation
  - Approximate inverse with segmentation operator
  - Variational regularisation using Mumford-Shah penalty
  - Level set approaches

A. K. Louis, *Feature reconstruction in inverse problems*, Inverse Problems, 27(6):065010, 2011

R. Ramlau and W. Ring, *A Mumford-Shah level-set approach for the inversion and segmentation of X-ray tomography data*, Journal of Computational Physics, 221:539–557, 2007.

K. Hohm et. al., *An algorithmic framework for Mumford-Shah regularization of inverse problems in imaging*, Inverse Problems, 31:115011, 2015.

S. Yoon et. al., *Simultaneous segmentation and reconstruction: A level set method approach for limited view computed tomography*, Medical Physics, 37(5): 2329–2340, 2010.

Task based reconstruction

Abstract functional analytic framework: Examples from imaging

- Edge reconstruction
  - Lambda tomography
  - Variational regularisation using TV-type of regularisers
  - Sparse regularisation w.r.t. learned or analytic (curvelets, shearlets, beamlets, bandlets, . . .) directional dictionaries
- Joint reconstruction and segmentation
  - Approximate inverse with segmentation operator
  - Variational regularisation using Mumford-Shah penalty
  - Level set approaches
- Joint reconstruction and pixel classification (labelling)
  - MAP estimation using Gauss-Markov-Potts priors

A. Mohammad-Djafar, *Gauss-Markov-Potts priors for images in computer tomography resulting to joint optimal reconstruction and segmentation*, International Journal of Tomography and Statistics, 11:76–92, 2009.

M. Romanov et. al., *Simultaneous tomographic reconstruction and segmentation with class priors*, Inverse Problems in Science and Engineering, 24:1432-1453, 2016.

Task based reconstruction

Abstract functional analytic framework: Examples from imaging

- Edge reconstruction
  - Lambda tomography
  - Variational regularisation using TV-type of regularisers
  - Sparse regularisation w.r.t. learned or analytic (curvelets, shearlets, beamlets, bandlets, ...) directional dictionaries
- Joint reconstruction and segmentation
  - Approximate inverse with segmentation operator
  - Variational regularisation using Mumford-Shah penalty
  - Level set approaches
- Joint reconstruction and pixel classification (labelling)
  - MAP estimation using Gauss-Markov-Potts priors
- Joint reconstruction and image registration
  - Optical flow based methods
  - LDDMM based diffeomorphic deformations
  - Optimal transport priors

M. Burger et. al., *On optical flow models for variational motion estimation* in Variational Methods, RICAM, pp. 226-252, 2016.

C. Chen and O. Öktem, *Indirect Image Registration with Large Diffeomorphic Deformations*, arXiv:1706.04048, 2017. Submitted to SIAM Journal on Imaging.

Task based reconstruction

Abstract functional analytic framework: Examples from imaging

- Edge reconstruction
  - Lambda tomography
  - Variational regularisation using TV-type of regularisers
  - Sparse regularisation w.r.t. learned or analytic (curvelets, shearlets, beamlets, bandlets, ...) directional dictionaries
- Joint reconstruction and segmentation
  - Approximate inverse with segmentation operator
  - Variational regularisation using Mumford-Shah penalty
  - Level set approaches
- Joint reconstruction and pixel classification (labelling)
  - MAP estimation using Gauss-Markov-Potts priors
- Joint reconstruction and image registration
  - Optical flow based methods
  - LDDMM based diffeomorphic deformations
  - Optimal transport priors

J. Karlsson and A. Ringh, *Generalized Sinkhorn iterations for regularizing inverse problems using optimal mass transport*, arXiv:1612.02273, 2017. Accepted for publication in SIAM Journal on Imaging.

Task based reconstruction

Functional analytic framework: Examples from imaging

State-of-the-art approaches based on variational formulation

Drawbacks

- Computationally unfeasible
- Can only handle 'simple' features that need further processing before usage in decision making.
- Single estimator of feature often inadequate for decision making, need to recover the feature along with its uncertainty.

Task based reconstruction
Computational framework

- Setting
  - $\mathbf{X}$ is X-valued random variable generating $x_{\text{true}} \in X$
  - $\mathbf{Y}$ is Y-valued random variable generating data: $\mathbf{Y} = \mathcal{A}(\mathbf{X}) + \mathbf{E}$
  - $\mathbf{Z}$ is Z-valued random variable generating task related feature.
  - Training data: Samples (x_i, y_i, z_i) of $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$
- Operators
  - Parametrised reconstruction: $\mathcal{R}_{\Theta_1} \colon Y \to X$
  - Parametrised task operator: $\mathcal{T}_{\Theta_2} \colon X \to Z$
  - Task based reconstruction: $\mathcal{T}_{\Theta_2} \circ \mathcal{R}_{\Theta_1} \colon Y \to Z$
- Distance notions:

$$\mathrm{d}_X \colon X \times X \to \mathbb{R} \quad \text{and} \quad \mathrm{d}_Z \colon Z \times Z \to \mathbb{R}$$

- Joint training: Set (Θ_1, Θ_2) by minimising joint 'expected' loss against training data

$$L(\Theta_1, \Theta_2) = \mathbb{E}_{(\mathbf{X},\mathbf{Y},\mathbf{Z})}\Big[C_1\, \mathrm{d}_X\big(\mathcal{R}_{\Theta_1}(\mathbf{Y}), \mathbf{X}\big) + C_2\, \mathrm{d}_Z\big(\mathcal{T}_{\Theta_2} \circ \mathcal{R}_{\Theta_1}(\mathbf{Y}), \mathbf{Z}\big)\Big]$$

Task based reconstruction
Computational framework

- Flexibility: Possible to jointly set Θ_1 (in parametrised reconstruction $\mathcal{R}_{\Theta_1}$) and Θ_2 (in task operator $\mathcal{T}_{\Theta_2}$).
- Computational complexity: Possible to apply task based reconstruction $\mathcal{T}_{\Theta_2} \circ \mathcal{R}_{\Theta_1}$ on large-scale problems.

Task based reconstruction
Computational framework

- Flexibility: Possible to jointly set Θ_1 (in parametrised reconstruction $\mathcal{R}_{\Theta_1}$) and Θ_2 (in task operator $\mathcal{T}_{\Theta_2}$).
- Computational complexity: Possible to apply task based reconstruction $\mathcal{T}_{\Theta_2} \circ \mathcal{R}_{\Theta_1}$ on large-scale problems.

- $\mathcal{R}_{\Theta_1} \colon Y \to X$ given by learned iterative reconstruction.
- $\mathcal{T}_{\Theta_2} \colon X \to Z$ typically given by off-the shelf machine learning method relevant for the task (without reconstruction step):
  - segmentation
  - image comparison
  - classification

Task based reconstruction
Computational framework

- Flexibility: Possible to jointly set Θ_1 (in parametrised reconstruction $\mathcal{R}_{\Theta_1}$) and Θ_2 (in task operator $\mathcal{T}_{\Theta_2}$).
- Computational complexity: Possible to apply task based reconstruction $\mathcal{T}_{\Theta_2} \circ \mathcal{R}_{\Theta_1}$ on large-scale problems.

- $\mathcal{R}_{\Theta_1} \colon Y \to X$ given by learned iterative reconstruction.
- $\mathcal{T}_{\Theta_2} \colon X \to Z$ typically given by off-the shelf machine learning method relevant for the task (without reconstruction step):
  - segmentation
  - image comparison
  - classification

Pre-training: If there are training data $\Sigma_1 \subset X \times Y$ and $\Sigma_2 \subset X \times Z$, then we can pre-train to learn preliminary values:

$$\Sigma_1 \implies \Theta_1 \in W_1 \implies \mathcal{R}_{\Theta_1} \colon Y \to X$$
$$\Sigma_2 \implies \Theta_2 \in W_2 \implies \mathcal{T}_{\Theta_2} \colon X \to Z$$

Examples of task based reconstruction

- Joint reconstruction and segmentation
- Joint reconstruction and image comparison
- Joint reconstruction and classification

Task based reconstruction
Joint reconstruction and segmentation

Task based reconstruction problem: Recover probability map of segmented part of the attenuation coefficient from a sinogram (tomographic data)

$$y = \mathcal{A}(x) + e$$

- Forward operator: 2D ray transform
- Geometry: Parallel beam, 183 lines/angle, 30 angles
- Noise: 0.1% additive Gaussian
- Image: 128×128 pixel
- Training data: 100 triplets (x_i, y_i, z_i) where z_i is the segmentation (binary image).
  Extended with data argumentation (± 5 pixel translation and $\pm 10^\circ$ rotation).

Task based reconstruction

Joint reconstruction and segmentation

- Learned reconstruction $\mathcal{R}_{\Theta_1} : Y \to X$
  - Learned primal dual
  - Loss: d_X is L^2-distance on X

J. Adler and O. Öktem, *Learned Primal-Dual Reconstruction*, arXiv:1707.06474. Submitted to IEEE Transactions in Medical Imaging, 2017.

Task based reconstruction
Joint reconstruction and segmentation

- Learned reconstruction $\mathcal{R}_{\Theta_1} : Y \to X$
  - Learned primal dual
  - Loss: d_X is L^2-distance on X
- Learned task operator $\mathcal{T}_{\Theta_2} : X \to Z$
  - $Z =$ probability distributions over binary images on Ω.
  - $Z =$ grey-scale images with values in $[0, 1]$ that gives probability that point is in segmented object.
  - $\mathcal{T}_{\Theta_2}$ given by 'off the shelf' U-net convolutional neural net architecture for segmentation
  - Loss: d_Z is cross entropy

J. Adler and O. Öktem, *Learned Primal-Dual Reconstruction*, arXiv:1707.06474. Submitted to IEEE Transactions in Medical Imaging, 2017.

Ronneberger O. et. al., *U-Net: Convolutional Networks for Biomedical Image Segmentation*. MICCAI 2015, Springer, LNCS 9351: 234–241, 2015.

Task based reconstruction

Joint reconstruction and segmentation

- Learned reconstruction $\mathcal{R}_{\Theta_1} \colon Y \to X$
  - Learned primal dual
  - Loss: d_X is L^2-distance on X
- Learned task operator $\mathcal{T}_{\Theta_2} \colon X \to Z$
  - $Z =$ probability distributions over binary images on Ω.
  - $Z =$ grey-scale images with values in $[0, 1]$ that gives probability that point is in segmented object.
  - $\mathcal{T}_{\Theta_2}$ given by 'off the shelf' U-net convolutional neural net architecture for segmentation
  - Loss: d_Z is cross entropy
- Task based reconstruction $\mathcal{T}_{\Theta_2} \circ \mathcal{R}_{\Theta_1} \colon Y \to Z$
  - No pre-training
  - Joint loss: $L(\Theta_1, \Theta_2) = C\, L_X(\Theta_1) + L_Z(\Theta_2)$

Phantom

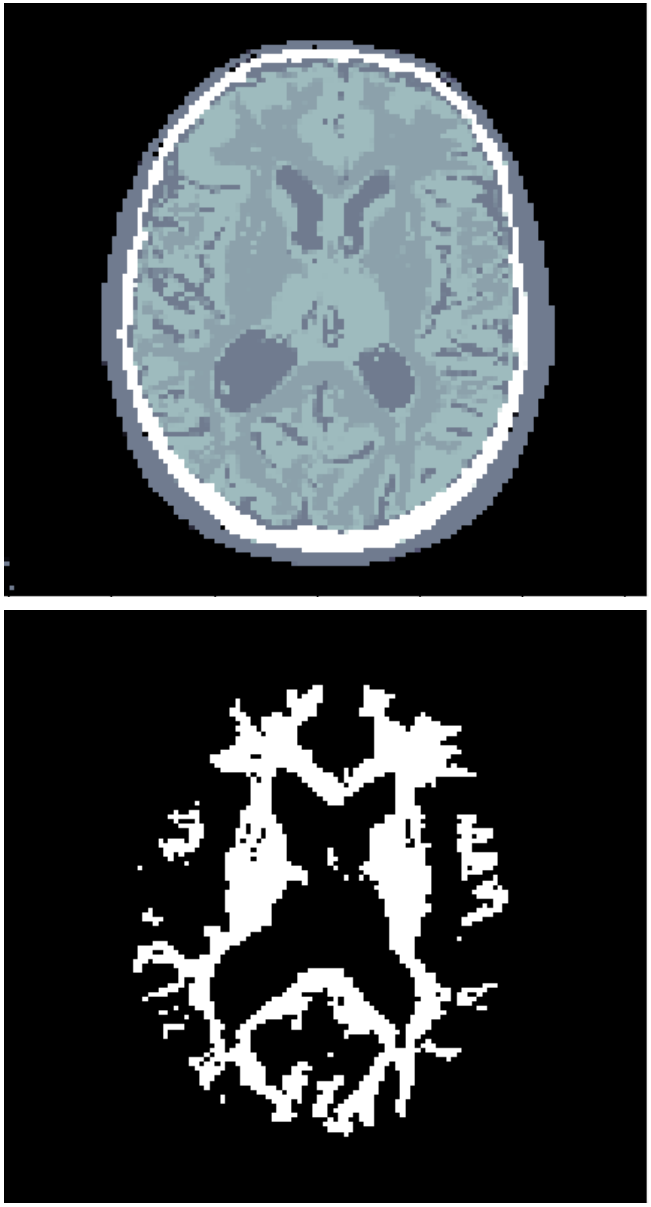

Joint rec. + seg.

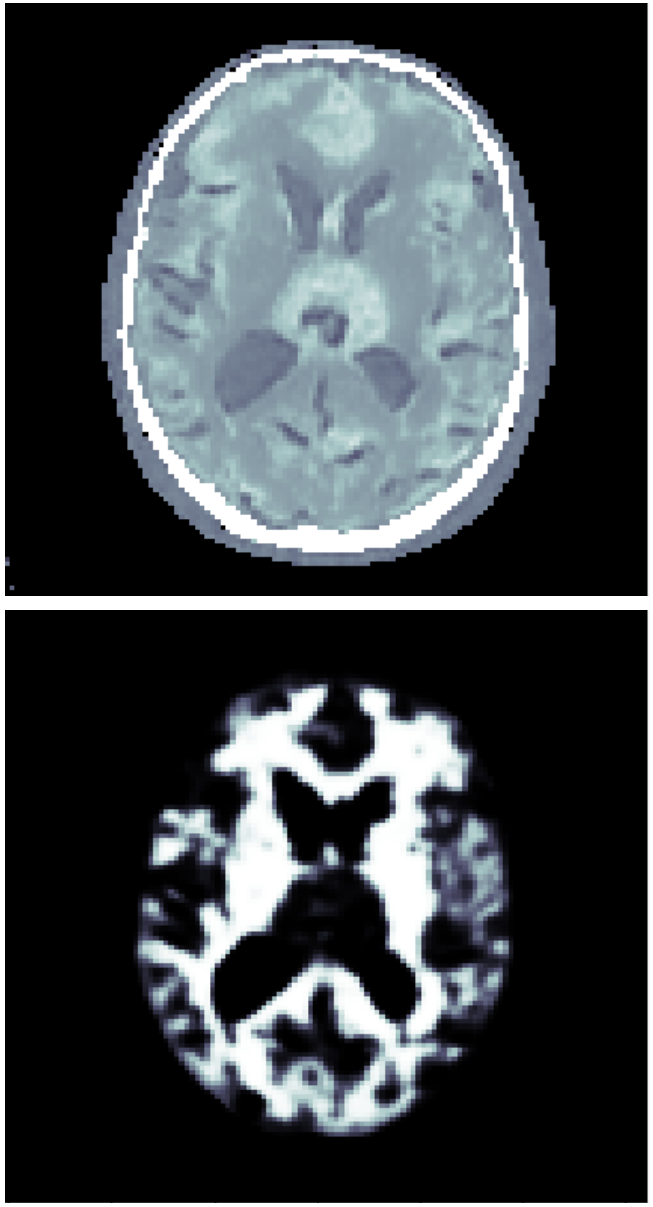

Task based reconstruction

Joint reconstruction and image comparison

Task based reconstruction problem: Recover the difference map for the attenuation coefficient in two images from their sinograms (tomographic data)

$$y_1 = \mathcal{A}(x_1) + e_1 \quad \text{and} \quad y_2 = \mathcal{A}(x_2) + e_2$$

- Forward operators: 2D ray transform
- Geometries: Parallel beam, 183 lines/angle, 30 angles
- Noise: 0.1% additive Gaussian
- Images: 128×128 pixel
- Training data: Triplets $\left((x_i^1, x_i^2), (y_i^1, y_i^2), z_i\right)$ where $z_i := x_i^1 - x_i^2$.

Task based reconstruction

Joint reconstruction and image comparison

- Learned reconstruction $\mathcal{R}_{\Theta_1} \colon Y \times Y \to X \times X$
  - Learned primal dual in each channel
  - Loss: d_X is sum of channel-wise L^2-distances on X.

J. Adler and O. Öktem, *Learned Primal-Dual Reconstruction*, arXiv:1707.06474. Submitted to IEEE Transactions in Medical Imaging, 2017.

Task based reconstruction

Joint reconstruction and image comparison

- Learned reconstruction $\mathcal{R}_{\Theta_1} \colon Y \times Y \to X \times X$
  - Learned primal dual in each channel
  - Loss: d_X is sum of channel-wise L^2-distances on X.
- Task operator $\mathcal{T}_{\Theta_2} \colon X \times X \to Z$
  - Z real-valued functions (difference images) defined on Ω.
  - $\mathcal{T}_{\Theta_2}(x_1, x_2) =$ U-net that maps pair of images in X to an image in Z (no probabilities).
  - Loss: d_Z is L^2-distance

Task based reconstruction

Joint reconstruction and image comparison

- Learned reconstruction $\mathcal{R}_{\Theta_1} \colon Y \times Y \to X \times X$
  - Learned primal dual in each channel
  - Loss: d_X is sum of channel-wise L^2-distances on X.
- Task operator $\mathcal{T}_{\Theta_2} \colon X \times X \to Z$
  - Z real-valued functions (difference images) defined on Ω.
  - $\mathcal{T}_{\Theta_2}(x_1, x_2) =$ U-net that maps pair of images in X to an image in Z (no probabilities).
  - Loss: d_Z is L^2-distance
- Task based reconstruction $\mathcal{T}_{\Theta_2} \circ \mathcal{R}_{\Theta_1} \colon Y \times Y \to Z$
  - No pre-training
  - Joint loss: $L(\Theta_1, \Theta_2) = C\, L_X(\Theta_1) + L_Z(\Theta_1, \Theta_2)$ with

$$L_Z(\Theta_1, \Theta_2) := \mathbb{E}\Big[\big\|\mathcal{T}_{\Theta_2}\big(\mathcal{R}_{\Theta_1}(\mathbf{Y}_1, \mathbf{Y}_2)\big) - (\mathbf{X}_1 - \mathbf{X}_2)\big\|_2\Big]$$

Phantom(s)

Joint rec. + difference

Task based reconstruction

Joint reconstruction and classification of MNIST

Task based reconstruction problem: Recover probabilities that MNIST image is a '0','1', ..., '9' from its sinogram (tomographic data)

$$y = \mathcal{A}(x) + e$$

- Forward operator: Exponential of 2D ray transform (non-linear)
- Geometry: Parallel beam, 25 lines/angle, 5 angles
- Noise: Poisson noise corresponding to 60 photons measurement, i.e., $\mathbf{Y} \sim \frac{1}{60}\,\text{Poisson}\big(60 \cdot \mathcal{A}(x_{\text{true}})\big)$
- Image: 28×28 pixel
- Training data: Triplets (x_i, y_i, z_i) where $z_i \in \{\text{'0'}, \text{'1'}, \ldots, \text{'9'}\}$.

Task based reconstruction

Joint reconstruction and classification of MNIST

- Learned reconstruction $\mathcal{R}_{\Theta_1} : Y \to X$
  - Learned gradient descent scheme
  - Loss: d_X is L^2-distance on X

J. Adler and O. Öktem, *Solving ill-posed inverse problems using iterative deep neural networks*, arXiv:1704.04058. Accepted for publication in Inverse Problems, 2017.

Task based reconstruction

Joint reconstruction and classification of MNIST

- Learned reconstruction $\mathcal{R}_{\Theta_1} \colon Y \to X$
  - Learned gradient descent scheme
  - Loss: d_X is L^2-distance on X
- Learned task operator $\mathcal{T}_{\Theta_2} \colon X \to Z$ (classifier)
  - $Z =$ probability distributions over $\{\text{'0'}, \text{'1'}, \ldots, \text{'9'}\}$
  - Given by 'off the shelf' convolutional neural net classifier with 3 convolutional layers, each followed by 2×2 max pooling for segmentation
  - Loss: d_Z is cross entropy

Task based reconstruction

Joint reconstruction and classification of MNIST

- Learned reconstruction $\mathcal{R}_{\Theta_1} \colon Y \to X$
  - Learned gradient descent scheme
  - Loss: d_X is L^2-distance on X
- Learned task operator $\mathcal{T}_{\Theta_2} \colon X \to Z$ (classifier)
  - $Z =$ probability distributions over $\{\text{'0'}, \text{'1'}, \ldots, \text{'9'}\}$
  - Given by 'off the shelf' convolutional neural net classifier with 3 convolutional layers, each followed by 2×2 max pooling for segmentation
  - Loss: d_Z is cross entropy
- Task based reconstruction $\mathcal{T}_{\Theta_2} \circ \mathcal{R}_{\Theta_1} \colon Y \to Z$
  - Both networks pre-trained
    - Learned reconstruction pre-trained
    - Task operator (classifier) pre-trained until 97.5% accuracy
  - Joint loss: $L(\Theta_1, \Theta_2) = L_Z(\Theta_2)$.

Task based reconstruction

Joint reconstruction and classification

Method	Classification accuracy
Sequential pipeline	93.35%
Joint pipeline	96.60%
Classification on clean images	97.50%

- Sequential pipeline: Reconstruction followed by classification
  - Reconstruction pre-trained with L^2-loss
  - Classifier pre-trained with cross-entropy loss
- Joint pipeline: Reconstruction and classification performed jointly
  - Reconstruction pre-trained with L^2-loss
  - Classifier pre-trained with cross-entropy loss
  - Reconstruction and classifier jointly trained with cross-entropy loss.

FBP

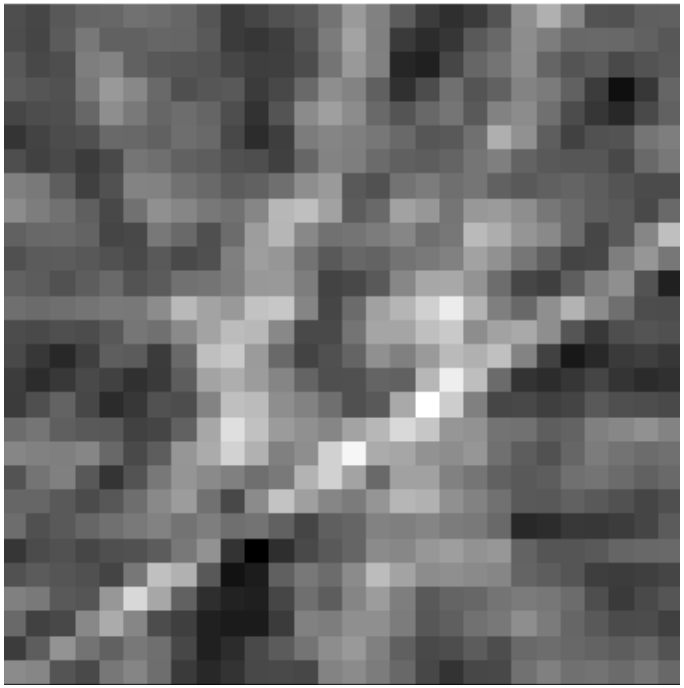

FBP

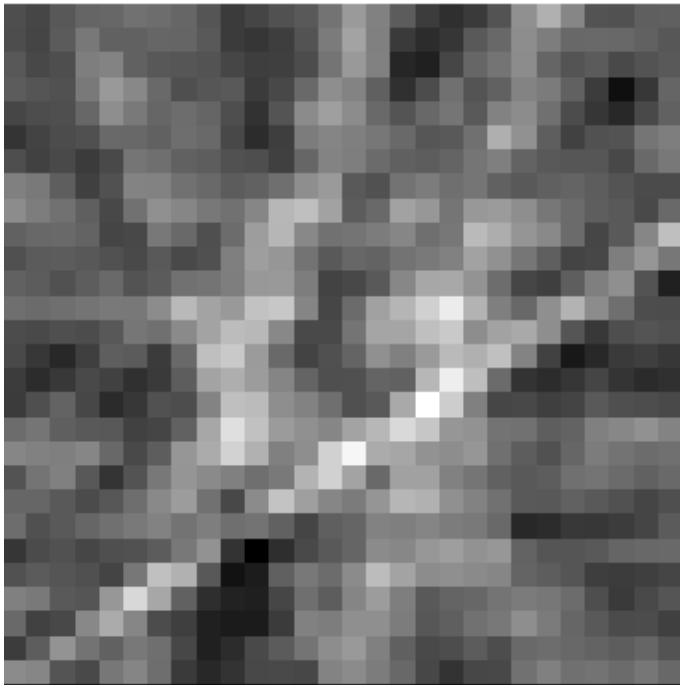

Class	Prob	Class	Prob
‘0’	0.00%	‘5’	0.01%
‘1’	0.00%	‘6’	0.00%
‘2’	0.00%	‘7’	0.00%
‘3’	0.00%	‘8’	99.99%
‘4’	0.00%	‘9’	0.00%

FBP

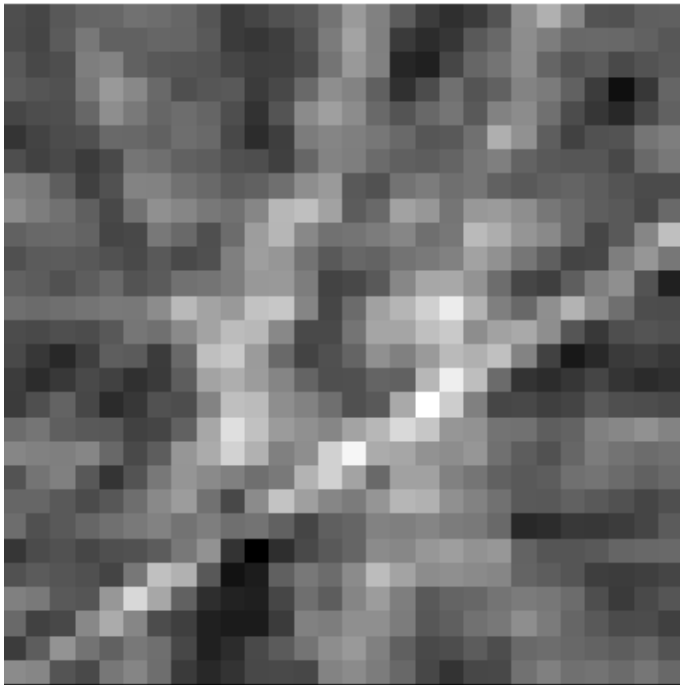

Class	Prob	Class	Prob
‘0’	0.00%	‘5’	0.01%
‘1’	0.00%	‘6’	0.00%
‘2’	0.00%	‘7’	0.00%
‘3’	0.00%	‘8’	99.99%
‘4’	0.00%	‘9’	0.00%

Task reco.: image

Class	Prob	Class	Prob
‘0’	0.00%	‘5’	0.00%
‘1’	0.00%	‘6’	0.00%
‘2’	0.00%	‘7’	0.00%
‘3’	0.00%	‘8’	0.00%
‘4’	99.70%	‘9’	0.30%

FBP

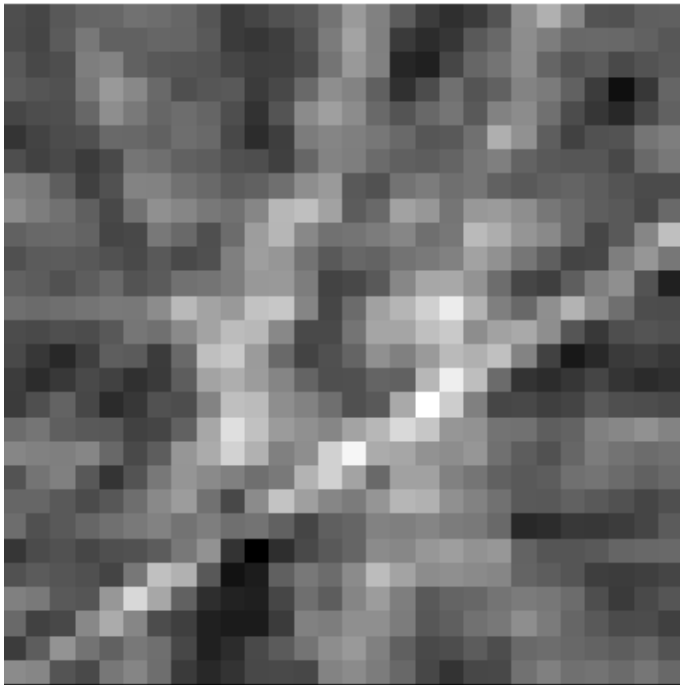

Class	Prob	Class	Prob
‘0’	0.00%	‘5’	0.01%
‘1’	0.00%	‘6’	0.00%
‘2’	0.00%	‘7’	0.00%
‘3’	0.00%	‘8’	99.99%
‘4’	0.00%	‘9’	0.00%

Task reco.: image

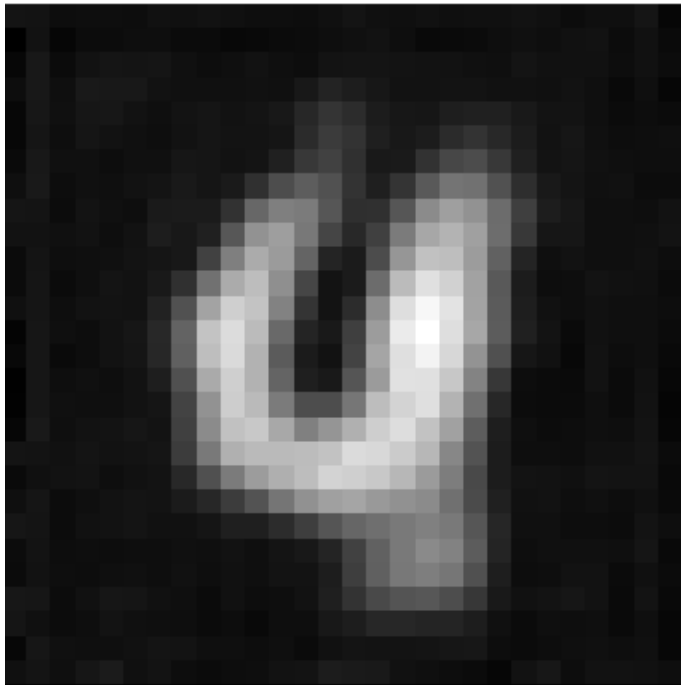

Class	Prob	Class	Prob
‘0’	0.00%	‘5’	0.00%
‘1’	0.00%	‘6’	0.00%
‘2’	0.00%	‘7’	0.00%
‘3’	0.00%	‘8’	0.00%
‘4’	99.70%	‘9’	0.30%

Phantom

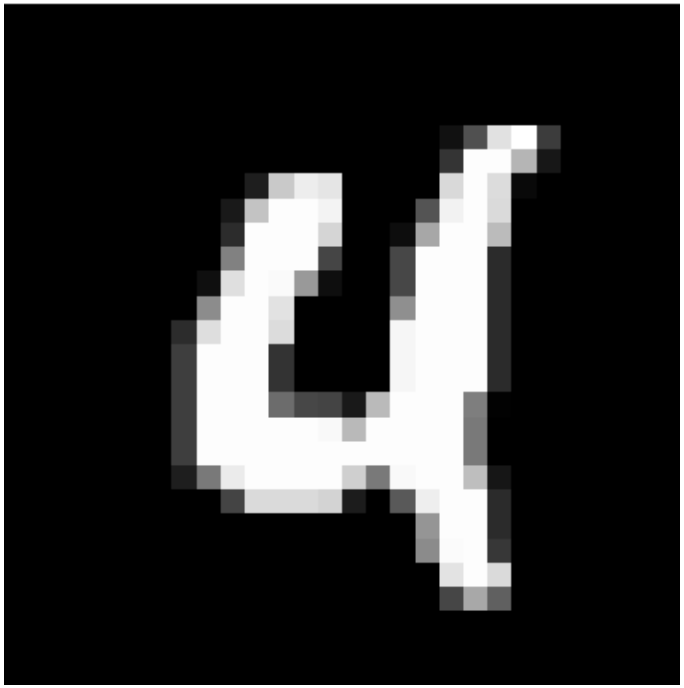

Task reco.: image

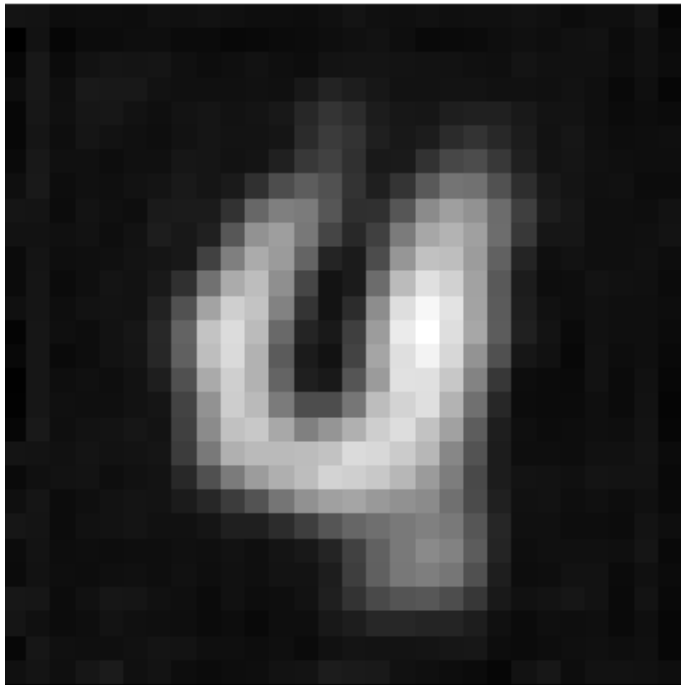

True class: ‘4’

Class	Prob	Class	Prob
‘0’	0.00%	‘5’	0.00%
‘1’	0.00%	‘6’	0.00%
‘2’	0.00%	‘7’	0.00%
‘3’	0.00%	‘8’	0.00%
‘4’	99.70%	‘9’	0.30%

FBP

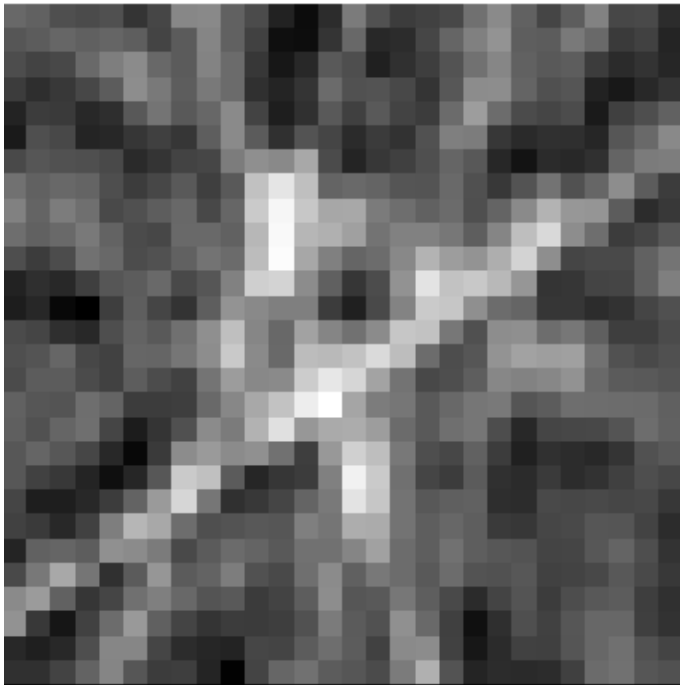

FBP

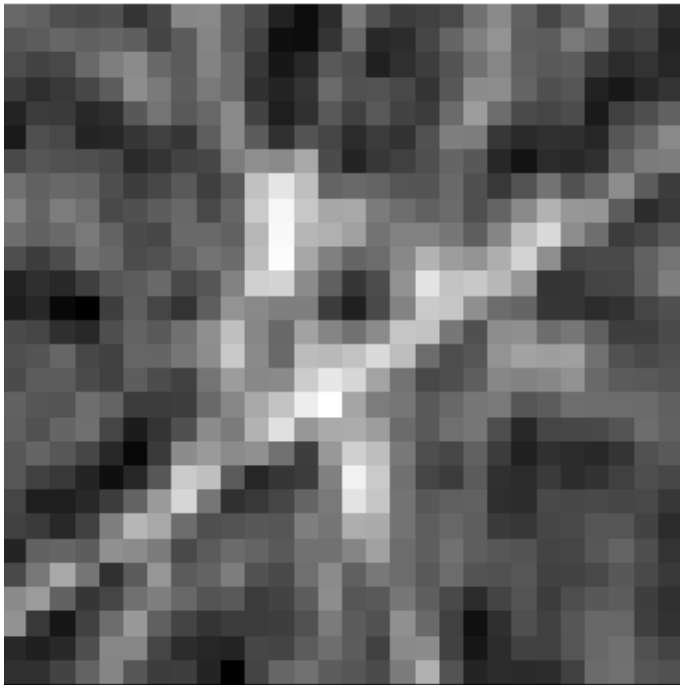

Class	Prob	Class	Prob
‘0’	0.00%	‘5’	0.01%
‘1’	0.00%	‘6’	0.00%
‘2’	0.00%	‘7’	31.07%
‘3’	0.00%	‘8’	0.00%
‘4’	68.93%	‘9’	0.00%

FBP

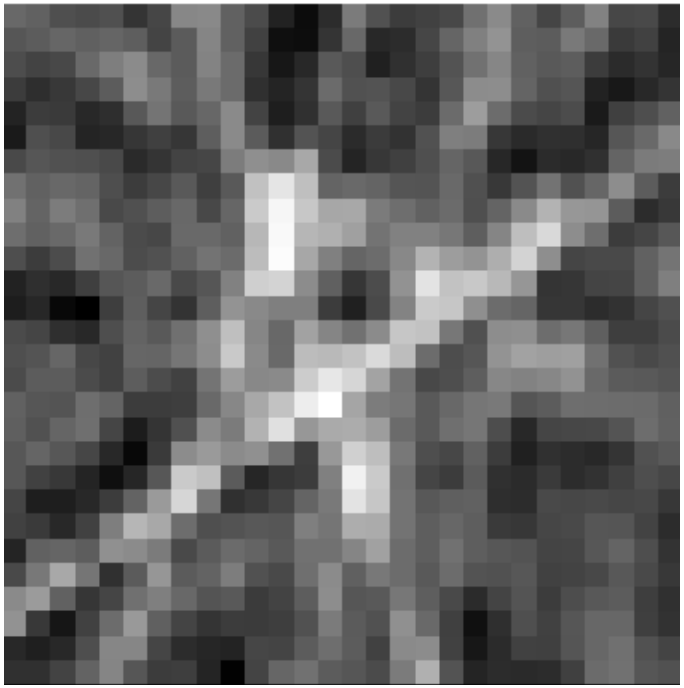

Class	Prob	Class	Prob
‘0’	0.00%	‘5’	0.01%
‘1’	0.00%	‘6’	0.00%
‘2’	0.00%	‘7’	31.07%
‘3’	0.00%	‘8’	0.00%
‘4’	68.93%	‘9’	0.00%

Task reco.: image

Class	Prob	Class	Prob
‘0’	0.00%	‘5’	0.00%
‘1’	0.00%	‘6’	0.00%
‘2’	0.00%	‘7’	0.00%
‘3’	0.00%	‘8’	0.00%
‘4’	99.97%	‘9’	0.03%

FBP

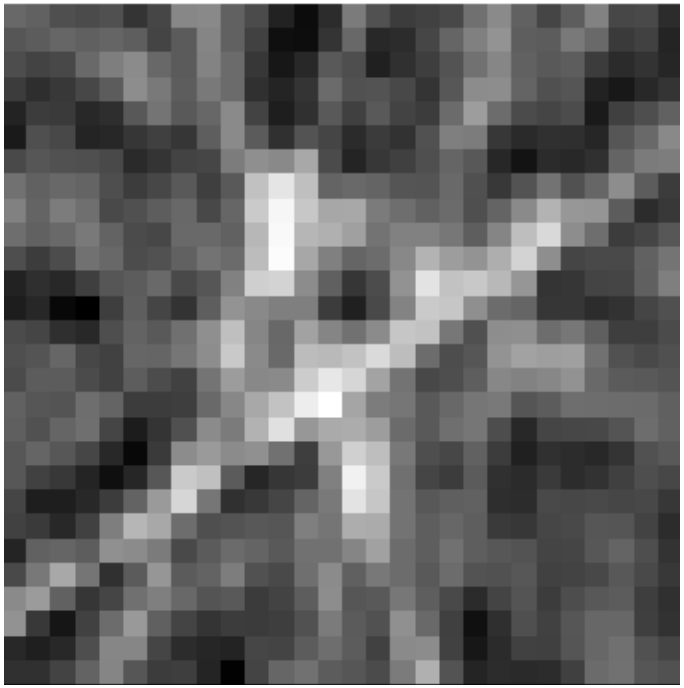

Class	Prob	Class	Prob
‘0’	0.00%	‘5’	0.01%
‘1’	0.00%	‘6’	0.00%
‘2’	0.00%	‘7’	31.07%
‘3’	0.00%	‘8’	0.00%
‘4’	68.93%	‘9’	0.00%

Task reco.: image

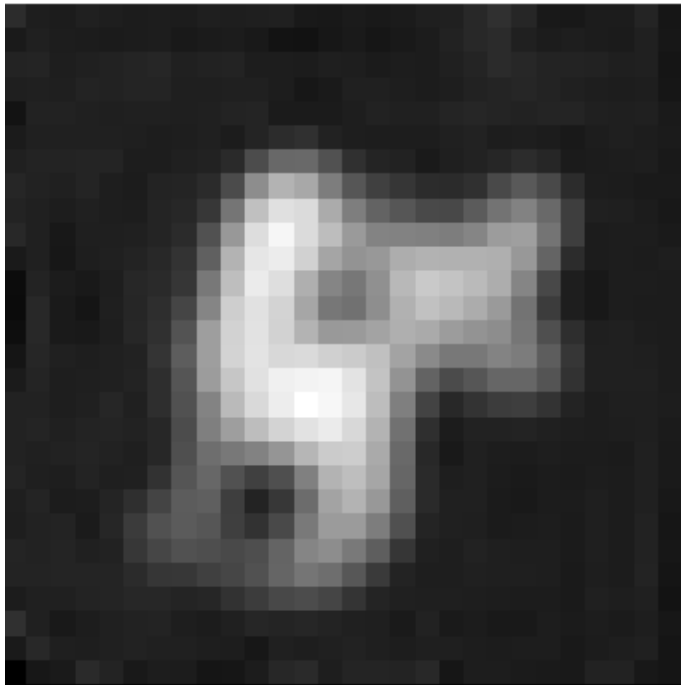

Class	Prob	Class	Prob
‘0’	0.00%	‘5’	0.00%
‘1’	0.00%	‘6’	0.00%
‘2’	0.00%	‘7’	0.00%
‘3’	0.00%	‘8’	0.00%
‘4’	99.97%	‘9’	0.03%

Phantom

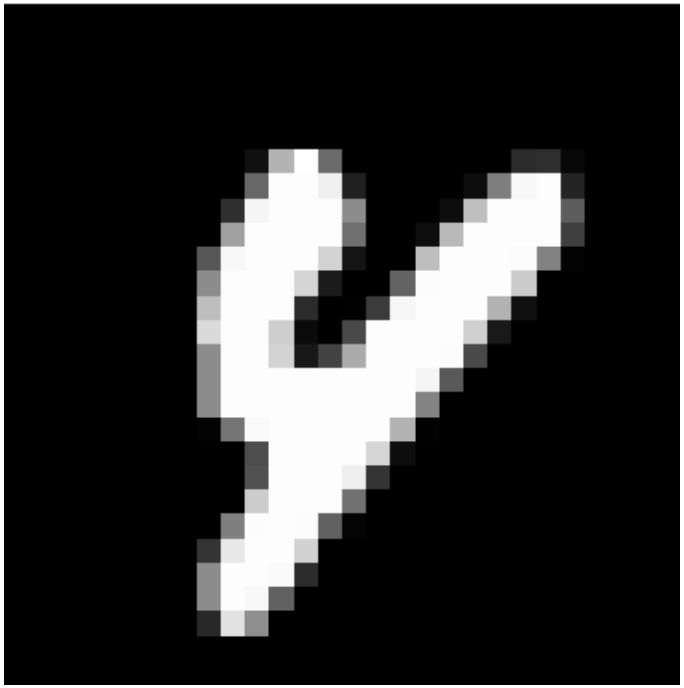

Task reco.: image

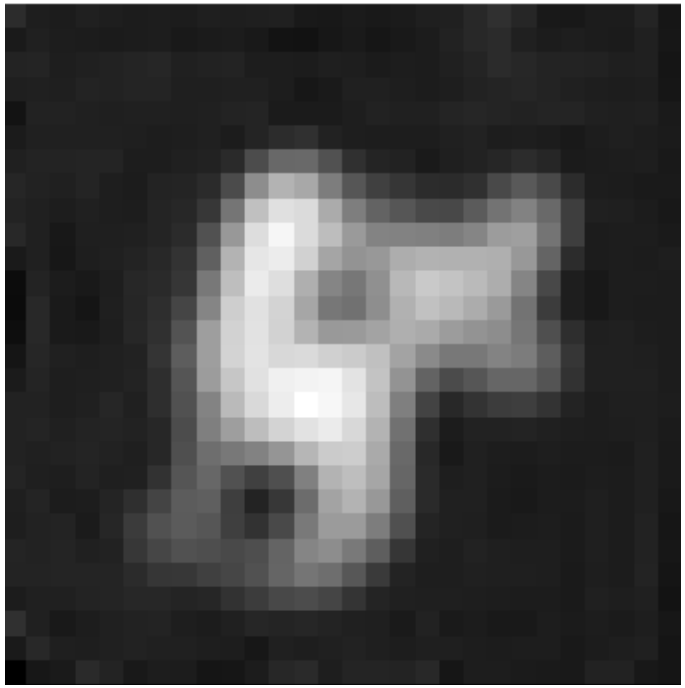

True class: ‘4’

Class	Prob	Class	Prob
‘0’	0.00%	‘5’	0.00%
‘1’	0.00%	‘6’	0.00%
‘2’	0.00%	‘7’	0.00%
‘3’	0.00%	‘8’	0.00%
‘4’	99.97%	‘9’	0.03%

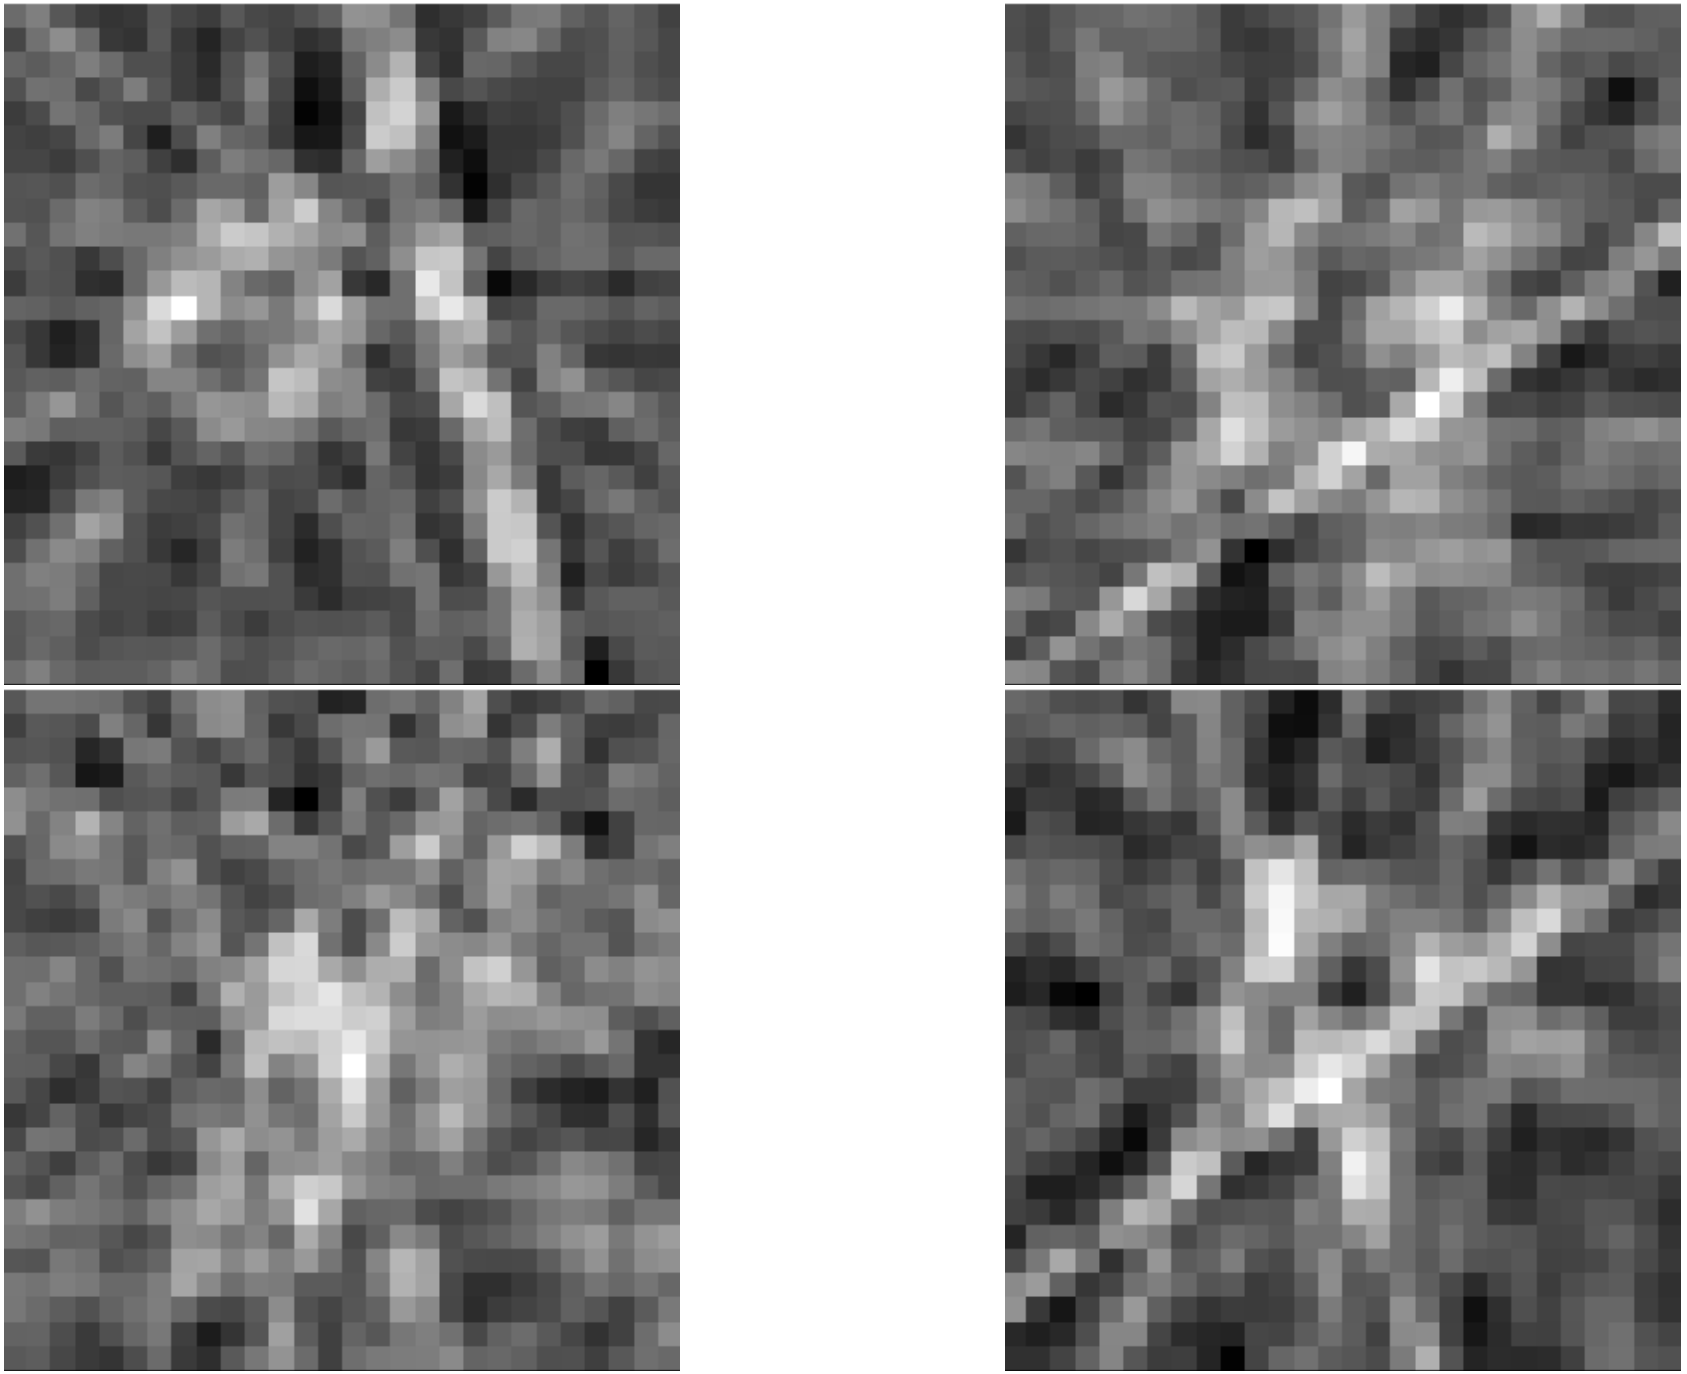

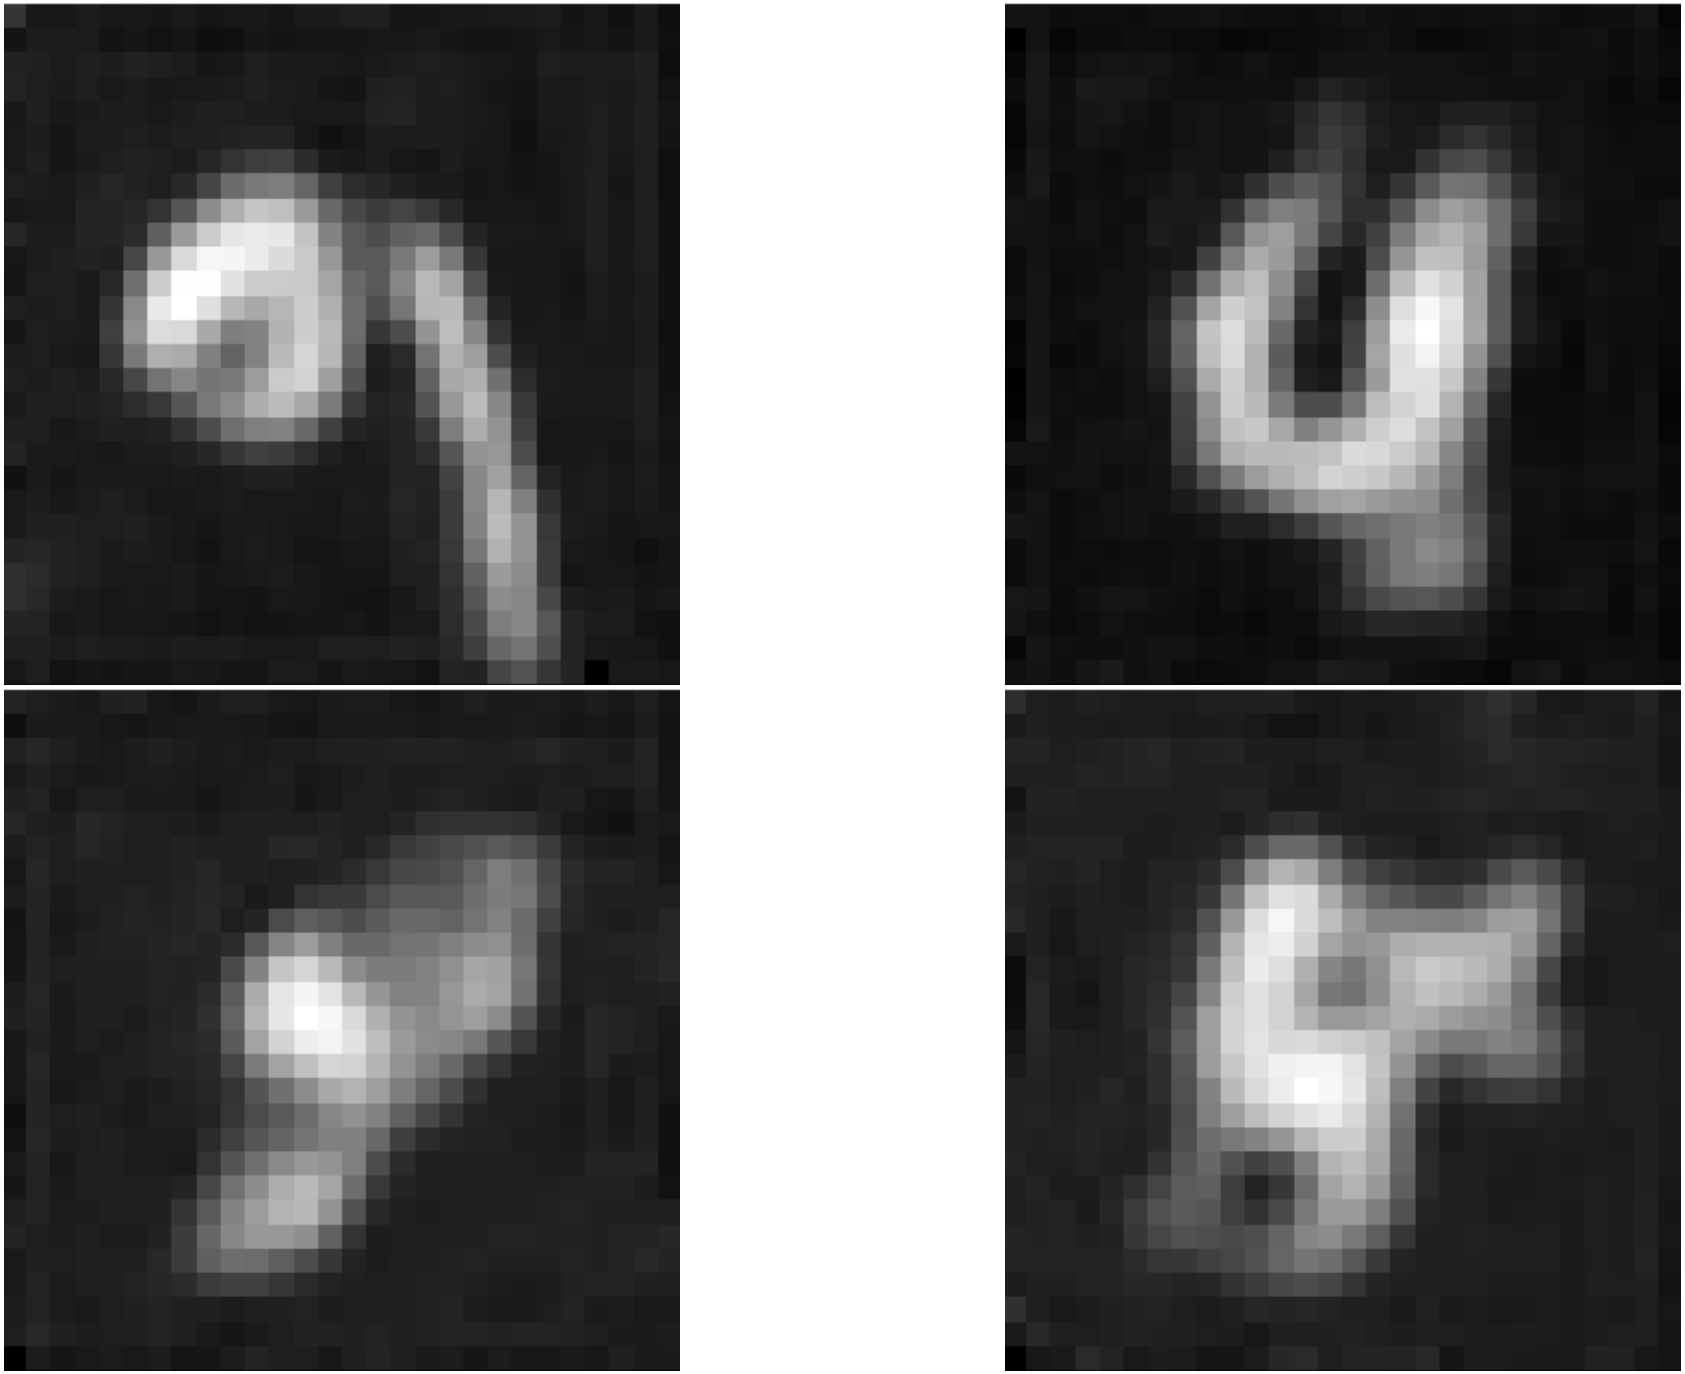

Thank you for your attention!