Paired autoencoders for inverse problems

Emma Hart, Julianne Chung, Matthias Chung

Georgia Tech SIAM Seminar

December 6, 2024

$\mathbf{b} = A(\mathbf{x}) + \boldsymbol{\epsilon}$

- Noisy observations $\mathbf{b} \in \mathbb{R}^m$
- Known forward process $A : \mathbb{R}^n \to \mathbb{R}^m$
- Parameters of interest $\mathbf{x} \in \mathbb{R}^n$
- Noise $\boldsymbol{\epsilon} \in \mathbb{R}^m$

Challenges

- Large scale
- Ill-posedness
- Uncertainty quantification

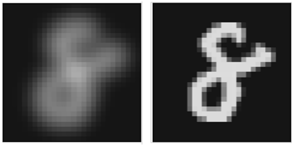

Deblurring Problem

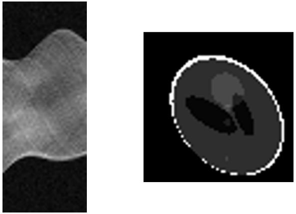

Tomography Problem

Previous Works

- Full inversion, surrogate modeling [Kulkarni et al. 2016]
- Regularization [Afkham et al. 2021; Li et al. 2020]
- Uncertainty quantification [Goh et al. 2019; Lan et al. 2022]

[Arridge et al. 2019], [Bai et al. 2020], [Lucas et al. 2018]

Encoder Decoder Networks

Supervised learning technique, popular network architecture in a variety of machine learning tasks

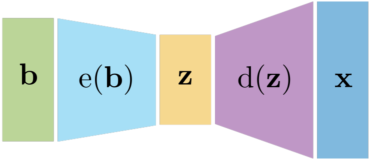

- $\mathbf{b} \in \mathbb{R}^m$
- $\mathbf{z} \in \mathbb{R}^r$ with $0 < r < min(m, n)$
- $\mathbf{x} \in \mathbb{R}^n$

- $e : \mathbb{R}^n \to \mathbb{R}^r$
- $d : \mathbb{R}^r \to \mathbb{R}^m$

Self-supervised learning technique, often used in dimensionality reduction and denoising applications

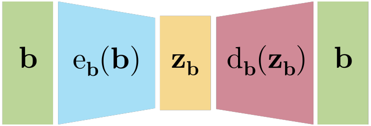

- $\mathbf{b} \in \mathbb{R}^m$
- $\mathbf{z_b} \in \mathbb{R}^r$ with $0 < r < m$
- $e_\mathbf{b} : \mathbb{R}^m \to \mathbb{R}^r$
- $d_\mathbf{b} : \mathbb{R}^r \to \mathbb{R}^m$

Paired Autoencoders for Inference and Regularization (PAIR)

Key ideas

- use **self-supervised** learning to create an autoencoder for targets, **x**
- use **self-supervised** learning to create an autoencoder for inputs, **b**
- use **supervised** learning to find a forward and/or inverse mapping between latent spaces

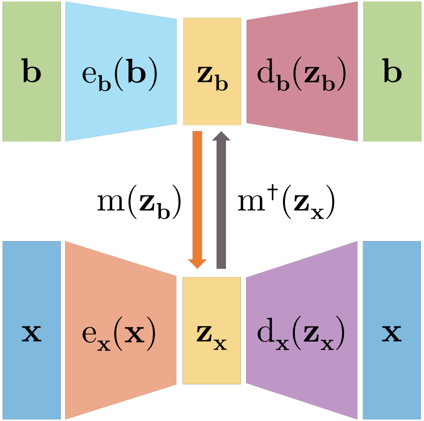

Works with similar paired structures: [Kun et al. 2015], [Feng et al. 2023]

Theory for Linear Autoencoders

Consider

$$\mathbf{z}_\mathbf{x} = e_\mathbf{x}(\mathbf{x}) = \mathbf{E}\mathbf{x} \quad \text{with} \quad \mathbf{E} \in \mathbb{R}^{r \times n}$$

and

$$\mathbf{x} \approx d_\mathbf{x}(\mathbf{z}_\mathbf{x}) = \mathbf{D}\mathbf{z}_\mathbf{x} \quad \text{with} \quad \mathbf{D} \in \mathbb{R}^{n \times r}$$

Then, define a *linear autoencoder*

$$(d_\mathbf{x} \circ e_\mathbf{x})(\mathbf{x}) = \underbrace{\mathbf{D}\mathbf{E}}_{=\mathbf{Y}}\mathbf{x} \equiv \mathbf{Y}\mathbf{x}$$

Bayes Risk Minimization

Let X be a random variable with a given probability distribution. An optimal linear autoencoder is given by

$$(\widehat{\mathsf{E}}, \widehat{\mathsf{D}}) = \underset{\mathsf{E},\mathsf{D}}{\arg\min}\ \mathbb{E}\, \|\mathsf{D}\mathsf{E}X - X\|_2^2$$

which simplifies to

$$\widehat{\mathsf{Y}} = \underset{rank(\mathsf{Y}) \leq r}{\arg\min}\ f(\mathsf{Y}) = \mathbb{E}\, \|\mathsf{Y}X - X\|_2^2 = \mathbb{E}\, \|(\mathsf{Y} - \mathsf{I})X\|_2^2$$

Given random variable X with symmetric positive definite second moment

- $\mathbb{E}XX^\top = \mathbf{\Gamma} = \mathbf{L}\mathbf{L}^\top$

Optimization Problem

$$\min_{rank(\mathbf{Y})\leq r} f(\mathbf{Y}) = \mathbb{E}\, tr\left(X^\top(\mathbf{Y}-\mathbf{I})^\top(\mathbf{Y}-\mathbf{I})X\right) = \mathbb{E}\, tr\left((\mathbf{Y}-\mathbf{I})^\top(\mathbf{Y}-\mathbf{I})XX^\top\right)$$

$$= tr\left((\mathbf{Y}-\mathbf{I})^\top(\mathbf{Y}-\mathbf{I})\underbrace{\mathbb{E}XX^\top}_{\mathbf{L}\mathbf{L}^\top}\right) = tr\left(\mathbf{L}^\top(\mathbf{Y}-\mathbf{I})^\top(\mathbf{Y}-\mathbf{I})\mathbf{L}\right) = \|(\mathbf{Y}-\mathbf{I})\mathbf{L}\|_F^2$$

Bayes Risk Minimization

Theorem

Let matrix $\mathbf{L} \in \mathbb{R}^{n\times n}$ have full rank with SVD given by $\mathbf{L} = \mathbf{U_L \Sigma_L V_L^\top}$. Then

$$\widehat{\mathbf{Y}} = \mathbf{U}_{\mathbf{L},r}\mathbf{U}_{\mathbf{L},r}^\top,$$

where $\mathbf{U}_{\mathbf{L},r}$ contains the first r columns of orthogonal matrix $\mathbf{U_L}$, is a solution to the minimization problem

$$\min_{rank(\mathbf{Y})\leq r} \|\mathbf{YL} - \mathbf{L}\|_{\mathrm{F}}^2,$$

having a minimal $\|\mathbf{Y}\|_{\mathrm{F}}$. This solution is unique if and only if either $r \geq n$ or $1 \leq r < n$ and $\sigma_r(\mathbf{L}) > \sigma_{r+1}(\mathbf{L})$. [Friedland and Torokhti 2007], [Chung and Chung 2017]

Bayes Risk Minimization Summary

- The low rank solution $\widehat{\mathbf{Y}} = \mathbf{U}_{\mathbf{L},r}\mathbf{U}_{\mathbf{L},r}^\top$ is unique for given conditions
- The decomposition into encoder $\widehat{\mathbf{E}}$ and decoder $\widehat{\mathbf{D}}$ is *not* unique, since

$$\widehat{\mathbf{Y}} = \underbrace{\mathbf{U}_{\mathbf{L},r}\mathbf{K}}_{\widehat{\mathbf{D}}}\underbrace{\mathbf{K}^{-1}\mathbf{U}_{\mathbf{L},r}^\top}_{\widehat{\mathbf{E}}}$$

  for any invertible $r \times r$ matrix $\mathbf{K}$

Empirical Bayes Risk

- Work directly with samples of a fixed distribution
- Realizations $\mathbf{x}_1, \ldots, \mathbf{x}_N$ of random variable X stored as

$$\mathbf{X} = [\mathbf{x}_1, \ldots, \mathbf{x}_N] \in \mathbb{R}^{n \times N}$$

One optimal choice of encoder and decoder

$$\widehat{\mathbf{E}} = \mathbf{K}^{-1}\mathbf{U}_{\mathbf{X},r}^{\top} \quad \text{and} \quad \widehat{\mathbf{D}} = \mathbf{U}_{\mathbf{X},r}\mathbf{K}$$

which is a low-rank SVD approximation of $\mathbf{X}$

Linear Mapping between Latent Spaces

- *Inverse mapping:* Consider

$$\mathrm{Z_X} = \begin{bmatrix} | & & | \\ e_{\mathrm{x}}(\mathrm{x}_1) & \dots & e_{\mathrm{x}}(\mathrm{x}_N) \\ | & & | \end{bmatrix} \quad \text{and} \quad \mathrm{Z_B} = \begin{bmatrix} | & & | \\ e_{\mathrm{b}}(\mathrm{b}_1) & \dots & e_{\mathrm{b}}(\mathrm{b}_M) \\ | & & | \end{bmatrix}$$

  then, using empirical Bayes risk minimization

$$\widehat{\mathrm{M}}^\dagger = \underset{\mathrm{M}^\dagger}{\arg\min} \left\| \mathrm{M}^\dagger \mathrm{Z_B} - \mathrm{Z_X} \right\|_F^2 = \mathrm{Z_X} \mathrm{Z_B}^\dagger$$

- *Forward mapping:* Analogously,

$$\widehat{\mathrm{M}} = \mathrm{Z_B} \mathrm{Z_X}^\dagger$$

[Feng et al. 2023]

Computed Tomography Example with Shepp Logan Phantoms

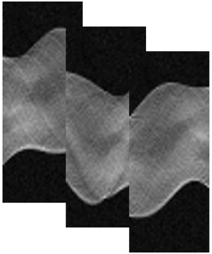

Noisy Sinogram Inputs, $\mathbf{b}$

Shepp Logan Targets, $\mathbf{x}$

Results from Linear PAIR

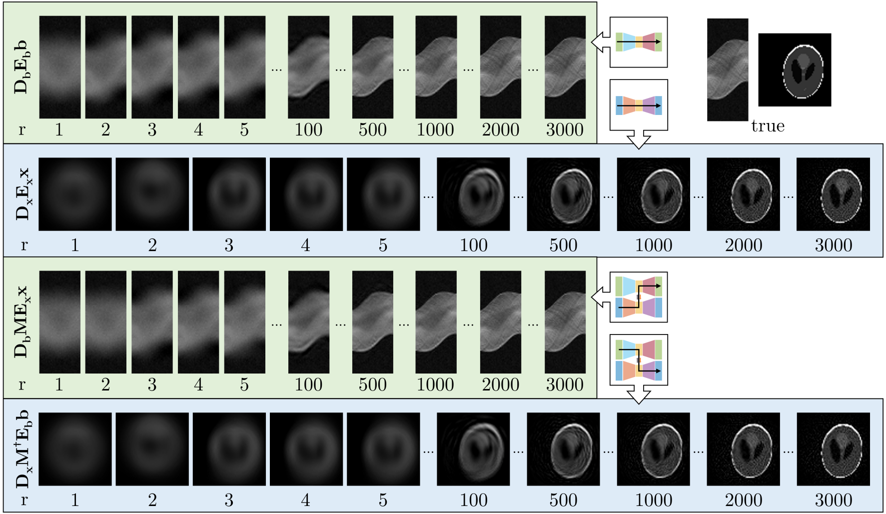

Comparison of Linear Techniques

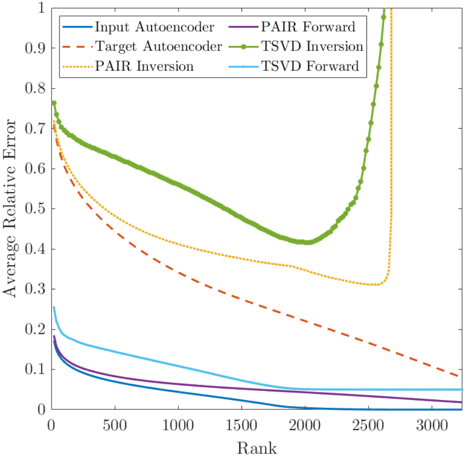

- Input Autoencoder: $\|D_b E_b b - b\|$
- Target Autoencoder: $\|D_x E_x x - x\|$
- PAIR Inversion: $\|D_x M^{\dagger} E_b b - x\|$
- PAIR Forward: $\|D_b M E_x x - b\|$
- TSVD Inversion: $\|V_{A,r} \Sigma_{A,r}^{-\top} U_{A,r}^{\top} b - x\|$
- TSVD Forward: $\|U_{A,r} \Sigma_{A,r} V_{A,r}^{\top} x - b\|$

Deblurring Example with MNIST Digits

Blurry Digit Inputs, **b**

Clear Digit Targets, **x**

Nonlinear PAIR for MNIST Deblurring

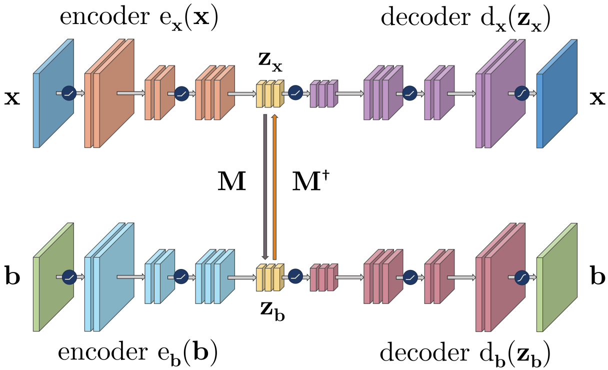

Results: MNIST Testing Data

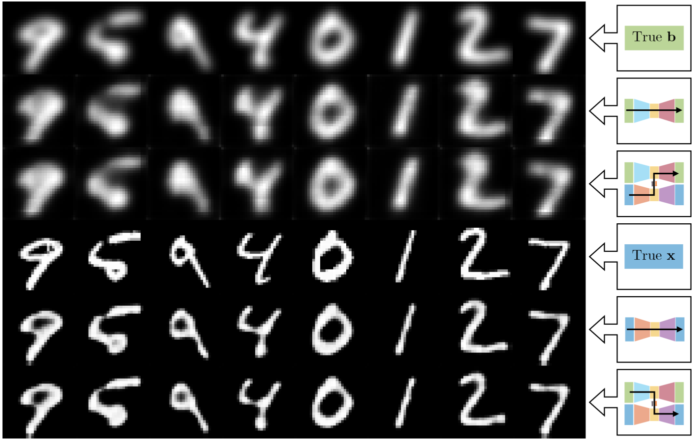

Results: Other Similar Images

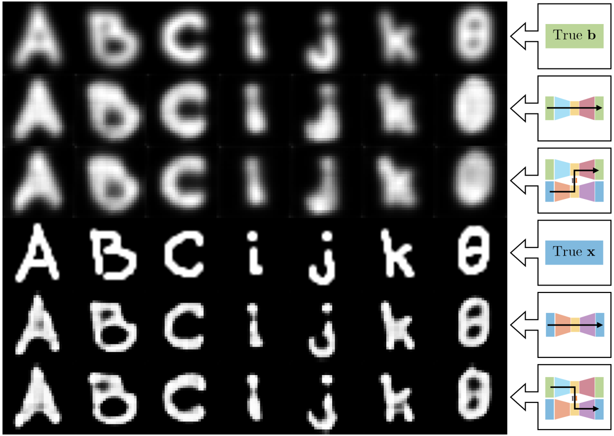

Advantages and Disadvantages

- PAIR can outperform existing methods when (# paired training images) is limited, but (# unpaired images) is abundant
- Fully supervised approaches can achieve more accurate results, but can take longer to converge

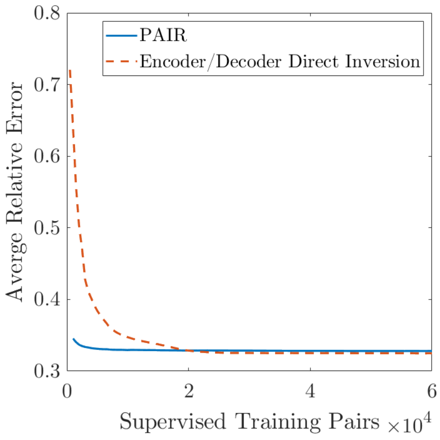

Out of Distribution Detection

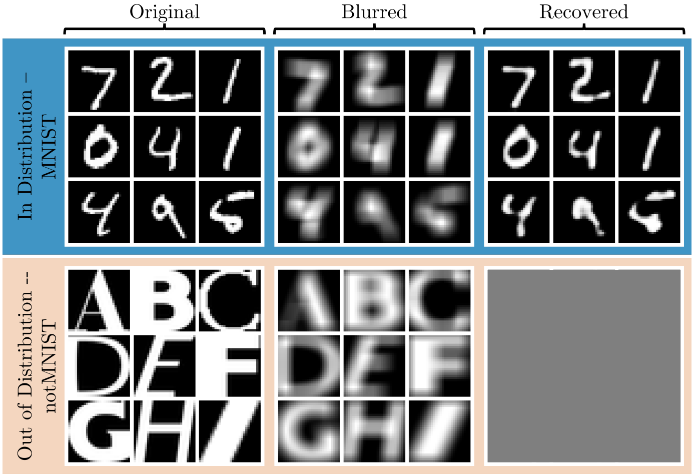

Out of Distribution Metrics

- PAIR offers some cheaply computable metrics to help predict if a new sample is “in distribution” of training data

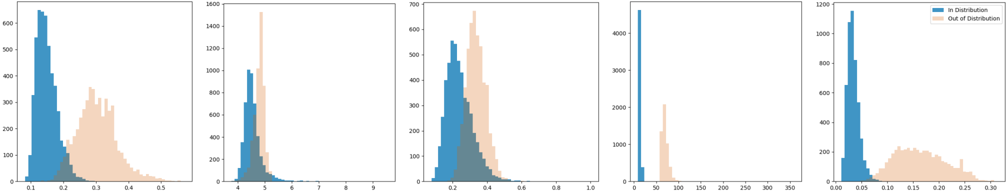

$\frac{\|(d_{\mathrm{b}} \circ e_{\mathrm{b}})(\mathrm{b})\|_2}{\|\mathrm{b}\|_2}$ $\frac{\|d_{\mathrm{b}}(\mathrm{M} e_{\mathrm{x}}(\mathrm{x}_{\mathrm{pred}})) - \mathrm{b}\|_2}{\|\mathrm{b}\|_2}$ $\frac{\|(d_{\mathrm{x}} \circ e_{\mathrm{x}})(\mathrm{x}_{\mathrm{pred}}) - \mathrm{x}_{\mathrm{pred}}\|_2}{\|\mathrm{x}_{\mathrm{pred}}\|_2}$ $\frac{\|\mathrm{M}^\dagger e_{\mathrm{b}}(\mathrm{b}) - e_{\mathrm{x}}(\mathrm{x}_{\mathrm{pred}})\|_2}{\|e_{\mathrm{x}}(\mathrm{x}_{\mathrm{pred}})\|_2}$ $\frac{\|\mathrm{M} e_{\mathrm{x}}(\mathrm{x}_{\mathrm{pred}}) - e_{\mathrm{b}}(\mathrm{b})\|_2}{\|e_{\mathrm{b}}(\mathrm{b})\|_2}$

Out of Distribution Detection

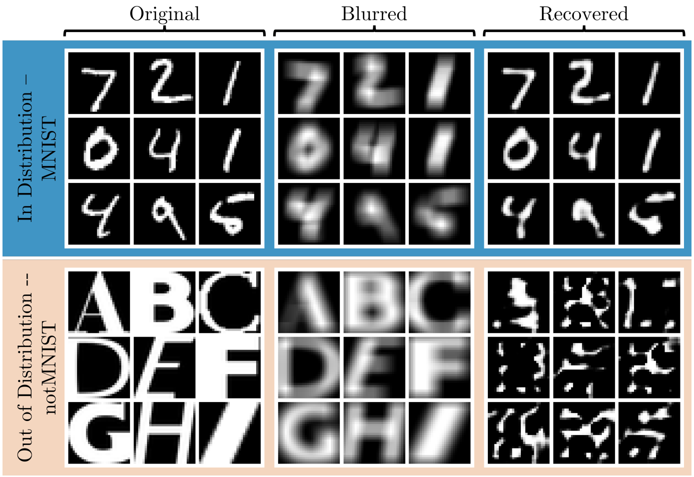

Continuing Work: Refining Solution with Latent Space Regularization

- When our out of distribution metrics are high, this may indicate we need to refine our solution
- We can still use the parametrization we found and leverage our forward model:

$$\underset{\mathbf{z} \in \mathcal{Z}_X}{\arg\min} \ \frac{1}{2} \left\| A\left(D_X(\mathbf{z})\right) - \mathbf{b} \right\|^2 + \frac{\alpha}{2} \left\| \mathbf{z} - \mathbf{z}^\star \right\|^2$$

https://arxiv.org/abs/2405.13220

Conclusions and Future Work

Conclusions

- Autoencoders can be used for dimension reduction in inverse problems
  - Self-supervised learning of inputs and targets
  - Supervised learning for mapping between latent spaces
- Theory for linear autoencoders and linear mappings
- Numerical results are promising, especially when paired data is limited

Future work

- Uncertainty quantification with variational autoencoders
- Explore new regularization/priors
- Generalization to other problems/data with exclusions

Acknowledgements

Thank you!

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Department of Energy Computational Science Graduate Fellowship under Award Number DE-SC0024386.

Experiment Design

- $60,000$ 28×28 pixel handwritten MNIST images ($50,000$ training and $10,000$ testing)
- Both convolutional neural network (CNN) autoencoders, each with 236 parameters
  - 2 layer encoder (77 parameters), 3 layer decoder (159 parameters)
  - 3×3 kernel
  - ReLU activation at each inner layer, sigmoid at output layer
  - Adam optimization
  - mean squared error loss
- Latent space with dimension $7 \times 7 \times 3$

Bibliography

[1] B. M. Afkham, J. Chung, and M. Chung. **"Learning regularization parameters of inverse problems via deep neural networks".** In: *Inverse Problems* 37.10 (2021), p. 105017 (cit. on p. 3).

[2] S. Arridge et al. **"Solving inverse problems using data-driven models".** In: *Acta Numerica* 28 (2019), pp. 1–174 (cit. on p. 3).

[3] Y. Bai et al. **"Deep learning methods for solving linear inverse problems: Research directions and paradigms".** In: *Signal Processing* 177 (2020), p. 107729. ISSN: 0165-1684. DOI: https://doi.org/10.1016/j.sigpro.2020.107729. URL: https://www.sciencedirect.com/science/article/pii/S0165168420302723 (cit. on p. 3).

[4] J. Chung and M. Chung. **"Optimal Regularized Inverse Matrices for Inverse Problems".** In: *SIAM Journal on Matrix Analysis and Applications* 38.2 (2017), pp. 458–477. DOI: 10.1137/16M1066531 (cit. on p. 10).

[5] Y. Feng et al. ***Simplifying Full Waveform Inversion via Domain-Independent Self-Supervised Learning.*** 2023. arXiv: 2305.13314 [physics.geo-ph] (cit. on pp. 6, 13).

[6] S. Friedland and A. Torokhti. **"Generalized Rank-Constrained Matrix Approximations".** In: *SIAM Journal on Matrix Analysis and Applications* 29.2 (2007), pp. 656–659. DOI: 10.1137/06065551 (cit. on p. 10).

[7] H. Goh et al. **"Solving Bayesian inverse problems via variational autoencoders".** In: *arXiv preprint arXiv:1912.04212* (2019) (cit. on p. 3).

[8] K. Kulkarni et al. **"ReconNet: Non-Iterative Reconstruction of Images From Compressively Sensed Measurements".** In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2016 (cit. on p. 3).

[9] Z. Kun et al. **"Coupled Deep Autoencoder for Single Image Super-Resolution".** In: *IEEE Transactions on Cybernetics* 47 (Dec. 2015), pp. 1–11. DOI: 10.1109/TCYB.2015.2501373 (cit. on p. 6).

[10] S. Lan, S. Li, and B. Shahbaba. **"Scaling Up Bayesian Uncertainty Quantification for Inverse Problems Using Deep Neural Networks".** In: *SIAM/ASA Journal on Uncertainty Quantification* 10.4 (2022), pp. 1684–1713. DOI: 10.1137/21M1439456 (cit. on p. 3).

[11] H. Li et al. **"NETT: Solving inverse problems with deep neural networks".** In: *Inverse Problems* 36.6 (2020), p. 065005 (cit. on p. 3).

[12] A. Lucas et al. **"Using deep neural networks for inverse problems in imaging: beyond analytical methods".** In: *IEEE Signal Processing Magazine* 35.1 (2018), pp. 20–36 (cit. on p. 3).