Paired autoencoders for inverse problems

Emma Hart, Julianne Chung, Matthias Chung

Emory University
Inverse Problems Research Group

April 24, 2024

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

- Input observations $\mathbf{b} \in \mathbb{R}^m$
- Known forward process $A : \mathbb{R}^n \to \mathbb{R}^m$
- Target reconstructions $\mathbf{x} \in \mathbb{R}^n$
- Noise $\boldsymbol{\epsilon} \in \mathbb{R}^m$

Challenges

- Large scale
- Ill-posedness
- Uncertainty quantification

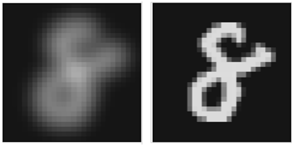

Deblurring Problem

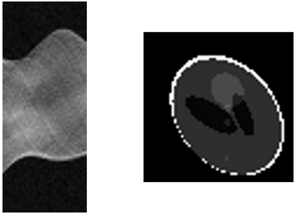

Tomography Problem

Previous Works

- Full inversion, surrogate modeling [Kulkarni et al. 2016]
- Prior learning [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

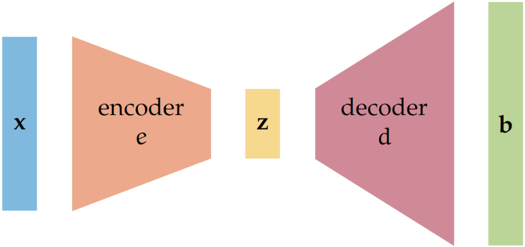

- $\mathbf{x} \in \mathbb{R}^n$
- $\mathbf{z} \in \mathbb{R}^r$ with $0 < r < min(m, n)$
- $\mathbf{b} \in \mathbb{R}^m$

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

Autoencoders

Unsupervised learning technique, often used in dimensionality reduction and denoising applications

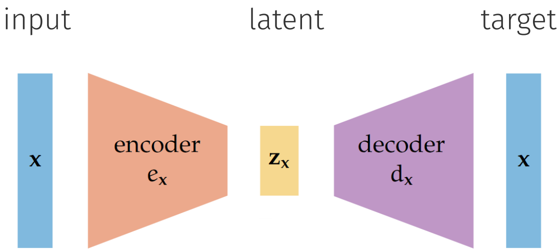

- $\mathbf{x} \in \mathbb{R}^n$
- $\mathbf{z}_{\mathbf{x}} \in \mathbb{R}^r$ with $0 < r < n$
- $e_{\mathbf{x}} : \mathbb{R}^n \to \mathbb{R}^r$
- $d_{\mathbf{x}} : \mathbb{R}^r \to \mathbb{R}^n$

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

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

which simplifies to

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

Bayes Risk Minimization

Given random variable X with

- $\mathbb{E}X = \boldsymbol{\mu}$
- $\mathbb{E}XX^\top = \mathbf{L}\mathbf{L}^\top + \boldsymbol{\mu}\boldsymbol{\mu}^\top = \mathbf{N}\mathbf{N}^\top$, with $\mathbf{N} = [\mathbf{L} \;\; \boldsymbol{\mu}]$

Optimization Problem

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

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

Bayes Risk Minimization

Theorem

Let matrix $\mathbf{N} \in \mathbb{R}^{n\times(n+1)}$ have full row rank with SVD given by $\mathbf{N} = \mathbf{U}_{\mathbf{N}}\Sigma_{\mathbf{N}}\mathbf{V}_{\mathbf{N}}^\top$. Then

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

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

$$\min_{rank(\mathbf{Y})\leq r} \|\mathbf{Y}\mathbf{N} - \mathbf{N}\|_{\mathsf{F}}^2,$$

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

Bayes Risk Minimization Summary

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

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

  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 $\mathbf{X}$

Paired Autoencoders for Inference and Regularization (PAIR)

Key ideas

- use **unsupervised** learning to create an autoencoder for targets, $\mathbf{x}$
- use **unsupervised** learning to create an autoencoder for inputs, $\mathbf{b}$
- use **supervised** learning to find a forward and/or inverse mapping between latent spaces

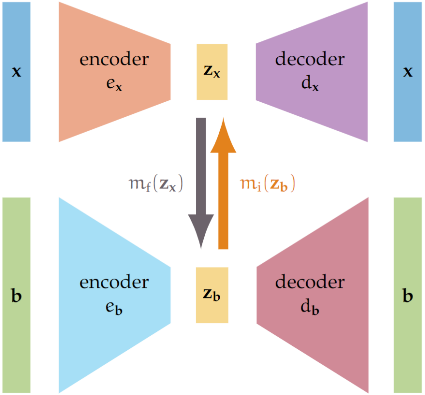

[Kun et al. 2015]

Linear Mapping between Latent Spaces

- *Inverse mapping:* Consider

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

  then, using empirical Bayes risk minimization

$$\mathbf{M}_i = \arg\min_{\mathbf{M}} \|\mathbf{M}\mathbf{Z_B} - \mathbf{Z_X}\|_{\mathrm{F}}^2 = \mathbf{Z_X}\mathbf{Z_B}^{\dagger}$$

- *Forward mapping:* Analogously,

$$\mathbf{M}_f = \mathbf{Z_B}\mathbf{Z_X}^{\dagger}$$

Computed Tomography Example with Shepp Logan Phantoms

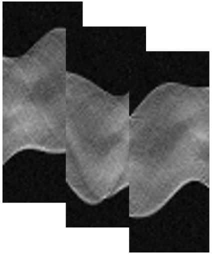

Noisy Sinogram Inputs, $\mathbf{b}$

Shepp Logan Targets, $\mathbf{x}$

Results from Linear PAIR

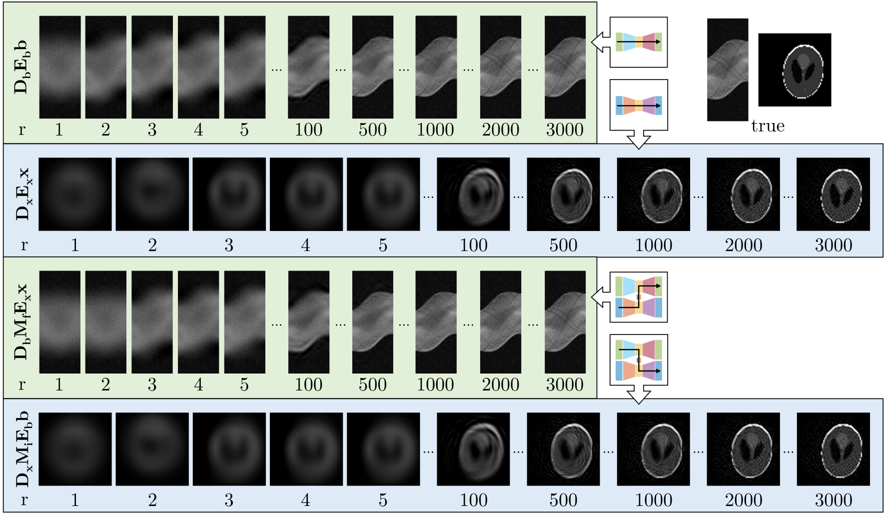

Comparison of Linear Techniques

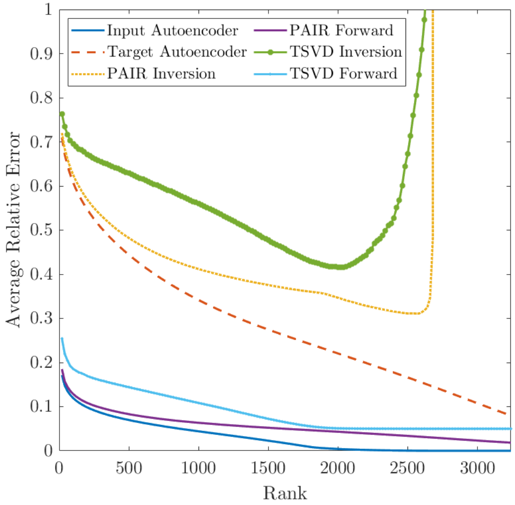

Deblurring Example with MNIST Digits

Blurry Digit Inputs, $\mathbf{b}$

Clear Digit Targets, $\mathbf{x}$

Nonlinear PAIR for MNIST Deblurring

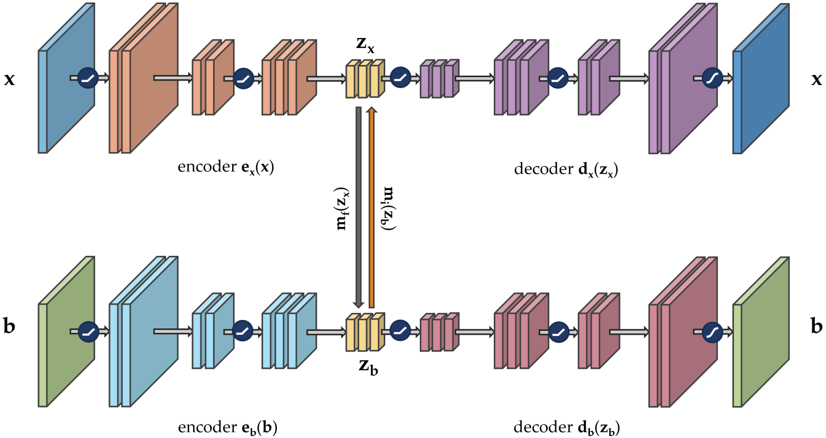

Experiment Methodologies

- $60,000$ 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 $7 \times 7 \times 3$

Results: Within Distribution

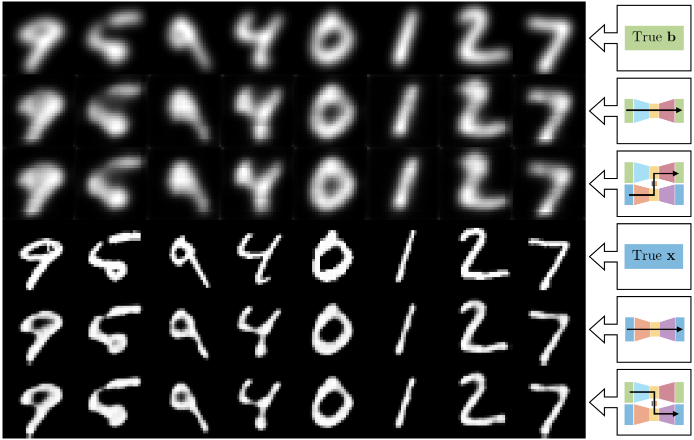

Results: Out of Distribution

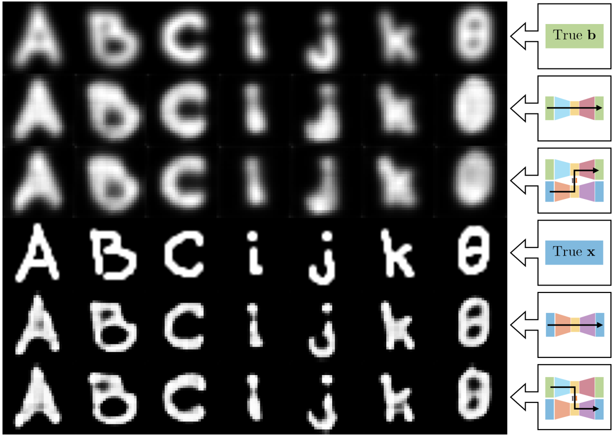

Advantages and Disadvantages

- PAIR can outperform existing methods when (# paired training images) is limited, but (# unpaired images) is abundant
- Otherwise, fully supervised approaches can achieve more accurate results, but are expensive

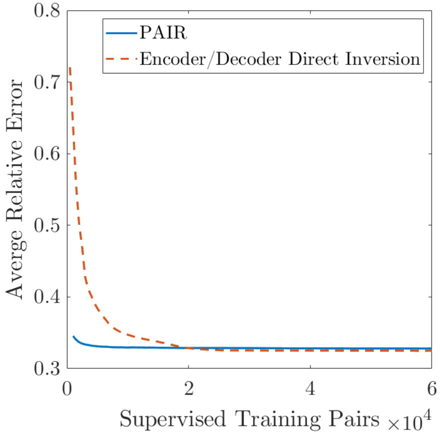

Conclusions and Future Work

Conclusions

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

Future work

- Approximating the adjoint
- Define new regularization/priors: approximate mean or prior covariance
- Surrogate Models: create a reduced model with a snapshot matrix for forward propagation of dynamical systems

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. 9).

[5] 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. 9).

References iii

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

[7] 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).

[8] 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. 12).

[9] 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).

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

[11] 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).