Deep Learning for Signal Processing and Medical Applications

Carlos Fernandez-Granda
www.cims.nyu.edu/~cfgranda

Acknowledgements

The projects described in this presentation are funded by NIH and NSF

Blind denoising of natural images

Bias-free CNNs

Wiener filtering

CNNs learn adaptive filters

CNNs learn unions of subspaces

Quantitative magnetic-resonance imaging

Early diagnostics of Alzheimer's disease

Quantitative rehabilitation of stroke patients

Data-driven estimation of sinusoid frequencies

Blind denoising of natural images

- Bias-free CNNs
- Wiener filtering
- CNNs learn adaptive filters
- CNNs learn unions of subspaces

Quantitative magnetic-resonance imaging

Early diagnostics of Alzheimer's disease

Quantitative rehabilitation of stroke patients

Data-driven estimation of sinusoid frequencies

Acknowledgements

Joint work with Zahra Kadkhodaie, Sreyas Mohan, and Eero Simoncelli

Image denoising

Goal: Estimate image from noisy data

Popular model: Additive Gaussian noise

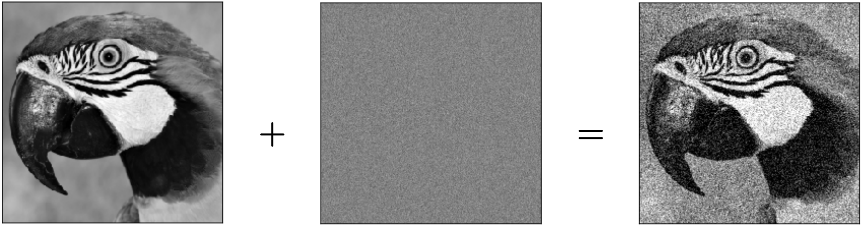

Blind denoising: Noise level is unknown

Image denoising

Goal: Estimate image from noisy data

Popular model: Additive Gaussian noise

$+$

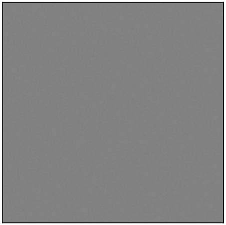

$=$

Blind denoising: Noise level is unknown

Image denoising

Goal: Estimate image from noisy data

Popular model: Additive Gaussian noise

$+$

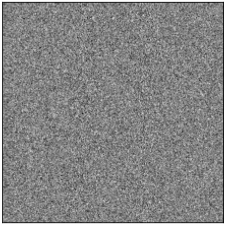

$=$

Blind denoising: Noise level is unknown

Deep learning for blind image denoising

- Gather dataset of natural images
- Add noise from a range of noise levels
- Train convolutional neural network (CNN) to estimate clean image minimizing mean squared error
- Works very well for additive Gaussian noise (state of the art)

Generalization across noise levels

What if we test on noise level not seen during training?

Training data
(low noise)

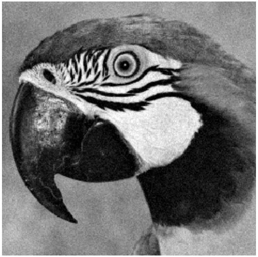

Test image
(high noise)

Generalization across noise levels

What if we test on noise level not seen during training?

Training data
(low noise)

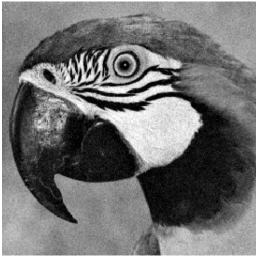

Test image
(high noise)

CNN

Blind denoising of natural images

- Bias-free CNNs
- Wiener filtering
- CNNs learn adaptive filters
- CNNs learn unions of subspaces

Quantitative magnetic-resonance imaging

Early diagnostics of Alzheimer's disease

Quantitative rehabilitation of stroke patients

Data-driven estimation of sinusoid frequencies

First-order Taylor expansion

Let f be the function learned by a CNN trained for denoising

The first-order Taylor expansion for a fixed input y is exact

$$\begin{aligned}\hat{x} = f(y) &= W_L R(\dots W_2 R(W_1 y + b_1) + b_2 \dots) + b_L \\ &= A_y y + b_y\end{aligned}$$

W_1, W_2, …, W_L are weight matrices
b_1, b_2, …, b_L are bias vectors

Residual and net bias

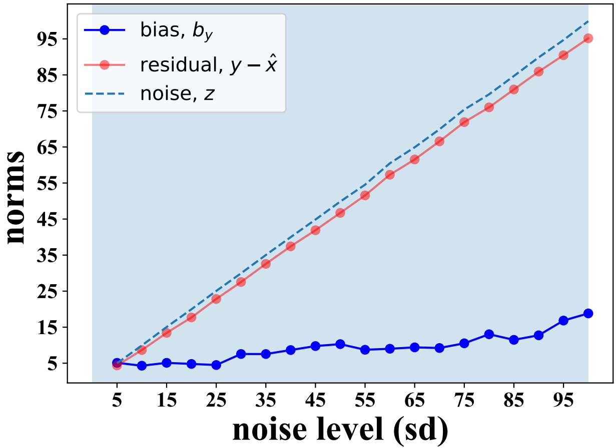

Residual and net bias

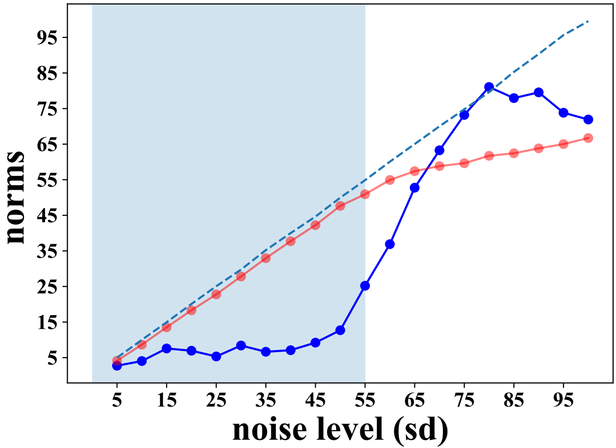

Residual and net bias

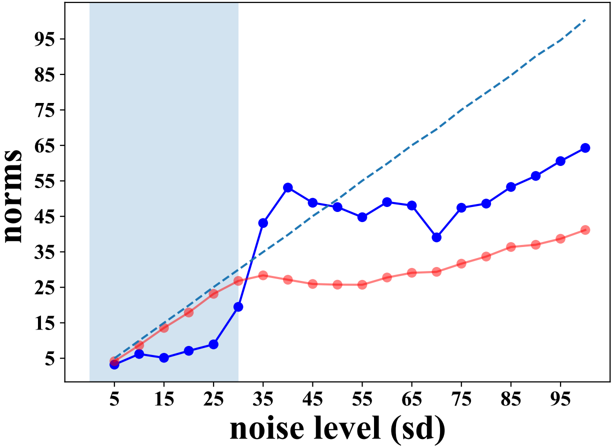

Bias-free networks

Within training range, learned net bias is small

Out of the range, it explodes, coinciding with dramatic performance loss

Net bias seems to overfit trained noise levels

Bias-free networks

Within training range, learned net bias is small

Out of the range, it explodes, coinciding with dramatic performance loss

Net bias seems to overfit trained noise levels

This motivates *removing all additive constants*

$$f(y) = W_L R(\ldots W_2 R(W_1 y + b_1) + b_2 \ldots) + b_L$$

Bias-free networks

Within training range, learned net bias is small

Out of the range, it explodes, coinciding with dramatic performance loss

Net bias seems to overfit trained noise levels

This motivates *removing all additive constants*

$$f(y) = W_L R(\ldots W_2 R(W_1 y + \cancel{b_1}) + \cancel{b_2} \ldots) + \cancel{b_L}$$

It works

Training data (low noise)

Test image (high noise)

CNN

It works

Training data (low noise)

Test image (high noise)

CNN

Bias-free CNN

DnCNN [Zhang *et al* 2016] vs bias-free DnCNN

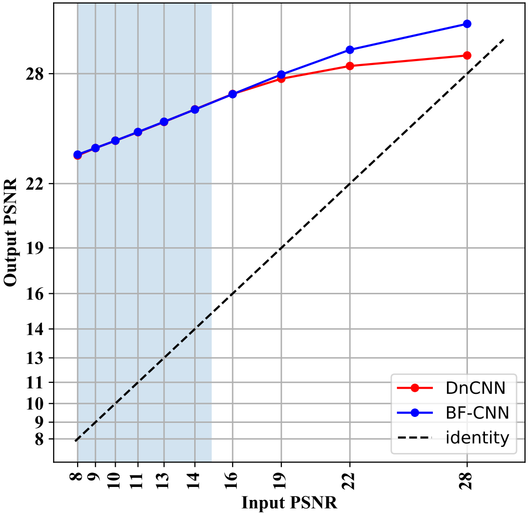

DnCNN [Zhang *et al* 2016] vs bias-free DnCNN

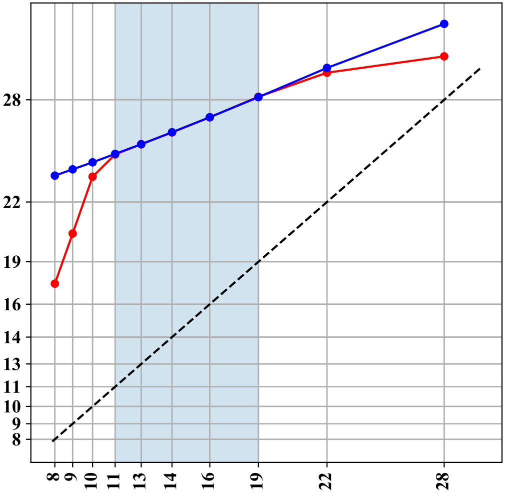

DnCNN [Zhang *et al* 2016] vs bias-free DnCNN

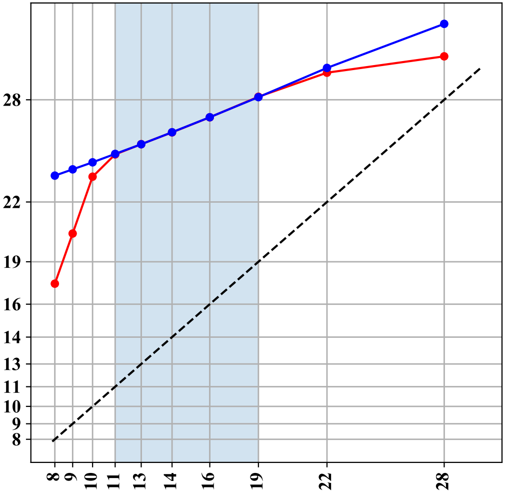

DnCNN [Zhang *et al* 2016] vs bias-free DnCNN

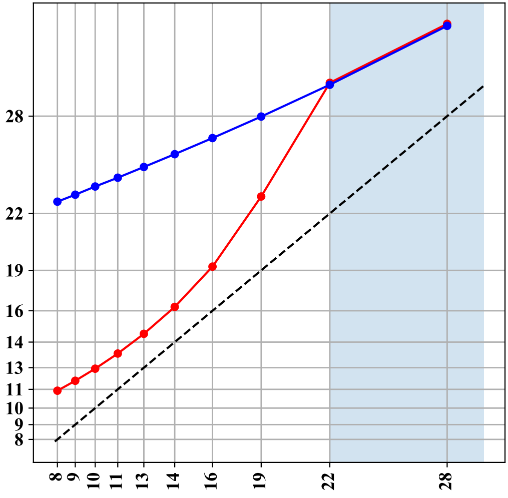

DnCNN [Zhang *et al* 2016] vs bias-free DnCNN

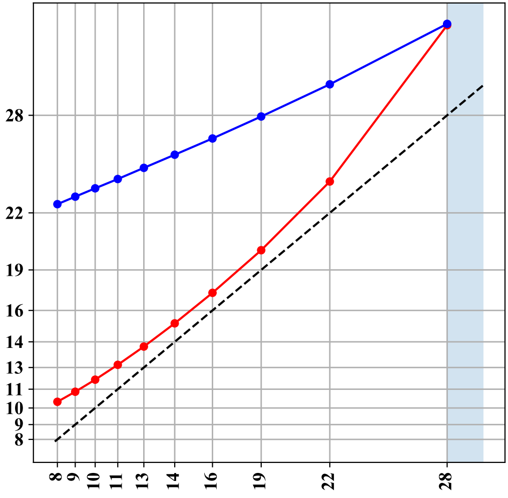

DnCNN [Zhang *et al* 2016] vs bias-free DnCNN

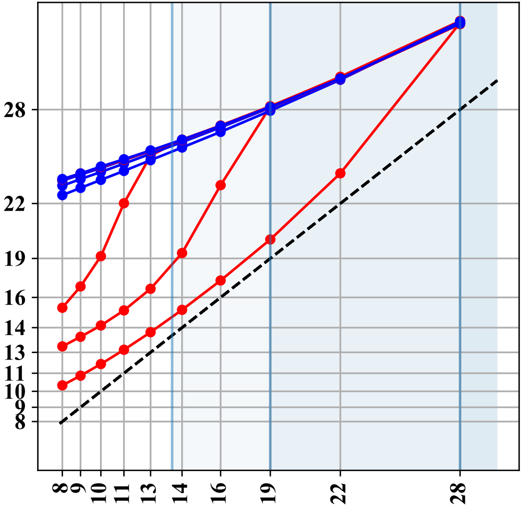

DenseNet [Huang *et al* 2017] vs bias-free DenseNet

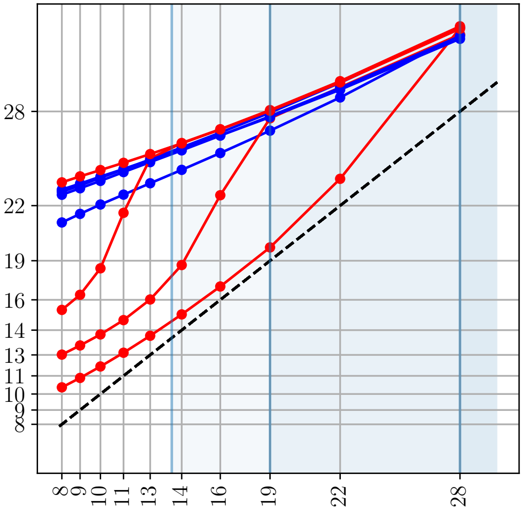

UNet [Ronneberger *et al* 2015] vs bias-free UNet

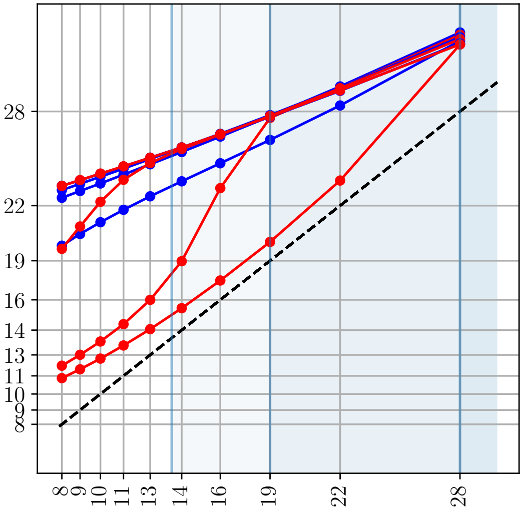

Recurrent CNN [Zhang *et al* 2018] vs bias-free recurrent CNN

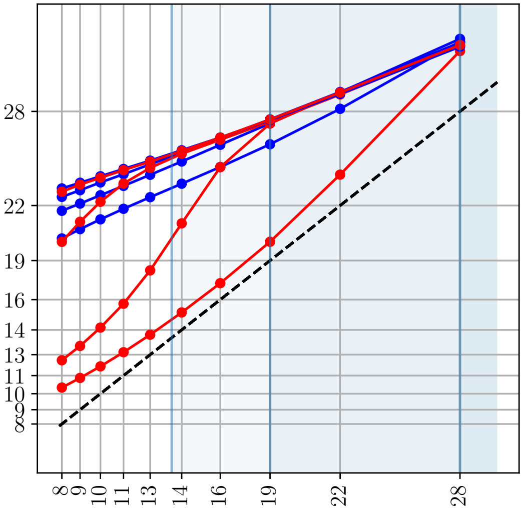

Blind denoising of natural images

- Bias-free CNNs
- Wiener filtering
- CNNs learn adaptive filters
- CNNs learn unions of subspaces

Quantitative magnetic-resonance imaging

Early diagnostics of Alzheimer's disease

Quantitative rehabilitation of stroke patients

Data-driven estimation of sinusoid frequencies

Linear estimation

Linear regression from pixels to pixels is intractable ($10^4 \times 10^4$ matrix!)

No need: covariance between pixels is translation invariant

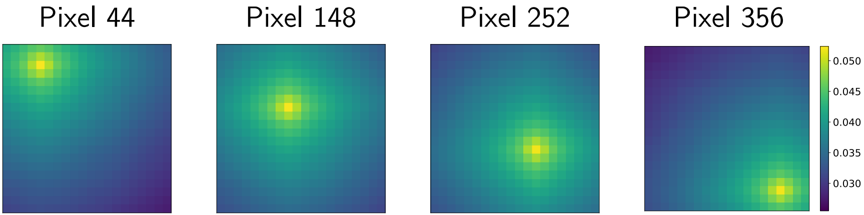

Linear estimator can be parameterized by a convolutional filter

Wiener filter [Wiener 1950]

Filter w that achieves optimal mean squared error

Random vectors: x (image), z (noise), $y := x + z$ (data)

Fourier transform is an orthogonal transformation so

$$\mathrm{E}\left(||x - w * y||_2^2 \right) = \mathrm{E}\left(||\hat{x} - \hat{w} \circ \hat{y}||_2^2 \right)$$

Wiener filter [Wiener 1950]

Filter w that achieves optimal mean squared error

Random vectors: x (image), z (noise), $y := x + z$ (data)

Fourier transform is an orthogonal transformation so

$$\begin{aligned} \mathrm{E}\left(\left|\left| x - w * y \right|\right|_2^2 \right) &= \mathrm{E}\left(\left|\left| \hat{x} - \hat{w} \circ \hat{y} \right|\right|_2^2 \right) \\ &= \sum_k \mathrm{E}\left(\left(\hat{x}_k - \hat{w}_k \hat{y}_k \right)^2 \right) \end{aligned}$$

We can estimate each Fourier coefficient separately

Wiener filter

If x and z are independent, and z is i.i.d. with variance σ^2

$$\hat{w}_k^{\text{opt}} := \arg\min_{\hat{w}} \mathrm{E}\left((\hat{x}_k - \hat{w}_k \hat{y}_k)^2 \right)$$
$$= \frac{\mathrm{E}\left(|\hat{x}_k|^2 \right)}{\mathrm{E}\left(|\hat{x}_k|^2 \right) + n\sigma^2}$$

Depends on spectral statistics of natural images and on noise level σ^2 (n is the number of pixels)

Image data: Mean square of Fourier coefficients

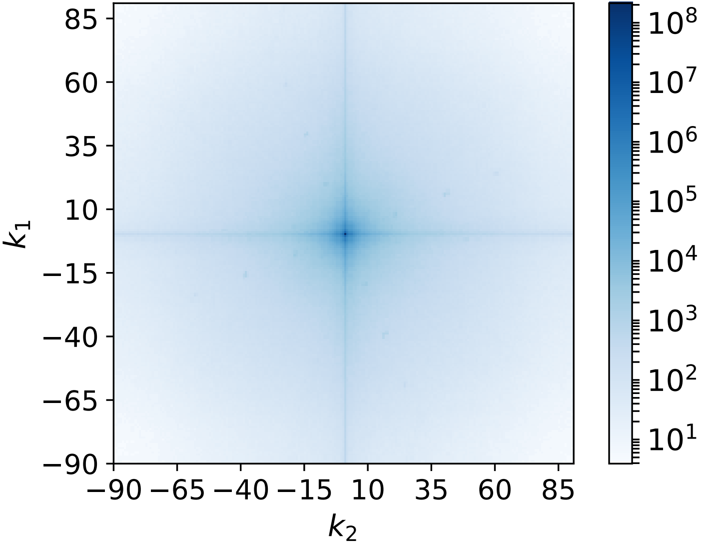

Wiener filter: $\sigma = 0.04$

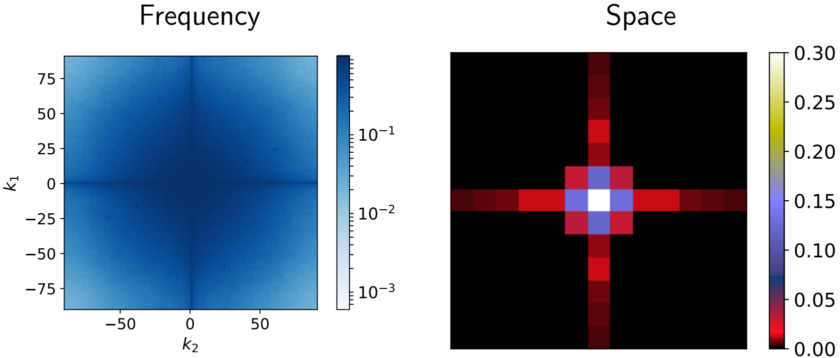

Wiener filter: $\sigma = 0.1$

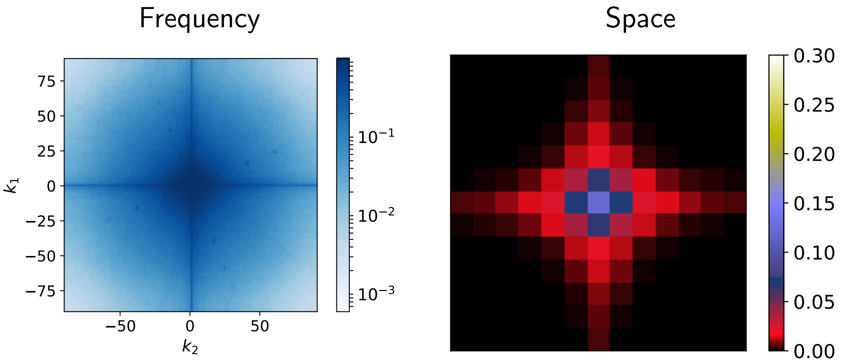

Wiener filter: $\sigma = 0.2$

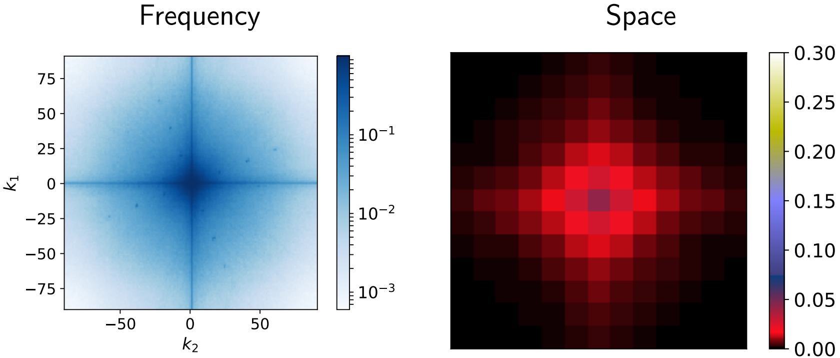

Wiener filter

Two perspectives:

1. *Image domain*: Weighted average of nearby pixels
2. *Frequency domain*: Weighted projection onto low-pass 2D sinusoids

Blind denoising of natural images

- Bias-free CNNs
- Wiener filtering
- CNNs learn adaptive filters
- CNNs learn unions of subspaces

Quantitative magnetic-resonance imaging

Early diagnostics of Alzheimer's disease

Quantitative rehabilitation of stroke patients

Data-driven estimation of sinusoid frequencies

Wiener filter

Image domain: Weighted average of nearby pixels

Problem: Same average for each pixel

Blurs edges and other features

Previous solution:
Adapt filter locally (e.g. bilateral filter [Tomasi and Manduchi 1998])

Bias-free CNN is locally linear

$$f(y) = W_L\, R\, W_{L-1} ... R\, W_1 y = A_y y$$

Rows interpreted as filters

Estimate at pixel i:

$$f_{\text{BF}}(y)(i) = (A_y y)(i) =< i\text{th row of } A_y, y >$$

Low noise

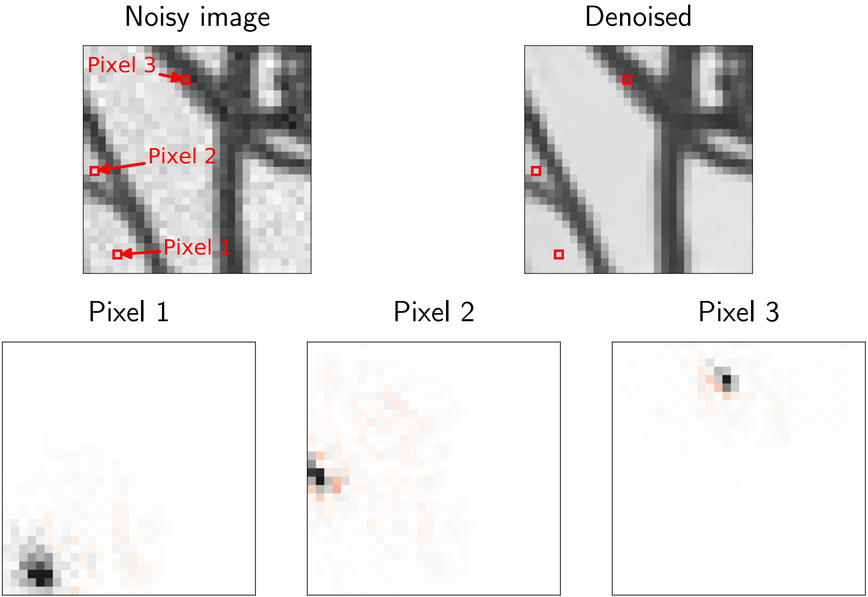
Noisy image
Denoised
Pixel 3
Pixel 2
Pixel 1
Pixel 1
Pixel 2
Pixel 3

Medium noise

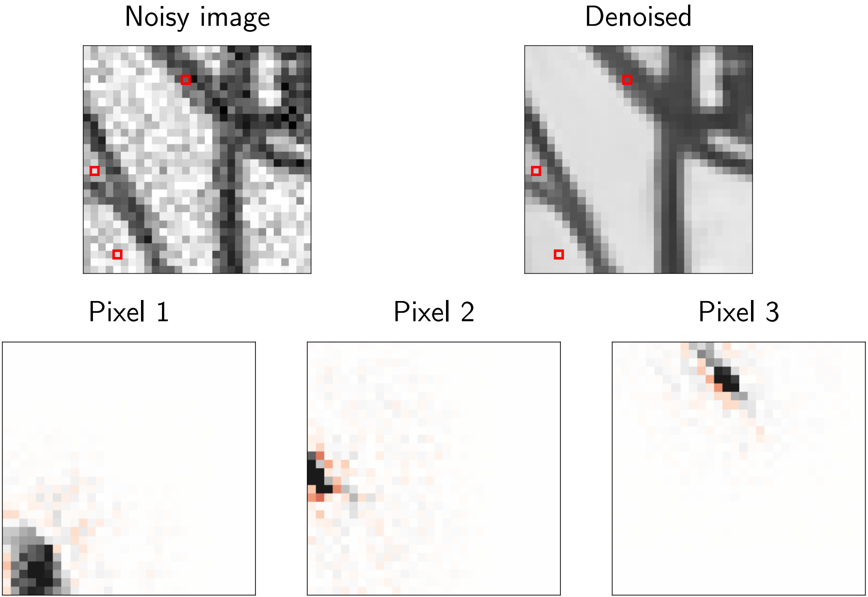

High noise

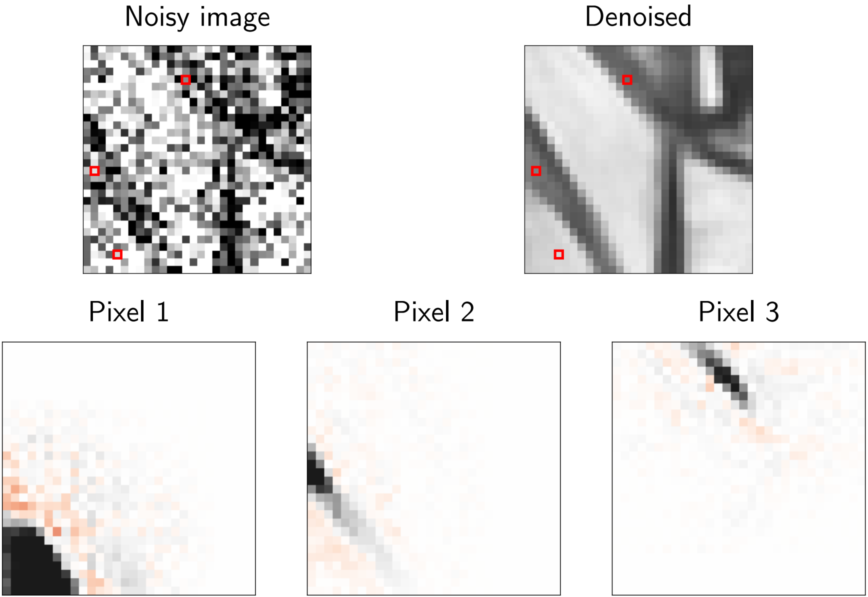

Conclusion

BF-CNN implicitly learns filters *adapted to image structure and noise*!

Blind denoising of natural images

- Bias-free CNNs
- Wiener filtering
- CNNs learn adaptive filters
- CNNs learn unions of subspaces

Quantitative magnetic-resonance imaging

Early diagnostics of Alzheimer's disease

Quantitative rehabilitation of stroke patients

Data-driven estimation of sinusoid frequencies

Wiener filter

Frequency domain: Approximate projection onto low-pass 2D sinusoids

Problem: Same projection for each image

Blurs edges and other features

Projection onto union of subspaces

Previous methodology:

1. Learn/design overcomplete dictionary of basis functions
2. Select sparse subset for each image/patch through thresholding/optimization
3. Project on span of sparse subset

Projection onto union of low-dimensional subspaces

Bias-free CNN is locally linear

$$f(y) = W_L\, R\, W_{L-1} ... R\, W_1 y = A_y y$$

SVD analysis

$$A_y = U\,S\,V^T$$

Empirical observations:

- Matrix is approximately symmetric $U \approx V$
- Matrix is approximately low-rank

Singular values

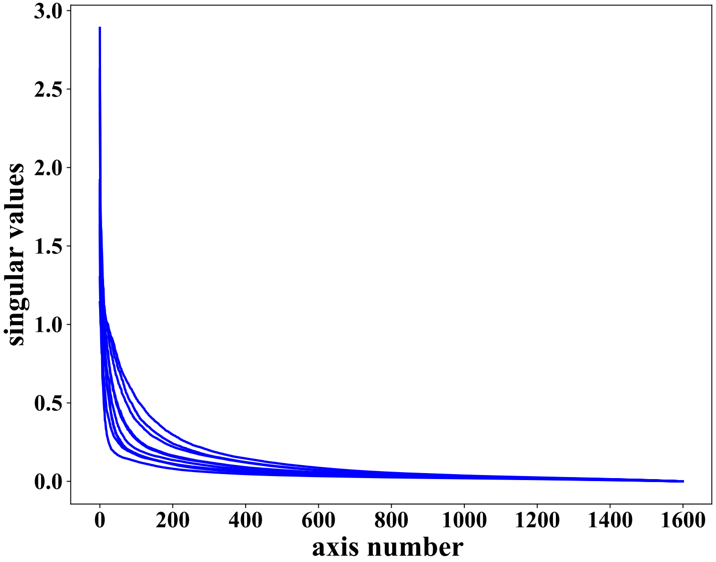

Singular vectors computed from noisy image

Clean image

Large singular values

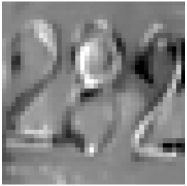

Small singular values

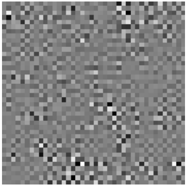

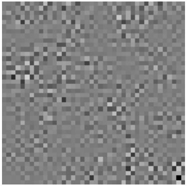

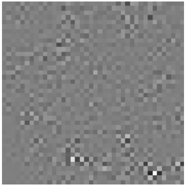

Dimensionality of learned subspace

Approximate dimensionality = sum of squared singular values

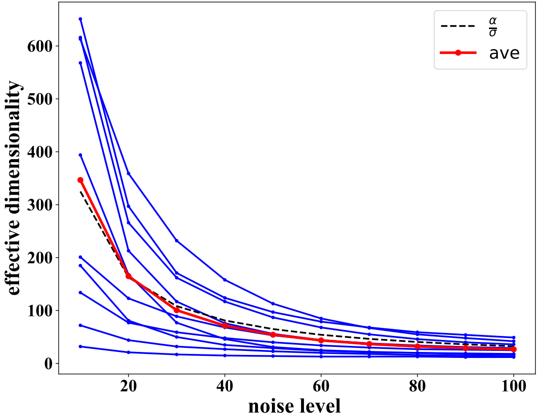

Subspaces are approximately nested

Conclusion

BF-CNN implicitly learns to project onto union of subspaces adapted to image features and noise!

For more information

Robust and interpretable blind image denoising via bias-free convolutional neural networks
S. Mohan, Z. Kadkhodaie, E. Simoncelli, C. Fernandez-Granda

Directions for future research

Properties of the learned representation in frequency estimation

Why does bias hinder generalization across noise levels?

Linear-algebraic analysis is completely empirical and very local

How are these adaptive filters / unions of subspaces learned?

How do the learned mechanisms vary as we change the input?

Blind denoising of natural images

Bias-free CNNs

Wiener filtering

CNNs learn adaptive filters

CNNs learn unions of subspaces

Quantitative magnetic-resonance imaging

Early diagnostics of Alzheimer's disease

Quantitative rehabilitation of stroke patients

Data-driven estimation of sinusoid frequencies

Quantitative magnetic-resonance imaging

Collaboration with the Department of Radiology of the NYU School of Medicine

Joint work with Jakob Assländer, Brett Bernstein, Quentin Duchemin, Cem Gutelkin, Vlad Kobzar, Sylvain Lannuzel, and Sunli Tang

Magnetic-resonance imaging (MRI)

- Hydrogen nuclei absorb/emit radio-frequency energy when placed in magnetic field
- Measured signal depends on relaxation parameters T_1 and T_2 of biological tissues

Traditional contrast-based MRI

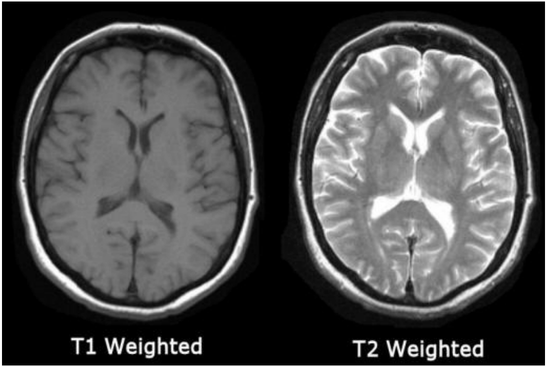

Not quantitative!

Difficult to reproduce/compare

Quantitative MRI via fingerprinting

Radio-frequency pulses are designed to produce irregular magnetization signals (fingerprints) encoding relaxation parameters

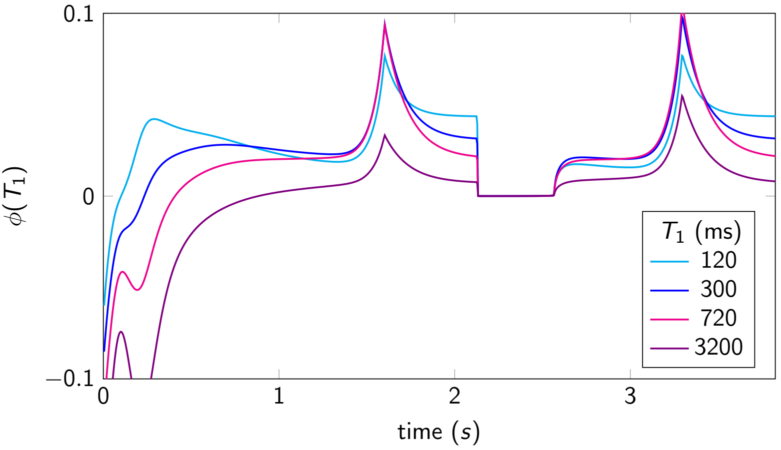

Multicompartment magnetic resonance fingerprinting

- Assumption in MRF: One tissue per voxel
- Problematic at tissue boundaries
- Ignores sub-voxel structure

Additive model

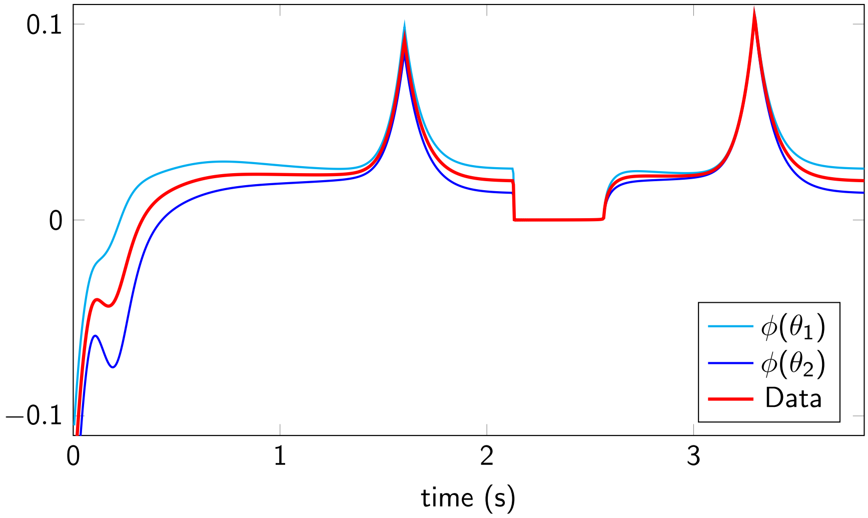

Correlation structure

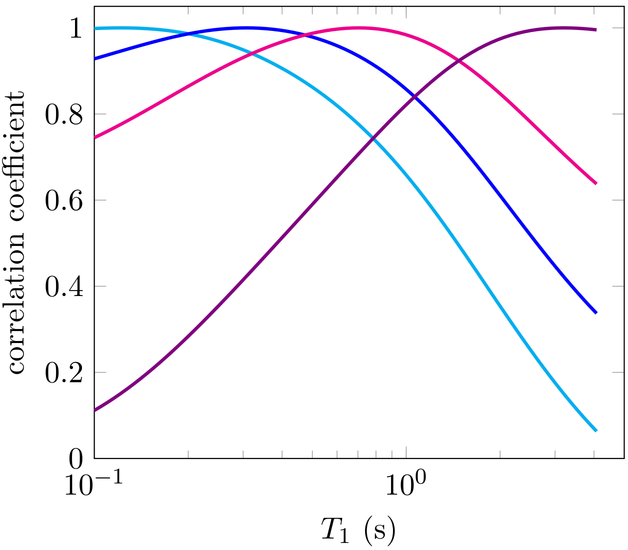

Multicompartment MRF via ℓ_1-norm regularization

- Fast-thresholding methods don't work
- We use an efficient interior-point solver
- Solving sequence of reweighted problems improves the solution

Drawback: Very slow

Validation with phantom

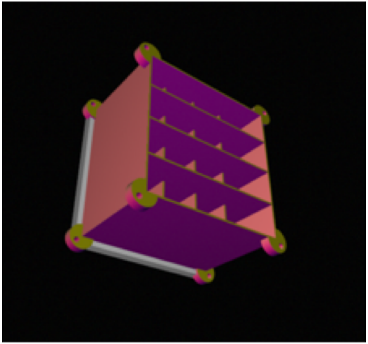

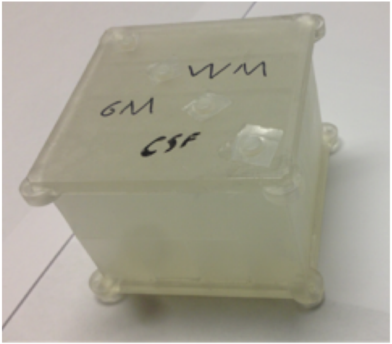

Validation with phantom

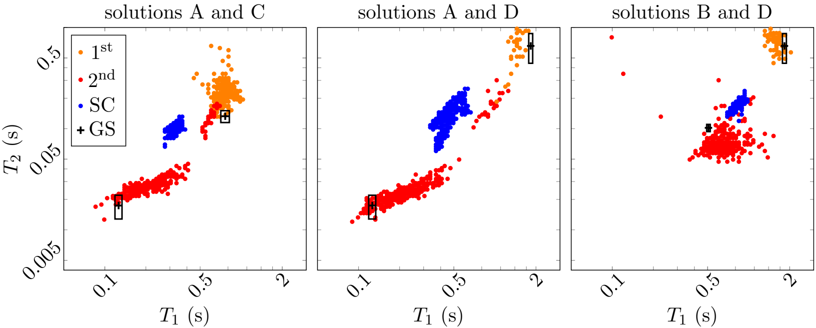

Current research

Goal: Fast multicompartment MRF for non-additive model

- Measurement design via ODE-constrained optimization
- Parameter estimation using a feedforward deep neural network trained on simulated data

Current research

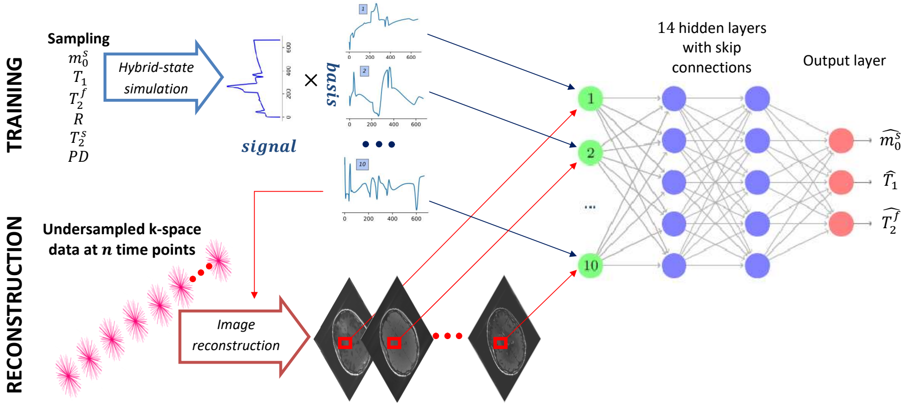

For more information

Multi-Compartment MR Fingerprinting via Reweighted-l1-norm Regularization. S. Tang, J. Asslaender, L. Tanenbaum, R. Lattanzi, M. Cloos, F. Knoll, C. Fernandez-Granda. ISMRM 2017

Multicompartment magnetic resonance fingerprinting. S. Tang, C. Fernandez-Granda, S. Lannuzel, B. Bernstein, R. Lattanzi, M. Cloos, F. Knoll and J. Asslaender. Inverse Problems 34 (9) 4005. 2018

Hybrid-State Free Precession for Measuring Magnetic Resonance Relaxation Times in the Presence of B0 Inhomogeneities. V. Kobzar, C. Fernandez-Granda, J. Asslaender. ISMRM 2019

Blind denoising of natural images

- Bias-free CNNs
- Wiener filtering
- CNNs learn adaptive filters
- CNNs learn unions of subspaces

Quantitative magnetic-resonance imaging

Early diagnostics of Alzheimer's disease

Quantitative rehabilitation of stroke patients

Data-driven estimation of sinusoid frequencies

Early diagnostics of Alzheimer's disease

Joint work with Sheng Liu, Narges Razavian, and Chhavi Yadav

Early diagnostics of Alzheimer's disease

Goal: Distinguish between three classes

1. Normal
2. Mild cognitive impairment
3. Mild Alzheimer's

Data: Structural MRI (T1) from Alzheimer's Disease Neuroimaging Initiative

Preprocessing: Images are registered to a template

Demographics

Split	Class	Num. subjects	Num. Scans	Mean Age (std)
Train	CN	140	567	77.0 (5.4)
	MCI	248	840	75.9 (7.3)
	AD	193	527	76.7 (7.4)
Val	CN	33	126	77.2 (5.6)
	MCI	39	138	73.3 (7.2)
	AD	41	124	76.1 (8.3)
Test	CN	24	105	79.0 (6.1)
	MCI	43	140	76.7 (6.5)
	AD	45	135	76.4 (5.1)

Simple biomarker (normalized volumes)

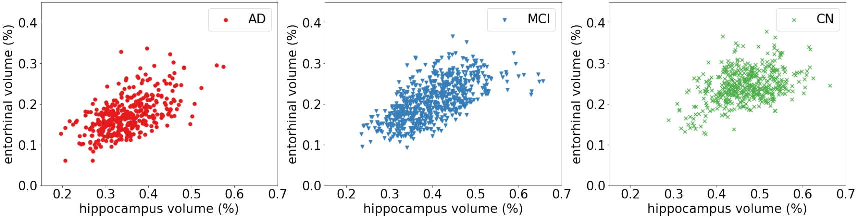

Accuracy: around 50%

Methodology

3D convolutional neural network

Main insights, performance is improved by:

- Using small (1x1) filter sizes in first layer
- Widening the network (as opposed to deepening)
- Using instance normalization instead of batch normalization
- Encoding age using a sinusoidal embedding

Architecture

Block	Layer	Type	Output size
	Inputs		$96 \times 96 \times 96$
1	Conv3D	k1-c4$\cdot f$-p0-s1-d1	$96 \times 96 \times 96$
	InstanceNorm3D		
	ReLU		
	MaxPool3D	k3-s2	$47 \times 47 \times 47$
2	Conv3D	k3-c32$\cdot f$-p0-s1-d2	$43 \times 43 \times 43$
	InstanceNorm3D		
	ReLU		
	MaxPool3D	k3-s2	$21 \times 21 \times 21$
3	Conv3D	k5-c64$\cdot f$-p2-s1-d2	$17 \times 17 \times 17$
	InstanceNorm3D		
	ReLU		
	MaxPool3D	k3-s2	$8 \times 8 \times 8$
4	Conv3D	k3-c64$\cdot f$-p1-s1-d2	$6 \times 6 \times 6$
	InstanceNorm3D		
	ReLU		
	MaxPool3D	k5-s2	$5 \times 5 \times 5$
FC1		1024	
FC2		3	
Softmax		3	

Results

Method	Accuracy	Balanced Acc	Micro-AUC	Macro-AUC
ResNet-18	50.8%	-	-	-
ResNet-18 pretrained	56.8%	-	-	-
ResNet-18 3D	$52.4 \pm 1.8\%$	53.1%	-	-
ResNet-18 3D	$50.1 \pm 1.1\%$	$51.3 \pm 1.0\%$	$71.2 \pm 0.4\%$	$72.4 \pm 0.7\%$
AlexNet 3D	$57.2 \pm 0.5\%$	$56.2 \pm 0.8\%$	$75.1 \pm 0.4\%$	$74.2 \pm 0.5\%$
proposed	$66.9 \pm 1.2\%$	$67.9 \pm 1.1\%$	$82.0 \pm 0.7\%$	$78.5 \pm 0.7\%$
proposed + Age	$\mathbf{68.2 \pm 1.1\%}$	$\mathbf{70.0 \pm 0.8\%}$	$\mathbf{82.0 \pm 0.2\%}$	$\mathbf{80.0 \pm 0.5\%}$

Australian Imaging, Biomarkers and Lifestyle dataset

Method	Accuracy	Balanced Acc	Micro-AUC	Macro-AUC
proposed on ADNI	$66.9 \pm 1.2\%$	$67.9 \pm 1.1\%$	$82.0 \pm 0.7\%$	$78.5 \pm 0.7\%$
proposed on AIBL	$63.6 \pm 0.7\%$	$65.7 \pm 1.1\%$	$90.0 \pm 0.6\%$	$82.1 \pm 0.7\%$

Visualization of gradient with respect to input

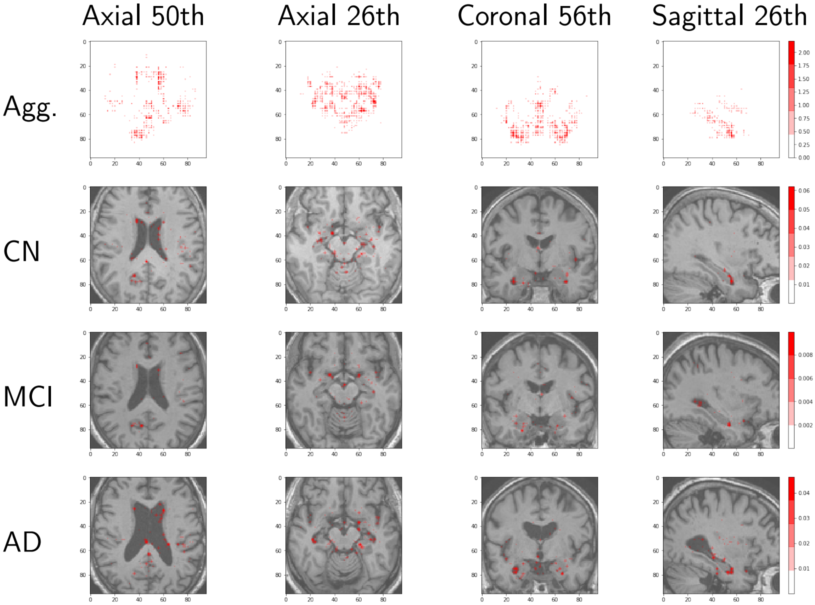

Blind denoising of natural images

Bias-free CNNs

Wiener filtering

CNNs learn adaptive filters

CNNs learn unions of subspaces

Quantitative magnetic-resonance imaging

Early diagnostics of Alzheimer's disease

Quantitative rehabilitation of stroke patients

Data-driven estimation of sinusoid frequencies

Quantitative rehabilitation of stroke patients

Collaboration with the Mobilis lab at the Department of Neurology of the NYU School of Medicine

Joint work with Aakash Kaku, Avinash Parnandi, and Heidi Schambra

Quantitative rehabilitation of stroke patients

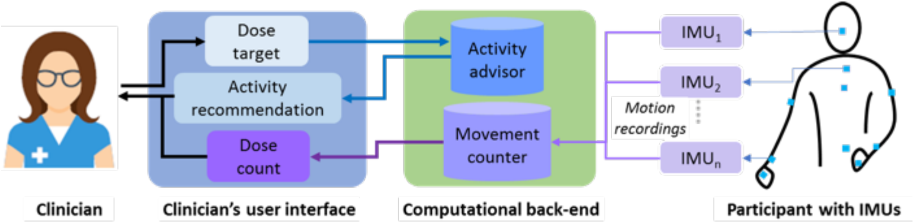

Goal: Automatic identification/counting of basic upper body movements

Data: 100 dimensional time series (accelerations, rotations)


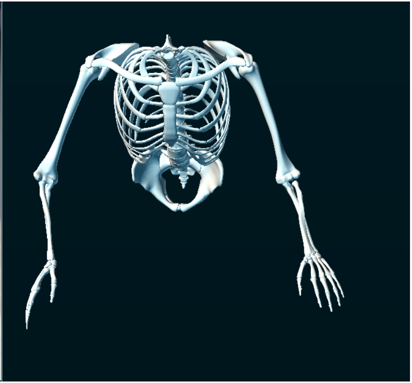

Methodology

- Deep convolutional neural networks achieves great results for fixed group of patients
- To be clinically practical we need to generalize to new patients
- Promising results by normalizing features (*instance normalization*)

Blind denoising of natural images

- Bias-free CNNs
- Wiener filtering
- CNNs learn adaptive filters
- CNNs learn unions of subspaces

Quantitative magnetic-resonance imaging

Early diagnostics of Alzheimer's disease

Quantitative rehabilitation of stroke patients

Data-driven estimation of sinusoid frequencies

Data-driven estimation of sinusoid frequencies

Joint work with Brett Bernstein, Gautier Izacard, and Sreyas Mohan

Frequency estimation (aka super-resolution of line spectra)

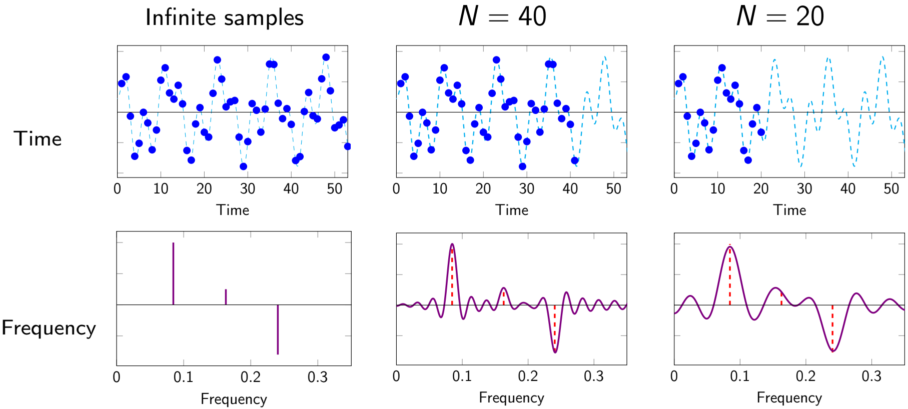

Traditional methodology

- Linear estimation (periodogram)
- Parametric methods based on eigendecomposition of sample covariance matrix (MUSIC, ESPRIT, matrix pencil)
- Sparsity-based methods

Learning-based approach

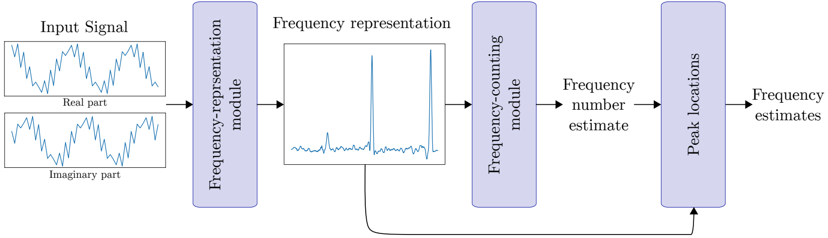

Frequency-representation module

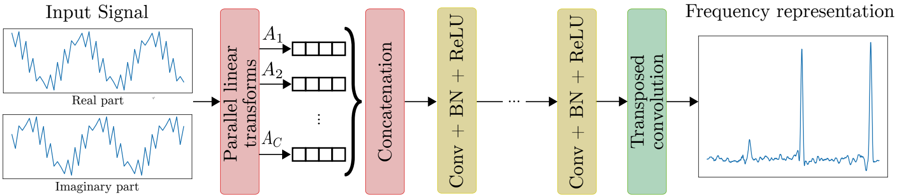

Fourier transform of learned transformations

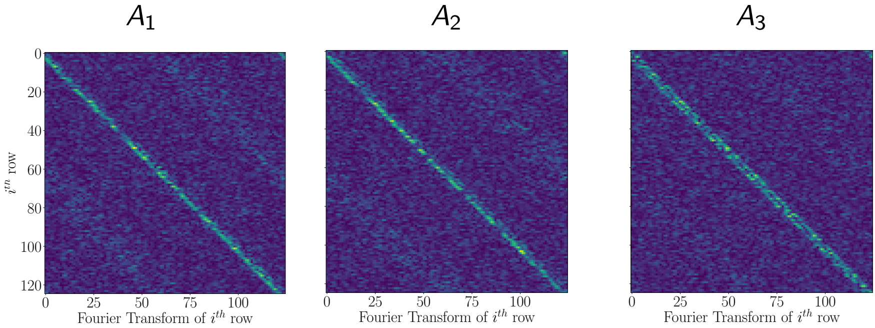

Comparison to state of the art

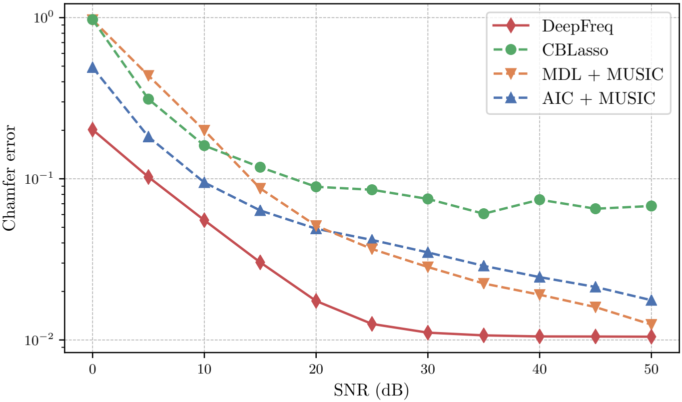

For more information

A Learning-Based Framework for Line-Spectra Super-resolution. G. Izacard, B. Bernstein, C. Fernandez-Granda. ICASSP 2019

Data-driven Estimation of Sinusoid Frequencies. G. Izacard, S. Mohan, C. Fernandez-Granda. NeurIPS 2019