Compressed Sensing and Machine Learning for Radar Imaging

Müjdat Çetin

Associate Professor, Department of Electrical and Computer Engineering
Interim Director, Goergen Institute for Data Science

University of Rochester, Rochester, NY

Contributors: Burak Alver, Ammar Saleem, Sadegh Samadi, Abdurrahim Soğanlı, Emre Güven, Alper Güngör, W. Clem Karl.

Computational Imaging

Computed tomography

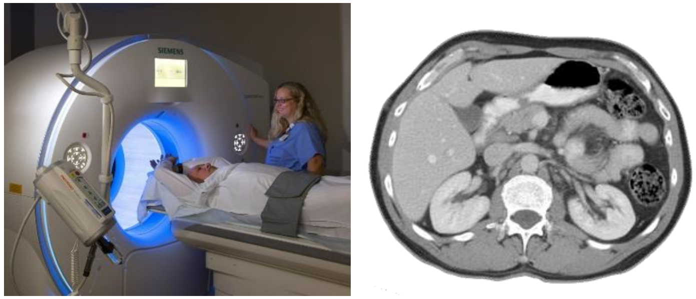

Seismic imaging

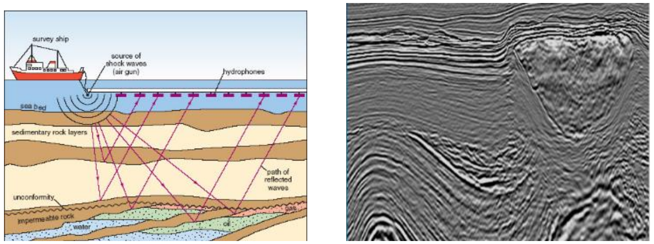

Coded-aperture imaging

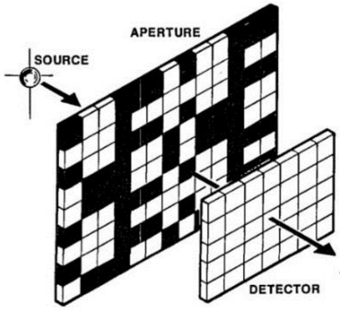

Computational microscopy

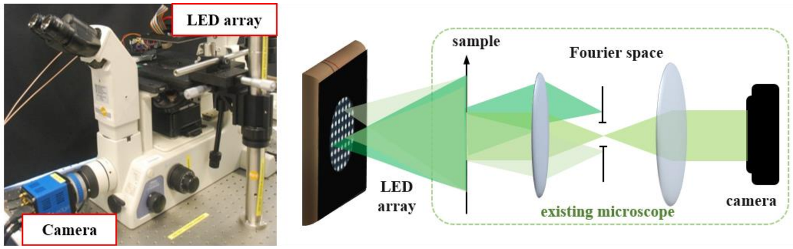

Light-field / plenoptic imaging

Single-pixel imaging

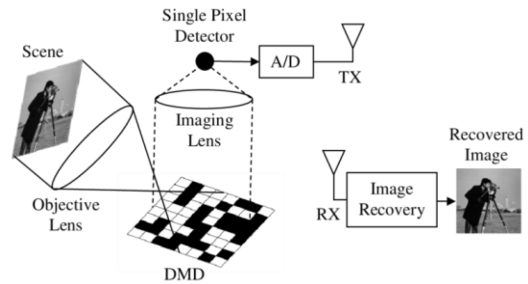

Radar Imaging Basics

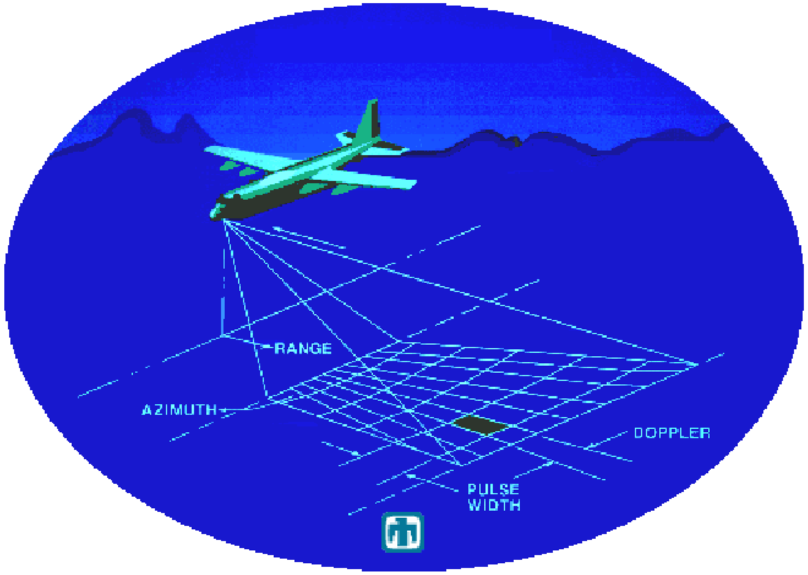

- All-weather
- Day and night operation
- Superposition of response from scatterers – tomographic measurements

- Synthetic aperture radar (SAR)
- Computational imaging problem:
  Obtain a spatial map of reflectivity from radar returns

Outline

- Sparsity and compressed sensing for radar imaging
- Machine learning for radar imaging

Initial motivation for our work

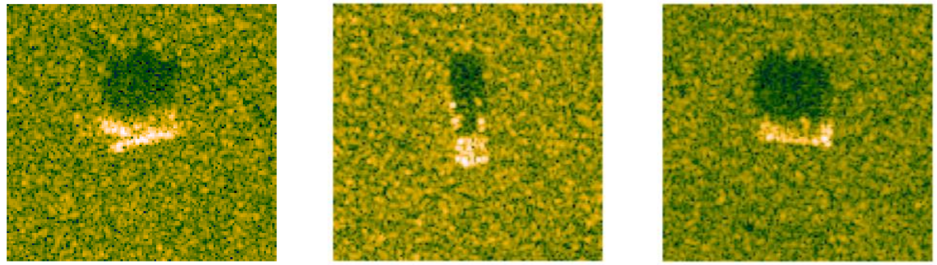

Some challenges for automatic decision-making from SAR images:

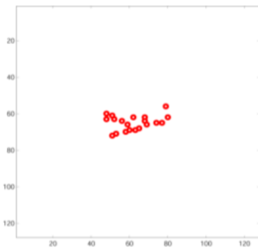

- Accurate localization of **dominant scatterers**
  - Limited resolution
  - Clutter and artifact energy

- **Region separability**
  - Speckle
  - Object boundaries
- Low SNR, limited apertures

Sparsity and Compressibility of Signals

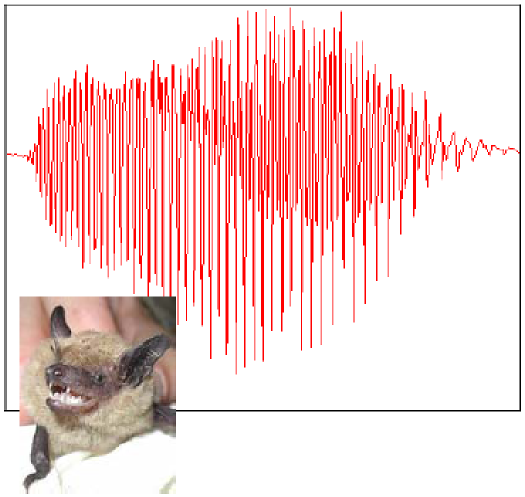

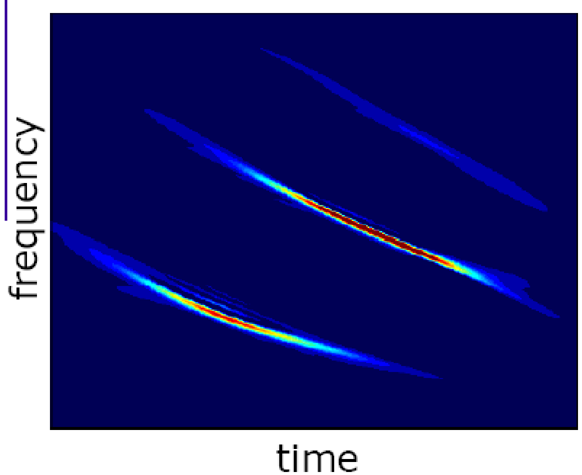

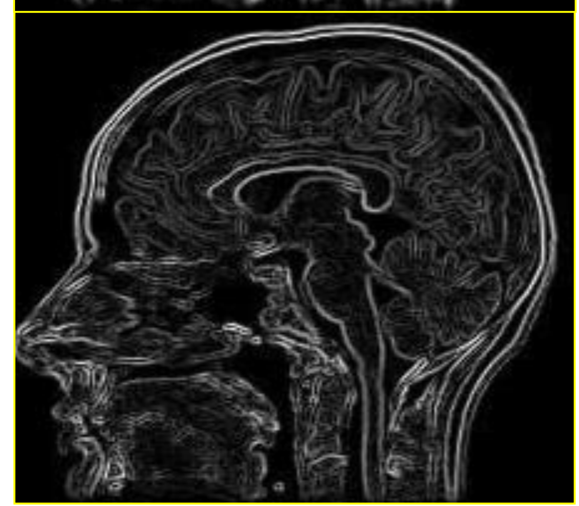

Underdetermined Linear Inverse Problems

$\mathbf{y}$ = $\mathbf{A}$ $\boldsymbol{\Psi}$ $\boldsymbol{\alpha}$

Observation Matrix

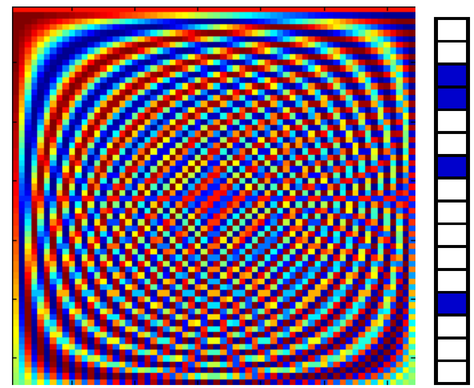

L non-zero elements

Observed data

$$\mathbf{A}: \quad M \times N$$

$$N > M > L$$

$\mathbf{f}$

- A simple example:
  - $\mathbf{A}$: "band-limited" DFT operator
  - $\boldsymbol{\Psi}$: identity matrix

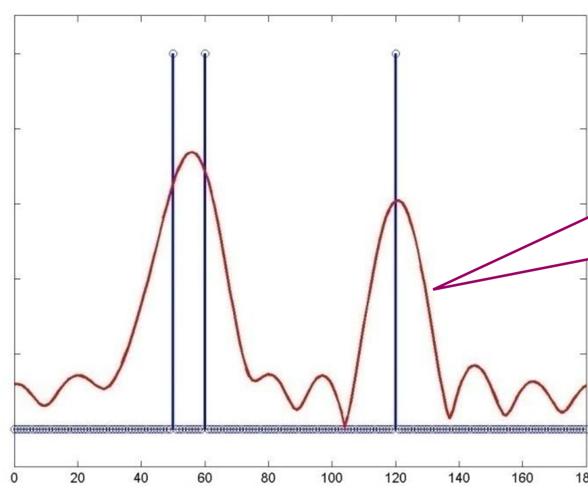

Conventional signal processing result

Underdetermined Linear Inverse Problems, Sparsity, Compressed Sensing

- Basic problem: find an estimate of $\mathbf{f}^*$, where

$$\mathbf{y} = \mathbf{A}\mathbf{f}^* \qquad (\mathbf{A} : M \times N,\ \ N > M) \qquad \left[\ \longleftarrow \mathbf{A} \longrightarrow\ \right]$$

- Underdetermined -- non-uniqueness of solutions
- Additional information/constraints needed for a unique solution
- If we know $\mathbf{f}^*$ is sparse (*i.e.,* has few non-zero elements)?

$$\hat{\mathbf{f}}_{\ell_0} = \arg\min \ \|\mathbf{f}\|_0^0 \ \text{ subject to } \mathbf{y} = \mathbf{A}\mathbf{f}$$

Number of non-zero elements in f

- Intractable combinatorial optimization problem
- **Past work on sparse signal representation (including ours) has produced principled and feasible alternatives**
  - l_p relaxations or greedy methods

SAR Ground-plane Geometry

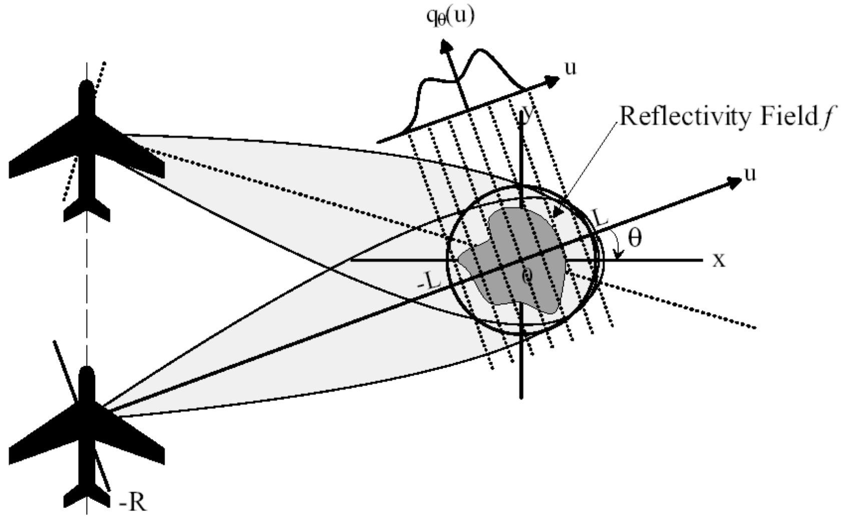

- Scalar 2-D complex reflectivity field $f(x,y)$
- Transmitted chirp signal: $s(t) = \Re\left[e^{j(\omega_0 t + \alpha t^2)}\right], \quad |t| \le \frac{T_p}{2}$
- Received, demodulated return from circular patch:

$$r_\theta(t) = \overbrace{\int_{|u|\le L} \underbrace{q_\theta(u)}_{\text{Projection of field } f(x,y)} \exp\left\{-j\frac{2}{c}\left[\omega_0 + 2\alpha\left(t - \frac{2R}{c}\right)\right]u\right\} du}^{\text{Band-limited Fourier transform of } q_\theta(u)}$$

SAR Observation Model

- Observations are related to projections of the field:

$$r_\theta(t) = \int_{|u|\le L} \underbrace{q_\theta(u)}_{\substack{\text{Projection of field} \\ f(x,y) \text{ at angle } \theta}} \exp\left\{ \underbrace{-j\frac{2}{c}\left[\omega_0 + 2\alpha\left(t - \frac{2R}{c}\right)\right]}_{\substack{\Omega(t) \\ \text{Spatial frequency}}} u \right\} du$$

- SAR observations are band-limited slices from the 2-D Fourier transform of the reflectivity field:

$$\begin{aligned} r_\theta(t) &= \iint_{x^2+y^2\le L^2} f(x,y) \exp\left\{-j\Omega(t)\left(x\cos\theta + y\sin\theta\right)\right\} \, dx\, dy \\ &= F\left[\Omega(t)\cos\theta, \Omega(t)\sin\theta\right] \end{aligned}$$

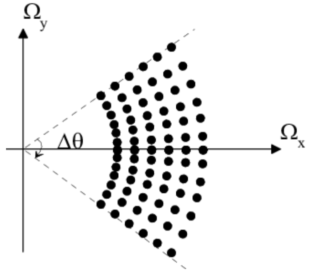

- **Discrete tomographic SAR observation model:**

(combining all measurements)

$$\boxed{\underbrace{\mathbf{y}}_{\text{Observed data}} = \underbrace{\mathbf{A}}_{\text{SAR Forward Model}}\underbrace{\mathbf{f}}_{\text{Unknown field}} + \underbrace{\mathbf{n}}_{\text{Noise}}}$$

Conventional Image Formation

- Given SAR returns, create an estimate of the reflectivity field **f**

Support of observed data in the spatial frequency domain

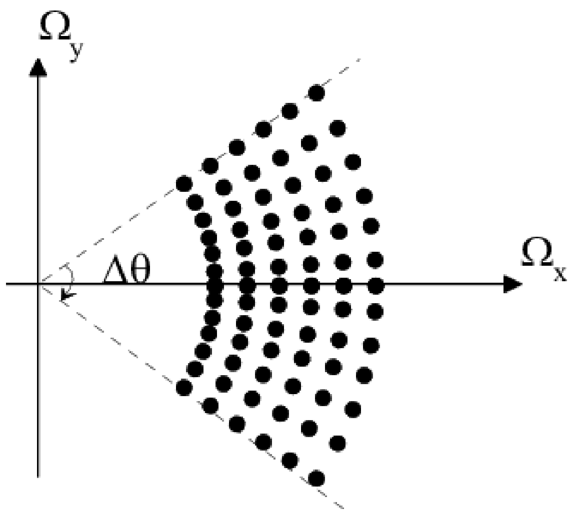

Sample Conventional Image

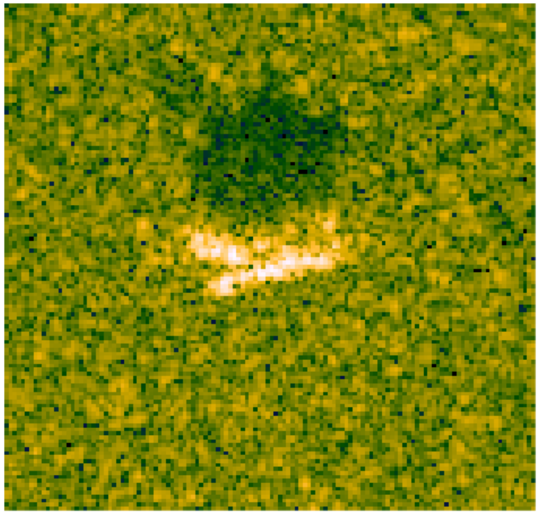

Polar format algorithm:

- **Each pulse gives slice of 2-D Fourier transform of field**
- Polar to rectangular resampling
- 2-D inverse DFT

Sparsity-Driven Radar Imaging – basic version

$$J(\mathbf{f}) = \|\mathbf{y} - \mathbf{A}\mathbf{f}\|_2^2 + \lambda \|\mathbf{L}|\mathbf{f}|\|_p^p$$

- Bayesian interpretation: MAP estimation problem with heavy-tailed priors $p(\mathrm{f} \mid \mathrm{y}) \propto p(\mathrm{y} \mid \mathrm{f})\, p(\mathrm{f})$
- Complex-valued data and image
- Magnitude of complex-valued field admits sparse representation
- No informative prior on reflectivity phase
- Typical choices for L:
  - identity (***point-enhanced imaging***)
  - gradient (***region-enhanced imaging***)

- Optimization problem structure is different from common sparse representation problems

One way to solve the optimization problem

- **Alternating Direction Method of Multipliers (ADMM)**
- An augmented Lagrangian method developed in 1970s, with roots in 1950s -- rediscovered recently!
- Contains ideas involving dual decomposition, method of multipliers, proximal methods, variable splitting
- **Enables decoupling terms related to data and priors**
- Suited to distributed optimization

A basic ADMM for l_1 minimization

- Cost function: $J(\mathbf{f}) = \frac{1}{2}\|\mathbf{y} - \mathbf{Af}\|_2^2 + \lambda\|\mathbf{f}\|_1$
- Augmented Lagrangian with variable splitting:

$$L_\rho(\mathbf{f}, \mathbf{g}, \mathbf{u}) = \frac{1}{2}\|\mathbf{y} - \mathbf{Af}\|_2^2 + \lambda\|\mathbf{g}\|_1 + \rho\mathbf{u}^T(\mathbf{f} - \mathbf{g}) + \frac{\rho}{2}\|\mathbf{f} - \mathbf{g}\|_2^2$$

- Iterative solution:

Data $$\mathbf{f}^{k+1} = \left(\mathbf{A}^T\mathbf{A} + \rho\mathbf{I}\right)^{-1}\left(\mathbf{A}^T\mathbf{y} + \rho\left(\mathbf{g}^k - \mathbf{u}^k\right)\right)$$

Prior $$\mathbf{g}^{k+1} = S_{\lambda/\rho}\left(\mathbf{f}^{k+1} + \mathbf{u}^k\right)$$

$$\mathbf{u}^{k+1} = \mathbf{u}^k + \left(\mathbf{f}^{k+1} - \mathbf{g}^{k+1}\right)$$

Point-Enhanced Imaging

Synthetic scene

Original　　　Conventional　　　Point-Enhanced

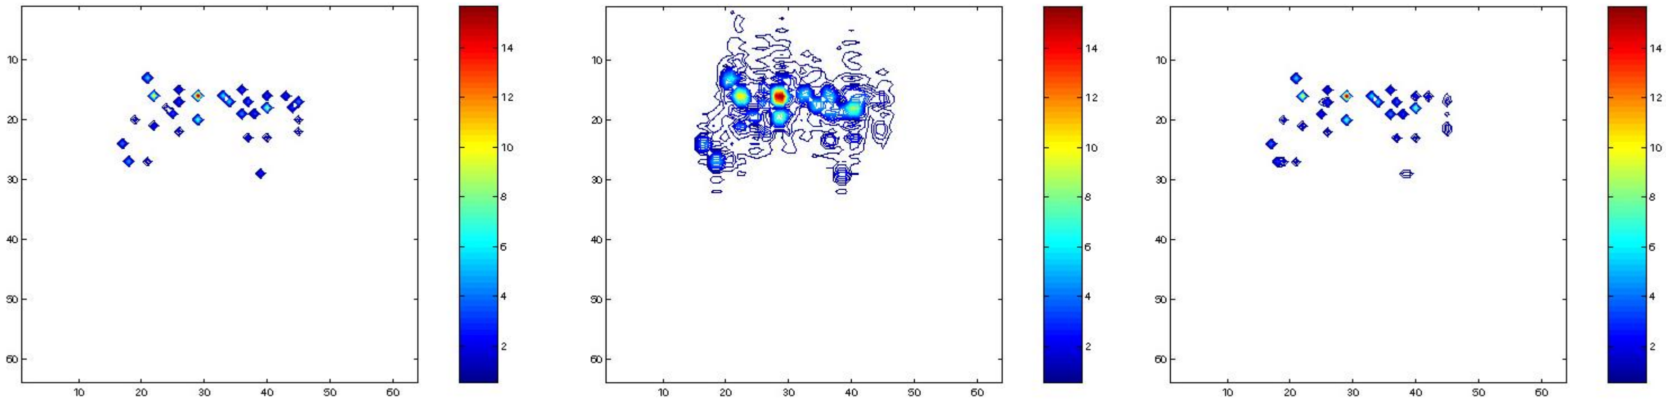

Sparsity-Driven SAR Imaging Results

Conventional

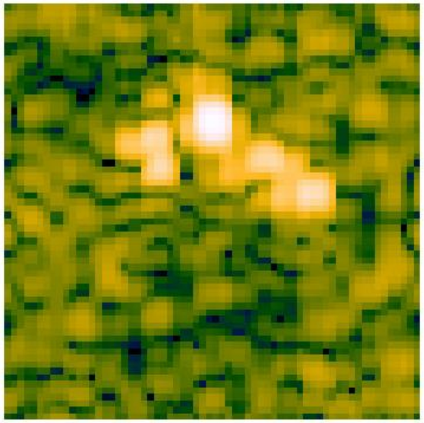

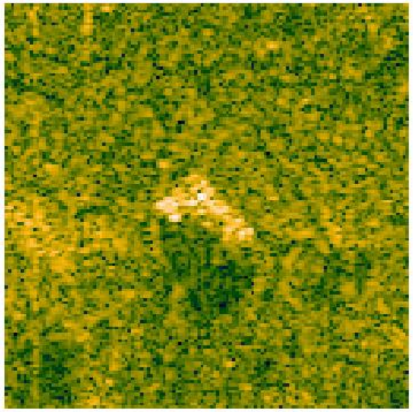

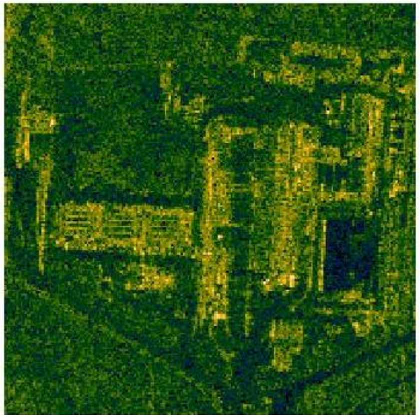

Sparsity-driven

Point-enhanced

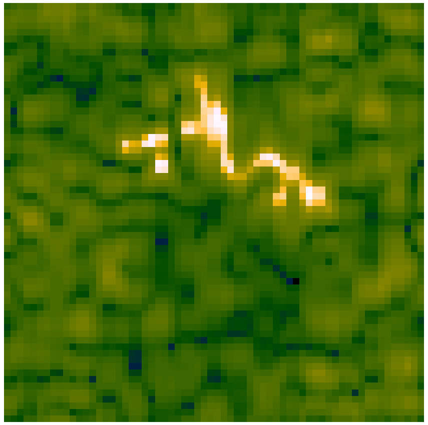

Region-enhanced

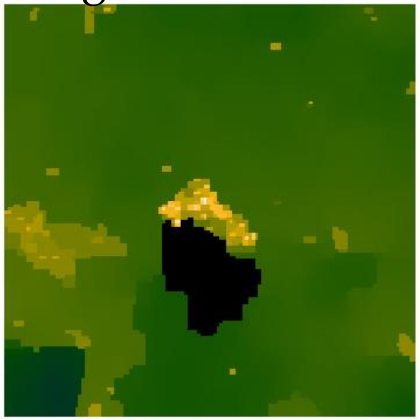

Wavelet dictionary

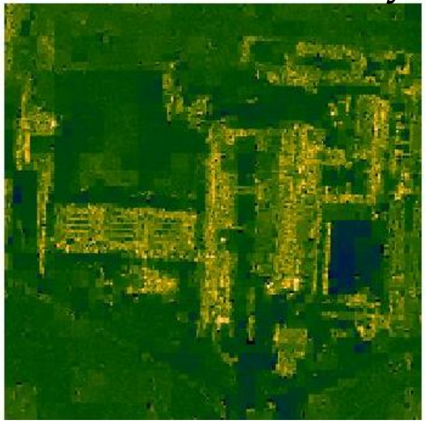

Sparsity-Driven SAR Imaging Results

Conventional

Sparsity-driven, ADMM-based

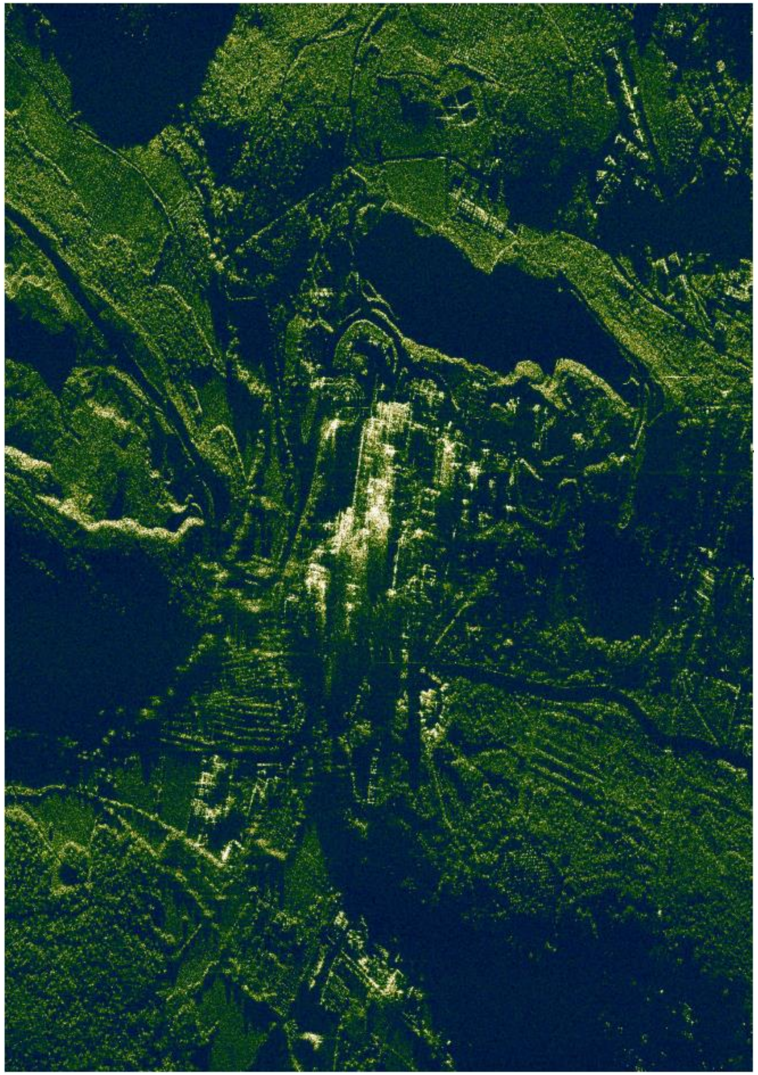

Dictionary Learning for Sparsity-Driven SAR Imaging

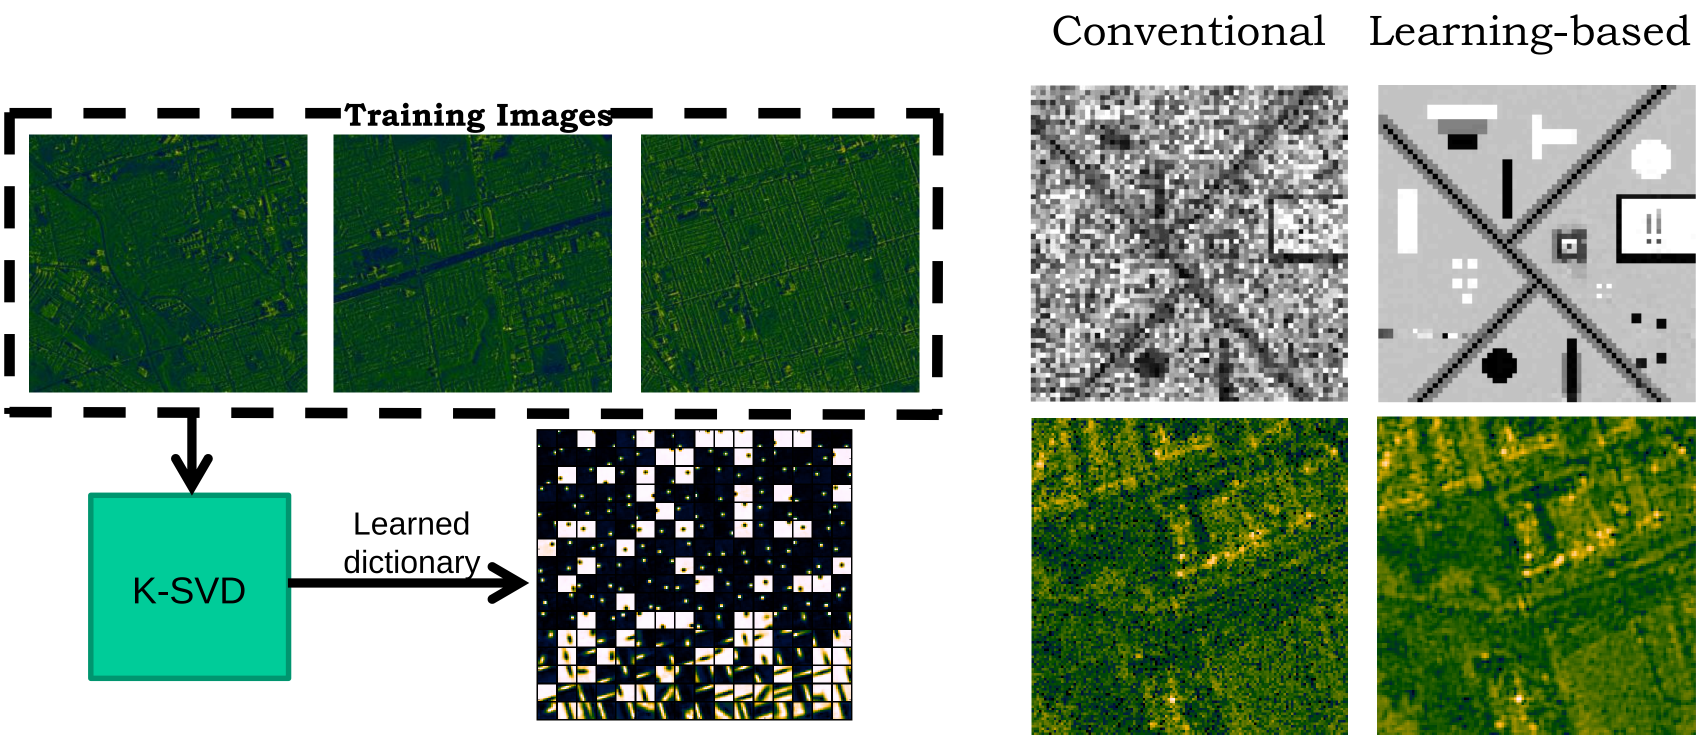

Deep Learning-based Priors for SAR Imaging

- A new SAR image reconstruction framework utilizing Plug-and-Play (PnP) priors
  - Optimization-based reconstruction - regularized inversion, MAP estimation
  - Decoupling the data and the prior through ADMM
  - Deep learning-based prior

Problem Formulation

- Discretized SAR observation model:

$$\mathbf{y} = \mathbf{A}\mathbf{f} + \mathbf{n}$$

- Retrieve **f** using a regularized cost function:

$$\hat{\mathbf{f}} = \underset{\mathbf{f}}{\operatorname{argmin}}\{\mathfrak{D}(\mathbf{f}) + \lambda\mathfrak{R}(\mathbf{f})\}$$

where $\mathfrak{D}(\mathbf{f}) = \|\mathbf{y} - \mathbf{A}\mathbf{f}\|_2^2$ and $\mathfrak{R}(\mathbf{f})$ is the regularizer

Problem Formulation

- Prior information or regularization constraints on the magnitude of $\mathbf{f}$
- Rewrite $\mathbf{f} = \boldsymbol{\Theta}|\mathbf{f}|$ where is $\boldsymbol{\Theta}$ a diagonal matrix containing the phase of $\mathbf{f}$ in the form $e^{j\phi(\mathbf{f})}$
- Cost function becomes:

$$\left\{|\hat{\mathbf{f}}|, \widehat{\boldsymbol{\Theta}}\right\} = \underset{|\mathbf{f}|,\boldsymbol{\Theta}}{\operatorname{argmin}} \|\mathbf{y} - \mathbf{A}\boldsymbol{\Theta}|\mathbf{f}|\|_2^2 + \lambda\Re(\mathbf{f})$$

Variable Splitting and ADMM

- Introduce an auxiliary variable with a constraint:

$$\{|\hat{\mathbf{f}}|, \widehat{\boldsymbol{\Theta}}, \hat{\mathbf{h}}\} = \underset{|\mathbf{f}|,\boldsymbol{\Theta},\mathbf{h}}{\operatorname{argmin}} \|\mathbf{y} - \mathbf{A}\boldsymbol{\Theta}|\mathbf{f}|\|_2^2 + \lambda\Re(\mathbf{h})$$

$$s.t. \; |\mathbf{f}| - \mathbf{h} = 0$$

- Augmented Lagrangian (in scaled form):

$$\{|\hat{\mathbf{f}}|, \widehat{\boldsymbol{\Theta}}, \hat{\mathbf{h}}, \hat{\mathbf{u}}\} = \underset{|\mathbf{f}|,\boldsymbol{\Theta},\mathbf{h},\mathbf{u}}{\operatorname{argmin}} \|\mathbf{y} - \mathbf{A}\boldsymbol{\Theta}|\mathbf{f}|\|_2^2 + \lambda\Re(\mathbf{h})$$

$$+ \frac{\rho}{2}\||\mathbf{f}| - \mathbf{h} + \mathbf{u}\|_2^2 + \frac{\rho}{2}\|\mathbf{u}\|_2^2$$

Variable Splitting and ADMM - details

- Let $\boldsymbol{\theta} \in \mathbb{C}^{N\times 1}$ be a vector containing the diagonal elements of the phase matrix $\boldsymbol{\Theta}$
- Invoke the constraint that the magnitudes of the elements of $\boldsymbol{\theta}$ should be 1, since they contain phases in the form $e^{j\phi(\mathbf{f})}$
- Let $\mathbf{B}$ be a matrix whose diagonal elements contain the reflectivity magnitudes
- Let $\tilde{\mathbf{f}} = \hat{\mathbf{h}} - \mathbf{u}$ and $\tilde{\mathbf{h}} = |\hat{\mathbf{f}}| + \mathbf{u}$

Variable Splitting and ADMM

- Each iteration of the ADMM algorithm performs the following steps enabling the use of a Plug-and-Play (PnP) prior approach:

Data $$\widehat{\boldsymbol{\theta}} = \underset{\boldsymbol{\theta}}{\operatorname{argmin}} \|\mathbf{y} - \mathbf{A}\mathbf{B}\boldsymbol{\theta}\|_2^2 + \lambda_{\boldsymbol{\theta}} \sum_{i=1}^{N} (|\boldsymbol{\theta}_i| - 1)^2$$

Data $$|\hat{\mathbf{f}}| = \underset{|\mathbf{f}|}{\operatorname{argmin}} \|\mathbf{y} - \mathbf{A}\boldsymbol{\Theta}|\mathbf{f}|\|_2^2 + \frac{\rho}{2} \left\| |\mathbf{f}| - \tilde{\mathbf{f}} \right\|_2^2$$

Prior $$\hat{\mathbf{h}} = \underset{\mathbf{h}}{\operatorname{argmin}}\, \lambda \Re(\mathbf{h}) + \frac{\rho}{2} \left\| \tilde{\mathbf{h}} - \mathbf{h} \right\|_2^2$$

$$\hat{\mathbf{u}} = \mathbf{u} + |\hat{\mathbf{f}}| - \hat{\mathbf{h}}$$

where $\lambda_{\boldsymbol{\theta}}$ is a hyperparameter

Convolutional Neural Network (CNN)-based Prior

- Architecture: modified version of the network used in [1]
- 20 convolutional modules
- First and even numbered modules: 64 3×3 filters with padding 1, stride 1
- Remaining modules: 64 5×5 filters with padding 2, stride 1
- Each module has batch normalization and ReLU layers

[1] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, "Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising," IEEE Transactions on Image Processing, 26(7):3142-3155, 2017.

Training the CNN - Synthetic Scenes

Using 64×64 ground truth images for CNN training:

- Add random phase to training images and obtain the phase histories.
- Apply complex-valued additive noise to the phase histories.
  - Magnitude uniformly distributed over $[0, \sigma_y]$, where σ_y is the standard deviation of the magnitude of the phase history data
  - Phase uniformly distributed over $[-\pi, \pi]$
- Perform conventional reconstructions.
- Extract 16×16 overlapping patches from these conventional images and their corresponding ground truths, to construct input-output pairs.
- Augment the pairs of images through rotation by $[90°, 180°, 270°]$.
- Train the network using these augmented pairs of images, with image reconstructed from noisy data as input and ground truth image as output.

Synthetic Data Experiments --Training Set

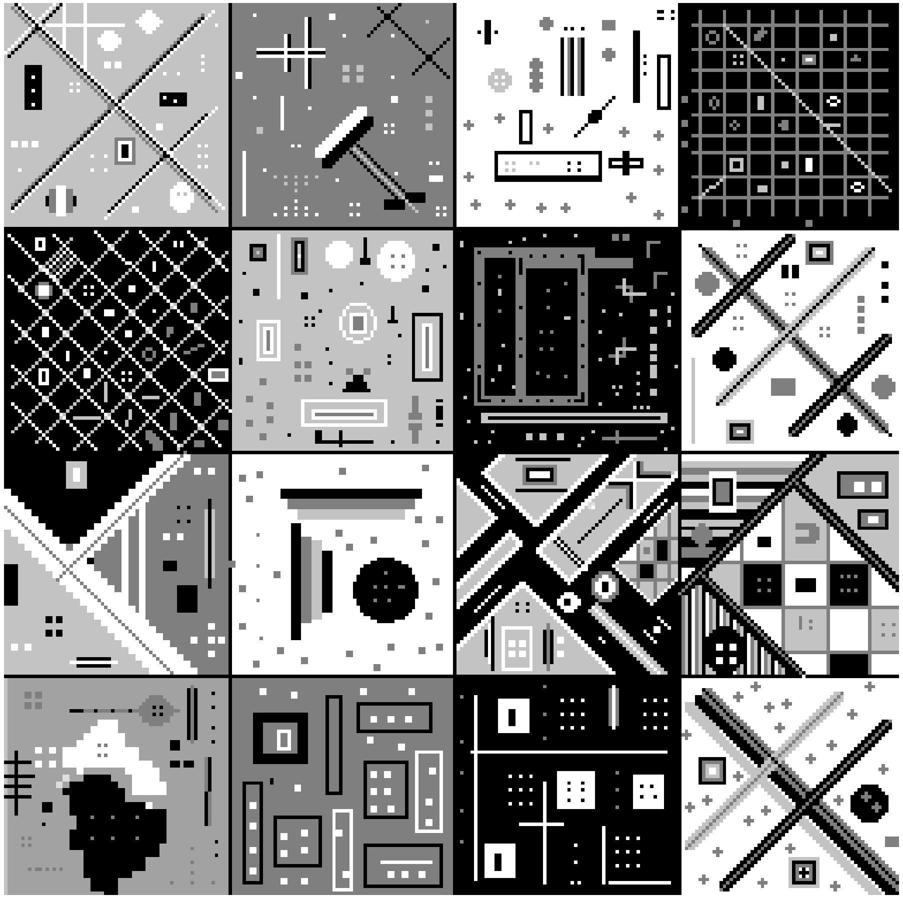

Test Set

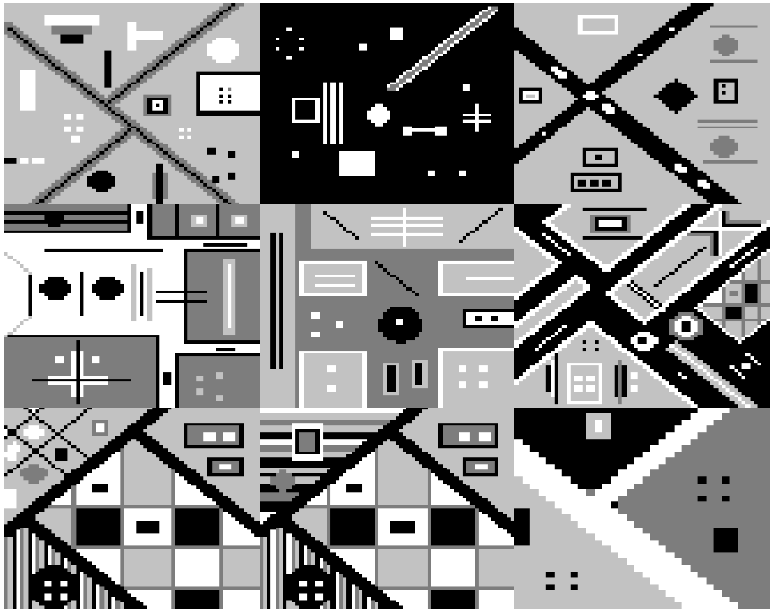

- 5 different phase history data availability levels: 100%, 87.89%, 76.56%, 56.25%, 25%
- 2 different noise levels ($\sigma_n \in \{0.1, 1\}\sigma_y$)
- Rectangular band-limitation for data reduction

Qualitative Results: Image 7, $\sigma_n = 0.1\sigma_y$ and **full data**

Ground truth

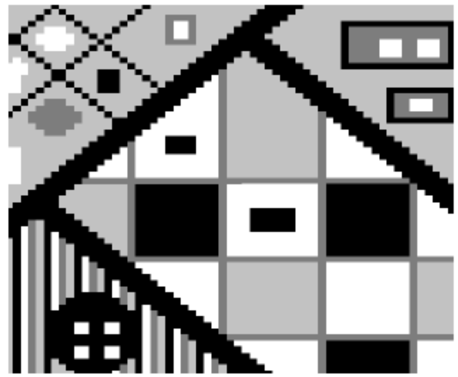

FFT-based

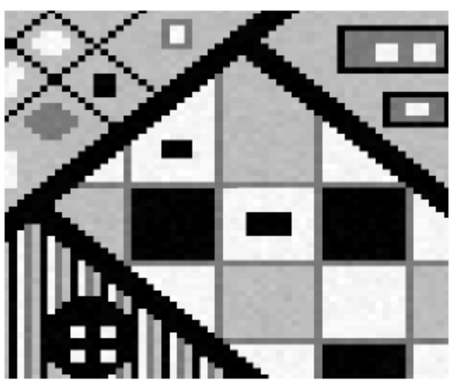

FE-based [3]

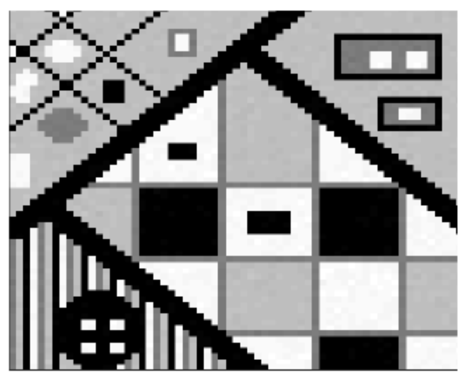

PnP-based

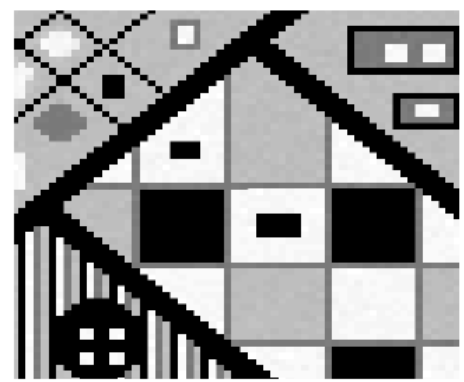

Qualitative Results: Image 7, $\sigma_n = 0.1\sigma_y$ and **87.89% data**

Ground truth

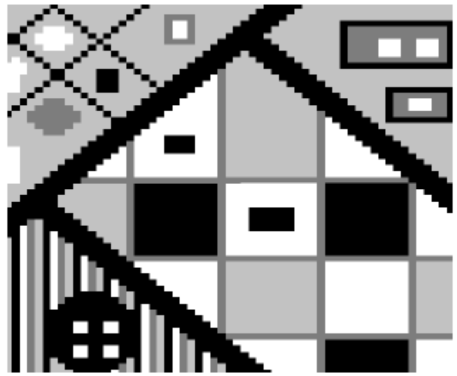

FFT-based

FE-based [3]

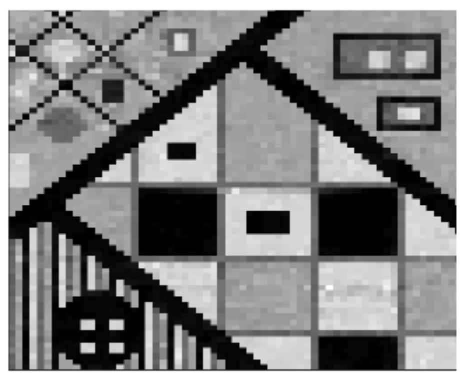

PnP-based

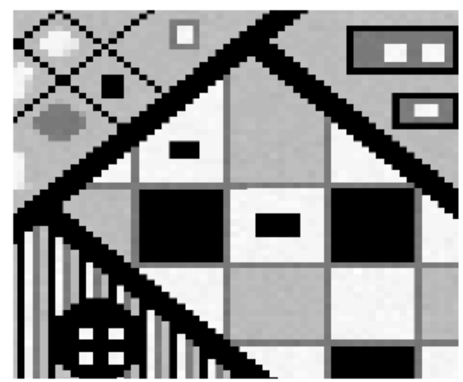

Qualitative Results: Image 7, $\sigma_n = 0.1\sigma_y$ and **76.56% data**

Ground truth

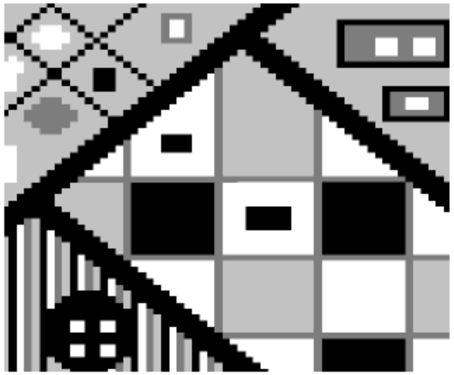

FFT-based

FE-based [3]

PnP-based

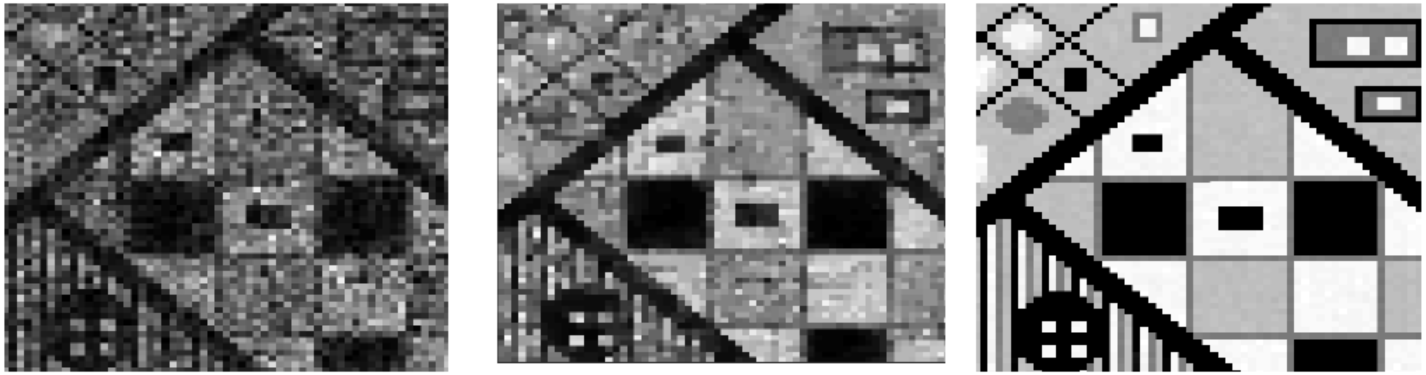

Qualitative Results: Image 7, $\sigma_n = 0.1\sigma_y$ and **56.25% data**

Ground truth

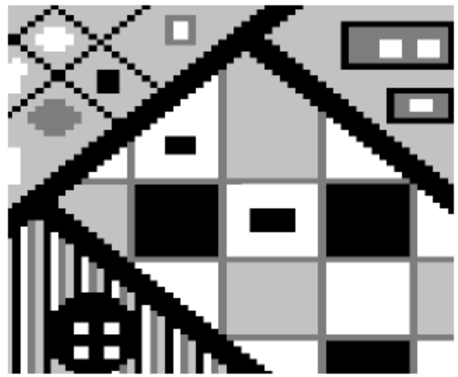

FFT-based

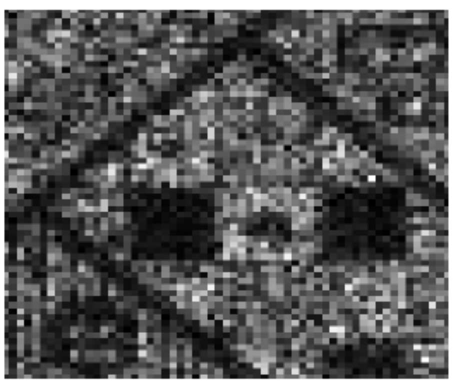

FE-based [3]

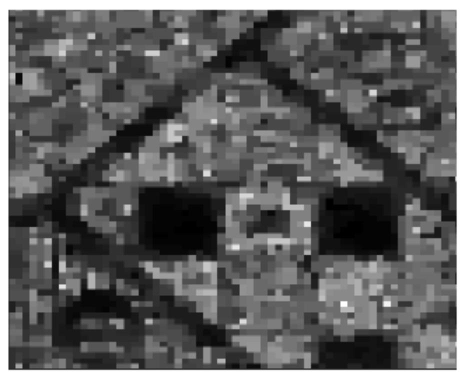

PnP-based

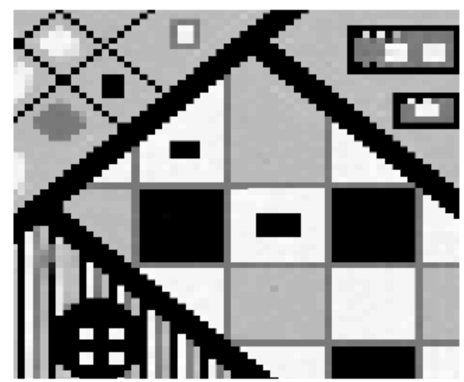

Qualitative Results: Image 7, $\sigma_n = 0.1\sigma_y$ and **25% data**

Ground truth

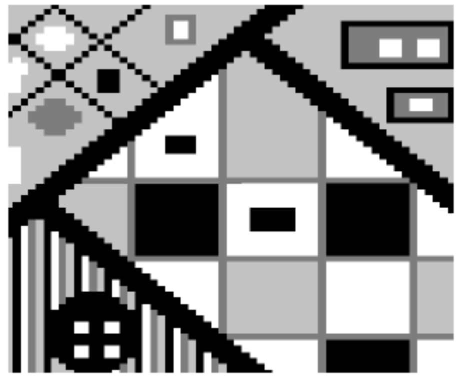

FFT-based

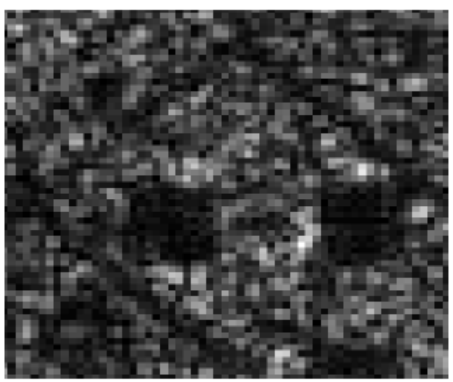

FE-based [3]

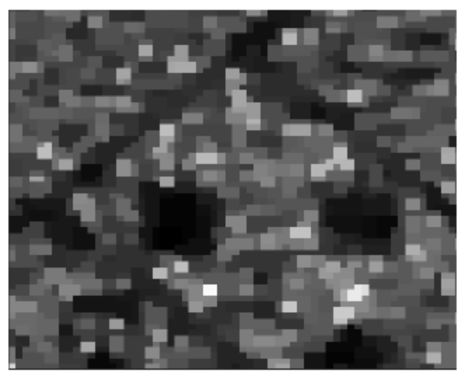

PnP-based

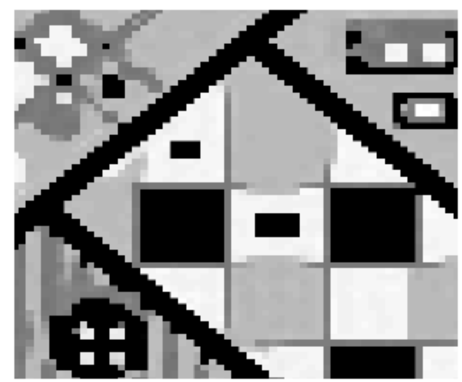

Quantitative Results

Table: SNRs for selected images at various noise & data availability levels

Available Data	Method	SNR (dB)	
		$0.1\,\sigma_y$, Image 7	σ_y, Image 2
100%	FFT-based Reconstruction	36.539	14.980
	Feature-enhanced Regularization [3]	36.968	15.069
	Proposed Framework	**38.075**	**23.235**
87.89%	FFT-based Reconstruction	10.960	8.883
	Feature-enhanced Regularization [3]	22.484	13.502
	Proposed Framework	**36.754**	**22.460**
76.56%	FFT-based Reconstruction	8.659	6.878
	Feature-enhanced Regularization [3]	12.402	10.518
	Proposed Framework	**35.465**	**18.730**
56.25%	FFT-based Reconstruction	6.393	4.432
	Feature-enhanced Regularization [3]	7.886	5.568
	Proposed Framework	**25.199**	**9.604**
25%	FFT-based Reconstruction	3.951	2.481
	Feature-enhanced Regularization [3]	4.692	2.770
	Proposed Framework	**13.579**	**2.955**

Preliminary Results on Real Scenes from TerraSAR-X

- Training based on the Netherlands *Rotterdam Harbor Staring Spotlight SAR image* (1041 × 1830)
  - Split into 448 non-overlapping 64 × 64 "windows"
  - 1075648 overlapping 16 × 16 patches extracted from windows
  - Patches augmented with rotations of 90°, 180°, 270°
- Test set: 751 selected windows extracted from the *Panama High Resolution Spotlight SAR image* (2375 × 3375)

Qualitative Results: Image 608, $\sigma_n = \sigma_y$, **full data**

Reference image

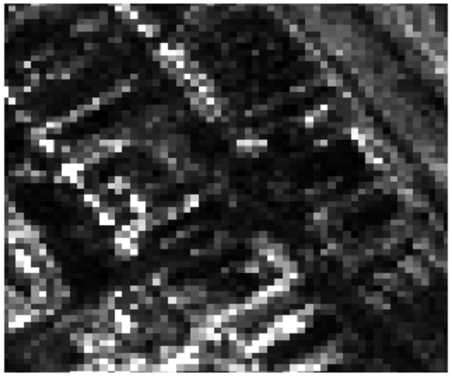

FFT-based

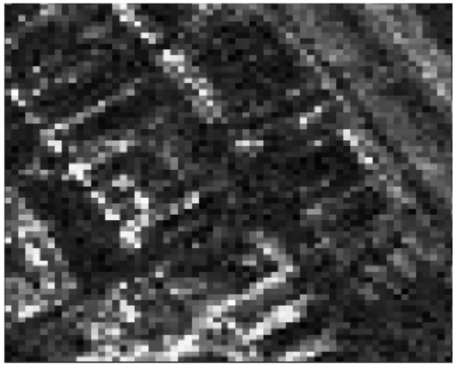

FE-based [3]

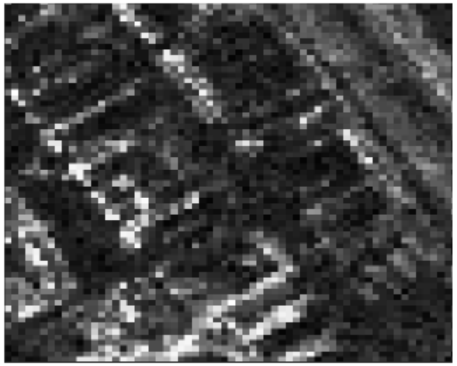

PnP-based

Qualitative Results: Image 608, $\sigma_n = \sigma_y$, **87.89% data**

Reference image

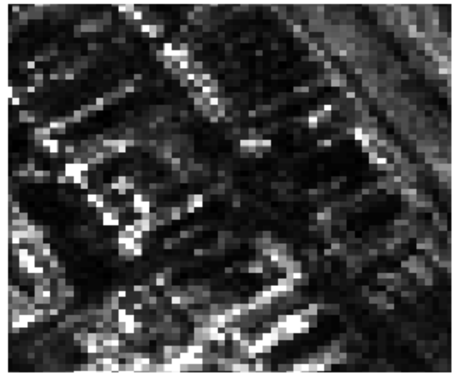

FFT-based

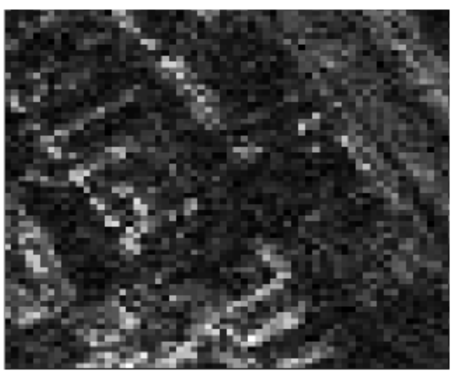

FE-based [3]

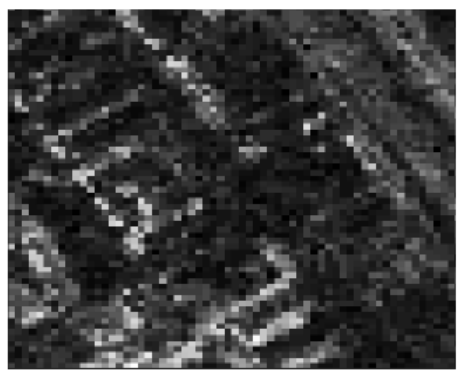

PnP-based

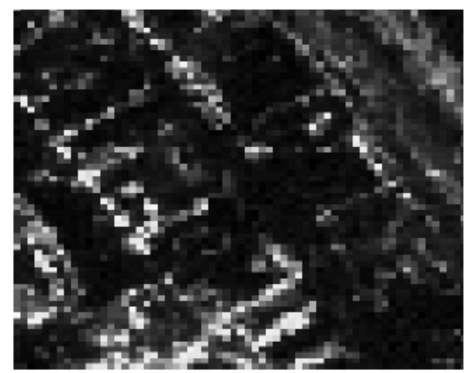

Qualitative Results: Image 608, $\sigma_n = \sigma_y$, **76.56% data**

Reference image

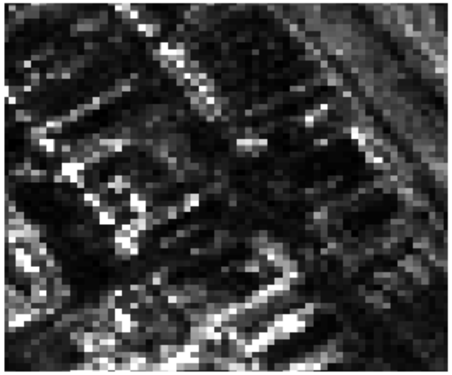

FFT-based

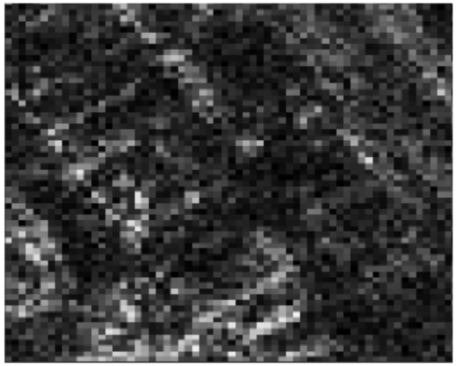

FE-based [3]

PnP-based

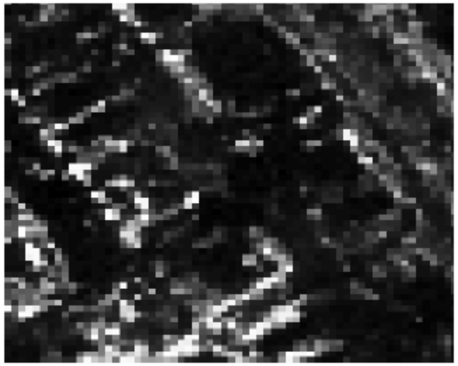

Qualitative Results: Image 608, $\sigma_n = \sigma_y$, **56.25% data**

Reference image

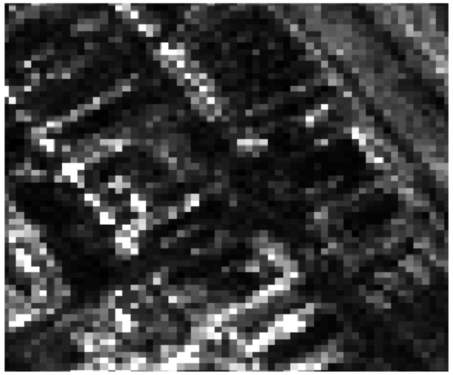

FFT-based

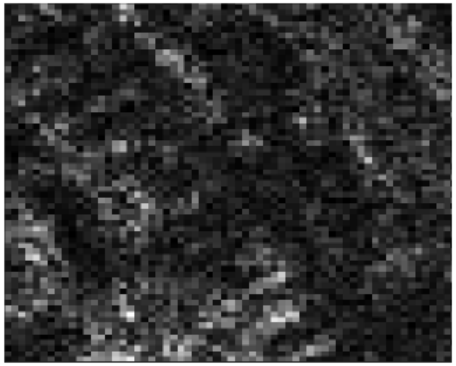

FE-based [3]

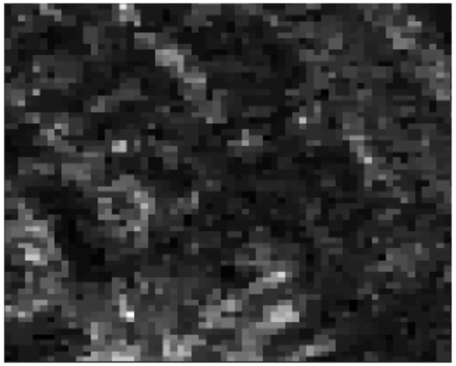

PnP-based

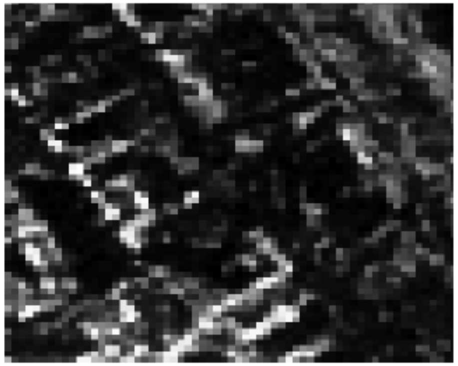

Qualitative Results: Image 608, $\sigma_n = \sigma_y$, **25% data**

Reference image

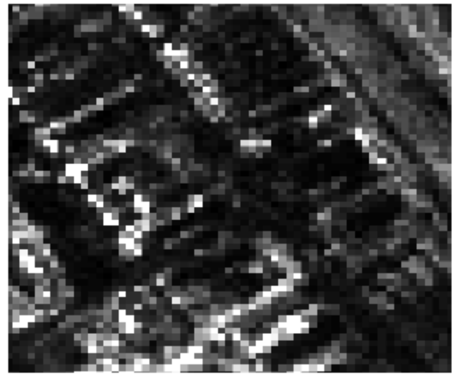

FFT-based

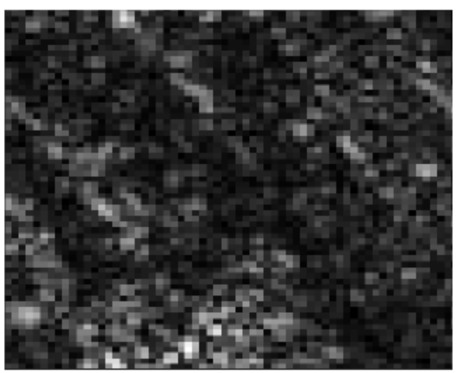

FE-based [3]

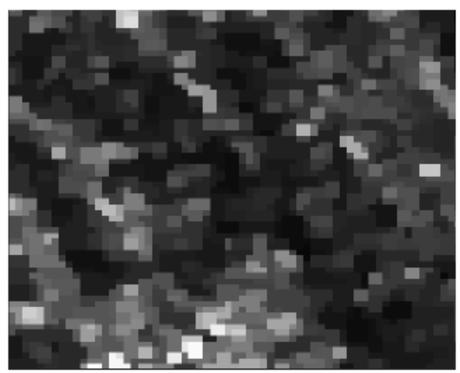

PnP-based

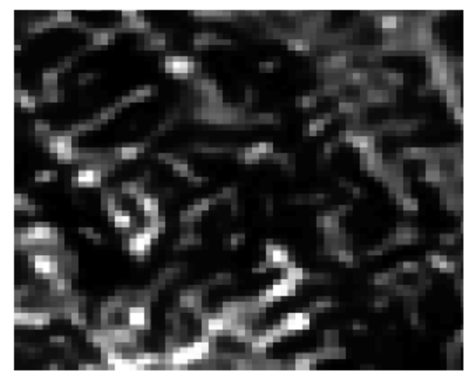

Conclusion

- A line of inquiry that lies at the intersection of several domains:
  - Radar sensing
  - Computational imaging
  - Signal representation, compressed sensing
  - Machine learning
- Sparsity is a useful asset for radar imaging especially in nonconventional data collection scenarios (e.g., when the data are sparse, irregular, limited)
- Deep learning methods may have the potential to learn complicated spatial patterns and enable their incorporation as priors into computational radar imaging

IEEE Computational Imaging Technical Committee

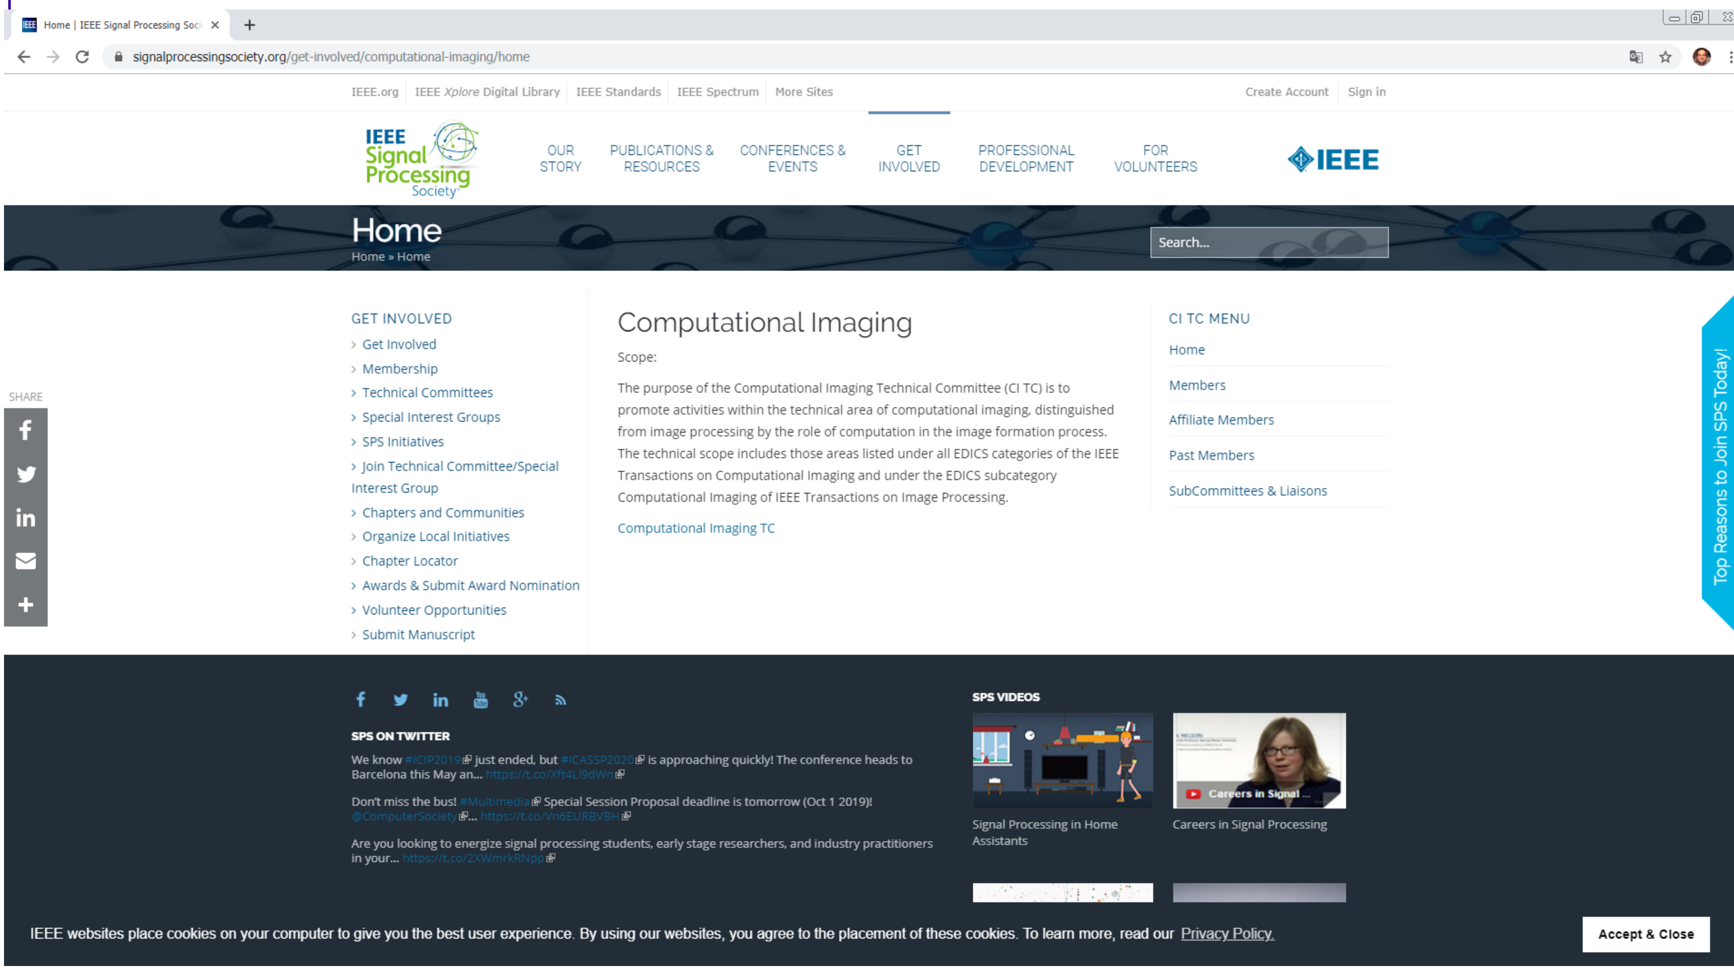

IEEE Transactions on Computational Imaging

About TCI | IEEE Signal Processing

signalprocessingsociety.org/publications-resources/ieee-transactions-computational-imaging/about-tci

IEEE.org | IEEE *Xplore* Digital Library | IEEE Standards | IEEE Spectrum | More Sites | Create Account | Sign in

IEEE Signal Processing Society

OUR STORY | PUBLICATIONS & RESOURCES | CONFERENCES & EVENTS | GET INVOLVED | PROFESSIONAL DEVELOPMENT | FOR VOLUNTEERS

IEEE Transactions on Computational Imaging

Home » About TCI

Search...

TCI MENU

- TCI Home
- About TCI
- Technical Liaison Committee/Editorial Board
- EDICS
- View on IEEE Xplore
  - Top Accessed Articles
  - Early Access

Sign up for TOC Alerts

FOR AUTHORS

- Information for Authors
- Special Issue Deadlines
- IEEE Author Center
- IEEE Copyright Form
- LaTeX and MS Word Templates

Submit a Manuscript

SHARE

IEEE TRANSACTIONS ON

COMPUTATIONAL IMAGING

About TCI

Special Announcement

We are pleased to announce that, as of October 2018, the IEEE Transactions on Computational Imaging has formally been accepted for indexing by the Clarivate Analytics Web of Science.

Articles published in TCI as of March 2016 will be covered in the following Clarivate Analytics products:

- Science Citation Index Expanded (aka SciSearch®)
- Journal Citation Reports/Science Edition
- Current Contents®/Engineering Computing and Technology

The initial **impact factor** for TCI, released by Clarivate in June 2019, is 4.546.

To access indexing data or learn more about the indexing platform, please visit the Clarivate Analytics Web of Science website, or contact the IEEE TCI administrator.

Scope

The IEEE Transactions on Computational Imaging will publish articles where computation plays an integral role in the image formation process. Papers will cover all areas of computational imaging ranging from fundamental theoretical methods to the latest innovative computational imaging system designs. Topics of interest will include advanced algorithms and mathematical techniques, model-based data inversion, methods for image and signal recovery from sparse and incomplete data, techniques for non-traditional sensing of image data, methods for dynamic information acquisition and extraction from imaging sensors,

TOC

Top Reasons to Join SPS Today!