Sparsity Driven Latent Space Sampling for Generative Prior based Compressive Sensing

Vinayak Killedar, Praveen Kumar Pokala and
Chandra Sekhar Seelamantula

Department of Electrical Engineering
Indian Institute of Science, Bangalore - 560012, India

International Conference on Acoustics, Speech and Signal Processing
8 June, 2021

SPECTRUM LAB

Overview

Linear Inverse Problem

Compressive Sensing:

Linear inverse problem with the measurement model:

$$\boldsymbol{y} = \mathrm{A}\boldsymbol{x}, \tag{1}$$

where $\mathrm{A} \in \mathbb{R}^{m \times n}, \boldsymbol{x} \in \mathbb{R}^{n}, \boldsymbol{y} \in \mathbb{R}^{m}$ and $m \ll n$.

$$\mathrm{P0}: \quad \min_{\boldsymbol{x}} \quad \|\boldsymbol{x}\|_0, \quad \text{s.t.} \quad \boldsymbol{y} = \mathrm{A}\boldsymbol{x} \tag{2}$$

Compressive Sensing using Generative Models (CSGM)

Recover $\boldsymbol{x} \in \mathbb{R}^n$ from $\boldsymbol{y} \in \mathbb{R}^m$, such that $m \ll n$. The measurement model is given as:

$$\boldsymbol{y} = \mathrm{A}\boldsymbol{x}, \quad \boldsymbol{x} = \mathrm{G}_{\boldsymbol{\theta}}(\boldsymbol{z}), \tag{3}$$

where $\mathrm{A} \in \mathbb{R}^{m \times n}$ satisfies Set - Restricted Eigenvalue Condition (S-REC), $\mathrm{G}_{\boldsymbol{\theta}}$ is the generator model with latent input $\boldsymbol{z} \in \mathbb{R}^k$ and parameter $\boldsymbol{\theta}$.

Definition: $\mathcal{S}$-REC

[Bora *et al.*, ICML, 2017] Let $\mathcal{S} \subseteq \mathbb{R}^n$. For some parameters $\gamma > 0, \delta \geq 0$, a matrix $\mathrm{A} \in \mathbb{R}^{m \times n}$ is said to satisfy the $\mathcal{S}$-REC$(\mathcal{S}, \gamma, \delta)$ if $\forall \boldsymbol{x}_1, \boldsymbol{x}_2 \in \mathcal{S}$,

$$\|\mathrm{A}(\boldsymbol{x}_1 - \boldsymbol{x}_2)\| \geq \gamma \|\boldsymbol{x}_1 - \boldsymbol{x}_2\| - \delta \tag{4}$$

Motivation

- Natural images do not lie on a single connected manifold, but rather on a union-of-submanifolds.
- A single generator cannot correctly model a distribution that lies on disconnected manifolds [Khayatkhoei *et al.*, NeuRIPS, 2018].
- Disconnected latent space model is considered for data lying on a disconnected manifold.

Contribution

- We propose a union-of-submanifolds to model the true data distribution.
- An optimization framework, namely, proximal meta-learning (PML) algorithm to promote sparsity into the latent variable.
- Sample complexity bounds for the proposed union-of-submanifolds model.
- At higher compression ratio, the performance of SDLSS is superior to state-of-art deep compressed sensing (DCS) [Wu *et al.*, ICML, 2019]

Union-of-Submanifolds

- The latent variable $\boldsymbol{z} \in \mathbb{R}^k$ is assumed to s-sparse.
- Sparsity assumption on the input latent space $\mathbb{Z}$ divides the latent space into $\binom{k}{s}$ subspaces $\mathcal{W}_i$.
- The generator model $\mathrm{G}_{\boldsymbol{\theta}}$ transforms each subspace $\mathcal{W}_i$ to sub-manifold $\mathcal{S}_i$.

Union-of-Submanifolds

The domain of reconstruction is a union-of-submanifolds:

$$\mathcal{S}_{s,\mathrm{G}_{\boldsymbol{\theta}}} = \bigcup_i \mathcal{S}_i, \tag{5}$$

where $\mathcal{S}_i$ is the submanifold generated by $\mathrm{G}_{\boldsymbol{\theta}}(\mathrm{z})$, with $\|\mathrm{z}\|_0 \leq s$.

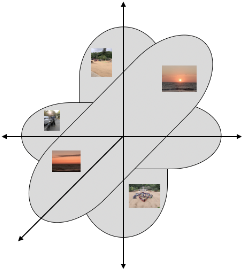

Figure: Union-of-Submanifolds

Sparsity Driven Latent Space Sampling (SDLSS)

Measurement model:

$$\boldsymbol{y} = \mathrm{A}\boldsymbol{x} \quad \text{s.t} \quad \boldsymbol{x} = \mathrm{G}_{\boldsymbol{\theta}}(\boldsymbol{z}), \quad \|\boldsymbol{z}\|_0 \leq s. \tag{6}$$

$\boldsymbol{x}$ is assumed to lie in a union of sub-manifolds $\mathcal{S}_{s,\mathrm{G}_{\boldsymbol{\theta}}}$.

Optimization problem:

$$\min_{\boldsymbol{z},\boldsymbol{\theta}} \|\boldsymbol{z}\|_0 \quad \text{s.t.} \quad \|\boldsymbol{y} - \mathrm{A}\mathrm{G}_{\boldsymbol{\theta}}(\boldsymbol{z})\|_2 \leq \epsilon, \tag{7}$$

with A satisfying set-restricted eigenvalue condition (S-REC) and $\boldsymbol{z} \in \mathbb{R}^k$ is the latent space with sparsity at most s.

Sparsity Driven Latent Space Samping (SDLSS)

Optimization cost:

$$\min_{z,\boldsymbol{\theta}} \left(\mathcal{L}_{\mathrm{G}} + \mathcal{L}_{\mathrm{A}}\right), \tag{8}$$

where

$$\mathcal{L}_{\mathrm{G}} = \mathbb{E}_{\boldsymbol{z}}\{E_{\boldsymbol{\theta}}(\boldsymbol{y},\boldsymbol{z}) = \|\boldsymbol{y} - \mathrm{A}\mathrm{G}_{\boldsymbol{\theta}}(\boldsymbol{z})\|_2^2 + \|\boldsymbol{z}\|_0\}$$
$$\mathcal{L}_{\mathrm{A}} = \mathbb{E}_{\boldsymbol{x}_1,\boldsymbol{x}_2}\{(\|\mathrm{A}\boldsymbol{x}_1 - \boldsymbol{x}_2\|_2 + \delta - \gamma\|\boldsymbol{x}_1 - \boldsymbol{x}_2\|_2)^2\}$$

$\mathcal{L}_{\mathrm{G}}$ and $\mathcal{L}_{\mathrm{A}}$ represent the measurement loss and $\mathcal{S} - REC$ loss, respectively.

Enforcing Sparsity

- The proximal step is introduced to enforce sparsity in the latent space via meta-learning [Finn *et al.*, ICML, 2017].

$$E_\theta(\boldsymbol{y}_i, \boldsymbol{z}_i) = \|\boldsymbol{y}_i - \mathrm{AG}_{\boldsymbol{\theta}}(\boldsymbol{z}_i)\|_2^2 + \|\boldsymbol{z}_i\|_0, \tag{9}$$

$$\begin{aligned}\hat{\boldsymbol{z}}_i &= \arg\min_{\boldsymbol{z}_i} E_\theta(\boldsymbol{y}_i, \boldsymbol{z}_i), \\ &= \mathrm{P}_s\left(\boldsymbol{z}_i - \beta \nabla_{\boldsymbol{z}} f(\boldsymbol{y}_i, \boldsymbol{z}_i)\right),\end{aligned}$$

where $f(\boldsymbol{y}_i, \boldsymbol{z}_i) = \|\boldsymbol{y}_i - \mathrm{AG}_{\boldsymbol{\theta}}(\boldsymbol{z}_i)\|_2^2$, β is the learning rate, and $\mathrm{P}_s(\mathrm{u})$ is the hard-thesholding operator that sets all but the largest (in magnitude) s elements of u to 0.

SDLSS Algorithm with PML

Algorithm 1 :Sparsity Driven Latent Sampling (SDLSS) for Generative Prior

Input: data $= \{\boldsymbol{x}_i\}_{i=1}^{N}$, sensing matrix A, generator $\mathrm{G}_{\boldsymbol{\theta}}$, learning rate α, number of latent optimization steps T, measurement error threshold ϵ and sparsity factor s.

repeat

  Initialize generator network parameter $\boldsymbol{\theta}$.

  for $i = 1$ **to** N **do**

    Measure the signal $\mathrm{y}_i = \mathrm{A}\mathrm{x}_i$

    Sample $\mathrm{z} \sim \mathcal{N}(0, I)$

    for $t = 1$ **to** T **do**

      $\hat{\mathrm{z}}_i = \mathrm{P}_s\left(\mathrm{z}_i - \beta \nabla_{\mathrm{z}} f(\mathrm{y}_i, \mathrm{z}_i)\right)$

    end for

  end for

  $\mathcal{L} = \mathcal{L}_{\mathrm{G}} + \mathcal{L}_{\mathrm{A}}$

  Update $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} - \alpha \frac{\partial \mathcal{L}}{\partial \boldsymbol{\theta}}$

until $\|\boldsymbol{y} - \mathrm{A}\mathrm{G}_{\boldsymbol{\theta}}(\hat{\boldsymbol{z}})\|_2 \leq \epsilon$

Theorem

Let $\mathrm{G}_{\boldsymbol{\theta}} : \mathbb{R}^k \rightarrow \mathbb{R}^n$ be a generative model with d-layers with ReLU activation and at most h neurons in each layer. Let $\mathrm{A} \in \mathbb{R}^{m \times n}$ be a random Gaussian matrix with IID entries such that $\mathrm{A}_{i,j} \sim \mathcal{N}(0, \frac{1}{m})$ and satisfies $\mathcal{S}$-REC$(\mathcal{S}_{s,\mathrm{G}_{\boldsymbol{\theta}}}, (1-\alpha), 0)$. For any $\boldsymbol{x} \in \mathbb{R}^n$, if the number of measurements given by $\boldsymbol{y} = \mathrm{A}\boldsymbol{x} \in \mathbb{R}^m$ is $\mathcal{O}(sd \log \frac{kh}{s})$, where s is sparsity of the latent variable $\boldsymbol{z}$, then with probability $1 - \exp^{(-\Omega(m))}$:

$$\|\mathrm{G}_{\boldsymbol{\theta}}(\hat{\boldsymbol{z}}) - \boldsymbol{x}\|_2 \leq C \min_{\boldsymbol{z} \in \mathbb{R}^k, \|\boldsymbol{z}\|_0 \leq s} \|\mathrm{G}_{\boldsymbol{\theta}}(\boldsymbol{z}) - \boldsymbol{x}\|_2 + 2\epsilon, \tag{10}$$

where $\hat{\boldsymbol{z}}$ minimizes $\|\boldsymbol{y} - \mathrm{A}\mathrm{G}_{\boldsymbol{\theta}}(\boldsymbol{z})\|_2^2$ to within additive ϵ of the optimum, C is a constant, Ω is the asymptotic lower bound, and $0 < \alpha < 1$.

MNIST:Effect of Sparsity

- Reconstruction error on test data as a function of sparsity; $m = 10$ and latent dimension $k = 784$.

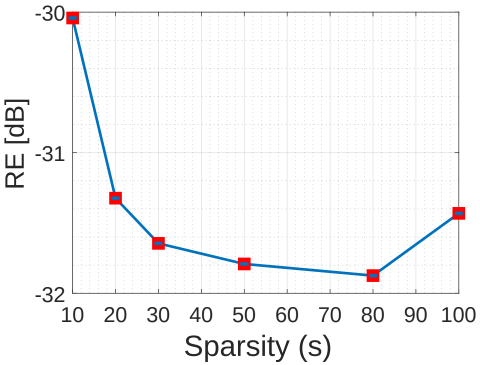

MNIST:PSNR

- PSNR evaluation of SDLSS and DCS as a function of measurements m for the latent dimension $k = 100$ and sparsity $s = 80$.

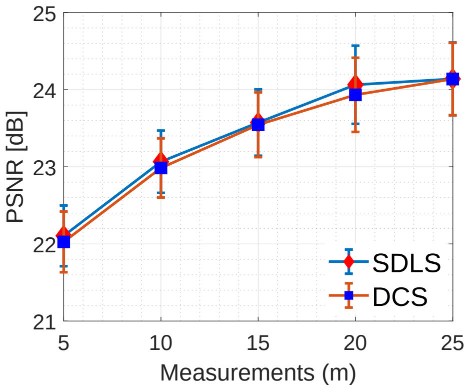

Results:MNIST

Ground truth

DCS
PSNR = 15.4 dB

SDLSS
PSNR = 15.6 dB

- Deep Compressed Sensing (DCS) [Wu *et al.*, ICML, 2019]

Results:MNIST

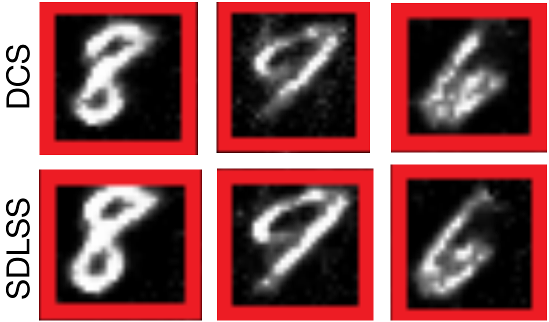

References

A. Bora, A. Jalal, E. Price, and A. G. Dimakis
Compressed Sensing using Generative Models
in Proc. International Conference on Machine Learning, pp. 537–546, 2017.

V. Shah, and C. Hegde
Solving Linear Inverse Problems using GAN Priors: An Algorithm with Provable Guarantees
in Proc. International Conference on Acoustics, Speech, and Signal Processing, pp. 4609–4613, 2018.

P. Hand, O. Leong, and V. Voroninski
Global Guarantees for Enforcing Deep Generative Priors by Empirical Risk
IEEE Transactions on Information Theory, pp. 401-418, 2020.

Y. Wu, M. Rosca, and T. Lillicrap
Deep Compressed Sensing
in Proc. International Conference on Machine Learning, pp. 1349–1353, 2019.

C. Finn, P. Abbeel, and S. Levine
Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks
in Proc. International Conference on Machine Learning, pp. 1126–1135, 2017.

M. Khayatkhoei, M.K. Singh, and A. Elgammal
Disconnected Manifold Learning for Generative Adversarial Networks
in Proc. Advances in Neural Information Processing Systems, pp. 7343–7353, 2018.

Acknowledgment

- Science and Engineering Research Board (SERB), Core Research Grant (CRG) and MATRICS grant.
- Robert Bosch Center for Cyber Physical Systems (RBCCPS).

The End