SOLVING LINEAR INVERSE PROBLEMS USING GAN PRIORS, AN ALGORITHM WITH PROVABLE GUARANTEES

Viraj Shah and Chinmay Hegde
Department of Electrical and Computer Engineering
Iowa State University

Linear Inverse Problems: Introduction

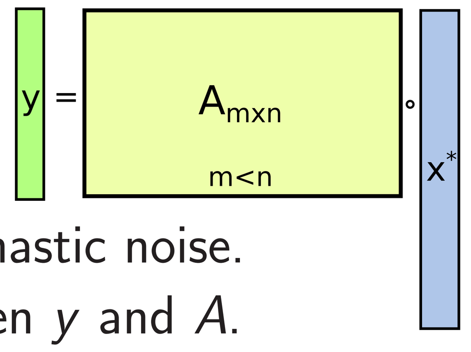

- Problems of the form $\mathbf{y} = \mathbf{A}\mathbf{x}^* + \mathbf{e}$, where,
  - $x^* \in \mathbb{R}^n$ is the target signal or image,
  - $A \in \mathbb{R}^{m \times n}$ is the linear operator,
  - $y \in \mathbb{R}^m$ is measurements, $e \in \mathbb{R}^m$ is stochastic noise.
- **Aim**: To recover the unknown signal x^* given y and A.
- Such problems arise in diverse fields such as computational imaging, optics, and astrophysics.

Linear Inverse Problems in Signal and Image Processing

- **Denoising:** it is the simplest case, with A being identity.
- **Super-resolution:** A represents a low-pass filter+ down-sampling.
- **Image inpainting:** A is pixel-wise selection operator.
- **Compressive sensing:** A is fat random matrix with $m < n$.
- In most cases, $\mathbf{m} < \mathbf{n} \implies$ N(equations) < N(variables) $\implies$ **ill-posed** $\implies$ infinite many solutions are possible, but only few of them are the required 'natural' signals (or images).

Common Solution to Linear Inverse Problems

- Restrict the solution space using a 'natural signals prior' as a constraint, which results in constrained optimization problem:

$$\widehat{x} = \operatorname{argmin}\ f(y; Ax), \quad \text{s. t.} \quad x \in \mathcal{S}, \qquad (1)$$

- $f(\cdot)$ is the loss function, and,
- set $\mathcal{S} \subseteq \mathbb{R}^n$ captures a structure that x is *a priori* assumed to obey.

Sparsity Prior and its Limitations

- $\mathcal{S}$ is defined as the set of sparse vectors; basic assumption is that the 'natural' signals are sparse in some basis. However,
- It suffers from poor discriminatory capacity.
- Not all sparse vectors are 'natural' images.
- Performs very poorly for $m << n$.
- Nature exhibits far richer nonlinear structure than sparsity alone.

Our Approach: GAN Priors with Projected Gradient Descent

- **Learned GAN Prior:** Learns the structure of the 'natural' signals from training data using Generative Adversarial Networks (GAN)[1].
- $G(\cdot)$ maps latent variable $z \in \mathbb{R}^k$ to the ambient signals $x \in \mathbb{R}^n$.
- **Key assumption:** the generator $G(\cdot)$ well-approximates the set $\mathcal{S}$.
- Substituting $x = G(z)$ in Eq.(1), the resulting problem[1] is the optimization in the latent space (over z) $\implies$ can stuck in local minima.
- Thus, we advocate **Projected Gradient Descent** to solve the Eq.(1) directly in ambient space (over x).

Our Algorithm: Projected Gradient Descent on GAN (PGD-GAN)

- We aim to solve for $\widehat{x}$, given y and A:

$$\widehat{x} = \operatorname*{argmin}_{x \in \mathcal{S}} \|y - Ax\|_2^2. \qquad (2)$$
$$\text{with } \mathcal{S} \equiv span(G).$$

- We train the generator $G : \mathbb{R}^k \to \mathbb{R}^n$:

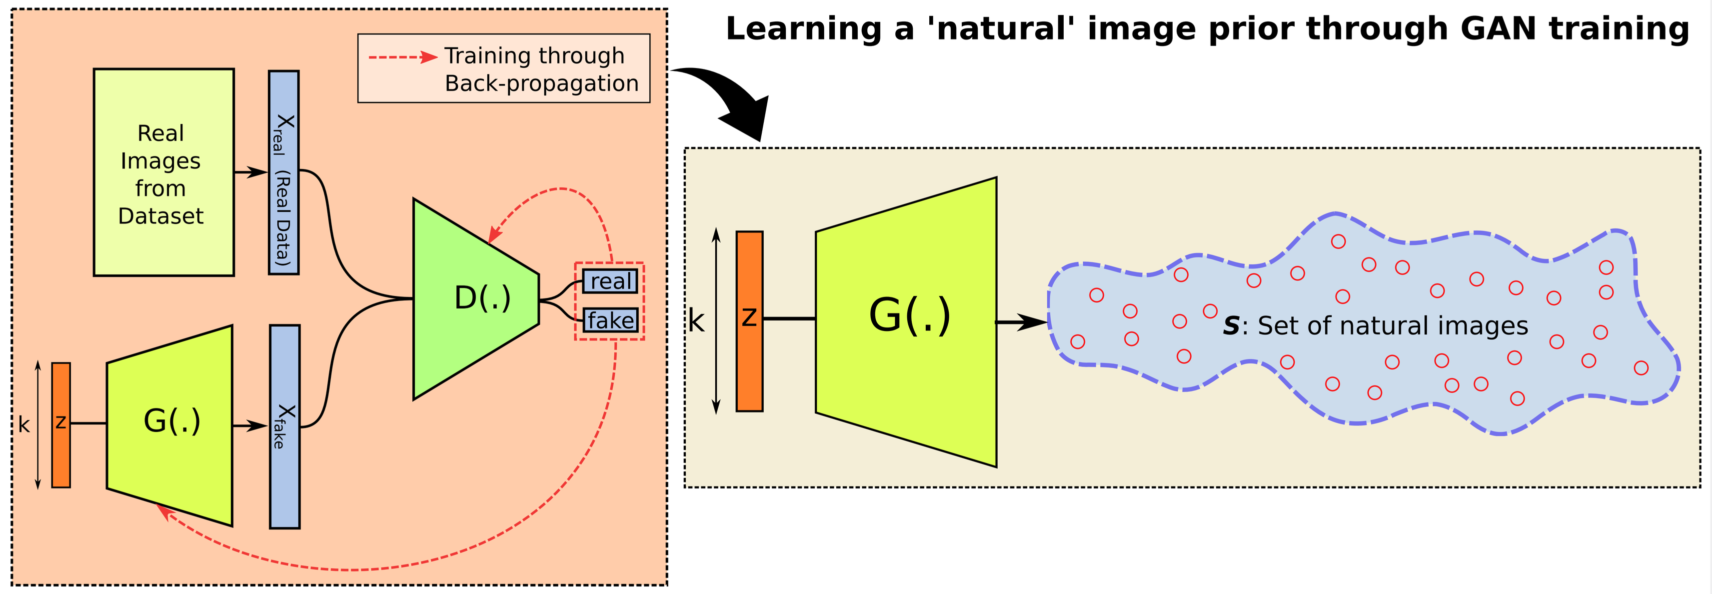

- With pre-trained generator(G), apply Projected Gradient Descent in 2 steps:

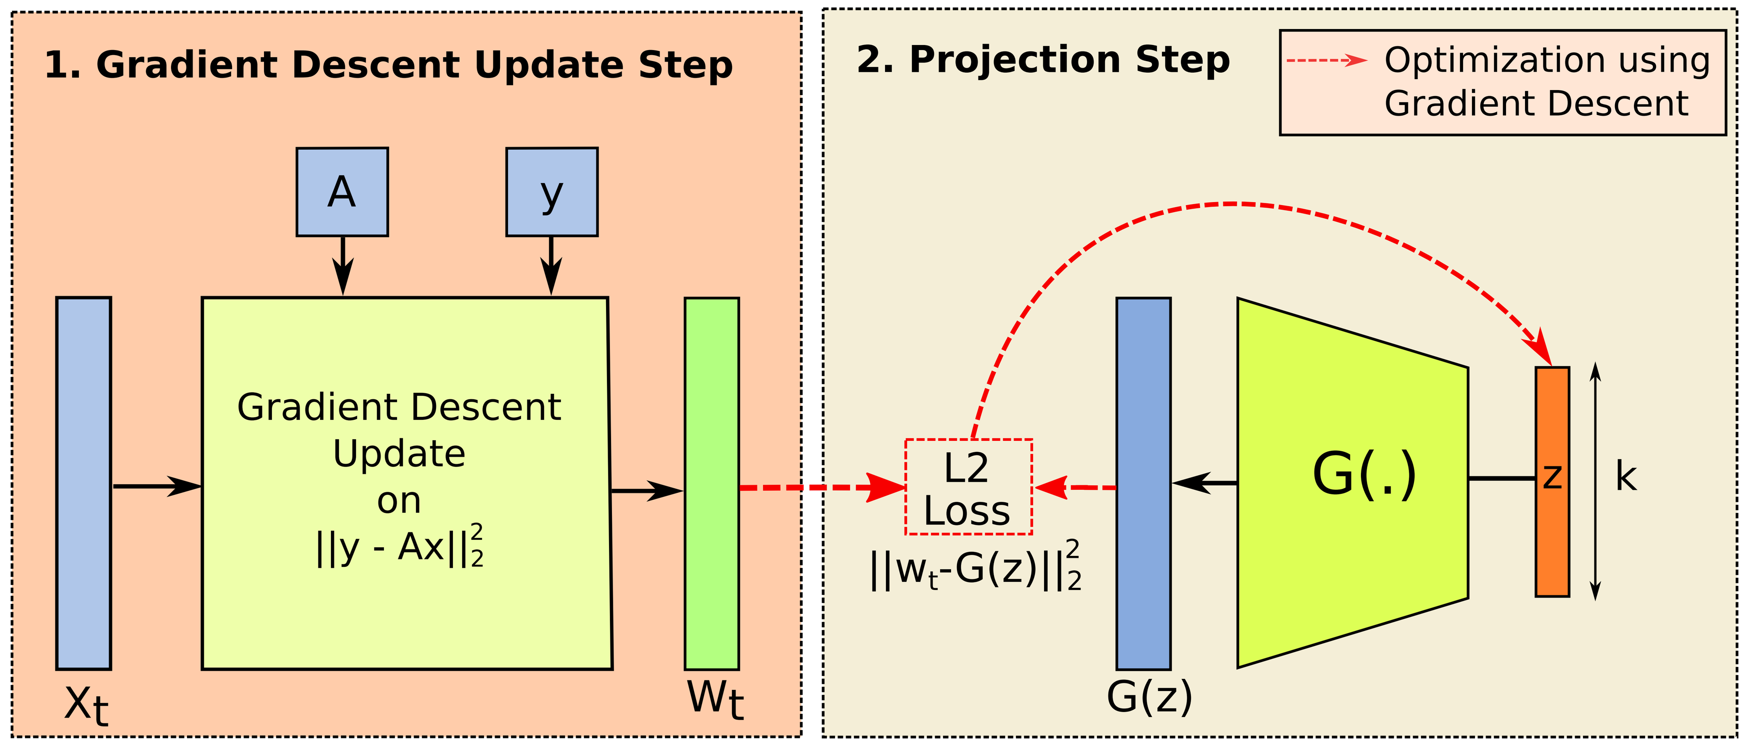

Calculate an estimate w_t using gradient descent update rule on the loss function,

$$f(x) := \|y - Ax\|_2^2.$$

Thus, update at t^{th} iteration is,

$$w_t \leftarrow x_t + \eta A^T(y - Ax_t),$$

Find an image from the span of the generator which is closest to our current estimate w_t.

$$\mathcal{P}_G(w_t) := G\left(\operatorname*{argmin}_z f_{in}(z)\right),$$

where, $f_{in}(z) := \|w_t - G(z)\|_2^2$.

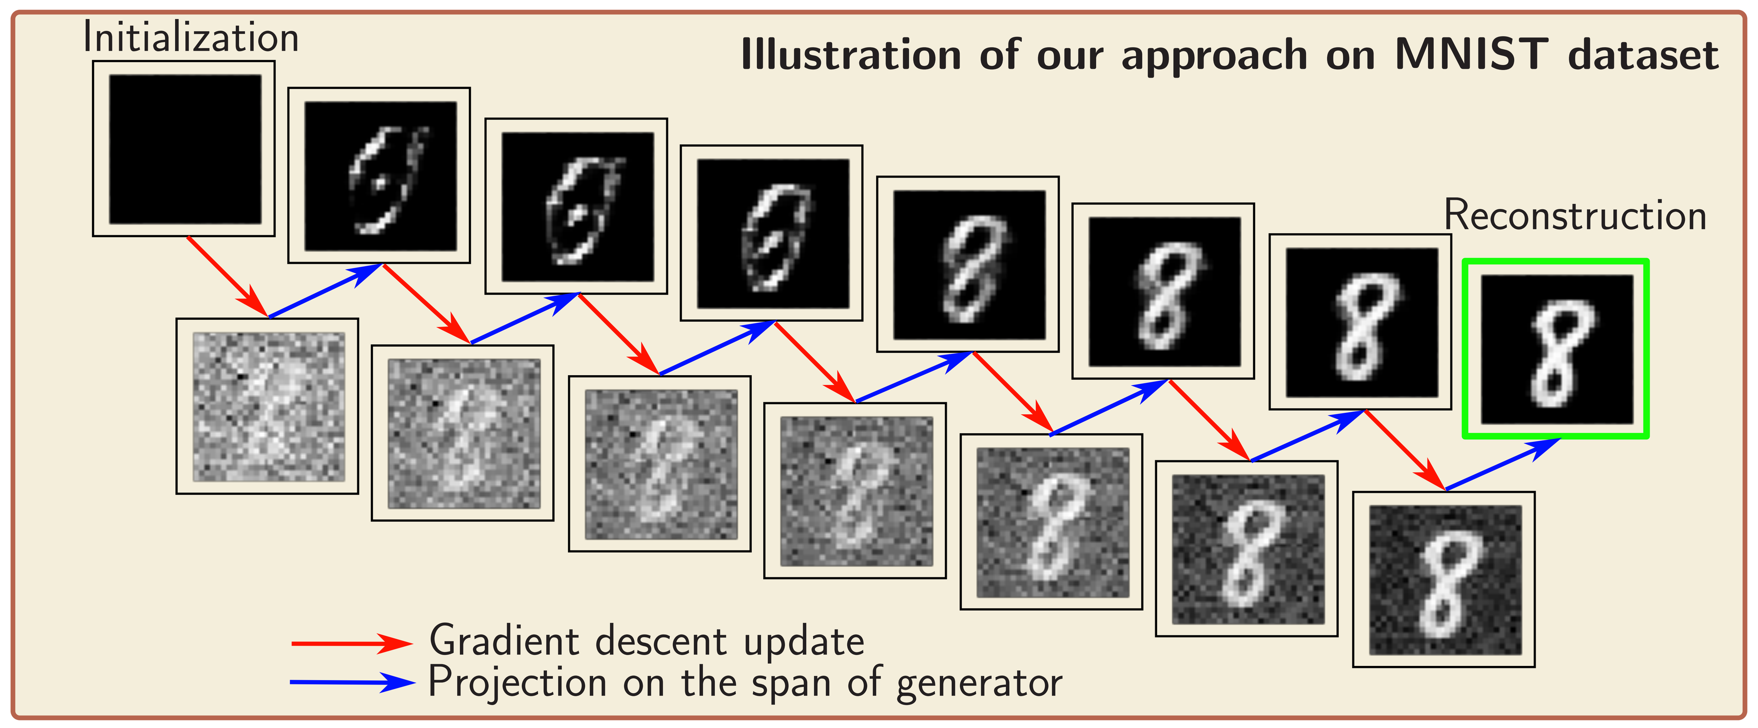

Proof of Convergence for our algorithm

- We need a condition on A to ensure that A preserves the uniqueness of x.
- No sparsity prior $\implies$ Restricted Isometry Property (RIP) can't be used.
- We use slightly modified version of $S-REC$ **Set Restricted Eigenvalue Condition** as defined in [1]:
- **Def. 1:** Let $\mathcal{S} \in \mathbb{R}^n$. $A \in \mathbb{R}^{m \times n}$. For parameters $\gamma > 0$, $\delta \geq 0$, A satisfies the S-REC($\mathcal{S}, \gamma, \delta$) if, $\|A(x_1 - x_2)\|^2 \geq \gamma\|x_1 - x_2\|^2 - \delta$, for $\forall x_1, x_2 \in \mathcal{S}$.
- Gaussian A satisfies the $S-REC$ condition for sufficiently large m [1].
- Given that $S-REC$ is satisfied, we prove that the sequence (x_t) defined by the PGD-GAN with $y = Ax^*$ converges to x^* with high probability.

Experimental Results

on MNIST dataset

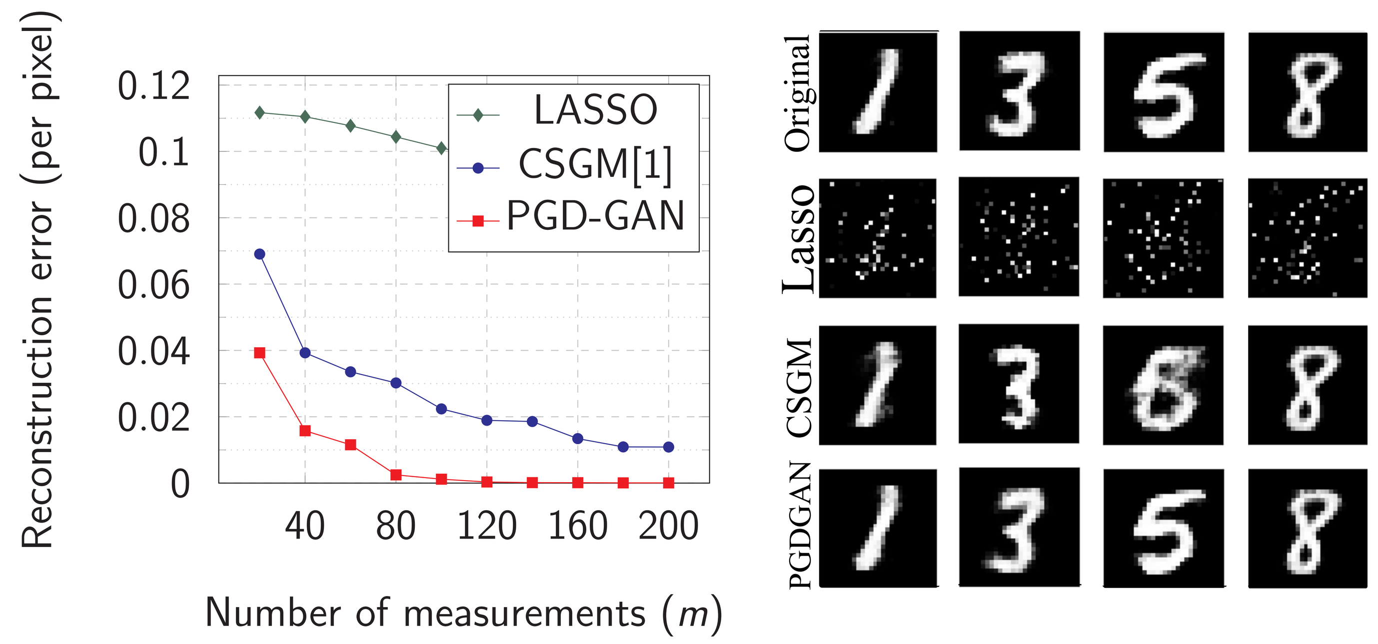

on celebA dataset

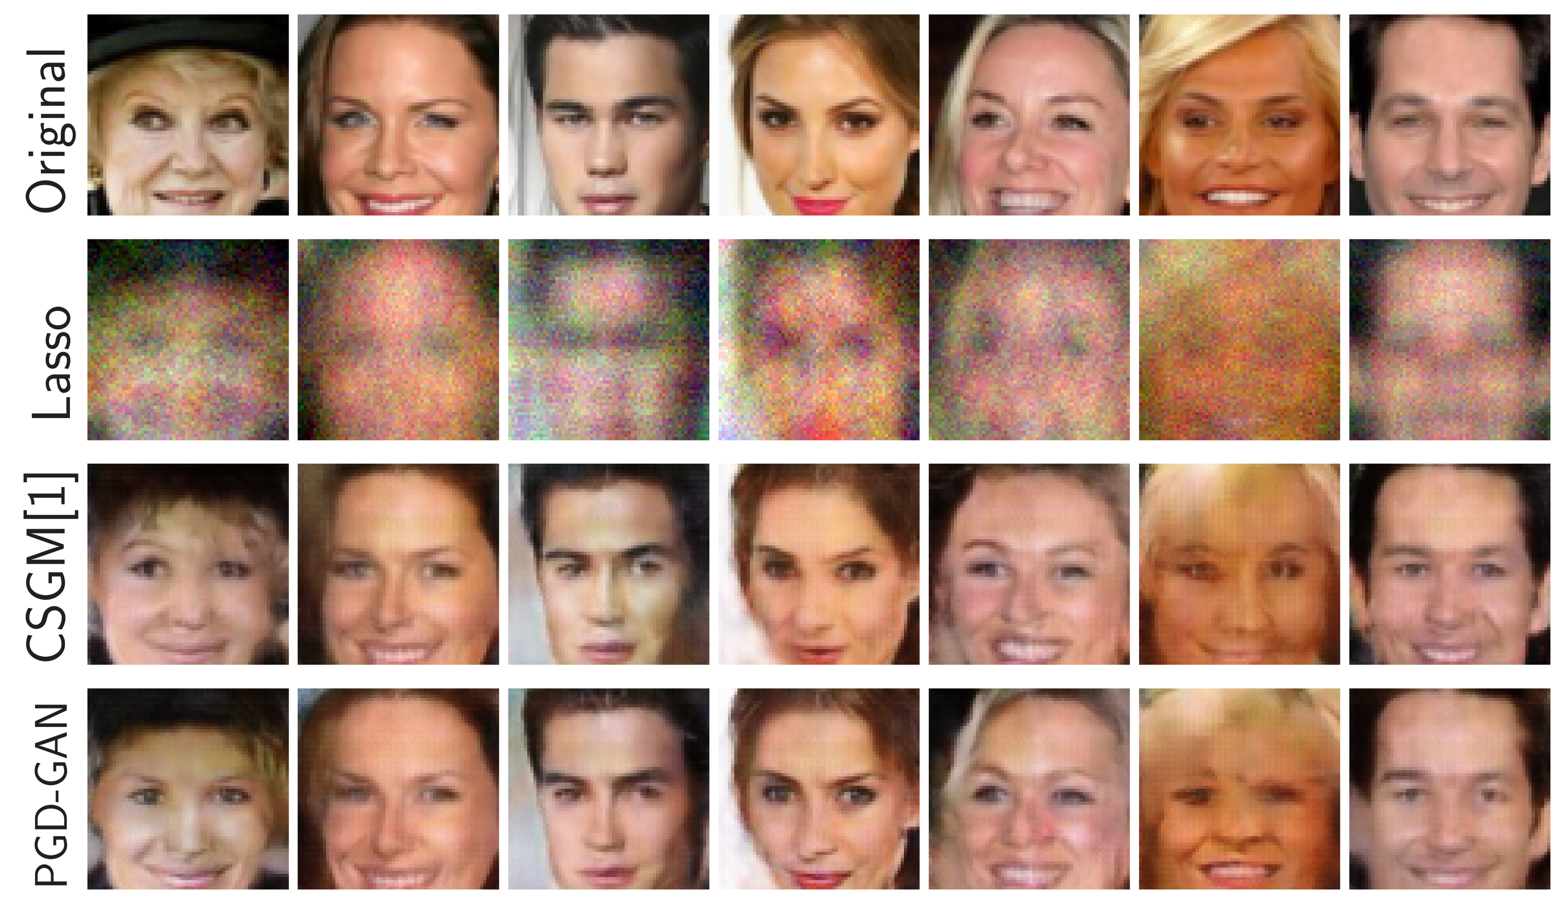

Acknowledgments

This work was supported in part by grants from the National Science Foundation and NVIDIA.

References

[1] A. Bora, A. Jalal, E. Price, and A. Dimakis, "Compressed Sensing using Generative Models," Proc. Int. Conf. Machine Learning (ICML), 2017.