Deep Image Prior - Dmitry Ulyanov

Guillaume De Gani

Motivation

- Inverse Problems are frequently encountered in signal analyses (ex: MRI, Artefact Removal, Super Resolution …).

- Finding Priors is a complicated process two methods exist finding an Explicit Prior or finding a Learned Prior.

- Note this paper doesn’t intend to compete with State Of the Art.

- Showing That the model of the structure can give a strong prior without learning.

State of the art

Classic Approach

- SRResNet (Super Resolution) [1]
- SRGan (Super Resolution) [1]
- ESPCNN (Super Resolution) [2]
- LapSRN (Super Resolution) [3]
- Shepard Network (Inpainting) [4]
- Convolutional sparse coding (Inpainting) [5]
- DnCNN (Denoising) [6]

Mathematical Approach

- BiCubic (Super Resolution)
- CBM3D (Denoising)
- LSSC(Denoising)
- NCSR (Denoising)
- WNNM (Denoising)
- Diffusion Based (Inpainting)

Modeling Invers Tasks

$x^* \longrightarrow$ Restored Image

$x_0 \longrightarrow$ Corrupted Observation

$x \longrightarrow$ Image

Optimizing The Image space :

$$x^* = \min_x E(x; x_0) + R(x)$$

$f_\theta \longrightarrow$ Generative model with θ as parameters

$z \longrightarrow$ Fixed Input

Optimizing The Parameter space :

$$\boxed{\theta^* = arg\min_\theta E(f_\theta(z); x_0)}$$

Denoising

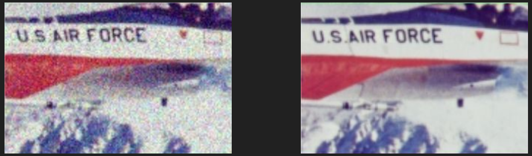

$\Rightarrow \quad ||x - x_0||^2$

Super Resolution

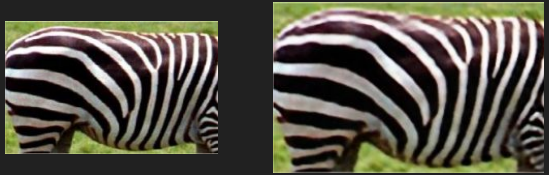

$\Rightarrow \quad ||d(x) - x_0||^2$

Inpainting

$\Rightarrow \quad ||(x - x_0) \odot m||^2$

Step by Step process

Initialization :

$$x_0 \in \mathbf{R}^{C \times H \times W}$$

$$f_\theta : \mathbf{R}^{C' \times H' \times W'} \rightarrow \mathbf{R}^{C \times H \times W}$$

$$z \in \mathbf{R}^{C' \times H' \times W'} \sim U(a, b)$$

Problem :

$$\min_\theta ||f_\theta(z) - x_0||^2$$

Method :

$$\theta_{k+1} := \theta_k - \alpha \frac{\delta ||f_\theta(z) - x_0||^2}{\delta\theta}$$

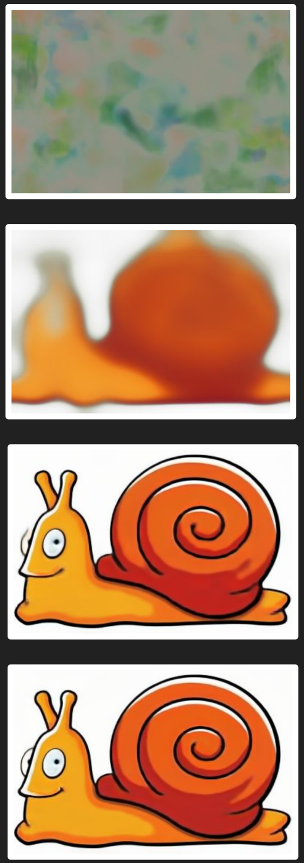

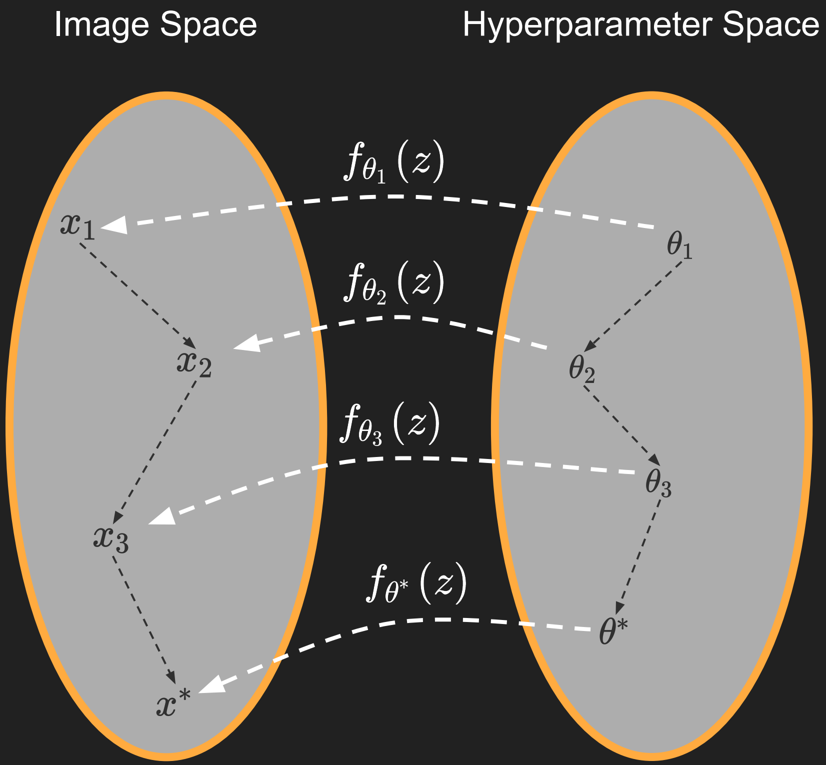

Overfitting And Convergence

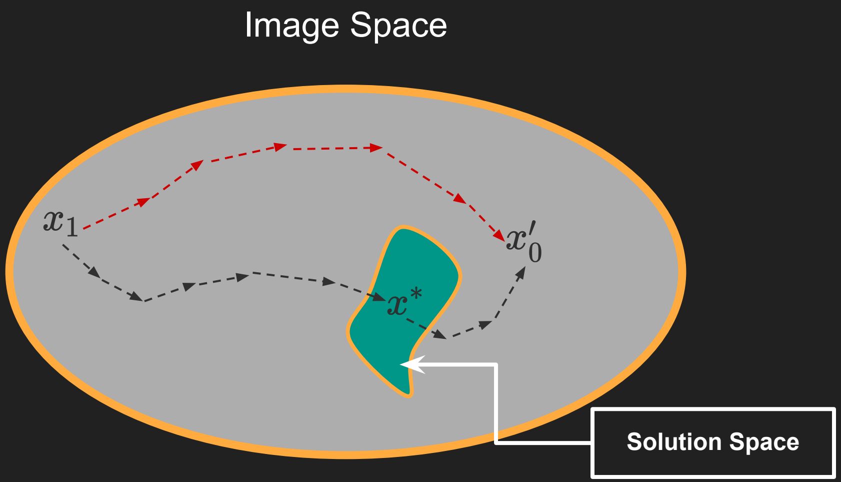

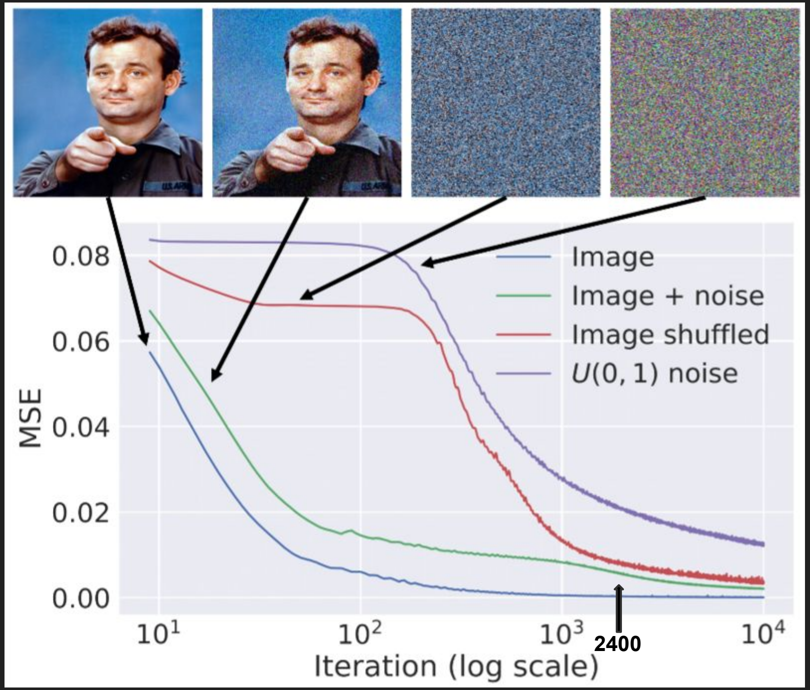

Two path possible:

- Red path overfits without passing thru the solution space
- Black path reaches and optimal solution x* before overfitting

Implementation

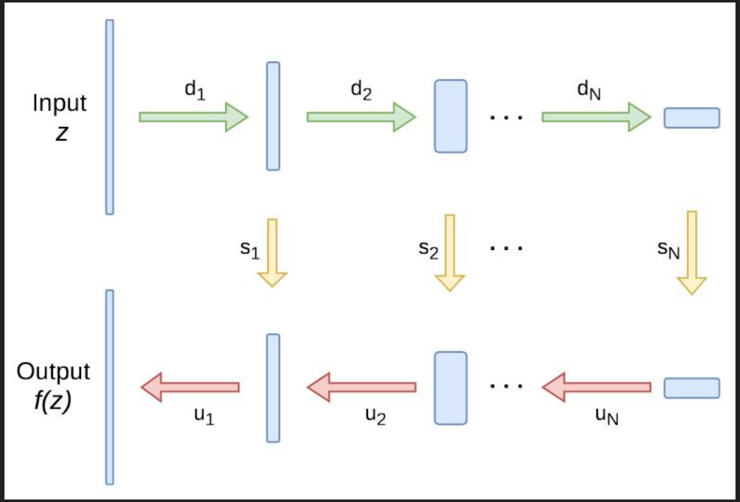

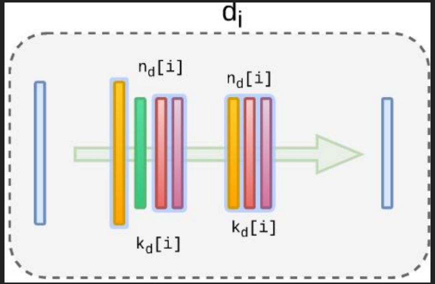

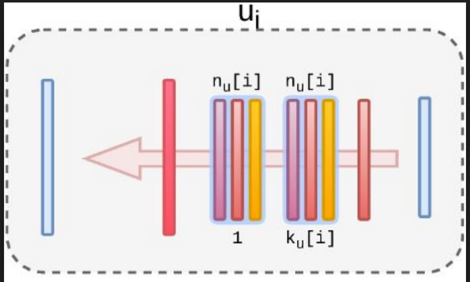

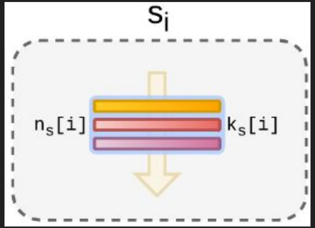

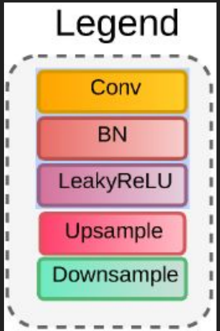

$z \in \mathbb{R}^{3 \times W \times H} \sim U(0, \frac{1}{10})$

$n_u = n_d = [8,\ 16,\ 32,\ 64,\ 128]$

$k_u = k_d = [3,\ 3,\ 3,\ 3,\ 3]$

$n_s = [0,\ 0,\ 0,\ 4,\ 4]$

$k_s = [\mathrm{NA},\ \mathrm{NA},\ \mathrm{NA},\ 1,\ 1]$

$\sigma_p = \frac{1}{30}$

num_iter = 2400

LR = 0.01

upsampling = bilinear

Results

$$PSNR(\hat{x}, x^*) = 10\log_{10}\left(\frac{MAX(\hat{x})^2}{MSE(\hat{x}, x^*)}\right)$$
$$= 20\log_{10}(MAX(\hat{x})) - 10\log_{10} MSE(\hat{x}, x^*)$$

	Deep Image Prior	CMB3D	Non-local Means	SRResNET	Bicubic	LapSRN	Papyan & al
Denoising on Standard Dataset	31	31,42	30,26				
Super Resolution x4 Set 5	28,89			32	28,44		
Super Resolution x4 Set 14	27			28,53	26,05		
Super Resolution x8 Set 5	25,88				24,37	26,1	
Super Resolution x8 Set 14	24,14				23,09	24,35	
Inpaint on Standard Dataset	33,48						31,1

Advantages

Doesn’t Require Training

Fairly close to state of the art

Simple to implement

Disadvantages

Doesn't necessarily converge

Can be fairly slow

Requires a models

Current Advancements

Deep Image Prior

- Early Stopping (Dec 2021) [12]
- Neural architecture search using Genetic Algorithme (Jan 2020) [13]
- Bayesian Perspective for Deep Image prior (apr 2019) [14]

Inverse Image

- Controllable confidence-based image **denoising** (jan 2021) [15]
- Image **Inpainting** with Partial Convolutions (Apr 2018) [16]
- Image Fine-grained **Inpainting** (feb 2020) [17]
- Image Restoration Using Swin Transformer (Aug 2021) [18]

Possible advancements

Pre Training

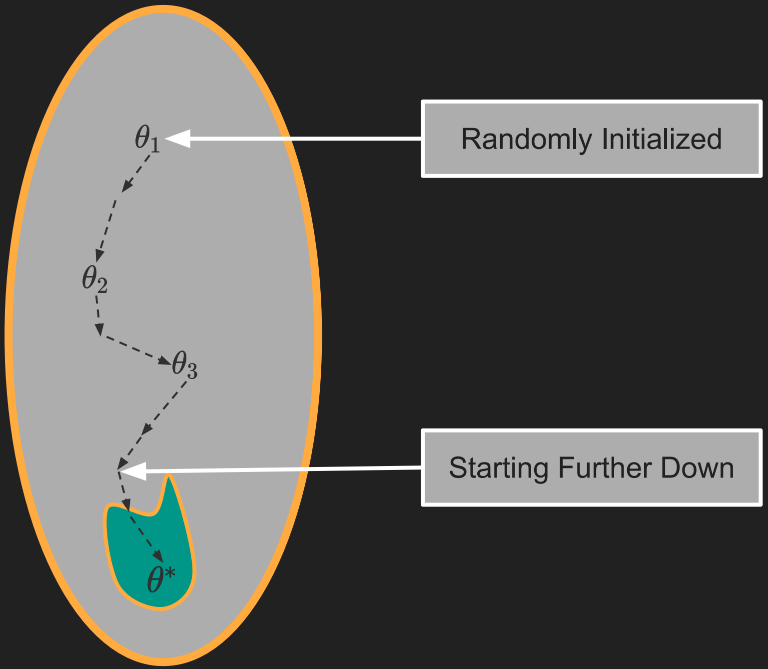

Hyperparameter Space

Model Pre-Selection

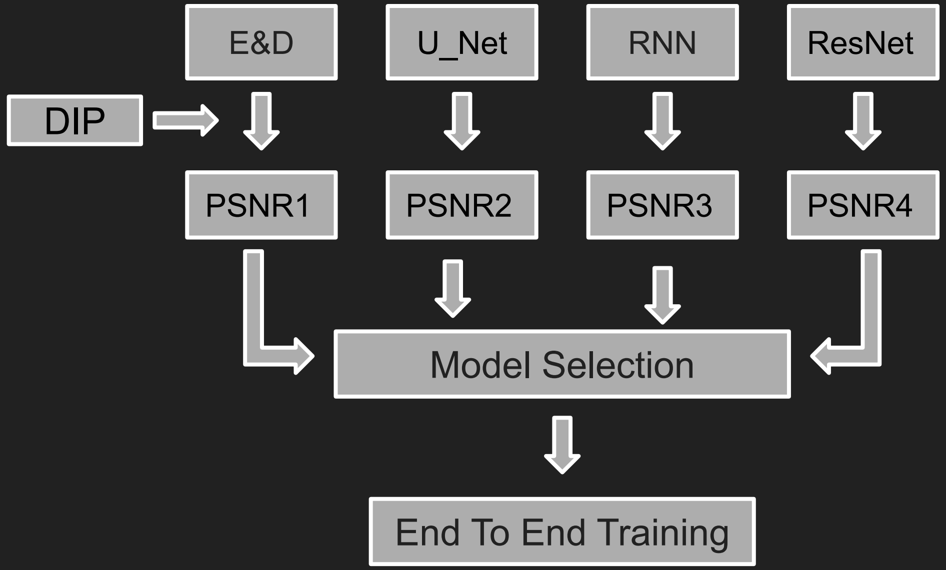

Conclusion

Inverse Problems are frequently encountered

Not necessarily a better solution that state of the art

Models contain a lot of information