
Deep Image Prior & Neural Fields

Representing images by neural networks

Supervised Learning

- Empirical distribution of training data $(x, y) \sim \hat{p}(x, y)$

$$(\mathbb{X}, \mathbb{Y}) = \{(x^{(1)}, y^{(1)}), \dots, (x^{(m)}, y^{(m)})\}$$

$$\psi_{ML} = \arg \max_{\psi} q(\mathbb{Y} | \mathbb{X}, \psi)$$

$$= \arg \max_{\psi} \sum_{i=1}^m \log q(y^{(i)} | x^{(i)}, \psi)$$

$$\psi_{ML} = \arg \max_{\psi} \mathbb{E}_{(x, y) \sim \hat{p}(x, y)} \log q(y | x, \psi)$$

Supervised Learning

- Training data: $(x, y) \sim \hat{p}(x, y)$
- Choose distribution $q(y|x; \psi)$
- NN represents a function $f_{\theta}(x)$, which is not a direct prediction of y , but of ψ (e.g. mean)
- ML principle give us the cost function as
$$-\log q(y|x; f_{\theta}(x))$$

Example – Gaussian Model

- Training data distribution: $p(x, y) = p(y|x)p(x)$
- Estimator of mean y : $f(x)$
- Model distribution: $q(y|x) \equiv N(y|f(x), \sigma^2)$

$$\hat{f}(x) = \arg \min_{f(x)} \mathbb{E}_{p(x,y)} [-\log q(y|f(x))]$$

$$L = \int \left[\int (y - f(x))^2 p(y|x) dy \right] p(x) dx + c$$

$$\frac{\partial L}{\partial f} \propto \int (y - f(x)) p(y|x) dy = 0$$

$$\hat{f}(x) = \int y p(y|x) dy = \mathbb{E}_{p(y|x)} [y]$$

Example – Laplace Model

- Training data distribution: $p(x, y) = p(y|x)p(x)$
- Estimator of mean y : $f(x)$
- Model distribution: $q(y|x) \equiv \text{Laplace}(y|f(x), \gamma)$

$$\hat{f}(x) = \arg \min_{f(x)} \mathbb{E}_{p(x,y)} [-\log q(y|f(x))]$$

$$L = \int \left[\int |y - f(x)| p(y|x) dy \right] p(x) dx + c$$


$$\frac{\partial L}{\partial f} \propto \int \text{sign}(y - f(x)) p(y|x) dy = 0$$

$$\hat{f}(x) = m \quad \dots \text{median: } P(y \leq m|x) = P(y \geq m|x)$$

Deep Image Prior

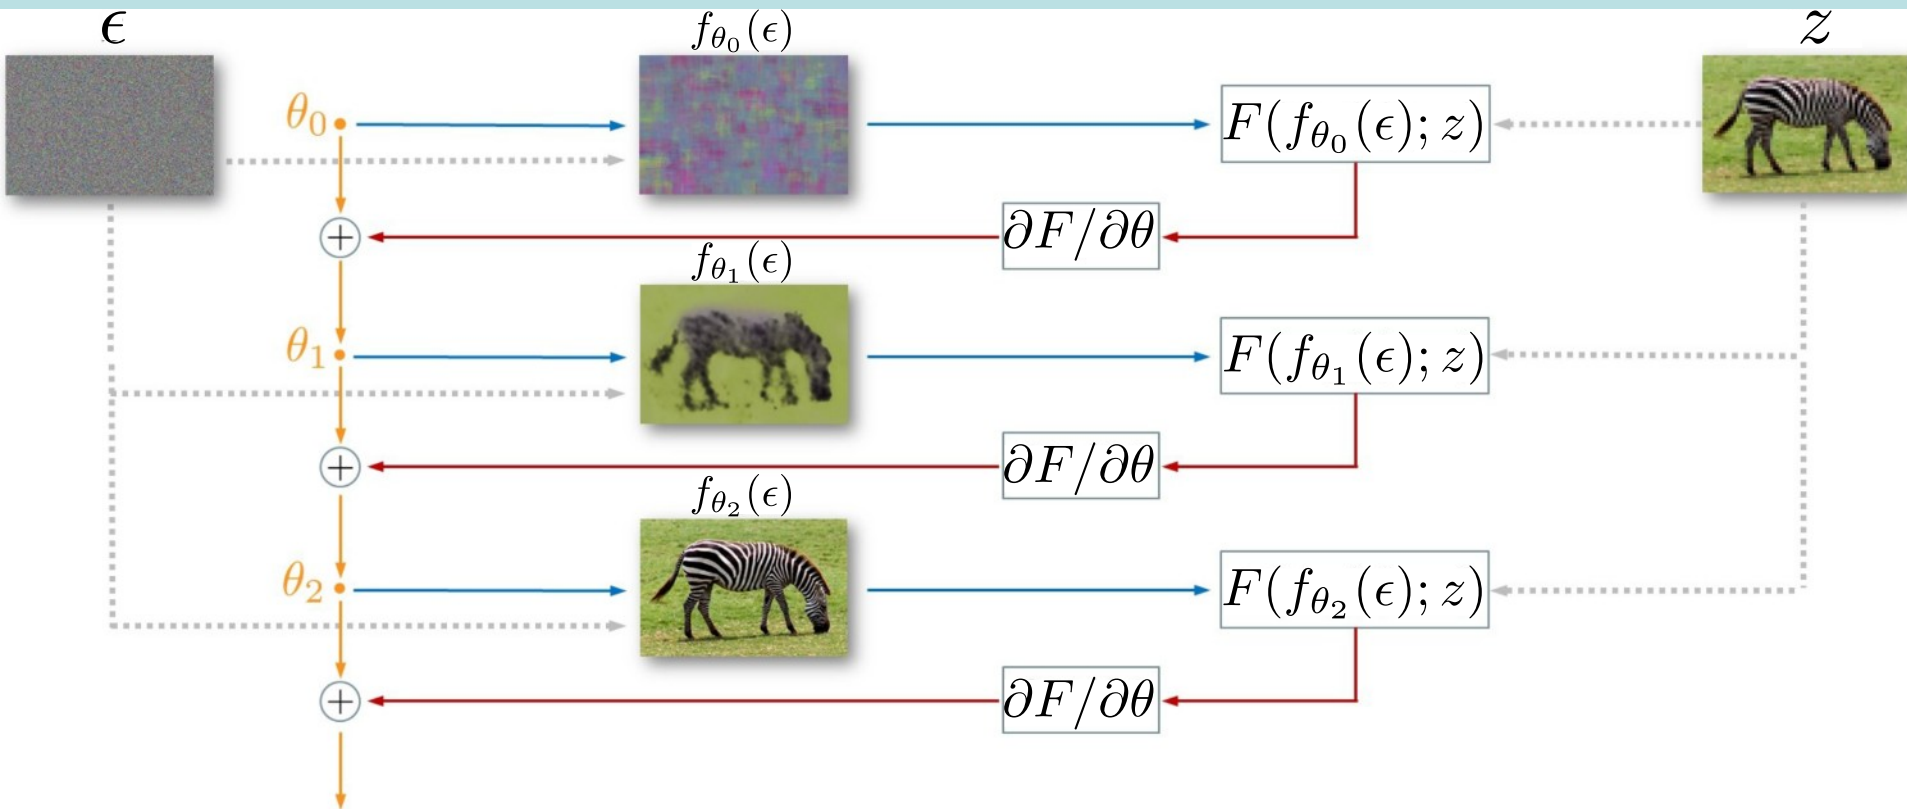
- Denoising $z = u + n$ $\min_u F(u; z)$
- no prior $F(u; z) = \|u - z\|^2$
- with prior $F(u; z) = \|u - z\|^2 + R(u)$
- with Deep Image Prior $\min_{\theta} F(f_{\theta}(\epsilon); z)$
 $F(f_{\theta}(\epsilon); z) = \|f_{\theta}(\epsilon) - z\|^2$

- $f_{\theta}(\epsilon)$ is U-Net architecture




Zero-shot learning

- Gradient descent



Denoising example



(a) GT

(b) Input

(c) DIP

(d) CBM3D

(e) NLM

Super-resolution example



(a) HR image

(b) Bicubic upsampling

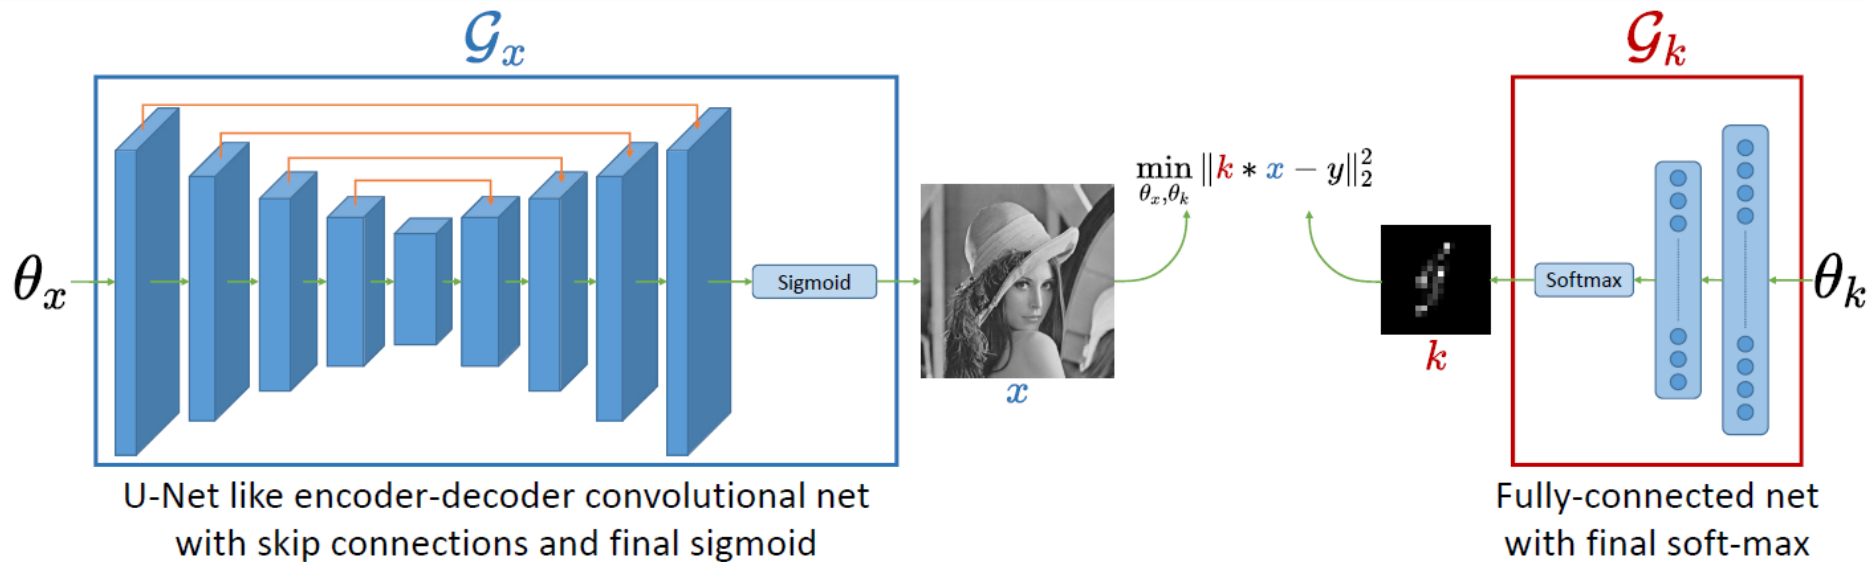
(c) No prior

(d) TV prior

(e) Deep image prior

Blind Deconvolution


$$\min_{\theta_x, \theta_k} \|k * x - y\|_2^2, \quad \text{s.t. } 0 \leq x \leq 1 \text{ and } 0 \leq k, \|k\|_1 = 1, \quad \text{where } x = \mathcal{G}_x(\theta_x), k = \mathcal{G}_k(\theta_k)$$



Implicit Neural Representation


- also called Neural Fields
- represent an image(object) function with MLP

$$f(x) : \mathbb{R}^m \rightarrow \mathbb{R}^n$$



– e.g. color image $f(x) : \mathbb{R}^2 \rightarrow \mathbb{R}^3$

Zero-shot learning





- e.g. fitting image $u(x)$:

$$\min_{\theta} \sum_i \|f_\theta(x_i) - u(x_i)\|^2$$




- or more complex losses

Neural Fields

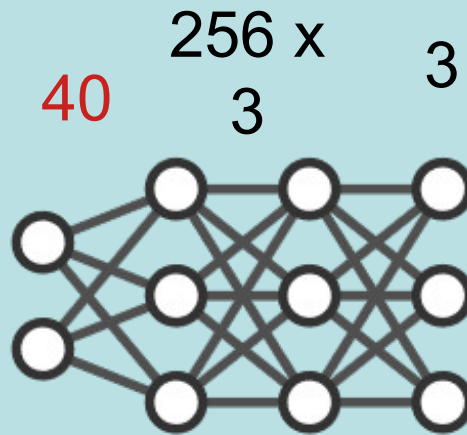
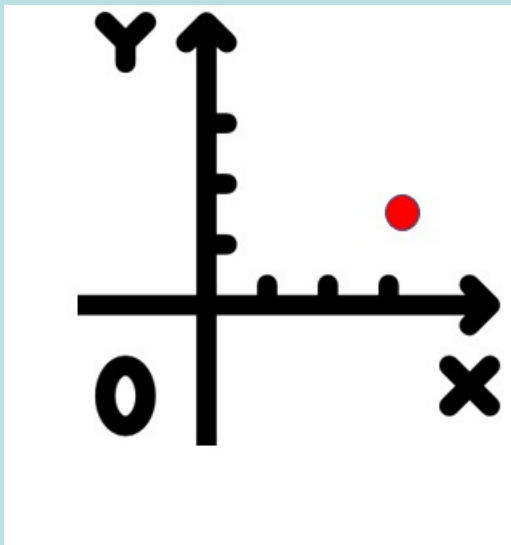
- Pros:
 - Compressed representation
 - Analytic derivatives $\nabla f(x)$
using back-propagation
- Cons:
 - train for every case



- ReLU



- Positional encoding



$$\gamma(\mathbf{x}) = [\gamma_1(\mathbf{x}), \gamma_2(\mathbf{x}), \dots, \gamma_m(\mathbf{x})]$$


$$\gamma_{(2i)}(x) = \sin(2^{i-1}\pi x),$$

$$\gamma_{(2i+1)}(x) = \cos(2^{i-1}\pi x)$$


SIREN

A) Network Φ

ReLU



SIREN



Derivatives

1st (Φ')

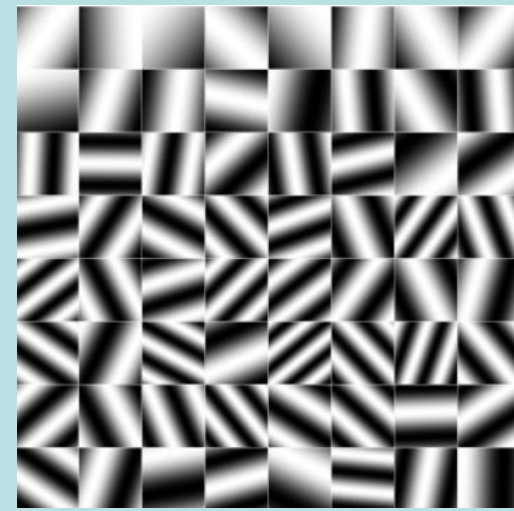
2nd (Φ'')





$$\text{ReLU}(\mathbf{W}\mathbf{x} + \mathbf{b})$$

$$\sin(\mathbf{W}\mathbf{x} + \mathbf{b})$$


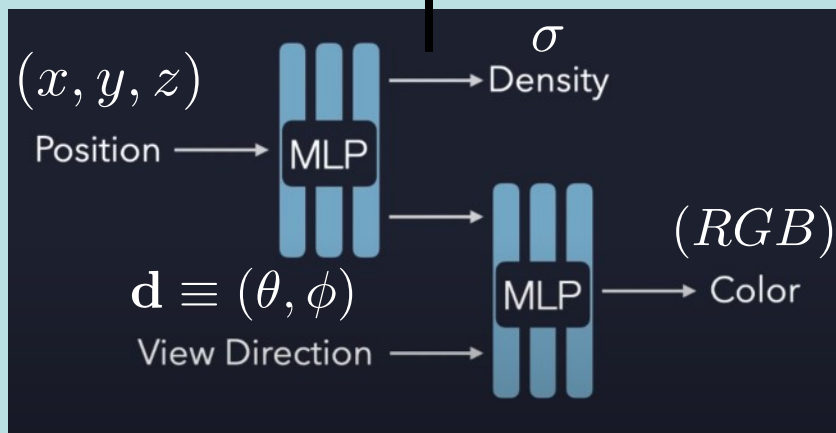

[Sitzmann et al. 2021]




- Sinus activation



NeRF




$$C(\mathbf{r}) = \int_{t_n}^{t_f} T(t) \sigma(\mathbf{r}(t)) \mathbf{c}(\mathbf{r}(t), \mathbf{d}) dt$$



Ground Truth



Complete Model



No View Dependence

NVIDIA Instant NeRF



Input Images

NVIDIA NGP
Instant NeRF



CURE

- Video

