

# Bayesian Imaging with Plug & Play Priors

## implicit & explicit cases

Andrés Almansa

Joint work with:

Mario González, Rémi Laumont, Valentin De Bortoli, Mauricio Delbracio, Julie Delon, Alain Durmus, José Lezama, Pablo Musé, Marcelo Pereyra.



Université  
de Paris



Preprint and code available here:

<http://up5.fr/jpmap>


<https://arxiv.org/abs/2103.04715>

July 27th 2021 - Bath LMS Workshop  
Analytic and Geometric Approaches to Machine Learning

- 1 Introduction
  - Inverse problems in Imaging
- 2 Implicitly decoupled methods
  - Proximal-based (PnP-ADMM)
  - Tweedie-based (PnP-SGD, PnP-ULA)
- 3 Explicitly decoupled methods
  - Variational AutoEncoder Priors
  - Joint Posterior Maximization with AutoEncoding Prior Denoising Criterion
  - Continuation Scheme
  - Experiments

# Inverse Problems in Imaging

Estimate clean image  $\mathbf{x} \in \mathbb{R}^d$   
from noisy, degraded measurements  $\mathbf{y} \in \mathbb{R}^m$ .



Known degradation model (usually log-concave):

$$p_{Y|X}(\mathbf{y} | \mathbf{x}) \propto e^{-F(\mathbf{x}, \mathbf{y})} \quad \text{where} \quad F(\mathbf{x}, \mathbf{y}) = \frac{1}{2\sigma^2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|^2. \quad (1)$$

# Inverse Problems in Imaging

Estimate clean image  $\mathbf{x} \in \mathbb{R}^d$

from noisy, degraded measurements  $\mathbf{y} \in \mathbb{R}^m$ .

Known degradation model (usually log-concave):

$$p_{Y|X}(\mathbf{y} | \mathbf{x}) \propto e^{-F(\mathbf{x}, \mathbf{y})} \quad \text{where} \quad F(\mathbf{x}, \mathbf{y}) = \frac{1}{2\sigma^2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|^2. \quad (1)$$

## Variational/Bayesian Approach

Use image prior  $p_X(\mathbf{x}) \propto e^{-\lambda R(\mathbf{x})}$  to compute estimator

$$\hat{\mathbf{x}}_{\text{MAP}} = \arg \max_{\mathbf{x}} p_{X|Y}(\mathbf{x} | \mathbf{y}) = \arg \min_{\mathbf{x}} \{F(\mathbf{x}, \mathbf{y}) + \lambda R(\mathbf{x})\} \quad (2)$$

$$\hat{\mathbf{x}}_{\text{MMSE}} = \arg \min_{\mathbf{x}} \mathbb{E} \left[ \|\mathbf{X} - \mathbf{x}\|^2 \mid \mathbf{Y} = \mathbf{y} \right] \quad (3)$$

# Inverse Problems in Imaging

$$p_{Y|X}(\mathbf{y} | \mathbf{x}) \propto e^{-F(\mathbf{x}, \mathbf{y})} \quad \text{where} \quad F(\mathbf{x}, \mathbf{y}) = \frac{1}{2\sigma^2} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|^2. \quad (1)$$

## Variational/Bayesian Approach

Use image prior  $p_X(\mathbf{x}) \propto e^{-\lambda R(\mathbf{x})}$  to compute estimator


$$\hat{\mathbf{x}}_{\text{MAP}} = \arg \max_{\mathbf{x}} p_{X|Y}(\mathbf{x} | \mathbf{y}) = \arg \min_{\mathbf{x}} \{F(\mathbf{x}, \mathbf{y}) + \lambda R(\mathbf{x})\} \quad (2)$$

$$\hat{\mathbf{x}}_{\text{MMSE}} = \arg \min_{\mathbf{x}} \mathbb{E} \left[ \|X - \mathbf{x}\|^2 \mid Y = \mathbf{y} \right] \quad (3)$$

## Common explicit priors

- **Total Variation** (CHAMBOLLE, 2004; LOUCHET AND MOISAN, 2013; PEREYRA, 2016; RUDIN ET AL., 1992)
- **Gaussian Mixtures** (TEODORO ET AL., 2018; YU ET AL., 2011; ZORAN AND WEISS, 2011)


## Neural Networks for inverse problems:

*Two paradigms*

- Agnostic approach** : find a sufficient number of image pairs  $(\mathbf{x}^i, \mathbf{y}^i)$  and train a neural network  $f_\theta$  to invert  $A$  by minimizing the empirical risk  $\sum_i \|f_\theta(\mathbf{y}^i) - \mathbf{x}^i\|_2^2$ 
  - ✓ no need to model  $A$ ,  $\mathbf{n}$  nor prior for  $\mathbf{x}$
  - ✗ needs retraining if  $A$  or  $\mathbf{n}$  change

## Neural Networks for inverse problems:

Two paradigms



- **Agnostic approach** : find a sufficient number of image pairs  $(\mathbf{x}^i, \mathbf{y}^i)$  and train a neural network  $f_\theta$  to invert  $A$  by minimizing the empirical risk  $\sum_i \|f_\theta(\mathbf{y}^i) - \mathbf{x}^i\|_2^2$ 
  - ✓ no need to model  $A$ ,  $\mathbf{n}$  nor prior for  $\mathbf{x}$
  - ✗ needs retraining if  $A$  or  $\mathbf{n}$  change
- **Decoupled (plug & play) approach** : Model separately
  - 1 conditional density  $p_{Y|X}(\mathbf{y} | \mathbf{x})$   
(using physical model, calibration)
  - 2 prior model  $p_X(\mathbf{x})$   
(through NN learning)
  - 3 Use Bayes theorem to estimate  $\mathbf{x}$  via MAP or MMSE

## Neural Networks for inverse problems:

*Two paradigms*

- **Agnostic approach** : find a sufficient number of image pairs  $(\mathbf{x}^i, \mathbf{y}^i)$  and train a neural network  $f_\theta$  to invert  $A$  by minimizing the empirical risk  $\sum_i \|f_\theta(\mathbf{y}^i) - \mathbf{x}^i\|_2^2$ 
  - ✓ no need to model  $A$ ,  $\mathbf{n}$  nor prior for  $\mathbf{x}$
  - ✗ needs retraining if  $A$  or  $\mathbf{n}$  change
- **Decoupled (plug & play) approach** : Model separately
  - ① conditional density  $p_{\mathbf{Y}|\mathbf{X}}(\mathbf{y}|\mathbf{x})$   
(using physical model, calibration)
  - ② prior model  $p_{\mathbf{X}}(\mathbf{x})$   
(through NN learning)
  - ③ Use Bayes theorem to estimate  $\mathbf{x}$  via MAP or MMSE
    - ✓ uses all available modeling information
    - ✓ train once, use for many inverse problems
    - ▲ difficult to learn  $p_{\mathbf{X}}(\mathbf{x})$  directly
    - ▲ Non-convex optimization



## Neural Networks for inverse problems:

*Implicitly decoupled approach*

Solve the optimization problem

$$\hat{\mathbf{x}}_{\text{MAP}} = \arg \max_{\mathbf{x}} p_{\mathbf{X}|\mathbf{Y}}(\mathbf{x} | \mathbf{y}) = \arg \min_{\mathbf{x}} \{F(\mathbf{x}, \mathbf{y}) + \lambda R(\mathbf{x})\}$$

via ADMM splitting (RYU ET AL., 2019)

- 1  $\mathbf{v}_{k+1} = \arg \min_{\mathbf{v}} R(\mathbf{v}) + \frac{1}{2\delta^2} \|\mathbf{v} - (\mathbf{x}_k - \mathbf{u}_k)\|^2$
- 2  $\mathbf{x}_{k+1} = \arg \min_{\mathbf{x}} F(\mathbf{x}, \mathbf{y}) + \frac{\lambda}{2\delta^2} \|\mathbf{x} - (\mathbf{v}_{k+1} - \mathbf{u}_k)\|^2$
- 3  $\mathbf{u}_{k+1} = \mathbf{u}_k + \mathbf{v}_{k+1} - \mathbf{x}_{k+1}$

$R$  is unknown but we can use a train a neural network to approximate the  $\delta$ -denoising problem in step 1:

$$D_{\delta}(\tilde{\mathbf{x}}) = \arg \min_{\mathbf{v}} R(\mathbf{v}) + \frac{1}{2\delta^2} \|\mathbf{v} - \tilde{\mathbf{x}}\|^2$$

## Neural Networks for inverse problems:

*Implicitly decoupled approach*

Solve the optimization problem via ADMM splitting

$$\hat{\mathbf{x}}_{\text{MAP}} = \arg \max_{\mathbf{x}} p_{X|Y}(\mathbf{x} | \mathbf{y}) = \arg \min_{\mathbf{x}} \{F(\mathbf{x}, \mathbf{y}) + \lambda R(\mathbf{x})\}$$

$R$  is unknown but a NN approximates its proximal operator:

$$D_{\delta}(\tilde{\mathbf{x}}) = \arg \min_{\mathbf{v}} R(\mathbf{v}) + \frac{1}{2\delta^2} \|\mathbf{v} - \tilde{\mathbf{x}}\|^2$$

**Challenges**

- NN training produces an approximate MMSE rather than a MAP estimator for  $D_{\delta}$
- Convergence guarantees?

## Neural Networks for inverse problems:

*Implicitly decoupled approach*

Solve the optimization problem via ADMM splitting (RYU ET AL., 2019)

$$\hat{\mathbf{x}}_{\text{MAP}} = \arg \max_{\mathbf{x}} p_{X|Y}(\mathbf{x} | \mathbf{y}) = \arg \min_{\mathbf{x}} \{F(\mathbf{x}, \mathbf{y}) + \lambda R(\mathbf{x})\}$$

## Assumption (A)

- 1 ✓  $D_{\delta} - I$  is  $L$ -Lipschitz with  $L \in (0, 1)$
- 2 ✗  $F(\cdot, \mathbf{y})$  is  $\mu$ -strongly convex
- 3 ✗  $\lambda < \frac{\sigma^2 \mu (1 + L - 2L^2)}{L} \xrightarrow{L \rightarrow 1^-} 0$

## Theorem (RYU ET AL. (2019))

Under assumption A, the Plug & Play ADMM algorithm converges to a fixed point.

# Tweedie's formula (EFRON, 2011)

Alternative link between denoiser  $D_\delta$  and prior  $p_X$ :

Tweedie's formula (EFRON, 2011)

If  $X \sim p_X$ ,  $N \sim \mathcal{N}(0, \delta^2 Id)$  and  $D_\delta(\mathbf{y}) = \mathbb{E}[X | X + N = \mathbf{y}]$   
then

$$(D_\delta - Id)(\mathbf{x}) = \delta^2 \nabla \log p_X^\delta(\mathbf{x})$$

with  $p_X^\delta := p_X * g_\delta$ , ( $g_\delta$  Gaussian of variance  $\delta^2$ ).

Instead of maximizing the true posterior  $\pi(\mathbf{x}) \propto p(\mathbf{y}|\mathbf{x})p_X(\mathbf{x})$   
we maximize the approximate posterior  $\pi^\delta(\mathbf{x}) \propto p(\mathbf{y}|\mathbf{x})p_X^\delta(\mathbf{x})$   
*i.e.* we minimize

$$E^\delta(\mathbf{x}) = -\log \pi^\delta(\mathbf{x}) = F(\mathbf{x}, \mathbf{y}) + R^\delta(\mathbf{x}) \quad \text{with } R^\delta(\mathbf{x}) = -\log p_X^\delta(\mathbf{x})$$

whose gradient writes exactly:

$$\nabla E^\delta(\mathbf{x}) = \nabla F(\mathbf{x}, \mathbf{y}) - \frac{1}{\delta^2} (D^\delta - Id)(\mathbf{x})$$

# PnP-SGD

Recall: we want to minimize

$$E^\delta(\mathbf{x}) = -\log \pi^\delta(\mathbf{x}) = F(\mathbf{x}, \mathbf{y}) + R^\delta(\mathbf{x}) \quad \text{with } R^\delta(\mathbf{x}) = -\log p_X^\delta(\mathbf{x})$$

whose gradient writes :

$$\nabla E^\delta(\mathbf{x}) = \nabla F(\mathbf{x}, \mathbf{y}) - \frac{1}{\delta^2}(D^\delta - Id)(\mathbf{x})$$

PnP-SGD convergence (LAUMONT ET AL., 2021)

Under mild assumptions, the PnP Stochastic Gradient descent

$$X_{k+1} = X_k - \gamma_k \nabla E^\delta(X_k) + \gamma_k Z_k$$

converges to a critical point of  $\pi^\delta$

# PnP-ULA

Recall the smooth posterior  $\pi^\delta$  satisfies

$$E^\delta(\mathbf{x}) = -\log \pi^\delta(\mathbf{x}) = F(\mathbf{x}, \mathbf{y}) + R^\delta(\mathbf{x}) \quad \text{with } R^\delta(\mathbf{x}) = -\log p_X^\delta(\mathbf{x})$$

whose gradient writes exactly:

$$\nabla E^\delta(\mathbf{x}) = \nabla F(\mathbf{x}, \mathbf{y}) - \frac{1}{\delta^2}(D^\delta - Id)(\mathbf{x})$$

A small variation of PnP-SGD provides a posterior sampling scheme

## PnP-ULA convergence (LAUMONT ET AL., 2021)


Under mild assumptions (including  $Lip(I - D_\varepsilon) < 1$ ), the (Plug & Play Unadjusted Langevin Algorithm)

$$X_{k+1} = X_k - \gamma \nabla E^\delta(X_k) + \sqrt{2\gamma} Z_k$$

provides a set of samples of the posterior  $\pi^\delta \approx \pi$  satisfying the following non-asymptotic error bound :

$$\left| \frac{1}{n} \sum_{k=1}^n \mathbb{E} [h(X_k)] - \int_{\mathbb{R}^d} h(\mathbf{y}) \pi^\delta(\mathbf{y}) d\mathbf{y} \right| \leq \left( C_1 + C_2 \left( \sqrt{\gamma} + \frac{1}{n\gamma} + C_M \right) \right)$$

observation  
(80% missing pixels)



PSNR=7.45.

PnP-ULA  
posterior mean




PSNR=31.51.

PnP-ADMM.



PSNR=30.06.



posterior std

## Conclusion on PnP-ULA

- ✓ PnP-ULA works remarkably well for point estimation, posterior sampling, uncertainty estimation
- ✓ PnP-ULA provides convergence guarantees under realistic conditions
- ✗ quite slow (20 minutes ... 2 days)
- ✗ prior  $p_X$  unknown (required for some uncertainty estimation techniques)



- 1 Introduction
  - Inverse problems in Imaging
- 2 Implicitly decoupled methods
  - Proximal-based (PnP-ADMM)
  - Tweedie-based (PnP-SGD, PnP-ULA)
- 3 Explicitly decoupled methods
  - Variational AutoEncoder Priors
  - Joint Posterior Maximization with AutoEncoding Prior
  - Denoising Criterion
  - Continuation Scheme
  - Experiments

## Explicitly decoupled approach (MAP- $\mathbf{x}$ ):

*How to use neural networks to learn the prior  $p_{\mathbf{X}}(\mathbf{x})$  ?*


Generative Adversarial Networks (GANs) (ARJOVSKY AND BOTTOU, 2017; GOODFELLOW ET AL., 2014)

Learn a generator function  $G$  that maps

$$\mathbf{z} \sim \mathcal{N}(0, Id)$$

to

$$\mathbf{x} = G(\mathbf{z}) \sim p_{\mathbf{X}}$$



## Explicitly decoupled approach (MAP- $\mathbf{x}$ ):

*How to use neural networks to learn the prior  $p_{\mathbf{x}}(\mathbf{x})$  ?*


Generative Adversarial Networks (GANs) (ARJOVSKY AND BOTTOU, 2017; GOODFELLOW ET AL., 2014)

Learn a generator function  $G$  that maps

$$\mathbf{z} \sim \mathcal{N}(0, Id)$$

to

$$\mathbf{x} = G(\mathbf{z}) \sim p_{\mathbf{x}}$$



MAP- $\mathbf{x}$  Following PAMAMAKARIOS ET AL. (2019, SECTION 5), the push-forward measure  $p_{\mathbf{x}} = G\#p_{\mathbf{z}}$  can be developed as

$$p_{\mathbf{x}}(\mathbf{x}) = \frac{p_{\mathbf{z}}(G^{-1}(\mathbf{x}))}{\sqrt{\det S(G^{-1}(\mathbf{x}))}} \delta_{\mathcal{M}}(\mathbf{x})$$

where

$$S = \left( \frac{\partial G}{\partial \mathbf{z}} \right)^T \left( \frac{\partial G}{\partial \mathbf{z}} \right)$$

$$\mathcal{M} = \{ \mathbf{x} : \exists \mathbf{z}, \mathbf{x} = G(\mathbf{z}) \}$$

## Explicitly decoupled approach (MAP-x):

*How to use neural networks to learn the prior  $p_X(\mathbf{x})$  ?*

**Generative Adversarial Networks (GANs)** (ARJOVSKY AND BOTTOU, 2017; GOODFELLOW ET AL., 2014) Learn a generator function  $G$  that maps

$$\mathbf{z} \sim \mathcal{N}(0, Id) \quad \text{to} \quad \mathbf{x} = G(\mathbf{z}) \sim p_X$$

**MAP-x** Following PAPAMAKARIOS ET AL. (2019, SECTION 5), the push-forward measure  $p_X = G\#p_Z$  can be developed as

$$p_X(\mathbf{x}) = \frac{p_Z(G^{-1}(\mathbf{x}))}{\sqrt{\det S(G^{-1}(\mathbf{x}))}} \delta_{\mathcal{M}}(\mathbf{x})$$

where

$$S = \left( \frac{\partial G}{\partial \mathbf{z}} \right)^T \left( \frac{\partial G}{\partial \mathbf{z}} \right)$$
$$\mathcal{M} = \{ \mathbf{x} : \exists \mathbf{z}, \mathbf{x} = G(\mathbf{z}) \}$$

$\mathbf{x}$ -optimization required to obtain  $\hat{\mathbf{x}}_{\text{MAP}}$  becomes intractable due to:

- computation of  $S$  and  $\det S$ ,
- inversion of  $G$ , and
- hard constraint  $\mathbf{x} \in \mathcal{M}$

## Explicitly decoupled approach (MAP-z):

Instead of solving the  $\mathbf{x}$ -optimisation problem:

$$\hat{\mathbf{x}}_{\text{MAP}} = \arg \max_{\mathbf{x}} p_{Y|X}(\mathbf{y} | \mathbf{x}) p_X(\mathbf{x}) = \arg \min_{\mathbf{x}} \{F(\mathbf{x}, \mathbf{y}) + R(\mathbf{x})\}$$

BORA ET AL. (2017) propose to optimize over  $\mathbf{z}$

$$\begin{aligned}\hat{\mathbf{z}} &= \arg \max_{\mathbf{z}} \{p_{Y|X}(\mathbf{y} | G(\mathbf{z})) p_Z(\mathbf{z})\} \\ &= \arg \min_{\mathbf{z}} \left\{ F(G(\mathbf{z}), \mathbf{y}) + \frac{1}{2} \|\mathbf{z}\|^2 \right\}\end{aligned}$$

$$\hat{\mathbf{x}}_{\mathbf{z}\text{-MAP}} = G(\hat{\mathbf{z}})$$

$\hat{\mathbf{x}}_{\mathbf{z}\text{-MAP}} (\neq \hat{\mathbf{x}}_{\text{MAP}})$  but it maximizes the latent posterior:

$$\hat{\mathbf{x}}_{\mathbf{z}\text{-MAP}} = G \left( \arg \max_{\mathbf{z}} \{p_{Z|Y}(\mathbf{z} | \mathbf{y})\} \right)$$

## Explicitly decoupled approach (MAP-z):

$\hat{\mathbf{x}}_{\mathbf{z}-\text{MAP}}$  ( $\neq \hat{\mathbf{x}}_{\text{MAP}}$ ) maximizes the latent posterior:

$$\begin{aligned}\hat{\mathbf{x}}_{\mathbf{z}-\text{MAP}} &= \text{G} \left( \arg \max_{\mathbf{z}} \{ p_{\mathbf{Z}|\mathbf{Y}}(\mathbf{z} | \mathbf{y}) \} \right) \\ &= \text{G} \left( \arg \min_{\mathbf{z}} \left\{ F(\text{G}(\mathbf{z}), \mathbf{y}) + \frac{1}{2} \|\mathbf{z}\|^2 \right\} \right)\end{aligned}$$

### Challenges

- Nonconvex optimization using gradient descent
- may get stuck in spurious local minima

Common solution: Splitting + continuation scheme

## MAP-z splitting and continuation scheme.

$$\hat{\mathbf{x}}_{\beta} = \arg \min_{\mathbf{x}} \min_{\mathbf{z}} \underbrace{\left\{ F(\mathbf{x}, \mathbf{y}) + \frac{\beta}{2} \|\mathbf{x} - G(\mathbf{z})\|^2 + \frac{1}{2} \|\mathbf{z}\|^2 \right\}}_{J_{1,\beta}(\mathbf{x}, \mathbf{z})}$$

$$\hat{\mathbf{x}}_{\text{MAP-z}} = \lim_{\beta \rightarrow \infty} \hat{\mathbf{x}}_{\beta}.$$

---

**Algorithm 1.1** MAP-z splitting

---

**Require:** Measurements  $\mathbf{y}$ , Initial condition  $\mathbf{x}_0$

**Ensure:**  $\hat{\mathbf{x}} = G(\arg \max_{\mathbf{z}} p_{Z|Y}(\mathbf{z} | \mathbf{y}))$

```
1: for  $k := 0$  to  $k_{\max}$  do
2:    $\beta := \beta_k$ 
3:   for  $n := 0$  to maxiter do
4:      $\mathbf{z}_{n+1} := \arg \min_{\mathbf{z}} J_{1,\beta}(\mathbf{x}_n, \mathbf{z})$            // Nonconvex
5:      $\mathbf{x}_{n+1} := \arg \min_{\mathbf{x}} J_{1,\beta}(\mathbf{x}, \mathbf{z}_{n+1})$        // Quadratic
6:   end for
7:    $\mathbf{x}_0 := \mathbf{x}_{n+1}$ 
8: end for
9: return  $\mathbf{x}_{n+1}$ 
```

---

Non-convex step 4: Use a local quadratic approximation (VAE encoder) ...

# VAEs and Joint Posterior


## Generative Adversarial Networks (GANs) (GOODFELLOW ET AL., 2014)

Learn a generator function  $G$  that maps

$$\mathbf{z} \sim \mathcal{N}(0, Id)$$

to

$$\mathbf{x} = G(\mathbf{z}) \sim p_{\mathbf{x}}$$





# VAEs and Joint Posterior


## Generative Adversarial Networks (GANs) (GOODFELLOW ET AL., 2014)

Learn a generator function  $G$  that maps


$$z \sim \mathcal{N}(0, Id)$$

to

$$x = G(z) \sim p_x$$



## Variational AutoEncoders (VAEs) (KINGMA AND WELLING, 2013)



Generative model:


$$p_{x|z}(x|z) = p_\theta(x|z) = \mathcal{N}(x; \mu_\theta(z), \gamma Id)$$

Approximate inverse:

$$p_{z|x}(z|x) \approx q_\phi(z|x) = \mathcal{N}(z; \mu_\phi(x), \Sigma_\phi(x))$$

# VAEs and Joint Posterior

## Variational AutoEncoders (VAEs) (KINGMA AND WELLING, 2013)



Generative model:

$$p_{\mathbf{x}|\mathbf{z}}(\mathbf{x}|\mathbf{z}) = p_{\theta}(\mathbf{x}|\mathbf{z}) = \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_{\theta}(\mathbf{z}), \gamma Id)$$

Approximate inverse:

$$p_{\mathbf{z}|\mathbf{x}}(\mathbf{z}|\mathbf{x}) \approx q_{\phi}(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\mathbf{z}; \boldsymbol{\mu}_{\phi}(\mathbf{x}), \boldsymbol{\Sigma}_{\phi}(\mathbf{x}))$$

Learning: Maximize the averaged *Evidence Lower BOund (ELBO)* for  $\mathbf{x} \in \mathcal{D}$

$$\mathcal{L}_{\theta, \phi}(\mathbf{x}) = \mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{x})}[\log p_{\theta}(\mathbf{x}|\mathbf{z})] - KL(q_{\phi}(\mathbf{z}|\mathbf{x}) || p_{\mathbf{z}}(\mathbf{z})) \leq \log p_{\theta}(\mathbf{x}).$$

# VAEs and Joint Posterior

**Variational AutoEncoders (VAEs)** (KINGMA AND WELLING, 2013)

Generative model:

$$p_{X|Z}(\mathbf{x} | \mathbf{z}) = p_{\theta}(\mathbf{x} | \mathbf{z}) = \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_{\theta}(\mathbf{z}), \gamma Id)$$

Joint density:

$$p_{X,Z}(\mathbf{x}, \mathbf{z}) = p_{\theta}(\mathbf{x} | \mathbf{z}) p_Z(\mathbf{z})$$

Approximate inverse:

$$p_{Z|X}(\mathbf{z} | \mathbf{x}) \approx q_{\phi}(\mathbf{z} | \mathbf{x}) = \mathcal{N}(\mathbf{z}; \boldsymbol{\mu}_{\phi}(\mathbf{x}), \boldsymbol{\Sigma}_{\phi}(\mathbf{x}))$$

Approximate joint density:

$$\tilde{p}_{X,Z}(\mathbf{x}, \mathbf{z}) := q_{\phi}(\mathbf{z} | \mathbf{x}) p_X(\mathbf{x}) \approx p_{X,Z}(\mathbf{x}, \mathbf{z})$$

# VAEs and Joint Posterior

**Variational AutoEncoders (VAEs)** (KINGMA AND WELLING, 2013)

Generative model:  $p_{X|Z}(x|z) = p_{\theta}(x|z) = \mathcal{N}(x; \mu_{\theta}(z), \gamma Id)$

Joint density:  $p_{X,Z}(x, z) = p_{\theta}(x|z) p_Z(z)$

Approximate inverse:  $p_{Z|X}(z|x) \approx q_{\phi}(z|x) = \mathcal{N}(z; \mu_{\phi}(x), \Sigma_{\phi}(x))$

Approximate joint density:  $\tilde{p}_{X,Z}(x, z) := q_{\phi}(z|x) p_X(x) \approx p_{X,Z}(x, z)$

Joint Posterior: (log-quadratic in  $x$ )

$$\begin{aligned} J_1(x, z) &:= -\log p_{X,Z|Y}(x, z|y) \\ &= -\log p_{Y|X,Z}(y|x, z) p_{\theta}(x|z) p_Z(z) \\ &= F(x, y) + \underbrace{\frac{1}{2\gamma} \|x - \mu_{\theta}(z)\|^2}_{H_{\theta}(x,z)} + \frac{1}{2} \|z\|^2. \end{aligned} \quad (4)$$

## VAEs and Joint Posterior

**Variational AutoEncoders (VAEs)** (KINGMA AND WELLING, 2013)Generative model:  $p_{X|Z}(x|z) = p_{\theta}(x|z) = \mathcal{N}(x; \mu_{\theta}(z), \gamma Id)$ Joint density:  $p_{X,Z}(x, z) = p_{\theta}(x|z) p_Z(z)$ Approximate inverse:  $p_{Z|X}(z|x) \approx q_{\phi}(z|x) = \mathcal{N}(z; \mu_{\phi}(x), \Sigma_{\phi}(x))$ Approximate joint density:  $\tilde{p}_{X,Z}(x, z) := q_{\phi}(z|x) p_X(x) \approx p_{X,Z}(x, z)$ Joint Posterior: (log-quadratic in  $x$ )

$$\begin{aligned}
 J_1(x, z) &:= -\log p_{X,Z|Y}(x, z | y) \\
 &= -\log p_{Y|X,Z}(y | x, z) p_{\theta}(x|z) p_Z(z) \\
 &= F(x, y) + \underbrace{\frac{1}{2\gamma} \|x - \mu_{\theta}(z)\|^2}_{H_{\theta}(x,z)} + \frac{1}{2} \|z\|^2.
 \end{aligned} \tag{4}$$

Approximate Joint Posterior: (log-quadratic in  $z$ )

$$\begin{aligned}
 J_2(x, z) &:= -\log p_{Y|X,Z}(y | x, z) q_{\phi}(z|x) p_X(x) \\
 &= F(x, y) + \underbrace{\frac{1}{2} \|\Sigma_{\phi}^{-1/2}(x)(z - \mu_{\phi}(x))\|^2}_{K_{\phi}(x,z)} + C(x) - \log p_X(x).
 \end{aligned} \tag{5}$$

# Joint Posterior Maximization - Alternate Convex Search

---

**Algorithm 2.1** Joint posterior maximization - exact case

---

**Require:** Measurements  $\mathbf{y}$ , Autoencoder parameters  $\theta, \phi$ , Initial condition  $\mathbf{x}_0$

**Ensure:**  $\hat{\mathbf{x}}, \hat{\mathbf{z}} = \arg \max_{\mathbf{x}, \mathbf{z}} p_{X,Z|Y}(\mathbf{x}, \mathbf{z} | \mathbf{y})$


- 1: **for**  $n := 0$  **to** maxiter **do**
  - 2:    $\mathbf{z}_{n+1} := \arg \min_{\mathbf{z}} J_2(\mathbf{x}_n, \mathbf{z}) = \boldsymbol{\mu}_\phi(\mathbf{x}_n)$  // Quadratic approx
  - 3:    $\mathbf{x}_{n+1} := \arg \min_{\mathbf{x}} J_1(\mathbf{x}, \mathbf{z}_{n+1})$  // Quadratic
  - 4: **end for**
  - 5: **return**  $\mathbf{x}_{n+1}, \mathbf{z}_{n+1}$
- 

## Proposition

If the encoder approximation is exact ( $J_2 = J_1$ ) then

- $J_1$  is biconvex, and following GORSKI ET AL. (2007):
- Algorithm 2.1 is an Alternate Convex Search
- Algorithm 2.1 converges to a critical point

# JPMAP - Accuracy of encoder approximation



Contour plots of  $-\log p_{Z|X}(z|x)$  and  $-\log q_{\phi}(z|x)$  for a fixed  $x$  and for a random 2D subspace in the  $z$  domain.

## JPMAP - Accuracy of encoder approximation




Figure 1. *Encoder approximation*: (a) Contour plots of  $-\log p_\theta(\mathbf{x}|\mathbf{z}) + \frac{1}{2}\|\mathbf{z}\|^2$  and  $-\log q_\phi(\mathbf{z}|\mathbf{x})$  for a fixed  $\mathbf{x}$  and for a random 2D subspace in the  $\mathbf{z}$  domain (the plot shows  $\pm 2\boldsymbol{\Sigma}_\phi^{1/2}$  around  $\boldsymbol{\mu}_\phi$ ). Observe the relatively small gap between the true posterior  $p_\theta(\mathbf{z}|\mathbf{x})$  and its variational approximation  $q_\phi(\mathbf{z}|\mathbf{x})$ . This figure shows some evidence of partial  $\mathbf{z}$ -convexity of  $J_1$  around the minimum of  $J_2$ , but it does not show how far is  $\mathbf{z}^1$  from  $\mathbf{z}^2$ . (b) Decoded exact optimum  $\mathbf{x}_1 = \boldsymbol{\mu}_\theta \left( \arg \max_{\mathbf{z}} p_\theta(\mathbf{x}|\mathbf{z}) e^{\frac{1}{2}\|\mathbf{z}\|^2} \right)$ . (c) Decoded approximate optimum  $\mathbf{x}_2 = \boldsymbol{\mu}_\theta \left( \arg \max_{\mathbf{z}} q_\phi(\mathbf{z}|\mathbf{x}) \right)$ . (d) Difference between (b) and (c)



# Joint Posterior Maximization - approximate case

---

**Algorithm 2.2** Joint posterior maximization - approximate case

---

**Require:** Measurements  $\mathbf{y}$ , Autoencoder parameters  $\theta, \phi$ , Initial conditions  $\mathbf{x}_0, \mathbf{z}_0$

**Ensure:**  $\hat{\mathbf{x}}, \hat{\mathbf{z}} = \arg \max_{\mathbf{x}, \mathbf{z}} p_{X, Z | Y}(\mathbf{x}, \mathbf{z} | \mathbf{y})$

```
1: for  $n := 0$  to maxiter do
2:    $\mathbf{z}^1 := \arg \min_{\mathbf{z}} J_2(\mathbf{x}_n, \mathbf{z}) = \boldsymbol{\mu}_\phi(\mathbf{x}_n)$  // Quadratic approx
3:    $\mathbf{z}^2 := \text{GD}_{\mathbf{z}} J_1(\mathbf{x}_n, \mathbf{z})$ , starting from  $\mathbf{z} = \mathbf{z}^1$ 
4:    $\mathbf{z}^3 := \text{GD}_{\mathbf{z}} J_1(\mathbf{x}_n, \mathbf{z})$ , starting from  $\mathbf{z} = \mathbf{z}_n$ 
5:   for  $i := 1$  to 3 do
6:      $\mathbf{x}^i := \arg \min_{\mathbf{x}} J_1(\mathbf{x}, \mathbf{z}^i)$  // Quadratic
7:   end for
8:    $i^* := \arg \min_{i \in \{1, 2, 3\}} J_1(\mathbf{x}^i, \mathbf{z}^i)$ 
9:    $(\mathbf{x}_{n+1}, \mathbf{z}_{n+1}) := (\mathbf{x}^{i^*}, \mathbf{z}^{i^*})$ 
10: end for
11: return  $\mathbf{x}_{n+1}, \mathbf{z}_{n+1}$ 
```

---

## Joint Posterior Maximization - approximate case

---

**Algorithm 2.3** Joint posterior maximization - approximate case (faster version)

---

**Require:** Measurements  $\mathbf{y}$ , Autoencoder parameters  $\theta$ ,  $\phi$ , Initial condition  $\mathbf{x}_0$ , iterations

$$n_1 \leq n_2 \leq n_{\max}$$

**Ensure:**  $\hat{\mathbf{x}}, \hat{\mathbf{z}} = \arg \max_{\mathbf{x}, \mathbf{z}} p_{X,Z|Y}(\mathbf{x}, \mathbf{z} | \mathbf{y})$ 

```

1: for  $n := 0$  to  $n_{\max}$  do
2:   done := FALSE
3:   if  $n < n_1$  then
4:      $\mathbf{z}^1 := \arg \min_{\mathbf{z}} J_2(\mathbf{x}_n, \mathbf{z}) = \boldsymbol{\mu}_\phi(\mathbf{x}_n)$  // Quadratic approx
5:      $\mathbf{x}^1 := \arg \min_{\mathbf{x}} J_1(\mathbf{x}, \mathbf{z}^1)$  // Quadratic
6:     if  $J_1(\mathbf{x}^1, \mathbf{z}^1) < J_1(\mathbf{x}_n, \mathbf{z}_n)$  then
7:        $i^* := 1$  // Faster alternative while  $J_2$  is good enough
8:       done := TRUE
9:     end if
10:  end if
11:  if not done and  $n < n_2$  then
12:     $\mathbf{z}^1 := \arg \min_{\mathbf{z}} J_2(\mathbf{x}_n, \mathbf{z}) = \boldsymbol{\mu}_\phi(\mathbf{x}_n)$  // Quadratic approx
13:     $\mathbf{z}^2 := \text{GD}_{\mathbf{z}} J_1(\mathbf{x}_n, \mathbf{z})$ , starting from  $\mathbf{z} = \mathbf{z}^1$ 
14:     $\mathbf{x}^2 := \arg \min_{\mathbf{x}} J_1(\mathbf{x}, \mathbf{z}^2)$  // Quadratic
15:    if  $J_1(\mathbf{x}^2, \mathbf{z}^2) < J_1(\mathbf{x}_n, \mathbf{z}_n)$  then
16:       $i^* := 2$  //  $J_2$  init is good enough
17:      done := TRUE
18:    end if
19:  end if
20:  if not done then
21:     $\mathbf{z}^3 := \text{GD}_{\mathbf{z}} J_1(\mathbf{x}_n, \mathbf{z})$ , starting from  $\mathbf{z} = \mathbf{z}_n$ 
22:     $\mathbf{x}^3 := \arg \min_{\mathbf{x}} J_1(\mathbf{x}, \mathbf{z}^3)$  // Quadratic
23:     $i^* := 3$ 
24:  end if
25:   $(\mathbf{x}_{n+1}, \mathbf{z}_{n+1}) := (\mathbf{x}^{i^*}, \mathbf{z}^{i^*})$ 
26: end for
27: return  $\mathbf{x}_{n+1}, \mathbf{z}_{n+1}$ 

```



---

# JPMAP - Effectiveness of the encoder initialization

Trajectories of  $\text{GD}_z J_1(\mathbf{x}_0, \mathbf{z})$ , starting from  $\mathbf{z} = \mathbf{z}_0$

Thick blue curve:  $\mathbf{z}_0 = \arg \min_z J_2(\mathbf{x}_0, \mathbf{z}) = \mu_\phi(\mathbf{x}_0)$

Thin curves: random initializations  $\mathbf{z}_0 \sim \mathcal{N}(0, Id)$



# JPMAP - Convergence

If we use ELU activations then the following assumption is verified:

## Assumption (2)

$J_1(\cdot, \mathbf{z})$  is convex and admits a minimizer for any  $\mathbf{z}$ . Moreover,  $J_1$  is coercive and continuously differentiable.


## Proposition (Convergence of Algorithm 2.3)

Let  $\{(\mathbf{x}_n, \mathbf{z}_n)\}$  be a sequence generated by Algorithm 2.3. Under Assumption 2 we have that:

- 1 The sequence  $\{J_1(\mathbf{x}_n, \mathbf{z}_n)\}$  converges monotonically when  $n \rightarrow \infty$ .
- 2 The sequence  $\{(\mathbf{x}_n, \mathbf{z}_n)\}$  has at least one accumulation point.
- 3 All accumulation points of  $\{(\mathbf{x}_n, \mathbf{z}_n)\}$  are stationary points of  $J_1$  and they all have the same function value.

# Denosing Criterion to train VAEs (IM ET AL., 2017)

## Variational AutoEncoders (VAEs) (KINGMA AND WELLING, 2013)




**Learning:** Maximize the averaged *Evidence Lower Bound* (ELBO) for  $\mathbf{x} \in \mathcal{D}$

$$\mathcal{L}_{\theta, \phi}(\mathbf{x}) = \mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{x})}[\log p_{\theta}(\mathbf{x}|\mathbf{z})] - KL(q_{\phi}(\mathbf{z}|\mathbf{x}) \parallel p_{\mathbf{z}}(\mathbf{z})) \leq \log p_{\theta}(\mathbf{x}).$$

# Denoising Criterion to train VAEs (IM ET AL., 2017)

## Variational AutoEncoders (VAEs) (KINGMA AND WELLING, 2013)



**Learning:** Maximize the averaged *Evidence Lower Bound (ELBO)* for  $\mathbf{x} \in \mathcal{D}$


$$\mathcal{L}_{\theta, \phi}(\mathbf{x}) = \mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{x})}[\log p_{\theta}(\mathbf{x}|\mathbf{z})] - KL(q_{\phi}(\mathbf{z}|\mathbf{x}) \parallel p_{\mathbf{z}}(\mathbf{z})) \leq \log p_{\theta}(\mathbf{x}).$$

**Problem:**  $\mu_{\phi}(\mathbf{x})$  only trained for  $\mathbf{x} \in \mathcal{D}$  or  $\mathbf{x} \in \mathcal{M} = \mu_{\theta}(\mathbb{R}^m)$ .

**But:** Step 2 in the algorithm evaluates  $\mu_{\phi}(\mathbf{x}_n)$  for degraded  $\mathbf{x}_n \notin \mathcal{M}$

# Denosing Criterion to train VAEs (IM ET AL., 2017)

## Variational AutoEncoders (VAEs) (KINGMA AND WELLING, 2013)



**Learning:** Maximize the averaged *Evidence Lower Bound (ELBO)* for  $x \in \mathcal{D}$

$$\mathcal{L}_{\theta, \phi}(\mathbf{x}) = \mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{x})}[\log p_{\theta}(\mathbf{x}|\mathbf{z})] - \text{KL}(q_{\phi}(\mathbf{z}|\mathbf{x}) \parallel p_{\mathbf{z}}(\mathbf{z})) \leq \log p_{\theta}(\mathbf{x}).$$


**Denosing criterion:** Train on  $\tilde{\mathcal{D}}$  but still require  $\mu_{\theta}(\mu_{\phi}(\tilde{\mathbf{x}})) \approx \mathbf{x}$ .

$$\tilde{\mathcal{D}} = \{\tilde{\mathbf{x}} = \mathbf{x} + \sigma_{\text{DVAE}}\varepsilon : \mathbf{x} \in \mathcal{D} \text{ and } \varepsilon \sim \mathcal{N}(0, I)\}$$


Maximize the denosing ELBO

$$\tilde{\mathcal{L}}_{\theta, \phi}(\mathbf{x}) = \mathbb{E}_{p(\tilde{\mathbf{x}}|\mathbf{x})} \left[ \mathbb{E}_{q_{\phi}(\mathbf{z}|\tilde{\mathbf{x}})}[\log p_{\theta}(\mathbf{x}|\mathbf{z})] - \text{KL}(q_{\phi}(\mathbf{z}|\tilde{\mathbf{x}}) \parallel p_{\mathbf{z}}(\mathbf{z})) \right]$$


# Denosing criterion does not degrade generative model



(a) Denoising




(b) Compressed Sensing




(c) Inpainting

**Figure 1.** Evaluating the quality of the generative model as a function of  $\sigma_{\text{DVAE}}$ . On (a) Denoising (Gaussian noise  $\sigma = 150$ ), (b) Compressed Sensing ( $\sim 10.2\%$  measurements, noise  $\sigma = 10$ ) and (c) Inpainting (80% of missing pixels, noise  $\sigma = 10$ ). Results of both algorithms are computed on a batch of 50 images and initialising on ground truth  $\mathbf{x}^*$  (for CSGM we use  $\mathbf{z}_0 = \boldsymbol{\mu}_\phi(\mathbf{x}^*)$ ).




Optimal value of  $\sigma_{\text{DVAE}}$ 

(a) Denoising



(b) Compressed Sensing



(c) Inpainting

**Figure 2.** Evaluating the effectiveness of JP MAP vs CSGM as a function of  $\sigma_{\text{DVAE}}$  (same setup of Figure 1). Without a denoising criterion  $\sigma_{\text{DVAE}} = 0$  the JP MAP algorithm may provide wrong guesses  $\mathbf{z}^1$  when applying the encoder in step 2 of Algorithm 2.2. For  $\sigma_{\text{DVAE}} > 0$  however, the alternating minimization algorithm can benefit from the robust initialization heuristics provided by the encoder, and it consistently converges to a better local optimum than the simple gradient descent in CSGM.

# MAP-z as the limit case for $\beta \rightarrow \infty$

Two options for MAP-z estimator instead of the joint MAP-x-z

- 1 CSGM - gradient descent, may be stuck in local minima
- 2 Use Algorithm 2.3 to solve a series of joint MAP-x-z problems with increasing values of  $\beta = \frac{1}{\gamma} \rightarrow \infty$  as suggested in Algorithm 1.1.

Stopping criterion: Inequality constrained problem

$$\arg \min_{\mathbf{x}, \mathbf{z}: \|G(\mathbf{z}) - \mathbf{x}\|^2 \leq \varepsilon} F(\mathbf{x}, \mathbf{y}) + \frac{1}{2} \|\mathbf{z}\|^2.$$

The corresponding Lagrangian form is

$$\max_{\beta} \min_{\mathbf{x}, \mathbf{z}} F(\mathbf{x}, \mathbf{y}) + \frac{1}{2} \|\mathbf{z}\|^2 + \beta (\|G(\mathbf{z}) - \mathbf{x}\|^2 - \varepsilon)^+ \quad (6)$$

We use the exponential multiplier method (Tseng and Bertsekas, 1993) to guide the search for the optimal value of  $\beta$  (see Algorithm 2.4)

MAP-z as the limit case for  $\beta \rightarrow \infty$ 


---

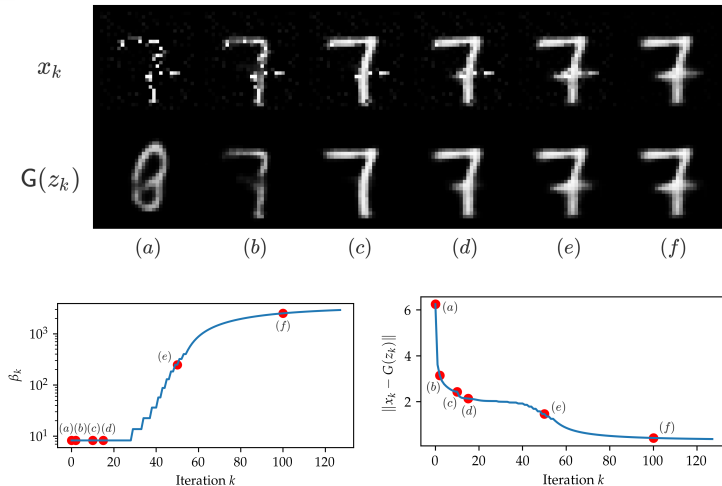
**Algorithm 2.4** MAP-z as the limit of joint MAP-x-z.

---

**Require:** Measurements  $\mathbf{y}$ , Tolerance  $\varepsilon$ , Rate  $\rho > 0$ , Initial  $\beta_0$ , Initial  $\mathbf{x}_0$ , Iterations  $0 \leq n_1 \leq n_2 \leq n_{\max}$



**Ensure:**  $\arg \min_{\mathbf{z}: \|\mathbf{G}(\mathbf{z}) - \mathbf{x}\|^2 \leq \varepsilon} F(\mathbf{x}, \mathbf{y}) + \frac{1}{2} \|\mathbf{z}\|^2$ .

- 1:  $\beta := \beta_0$
  - 2:  $\mathbf{x}^0, \mathbf{z}^0 :=$  Algorithm 2.3 starting from  $\mathbf{x} = \mathbf{x}_0$  with  $\beta, n_1, n_2, n_{\max}$ .
  - 3: converged := FALSE
  - 4:  $k := 0$
  - 5: **while not** converged **do**
  - 6:    $\mathbf{x}^{k+1}, \mathbf{z}^{k+1} :=$  Algorithm 2.3 starting from  $\mathbf{x} = \mathbf{x}^k$  with  $\beta$  and  $n_1 = n_2 = 0$
  - 7:    $C = \|\mathbf{G}(\mathbf{z}^{k+1}) - \mathbf{x}^{k+1}\|^2 - \varepsilon$
  - 8:    $\beta := \beta \exp(\rho C)$
  - 9:   converged := ( $C \leq 0$ )
  - 10:    $k := k + 1$
  - 11: **end while**
  - 12: **return**  $\mathbf{x}^k, \mathbf{z}^k$
-



MAP-z as the limit case for  $\beta \rightarrow \infty$ 

**Figure 5.** Evolution of Algorithm 2.4. In this inpainting example, JPMAP starts with the initialization in (a). During first iterations (b) – (d) where  $\beta_k$  is small,  $\mathbf{x}_k$  and  $G(\mathbf{z}_k)$  start loosely approaching each other



# Denoising experiments (MNIST)




# Compressed sensing experiments (MNIST)



# Inpainting experiments (MNIST)




Denoising experiment:  $\sigma = 110/255$ 




# Compressed sensing experiment: $m = 140$ random measurements




## Inpainting experiment: 80% missing pixels



# Inpainting experiment: 80% missing pixels $\sigma = 10/255$ (CelebA)



From top to bottom: original image  $\mathbf{x}^*$ , corrupted image  $\tilde{\mathbf{x}}$ , restored by CSGM, restored image  $\hat{\mathbf{x}}$  by our framework.

CelebA reconstructions  $\mu_{\theta}(\mu_{\phi}(\mathbf{x}))$ 

Reconstructions  $\mu_{\theta}(\mu_{\phi}(\mathbf{x}))$  (even columns) for some test samples  $\mathbf{x}$  (odd columns), showing the over-regularization of data manifold imposed by the trained VAE. As a consequence,  $-\log p_{Z|Y}(\mathbf{z} | \mathbf{y})$  does not have as many local minima and then a simple gradient

# Conclusion

- JPMAP avoids spurious local minima thanks to
  - Quasi bi-convex optimization
  - Encoder initialization
  - Denoising VAE
  - Splitting and continuation scheme
- JPMAP converges for all quadratic problems and regularisation parameters (unlike denoiser-based PnP approaches (RYU ET AL., 2019) that are more restrictive)
- Constraints
  - Fixed size
  - VAEs lag behind GANs

## Future work

- Use a more powerful VAE like NVAE (VAHDAT AND KAUTZ, 2020) or TwoStageVAE (DAI AND WIPF, 2019) or VDVAE.
- ... more to come ...

Preprint and code available here  
<http://up5.fr/jpmap>  
<https://arxiv.org/abs/2103.04715>

Thank you for your attention!

Questions? Comments

# Future Work & Open Questions

- Resizable explicit priors / regularizers ?
  - Fully convolutional generative models like Glow (KINGMA AND DHARIWAL, 2018) ? model guarantees after resize ?
  - Patch-based approach (HELMINGER ET AL., 2020; PROST ET AL., 2021)

$$R(x) = \sum_i r(p_i(x))$$

- Generalization of JPMAP to
  - invertible generative models other than VAE: Normalizing Flows?
  - posterior sampling
    - pCN (HOLDEN ET AL., 2021) uses generative model but not inverse
    - Other sampling schemes using the inverse should be faster ?
- How does PnP ULA compare to SRFlow (LUGMAYR ET AL., 2020) (NF  $f$  trained to learn the posterior of a particular inverse problem, i.e. if  $n \sim N(0, Id)$  then  $f(y, n) \sim p(x|y)$ ).
  - First comparative study in (ANDRIE ET AL., 2021)



# Patch-based regularization (PROST ET AL., 2021)

## Image inverse problem

$$y = Ax + \epsilon$$


$$\epsilon \sim \mathcal{N}(0, \sigma^2 I)$$

## Variational problem


$$\hat{x} = \arg \min_x \underbrace{F(x, y)}_{\frac{1}{2\sigma^2} \|Ax - y\|^2} + \lambda R(x)$$







## From local to global regularization







$$R(x) = \frac{1}{|\Omega_x|} \sum_{u \in \Omega_x} r_\theta(u)$$



## Results



-  Andrieu, Anna, Nando Fardijn, Paul Hagemann, Sebastian Heidenreich, Victor Soltwisch, and Gabriele Steidl (2021). “Invertible Neural Networks Versus MCMC for Posterior Reconstruction in Grazing Incidence X-Ray Fluorescence”. In: *SSVM 2021, LNCS*. Vol. 12679 LNCS, pp. 528–539. ISBN: 9783030755485. DOI: 10.1007/978-3-030-75549-2\_42. arXiv: 2102.03189 (cit. on p. 56).
-  Arjovsky, Martin and Léon Bottou (2017). “Towards Principled Methods for Training Generative Adversarial Networks”. In: *(ICLR) International Conference on Learning Representations*, pp. 1–17 (cit. on pp. 18–20).
-  Bora, Ashish, Ajil Jalal, Eric Price, and Alexandros G Dimakis (2017). “Compressed sensing using generative models”. In: *(ICML) International Conference on Machine Learning*. Vol. 2. JMLR. org, pp. 537–546. ISBN: 9781510855144. arXiv: arXiv:1703.03208v1 (cit. on p. 21).
-  Chambolle, A (2004). “An algorithm for total variation minimization and applications”. In: *Journal of Mathematical Imaging and Vision* 20, pp. 89–97. DOI: 10.1023/B:JMIV.0000011325.36760.1e (cit. on p. 5).
-  Dai, Bin and David Wipf (2019). “Diagnosing and Enhancing VAE Models”. In: *ICLR*, pp. 1–42. arXiv: 1903.05789 (cit. on p. 54).
-  Efron, Bradley (2011). “Tweedie’s Formula and Selection Bias”. In: *Journal of the American Statistical Association* 106.496, pp. 1602–1614. ISSN: 0162-1459. DOI: 10.1198/jasa.2011.tm11181 (cit. on p. 12).

-  González, Mario, Andrés Almansa, and Pauline Tan (2021). “Solving Inverse Problems by Joint Posterior Maximization with Autoencoding Prior”. In: *arXiv: 2103.01648*.
-  Goodfellow, Ian J., Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio (2014). “Generative Adversarial Networks”. In: *Advances in Neural Information Processing Systems 27*, pp. 2672–2680. ISSN: 10495258. arXiv: 1406.2661 (cit. on pp. 18–20, 24, 25).
-  Gorski, Jochen, Frank Pfeuffer, and Kathrin Klamroth (2007). “Biconvex sets and optimization with biconvex functions: a survey and extensions”. In: *Mathematical Methods of Operations Research* 66.3, pp. 373–407. ISSN: 1432-2994. DOI: 10.1007/s00186-007-0161-1 (cit. on p. 30).
-  Helminger, Leonhard, Michael Bernasconi, Abdelaziz Djelouah, Markus Gross, and Christopher Schroers (2020). *Blind Image Restoration with Flow Based Priors*. Tech. rep. arXiv: 2009.04583 (cit. on p. 56).
-  Holden, Matthew, Marcelo Pereyra, and Konstantinos C. Zygalakis (2021). “Bayesian Imaging With Data-Driven Priors Encoded by Neural Networks: Theory, Methods, and Algorithms”. In: pp. 1–22. arXiv: 2103.10182 (cit. on p. 56).
-  Im, Daniel Jiwoong, Sungjin Ahn, Roland Memisevic, and Yoshua Bengio (2017). “Denoising criterion for variational auto-encoding framework”. In: *31st AAAI Conference on Artificial Intelligence, AAAI 2017*. AAAI press, pp. 2059–2065. arXiv: 1511.06406 (cit. on pp. 37–39).

-  Kingma, Diederik P. and Prafulla Dhariwal (2018). “Glow: Generative Flow with Invertible 1x1 Convolutions”. In: *(NeurIPS) Advances in Neural Information Processing Systems 2018-Decem.2*, pp. 10215–10224. ISSN: 10495258. arXiv: 1807.03039 (cit. on p. 56).
-  Kingma, Diederik P and Max Welling (2013). “Auto-Encoding Variational Bayes”. In: *(ICLR) International Conference on Learning Representations. MI*, pp. 1–14. ISBN: 1312.6114v10. DOI: 10.1051/0004-6361/201527329. arXiv: 1312.6114 (cit. on pp. 25–29, 37–39).
-  Laumont, Rémi, Valentin de Bortoli, Andrés Almansa, Julie Delon, Alain Durmus, and Marcelo Pereyra (2021). “Bayesian imaging using Plug & Play priors: when Langevin meets Tweedie”. In: arXiv: 2103.04715 (cit. on pp. 13, 14).
-  Louchet, Cécile and Lionel Moisan (2013). “Posterior expectation of the total variation model: Properties and experiments”. In: *SIAM Journal on Imaging Sciences* 6.4, pp. 2640–2684. ISSN: 19364954. DOI: 10.1137/120902276 (cit. on p. 5).
-  Lugmayr, Andreas, Martin Danelljan, Luc Van Gool, and Radu Timofte (2020). “SRFlow: Learning the Super-Resolution Space with Normalizing Flow”. In: *(ECCV) European Conference on Computer Vision. Vol. 12350 LNCS*, pp. 715–732. ISBN: 9783030585570. DOI: 10.1007/978-3-030-58558-7\_42. arXiv: 2006.14200 (cit. on p. 56).

-  Papamakarios, George, Eric Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan (2019). “Normalizing Flows for Probabilistic Modeling and Inference”. In: arXiv: 1912.02762 (cit. on pp. 19, 20).
-  Pereyra, Marcelo (2016). “Proximal Markov chain Monte Carlo algorithms”. In: *Statistics and Computing* 26.4, pp. 745–760. ISSN: 0960-3174. DOI: 10.1007/s11222-015-9567-4. arXiv: 1306.0187 (cit. on p. 5).
-  Prost, Jean, Antoine Houdard, Andrés Almansa, and Nicolas Papadakis (2021). “Learning local regularization for variational image restoration”. In: pp. 1–12. arXiv: 2102.06155 (cit. on pp. 56, 57).
-  Rudin, Leonid I., Stanley Osher, and Emad Fatemi (1992). “Nonlinear total variation based noise removal algorithms”. In: *Physica D: Nonlinear Phenomena* 60.1-4, pp. 259–268. ISSN: 01672789. DOI: 10.1016/0167-2789(92)90242-F (cit. on p. 5).
-  Ryu, Ernest K., Jialin Liu, Sicheng Wang, Xiaohan Chen, Zhangyang Wang, and Wotao Yin (2019). “Plug-and-Play Methods Provably Converge with Properly Trained Denoisers”. In: *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, pp. 5546–5557. arXiv: 1905.05406 (cit. on pp. 9, 11, 53).
-  Teodoro, Afonso M., José M. Bioucas-Dias, and Mário A. T. Figueiredo (2018). *Scene-Adapted Plug-and-Play Algorithm with Guaranteed Convergence: Applications to Data Fusion in Imaging*. arXiv: 1801.00605 (cit. on p. 5).



Tseng, Paul and Dimitri P. Bertsekas (1993). “On the convergence of the exponential multiplier method for convex programming”. In: *Mathematical Programming* 60.1-3, pp. 1–19. ISSN: 00255610. DOI: 10.1007/BF01580598 (cit. on p. 42).



Vahdat, Arash and Jan Kautz (2020). “Nvae: A deep hierarchical variational autoencoder”. In: *Advances in Neural Information Processing Systems* 33 (cit. on p. 54).



Yu, Guoshen, Guillermo Sapiro, and Stéphane Mallat (2011). “Solving inverse problems with piecewise linear estimators: From Gaussian mixture models to structured sparsity”. In: *IEEE Transactions on Image Processing* 21.5, pp. 2481–2499 (cit. on p. 5).



Zoran, Daniel and Yair Weiss (2011). “From learning models of natural image patches to whole image restoration”. In: *2011 International Conference on Computer Vision*. IEEE, pp. 479–486. ISBN: 978-1-4577-1102-2. DOI: 10.1109/ICCV.2011.6126278 (cit. on p. 5).