

Filling the gaps: a story of priors and conditional probabilities

Thomas Moreau INRIA Saclay



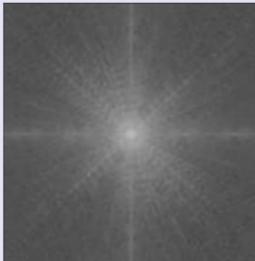
Inverse Problems

Imaging

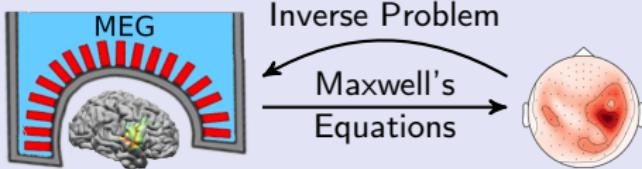


Super-Resolution, Inpainting,
Deblurring, ...

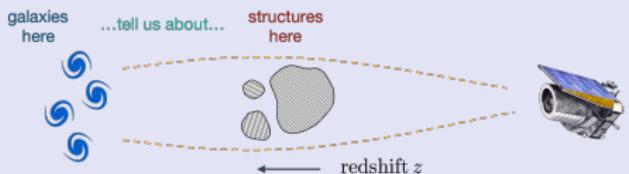
Neuroimaging – MRI



Neuroimaging – M/EEG



Astrophysics



Seismology – Prospection



Inverse Problems

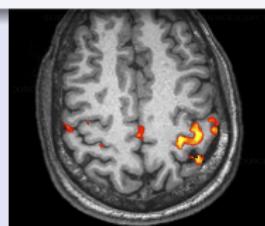
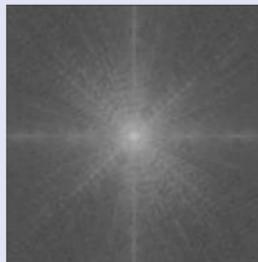
Imaging



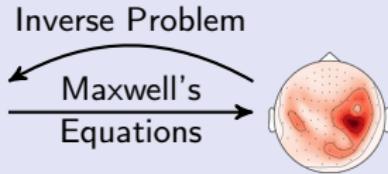
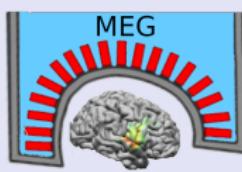
Super-Resolution
Deblurring



Neuroimaging



Neuroimaging – M/EEG

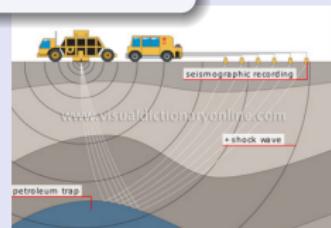


Forward model:

$$y = Ax + \varepsilon$$

Inverse problem:

find x from y



MAP estimate as a regularized regression problem

$$\mathbf{x}^*(\mathbf{y}; \theta) = \underset{\mathbf{x}}{\operatorname{argmin}} -\log p(\mathbf{x}|\mathbf{y}; \theta) = \underbrace{\frac{1}{2} \|\mathbf{y} - \mathbf{Ax}\|_2^2}_{-\log p(\mathbf{y}|\mathbf{x})} + \underbrace{\mathcal{R}(\mathbf{x}; \theta)}_{-\log p(\mathbf{x}; \theta)}$$

where \mathcal{R} encodes prior information to select a good/plausible solution

MAP estimate as a regularized regression problem

$$\mathbf{x}^*(\mathbf{y}; \theta) = \underset{\mathbf{x}}{\operatorname{argmin}} -\log p(\mathbf{x}|\mathbf{y}; \theta) = \underbrace{\frac{1}{2} \|\mathbf{y} - \mathbf{Ax}\|_2^2}_{-\log p(\mathbf{y}|\mathbf{x})} + \underbrace{\mathcal{R}(\mathbf{x}; \theta)}_{-\log p(\mathbf{x}; \theta)}$$

where \mathcal{R} encodes prior information to select a good/plausible solution

Common framework:

- ▶ **Efficient solvers:** Forward backward, ADMM, HQS ...
 - ⇒ But might require many iterations to get good results
- ▶ **Flexible:** Can choose many priors – *handpicked, learned, implicit, ...*
 - ⇒ Quality of the solution depends on the prior's choice $p(\mathbf{x}; \theta)$

Defining a the prior $p(\mathbf{x}; \theta)$

Explicitely: choose a log-prior \mathcal{R} promoting certain properties of the signal.

TV, sparsity, wavelets, dictionary, ...

Interpretable, no learning, convergence guarantees

Defining a the prior $p(\mathbf{x}; \theta)$

Explicitely: choose a log-prior \mathcal{R} promoting certain properties of the signal.

TV, sparsity, wavelets, dictionary, ...

Interpretable, no learning, convergence guarantees

Implicitely: we only need the prior's score $\nabla \mathcal{R}$ or its prox

Pnp/RED frameworks with denoisers

Linked to the score with the Tweedie's formula

Defining a the prior $p(\mathbf{x}; \theta)$

Explicitely: choose a log-prior \mathcal{R} promoting certain properties of the signal.

TV, sparsity, wavelets, dictionary, ...

Interpretable, no learning, convergence guarantees

Implicitely: we only need the prior's score $\nabla \mathcal{R}$ or its prox

Pnp/RED frameworks with denoisers

Linked to the score with the Tweedie's formula



x



$y = x + e$



$\text{prox}_{\text{TV}}(y)$



$\text{DRUNet}(y)$

Desirable properties of the prior

How to choose the structure of the prior $p(\mathbf{x}; \theta)$?

- ▶ **Adapted to the data:** capture the distribution of the data
- ▶ **Easy to learn:** few data, fast training, no overfitting
- ▶ **Fast to compute:** used in an iterative algorithm
- ▶ **Interpretable:** understand the role of the prior and its convergence guarantees
- ▶ **Adapt to the task:** allow to recover the information lost by **A**

Desirable properties of the prior

How to choose the structure of the prior $p(\mathbf{x}; \theta)$?

- ▶ **Adapted to the data:** capture the distribution of the data
- ▶ **Easy to learn:** few data, fast training, no overfitting
- ▶ **Fast to compute:** used in an iterative algorithm
- ▶ **Interpretable:** understand the role of the prior and its convergence guarantees
- ▶ **Adapt to the task:** allow to recover the information lost by **A**

The role of the structure in unsupervised prior learning

References

- ▶ Malézieux, B., Michel, F., **TM**, and Kowalski, M. (2024). [Where prior learning can and can't work in unsupervised inverse problems](#). Preprint

The goal of prior learning

The goal of prior learning is to **recover the information lost by the operator A** , while the data fidelity term ensures consistency with the observed measurements.

- ▶ **Data fidelity:** Ensures that the solution matches the known measurements (what is observed).
- ▶ **Prior:** Completes the missing information (what is lost by A), guiding the solution towards plausible images.

⇒ A well-chosen prior needs to recover $P_{\text{ker}A}(x)$ from the observations (either y or $A^\dagger y$).

Prior selection and learning

Supervised learning of the prior

Evaluate the quality of the prior on a dataset (\mathbf{y}, \mathbf{x})

$$\min_{\theta} \mathbb{E}_{(\mathbf{y}, \mathbf{x})} \frac{1}{2} \|\mathbf{x} - \mathbf{x}^*(\mathbf{y}; \theta)\|_2^2 \quad s.t. \quad \mathbf{x}^*(\mathbf{y}; \theta) = \operatorname{argmin}_{\mathbf{y}} -\log p(\mathbf{x}|\mathbf{y}; \theta)$$

Task Agnostic learning

Characterize the distribution of the data $p(\mathbf{x})$ by training a denoiser

$$\min_{\theta} \mathbb{E}_{\mathbf{x}} \frac{1}{2} \|\mathbf{x} - \mathbf{x}^*(\mathbf{x} + \epsilon; \theta)\|_2^2$$

Self-supervised learning

[Tachella et al., 2022]

Learn with incomplete measurements \mathbf{y} with equivariance, consistency, . . .

$$\min_{\theta} \mathbb{E}_{\mathbf{y}=(\mathbf{y}_1, \mathbf{y}_2)} \frac{1}{2} \|\mathbf{y}_2 - \mathbf{A} \mathbf{x}^*(\mathbf{y}_1; \theta)\|_2^2$$

Prior selection and learning

Supervised learning of the prior

Evaluate the quality of the prior on a dataset (\mathbf{y}, \mathbf{x})

$$\min_{\theta} \mathbb{E}_{(\mathbf{y}, \mathbf{x})} \frac{1}{2} \|\mathbf{x} - \mathbf{x}^*(\mathbf{y}; \theta)\|_2^2 \quad s.t. \quad \mathbf{x}^*(\mathbf{y}; \theta) = \operatorname{argmin}_{\mathbf{y}} -\log p(\mathbf{x}|\mathbf{y}; \theta)$$

Task Agnostic learning

Characterize the distribution of the data $p(\mathbf{x})$ by training a denoiser

$$\min_{\theta} \mathbb{E}_{\mathbf{x}} \frac{1}{2} \|\mathbf{x} - \mathbf{x}^*(\mathbf{x} + \epsilon; \theta)\|_2^2$$

Self-supervised learning

[Tachella et al., 2022]

Learn with incomplete measurements \mathbf{y} with equivariance, consistency, . . .

$$\min_{\theta} \mathbb{E}_{\mathbf{y}=(\mathbf{y}|_1, \mathbf{y}|_2)} \frac{1}{2} \|\mathbf{y}|_2 - \mathbf{A}|_2 \mathbf{x}^*(\mathbf{y}|_1; \theta)\|_2^2$$

Dictionary-based priors: a tool to study prior learning

We consider the problem of learning a dictionary D to solve inverse problems with a sparse prior with a synthesis formulation

$$x^*(y; D) = D \left(\operatorname*{argmin}_z \frac{1}{2} \|y - ADz\|_2^2 + \lambda \|z\|_1 \right)$$

- ▶ Explicit prior parameterization with the dictionary D
- ▶ Can generate data according to the model
- ▶ Can study the dynamic of learning D

Single measurement dictionary learning

With no extra constraint, if the dictionary is learned in an unsupervised manner, the dictionary cannot recover any information lost by \mathbf{A}

$$\mathbf{D} = \operatorname{argmin}_{\|\mathbf{D}\|_2 \leq 1} \operatorname{argmin}_{\mathbf{z}} \frac{1}{2} \|\mathbf{y} - \mathbf{A}\mathbf{D}\mathbf{z}\|_2^2 + \lambda \|\mathbf{z}\|_1 \implies \mathbf{D} \in \ker(\mathbf{A})^\perp$$

\Rightarrow The dictionary is null in the null space of \mathbf{A}

Therefore, the prior cannot help in solving the inverse problem.

Multiple measurement dictionary learning

When learning with multiple operators \mathbf{A}_i , the dictionary can recover some information lost by each operator

$$\mathbf{D} = \operatorname{argmin}_{\|\mathbf{D}\|_2 \leq 1} \sum_i \operatorname{argmin}_{\mathbf{z}_i} \frac{1}{2} \|\mathbf{y}_i - \mathbf{A}_i \mathbf{D} \mathbf{z}_i\|_2^2 + \lambda \|\mathbf{z}_i\|_1$$

Here, an interesting case is when the operators \mathbf{A}_i are *incomplete* but their union is complete, i.e. $\bigcap_i \ker(\mathbf{A}_i) = \{0\}$

In this case, the dictionary can recover information lost by each operator and perform well in solving the inverse problem.

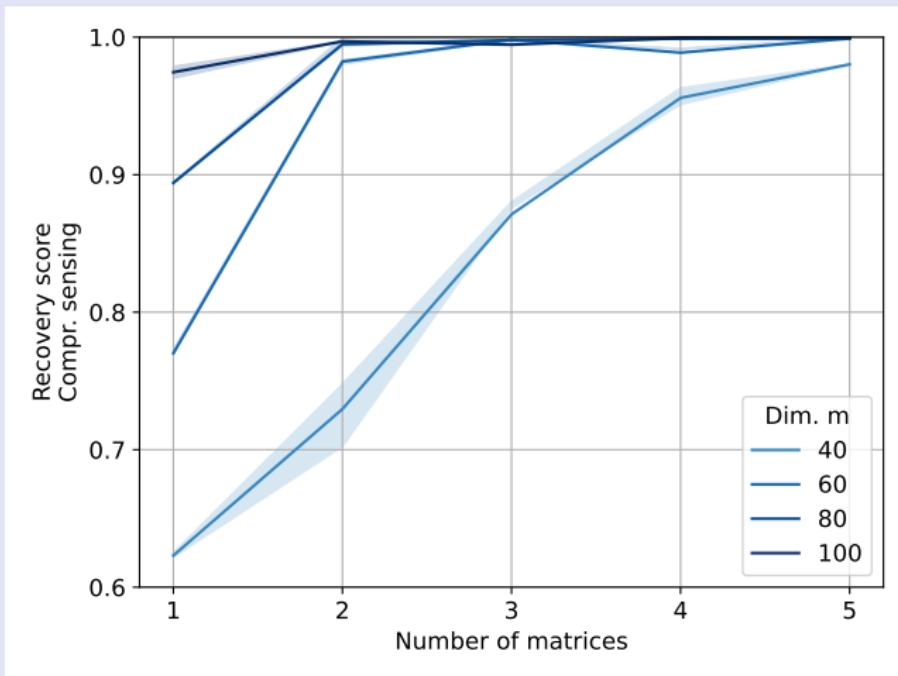
This is similar results as for unsupervised inverse problems training loss.

[Tachella et al., 2022]

Prior recovery with incomplete operators – Compressed Sensing

Recovery of a \mathbf{D} generated as a 100×100 normalized Gaussian dictionary

\mathbf{A}_i is a random $m \times 100$ sensing matrix, with Bernouilli-gaussian signals \mathbf{z}_i



Convolutional Priors in Inverse Problems

Convolutional structures are widely used as priors for inverse problems.

We can consider Convolutional Dictionary Learning as a simple model of convolutional priors:

$$\mathbf{x}^*(\mathbf{y}; \mathbf{D}) = \mathbf{D} * \operatorname{argmin}_{\mathbf{z}} \frac{1}{2} \|\mathbf{y} - \mathbf{A}(\mathbf{D} * \mathbf{z})\|_2^2 + \lambda \|\mathbf{z}\|_1$$

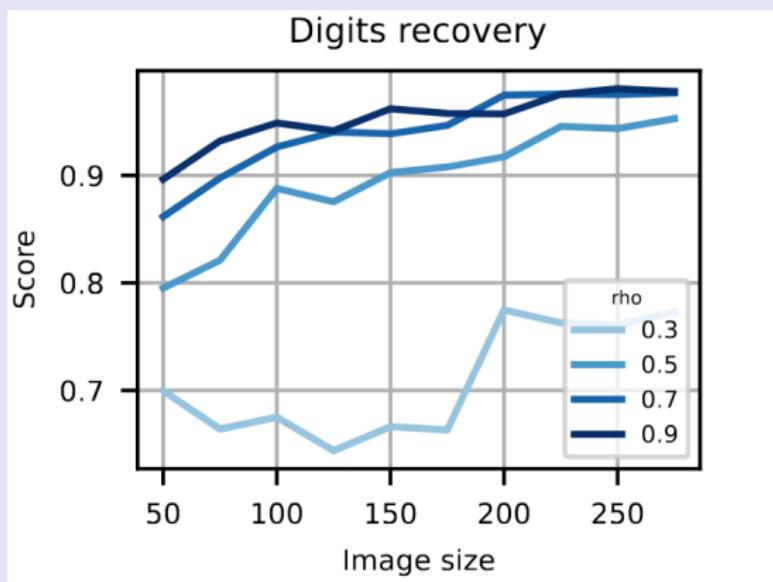
- ▶ Convolutional structure encodes translation invariance.
- ▶ The dictionary atoms capture local spatial patterns.
- ▶ Constrain local context to be similar to learned filters.

⇒ But how does this structure impact different inverse problems?

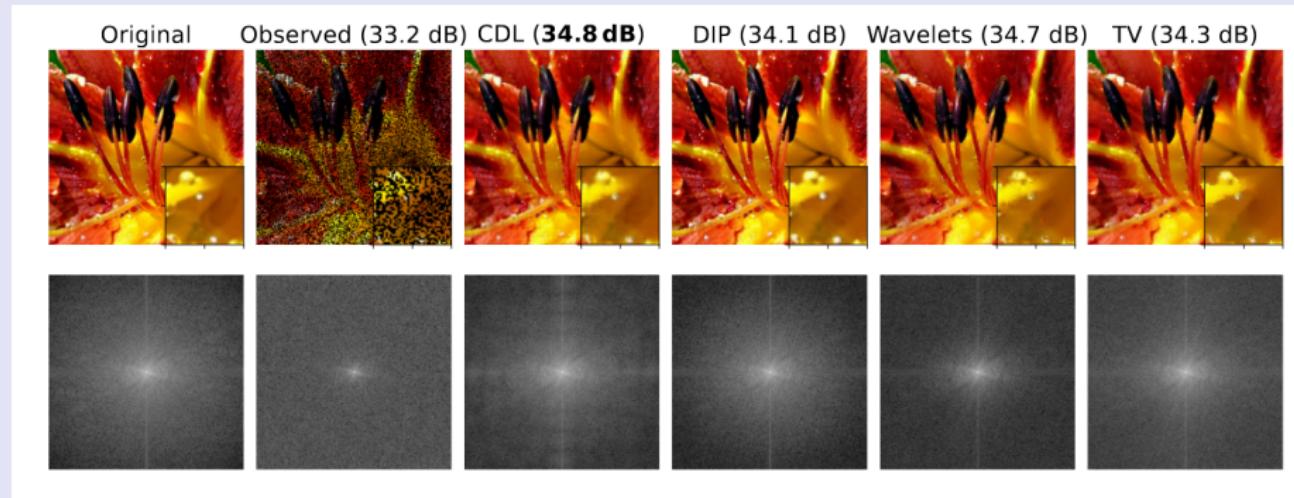
Inpainting: Leveraging Local Structure

Inpainting involves recovering missing pixels using surrounding information.

If the signal is stationary, convolutional prior on a single large measurement acts as on multiple measurements for each patch independently.



Inpainting: Leveraging Local Structure

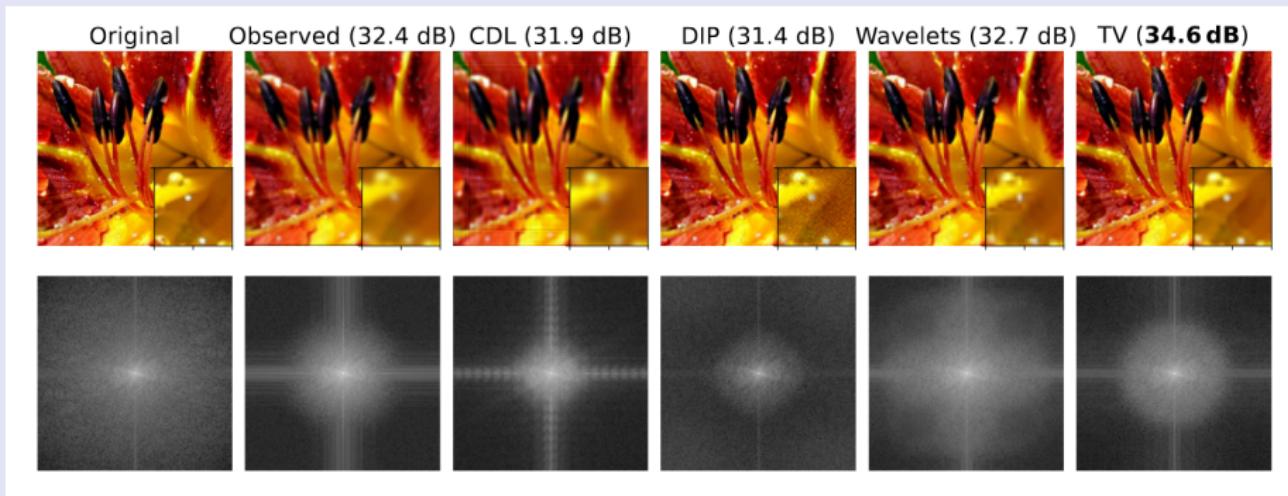


This structure is well adapted as it captures the distribution of local patterns, which is sufficient to fill in missing pixels.

Deblurring: The Challenge for Convolutional Priors

Deblurring aims to recover sharp images from blurred observations.

In this case, the convolutional structure is less adapted, as the kernel of the blur is aligned with the spectral structure of the prior.



For each patch in the image, the kernel is the same, so the convolutional structure does not provide the same advantage as in inpainting.

Takeaways on prior learning

- ▶ Dictionary-based priors are useful tools to study prior learning.
- ▶ Understanding the structure of the kernel of \mathbf{A} is crucial to design or learn a good prior.
- ▶ Our aim should be to find priors that links $P_{\text{Ker}\mathbf{A}}(\mathbf{x})$ to the observed part of the image.

⇒ Evaluating priors based on this criterion should help select better priors.

This is typically what is done with splitting loss or equivariant learning.

Unrolling dictionary-based denoiser

References

- ▶ Kowalski, M., Malézieux, B., **TM**, and Repetti, A. (2025). [Analysis and synthesis denoisers for forward-backward plug-and-play algorithms](#). In *SIIMS*

Proximal Gradient Descent: Iterate

$$\mathbf{x}^{k+1} = \text{prox}_{\rho\mathcal{R}}(\mathbf{x}^k - \rho \nabla f(\mathbf{x}^k))$$

for the original problem $\mathbf{x}^*(\mathbf{y}; \theta) = \underset{\mathbf{x}}{\operatorname{argmin}} \underbrace{\frac{1}{2} \|\mathbf{y} - \mathbf{Ax}\|_2^2}_{f(\mathbf{x})} + \mathcal{R}(\mathbf{x}; \theta)$

Proximal Gradient Descent: Iterate

$$\mathbf{x}^{k+1} = \text{prox}_{\rho\mathcal{R}}(\mathbf{x}^k - \rho \nabla f(\mathbf{x}^k))$$

for the original problem $\mathbf{x}^*(\mathbf{y}; \theta) = \underset{\mathbf{x}}{\text{argmin}} \underbrace{\frac{1}{2} \|\mathbf{y} - \mathbf{Ax}\|_2^2}_{f(\mathbf{x})} + \mathcal{R}(\mathbf{x}; \theta)$

Dictionary-based denoisers: take \mathcal{D} as

$$\mathcal{D}(\underset{\mathbf{z}}{\text{argmin}} \|\mathbf{x} - \mathcal{D}\mathbf{z}\|_2^2 + \lambda \|\mathbf{z}\|_1) \quad \text{or} \quad \underset{\mathbf{u}}{\text{argmin}} \|\mathbf{x} - \mathbf{u}\|_2^2 + \lambda \|\mathcal{F}\mathbf{u}\|_1$$

These denoisers are proximal operators, so the PnP algorithm converges

Proximal Gradient Descent: Iterate

$$\mathbf{x}^{k+1} = \text{prox}_{\rho \mathcal{R}}(\mathbf{x}^k - \rho \nabla f(\mathbf{x}^k))$$

for the original problem $\mathbf{x}^*(\mathbf{y}; \theta) = \underset{\mathbf{x}}{\text{argmin}} \underbrace{\frac{1}{2} \|\mathbf{y} - \mathbf{Ax}\|_2^2}_{f(\mathbf{x})} + \mathcal{R}(\mathbf{x}; \theta)$

Dictionary-based denoisers: take \mathcal{D} as

$$\mathcal{D}(\underset{\mathbf{z}}{\text{argmin}} \|\mathbf{x} - \mathcal{D}\mathbf{z}\|_2^2 + \lambda \|\mathbf{z}\|_1) \quad \text{or} \quad \underset{\mathbf{u}}{\text{argmin}} \|\mathbf{x} - \mathbf{u}\|_2^2 + \lambda \|\mathcal{F}\mathbf{u}\|_1$$

These denoisers are proximal operators, so the PnP algorithm converges

However, computing the proximal operator can be expensive

requires a sub-routine

⇒ Replace the proximal operator by an unrolled model
but keeping convergence guarantees?

Replace computing the proximal operator with L steps of a solver for the proximal operator with warm-starting

$$\begin{cases} \mathbf{x}^{k+1/2} = \mathbf{x}^k - \rho \nabla f(\mathbf{x}^k) \\ \mathbf{x}^{k+1}, \mathbf{u}^{k+1} = T^L(\mathbf{x}^{k+1/2}, \mathbf{u}^k) \end{cases}$$

- ▶ We show that for $L \rightarrow \infty$, the PnP algorithm converges
- ▶ For $L = 1$, the PnP algorithm also converges to the same solution
- ▶ For intermediate L , it is conjectured that it converges.
Only able to show convergence for a smoothed version of the problem

⇒ Warm-starting is key for the convergence here!

Warm-starting for bilevel optimization

Bilevel optimization problem:

$$\min_{\theta} F(\theta, x^*(\theta)) \quad \text{s.t.} \quad x^*(\theta) = \operatorname{argmin}_x f(x, \theta)$$

Warm-starting for bilevel optimization

Bilevel optimization problem:

$$\min_{\theta} F(\theta, x^*(\theta)) \quad \text{s.t.} \quad x^*(\theta) = \operatorname{argmin}_x f(x, \theta)$$

Unrolling replace $x^*(\theta)$ by $x^N(\theta)$ the output of N iterations of an algorithm

⇒ “Bilevel” convergence is hindered by the fixed precision of the unrolled network.

Warm-starting for bilevel optimization

Bilevel optimization problem:

$$\min_{\theta} F(\theta, x^*(\theta)) \quad \text{s.t.} \quad x^*(\theta) = \operatorname{argmin}_x f(x, \theta)$$

Unrolling replace $x^*(\theta)$ by $x^N(\theta)$ the output of N iterations of an algorithm

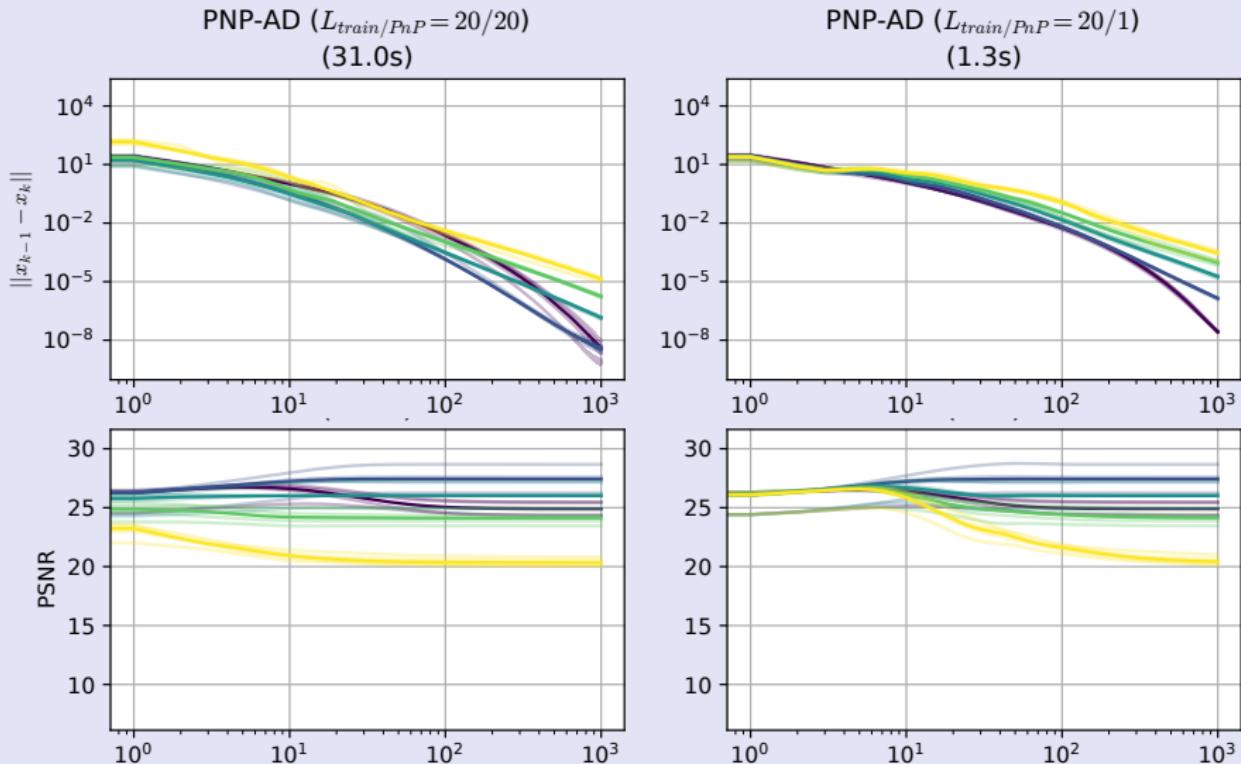
⇒ “Bilevel” convergence is hindered by the fixed precision of the unrolled network.

Idea: warm-start the unrolled algorithm with the previous value $x^N(\theta^{k-1})$.

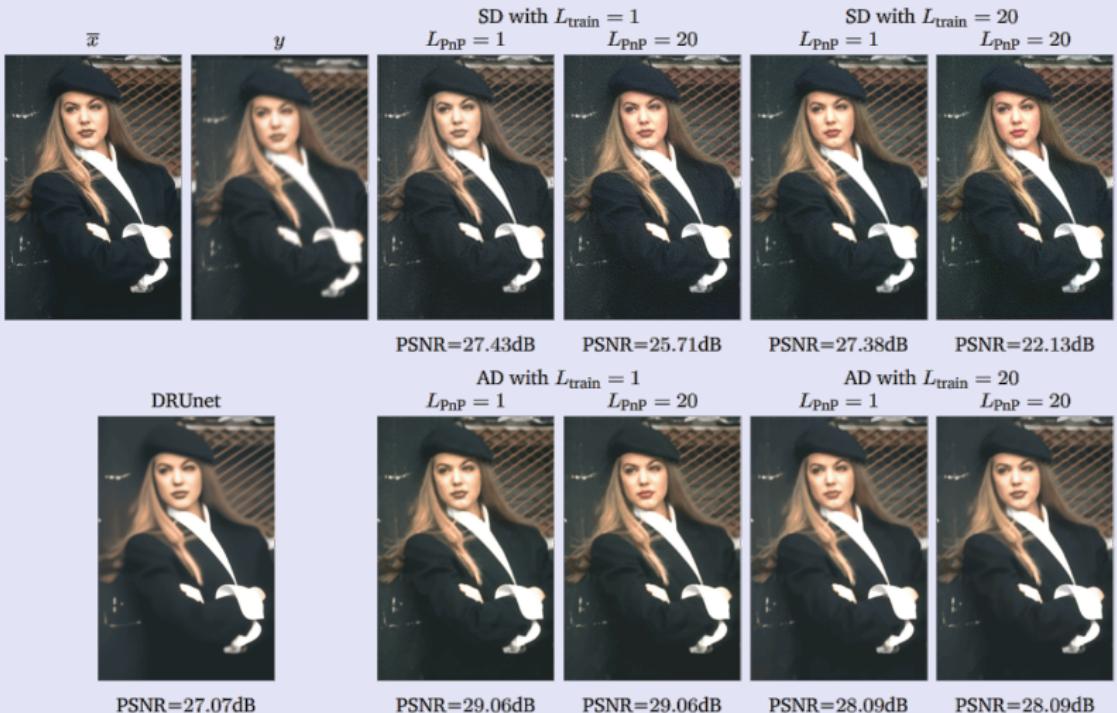
Key point for efficient stochastic bilevel solvers:

- ▶ Two loops: AMIGO [Arbel and Mairal 2021]
- ▶ One-loop: FSL/SOBA [Li et al., 2022, Dagreou et al. 2022]

Stability of the unrolled algorithm



Dictionary-based denoisers in PnP



Take-home message

- ▶ Dictionary-based priors are useful tools to study prior learning and interpretable.
- ▶ Warm-starting can be a key to get convergence with unrolled algorithms
- ▶ We obtain reasonable performance with few unrolled iterations for learning a dictionary-based denoiser

FiRe: Fixed-Point Restoration

References

- ▶ Terris, M., Kamilov, U., and TM (2025). FiRe: Fixed-points of Restoration Priors for Solving Inverse Problems. In *CVPR*

Using restoration networks as priors

Many efficient restoration networks have been proposed for various tasks in the last years.

- ▶ JPEG restoration: SCUNet [Zhang et al., 2023]
- ▶ Deblurring: Restormer [Zamir et al., 2022]
- ▶ Inpainting: LAMA [Suvorov et al., 2022]
- ▶ ...

These models are trained to solve $p(\mathbf{x}|\mathbf{y})$ for a specific degradation $\mathbf{y} = D(\mathbf{x})$.

A good network for deblurring should be able to recover high frequencies from a blurry image.

⇒ Can we use them to solve other inverse problems?

TL;DR: *yes, this is what is done with DRP, SHARP, ...*

[Hu et al., 2024a, 2024b]

Fixed-point of restoration networks

Observation: Denoisers are not *stable* when iterated

$X_k = \underbrace{D \circ D \cdots \circ D(X_0)}_{k \text{ times}}$ is not converging to a realistic image

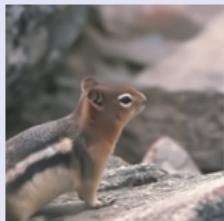
Here with **DRUNet** with $\sigma = 0.05$

[Zhang et al., 2021]

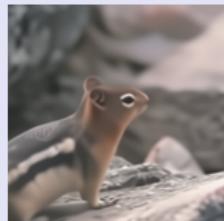
x_0



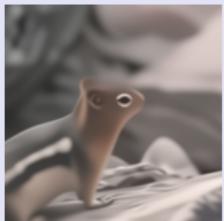
x_1



x_3



x_{10}



Fixed-point of restoration networks

Observation: Similar observation holds for restoration networks

Here with **SCUNet**

[Zhang et al., 2023]



Fixed-point of restoration networks

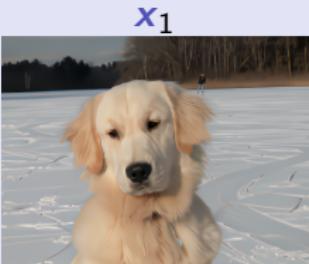
Observation: Similar observation holds for restoration networks

Here with **SCUNet**

[Zhang et al., 2023]



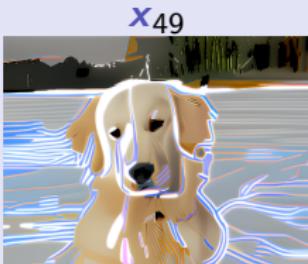
x_0



x_1



x_{20}



x_{49}

and Restformer

[Zamir et al., 2022]



Restoration models

Definition

A restoration model R^D adapted to a degradation D is model solving

$$R^D(D(\mathbf{x})) \approx \mathbf{x} \quad \text{for images} \quad \mathbf{x} \sim \mathcal{X} .$$

We have in mind degradation models $D(x)$ of the form

$$D(\mathbf{x}) = \mathbf{A}\mathbf{x} + \mathbf{w},$$

with \mathbf{A} a linear operator, and some noise $\mathbf{w} \sim \mathcal{W}$.

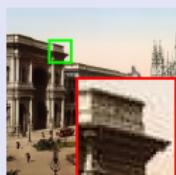
Restormer

Deblurring



SCUNet

JPEG Restoration



LAMA

inpainting



Training of restoration models

Direct restoration models are trained in a supervised manner, starting from clean images $x \sim \mathcal{X}$, and a degradation model D .

$$\mathcal{L}(\theta) = \mathbb{E}_{\mathbf{x}, w \sim \mathcal{W}} \left[\| \mathbf{R}_\theta^D(H\mathbf{x} + w) - \mathbf{x} \| \right].$$

Observation:

- ▶ Let $T = \mathbf{R}^D \circ D$, then for \mathbf{R}^D sufficiently well trained, we expect $T^2 \approx T$.
- ▶ T shows an idempotence property so if fixed-points exist, they should be for $\mathbf{R} \circ D$ and not \mathbf{R} .
- ▶ With this training, we expect that the original data \mathbf{x} are part of the fixed-point set of T .

Fixed-point of restoration networks

Observation: degradation + denoisers are *stable* when iterated

$X_k = \underbrace{D \circ D \cdots \circ D(X_0)}_{k \text{ times}}$ is not converging to a realistic image

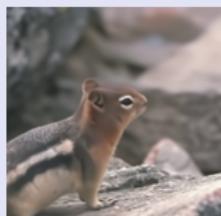
Here with **DRUNet** with $\sigma = 0.05$

[Zhang et al., 2021]

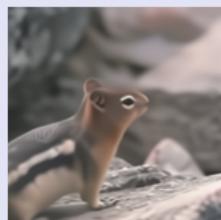
x_0



x_1



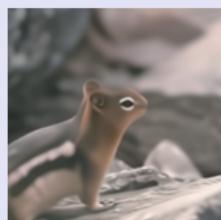
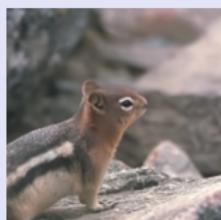
x_3



x_{10}



and for T :



Fixed-point of restoration networks

Observation: Similar observation holds for restoration networks



and for T :



Restoration models as projections

Assumption

Let $C = \{x \in \mathbb{R}^n; T(x) = x\}$, then $T = R \circ D$ can be expressed as a projection $T = \text{proj}_C$ onto a closed, prox-regular set C .

Proposition

Under our projection assumption, around any point \mathbf{x} where C is prox-regular, we have:

$$T(\mathbf{x}) = \mathbf{x} - \frac{1}{2} \nabla d_C^2(\mathbf{x}),$$

where: $d_C(\mathbf{x}) = \inf_{u \in C} \|\mathbf{x} - u\|$ is the distance to the set C .

\Rightarrow Define a prior $p(\mathbf{x}) \propto \exp\left(-\frac{1}{2}d_C^2(\mathbf{x})\right)$, promoting fixed points of T .

A anomaly detection view on restoration networks

T is a reconstruction network, that can be viewed as an auto-encoder with a fixed encoder D .

Anomaly detection literature:

[Liu and Paparrizos 2024]

- ▶ Train a reconstruction network on normal data,
- ▶ use large reconstruction error as an anomaly indicator.

$$\text{Anomaly score}(\mathbf{x}) = \|\mathbf{x} - T(\mathbf{x})\|^2 = d_C^2(\mathbf{x})$$

⇒ Characterize $p(\mathbf{x}|D(\mathbf{x}))$ for $\mathbf{x} \sim \mathcal{X}$ a natural image.

This is usually easier to learn than $p(\mathbf{x})$ directly.

FiRe-HQS Algorithm

We can define the FiRe-HQS algorithm as:

$$\begin{aligned} u_k &= x_k - \frac{\gamma}{2} \nabla d_C^2(x) \\ x_{k+1} &= \text{prox}_{\lambda f}(u_k). \end{aligned}$$

Note: This is also similar to SNORE for denoisers.

[Renaud et al. 2024]

Proposition

Under our projection assumption, the FiRe-HQS algorithm converges to a point x^* satisfying

$$x^* = \operatorname{argmin}_x \lambda f(x) + \frac{\gamma}{2} d_C^2(x),$$

Extended FiRe-HQS Algorithm

We can extend the FiRe-HQS algorithm to multiple restoration models $\{T_i = R_i \circ D_i\}_{i=1}^M$ as:

Input: Initial estimate x_0 , weights γ_n , regularization parameter λ .

for $k = 1, \dots, K$ **do**

for $n = 1, \dots, N$ **do**

 Select restoration model (R^n, D^n) ;

 Compute residual: $r_k^n = x_k - R^n(D^n(x_k))$;

$u_k = x_k - \sum_{n=1}^N \gamma_n r_k^n \quad \leftarrow \quad x_k - \gamma \nabla d_C^2(x_k)$;

$x_{k+1} = \text{prox}_{\lambda f}(u_k)$;

Output: Final estimate x_{k+1} .

⇒ We combine the strengths of multiple restoration models to remove artifacts from each other.

Key point: They are all trained to have the natural images as fixed-points.

Experimental Setting

We consider inverse problems $\mathbf{y} = \mathbf{Ax} + \epsilon$ and solve:

$$\mathbf{x}^* = \underset{\mathbf{x}}{\operatorname{argmin}} \underbrace{\lambda \frac{1}{2} \|\mathbf{Ax} - \mathbf{y}\|}_{f(\mathbf{x})} + \frac{\gamma}{2} \mathbb{E}_{\xi \sim \Xi} [d_{C_\xi}^2(\mathbf{x})]$$

for various implicit priors (recall that $\frac{1}{2} \nabla d_C(\mathbf{x}) = \mathbf{x} - \mathbf{R}(\mathbf{Hx} + \mathbf{w})$).

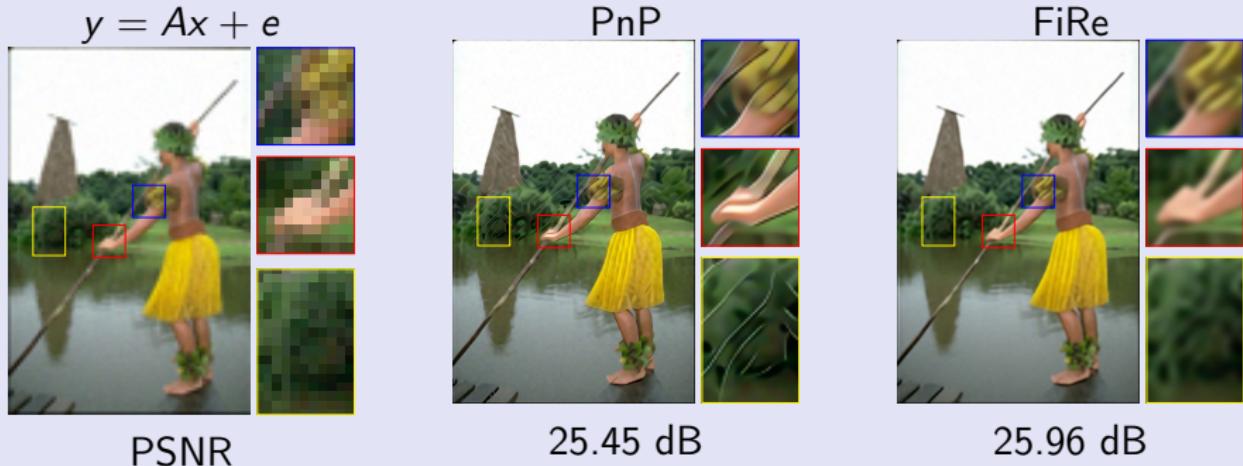
Restoration Models R:

- **DRUNet:** Gaussian denoising with $H = \text{Id}$ and $\mathbf{w} \sim \mathcal{N}(0, \sigma^2)$.
- **Restormer:** Gaussian or motion deblurring with H as Gaussian or motion blurs, $\mathbf{w} \sim \mathcal{N}(0, \sigma^2)$.
- **SCUNet:** Non-linear restoration with $H = \text{JPEG}_q$ ($q \in [20, 100]$), $\mathbf{w} \sim \mathcal{N}(0, \sigma^2)$.
- **SwinIR:** Super-resolution ($\times 2, \times 3$) with H as downsampling, $\mathbf{w} = 0$.
- **LAMA:** (a) Pretrained: H as a large mask, $\mathbf{w} = 0$. (b) Fine-tuned: H for random inpainting, $\mathbf{w} = 0$.

Results with a single prior

We first consider our FiRe approach with a **single** restoration model. We consider the $4 \times$ SR problem.

Case 1: prior is the SCUNet with noisy JPEG degradations.



Results with a single prior

(a) \mathbf{y}



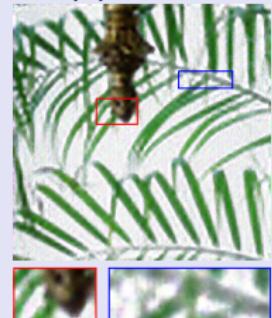
(b) $\mathbf{A}^\dagger \mathbf{y}$



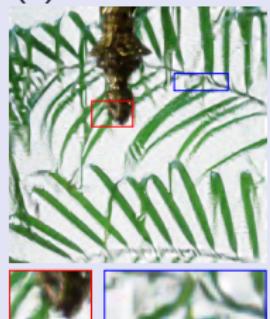
(c) DPIR



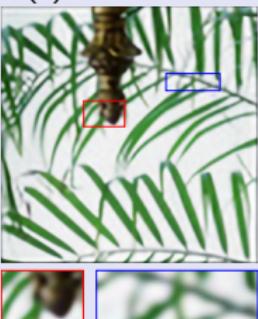
(d) LAMA



(e) Rest. Motion



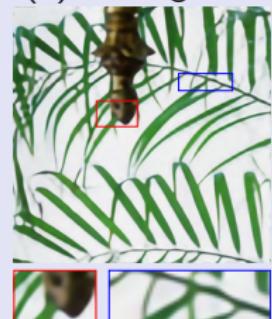
(f) SwinIR 2x



(g) SCUNet JPEG



(h) Rest. gauss.

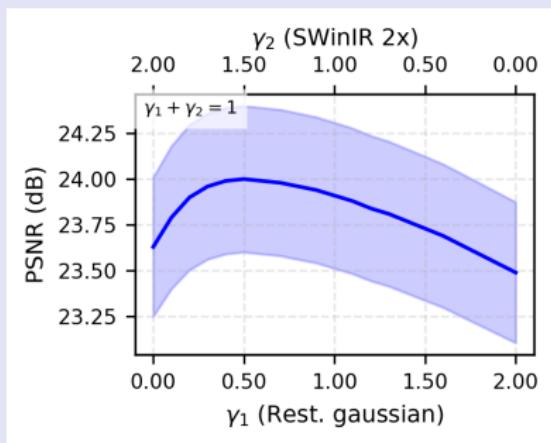
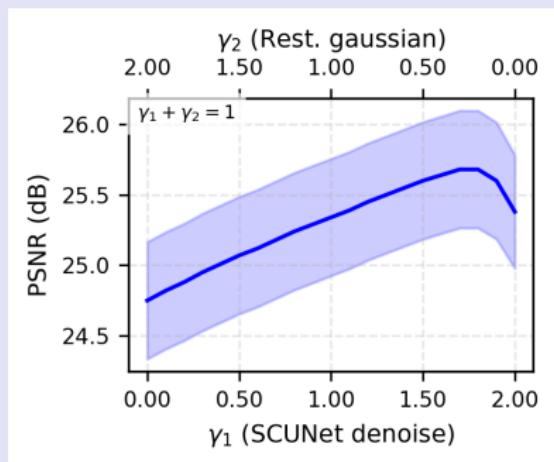


Combining priors

The FiRe framework allows us to combine priors as

$$\gamma_1 R_1(H_1 \mathbf{x} + w_1) + \gamma_2 R_2(H_2 \mathbf{x} + w_2)$$

Example:



γ vs PSNR within the reconstruction quality for two different problems. Left: Gaussian deblurring, right: SR \times 4. The γ_1 and γ_2 parameter control the strength of the associated prior.

Combining priors

Visual results:

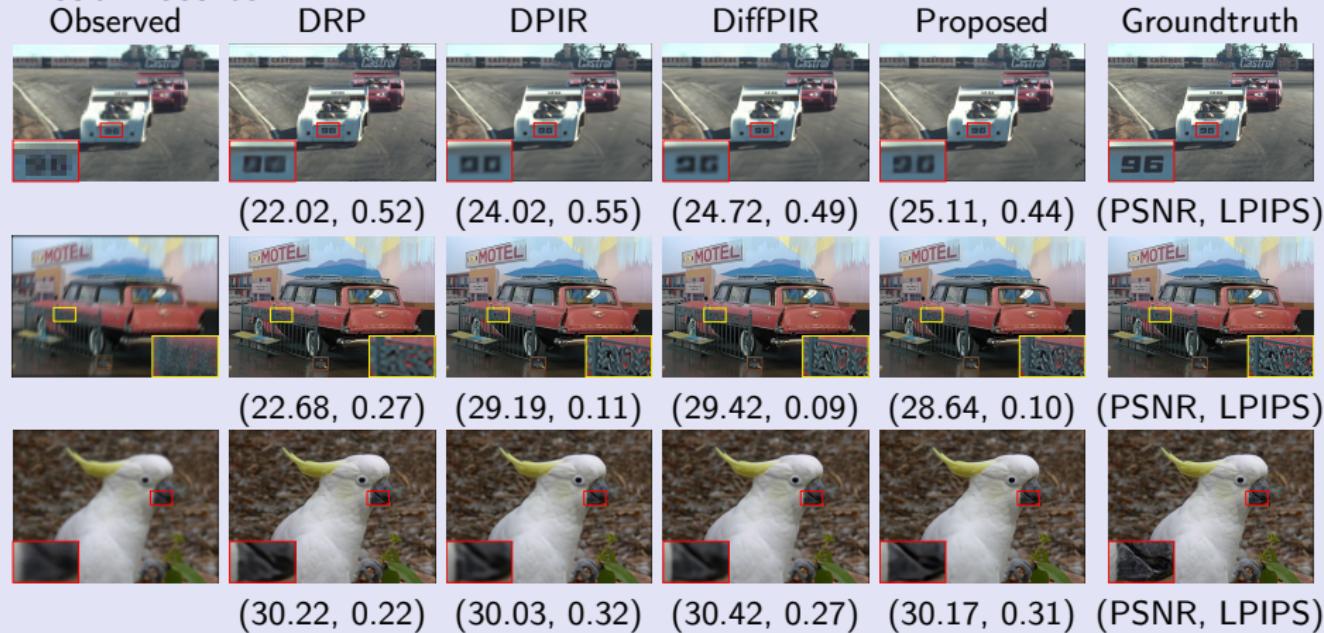


Image restoration with various algorithms. Top: SR $\times 4$ problem with $\sigma = 0.01$ on BSD20. Middle: Motion blur on Imnet100. Bottom: Gaussian deblurring with blur kernel of size 3 and $\sigma = 0.01$ on Imnet100.

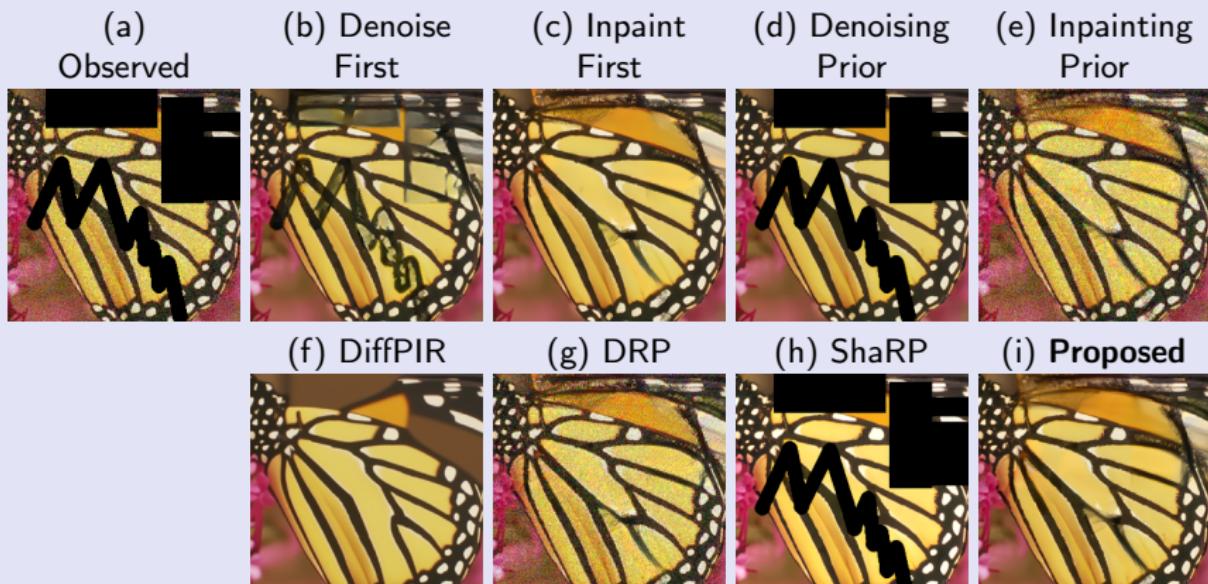
Conditioning the prior on the measurements

Recall the iteration (simplified)

$$u_k = R(H\mathbf{x}_k + w_k)$$

$$\mathbf{x}_{k+1} = \text{prox}_{\lambda f}(u_k).$$

Given that $f = \frac{1}{2}\|\mathbf{Ax} - \mathbf{y}\|^2$, one can set $H = \mathbf{A}$. Application to inpainting:



Take-home message

- ▶ Restoration networks can serve as interpretable priors.
- ▶ A key point is to consider the fixed-point set of the restoration network.
- ▶ Multiple priors can be combined efficiently using a stochastic approach.
- ▶ When adapted to the task, Fire can complete the missing information from $\ker(\textcolor{green}{A})$.

Reproducible method comparison with Benchopt



References

- ▶ **TM**, Massias, M., Gramfort, A., Ablin, P., Bannier, P.-A., Charlier, B., Dagréou, M., la Tour, T. D., Durif, G., Dantas, C. F., Klopfenstein, Q., Larsson, J., Lai, E., Lefort, T., Malézieux, B., Moufad, B., Nguyen, B. T., Rakotomamonjy, A., Ramzi, Z., Salmon, J., and Vaiter, S. (2022). [Benchopt: Reproducible, efficient and collaborative optimization benchmarks](#). In *NeurIPS*

Benchmarks and reproducibility

Benchmarks fueled AI progress



Benchmarks fueled AI progress



But sometime, it is not so clear which methods should be included:

- ▶ Different evaluation protocols
- ▶ Different implementations
- ▶ Hard to tune all methods
- ▶ ...

⇒ Many novel methods but unclear improvements

Making runnable benchmarks with `benchopt`



benchopt provides a framework to organize and run benchmarks

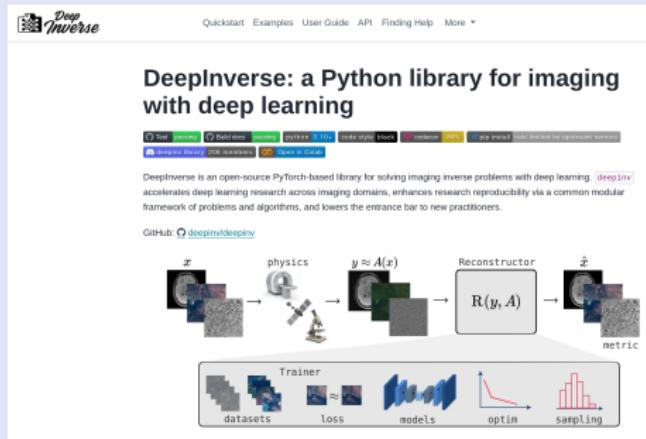
Examples of existing benchmarks:

- ▶ **NanoGPT optimization (GPT2)**
- ▶ **Image Classification (ResNet)**
- ▶ **Logistic regression**
- ▶ **Lasso**
- ▶ **Unsup. Domain Adaptation**
- ▶ **Bilevel Optimization**
- ▶ **Brain Computer Interface**
- ▶ **...**

Example: Benchmarking Inverse Problems solvers with Deepinv

Benchmarking various methods in a single repo:

- ▶ Imaging, MRI, CT, ...
- ▶ Direct, PnP, Variational, ...
- ▶ Centralized evaluation
- ▶ Clear rules on tuning the methods



⇒ **Goal:** Make it easy to add new methods and datasets

https://github.com/benchopt/benchmark_inverse_problems/

Thanks for your attention!

Slides are on my web page:

 tommoral.github.io

 @tommoral