FiRe: Fixed-points of Restoration Priors for Solving Inverse Problems

Supplementary Material

A. Convergence of stochastic algorithm

We now formalize the convergence result stated in Section 4.2 for algorithm (15). Our analysis relies on two key assumptions. First, we require unbiased gradient estimators with variance growth condition:

Assumption A.1. For all k, g_k in (15) is an unbiased estimator of $\frac{1}{2}\nabla d_C^2$. More precisely, assuming that the sequence $(x_k)_{k\in\mathbb{N}}$ is adapted to the filtration $\{\mathcal{F}_k\}_{k\geq 0}$, we assume that $\mathbb{E}_{A_k,e_k}[g_k|\mathcal{F}_k] = \frac{1}{2}\nabla d_C^2(x_k)$ where d_C is defined as in Proposition 3.2. Furthermore, we assume that there exists constants $A, B \geq 0$ such that

$$\mathbb{E}[\|g_k - \frac{1}{2}\nabla d_C^2(x_k)\|^2|\mathcal{F}_k] \leq A(d_C^2(x_k)) + B \quad (16)$$

holds almost surely for all $k \in \mathbb{N}$.

Second, we require a standard assumption on the step sizes:

Assumption A.2. The stepsize $(\gamma_k)_{k\in\mathbb{N}}$ satisfies $\sum_{k=0}^{\infty}\gamma_k = \infty$ and $\sum_{k=0}^{\infty}\gamma_k^2 < \infty$.

Under those assumptions, the proposed algorithm rewrites as a proximal stochastic gradient algorithm, and we can derive the following result.

Proposition A.3. Assume that Assumptions A.1 and A.2 hold, and define the residual function

$$F(x) = x - \text{prox}_{\lambda f}(x - \frac{1}{2}d_C^2(x)). \quad (17)$$

Then we have that $\mathbb{E}[F(x_k)] \underset{k\to\infty}{\longrightarrow} 0$.

Proof. First, we have that in our setting, f is convex. Moreover, we have that d_C^2 is bounded from below, and from Proposition 3.2, we have that ∇d_C^2 is Lipschitz. The result then follows from [19, Theorem Corollary 3.6]. □

B. Prior loss

To investigate the underlying prior associated with the different models, we investigate the quantity $d(y) = \|y - R(Hy + w)\|$ where y are images with various levels of degradation. More precisely, we set $y = \sigma_{\text{blur}} * x + \sigma_{\text{noise}} n$ where $n \sim \mathcal{N}(0, \text{Id})$; thus, as $\sigma_{\text{noise}} \to 0$ and $\sigma_{\text{blur}} \to 0$, y tends to a natural image. We plot values of d obtained for different models R in Figure 7. We observe that the lowest values of $d(y)$ are obtained for noiseless, smooth images, suggesting that the proposed priors tends to show a smoothing property and promotes image regularity.

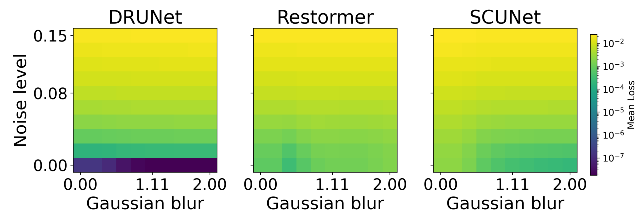

Figure 7. Average distance $\|y - R(H(y))\|_2$ for different models and different degradations for y. More precisely, y is a blurred and noisy version of x defined as $y = \sigma_{\text{blur}} * x + \sigma_{\text{noise}} n$.

C. Finetuning of restoration priors

In this section, we summarize the implementation details for the finetuned models used in our experiments. We consider 3 finetuned models: two versions of the Restormer model for image deblurring (Gaussian and motion), and random mask inpainting. All finetunings are performed on the training dataset from [39].

C.1. Deblurring models

The Restormer model from [37] is trained on real blur images with blur kernels difficult to simulate. In turn, computing the degradation set $\mathcal{D}$ associated to the model is not straightforward. Instead, we propose to finetune the model in two easily simulated setups.

Restormer Gaussian: Finetuned on Gaussian and diffraction blur removal with kernel size sampled uniformly at random $\sigma_{\text{blur}} \in [0.001, 4]$ and additive Gaussian noise with standard deviation sampled uniformly at random $\sigma \in [0.001, 0.1]$. We train on randomly cropped image patches of sizes 256^2 with batch size 8. Optimization is performed on L1 loss and uses Adam optimizer with default PyTorch parameters and learning rate 1e-4 for 90k steps.

Restormer Motion: Finetuned on motion blur removal from [30] with trajectory length scale 0.6, Gaussian Process standard deviation 1.0, and additive Gaussian noise sampled uniformly at random $\sigma \in [0.001, 0.1]$. Other training parameters are the same as above.

C.2. Inpainting models

The LAMA model [29] is trained on large blur kernels and performs well on large inpainting masks, but we observed suboptimal performance on binary random masks.

LAMA random inpainting: Finetuned for random inpainting with mask probability sampled uniformly at random $p \in [0.1, 0.9]$ (no noise is added during training). Only the last 4 convolutional layers are updated during training. We train on randomly cropped image patches of sizes 128^2 with batch size 64. Optimization is performed on L1 loss uses Adam optimizer with learning rate 1e-4 for 300k steps.

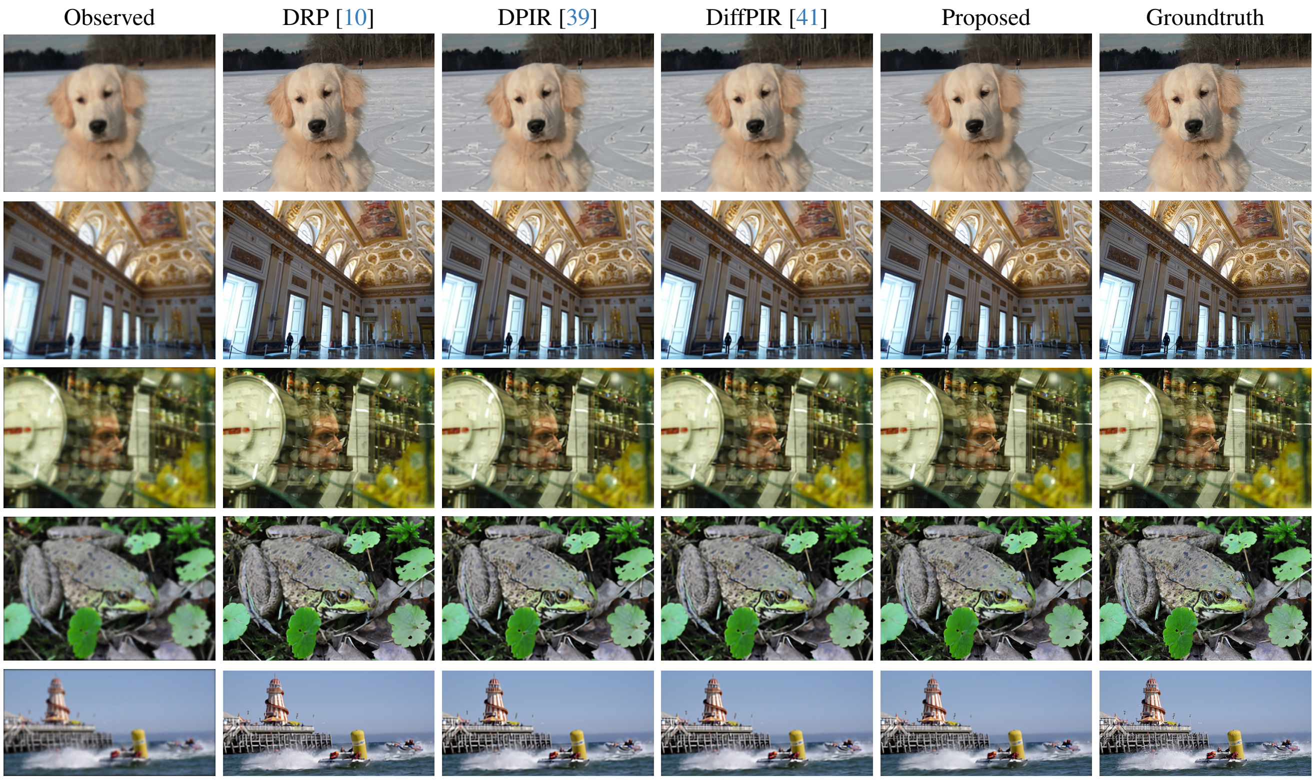

Figure 8. Gaussian deblurring on Imnet100, $\sigma = 0.01$.

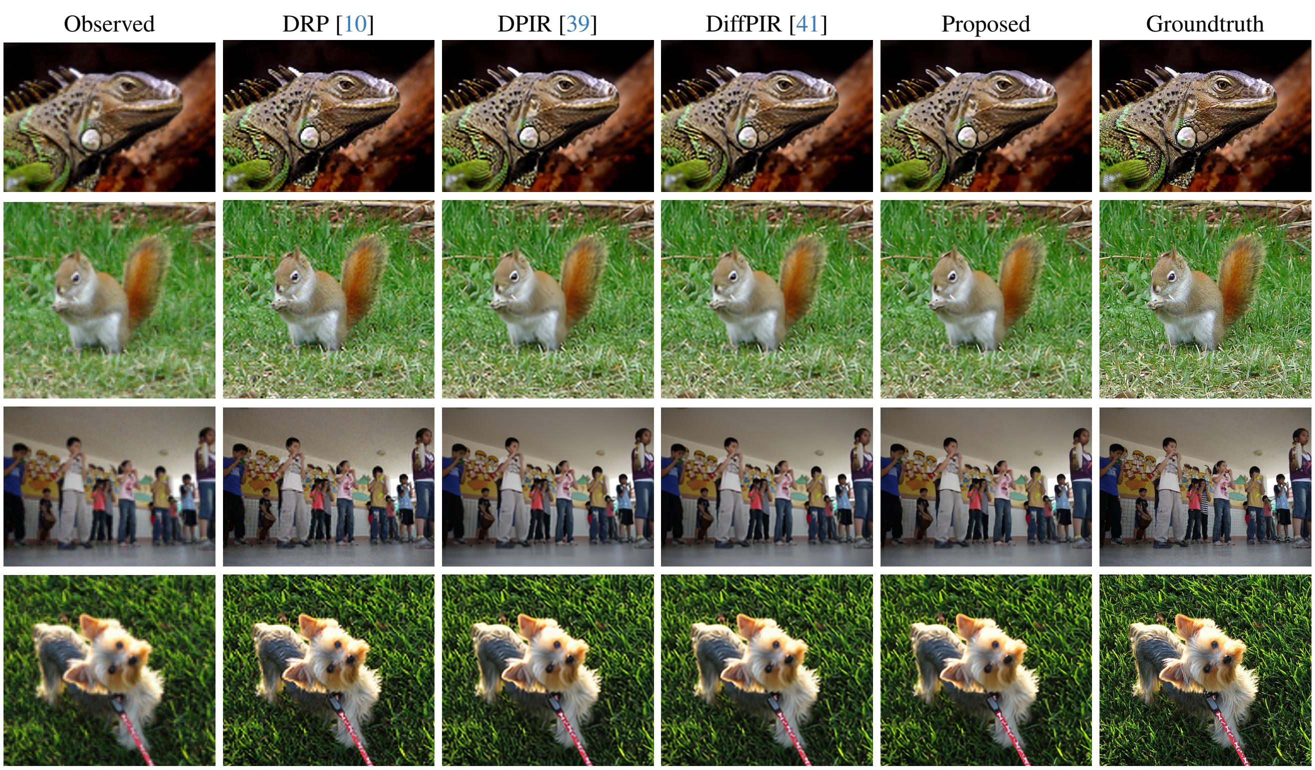

Figure 9. SRx4 on Imnet100, $\sigma = 0.01$.

D. Additional visual results

We provide further reconstruction results on the Gaussian deblurring problem in Figure 8 and on the SRx4 problem in Figure 9.

E. Influence of the degradation

The choice of degradation operator significantly impacts reconstruction quality. In Fig. 10, we give reconstruction metrics for varying degradation strengths using either a Gaussian Restormer prior and a DRUNet denoising prior. This shows that a minimum degradation is required to stabilize the reconstruction, while increasing beyond a certain threshold leads to excessive smoothing.

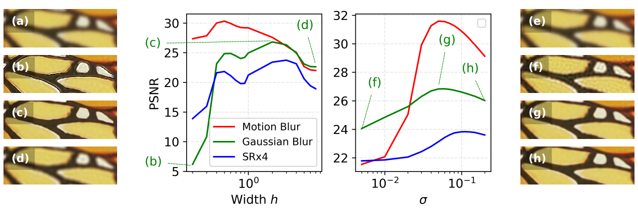

Figure 10. Influence of the degradation operator on the FiRe restoration algorithm on the Set3C algorithm for different problems. Left plot: FiRe with Restormer Gaussian prior. We run the algorithm with different kernel widths. Reconstructions (details) corresponding to points (b), (c), (d) on the plot are shown on the left. Right plot: same, but with a DRUNet denoiser restoration prior; (f), (g) and (h) panels show the associated reconstructions. (a) and (e) show the degraded measurements y.