Provable Probabilistic Imaging using Score-based Generative Priors

Yu Sun[1], Zihui Wu[2], Yifan Chen[3], Berthy Feng[2], and Katherine L. Bouman[2]

SCAN ME

1. Goal & Problem

Goal: Develop **provable posterior sampling** method for **ill-posed** imaging problems by combining ***physical models*** and ***generative models***.

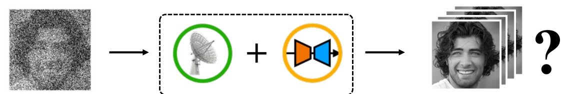

Plug-and-play priors (PnP) is a deterministic method that Integrates forward model and denoising priors to do **maximum a posterior (MAP)** estimation

$$\arg\max_{\boldsymbol{x}} \log \pi(\boldsymbol{x}) \qquad \pi \propto \ell(\boldsymbol{y}|\boldsymbol{x})p(\boldsymbol{x})$$

Likelihood Prior

Problem: PnP can only produce a single point estimate, failing to capture the entire distribution.

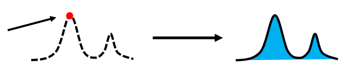

Strategy: Extend PnP for posterior sampling by connecting denoiser with score function.

2. Plug-and-Play Monte Carlo (PMC)

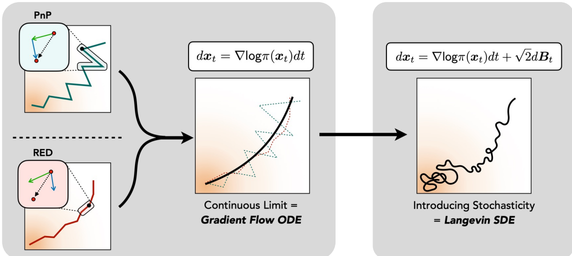

Contribution #1: We show the continuous limits of gradient-based PnP and RED converge to the ***gradient flow ODE*** associated with the posterior.

- We use ***Tweedie's formula*** to connect the minimum mean squared error (***MMSE***) denoiser with the ***image prior score***.

Contribution #2.(a): We develop PMC as sampling extensions of PnP/RED by discretizing the ***Langevin SDE*** using PnP/RED's discretization strategies.

Contribution #2.(b): We propose ***weighted annealing*** to accelerate the convergence of .

Algorithm 1 Plug-and-play Monte Carlo (PMC)

input: $\boldsymbol{x}_0 \in \mathbb{R}^n$, $\gamma > 0$, $\alpha_0 > 0$, and $\sigma_0 > 0$.
for $k = 0, 1, \ldots, N-1$ **do**
  $\boldsymbol{z}_k \leftarrow \mathcal{N}(0, I)$
  $\sigma_k, \alpha_k \leftarrow \mathsf{WeightedAnnealing}(\sigma_0, \alpha_0, k)$.
  switch *discretization*
    case PnP:
      $\mathcal{P}_k(\boldsymbol{x}_k) \leftarrow \nabla g(\boldsymbol{x}_k) - \alpha_k \mathcal{S}_\theta(\boldsymbol{x}_k - \gamma\nabla g(\boldsymbol{x}_k), \sigma_k)$
      $\boldsymbol{x}_{k+1} \leftarrow \boldsymbol{x}_k - \gamma\mathcal{P}_k(\boldsymbol{x}_k) + \sqrt{2\gamma}\boldsymbol{z}_k$
    case RED:
      $\mathcal{G}_k(\boldsymbol{x}_k) \leftarrow \nabla g(\boldsymbol{x}_k) - \alpha_k \mathcal{S}_\theta(\boldsymbol{x}_k, \sigma_k)$
      $\boldsymbol{x}_{k+1} \leftarrow \boldsymbol{x}_k - \gamma\mathcal{G}_k(\boldsymbol{x}_k) + \sqrt{2\gamma}\boldsymbol{z}_k$
end for

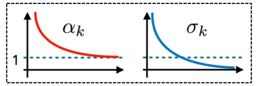

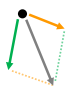

Fig. Illustration of update rule

$g(\boldsymbol{x}) = -\log \ell(\boldsymbol{y}|\boldsymbol{x})$

Enforcing **consistency with measurements**.

$\mathcal{S}_\theta(\boldsymbol{x}, \sigma)$

Enforcing **generative prior** learned from data.

3. Theoretical Guarantees

Theorem [Sun, Wu, Chen]. The law $(\nu_t)_{t\geq 0}$ for the continuous interpolation of $\{x_k\}_{k=0}^N$ in PMC satisfies

$$\frac{1}{N\gamma}\int_0^{N\gamma} \mathrm{FI}(\nu_t \,\|\, \pi)\, dt \leq \frac{4\mathrm{KL}(\nu_0 \,\|\, \pi)}{N\gamma} + \mathrm{error}(\gamma, \sigma, \epsilon)$$

N: number of iterations $\mathrm{KL}(\cdot\,|\,\pi)$: KL divergence $\mathrm{FI}(\cdot\,|\,\pi)$: Fisher information

γ: step-size σ: blur strength ϵ: score training error

Contribution #3: We show that PMC achieves ***O(1/N) stationarity*** w.r.t the ***true posterior*** in the joint presence of

- ***nonlinear*** forward model,
- ***imperfect*** score function (denoiser), and
- weighted ***annealing***.

4. Compressed Sensing (1:10)

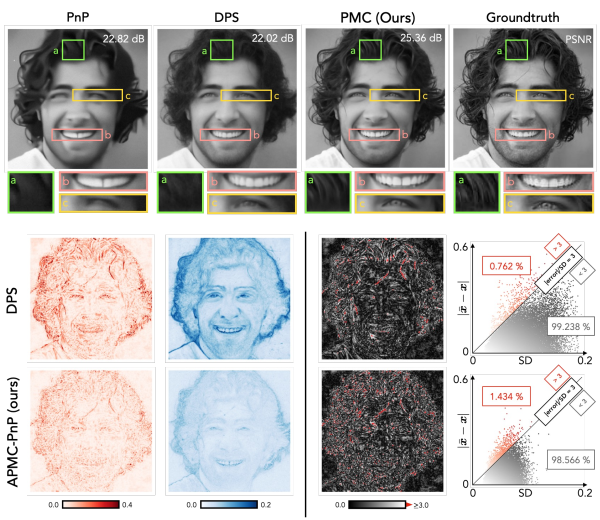

5. Nonlinear Black-hole Imaging

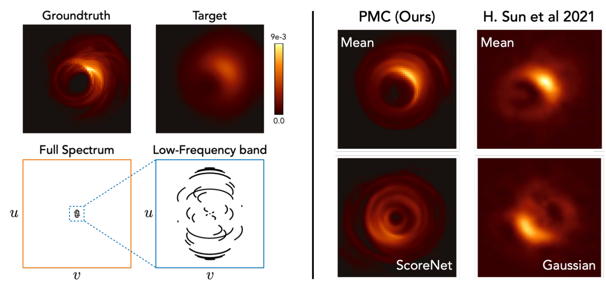

Real M87 Black-hole [NeurIPS 2024]

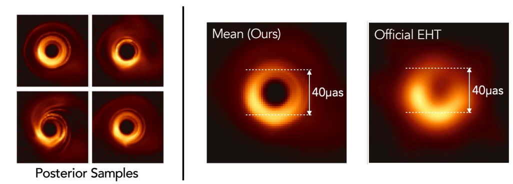

More Information!