InverseBench: Benchmarking Plug-and-Play Diffusion Priors for Inverse Problems in Physical Sciences

Hongkai Zheng[1,*], Wenda Chu[1,*], Bingliang Zhang[1,*], Zihui Wu[1,*], Austin Wang[1], Berthy T. Feng[1], Caifeng Zou[1], Yu Sun[2], Nikola Kovachki[3], Zachary E. Ross[1], Katherine L. Bouman[1], Yisong Yue[1]

[1] CMS, Caltech; [2] EE, JHU; [3] NVIDIA; * Equal contribution

Scan me for more details!

TL;DR

- Most methods using diffusion models for solving inverse problems claim to be general but are mostly tested on image restoration tasks.
- We benchmark **14 methods** on **five representative inverse problems.**
- We provide a highly **modular and scalable** codebase.

Background

- Imaging inverse problems

Recover an image $\widehat{\boldsymbol{x}}$ from $\boldsymbol{y} = \underbrace{A(\boldsymbol{x})}_{\text{forward model}} + \underbrace{\boldsymbol{e}}_{\text{noise}}$ so that $\widehat{\boldsymbol{x}} \approx \boldsymbol{x}$

- Plug-and-Play Diffusion Priors (PnPDP)

$$\underbrace{p(\boldsymbol{x}|\boldsymbol{y})}_{\text{Posterior}} \propto \underbrace{p(\boldsymbol{y}|\boldsymbol{x})}_{\text{Likelihood}} \underbrace{p(\boldsymbol{x})}_{\text{Prior}}$$

Physical model + Diffusion model = PnPDP method

Plug-and-Play combination

- Existing PnPDP methods

Category	Method	SVD	Pseudo inverse	Linear	Gradient
Linear guidance	DDRM (Kawar et al., 2022)	®	®	®	–
	DDNM (Wang et al., 2022)	–	®	®	–
	ΠGDM (Song et al., 2023a)	–	®	–	–
General guidance	DPS (Chung et al., 2023)	–	–	–	®
	LGD (Song et al., 2023b)	–	–	–	®
	DPG (Tang et al., 2023)	–	–	–	–
	SCG (Huang et al., 2024)	–	–	–	–
	EnKG (Zheng et al., 2024)	–	–	–	–
Variable-splitting	DiffPIR (Zhu et al., 2023)	–	–	–	®
	PnP-DM (Wu et al., 2024)	–	–	–	®
	DAPS (Zhang et al., 2024)	–	–	–	®
Variational Bayes	RED-diff (Mardani et al., 2023)	–	–	–	®
Sequential Monte Carlo	FPS (Dou & Song, 2024)	–	–	®	–
	MCGDiff (Cardoso et al., 2024)	®	®	®	–

®: Required

Benchmark Problems

- We consider five representative inverse problems in physical sciences

Problem	SVD	Linear	Complex domain	Closed-form forward	Gradient access	Noise type
Linear inverse scattering	✓	✓	✓	✓	✓	Gaussian
Compressed sensing MRI	✗	✓	✓	✓	✓	Real-world
Black hole imaging	✗	✗	✗	✓	✓	Non-additive
Full waveform inversion	✗	✗	✗	✗	✓ (slow)	Noise-free
Navier-Stokes equation	✗	✗	✗	✗	✗	Gaussian

Main Results

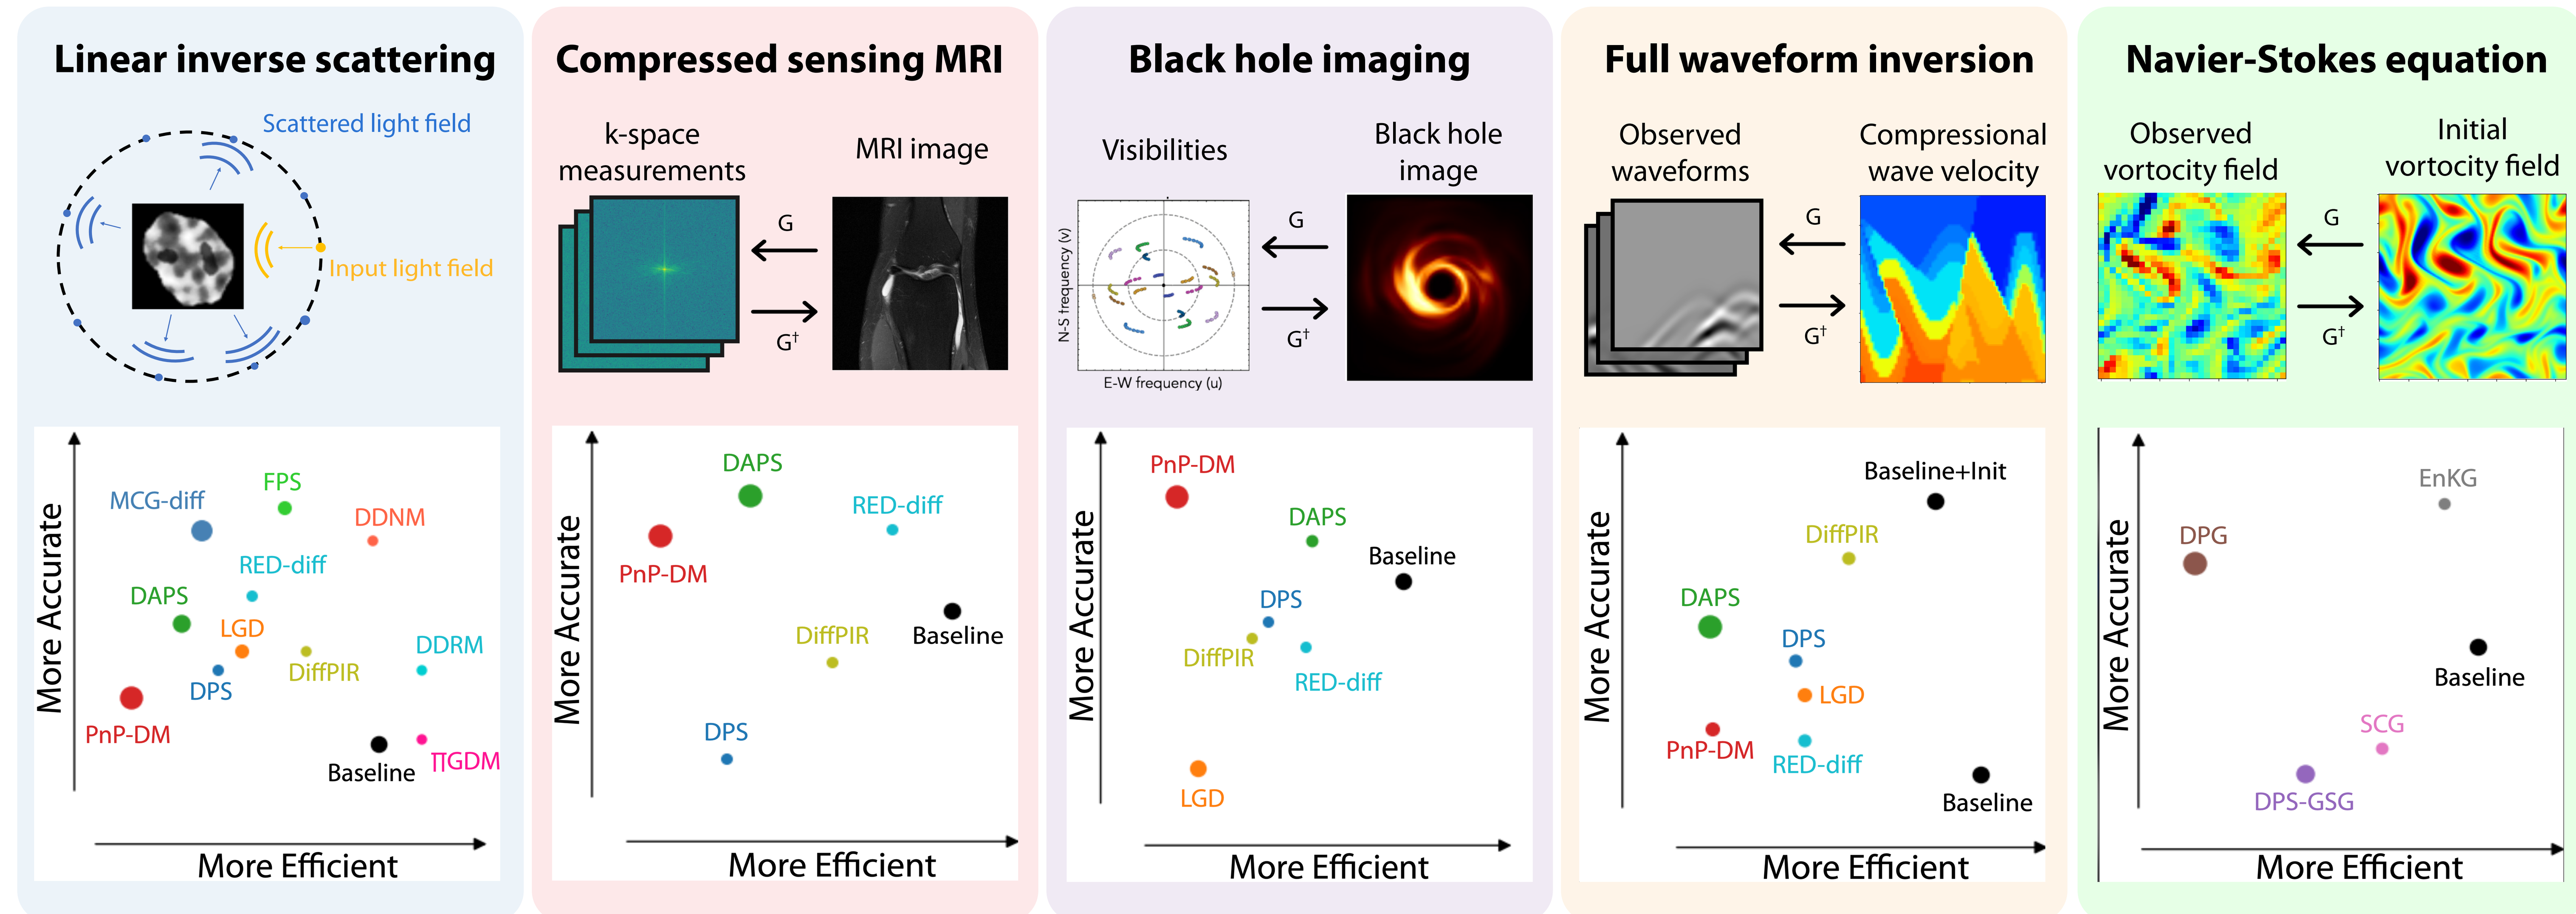

- Accuracy: PSNR; SSIM
- Efficiency: total compute; compute-bound when fully parallelized

Findings

- **PnPDP vs. traditional:** More accurate but less time- and memory-efficient.
- **PnPDP vs. each other:**
  - **With closed-form forward model:** methods that require more gradient calls, e.g. PnP-DM and DAPS, perform better.
  - **Without closed-form forward model:** methods that require more gradient calls can be very slow and unstable.

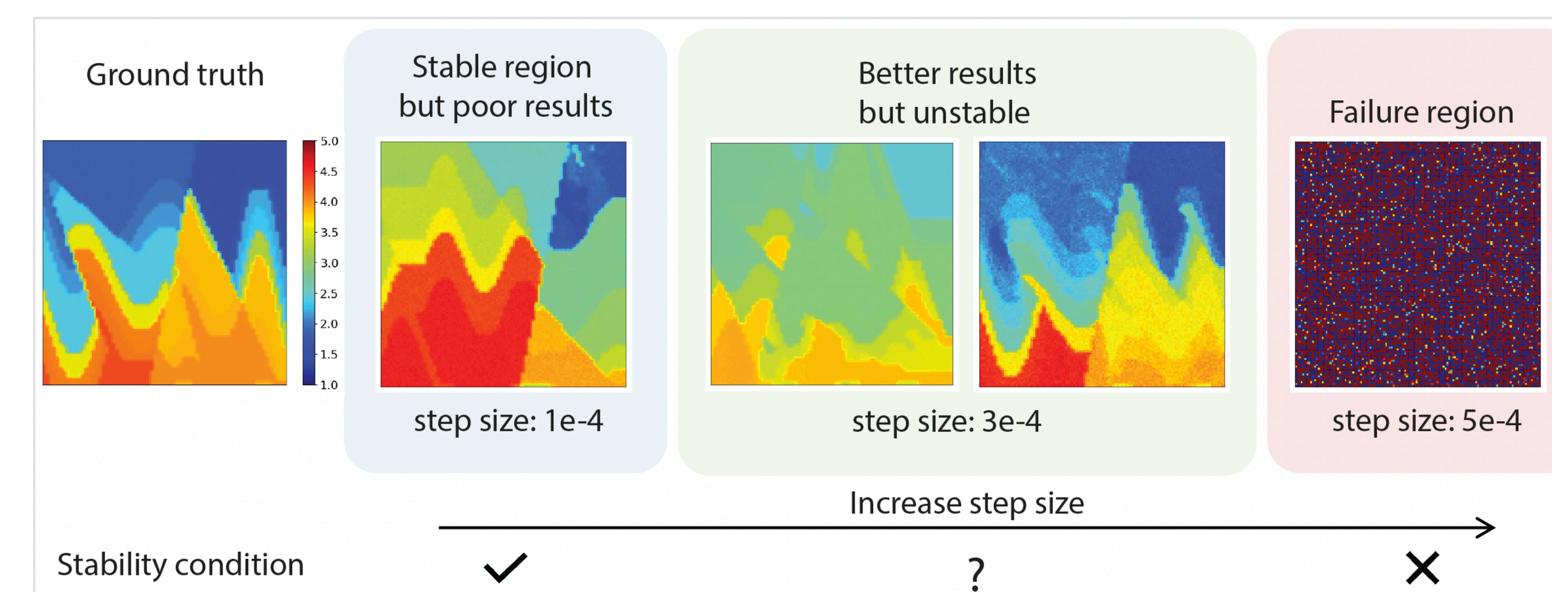

  - **Linear problem:** methods that leverage the linearity, e.g. DDNM, perform better.
- **Measurement sparsity:**
  - Higher measurement sparsity $\Longrightarrow$ larger advantage over baselines

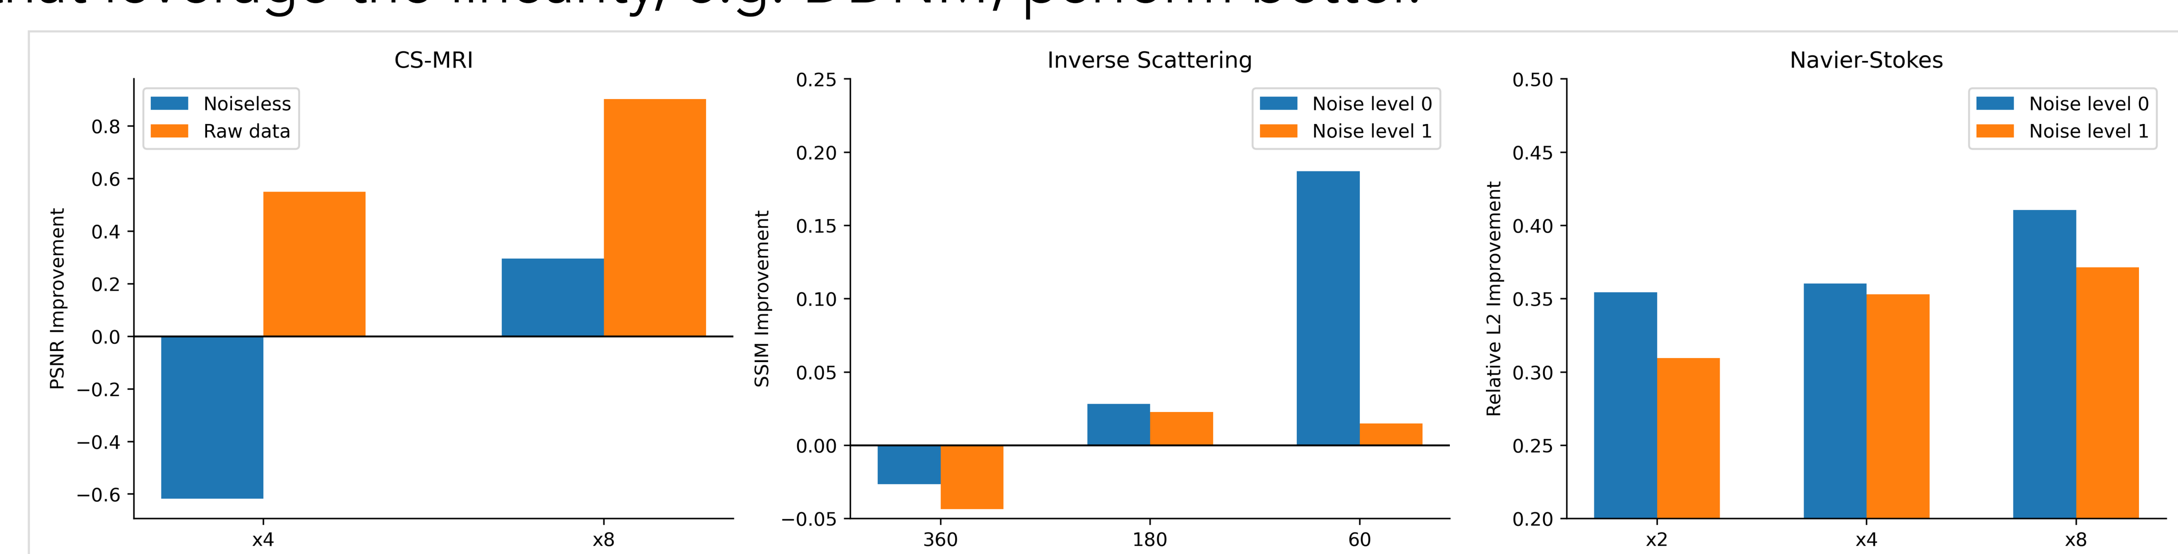

- **Forward model generalization:** PnPDP can generalize to different forward model configurations (e.g. different MRI sampling patterns) without additional hyper-parameter tuning.
- **OOD generalization:** PnPDP tends to bias toward the prior if the underlying source is out-of-distribution.