Uncertainty quantification for high-dimensional inverse problems with deep data-driven priors

Tobías I. Liaudat
IRFU, CEA Paris-Saclay

https://tobias-liaudat.github.io

Lumière conference, Institut Pascal, Saclay

16th January 2026

Inverse imaging problems

General inverse problem model

$$Y \sim P(\Phi(x)) \xrightarrow{\text{linear case}} y = \Phi x + n \qquad (1)$$

- $Y = y \in \mathbb{R}^M$: Observations/Measurements.
- $x \in \mathcal{X} \subset \mathbb{R}^N$: Signal/image to reconstruct from a given signal set $\mathcal{X}$.
- Φ : Forward (measurement) model including the deterministic physical part.
- P : Probabilistic model encompassing stochastic aspects of the observation y, e.g. noise n.
- **Objective:** estimate x from y given the model in Eq (1).

Inverse problem examples

General model

$$Y \sim P(\Phi(x)) \xrightarrow{\text{linear case}} y = \Phi x + n$$

Radio interferometric imaging

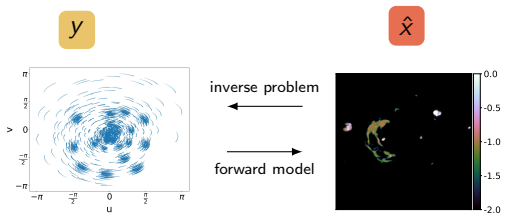

Magnetic Resonance Imaging

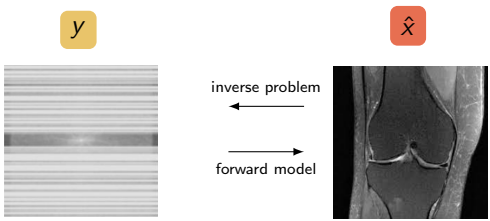

Other inverse problems: cosmological mass-mapping, PSF modelling, computed tomography imaging, deblurring, super-resolution, denoising, among others.

Ill-conditioned and ill-posed problems

The difficulty of inverse problems is that they are often **ill-conditioned** or **ill-posed** (in Hadamard's sense):

1. The solution may not exist.
2. The solution may not be unique.
3. The solution may not be stable.

We need to:

- Inject **prior information to regularise** the problem.
- **Quantify the uncertainty** of the reconstruction.

Bayesian inference provides a principled framework to address these two aspects.

Ill-conditioned and ill-posed problems

The difficulty of inverse problems is that they are often **ill-conditioned** or **ill-posed** (in Hadamard's sense):

1. The solution may not exist.
2. The solution may not be unique.
3. The solution may not be stable.

We need to:

- Inject **prior information to regularise** the problem.
- **Quantify the uncertainty** of the reconstruction.

Bayesian inference provides a principled framework to address these two aspects.

Ill-conditioned and ill-posed problems

The difficulty of inverse problems is that they are often **ill-conditioned** or **ill-posed** (in Hadamard's sense):

1. The solution may not exist.
2. The solution may not be unique.
3. The solution may not be stable.

We need to:

- Inject **prior information to regularise** the problem.
- **Quantify the uncertainty** of the reconstruction.

Bayesian inference provides a principled framework to address these two aspects.

Bayesian inference

Bayes' theorem

$$\underbrace{p(x \mid y, M)}_{\text{posterior}} = \frac{\overbrace{p(y \mid x, M)}^{\text{likelihood}}\ \overbrace{p(x \mid M)}^{\textit{prior}}}{\underbrace{p(y \mid M)}_{\text{evidence}}} = \frac{\overbrace{\mathcal{L}(x)}^{\text{likelihood}}\ \overbrace{\pi(x)}^{\text{prior}}}{\underbrace{\mathcal{Z}}_{\textit{evidence}}}$$

for a model M, observation y and signal/parameters x.

We often only require the unnormalised probability (disregarding $\mathcal{Z}$) to compute a point estimator or samples from the posterior distribution,

$$\overbrace{p(x \mid y, M)}^{\text{posterior}} \propto \overbrace{p(y \mid x, M)}^{\text{likelihood}}\ \overbrace{p(x \mid M)}^{\textit{prior}}$$

- We rely on *Markov Chain Monte Carlo* (MCMC) to estimate posterior samples,

Bayes' theorem

$$\underbrace{p(x \mid y, M)}_{\text{posterior}} = \frac{\overbrace{p(y \mid x, M)}^{\text{likelihood}}\ \overbrace{p(x \mid M)}^{\textit{prior}}}{\underbrace{p(y \mid M)}_{\text{evidence}}} = \frac{\overbrace{\mathcal{L}(x)}^{\text{likelihood}}\ \overbrace{\pi(x)}^{\text{prior}}}{\underbrace{\mathcal{Z}}_{\textit{evidence}}}$$

for a model M, observation y and signal/parameters x.

We often only require the unnormalised probability (disregarding $\mathcal{Z}$) to compute a point estimator or samples from the posterior distribution,

$$\overbrace{p(x \mid y, M)}^{\text{posterior}} \propto \overbrace{p(y \mid x, M)}^{\text{likelihood}}\ \overbrace{p(x \mid M)}^{\textit{prior}}$$

- We rely on *Markov Chain Monte Carlo* (MCMC) to estimate posterior samples,

Point estimates and priors

We select a point estimate to use as reconstruction, for example:

- Minimum mean squared error (MMSE): $\hat{x}_{M,\text{MMSE}} = \mathbb{E}[\, x \mid y, M \,]$ (posterior mean).
- *Maximum-a-posteriori* (MAP): $\hat{x}_{M,\text{MAP}} = \arg\max_{x \in \mathbb{R}^N} p(x \mid y, M)$ (posterior mode).

Then,

1. The likelihood is based on the physics of the inverse problem.
2. We choose the prior based on our previous knowledge of $\mathcal{X}$.
3. We usually characterise the high-dimensional posterior through posterior samples.

However, MCMC sampling can be prohibitively expensive in some settings...

Point estimates and priors

We select a point estimate to use as reconstruction, for example:

- Minimum mean squared error (MMSE): $\hat{x}_{M,\text{MMSE}} = \mathbb{E}[\, x \mid y, M \,]$ (posterior mean).
- *Maximum-a-posteriori* (MAP): $\hat{x}_{M,\text{MAP}} = \arg\max_{x \in \mathbb{R}^N} p(x \mid y, M)$ (posterior mode).

Then,

1. The likelihood is based on the physics of the inverse problem.
2. We choose the prior based on our previous knowledge of $\mathcal{X}$.
3. We usually characterise the high-dimensional posterior through posterior samples.

However, MCMC sampling can be prohibitively expensive in some settings...

Point estimates and priors

We select a point estimate to use as reconstruction, for example:

- Minimum mean squared error (MMSE): $\hat{x}_{M,\text{MMSE}} = \mathbb{E}[\, x \mid y, M \,]$ (posterior mean).
- *Maximum-a-posteriori* (MAP): $\hat{x}_{M,\text{MAP}} = \arg\max_{x \in \mathbb{R}^N} p(x \mid y, M)$ (posterior mode).

Then,

1. The likelihood is based on the physics of the inverse problem.
2. We choose the prior based on our previous knowledge of $\mathcal{X}$.
3. We usually characterise the high-dimensional posterior through posterior samples.

However, MCMC sampling can be prohibitively expensive in some settings...

Challenges

Challenges of high-dimensional Bayesian inference for inverse problems:

- High-dimensional parameter space, *i.e.*, $x \in \mathbb{R}^N$ with large N.
- Large data volume, *i.e.*, $y \in \mathbb{R}^M$ with large M.
- Computationally expensive forward model, $\Phi : \mathbb{R}^N \to \mathbb{R}^M$.

Additional challenges for inverse problems:

- The signals to reconstruct have more complex structure, $x \in \mathcal{X}$.
  - Instruments are getting more powerful.
  - Handcrafted priors (e.g., sparsity, smoothness) are not expressive enough.
- We need *scalable* ways to compare models without access to ground truth data.
  - Only having access to the observations y.

Challenges

Challenges of high-dimensional Bayesian inference for inverse problems:

- High-dimensional parameter space, *i.e.*, $x \in \mathbb{R}^N$ with large N.
- Large data volume, *i.e.*, $y \in \mathbb{R}^M$ with large M.
- Computationally expensive forward model, $\Phi : \mathbb{R}^N \to \mathbb{R}^M$.

Additional challenges for inverse problems:

- The signals to reconstruct have more complex structure, $x \in \mathcal{X}$.
  - Instruments are getting more powerful.
  - Handcrafted priors (e.g., sparsity, smoothness) are not expressive enough.
- We need *scalable* ways to compare models without access to ground truth data.
  - Only having access to the observations y.

Motivation: SKA's radio interferometer

Artist's impression of the Square Kilometre Array (SKA) Observatory.

SKA sites

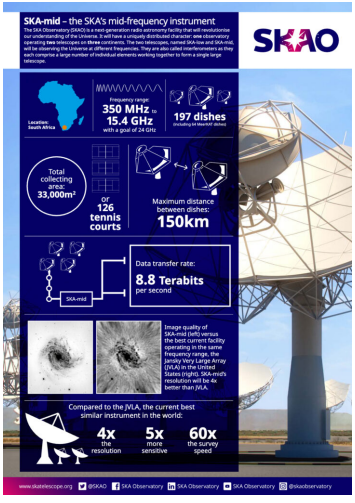

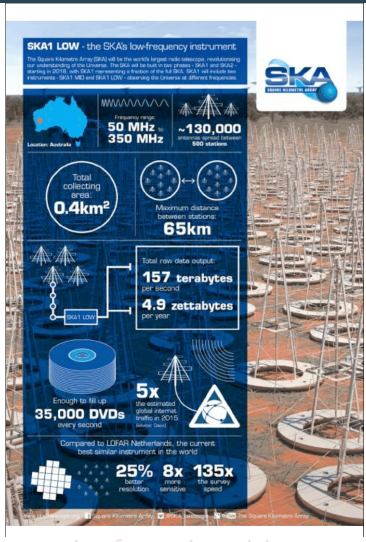

All 3 challenges: High-dimensional, large data volume, expensive forward model.

SKA sites

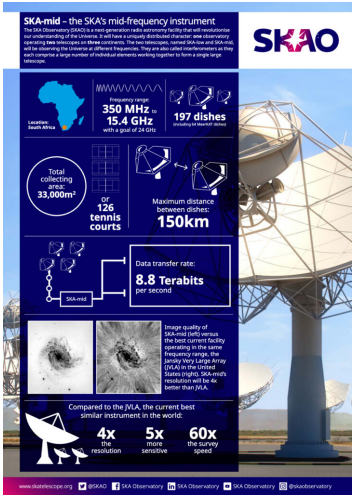

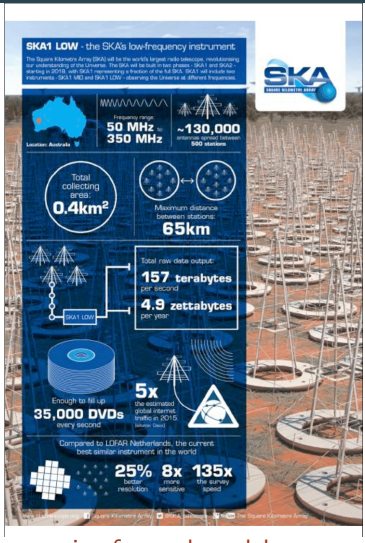

All 3 challenges: High-dimensional, large data volume, expensive forward model.

Goals

We want our methodology for inverse problems:

- **Computationally efficient** (optimisation over sampling)
- **Physics-informed** (robust and interpretable)
- **Expressive data-driven AI priors** (enhance reconstruction quality, reduce bias)
- **Quantify uncertainties** (for scientific inference)

Bonus:

- **Fast model comparison in high dimensions**

Goals

We want our methodology for inverse problems:

- **Computationally efficient** (optimisation over sampling)
- **Physics-informed** (robust and interpretable)
- **Expressive data-driven AI priors** (enhance reconstruction quality, reduce bias)
- **Quantify uncertainties** (for scientific inference)

Bonus:

- **Fast model comparison in high dimensions**

Outline

1. Physics-informed AI
2. Physics-informed AI + UQ
3. Physics-informed AI + UQ + Calibration
4. Bonus: Model comparison in high dimensions

Physics-informed AI

Physics-informed AI reconstructions

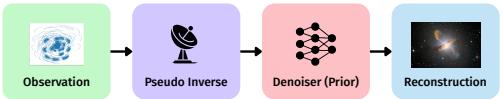

Learned post-processing

Physics-informed AI reconstructions

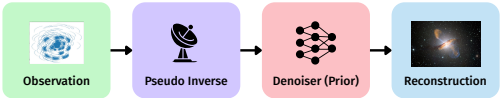

Learned post-processing

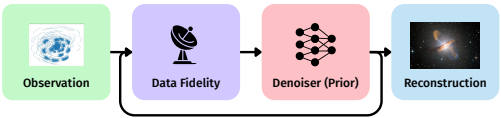

Plug-and-Play (PnP)

Physics-informed AI reconstructions

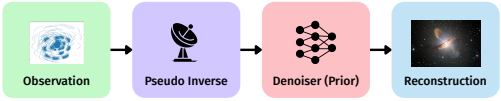

Learned post-processing

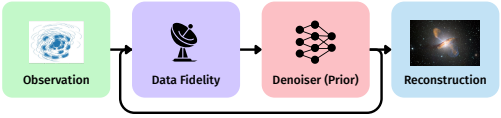

Plug-and-Play (PnP)

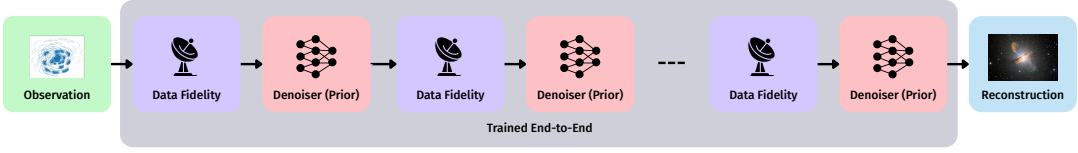

Unrolled ($n_{\text{unrolled}} \ll n_{\text{PnP}}$)

Physics-informed AI reconstructions

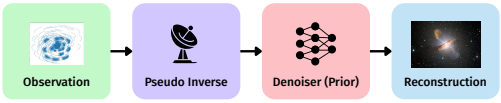

Learned post-processing

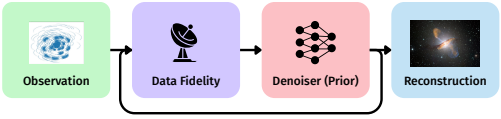

Plug-and-Play (PnP)

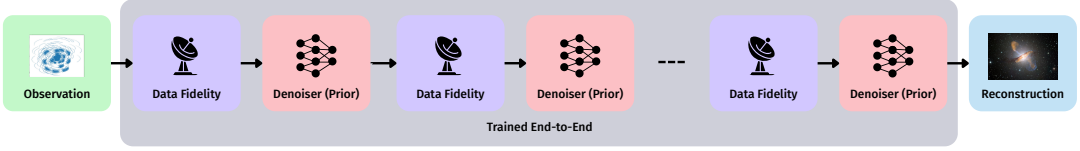

Unrolled ($n_{\text{unrolled}} \ll n_{\text{PnP}}$)

Based on the review (McEwen & Liaudat, Liaudat 2026) on **high-dimensional UQ with data-driven AI priors**

Learned post-processing

- Obtain a first reconstruction $\hat{x}_0$ using a physics-based method,*e.g.*, the pseudo-inverse or a classical method.
- Train a neural network to enhance $\hat{x}_0$, *i.e.*, $\hat{x} = \mathrm{NN}_\theta(\hat{x}_0)$.

$\rightarrow$ Fast, but limited by the initial reconstruction quality.

$\rightarrow$ No guarantee of data-consistency.

Plug-and-Play (PnP)

- Solve the inverse problem through an iterative scheme alternating between:
  - A data-consistency step based on the physics model, $\nabla p(y \mid x, M)$
  - A denoising step using a neural network, $D(\cdot)$, as image prior, $\nabla p(x \mid M)$.

x

→ Under some constraints on D we can guarantee convergence (Pesquet et al., 2021; Ryu et al., 2019).

→ An extension of a convex optimisation algorithm with a sparse wavelet prior.

→ More robust and interpretable than learned post-processing.

→ Large number of iterations until convergence.

The link between denoising and the prior is given by **Tweedie's formula**.

Plug-and-Play (PnP)

- Solve the inverse problem through an iterative scheme alternating between:
  - A data-consistency step based on the physics model, $\nabla p(y \mid x, M)$
  - A denoising step using a neural network, $D(\cdot)$, as image prior, $\nabla p(x \mid M)$.

  x

→ Under some constraints on D we can guarantee convergence (Pesquet et al., 2021; Ryu et al., 2019).

→ An extension of a convex optimisation algorithm with a sparse wavelet prior.

→ More robust and interpretable than learned post-processing.

→ Large number of iterations until convergence.

The link between denoising and the prior is given by **Tweedie's formula**.

Plug-and-Play (PnP)

- Solve the inverse problem through an iterative scheme alternating between:
  - A data-consistency step based on the physics model, $\nabla p(y \mid x, M)$
  - A denoising step using a neural network, $D(\cdot)$, as image prior, $\nabla p(x \mid M)$.

  x

→ Under some constraints on D we can guarantee convergence (Pesquet et al., 2021; Ryu et al., 2019).

→ An extension of a convex optimisation algorithm with a sparse wavelet prior.

→ More robust and interpretable than learned post-processing.

→ Large number of iterations until convergence.

The link between denoising and the prior is given by **Tweedie's formula**.

Plug-and-Play (PnP): Tweedie's formula

Tweedie's formula connects denoising with the gradient of the log-prior (Efron, 2011)

For a Gaussian denoising problem $z = x + \sigma\omega$ with $z \mid x \sim \mathcal{N}(x, \sigma^2 I)$ where $\omega \sim \mathcal{N}(0, I)$. We assume x has a marginal distribution $x \sim p_x(\cdot)$. Then, the MMSE denoiser $\mathbb{E}[x \mid z]$ satisfies:

$$\mathbb{E}[x \mid z] = z + \sigma^2 \nabla_z \log p_z(z),$$

where $p_z(z) = \int p_x(x)\varphi_\sigma(z - x)\mathrm{d}x$ is the marginal density of z and φ_σ is the Gaussian density of ω.

- We interpret a **denoising neural network as the gradient of the log prior** learned from a dataset of images.
- One of the basis of diffusion models.

Unrolled methods

- Unroll a fixed number of iterations n_{unrolled} of an iterative algorithm solving the inverse problem.
- Use a neural network to represent the implicit prior from data.
- Learn the parameters, *e.g.*, denoiser, step-sizes, etc.. end-to-end through backprop.

$\rightarrow$ Forces convergeence in n_{unrolled} iterations and faster than PnP since $n_{\text{unrolled}} \ll n_{\text{PnP}}$.

$\rightarrow$ End-to-end training adapts the model to the training set.

$\rightarrow$ Better performance but less robust than PnP.

Physics-informed AI + UQ

Direct UQ estimation

- Magnitude of residual: train a network to estimate residuals.
- Gaussian per pixel: train a network to estimate the standard deviation.
- Classification for regression ranges: train a classifier with softmax output to estimate distribution of pixel values.
- Pixelwise quantile regression: train network to estimate lower/upper quantiles for $1 - \alpha$ uncertainty level, using quantile (pinball) loss.

Heuristic $\rightarrow$ **no statistical guarantees.**

Direct UQ estimation

- Magnitude of residual: train a network to estimate residuals.
- Gaussian per pixel: train a network to estimate the standard deviation.
- Classification for regression ranges: train a classifier with softmax output to estimate distribution of pixel values.
- Pixelwise quantile regression: train network to estimate lower/upper quantiles for $1 - \alpha$ uncertainty level, using quantile (pinball) loss.

Heuristic $\rightarrow$ **no statistical guarantees**.

PnP UQ: convex probability concentration for UQ

Posterior credible region:

$$p(\mathbf{x} \in C_\alpha|\mathbf{y}) = \int_{\mathbf{x}\in\mathbb{R}^N} p(\mathbf{x}|\mathbf{y})\mathbb{1}_{C_\alpha}\mathrm{d}\mathbf{x} = 1-\alpha.$$

Consider the **highest posterior density (HPD) region**

$$C_\alpha^* = \left\{\mathbf{x} : -\log p(\mathbf{x}|\mathbf{y}) \leq \gamma_\alpha\right\}, \quad \text{with } \gamma_\alpha \in \mathbb{R}, \quad \text{and } p(\mathbf{x} \in C_\alpha^*|\mathbf{y}) = 1-\alpha \text{ holds.}$$

Bound of HPD region for log-concave distributions (Pereyra, 2017)

Suppose the posterior $\log p(\mathbf{x}|\mathbf{y}) \propto \log \mathcal{L}(\mathbf{x}) + \log \pi(\mathbf{x})$ is log-concave on $\mathbb{R}^N$. Then, for any $\alpha \in (4\exp[(-N/3)], 1)$, the HPD region C_α^* is contained by

$$\hat{C}_\alpha = \left\{\mathbf{x} : \log \mathcal{L}(\mathbf{x}) + \log \pi(\mathbf{x}) \leq \hat{\gamma}_\alpha = \log \mathcal{L}(\hat{\mathbf{x}}_{\text{MAP}}) + \log \pi(\hat{\mathbf{x}}_{\text{MAP}}) + \sqrt{N}\tau_\alpha + N\right\},$$

with a positive constant $\tau_\alpha = \sqrt{16\log(3/\alpha)}$ independent of $p(\mathbf{x}|\mathbf{y})$.

We only need to evaluate $\log \mathcal{L}(\mathbf{x}) + \log \pi(\mathbf{x})$ on the MAP estimation $\hat{\mathbf{x}}_{\text{MAP}}$!

PnP UQ: convex probability concentration for UQ

Posterior credible region:

$$p(\mathbf{x} \in C_\alpha|\mathbf{y}) = \int_{\mathbf{x}\in\mathbb{R}^N} p(\mathbf{x}|\mathbf{y})\mathbb{1}_{C_\alpha}\mathrm{d}\mathbf{x} = 1-\alpha.$$

Consider the **highest posterior density (HPD) region**

$$C_\alpha^* = \left\{\mathbf{x} : -\log p(\mathbf{x}|\mathbf{y}) \le \gamma_\alpha\right\}, \quad \text{with } \gamma_\alpha \in \mathbb{R}, \quad \text{and } p(\mathbf{x} \in C_\alpha^*|\mathbf{y}) = 1-\alpha \text{ holds.}$$

Bound of HPD region for log-concave distributions (Pereyra, 2017)

Suppose the posterior $\log p(\mathbf{x}|\mathbf{y}) \propto \log \mathcal{L}(\mathbf{x}) + \log \pi(\mathbf{x})$ is **log-concave** on $\mathbb{R}^N$. Then, for any $\alpha \in (4\exp[(-N/3)], 1)$, the HPD region C_α^* is contained by

$$\hat{C}_\alpha = \left\{\mathbf{x} : \log \mathcal{L}(\mathbf{x}) + \log \pi(\mathbf{x}) \le \hat{\gamma}_\alpha = \log \mathcal{L}(\hat{\mathbf{x}}_{\text{MAP}}) + \log \pi(\hat{\mathbf{x}}_{\text{MAP}}) + \sqrt{N}\tau_\alpha + N\right\},$$

with a positive constant $\tau_\alpha = \sqrt{16\log(3/\alpha)}$ independent of $p(\mathbf{x}|\mathbf{y})$.

We only need to evaluate $\log \mathcal{L}(\mathbf{x}) + \log \pi(\mathbf{x})$ on the MAP estimation $\hat{\mathbf{x}}_{\text{MAP}}$!

PnP UQ: convex probability concentration for UQ

Posterior credible region:

$$p(\mathbf{x} \in C_\alpha|\mathbf{y}) = \int_{\mathbf{x}\in\mathbb{R}^N} p(\mathbf{x}|\mathbf{y})\mathbb{1}_{C_\alpha}\mathrm{d}\mathbf{x} = 1-\alpha.$$

Consider the **highest posterior density (HPD) region**

$$C_\alpha^* = \left\{\mathbf{x} : -\log p(\mathbf{x}|\mathbf{y}) \le \gamma_\alpha\right\}, \quad \text{with } \gamma_\alpha \in \mathbb{R}, \quad \text{and } p(\mathbf{x} \in C_\alpha^*|\mathbf{y}) = 1-\alpha \text{ holds.}$$

Bound of HPD region for log-concave distributions (Pereyra, 2017)

Suppose the posterior $\log p(\mathbf{x}|\mathbf{y}) \propto \log \mathcal{L}(\mathbf{x}) + \log \pi(\mathbf{x})$ is **log-concave** on $\mathbb{R}^N$. Then, for any $\alpha \in (4\exp[(-N/3)], 1)$, the HPD region C_α^* is contained by

$$\hat{C}_\alpha = \left\{\mathbf{x} : \log \mathcal{L}(\mathbf{x}) + \log \pi(\mathbf{x}) \le \hat{\gamma}_\alpha = \log \mathcal{L}(\hat{\mathbf{x}}_{\mathrm{MAP}}) + \log \pi(\hat{\mathbf{x}}_{\mathrm{MAP}}) + \sqrt{N}\tau_\alpha + N\right\},$$

with a positive constant $\tau_\alpha = \sqrt{16\log(3/\alpha)}$ independent of $p(\mathbf{x}|\mathbf{y})$.

We only need to evaluate $\log \mathcal{L}(\mathbf{x}) + \log \pi(\mathbf{x})$ on the MAP estimation $\hat{\mathbf{x}}_{\mathrm{MAP}}$!

Leveraging the approximate HPD region for UQ

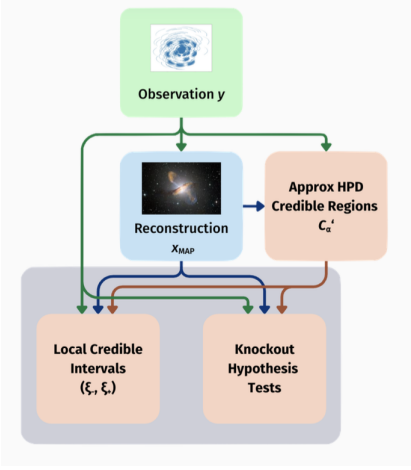

Img credit: Jason McEwen

Leveraging the approximate HPD region for UQ

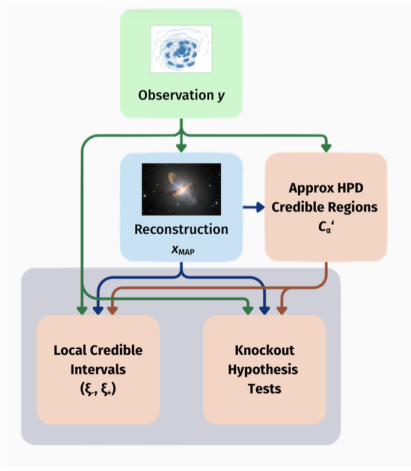

Img credit: Jason McEwen

UQ techniques:

1. Hypothesis test with significance α
   - e.g. with respect to a surrogate image with an inpainted structure.
2. Local credible intervals (LCI)
   - Test the approx HPD region for each pixel or super-pixel in the image.
3. Fast LCIs at different scales
   - Test the approx HPD region from the coefficients of a multi-resolution decomposition of the image.

Refs: Cai et al. (2018), Liaudat et al. (2024), and Pereyra (2017)

Hypothesis test

Hypothesis testing of physical structure (Cai et al., 2018, August; Pereyra, 2017)

1. Remove the structure of interest from the MAP estimate $\hat{\mathbf{x}}_{\text{MAP}}$.
2. Inpaint removed region to create a surrogate test image $\mathbf{x}_{\text{sur}}$.
3. Test if $\mathbf{x}_{\text{sur}} \in \hat{C}_\alpha$:
   - If no, *i.e.*, $\mathbf{x}_{\text{sur}} \notin \hat{C}_\alpha$, reject hypothesis at significance α and conclude the **structure is most likely physical**.
   - If yes, *i.e.*, $\mathbf{x}_{\text{sur}} \in \hat{C}_\alpha$, we cannot reject the hypothesis as uncertainty is too high to draw conclusions.

Learned data-driven convex regulariser

We use the neural-network-based convex regulariser $R : \mathbb{R}^N \mapsto \mathbb{R}$,
(Goujon et al., 2023; Liaudat et al., 2024)

$$R(\mathbf{x}) = \sum_{n=1}^{N_C} \sum_{k} \psi_n \left((\mathbf{h}_n * \mathbf{x})[k] \right),$$

- ψ_n are learned convex profile functions with Lipschitz continuous derivate
- There are N_C learned convolutional filters $\mathbf{h}_n$
- R is trained as a (multi-)gradient step denoiser

Properties:

1. **Convex + explicit potential** → leverage HPD approximation for UQ
2. **Smooth regulariser with known Lipschitz constant** → MAP convergence guarantees

Learned data-driven convex regulariser

We use the neural-network-based convex regulariser $R : \mathbb{R}^N \mapsto \mathbb{R}$,
(Goujon et al., 2023; Liaudat et al., 2024)

$$R(\mathbf{x}) = \sum_{n=1}^{N_C} \sum_{k} \psi_n \left((\mathbf{h}_n * \mathbf{x})[k] \right),$$

- ψ_n are learned convex profile functions with Lipschitz continuous derivate
- There are N_C learned convolutional filters $\mathbf{h}_n$
- R is trained as a (multi-)gradient step denoiser

Properties:

1. **Convex + explicit potential** $\rightarrow$ leverage HPD approximation for UQ
2. **Smooth regulariser with known Lipschitz constant** $\rightarrow$ MAP convergence guarantees

Reconstructed images

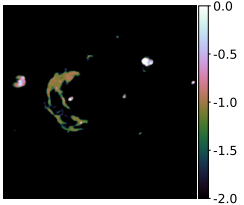

Ground truth

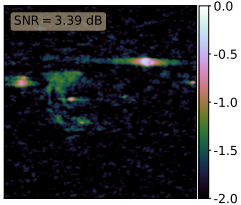

Dirty image

SNR = 3.39 dB

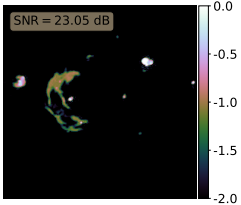

Reconstruction (classical)

SNR = 23.05 dB

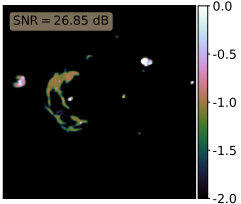

Reconstruction (learned)

SNR = 26.85 dB

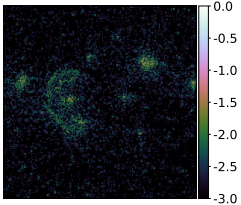

Error (classical)

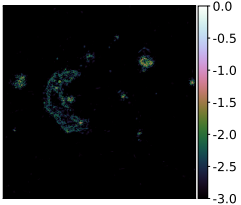

Error (learned)

Ref: Liaudat et al. (2024)

Hypothesis test

Scalable hypothesis testing for structure in the reconstruction

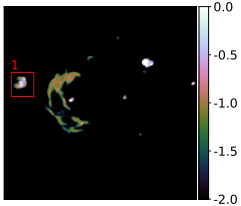

MAP reconstruction

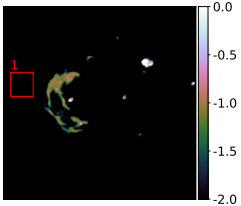

Inpainted surrogate

Is the blob physical? $\rightarrow$ Yes

Ref: Liaudat et al. (2024)

Fast Local Credible Intervals

Results for $M \approx 2.4 \times 10^5$ and $N = 256 \times 256$.

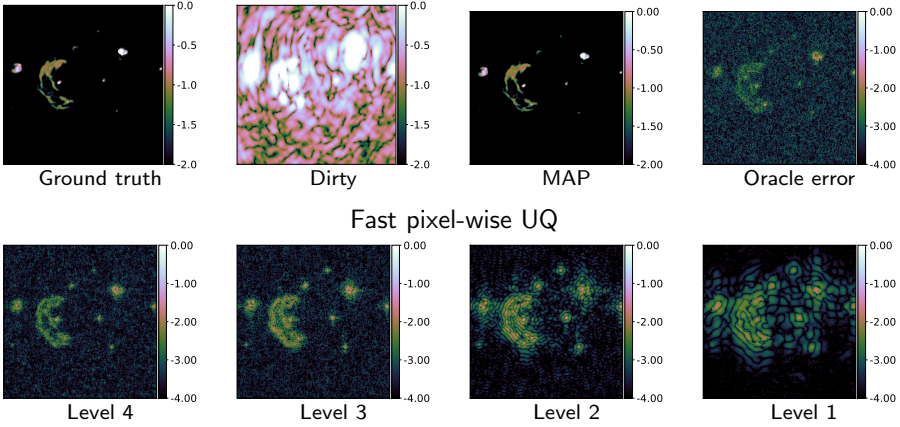

- Computation wall-clock time: MAP estimation $\rightarrow$ 137.0s, **fast pixel UQ $\rightarrow$ 1.84s**
- Measurement operator evaluations: 28 $\rightarrow$ 10^6 times lower than MCMC sampling.

QuantifAI code

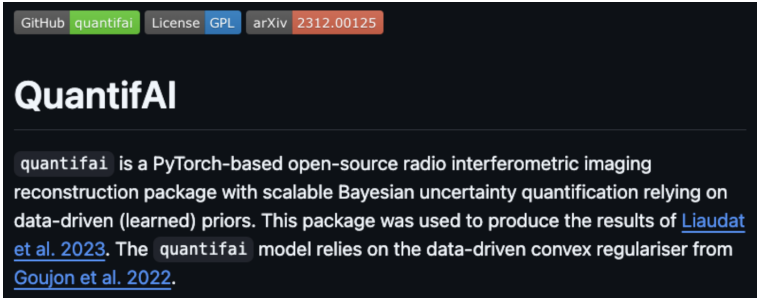

- QuantifAI GitHub code: github.com/astro-informatics/quantifai
- Entirely implemented in Pytorch: automatic differentiation + GPU acceleration
- Ref: Liaudat et al. (2024)

Unrolled generative UQ estimation

We focus on two main approaches for **generating approximate posterior samples** ($\hat{\mathbf{x}} \sim \hat{p}(\mathbf{x} \mid \mathbf{y})$) in a physics-informed manner:

- Denoising diffusion models.
- Regularized Conditional Generative Adversarial Networks (cGANs).

Extra approach exploiting data symmetries:

- Equivariant Bootstrap.

Denoising diffusion models

Denoising diffusion models (DDMs) (Ho et al., 2020; Song & Ermon, 2019; Song et al., 2020) are generative models that learn to sample from a data distribution, $x \sim p_x(\cdot)$ by reversing a gradual noising process.

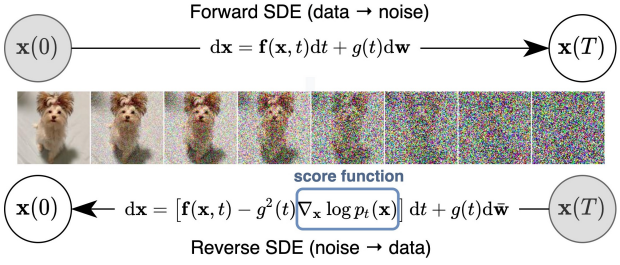

- Very expressive models generating high-quality posterior samples.

Denoising diffusion models for inverse problems

We can adapt DDMs to solve inverse problems by incorporating the physics model (likelihood) in the reverse diffusion process (Chung et al., 2022).
$\rightarrow$ Combine the expressive learnt prior with physical data-consistency.

- Highly expressive models.
- Likelihood is analytically intractable due to the dependence of the diffusion process on time so approximations need to be made.
- Computationally expensive due to the large number of neural network evaluations required.
  - We need $n_{\text{diffusion}}$ steps per sample produced.

See Daras et al. (2024) for a review on the subject.

Generative Adversarial Networks (GANs)

Standard GANs achieve high-fidelity image generation.

They had some challenges:

- Difficult to train.
- Mode collapse.

These were recently addressed:

- Wasserstein GAN loss (Arjovsky et al., 2017).
- Regularisation to avoid mode collapse (Bendel et al., 2023).

Regularized Conditional GANs

The rcGAN allows us to **generate approximate posterior samples** ($\hat{\mathbf{x}} \sim \hat{p}(\mathbf{x} \mid \mathbf{y})$) by conditioning the generator on the observations $\mathbf{y}$ (Bendel et al., 2023).
$\rightarrow$ The novelty is a regularisation term in the training that **rewards sampling diversity** and avoids mode collapse.

Main points of the approach:

- Under a simplifying Gaussian assumption, the first two moments of the approximated posterior (mean and covariance) match the true posterior.
- We add a conditioning on other variables like the pseudo-inverse.
- Extremely-fast sampling $\rightarrow$ One NN evaluation per sample.

Exploited this approach for:

- Radio interferometric imaging (Mars et al., Liaudat 2025).
- Mass-mapping in cosmology (Whitney et al., Liaudat 2025).

Regularized Conditional GANs

The rcGAN allows us to **generate approximate posterior samples** ($\hat{\mathbf{x}} \sim \hat{p}(\mathbf{x} \mid \mathbf{y})$) by conditioning the generator on the observations $\mathbf{y}$ (Bendel et al., 2023).
$\rightarrow$ The novelty is a regularisation term in the training that **rewards sampling diversity** and avoids mode collapse.

Main points of the approach:

- Under a simplifying Gaussian assumption, the first two moments of the approximated posterior (mean and covariance) match the true posterior.
- We add a conditioning on other variables like the pseudo-inverse.
- Extremely-fast sampling $\rightarrow$ One NN evaluation per sample.

Exploited this approach for:

- Radio interferometric imaging (Mars et al., Liaudat 2025).
- Mass-mapping in cosmology (Whitney et al., Liaudat 2025).

Regularized Conditional GANs

The rcGAN allows us to **generate approximate posterior samples** ($\hat{\mathbf{x}} \sim \hat{p}(\mathbf{x} \mid \mathbf{y})$) by conditioning the generator on the observations $\mathbf{y}$ (Bendel et al., 2023).
$\rightarrow$ The novelty is a regularisation term in the training that **rewards sampling diversity** and avoids mode collapse.

Main points of the approach:

- Under a simplifying Gaussian assumption, the first two moments of the approximated posterior (mean and covariance) match the true posterior.
- We add a conditioning on other variables like the pseudo-inverse.
- Extremely-fast sampling $\rightarrow$ One NN evaluation per sample.

Exploited this approach for:

- Radio interferometric imaging (Mars et al., Liaudat 2025).
- Mass-mapping in cosmology (Whitney et al., Liaudat 2025).

Regularised conditional GAN for RI imaging

Ref: Mars et al. (2025).

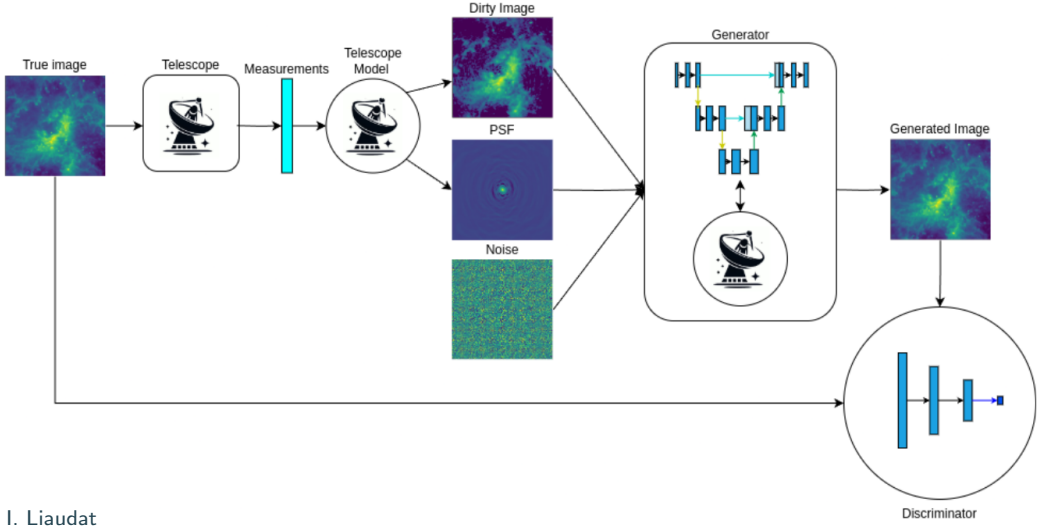

Equivariant bootstrap

Based on the work of Tachella and Pereyra (2023).

Given an observation model $y = Ax + n$, *e.g.*, RI imaging, group actions $\{T_g\}_{g\in\mathcal{G}}$ such that $T_g x \in \mathcal{X}$ (known symmetries) and a reconstruction method $\hat{x}(y) = f(y)$:

Procedure:
For $i = 1, \dots, N$:

1. Draw transform g_i from $\mathcal{G}$ and sample noise $n_i \sim \mathcal{N}(0, \sigma^2 I)$
2. Build bootsrap measurement $\tilde{y}_i = AT_{g_i}\hat{x}(y) + n_i := A_{g_i}\hat{x}(y) + n_i$
3. Reconstruct $\tilde{x}_i = T_{g_i}^{-1}\hat{x}(\tilde{y}_i)$
4. Collect error estimate $e_i = \|\hat{x}(y) - \tilde{x}_i\|^2$ and bootstrap samples $\tilde{x}_i$.

Pixel-wise UQ maps: From the collection of N bootstrap samples, $\{\tilde{x}_i\}_{i=1}^N$, we build confidence regions, $\mathcal{C}_\alpha$, for $x^\star$ (ground truth) using q_α the top α-quantile of the samples $\{|\hat{x}(y) - \tilde{x}_i|\}_{i=1}^N$, with $\mathcal{C}_\alpha = \{x : |x - \hat{x}(y)| < q_\alpha\}$.

Equivariant bootstrap

Based on the work of Tachella and Pereyra (2023).

Given an observation model $y = Ax + n$, *e.g.*, RI imaging, group actions $\{T_g\}_{g\in\mathcal{G}}$ such that $T_g x \in \mathcal{X}$ (known symmetries) and a reconstruction method $\hat{x}(y) = f(y)$:

Procedure:
For $i = 1, \dots, N$:

1. Draw transform g_i from $\mathcal{G}$ and sample noise $n_i \sim \mathcal{N}(0, \sigma^2 I)$
2. Build bootsrap measurement $\tilde{y}_i = AT_{g_i}\hat{x}(y) + n_i := A_{g_i}\hat{x}(y) + n_i$
3. Reconstruct $\tilde{x}_i = T_{g_i}^{-1}\hat{x}(\tilde{y}_i)$
4. Collect error estimate $e_i = \|\hat{x}(y) - \tilde{x}_i\|^2$ and bootstrap samples $\tilde{x}_i$.

Pixel-wise UQ maps: From the collection of N bootstrap samples, $\{\tilde{x}_i\}_{i=1}^N$, we build confidence regions, $\mathcal{C}_\alpha$, for $x^\star$ (ground truth) using q_α the top α-quantile of the samples $\{|\hat{x}(y) - \tilde{x}_i|\}_{i=1}^N$, with $\mathcal{C}_\alpha = \{x : |x - \hat{x}(y)| < q_\alpha\}$.

Equivariant bootstrap

Based on the work of Tachella and Pereyra (2023).

Given an observation model $y = Ax + n$, *e.g.*, RI imaging, group actions $\{T_g\}_{g\in\mathcal{G}}$ such that $T_g x \in \mathcal{X}$ (known symmetries) and a reconstruction method $\hat{x}(y) = f(y)$:

Procedure:
For $i = 1, \dots, N$:

1. Draw transform g_i from $\mathcal{G}$ and sample noise $n_i \sim \mathcal{N}(0, \sigma^2 I)$
2. Build bootsrap measurement $\tilde{y}_i = AT_{g_i}\hat{x}(y) + n_i := A_{g_i}\hat{x}(y) + n_i$
3. Reconstruct $\tilde{x}_i = T_{g_i}^{-1}\hat{x}(\tilde{y}_i)$
4. Collect error estimate $e_i = \|\hat{x}(y) - \tilde{x}_i\|^2$ and bootstrap samples $\tilde{x}_i$.

Pixel-wise UQ maps: From the collection of N bootstrap samples, $\{\tilde{x}_i\}_{i=1}^N$, we build confidence regions, $\mathcal{C}_\alpha$, for $x^\star$ (ground truth) using q_α the top α-quantile of the samples $\{|\hat{x}(y) - \tilde{x}_i|\}_{i=1}^N$, with $\mathcal{C}_\alpha = \{x : |x - \hat{x}(y)| < q_\alpha\}$.

Equivariant bootstrap

Main idea: Assuming that $\mathcal{X}$ is $\mathcal{G}$-invariant, we can have access to multiple virtual forward operators, $AT_{g_i} := A_{g_i}$. If T_{g_i} is properly chosen based on $\mathcal{X}$, A and $\hat{x}(\cdot)$, the composition AT_{g_i} can have different null spaces than A helping to probe the variability of the estimator $\hat{x}(y)$ and characterize its uncertainties with respect to $x^\star$ (ground truth).

Motivation:

- Unsupervised method → No ground truth or training required
- Independent of the reconstruction method and each sample trivially can run in parallel
- Well-suited to fast reconstruction methods, *e.g.* deterministic unrolled algorithms
- Carefully selected group transforms allow us to explore the big nullspace of the RI imaging forward operator and better characterise the errors

We applied this framework to the RI imaging problem (Cherif et al., Liaudat 2024).

Equivariant bootstrap

Main idea: Assuming that $\mathcal{X}$ is $\mathcal{G}$-invariant, we can have access to multiple virtual forward operators, $AT_{g_i} := A_{g_i}$. If T_{g_i} is properly chosen based on $\mathcal{X}$, A and $\hat{x}(\cdot)$, the composition AT_{g_i} can have different null spaces than A helping to probe the variability of the estimator $\hat{x}(y)$ and characterize its uncertainties with respect to $x^\star$ (ground truth).

Motivation:

- Unsupervised method $\rightarrow$ No ground truth or training required
- Independent of the reconstruction method and each sample trivially can run in parallel
- Well-suited to fast reconstruction methods, *e.g.* deterministic unrolled algorithms
- Carefully selected group transforms allow us to explore the big nullspace of the RI imaging forward operator and better characterise the errors

We applied this framework to the RI imaging problem (Cherif et al., Liaudat 2024).

UQ overview

1. Direct UQ estimation
   - Fast
   - Heuristic, no statistical guarantees
2. PnP UQ: convex probability concentration for UQ
   - Fast
   - Statistical guarantees under model convexity
   - Restricted to HPD-related UQ
3. Unrolled generative UQ estimation
   - rcGANs: fast, and Diffusion models: slow
   - Target posterior sampling but no statistical guarantees as approximations are made
4. Equivariant bootstrap
   - Method-agnostic (including deterministic methods) and unsupervised
   - Well adapted for fast reconstruction methods
   - Even if it shows that it reduces the error estimation biases
     $\rightarrow$ No statistical guarantees on the uncertainties

UQ overview

1. Direct UQ estimation
   - Fast
   - Heuristic, no statistical guarantees
2. PnP UQ: convex probability concentration for UQ
   - Fast
   - Statistical guarantees under model convexity
   - Restricted to HPD-related UQ
3. Unrolled generative UQ estimation
   - rcGANs: fast, and Diffusion models: slow
   - Target posterior sampling but no statistical guarantees as approximations are made
4. Equivariant bootstrap
   - Method-agnostic (including deterministic methods) and unsupervised
   - Well adapted for fast reconstruction methods
   - Even if it shows that it reduces the error estimation biases
     $\rightarrow$ No statistical guarantees on the uncertainties

UQ overview

1. Direct UQ estimation
   - Fast
   - Heuristic, no statistical guarantees
2. PnP UQ: convex probability concentration for UQ
   - Fast
   - Statistical guarantees under model convexity
   - Restricted to HPD-related UQ
3. Unrolled generative UQ estimation
   - rcGANs: fast, and Diffusion models: slow
   - Target posterior sampling but no statistical guarantees as approximations are made
4. Equivariant bootstrap
   - Method-agnostic (including deterministic methods) and unsupervised
   - Well adapted for fast reconstruction methods
   - Even if it shows that it reduces the error estimation biases
     $\rightarrow$ No statistical guarantees on the uncertainties

UQ overview

1. Direct UQ estimation
   - Fast
   - Heuristic, no statistical guarantees
2. PnP UQ: convex probability concentration for UQ
   - Fast
   - Statistical guarantees under model convexity
   - Restricted to HPD-related UQ
3. Unrolled generative UQ estimation
   - rcGANs: fast, and Diffusion models: slow
   - Target posterior sampling but no statistical guarantees as approximations are made
4. Equivariant bootstrap
   - Method-agnostic (including deterministic methods) and unsupervised
   - Well adapted for fast reconstruction methods
   - Even if it shows that it reduces the error estimation biases
     $\rightarrow$ No statistical guarantees on the uncertainties

Physics-informed AI + UQ + Calibration

Coverage test

We have studied ways to estimate uncertainty, but not if the uncertainty estimated is well-calibrated. Coverage plots allows us to validate the uncertainties.

- We want to compare the model's Bayesian probability with a frequentist interpretation.
- Compute a credible interval with the model.
- Check the empirical frequency at which the ground truth falls within the credible interval.

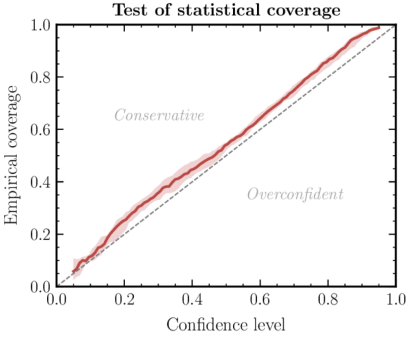

Calibration with conformal prediction

The estimated uncertainties can be calibrated using conformal prediction with **Risk-Controlling Prediction Sets (RCPS)** (Angelopoulos et al., 2022) by exploiting an exchangeable calibration dataset.

Given an estimator $\hat{f}(y)$, and lower and upper interval lengths, $\hat{l}(y)$ and $\hat{u}(y)$, we define the uncertainty interval for pixel (m, n) as:

$$\mathcal{T}_\lambda(y)_{(m,n)} = [\hat{f}(y)_{(m,n)} - \lambda\hat{l}(y)_{(m,n)}, \hat{f}(y)_{(m,n)} + \lambda\hat{u}(y)_{(m,n)}].$$

The RCPS procedure finds the smallest λ such that the interval $\mathcal{T}_\lambda(y)$ guarantees a desired risk level α (exploiting an upper confidence bound, *i.e.*, Hoeffding's).

Main points:

- **Distribution-free** uncertainty calibration with **statistical guarantees**.
- Guaranteed to be valid but not necessarily useful
  $\rightarrow$ We still need good initial uncertainties estimates.

We applied this calibration on top of the equivariant bootstrap UQ method (Cherif et al., Liaudat 2024).

MMGAN coverage test

Coverage plot and conformal prediction of MMGAN (Whitney et al., Liaudat 2025) for mass-mapping in cosmology.

- Extremely good coverage (without RCPS)
- Optimal coverage after RCPS calibration

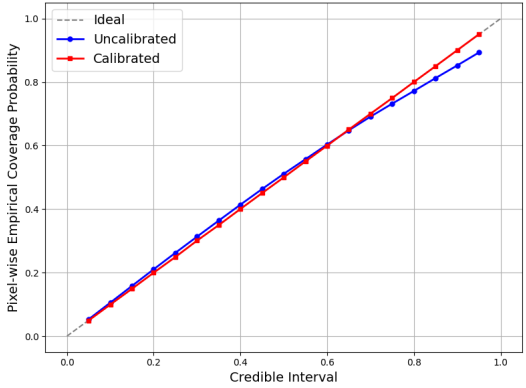

Summary

Summary

Inverse problem in imaging are often ill-posed and ill-conditioned and require
→ **Injecting prior information**, and **quantifying uncertainties** → **Bayesian framework**.

MCMC sampling is computationally prohibitive for many problems. Some goals are

- **Computationally efficient** (optimisation-based)
- **Physics-informed** (robust and interpretable)
- **Expressive data-driven AI priors** (enhance reconstruction quality and reduce bias)
- **Quantify uncertainties** (scientific inference)

PnP UQ with convexity (Liaudat et al., 2024) points towards these directions.
→ Unrolled methods are more expressive.

It is important to **validate the estimated uncertainties** (coverage test) and **calibrate** them if necessary (conformal prediction - RCPS).

Calibration helps to provide statistical guarantees on the model uncertainties.
→ A deterministic reconstruction mehod can be equipped with statistically valid UQ by exploiting equivariant bootstrap + RCPS as was done in (Cherif et al., Liaudat 2024).

References I

Angelopoulos, A. N., Kohli, A. P., Bates, S., Jordan, M., Malik, J., Alshaabi, T., Upadhyayula, S., & Romano, Y. (2022). Image-to-image regression with distribution-free uncertainty quantification and applications in imaging. In K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, & S. Sabato (Eds.), Proceedings of the 39th international conference on machine learning (pp. 717–730, Vol. 162). PMLR.

Arjovsky, M., Chintala, S., & Bottou, L. (2017).Wasserstein generative adversarial networks. International conference on machine learning, 214–223.

Bendel, M., Ahmad, R., & Schniter, P. (2023). A regularized conditional gan for posterior sampling in image recovery problems. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, & S. Levine (Eds.), Advances in neural information processing systems (pp. 68673–68684, Vol. 36). Curran Associates, Inc.

Cai, X., Pereyra, M., & McEwen, J. D. (2018).Uncertainty quantification for radio interferometric imaging: II. MAP estimation. Monthly Notices of the Royal Astronomical Society, 480(3), 4170–4182.

Cherif, M., Liaudat, T. I., Kern, J., Kervazo, C., & Bobin, J. (2024).Uncertainty quantification for fast reconstruction methods using augmented equivariant bootstrap: Application to radio interferometry. arXiv e-prints, Article arXiv:2410.23178, arXiv:2410.23178.

Chung, H., Kim, J., Mccann, M. T., Klasky, M. L., & Ye, J. C. (2022).Diffusion posterior sampling for general noisy inverse problems. arXiv preprint arXiv:2209.14687.

Daras, G., Chung, H., Lai, C.-H., Mitsufuji, Y., Ye, J. C., Milanfar, P., Dimakis, A. G., & Delbracio, M. (2024).A survey on diffusion models for inverse problems. arXiv preprint arXiv:2410.00083.

References II

Efron, B. (2011).Tweedie's formula and selection bias [PMID: 22505788]. Journal of the American Statistical Association, 106(496), 1602–1614.

Gneiting, T., & Raftery, A. E. (2007).Strictly proper scoring rules, prediction, and estimation. Journal of the American Statistical Association, 102(477), 359–378.

Goujon, A., Neumayer, S., Bohra, P., Ducotterd, S., & Unser, M. (2023).A neural-network-based convex regularizer for inverse problems. IEEE Transactions on Computational Imaging, 9, 781–795.

Ho, J., Jain, A., & Abbeel, P. (2020).Denoising diffusion probabilistic models. Advances in neural information processing systems, 33, 6840–6851.

Leiner, J., Duan, B., Wasserman, L., & Ramdas, A. (2025).Data fission: Splitting a single data point. Journal of the American Statistical Association, 120(549), 135–146.

Liaudat, T. I., Mars, M., Price, M. A., Pereyra, M., Betcke, M. M., & McEwen, J. D. (2024).Scalable bayesian uncertainty quantification with data-driven priors for radio interferometric imaging. RAS Techniques and Instruments, 3(1), 505–534.

Mars, M., Liaudat, T. I., Whitney, J. J., Betcke, M. M., & McEwen, J. D. (2025).Generative imaging for radio interferometry with fast uncertainty quantification. arXiv preprint arXiv:2507.21270.

McEwen, J. D., & Liaudat, T. I. (2026).High-Dimensional Uncertainty Quantification with Deep Data-Driven AI Priors. arXiv e-prints, Article arXiv:2601.xxxxx, arXiv:2601.xxxxx.

Monroy, B., Bacca, J., & Tachella, J. (2025).Generalized recorrupted-to-recorrupted: Self-supervised learning beyond gaussian noise. Proceedings of the Computer Vision and Pattern Recognition Conference, 28155–28164.

References III

Pereyra, M. (2017).Maximum-a-posteriori estimation with bayesian confidence regions. SIAM Journal on Imaging Sciences, 10(1), 285–302.

Pesquet, J.-C., Repetti, A., Terris, M., & Wiaux, Y. (2021).Learning maximally monotone operators for image recovery. SIAM Journal on Imaging Sciences, 14(3), 1206–1237.

Ryu, E., Liu, J., Wang, S., Chen, X., Wang, Z., & Yin, W. (2019).Plug-and-play methods provably converge with properly trained denoisers. International Conference on Machine Learning, 5546–5557.

Skilling, J. (2006).Nested sampling for general Bayesian computation. Bayesian Analysis, 1(4), 833–859.

Song, Y., & Ermon, S. (2019).Generative modeling by estimating gradients of the data distribution. Advances in neural information processing systems, 32.

Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., & Poole, B. (2020).Score-based generative modeling through stochastic differential equations. arXiv preprint arXiv:2011.13456.

Sprunck, T., Pereyra, M., & Liaudat, T. (2025).Bayesian model selection and misspecification testing in imaging inverse problems only from noisy and partial measurements. arXiv preprint arXiv:2510.27663.

Tachella, J., & Pereyra, M. (2023).Equivariant Bootstrapping for Uncertainty Quantification in Imaging Inverse Problems. arXiv e-prints, Article arXiv:2310.11838, arXiv:2310.11838.

Whitney, J. J., Liaudat, T. I., Price, M. A., Mars, M., & McEwen, J. D. (2025).Generative modelling for mass-mapping with fast uncertainty quantification. Monthly Notices of the Royal Astronomical Society, 542(3), 2464–2479.

Bayesian model selection and misspecification testing

Based on (Sprunck et al., Liaudat 2025) (arXiv:2510:27663).

- The Bayesian evidence under model $\mathcal{M}$ is $p_{\mathcal{M}}(y) = \int p_{\mathcal{M}}(y \mid x) p_{\mathcal{M}}(x) \mathrm{d}x$, and provides a principled way for model selection through the Bayes factor.
- The evidence is often intractable, but clever algorithms like Nested Sampling (Skilling, 2006) allow us to estimate it.

  $\rightarrow$ Still computationally expensive and typically only used for low-dimensional problems.

The goal of the method is:

1. Carry out **Bayesian model selection in high-dimensions**.
2. **Detect if the model works on an out-of-distribution (OOD) regime** $\rightarrow$ important for AI-based methods.
3. **Only work from a single observation** y**.**

Data fission

We will rely in the data fission procedure (Leiner et al., 2025) to create statistically independent data splits from a single observation y.

Data fission procedure for Gaussian noise model (Leiner et al., 2025)

We partition a single $\mathbf{y} = y$ from $\mathbf{y} \sim P(A(x_\star))$ into two synthetic measurements $\mathbf{y}^+$ and $\mathbf{y}^-$ that are conditionally independent given $x_\star$. Suppose that $\mathbf{y} \sim \mathcal{N}(A(x_\star), \Sigma)$ and let $\mathbf{w} \sim \mathcal{N}(0, \Sigma)$. Then, for any $\alpha \in (0, 1)$,

$$\begin{aligned} \mathbf{y}^+ &= f_\alpha^-(\mathbf{y}, \boldsymbol{w}) \coloneqq \mathbf{y} + c_\alpha \boldsymbol{w}\,, \\ \mathbf{y}^- &= f_\alpha^-(\mathbf{y}, \boldsymbol{w}) \coloneqq \mathbf{y} - \boldsymbol{w}/c_\alpha\,, \end{aligned} \tag{2}$$

with $c_\alpha = \sqrt{\alpha/(1-\alpha)}$.

The α parameter controls the information split between $\mathbf{y}^+$ and $\mathbf{y}^-$.

Equivalent splitting strategies are available for other noise models from the natural exponential family Monroy et al., 2025.

Proposed metric: Bayesian cross validation through data fission

We evaluate a model $\mathcal{M}$, comprising a prior and likelihood, by computing a summary of the form

$$\Psi(\mathcal{M}) = \mathbb{E}_{\mathbf{y}^+,\mathbf{y}^-}\left[S(p_{\mathcal{M}}(\mathbf{y}^+|\mathbf{y}^-, y^+)\right] \ , = \int S(p_{\mathcal{M}}(\mathbf{y}^+|\mathbf{y} = y^-), y^+) p_{\mathcal{M}}(y^-, y^+) \mathrm{d}y^- \, \mathrm{d}y^+ \, ,$$

where $S : \mathcal{P} \times \mathbb{R}^m \mapsto \mathbb{R}_+$ is a scoring rule (Gneiting & Raftery, 2007) that takes a predictive density $p \in \mathcal{P}$, with $\mathcal{P}$ being a probability measure, and a realization mapping it to a numerical assessment of that prediction.

For the previous data fission procedure we can write our metric as

$$\Phi(\mathcal{M}) = \mathbb{E}_{\mathbf{w}}\left[\mathbb{E}_{\mathbf{x}|f_\alpha^-(y,\mathbf{w}),\mathcal{M}}\left[\phi_{\mathcal{M}}(f_\alpha^+(y,\mathbf{w}),\mathbf{x}))\right]\right] = \int \phi_{\mathcal{M}}(f_\alpha^+(y,w),x)\, p_{\mathcal{M}}(x \mid f_\alpha^-(y,w)) p(w) \mathrm{d}x \mathrm{d}w,$$

where $\phi_{\mathcal{M}} : \mathbb{R}^m \times \mathbb{R}^n \mapsto \mathbb{R}_+$ quantifies the discrepancy between a possible x and y^+, leading to a scoring rule $\mathbb{E}_{\mathbf{x}|f_\alpha^-(y,\mathbf{w}),\mathcal{M}}\left[\phi_{\mathcal{M}}(f_\alpha^+(y,\mathbf{w}),\mathbf{x}))\right]$ for the prediction of $\mathbf{y}^+$ from y^-.

Intuition and link with the Bayesian evidence

For a given choice of scoring rule, *i.e.*, the log score (log likelihood), we obtain a regularised version of the Bayesian evidence,

$$\Phi^3_y(\mathcal{M}) = \mathbb{E}_{\mathbf{x}|y^-,\mathcal{M}}\left[p_{\mathcal{M}}(y^+|\mathbf{x})\right] = p_{\mathcal{M}}(y^+|y^-) = \int p_{\mathcal{M}}(y^+|x)p_{\mathcal{M}}(x|y^-)\mathrm{d}x, \tag{3}$$

which in the limit of $\alpha \to 0$ we obtain the Bayesian Evidence as follows,

$$\lim_{\alpha\to 0} \Phi^3_y(\mathcal{M}) = \mathbb{E}_{\mathbf{x}|\mathcal{M}}\left[p_{\mathcal{M}}(y|\mathbf{x})\right] = p_{\mathcal{M}}(y) = \int p_{\mathcal{M}}(y|x)p_{\mathcal{M}}(x)\mathrm{d}x, \tag{4}$$

due to $y^+ \to y$ and y^- tending to an independent noise realisation.

Likelihood-based metric for model selection

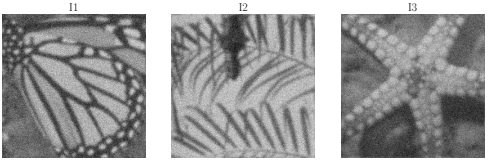

Examples of blurred measurements, generated by using the blur kernel $\kappa_{\mathcal{G}}(2)$.

	Single Shot	Few Shot
Ours (w. $\widehat{\Phi}^1$)	86.7%	100%
Bayes Res. [55]	40.0%	40.0%
EB Res. [55]	46.7%	60.0%

Table 1: Accuracy of likelihood model selection, using the the proposed summary $\widehat{\Phi}^1$ and two variants of the baseline method [55], from a single measurement (single shot) or three measurements (few shot).

- Misspecified likelihood scenario.

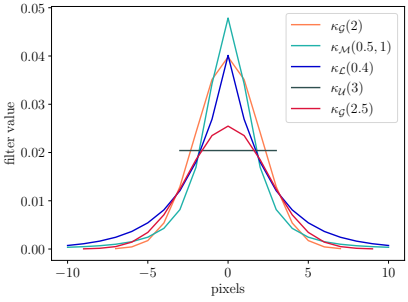

Profile of the considered blur kernels, their similarity makes model selection difficult.

Excellent performance selecting the correct model even for very similar blur kernels.

Posterior-based metric for model OOD detection

Detect if the model is working on out-of-distribution (OOD) data.

		$\sigma_\kappa = 0.5$	$\sigma_\kappa = 2$	$\sigma_\kappa = 5$
Type I Error	FFHQ	0%	6.7%	6.7%
	Celeb	6.7%	6.7%	6.7%
Power	Moderate OOD	86.7%	73.3%	60%
	Strong OOD	100%	100%	100%

Table 2: Type I error rate (incorrect rejection of ID samples from FFHQ, Celeb) and Power (correct rejection of moderate OOD (Met) and strong OOD (bedrooms, CBSD68, AFHQ) examples).

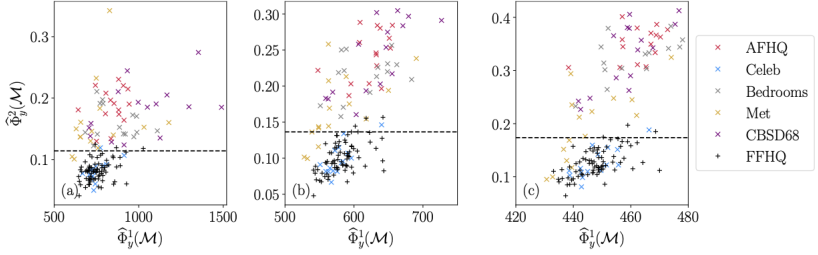

Posterior-based metric for model OOD detection

→ Impressive detection performance for only relying the observations y.