

Towards fast and scalable uncertainty quantification for scientific imaging

Tobías I. Liaudat
IRFU, CEA Paris-Saclay

<https://tobias-liaudat.github.io>


Deep CosmoStat Days


16th January 2025

Motivation: SKA's radio interferometer


SKA-mid – the SKA's mid-frequency instrument

The SKA Observatory (SKAO) is a next-generation radio astronomy facility that will revolutionize our understanding of the Universe. It will have a uniquely distributed character: one observatory operating five telescopes on three continents. The two telescopes, named SKA-low and SKA-mid, will be observing the Universe at different frequencies. They are also called interferometers as they each comprise a large number of individual elements working together to form a single large telescope.






Location: South Africa




Frequency range:
350 MHz to 15.4 GHz
with a goal of 24 GHz




197 dishes
(including 64 MeerKAT dishes)

Total collecting area:
33,000m²

or **126 tennis courts**



Maximum distance between dishes:
150km



Data transfer rate:
8.8 Terabits per second

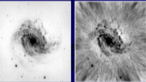



Image quality of SKA-mid (left) versus the best current facility operating in the same frequency range, the James Very Large Array (JVLA) in the United States (right). SKA-mid's resolution will be 4x better than JVLA.

Compared to the JVLA, the current best similar instrument in the world:


4x the resolution

5x more sensitive

60x the survey speed





www.skatelescope.org @SKAO SKA Observatory SKA Observatory SKA Observatory @skaoobservatory




SKA1 LOW - the SKA's low-frequency instrument

The Square Kilometre Array (SKA) will be the world's largest radio telescope, revolutionizing our understanding of the Universe. The SKA will be built in two phases: SKA1 and SKA2 - starting in 2018, with SKA1 representing a fraction of the full SKA. SKA1 will include two instruments: SKA1 MID and SKA1 LOW - observing the Universe at different frequencies.






Location: Australia




Frequency range:
50 MHz to 350 MHz




~130,000 antennas spread between **500 stations**

Total collecting area:
0.4km²




Maximum distance between stations:
65km



Total raw data output:
157 terabytes per second
4.9 zettabytes per year

Enough to fill up **35,000 DVDs** every second

5x the estimated global internet traffic in 2015 (Source: Cisco)

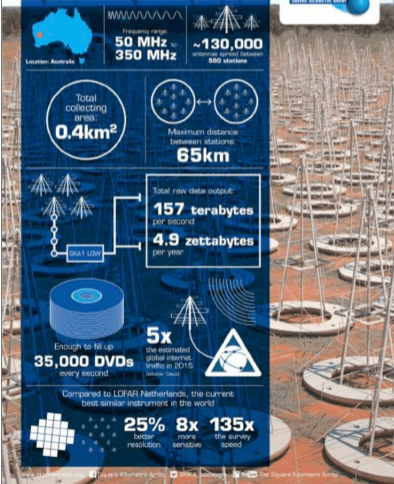



Compared to LOFAR Netherlands, the current best similar instrument in the world

25% better resolution

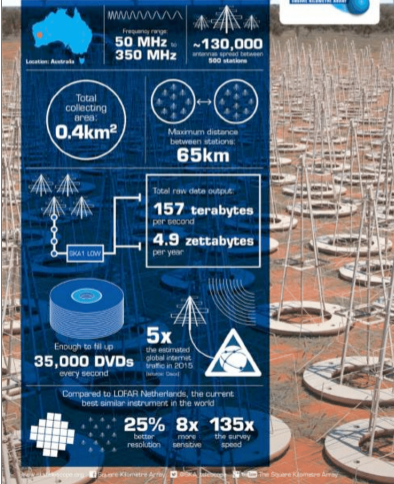
8x more sensitive

135x the survey speed





www.skatelescope.org @SKAO SKA Observatory SKA Observatory SKA Observatory @skaoobservatory



Radio interferometric imaging

Linear observational model

$$\mathbf{y} = \Phi \mathbf{x} + \mathbf{n}$$

$\mathbf{y} \in \mathbb{C}^M$: Observed Fourier coefficients

$\mathbf{n} \in \mathbb{C}^M$: Observational noise (assumed White and Gaussian)



$\mathbf{x} \in \mathbb{R}^N$: Sky intensity image

$\Phi \in \mathbb{C}^{M \times N}$: Linear measurement operator

- In its simplest case: FFT and Fourier mask

Due to \mathbf{n} and Φ the inverse problem is ill-posed


Goal: Estimate $\hat{\mathbf{x}}$ from \mathbf{y}



$\hat{\mathbf{x}}$

Uncertainty quantification: more than a point estimate

Image reconstruction: $\hat{\mathbf{x}}$



Is this blob *physical*?

- Is it a reconstruction artefact?
- Is it backed by the data?
- Can we base a scientific decision on this image?

Several reasons motivate us to develop uncertainty quantification (UQ) techniques for the reconstruction methods,

- Usual UQ techniques from the Bayesian framework rely on interrogating the posterior exploiting Bayes' theorem:

$$\underbrace{p(x | y, M)}_{\text{posterior}} \propto \underbrace{p(y | x, M)}_{\text{likelihood}} \underbrace{p(x | M)}_{\text{prior}}$$

For example, Cai et al. (2018a) applies this for radio imaging using a ℓ_1 regularised wavelet-based prior.

Sample from the posterior which is non-smooth to obtain


$$\{\mathbf{x}^{(j)}\}_{j=1}^K, \mathbf{x}^{(j)} \sim p(\mathbf{x}|\mathbf{y})$$

- **Proximal MCMC algorithm** (Pereyra, 2016) following Langevin dynamics

Is the problem solved?

Uncertainty quantification: more than a point estimate

Image reconstruction: $\hat{\mathbf{x}}$



Is this blob *physical*?

- Is it a reconstruction artefact?
- Is it backed by the data?
- Can we base a scientific decision on this image?

Several reasons motivate us to develop uncertainty quantification (UQ) techniques for the reconstruction methods,

- Usual UQ techniques from the Bayesian framework rely on interrogating the posterior exploiting Bayes' theorem:

$$\underbrace{p(x | y, M)}_{\text{posterior}} \propto \underbrace{p(y | x, M)}_{\text{likelihood}} \underbrace{p(x | M)}_{\text{prior}}$$

For example, Cai et al. (2018a) applies this for radio imaging using a ℓ_1 regularised wavelet-based prior.


Sample from the posterior which is non-smooth to obtain $\{\mathbf{x}^{(j)}\}_{j=1}^K$, $\mathbf{x}^{(j)} \sim p(\mathbf{x}|\mathbf{y})$

- **Proximal MCMC algorithm** (Pereyra, 2016) following Langevin dynamics

Is the problem solved?

Uncertainty quantification: more than a point estimate

Image reconstruction: \hat{x}



Is this blob *physical*?

- Is it a reconstruction artefact?
- Is it backed by the data?
- Can we base a scientific decision on this image?

Several reasons motivate us to develop uncertainty quantification (UQ) techniques for the reconstruction methods,

- Usual UQ techniques from the Bayesian framework rely on interrogating the posterior exploiting Bayes' theorem:

$$\underbrace{p(x | y, M)}_{\text{posterior}} \propto \underbrace{p(y | x, M)}_{\text{likelihood}} \underbrace{p(x | M)}_{\text{prior}}$$

For example, Cai et al. (2018a) applies this for radio imaging using a ℓ_1 regularised wavelet-based prior.


Sample from the posterior which is non-smooth to obtain $\{\mathbf{x}^{(j)}\}_{j=1}^K$, $\mathbf{x}^{(j)} \sim p(\mathbf{x}|\mathbf{y})$

- **Proximal MCMC algorithm** (Pereyra, 2016) following Langevin dynamics

Is the problem solved?

Uncertainty quantification: more than a point estimate

Image reconstruction: $\hat{\mathbf{x}}$



Is this blob *physical*?

- Is it a reconstruction artefact?
- Is it backed by the data?
- Can we base a scientific decision on this image?

Several reasons motivate us to develop uncertainty quantification (UQ) techniques for the reconstruction methods,

- Usual UQ techniques from the Bayesian framework rely on interrogating the posterior exploiting Bayes' theorem:

$$\underbrace{p(x | y, M)}_{\text{posterior}} \propto \underbrace{p(y | x, M)}_{\text{likelihood}} \underbrace{p(x | M)}_{\text{prior}}$$

For example, Cai et al. (2018a) applies this for radio imaging using a ℓ_1 regularised wavelet-based prior.

Sample from the posterior which is non-smooth to obtain $\{\mathbf{x}^{(j)}\}_{j=1}^K$, $\mathbf{x}^{(j)} \sim p(\mathbf{x}|\mathbf{y})$

- **Proximal MCMC algorithm** (Pereyra, 2016) following Langevin dynamics

Is the problem solved?

The problem is not solved

Based on: **Scalable Bayesian uncertainty quantification with data-driven priors for radio interferometric imaging** (Liaudat, *et al.*, 2024 (arXiv:2312.00125))

Difficulties in the high-dimensional setting:

1. Even if we know the likelihood, applying Φ is **computationally expensive**
2. Handcrafted priors like wavelets are **not expressive enough**
3. Sampling-based techniques are **prohibitively expensive** in this setting

How can we obtain information from the high-dimensional posterior $p(\mathbf{x}|\mathbf{y})$ without sampling from it?

If we restrict to **log-concave posteriors** something beautiful happens!

→ **A concentration phenomenon** (Pereyra, 2017)

log-concave posterior $p(\mathbf{x}|\mathbf{y}) = \exp[-f(\mathbf{x}) - g(\mathbf{x})]/Z \rightarrow$ convex potential $f(\mathbf{x}) + g(\mathbf{x})$

The problem is not solved

Based on: **Scalable Bayesian uncertainty quantification with data-driven priors for radio interferometric imaging** (Liaudat, *et al.*, 2024 (arXiv:2312.00125))

Difficulties in the high-dimensional setting:

1. Even if we know the likelihood, applying Φ is **computationally expensive**
2. Handcrafted priors like wavelets are **not expressive enough**
3. Sampling-based techniques are **prohibitively expensive** in this setting

How can we obtain information from the high-dimensional posterior $p(\mathbf{x}|\mathbf{y})$ without sampling from it?

If we restrict to **log-concave posteriors** something beautiful happens!

→ **A concentration phenomenon** (Pereyra, 2017)

log-concave posterior $p(\mathbf{x}|\mathbf{y}) = \exp[-f(\mathbf{x}) - g(\mathbf{x})]/Z \rightarrow$ convex potential $f(\mathbf{x}) + g(\mathbf{x})$

The problem is not solved

Based on: **Scalable Bayesian uncertainty quantification with data-driven priors for radio interferometric imaging** (Liaudat, *et al.*, 2024 (arXiv:2312.00125))

Difficulties in the high-dimensional setting:

1. Even if we know the likelihood, applying Φ is **computationally expensive**
2. Handcrafted priors like wavelets are **not expressive enough**
3. Sampling-based techniques are **prohibitively expensive** in this setting

How can we obtain information from the high-dimensional posterior $p(\mathbf{x}|\mathbf{y})$ without sampling from it?

If we restrict to **log-concave posteriors** something beautiful happens!

→ **A concentration phenomenon** (Pereyra, 2017)

log-concave posterior $p(\mathbf{x}|\mathbf{y}) = \exp[-f(\mathbf{x}) - g(\mathbf{x})]/Z \rightarrow$ convex potential $f(\mathbf{x}) + g(\mathbf{x})$

Highest posterior density region

Posterior credible region:

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

We consider the **highest posterior density (HPD) region**

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

Theorem 3.1 (Pereyra, 2017)

Suppose the posterior $p(\mathbf{x} | \mathbf{y}) = \exp[-f(\mathbf{x}) - g(\mathbf{x})]/Z$ 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

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

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

We only need to evaluate $f + g$ on the MAP estimation $\hat{\mathbf{x}}_{\text{MAP}}$!

Highest posterior density region

Posterior credible region:

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

We consider the **highest posterior density (HPD) region**

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

Theorem 3.1 (Pereyra, 2017)

Suppose the posterior $p(\mathbf{x} | \mathbf{y}) = \exp[-f(\mathbf{x}) - g(\mathbf{x})]/Z$ 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

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

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

We only need to evaluate $f + g$ on the MAP estimation $\hat{\mathbf{x}}_{\text{MAP}}$!

Highest posterior density region

Posterior credible region:

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

We consider the **highest posterior density (HPD) region**

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

Theorem 3.1 (Pereyra, 2017)

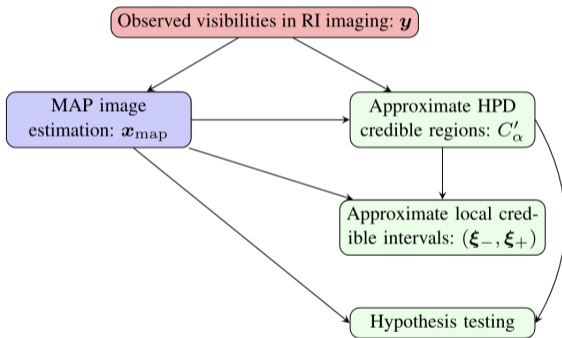
Suppose the posterior $p(\mathbf{x} | \mathbf{y}) = \exp[-f(\mathbf{x}) - g(\mathbf{x})]/Z$ is **log-concave** on \mathbb{R}^N . Then, for any $\alpha \in (4 \exp[(-N/3)], 1)$, the HPD region C_α^* is contained by

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

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

We only need to evaluate $f + g$ on the MAP estimation $\hat{\mathbf{x}}_{\text{MAP}}$!

MAP-based uncertainty quantification



Cai et al. (2018b)

UQ techniques:

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

Scalable Bayesian uncertainty quantification

1. **Scalability** → Need to rely on **optimisation sampling**, use the **MAP estimator**
2. **Uncertainty quantification** → Need the potential to be **convex** and **explicit**
3. **Good reconstruction** → Need to use **data-driven** (learned) approaches

The approach requires our prior to be convex and with an explicit potential

→ We **constrain our prior to be convex**, but we **gain an effortless UQ!**

We use the neural-network-based convex regulariser from Goujon et al. (2023):

- Shallow network using learned spline-based activation functions trained as a (multi-)gradient step denoiser.
- Properties: **Explicit cost**, **Convex**, **Smooth regulariser with known Lipschitz constant**.

Scalable Bayesian uncertainty quantification

1. **Scalability** → Need to rely on **optimisation sampling**, use the **MAP estimator**
2. **Uncertainty quantification** → Need the potential to be **convex** and **explicit**
3. **Good reconstruction** → Need to use **data-driven** (learned) approaches

The approach requires our prior to be convex and with an explicit potential

→ We **constrain our prior to be convex**, but we **gain an effortless UQ!**

We use the neural-network-based convex regulariser from Goujon et al. (2023):

- Shallow network using learned spline-based activation functions trained as a (multi-)gradient step denoiser.
- Properties: **Explicit cost**, **Convex**, **Smooth regulariser with known Lipschitz constant**.

Scalable Bayesian uncertainty quantification

1. **Scalability** → Need to rely on **optimisation sampling**, use the **MAP estimator**
2. **Uncertainty quantification** → Need the potential to be **convex** and **explicit**
3. **Good reconstruction** → Need to use **data-driven** (learned) approaches

The approach requires our prior to be convex and with an explicit potential

→ We **constrain our prior to be convex**, but we **gain an effortless UQ!**

We use the neural-network-based convex regulariser from Goujon et al. (2023):

- Shallow network using learned spline-based activation functions trained as a (multi-)gradient step denoiser.
- Properties: **Explicit cost**, **Convex**, **Smooth regulariser with known Lipschitz constant**.

Computing time and likelihood evaluations

Computation wall-clock times for the W28 image in seconds.

Models	MAP optim.	Posterior sampling	LCIs 8×8	Fast pixel UQ
Wavelet-based	0.94	36.0×10^3	149.7	—
QuantifAI	0.64	6.44×10^3	108.2	0.17

The number of measurement operator evaluations used by QuantifAI for the W28 image.

MCMC sampling	LCIs 8×8	LCIs 16×16	Fast pixel UQ
11×10^6	81.5×10^3	21.2×10^3	28

The fast pixel UQ is 10^6 and 10^3 times faster than the MCMC sampling and LCIs, respectively.

We rely on an iterative algorithm for the optimisation, which can be **computationally expensive**:

- We need to decrease the number of iteration to further accelerate

Computing time and likelihood evaluations

Computation wall-clock times for the W28 image in seconds.

Models	MAP optim.	Posterior sampling	LCIs 8×8	Fast pixel UQ
Wavelet-based	0.94	36.0×10^3	149.7	—
QuantifAI	0.64	6.44×10^3	108.2	0.17

The number of measurement operator evaluations used by QuantifAI for the W28 image.

MCMC sampling	LCIs 8×8	LCIs 16×16	Fast pixel UQ
11×10^6	81.5×10^3	21.2×10^3	28

The fast pixel UQ is 10^6 and 10^3 times faster than the MCMC sampling and LCIs, respectively.

We rely on an iterative algorithm for the optimisation, which can be **computationally expensive**:

- We need to decrease the number of iteration to further accelerate

Faster reconstruction: algorithm unrolling

Based on: **EVIL-Deconv: Efficient Variability-Informed Learned Deconvolution using Algorithm Unrolling** (Kern, Kervazo & Bobin, 2024 (submitted))

Main motivation:

1. **Reduce the number of iterations!**
2. Improve reconstruction performance

The main algorithm step which is unrolled for L steps

$$x_{l+1} = g_l(x_l + \Phi_l(M)(y - M * x_l))$$

- $\Phi_l(M)$: Learned preconditioning step based on CNNs with M being the PSF
- g_l : Learned proximal operator (denoiser) based on DRUNets

Everything trained on a supervised manner end-to-end for th L unrolled steps.

Faster reconstruction: algorithm unrolling

EVIL-Deconv results:

- Greatly reduced computation budget
- Great reconstruction quality (for in-distribution data)

EVIL-Deconv drawbacks:

- Lost interpretation of the reconstruction (which is helpful for UQ)
 - ▶ Is it the fixed point of an equation?
 - ▶ Is the reconstruction related to a posterior probability distribution?
- UQ is missing

These drawbacks limit its scientific application

Faster reconstruction: algorithm unrolling

EVIL-Deconv results:

- Greatly reduced computation budget
- Great reconstruction quality (for in-distribution data)

EVIL-Deconv drawbacks:

- Lost interpretation of the reconstruction (which is helpful for UQ)
 - ▶ Is it the fixed point of an equation?
 - ▶ Is the reconstruction related to a posterior probability distribution?
- UQ is missing

These drawbacks limit its scientific application

CARB: Conformalized Augmented Radio Bootstrap

Based on: **Uncertainty quantification for fast reconstruction methods using augmented equivariant bootstrap: Application to radio interferometry**

(Cherif, [Liaudat](#), Kern, Kervazo & Bobin, 2024 (arXiv:2410:23178))

Based on the equivariant Bootstrap framework 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 $T_g x \in \mathcal{X}$ and a reconstruction method $\hat{x}(y) = f(y)$ (e.g. EVIL-Deconv):

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 bootstrap 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 $e_i = \|\hat{x}(y) - \tilde{x}_i\|^2$

CARB: Conformalized Augmented Radio Bootstrap

Based on: **Uncertainty quantification for fast reconstruction methods using augmented equivariant bootstrap: Application to radio interferometry**

(Cherif, [Liaudat](#), Kern, Kervazo & Bobin, 2024 (arXiv:2410:23178))

Based on the equivariant Bootstrap framework 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 $T_g x \in \mathcal{X}$ and a reconstruction method $\hat{x}(y) = f(y)$ (e.g. EVIL-Deconv):

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 bootstrap 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 $e_i = \|\hat{x}(y) - \tilde{x}_i\|^2$

CARB: Conformalized Augmented Radio Bootstrap

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

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} and A , and A is not \mathcal{G} -equivariant, **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^* (ground truth).**

Motivation:

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

CARB: Conformalized Augmented Radio Bootstrap

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

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} and A , and A is not \mathcal{G} -equivariant, **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^* (ground truth).**

Motivation:

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

CARB: Conformalized Augmented Radio Bootstrap

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

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} and A , and A is not \mathcal{G} -equivariant, **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^* (ground truth).**

Motivation:

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

CARB method consists of:

1. Fast reconstruction algorithm (e.g. EVIL-Deconv),
2. Equivariant bootstrap framework,
3. Adapted group actions for the RI imaging problem,
4. Conformalisation procedure: Risk-Controlling Prediction Sets (RCPS) to have statistical guarantees on the coverage (Angelopoulos & Bates, 2023).

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

CARB method consists of:

1. Fast reconstruction algorithm (e.g. EVIL-Deconv),
2. Equivariant bootstrap framework,
3. Adapted group actions for the RI imaging problem,
4. Conformalisation procedure: Risk-Controlling Prediction Sets (RCPS) to have statistical guarantees on the coverage (Angelopoulos & Bates, 2023).


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

CARB: Conformalized Augmented Radio Bootstrap

The **group actions** we consider are:

1. Circular shift translations not exceeding 2 pixels,
2. Image flips over the horizontal and vertical axis,
3. Rotations of 90-degrees multiples,
4. Invertible 2D radially-symmetric filters in the specific form of low-shelving and high-shelving filters with varying cutoff frequencies.

Examples of (4) filter transformations:



The **final group action**: a **random composition** of the aforementioned transformations, where each transformation is applied with a given probability.


- It allows us to **significantly expand the number of possible group action**, helping to estimate uncertainties better.

CARB: Conformalized Augmented Radio Bootstrap

The **group actions** we consider are:

1. Circular shift translations not exceeding 2 pixels,
2. Image flips over the horizontal and vertical axis,
3. Rotations of 90-degrees multiples,
4. Invertible 2D radially-symmetric filters in the specific form of low-shelving and high-shelving filters with varying cutoff frequencies.

Examples of (4) filter transformations:



The final group action: a **random composition** of the aforementioned transformations, where each transformation is applied with a given probability.

- It allows us to **significantly expand the number of possible group action**, helping to estimate uncertainties better.

Numerical experiments: Quantitative UQ comparison

Length ratio: We compare the average ℓ_2 norm ratio between the confidence interval lengths and the ground truth image to study the **tightness of the error estimations**.

Empirical coverage: We empirically compute $\mathbb{E}_y\{P(\hat{I}_\alpha^-(y) < x^* - \hat{x}(y) < \hat{I}_\alpha^+(y))|y\}$ (error rate) over all pixels in the image dataset, where α (risk level) is set to 0.1, and $\hat{I}_\alpha^-(y), \hat{I}_\alpha^+(y)$ are the estimated lower and higher interval limits.

Method	Length	Coverage
Quantile Regression (QR)	0.15	14%
Conformalized QR	204.08	92%
Parametric Bootstrap	0.07	0%
Equivariant Bootstrap	0.13	7%
Augmented Radio Bootstrap	0.29	87%
CARB	0.34	91%

Tight intervals and very good coverage! Results showcase:

- the importance of selecting adapted group actions,
- the conformalisation is useful once the intervals are already good.

Numerical experiments: Quantitative UQ comparison

Length ratio: We compare the average ℓ_2 norm ratio between the confidence interval lengths and the ground truth image to study the **tightness of the error estimations**.


Empirical coverage: We empirically compute $\mathbb{E}_y\{P(\hat{I}_\alpha^-(y) < x^* - \hat{x}(y) < \hat{I}_\alpha^+(y))|y\}$ (error rate) over all pixels in the image dataset, where α (risk level) is set to 0.1, and $\hat{I}_\alpha^-(y), \hat{I}_\alpha^+(y)$ are the estimated lower and higher interval limits.

Method	Length	Coverage
Quantile Regression (QR)	0.15	14%
Conformalized QR	204.08	92%
Parametric Bootstrap	0.07	0%
Equivariant Bootstrap	0.13	7%
Augmented Radio Bootstrap	0.29	87%
CARB	0.34	91%

Tight intervals and very good coverage! Results showcase:


- the importance of selecting adapted group actions,
- the conformalisation is useful once the intervals are already good.

Numerical experiments: Qualitative UQ comparison




RI image reconstruction result for the unrolling algorithm.


Numerical experiments: Qualitative UQ comparison




(a) Oracle abs. residual




(b) Quantile Regression




(c) Conformalized QR




(d) Parametric bootstrap



(e) Equivariant bootstrap



(f) ARB



(g) CARB


Numerical experiments: Coverage plots

Coverage plots:

We estimate a confidence region for x^* , derived from the pivotal statistic $\|x^* - \hat{x}(Y)\|_2^2$ related to the estimation MSE.

We then compute the empirical coverage probabilities on the test set, as measured by the proportion of test images that lie within the confidence regions for a range of confidence levels.

We still need to validate the method on higher dimensions.



Coverage plots for equivariant bootstrap methods with different group actions.


Numerical experiments: Coverage plots

Coverage plots:

We estimate a confidence region for x^* , derived from the pivotal statistic $\|x^* - \hat{x}(Y)\|_2^2$ related to the estimation MSE.

We then compute the empirical coverage probabilities on the test set, as measured by the proportion of test images that lie within the confidence regions for a range of confidence levels.

We still need to validate the method on higher dimensions.



Coverage plots for equivariant bootstrap methods with different group actions.

Bonus: Can we go even faster?

Based on: **Generative imaging for radio interferometry with fast UQ**
(Mars, [Liaudat](#), Whitney, Betcke & McEwen, 2024 (in prep.))

Based on the regularised conditional GAN (rcGAN) proposed in Bendel et al., 2023 that is able to generate **approximate posterior samples**

Main points of the proposed approach:

- Builds from the conditional Wasserstein GAN (Adler and Öktem, 2018)
- Regularisation to **avoid mode collapse** and **reward sample diversity**.
- Under simplifying assumptions, the first two moments of the approximated posterior (mean and covariance) **match the true posterior**.
- We condition on the **dirty image** and the **PSF**.
- **Extremely-fast** reconstruction and sampling.

Bonus: Can we go even faster?

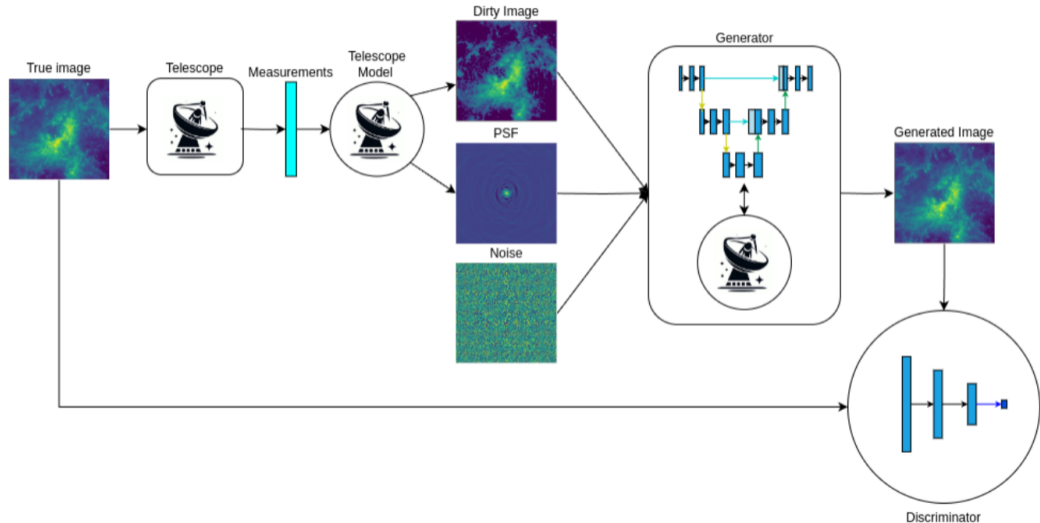
Based on: **Generative imaging for radio interferometry with fast UQ**
(Mars, [Liaudat](#), Whitney, Betcke & McEwen, 2024 (in prep.))

Based on the regularised conditional GAN (rcGAN) proposed in Bendel et al., 2023 that is able to generate **approximate posterior samples**

Main points of the proposed approach:


- Builds from the conditional Wasserstein GAN (Adler and Öktem, 2018)
- Regularisation to **avoid mode collapse** and **reward sample diversity**.
- Under simplifying assumptions, the first two moments of the approximated posterior (mean and covariance) **match the true posterior**.
- We condition on the **dirty image** and the **PSF**.
- **Extremely-fast** reconstruction and sampling.

Regularised conditional GAN for RI imaging and fast UQ



Regularised conditional GAN for RI imaging and fast UQ

Reconstruction of simulated MeerKAT observation of galaxies from Illustris TNG simulations.



High reconstruction PSNR and good correlation between the oracle error and the Std Dev.

A deeper validation of the produced samples is yet to be done.

Bonus n2: Regularised conditional GAN for mass-mapping in cosmology

Based on: **Generative modelling for mass-mapping with fast uncertainty quantification** (Whitney, Liaudat, Price, Mars & McEwen, 2024 (arXiv:2410.24197))

We adapted and applied the regularised conditional GAN to the dark matter mass-mapping problem in cosmology.

- We use the posterior mean as reconstruction and posterior samples for UQ.
- Training of the MMGAN for real data was done with mock COSMOS simulations.
- Validation using the κ TNG simulations. Results for simulations below:

	Pearson \uparrow	RMSE \downarrow	PSNR \uparrow
MMGAN (Ours)	0.727	0.0197	31.674
Kaiser-Squires	0.622	0.0229	30.387






UQ validation: Computed coverage probabilities using 100 maps at 90% and 95% confidence intervals obtaining **85% and 89% empirical coverage probability**.

Bonus n2: Regularised conditional GAN for mass-mapping in cosmology

We applied MMGAN to **real COSMOS field data** from HST.

We compare on real data with:

- DLPosterior (Remy, B. et al., 2023)
- Kaiser-Squires



We have explored different reconstruction methods **with UQ** for radio interferometric imaging (and other problems) exploiting different ML/AI tools:

1. In a Bayesian framework, favour optimisation and avoid sampling by approximating the HPD region while using learned data-driven priors.
2. Accelerate reconstruction with algorithm unrolling but loose interpretability.
3. The CARB method picks up the unrolled method and provides UQ in an unsupervised framework based on equivariant bootstrap.
4. The regularised conditional GAN trained on a supervised manner allows us to do instant (approximate) posterior sampling.

Questions?

We have explored different reconstruction methods **with UQ** for radio interferometric imaging (and other problems) exploiting different ML/AI tools:

1. In a Bayesian framework, favour optimisation and avoid sampling by approximating the HPD region while using learned data-driven priors.
2. Accelerate reconstruction with algorithm unrolling but loose interpretability.
3. The CARB method picks up the unrolled method and provides UQ in an unsupervised framework based on equivariant bootstrap.
4. The regularised conditional GAN trained on a supervised manner allows us to do instant (approximate) posterior sampling.

Questions?

References I

- Adler, J., & Öktem, O. (2018). Deep Bayesian Inversion. [arXiv e-prints](#), Article arXiv:1811.05910, arXiv:1811.05910.
- Angelopoulos, A. N., & Bates, S. (2023). Conformal prediction: A gentle introduction. [Foundations and Trends® in Machine Learning](#), *16*(4), 494–591.
- 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. (2018a). Uncertainty quantification for radio interferometric imaging – I. Proximal MCMC methods. [Monthly Notices of the Royal Astronomical Society](#), *480*(3), 4154–4169.
- Cai, X., Pereyra, M., & McEwen, J. D. (2018b). 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.
- 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.
- Kern, J., Bobin, J., & Kervazo, C. (2024). Evil-deconv: Efficient variability-informed learned deconvolution using algorithm unrolling. [Submitted](#).

- Liaudat, T. I., Mars, M., Price, M. A., Pereyra, M., Betcke, M. M., & McEwen, J. D. (2023). Scalable Bayesian uncertainty quantification with data-driven priors for radio interferometric imaging. [arXiv e-prints, Article arXiv:2312.00125](#), arXiv:2312.00125.
- Pereyra, M. (2016). Proximal markov chain monte carlo algorithms. *Statistics and Computing*, 26(4), 745–760.
- Pereyra, M. (2017). Maximum-a-posteriori estimation with bayesian confidence regions. *SIAM Journal on Imaging Sciences*, 10(1), 285–302.
- Remy, B., Lanusse, F., Jeffrey, N., Liu, J., Starck, J.-L., Osato, K., & Schrabback, T. (2023). Probabilistic mass-mapping with neural score estimation. *A&A*, 672, A51.
- 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. (2024). Generative modelling for mass-mapping with fast uncertainty quantification. [arXiv e-prints, Article arXiv:2410.24197](#), arXiv:2410.24197.

Questions?