

Data-Driven Generative Priors for Radio Interferometric Imaging

Learning the sky prior

Massimo De Santis · in collaboration with Chris Finlay

THE PROBLEM


An interferometer never sees the whole Fourier plane

- Each antenna pair measures one point in the uv-plane
- Earth rotation fills in arcs, but never everything
- MeerKAT, 2 h track, 64 x 64 grid at 2" pixels:

24% of uv cells retained · 76% lost


Infinitely many sky images fit the same visibilities.

Reconstruction is impossible without a prior.



Missing samples become structured artefacts

Dirty beam (PSF): real MeerKAT uv coverage vs. random-dropout mask



The dirty beam is the FFT of the sampling mask.

10.6%

peak sidelobe, structured concentric rings

Every bright source scatters power across the whole image, on top of the faint ones.

Every reconstruction is a prior in disguise

Hand-crafted

- Sparsity, smoothness, positivity
- Robust, well understood, fast
- Encode generic image properties
- Nothing about what skies look like

Learned

- Train a generative model on realistic skies
- The model becomes $p(x_{\text{sky}})$
- Restricts solutions to plausible skies
- Risk: bias, hallucination, mismatch

Both terms live in one posterior

$$p(\mathbf{x}_{\text{sky}} | y) \propto p(y | \mathbf{x}_{\text{sky}}) \times p(\mathbf{x}_{\text{sky}})$$

posterior


likelihood : the physics

prior : the learned model

The measurement equation supplies the data-fidelity term. The generative model supplies the prior. The objective is a single negative log-posterior.

Division of labour: generative prior (this work) · inference engine, TABASCAL (C. Finlay)

We search in latent space, not pixel space



gradient flows back through A, the decoder, and the asinh transform, only z is updated

- 4096 pixels → ~32 latent numbers: the prior is built into the parameterisation
- MAP estimate today; the same objective supports Langevin / HMC posterior sampling

Staged

1

Toy skies, known truth

Elliptical Gaussians with a ground-truth catalogue

2

Realistic sampling

Real MeerKAT-derived uv coverage

3

Science-level evaluation

Source detection, not pixel error

4

Then scale the prior


Score-based / diffusion models next

Why a VAE first?

- Fast, stable, easy to inspect
- Gives a low-dimensional inference problem
- Enough to answer: does a learned prior help at all?
- A control for everything downstream

Synthetic skies with a ground-truth catalogue


elliptical-gaussian-v5: example images [Jy/pixel] ($N_{\text{SOURCES_RANGE}}=(3, 8)$, $S \sim S^{-1.5}$ over $(0.05, 5.0)$ Jy)



- 3–8 elliptical Gaussians per field
- Integrated flux S in Jy, power law $dN/dS \propto S^{-1.5}$
- $S \in [0.05, 5]$ Jy, two decades of dynamic range
- FWHM axes + position angle, east of north
- Catalogue stored per source: the truth we score against

*Conventions match PyBDSF and OSKAR,
validated by round-trip to 0.08° in PA*

Corruption from the real array, not a random mask



- 64 antennas, 2016 baselines
- 2 h track, 60 s integrations
- ~244k uv points, gridded
- Hard radial cutoff, dense core

Random dropout removes low frequencies
the real array never loses, it is the harsher,
less realistic corruption.

Pixel metrics point the wrong way

For blur-dominated reconstructions, RMSE and SSIM are negatively correlated with photometric accuracy.

-0.35

RMSE vs amplitude accuracy (VAE direct)

-0.31


SSIM vs amplitude accuracy (VAE direct)

0.93

RMSE vs SSIM — redundant with each other

**The best-looking reconstruction can be the photometrically worst one.
If we had evaluated with pixel error, we would have drawn the opposite conclusion.**

We score reconstructions the way an astronomer would



reconstruct → run PyBDSF → match to catalogue → score

- Amplitude error (%)
- Position error / source size
- Half-maximum ellipse IoU
- Completeness
- False detections

*Same detection threshold for every method,
so differences reflect reconstruction only*

The learned prior wins decisively on photometry

method	amp error	IoU	completeness	false det.
clean (ceiling)	0.4 %	0.95	0.98	2
dirty image	56.7 %	0.64	0.78	659
Gaussian smoothing	63.9 %	0.64	0.72	445
VAE encode–decode	60.4 %	0.28	0.40	417
latent optimisation	3.5 %	0.76	0.77	69

median values; latent optimisation initialised at $z = 0$

3.5 % amplitude error against 57–64 % for every baseline, and 6× fewer false detections than the dirty image.

Where did the sources go?



median-ranked field, (not a hand-picked example)

Two sources went in. The dirty image and the smoothed version have lost them both in the sidelobes.
The latent reconstruction puts them back in the right place, at more or less the right size and the right brightness.

How much of this is luck?

We changed one thing at a time, with everything else held fixed.

Latent initialisation

- Encoder init extrapolates on dirty data: $\|z\| \approx 9.9$ vs 5.7 expected
- Overshoot scales with source brightness
- $z = 0$ reaches in 500 steps what the encoder start needs 10,000 for

completeness 0.65 → 0.77 |amp err| 13.4 % → 3.5 %

LIMITATIONS

What this does not yet show

The prior

A VAE blurs, has limited capacity, and gives no uncertainty. It is a control for now, not a proposal.

The measurement

Gridded visibilities, binary weighting, no thermal noise yet, no primary beam or calibration.

The sky

Elliptical Gaussians. Real fields are point-source dominated with diffuse structure.

The inference

A single MAP estimate. No credible intervals, no test for multimodality.

Where this goes next

1 Thermal noise, properly weighted

Realistic thermal noise, weighted by how well each uv cell was measured

2 Visibility-space likelihood

Noise is independent there; the regularisation weight stops being a free parameter

3 Score-based / diffusion prior

Expressiveness, and posterior sampling instead of a point estimate

4 TABASCAL integration

Joint calibration, RFI and sky inference with a learned sky prior

The longer game

Different priors for different astrophysical components ,
compact, diffuse, etc.

Not just a reconstructed sky, but a decomposition of it.