VROOM-SBI

Amortized Bayesian QU-Fitting for Survey-Scale Faraday Analysis

Arpan Pal[1,2,3] **Preshanth Jagannathan**[1]

[1]National Radio Astronomy Observatory, Socorro, NM [2]National Centre for Radio Astrophysics, TIFR, Pune [3]Fulbright-Nehru Fellowship, USIEF

QU-fitting without the sampler

QU-fitting gives calibrated posteriors and model comparison, but costs 10^4 to 10^6 forward-model calls per sightline. VROOM-SBI learns the posterior once from simulations, then answers each new sightline in one forward pass.

Forward model

$$P(\lambda^2) = Q + iU = \sum_{j=1}^{N} P_j(\lambda^2), \quad \varphi_1 > \varphi_2 > \ldots$$

Faraday-thin

$p_0\, e^{2i(\chi_0+\varphi\lambda^2)}$

Burn slab

$p_0\, \mathrm{sinc}(\Delta\varphi\lambda^2)\, e^{2i(\chi_0+\varphi_c\lambda^2)}$

External dispersion

$p_0\, e^{-2\sigma_\varphi^2\lambda^4}\, e^{2i(\chi_0+\varphi\lambda^2)}$

Internal dispersion

$p_0\, \frac{1-e^{-S}}{S}, S = 2\sigma_\varphi^2\lambda^4 - 2i\varphi\lambda^2$

Neural posterior estimation

1 Draw $\theta \sim p(\theta)$
scrambled Sobol

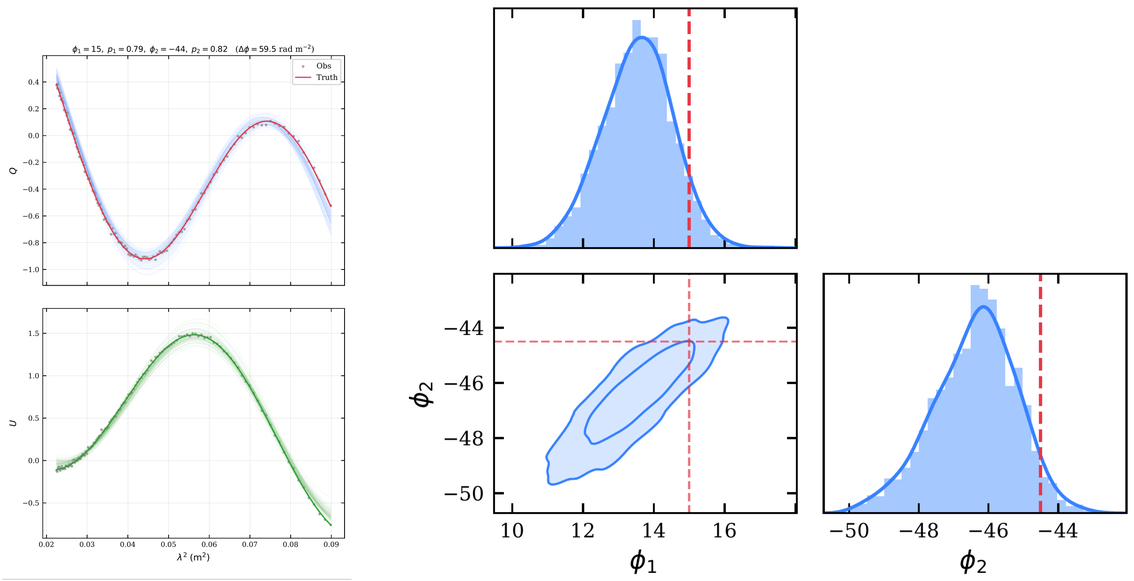
2 Simulate $x = [Q/I, U/I, w]$
noise, RFI flags, channel weights

3 Train flow $q_w(\theta \mid x)$
once per model and band

4 Infer: one forward pass
5,000 samples in < 0.1 s on GPU

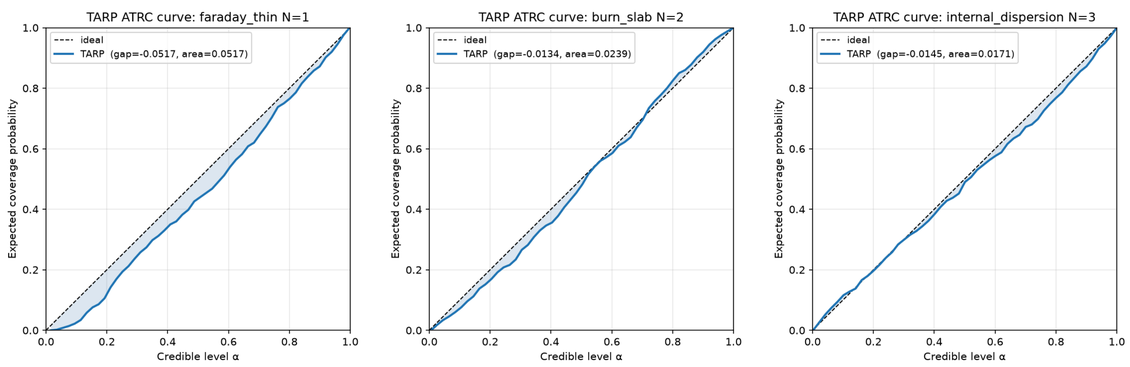
$$\mathcal{L}(w) = -\frac{1}{N}\sum_{i=1}^{N} \log q_w(\theta_i \mid x_i)$$

Two components from one spectrum



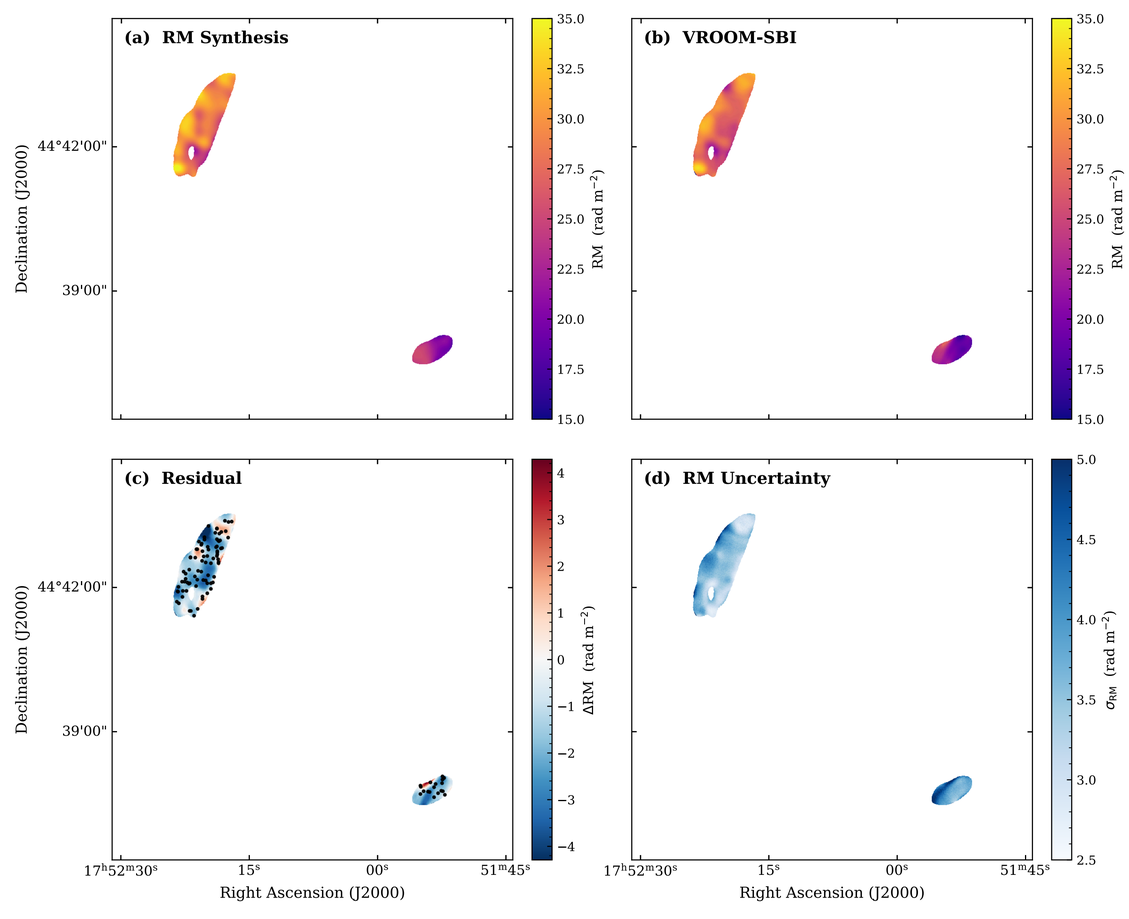
$\varphi_1 = +15, \varphi_2 = -44.5$ rad m^{-2} ($\Delta\varphi = 59.5$). Left: Q, U with posterior draws. Right: joint posterior of φ_1, φ_2; red dashed is truth. Over 256 test cases both RMs are recovered to within 2 rad m^{-2}.

Joint calibration



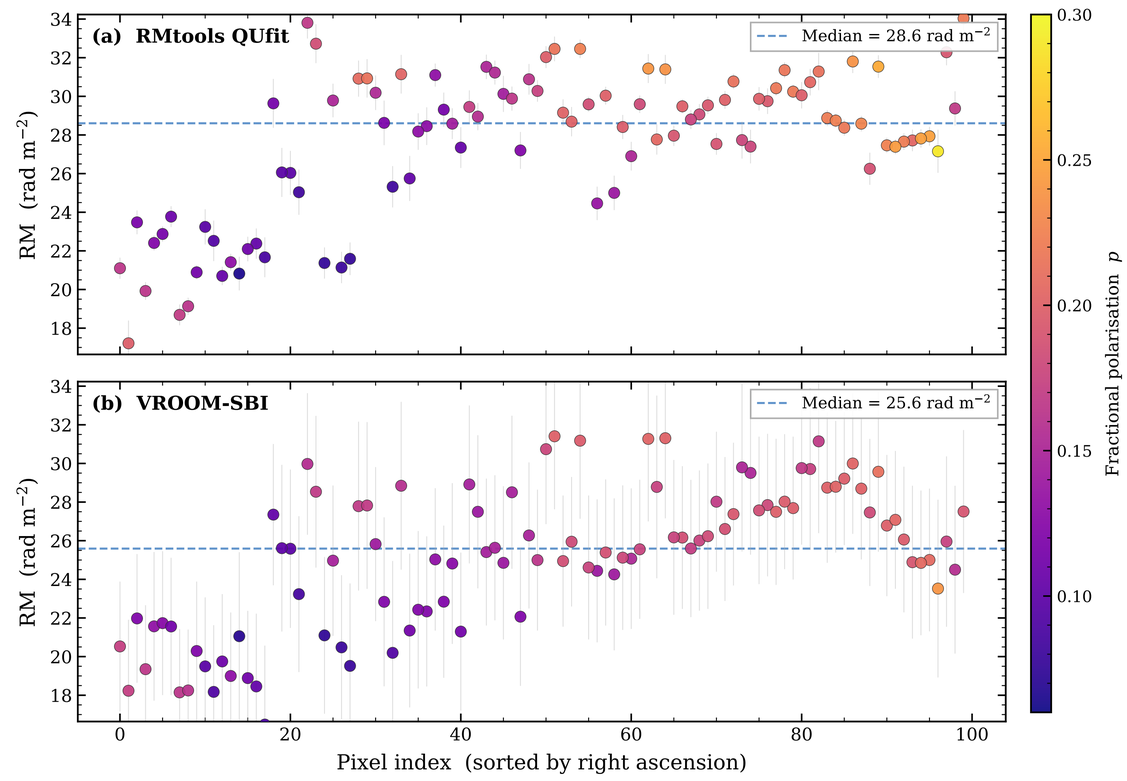
TARP: 15 of 20 posteriors calibrated (4 models, $N = 1$ to 5). Worst case within 5% of the diagonal.

MACS J1752+4440, VLA L-band



(a) RM synthesis (b) VROOM-SBI posterior mean (c) residual (d) posterior RM uncertainty.

Against RM-Tools QU-fitting



100 stratified pixels. VROOM-SBI sits a median 2.3 rad m^{-2} below QU-fitting.

Cost

	per pixel	35,700 pixels
VROOM-SBI, GPU	0.21 s	2.1 h
VROOM-SBI, CPU	0.59 s	5.9 h
QU-fitting, dynesty	1.5 to 2 min	900 to 1,200 h

$$N_{\mathrm{break}} \approx \frac{T_{\mathrm{train}}}{t_{\mathrm{QU}}}$$

One $N = 1$ posterior pays for itself after 165 to 220 pixels.

Code github.com/skunkworks-ra/vroom-sbi
Models huggingface.co/skunkworks-ra/vroomsbi
arXiv arxiv.org/abs/2605.27538
AJ Pal & Jagannathan 2026, AJ 172, 228