Radio Shape Measurement using Deep-Learning

ARGOS-TITAN-TOSCA Workshop 2025

Priyamvad Tripathi

07 July, 2025

Observatoire de la Côte d'Azur

Table of contents

DeepShape

Shape measurement of isolated radio galaxies in the image domain using a supervised deep-learning framework, dubbed DeepShape.

Introduction

Shape measurement of isolated radio galaxies in the image domain using a supervised deep-learning framework, dubbed DeepShape. DeepShape is made up of two independent modules:

1. Plug-and-play (PnP) image reconstruction algorithm [Zhang et al., 2017]

Shape measurement of isolated radio galaxies in the image domain using a supervised deep-learning framework, dubbed DeepShape. DeepShape is made up of two independent modules:

1. Plug-and-play (PnP) image reconstruction algorithm [Zhang et al., 2017]
2. Shape measurement network that predicts shape using PSF-reconstructed image pairs [Tripathi et al., 2024]

Image Reconstruction

Inverse problem: $i_D = Hi_T + \bar{n}$

HQS objective: $\mathcal{L}_{\mu,\lambda}(i,z) = \frac{1}{2\sigma^2}\|i_D - Hi\|^2 + \lambda\Phi(z) + \frac{\mu}{2}\|i - z\|^2$

Inverse problem: $i_D = Hi_T + \bar{\boldsymbol{n}}$
HQS objective: $\mathcal{L}_{\mu,\lambda}(i,z) = \frac{1}{2\sigma^2}\|i_D - Hi\|^2 + \lambda\Phi(z) + \frac{\mu}{2}\|i - z\|^2$

Data step

$$\begin{aligned} i_{k+1} &= \arg\min_i \|i_D - Hi\|^2 + \mu_k\sigma^2\|i - z_k\|^2 \\ &= (H^\dagger H + \mu_k\sigma^2\mathbb{I})^{-1}(H^\dagger i_D + \mu_k\sigma^2 z_k) \end{aligned}$$

Inverse problem: $i_D = H i_T + \bar{\boldsymbol{n}}$

HQS objective: $\mathcal{L}_{\mu,\lambda}(i,z) = \frac{1}{2\sigma^2}\|i_D - Hi\|^2 + \lambda\Phi(z) + \frac{\mu}{2}\|i-z\|^2$

Data step

$$\begin{aligned} i_{k+1} &= \arg\min_i \|i_D - Hi\|^2 + \mu_k\sigma^2\|i - z_k\|^2 \\ &= (H^\dagger H + \mu_k\sigma^2\mathbb{I})^{-1}(H^\dagger i_D + \mu_k\sigma^2 z_k) \end{aligned}$$

Regularization step

$$\begin{aligned} z_{k+1} &= \arg\min_z \frac{1}{2\tau_k^2}\|i_{k+1} - z\|^2 + \Phi(z) \\ &= DN_{\hat{\boldsymbol{\theta}}}(i_{k+1};\tau_k), \quad \tau_k = \sqrt{\lambda/\mu_k} \end{aligned}$$

Inverse problem: $i_D = Hi_T + \bar{\boldsymbol{n}}$
HQS objective: $\mathcal{L}_{\mu,\lambda}(i,z) = \frac{1}{2\sigma^2}\|i_D - Hi\|^2 + \lambda\Phi(z) + \frac{\mu}{2}\|i - z\|^2$

Data step

$$\begin{aligned} i_{k+1} &= \arg\min_i \|i_D - Hi\|^2 + \mu_k\sigma^2\|i - z_k\|^2 \\ &= (H^\dagger H + \mu_k\sigma^2\mathbb{I})^{-1}(H^\dagger i_D + \mu_k\sigma^2 z_k) \end{aligned}$$

Regularization step

$$\begin{aligned} z_{k+1} &= \arg\min_z \frac{1}{2\tau_k^2}\|i_{k+1} - z\|^2 + \Phi(z) \\ &= DN_{\hat{\boldsymbol{\theta}}}(i_{k+1}; \tau_k), \quad \tau_k = \sqrt{\lambda/\mu_k} \end{aligned}$$

Missing frequencies recovered via prior: $F[i_{k+1}](\tilde{n}) = F[z_k](\tilde{n})$

Image Reconstruction

Inverse problem: $i_D = Hi_T + \bar{\boldsymbol{n}}$
HQS objective: $\mathcal{L}_{\mu,\lambda}(i,z) = \frac{1}{2\sigma^2}\|i_D - Hi\|^2 + \lambda\Phi(z) + \frac{\mu}{2}\|i - z\|^2$

Data step

$$
\begin{aligned}
i_{k+1} &= \arg\min_{i} \|i_D - Hi\|^2 + \mu_k\sigma^2\|i - z_k\|^2 \\
&= (H^\dagger H + \mu_k\sigma^2\mathbb{I})^{-1}(H^\dagger i_D + \mu_k\sigma^2 z_k)
\end{aligned}
$$

Regularization step

$$
\begin{aligned}
z_{k+1} &= \arg\min_{z} \frac{1}{2\tau_k^2}\|i_{k+1} - z\|^2 + \Phi(z) \\
&= DN_{\hat{\boldsymbol{\theta}}}(i_{k+1}; \tau_k), \quad \tau_k = \sqrt{\lambda/\mu_k}
\end{aligned}
$$

Missing frequencies recovered via prior: $F[i_{k+1}](\tilde{n}) = F[z_k](\tilde{n})$

We use Plug-and-Play (PnP) scheme rather than unfolding— allowing variable PSFs.

Shape Measurement Network

- Feature extraction via $E(2)$-equivariant CNN [Weiler and Cesa, 2019]
- Convolution kernels are expressed in a steerable $E(2)$ basis:

$$k(\mathbf{x}|\mathbf{w}) = \sum_{\ell=1}^{8} w_\ell(r)\, Y_\ell(\alpha), \quad Y_\ell(\alpha) = e^{i\ell\alpha},\ \mathbf{x} = (r, \alpha) \tag{1}$$

Shape Measurement Network

- Feature extraction via $E(2)$-equivariant CNN [Weiler and Cesa, 2019]
- Convolution kernels are expressed in a steerable $E(2)$ basis:

$$k(\mathbf{x}|\mathbf{w}) = \sum_{\ell=1}^{8} w_\ell(r)\, Y_\ell(\alpha), \quad Y_\ell(\alpha) = e^{i\ell\alpha},\ \mathbf{x} = (r, \alpha) \tag{1}$$

- PSF is compressed via a pre-trained CNN autoencoder

Shape Measurement Network

- Feature extraction via $E(2)$-equivariant CNN [Weiler and Cesa, 2019]
- Convolution kernels are expressed in a steerable $E(2)$ basis:

$$k(\boldsymbol{x}|\boldsymbol{w}) = \sum_{\ell=1}^{8} w_\ell(r)\, Y_\ell(\alpha), \quad Y_\ell(\alpha) = e^{i\ell\alpha},\ \boldsymbol{x} = (r, \alpha) \tag{1}$$

- PSF is compressed via a pre-trained CNN autoencoder
- Encoded PSF + extracted features $\rightarrow$ dense layers $\rightarrow$ shape prediction

Network Architecture

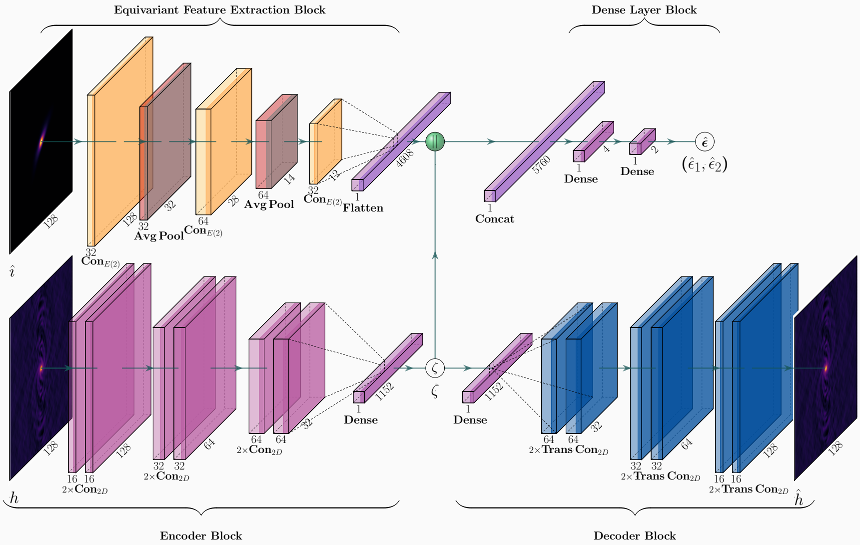

Experiments

Simulation Setup

- **Galaxy Population:** T-RECS SFG wide catalog [Bonaldi et al., 2018], flux range $50\ \mu\text{Jy} \leq I < 200\ \mu\text{Jy}$
- **Morphology:** Sérsic profile

$$I(r) = I_e \exp\left\{-b_{n_s}\left[\left(\frac{r}{r_e}\right)^{1/n_s} - 1\right]\right\}, \quad n_s \sim \mathcal{U}(0.7, 2)$$

- **Image Domain:** Isolated galaxies on 128×128 grid, pointing centered on galaxy
- **Visibility Sampling:** SKA-MID A4, $\nu_0 = 1.4\text{GHz}$, 8 hr obs, $\Delta t = 300\text{s} \Rightarrow 1.8\text{M}$ visibilities
- **Noise Model:** Additive Gaussian noise $\mathcal{N}(0, \sigma_V)$, with

$$\sigma_V = \frac{2k_B T_{\text{sys}}}{\eta_S A_{\text{eff}} \sqrt{2\tau_{\text{int}} \Delta\nu}} \approx 0.5\ \text{mJy}$$

- **Imaging:** Gridding via Briggs weighting ($R = -0.5$), resulting in RMS noise $\sigma \approx 0.71\ \mu\text{Jy}$

HQS-PnP vs MS-CLEAN

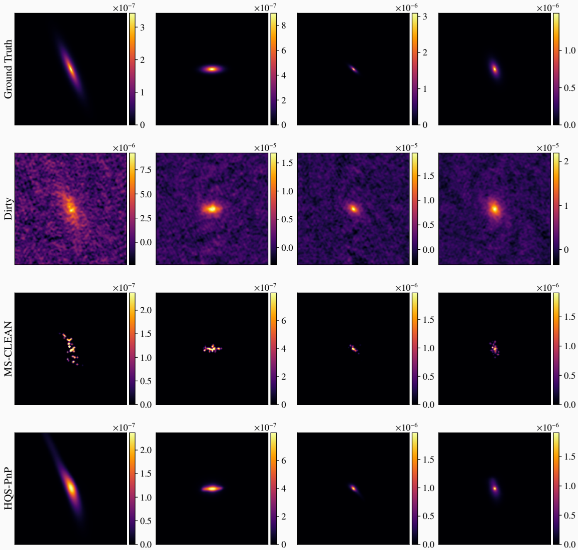

Image Residuals

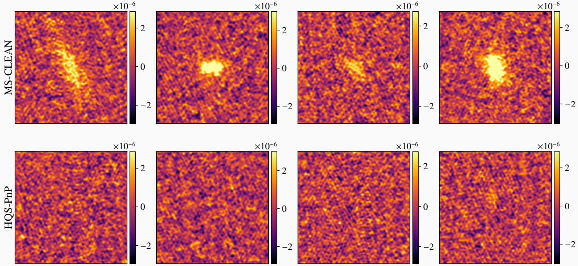

- Residuals: $i_R = F^{\dagger}G(v - G^{\dagger}F\hat{\imath})$

Image Residuals

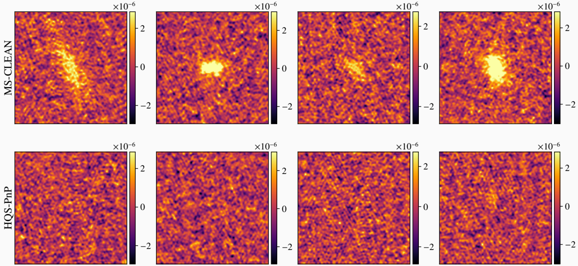

- Residuals: $i_R = F^\dagger G(v - G^\dagger F\hat{\imath})$
- If i_R shows structure, further deconvolution is required — major-minor loops — common in wide-field imaging

Image Residuals

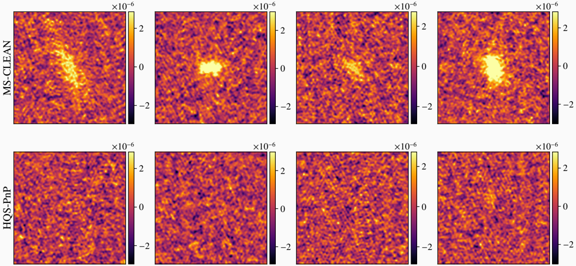

- Residuals: $i_R = F^\dagger G(v - G^\dagger F\hat{\imath})$
- If i_R shows structure, further deconvolution is required — major-minor loops — common in wide-field imaging
- **HQS-PnP:** Residuals appear noise-like $\Rightarrow$ no extra correction needed

Image Residuals

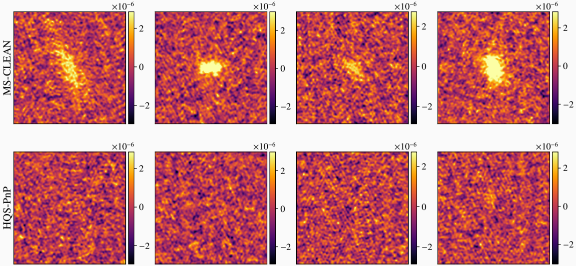

- Residuals: $i_R = F^\dagger G(v - G^\dagger F\hat{\imath})$
- If i_R shows structure, further deconvolution is required — major-minor loops — common in wide-field imaging
- **HQS-PnP:** Residuals appear noise-like ⇒ no extra correction needed
- **MS-CLEAN:** Residuals contain structure ⇒ further deconvolution could improve results, but adds computational cost

Image Residuals

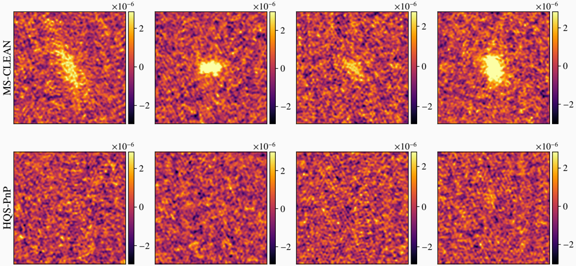

- Residuals: $i_R = F^\dagger G(v - G^\dagger F\hat{\imath})$
- If i_R shows structure, further deconvolution is required — major-minor loops — common in wide-field imaging
- **HQS-PnP:** Residuals appear noise-like $\Rightarrow$ no extra correction needed
- **MS-CLEAN:** Residuals contain structure $\Rightarrow$ further deconvolution could improve results, but adds computational cost
- Note: Our analysis is at the pointing center $\Rightarrow$ minimal w-term

Dependence on Peak S/N

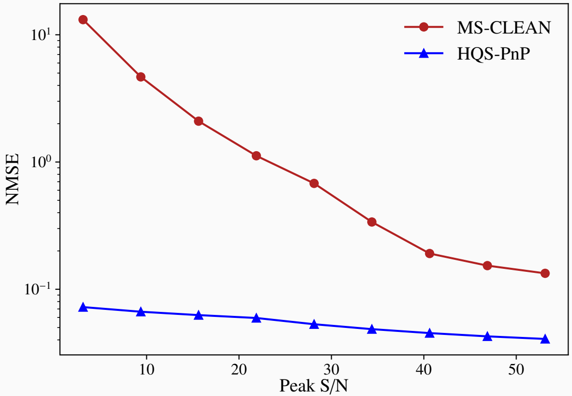

Figure 1: Variation of NMSE $\equiv \frac{\sum_p \|\hat{\imath}^p - i_T^p\|_2^2}{\sum_p \|i_T^p\|_2^2}$, as a function of peak S/N. The statistics were calculated based on a test set containing 20,000 images split into 9 peak S/N bins.

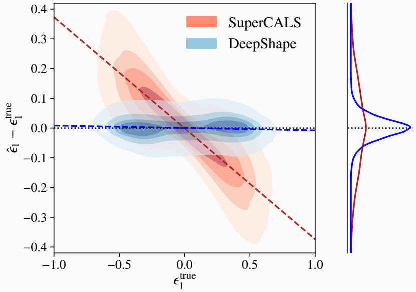

Figure 2: Shape residuals from three methods: DeepShape, SuperCALS, and RadioLensfit. Contours via KDE; dashed lines show best-fit trends.

- **SuperCALS** [Harrison et al., 2020]: Image-domain method; uses MS-CLEAN for deconvolution and calibrates shape estimates by injecting known model sources and correcting measured biases.

Shape Measurement

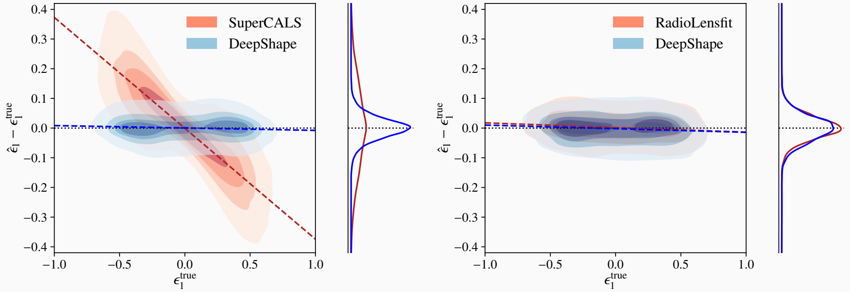

Figure 2: Shape residuals from three methods: DeepShape, SuperCALS, and RadioLensfit. Contours via KDE; dashed lines show best-fit trends.

- **SuperCALS** [Harrison et al., 2020]: Image-domain method; uses MS-CLEAN for deconvolution and calibrates shape estimates by injecting known model sources and correcting measured biases.
- **RadioLensfit** [Rivi et al., 2016]: Visibility-domain parametric model-fitting using exponential brightness profile; ellipticity estimated by marginalizing the likelihood over nuisance parameters.

Table 1: Comparison metrics for the ellipticity measurements from different methods.

Ngal	Method	ρ_1	$\hat{m}_1$ $[10^{-3}]$	$\hat{c}_1$ $[10^{-4}]$	$\sqrt{\langle(\hat{\epsilon}_1 - \epsilon_1^T)^2\rangle}$
20000	SuperCALS	0.939	-373.5 ± 1.6	-7.1 ± 5.5	0.128
	DeepShape	**0.993**	$\mathbf{-8.1} \pm 0.9$	$\mathbf{3.6} \pm 2.9$	**0.041**
1800	RadioLensfit	**0.992**	-15.7 ± 2.9	$\mathbf{19.3} \pm 9.5$	**0.040**
	DeepShape	0.990	$\mathbf{-12.2} \pm 3.3$	-20.2 ± 10.6	0.045

Table 1: Comparison metrics for the ellipticity measurements from different methods.

Ngal	Method	ρ_1	$\hat{m}_1$ $[10^{-3}]$	$\hat{c}_1$ $[10^{-4}]$	$\sqrt{\langle(\hat{\epsilon}_1 - \epsilon_1^T)^2\rangle}$
20000	SuperCALS	0.939	-373.5 ± 1.6	-7.1 ± 5.5	0.128
	DeepShape	**0.993**	$\mathbf{-8.1} \pm 0.9$	$\mathbf{3.6} \pm 2.9$	**0.041**
1800	RadioLensfit	**0.992**	-15.7 ± 2.9	$\mathbf{19.3} \pm 9.5$	**0.040**
	DeepShape	0.990	$\mathbf{-12.2} \pm 3.3$	-20.2 ± 10.6	0.045

Table 2: Computational time

Method	SuperCALS	RadioLensfit	DeepShape
Prediction time	**0.18 sec**	4 min	0.22 sec

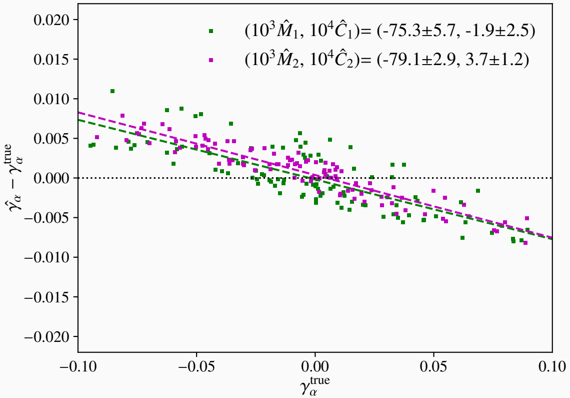

Figure 3: Shear measurement residuals as a function of input shear. Each point represents the average of shape measurements from DeepShape derived from 10,000 galaxies with varying properties, but sharing the same shear.

Quality cuts

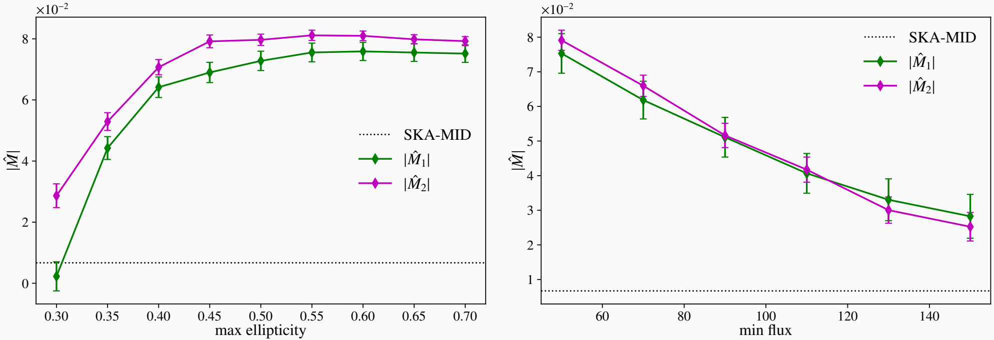

Figure 4: Estimated multiplicative shear bias (M_α) improves when applying quality cuts on galaxy ensembles based on ellipticity and flux.

Wide Field measurements

Overview

- The SKA-MID primary beam will be around 1 deg^2 containing upto 10,000 galaxies.
- Galaxies are not localized in the visibility domain, making source separation complicated

Overview

- The SKA-MID primary beam will be around 1 deg^2 containing upto 10,000 galaxies.
- Galaxies are not localized in the visibility domain, making source separation complicated
- Model-based shape measurement methods [Chang et al., 2004, Rivi et al., 2016] jointly fill all sources in the FoV

Overview

- The SKA-MID primary beam will be around 1 deg^2 containing upto 10,000 galaxies.
- Galaxies are not localized in the visibility domain, making source separation complicated
- Model-based shape measurement methods [Chang et al., 2004, Rivi et al., 2016] jointly fill all sources in the FoV
- Another method would be to create small facets around each source in the FoV

- Visibilities w.r.t. pointing center $(l_0, m_0) = (0, 0)$:

$$V(u, v, w) = \iint I(l - l_0, m - m_0)\, e^{-2\pi j(u(l-l_0)+v(m-m_0)+w(n-n_0-1))}\, \mathrm{d}l\, \mathrm{d}m$$

with $n = \sqrt{1 - l^2 - m^2}$

- Visibilities w.r.t. pointing center $(l_0, m_0) = (0, 0)$:

$$V(u, v, w) = \iint I(l - l_0, m - m_0)\, e^{-2\pi j(u(l-l_0)+v(m-m_0)+w(n-n_0-1))}\, \mathrm{d}l\, \mathrm{d}m$$

with $n = \sqrt{1 - l^2 - m^2}$

- To center a facet at (l_g, m_g), visibilities are phase-rotated by $e^{-2\pi j\phi}$ with $\phi = ul_g + vm_g + wn_g$:

$$V'(u, v, w) = \iint I(l - l_g, m - m_g)\, e^{-2\pi j(u(l-l_g)+v(m-m_g)+w(n-n_g-1))}\, \mathrm{d}l\, \mathrm{d}m$$

Faceting

- Visibilities w.r.t. pointing center $(l_0, m_0) = (0, 0)$:

$$V(u, v, w) = \iint I(l - l_0, m - m_0)\, e^{-2\pi j(u(l-l_0)+v(m-m_0)+w(n-n_0-1))}\, \mathrm{d}l\, \mathrm{d}m$$

  with $n = \sqrt{1 - l^2 - m^2}$

- To center a facet at (l_g, m_g), visibilities are phase-rotated by $e^{-2\pi j\phi}$ with $\phi = ul_g + vm_g + wn_g$:

$$V'(u, v, w) = \iint I(l - l_g, m - m_g)\, e^{-2\pi j(u(l-l_g)+v(m-m_g)+w(n-n_g-1))}\, \mathrm{d}l\, \mathrm{d}m$$

- The (u, v, w) coordinates are rotated to make the facet plane tangent to the celestial sphere:

$$(u', v', w')^\top = R(\alpha_g, \delta_g) R^\top(\alpha_0, \delta_0)(u, v, w)^\top$$

  where

$$R(\alpha, \delta) = \begin{bmatrix} -\sin\alpha & \cos\alpha & 0 \\ -\sin\delta\cos\alpha & -\sin\delta\sin\alpha & \cos\delta \\ \cos\delta\cos\alpha & \cos\delta\sin\alpha & \sin\delta \end{bmatrix}$$

Facet Plane

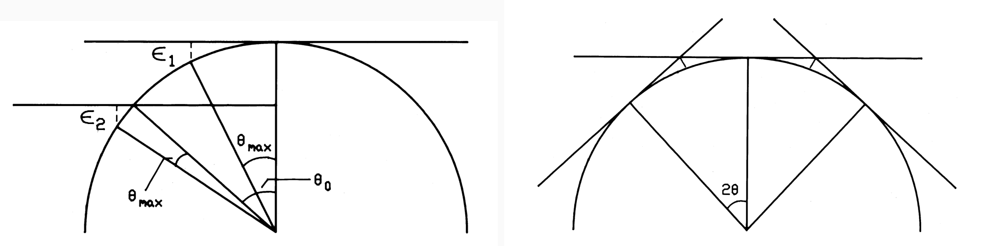

Figure 5: *Left:* Phase shifting alone limits the field of view. *Right:* Rotating (u, v, w) aligns facets with the sky, forming a polyhedral tiling. Source: [Taylor et al., 1999].

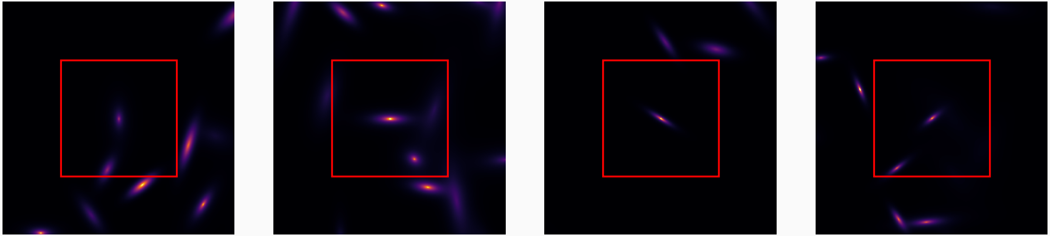

Figure 6: The far sidelobes of the PSF can pick up distant sources, causing noise-like contamination in the dirty image. The dirty image needs to be large enough so that these sources can be properly deconvolved.

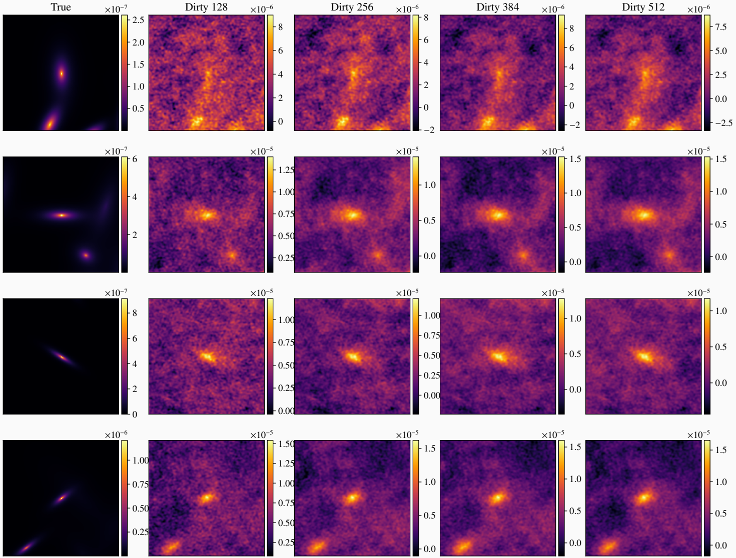
True
Dirty 128
Dirty 256
Dirty 384
Dirty 512

Far Sidelobe Confusion Noise

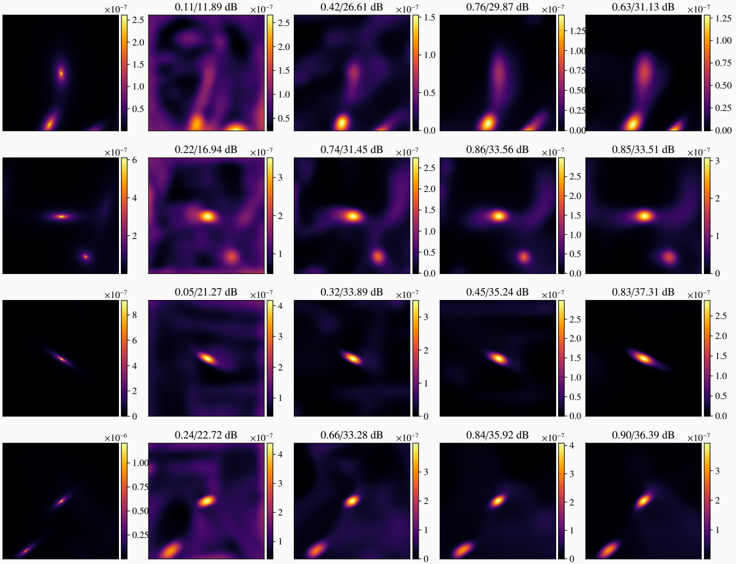

Galaxy Deblending

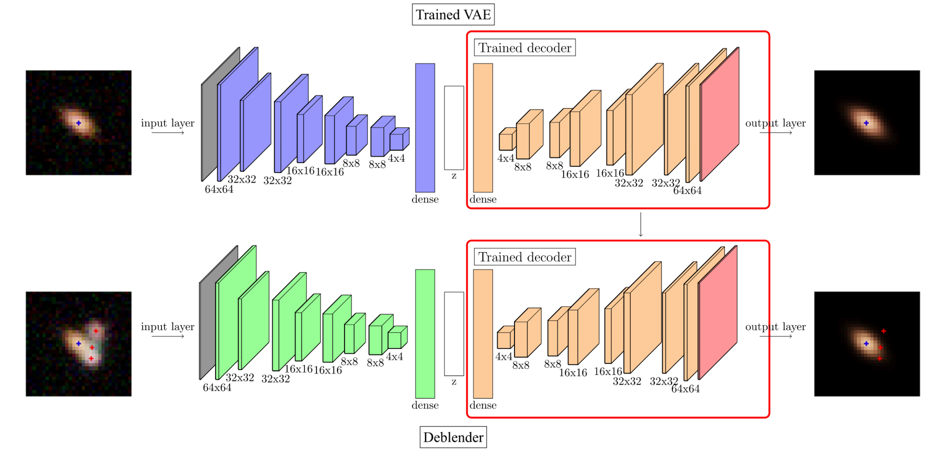

Figure 7: The VAE is first trained to generate centered, isolated galaxy images. This is followed by retraining the encoder to find the latent space representation of the central galaxy in postage stamps containing a blend.

Theory

- A VAE learns to encode observed data x into latent variables z, and decode z back to x, in order to maximize the marginal likelihood $\log p(x)$.

- A VAE learns to encode observed data x into latent variables z, and decode z back to x, in order to maximize the marginal likelihood $\log p(x)$.
- In practice, this is done by maximizing the evidence lower bound (ELBO):

$$\log p(x) \geq \underbrace{\mathbb{E}_{q(z|x)}[\log p(x|z)]}_{\text{(-ve) Recon loss}} \quad - \quad \underbrace{\beta\,\mathrm{KL}(q(z|x) \,\|\, p(z))}_{\text{(-ve) KL loss}} \tag{2}$$

- The reconstruction term ensures data fidelity, while the KL term (second term) regularizes the latent space by encouraging the approximate posterior $q(z|x)$ to align with the prior $p(z)$.

- A VAE learns to encode observed data x into latent variables z, and decode z back to x, in order to maximize the marginal likelihood $\log p(x)$.
- In practice, this is done by maximizing the evidence lower bound (ELBO):

$$\log p(x) \geq \underbrace{\mathbb{E}_{q(z|x)}[\log p(x|z)]}_{\text{(-ve) Recon loss}} - \underbrace{\beta\,\mathrm{KL}(q(z|x)\,\|\,p(z))}_{\text{(-ve) KL loss}} \tag{2}$$

- The reconstruction term ensures data fidelity, while the KL term (second term) regularizes the latent space by encouraging the approximate posterior $q(z|x)$ to align with the prior $p(z)$.
- The deblender is trained using:

$$\hat{\phi}_2 = \underset{\phi_2}{\operatorname{argmin}}\ \mathrm{KL}(q_{\phi_1}(z_1|x_1)\,\|\,q_{\phi_2}(z_2|x_2)) + \mathrm{KL}(q_{\phi_2}(z_2|x_2)\,\|\,q_{\phi_1}(z_1|x_1)) \tag{3}$$

Galaxy Deblending

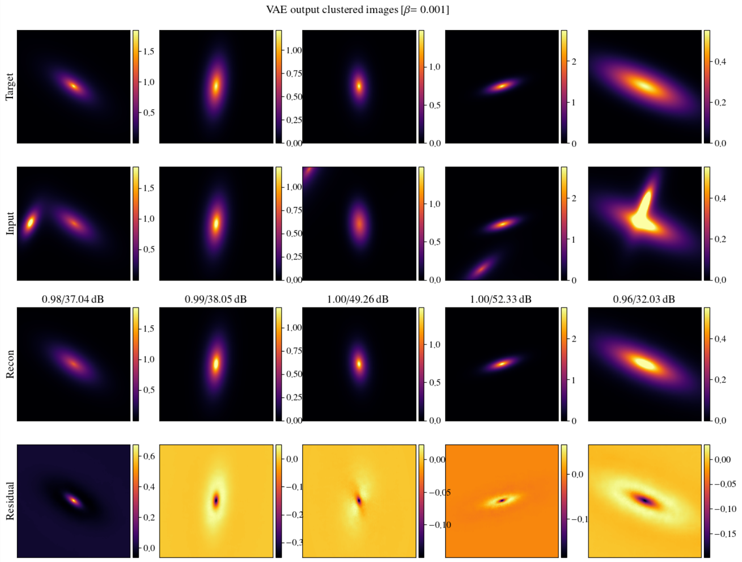

Galaxy Deblending

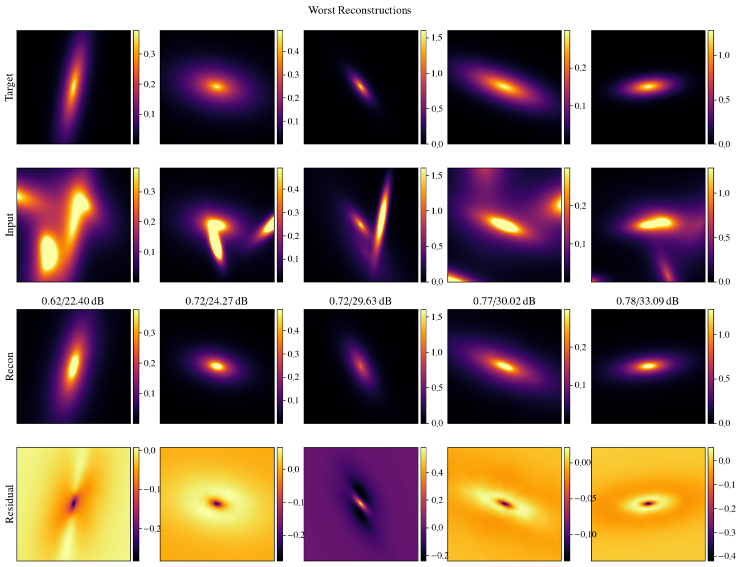

Galaxy Deblending

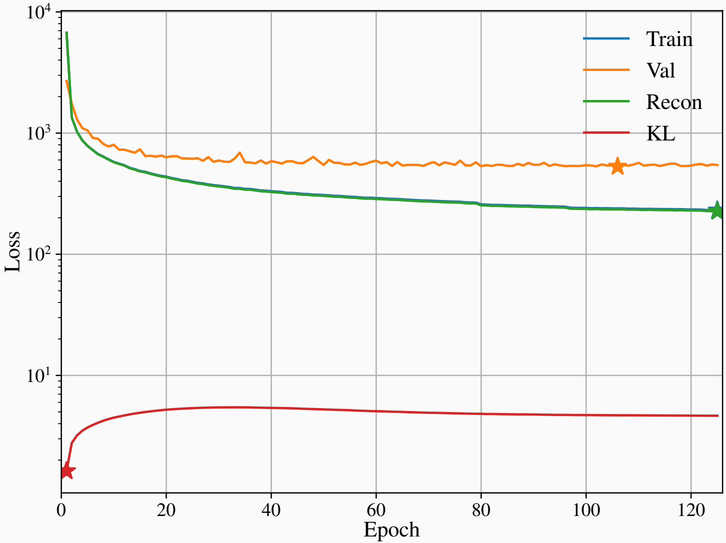

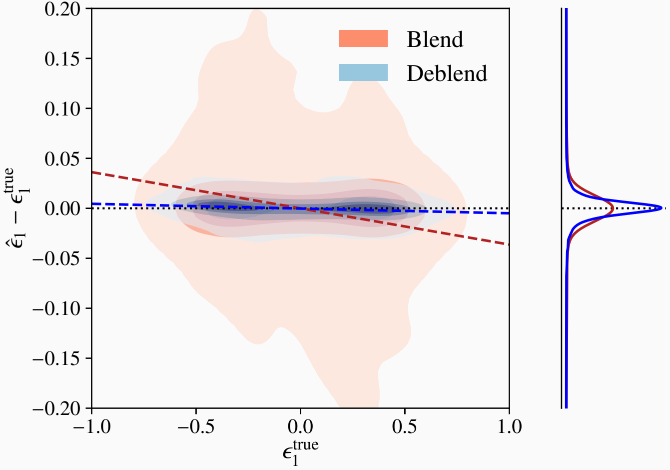

Blended:

$m_1 = (-36.2 \pm 2.2) \times 10^{-3}$

$c_1 = (-7.1 \pm 7.7) \times 10^{-4}$

Deblended:

$m_1 = (-4.7 \pm 0.4) \times 10^{-3}$

$c_1 = (6.9 \pm 1.3) \times 10^{-4}$

Deblending significantly reduces multiplicative bias.

Figure 8: Shape measurement comparison for blended (red) and deblended (blue) stamps using GalSim.

Thank you for your time

Training Details

- DRUNET is initialized using pre-trained weights provided by the DeepInverse library. These weights are subsequently fine-tuned on a training set comprising 250 000 galaxy image pairs, each consisting of a noiseless and a corresponding noisy image.
- Autoencoder is trained using a training set of 100,000 PSFs
- DeepShape is trained using a two-step procedure
  1. The network is pre-trained using a training set of 250,000 true noiseless galaxy images
  2. The network is fine-tuned using a training set of 250,000 reconstructed galaxy image-PSF pairs.

RadioLensfit [Rivi et al., 2016]

- Works using visibilities
- Galaxy brightness profile: $I(r) = I_0 \exp(-r/\alpha)$,
- Transformation matrix $\mathbf{A}$ with ellipticity parameters $\mathbf{e} = (e_1, e_2)$ such that:

$$\begin{pmatrix} l_r \\ m_r \end{pmatrix} = \mathbf{A}x = \begin{pmatrix} 1-e_1 & -e_2 \\ -e_2 & 1+e_1 \end{pmatrix} \times \begin{pmatrix} l \\ m \end{pmatrix}$$

- Observed visibility due to a galaxy at point $\mathbf{k} = (u, v)$ can be given by:

$$V_s(u,v) = \frac{2\pi\alpha^2 I_0}{|\mathbf{A}|(1 + 4\pi^2\alpha^2|\mathbf{A}^{-\intercal}k|)^{3/2}} \times \exp 2\pi i \mathbf{k}^\intercal \mathbf{x_0} \tag{4}$$

- Perform a Bayesian marginalization of the likelihood over I_0, α and source centroid position $\mathbf{x_0} = (l_0, m_0) \Rightarrow P(\mathbf{A}|D)$

SuperCLASS calibration [Harrison et al., 2020]

- Reconstruct image by deconvolving the PSF from the dirty image and estimate ellipticity $\boldsymbol{\epsilon}_k^{\text{calc}}$
- In the residual image, inject model sources with the same size and flux properties, but known ellipticity $\epsilon_i^{\text{inp}} = \{0,\ \pm 0.2375,\ \pm 0.475,\ \pm 0.7125,\ \pm 0.95\}$
- Simulate visibilities $\Rightarrow$ Dirty Image $\Rightarrow$ Reconstructed Image $\Rightarrow$ Measure ellipticity $\boldsymbol{\epsilon}^{\text{obs}}$
- Fit second order 2D polynomial $b_k(\epsilon_1^{\text{inp}}, \epsilon_2^{\text{inp}}) = \epsilon_1^{\text{obs}} - \epsilon_1^{\text{inp}}$
- Calibrate observed ellipticities using $\epsilon_{1,k}^{\text{SC}} = \epsilon_{1,k}^{\text{calc}} - b_k(\epsilon_{1,k}^{\text{calc}}, \epsilon_{2,k}^{\text{calc}})$
- Repeat for ϵ_2

Reconstruction using tanh scaling

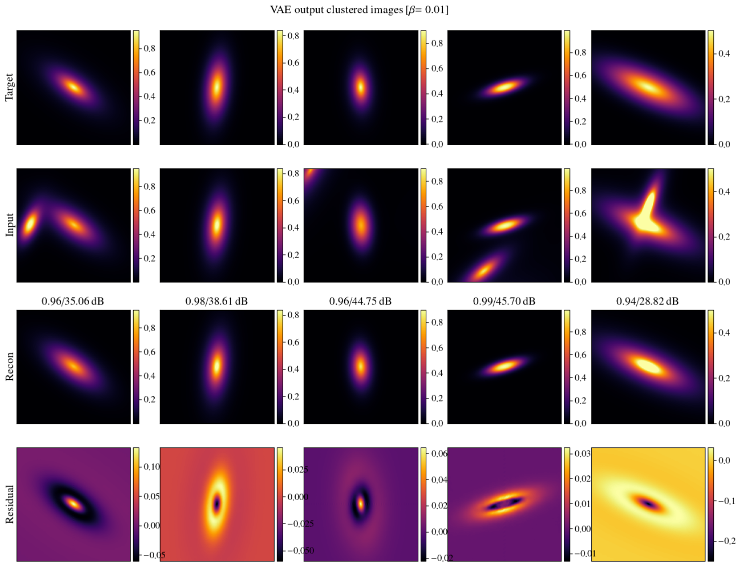

Reconstruction using tanh scaling

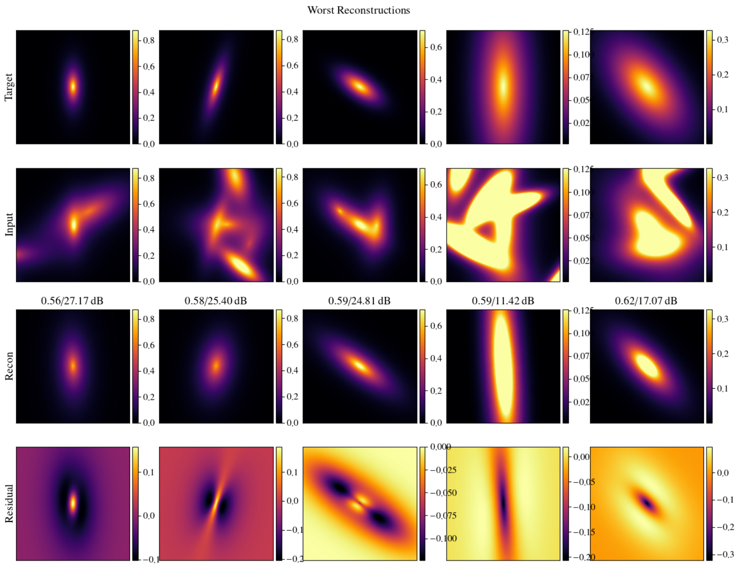

Arcelin, B., Doux, C., Aubourg, E., Roucelle, C., and Collaboration), T. L. D. E. S. (2020).
Deblending galaxies with variational autoencoders: A joint multiband, multi-instrument approach.
Monthly Notices of the Royal Astronomical Society, 500(1):531–547.

Bonaldi, A., Bonato, M., Galluzzi, V., Harrison, I., Massardi, M., Kay, S., De Zotti, G., and Brown, M. L. (2018).
The tiered radio extragalactic continuum simulation (t-recs).
MNRAS, 482(1):2–19.

Chang, T.-C., Refregier, A., and Helfand, D. J. (2004).
Weak lensing by large-scale structure with the FIRST radio survey.
The Astrophysical Journal, 617(2):794–810.

References II

Harrison, I., Brown, M. L., Tunbridge, B., Thomas, D. B., Hillier, T., Thomson, A. P., Whittaker, L., Abdalla, F. B., Battye, R. A., Bonaldi, A., Camera, S., Casey, C. M., Demetroullas, C., Hales, C. A., Jackson, N. J., Kay, S. T., Manning, S. M., Peters, A., Riseley, C. J., and Watson, R. A. (2020).
Superclass – iii. weak lensing from radio and optical observations in data release 1.
MNRAS, 495(2):1737–1759.

Rivi, M., Miller, L., Makhathini, S., and Abdalla, F. B. (2016).
Radio weak lensing shear measurement in the visibility domain – i. methodology.
MNRAS, 463(2):1881–1890.

Taylor, G. B., Carilli, C. L., and Perley, R. A., editors (1999).
Synthesis Imaging in Radio Astronomy II, volume 180 of *Astronomical Society of the Pacific Conference Series*.

Tripathi, P., Wang, S., Prunet, S., and Ferrari, A. (2024).
Shape measurement of radio galaxies using equivariant cnns.
In *EUSIPCO*, pages 2377–2381.

Weiler, M. and Cesa, G. (2019).
General e (2)-equivariant steerable cnns.
Advances in neural information processing systems, 32.

Zhang, K., Zuo, W., Gu, S., and Zhang, L. (2017).
Learning deep cnn denoiser prior for image restoration.