TITAN - ARGOS - TOSCA Meeting

07/07/2025

📍 **Heraklion**, Greece

Deep and Sparse Denoising of High-z Spectral Data Cubes:

*Benchmarking **U-Net**, **Wavelets**, and **BSS** – from simulations to observations*

Arnab Lahiry, Tanio Díaz-Santos, Jean-Luc Starck

Niranjan Chandra Roy, Daniel Anglés-Alcázar,

Grigorios Tsagkatakis, Panagiotis Tsakalides

M O T I V A T I O N

Challenges with high-redshift IFUs

- **Low SNR**: Weak signals dominated by noise due to large cosmological distances.
- **Instrumental Noise**: Detector artifacts and sky subtraction errors.
- **Convolution Effects**: PSF and beam smearing may distort flux values affecting flux conservation.
- **Denoising Bias**: Aggressive methods may lead to over-smoothing and can suppress real signals or overestimate "not-real" signal
- **Lack of a large sample space**

A I M S

To **analyse and compare denoising methods** for **spectral cubes** across a **broad parameter space**, focusing on:

- **Noise characteristics:** Spatially correlated Gaussian noise (ALMA)
- **Noise levels:** Varying signal-to-noise ratios
- **Spatial resolutions:** Resolved or unresolved by the synthesised beam

And understand how each method performs under different conditions and identify the **optimal approach** for **flux conservation** and **denoising,** for specific datasets and noise characteristics, with application to observational data

Datasets

- Toy models of rotating galaxies
- Mock IFU data from cosmological simulations
- Observational data (ALMA)

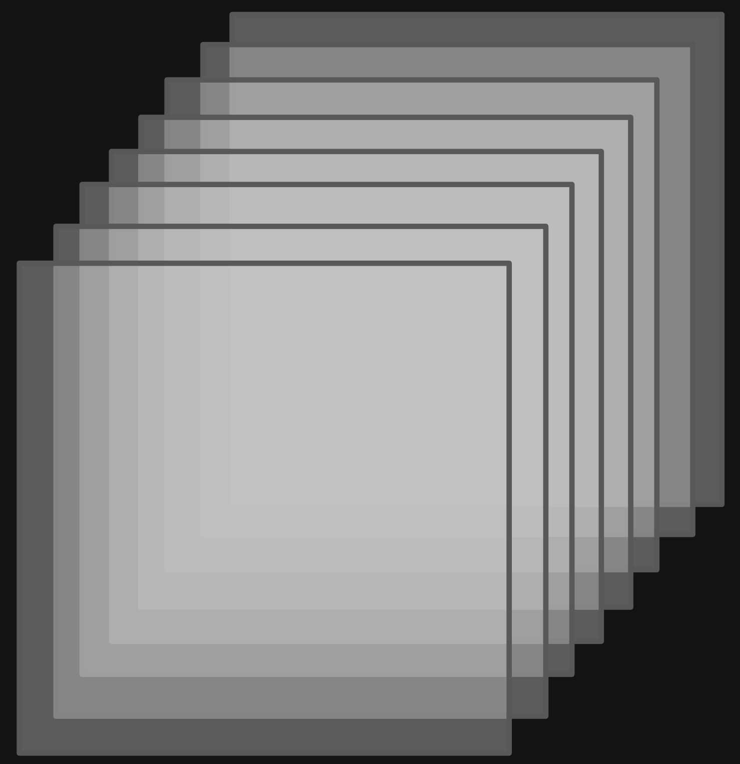

Spectral cubes

multiple spectral observation of the same spatial area, where each *(x,y)* spatial point corresponds to a spectrum

Constructing toy data of rotating galaxies

Describing the specific intensity profile in 3D space

$$I_\nu(x, y, z) = I_e \exp\left[-b_n\left(\left(\frac{\sqrt{x^2+y^2}}{R_e}\right)^{1/n} - 1\right)\right] \cdot \exp\left(-\frac{|z|}{h_z}\right)$$

Sérsic profile

Exponential profile

Constructing toy data of rotating galaxies

Describing the specific intensity profile in 3D space

Sérsic profile

$$I_\nu(x,y,z) = I_e \exp\left[-b_n\left(\left(\frac{\sqrt{x^2+y^2}}{R_e}\right)^{1/n}-1\right)\right]\cdot\exp\left(-\frac{|z|}{h_z}\right)$$

Exponential profile

Specific Intensity at a given location in 3D space

Effective Specific Intensity: Specific intensity contained within the half-light radius

Constant depending on n, ensuring that the effective radius encloses half of the total light

Effective/half-light radius: radius at which half of the total light of the galaxy is contained

Sérsic index: determines the shape of the profile

Scale height: determines how the flux density varies above or below the galactic mid-plane.

Integrated along Z-axis
(Face-on)

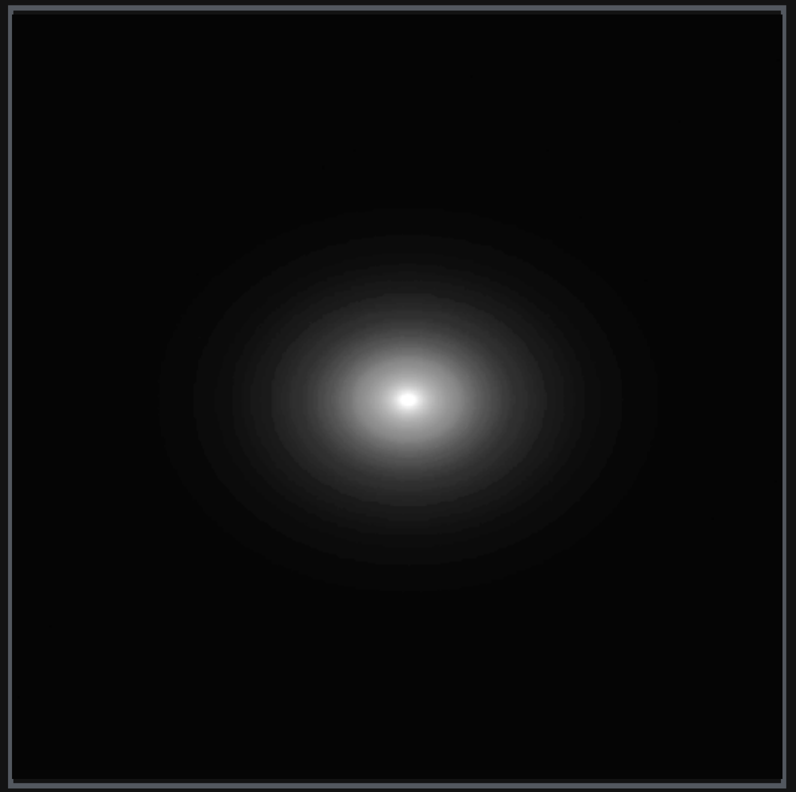

Integrated along Y-axis
(Edge-on)

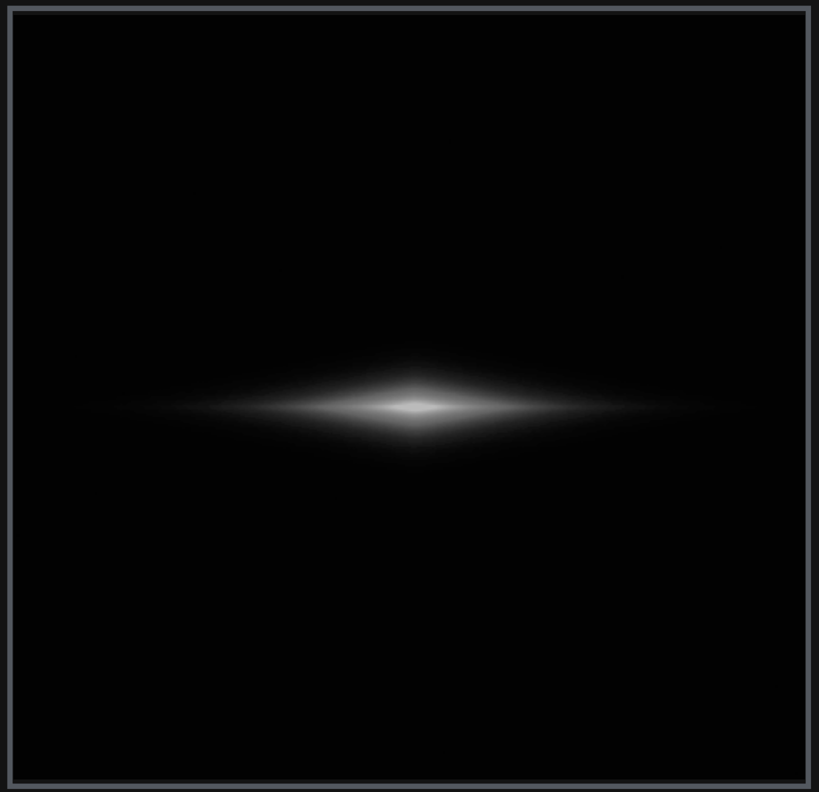

Constructing toy data of rotating galaxies

- Rotation velocity vectors $(\mathbf{v_x}, \mathbf{v_y}, \mathbf{v_z})$ calculated **in the plane of the disk**

$$v = v_0 \times 1.022 \times \left(\frac{R}{R_0}\right)^{0.0803}$$

- The entire system is rotated and **Z axis is chosen as the line of sight**
- N_z 2D projections are made along line of sight based on $\mathbf{v_z}$ bins
- Final spectral cube is convolved with **2D circular Gaussian beam**
- Galaxies smaller than beam FWHM are **unresolved**, and larger than FWHM are **resolved**
- **Gaussian noise** is overlaid onto the cube

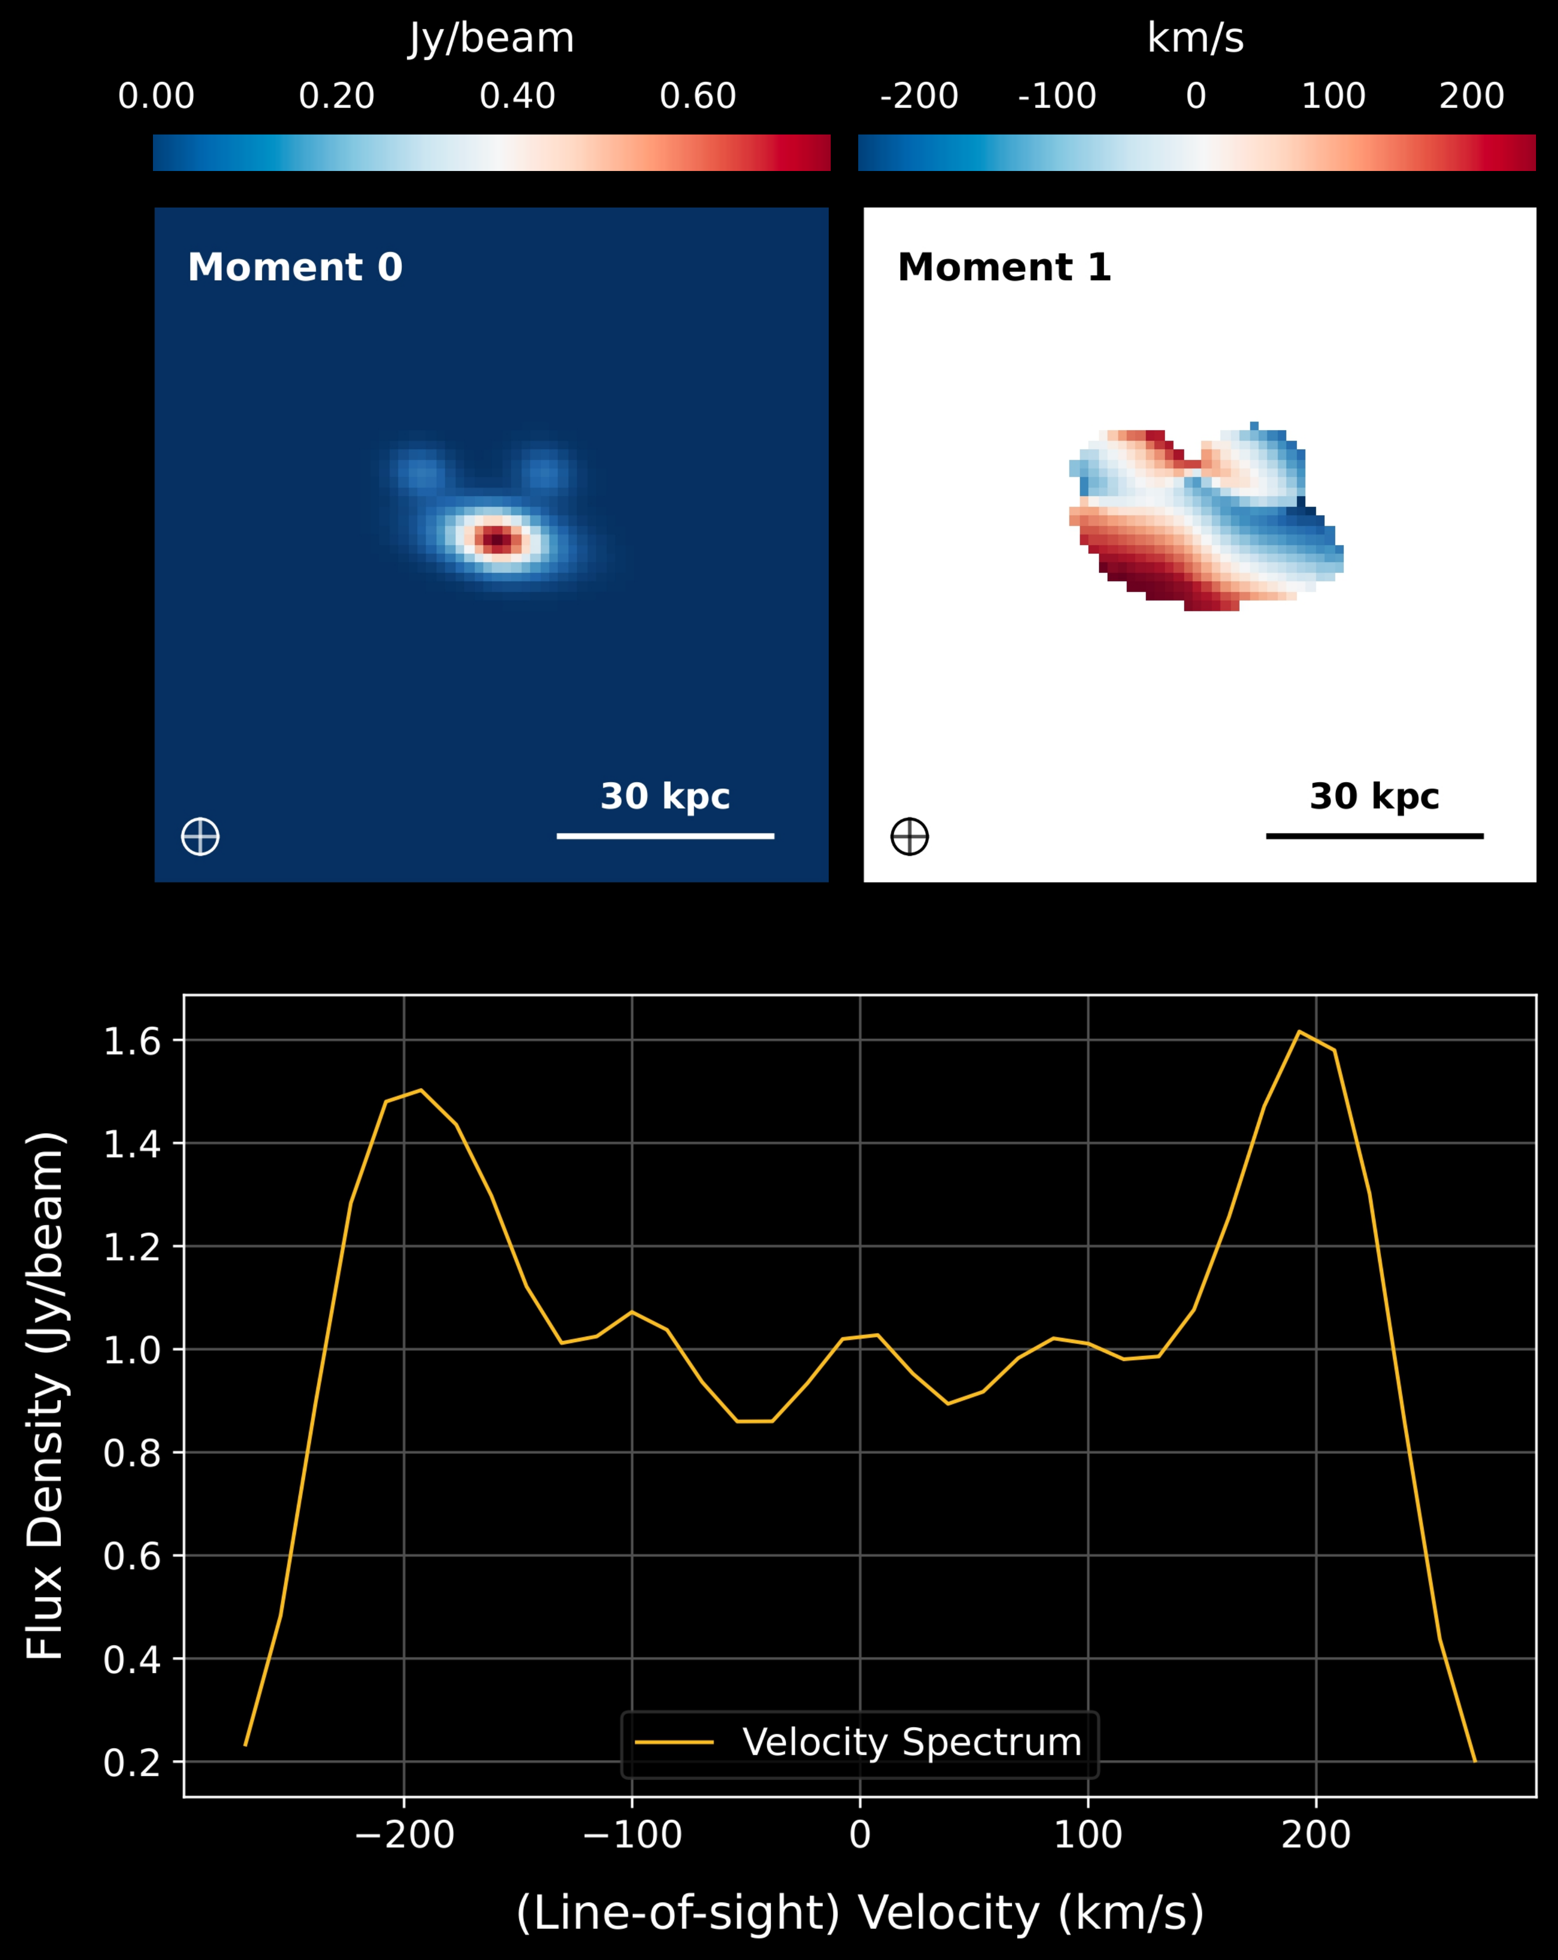

Resolved source with multiple satellites showing double-horned spectra with well-resolved kinematics

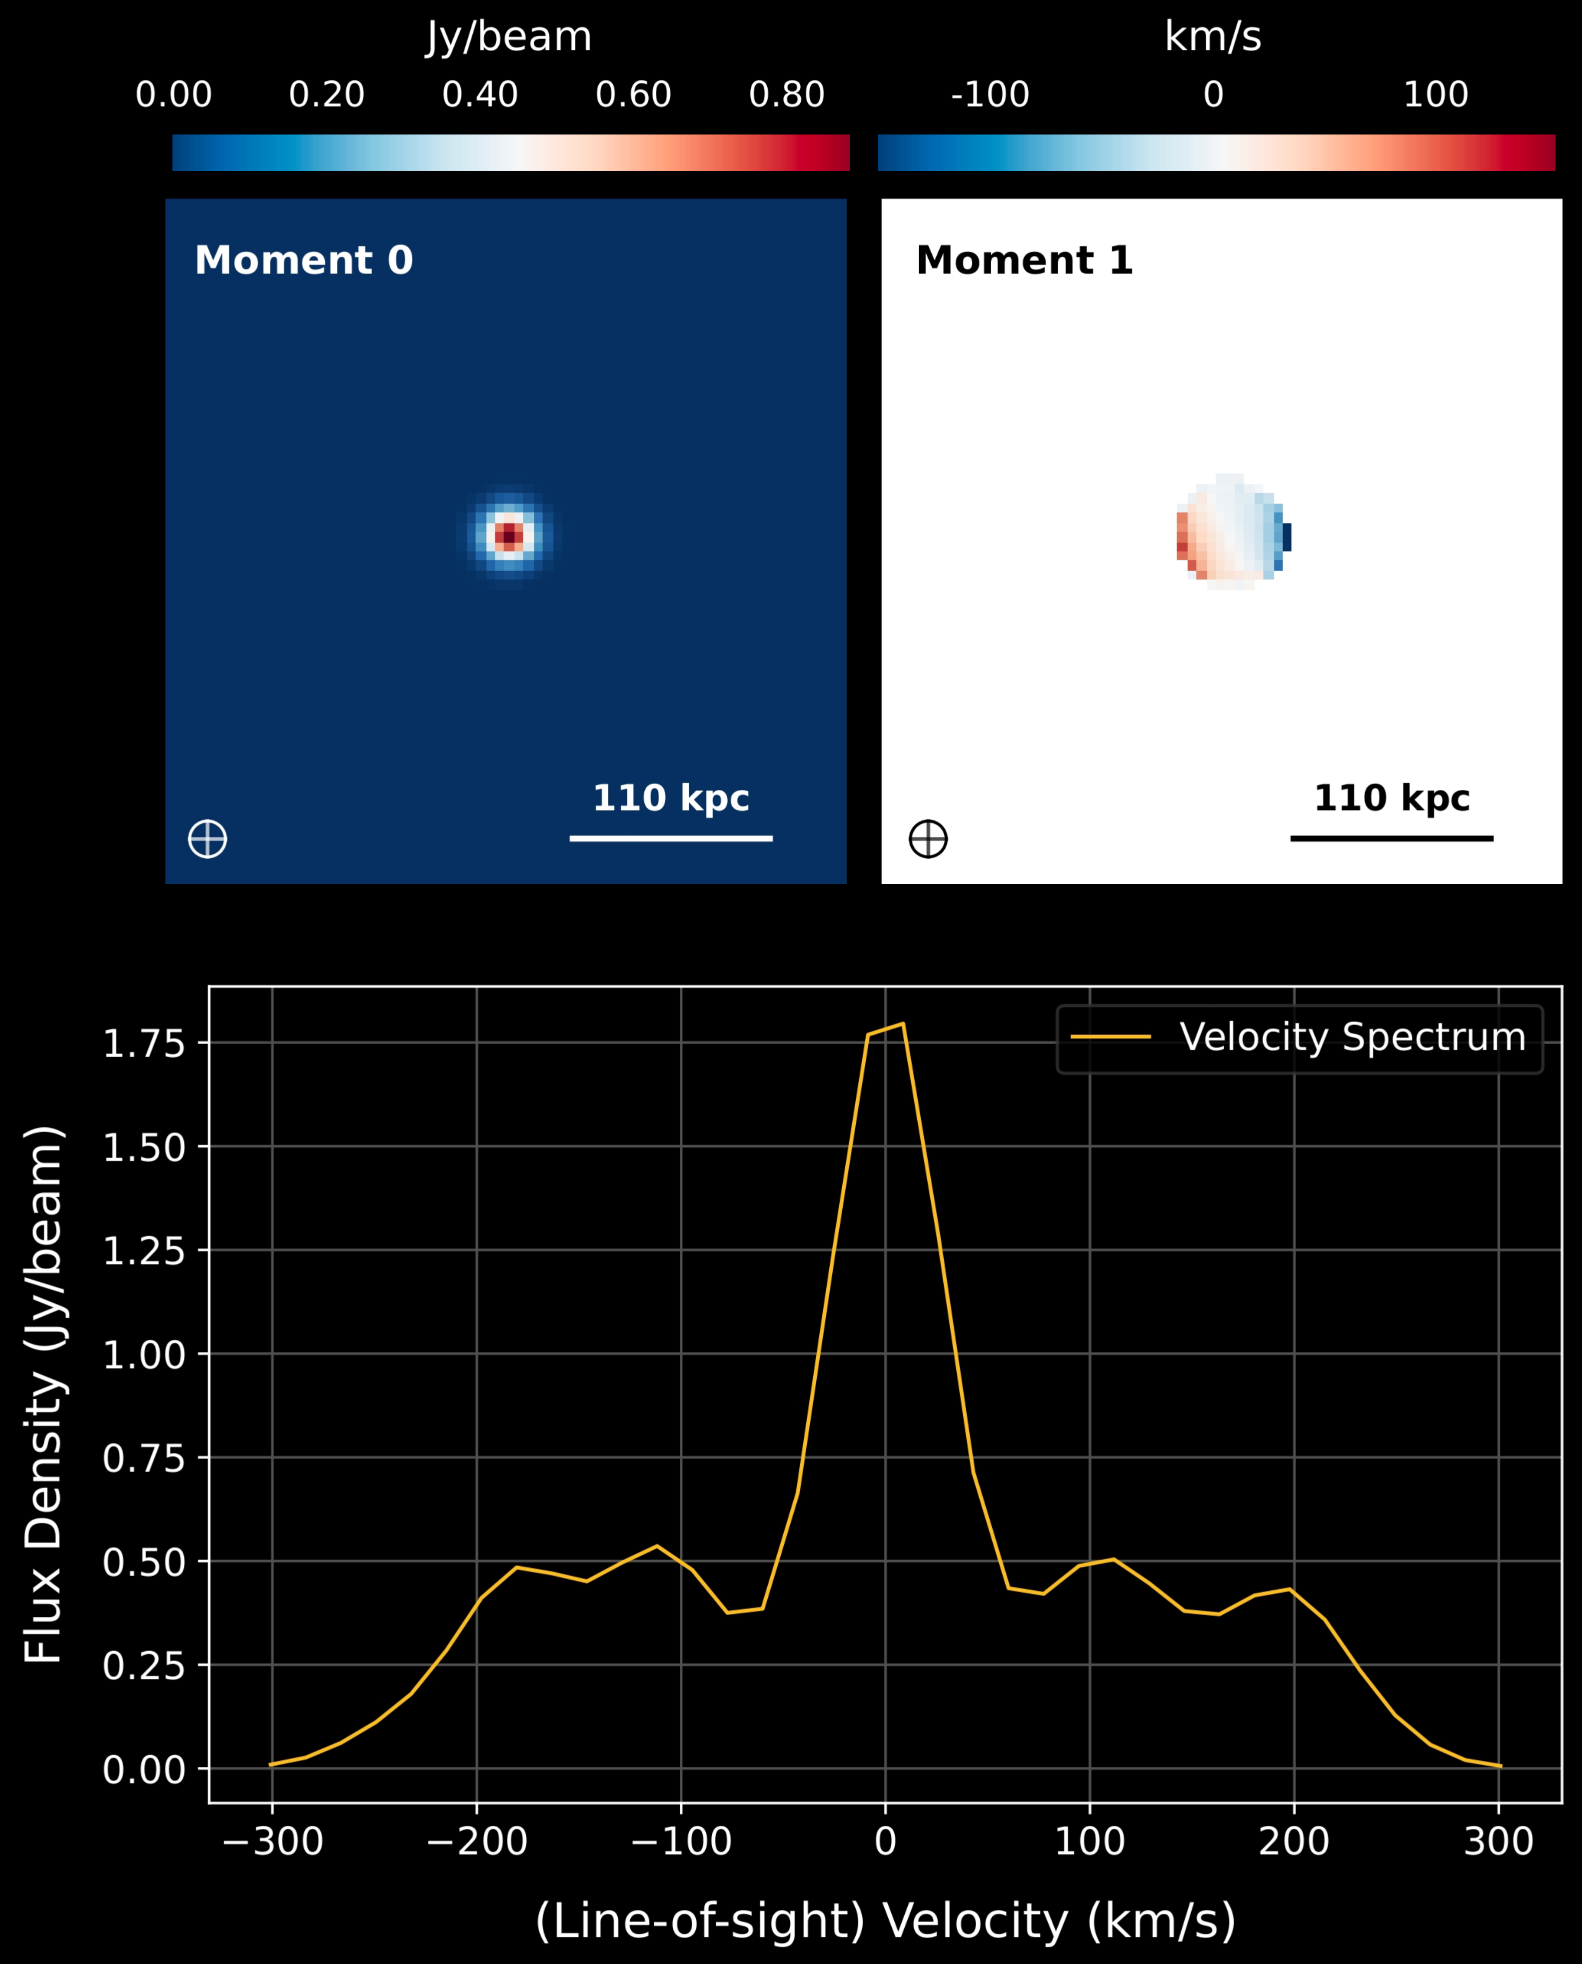

Unresolved source with compact emission with peak at systemic velocity and improperly-resolved kinematics

Constructing toy data of rotating galaxies

Spatially correlated **Gaussian noise** (convolved with the beam) is overlaid onto the cube

Focusing on very high noise regime!

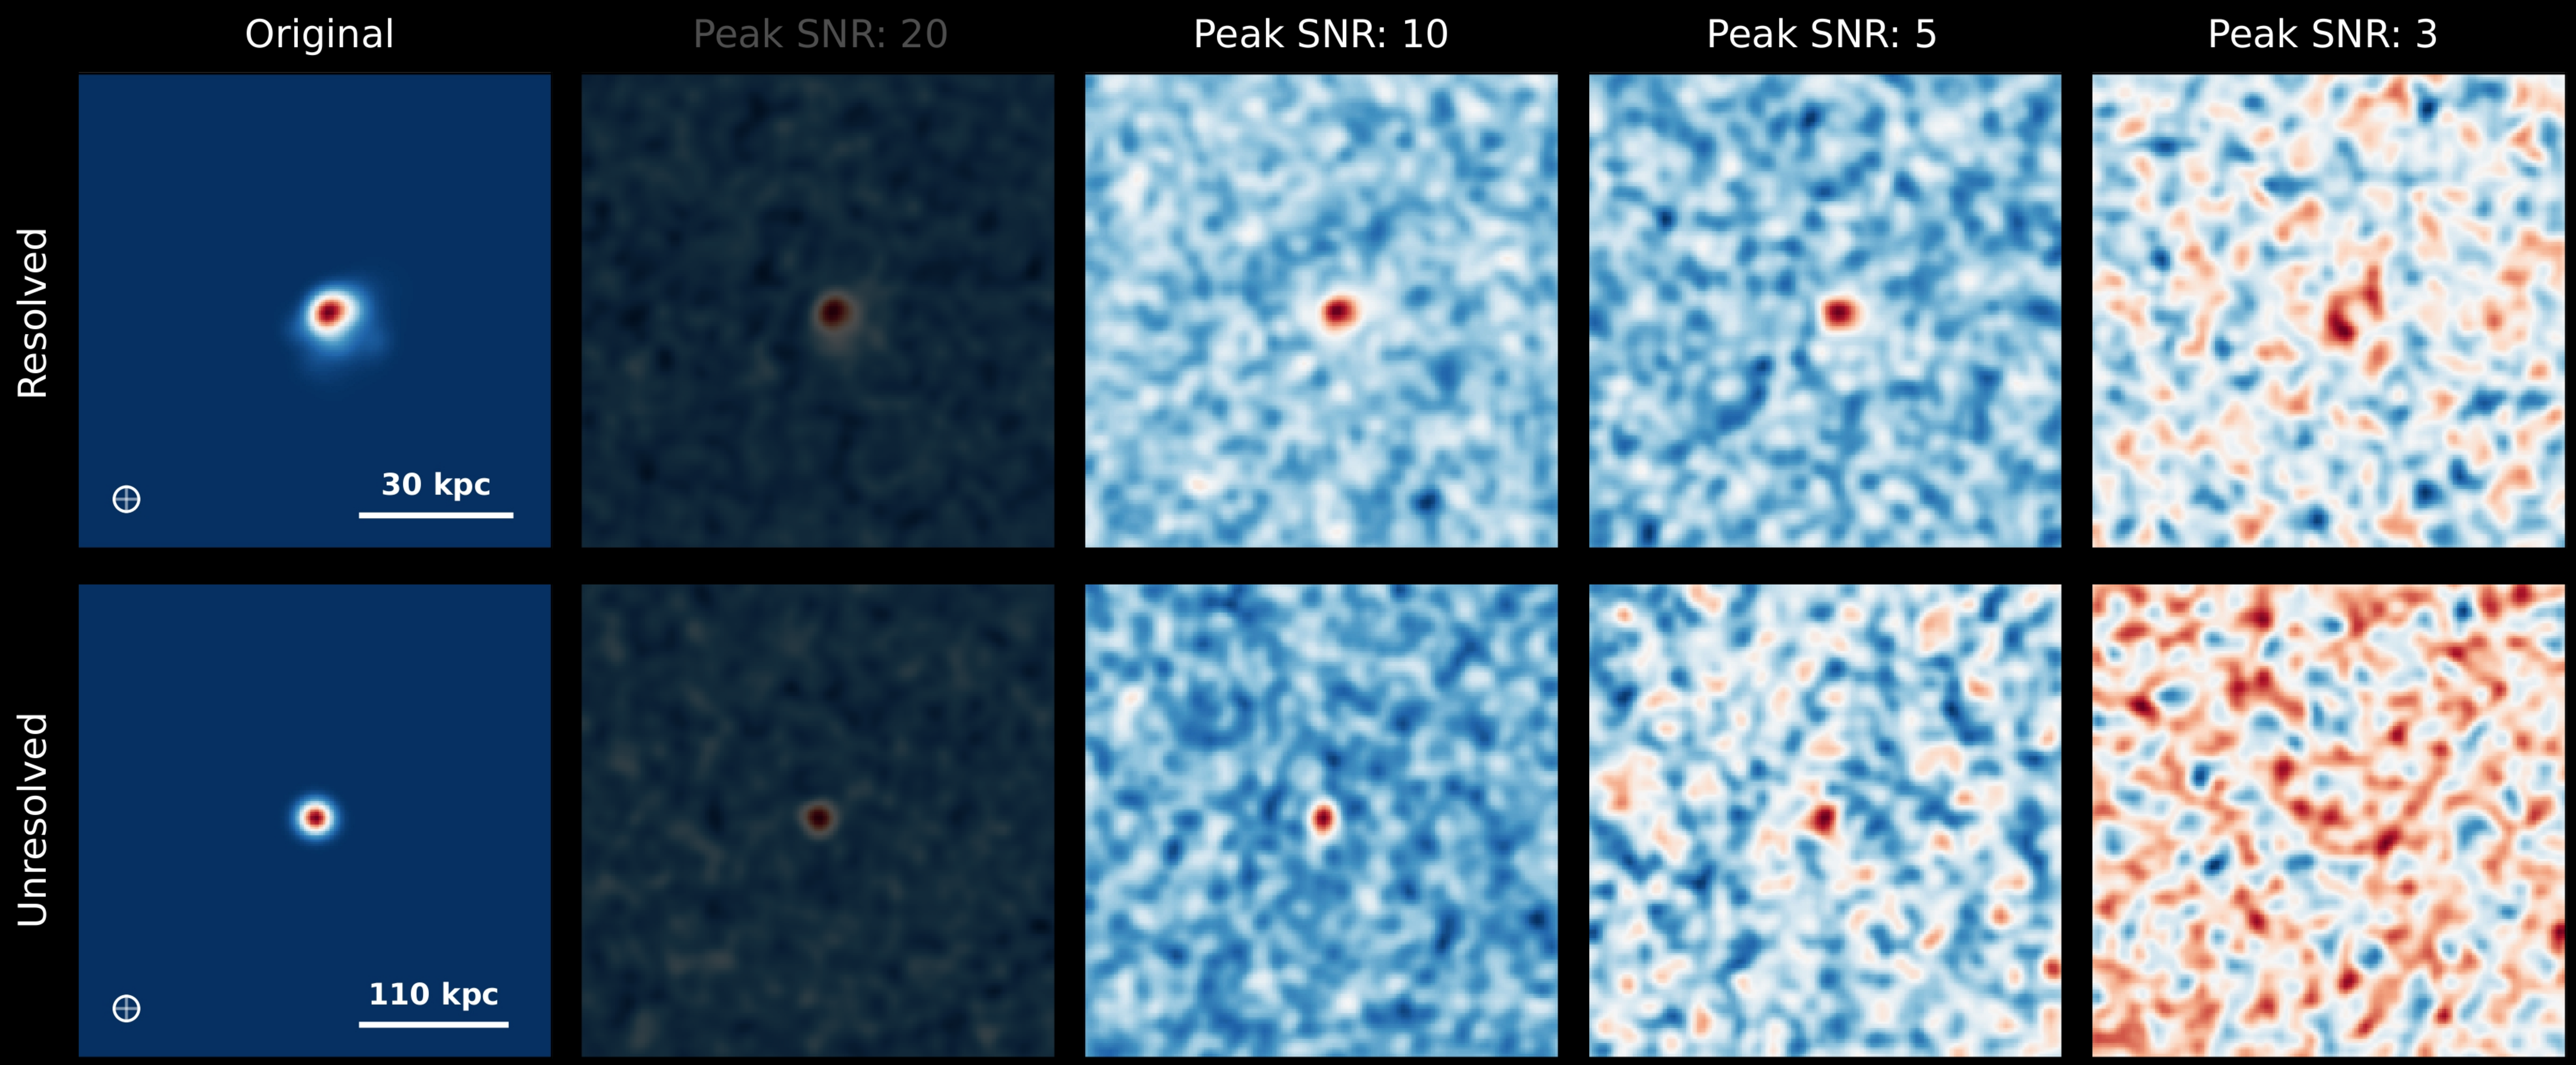

Datasets

- Toy models of rotating galaxies
- Mock IFU data from cosmological simulations
- Observational data (ALMA)

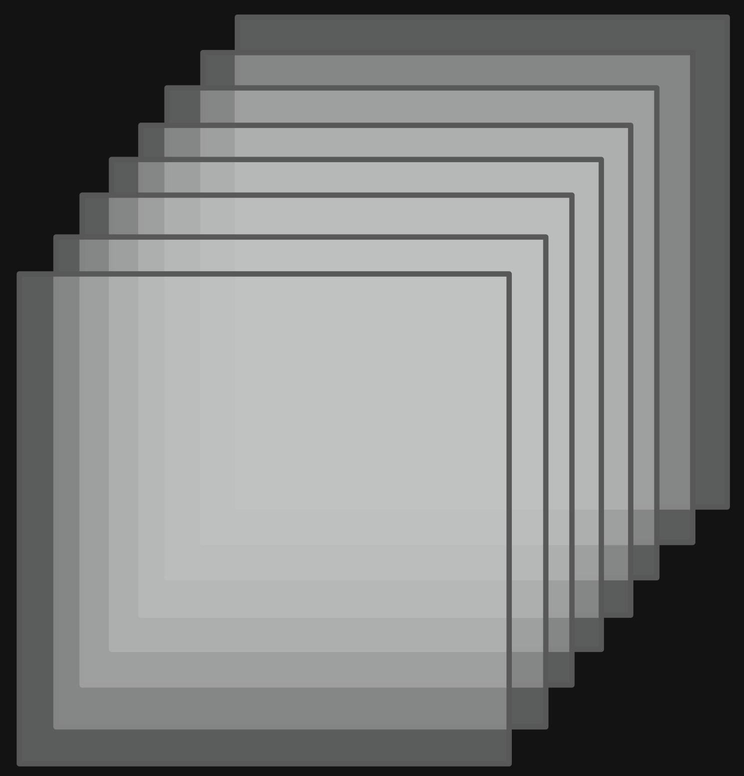

Spectral cubes

multiple spectral observation of the same spatial area, where each *(x,y)* spatial point corresponds to a spectrum

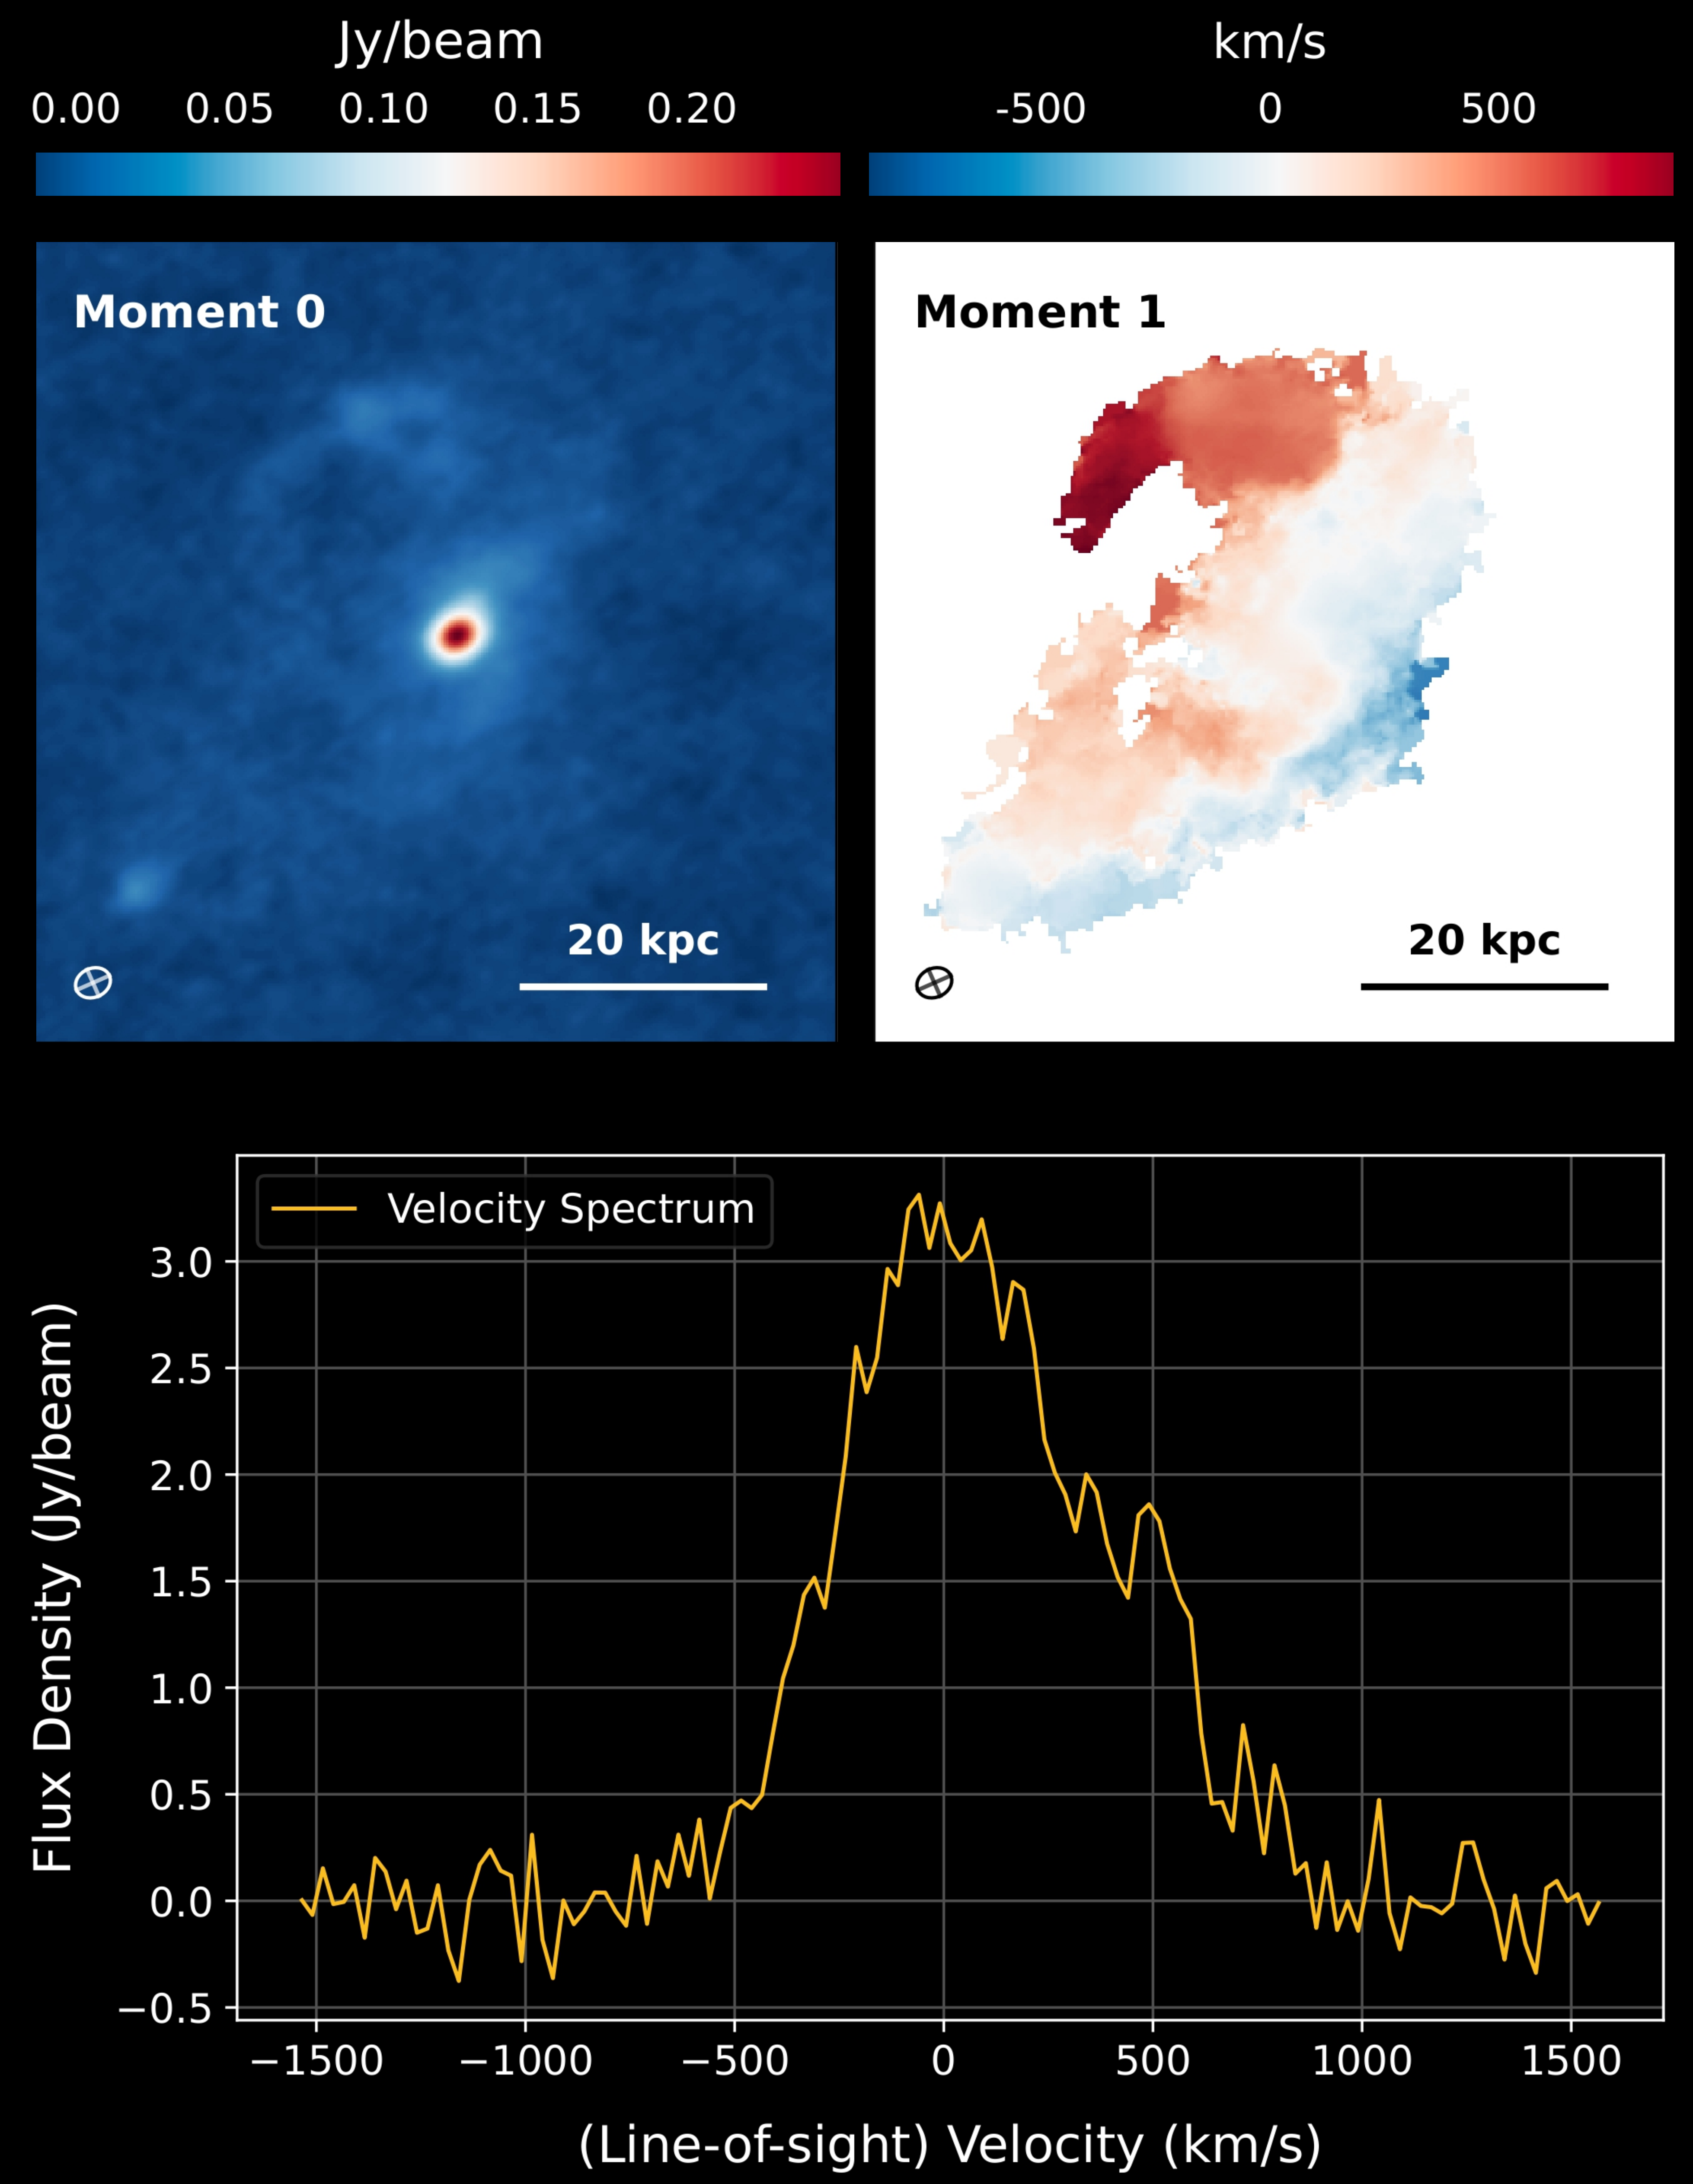

Observational Reference:

W2246–0526

1. **W2246–0526 Overview:** An exceptionally luminous, hot dust-obscured galaxy (Hot DOG) at $z \approx 4.6$, powered by a buried AGN with $L_{\rm bol} \sim 3.5 \times 10^{14} L_\odot$ (Díaz-Santos et al. 2021; Tsai et al. 2015).
2. **Highly Turbulent ISM:** ALMA [C II] maps reveal extremely broad ($\mathrm{FWHM} \geq 500\,\mathrm{km\,s^{-1}}$) and uniform line emission over ~2.5 kpc, indicating a turbulent interstellar medium without a settled disk (Díaz-Santos et al. 2015).
3. **Merger and Outflows:** Deep ALMA imaging shows W2246 interacting with at least three companion galaxies through dusty tidal bridges, undergoing a multiple-merger event and driving large-scale outflows (Díaz-Santos et al. 2021; Fernández Aranda et al. 2025).
4. **Challenges in Observations:** Compact, dust-enshrouded starburst regions and AGN make the [C II] line faint, extended, and embedded in strong dust continuum, compounded by low flux, beam smearing, and high noise at high redshift.
5. **Need for Denoising:** Advanced denoising techniques are critical to recover faint astrophysical signals, improve SNR, and enable reliable flux, kinematics, and morphological analysis of both the host and its extended structures.

Datasets

- Toy models of rotating galaxies
- Mock IFU data from cosmological simulations
- Observational data (ALMA)

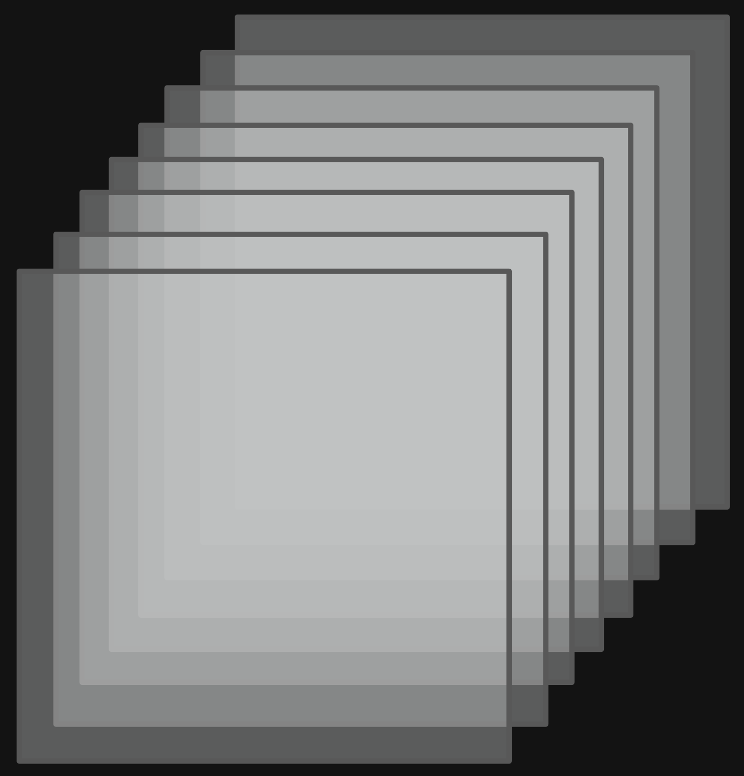

Spectral cubes
multiple spectral observation of the same spatial area, where each *(x,y)* spatial point corresponds to a spectrum

Pre-processing mock IFU cubes from FIRE

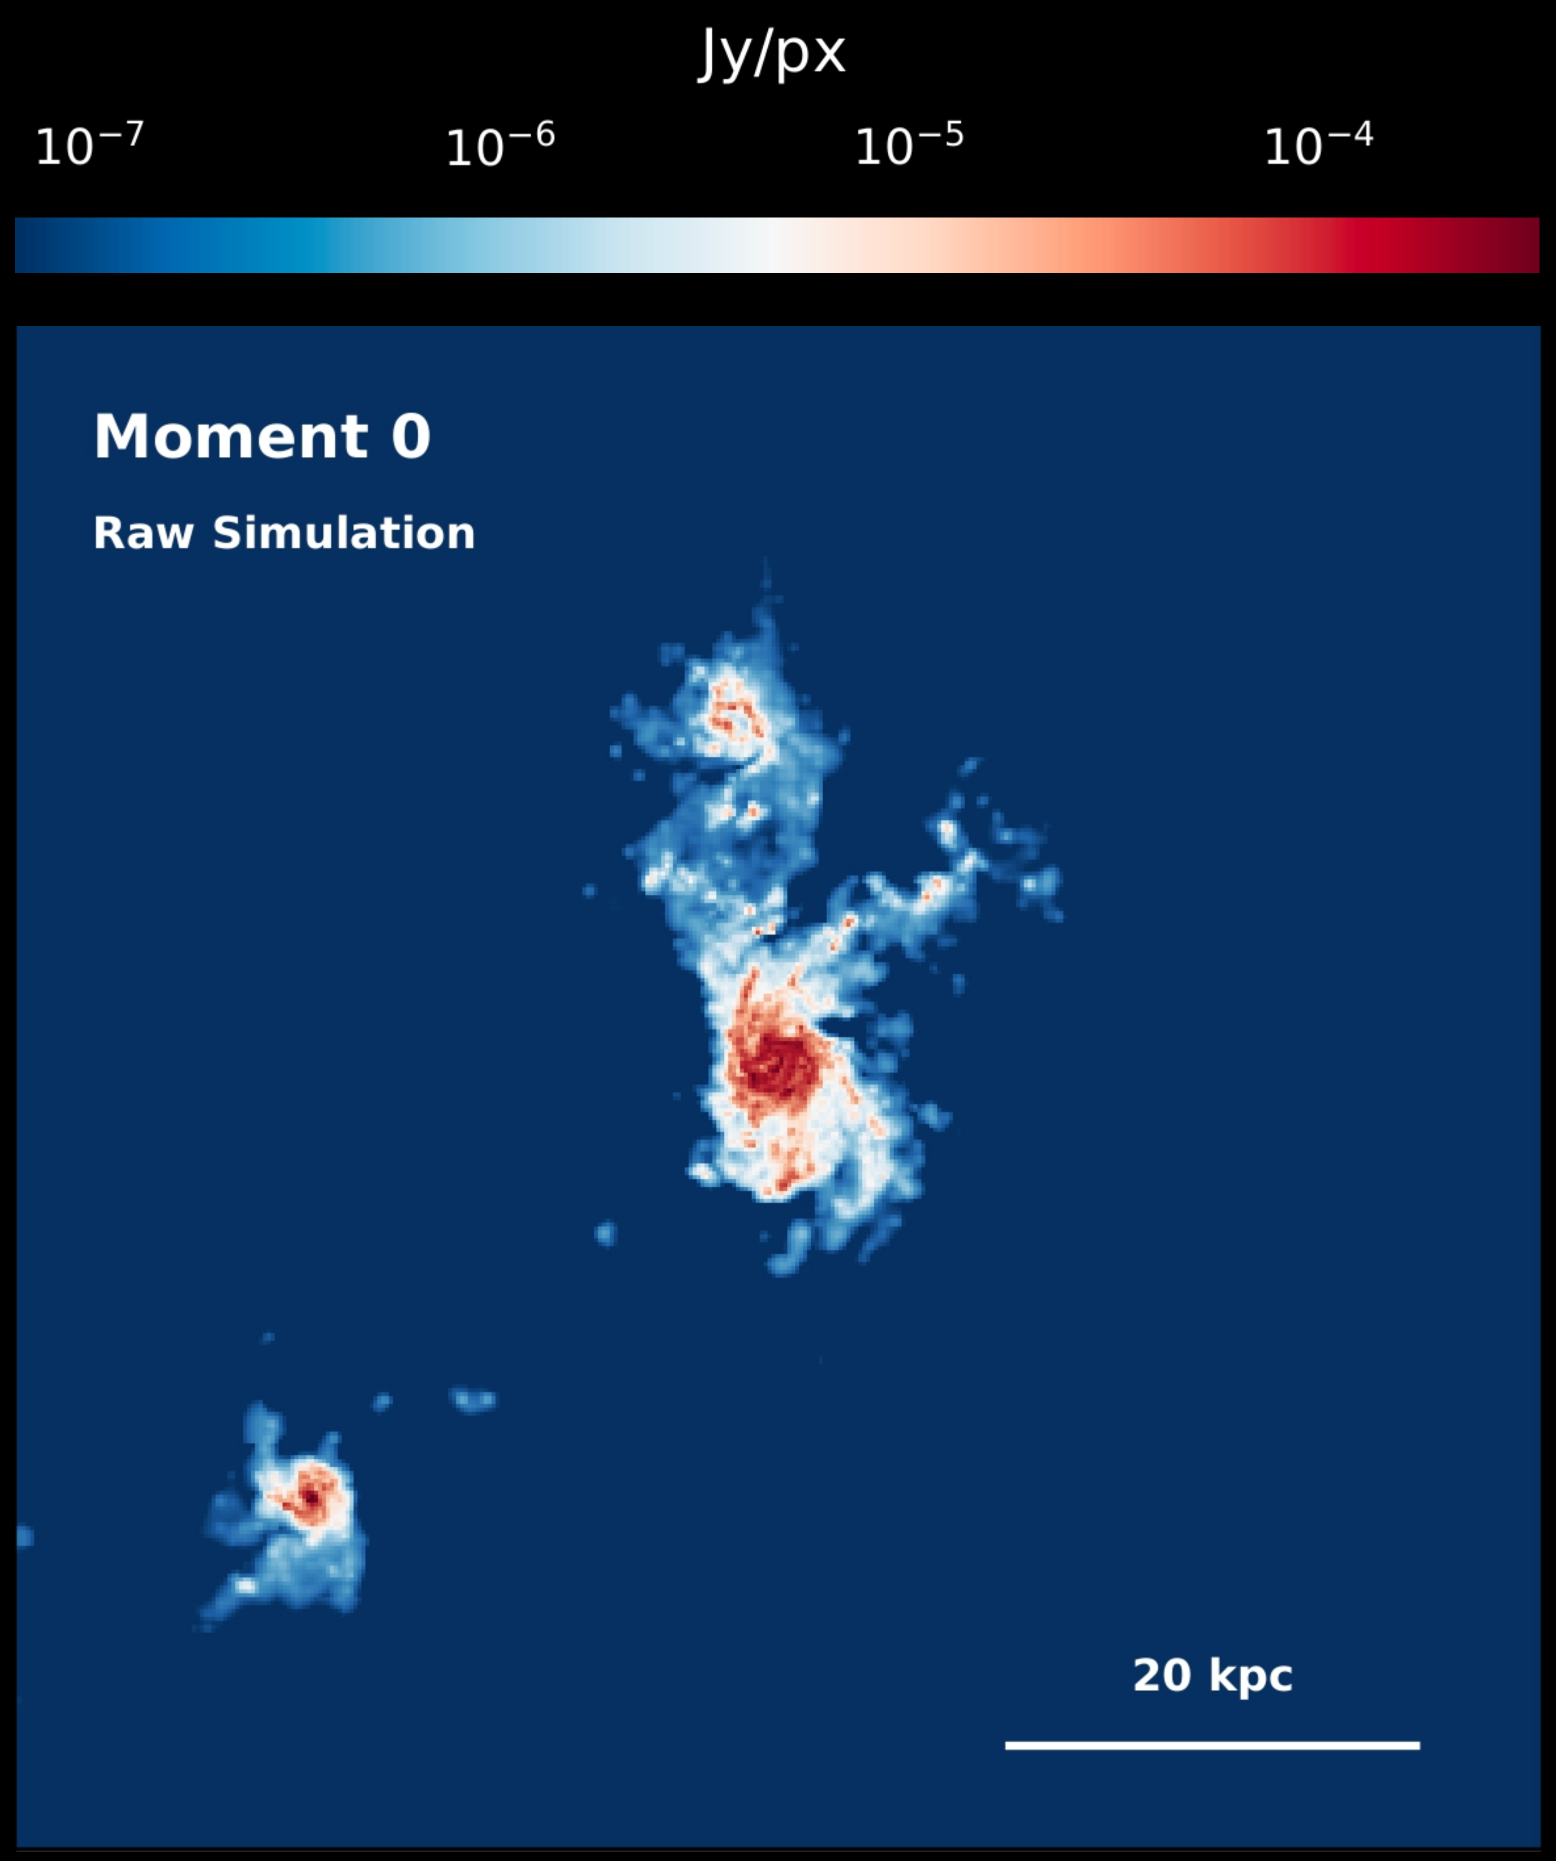

W2246-like system

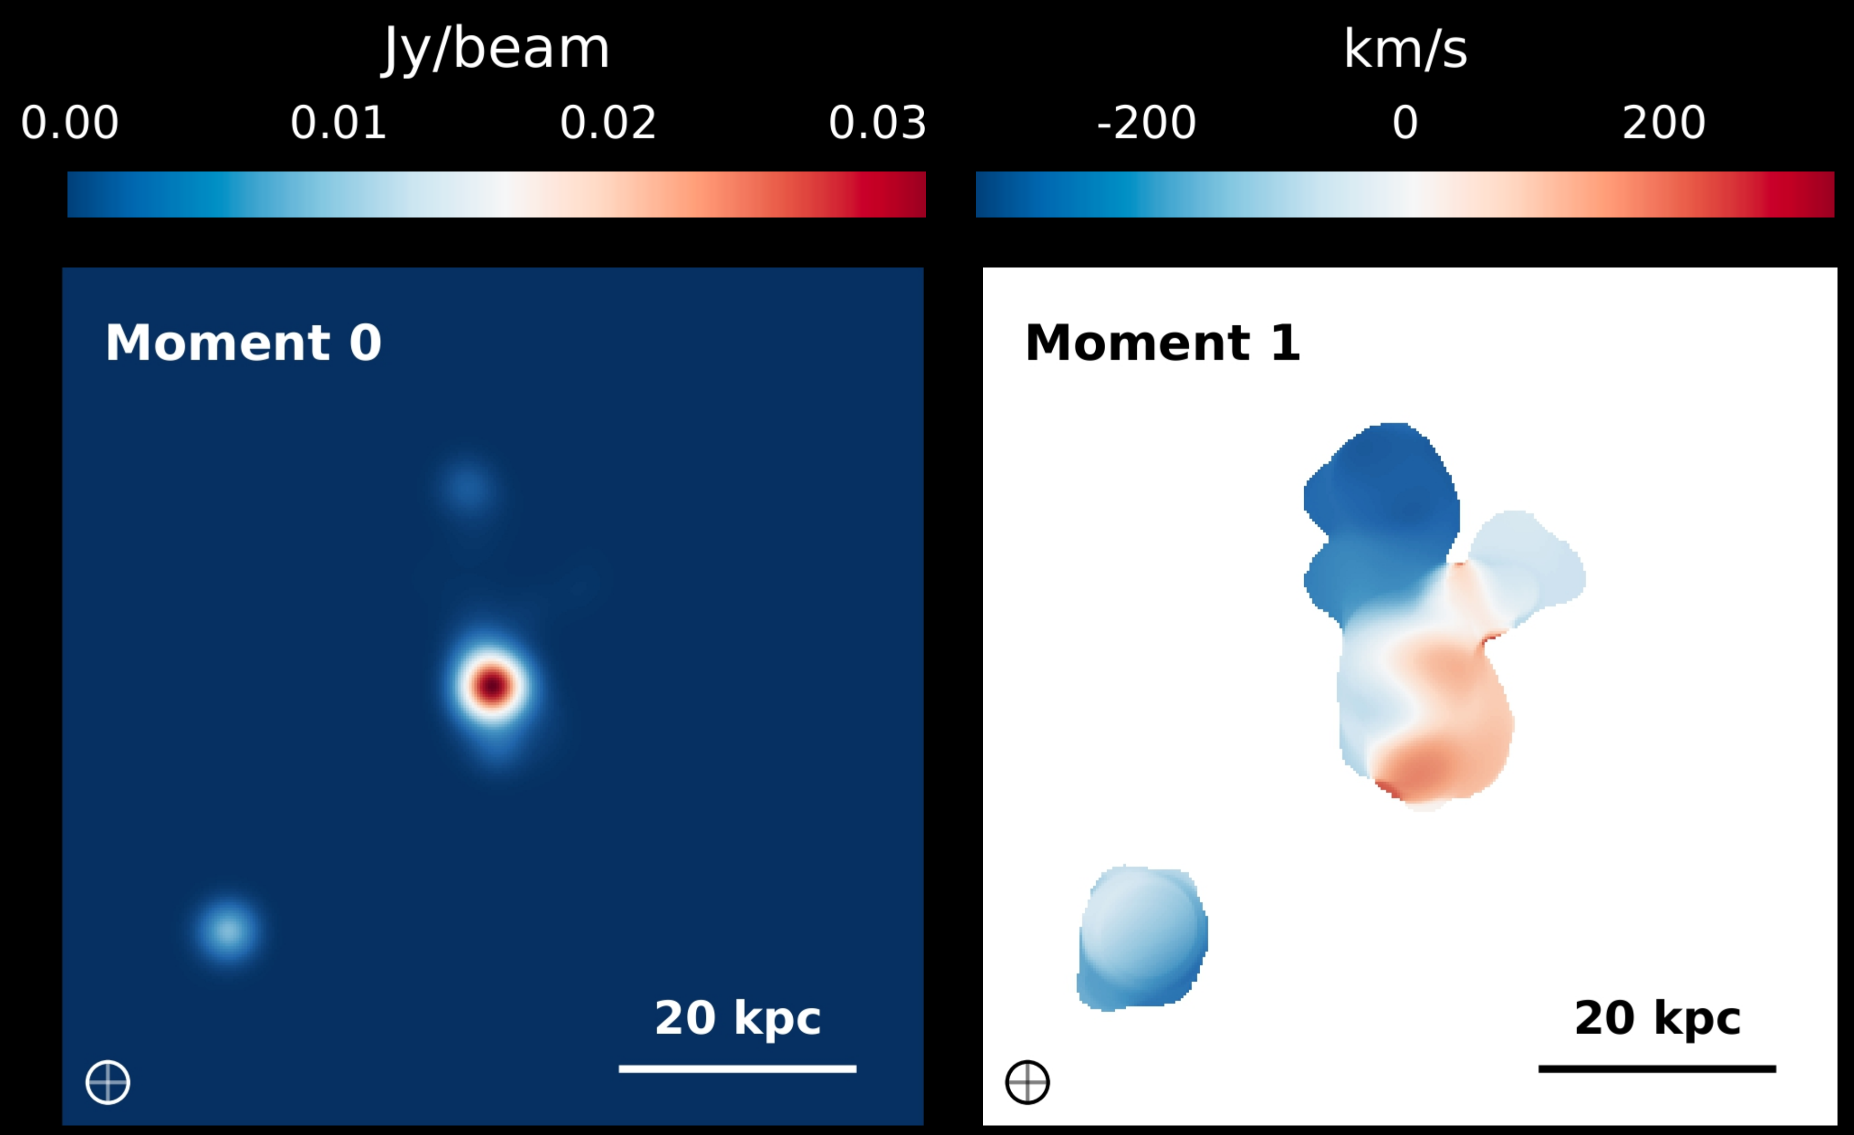

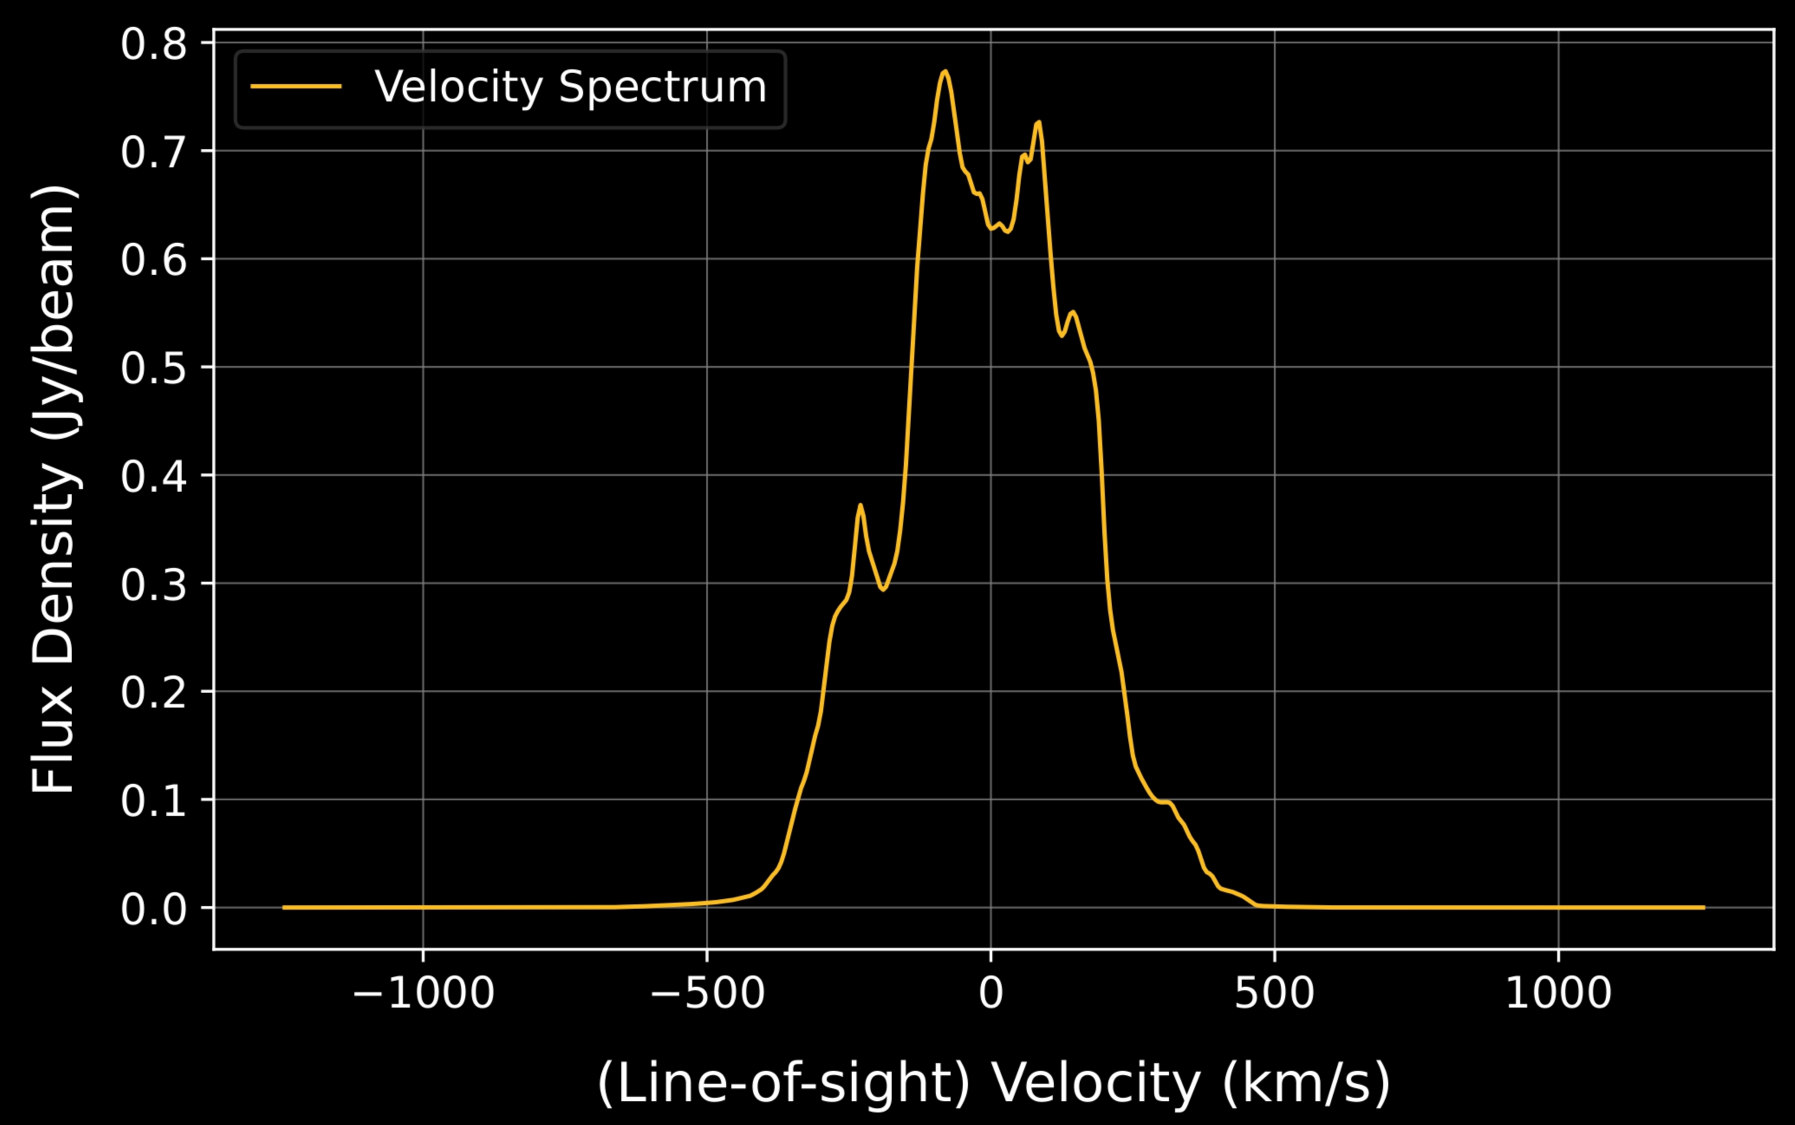

A multiple-merger system was discovered in the FIRE simulations - simulated at redshift 4.5, central galaxy with multiple satellites and visible bridges and streams of gas

Aim: To pre-process the raw, highly resolved mock IFU cube and convolve it with an appropriate beam to make the cube as close as possible to real observations (ALMA)

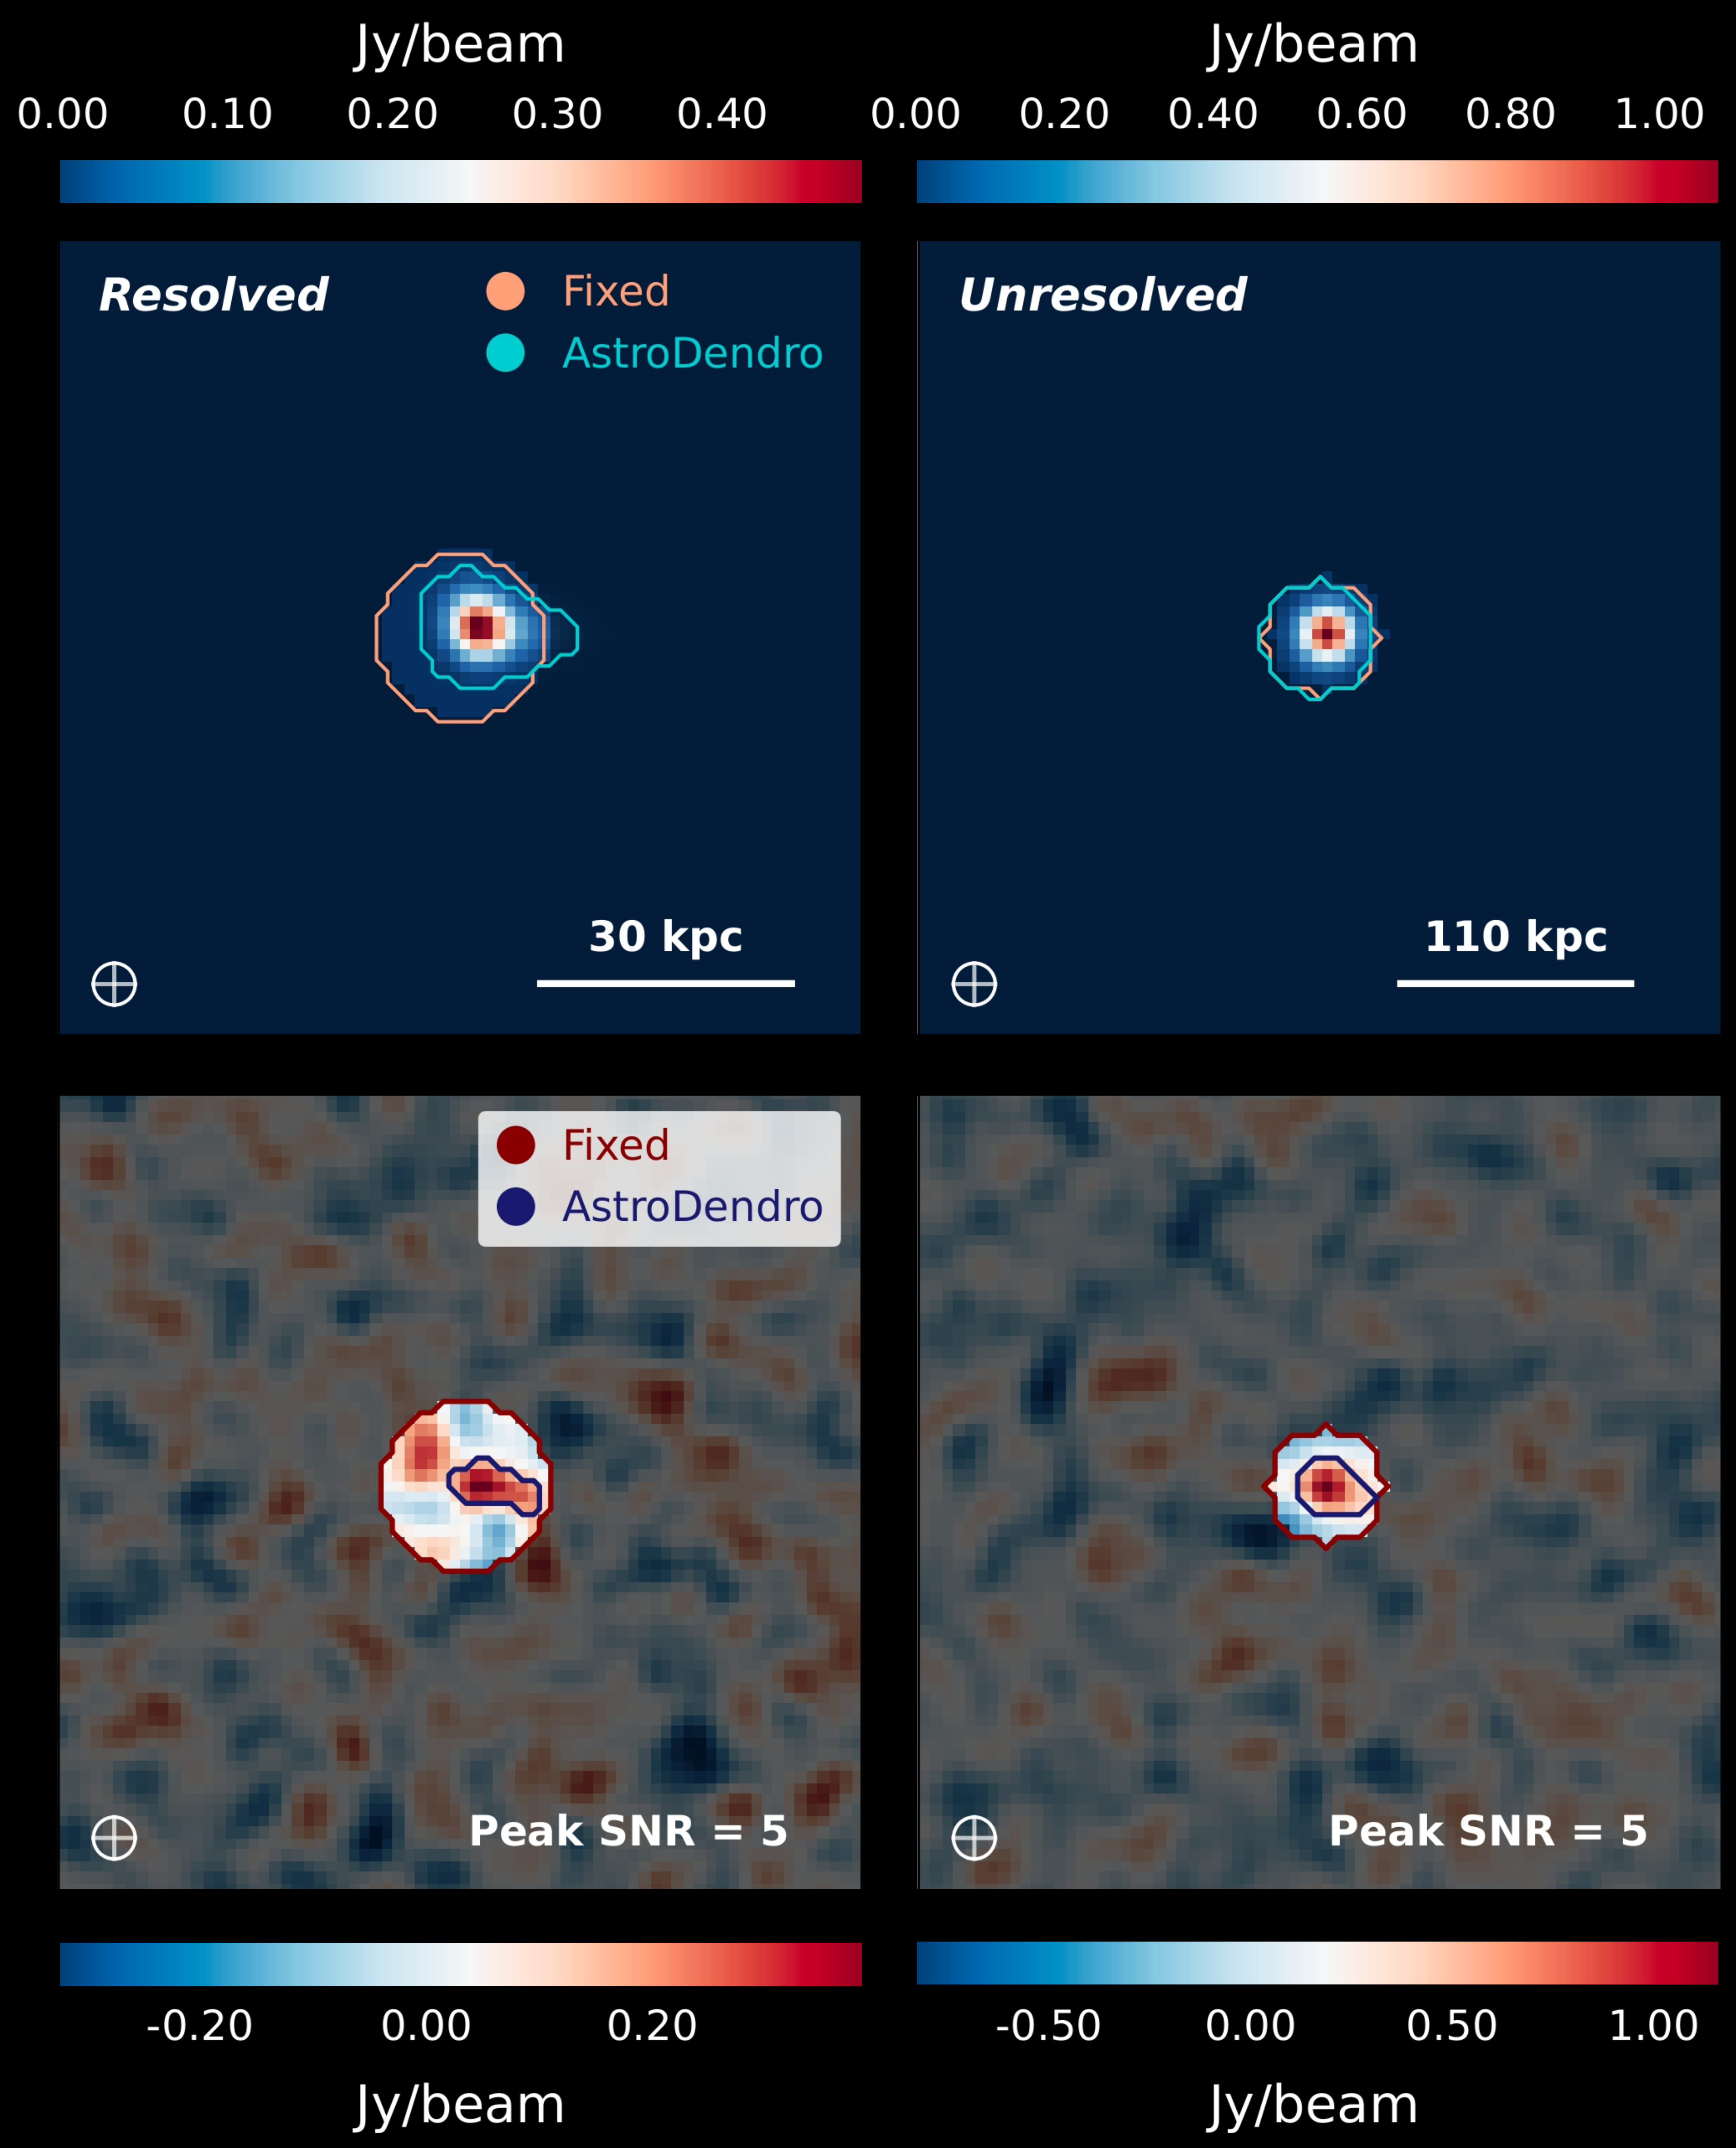

Identifying emission regions

Fixed masks focusing on the central galaxies are constructed as a function of the effective diameter of the galaxy and the beam FWHM (both in pixels)

$$D_{\mathrm{ap}} = \begin{cases} 2 \times D_e & \text{if } D_e > \mathrm{FWHM}_{\mathrm{beam}} \\ 2 \times \mathrm{FWHM}_{\mathrm{beam}} & \text{if } D_e \leq \mathrm{FWHM}_{\mathrm{beam}} \end{cases}$$

Accurate emission masks constructed using **astrodendro**, which identifies regions with strong emission hierarchically in the whole cube

< However, astrodendro masks are unable to get accurate apertures on data with very high levels of correlated noise, hence we use fixed masks for our comparative analysis

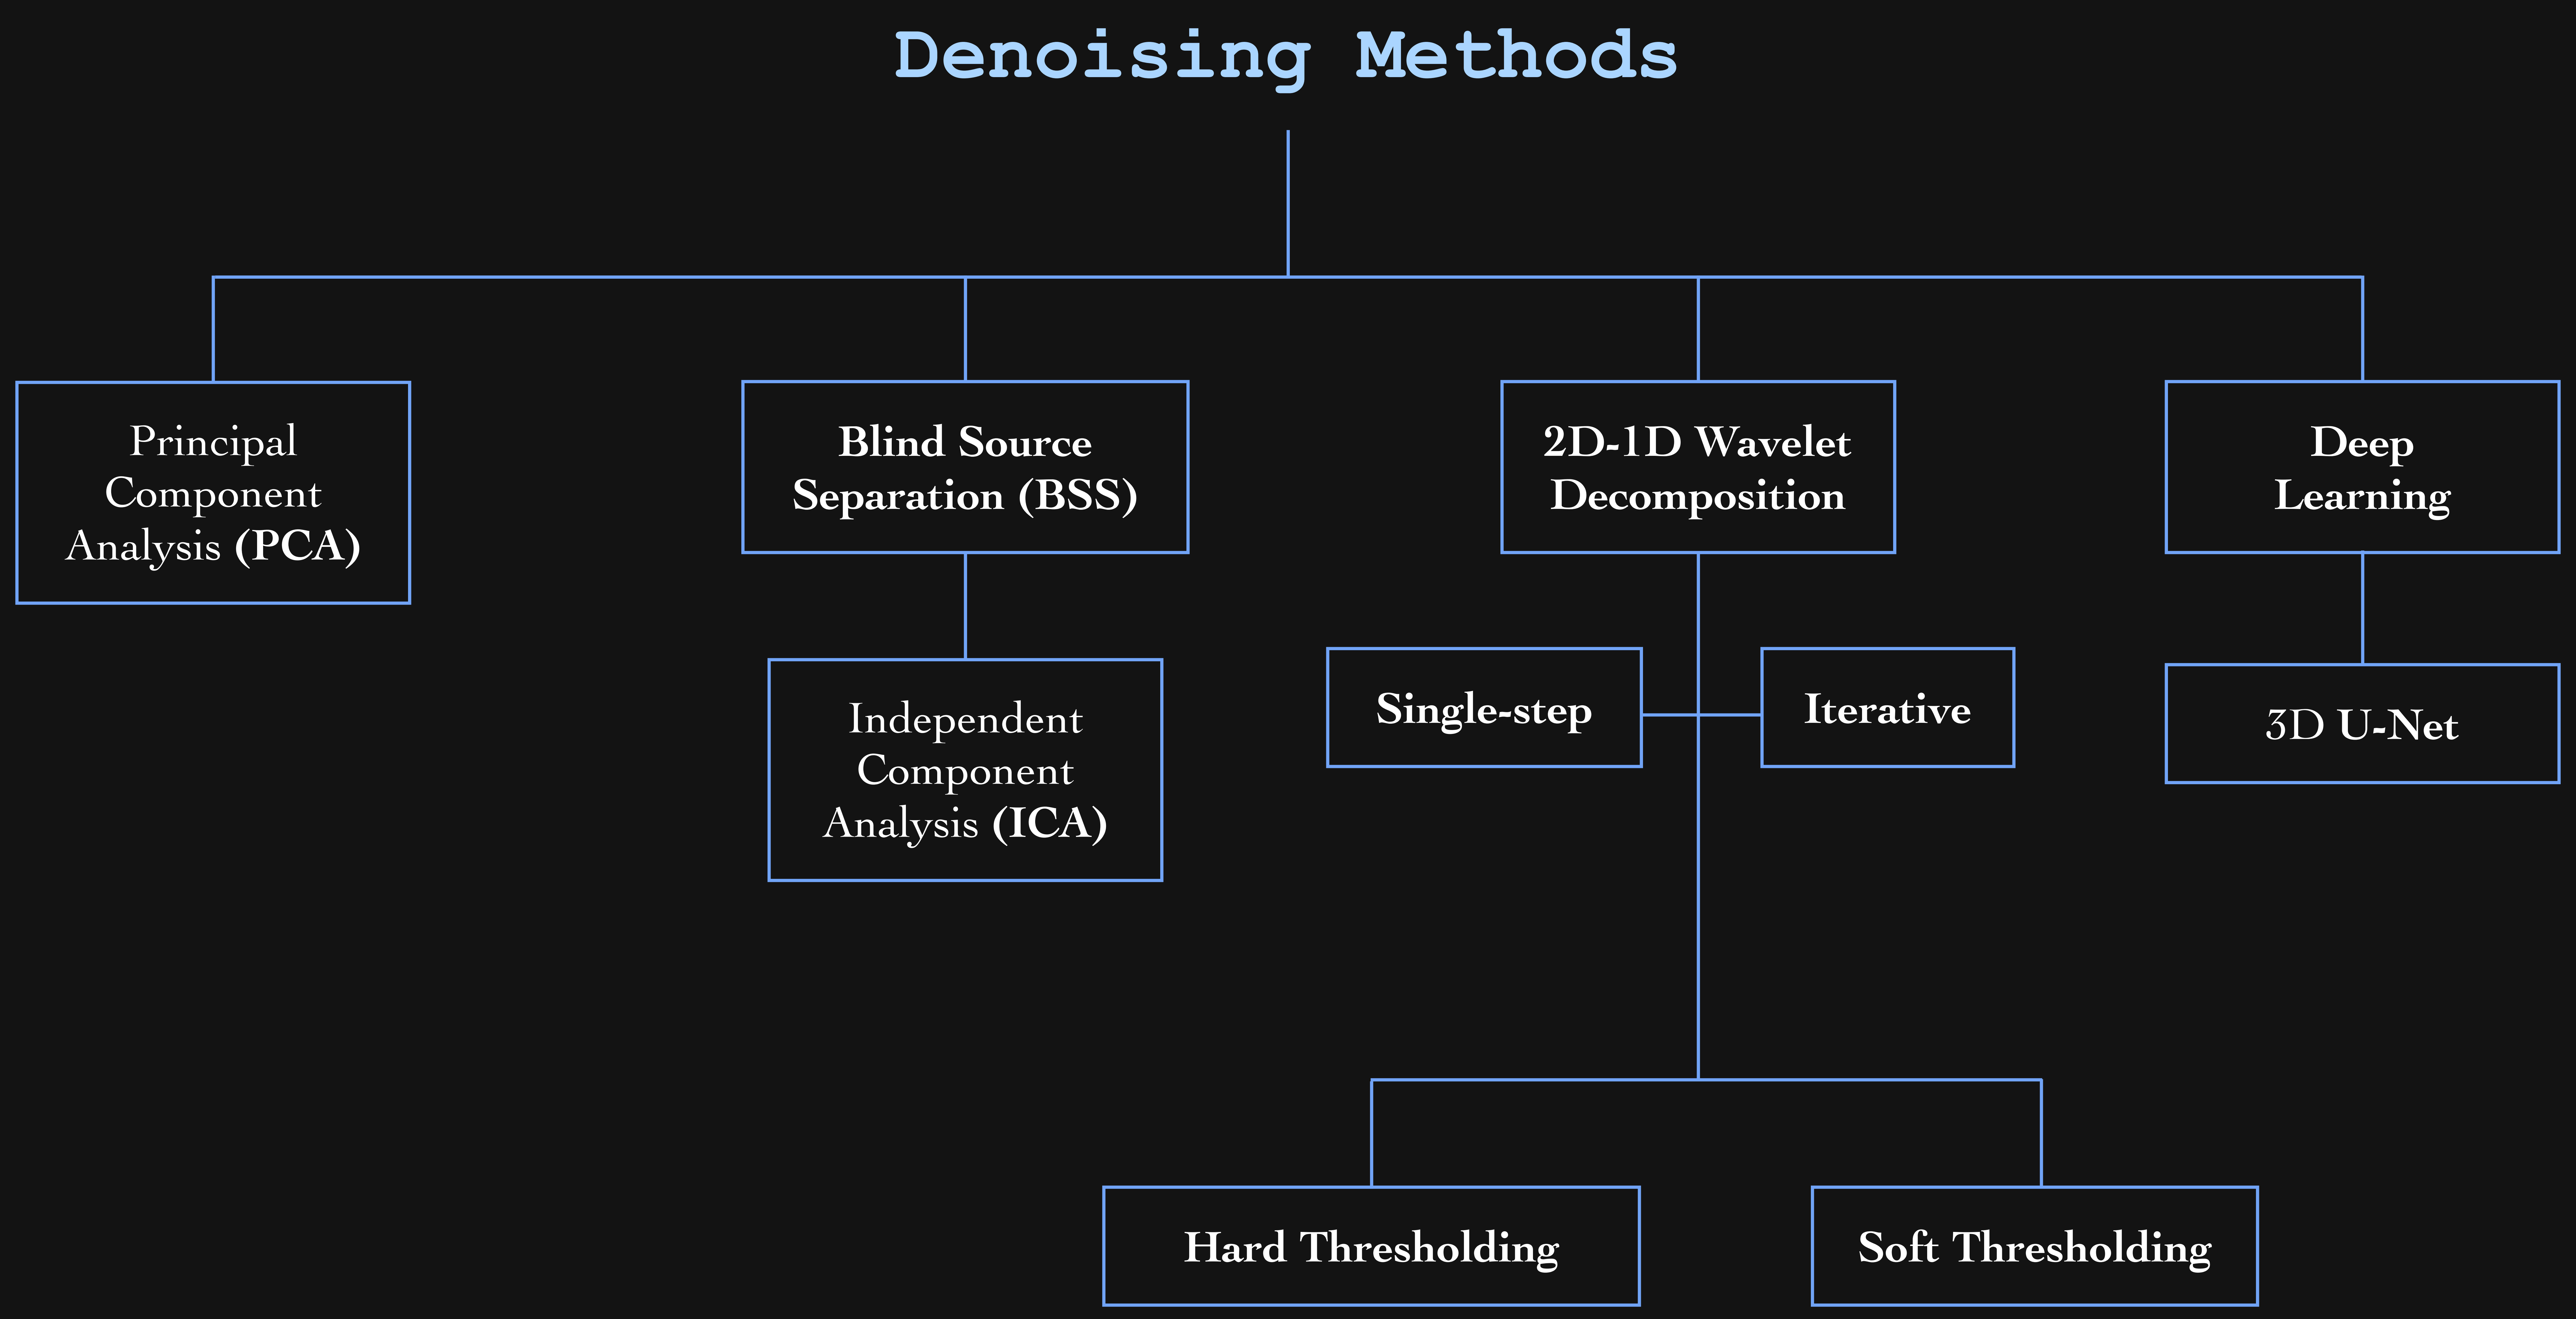
Denoising Methods
Principal Component Analysis (PCA)
Blind Source Separation (BSS)
Independent Component Analysis (ICA)
2D-1D Wavelet Decomposition
Single-step
Iterative
Hard Thresholding
Soft Thresholding
Deep Learning
3D U-Net

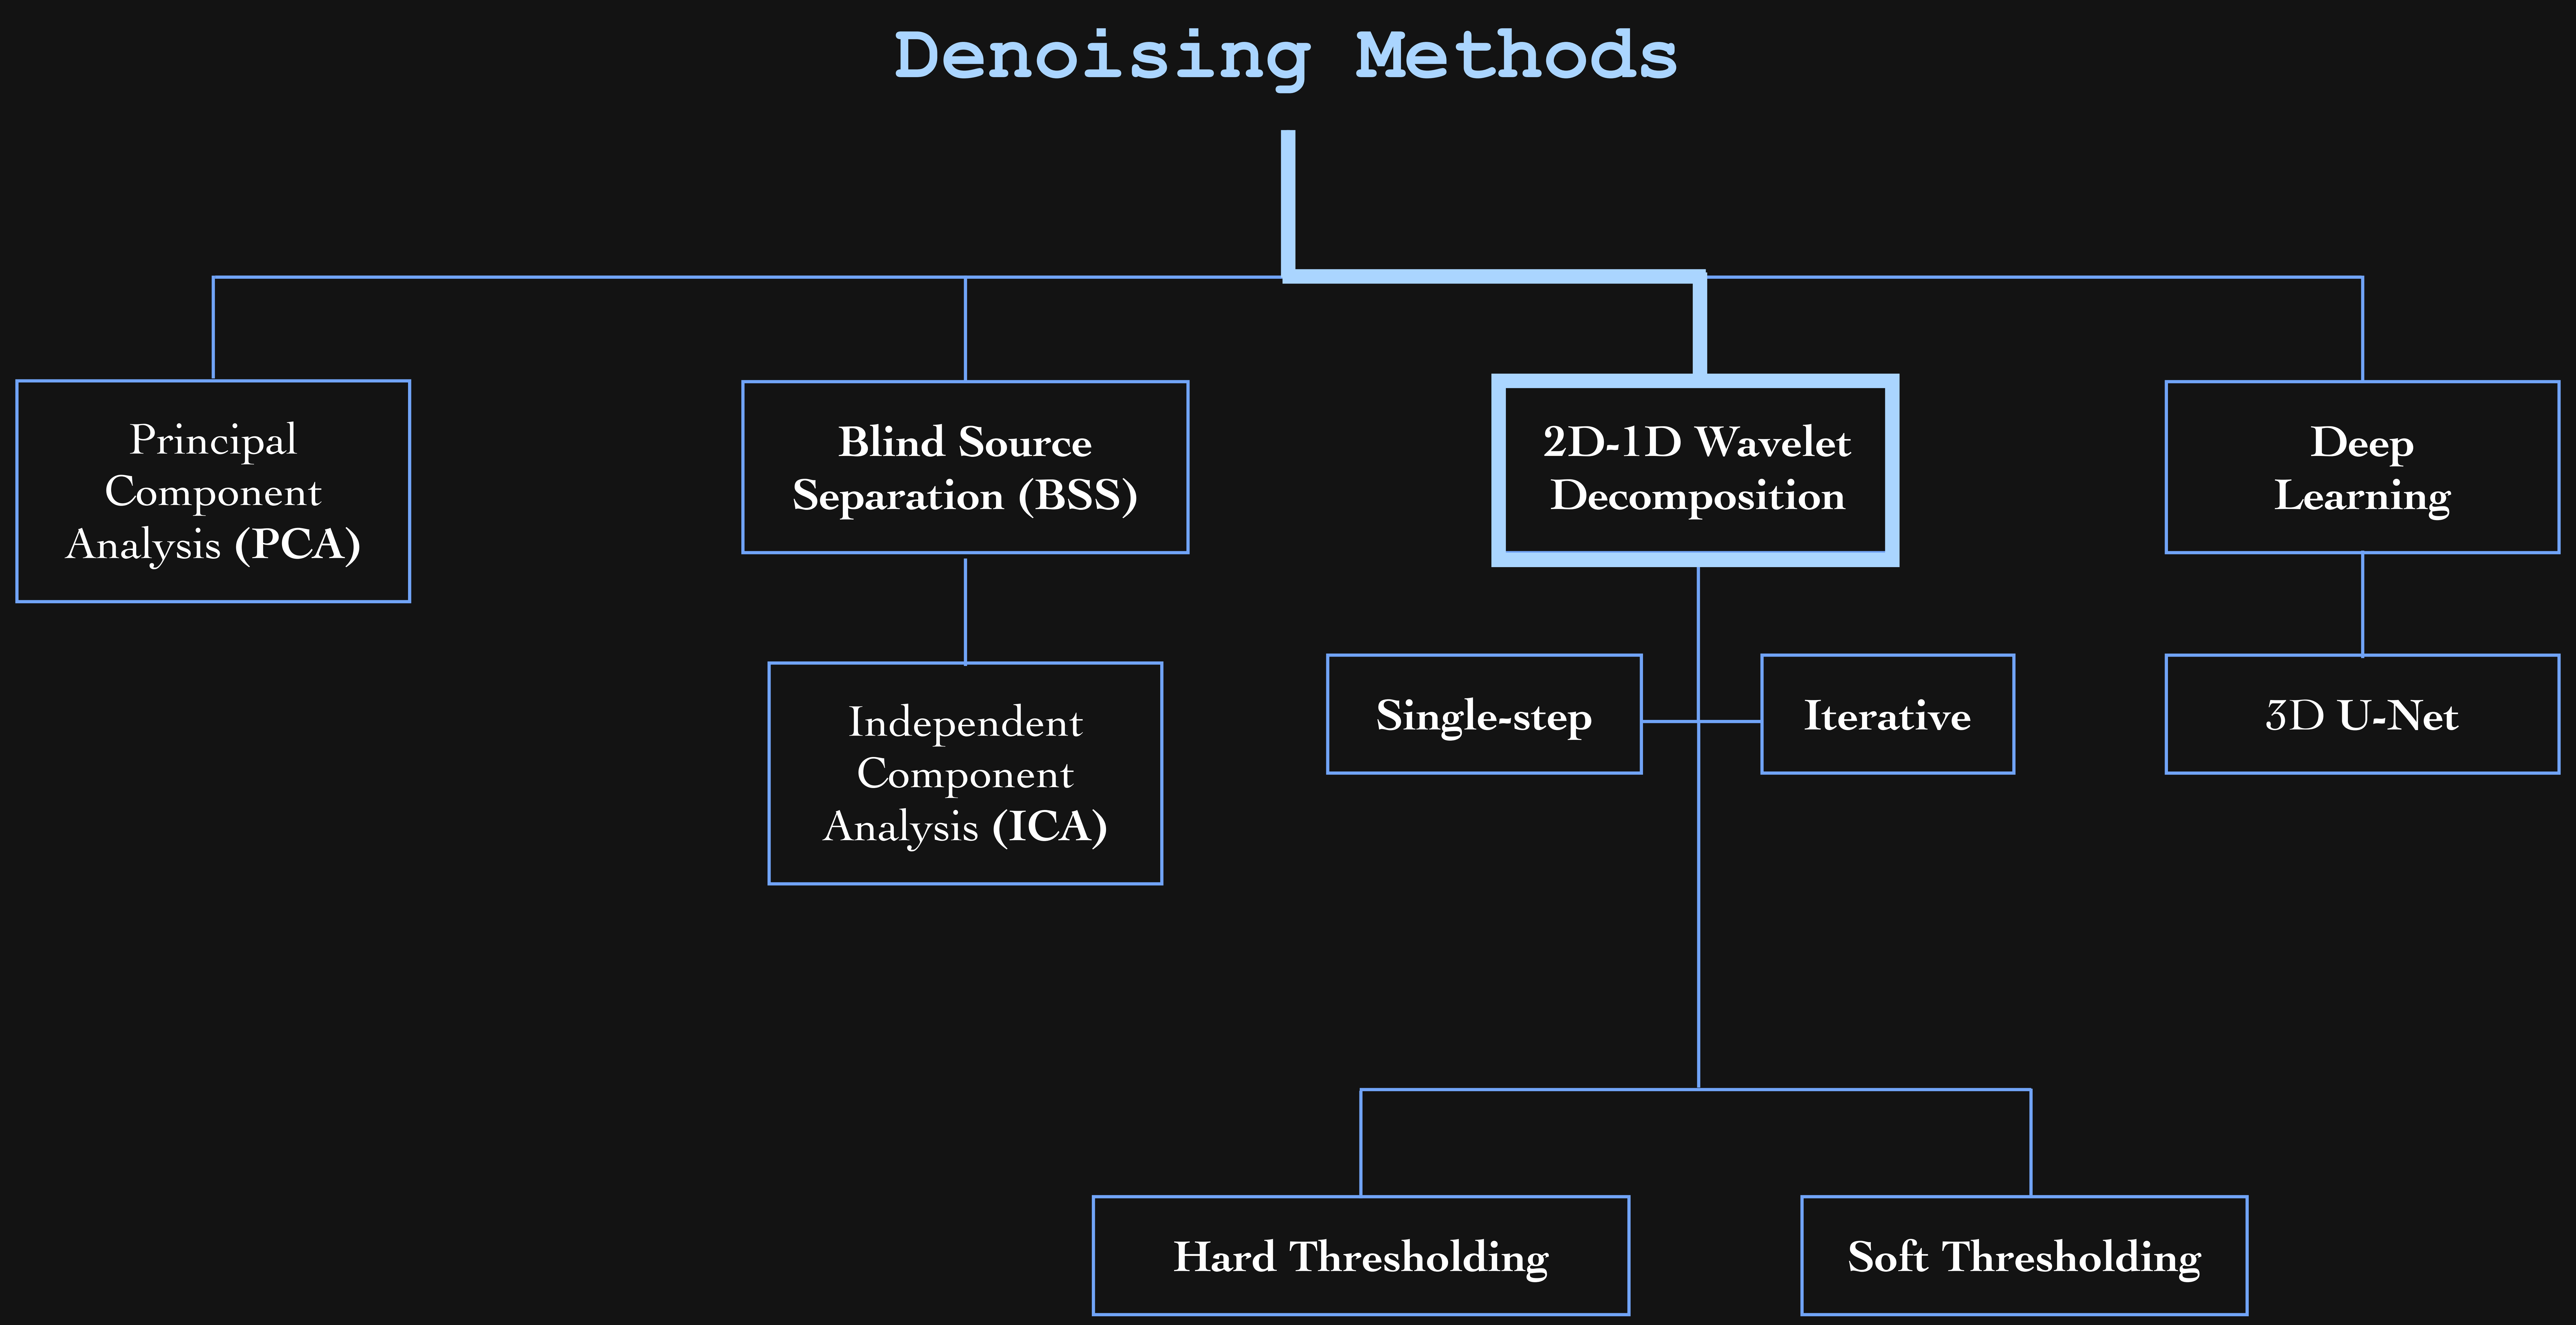
Denoising Methods
Principal Component Analysis (PCA)
Blind Source Separation (BSS)
2D-1D Wavelet Decomposition
Deep Learning
Independent Component Analysis (ICA)
Single-step
Iterative
3D U-Net
Hard Thresholding
Soft Thresholding

2D-1D wavelet decomposition

Each spectral slice… …is decomposed into n_{2d} scales

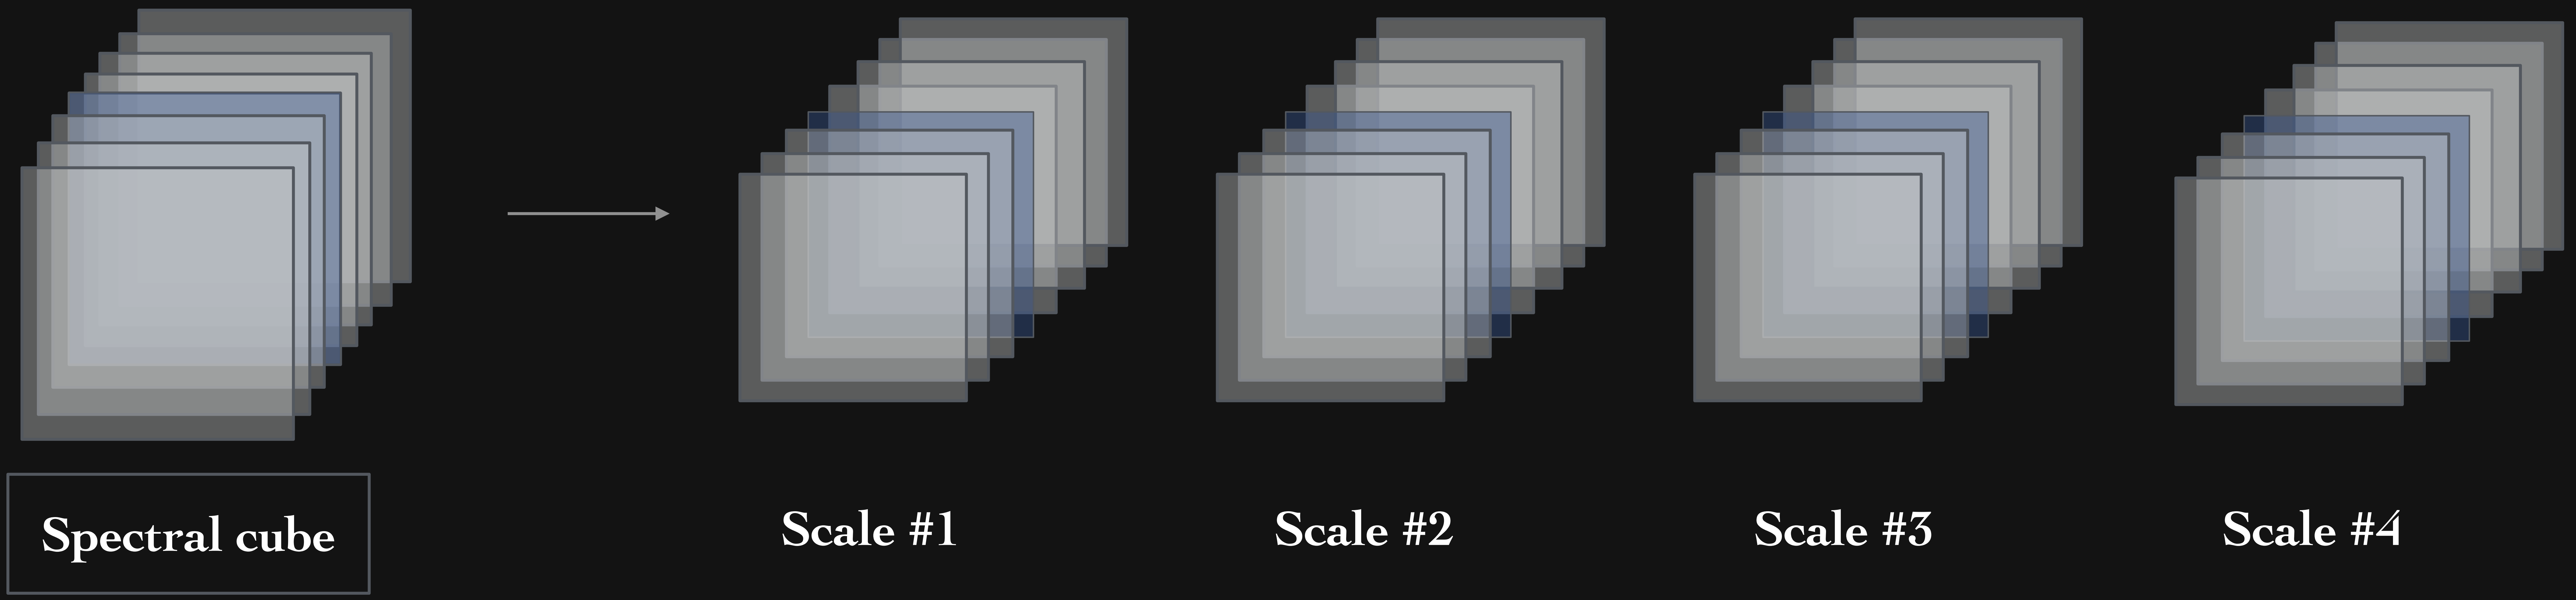

2D decomposition: Starlet transform on each spectral slice : undecimated and non orthogonal (dimensions of each slice are preserved)

2D-1D wavelet decomposition

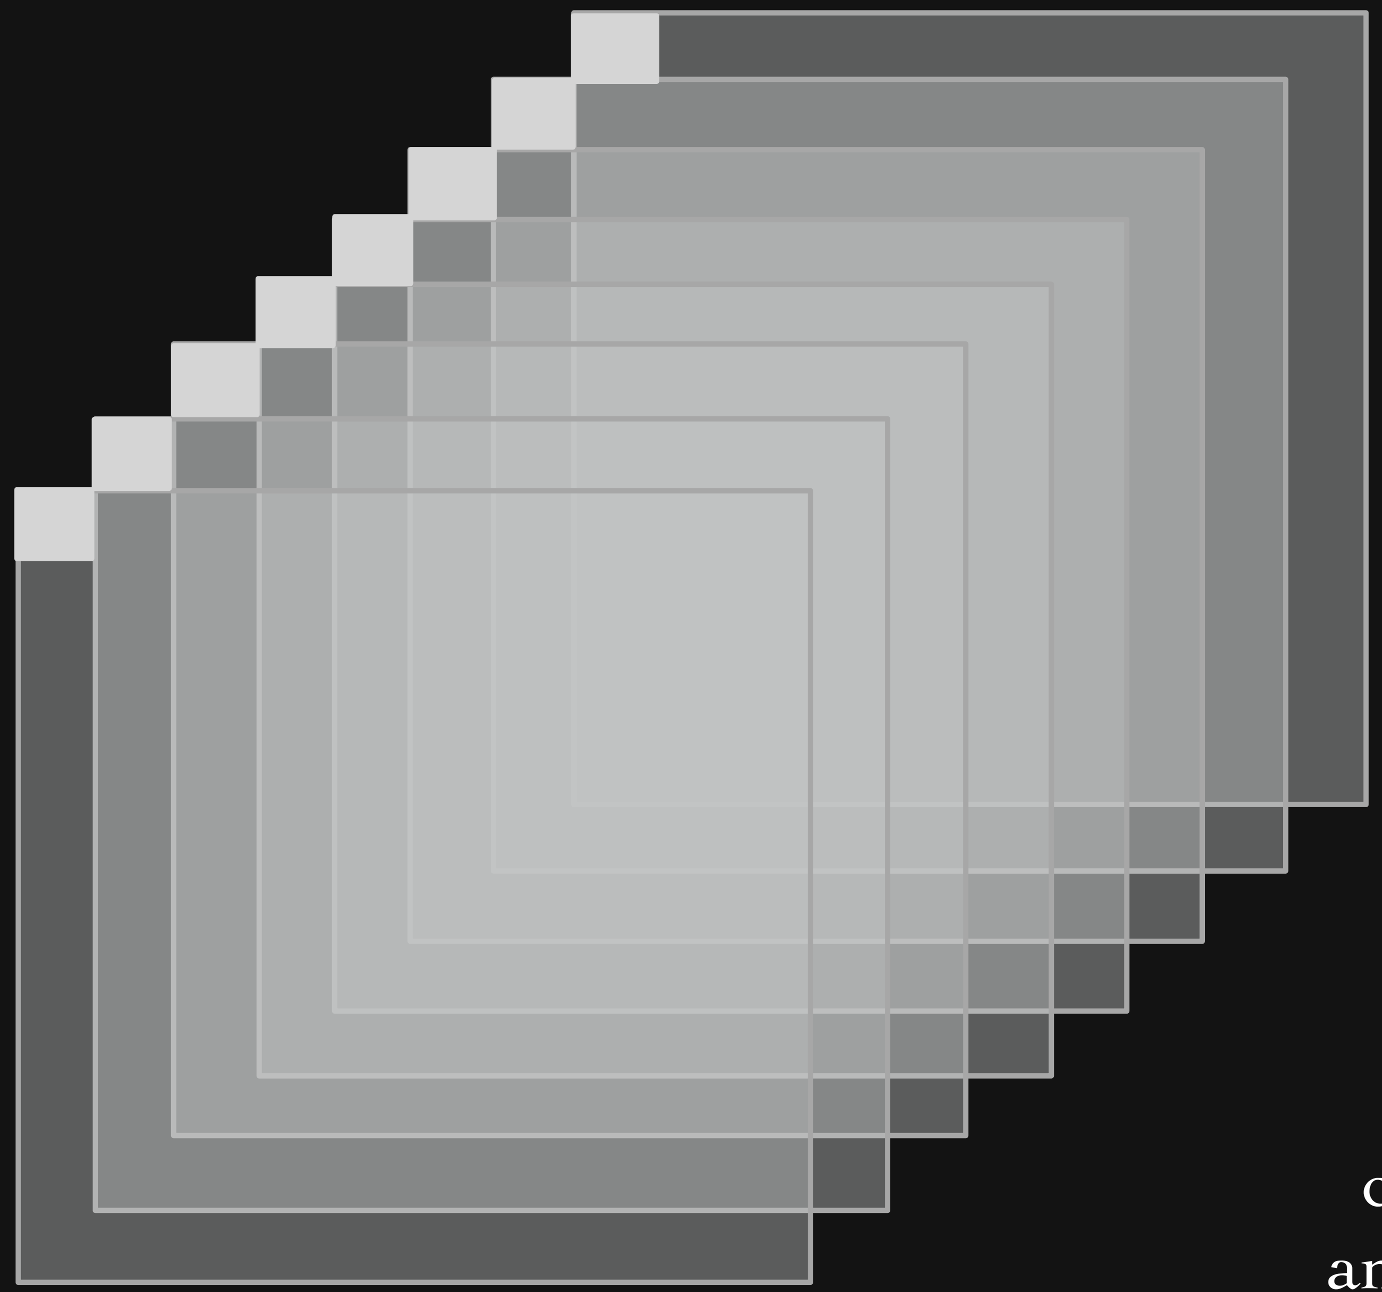

Each *(x,y)* spatial point corresponds to an element of the spectrum

Cube at one 2D scale

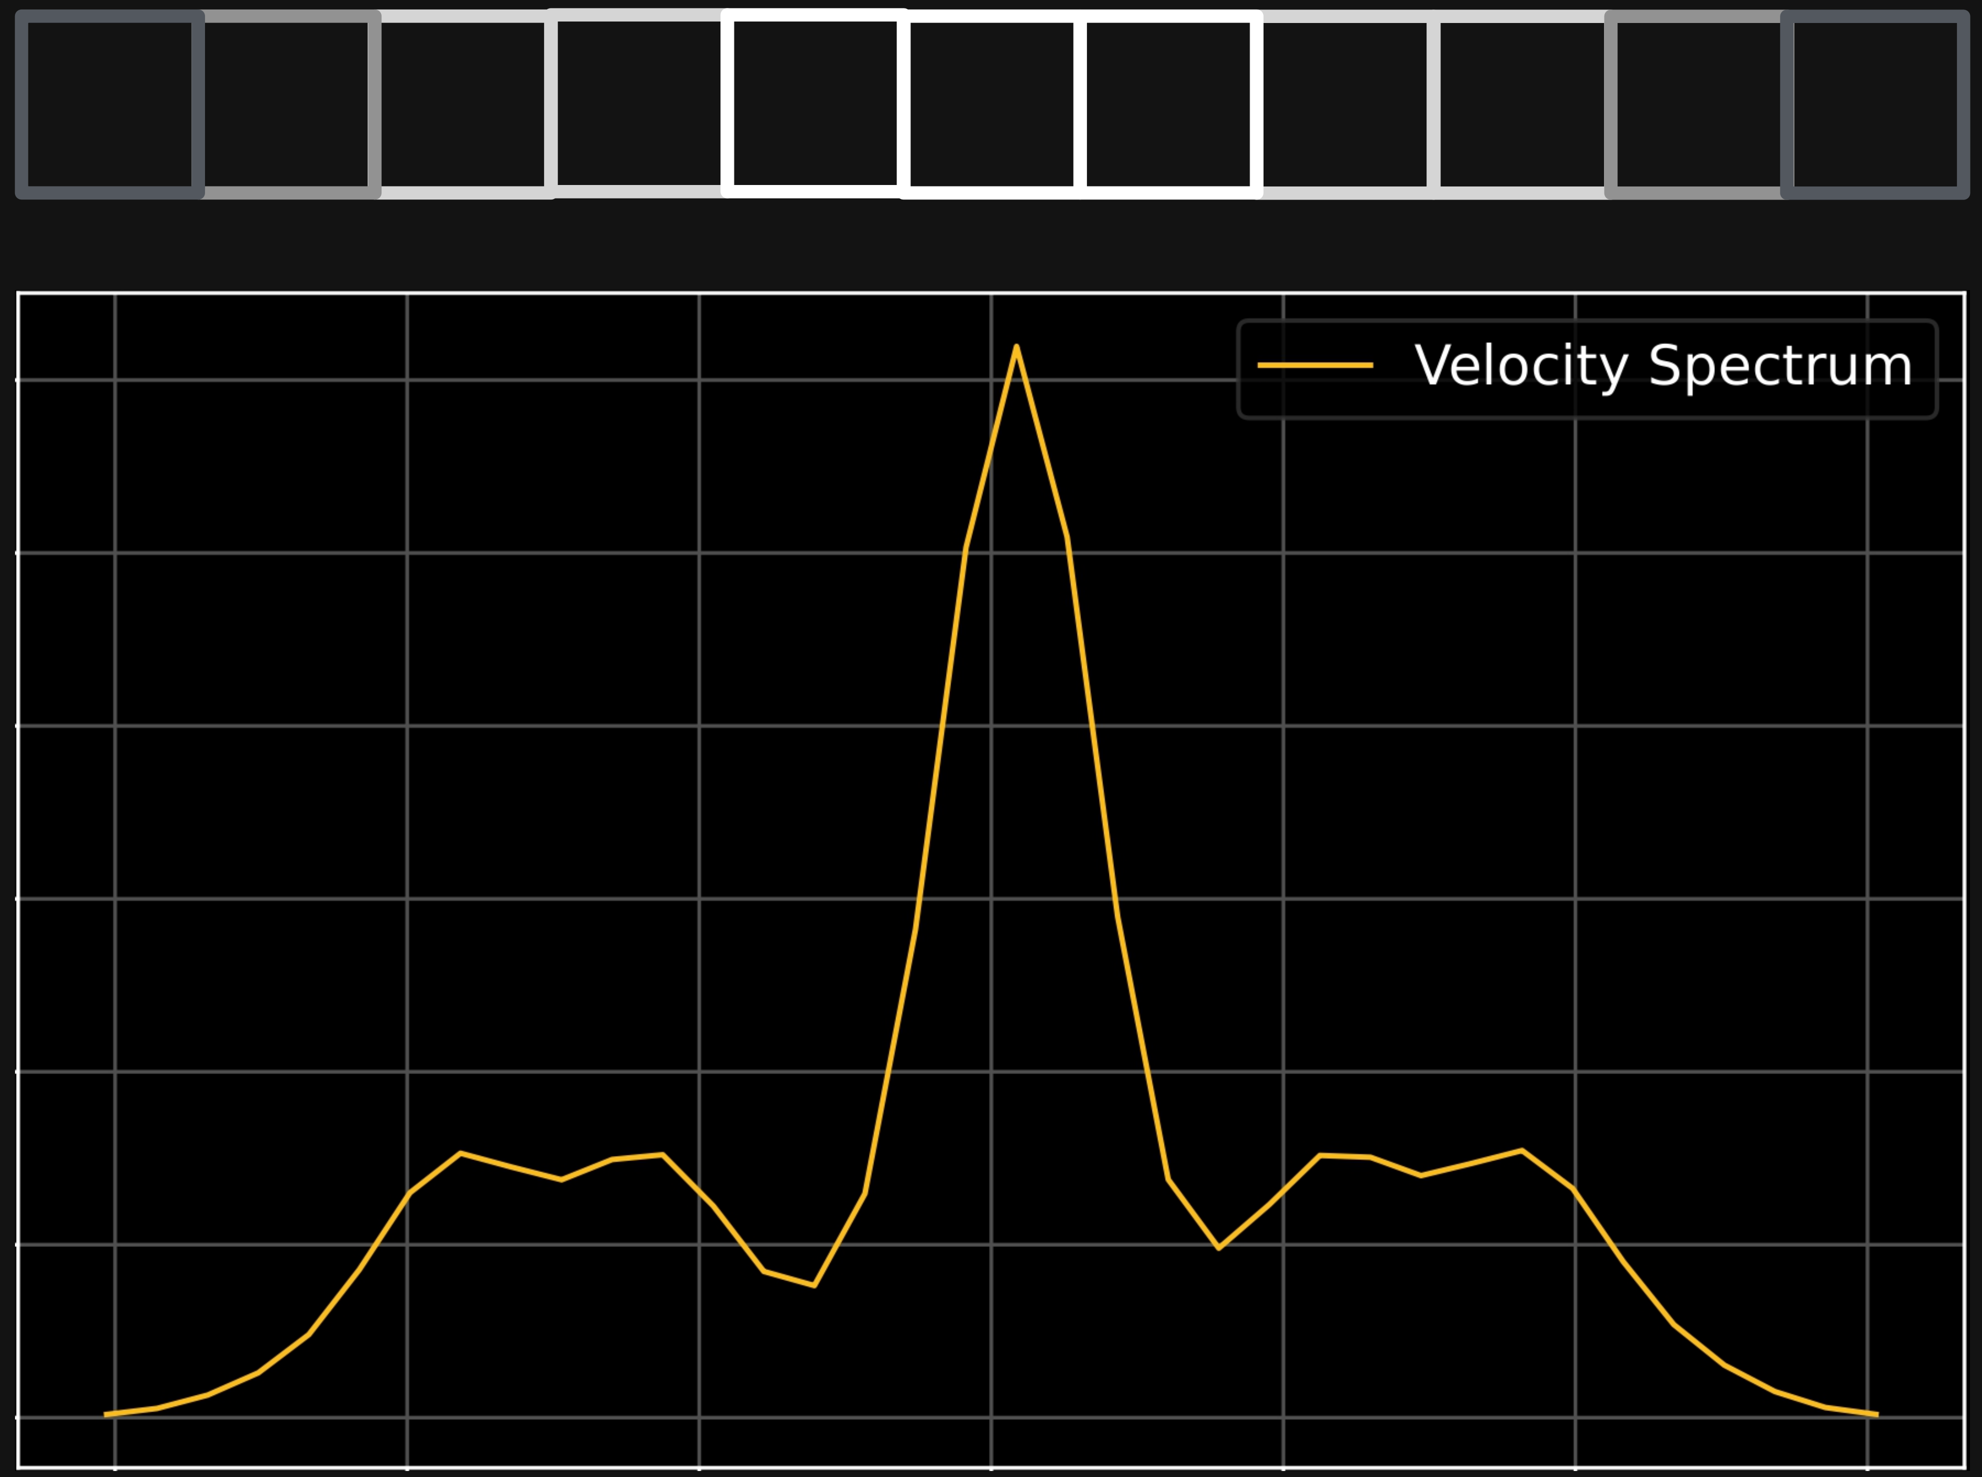

Example velocity spectrum for one spaxel (each element of 1 D array is the flux density at an (x,y) spaxel

2D-1D wavelet decomposition

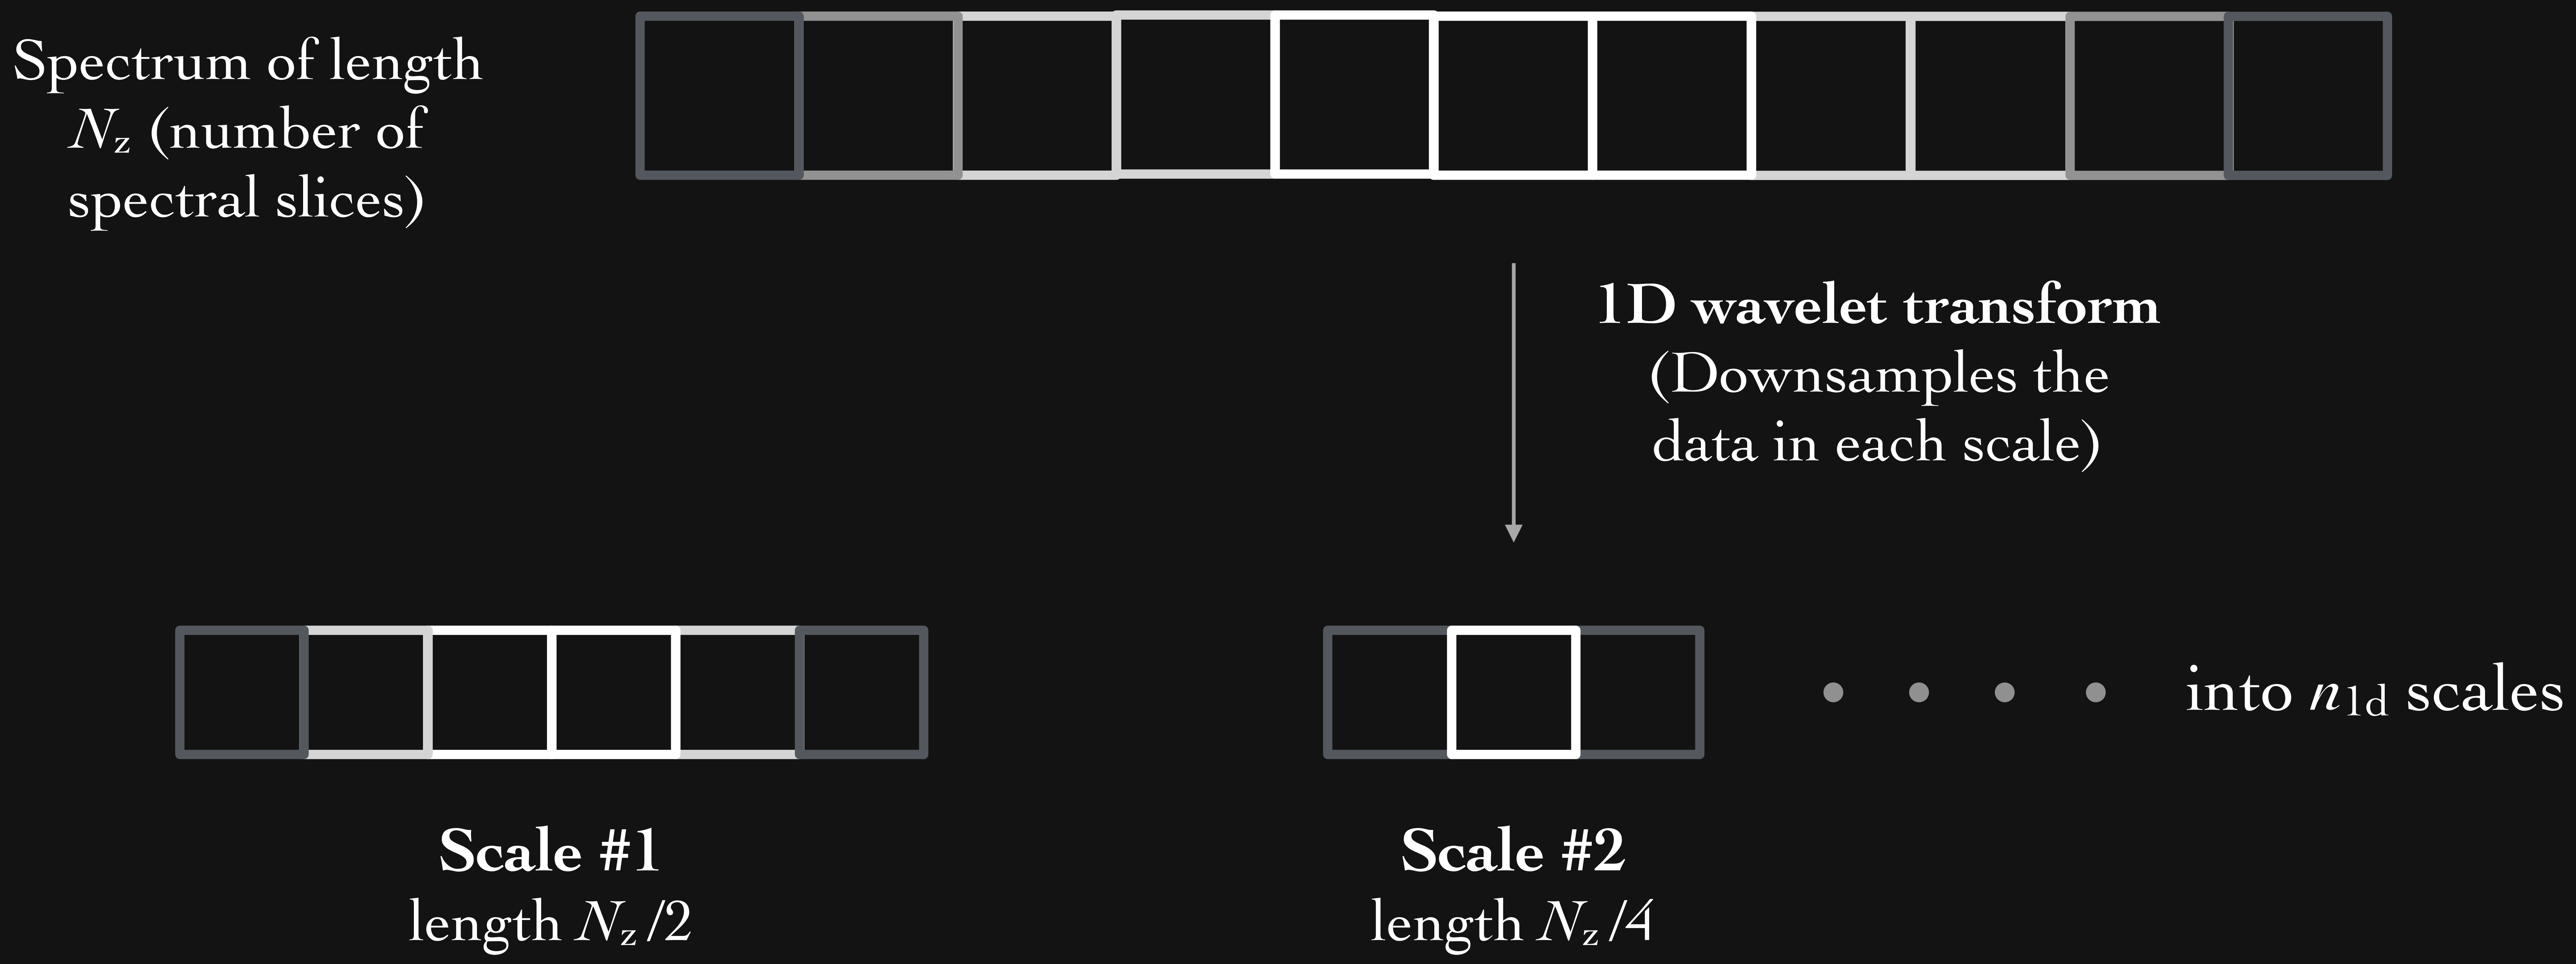

1D decomposition: The spectra associated with each spaxel is decomposed with a 1D wavelet transform

2D-1D wavelet decomposition

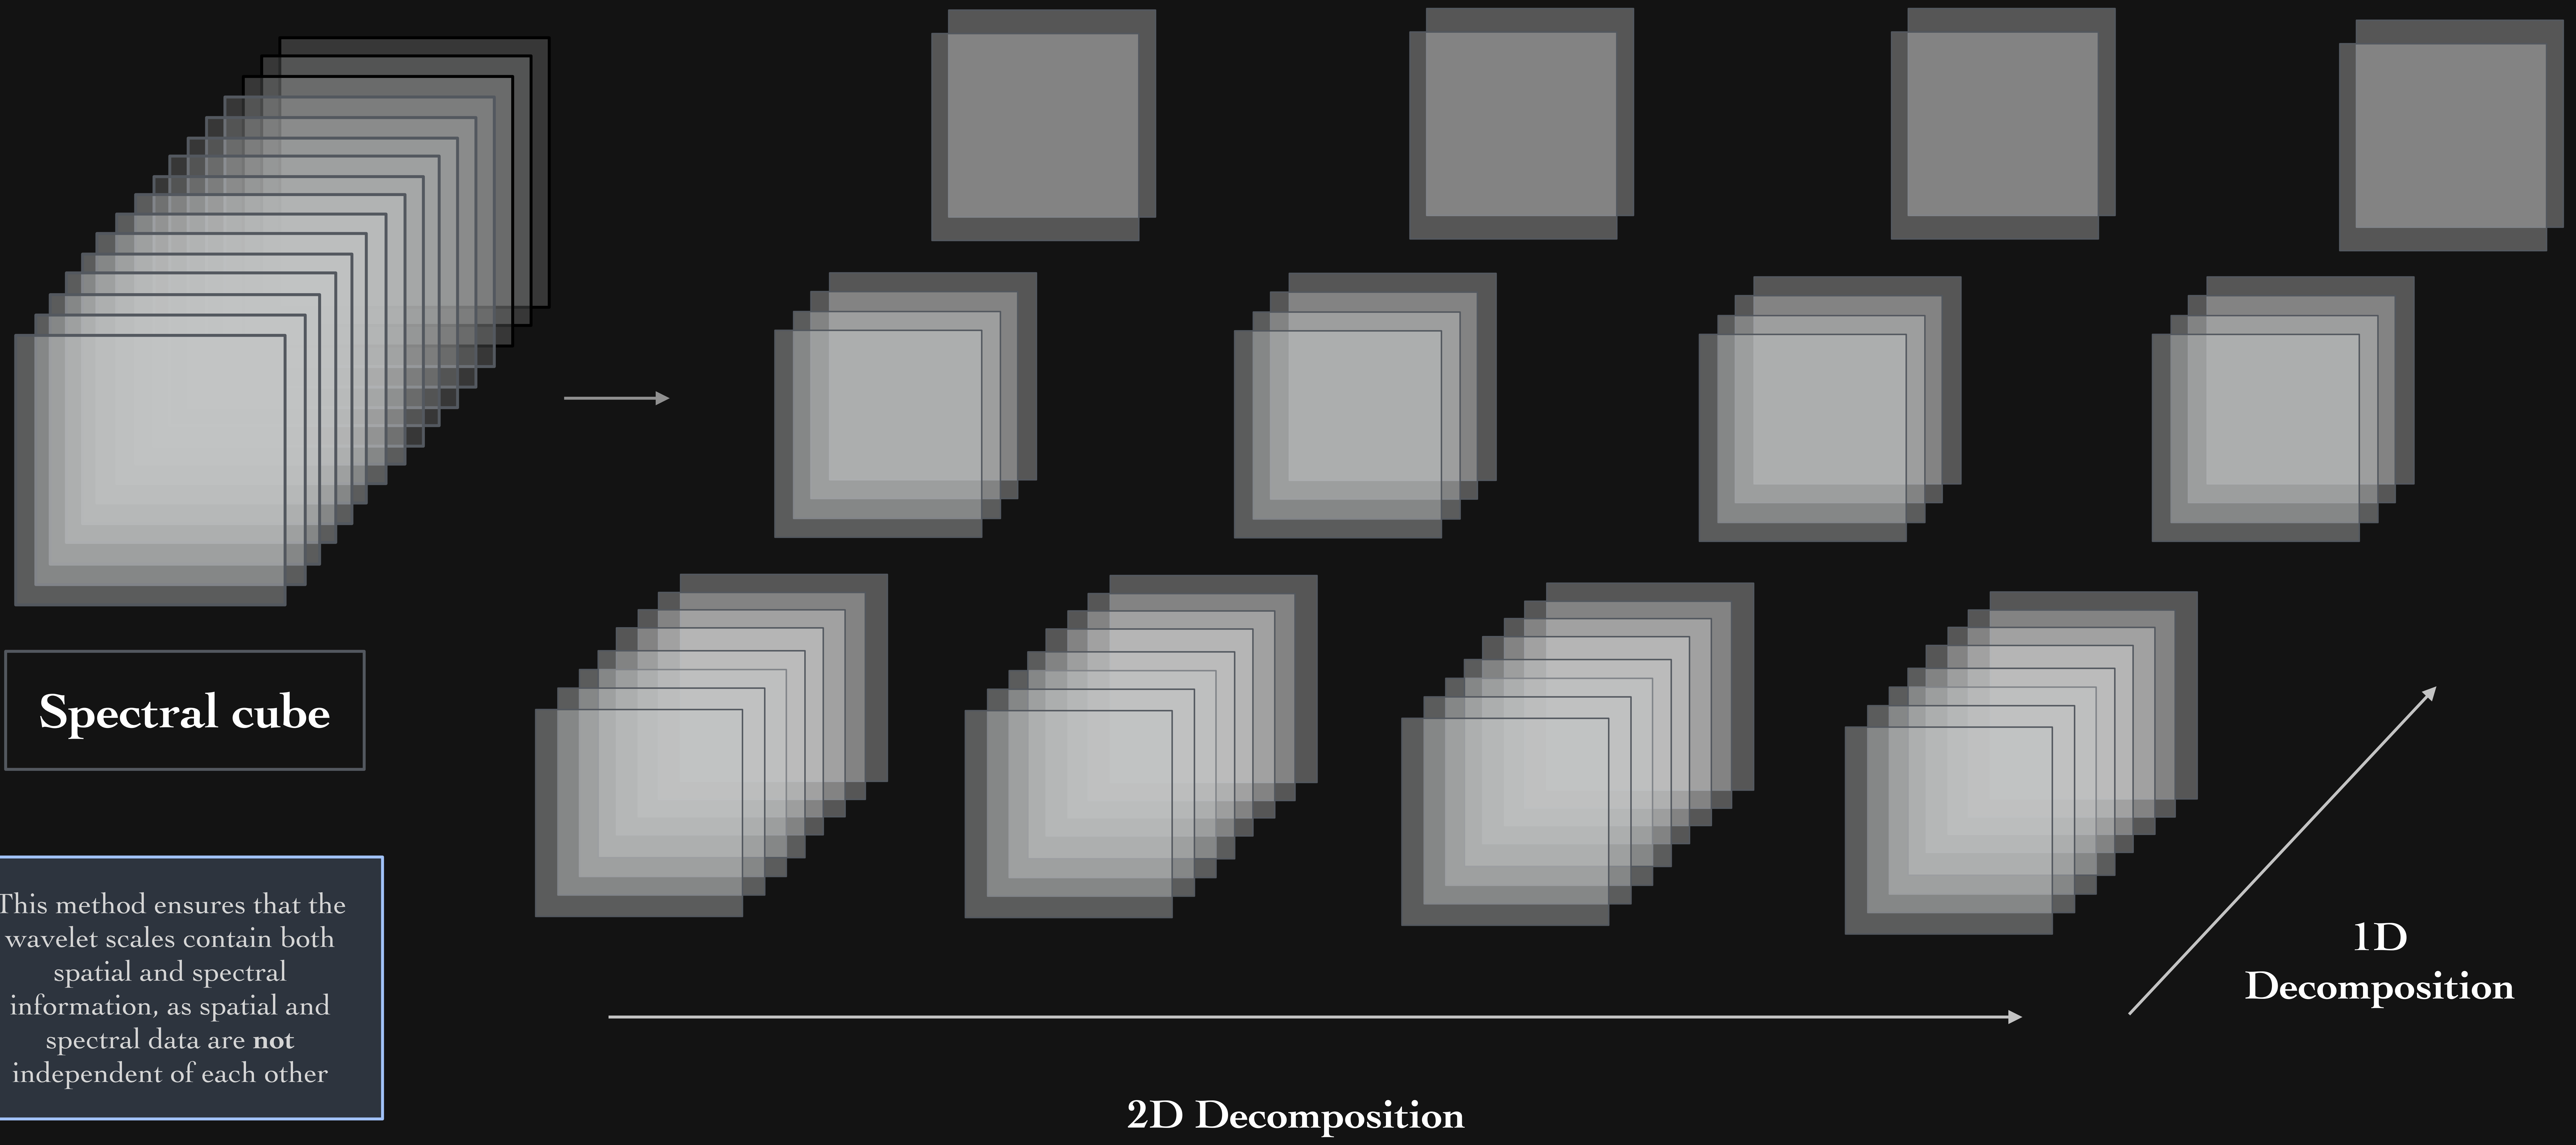

This method ensures that the wavelet scales contain both spatial and spectral information, as spatial and spectral data are **not** independent of each other

Denoising Methods

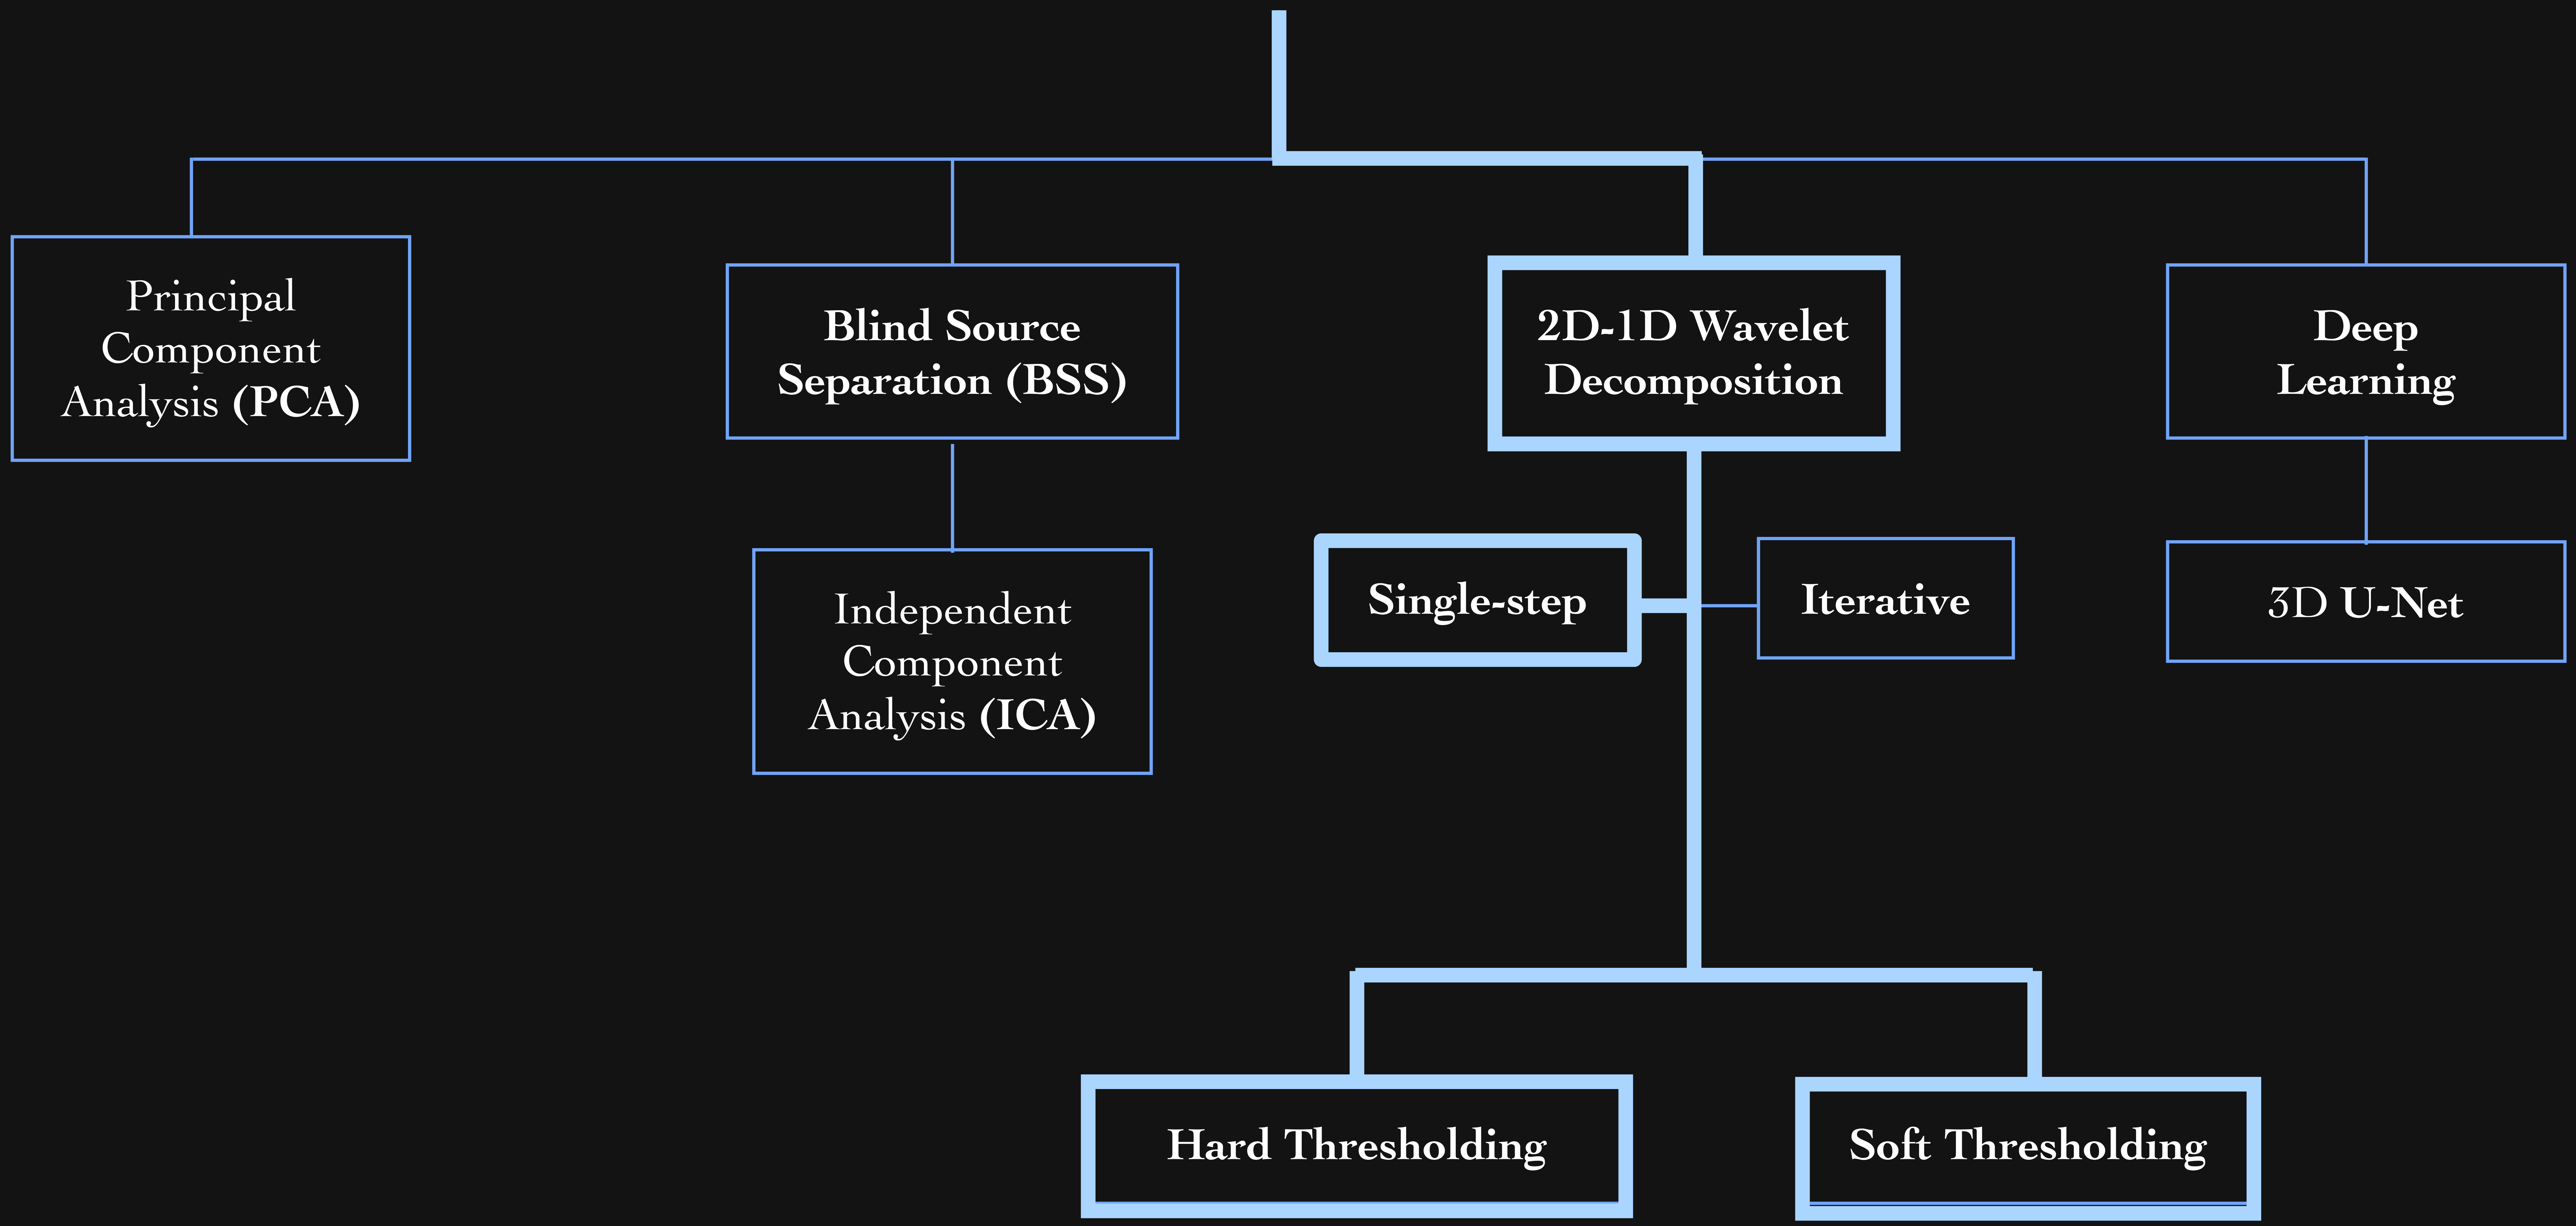

Single-Step Hard & Soft Thresholding

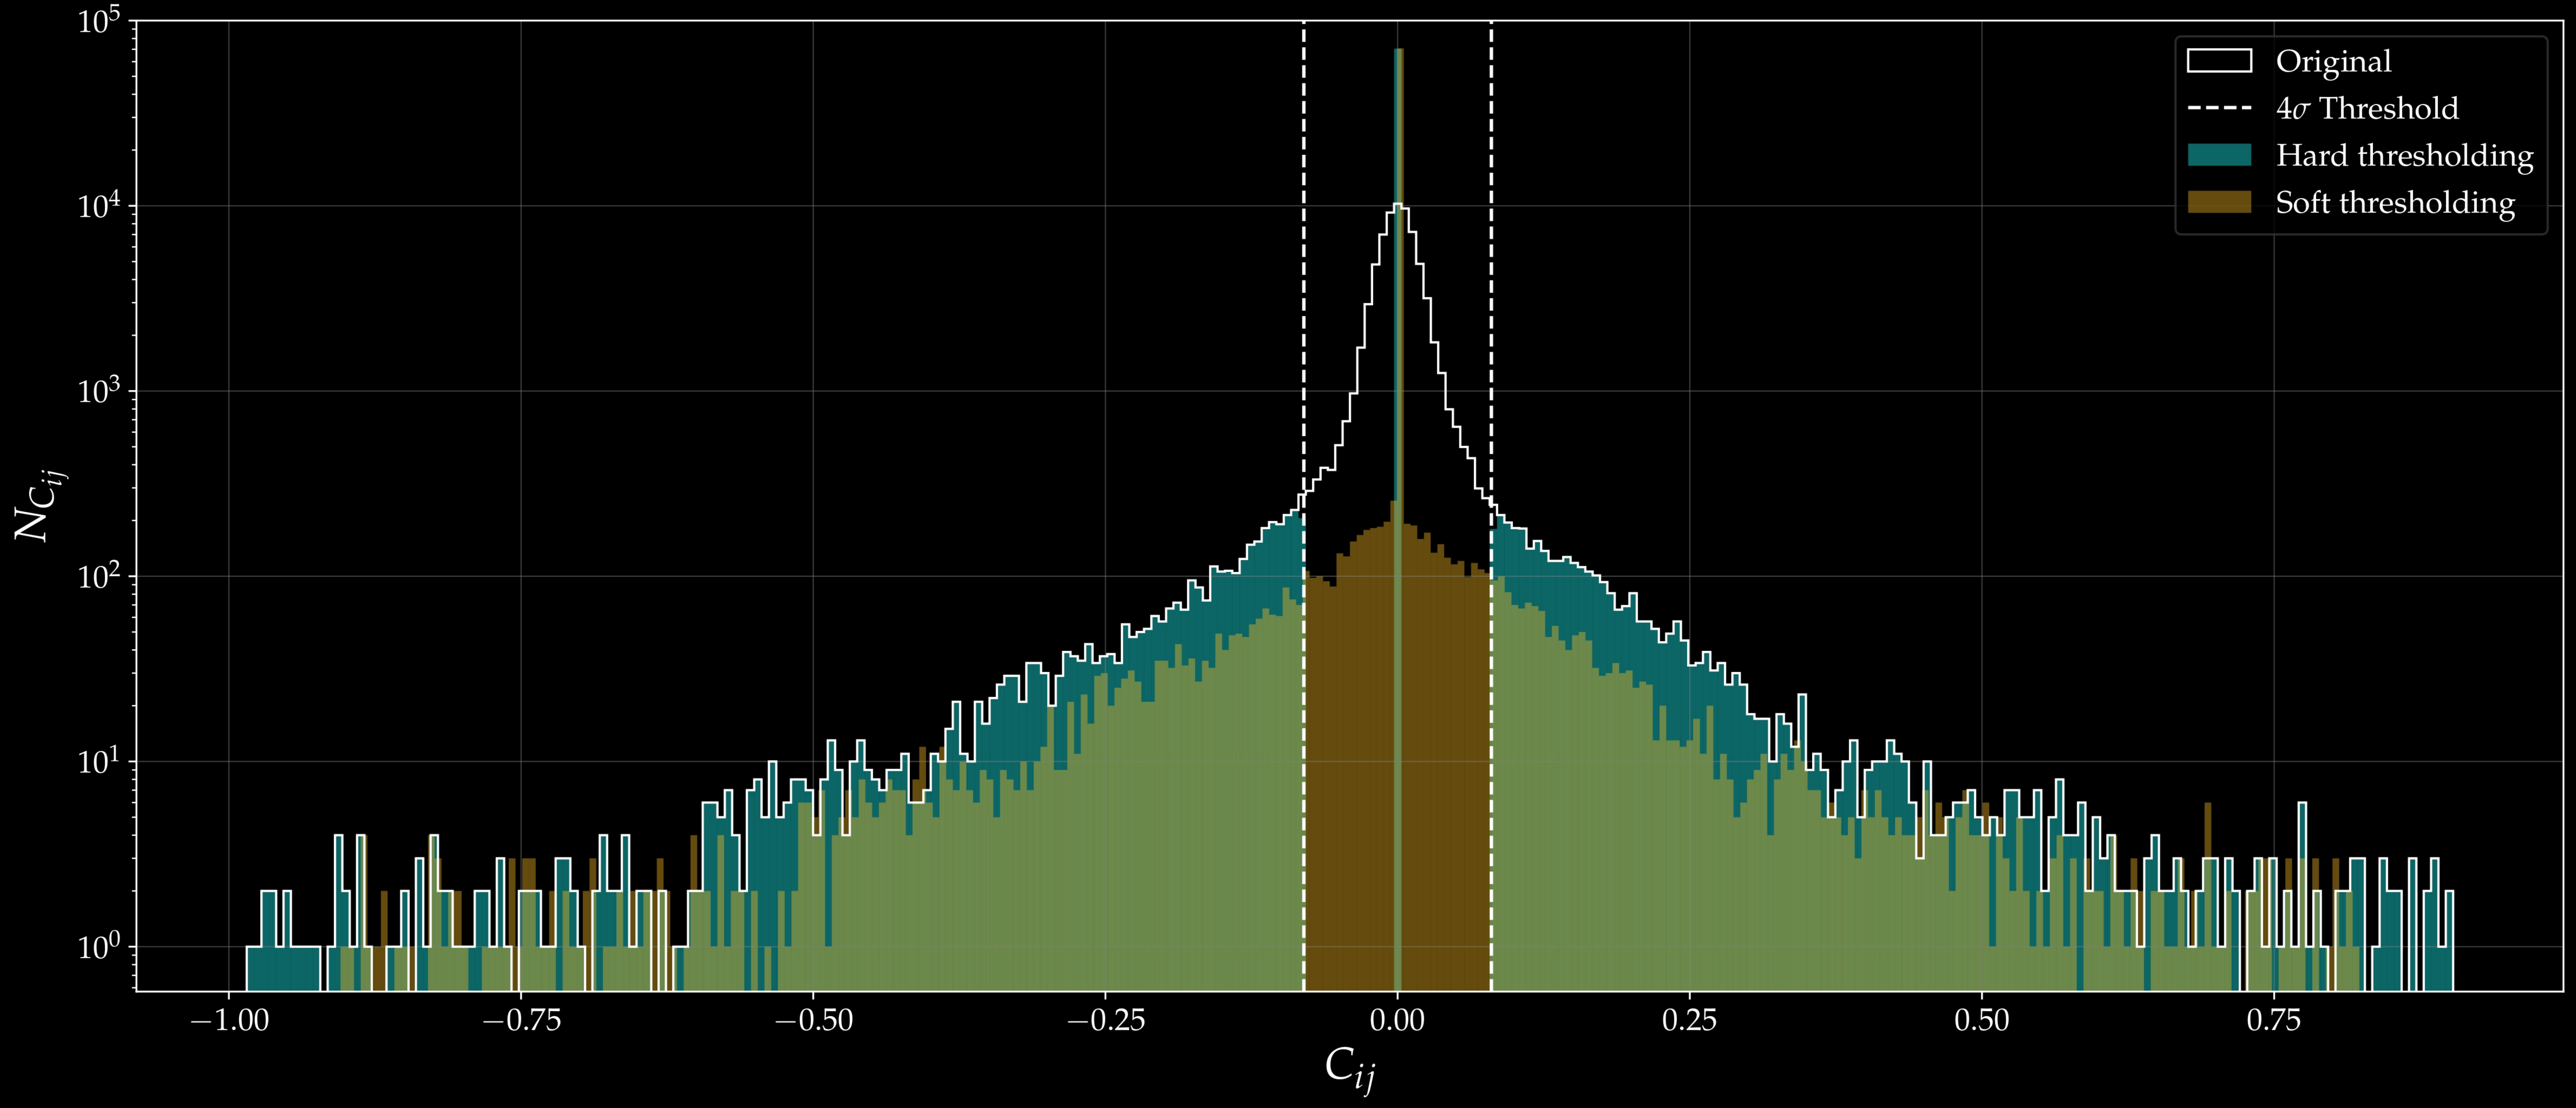

$$\mathcal{T}_{\text{hard}}(\alpha, \lambda) = \begin{cases} \alpha, & \text{if } |\alpha| \geq \lambda \\ 0, & \text{if } |\alpha| < \lambda \end{cases}$$

$$\mathcal{T}_{\text{soft}}(\alpha, \lambda) = \text{sign}(\alpha) \cdot \max(|\alpha| - \lambda, 0)$$

Denoising Methods

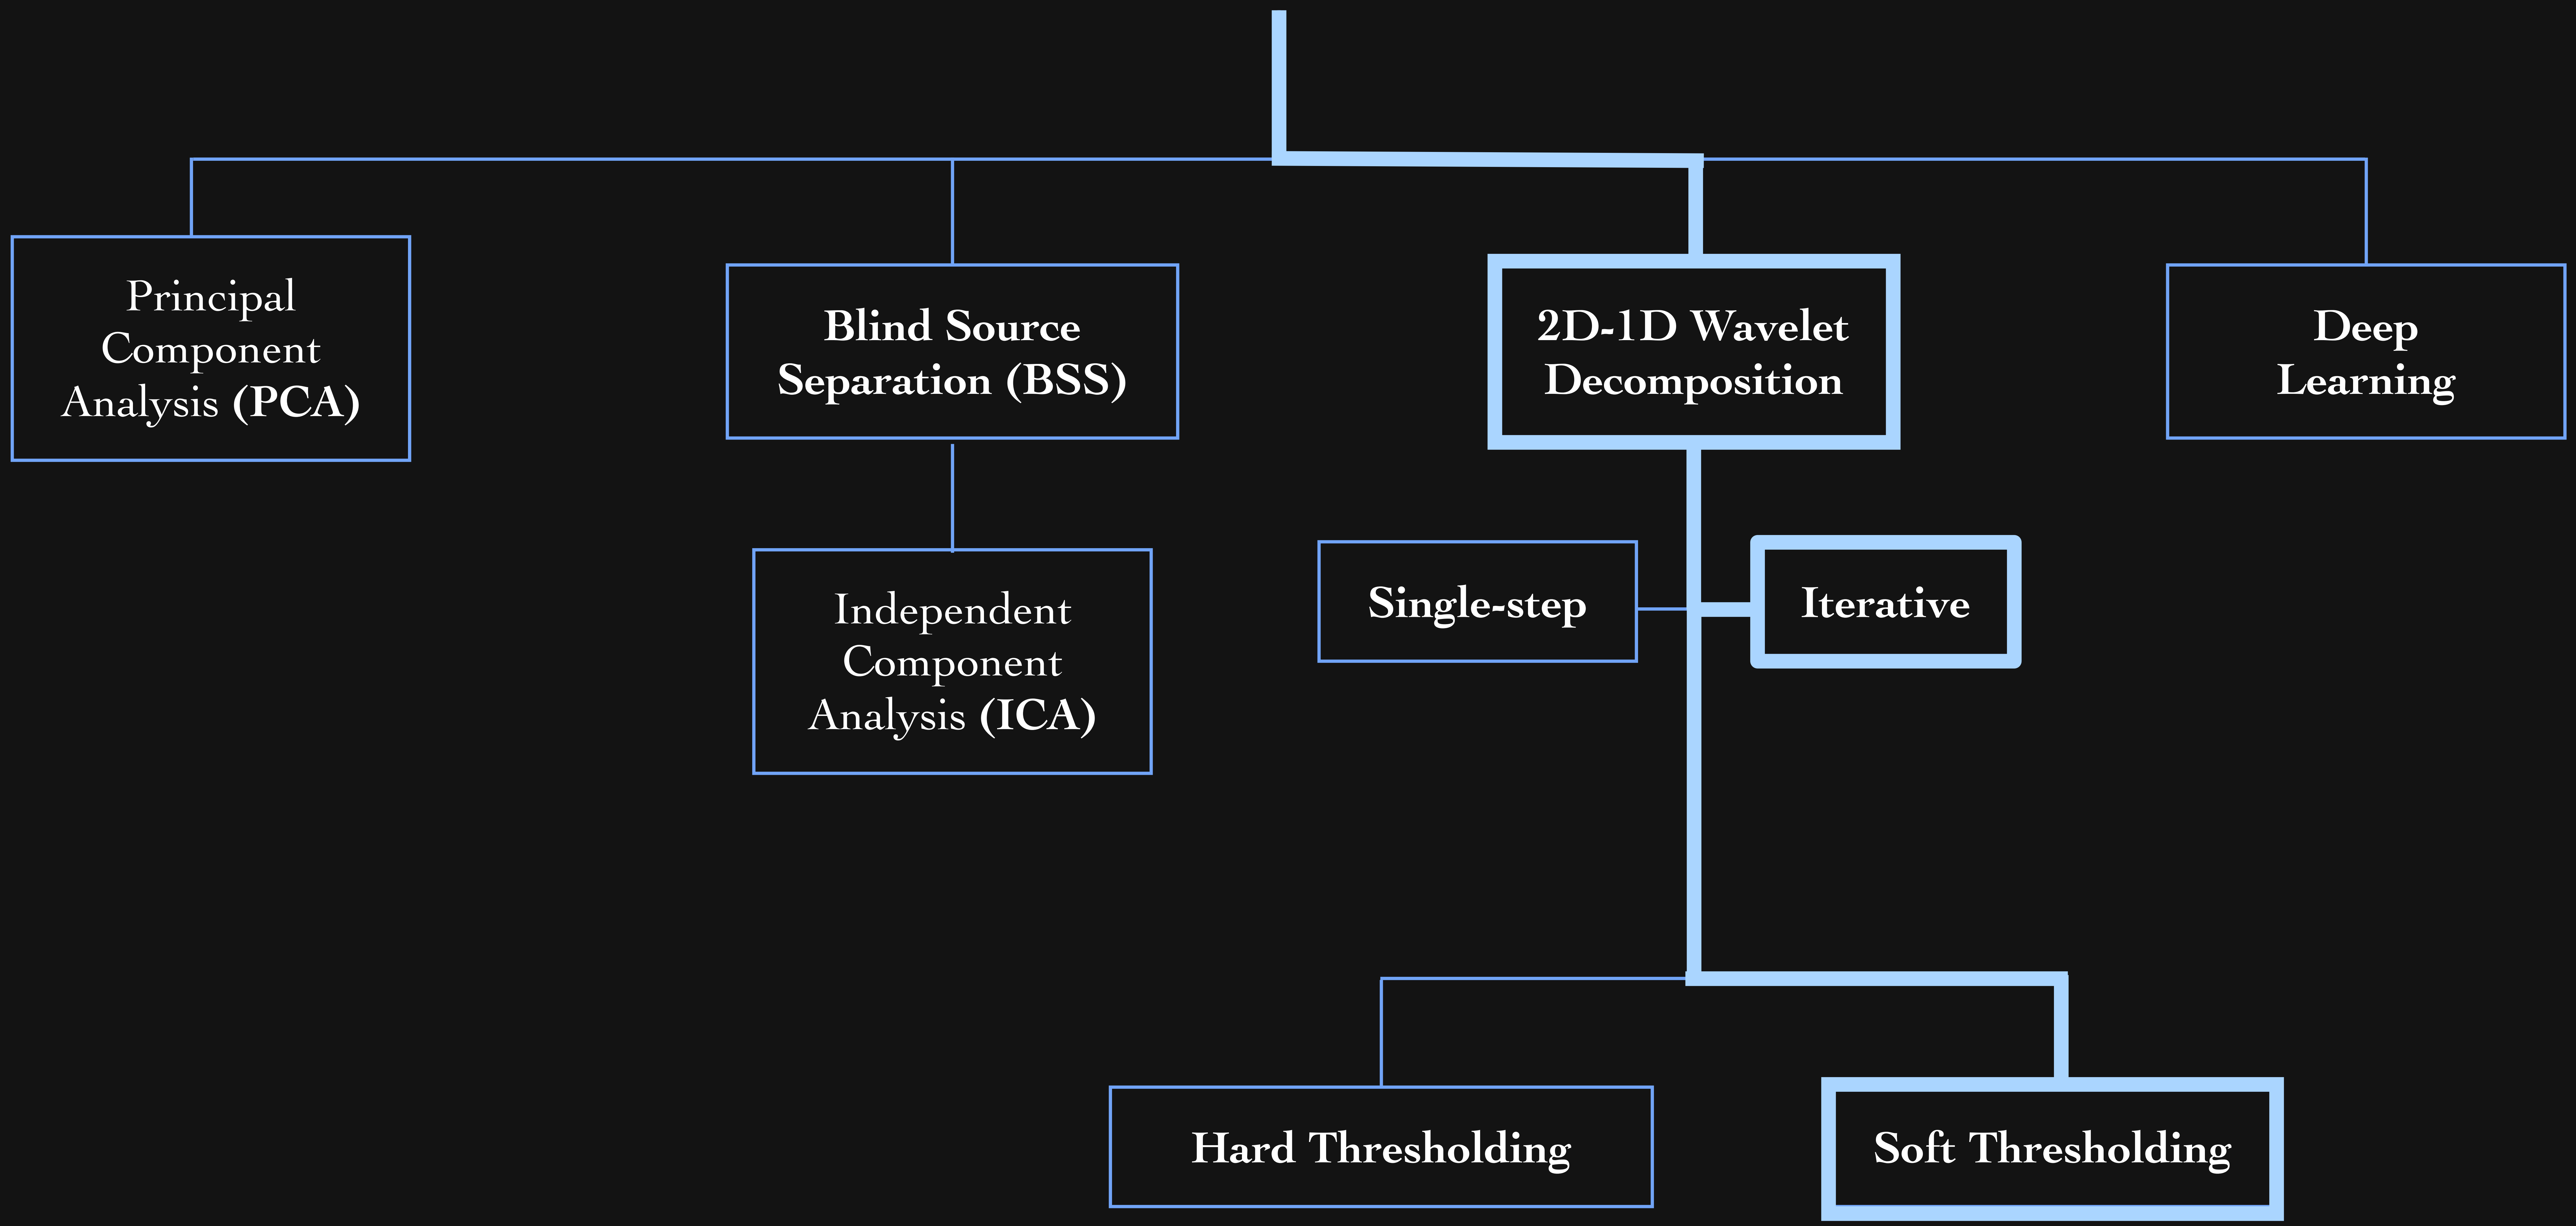

Re-weighted Iterative Soft Thresholding

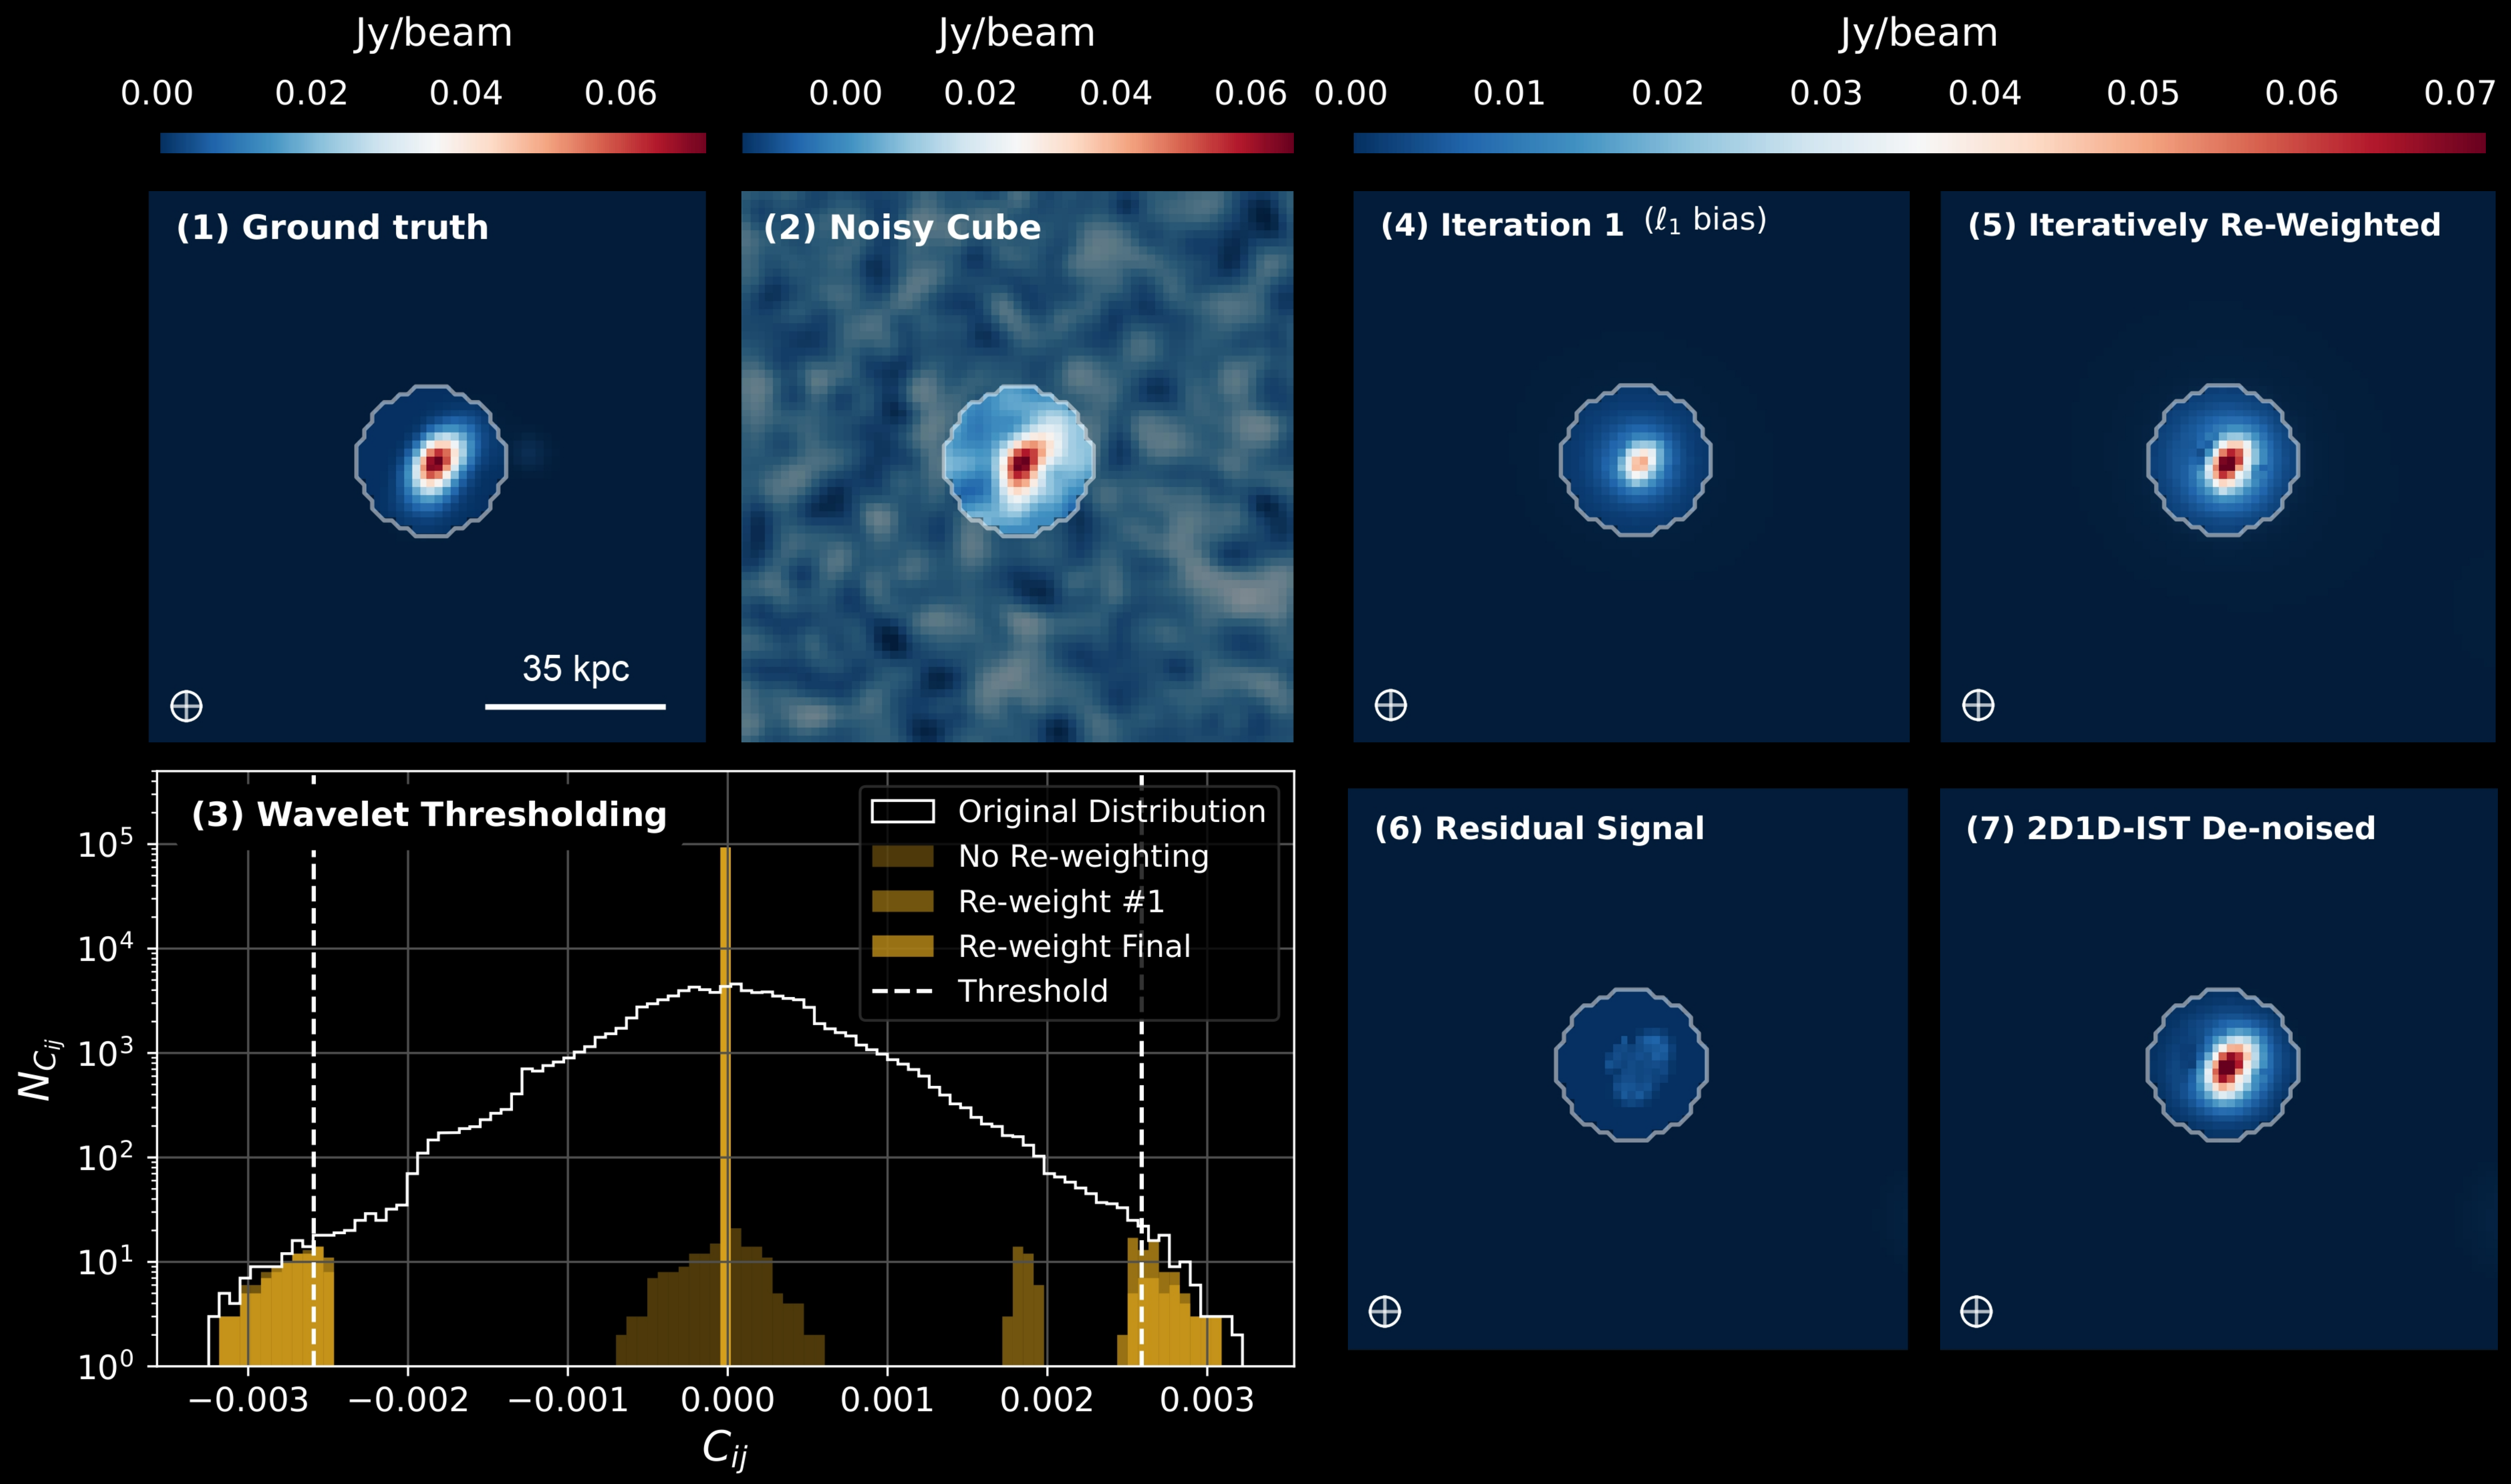

In subsequent iterations, **weights are calculated** as a function of the closeness of the coefficient magnitudes of the previous iteration to the threshold value

$$\mathcal{W}_{ij}^{(n)} = \begin{cases} \dfrac{\lambda \sigma_{\alpha^{(n-1)}}}{|\alpha_{ij}^{(n-1)}| + \epsilon}, & \text{if } |\alpha_{ij}^{(n)}| \geq \lambda \sigma_{\alpha}^{(n)} \\ 1, & \text{otherwise} \end{cases}$$

Re-weighted thresholding is applied in the next iteration after a gradient step, which pushes the data closer to the input

Coefficients that are closer to the threshold are thresholded more aggressively (higher w) and shrunken more towards 0 and the remaining coefficients are shrunken less

The residual is wavelet-transformed, and the final learned weights from the reweighting phase are applied to perform **weighted soft thresholding on the residual coefficients**. The resulting denoised residual is reconstructed and added to the previous output

$$\boldsymbol{X}^{(n)} = \text{prox}_{\max}\left(\boldsymbol{\Phi}\left(\mathcal{T}_{\lambda,\mathcal{W}^{(n-1)}}\left(\boldsymbol{\Phi}^T \boldsymbol{Y}\right)\right)\right)$$

$$\boldsymbol{X}^{(n+1)} = \text{prox}_{\max}\left(\boldsymbol{X}^{(n)} + \boldsymbol{\Phi}\left(\mathcal{T}_{\lambda,\mathcal{W}^{(\text{final})}}\left(\boldsymbol{\Phi}^T(\boldsymbol{Y} - \boldsymbol{X}^{(n)})\right)\right)\right)$$

Denoising Methods

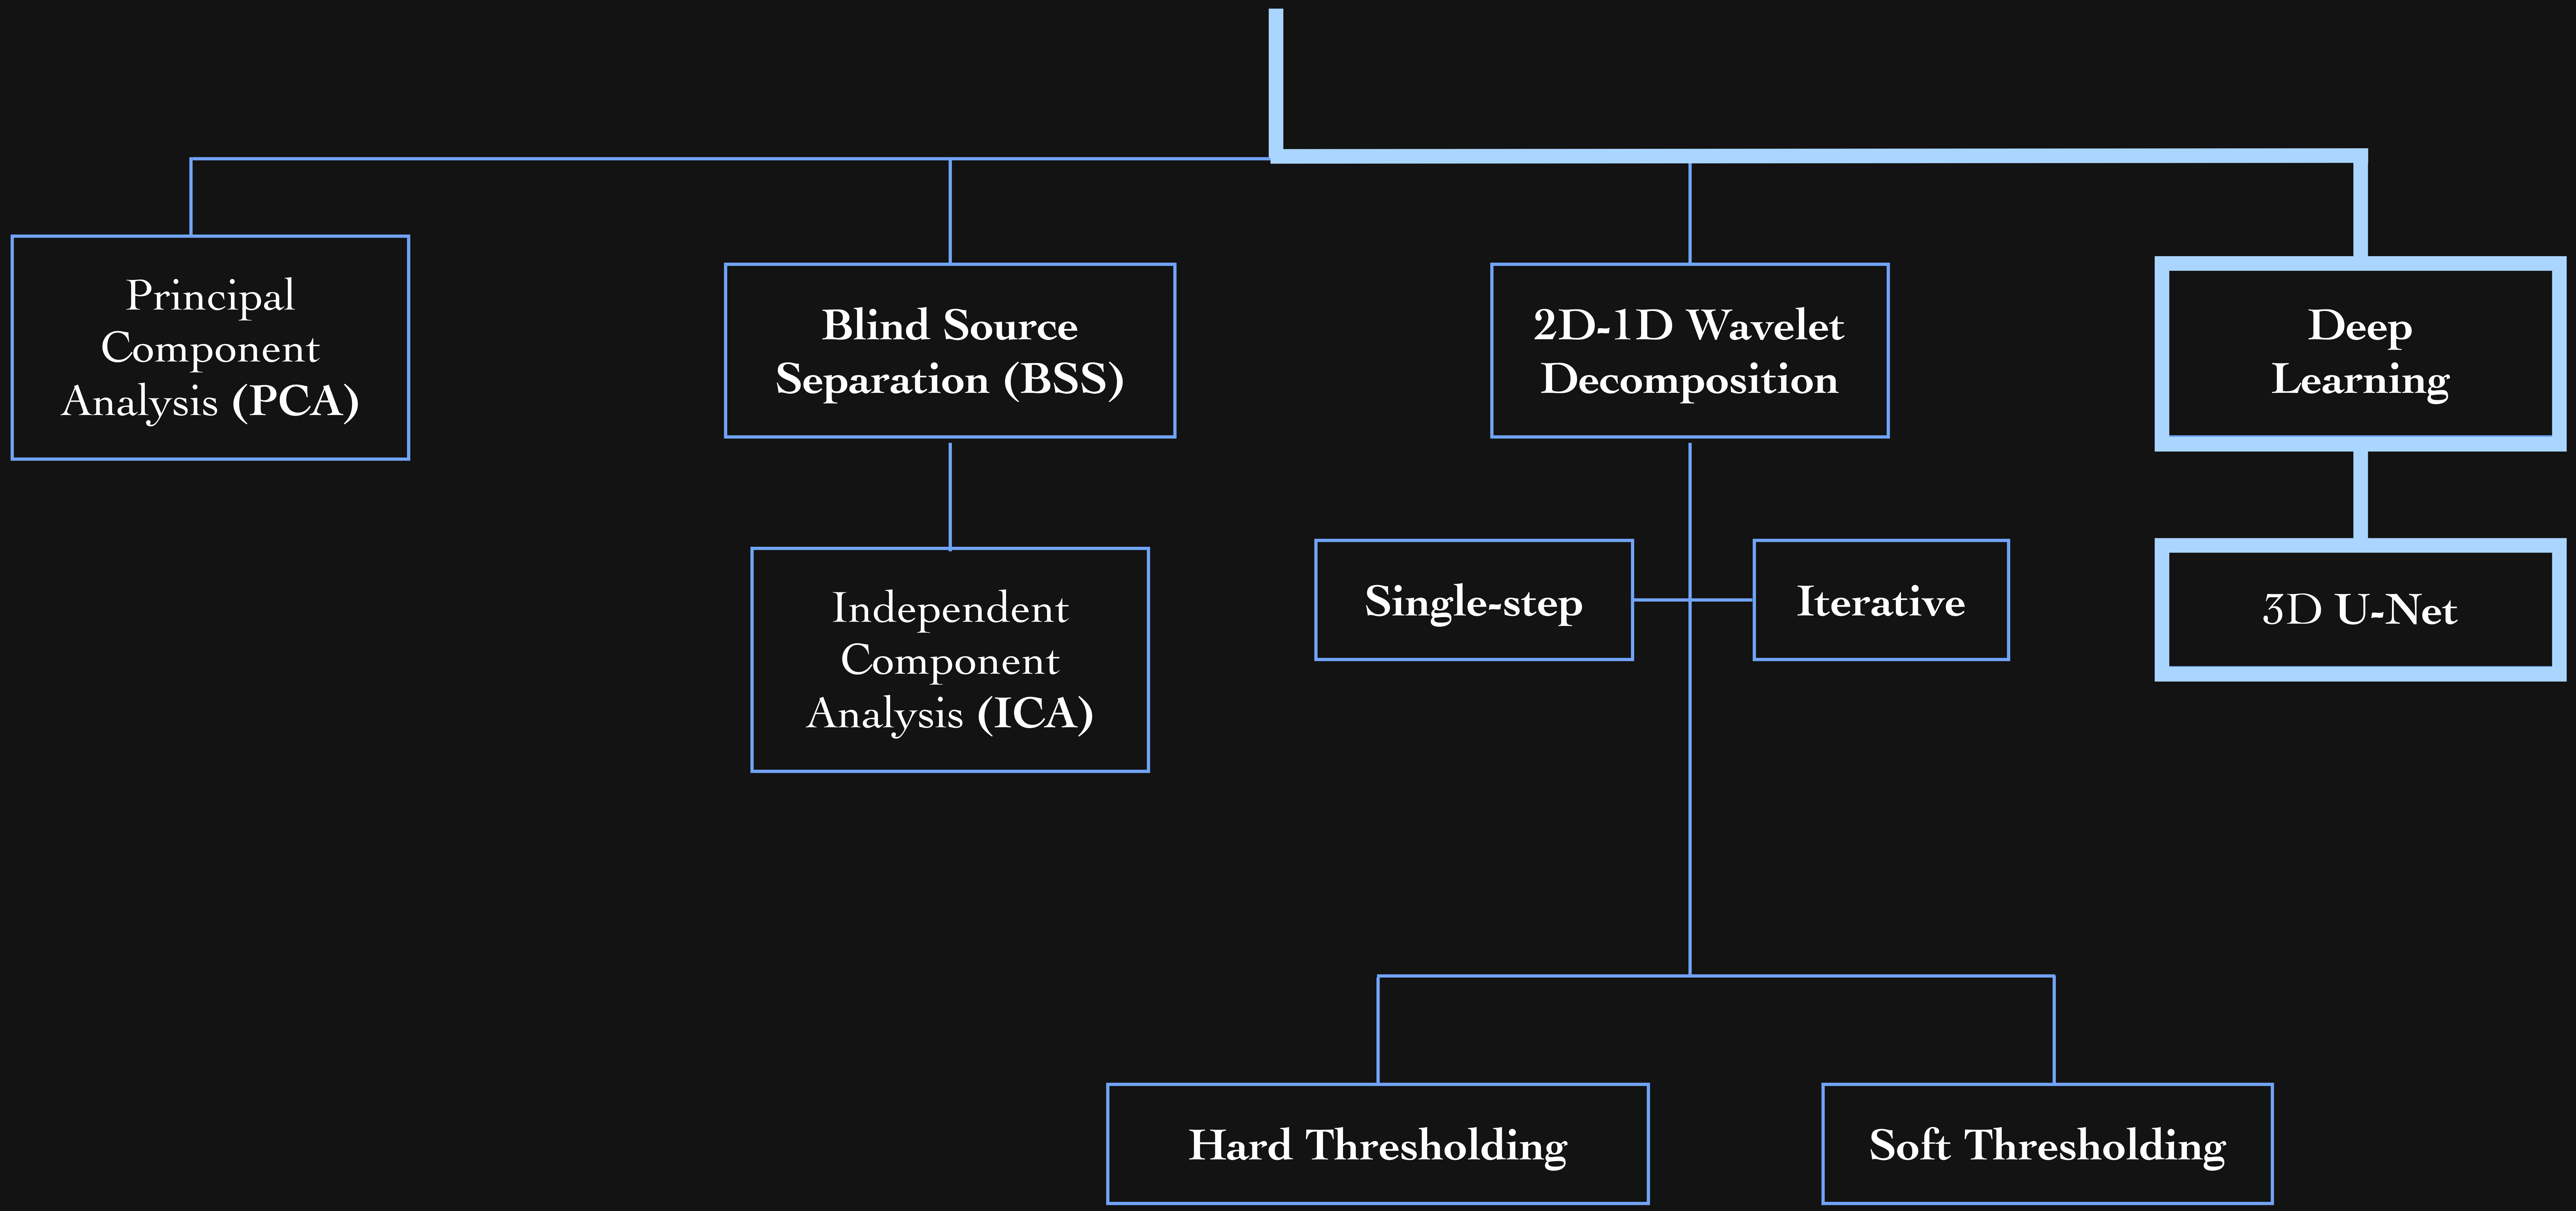

U-Net3D

Architecture

Encoder ***LEVEL 1***

Conv3D (input → F), *3×3×3, ReflectPad*

Conv3D (F → F), *3×3×3, ReflectPad*

AvgPool3D, *2×2×2*

Encoder ***LEVEL 2***

Conv3D (F → 2F), *3×3×3, ReflectPad*

Conv3D (2F → 2F), *3×3×3, ReflectPad*

AvgPool3D, *2×2×2*

Encoder ***LEVEL 3***

Conv3D (2F → 4F), *3×3×3, ReflectPad*

Conv3D (4F → 4F), *3×3×3, ReflectPad*

AvgPool3D, *2×2×2*

Encoder ***LEVEL 4***

Conv3D (4F → 8F), *3×3×3, ReflectPad*

Conv3D (8F → 8F), *3×3×3, ReflectPad*

AvgPool3D, *2×2×2*

BOTTLENECK

Conv3D (8F → 16F), *3×3×3, ReflectPad*

Conv3D (16F → 16F), *3×3×3, ReflectPad*

Skip connection ***(Level 1)***

Skip connection ***(Level 2)***

Skip connection ***(Level 3)***

Skip connection ***(Level 4)***

Decoder ***LEVEL 1***

ConvTrans3D (16F → 8F), *2×2×2 [****Concat Level 4****]*

Conv3D (16F → 8F), *3×3×3, ReflectPad*

Conv3D (8F → 8F), *3×3×3, ReflectPad*

Decoder ***LEVEL 2***

ConvTrans3D (8F → 4F), *2×2×2 [****Concat Level 3****]*

Conv3D (8F → 4F), *3×3×3, ReflectPad*

Conv3D (4F → 4F), *3×3×3, ReflectPad*

Decoder ***LEVEL 3***

ConvTrans3D (4F → 2F), *2×2×2 [****Concat Level 2****]*

Conv3D (4F → 2F), *3×3×3, ReflectPad*

Conv3D (2F → 2F), *3×3×3, ReflectPad*

Decoder ***LEVEL 4***

ConvTrans3D (2F → 1F), *2×2×2 [****Concat Level 1****]*

Conv3D (2F → 1F), *3×3×3, ReflectPad*

Conv3D (1F → 1F), *3×3×3, ReflectPad*

Analysis & Results


PUNCH IT CHEWIE,
LET'S MAKE THE JUMP TO LIGHT MODE

RMSE reduction & Flux conservation

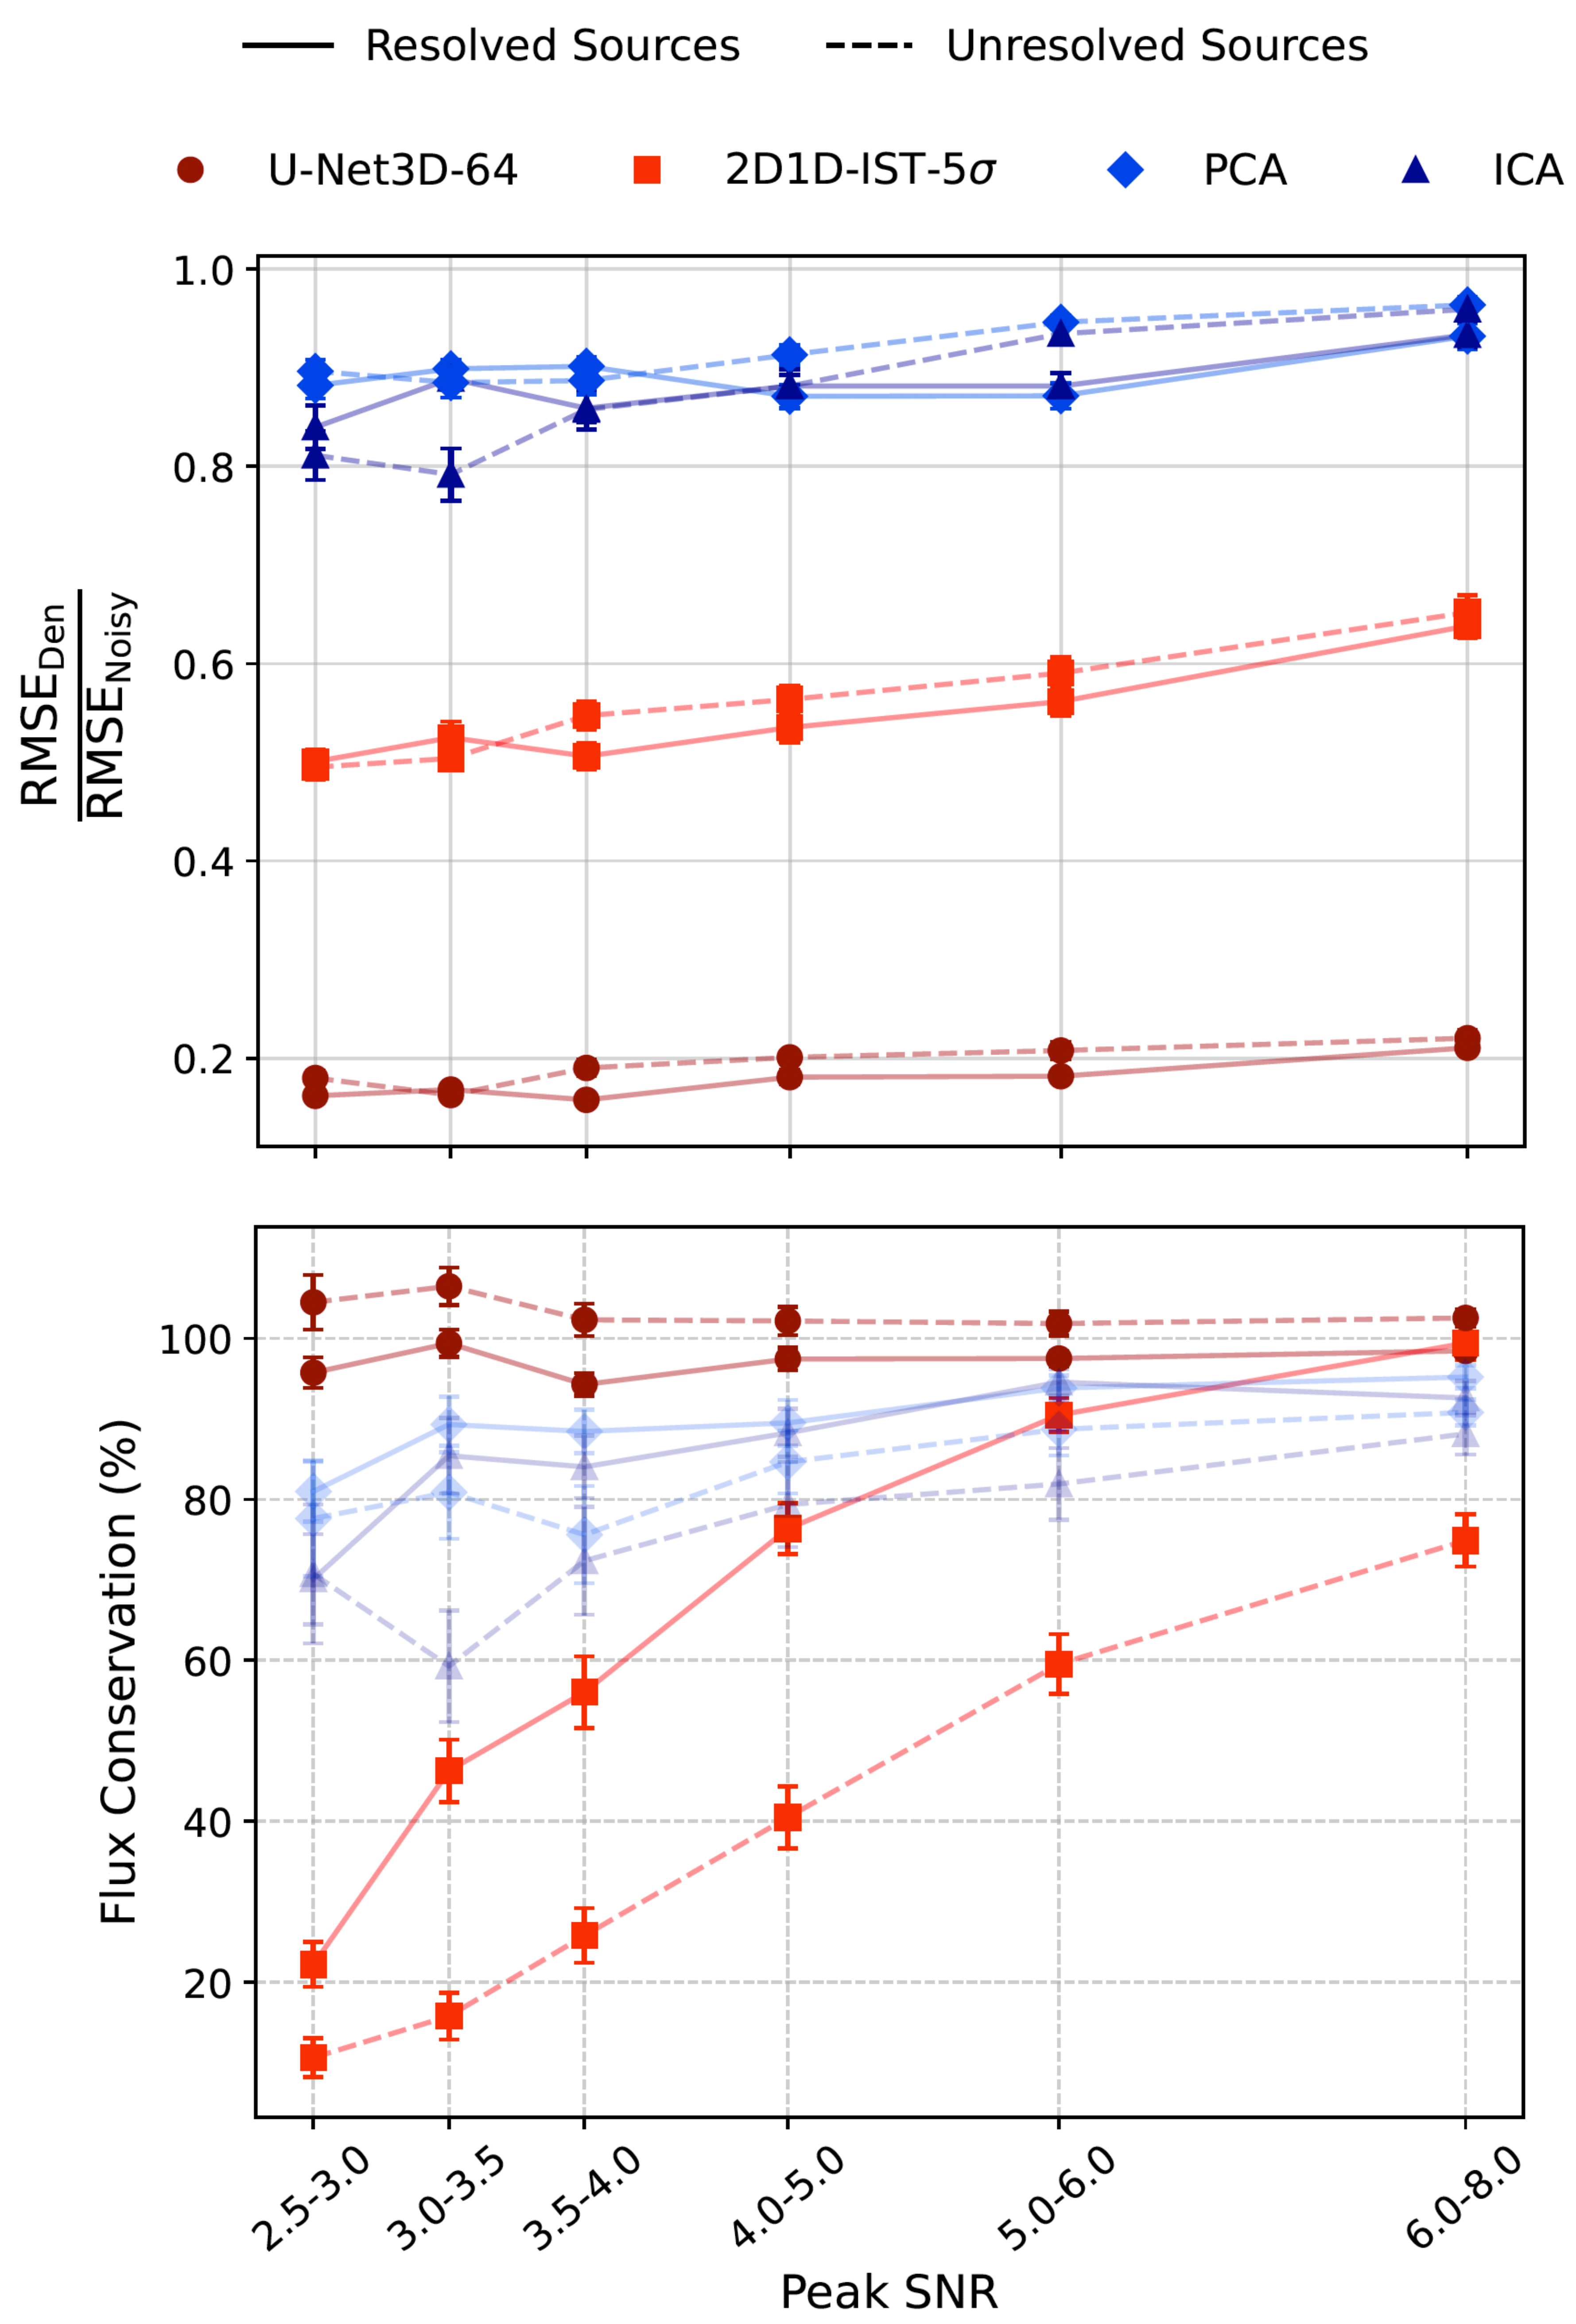

- **PCA & ICA:**
  - RMSE ratio after denoising remains close to 1
    almost no noise suppression.
  - Flux appears well-conserved, but mainly because noise is not effectively removed — giving a **misleading impression of good performance.**
  - **Ineffective for correlated 3D noise** in spectral cubes.

- **2D1D IST:**
  - RMSE reduced by ~40-50%
    Significant noise reduction compared to PCA/ICA.
  - **Flux is underestimated at very low SNRs**
    (due to conservative thresholding) but improves as SNR increases.
  - **Better balance between noise suppression and signal preservation** thanks to its multi-scale representation.

- **U-Net3D-64:**
  - Best overall performance: RMSE reduced by ~80%
    Superior noise suppression even at low SNRs.
  - Flux conservation is excellent for **resolved** sources.
  - For **unresolved** sources at very low SNR, U-Net tends to **overestimate flux by ~5-10%**, which improves at higher SNRs. Likely due to the **U-Net learning to recover compact signal patterns more aggressively** than real emission.

Spectral Feature Retention

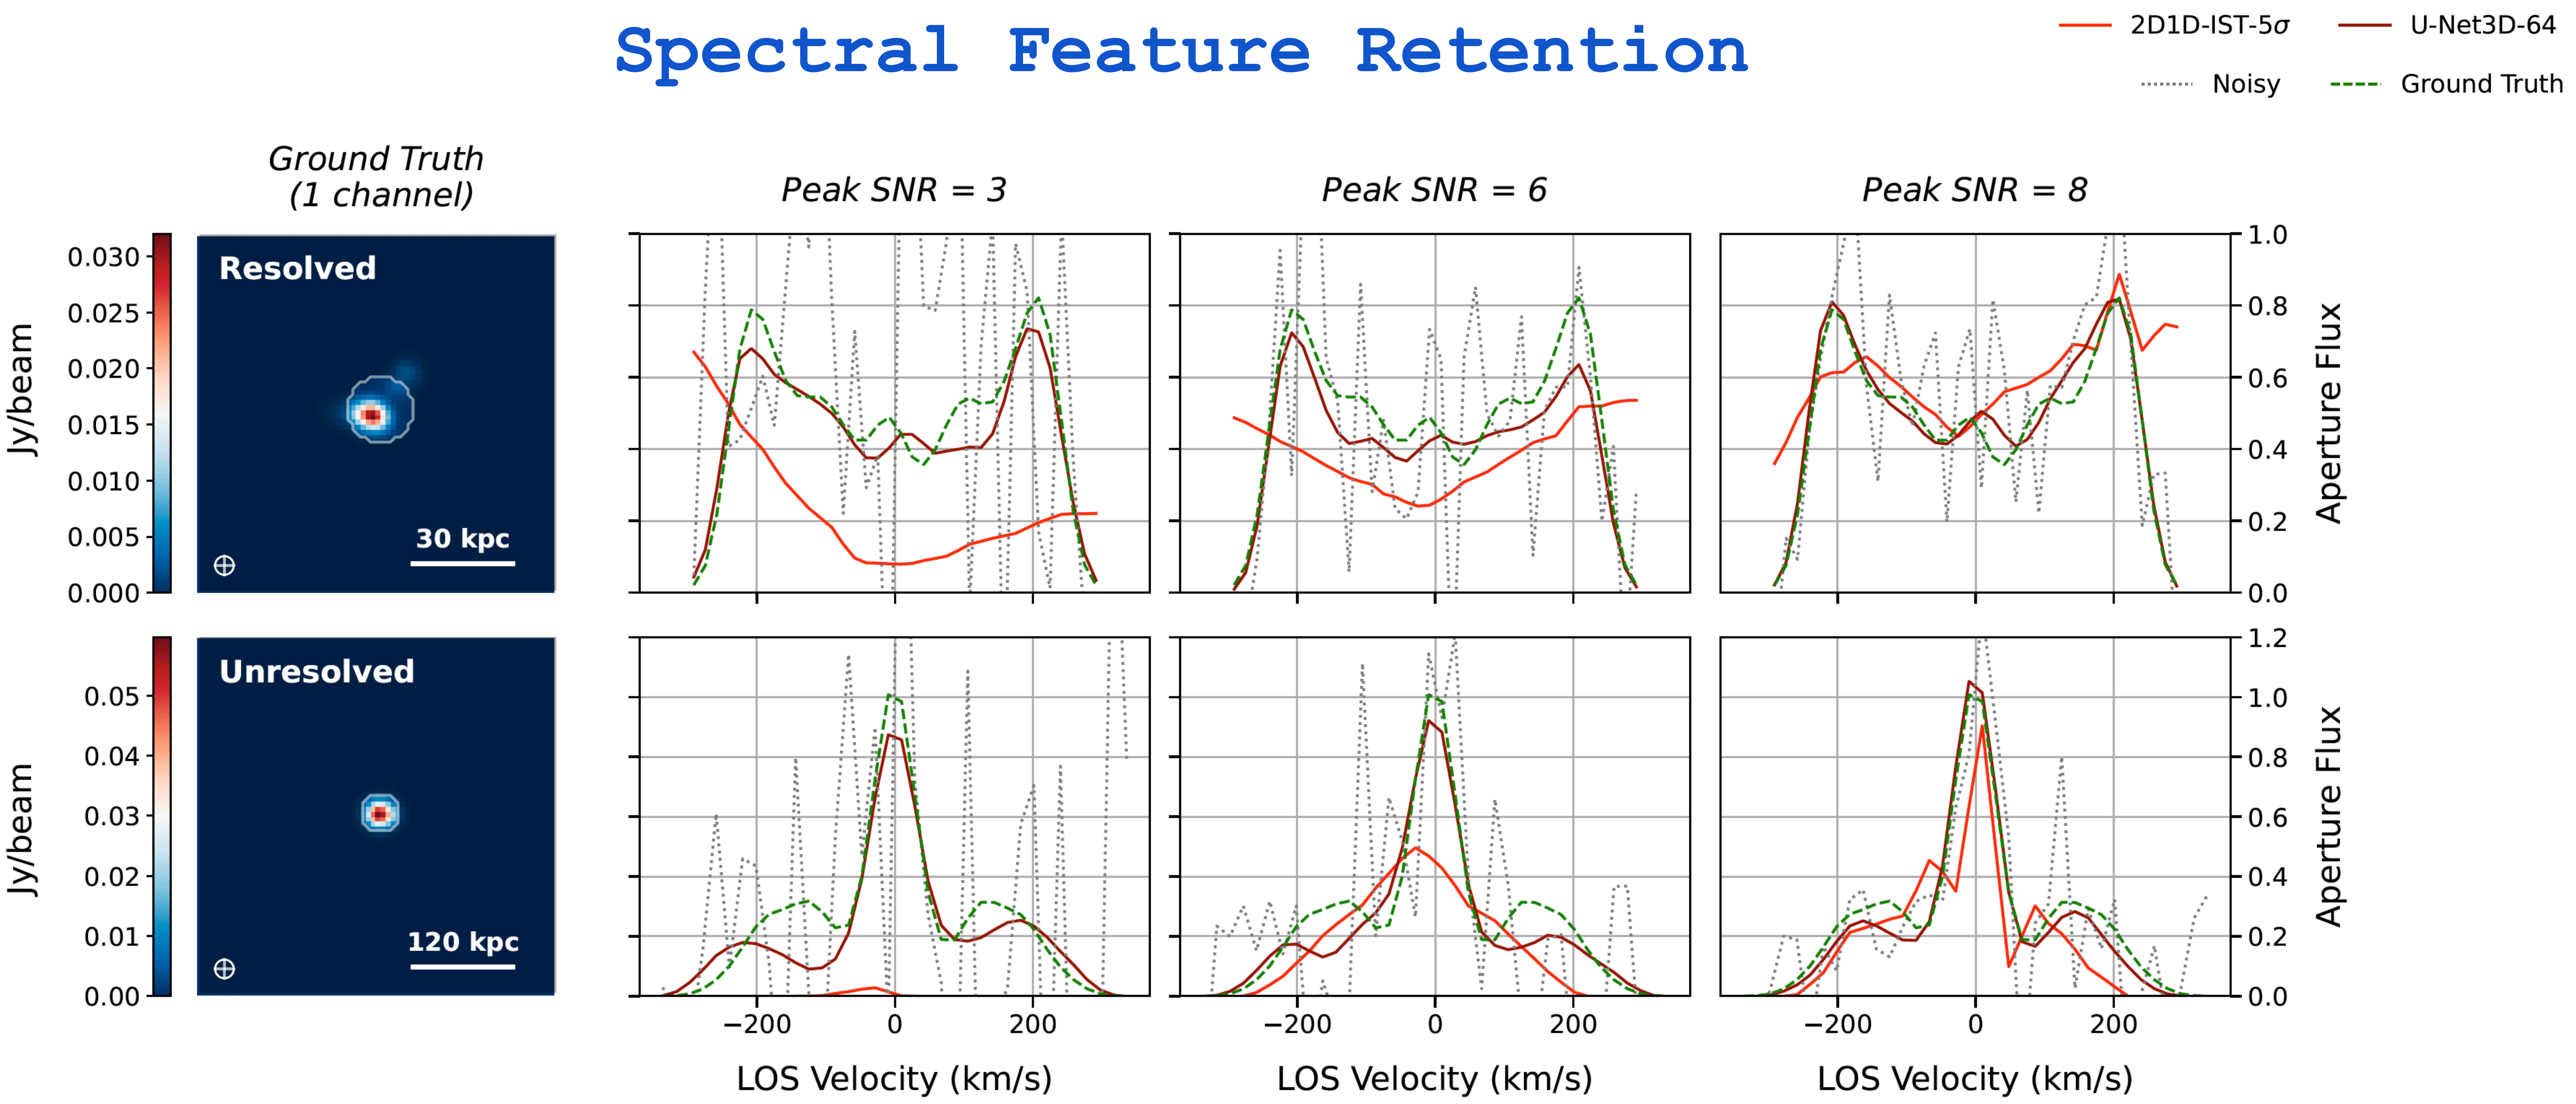

- **U-Net3D-64:**

  For both resolved and unresolved sources, even at low peak-SNR, the U-net is able to recover accurate spectral structures, and model kinematic features, getting better for high peak-SNR.

- **2D1D IST:**

  For both resolved and unresolved sources, reconstruction is poor at low peak-SNRs, but for higher peak-SNR, compact spectral shape of unresolved source is well preserved. For resolved, the shape is recovered at low peak-SNRs, but the overall spectrum (especially in channels with low flux) are not accurately recovered.

Deep learning challenges: Hallucinations

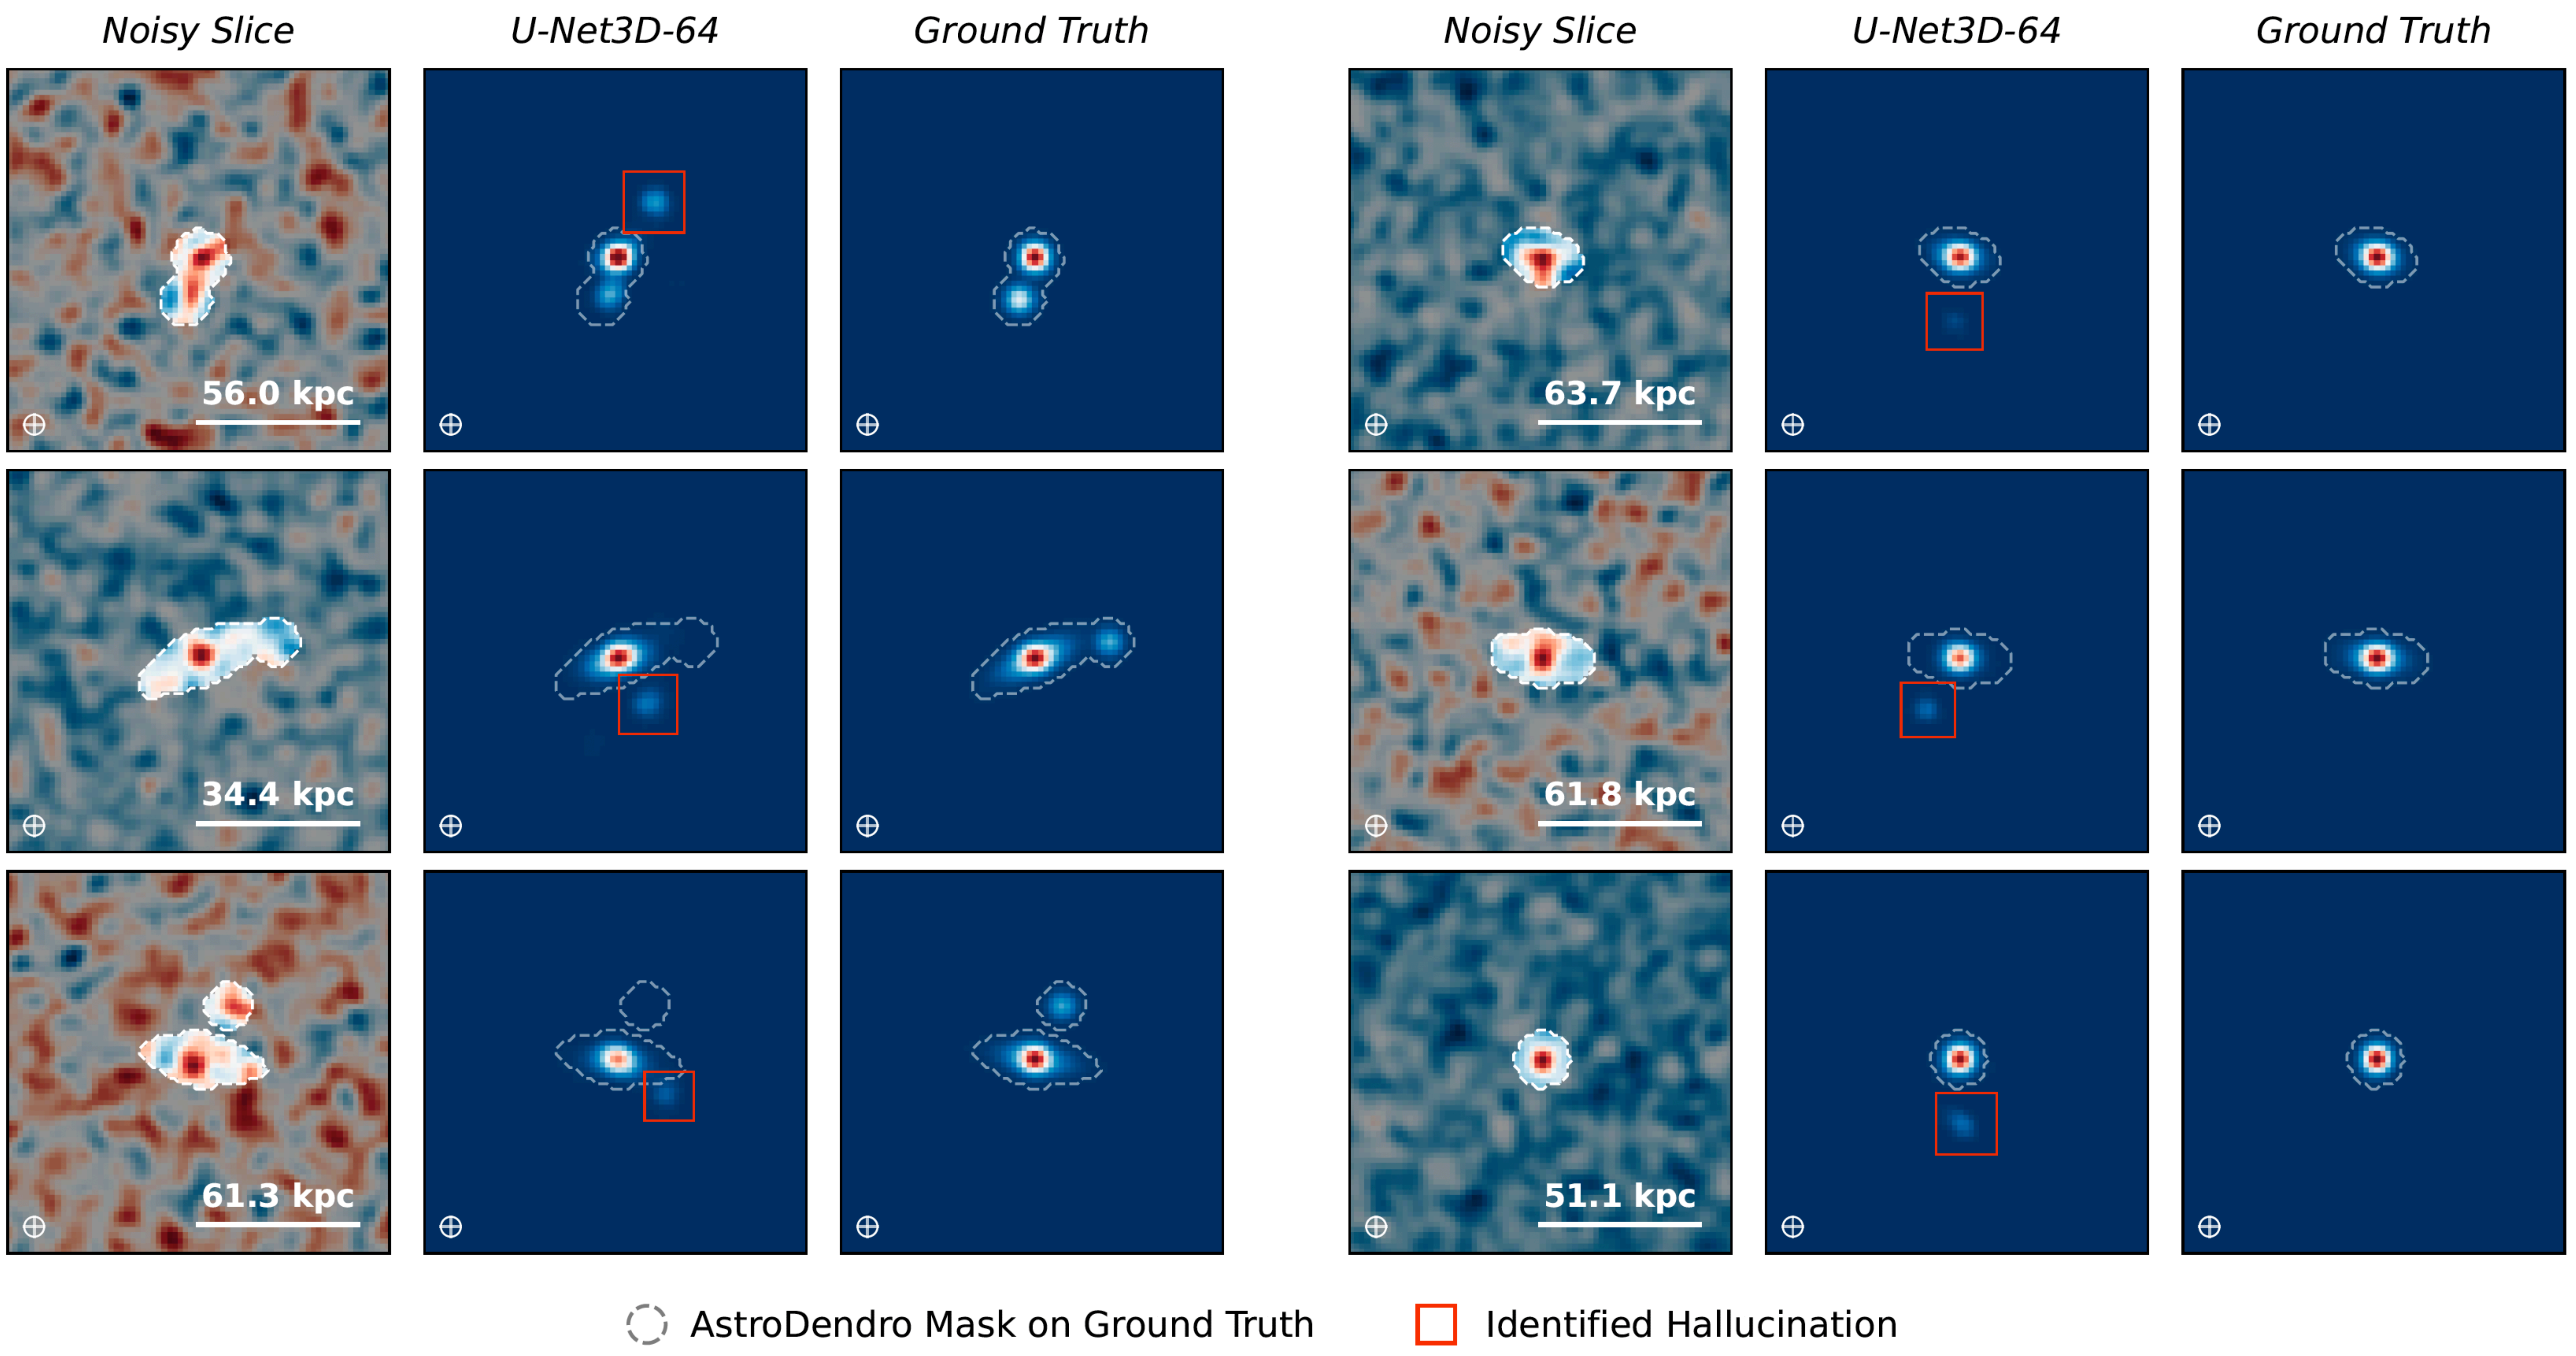

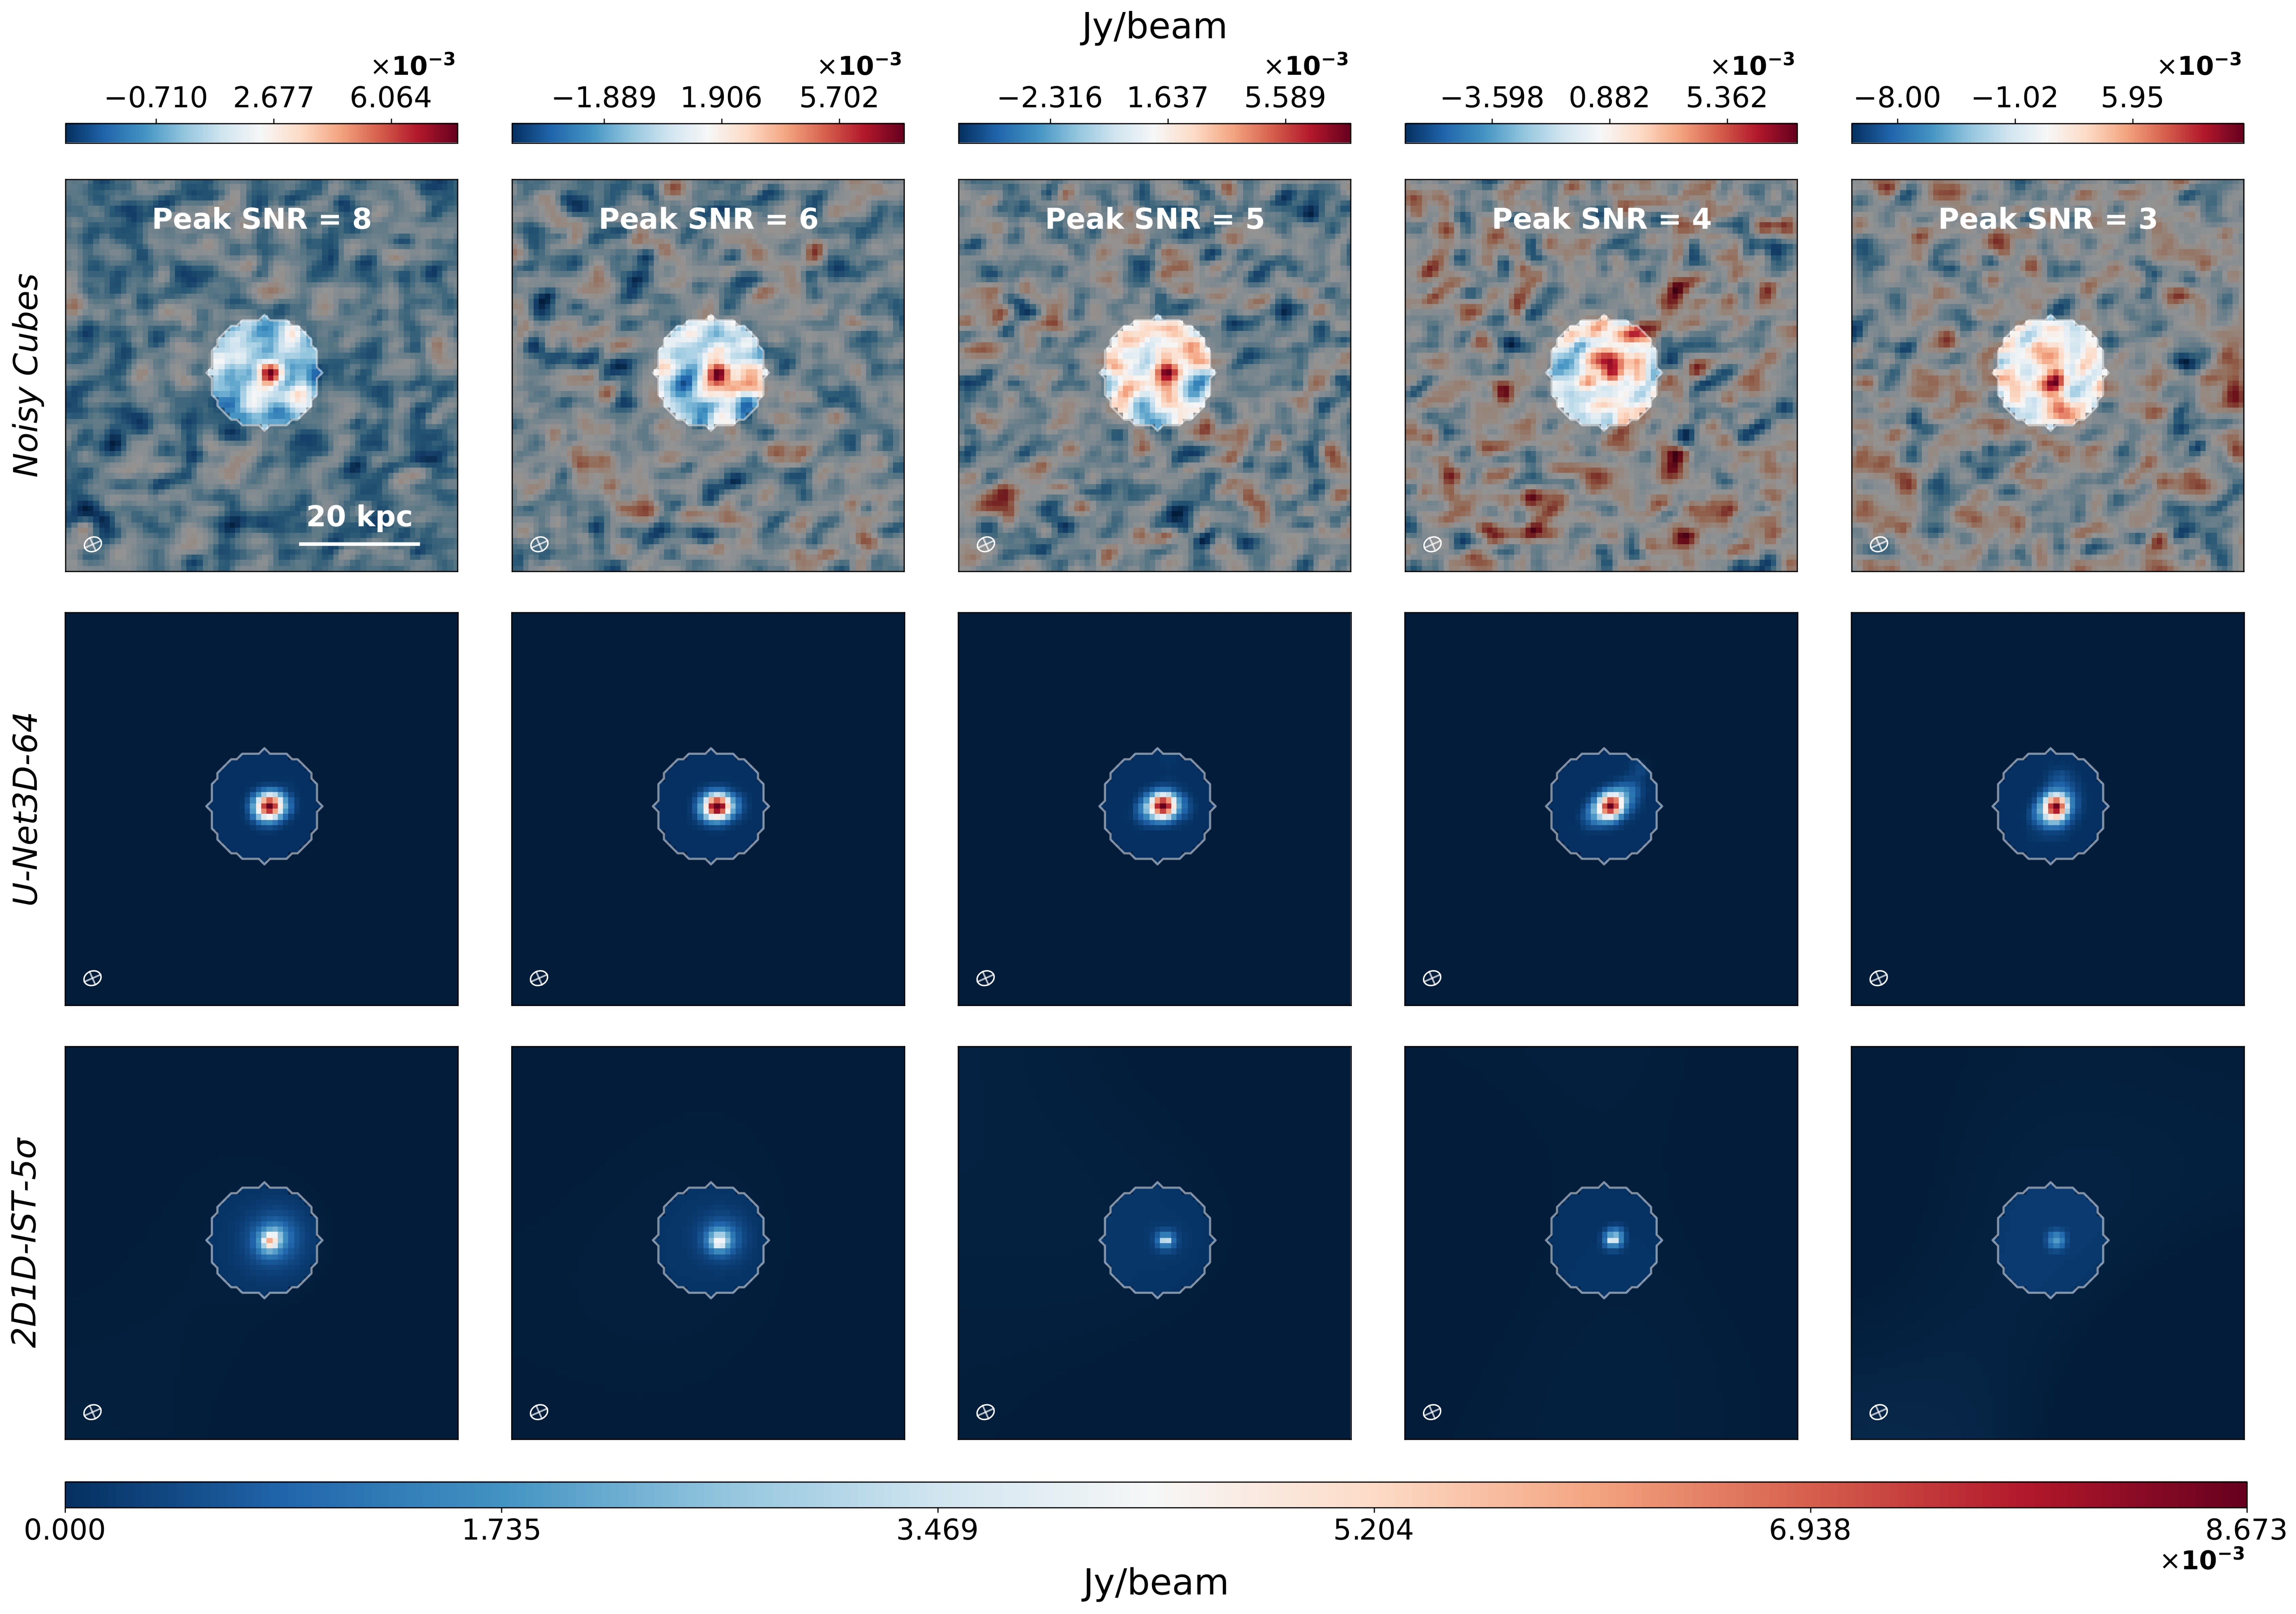

Application of methodology to **FIRE mock IFU & W2246 *CII IFU***

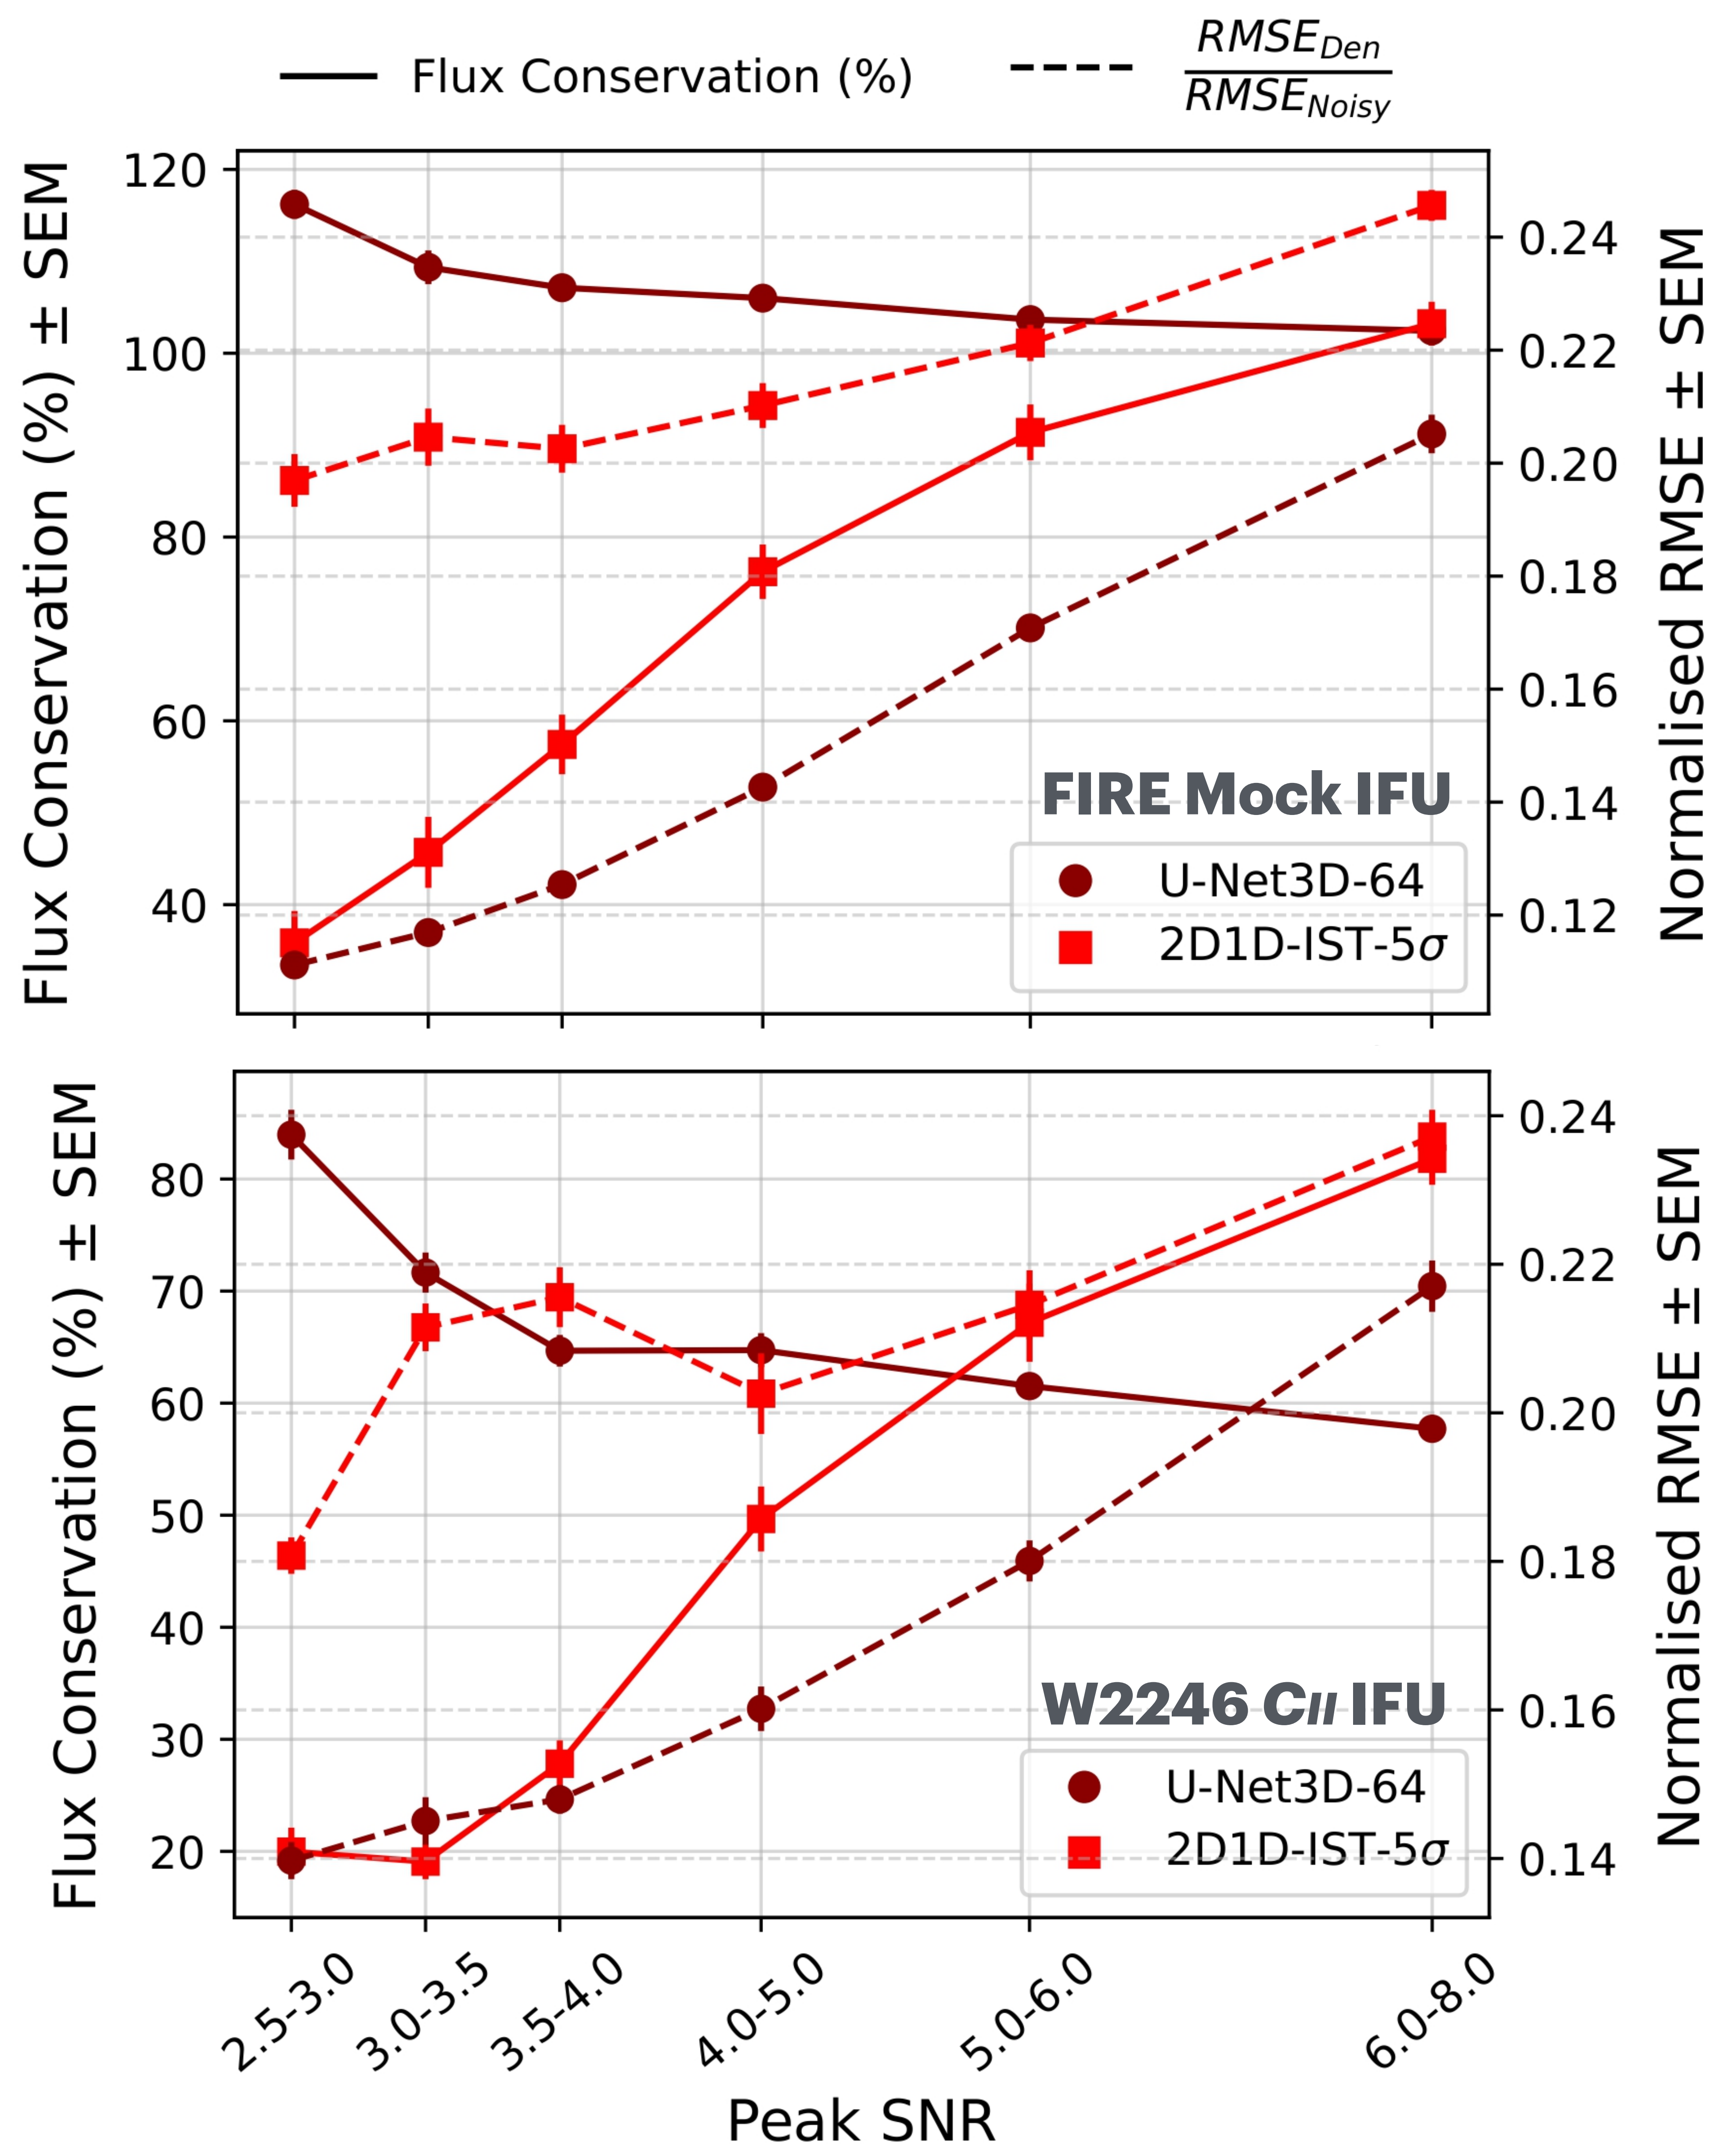

- **FIRE Mock IFU**
  - The RMSE ratio for both U-Net and 2D1D Wavelet-IST show expected trends.
  - The flux conservation trends are in line with what is expected for an unresolved central galaxy in the toy cubes. **Flux is very well estimated for higher peak-SNRs but for low peak-SNRs there is an overestimation of flux by ~ 10-15%**
  - **U-Net trained completely on toy cubes show promising generalisation to cubes with sophisticated modelling in cosmological simulations**

- **W2246 *CII* IFU**
  - The RMSE and flux conservation trends are similar to the previous case and what is expected for unresolved central galaxy BUT:
    - **The total flux as a whole is underestimated** in this case, despite systematic overestimation at low peak-SNRs
    - **W2246 has significant diffuse emission and faint morphological features** - tidal tails, bridges - no priors on the U-Net as it is trained on toy cubes
  - U-Net trained completely on toy cubes show promising results as it is **able to recover at least >50% of the total flux of a real ALMA observed high-redshift galaxy!**

Conclusions

- **Developed a multi-tiered evaluation framework**: from synthetic toy cubes → FIRE simulations → real ALMA (*W2246–0526*).
- **Created configurable toy datasets with controlled morphologies & noise**, ideal for training & benchmarking denoisers.
- Benchmarked 4 methods (**PCA, ICA, 2D1D-Wavelet-IST, 3D U-Net**) — highlighting their strengths & weaknesses:
  - **PCA/ICA**: limited; fail at noise suppression in low SNR.
  - **Wavelet-IST**: robust, interpretable, preserves flux well in medium-high SNR, struggles with faint diffuse emission.
  - **3D U-Net**: best overall performance; recovers structure even in low SNR, generalizes to real data — but tends to hallucinate features & overestimate flux in some cases.
- Demonstrated that **our toy dataset is powerful for training deep denoisers and can serve as a robust first stage in transfer learning**.

Future work:

- Use more **realistic priors in training set** (e.g., feedback & dust).
- Develop **uncertainty-aware & probabilistic denoisers** to quantify confidence.
- **Combine wavelets & deep learning** (e.g., **Learnlets**) for hybrid methods.
- Incorporate **regularization & loss terms to reduce flux bias & hallucinations**.
- Implement **hallucination detection & statistical uncertainty bounds** (e.g., **quantile regression and conformal predictions**).

(& thank you
for your time)