





Latent-space field tension for astrophysical component detection

An application to X-ray imaging

Matteo Guardiani^{1,2*}, Vincent Eberle^{1,2}, Margret Westerkamp^{1,2}, Julian Rüstig^{2,5},
Philipp Frank^{1,3}, and Torsten Enßlin^{1,2,4,5}

¹ Max Planck Institute for Astrophysics, Karl-Schwarzschild-Straße 1, 85748 Garching, Germany

² Ludwig-Maximilians-Universität München, Geschwister-Scholl-Platz 1, 80539 Munich, Germany

³ Kavli Institute for Particle Astrophysics & Cosmology, Stanford University, Stanford, CA 94305, USA

⁴ Excellence Cluster ORIGINS, Boltzmannstraße 2, 85748 Garching, Germany

⁵ Deutsches Zentrum für Astrophysik, Postplatz 1, 02826 Görlitz, Germany

Received 25 June 2025 / Accepted 24 September 2025

ABSTRACT

Modern observatories are designed to deliver increasingly detailed views of astrophysical signals. To fully realize the potential of these observations, principled data-analysis methods are required to effectively separate and reconstruct the underlying astrophysical components from data corrupted by noise and instrumental effects. In this work, we introduce a novel multifrequency Bayesian model of the sky emission field that leverages latent-space tension as an indicator of model misspecification, enabling an automated separation of diffuse, point-like, and extended astrophysical emission components across wavelength bands. Deviations from latent-space prior expectations are used as diagnostics for model misspecification, thus systematically guiding the introduction of new sky components, such as point-like and extended sources. We demonstrate the effectiveness of this method on synthetic multifrequency imaging data and apply it to observational X-ray data from the eROSITA Early Data Release (EDR) of the SN1987A region in the Large Magellanic Cloud (LMC). Our results highlight the method's capability to reconstruct astrophysical components with a high accuracy, achieving sub-pixel localization of point sources, robust separation of extended emission, and detailed uncertainty quantification. The developed methodology offers a general and well-founded framework applicable to a wide variety of astronomical datasets, and is therefore well suited to support the analysis needs of next-generation multiwavelength and multimessenger surveys.

Key words. methods: data analysis – techniques: image processing – techniques: photometric – ISM: supernova remnants – galaxies: structure – X-rays: general

1. Introduction

Astronomical observatories, however precise, inevitably imprint instrumental distortions onto the measured signals, corrupting the observational data. These systematic distortions – that stem both from the fundamental physics of measurement processes and from practical challenges inherent to observing faint and distant sources – degrade the astrophysical signals of interest. While the construction of more powerful instruments can partly alleviate these issues, such endeavors demand substantial investments of time and resources (Enßlin 2025). Moreover, the underlying challenge still remains: astronomical data typically contain more information than currently exploited by standard analysis methods. Therefore, it is essential to develop advanced methodologies capable of maximally extracting the information encoded in available observations.

A central challenge is that observational data consist of superpositions of multiple astrophysical components – including point sources (such as stars, quasars, and supernovae), extended compact objects (nebulae, supernova remnants), diffuse emission (intracluster gas, Galactic cirrus), and observational backgrounds – all blended together by the instrument response. Accurately characterizing each component requires

separating and reconstructing them from this mixture, but the task is complicated by shot noise, particularly in low-count X-ray and γ -ray regimes, as well as by instrumental blurring. The point-spread function (PSF), for example, describes how photons are redistributed by the optics across neighboring pixels, effectively convolving both diffuse and point-like emission and producing apparent superpositions. The redistribution matrix file (RMF; see Appendix D) further accounts for spectral blurring. At the same time, intrinsic astrophysical superposition arises when transparent structures along the line of sight overlap (for example, nebulae, supernova remnants, and diffuse Galactic foregrounds). Multifrequency or multiband observations exacerbate the complexity, as overlapping spectral features from distinct phenomena become intertwined spatially and spectrally.

Separating these mixed contributions to recover the underlying physical components is thus an ill-posed inverse problem: the data alone are insufficient to uniquely disentangle them, and additional prior knowledge is required. Here, we use the term “source separation” to mean the recovery of diffuse, compact, and point-like emission components from their superposed appearance in the data, while consistently accounting for both astrophysical and instrumental mixing. This capability is crucial for interpreting multifrequency survey data and for identifying new astrophysical phenomena. We present the formal inverse problem and its Bayesian solution in Section 2.

* Corresponding author: matteani@mpa-garching.mpg.de

A number of methods have been developed to extract sources and separate components in astronomical imaging, each with important limitations. Heuristic source finders like SExtractor (Bertin & Arnouts 1996; Bertin et al. 2020) or ProFound (Robotham et al. 2018) are widely used for their speed and simplicity. However they require manual parameter tuning, struggle to deal with complicated detector responses or noise, and provide limited uncertainty quantification. While our application in this work focuses on X-ray observations, similar challenges in source separation and image reconstruction arise in other wavebands, such as radio or infrared. In radio interferometry, the de facto standard deconvolution algorithm is CLEAN (Högbom 1974), which iteratively subtracts PSF-scaled point sources. CLEAN effectively assumes all emission is point-like; as a result, extended and diffuse structures are poorly recovered. Multi-scale extensions of CLEAN (Rau & Cornwell, 2011) – for example, versions that incorporate flux on multiple angular scales – offer an improved recovery for diffuse features, but still often struggle with complex, multicomponent scenes and require careful tuning of scale parameters.

In recent years, machine learning methods have also been applied to component separation and source detection. Deep learning models – typically convolutional neural networks (CNNs) – can learn to identify and remove foreground sources from complex backgrounds given sufficient training data. For instance, Madarász et al. (2025) demonstrate a U-net based network that subtracts compact sources from Herschel far-infrared maps, significantly improving the analysis of extended emission in star-forming regions. Similarly, Vafaei Sadr et al. (2019) use a CNN-based architecture to detect point sources in radio data and Casas et al. (2022) incorporate spectral training into their CNN architecture to detect point sources in microwave simulations. While such approaches show promise in specific cases, they usually require supervised training with simulated data, for example by injecting synthetic sources into real backgrounds as training pairs. This means the performance can be sensitive to the realism of the simulations and the method may not generalize outside the conditions on which it was trained (different instruments, noise levels, or source populations). Moreover, purely data-driven networks do not inherently provide uncertainty estimates or physical interpretability, and they tend to treat instrument characteristics as fixed nuisances – often one must train a new model for each telescope or wavelength.

More sophisticated approaches recast component separation as a probabilistic inference problem, allowing a principled treatment of noise, instrument effects by incorporating prior knowledge. Early examples include maximum-likelihood or maximum-entropy reconstructions for high-energy instruments (Skilling & Bryan 1984; Strong 2003), and sparse regularization techniques (wavelet thresholding, compressed sensing, and Gaussian mixture models) to de-noise and deconvolve astronomical images (Maisinger et al. 2004; Wiaux et al. 2009; Donath et al. 2024). These methods have seen success in specific contexts – for example, wavelet-based filtering has been used to detect faint sources in Planck and Fermi-LAT maps (Principe et al. 2018; Herranz et al. 2002) – but they often require choosing a suitable basis or regularization parameter by hand.

Given the shortcomings of existing techniques, there is a need for a general and robust framework to automatically separate astrophysical components across multiple wavelengths, without extensive manual tuning or instrument-specific tailoring. Bayesian component separation methods aim to incorporate physical priors more transparently. For instance, in radio astronomy, *resolve* and *fast resolve* (Arras et al. 2021; Roth et al.

2024) apply hierarchical Bayesian modeling within the framework of information field theory (IFT) (Enßlin et al. 2009; Enßlin 2013, 2019) – a Bayesian statistical field theory for signal inference – to surpass the limitations of traditional deconvolution algorithms such as CLEAN. However, these methods currently struggle to model point-like sources. In the high-energy regime, Guglielmetti et al. (2009) applied a Bayesian forward modeling approach to separate diffuse and point-like components in ROSAT data, though their spline-based background model was limited to a small number of sampling points. A significant step forward was achieved by the D3PO algorithm (Denosing, Deconvolution, and Decomposition for Photon Observations; Selig & Enßlin 2015), which leverages prior knowledge of the diffuse background’s spatial correlations and the statistical distribution of point-source brightnesses. D3PO enables simultaneous denoising, PSF deconvolution, and component-wise reconstruction in X-ray and γ -ray observations. Since then, similar techniques based on IFT have been successfully employed to de-noise, deconvolve, and decompose the high-energy sky (Scheel-Platz et al. 2023; Westerkamp et al. 2024; Eberle et al. 2025). In this work, we extend the D3PO framework and the broader IFT-based toolset by introducing a novel approach that achieves automated component detection and separation through latent-space field analysis.

Our method leverages the statistical information encoded in the latent random fields of hierarchical Bayesian forward models, while explicitly incorporating instrumental effects – such as the point-spread function (PSF), exposure variations, and noise characteristics – into a unified model of the observation. Within this framework, localized tension in the latent fields, defined as statistically significant deviations from their prior expectations, serves as a diagnostic indicator for the presence of astrophysical components.

Unlike traditional workflows that rely on sequential deconvolution, background subtraction, or source-detection stages, our approach jointly infers the most probable decomposition of the data into multiple astrophysical components, conditioned on the physical priors and the observational model. Because the instrument response is cleanly separated from the signal model, the method remains agnostic to the specific telescope or survey, and can be applied across different observational settings and wavelengths without retraining or tuning.

Crucially, the IFT framework enables reconstruction of both the signal fields and their spatial and spectral correlation structures via hierarchical modeling of covariance hyperparameters. This allows the algorithm to learn, directly from the data, the morphology of diffuse emission (e.g., smoothly varying background versus small-scale fluctuations) as well as the spectral coherence of components across multiple bands. While the method introduced here is fully automated for point-like source detection, it also supports the identification and modeling of more extended or compact sources through the same latent-space diagnostics.

The proposed approach offers several key advantages. Generality and scalability are achieved using fields, which are resolution independent and by formulating the inference through efficient solvers (building on the NIFTy library (Selig et al. 2013; Arras et al. 2019; Edenhofer et al. 2024)) that can handle fairly large imaging data sets and arbitrarily gridded domains. The method natively supports multifrequency (multiband) data, solving for a consistent set of component maps across all input images – this improves separation of components that are spectrally mixed in conventional analyses, and enables cross-band

information transfer, for example by reinforcing a source detection using its signal in another band with a higher signal-to-noise ratio (S/N). By jointly modeling diffuse, point-like, and extended structures, the algorithm can disentangle overlapping sources of different morphology, yielding more accurate flux estimates and cleaner separation than single-technique methods. Another important benefit is in the low-count or low-surface-brightness regime: by combining data across frequencies and leveraging prior knowledge of correlation structures, the presented IFT-based reconstruction boosts the detectability of faint features that would be missed by threshold-based detectors. We demonstrate that our method produces de-noised, de-convolved, and decomposed images (see Eberle et al. 2025) that closely approximate the true sky components, while maintaining uncertainty quantification for each reconstructed field. In summary, this work introduces a powerful new tool for astrophysical imaging: a multicomponent, multiwavelength Bayesian reconstruction technique that separates superposed sources in a principled, instrument-agnostic way, pushing the state of the art in both sensitivity and fidelity of component recovery.

The rest of this work is structured as follows: in Section 2.1, we introduce the Bayesian inference framework and describe the hierarchical modeling techniques used for astrophysical component separation. We begin with a one-dimensional toy model to illustrate the method’s core principles in Section 2.2 and then extend the approach to two-dimensional imaging data in Section 2.3. Sections 2.4 and 2.5 present the extension to multifrequency data, including our prior modeling of diffuse, point-like, and extended sources across wavelength bands. In Section 2.6, we apply our method to real SRG/eROSITA X-ray observations of the LMC SN1987A region, demonstrating its ability to reconstruct and disentangle complex astrophysical structures. Finally, Section 3 summarizes our findings and outlines directions for future work.

2. Methods and applications

In order to accurately reconstruct astronomical signals from noisy and systematically biased data, we need to make assumptions on the properties of the signal and on the behavior of the observational instruments that collect the data. In this work, we formulate these assumptions using the principles of Bayesian inference. Specifically, since the signals we are interested in are well suited to be mathematically described by scalar or tensor fields, we use the IFT formalism.

2.1. Bayesian astrophysical imaging

The ultimate goal of astrophysical imaging is to infer the spatial and temporal distributions of physical observables of distant sources. These observables include – but are not limited to – electromagnetic radiation from astrophysical sources, its polarization, and velocity fields. IFT allows to reconstruct the posterior distribution of high-dimensional and continuous fields from finite-dimensional data leveraging Bayes’ theorem

$$\mathcal{P}(s|d) = \frac{\mathcal{P}(d|s)\mathcal{P}(s)}{\mathcal{P}(d)}, \quad (1)$$

where we have indicated the signal field with s and the data with d . We note that the signal field s could depend on position, time, energy, polarization, and other physical dimensions, and may consist of multiple components.

2.1.1. Sky brightness prior

In the context of imaging, we can, for instance, choose to encode in the prior $\mathcal{P}(s)$ our assumptions about the surface brightness distribution across the observed patch of sky. In this work, we address the problem of designing a prior for the sky brightness distribution across different wavelengths. Following the approach outlined by Westerkamp et al. (2024); Eberle et al. (2025), we formulate these priors in the language of Bayesian hierarchical forward models.

To provide a standardized interface for model implementation and inference, we formulate our models as generative models. This introduces a latent Hilbert space Ω_ξ and a deterministic mapping $\xi \mapsto s$ into the target space Ω_s , where the signal s is defined. The mapping is constructed so that the transformed variables in Ω_s follow the desired prior distribution $\mathcal{P}(s)$. In practice, Ω_s provides a finite-dimensional parametrization of the more general Hilbert space in which the theoretical signal field s resides. This enables practical inference while maintaining a well-defined connection to the underlying infinite-dimensional problem.

The latent variables ξ are assumed to follow a standard multivariate normal distribution, which makes sampling and gradient-based optimization straightforward. As emphasized by Knollmüller & Enßlin (2020); Frank et al. (2021), this construction greatly facilitates approximate inference in high-dimensional posteriors.

In general, the latent space Ω_ξ itself has no direct physical meaning; it is introduced as a mathematically convenient representation. Similar constructs are common in machine learning (ML), where latent spaces provide coordinate systems that capture hidden structure in complex data distributions (for instance, in variational autoencoders or transformers). The key difference is that, in generic ML applications, the generative mapping is usually learned from data, whereas in our Bayesian forward models it is parametrized explicitly. This means that, once the generative mapping is specified, the role of the latent variables can be interpreted in terms of the corresponding prior. For example, the “correlated field model” introduced in Section 2.4.1 defines a transformation from Gaussian white noise in Ω_ξ into a Gaussian process prior – a smooth family of functions with interpretable correlation structure. This explicit construction makes it easier to connect the latent representation to astrophysically meaningful priors. We will exploit this property later to define our notion of latent-space model tension.

This framework underpins the central challenge addressed in this work: detecting, separating, and reconstructing astrophysical signal components. Because the observed data consists of a noisy and systematically biased superpositions of multiple components, reconstructing and separating them is an inherently ill-posed problem. That is, an infinite number of signal field configurations are consistent with the same data realization. As a result, the component separation task becomes entirely prior-driven, since the data alone cannot discriminate between different plausible solutions. This motivates the need for prior models that better encode our knowledge of complex astrophysical signal distributions.

2.1.2. Likelihood

Having introduced the role of the prior $\mathcal{P}(s)$ for a given signal of interest s , we now turn to the likelihood $\mathcal{P}(d|s)$, which connects the signal to the observed data d . In this work, we define the signal space to encompass physically plausible configurations of

the observed astrophysical fields, making the prior independent of any instrumental effects introduced during measurement. To model how the signal s gives rise to the observed data d , we account for the instrument's response function,

$$R(s) := \langle d \rangle_{(d|s)},$$

which describes the expected data measurement corresponding to a given signal configuration under the assumption that the noise has zero mean. This involves constructing a forward model that accurately represents how the signal is processed by the instrument. For efficient signal reconstruction, the response model must be both computationally scalable and mathematically differentiable. We denote the instrument's response operator with R . For X-ray observatories, R is typically approximated by chaining three operators: the point spread function (PSF) operator O to an exposure operator E and a masking operator M , as discussed in Eberle et al. (2025),

$$R = M \circ E \circ O.$$

While the PSF represents the conditional probability distribution of measuring a photon at a specific location in the field of view given its true origin, the exposure encodes the observation time, effective area, and vignetting, and the masking operator M accounts for detector characteristics such as dead pixels or regions obscured by hardware components. Additional effects such as pile-up, charge-coupled device (CCD) contamination, and energy-dependent redistribution can also influence the instrument response. In principle, including these effects in the response operator R would increase the accuracy of the reconstructed signal s , but at the cost of greater computational complexity. As this work focuses on developing efficient prior models, we restrict our analysis to the systematics listed above and neglect the response matrix file (RMF) in the forward response model. As shown in Appendix D, the eROSITA RMF is nearly diagonal in the 0.2–4.5 KeV band, with crosstalk between adjacent bins at the percent level or below. This spectral bleeding is therefore negligible compared to statistical uncertainties and PSF-induced mixing. We note that, although in principle any effect with a mathematical description can be incorporated into the forward response model, the feasibility of doing so depends on computational cost. Since the forward model is called repeatedly within the inference scheme, each evaluation needs to remain inexpensive in order to keep overall runtimes tractable. As a rule of thumb, maintaining inference times of order hours to days requires that a single call to the forward model completes within seconds.

Importantly, the response model R is independent of the astrophysical signal s , so prior models developed for s are, in principle, transferable between instruments once R is specified. In practice, this requires an efficient implementation of R , since the forward model is evaluated repeatedly during inference. However, for typical X-ray imaging tasks such implementations are readily achievable (for example, via FFT-based convolutions and sparse operators). This has already been demonstrated beyond eROSITA, for example in spatio-spectral IFT imaging of SN 1006 with multiple Chandra/ACIS observations (Westerkamp et al. 2024).

To implement the instrumental response R into our forward models, we make use of the universal Bayesian imaging kit (J-UBIK; Eberle et al. 2024). For photon-count instruments like eROSITA, the likelihood of a specific number of counts d_i in a

pixel $i \in \{1, \dots, N\}$ given an observed rate $\lambda_i = (R(s))_i$ is given by the Poisson distribution

$$\mathcal{P}(d_i|s) = \frac{\lambda_i^{d_i}}{d_i!} e^{-\lambda_i}.$$

Finally, since – conditional to a fixed signal – the number of recorded photons in each pixel i is assumed to be statistically independent from each other, we can write the combined likelihood of all pixels as

$$\mathcal{P}(d|s) = \prod_{i=1}^N \mathcal{P}(d_i|s) = \prod_{i=1}^N \frac{\lambda_i^{d_i}}{d_i!} e^{-\lambda_i}. \quad (2)$$

For practical purposes, we work with the so-called information Hamiltonian of the likelihood, defined as

$$\mathcal{H}(d|s) := -\log \mathcal{P}(d|s) = \sum_{i=1}^N [\lambda_i(s) - d_i \log \lambda_i(s) + \log d_i!],$$

where the final term, $\log d_i!$, is independent of the signal s and can therefore be neglected during inference. This formulation is proportional to the well-known Cash statistic (Cash 1979), commonly used to deal with Poisson-distributed data from astrophysical observations.

2.1.3. Posterior

Once the prior and likelihood distributions for the fields of interest are specified, Bayes' theorem for fields (Eq. (1)) can be used to determine the posterior probability distribution of the signal s , conditioned on the observation d . However, directly evaluating the posterior can be computationally prohibitive. The main challenge lies in calculating the evidence term

$$\mathcal{P}(d) = \int_{\Omega_s} \mathcal{P}(d|\tilde{s}) \mathcal{P}(\tilde{s}) \mathcal{D}\tilde{s},$$

which requires integrating over the parameter space Ω_s of the signal field s . Here, $\mathcal{D}\tilde{s}$ denotes the functional volume element, that is, integration over all possible configurations of the signal field. For the astrophysical applications discussed in this work, Ω_s is multi-million dimensional, making the evaluation of this integral infeasible with direct methods.

To overcome this computational barrier, we employ variational inference (VI), which approximates the posterior distribution efficiently and scales to high-dimensional parameter spaces. In VI, rather than directly evaluating the evidence $\mathcal{P}(d)$, we approximate the posterior distribution using a family of distributions $Q_\alpha(s|d)$, parameterized by the variational parameters α . This approximation is achieved by minimizing the Kullback-Leibler divergence (KL)

$$\mathcal{D}_{\text{KL}}(Q_\alpha \| \mathcal{P}) := \int_{\Omega_s} Q_\alpha(\tilde{s}|d) \log \frac{Q_\alpha(\tilde{s}|d)}{\mathcal{P}(\tilde{s}|d)} \mathcal{D}\tilde{s},$$

with respect to α . In this work, we approximate the posterior distribution $Q_\alpha(s|d)$ using geometric variational inference (geoVI, Frank et al. 2021). Following metric Gaussian variational inference (MGVI, Knollmüller & Enßlin 2020), geoVI extends traditional variational inference by leveraging the Fisher information metric to better capture the local geometry of the posterior distribution. The Fisher metric reflects the curvature of




Fig. 1. 1D example of the automatic component modeling method. *The top panel* displays the background b (in blue) and the total signal s (in orange) which is the superposition of the background b and the foreground line f signal components. *The bottom panel* shows the signal response (solid green line), obtained by convolving the total signal s with a Gaussian PSF. Additionally, it displays a synthetic data realization of the signal response (in red), along with the 2σ shot noise level contours associated with the background (dotted green lines). In both panels the vertical dashed-dotted gray lines represent the ground truth line component f locations. We note that a large fraction of the line signal f is buried beneath the background noise level.

both the likelihood and prior, providing a natural measure of the posterior’s geometric structure. By utilizing this metric, geoVI constructs a local isometry – a transformation that maps the curved parameter space to a Euclidean space while preserving its intrinsic geometric properties. In this transformed space, the posterior distribution is closer to Gaussian, enabling a more accurate Gaussian variational approximation. This approach allows geoVI to represent non-Gaussian posteriors with higher fidelity compared to MGVI, improving the quality of inference for complex Bayesian models. The ability to exploit the posterior’s geometry is particularly valuable for the astrophysical component separation problem addressed in this work, for which accurate modeling and inference are critical.

2.2. A one-dimensional example

In this work, we aim to exploit latent-space information to improve component separation and detect faint foreground structures from noisy data. To illustrate the core ideas, we begin with a simple one-dimensional example. Assume the total signal field of interest $s(x)$ consists of a positive, diffuse, correlated

background component $b(x)$, and a localized, uncorrelated line component $f(x)$

$$s(x) = b(x) + f(x).$$

The task is to reconstruct the total signal s and, importantly, separate the background emission b from the line component f from noisy data. We show an example of this task in Fig. 1.

To model the background signal $b(x)$ we use a lognormal process

$$b(x) = e^{\tau(x)}, \quad (3)$$

where $\tau(x)$ is the Gaussian process

$$\tau \sim p(\tau|T) = \mathcal{G}(\tau, T) := \frac{1}{\sqrt{|2\pi T|}} \exp\left(-\frac{1}{2}\tau^\dagger T^{-1}\tau\right),$$

with covariance T . The lognormal distribution of b ensures positivity and naturally captures variations on logarithmic scales,

which are characteristic of astrophysical emission signals. Following Arras et al. (2022), and assuming prior spatial homogeneity and isotropy, we adopt a flexible yet computationally efficient parametrization of the covariance T in harmonic space, known as the correlated field model. The generative model for $\tau(x)$ can be written as

$$\tau_x = \mathcal{F}_{xk} \sqrt{\hat{T}_{kk'}} \xi_{k'}, \quad (4)$$

where ξ_k denotes the latent-space excitations of the Gaussian process in harmonic space Ω_k , indexed by modes $k \in \Omega_k$, and \hat{T} is the Fourier transform of the covariance T . In Eq. (4), we have employed Einstein's summation and integration convention over repeated indices, which we will continue to adopt throughout this work. The Fourier transform operator \mathcal{F}_{xk} establishes the connection between field's support in position space Ω_x and in harmonic space Ω_k . More generally, for a field defined on vector spaces with coordinates \mathbf{x} and \mathbf{k} , the (inverse) Fourier transform operator acts as

$$\mathcal{F}_{\mathbf{xk}} \tau_{\mathbf{k}} := \frac{1}{(2\pi)^D} \int_{\Omega_k} e^{i\mathbf{k}\cdot\mathbf{x}} \tau_{\mathbf{k}} d\mathbf{k} = \tau_{\mathbf{x}},$$

where D is the number of dimensions of Ω_x and Ω_k . It maps fields defined in the spatial domain to their representation in the frequency domain, enabling efficient manipulation and analysis of spatial correlations, such as those encoded in the covariance structure T . In particular, the assumption of statistical homogeneity and isotropy diagonalizes the Fourier-transformed covariance

$$\hat{T}_{kk'} = (2\pi)^D \delta(\mathbf{k} - \mathbf{k}') P_T(|\mathbf{k}|),$$

where $P_T(|\mathbf{k}|)$ is the prior isotropic power spectrum of the process. In our standardized generative models, the latent-space excitations ξ are standard-normally distributed

$$\xi \sim \mathcal{P}(\xi) = \mathcal{G}(\xi, \mathbf{1}).$$

To model the line signal $f(x)$ we use a sum of Gaussian profiles centered at N uniformly sampled $x_i \in \Omega_x$ locations (in the example in Fig. 1, $N = 30$)

$$f(x) = \sum_{i=1}^N \frac{f_i}{\sqrt{2\pi} \sigma_f} \exp\left(-\frac{(x-x_i)^2}{2\sigma_f^2}\right)$$

with fixed variance σ_f^2 and intensity f_i sampled from an inverse-Gamma distribution

$$f_i \sim \mathcal{P}(f_i|\alpha, q) = \frac{q^\alpha}{\Gamma(\alpha)} f_i^{-\alpha-1} e^{-\frac{q}{f_i}},$$

where α is known as the shape parameter and q controls the scale of the intensity.

To mimic a very rudimentary astronomical response, the total signal s is then convolved with a Gaussian PSF and multiplied by an exposure time E , constant across the spatial dimension. We refer to the resulting integrated flux, from which the Poisson counts are drawn, as the signal response, shown in Fig. 1 together with the synthetic data realization.

Separating the two components is challenging, as the data only constrains their sum. This task is further complicated by noise and the unknown nature of the line component f . In particular, neither the number nor the locations of the lines are known

a priori, and many may lie below the noise level, making them effectively indistinguishable from the diffuse background b . In order to circumvent this problem, we need to introduce missing information encoded as prior knowledge. Specifically, we impose smoothness assumptions on the background component through a prior on the power spectrum of the Gaussian field τ . This assumption is crucial for separating f from b , since the presence of sharp, localized lines explicitly violates the prior expectation of smoothness.

To begin reconstruction, we first infer the signal assuming it is fully described by the diffuse background model $b = e^\tau$ alone. Learning b implicitly involves inferring its correlation structure, that is, the power spectrum of τ (Arras et al. 2022). However, if the data contains signatures of unresolved or point-like features, the background model – built on the assumption of homogeneity and isotropy – will absorb these features into the continuous background field. This will have an impact on the inferred correlation structure. As a result, the posterior correlation structure will no longer reflect the original prior assumptions, but instead encode inhomogeneities and anisotropies introduced by the unmodeled foreground components. This leads to a degraded fit: the reconstructed background appears artificially rough, as it tries to explain both the diffuse and line-like components under the assumption of smoothness. The result is a contamination of the inferred power spectrum by white-noise-like features, typically visible as a flattening or kink at higher k modes. We illustrate such a background-only fit in Fig. 2.

While this fit may appear unsatisfactory due to the issues outlined above, it provides a powerful insight: by analyzing the latent excitation field ξ , we can detect where the current model assumptions (a single background with long-ranged correlations) break down. Recall that the log-normal field $b(x) = e^{\tau(x)}$ is constructed via a Gaussian process, which in turn is the convolution of white Gaussian excitations ξ_x with a correlation kernel defined by \sqrt{T} . If the data conforms well to the prior – that is, if the signal correlation structure was perfectly represented by the kernel – these latent-space excitations should remain close to zero, with a variance of one. However, if the data includes features that are incompatible with the assumed smoothness – such as uncorrelated line emission – the excitations will grow in magnitude at those locations. Thus, the excitation field ξ contains implicit information about violations of the background prior assumptions. In this way, we can use the results of a background-only model fit as a diagnostic in the latent space to identify and localize the line component f . Specifically, since we are interested in finding the locations at which possible model misspecification might happen, we look at the excitations in position space $\xi_x = \mathcal{F}_{xk} \xi_k$.

Keeping in mind that the excitations are a priori standard normally distributed, we can define a detection threshold σ_{thresh} in position space above which we consider a possible line detection. We present a deeper theoretical analysis of the latent-space representation of the signal in Appendix A. As discussed there, the detection threshold is a measure of model tension with the prior. The posterior mean $\langle \xi_x \rangle_{(\xi|d)}$ typically shows reduced variance compared to the prior, but exhibits spatial correlations. Smooth structures in the field that might have been excited by a single large ξ_x can be equally well represented in the posterior mean by several nearby correlated ξ_x s that are each less strongly excited. Noise, the PSF, and the averaging inherent in the posterior mean further smooth and dilute such features.

For this reason, we calibrate the threshold empirically, by using synthetic reconstructions to balance false-positive and false-negative detections. For the examples at hand, we found




Fig. 2. Background-only reconstruction. *The top panel* displays the posterior mean of the background-only reconstruction (solid blue line) in comparison to the ground truth background signal response (dashed black line). *The bottom panel* shows corresponding latent position-space excitations ξ (purple solid line). The detection level for the line component was set to $\sigma_{\text{thresh}} = 0.6$ in prior units and it is displayed in red (solid horizontal line). The detected line positions for the given threshold σ_{thresh} are displayed as dashed vertical gray lines in both panels.

that a value of $\sigma_{\text{thresh}} = 0.6$ provides a good compromise: values above this level indicate potential model misspecification such as unresolved line components. A lower threshold would make detection less conservative but increase false positives, while a higher threshold would risk missing faint true components.

After identifying potential model misspecification locations, we iteratively refine the model by introducing line components centered at these positions. The validity of the extended model can be assessed by monitoring the evidence lower bound (ELBO), which is computed using NIFTy and provides a principled criterion to determine whether the inclusion of additional components improves the overall fit. This procedure can be repeated until no further excitations exceed the detection threshold. The result of this process is shown in Fig. 3.

While lowering the detection threshold may help recover additional true line components, it also increases the likelihood of introducing spurious ones in regions without actual signal. Such false positives typically exhibit lower inferred intensity and larger posterior uncertainty, allowing them to be filtered out in a post-processing step. The efficiency of this detection process can be significantly enhanced by analyzing the response of the correlated field to a line-like perturbation in the data. This is

discussed in Appendix A, with a practical demonstration provided in Sect. 2.6.

Overall, we find that the final reconstruction not only closely matches the ground truth signal across nearly all locations, but also provides a clear separation between the background component b and the line component f . Moreover, with the chosen detection threshold, our method successfully recovers point sources below the 2σ background shot-noise level.

2.3. An imaging example

We now want to apply the same technique to imaging. Specifically, we focus on the task of separating diffuse and point-like emission, crucial for astrophysical imaging. We start by showing how to use the method on a synthetic-data example.

2.3.1. Imaging setup

For instance, we consider the synthetic example shown in Fig. 4. This example should entail all the main ingredients of an astrophysical imaging task with count data. For the diffuse emission field model $b(\mathbf{x})$, we use a two-dimensional extension of the

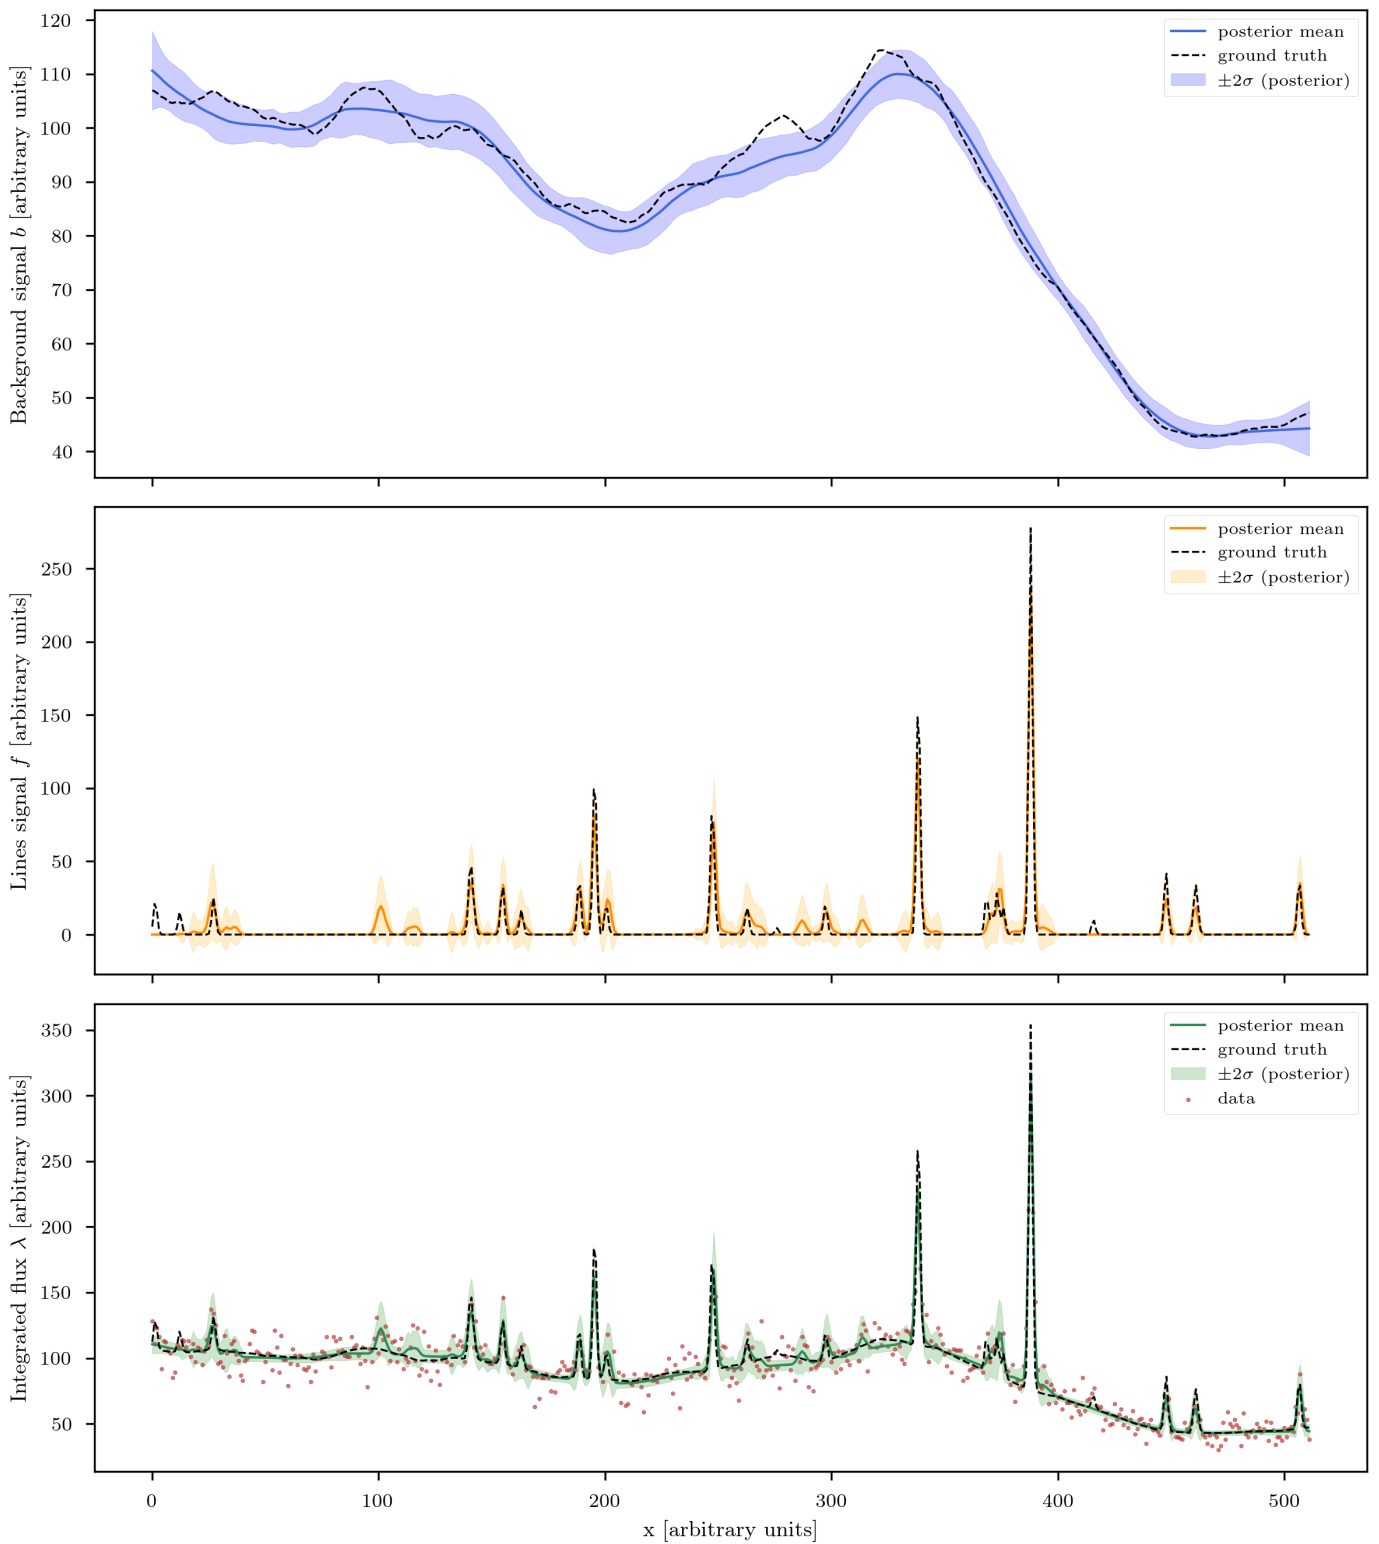


Fig. 3. Results of one-dimensional automatic component modeling method. *The top panel* displays the posterior mean of the background component b (solid blue line) in comparison to the ground truth background signal (dashed black line). *The central panel* depicts the posterior mean of the line component f (solid orange line) in comparison to the ground truth line signal (dashed black line). *The bottom panel* shows the posterior mean of the total signal s (solid green line) in comparison to the ground truth signal (dashed black line). The data is also shown in red. Shaded regions denote the $\pm 2\sigma$ posterior contours for each component in the respective panels.




Fig. 4. Setup of the synthetic component separation problem for imaging. *The top row* displays a realization of the diffuse emission component (left), the point-source component (middle), and the convolved point-source component (right), which has undergone Gaussian PSF convolution and multiplication by a spatially constant exposure of 80 000 s. *The bottom row* presents the corresponding sky emission, obtained as the sum of the diffuse and point-source components (left), the convolved and exposure-multiplied sky emission (middle), and data generated by a Poisson realization of the observed simulated sky (right).

lognormal correlated field model from Eq. (3), following the approach used in real imaging scenarios by Arras et al. (2022); Scheel-Platz et al. (2023). For the point-source emission field model we use a product of inverse-Gamma distributions

$$\mathcal{P}(f|\alpha, q) = \prod_{x_i} \frac{q^\alpha}{\Gamma(\alpha)} f_{x_i}^{-\alpha-1} e^{-\frac{q}{f_{x_i}}},$$

where $x_i \in \{0, \dots, N\}$ represent the regular grid points onto which the point-source field f is discretized, with N being the total number of pixels. We note that the inverse-Gamma prior can be made resolution-independent for $\alpha = 1.5$, since coarsening the resolution of the point-source field in this case still yields a field that is inverse-Gamma distributed, albeit with rescaled parameters¹. Following the one-dimensional example from Sect. 2.2, we construct the total photon flux by summing the PSF-convolved point-source and diffuse fields, and multiplying the result by a constant exposure time of $E = 80$ ks, assumed for simplicity to be uniform across the field of view. From this resulting flux field, we generate a Poissonian realization, which serves as the synthetic data shown in Fig. 4.

2.3.2. Imaging example inference

Our goal is to perform inference in this two-dimensional setting in a manner analogous to Sect. 2.2. As discussed in Sect. 2.3.1,

¹ This property was discussed by Selig & Enßlin (2015) and Guglielmetti et al. (2009), and rigorously analyzed by Girón (2001), where it is shown that under certain scaling relations, the inverse-Gamma distribution is stable under discretization changes.

we assume in this example that each pixel contains both a point-source component and a diffuse emission component. This is different from the one-dimensional case, where only a finite set of line components was present at random locations and most pixels contained only diffuse background.

In real astronomical data, we do not know a priori whether a point source is present in every pixel. Moreover, the true positions of point sources rarely align with pixel centers. To account for this, we model the point-source field in continuous space, representing each source explicitly by its position and flux. The point-source field is then given by

$$f(\mathbf{x}) = \sum_{i=0}^{N_{\text{ps}}} \delta(\mathbf{x} - \mathbf{x}_i) f_i, \quad (5)$$

where \mathbf{x} denotes the position of the i -th source and f_i its corresponding flux, with $i \in \{0, \dots, N_{\text{ps}}\}$. Because \mathbf{x} lies in continuous space, this representation is not restricted to pixel centers, enabling sub-pixel localization of sources from the data. A straightforward way to include this model in the likelihood would be to evaluate the PSF at each position \mathbf{x} and scale it by f_i . However, this approach can become computationally expensive – especially when the PSF varies significantly across the field of view (Eberle et al. 2024) – as it requires evaluating the PSF at updated positions during inference.

To mitigate this challenge, we adopt an alternative but equivalent strategy. Instead of evaluating the PSF directly at every source position, we bilinearly interpolate the point-source fluxes onto a regular pixel grid and subsequently convolve the resulting image with the local PSF. This procedure is equivalent in




Fig. 5. Interpolation of the point-source field flux onto a regular pixel grid. *The left panel* shows the interpolation result for two, unit-flux point sources located at $(0.5, 0.5)$ (center of the pixel) and at $(2.75, 2.6)$ on a 4×4 pixel grid. The location of the point sources is marked with a red cross. *The right panel* displays a random realization of 1000 uniformly distributed point sources on a 128×128 pixel grid with flux values uniformly distributed between 0 and 100. A white rectangle in the lower-left corner indicates a region of the same size as that shown in the left panel, providing a spatial scale reference.




Fig. 6. Result of background-only reconstruction. *Left:* posterior mean of the reconstructed background field after two geoVI iterations. Artifacts at the point-source length scale are visible (low-flux coronas around the point sources, wiggles in the diffuse field). *Center:* position-space excitation field ξ after two geoVI iterations. Detected point-source candidates above the $\xi = 0.65$ threshold are circled in lime. *Right:* reconstructed spatial power spectrum. The posterior mean is shown in blue, posterior samples in gray, and the ground truth in lime.

the case of spatially invariant PSFs, provided the PSF does not vary significantly within a pixel. For spatially varying PSFs, the situation is more complex, as the PSF itself may already be interpolated from a finite sampling. Nonetheless, our approach retains sub-pixel accuracy in source positioning and enables efficient convolution, substantially accelerating the likelihood evaluation. Figure 5 shows both an example of the interpolation procedure and a random realization of 100 point sources with uniformly distributed brightness, illustrating the structure of the prior over point-source fields under this model.

In this two-dimensional example, the inference must disentangle two overlapping contributions in every pixel: one from the diffuse background field and one from the latent point-source field. This is substantially more challenging than in the one-dimensional case, where only a finite set of localized line components had to be identified. Here, the algorithm must decide

for every pixel how much of the measured flux originates from the diffuse background and how much from point sources, which are not known in advance. The difficulty is further compounded by the fact that the contribution of a point source depends continuously on its true position on the sky, not only on its intrinsic brightness, and these positions are generally not aligned with pixel centers. If we now perform a geoVI data fit using our diffuse emission background field $b(\mathbf{x})$, we obtain a first estimate of the total sky emission field $s(\mathbf{x})$. The result of this procedure is shown in Fig. 6. Similarly to the one-dimensional example, we note that the reconstruction exhibits clear artifacts at the locations in which the diffuse field $b(\mathbf{x})$ is used to model point-source flux $f(\mathbf{x})$. Moreover, since the diffuse field model is forced to explain point-like structures, the inferred spatial power spectrum $\tau_{\mathbf{x}}$ exhibits excess power at high-frequency modes (see right panel of Fig. 6). In particular, a pronounced kink around $k \sim 35$




Fig. 7. Results of synthetic component separation example for imaging. *The top row* shows the reconstructed (posterior mean) diffuse-field emission (left), the reconstructed point-source component (middle), and the relative uncertainty of the sky reconstruction (right), defined as the posterior sample standard deviation divided by the posterior mean. Notably, the relative uncertainty is highest at locations containing point sources, as individual posterior samples place them at slightly different positions due to the inferred positional uncertainty. *The bottom row* presents the reconstructed sky emission as the sum of the diffuse and point-source components (left), the convolved and exposure-multiplied reconstructed sky emission (middle), and a comparison of the reconstructed spatial power spectrum (posterior mean in blue, posterior samples in gray) with the ground truth power spectrum (in lime, right).

– corresponding to the typical size of the reconstructed point sources – drives characteristic artifacts in the background reconstruction, such as the low-flux coronas surrounding the point sources visible in the left panel of Fig. 6. As a consequence, the reconstructed field (left panel in Fig. 6) will appear rougher than the original signal, since our prior on the diffuse field assumes statistical homogeneity.

By inspecting the latent-space excitations ξ in position space, we can identify locations of model misspecification – specifically, where a point source should be injected into the model. In this synthetic reconstruction example, we set a threshold of $\bar{\xi} = 0.65$. At locations where the posterior mean of ξ exceeds $\bar{\xi}$, we inject a point source component modeled as described in Eq. (5). We initialize the point source position at the detected location and set its flux to the integral of the diffuse field over one PSF length. This procedure is iteratively repeated to detect additional point sources while simultaneously optimizing their fluxes and positions as free parameters, jointly with the background field. After each detected point source is introduced, the diffuse latent field ξ is locally reset to standard-normal noise to ensure that overlapping flux is reattributed in subsequent updates.

At convergence – when the KL divergence is minimized – we obtain the reconstruction shown in Fig. 7. This required about 50 global geoVI iterations and completed in less than an hour on a 2020 M1 MacBook Pro. Examining this reconstruction, we find that the signal is accurately recovered down to the noise level. For additional reconstruction metrics, see Appendix B.1.

Moreover, the method successfully disentangles and reconstructs both the diffuse and point-source emission components. Importantly, the reconstruction model differs from the data-generating inverse-Gamma field. This mismatch illustrates that our method can generalize beyond its prior assumptions, making it robust to moderate model misspecification. This highlights the robustness of our method in realistic scenarios where model misspecification is inevitable.

2.4. A multifrequency sky model

An important aspect of this work is the construction of a sky model that captures structure coherently across both space and frequency. As more astronomical instruments operate across increasingly broad wavelength ranges, it becomes ever more important to leverage the rich correlations encoded in multi-wavelength and multimessenger observations. Modeling the sky independently at each frequency discards valuable physical information and leads to suboptimal inferences. Astrophysical processes – such as synchrotron emission, thermal bremsstrahlung, dust scattering, or line emission – imprint coherent patterns across both the spatial and spectral domains. Properly modeling these correlations allows us to regularize reconstructions, enhance sensitivity, and constrain astrophysical parameters more robustly. To capture these effects while maintaining computational tractability, we extend our generative modeling framework into the frequency domain. In this section, we present a model

that jointly describes spatial and spectral correlations in both diffuse, point-source, and extended-source emission, forming a consistent multifrequency sky prior.

2.4.1. Diffuse emission

In Sect. 2.2, we introduced a one-dimensional Gaussian process model for diffuse flux, based on the ‘‘correlated field model’’ developed by Arras et al. (2022). In their formulation, the correlation structure of the emission field is described via its amplitude spectrum along an axis i , which may represent space, time, or energy

$$A_{kk'}^{(i)} := (2\pi)^{D^{(i)}} \delta(k - k') \sqrt{p_{T^{(i)}}(|k|)},$$

where $p_{T^{(i)}}(|k|)$ is the power spectrum along the (i) direction and $D^{(i)}$ is the dimensionality of the i -th space. The full field’s correlation structure is then written as an outer product over all axes (e.g., spatial x and spectral ν)

$$A = \bigotimes_{i \in \{x, \nu\}} A_{kk'}^{(i)}.$$

This formulation is flexible and can describe many classes of diffuse emission. However, it implicitly assumes independence between axes, which may limit its ability to capture, for example, spatial variations in energy-dependent flux. Such correlations are especially important in X-ray and radio astronomy, where non-thermal emission from active galactic nuclei (AGN) – due to synchrotron radiation or inverse Compton scattering – often follows power-law spectra (Rybicki & Lightman 1986). To better capture these correlations, we introduce alternative priors that incorporate astrophysically motivated spectral and spatial dependencies. Specifically, we model the sky brightness at location \mathbf{x} and frequency ν as

$$I^{\text{diff}}(\mathbf{x}, \nu) = I^{\text{diff}}(\mathbf{x}, \nu_{\text{ref}}) \left(\frac{\nu}{\nu_{\text{ref}}} \right)^{\alpha(\mathbf{x})} I_{\delta}(\mathbf{x}, \nu), \quad (6)$$

where $I(\mathbf{x}, \nu_{\text{ref}})$ represents the spatial brightness distribution at reference frequency ν_{ref} , and $\alpha(\mathbf{x})$ is the spatially varying spectral index. The deviation field $I_{\delta}(\mathbf{x}, \nu)$ captures departures from the idealized power-law, due to effects such as thermal emission, spectral curvature, absorption features, or intrinsic source variability. In this model, $\alpha(\mathbf{x})$ varies spatially, while $I_{\delta}(\mathbf{x}, \nu)$ is a fully space- and frequency-dependent field. The multiplicative nature of I_{δ} with both the reference brightness and the power-law components introduces degeneracies. To break these, we enforce $I_{\delta}(\mathbf{x}, \nu_{\text{ref}}) \equiv \mathbb{1}$ and further remove degeneracy with the power-law term by parameterizing

$$I_{\delta}(\mathbf{x}, \nu) := e^{\tilde{\delta}(\mathbf{x}, \nu)} \left(\frac{\tilde{\nu}}{\nu_{\text{ref}}} \right)^{-\tilde{\delta}(\mathbf{x})}, \quad (7)$$

where $\tilde{\delta}(\mathbf{x}, \nu)$ is modeled as a Gaussian Markov process in the frequency direction – typically a Wiener process, an integrated Wiener process, or their combination. The term

$$\tilde{\delta}(\mathbf{x}) := \frac{\int_{\Omega_{\nu}} \tilde{\delta}(\mathbf{x}, \tilde{\nu}) \log \frac{\tilde{\nu}}{\nu_{\text{ref}}} d\tilde{\nu}}{\int_{\Omega_{\nu}} \left(\log \frac{\tilde{\nu}}{\nu_{\text{ref}}} \right)^2 d\tilde{\nu}}, \quad (8)$$

where Ω_{ν} is the frequency domain of I^{diff} , removes the average slope of $\tilde{\delta}$, thereby eliminating degeneracy with the spectral

index. (See Appendix C for derivation). We model the logarithm of the reference frequency brightness distribution $\log I^{\text{diff}}(\mathbf{x}, \nu_{\text{ref}})$ and the spectral index $\alpha(\mathbf{x})$ as correlated fields to learn their spatial and spectral covariance structure. In logarithmic form, the total brightness is expressed as

$$\log I^{\text{diff}}(\mathbf{x}, \nu) = \log I^{\text{diff}}(\mathbf{x}, \nu_{\text{ref}}) + \alpha(\mathbf{x}) \log \left(\frac{\nu}{\nu_{\text{ref}}} \right) + \log I_{\delta}(\mathbf{x}, \nu).$$

In this formulation, all components – spatial brightness, spectral index, and deviations – are parametrized in harmonic space, while the total brightness field $\log I^{\text{diff}}(\mathbf{x}, \nu)$ remains defined in real space and is reconstructed via an inverse Fourier transform

$$\log I^{\text{diff}}(\mathbf{x}, \nu) := \mathcal{F}_{\mathbf{x}\mathbf{k}} \left[\frac{\mathcal{A}_{\mathbf{k}}^{\text{spat}}}{\int \mathcal{A}_{\mathbf{k}}^{\text{spat}} d\tilde{\mathbf{k}}} \sigma^{\text{Iref}} \xi_{\mathbf{k}}^{\text{Iref}} + \frac{\mathcal{A}_{\mathbf{k}}^{\text{spec}}}{\int \mathcal{A}_{\mathbf{k}}^{\text{spec}} d\tilde{\mathbf{k}}} \left(\sigma^{\text{spec}} \xi_{\mathbf{k}}^{\text{spec}} \log \frac{\nu}{\nu_{\text{ref}}} + \delta_{\mathbf{k}\nu} \right) \right] + \alpha_0 \log \frac{\nu}{\nu_{\text{ref}}} + \log I_0. \quad (9)$$

This construction defines a hierarchical, nonstationary prior over $\log I^{\text{diff}}(\mathbf{x}, \nu)$, capturing both spatial and spectral correlations with explicit control over power-law structure and deviations. Here, $\mathcal{A}_{\mathbf{k}}^{\text{spat}}$ and $\mathcal{A}_{\mathbf{k}}^{\text{spec}}$ are the amplitude spectra encoding spatial and spectral correlations, respectively; $\xi_{\mathbf{k}}^{\text{spat}}$ and $\xi_{\mathbf{k}}^{\text{spec}}$ are the corresponding Gaussian field excitations. The parameters σ^{Iref} and σ^{spec} set the amplitudes of the fluctuations of the spatial brightness and spectral index. Lastly, $\log I_0$ and α_0 represent global offsets in brightness and spectral slope.

2.4.2. Point source emission

To further extend our prior modeling of the sky across multiple wavelengths, we now incorporate spectral information into our point-source emission model. We begin by introducing the point source brightness from Eq. (5) at the reference frequency ν_{ref} as

$$I^{\text{ps}}(\mathbf{x}, \nu_{\text{ref}}) = \sum_{i=0}^{N_{\text{ps}}} I_i^{\text{ps}}(\mathbf{x}, \nu_{\text{ref}}) := \sum_{i=0}^{N_{\text{ps}}} \delta(\mathbf{x} - \mathbf{x}_i) f_i,$$

where f_i denotes the flux of the i^{th} point source, typically drawn from an inverse-Gamma distribution to account for the heavy-tailed nature of astrophysical source populations. The extension to multiple frequencies is then natural. Drawing inspiration from the diffuse emission model in Eq. (6), we define the point source field across frequency as

$$I^{\text{ps}}(\mathbf{x}, \nu) = \sum_{i=0}^{N_{\text{ps}}} I_i^{\text{ps}}(\mathbf{x}, \nu_{\text{ref}}) \left(\frac{\nu}{\nu_{\text{ref}}} \right)^{\alpha_i^{\text{ps}}} I_{\delta, i}^{\text{ps}}(\nu). \quad (10)$$

Here, $\alpha_i^{\text{ps}}(\mathbf{x})$ is the spectral index of the i^{th} source, and $I_{\delta, i}^{\text{ps}}(\nu)$ captures deviations from a strict power-law spectrum. In contrast to the diffuse component, these spectral parameters are assumed to be uncorrelated across sources. This independence reflects the physical reality that different point sources often originate from distinct emission mechanisms or environments. As before, we address the degeneracy between the power-law component and its frequency-dependent deviation field by subtracting the average slope from $\tilde{\delta}_i(\nu)$, using the term $\tilde{\delta}$ defined in Eq. (8). This ensures that α_i^{ps} remains identifiable even in the presence of broad spectral deviations.




Fig. 8. Prior samples from the multifrequency sky model in Eq. (13). Each panel shows one realization of the sky brightness across three frequency channels (columns), with $N_{\text{ps}} = 1000$ point sources and $N_{\text{es}} = 3$ extended sources. The samples illustrate the interplay between diffuse, point-like, and extended components, all drawn independently from their respective priors. Each image is rendered as an RGB composite using three selected frequency bands, with intensity displayed on a logarithmic scale to enhance the visibility of both faint and bright features.

2.4.3. Extended source emission

So far we have built up a model to describe the distribution of diffuse and point-source emission across multiple wavelengths. However, many observed fields contain additional astrophysical structures – such as radio galaxies, galaxy clusters, or supernova remnants – that exhibit spatial extents and spectral behaviors distinct from the surrounding background or compact sources. To accurately capture these components, and to separate them cleanly from diffuse and point-source fields, we introduce a third component: the extended source field. These sources are neither spatially unresolved nor fully diffuse, and thus require a dedicated prior that reflects their intermediate nature. Following a structure analogous to the point-source model described above, we represent the extended-source emission as

$$I^{\text{es}}(\mathbf{x}, \nu) = \sum_{i=1}^{N_{\text{es}}} I_i^{\text{es}}(\mathbf{x}, \nu_{\text{ref}}) \left(\frac{\nu}{\nu_{\text{ref}}} \right)^{\alpha_i^{\text{es}}(\mathbf{x})} I_{\delta,i}^{\text{es}}(\mathbf{x}, \nu), \quad (11)$$

where $I_i^{\text{es}}(\mathbf{x}, \nu_{\text{ref}})$ denotes the spatial brightness distribution of the i -th extended source at the reference frequency ν_{ref} , with $i \in \{1, \dots, N_{\text{es}}\}$ and N_{es} denoting the number of extended sources in the field. Unlike point sources, which are modeled as Dirac delta functions, the extended sources are assumed to have nonzero spatial extent. Specifically, we follow the approach introduced in Rüstig et al. (2024) for modeling the spatial distributions of strongly lensed galaxies at the reference frequency ν_{ref}

$$I_i^{\text{es}}(\mathbf{x}, \nu_{\text{ref}}) = \gamma_i(\mathbf{x}, \boldsymbol{\theta}_i) I_i^{\text{es}}(\mathbf{x}), \quad (12)$$

where $\gamma_i(\mathbf{x}, \boldsymbol{\theta}_i)$ is a parametric envelope function depending on shape parameters $\boldsymbol{\theta}_i$, such as the mean (source location) and covariance (extent and ellipticity), typically instantiated as a Gaussian profile. The underlying $I_i^{\text{es}}(\mathbf{x})$ field is then modeled as a correlated Gaussian process, similar to our diffuse emission model. The spectral index $\alpha_i^{\text{es}}(\mathbf{x})$ governs the power-law scaling of each source's brightness with frequency, and may vary independently from source to source. Spectral deviations from the pure power-law form – arising from physical phenomena such as spectral curvature, absorption, or thermal emission – are described by the term $I_{\delta,i}^{\text{es}}(\mathbf{x}, \nu)$, defined analogously to Eq. (7). As with point sources, we assume the spectral parameters of

extended sources are spatially uncorrelated between sources, while allowing for internal spectral correlations across frequency within each individual source.

2.4.4. Multifrequency sky model

Taking all components into consideration, we can now assemble a comprehensive prior model for the sky brightness distribution across multiple wavelengths. This model remains highly flexible, yet explicitly incorporates spatial and spectral correlations in a physics-informed manner. Such a construction allows for realistic astrophysical priors while preserving computational tractability, and enables the reconstruction of signals with strong spectral-spatial entanglement – often present in real multifrequency observations.

Given a field of view containing N_{ps} point sources and N_{es} extended sources, the total sky brightness is modeled as the sum of three emission components

$$I^{\text{sky}}(\mathbf{x}, \nu) = I^{\text{diff}}(\mathbf{x}, \nu) + I^{\text{ps}}(\mathbf{x}, \nu) + I^{\text{es}}(\mathbf{x}, \nu). \quad (13)$$

Here, I^{diff} captures spatially correlated diffuse emission with spectral coherence, I^{ps} accounts for unresolved point sources with independent power-law scaling, and I^{es} models spatially extended objects with richer morphologies and spectral structure. Figure 8 shows representative prior samples drawn from this model. Visually, they resemble realistic astrophysical fields, capturing the spatial and spectral diversity found in broadband sky surveys. The model's flexibility is reflected in the diverse morphologies and spectral behaviors of sources. For example, in the left panel of Fig. 8, the sources exhibit markedly different spatial and spectral correlation lengths, while in the center panel, two of the extended sources are invisible due to their low brightness relative to the diffuse and point-source background.

2.5. Synthetic multifrequency imaging example

Now that we have defined a multifrequency sky model, we extend the synthetic example from Sect. 2.3 into the spectral domain. To this end, we use the sky distribution from that example as the reference-frequency brightness distribution in the




Fig. 9. Setup of the synthetic multifrequency component separation example in imaging. *Top row:* synthetic sky emission realizations for each frequency channel – $\nu_{\text{low}} = 0.2\text{--}1.0\text{ keV}$, $\nu_{\text{ref}} = 1.0\text{--}2.0\text{ keV}$, and $\nu_{\text{hi}} = 2.0\text{--}4.5\text{ keV}$ – shown in red, green, and blue, respectively. *Bottom row:* corresponding synthetic observations after convolution with a Gaussian PSF and multiplication by a spatially constant exposure time of 80 000 s.

multifrequency model of Eq. (13), and generate a synthetic realization for the associated spectral components. For simplicity, this example omits extended source emission; that is, we set $I^{\text{es}}(\mathbf{x}, \nu) \equiv 0$.

We divide the spectral domain into three nonoverlapping bands: $\nu_{\text{low}} = 0.2\text{ keV}$ to 1.0 keV , $\nu_{\text{ref}} = 1.0\text{ keV}$ to 2.0 keV , and $\nu_{\text{hi}} = 2.0\text{ keV}$ to 4.5 keV . The setup of this inference problem – including the ground-truth signal and the synthetic data across the three bands – is shown in Fig. 9. To reconstruct the sky brightness, we proceed similarly to Sect. 2.3, now using the full multifrequency diffuse model $I^{\text{diff}}(\mathbf{x}, \nu)$. For the frequency-dependent deviation field $\delta(\mathbf{x}, \nu)$ in Eq. (7), we model spectral variability as a Wiener process

$$\tilde{\delta}(\mathbf{x}, \nu) := \sigma_{\text{dev}} \int_{\Omega} \xi_{\mathbf{x}, \tilde{\nu}}^{\text{dev}} d\tilde{\nu},$$

where again $\mathcal{P}(\xi_{\mathbf{x}, \tilde{\nu}}^{\text{dev}}) = \mathcal{G}(\xi_{\mathbf{x}, \tilde{\nu}}^{\text{dev}}, \mathbb{1})$ is a white Gaussian field, and σ_{dev} steers the amplitude of the spectral deviations. Compared to the single-frequency example in Sect. 2.3, this multifrequency setup provides significantly more constraining power. In particular, spectral information greatly improves the ability to separate overlapping astrophysical components. This reflects the physical reality that it is extremely unlikely for two unrelated sources to share identical spectral characteristics. This additional information is encoded in the richer set of latent fields, including the spatial excitation $\xi_{\mathbf{x}}^{\text{spat}}$, the spectral-index excitation $\xi_{\mathbf{x}}^{\text{spec}}$, and the frequency deviations excitation field $\xi_{\mathbf{x}, \nu}^{\text{dev}}$.

To make use of the excitation fields for identifying potential model misspecification – such as locations where a point-source component may be missing – we seek a way to combine all latent-space excitations into a single scalar diagnostic. Specifically, we are interested in locations where some or all of

the posterior excitation fields deviate significantly from their prior distribution, as such deviations may indicate unmodeled structure in the data. We define the excitation-norm diagnostic as

$$r(\mathbf{x}) := \sqrt{\left(\xi_{\mathbf{x}}^{\text{spat}}\right)^2 + \left(\xi_{\mathbf{x}}^{\text{spec}}\right)^2 + \sum_{\tilde{\nu}} \left(\xi_{\mathbf{x}, \tilde{\nu}}^{\text{dev}}\right)^2}.$$

Since all excitation fields are a priori white Gaussian fields, the quantity $r(\mathbf{x})$ follows a Chi distribution with $N_{\nu} + 1$ degrees of freedom, where N_{ν} is the number of frequency channels in the model. Posterior deviations from this distribution can be used to flag candidate locations for inserting additional model components, such as unresolved point sources. When we reconstruct the sky signal shown in Fig. 9 using only the diffuse component $I^{\text{diff}}(\mathbf{x}, \nu)$, the resulting posterior excitation fields – together with the excitation-norm diagnostic – are shown in Fig. 10.

Investigating the reconstructed excitation-norm diagnostic field, we find numerous locations suggestive of model misspecification. In particular, because the diffuse model is forced to explain point-like features present in the data, the inferred spectral index map $\alpha(\mathbf{x})$ shows artificially sharp structures which resemble point sources. To address this, we use the excitation-norm diagnostic $r(\mathbf{x})$ to identify candidate locations for inserting point-source components. Specifically, we select all positions where $r(\mathbf{x}) > \bar{r}$, with a detection threshold of $\bar{r} = 0.6$. At these positions, we introduce point-source model components $I^{\text{ps}}(\mathbf{x}, \nu)$, drawn from the model described in Eq. (10).

The resulting improved sky reconstruction – with the model now including both diffuse and point-source components – is shown in Fig. 11. As expected, the additional information contained in the multifrequency data improves the reconstruction




Fig. 10. Result of diffuse-only reconstruction. *Top row:* posterior mean of the spatial excitation field $\xi_{\mathbf{x}}^{\text{spat}}$ (left), the spectral-index excitation field $\xi_{\mathbf{x}}^{\text{spec}}$ (center), and the inferred spectral index field $\alpha(\mathbf{x})$ (right) after reconstructing the sky using only the diffuse component. *Bottom row:* posterior mean of the spectral deviation excitation field $\xi_{\mathbf{x},\nu}^{\text{dev}}$ at ν_{low} (left), at ν_{hi} (center), and the excitation-norm diagnostic $r(\mathbf{x})$ (right), computed from the displayed excitation fields. Candidate point-source locations identified using the excitation-norm threshold $\bar{r} = 0.6$ are highlighted with lime-colored circles in the $r(\mathbf{x})$ panel.




Fig. 11. Results of the synthetic multifrequency imaging reconstruction example. *Left:* exposure-corrected synthetic data. *Middle:* ground-truth sky emission. *Right:* posterior mean of the reconstructed sky. All images combine the three energy bands into an RGB composite, where red corresponds to the ν_{low} band, green to ν_{ref} , and blue to ν_{hi} .

and allows for more accurate separation into distinct components. We note that our reconstruction model – where only explicitly detected point sources are modeled – differs from the synthetic data-generating model, which assumes a point source at every pixel. Consequently, the reconstructed diffuse background $I^{\text{diff}}(\mathbf{x}, \nu)$ must absorb the flux from undetected point sources. From a statistical perspective, if numerous low-flux sources contribute significantly to the inverse-Gamma distribution’s probability mass, their combined flux may introduce an

overall offset in the expected photon rate. This additional flux must then be compensated for by the diffuse background component, potentially leading to its overestimation. However, this scenario is somewhat artificial: in real observations, the fraction of background flux originating from unresolved sources is typically unknown, as it may also be genuinely diffuse or absorbed by intervening material, such as dust. For additional diagnostics and a quantitative evaluation of the reconstruction accuracy, we refer to Appendix B.2.

2.6. LMC SN1987A eROSITA reconstruction

We now demonstrate the applicability of our method using real X-ray data from the early data release (EDR) of the SRG/eROSITA observatory.

2.6.1. Data

To illustrate the capabilities of our method in a realistic astrophysical scenario, we apply it to an eROSITA pointed observation of SN1987A in the Large Magellanic Cloud (LMC). SN1987A, a well-studied supernova remnant, offers an ideal test case because of its complex combination of diffuse emission, compact sources, and extended structures, including shocks, interaction regions, and star-forming regions. We processed the raw photon-count data using the eROSITA Science Analysis Software System (eSASS), correcting for instrumental exposure, vignetting, and detector characteristics. The observations were centered at a celestial position of RA = 83.78°, Dec = -69.31° and binned into three distinct energy bands following the procedure outlined by Eberle et al. (2025): $\nu_{\text{low}} = 0.2\text{--}1.0\text{ keV}$, $\nu_{\text{ref}} = 1.0\text{--}2.0\text{ keV}$, and $\nu_{\text{hi}} = 2.0\text{--}4.5\text{ keV}$. For the present analysis, we utilized data from eROSITA telescope modules (TMs) TM1, TM3, and TM4. This selection was motivated by computational efficiency: each TM has a distinct instrumental response, so including all five available TMs would significantly increase computational costs. While using more TMs would improve the signal-to-noise ratio and could potentially enhance reconstruction accuracy and point-source detectability, we found that restricting the analysis to three modules offers a favorable trade-off between fidelity and runtime. Moreover, the primary objective of this application is not to achieve maximum reconstruction fidelity, but rather to showcase the latent-space tension diagnostic on real observational data. The resulting dataset, shown in the top-left panel of Fig. 13, provides the input to our inference.

2.6.2. Reconstruction

We initially reconstructed the sky emission using our multifrequency diffuse emission model (Eq. (6)). We implemented the multifrequency emission model and the eROSITA response using the J-UBIK package (Eberle et al. 2024). This first step identifies regions where the diffuse-only assumption is inadequate. To diagnose potential model misspecifications, we compute the excitation-norm diagnostic, $r(\mathbf{x})$, derived from the latent-space excitations. Elevated values in this diagnostic map, as shown in Fig. 12, highlight discrepancies between the assumed smooth, diffuse background emission model and actual astrophysical structures in the observed data.

The excitation-norm diagnostic clearly identifies two prominent extended X-ray sources: the Tarantula Nebula and 30 Doradus C. The accurate recovery of these well-known astrophysical features confirms the effectiveness of the latent-space diagnostic as a robust indicator for refining the emission model. Consequently, we enhance our reconstruction by explicitly incorporating two extended-source model components at these identified positions, modeled according to Eq. (12). Currently, extended sources are identified manually. Automating this step would reduce subjective bias, improve reproducibility, and streamline the analysis pipeline. Leveraging neural-network segmentation of radio data has yielded promising extended-source component separation. Future work might include supplementing these methods with latent-space information. We




Fig. 12. Excitation-norm diagnostic map highlighting regions of potential model misspecification after the initial diffuse-only reconstruction of the SN1987A eROSITA dataset. Bright areas represent elevated latent-space excitations, clearly identifying well-known extended sources such as the Tarantula Nebula and 30 Doradus C (highlighted with red boxes). These features motivate the explicit modeling of extended-source components.

note that explicitly modeling these extended-source components substantially reduces residual latent-space excitations, thereby enhancing the physical realism and interpretability of the reconstruction.

To speed up the inference and source detection processes, we initiate our final reconstruction incorporating the two extended source models by reconstructing the sky only at the reference frequency ν_{ref} ; that is, $I(\mathbf{x}, \nu_{\text{ref}}) = I^{\text{diff}}(\mathbf{x}, \nu_{\text{ref}}) + I^{\text{es}}(\mathbf{x}, \nu_{\text{ref}})$. We utilize latent-space tension at the reference frequency to iteratively introduce reference-frequency point sources, as described in Sect. 2.3. Upon convergence of this reconstruction, we perform a multifrequency reconstruction using the full model described by Eq. (13). Due to the clean factorization of our prior into a spatial component defined at the reference frequency and a spatio-spectral component, no intermediate transition step (such as the transition model implemented by Westerkamp et al. 2024) is required, significantly accelerating the reconstruction process. The inference proceeds iteratively, simultaneously optimizing for diffuse, point-source, and extended-source components until convergence. Point sources are introduced wherever significant features appear in the excitation-norm diagnostic map, as detailed in Sect. 2.5. To enhance the efficiency and accuracy of point-source detection, we apply a matched filter to the excitation-norm diagnostic map (see Appendix A.3 for details).

As detection criteria, we incorporate a point-like source into our forward model at locations where the excitation-norm diagnostic exceeds a chosen threshold (0.6 for unfiltered and 0.3 for filtered diagnostics). Convergence is defined as stabilization of the Kullback-Leibler divergence difference between consecutive iterations, accompanied by the absence of significant features in the excitation-norm diagnostic map over ten consecutive iterations. For the results presented here, the full multifrequency reconstruction required about 40 global geoVI iterations and took roughly 6 hours to complete on a single AMD Instinct




Fig. 13. Reconstruction results from the eROSITA early data release (EDR) observation of SN1987A in the LMC. Top row (left to right): exposure- and effective-area-corrected observed data, reconstructed extended emission associated with 30 Doradus C, and reconstructed extended emission from the Tarantula Nebula. Bottom row (left to right): reconstructed total sky signal (including all modeled components), reconstructed diffuse background emission (excluding point and extended sources), and reconstructed point-source emission field. The point source field has been convolved with a Gaussian kernel with a standard deviation of $4''$. All panels employ the same RGB color scheme, representing intensities in the three energy bands: red ($\nu_{\text{low}} = 0.2\text{--}1.0\text{ keV}$), green ($\nu_{\text{ref}} = 1.0\text{--}2.0\text{ keV}$), and blue ($\nu_{\text{hi}} = 2.0\text{--}4.5\text{ keV}$).

MI300A GPU. We expect further improvements in runtime with more optimized implementations of the response operator. We present the results of applying this method in the following section.

2.6.3. Results

The main results from our multicomponent reconstruction of the eROSITA early data release (EDR) observation of SN1987A in the LMC are presented in Fig. 13. All astrophysical source components are clearly separated, and no residual point-like structures remain visible within the reconstructed diffuse background. Given that we model the sky emission as a spatio-spectral field, we can extract the model prediction for the sky (or any component's) emission at any given wavelength. We show an example of the predicted sky photon flux at 2.4 keV in Fig. B.4 in Appendix B.3. Thanks to the explicit modeling of the extended sources 30 Doradus C and the Tarantula Nebula, we are able to recover their spatial and spectral correlation structures from eROSITA data with unprecedented detail. For a direct comparison with previous IFT-based reconstructions and Chandra observations covering the same region, we refer the reader to Eberle et al. (2025). Our approach enables sub-pixel localization of point-like structures and improves over earlier inverse-Gamma-based models by explicitly modeling source positions and allowing for cleaner disentanglement of

components. A complete list of the detected point-like sources, along with additional posterior diagnostics – including noise-weighted residual analysis – is provided in Appendix B.3. Lastly, we emphasize that our Bayesian inference framework not only yields a single reconstruction of the modeled signal fields but also quantifies uncertainties through posterior sampling. These posterior samples comprehensively characterize the inferred probability distribution of the reconstructed signals.

3. Conclusions

In this work, we presented a novel method for astrophysical component detection and separation in single- and multiband imaging data. Our approach combines a new multifrequency sky model with a latent-space field-tension diagnostic to automatically identify and disentangle point-like and extended emission from diffuse background structures. The model was validated on synthetic data, demonstrating accurate component separation and denoising. In addition to imaging applications, a one-dimensional variant of the method can be used, for instance, for line detection in noisy spectra.

We further applied the method to SRG/eROSITA X-ray observations of SN1987A in the LMC, successfully reconstructing diffuse, extended, and point-source components potentially with sub-pixel accuracy. Extended sources like the Tarantula

Nebula and 30 Doradus C were explicitly modeled, enabling detailed spatial and spectral reconstruction. By jointly analyzing multifrequency data, we recover spatially resolved spectral features and improve physical interpretability. Since the sky emission model is fully decoupled from the instrument response, the method is applicable to any telescope or wavelength band, making it well suited for future multimessenger analyses.

Looking ahead, IFT-based radio imaging methods such as `resolve` and `fast resolve` could benefit from the automatic point-source modeling introduced here. The Bayesian formulation naturally provides uncertainty quantification through posterior samples, enabling robust flux estimation and residual diagnostics. The multifrequency sky model and latent-space detection tools will be made publicly available in J-UBIK².

4. Future work

Several enhancements to our method are currently under development to extend its performance and applicability. A key priority is calibrating point-source detection to achieve catalog-grade completeness and purity, validated through simulations. We plan to refine detection diagnostics, incorporating matched-filter techniques to improve sensitivity, particularly for faint and crowded fields.

Automating extended-source identification is another important goal. Leveraging latent-space diagnostics, we will develop criteria for algorithmically detecting and modeling extended astrophysical sources, removing manual intervention.

Instrumental modeling improvements, including continuous spatial interpolation of the PSF for the point sources, will be integrated into the forward model, ensuring more precise reconstructions. Additionally, we plan to explore broader astrophysical applications, such as cosmic microwave background source detection, spectral-line identification in spectroscopic data, exoplanet transit searches in multimessenger data, and more. These ongoing developments aim to provide a robust, versatile, and fully automated Bayesian imaging tool, broadly applicable across astronomical datasets.

Acknowledgements. M.G., V.E., and M.W. acknowledge financial support from the German Aerospace Center and Federal Ministry of Education and Research through the project *Universal Bayesian Imaging Kit – Information Field Theory for Space Instrumentation* (Förderkennzeichen 500O2103). M.G. acknowledges support from the European Union-funded project *mw-at1as* under grant agreement No. 101166905. V.E. and P.F. acknowledge funding through the German Federal Ministry of Education and Research for the project *ErUM-IFT: Informationsfeldtheorie für Experimente an Großforschungsanlagen* (Förderkennzeichen: 05D23EO1). This work makes use of data from eROSITA, the soft X-ray instrument aboard SRG, a joint Russian-German science mission supported by the Russian Space Agency (Roskosmos), in the interests of the Russian Academy of Sciences represented by its Space Research Institute (IKI), and the Deutsches Zentrum für Luft- und Raumfahrt (DLR). The SRG spacecraft was built by Lavochkin Association (NPOL) and its subcontractors, and is operated by NPOL with support from the Max Planck Institute for Extraterrestrial Physics (MPE). The development and construction of the eROSITA X-ray instrument was led

by MPE, with contributions from the Dr. Karl Remeis Observatory Bamberg & ECAP (FAU Erlangen-Nuernberg), the University of Hamburg Observatory, the Leibniz Institute for Astrophysics Potsdam (AIP), and the Institute for Astronomy and Astrophysics of the University of Tübingen, with the support of DLR and the Max Planck Society. The Argelander Institute for Astronomy of the University of Bonn and the Ludwig Maximilians Universität Munich also participated in the science preparation for eROSITA. The eROSITA data shown here was processed using the eSASS software system developed by the German eROSITA consortium. M.G. thanks Jakob Roth for valuable discussions. The authors also thank the anonymous referee for their constructive comments, which helped improve the clarity of this work.

References

- Arras, P., Baltac, M., Ensslin, T. A., et al. 2019, *Astrophysics Source Code Library* [[record ascl:1903.008](#)]
- Arras, P., Bester, H. L., Perley, R. A., et al. 2021, *A&A*, **646**, A84
- Arras, P., Frank, P., Haim, P., et al. 2022, *Nat. Astron.*, **6**, 259
- Bertin, E., & Arnouts, S. 1996, *Astron. Astrophys. Suppl. Ser.*, **117**, 393
- Bertin, E., Schefer, M., Apostolakis, N., et al. 2020, *ASP Conf. Ser.*, **527**, 461
- Casas, J. M., González-Nuevo, J., Bonavera, L., et al. 2022, *A&A*, **658**, A110
- Cash, W. 1979, *ApJ*, **228**, 939
- Donath, A., Siemiginowska, A., Kashyap, V. L., van Dyk, D. A., & Burke, D. 2024, *AJ*, **168**, 182
- Eberle, V., Guardiani, M., Westerkamp, M., et al. 2024, J-UBIK: The JAX-accelerated Universal Bayesian Imaging Kit [[arXiv:2409.10381](#)]
- Eberle, V., Guardiani, M., Westerkamp, M., et al. 2025, <https://doi.org/10.5281/zenodo.16918522>
- Edenhofer, G., Frank, P., Roth, J., et al. 2024, *J. Open Source Softw.*, **9**, 6593
- Enßlin, T. 2013, *AIP Conf. Ser.* **1553**, 184
- Enßlin, T. 2025, *Quantifying Imperfect Cognition via Achieved Information Gain* (Hoboken: Wiley)
- Enßlin, T. A. 2019, *Annalen der Physik*, **531**, 1800127
- Enßlin, T. A., Frommert, M., & Kitaura, F. S. 2009, *Phys. Rev. D*, **80**, 105005
- Frank, P., Leike, R., & Enßlin, T. A. 2021, *Entropy*, **23**, 853
- Girón, F. J. C. C. d. 2001, *RACSAM*, **95**, 39
- Guglielmetti, F., Fischer, R., & Dose, V. 2009, *MNRAS*, **396**, 165
- Herranz, D., Gallegos, J., Sanz, J. L., & Martínez-González, E. 2002, *MNRAS*, **334**, 533
- Högbom, J. A. 1974, *A&AS*, **15**, 417
- Knollmüller, J., & Enßlin, T. A. 2020, *Metric Gaussian Variational Inference* [[arXiv:1901.11033](#)]
- Madarász, M., Marton, G., Gezer, I., et al. 2025, *A&A*, **696**, A37
- Maisinger, K., Hobson, M. P., & Lasenby, A. N. 2004, *MNRAS*, **347**, 339
- Principe, G., Malyshev, D., Ballet, J., & Funk, S. 2018, *A&A*, **618**, A22
- Rau, U., & Cornwell, T. J. 2011, *A&A*, **532**, A71
- Robotham, A. S. G., Davies, L. J. M., Driver, S. P., et al. 2018, *MNRAS*, **476**, 3137
- Roth, J., Frank, P., Bester, H. L., et al. 2024, *A&A*, **690**, A387
- Rybicki, G. B., & Lightman, A. P. 1986, *Radiative Processes in Astrophysics* (Hoboken: Wiley)
- Rüstig, J., Guardiani, M., Roth, J., Frank, P., & Enßlin, T. 2024, *A&A*, **682**, A146
- Scheel-Platz, L. I., Knollmüller, J., Arras, P., et al. 2023, *A&A*, **680**, A2
- Selig, M., & Enßlin, T. A. 2015, *A&A*, **574**, A74
- Selig, M., Bell, M. R., Junklewitz, H., et al. 2013, *A&A*, **554**, A26
- Skilling, J., & Bryan, R. K. 1984, *MNRAS*, **211**, 111
- Strong, A. W. 2003, *A&A*, **411**, L127
- Vafaei Sadr, A., Vos, E. E., Bassett, B. A., et al. 2019, *MNRAS*, **484**, 2793
- Westerkamp, M., Eberle, V., Guardiani, M., et al. 2024, *A&A*, **684**, A155
- Wiaux, Y., Jacques, L., Puy, G., Scaife, A. M. M., & Vanderghynst, P. 2009, *MNRAS*, **395**, 1733

² <https://github.com/NIFTy-PPL/J-UBIK>

Appendix A: Latent-space point source representation

In this appendix, we provide a deeper theoretical understanding of how point sources, when incorrectly modeled by a purely diffuse sky component, manifest in the latent-space excitations. This insight is crucial for interpreting the excitation maps used in our component-separation method.

A.1. Gaussian correlated field

Consider a one-dimensional sky field s composed of a diffuse component and N_{ps} point sources

$$s_x^{\text{gt}} := \mathcal{F}_{xk} \mathcal{A}_k^{\text{gt}} \xi_k^{\text{gt}} + \sum_{i=1}^{N_{\text{ps}}} b_i \delta(x - x_i),$$

where we define the fixed amplitude spectrum as $\mathcal{A}_k := \sqrt{\hat{T}_k}$. Suppose we attempt to fit this sky using only a diffuse component

$$s_x := \mathcal{F}_{xk} \mathcal{A}_k \xi_k.$$

We then aim to minimize the functional

$$\mathcal{L} := \int_{\Omega_x} (s_x^{\text{gt}} - s_x)^2 d\tilde{x} + \int_{\Omega_x} \xi_x^2 d\tilde{x}, \quad (\text{A.1})$$

where minimizing \mathcal{L} ensures both data fidelity (matching observations) and regularization (enforcing the standard-normal prior on ξ_x). Invoking Plancherel's theorem we can rewrite the integral in Fourier space

$$\begin{aligned} \mathcal{L} &= \frac{1}{(2\pi)^D} \left[\int_{\Omega_k} \mathcal{F}_{\tilde{k}\tilde{x}} (s_x^{\text{gt}} - s_x)^2 d\tilde{k} + \int_{\Omega_k} \mathcal{F}_{\tilde{k}\tilde{x}} \xi_x^2 d\tilde{k} \right] \\ &= \frac{1}{(2\pi)^D} \left[\int_{\Omega_k} \mathcal{F}_{\tilde{k}\tilde{x}} \left(\mathcal{F}_{\tilde{x}\tilde{k}} \mathcal{A}_k^{\text{gt}} \xi_k^{\text{gt}} + \sum_{i=1}^{N_{\text{ps}}} b_i \delta(\tilde{x} - x_i) - \mathcal{F}_{\tilde{x}\tilde{k}} \mathcal{A}_k \xi_k \right)^2 d\tilde{k} \right. \\ &\quad \left. + \int_{\Omega_k} \mathcal{F}_{\tilde{k}\tilde{x}} \xi_x^2 d\tilde{k} \right]. \end{aligned}$$

This loss is minimized for

$$\begin{aligned} 0 &\stackrel{!}{=} \frac{\delta \mathcal{L}}{\delta \xi_k'} \\ &= \frac{2}{(2\pi)^D} \int_{\Omega_k} \delta \xi_k' \left[-\mathcal{A}_k^{\text{gt}} \mathcal{A}_k \xi_k^{\text{gt}} - \mathcal{A}_k \sum_{i=1}^{N_{\text{ps}}} b_i \mathcal{F}_{\tilde{k}\tilde{x}} \delta(\tilde{x} - x_i) \right. \\ &\quad \left. + (\mathcal{A}_k^2 + 1) \xi_k \right] d\tilde{k} \end{aligned}$$

which, to hold for all $\delta \xi_k'$, implies

$$\xi_k = \frac{\mathcal{A}_k \mathcal{A}_k^{\text{gt}} \xi_k^{\text{gt}}}{\mathcal{A}_k^2 + 1} + \frac{\mathcal{A}_k}{\mathcal{A}_k^2 + 1} \sum_{i=1}^{N_{\text{ps}}} b_i \mathcal{F}_{\tilde{k}\tilde{x}} \delta(\tilde{x} - x_i). \quad (\text{A.2})$$

The first term in Eq. (A.2) directly relates the inferred excitations to the ground truth excitations scaled by the amplitude spectrum. Specifically, if a given Fourier mode k has a high amplitude \mathcal{A}_k , the inferred excitation ξ_k is closely aligned with the ground truth excitation ξ_k^{gt} , since the likelihood-like term (first term in

the sum in Eq. (A.1)) dominates over the prior-like one (second term in Eq. (A.1)). Conversely, if the amplitude is low, the prior enforcing small excitations dominates, thereby pushing the inferred excitation closer to zero and diminishing the influence of the ground truth excitation. The second term explicitly represents the latent-space excitations induced by point sources. Each point source creates excitations that appear as a superposition of monochromatic waves, whose frequencies are determined by the positions of the point sources in position-space. Crucially, these waves are scaled by the amplitude spectrum \mathcal{A}_k , implying that the latent-space signature of a point source is modulated by the correlation structure assumed in the diffuse field model. Understanding the meaning of the latent-space excitations created by a point source is not just an interesting mathematical exercise. In fact, convolving the real-space excitations ξ_x with the inverse Fourier transform of the second term template enhances point-source visibility.

A.2. Lognormal correlated field

In the case, like the ones analyzed in this work, in which we want to model a lognormal field with point sources, the situation becomes more complicated. The ground truth sky becomes

$$s_x^{\text{gt}} := e^{\mathcal{F}_{xk} \mathcal{A}_k \xi_k^{\text{gt}}} + \sum_{i=1}^{N_{\text{ps}}} b_i \delta(x - x_i),$$

while the reconstruction sky is

$$s_x := e^{\mathcal{F}_{xk} \mathcal{A}_k \xi_k}.$$

In fact the functional that needs to be minimized in this case becomes

$$\mathcal{L} = \int_{\Omega_x} \left[\left(e^{\mathcal{F}_{xk} \mathcal{A}_k \xi_k^{\text{gt}}} + \sum_{i=1}^{N_{\text{ps}}} b_i \delta(\tilde{x} - x_i) - e^{\mathcal{F}_{xk} \mathcal{A}_k \xi_k} \right)^2 + \xi_x^2 \right] d\tilde{x}.$$

This requires solving a system of transcendental equations. We defer the analysis of this system to future work. In the linear approximation, we can use the point source filter derived in the Gaussian case.

A.3. Latent-space point source filtering

Motivated by the theoretical considerations presented above, we now develop a filtering approach for latent-space excitations induced by point sources. Filtering latent-space excitations improves sensitivity to faint point sources and significantly reduces false-positive detections, making our approach particularly effective in fields with multiple different source components.

In the Gaussian field case, as discussed previously, the optimal latent-space filter is given by the inverse Fourier transform of the point-source-induced excitation pattern (second term in Eq. (A.2)), modulated by the amplitude spectrum

$$f_x^{\text{lin}} := \mathcal{F}_{x\tilde{k}} \left[\frac{\mathcal{A}_k}{\mathcal{A}_k^2 + 1} \mathcal{F}_{\tilde{k}\tilde{x}} \delta(\tilde{x} - x_0) \right],$$

where x_0 is the center of the filter. where the filter center is positioned at x_0 . Here, the delta function in real space transforms into a complex exponential in Fourier space, clearly showing that the

filter aligns with monochromatic waves modulated by the amplitude spectrum. On a grid, the specific shape of the filter then depends on the choice of discrete harmonic transform used.

In the lognormal scenario, it is not straightforward to analytically determine the latent-space response to a point source due to the complexity of the lognormal correlation structure. Moreover, the exact latent-space response depends on the accuracy of the posterior approximation – specifically, how effectively the KL divergence has been minimized at the current iteration.

To overcome these complexities, we empirically construct a heuristic matched filter based on the latent-space response observed in synthetic simulations, such as those presented in Sec. 2.5. We define this heuristic filter as

$$f_{\mathbf{x}}^{\text{heur}} := \begin{cases} 1 & \text{if } |\mathbf{x}| \leq r_{\text{ps}}, \\ -1 & \text{if } r_{\text{ps}} < |\mathbf{x}| \leq r_{\text{max}}, \\ 0 & \text{otherwise.} \end{cases}$$

This filter assigns a positive weight within a radius r_{ps} around the source center, capturing the primary excitation region, and a negative weight in the surrounding annulus up to a radius r_{max} , accounting for typical overshoot patterns observed in the latent-space excitation field, as seen for example in Fig. 6 and Fig. 10. Exploring rigorous analytical treatments or data-driven methods (e.g., neural-network-based latent-space filters) represents promising future directions to further refine point-source filtering for lognormal correlated fields.

Appendix B: Diagnostics

We provide two reconstruction metrics to quantitatively assess the fidelity of the reconstructions presented throughout this work. We define the posterior mean of the noise weighted residuals of a signal field s as

$$\text{nwr} := \left\langle \frac{R(s) - d}{\sqrt{R(s)}} \right\rangle_{\mathcal{P}(s|d)} \approx \left\langle \frac{R(s) - d}{\sqrt{R(s)}} \right\rangle_{\mathcal{Q}(s|d)}, \quad (\text{B.1})$$

where the expectation value is evaluated as an average over approximate geoVI posterior $\mathcal{Q}(s|d)$ samples. Furthermore, for synthetic data simulations in which the ground truth signal is known, we can define the uncertainty weighted residuals as

$$\text{uwr} := \frac{\langle s \rangle_{\mathcal{P}(s|d)} - s_{\text{gt}}}{\sqrt{\left\langle \left(s - \langle s \rangle_{\mathcal{P}(s|d)} \right)^2 \right\rangle_{\mathcal{P}(s|d)}}} \approx \frac{\langle s \rangle_{\mathcal{Q}(s|d)} - s_{\text{gt}}}{\sqrt{\left\langle \left(s - \langle s \rangle_{\mathcal{Q}(s|d)} \right)^2 \right\rangle_{\mathcal{Q}(s|d)}}}, \quad (\text{B.2})$$

where s_{gt} is the ground truth signal and again the same holds for the expectation values over the approximate posterior $\mathcal{Q}(s|d)$.

B.1. Imaging example

Figure B.1 shows both residual metrics for the single-frequency imaging example, showing good reconstruction fidelity.

B.2. Synthetic multifrequency imaging example

In Fig. B.2, we present the ground truth and reconstructed spectral index maps and spectral deviation fields, demonstrating that our method accurately captures both large-scale and small-scale spectral features. The reconstructed spectral index closely matches the ground truth, confirming that the method reliably

separates spectral variations from the primary power-law trend. Similarly, the spectral deviations reconstruction effectively captures localized deviations from the power-law assumption.

Figure B.3 shows nwr and uwr across three frequency bands for the multifrequency example, confirming accurate reconstruction across the entire spectral range.

B.3. LMC SN1987A eROSITA reconstruction

We present additional diagnostics of the eROSITA LMC SN1987A reconstruction. Figure B.5 shows the posterior mean of the reconstructed sky photon flux for each energy band, along with the posterior mean maps of the background spectral index and spectral deviations. Figure B.6 presents the corresponding noise-weighted residuals, computed according to Eq. (B.1). We also show the ability of the model to predict the photon flux distribution at unobserved wavelengths in Fig. B.4.

Table B.1 lists all detected point-like sources. In a post-processing step, we removed all point-source candidates (48 in total) for which the posterior standard deviation of the flux density exceeded 1.5 times the posterior mean, or for which the positional uncertainty exceeded two pixels. This procedure eliminates poorly constrained point-source components that do not meaningfully contribute to the reconstruction. We emphasize that this list does not constitute a validated astrophysical source catalog. The detection algorithm may flag sharp edges or residual structures – often arising from limitations in the PSF model – as spurious point sources. These residuals primarily stem from incomplete or imperfect knowledge of the instrument, such as inaccuracies in the calibration database and corresponding response implementation. For further discussion of these systematics, see Eberle et al. (2025). We visually inspected the detected sources and flagged those likely attributable to such systematics in the rightmost column of Table B.1.

Appendix C: Slope deviations degeneracy

The logarithm of the spectral deviations term $\log I_{\delta}(\mathbf{x}, \nu)$ from Eq. (7) reads

$$\log I_{\delta}(\mathbf{x}, \nu) = \tilde{\delta}(\mathbf{x}, \nu) - \bar{\delta}(\mathbf{x}, \nu) \log \frac{\nu}{\nu_{\text{ref}}}.$$

Suppose that the deviation field $\tilde{\delta}(\mathbf{x}, \nu)$ contains a slope-like term in addition to residual structure

$$\tilde{\delta}(\mathbf{x}, \nu) := \alpha'(\mathbf{x}) \log \frac{\nu}{\nu_{\text{ref}}} + \tilde{\delta}_R(\mathbf{x}, \nu),$$

where $\alpha'(\mathbf{x})$ is a spurious spectral slope introduced by the stochastic deviation process, and $\tilde{\delta}_R(\mathbf{x}, \nu)$ captures variations orthogonal to the power-law subspace. The term $\alpha'(\mathbf{x}) \log(\nu/\nu_{\text{ref}})$ is degenerate with the main power-law component parametrized by $\alpha(\mathbf{x}) \log(\nu/\nu_{\text{ref}})$ in Eq. (7).

To eliminate this degeneracy, we subtract the empirical average slope $\bar{\delta}(\mathbf{x})$ defined in Eq. (8), which removes the power-law-like component in $\tilde{\delta}(\mathbf{x}, \nu)$

$$\begin{aligned} \bar{\delta}(\mathbf{x}) &= \frac{\int_{\Omega_{\nu}} \delta(\mathbf{x}, \tilde{\nu}) \log \frac{\tilde{\nu}}{\nu_{\text{ref}}} d\tilde{\nu}}{\int_{\Omega_{\nu}} \left(\log \frac{\tilde{\nu}}{\nu_{\text{ref}}} \right)^2 d\tilde{\nu}} \\ &= \alpha'(\mathbf{x}) + \frac{\int_{\Omega_{\nu}} \tilde{\delta}_R(\mathbf{x}, \tilde{\nu}) \log \frac{\tilde{\nu}}{\nu_{\text{ref}}} d\tilde{\nu}}{\int_{\Omega_{\nu}} \left(\log \frac{\tilde{\nu}}{\nu_{\text{ref}}} \right)^2 d\tilde{\nu}} \\ &=: \alpha'(\mathbf{x}) + \alpha_R(\mathbf{x}). \end{aligned}$$




Fig. B.1. Imaging example residuals (for the synthetic data example shown in Fig. 4). *Left panel:* noise-weighted residuals (nwr); *middle panel:* uncertainty-weighted residuals (uwr); *right panel:* reconstructed flux vs. ground-truth flux comparison. The nwr and uwr share the same color bar.




Fig. B.2. Spectral reconstruction from the synthetic multifrequency imaging example from Sec. 2.5. The top row shows the ground truth spectral index map $\alpha(\mathbf{x})$ (left panels) and spectral deviations $\log I_\delta(\mathbf{x}, \nu)$ as defined in Eq. (6). The central panels represent the spectral fluctuations at ν_{low} and the right ones at ν_{hi} . The bottom row presents the corresponding posterior mean reconstructions. These highlight the method’s ability to recover both large-scale spectral slopes and localized deviations across the frequency range.

The resulting expression for $\tilde{\delta}(\mathbf{x}, \nu)$ thus becomes

$$\tilde{\delta}(\mathbf{x}, \nu) = \tilde{\delta}_R(\mathbf{x}, \nu) - \alpha_R(\mathbf{x}) \log \frac{\nu}{\nu_{\text{ref}}}.$$

This formulation ensures that any spectral slope contained in the stochastic deviation field is projected out, preserving the interpretability and identifiability of the true spectral index $\alpha(\mathbf{x})$ in Eq. (7). In the ideal case where $\tilde{\delta}_R(\mathbf{x}, \nu)$ contains no power-law-like contribution, the correction term $\alpha_R(\mathbf{x})$ vanishes, and $\tilde{\delta}(\mathbf{x}) = \alpha'(\mathbf{x})$ as expected.

Appendix D: Impact of spectral blurring (RMF)

In addition to the spatial blurring described by the point-spread function (PSF), X-ray detectors also introduce spectral blurring, quantified by the redistribution matrix file (RMF). The RMF encodes the probability distribution of detected photon energies for each incident photon energy and is therefore the spectral analog of the PSF. This redistribution leads to “spectral bleeding,” i.e. a fraction of photons being registered in energy channels offset from their true energy.

For SRG/eROSITA, this effect is negligible in the context of our analysis. Figure D.1 shows the TM1 RMF extracted from




Fig. B.3. Residual diagnostics for the synthetic multifrequency imaging example in Sec. 2.5. *Top row:* noise-weighted residuals for the three energy bands (ν_{low} , ν_{ref} , ν_{hi}) shown column-wise. *Bottom row:* uncertainty-weighted residuals for the same bands, illustrating agreement within the inferred posterior uncertainties.

the `tm1_rm_f_141103v02.fits` calibration file in the eSASS CALDB. The left panel displays the full-resolution response in the 0.2 keV to 4.5 keV band, which follows the $E_{\text{true}} = E_{\text{det}}$ diagonal closely with only minor off-diagonal leakage. The right panel shows the same response integrated into the three broad bands used in this work (0.2 keV to 1.0 keV, 1.0 keV to 2.0 keV, and 2.0 keV to 4.5 keV). In this representation, the response matrix is nearly diagonal: more than 96% of the flux remains in the correct band, with cross-talk between adjacent bins at the percent level or below.

Since our method already accounts for statistical noise and PSF-induced spatial mixing, and given that the focus of this work lies on developing efficient prior models, we neglect RMF effects in the forward model. Including them would increase computational cost without altering our conclusions, as the level of spectral bleeding is subdominant to statistical uncertainties and astrophysical variability. For applications requiring higher spectral precision—such as detailed line studies—an explicit treatment of the RMF could be incorporated into our framework in a straightforward manner.




Fig. B.4. Posterior mean reconstruction of the eROSITA LMC SN1987A sky photon flux at an unobserved intermediate energy band (here, 2.4 keV), predicted from the joint multifrequency model in Eq. (13). This demonstrates the model's ability to interpolate the sky signal across energy bands based on the inferred spectral correlation structure.

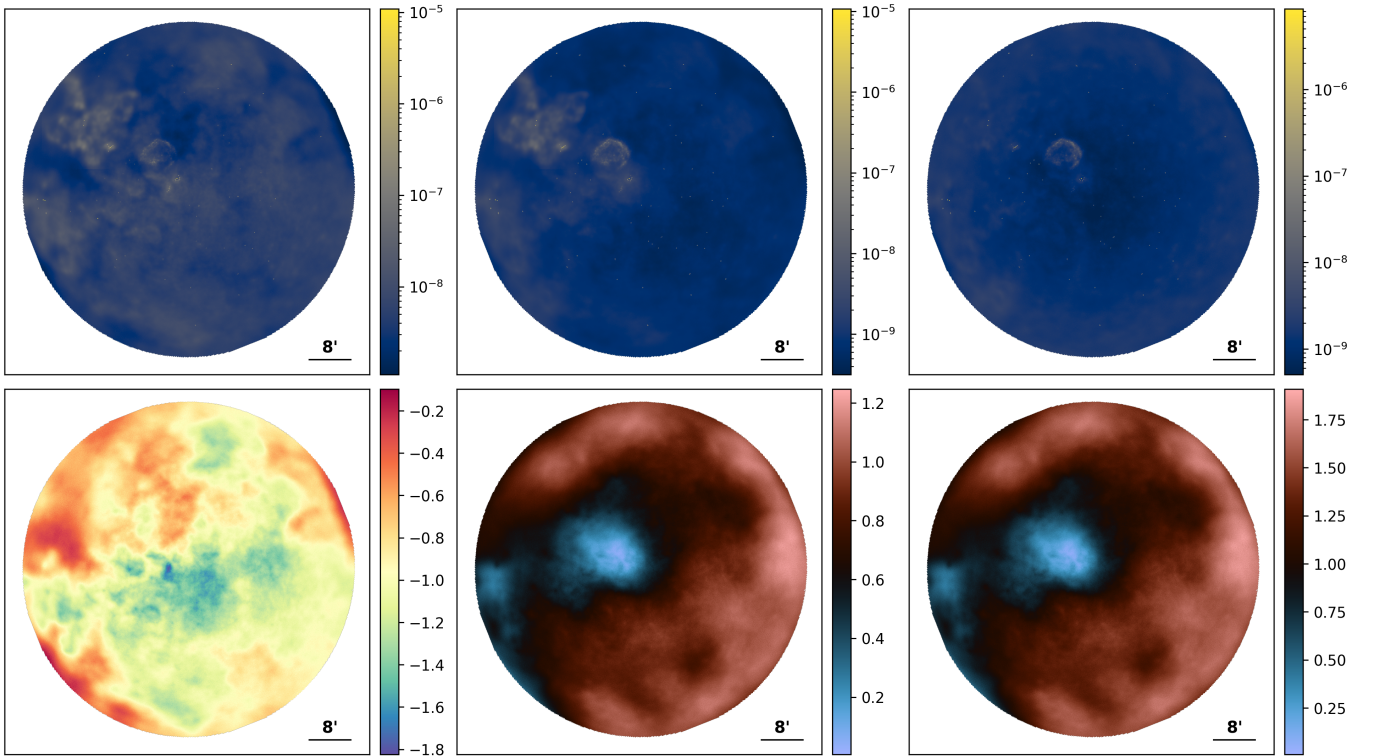


Fig. B.5. eROSITA LMC SN1987A reconstruction. The top row shows the reconstructed sky photon flux in the three energy bands, from left to right: $\nu_{\text{low}} = 0.2 - 1.0$, keV, $\nu_{\text{ref}} = 1.0 - 2.0$, keV, and $\nu_{\text{hi}} = 2.0 - 4.5$, keV. The bottom row displays the posterior mean of the reconstructed spectral index map $\alpha(\mathbf{x})$ and spectral deviation map $\log I_{\delta}(\mathbf{x}, \nu)$ (see Eq. (6)) for the diffuse background component.




Fig. B.6. eROSITA LMC SN1987A data noise-weighted residuals. The columns represent, from left to right, the nwr in the three energy channels ($\nu_{\text{low}} = 0.2 - 1.0$, keV, $\nu_{\text{ref}} = 1.0 - 2.0$, keV, and $\nu_{\text{hi}} = 2.0 - 4.5$, keV), while the rows correspond to the three telescope modules used in the reconstruction (TM1, TM3, TM4).




Fig. D.1. Response matrix file (RMF) of the eROSITA telescope for telescope module TM1. Left: Full-resolution RMF in the 0.2 keV to 4.5 keV range, showing the probability distribution of detected energies (horizontal axis) for incident photon energies (vertical axis). The dashed white line marks the ideal one-to-one relation $E_{\text{true}} = E_{\text{det}}$. Right: The same response integrated into three broad bands (0.2 keV to 1.0 keV, 1.0 keV to 2.0 keV, and 2.0 keV to 4.5 keV), yielding a 3×3 representation of the energy redistribution. Cell values indicate normalized detection probabilities, and grid lines mark the bin boundaries. The nearly diagonal structure demonstrates that spectral bleeding is negligible for the purposes of this work.

Table B.1. Detected point sources in the SN 1987A field.

Positions (px)	log flux density (0.2–1 keV)	log flux density (1–2 keV)	log flux density (2–4.5 keV)	Possible artifact
108.37 ± 0.01, 488.84 ± 0.01	−8.92 ± 0.01	−8.40 ± 0.01	−8.63 ± 0.01	0
483.27 ± 0.01, 549.33 ± 0.01	−8.50 ± 0.00	−9.17 ± 0.00	−10.12 ± 0.01	0
297.25 ± 0.05, 635.97 ± 0.11	−11.04 ± 0.08	−10.34 ± 0.08	−10.47 ± 0.08	0
113.28 ± 0.09, 486.88 ± 0.08	−10.21 ± 0.03	−11.78 ± 0.13	−10.76 ± 0.09	0
230.08 ± 0.15, 892.54 ± 0.24	−9.52 ± 0.05	−11.12 ± 0.08	−12.40 ± 0.24	0
299.79 ± 0.14, 638.38 ± 0.14	−12.32 ± 0.18	−11.45 ± 0.06	−11.66 ± 0.24	0
260.80 ± 0.02, 437.37 ± 0.03	−10.08 ± 0.02	−12.01 ± 0.02	−13.90 ± 0.21	0
221.13 ± 0.05, 698.05 ± 0.06	−12.70 ± 0.14	−12.02 ± 0.05	−12.13 ± 0.09	0
223.38 ± 0.10, 711.39 ± 0.09	−12.53 ± 0.10	−12.17 ± 0.05	−12.99 ± 0.21	0
759.08 ± 0.05, 492.39 ± 0.08	−12.66 ± 0.04	−13.45 ± 0.07	−13.51 ± 0.11	0
478.26 ± 0.14, 546.88 ± 0.20	−12.46 ± 0.13	−13.25 ± 0.08	−15.12 ± 0.39	0
303.33 ± 0.31, 640.88 ± 0.38	−13.36 ± 0.18	−13.42 ± 0.17	−14.55 ± 0.76	0
116.31 ± 0.30, 451.96 ± 1.19	−13.38 ± 0.35	−13.66 ± 0.07	−14.82 ± 0.36	1
701.40 ± 0.04, 397.79 ± 0.04	−13.50 ± 0.09	−14.59 ± 0.24	−13.93 ± 0.21	0
92.46 ± 0.33, 490.79 ± 0.59	−13.87 ± 0.50	−13.27 ± 0.12	−14.87 ± 0.70	1
293.05 ± 0.34, 634.29 ± 0.42	−12.97 ± 0.25	−13.44 ± 0.26	−15.58 ± 0.36	0
364.48 ± 0.08, 589.42 ± 0.08	−15.61 ± 0.37	−13.53 ± 0.04	−13.04 ± 0.06	0
740.36 ± 0.27, 317.60 ± 0.12	−13.78 ± 0.21	−14.35 ± 0.27	−14.18 ± 0.29	0
635.14 ± 0.33, 867.83 ± 0.16	−12.28 ± 0.06	−14.36 ± 0.23	−15.99 ± 0.95	0
279.93 ± 0.68, 899.75 ± 0.58	−14.91 ± 0.35	−14.28 ± 0.18	−14.04 ± 0.42	0
95.24 ± 0.24, 465.33 ± 0.90	−15.48 ± 1.61	−13.28 ± 0.09	−14.80 ± 0.89	1
687.02 ± 0.21, 754.48 ± 0.16	−14.11 ± 0.19	−14.49 ± 0.14	−15.09 ± 0.54	0
535.17 ± 0.62, 137.99 ± 0.51	−14.69 ± 0.74	−14.26 ± 0.08	−14.81 ± 0.62	0
344.52 ± 1.29, 894.66 ± 0.49	−14.64 ± 0.53	−14.53 ± 0.22	−14.83 ± 0.50	0
449.49 ± 0.25, 266.39 ± 0.20	−14.54 ± 0.14	−14.88 ± 0.07	−14.89 ± 0.22	0
485.93 ± 0.27, 542.76 ± 0.59	−13.79 ± 0.33	−14.75 ± 0.26	−15.85 ± 0.64	1
754.00 ± 0.59, 267.49 ± 0.32	−14.68 ± 0.44	−14.77 ± 0.17	−15.23 ± 0.60	0
314.23 ± 0.18, 321.66 ± 0.27	−13.95 ± 0.29	−14.93 ± 0.13	−15.78 ± 0.60	0
724.84 ± 0.22, 786.80 ± 0.14	−12.93 ± 0.09	−14.97 ± 0.16	−16.82 ± 0.53	0
184.75 ± 0.23, 456.68 ± 0.60	−16.04 ± 0.61	−14.58 ± 0.15	−14.35 ± 0.47	0

Table B.1. continued.

Positions (px)	log flux density (0.2–1 keV)	log flux density (1–2 keV)	log flux density (2–4.5 keV)	Possible artifact
264.24 ± 0.16, 835.43 ± 0.62	-15.50 ± 0.93	-14.75 ± 0.23	-14.95 ± 0.68	0
886.52 ± 0.98, 666.21 ± 0.66	-15.40 ± 0.77	-14.92 ± 0.46	-14.98 ± 0.40	0
290.83 ± 0.34, 625.30 ± 0.34	-14.16 ± 0.29	-14.82 ± 0.43	-16.26 ± 0.47	0
340.25 ± 0.13, 592.13 ± 0.26	-16.22 ± 0.46	-14.85 ± 0.12	-14.42 ± 0.20	0
318.82 ± 0.21, 509.42 ± 0.41	-15.25 ± 0.15	-15.14 ± 0.14	-15.15 ± 0.32	0
490.64 ± 0.33, 551.07 ± 0.50	-13.37 ± 0.25	-15.14 ± 0.56	-16.89 ± 0.92	1
538.03 ± 0.19, 443.07 ± 0.21	-15.37 ± 0.40	-15.37 ± 0.08	-15.26 ± 0.35	0
503.63 ± 0.39, 790.24 ± 0.19	-15.77 ± 0.71	-15.16 ± 0.14	-15.08 ± 0.51	0
337.48 ± 0.50, 332.43 ± 0.34	-15.44 ± 0.72	-15.21 ± 0.25	-15.41 ± 0.60	0
521.24 ± 0.34, 243.54 ± 0.60	-15.38 ± 0.70	-15.23 ± 0.13	-15.46 ± 0.52	0
674.09 ± 0.36, 361.87 ± 0.34	-15.07 ± 0.43	-15.46 ± 0.23	-15.75 ± 0.52	0
656.76 ± 0.46, 222.89 ± 0.62	-15.75 ± 0.90	-15.27 ± 0.27	-15.57 ± 0.58	0
630.54 ± 0.25, 496.85 ± 0.26	-15.31 ± 0.42	-15.57 ± 0.19	-15.83 ± 0.54	0
598.66 ± 0.99, 800.86 ± 0.23	-15.70 ± 0.51	-15.21 ± 0.28	-15.86 ± 0.64	0
694.77 ± 0.29, 353.21 ± 0.26	-15.04 ± 0.50	-15.72 ± 0.26	-16.20 ± 0.74	0
291.61 ± 0.38, 639.01 ± 0.20	-13.37 ± 0.18	-15.95 ± 0.05	-17.50 ± 0.52	0
308.93 ± 0.62, 300.97 ± 0.58	-16.08 ± 1.17	-15.33 ± 0.39	-15.74 ± 0.50	0
615.64 ± 0.30, 354.15 ± 0.32	-16.20 ± 0.91	-15.72 ± 0.45	-15.32 ± 0.44	0
258.47 ± 0.21, 487.28 ± 0.94	-15.92 ± 0.59	-15.57 ± 0.17	-15.79 ± 0.61	0
479.93 ± 0.47, 542.94 ± 0.63	-13.83 ± 0.41	-15.57 ± 1.01	-17.67 ± 2.04	1
489.71 ± 1.05, 545.97 ± 1.06	-13.94 ± 0.43	-15.87 ± 0.87	-17.44 ± 1.26	1
374.77 ± 0.34, 545.10 ± 0.21	-15.89 ± 0.40	-15.74 ± 0.36	-15.82 ± 0.19	0
614.23 ± 0.26, 669.74 ± 0.58	-15.88 ± 0.58	-15.81 ± 0.23	-15.76 ± 0.52	0
350.40 ± 0.23, 600.40 ± 0.51	-16.69 ± 0.69	-15.55 ± 0.15	-15.22 ± 0.53	0
746.26 ± 0.29, 471.30 ± 0.52	-16.35 ± 0.83	-15.72 ± 0.12	-15.61 ± 0.50	0
469.59 ± 1.40, 865.06 ± 1.28	-16.22 ± 1.13	-15.77 ± 0.41	-15.96 ± 0.86	0
324.28 ± 0.36, 483.49 ± 0.29	-17.04 ± 1.42	-15.81 ± 0.22	-15.62 ± 0.45	0
629.44 ± 0.52, 648.48 ± 0.64	-17.10 ± 1.06	-15.88 ± 0.29	-15.66 ± 0.72	0
287.73 ± 0.39, 644.59 ± 0.39	-13.44 ± 0.36	-16.56 ± 0.92	-18.41 ± 1.27	0
718.66 ± 0.33, 565.17 ± 0.69	-16.74 ± 1.27	-15.90 ± 0.09	-16.27 ± 0.78	0
715.98 ± 0.42, 487.18 ± 0.29	-17.20 ± 1.46	-15.84 ± 0.09	-15.89 ± 0.51	0
517.54 ± 0.27, 473.98 ± 0.45	-16.34 ± 0.84	-16.19 ± 0.16	-16.48 ± 0.56	0
419.68 ± 0.38, 595.41 ± 0.18	-13.58 ± 0.20	-16.69 ± 0.34	-18.70 ± 0.90	0
654.17 ± 0.90, 498.55 ± 0.48	-15.15 ± 0.54	-16.69 ± 0.35	-17.52 ± 0.54	0
653.14 ± 1.17, 281.40 ± 1.08	-16.68 ± 0.96	-16.29 ± 0.35	-16.61 ± 0.89	0
331.73 ± 0.28, 562.60 ± 0.45	-17.77 ± 1.52	-16.28 ± 0.49	-15.54 ± 0.57	0
580.55 ± 0.94, 576.01 ± 0.44	-15.13 ± 0.33	-16.70 ± 0.43	-17.89 ± 0.87	0
410.25 ± 0.48, 672.95 ± 0.64	-18.12 ± 1.91	-16.17 ± 0.40	-15.49 ± 0.48	0
462.50 ± 0.61, 519.71 ± 0.80	-14.24 ± 0.73	-16.70 ± 0.38	-18.69 ± 1.32	0
537.24 ± 0.50, 373.28 ± 0.55	-17.56 ± 1.56	-16.25 ± 0.30	-16.15 ± 0.43	0
565.37 ± 0.21, 354.42 ± 0.77	-17.55 ± 1.52	-16.00 ± 0.13	-16.42 ± 0.91	0
459.84 ± 0.18, 515.75 ± 0.67	-13.95 ± 0.19	-16.91 ± 0.57	-18.89 ± 1.02	0
443.61 ± 0.73, 323.36 ± 1.24	-16.59 ± 0.74	-16.69 ± 0.64	-16.92 ± 0.96	0
280.62 ± 0.41, 549.14 ± 0.91	-17.74 ± 1.69	-16.71 ± 0.74	-15.75 ± 0.50	0
462.26 ± 0.82, 525.21 ± 0.47	-14.05 ± 0.34	-17.01 ± 0.62	-19.13 ± 1.29	0
660.01 ± 0.77, 575.79 ± 0.59	-17.06 ± 0.91	-16.69 ± 0.63	-16.87 ± 0.59	0
574.04 ± 1.40, 385.20 ± 0.80	-18.04 ± 1.04	-16.99 ± 0.55	-15.53 ± 0.42	0
487.49 ± 0.26, 557.36 ± 0.26	-14.03 ± 0.29	-17.25 ± 0.59	-19.06 ± 1.00	1
488.25 ± 0.68, 424.83 ± 0.92	-16.72 ± 0.66	-17.02 ± 0.35	-17.16 ± 0.87	0
547.73 ± 0.56, 365.26 ± 0.65	-17.00 ± 1.31	-16.86 ± 0.61	-17.18 ± 0.78	0
455.83 ± 1.06, 511.60 ± 0.51	-15.17 ± 0.63	-17.22 ± 0.38	-18.84 ± 0.83	0
459.63 ± 1.76, 521.05 ± 0.92	-14.81 ± 0.97	-17.34 ± 0.57	-19.30 ± 1.70	0
544.68 ± 0.62, 472.26 ± 1.10	-14.98 ± 0.35	-17.59 ± 0.95	-19.08 ± 1.31	0
597.97 ± 0.58, 697.99 ± 0.61	-17.03 ± 1.22	-17.51 ± 0.97	-17.98 ± 0.80	0

Notes. Positions are $x \pm \sigma_x$, $y \pm \sigma_y$ in pixel coordinates. Logarithmic photon flux densities are measured in $\text{s}^{-1} \text{arcsec}^{-1}$ and reported with their associated 1σ posterior uncertainties. A value of 1 in the final column marks sources manually flagged as possible artifacts.