Denoising, deconvolving, and decomposing photon observations[★]

Derivation of the D³PO algorithm

Marco Selig^{1,2} and Torsten A. Enßlin^{1,2}

¹ Max-Planck Institut für Astrophysik, Karl-Schwarzschild-Straße 1, 85748 Garching, Germany
e-mail: mselig@mpa-garching.mpg.de

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

Received 7 November 2013 / Accepted 22 October 2014

ABSTRACT

The analysis of astronomical images is a non-trivial task. The D³PO algorithm addresses the inference problem of denoising, deconvolving, and decomposing photon observations. Its primary goal is the simultaneous but individual reconstruction of the diffuse and point-like photon flux given a single photon count image, where the fluxes are superimposed. In order to discriminate between these morphologically different signal components, a probabilistic algorithm is derived in the language of information field theory based on a hierarchical Bayesian parameter model. The signal inference exploits prior information on the spatial correlation structure of the diffuse component and the brightness distribution of the spatially uncorrelated point-like sources. A maximum a posteriori solution and a solution minimizing the Gibbs free energy of the inference problem using variational Bayesian methods are discussed. Since the derivation of the solution is not dependent on the underlying position space, the implementation of the D³PO algorithm uses the NIFTY package to ensure applicability to various spatial grids and at any resolution. The fidelity of the algorithm is validated by the analysis of simulated data, including a realistic high energy photon count image showing a 32×32 arcmin² observation with a spatial resolution of 0.1 arcmin. In all tests the D³PO algorithm successfully denoised, deconvolved, and decomposed the data into a diffuse and a point-like signal estimate for the respective photon flux components.

Key words. methods: data analysis – methods: numerical – methods: statistical – techniques: image processing – gamma rays: general – X-rays: general

1. Introduction

An astronomical image might display multiple superimposed features, such as point sources, compact objects, diffuse emission, or background radiation. The raw photon count images delivered by high energy telescopes are far from perfect; they suffer from shot noise and distortions caused by instrumental effects. The analysis of such astronomical observations demands elaborate denoising, deconvolution, and decomposition strategies.

The data obtained by the detection of individual photons is subject to Poissonian shot noise which is more severe for low count rates. This causes difficulty for the discrimination of faint sources against noise, and makes their detection exceptionally challenging. Furthermore, uneven or incomplete survey coverage and complex instrumental response functions leave imprints in the photon data. As a result, the data set might exhibit gaps and artificial distortions rendering the clear recognition of different features a difficult task. Point-like sources are especially afflicted by the instrument's point spread function (PSF) that smooths them out in the observed image, and therefore can cause fainter ones to vanish completely in the background noise.

In addition to such noise and convolution effects, it is the superposition of the different objects that makes their separation ambiguous, if possible at all. In astrophysics, photon emitting

objects are commonly divided into two morphological classes, diffuse sources and point sources. Diffuse sources span rather smoothly across large fractions of an image, and exhibit apparent internal correlations. Point sources, on the contrary, are local features that, if observed perfectly, would only appear in one pixel of the image. In this work, we will not distinguish between diffuse sources and background, both are diffuse contributions. Intermediate cases, which are sometimes classified as extended or compact sources, are also not considered here.

The question arises of how to reconstruct the original source contributions, the individual signals, that caused the observed photon data. This task is an ill-posed inverse problem without a unique solution. There are a number of heuristic and probabilistic approaches that address the problem of denoising, deconvolution, and decomposition in partial or simpler settings.

SEXTRACTOR (Bertin & Arnouts 1996) is one of the heuristic tools and the most prominent for identifying sources in astronomy. Its popularity is mostly based on its speed and easy operability. However, SEXTRACTOR produces a catalog of fitted sources rather than denoised and deconvolved signal estimates. CLEAN (Högbom 1974) is commonly used in radio astronomy and attempts a deconvolution assuming there are only contributions from point sources. Therefore, diffuse emission is not optimally reconstructed in the analysis of real observations using CLEAN.

Multiscale extensions of CLEAN (Cornwell 2008; Rau & Cornwell 2011) improve on this, but are also not perfect

[★] A copy of the code is available at the CDS via anonymous ftp to [cdsarc.u-strasbg.fr](ftp://cdsarc.u-strasbg.fr) (130.79.128.5) or via <http://cdsarc.u-strasbg.fr/viz-bin/qcat?J/A+A/574/A74>

(Junklewitz et al. 2014). Decomposition techniques for diffuse backgrounds, based on the analysis of angular power spectra have recently been proposed by Hensley et al. (2013).

Inference methods, in contrast, investigate the probabilistic relation between the data and the signals. Here, the signals of interest are the different source contributions. Probabilistic approaches allow a transparent incorporation of model and a priori assumptions, but often result in computationally heavier algorithms.

As an initial attempt, a maximum likelihood analysis was proposed by Valdes (1982). In later work, maximum entropy (Strong 2003) and minimum χ^2 methods (e.g., Bouchet et al. 2013) were applied to the INTEGRAL/SPI data reconstructing a single signal component, though. On the basis of sparse regularization a number of techniques exploiting waveforms (based on the work by Haar 1910, 1911) have proven successful in performing denoising and deconvolution tasks in different settings (González-Nuevo et al. 2006; Willett & Nowak 2007; Dupé et al. 2009, 2011; Figueiredo & Bioucas-Dias 2010). For example, Schmitt et al. (2010, 2012) analyzed simulated (single and multi-channel) data from the Fermi γ -ray space telescope focusing on the removal of Poisson noise and deconvolution or background separation. Furthermore, a (generalized) morphological component analysis denoised, deconvolved and decomposed simulated radio data assuming Gaussian noise statistics (Bobin et al. 2007; Chapman et al. 2013).

Still in the regime of Gaussian noise, Giovannelli & Coulaï (2005) derived a deconvolution algorithm for point and extended sources minimizing regularized least squares. They introduce an efficient convex regularization scheme at the price of a priori unmotivated fine tuning parameters. The fast algorithm PowellSnakes I/II by Carvalho et al. (2009, 2012) is capable of analyzing multi-frequency data sets and detecting point-like objects within diffuse emission regions. It relies on matched filters using PSF templates and Bayesian filters exploiting, among others, priors on source position, size, and number. PowellSnakes II has been successfully applied to the Planck data (Planck Collaboration VII 2011).

The approach closest to ours is the background-source separation technique used to analyze the ROSAT data (Guglielmetti et al. 2009). This Bayesian model is based on a two-component mixture model that reconstructs extended sources and (diffuse) background concurrently. The latter is, however, described by a spline model with a small number of spline sampling points.

The strategy presented in this work aims at the simultaneous reconstruction of two signals, the diffuse and point-like photon flux. Both fluxes contribute equally to the observed photon counts, but their morphological imprints are very different. The proposed algorithm, derived in the framework of information field theory (IFT; Enßlin et al. 2009; Enßlin 2013), therefore incorporates prior assumptions in form of a hierarchical parameter model. The fundamentally different morphologies of diffuse and point-like contributions reflected in different prior correlations and statistics. The exploitation of these different prior models is crucial to the signal decomposition. In this work, we exclusively consider Poissonian noise, in particular, but not exclusively, in the low count rate regime, where the signal-to-noise ratio becomes challengingly low. The D³PO algorithm presented here targets the simultaneous denoising, deconvolution, and decomposition of photon observations into two signals, the diffuse and point-like photon flux. This task includes the reconstruction of the harmonic power spectrum of the diffuse component from the data themselves. Moreover, the proposed algorithm provides a posteriori uncertainty information on both inferred signals.

The fluxes from diffuse and point-like sources contribute equally to the observed photon counts, but their morphological imprints are very different. The proposed algorithm, derived in the framework of information field theory (IFT; Enßlin et al. 2009; Enßlin 2013, 2014), therefore incorporates prior assumptions in form of a hierarchical parameter model. The fundamentally different morphologies of diffuse and point-like contributions reflected in different prior correlations and statistics. The exploitation of these different prior models is the key to the signal decomposition.

The diffuse and point-like signal are treated as two separate signal fields. A signal field represents an original signal appearing in nature; e.g., the physical photon flux distribution of one source component as a function of real space or sky position. In theory, a field has infinitely many degrees of freedom being defined on a continuous position space. In computational practice, however, a field needs of course to be defined on a finite grid. It is desirable that the signal field is reconstructed independently from the grid's resolution, except for potentially unresolvable features¹. We note that the point-like signal field hosts one point source in every pixel, however, most of them might be invisibly faint. Hence, a complicated determination of the number of point sources, as many algorithms perform, is not required in our case.

The derivation of the algorithm makes use of a wide range of Bayesian methods that are discussed below in detail with regard to their implications and applicability. For now, let us consider an example to emphasize the range and performance of the D³PO algorithm. Figure 1 illustrates a reconstruction scenario in one dimension, where the coordinate could be an angle or position (or time, or energy) in order to represent a 1D sky (or a time series, or an energy spectrum). The numerical implementation uses the NIFTY² package (Selig et al. 2013). NIFTY permits an algorithm to be set up abstractly, independent of the finally chosen topology, dimension, or resolution of the underlying position space. In this way, a 1D prototype code can be used for development, and then just be applied in 2D, 3D, or even on the sphere.

The remainder of this paper is structured as follows. Section 2 discusses the inference on photon observations; i.e., the underlying model and prior assumptions. The D³PO algorithm solving this inference problem by denoising, deconvolution, and decomposition is derived in Sect. 3. In Sect. 4 the algorithm is demonstrated in a numerical application on simulated high energy photon data. We conclude in Sect. 5.

2. Inference on photon observations

2.1. Signal inference

Here, a signal is defined as an unknown quantity of interest that we want to learn about. The most important information source on a signal is the data obtained in an observation to measure the signal. Inferring a signal from an observational data set poses

¹ If the resolution of the reconstruction were increased gradually, the diffuse signal field might exhibit more and more small scale features until the information content of the given data is exhausted. From this point on, any further increase in resolution would not change the signal field reconstruction significantly. In a similar manner, the localization accuracy and number of detections of point sources might increase with the resolution until all relevant information of the data was captured. All higher resolution grids can then be regarded as acceptable representations of the continuous position space.

² NIFTY homepage <http://www.mpa-garching.mpg.de/ift/nifty/>

Fig. 1. Illustration of a 1D reconstruction scenario with 1024 pixels. Panel **a**) shows the superimposed diffuse and point-like signal components (green solid line) and its observational response (gray contour). Panel **b**) shows again the signal response representing noiseless data (gray contour) and the generated Poissonian data (red markers). Panel **c**) shows the reconstruction of the point-like signal component (blue solid line), the diffuse one (orange solid line), its 2σ reconstruction uncertainty interval (orange dashed line), and again the original signal response (gray contour). The point-like signal comprises 1024 point-sources of which only five are not too faint. Panel **d**) shows the reproduced signal response representing noiseless data (black solid line), its 2σ shot noise interval (black dashed line), and again the data (gray markers).

a fundamental problem because of the presence of noise in the data and the ambiguity that several possible signals could have produced the same data, even in the case of negligible noise.

For example, given some image data like photon counts, we want to infer the underlying photon flux distribution. This physical flux is a continuous scalar field that varies with respect to time, energy, and observational position. The measured photon count data, however, is restricted by its spatial and energy binning, as well as its limitations in energy range and observation time. Basically, all data sets are finite for practical reasons, and therefore cannot capture all of the infinitely many degrees of freedom of the underlying continuous signal field.

There is no exact solution to such signal inference problems, since there might be (infinitely) many signal field configurations that could lead to the same data. This is why a probabilistic data analysis, which does not pretend to calculate the correct field configuration but provides expectation values and uncertainties of the signal field, is appropriate for signal inference.

Given a data set $\mathbf{d}$, the a posteriori probability distribution $P(\mathbf{s}|\mathbf{d})$ judges how likely a potential signal $\mathbf{s}$ is. This posterior is given by Bayes' theorem,

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

as a combination of the likelihood $P(\mathbf{d}|\mathbf{s})$, the signal prior $P(\mathbf{s})$, and the evidence $P(\mathbf{d})$, which serves as a normalization. The likelihood characterizes how likely it is to measure data set $\mathbf{d}$ from a given signal field $\mathbf{s}$. It covers all processes that are relevant for the measurement of $\mathbf{d}$. The prior describes the knowledge about $\mathbf{s}$ without considering the data, and should, in general, be less restrictive than the likelihood.

IFT is a Bayesian framework for the inference of signal fields exploiting mathematical methods for theoretical physics. A signal field, $\mathbf{s} = s(x)$, is a function of a continuous position x in some position space Ω . In order to avoid a dependence of the reconstruction on the partition of Ω , the according calculus regarding fields is geared to preserve the continuum limit (Enßlin 2013, 2014; Selig et al. 2013). In general, we are interested in the a posteriori mean estimate $\mathbf{m}$ of the signal field given the data, and its (uncertainty) covariance $\mathbf{D}$, defined as

$$\mathbf{m} = \langle \mathbf{s} \rangle_{(\mathbf{s}|\mathbf{d})} = \int \mathcal{D}\mathbf{s} \, \mathbf{s} \, P(\mathbf{s}|\mathbf{d}), \quad (2)$$

$$\mathbf{D} = \langle (\mathbf{m} - \mathbf{s})(\mathbf{m} - \mathbf{s})^\dagger \rangle_{(\mathbf{s}|\mathbf{d})}, \quad (3)$$

where $\dagger$ denotes adjunction and $\langle \cdot \rangle_{(\mathbf{s}|\mathbf{d})}$ the expectation value with respect to the posterior probability distribution $P(\mathbf{s}|\mathbf{d})$ ³.

In the following, the posterior of the physical photon flux distribution of two morphologically different source components given a data set of photon counts is build up piece by piece according to Eq. (1).

2.2. Poissonian likelihood

The images provided by astronomical high energy telescopes typically consist of integer photon counts that are binned spatially into pixels. Let d_i be the number of detected photons, also called events, in pixel i , where $i \in \{1, \dots, N_{\text{pix}}\} \subset \mathbb{N}$.

The kind of signal field we would like to infer from such data is the causative photon flux distribution. The photon

flux, $\rho = \rho(x)$, is defined for each position x on the observational space Ω . In astrophysics, this space Ω is typically the $\mathcal{S}^2$ sphere representing an all-sky view, or a region within $\mathbb{R}^2$ representing an approximately plane patch of the sky. The flux ρ might express different morphological features, which can be classified into a diffuse and point-like component. The exact definitions of the diffuse and point-like flux should be specified a priori, without knowledge of the data, and are addressed in Sects. 2.3.1 and 2.3.3, respectively. At this point it shall suffice to say that the diffuse flux varies smoothly on large spatial scales, while the flux originating from point sources is fairly local. These two flux components are superimposed,

$$\rho = \rho_{\text{diffuse}} + \rho_{\text{point-like}} = \rho_0 (e^s + e^u), \quad (4)$$

where we introduced the dimensionless diffuse and point-like signal fields, $\mathbf{s}$ and $\mathbf{u}$, and the constant ρ_0 which absorbs the physical dimensions of the photon flux; i.e., events per area per energy and time interval. The exponential function in Eq. (4) is applied componentwise. In this way, we naturally account for the strict positivity of the photon flux at the price of a non-linear change of variables, from the flux to its natural logarithm.

A measurement apparatus observing the photon flux ρ is expected to detect a certain number of photons λ . This process can be modeled by a linear response operator $\mathbf{R}_0$ as follows,

$$\lambda = \mathbf{R}_0 \rho = \mathbf{R} (e^s + e^u), \quad (5)$$

where $\mathbf{R} = \mathbf{R}_0 \rho_0$. This reads for pixel i ,

$$\lambda_i = \int_{\Omega} dx \, R_i(x) (e^{s(x)} + e^{u(x)}). \quad (6)$$

The response operator $\mathbf{R}_0$ comprises all aspects of the measurement process; i.e., all instrument response functions. This includes the survey coverage, which describes the instrument's overall exposure to the observational area, and the instrument's PSF, which describes how a point source is imaged by the instrument.

The superposition of different components and the transition from continuous coordinates to some discrete pixelization, cf. Eq. (6), cause a severe loss of information about the original signal fields. In addition to that, measurement noise distorts the signal's imprint in the data. The individual photon counts per pixel can be assumed to follow a Poisson distribution $\mathcal{P}$ each. Therefore, the likelihood of the data $\mathbf{d}$ given an expected number of events λ is modeled as a product of statistically independent Poisson processes,

$$P(\mathbf{d}|\lambda) = \prod_i \mathcal{P}(d_i, \lambda_i) = \prod_i \frac{1}{d_i!} \lambda_i^{d_i} e^{-\lambda_i}. \quad (7)$$

The Poisson distribution has a signal-to-noise ratio of $\sqrt{\lambda}$ which scales with the expected number of photon counts. Therefore, Poissonian shot noise is most severe in regions with low photon fluxes. This makes the detection of faint sources in high energy astronomy a particularly challenging task, as X- and γ -ray photons are sparse.

The likelihood of photon count data given a two component photon flux is hence described by the Eqs. (5) and (7). Rewriting this likelihood $P(\mathbf{d}|\mathbf{s}, \mathbf{u})$ in form of its negative logarithm yields

³ This expectation value is computed by a path integral, $\int \mathcal{D}\mathbf{s}$, over the complete phase space of the signal field $\mathbf{s}$; i.e., all possible field configurations.

the information Hamiltonian $H(\mathbf{d}|\mathbf{s}, \mathbf{u})^4$,

$$H(\mathbf{d}|\mathbf{s}, \mathbf{u}) = -\log P(\mathbf{d}|\mathbf{s}, \mathbf{u}) \quad (8)$$

$$= H_0 + \mathbf{1}^\dagger \boldsymbol{\lambda} - \mathbf{d}^\dagger \log(\boldsymbol{\lambda}) \quad (9)$$

$$= H_0 + \mathbf{1}^\dagger \mathbf{R}(\mathbf{e}^{\mathbf{s}} + \mathbf{e}^{\mathbf{u}}) - \mathbf{d}^\dagger \log(\mathbf{R}(\mathbf{e}^{\mathbf{s}} + \mathbf{e}^{\mathbf{u}})), \quad (10)$$

where the ground state energy H_0 comprises all terms constant in $\mathbf{s}$ and $\mathbf{u}$, and $\mathbf{1}$ is a constant data vector being 1 everywhere.

2.3. Prior assumptions

The diffuse and point-like signal fields, $\mathbf{s}$ and $\mathbf{u}$, contribute equally to the likelihood defined by Eq. (10), and thus leaving it completely degenerate. On the mere basis of the likelihood, the full data set could be explained by the diffuse signal alone, or only by point-sources, or any other conceivable combination. In order to downweight intuitively implausible solutions, we introduce priors. The priors discussed in the following address the morphology of the different photon flux contributions, and define diffuse and point-like in the first place. These priors aid the reconstruction by providing some remedy for the degeneracy of the likelihood. The likelihood describes noise and convolution properties, and the prior describe the individual morphological properties. Therefore, the denoising and deconvolution of the data towards the total photon flux $\boldsymbol{\rho}$ is primarily likelihood driven, but for a decomposition of the total photon flux into $\boldsymbol{\rho}^{(s)}$ and $\boldsymbol{\rho}^{(u)}$, the signal priors are imperative.

2.3.1. Diffuse component

The diffuse photon flux, $\boldsymbol{\rho}^{(s)} = \rho_0 \mathbf{e}^{\mathbf{s}}$, is strictly positive and might vary in intensity over several orders of magnitude. Its morphology shows cloudy patches with smooth fluctuations across spatial scales; i.e., one expects similar values of the diffuse flux in neighboring locations. In other words, the diffuse component exhibits spatial correlations. A log-normal model for $\boldsymbol{\rho}^{(s)}$ satisfies those requirements according to the maximum entropy principle (Oppermann et al. 2013; Kinney 2014). If the diffuse photon flux follows a multivariate log-normal distribution, the diffuse signal field $\mathbf{s}$ obeys a multivariate Gaussian distribution $\mathcal{G}$,

$$P(\mathbf{s}|\mathbf{S}) = \mathcal{G}(\mathbf{s}, \mathbf{S}) = \frac{1}{\sqrt{\det[2\pi\mathbf{S}]}} \exp\left(-\frac{1}{2}\mathbf{s}^\dagger \mathbf{S}^{-1} \mathbf{s}\right), \quad (11)$$

with a given covariance $\mathbf{S} = \langle \mathbf{s} \mathbf{s}^\dagger \rangle_{(\mathbf{s}|\mathbf{S})}$. This covariance describes the strength of the spatial correlations, and thus the smoothness of the fluctuations.

A convenient parametrization of the covariance $\mathbf{S}$ can be found, if the signal field $\mathbf{s}$ is a priori not known to distinguish any position or orientation axis; i.e., its correlations only depend on relative distances. This is equivalent to assume $\mathbf{s}$ to be statistically homogeneous and isotropic. Under this assumption, $\mathbf{S}$ is diagonal in the harmonic basis⁵ of the position space Ω such that

$$\mathbf{S} = \sum_k \mathbf{e}^{\tau_k} \mathbf{S}_k, \quad (12)$$

where τ_k are spectral parameters and $\mathbf{S}_k$ are projections onto a set of disjoint harmonic subspaces of Ω . These subspaces are commonly denoted as spectral bands or harmonic modes. The set of spectral parameters, $\boldsymbol{\tau} = \{\tau_k\}_k$, is then the logarithmic power spectrum of the diffuse signal field $\mathbf{s}$ with respect to the chosen harmonic basis denoted by k .

However, the diffuse signal covariance is in general unknown a priori. This requires the introduction of another prior for the covariance, or for the set of parameters $\boldsymbol{\tau}$ describing it adequately. This approach of hyperpriors on prior parameters creates a hierarchical parameter model.

2.3.2. Unknown power spectrum

The lack of knowledge of the power spectrum, requires its reconstruction from the same data the signal is inferred from (Wandelt et al. 2004; Jasche et al. 2010; Enßlin & Frommert 2011; Jasche & Wandelt 2013). Therefore, two a priori constraints for the spectral parameters $\boldsymbol{\tau}$, which describe the logarithmic power spectrum, are incorporated in the model.

The power spectrum is unknown and might span over several orders of magnitude. This implies a logarithmically uniform prior for each element of the power spectrum, and a uniform prior for each spectral parameter τ_k , respectively. We initially assume independent inverse-Gamma distributions $\mathcal{I}$ for the individual elements,

$$P(\mathbf{e}^{\boldsymbol{\tau}}|\boldsymbol{\alpha}, \mathbf{q}) = \prod_k \mathcal{I}(\mathbf{e}^{\tau_k}, \alpha_k, q_k) \quad (13)$$

$$= \prod_k \frac{q_k^{\alpha_k-1}}{\Gamma(\alpha_k-1)} \mathbf{e}^{-(\alpha_k \tau_k + q_k \mathbf{e}^{-\tau_k})}, \quad (14)$$

and hence

$$P_{\text{un}}(\boldsymbol{\tau}|\boldsymbol{\alpha}, \mathbf{q}) = \prod_k \mathcal{I}(\mathbf{e}^{\tau_k}, \alpha_k, q_k) \left| \frac{d\mathbf{e}^{\tau_k}}{d\tau_k} \right| \quad (15)$$

$$\propto \exp\left(-(\alpha-1)^\dagger \boldsymbol{\tau} - \mathbf{q}^\dagger \mathbf{e}^{-\boldsymbol{\tau}}\right), \quad (16)$$

where $\boldsymbol{\alpha} = \{\alpha_k\}_k$ and $\mathbf{q} = \{q_k\}_k$ are the shape and scale parameters, and Γ denotes the Gamma function. In the limit of $\alpha_k \rightarrow 1$ and $q_k \rightarrow 0 \forall k$, the inverse-Gamma distributions become asymptotically flat on a logarithmic scale, and thus P_{un} constant⁶. Small non-zero scale parameters, $0 < q_k$, provide lower limits for the power spectrum that, in practice, lead to more stable inference algorithms.

So far, the variability of the individual elements of the power spectrum is accounted for, but the question about their correlations has not been addressed. Empirically, power spectra of a diffuse signal field do not exhibit wild fluctuation or change drastically over neighboring modes. They rather show some sort of spectral smoothness. Moreover, for diffuse signal fields that were shaped by local and causal processes, we might expect a finite correlation support in position space. This translates into a smooth power spectrum. In order to incorporate spectral smoothness, we employ a prior introduced by Enßlin & Frommert (2011); Oppermann et al. (2013). This prior is based on the second logarithmic derivative of the spectral parameters $\boldsymbol{\tau}$, and favors power spectra that obey a power law. It reads

$$P_{\text{sm}}(\boldsymbol{\tau}|\boldsymbol{\sigma}) \propto \exp\left(-\frac{1}{2}\boldsymbol{\tau}^\dagger \mathbf{T} \boldsymbol{\tau}\right), \quad (17)$$

⁴ Throughout this work we define $H(\cdot) = -\log P(\cdot)$, and absorb constant terms into a normalization constant H_0 in favor of clarity.

⁵ The basis in which the Laplace operator is diagonal is denoted harmonic basis. If Ω is a n -dimensional Euclidean space $\mathbb{R}^n$ or Torus $\mathcal{T}^n$, the harmonic basis is the Fourier basis; if Ω is the $\mathcal{S}^2$ sphere, the harmonic basis is the spherical harmonics basis.

⁶ If $P(\tau_k = \log z) = \text{const.}$, then a substitution yields $P(z) = P(\log z) |d(\log z)/dz| \propto z^{-1} \sim \mathcal{I}(z, \alpha \rightarrow 1, q \rightarrow 0)$.

with

$$\tau^\dagger \mathbf{T} \tau = \int d(\log k) \frac{1}{\sigma_k^2} \left(\frac{\partial^2 \tau_k}{\partial (\log k)^2} \right)^2, \quad (18)$$

where $\sigma = \{\sigma_k\}_k$ are Gaussian standard deviations specifying the tolerance against deviation from a power-law behavior of the power spectrum. A choice of $\sigma_k = 1 \forall k$ would typically allow for a change in the power law's slope of 1 per e-fold in k . In the limit of $\sigma_k \rightarrow \infty \forall k$, no smoothness is enforced upon the power spectrum.

The resulting prior for the spectral parameters is given by the product of the priors discussed above,

$$P(\tau|\alpha, \mathbf{q}, \sigma) = P_{\text{un}}(\tau|\alpha, \mathbf{q}) P_{\text{sm}}(\tau|\sigma). \quad (19)$$

The parameters α , $\mathbf{q}$ and σ are considered to be given as part of the hierarchical Bayesian model, and provide a flexible handle to model our knowledge on the scaling and smoothness of the power spectrum.

2.3.3. Point-like component

The point-like photon flux, $\rho^{(u)} = \rho_0 e^u$, is supposed to originate from very distant astrophysical sources. These sources appear morphologically point-like to an observer because their actual extent is negligible given the extreme distances. This renders point sources to be spatially local phenomena. The photon flux contributions of neighboring point sources can (to zeroth order approximation) be assumed to be statistically independent of each other. Even if the two sources are very close on the observational plane, their physical distance might be huge. Even in practice, the spatial cross-correlation of point sources is negligible. Therefore, statistically independent priors for the photon flux contribution of each point-source are assumed in the following.

Because of the spatial locality of a point source, the corresponding photon flux signal is supposed to be confined to a single spot, too. If the point-like signal field, defined over a continuous position space Ω , is discretized properly⁷, this spot is sufficiently identified by an image pixel in the reconstruction. A discretization, $\rho(x \in \Omega) \rightarrow (\rho_x)_x$, is an inevitable step since the algorithm is to be implemented in a computer environment anyway. Nevertheless, we have to ensure that the a priori assumptions do not depend on the chosen discretization but satisfy the continuous limit.

Therefore, the prior for the point-like signal component factorizes spatially,

$$P(\rho^{(u)}) = \prod_x P(\rho_x^{(u)}), \quad (20)$$

but the functional form of the priors are yet to be determined. This model allows the point-like signal field to host one point source in every pixel. Most of these point sources are expected to be invisibly faint contributing negligibly to the total photon flux. However, the point sources which are just identifiable from the data are pinpointed in the reconstruction. In this approach, there is no necessity for a complicated determination of the number and position of sources.

⁷ The numerical discretization of information fields is described in great detail in Selig et al. (2013).

For the construction of a prior, that the photon flux is a strictly positive quantity also needs to be considered. Thus, a simple exponential prior,

$$P(\rho_x^{(u)}) \propto \exp(-\rho_x^{(u)}/\rho_0), \quad (21)$$

has been suggested (e.g., Guglielmetti et al. 2009). It has the advantage of being (easily) analytically treatable, but its physical implications are questionable. This distribution strongly suppresses high photon fluxes in favor of lower ones. The maximum entropy prior, which is also often applied, is even worse because it corresponds to a brightness distribution⁸,

$$P(\rho_x^{(u)}) \propto (\rho_x^{(u)}/\rho_0)^{(-\rho_x^{(u)}/\rho_0)}. \quad (22)$$

The following (rather crude) consideration might motivate a more astrophysical prior. Say the universe hosts a homogeneous distribution of point sources. The number of point sources would therefore scale with the observable volume; i.e., with distance cubed. Their apparent brightness, which is reduced because of the spreading of the light rays; i.e., a proportionality to the distance squared. Consequently, a power-law behavior between the number of point sources and their brightness with a slope $\beta = \frac{3}{2}$ is to be expected (Fomalont 1968; Malyshev & Hogg 2011). However, such a plain power law diverges at 0, and is not necessarily normalizable. Furthermore, Galactic and extragalactic sources cannot be found in arbitrary distances owing to the finite size of the Galaxy and the cosmic (past) light cone. Imposing an exponential cut-off above 0 onto the power law yields an inverse-Gamma distribution, which has been shown to be an appropriate prior for point-like photon fluxes (Guglielmetti et al. 2009; Carvalho et al. 2009, 2012).

The prior for the point-like signal field is therefore derived from a product of independent inverse-Gamma distributions⁹,

$$P(\rho^{(u)}|\beta, \eta) = \prod_x \mathcal{I}(\rho_x^{(u)}, \beta_x, \rho_0 \eta_x) \quad (23)$$

$$= \prod_x \frac{(\rho_0 \eta_x)^{\beta_x-1}}{\Gamma(\beta_x-1)} (\rho_x^{(u)})^{-\beta_x} \exp\left(-\frac{\rho_0 \eta_x}{\rho_x^{(u)}}\right), \quad (24)$$

yielding

$$P(\mathbf{u}|\beta, \eta) = \prod_x \mathcal{I}(\rho_0 e^{u_x}, \beta_x, \rho_0 \eta_x) \left| \frac{d\rho_0 e^{u_x}}{du_x} \right| \quad (25)$$

$$\propto \exp\left(-(\beta-1)^\dagger \mathbf{u} - \eta^\dagger e^{-\mathbf{u}}\right), \quad (26)$$

where $\beta = \{\beta_x\}_x$ and $\eta = \{\eta_x\}_x$ are the shape and scale parameters. The latter is responsible for the cut-off of vanishing fluxes, and should be chosen adequately small in analogy to the spectral scale parameters $\mathbf{q}$. The determination of the shape parameters is more difficile. The geometrical argument above suggests a universal shape parameter, $\beta_x = \frac{3}{2} \forall x$. A second argument for this value results from demanding a priori independence of the discretization. If we choose a coarser resolution that would add up the flux from two point sources at merged pixels, then our prior should still be applicable. The universal value of $\frac{3}{2}$ indeed fulfills this requirement as shown in Appendix A. There it is also shown that η has to be chosen resolution dependent, though.

⁸ The so-called maximum entropy regularization $\sum_x (\rho_x^{(u)}/\rho_0) \log(\rho_x^{(u)}/\rho_0)$ of the log-likelihood can be regarded as log-prior, cf. Eqs. (20) and (22).

⁹ A possible extension of this prior model that includes spatial correlations would be an inverse-Wishart distribution for $\text{diag}[\rho^{(u)}]$.

Fig. 2. Graphical model of the model parameters α , q , σ , β , and η , the logarithmic spectral parameters τ , the diffuse signal field s , the point-like signal field u , the total photon flux ρ , the expected number of photons λ , and the observed photon count data d .

2.4. Parameter model

Figure 2 gives an overview of the parameter hierarchy of the suggested Bayesian model. The data d is given, and the diffuse signal field s and the point-like signal field u shall be reconstructed from that data. The logarithmic power spectrum τ is a set of nuisance parameters that also need to be reconstructed from the data in order to accurately model the diffuse flux contributions. The model parameters form the top layer of this hierarchy and are given to the reconstruction algorithm. This set of model parameters can be boiled down to five scalars, namely α , q , σ , β , and η , if one defines $\alpha = \alpha 1$, etc. The incorporation of the scalars in the inference is possible in theory, but this would increase the computational complexity dramatically.

We discussed reasonable values for these scalars to be chosen a priori. If additional information sources, such as theoretical power spectra or object catalogs, are available the model parameters can be adjusted accordingly. In Sect. 4, different parameter choices for the analysis of simulated data are investigated.

3. Denoising, deconvolution, and decomposition

The likelihood model, describing the measurement process, and the prior assumptions for the signal fields and the power spectrum of the diffuse component yield a well-defined inference problem. The corresponding posterior is given by

$$P(s, \tau, u|d) = \frac{P(d|s, u) P(s|\tau) P(\tau|\alpha, q, \sigma) P(u|\beta, \eta)}{P(d)}, \quad (27)$$

which is a complex form of Bayes' theorem (1).

Ideally, we would now calculate the a posteriori expectation values and uncertainties according to Eqs. (2) and (3) for the diffuse and point-like signal fields, s and u , as well as for the logarithmic spectral parameters τ . However, an analytical evaluation of these expectation values is not possible because of the complexity of the posterior.

The posterior is non-linear in the signal fields and, except for artificially constructed data, non-convex. It, however, is more flexible and therefore allows for a more comprehensive description of the parameters to be inferred (Kirkpatrick et al. 1983; Geman & Geman 1984).

Numerical approaches involving Markov chain Monte Carlo methods (Metropolis & Ulam 1949; Metropolis et al. 1953)

are possible, but hardly feasible because of the huge parameter phase space. Nevertheless, similar problems have been addressed by elaborate sampling techniques (Wandelt et al. 2004; Jasche et al. 2010; Jasche & Kitaura 2010; Jasche & Wandelt 2013).

Here, two approximative algorithms with lower computational costs are derived. The first one uses the maximum a posteriori (MAP) approximation, the second one minimizes the Gibbs free energy of an approximate posterior ansatz in the spirit of variational Bayesian methods. The fidelity and accuracy of these two algorithms are compared in a numerical application in Sect. 4.

3.1. Posterior maximum

The posterior maximum and mean coincide, if the posterior distribution is symmetric and single peaked. In practice, this often holds – at least in good approximation –, so that the maximum a posteriori approach can provide suitable estimators. This can either be achieved using a δ -distribution at the posterior's mode,

$$\langle s \rangle_{(s|d)} \stackrel{\text{MAP-}\delta}{\approx} \int \mathcal{D}s \, s \, \delta(s - s_{\text{mode}}), \quad (28)$$

or using a Gaussian approximation around this point,

$$\langle s \rangle_{(s|d)} \stackrel{\text{MAP-G}}{\approx} \int \mathcal{D}s \, s \, \mathcal{G}(s - s_{\text{mode}}, \mathbf{D}_{\text{mode}}). \quad (29)$$

Both approximations require us to find the mode, which is done by extremizing the posterior.

Instead of the complex posterior distribution, it is convenient to consider the information Hamiltonian, defined by its negative logarithm,

$$H(s, \tau, u|d) = -\log P(s, \tau, u|d) \quad (30)$$

$$\begin{aligned} &= H_0 + 1^\dagger \mathbf{R} (e^s + e^u) - d^\dagger \log(\mathbf{R} (e^s + e^u)) \\ &+ \frac{1}{2} \log(\det[\mathbf{S}]) + \frac{1}{2} s^\dagger \mathbf{S}^{-1} s \\ &+ (\alpha - 1)^\dagger \tau + q^\dagger e^{-\tau} + \frac{1}{2} \tau^\dagger \mathbf{T} \tau \\ &+ (\beta - 1)^\dagger u + \eta^\dagger e^{-u}, \end{aligned} \quad (31)$$

where all terms constant in s , τ , and u have been absorbed into a ground state energy H_0 , cf. Eqs. (7), (11), (19), and (26), respectively.

The MAP solution, which maximizes the posterior, minimizes the Hamiltonian. This minimum can thus be found by taking the first (functional) derivatives of the Hamiltonian with respect to s , τ , and u and equating them with zero. Unfortunately, this yields a set of implicit, self-consistent equations rather than an explicit solution. However, these equations can be solved by an iterative minimization of the Hamiltonian using a steepest descent method for example, see Sect. 3.4 for details.

In order to better understand the structure of the MAP solution, we consider the minimum $(s, \tau, u) = (\mathbf{m}^{(s)}, \tau^*, \mathbf{m}^{(u)})$. The resulting filter formulas for the diffuse and point-like signal field read

$$\left. \frac{\partial H}{\partial s} \right|_{\min} = \mathbf{0} = (1 - d/I)^\dagger \mathbf{R} * e^{\mathbf{m}^{(s)}} + \mathbf{S}^{*-1} \mathbf{m}^{(s)}, \quad (32)$$

$$\left. \frac{\partial H}{\partial u} \right|_{\min} = \mathbf{0} = (1 - d/I)^\dagger \mathbf{R} * e^{\mathbf{m}^{(u)}} + \beta - 1 - \eta * e^{-\mathbf{m}^{(u)}}, \quad (33)$$

with

$$l = \mathbf{R} \left(e^{m^{(s)}} + e^{m^{(u)}} \right), \quad (34)$$

$$\mathbf{S}^* = \sum_k e^{\tau_k^*} \mathbf{S}_k. \quad (35)$$

Here, $*$ and $/$ denote componentwise multiplication and division, respectively. The first term in Eqs. (32) and (33), which comes from the likelihood, vanishes in case $l = d$. We note that $l = \lambda_{\min}$ describes the most likely number of photon counts, not the expected number of photon counts $\lambda = \langle d \rangle_{(d|s,u)}$, cf. Eqs. (5) and (7). Disregarding the regularization by the priors, the solution would overfit; i.e., noise features are partly assigned to the signal fields in order to achieve an unnecessarily close agreement with the data. However, the a priori regularization suppresses this tendency to some extent.

The second derivative of the Hamiltonian describes the curvature around the minimum, and therefore approximates the (inverse) uncertainty covariance,

$$\left. \frac{\partial^2 H}{\partial s \partial s^\dagger} \right|_{\min} \approx \mathbf{D}^{(s)-1}, \quad \left. \frac{\partial^2 H}{\partial u \partial u^\dagger} \right|_{\min} \approx \mathbf{D}^{(u)-1}. \quad (36)$$

The closed form of $\mathbf{D}^{(s)}$ and $\mathbf{D}^{(u)}$ is given explicitly in Appendix B.

The filter formula for the power spectrum, which is derived from a first derivative of the Hamiltonian with respect to τ , yields

$$e^{\tau^*} = \frac{\mathbf{q} + \frac{1}{2} \left(\text{tr} \left[\mathbf{m}^{(s)} \mathbf{m}^{(s)\dagger} \mathbf{S}_k^{-1} \right] \right)_k}{\gamma + \mathbf{T} \tau^*}, \quad (37)$$

where $\gamma = (\alpha - 1) + \frac{1}{2} \left(\text{tr} \left[\mathbf{S}_k \mathbf{S}_k^{-1} \right] \right)_k$. This formula is in accordance with the results by Enßlin & Frommert (2011); Oppermann et al. (2013). It has been shown by the former authors that such a filter exhibits a perception threshold; i.e., on scales where the signal-response-to-noise ratio drops below a certain bound the reconstructed signal power becomes vanishingly low. This threshold can be cured by a better capture of the a posteriori uncertainty structure.

3.2. Posterior approximation

In order to overcome the analytical infeasibility as well as the perception threshold, we seek an approximation to the true posterior. Instead of approximating the expectation values of the posterior, approximate posteriors are investigated in this section. In case the approximation is good, the expectation values of the approximate posterior should then be close to the real ones.

The posterior given by Eq. (27) is inaccessible because of the entanglement of the diffuse signal field s , its logarithmic power spectrum τ , and the point-like signal field u . The involvement of τ can be simplified by a mean field approximation,

$$P(s, \tau, u|d) \approx Q = Q_s(s, u|\mu, d) Q_\tau(\tau|\mu, d), \quad (38)$$

where μ denotes an abstract mean field mediating some information between the signal field tuple (s, u) and τ that are separated by the product ansatz in Eq. (38). This mean field is fully determined by the problem, as it represents effective (rather than additional) degrees of freedom. It is only needed implicitly for the derivation, an explicit formula can be found in Appendix C.3, though.

Since the a posteriori mean estimates for the signal fields and their uncertainty covariances are of primary interest, a Gaussian

approximation for Q_s that accounts for correlation between s and u would be sufficient. Hence, our previous approximation is extended by setting

$$Q_s(s, u|\mu, d) = \mathcal{G}(\varphi, \mathbf{D}), \quad (39)$$

with

$$\varphi = \begin{pmatrix} s - m^{(s)} \\ u - m^{(u)} \end{pmatrix}, \quad \mathbf{D} = \begin{pmatrix} \mathbf{D}^{(s)} & \mathbf{D}^{(su)} \\ \mathbf{D}^{(su)\dagger} & \mathbf{D}^{(u)} \end{pmatrix}. \quad (40)$$

This Gaussian approximation is also a convenient choice in terms of computational complexity because of its simple analytic structure.

The goodness of the approximation $P \approx Q$ can be quantified by an information theoretical measure, see Appendix C.1. The Gibbs free energy of the inference problem,

$$G = \langle H(s, \tau, u|d) \rangle_Q - \langle -\log Q(s, \tau, u|d) \rangle_Q, \quad (41)$$

which is equivalent to the Kullback-Leibler divergence $D_{\text{KL}}(Q, P)$, is chosen as such a measure (Enßlin & Weig 2010).

In favor of comprehensibility, we suppose the solution for the logarithmic power spectrum τ^* is known for the moment. The Gibbs free energy is then calculated by plugging in the Hamiltonian, and evaluating the expectation values¹⁰,

$$G = G_0 + \langle H(s, u|d) \rangle_{Q_s} - \frac{1}{2} \log(\det[\mathbf{D}]) \quad (42)$$

$$\begin{aligned} &= G_1 + l^\dagger l - d^\dagger \left\{ \log(l) - \sum_{v=2}^{\infty} \frac{(-1)^v}{v} \langle (\lambda/l - 1)^v \rangle_{Q_s} \right\} \\ &\quad + \frac{1}{2} m^{(s)\dagger} \mathbf{S}^{*-1} m^{(s)} + \frac{1}{2} \text{tr}[\mathbf{D}^{(s)} \mathbf{S}^{*-1}] \\ &\quad + (\beta - 1)^\dagger m^{(u)} + \eta^\dagger e^{-m^{(u)} + \frac{1}{2} \hat{\mathbf{D}}^{(u)}} \\ &\quad - \frac{1}{2} \log(\det[\mathbf{D}]), \end{aligned} \quad (43)$$

with

$$\lambda = \mathbf{R}(e^s + e^u), \quad (44)$$

$$l = \langle \lambda \rangle_{Q_s} = \mathbf{R} \left(e^{m^{(s)} + \frac{1}{2} \hat{\mathbf{D}}^{(s)}} + e^{m^{(u)} + \frac{1}{2} \hat{\mathbf{D}}^{(u)}} \right), \quad (45)$$

$$\mathbf{S}^* = \sum_k e^{\tau_k^*} \mathbf{S}_k, \text{ and} \quad (46)$$

$$\hat{\mathbf{D}} = \text{diag}[\mathbf{D}]. \quad (47)$$

Here, G_0 and G_1 carry all terms independent of s and u . In comparison to the Hamiltonian given in Eq. (31), there are a number of correction terms that now also consider the uncertainty covariances of the signal estimates properly. For example, the expectation values of the photon fluxes differ comparing l in Eqs. (34) and (45) where it now describes the expectation value of λ over the approximate posterior. In case $l = \lambda$ the explicit sum in Eq. (43) vanishes. Since this sum includes powers of $\langle \lambda^{v>2} \rangle_{Q_s}$

¹⁰ The second likelihood term in Eq. (43), $d^\dagger \log(\lambda)$, is thereby expanded according to

$$\begin{aligned} \log(x) &= \log \langle x \rangle - \sum_{v=2}^{\infty} \frac{(-1)^v}{v} \left\langle \left(\frac{x}{\langle x \rangle} - 1 \right)^v \right\rangle \\ &\approx \log \langle x \rangle + \mathcal{O}(\langle x^2 \rangle), \end{aligned}$$

under the assumption $x \approx \langle x \rangle$.

its evaluation would require all entries of $\mathbf{D}$ to be known explicitly. In order to keep the algorithm computationally feasible, this sum shall hereafter be neglected. This is equivalent to truncating the corresponding expansion at second order; i.e., $\nu = 2$. It can be shown that, in consequence of this approximation, the cross-correlation $\mathbf{D}^{(su)}$ equals zero, and $\mathbf{D}$ becomes block diagonal.

Without these second order terms, the Gibbs free energy reads

$$\begin{aligned} G = & G_1 + 1^\dagger \mathbf{l} - d^\dagger \log(\mathbf{l}) \\ & + \frac{1}{2} \mathbf{m}^{(s)\dagger} \mathbf{S}^{\star-1} \mathbf{m}^{(s)} + \frac{1}{2} \text{tr} [\mathbf{D}^{(s)} \mathbf{S}^{\star-1}] \\ & + (\beta - 1)^\dagger \mathbf{m}^{(u)} + \eta^\dagger e^{-\mathbf{m}^{(u)} + \frac{1}{2} \hat{\mathbf{D}}^{(u)}} \\ & - \frac{1}{2} \log(\det[\mathbf{D}^{(s)}]) - \frac{1}{2} \log(\det[\mathbf{D}^{(u)}]). \end{aligned} \quad (48)$$

Minimizing the Gibbs free energy with respect to $\mathbf{m}^{(s)}$, $\mathbf{m}^{(u)}$, $\mathbf{D}^{(s)}$, and $\mathbf{D}^{(u)}$ would optimize the fitness of the posterior approximation $P \approx Q$. Filter formulas for the Gibbs solution can be derived by taking the derivative of G with respect to the approximate mean estimates,

$$\frac{\partial G}{\partial \mathbf{m}^{(s)}} = \mathbf{0} = (1 - d/l)^\dagger \mathbf{R} * e^{\mathbf{m}^{(s)} + \frac{1}{2} \hat{\mathbf{D}}^{(s)}} + \mathbf{S}^{\star-1} \mathbf{m}^{(s)}, \quad (49)$$

$$\begin{aligned} \frac{\partial G}{\partial \mathbf{m}^{(u)}} = \mathbf{0} = & (1 - d/l)^\dagger \mathbf{R} * e^{\mathbf{m}^{(u)} + \frac{1}{2} \hat{\mathbf{D}}^{(u)}} \\ & + \beta - 1 - \eta * e^{-\mathbf{m}^{(u)} + \frac{1}{2} \hat{\mathbf{D}}^{(u)}}. \end{aligned} \quad (50)$$

This filter formulas again account for the uncertainty of the mean estimates in comparison to the MAP filter formulas in Eqs. (32) and (33). The uncertainty covariances can be constructed by either taking the second derivatives,

$$\frac{\partial^2 G}{\partial \mathbf{m}^{(s)} \partial \mathbf{m}^{(s)\dagger}} \approx \mathbf{D}^{(s)-1}, \quad \frac{\partial^2 G}{\partial \mathbf{m}^{(u)} \partial \mathbf{m}^{(u)\dagger}} \approx \mathbf{D}^{(u)-1}, \quad (51)$$

or setting the first derivatives of G with respect to the uncertainty covariances equal to zero matrices,

$$\frac{\partial G}{\partial D_{xy}^{(s)}} = 0, \quad \frac{\partial G}{\partial D_{xy}^{(u)}} = 0. \quad (52)$$

The closed form of $\mathbf{D}^{(s)}$ and $\mathbf{D}^{(u)}$ is given explicitly in Appendix B.

So far, the logarithmic power spectrum τ^* , and with it $\mathbf{S}^*$, have been supposed to be known. The mean field approximation in Eq. (38) does not specify the approximate posterior $Q_\tau(\tau|\mu, d)$, but it can be retrieved by variational Bayesian methods (Jordan et al. 1999; Wingate & Weber 2013), according to the procedure detailed in Appendix C.2. The subsequent Appendix C.3 discusses the derivation of an solution for τ by extremizing Q_τ . This result, which was also derived in Oppermann et al. (2013), applies to the inference problem discussed here, yielding

$$e^{\tau^*} = \frac{\mathbf{q} + \frac{1}{2} \left(\text{tr} \left[\left(\mathbf{m}^{(s)} \mathbf{m}^{(s)\dagger} + \mathbf{D}^{(s)} \right) \mathbf{S}_k^{-1} \right] \right)_k}{\gamma + \mathbf{T} \tau^*}. \quad (53)$$

Again, this solution includes a correction term in comparison to the MAP solution in Eq. (37). Since $\mathbf{D}^{(s)}$ is positive definite, it contributes positive to the (logarithmic) power spectrum, and therefore reduces the possible perception threshold further.

We note that this is a minimal Gibbs free energy solution that maximizes Q_τ . A proper calculation of $\langle \tau \rangle_{Q_\tau}$ might include further correction terms, but their derivation is not possible in closed form. Moreover, the above used diffuse signal covariance $\mathbf{S}^{\star-1}$ should be replaced by $\langle \mathbf{S}^{-1} \rangle_{Q_\tau}$, adding further correction terms to the filter formulas.

In order to keep the computational complexity on a feasible level, all these higher order corrections are not considered here. The detailed characterization of their implications and implementation difficulties is left for future investigation.

3.3. Physical flux solution

To perform calculations on the logarithmic fluxes is convenient for numerical reasons, but it is the physical fluxes that are actually of interest to us. Given the chosen approximation, we can compute the posterior expectation values of the diffuse and point-like photon flux, $\rho^{(s)}$ and $\rho^{(u)}$, straight forwardly,

$$\langle \rho^{(\cdot)} \rangle_P \stackrel{\text{MAP-}\delta}{\approx} \langle \rho^{(\cdot)} \rangle_\delta = \rho_0 e^{\mathbf{m}_{\text{mode}}^{(\cdot)}}, \quad (54)$$

$$\stackrel{\text{MAP-}\mathcal{G}}{\approx} \langle \rho^{(\cdot)} \rangle_{\mathcal{G}} = \rho_0 e^{\mathbf{m}_{\text{mode}}^{(\cdot)} + \frac{1}{2} \hat{\mathbf{D}}_{\text{mode}}^{(\cdot)}}, \quad (55)$$

$$\stackrel{\text{Gibbs}}{\approx} \langle \rho^{(\cdot)} \rangle_Q = \rho_0 e^{\mathbf{m}_{\text{mean}}^{(\cdot)} + \frac{1}{2} \hat{\mathbf{D}}_{\text{mean}}^{(\cdot)}}, \quad (56)$$

in accordance with Eqs. (28), (29), or (38), respectively. Those solutions differ from each other in terms of the involvement of the posterior's mode or mean, and in terms of the inclusion of the uncertainty information, see subscripts.

In general, the mode approximation holds for symmetric, single peaked distributions, but can perform poorly in other cases (e.g., Enßlin & Frommert 2011). The exact form of the posterior considered here is highly complex because of the many degrees of freedom. In a dimensionally reduced frame, however, the posterior appears single peaked and exhibits a negative skewness¹¹. Although this is not necessarily generalizable, it suggest a superiority of the posterior mean compared to the MAP because of the asymmetry of the distribution. Nevertheless, the MAP approach is computationally cheaper compared to the Gibbs approach that requires permanent knowledge of the uncertainty covariance.

The uncertainty of the reconstructed photon flux can be approximated as for an ordinary log-normal distribution,

$$\langle \rho^{(\cdot)2} \rangle_P - \langle \rho^{(\cdot)} \rangle_P^2 \stackrel{\text{MAP}}{\approx} \langle \rho^{(\cdot)} \rangle_{\mathcal{G}}^2 \left(e^{\hat{\mathbf{D}}_{\text{mode}}^{(\cdot)}} - 1 \right), \quad (57)$$

$$\stackrel{\text{Gibbs}}{\approx} \langle \rho^{(\cdot)} \rangle_Q^2 \left(e^{\hat{\mathbf{D}}_{\text{mean}}^{(\cdot)}} - 1 \right), \quad (58)$$

where the square root of the latter term would describe the relative uncertainty.

3.4. Imaging algorithm

The problem of denoising, deconvolving, and decomposing photon observations is a non-trivial task. Therefore, this section discusses the implementation of the D³PO algorithm given the two sets of filter formulas derived in Sects. 3.1 and 3.2, respectively.

The information Hamiltonian, or equivalently the Gibbs free energy, are scalar quantities defined over a huge phase space

¹¹ For example, the posterior $P(s|d)$ for a one-dimensional diffuse signal is proportional to $\exp(-\frac{1}{2}s^2 + ds - \exp(s))$, whereby all other parameters are fixed to unity. Analogously, $P(u|d) \propto \exp(du - 2 \cosh(u))$.

of possible field and parameter configurations including, among others, the elements of $\mathbf{m}^{(s)}$ and $\mathbf{m}^{(u)}$. If we only consider those, and no resolution refinement from data to signal space, two numbers need to be inferred from one data value each. Including τ and the uncertainty covariances $\mathbf{D}^{(s)}$ and $\mathbf{D}^{(u)}$ in the inference, the problem of underdetermined degrees of freedom gets worse. This is reflected in the possibility of a decent number of local minima in the non-convex manifold landscape of the codomain of the Hamiltonian, or Gibbs free energy, respectively (Kirkpatrick et al. 1983; Geman & Geman 1984; Giovannelli & Coulaos 2005). The complexity of the inference problem goes back to the, in general, non-linear entanglement between the individual parameters.

The D³PO algorithm is based on an iterative optimization scheme, where certain subsets of the problem are optimized alternately instead of the full problem at once. Each subset optimization is designed individually, see below. The global optimization cycle is in some degree sensitive to the starting values because of the non-convexity of the considered potential; i.e., the information Hamiltonian or Gibbs free energy, respectively. We can find such appropriate starting values by solving the inference problem in a reduced frame in advance, see below. So far, a step-by-step guide of the algorithm looks like the following.

1. Initialize the algorithm with primitive starting values; e.g., $m_x^{(s)} = m_x^{(u)} = 0$, $D_{xy}^{(s)} = D_{xy}^{(u)} = \delta_{xy}$, and $\tau_k^* = \log(k^{-2})$. – Those values are arbitrary. Although the optimization is rather insensitive to them, inappropriate values can cripple the algorithm for numerical reasons because of the high non-linearity of the inference problem.
2. Optimize $\mathbf{m}^{(s)}$, the diffuse signal field, coarsely. – The preliminary optimization shall yield a rough estimate of the diffuse only contribution. This can be achieved by reconstructing a coarse screened diffuse signal field that only varies on large scales; i.e., limiting the bandwidth of the diffuse signal in its harmonic basis. Alternatively, obvious point sources in the data could be masked out by introducing an artificial mask into the response, if feasible.
3. Optimize $\mathbf{m}^{(u)}$, the point-like signal field, locally. – This initial optimization shall approximate the brightest, most obvious, point sources that are visible in the data image by eye. Their current disagreement with the data dominates the considered potential, and introduces some numerical stiffness. The gradient of the potential can be computed according to Eqs. (33) or (50), and its minima will be at the expected position of the brightest point source which has not been reconstructed, yet. It is therefore very efficient to increase $\mathbf{m}^{(u)}$ at this location directly until the sign of the gradient flips, and repeat this procedure until the obvious point sources are fit.
4. Optimize $\mathbf{m}^{(u)}$, the point-like signal field. – This task can be done by a steepest descent minimization of the potential combined with a line search following the Wolfe conditions (Nocedal & Wright 2006). The potentials can be computed according to Eqs. (31) or (43) neglecting terms independent of $\mathbf{m}^{(u)}$, and the gradient according to Eqs. (33) or (50). A more sophisticated minimization scheme, such as a non-linear conjugate gradient (Shewchuk 1994), is conceivable but would require the application of the full Hessian, cf. step 5. In the first run, it might be sufficient to restrict the optimization to the locations identified in step 3.
5. Update $\hat{\mathbf{D}}^{(u)}$, the point-like uncertainty variance, in case of a Gibbs approach. – It is not feasible to compute the full uncertainty covariance $\mathbf{D}^{(u)}$ explicitly in order to extract its diagonal. A more elegant way is to apply a probing technique

relying on the application of $\mathbf{D}^{(u)}$ to random fields ξ that project out the diagonal (Hutchinson 1989; Selig et al. 2012). The uncertainty covariance is given as the inverse Hessian by Eqs. (36) or (51), and should be symmetric and positive definite. For that reason, it can be applied to a field using a conjugate gradient (Shewchuk 1994); i.e., solving $(\mathbf{D}^{(u)})^{-1}\mathbf{y} = \xi$ for $\mathbf{y}$. However, if the current phase space position is far away from the minimum, the Hessian is not necessarily positive definite. One way to overcome this temporal instability, would be to introduce a Levenberg damping in the Hessian (inspired by Transtrum et al. 2010; Transtrum & Sethna 2012).

6. Optimize $\mathbf{m}^{(s)}$, the diffuse signal field. – An analog scheme as in step 4 using steepest descent and Wolfe conditions is effective. The potentials can be computed according to Eqs. (31) or (43) neglecting terms independent of $\mathbf{m}^{(s)}$, and the gradient according to Eqs. (32) or (49), respectively. It has proven useful to first ensure a convergence on large scales; i.e., small harmonic modes k . This can be done repeating steps 6–8 for all $k < k_{\max}$ with growing $k_{\max}$ using the corresponding projections $\mathbf{S}_k$.
7. Update $\hat{\mathbf{D}}^{(s)}$, the diffuse uncertainty variance, in case of a Gibbs approach in analogy to step 5.
8. Optimize τ^* , the logarithmic power spectrum. – This is done by solving Eqs. (37) or (53). The trace term can be computed analog to the diagonal; e.g., by probing. Given this, the equation can be solved efficiently by a Newton-Raphson method.
9. Repeat the steps 4 to 8 until convergence. – This scheme will take several cycles until the algorithm reaches the desired convergence level. Therefore, it is not required to achieve a convergence to the final accuracy level in all subsets in all cycles. It is advisable to start with weak convergence criteria in the first loop and increase them gradually.

A few remarks are in order.

The phase space of possible signal field configurations is tremendously huge. It is therefore impossible to judge if the algorithm has converged to the global or some local minima, but this does not matter if both yield reasonable results that do not differ substantially.

In general, the converged solution is also subject to the choice of starting values. Solving a non-convex, non-linear inference problem without proper initialization can easily lead to nonsensical results, such as fitting (all) diffuse features by point sources. Therefore, the D³PO algorithm essentially creates its own starting values executing the initial steps 1 to 3. The primitive starting values are thereby processed to rough estimates that cover coarsely resolved diffuse and prominent point-like features. These estimates serve then as actual starting values for the optimization cycle.

Because of the iterative optimization scheme starting with the diffuse component in step 2, the algorithm might be prone to explaining some point-like features by diffuse sources. Starting with the point-like component instead would give rise to the opposite bias. To avoid such biases, it is advisable to restart the algorithm partially. To be more precise, we propose to discard the current reconstruction of $\mathbf{m}^{(u)}$ after finishing step 8 for the first time, then start the second iteration again with step 3, and to discard the current $\mathbf{m}^{(s)}$ before step 6.

The above scheme exploits a few numerical techniques, such as probing or Levenberg damping, that are described in great detail in the given references. The code of our implementation of the D³PO algorithm will be made public in the future under <http://www.mpa-garching.mpg.de/ift/d3po/>.

Fig. 3. Illustration of the data and noiseless, but reconvolved, signal responses of the reconstructions. Panel **a**) shows the data from a mock observation of a $32 \times 32 \text{ arcmin}^2$ patch of the sky with a resolution of 0.1 arcmin corresponding to a total of 102 400 pixels. The data had been convolved with a Gaussian-like PSF ($FWHM \approx 0.2 \text{ arcmin} = 2 \text{ pixels}$, finite support of 1.1 arcmin = 11 pixels) and masked because of an uneven exposure. Panel **b**) shows the centered convolution kernel. Panel **c**) shows the exposure mask. The bottom panels show the reconvolved signal response $\mathbf{R}(\rho)$ of a reconstruction using a different approach each, namely **d**) MAP- δ , **e**) MAP- $\mathcal{G}$, and **f**) Gibbs. All reconstructions shown here and in the following figures used the same model parameters: $\alpha = 1$, $q = 10^{-12}$, $\sigma = 10$, $\beta = \frac{3}{2}$, and $\eta = 10^{-4}$.

4. Numerical application

Exceeding the simple 1D scenario illustrated in Fig. 1, the D³PO algorithm is now applied to a realistic, but simulated, data set. The data set represents a high energy observation with a field of view of $32 \times 32 \text{ arcmin}^2$ and a resolution of 0.1 arcmin; i.e., the photon count image comprises 102 400 pixels. The instrument response includes the convolution with a Gaussian-like PSF with a FWHM of roughly 0.2 arcmin, and an uneven survey mask attributable to the inhomogeneous exposure of the virtual instrument. The data image and those characteristics are shown in Fig. 3.

In addition, the top panels of Fig. 3 show the reproduced signal responses of the reconstructed (total) photon flux. The reconstructions used the same model parameters, $\alpha = 1$, $q = 10^{-12}$, $\sigma = 10$, $\beta = \frac{3}{2}$, and $\eta = 10^{-4}$ in a MAP- δ , MAP- $\mathcal{G}$ and a Gibbs approach, respectively. They all show a very good agreement with the actual data, and differences are barely visible by eye. We note that only the quality of denoising is visible, since the signal response shows the convolved and superimposed signal fields.

The diffuse contribution to the deconvolved photon flux is shown Fig. 4 for all three estimators, cf. Eqs. (54) to (56). There, all point-like contributions as well as noise and instrumental

effects have been removed presenting a denoised, deconvolved and decomposed reconstruction result for the diffuse photon flux. Figure 4 also shows the absolute difference to the original flux. Although the differences in the MAP estimators are insignificant, the Gibbs solution seems to be slightly better.

In order to have a quantitative statement about the goodness of the reconstruction, we define a relative residual error $\epsilon^{(s)}$ for the diffuse contribution as follows,

$$\epsilon^{(s)} = \left| \rho^{(s)} - \langle \rho^{(s)} \rangle \right|_2 \left| \rho^{(s)} \right|_2^{-1}, \quad (59)$$

where $|\cdot|_2$ is the Euclidean L^2 -norm. For the point-like contribution, however, we have to consider an error in brightness and position. For this purpose we define,

$$\epsilon^{(u)} = \int_1^N dn \left| \mathbf{R}_{\text{PSF}}^n \rho^{(u)} - \mathbf{R}_{\text{PSF}}^n \langle \rho^{(u)} \rangle \right|_2 \left| \mathbf{R}_{\text{PSF}}^n \rho^{(u)} \right|_2^{-1}, \quad (60)$$

where $\mathbf{R}_{\text{PSF}}$ is a (normalized) convolution operator, such that $\mathbf{R}_{\text{PSF}}^N$ becomes the identity for large N . These errors are listed in Table 1. When comparing the MAP- δ and MAP- $\mathcal{G}$ approach, the incorporation of uncertainty corrections seems to improve the results slightly. The full regularization treatment within the Gibbs approach outperforms MAP solutions in terms of the chosen error measure $\epsilon^{(\cdot)}$. For a discussion of how such measures

Fig. 4. Illustration of the diffuse reconstruction. The *top panels* show the denoised and deconvolved diffuse contribution $\langle \rho^{(s)} \rangle / \rho_0$ reconstructed using a different approach each, namely **d)** MAP- δ , **e)** MAP- $\mathcal{G}$, and **f)** Gibbs. The *bottom panels d) to f)* show the difference between the originally simulated signal and the respective reconstruction.

Table 1. Overview of the relative residual errors in the photon flux reconstructions for the respective approaches, all using the same model parameters, cf. text.

MAP- δ	MAP- $\mathcal{G}$	Gibbs
$\epsilon^{(s)} = 4.442\%$	$\epsilon^{(s)} = 4.441\%$	$\epsilon^{(s)} = 2.078\%$
$\epsilon^{(u)} = 1.540\%$	$\epsilon^{(u)} = 1.540\%$	$\epsilon^{(u)} = 1.089\%$

can change the view on certain Bayesian estimators, we refer to the work by [Burger & Lucka \(2014\)](#).

Figure 5 illustrates the reconstruction of the diffuse signal field, now in terms of logarithmic flux. The original and the reconstructions agree well, and the strongest deviations are found in the areas with low amplitudes. With regard to the exponential ansatz in Eq. (4), it is not surprising that the inference on the signal fields is more sensitive to higher values than to lower ones. For example, a small change in the diffuse signal field, $s \rightarrow (1 \pm \epsilon)s$, translates into a factor in the photon flux, $\rho^{(s)} \rightarrow \rho^{(s)} e^{\pm \epsilon s}$, that scales exponentially with the amplitude of the diffuse signal field. The Gibbs solution shows less deviation from the original signal than the MAP solution. Since the latter lacks the regularization by the uncertainty covariance it exhibits a stronger tendency to overfitting compared to the former. This includes overestimates in noisy regions with low flux intensities, as well as underestimates at locations where point-like contributions dominate the total flux.

The reconstruction of the power spectrum, as shown in Fig. 6, gives further indications of the reconstruction quality of the diffuse component. The simulation used a default power spectrum of

$$\exp(\tau_k) = 42(k+1)^{-7}. \quad (61)$$

This power spectrum was on purpose chosen to deviate from a strict power law supposed by the smoothness prior.

From Fig. 6 it is apparent that the reconstructed power spectra track the original well up to a harmonic mode k of roughly 0.4 arcmin^{-1} . Beyond that point, the reconstructed power spectra fall steeply until they hit a lower boundary set by the model parameter $q = 10^{-12}$. This drop-off point at 0.4 arcmin^{-1} corresponds to a physical wavelength of roughly 2.5 arcmin , and thus (half-phase) fluctuations on a spatial distances below 1.25 arcmin . The Gaussian-like PSF of the virtual observatory has a finite support of 1.1 arcmin . The lack of reconstructed power indicates that the algorithm assigns features on spatial scales smaller than the PSF support preferably to the point-like component. This behavior is reasonable because solely the point-like signal can cause PSF-like shaped imprints in the data image. However, there is no strict threshold in the distinction between the components on the mere basis of their spatial extend. We rather observe a continuous transition from assigning flux to the diffuse component to assigning it to the point-like component while reaching smaller spatial scales because strict boundaries are blurred out under the consideration of noise effects.

Fig. 5. Illustration of the reconstruction of the diffuse signal field $s = \log \rho^{(s)}$ and its uncertainty. The top panels show diffuse signal fields. Panel **a**) shows the original simulation s , panel **b**) the reconstruction $\mathbf{m}_{\text{mode}}^{(s)}$ using a MAP approach, and panel **c**) the reconstruction $\mathbf{m}_{\text{mean}}^{(s)}$ using a Gibbs approach. The panels **d**) and **e**) show the differences between original and reconstruction. Panel **f**) shows the relative difference. The panels **g**) and **h**) show the relative uncertainty of the above reconstructions.

The differences between the reconstruction using a MAP and a Gibbs approach are subtle. The difference in the reconstruction formulas given by Eqs. (37) and (53) is an additive trace term involving $\mathbf{D}^{(s)}$, which is positive definite. Therefore, a reconstructed power spectrum regularized by uncertainty corrections is never below the one with out given the same $\mathbf{m}^{(s)}$. However, the reconstruction of the signal field follows different filter formulas, respectively. Since the Gibbs approach considers the uncertainty covariance $\mathbf{D}^{(s)}$ properly in each cycle, it can present a more conservative solution. The drop-off point is apparently at higher k for the MAP approach, leading to higher power on scales between roughly 0.3 and 0.7 arcmin⁻¹. In turn, the MAP

solution tends to overfit by absorbing some noise power into $\mathbf{m}^{(s)}$ as discussed in Sect. 3. Thus, the higher MAP power spectrum in Fig. 6 seems to be caused by a higher level of noise remnants in the signal estimate.

The influence of the choice of the model parameter σ is also shown in Fig. 6. Neither a smoothness prior with $\sigma = 10$, nor a weak one with $\sigma = 1000$ influences the reconstruction of the power spectrum substantially in this case¹². The latter choice,

¹² For a discussion of further log-normal reconstruction scenarios please refer to the work by Oppermann et al. (2013).

Fig. 6. Illustration of the reconstruction of the logarithmic power spectrum τ . Both panels show the default power spectrum (black dashed line), and the simulated realization (black dotted line), as well as the reconstructed power spectra using a MAP (orange solid line), plus second order corrections (orange dashed line), and a Gibbs approach (blue solid line). Panel **a**) shows the reconstruction for a chosen σ parameter of 10, panel **b**) for a σ of 1000.

however, exhibits some more fluctuations in order to better track the concrete realization.

The results for the reconstruction of the point-like component are illustrated in Fig. 7. Overall, the reconstructed point-like signal field and the corresponding photon flux are in good agreement with the original ones. The point-sources have been located with an accuracy of ± 0.1 arcmin, which is less than the FWHM of the PSF. The localization tends to be more precise for higher flux values because of the higher signal-to-noise ratio. The reconstructed intensities match the simulated ones well, although the MAP solution shows a spread that exceeds the expected shot noise uncertainty interval. This is again an indication of the overfitting known for MAP solutions. Moreover, neither reconstruction shows a bias towards higher or lower fluxes.

The uncertainty estimates for the point-like photon flux $\rho^{(u)}$ obtained from $\mathbf{D}^{(u)}$ according to Eqs. (57) and (58) are, in general, consistent with the deviations from the original and the shot noise uncertainty, cf. Fig. 7. They show a reasonable scaling being higher for lower fluxes and vice versa. However, some uncertainties seem to be underestimated. There are different reasons for this.

On the one hand, the Hessian approximation for $\mathbf{D}^{(u)}$ in Eqs. (36) or (51) is in individual cases in so far poor as that the curvature of the considered potential does not describe the uncertainty of the point-like component adequately. The data admittedly constrains the flux intensity of a point source sufficiently, especially if it is a bright one. However, the rather narrow dip in the manifold landscape of the considered potential can be asymmetric, and thus not always well described by the quadratic approximation of Eqs. (36) or (51), respectively.

On the other hand, the approximation leading to vanishing cross-correlation $\mathbf{D}^{(su)}$, takes away the possibility of communicating uncertainties between diffuse and point-like components. However, omitting the used simplification or incorporating higher order corrections would render the algorithm too computationally expensive. The fact that the Gibbs solution, which takes $\mathbf{D}^{(u)}$ into account, shows improvements backs up this argument.

The reconstructions shown in Figs. 5 and 7 used the model parameters $\sigma = 10$, $\beta = \frac{3}{2}$, and $\eta = 10^{-4}$. In order to reflect the influence of the choice of σ , β , and η , Table 2 summarizes the results from several reconstructions carried out with varying model parameters. Accordingly, the best parameters seem to be $\sigma = 10$, $\beta = \frac{5}{4}$, and $\eta = 10^{-4}$, although we caution that the total error is

difficile to determine as the residual errors, $\epsilon^{(s)}$ and $\epsilon^{(u)}$, are defined differently. Although the errors vary significantly, 2–15% for $\epsilon^{(s)}$, we like to stress that the model parameters were changed drastically, partly even by orders of magnitude. The impact of the prior clearly exists, but is moderate. We note that the case of $\sigma \rightarrow \infty$ corresponds to neglecting the smoothness prior completely. The $\beta = 1$ case that corresponds to a logarithmically flat prior on $\mathbf{u}$ showed a tendency to fit more noise features by point-like contributions.

In summary, the D³PO algorithm is capable of denoising, deconvolving and decomposing photon observations by reconstructing the diffuse and point-like signal field, and the logarithmic power spectrum of the former. The reconstruction using MAP and Gibbs approaches perform flawlessly, except for a little underestimation of the uncertainty of the point-like component. The MAP approach shows signs of overfitting, but those are not overwhelming. Considering the simplicity of the MAP approach that goes along with a numerically faster performance, this shortcoming seems acceptable.

Because of the iterative scheme of the algorithm, a combination of the MAP approach for the signal fields and a Gibbs approach for the power spectrum is possible.

5. Conclusions and summary

The D³PO algorithm for the denoising, deconvolving and decomposing photon observations has been derived. It allows for the simultaneous but individual reconstruction of the diffuse and point-like photon fluxes, as well as the harmonic power spectrum of the diffuse component, from a single data image that is exposed to Poissonian shot noise and effects of the instrument response functions. Moreover, the D³PO algorithm can provide a posteriori uncertainty information on the reconstructed signal fields. With these capabilities, D³PO surpasses previous approaches that address only subsets of these complications.

The theoretical foundation is a hierarchical Bayesian parameter model embedded in the framework of IFT. The model comprises a priori assumptions for the signal fields that account for the different statistics and correlations of the morphologically different components. The diffuse photon flux is assumed to obey multivariate log-normal statistics, where the covariance is described by a power spectrum. The power spectrum is a priori unknown and reconstructed from the data along with the signal. Therefore, hyperpriors on the (logarithmic) power spectra have

Fig. 7. Illustration of the reconstruction of the point-like signal field $u = \log \rho^{(u)}$ and its uncertainty. The *top panels* show the location (markers) and intensity (gray scale) of the point-like photon fluxes, underlaid is the respective diffuse contribution (contours) to guide the eye, cf. Fig. 4. Panel **a**) shows the original simulation, panel **b**) the reconstruction using a MAP approach, and panel **c**) the reconstruction using a Gibbs approach. The *bottom panels* **d**) and **e**) show the match between original and reconstruction in absolute and relative fluxes, the 2σ shot noise interval (gray contour), as well as some reconstruction uncertainty estimate (error bars).

been introduced, including a spectral smoothness prior (Enßlin & Frommert 2011; Oppermann et al. 2013). The point-like photon flux, in contrast, is assumed to factorize spatially in independent inverse-Gamma distributions implying a (regularized) power-law behavior of the amplitudes of the flux.

An adequate description of the noise properties in terms of a likelihood, here a Poisson distribution, and the incorporation of all instrumental effects into the response operator renders the denoising and deconvolution task possible. The strength of the proposed approach is the performance of the additional decomposition task, which especially exploits the a priori description of diffuse and point-like. The model comes down to five scalar parameters, for which all a priori defaults can be motivated, and of which none is driving the inference predominantly.

We discussed maximum a posteriori (MAP) and Gibbs free energy approaches to solve the inference problem. The derived solutions provide optimal estimators that, in the considered examples, yielded equivalently excellent results. The Gibbs solution slightly outperforms MAP solutions (in terms of the considered L^2 -residuals) thanks to the full regularization treatment, however, for the price of a computationally more expensive optimization. Which approach is to be preferred in general might

depend on the concrete problem at hand and the trade-off between reconstruction precision against computational effort.

The performance of the D³PO algorithm has been demonstrated in realistic simulations carried out in 1D and 2D. The implementation relies on the NIFTY package (Selig et al. 2013), which allows for the application regardless of the underlying position space.

In the 2D application example, a high energy observation of a 32×32 arcmin² patch of a simulated sky with a 0.1 arcmin resolution has been analyzed. The D³PO algorithm successfully denoised, deconvolved and decomposed the data image. The analysis yielded a detailed reconstruction of the diffuse photon flux and its logarithmic power spectrum, the precise localization of the point sources and accurate determination of their flux intensities, as well as a posteriori estimates of the reconstructed fields.

The D³PO algorithm should be applicable to a wide range of inference problems appearing in astronomical imaging and related fields. Concrete applications in high energy astrophysics, for example, the analysis of data from the *Chandra* X-ray observatory or the Fermi γ -ray space telescope, are currently considered by the authors. In this regard, the public release of the D³PO code is planned.

Table 2. Overview of the relative residual error in the photon flux reconstructions for a MAP- δ approach with varying model parameters σ , β , and η .

$\alpha = 1$	$q = 10^{-12}$	$\sigma = 1$	$\sigma = 10$	$\sigma = 100$	$\sigma = 1000$	$\sigma \rightarrow \infty$
$\beta = 1$	$\eta = 10^{-6}$	$\epsilon^{(s)} = 0.06710$ $\epsilon^{(u)} = 0.02000$	$\epsilon^{(s)} = 0.05406$ $\epsilon^{(u)} = 0.01941$	$\epsilon^{(s)} = 0.05323$ $\epsilon^{(u)} = 0.01602$	$\epsilon^{(s)} = 0.05383$ $\epsilon^{(u)} = 0.01946$	$\epsilon^{(s)} = 0.05359$ $\epsilon^{(u)} = 0.01898$
$\beta = \frac{5}{4}$	$\eta = 10^{-6}$	$\epsilon^{(s)} = 0.02874$ $\epsilon^{(u)} = 0.01207$	$\epsilon^{(s)} = 0.01929$ $\epsilon^{(u)} = 0.01102$	$\epsilon^{(s)} = 0.01974$ $\epsilon^{(u)} = 0.01090$	$\epsilon^{(s)} = 0.02096$ $\epsilon^{(u)} = 0.01123$	$\epsilon^{(s)} = 0.01991$ $\epsilon^{(u)} = 0.01104$
$\beta = \frac{3}{2}$	$\eta = 10^{-6}$	$\epsilon^{(s)} = 0.05890$ $\epsilon^{(u)} = 0.02741$	$\epsilon^{(s)} = 0.02237$ $\epsilon^{(u)} = 0.01343$	$\epsilon^{(s)} = 0.02318$ $\epsilon^{(u)} = 0.01346$	$\epsilon^{(s)} = 0.02238$ $\epsilon^{(u)} = 0.01342$	$\epsilon^{(s)} = 0.02344$ $\epsilon^{(u)} = 0.01351$
$\beta = \frac{7}{4}$	$\eta = 10^{-6}$	$\epsilon^{(s)} = 0.10864$ $\epsilon^{(u)} = 0.04840$	$\epsilon^{(s)} = 0.04304$ $\epsilon^{(u)} = 0.02767$	$\epsilon^{(s)} = 0.03234$ $\epsilon^{(u)} = 0.02142$	$\epsilon^{(s)} = 0.03248$ $\epsilon^{(u)} = 0.02143$	$\epsilon^{(s)} = 0.03263$ $\epsilon^{(u)} = 0.02167$
$\beta = 2$	$\eta = 10^{-6}$	$\epsilon^{(s)} = 0.11870$ $\epsilon^{(u)} = 0.05360$	$\epsilon^{(s)} = 0.04614$ $\epsilon^{(u)} = 0.02926$	$\epsilon^{(s)} = 0.04527$ $\epsilon^{(u)} = 0.02924$	$\epsilon^{(s)} = 0.04522$ $\epsilon^{(u)} = 0.02926$	$\epsilon^{(s)} = 0.04500$ $\epsilon^{(u)} = 0.02915$
$\beta = 1$	$\eta = 10^{-4}$	$\epsilon^{(s)} = 0.06660$ $\epsilon^{(u)} = 0.02157$	$\epsilon^{(s)} = 0.05474$ $\epsilon^{(u)} = 0.01903$	$\epsilon^{(s)} = 0.05377$ $\epsilon^{(u)} = 0.01657$	$\epsilon^{(s)} = 0.05474$ $\epsilon^{(u)} = 0.01986$	$\epsilon^{(s)} = 0.05423$ $\epsilon^{(u)} = 0.02055$
$\beta = \frac{5}{4}$	$\eta = 10^{-4}$	$\epsilon^{(s)} = 0.02874$ $\epsilon^{(u)} = 0.01207$	$\epsilon^{(s)} = 0.01929$ $\epsilon^{(u)} = 0.01100$	$\epsilon^{(s)} = 0.01974$ $\epsilon^{(u)} = 0.01103$	$\epsilon^{(s)} = 0.02096$ $\epsilon^{(u)} = 0.01123$	$\epsilon^{(s)} = 0.01991$ $\epsilon^{(u)} = 0.01102$
$\beta = \frac{3}{2}$	$\eta = 10^{-4}$	$\epsilon^{(s)} = 0.05890$ $\epsilon^{(u)} = 0.02743$	$\epsilon^{(s)} = 0.02237$ $\epsilon^{(u)} = 0.01343$	$\epsilon^{(s)} = 0.02318$ $\epsilon^{(u)} = 0.01346$	$\epsilon^{(s)} = 0.02238$ $\epsilon^{(u)} = 0.01340$	$\epsilon^{(s)} = 0.02344$ $\epsilon^{(u)} = 0.01352$
$\beta = \frac{7}{4}$	$\eta = 10^{-4}$	$\epsilon^{(s)} = 0.10864$ $\epsilon^{(u)} = 0.04840$	$\epsilon^{(s)} = 0.04304$ $\epsilon^{(u)} = 0.02766$	$\epsilon^{(s)} = 0.03234$ $\epsilon^{(u)} = 0.02145$	$\epsilon^{(s)} = 0.03248$ $\epsilon^{(u)} = 0.02142$	$\epsilon^{(s)} = 0.03263$ $\epsilon^{(u)} = 0.02166$
$\beta = 2$	$\eta = 10^{-4}$	$\epsilon^{(s)} = 0.11870$ $\epsilon^{(u)} = 0.05358$	$\epsilon^{(s)} = 0.04614$ $\epsilon^{(u)} = 0.02926$	$\epsilon^{(s)} = 0.04527$ $\epsilon^{(u)} = 0.02926$	$\epsilon^{(s)} = 0.04522$ $\epsilon^{(u)} = 0.02927$	$\epsilon^{(s)} = 0.04500$ $\epsilon^{(u)} = 0.02916$
$\beta = 1$	$\eta = 10^{-2}$	$\epsilon^{(s)} = 0.07271$ $\epsilon^{(u)} = 0.02252$	$\epsilon^{(s)} = 0.06209$ $\epsilon^{(u)} = 0.02047$	$\epsilon^{(s)} = 0.06192$ $\epsilon^{(u)} = 0.02109$	$\epsilon^{(s)} = 0.06291$ $\epsilon^{(u)} = 0.01764$	$\epsilon^{(s)} = 0.06265$ $\epsilon^{(u)} = 0.02068$
$\beta = \frac{5}{4}$	$\eta = 10^{-2}$	$\epsilon^{(s)} = 0.02335$ $\epsilon^{(u)} = 0.01139$	$\epsilon^{(s)} = 0.01934$ $\epsilon^{(u)} = 0.01112$	$\epsilon^{(s)} = 0.02042$ $\epsilon^{(u)} = 0.01097$	$\epsilon^{(s)} = 0.01999$ $\epsilon^{(u)} = 0.01124$	$\epsilon^{(s)} = 0.01930$ $\epsilon^{(u)} = 0.01102$
$\beta = \frac{3}{2}$	$\eta = 10^{-2}$	$\epsilon^{(s)} = 0.05999$ $\epsilon^{(u)} = 0.02745$	$\epsilon^{(s)} = 0.02227$ $\epsilon^{(u)} = 0.01341$	$\epsilon^{(s)} = 0.02347$ $\epsilon^{(u)} = 0.01356$	$\epsilon^{(s)} = 0.02266$ $\epsilon^{(u)} = 0.01332$	$\epsilon^{(s)} = 0.02274$ $\epsilon^{(u)} = 0.01351$
$\beta = \frac{7}{4}$	$\eta = 10^{-2}$	$\epsilon^{(s)} = 0.10715$ $\epsilon^{(u)} = 0.04833$	$\epsilon^{(s)} = 0.04304$ $\epsilon^{(u)} = 0.02766$	$\epsilon^{(s)} = 0.03254$ $\epsilon^{(u)} = 0.02140$	$\epsilon^{(s)} = 0.03264$ $\epsilon^{(u)} = 0.02144$	$\epsilon^{(s)} = 0.03258$ $\epsilon^{(u)} = 0.02163$
$\beta = 2$	$\eta = 10^{-2}$	$\epsilon^{(s)} = 0.12496$ $\epsilon^{(u)} = 0.05361$	$\epsilon^{(s)} = 0.04614$ $\epsilon^{(u)} = 0.02927$	$\epsilon^{(s)} = 0.04497$ $\epsilon^{(u)} = 0.02915$	$\epsilon^{(s)} = 0.04528$ $\epsilon^{(u)} = 0.02914$	$\epsilon^{(s)} = 0.04500$ $\epsilon^{(u)} = 0.02915$
$\beta = 1$	$\eta = 1$	$\epsilon^{(s)} = 0.15328$ $\epsilon^{(u)} = 0.03250$	$\epsilon^{(s)} = 0.14544$ $\epsilon^{(u)} = 0.03291$	$\epsilon^{(s)} = 0.14138$ $\epsilon^{(u)} = 0.02905$	$\epsilon^{(s)} = 0.14181$ $\epsilon^{(u)} = 0.03087$	$\epsilon^{(s)} = 0.14185$ $\epsilon^{(u)} = 0.02876$
$\beta = \frac{5}{4}$	$\eta = 1$	$\epsilon^{(s)} = 0.15473$ $\epsilon^{(u)} = 0.03217$	$\epsilon^{(s)} = 0.14406$ $\epsilon^{(u)} = 0.03166$	$\epsilon^{(s)} = 0.14357$ $\epsilon^{(u)} = 0.03089$	$\epsilon^{(s)} = 0.14465$ $\epsilon^{(u)} = 0.03101$	$\epsilon^{(s)} = 0.13964$ $\epsilon^{(u)} = 0.03160$
$\beta = \frac{3}{2}$	$\eta = 1$	$\epsilon^{(s)} = 0.15360$ $\epsilon^{(u)} = 0.03262$	$\epsilon^{(s)} = 0.14216$ $\epsilon^{(u)} = 0.03063$	$\epsilon^{(s)} = 0.14248$ $\epsilon^{(u)} = 0.02534$	$\epsilon^{(s)} = 0.14208$ $\epsilon^{(u)} = 0.02872$	$\epsilon^{(s)} = 0.14233$ $\epsilon^{(u)} = 0.03095$
$\beta = \frac{7}{4}$	$\eta = 1$	$\epsilon^{(s)} = 0.15206$ $\epsilon^{(u)} = 0.03262$	$\epsilon^{(s)} = 0.14156$ $\epsilon^{(u)} = 0.03065$	$\epsilon^{(s)} = 0.13772$ $\epsilon^{(u)} = 0.03174$	$\epsilon^{(s)} = 0.14160$ $\epsilon^{(u)} = 0.03141$	$\epsilon^{(s)} = 0.14390$ $\epsilon^{(u)} = 0.03178$
$\beta = 2$	$\eta = 1$	$\epsilon^{(s)} = 0.06421$ $\epsilon^{(u)} = 0.02043$	$\epsilon^{(s)} = 0.05479$ $\epsilon^{(u)} = 0.01966$	$\epsilon^{(s)} = 0.05365$ $\epsilon^{(u)} = 0.01676$	$\epsilon^{(s)} = 0.05499$ $\epsilon^{(u)} = 0.02070$	$\epsilon^{(s)} = 0.05429$ $\epsilon^{(u)} = 0.01996$

Notes. The parameters α and q were fixed. The best and worst residuals are printed in bold face.

Acknowledgements. We thank Niels Oppermann, Henrik Junklewitz and two anonymous referees for the insightful discussions and productive comments. Furthermore, we thank the DFG Forschergruppe 1254 “Magnetisation of Interstellar and Intergalactic Media: The Prospects of Low-Frequency Radio Observations” for travel support in order to present this work at their annual meeting in 2013. Some of the results in this publication have been derived using the NIFTY package (Selig et al. 2013). This research has made use of NASA’s Astrophysics Data System.

Appendix A: Point source stacking

In Sect. 2.3.3, a prior for the point-like signal field has been derived under the assumption that the photon flux of point sources is independent between different pixels and identically inverse-Gamma distributed,

$$\rho_x^{(u)} \sim \mathcal{I}\left(\rho_x^{(u)}, \beta = \frac{3}{2}, \rho_0 \eta\right) \quad \forall x, \quad (\text{A.1})$$

with the shape and scale parameters, β and η . It can be shown that, for $\beta = \frac{3}{2}$, the sum of N such variables still obeys an inverse-Gamma distribution,

$$\rho_N^{(u)} = \sum_x^N \rho_x^{(u)} \quad (\text{A.2})$$

$$\rho_N^{(u)} \sim \mathcal{I}\left(\rho_N^{(u)}, \beta = \frac{3}{2}, N^2 \rho_0 \eta\right). \quad (\text{A.3})$$

For a proof see Giron (2001).

In the case of $\beta = \frac{3}{2}$, the power-law behavior of the prior becomes independent of the discretization of the continuous position space. This means that the slope of the distribution of $\rho_x^{(u)}$ remains unchanged notwithstanding that we refine or coarsen the resolution of the reconstruction. However, the scale parameter η needs to be adapted for each resolution; i.e., $\eta \rightarrow N^2 \eta$ if N pixels are merged.

Appendix B: Covariance and curvature.

The covariance $\mathbf{D}$ of a Gaussian $\mathcal{G}(s - \mathbf{m}, \mathbf{D})$ describes the uncertainty associated with the mean $\mathbf{m}$ of the distribution. It can be computed by second moments or cumulants according to Eq. (3), or in this Gaussian case as the inverse Hessian of the corresponding information Hamiltonian,

$$\left. \frac{\partial^2 H}{\partial s \partial s^\dagger} \right|_{s=\mathbf{m}} = \frac{\partial^2}{\partial s \partial s^\dagger} \left(\frac{1}{2} (s - \mathbf{m})^\dagger \mathbf{D}^{-1} (s - \mathbf{m}) \right) \Big|_{s=\mathbf{m}} = \mathbf{D}^{-1}. \quad (\text{B.1})$$

In Sect. 3, uncertainty covariances for the diffuse signal field s and the point-like signal field u have been derived that are here given in closed form.

The MAP uncertainty covariances introduced in Sect. 3.1 are approximated by inverse Hessians. According to Eq. (36), they read

$$D_{xy}^{(s)-1} \approx \left\{ \sum_i \left(1 - \frac{d_i}{l_i} \right) R_{ix} e^{m_x^{(s)}} \right\} \delta_{xy} + \sum_i \frac{d_i}{l_i^2} \left(R_{ix} e^{m_x^{(s)}} \right) \left(R_{iy} e^{m_y^{(s)}} \right) + S_{xy}^{\star -1}, \quad (\text{B.2})$$

and

$$D_{xy}^{(u)-1} \approx \left\{ \sum_i \left(1 - \frac{d_i}{l_i} \right) R_{ix} e^{m_x^{(u)}} + \eta e^{-m_x^{(u)}} \right\} \delta_{xy} + \sum_i \frac{d_i}{l_i^2} \left(R_{ix} e^{m_x^{(u)}} \right) \left(R_{iy} e^{m_y^{(u)}} \right), \quad (\text{B.3})$$

with

$$l_i = \int dx R_{ix} (e^{m_x^{(s)}} + e^{m_x^{(u)}}). \quad (\text{B.4})$$

The corresponding covariances derived in the Gibbs approach according to Eq. (51), yield

$$D_{xy}^{(s)-1} \approx \left\{ \sum_i \left(1 - \frac{d_i}{l_i} \right) R_{ix} e^{m_x^{(s)} + \frac{1}{2} D_{xx}^{(s)}} \right\} \delta_{xy} + \sum_i \frac{d_i}{l_i^2} \left(R_{ix} e^{m_x^{(s)} + \frac{1}{2} D_{xx}^{(s)}} \right) \left(R_{iy} e^{m_y^{(s)} + \frac{1}{2} D_{yy}^{(s)}} \right) + S_{xy}^{\star -1}, \quad (\text{B.5})$$

and

$$D_{xy}^{(u)-1} \approx \left\{ \sum_i \left(1 - \frac{d_i}{l_i} \right) R_{ix} e^{m_x^{(u)} + \frac{1}{2} D_{xx}^{(u)}} + \eta e^{-m_x^{(u)} + \frac{1}{2} D_{xx}^{(u)}} \right\} \delta_{xy} + \sum_i \frac{d_i}{l_i^2} \left(R_{ix} e^{m_x^{(u)} + \frac{1}{2} D_{xx}^{(u)}} \right) \left(R_{iy} e^{m_y^{(u)} + \frac{1}{2} D_{yy}^{(u)}} \right), \quad (\text{B.6})$$

with

$$l_i = \int dx R_{ix} \left(e^{m_x^{(s)} + \frac{1}{2} D_{xx}^{(s)}} + e^{m_x^{(u)} + \frac{1}{2} D_{xx}^{(u)}} \right). \quad (\text{B.7})$$

They are identical up to the $+\frac{1}{2}D_{xx}$ terms in the exponents. On the one hand, this reinforces the approximations done in Sect. 3.2. On the other hand, this shows that higher order correction terms might alter the uncertainty covariances further, cf. Eq. (43). The concrete impact of these correction terms is difficult to judge, since they introduce terms involving D_{xy} that couple all elements of $\mathbf{D}$ in an implicit manner.

We note that the inverse Hessian describes the curvature of the potential, its interpretation as uncertainty is, strictly speaking, only valid for quadratic potentials. However, in most cases it is a sufficient approximation.

The Gibbs approach provides an alternative by equating the first derivative of the Gibbs free energy with respect to the covariance with zero. Following Eq. (52), the covariances read

$$D_{xy}^{(s)-1} = \left\{ \sum_i \left(1 - \frac{d_i}{l_i} \right) R_{ix} e^{m_x^{(s)} + \frac{1}{2} D_{xx}^{(s)}} \right\} \delta_{xy} \quad (\text{B.8})$$

$$+ S_{xy}^{\star -1}, \quad (\text{B.9})$$

and

$$D_{xy}^{(u)-1} = \left\{ \sum_i \left(1 - \frac{d_i}{l_i} \right) R_{ix} e^{m_x^{(u)} + \frac{1}{2} D_{xx}^{(u)}} \right. \quad (\text{B.10})$$

$$\left. + \eta e^{-m_x^{(u)} + \frac{1}{2} D_{xx}^{(u)}} \right\} \delta_{xy}. \quad (\text{B.11})$$

Compared to the above solutions, there is one term missing indicating that they already lack first order corrections. For this reasons, the solutions obtained from the inverse Hessians are used in the D³PO algorithm.

Appendix C: Posterior approximation

C.1. Information theoretical measure

If the full posterior $P(z|\mathbf{d})$ of an inference problem is so complex that an analytic handling is infeasible, an approximate posterior Q might be used instead. The fitness of such an approximation can be quantified by an asymmetric measure for which different terminologies appear in the literature.

First, the Kullback-Leibler divergence,

$$D_{\text{KL}}(Q, P) = \int \mathcal{D}z Q(z|\mathbf{d}) \log \frac{Q(z|\mathbf{d})}{P(z|\mathbf{d})} \quad (\text{C.1})$$

$$= \left\langle \log \frac{Q(z|\mathbf{d})}{P(z|\mathbf{d})} \right\rangle_Q, \quad (\text{C.2})$$

defines mathematically an information theoretical distance, or divergence, which is minimal if a maximal cross information between P and Q exists (Kullback & Leibler 1951).

Second, the information entropy,

$$S_E(Q, P) = - \int \mathcal{D}z P(z|\mathbf{d}) \log \frac{P(z|\mathbf{d})}{Q(z|\mathbf{d})} \quad (\text{C.3})$$

$$= \left\langle - \log \frac{P(z|\mathbf{d})}{Q(z|\mathbf{d})} \right\rangle_P \quad (\text{C.4})$$

$$= -D_{\text{KL}}(P, Q),$$

is derived under the maximum entropy principle (Jaynes 1957a,b) from fundamental axioms demanding locality, coordinate invariance and system independence (see e.g. Caticha 2008, 2011).

Third, the (approximate) Gibbs free energy (Enßlin & Weig 2010),

$$G = \langle H(\mathbf{z}|\mathbf{d}) \rangle_Q - S_B(Q) \quad (\text{C.5})$$

$$= \langle -\log P(\mathbf{z}|\mathbf{d}) \rangle_Q - \langle -\log Q(\mathbf{z}|\mathbf{d}) \rangle_Q \quad (\text{C.6})$$

$$= D_{\text{KL}}(Q, P),$$

describes the difference between the internal energy $\langle H(\mathbf{z}|\mathbf{d}) \rangle_Q$ and the Boltzmann-Shannon entropy $S_B(Q) = S_E(1, Q)$. The derivation of the Gibbs free energy is based on the principles of thermodynamics¹³.

The Kullback-Leibler divergence, information entropy, and the Gibbs free energy are equivalent measures that allow one to assess the approximation $Q \approx P$. Alternatively, a parametrized proposal for Q can be pinned down by extremizing the measure of choice with respect to the parameters.

C.2. Calculus of variations

The information theoretical measure can be interpreted as an action to which the principle of least action applies. This concept is the basis for variational Bayesian methods (Jordan et al. 1999; Wingate & Weber 2013), which enable among others the derivation of approximate posterior distributions.

We suppose that $\mathbf{z}$ is a set of multiple signal fields, $\mathbf{z} = \{\mathbf{z}^{(i)}\}_{i \in \mathbb{N}}$, $\mathbf{d}$ a given data set, and $P(\mathbf{z}|\mathbf{d})$ the posterior of interest. In practice, such a problem is often addressed by a mean field approximation that factorizes the variational posterior Q ,

$$P(\mathbf{z}|\mathbf{d}) \approx Q = \prod_i Q_i(\mathbf{z}^{(i)}|\boldsymbol{\mu}, \mathbf{d}). \quad (\text{C.7})$$

Here, the mean field $\boldsymbol{\mu}$, which mimics the effect of all $\mathbf{z}^{(i \neq j)}$ onto $\mathbf{z}^{(j)}$, has been introduced. The approximation in Eq. (C.7) shifts any possible entanglement between the $\mathbf{z}^{(i)}$ within P into the dependence of $\mathbf{z}^{(i)}$ on $\boldsymbol{\mu}$ within Q_i . Hence, the mean field $\boldsymbol{\mu}$ is well determined by the inference problem at hand, as demonstrated in the subsequent Sect. C.3. We note that $\boldsymbol{\mu}$ represents effective rather than additional degrees of freedom.

Following the principle of least action, any variation of the Gibbs free energy must vanish. We consider a variation $\delta_j = \delta/\delta Q_j(\mathbf{z}^{(j)}|\boldsymbol{\mu}, \mathbf{d})$ with respect to one approximate posterior $Q_j(\mathbf{z}^{(j)}|\boldsymbol{\mu}, \mathbf{d})$. It holds,

$$\frac{\delta Q_i(\mathbf{z}^{(i)}|\boldsymbol{\mu}, \mathbf{d})}{\delta Q_j(\mathbf{z}^{(j)}|\boldsymbol{\mu}, \mathbf{d})} = \delta_{ij} \delta(\mathbf{z}^{(i)} - \mathbf{z}^{(j)}). \quad (\text{C.8})$$

Computing the variation of the Gibbs free energy yields

$$\delta_j G = 0 = \frac{\delta}{\delta Q_j(\mathbf{z}^{(j)}|\boldsymbol{\mu}, \mathbf{d})} \left\{ \langle H(\mathbf{z}|\mathbf{d}) \rangle_Q - \langle -\log Q \rangle_Q \right\} \quad (\text{C.9})$$

$$= \frac{\delta}{\delta Q_j(\mathbf{z}^{(j)}|\boldsymbol{\mu}, \mathbf{d})} \left\{ \langle H(\mathbf{z}|\mathbf{d}) \rangle_Q + \sum_i \langle \log Q_i(\mathbf{z}^{(i)}|\boldsymbol{\mu}, \mathbf{d}) \rangle_{Q_i} \right\} \quad (\text{C.10})$$

$$\begin{aligned} &= \frac{\delta}{\delta Q_j(\mathbf{z}^{(j)}|\boldsymbol{\mu}, \mathbf{d})} \int \mathcal{D}\mathbf{z}^{(j)} Q_j(\mathbf{z}^{(j)}|\boldsymbol{\mu}, \mathbf{d}) \\ &\quad \times \left\{ \langle H(\mathbf{z}|\mathbf{d}) \rangle_{\prod_{i \neq j} Q_i} + \log Q_j(\mathbf{z}^{(j)}|\boldsymbol{\mu}, \mathbf{d}) \right\} \\ &\quad + \underbrace{\frac{\delta}{\delta Q_j(\mathbf{z}^{(j)}|\boldsymbol{\mu}, \mathbf{d})} \sum_{i \neq j} \dots}_{=0} \\ &= \int \mathcal{D}\mathbf{z}^{(j)} \delta(\mathbf{z}^{(i)} - \mathbf{z}^{(j)}) \left\{ \langle H(\mathbf{z}|\mathbf{d}) \rangle_{\prod_{i \neq j} Q_i} \right. \\ &\quad \left. + \log Q_j(\mathbf{z}^{(j)}|\boldsymbol{\mu}, \mathbf{d}) + 1 \right\} \\ &= \left\langle H(\mathbf{z}|\mathbf{d}) \right|_{\mathbf{z}^{(j)}} \Big|_{\prod_{i \neq j} Q_i} + \log Q_j(\mathbf{z}^{(j)}|\boldsymbol{\mu}, \mathbf{d}) + \text{const.} \quad (\text{C.11}) \end{aligned}$$

This defines a solution for the approximate posterior Q_j , where the constant term in Eq. (C.11) ensures the correct normalization¹⁴ of Q_j ,

$$Q_j(\mathbf{z}^{(j)}|\boldsymbol{\mu}, \mathbf{d}) \propto \exp \left(- \left\langle H(\mathbf{z}|\mathbf{d}) \right|_{\mathbf{z}^{(j)}} \right)_{\prod_{i \neq j} Q_i}. \quad (\text{C.12})$$

Although the parts $\mathbf{z}^{(i \neq j)}$ are integrated out, Eq. (C.12) is no marginalization since the integration is performed on the level of the (negative) logarithm of a probability distribution. The success of the mean field approach might be that this integration is often more well-behaved in comparison to the corresponding marginalization. However, the resulting equations for the Q_i depend on each other, and thus need to be solved self-consistently.

A maximum a posteriori solution for $\mathbf{z}^{(j)}$ can then be found by minimizing an effective Hamiltonian,

$$\underset{\mathbf{z}^{(j)}}{\text{argmax}} P(\mathbf{z}|\mathbf{d}) = \underset{\mathbf{z}^{(j)}}{\text{argmin}} H(\mathbf{z}|\mathbf{d}) \quad (\text{C.13})$$

$$\approx \underset{\mathbf{z}^{(j)}}{\text{argmin}} \left\langle H(\mathbf{z}|\mathbf{d}) \right|_{\mathbf{z}^{(j)}} \Big|_{\prod_{i \neq j} Q_i}. \quad (\text{C.14})$$

Since the posterior is approximated by a product, the Hamiltonian is approximated by a sum, and each summand depends on solely one variable in the partition of the latent variable $\mathbf{z}$.

C.3. Example

In this section, the variational method is demonstrated with an exemplary posterior of the following form,

$$P(\mathbf{s}, \boldsymbol{\tau}|\mathbf{d}) = \frac{P(\mathbf{d}|\mathbf{s})}{P(\mathbf{d})} P(\mathbf{s}|\boldsymbol{\tau}) P(\boldsymbol{\tau}) \quad (\text{C.15})$$

$$= \frac{P(\mathbf{d}|\mathbf{s})}{P(\mathbf{d})} \mathcal{G}(\mathbf{s}, \mathbf{S}) P_{\text{un}}(\boldsymbol{\tau}|\boldsymbol{\alpha}, \mathbf{q}) P_{\text{sm}}(\boldsymbol{\tau}|\boldsymbol{\sigma}), \quad (\text{C.16})$$

¹³ In Eq. (C.5), a unit temperature is implied, see discussion by Enßlin & Weig (2010, 2012); Iatsenko et al. (2012).

¹⁴ The normalization could be included by usage of Lagrange multipliers; i.e., by adding a term $\sum_i \lambda_i (1 - \int \mathcal{D}\mathbf{z}^{(i)} Q_i(\mathbf{z}^{(i)}|\boldsymbol{\mu}, \mathbf{d}))$ to the Gibbs free energy in Eq. (C.9).

Fig. C.1. Graphical model for the variational method applied to the example posterior in Eq. (C.15). Panel **a**) shows the graphical model without, and panel **b**) with the mean field μ .

where $P(d|s)$ stands for an arbitrary likelihood describing how likely the data d can be measured from a signal s , and $S = \sum_k e^{\tau_k} S_k$ for a parametrization of the signal covariance. This posterior is equivalent to the one derived in Sect. 2 in order to find a solution for the logarithmic power spectrum τ . Here, any explicit dependence on the point-like signal field u is veiled in favor of clarity.

The corresponding Hamiltonian reads

$$H(s, \tau|d) = -\log P(s, \tau|d) \quad (C.17)$$

$$= H_0 + \frac{1}{2} \sum_k \left(\varrho_k \tau_k + \text{tr} \left[s s^\dagger S_k^{-1} \right] e^{-\tau_k} \right) \quad (C.18)$$

$$+ (\alpha - 1)^\dagger \tau + q^\dagger e^{-\tau} + \frac{1}{2} \tau^\dagger T \tau,$$

where $\varrho_k = \text{tr} [S_k S_k^{-1}]$ and all terms constant in τ , including the likelihood $P(d|s)$, have been absorbed into H_0 .

For an arbitrary likelihood it might not be possible to marginalize the posterior over s analytically. However, an integration of the Hamiltonian over s might be feasible since the only relevant term is quadratic in s . As, on the one hand, the prior $P(s|\tau)$ is Gaussian and, on the other hand, a posterior mean m and covariance D for the signal field s suffice, cf. Eqs. (2) and (3), we assume a Gaussian approximation for Q_s ; i.e., $Q_s = \mathcal{G}(s - m, D)$.

We now introduce a mean field approximation, denoted by μ , by changing the causal structure as depicted in Fig. C.1. With the consequential approximation of the posterior,

$$P(s, \tau|d) \approx \mathcal{G}(s - m, D) Q_\tau(\tau|\mu, d), \quad (C.19)$$

we can calculate the effective Hamiltonian for τ as

$$\left\langle H(s, \tau|d) \right\rangle_{Q_s} = H_0 + \gamma^\dagger \tau + \frac{1}{2} \tau^\dagger T \tau + q^\dagger e^{-\tau} \quad (C.20)$$

$$+ \frac{1}{2} \sum_k \text{tr} \left[\left\langle s s^\dagger \right\rangle_{Q_s} S_k^{-1} \right] e^{-\tau_k}$$

$$= H_0 + \gamma^\dagger \tau + \frac{1}{2} \tau^\dagger T \tau + q^\dagger e^{-\tau} \quad (C.21)$$

$$+ \frac{1}{2} \sum_k \text{tr} \left[(m m^\dagger + D) S_k^{-1} \right] e^{-\tau_k},$$

where $\gamma = (\alpha - 1) + \frac{1}{2} \varrho$.

The nature of the mean field μ can be derived from the coupling term in Eq. (C.18) that ensures an information flow between s and τ ,

$$\mu = \left(\left\langle \text{tr} \left[s s^\dagger S_k^{-1} \right] \right\rangle_{Q_s} \right)_{Q_\tau} = \left(\text{tr} \left[(m m^\dagger + D) S_k^{-1} \right] \right)_{Q_\tau} \quad (C.22)$$

Hence, the mean field effect on τ_k is given by the above trace, and the mean field effect on s is described by $\langle S^{-1} \rangle_{Q_\tau}$.

Extremizing Eq. (C.21) yields

$$e^\tau = \frac{q + \frac{1}{2} \left(\text{tr} \left[(m m^\dagger + D) S_k^{-1} \right] \right)_k}{\gamma + T \tau}. \quad (C.23)$$

This formula is in agreement with the critical filter formula (Enßlin & Frommert 2011; Oppermann et al. 2013). In case a Gaussian likelihood and no smoothness prior is assumed, it is the exact maximum of the true posterior with respect to the (logarithmic) power spectrum.

References

- Bertin, E., & Arnouts, S. 1996, A&AS, 117, 393
- Bobin, J., Starck, J.-L., Fadili, J. M., Moudden, Y., & Donoho, D. 2007, IEEE Trans. Image Proc., 16, 2675
- Bouchet, L., Amestoy, P., Buttari, A., Rouet, F.-H., & Chauvin, M. 2013, Astronomy and Computing, 1, 59
- Burger, M., & Lucka, F. 2014, Inverse Problems, 30, 114004
- Carvalho, P., Rocha, G., & Hobson, M. P. 2009, MNRAS, 393, 681
- Carvalho, P., Rocha, G., Hobson, M. P., & Lasenby, A. 2012, MNRAS, 427, 1384
- Caticha, A. 2008 [arXiv:0808.0012]
- Caticha, A. 2011, in Am. Inst. Phys. Conf. Ser., eds. A. Mohammad-Djafari, J.-F. Bercher, & P. Bessière, AIP Conf. Ser., 1305, 20
- Chapman, E., Abdalla, F. B., Bobin, J., et al. 2013, MNRAS, 429, 165
- Cornwell, T. J. 2008, IEEE J. Select. Top. Sign. Process., 2, 793
- Dupé, F.-X., Fadili, J. M., & Starck, J.-L. 2009, IEEE Trans. Im. Process., 18, 310
- Dupé, F.-X., Fadili, J., & Starck, J.-L. 2011 [arXiv:1103.2213]
- Enßlin, T. 2013, in AIP Conf. Ser., Am. Inst. Phys. Conf. Ser., 1553, ed. U. von Toussaint, 184
- Enßlin, T. 2014, AIP Conf. Proc., 1636, 49
- Enßlin, T. A., & Frommert, M. 2011, Phys. Rev. D, 83, 105014
- Enßlin, T. A., & Weig, C. 2010, Phys. Rev. E, 82, 1112
- Enßlin, T. A., & Weig, C. 2012, Phys. Rev. E, 85, 3102
- Enßlin, T. A., Frommert, M., & Kitaura, F. S. 2009, Phys. Rev. D, 80, 5005
- Figueiredo, M. A. T., & Bioucas-Dias, J. M. 2010, IEEE Trans. Image Process., 19, 3133
- Fomalont, E. B. 1968, Bull. Astron. Inst. Netherlands, 20, 69
- Geman, S., & Geman, D. 1984, IEEE Trans. Pattern Anal. Mach. Intell., 6, 721
- Giovannelli, J.-F., & Coulais, A. 2005, A&A, 439, 401
- Giron, Francisco Javier, C. C. d. 2001, RACSAM, 95, 39
- González-Nuevo, J., Argüeso, F., Lopez-Caniago, M., Toffolatti, L., et al. 2006, Not. Roy. Astron. Soc., 1603
- Guglielmetti, F., Fischer, R., & Dose, V. 2009, MNRAS, 396, 165
- Haar, A. 1910, Math. Ann., 69, 331
- Haar, A. 1911, Math. Ann., 71, 38
- Hensley, B. S., Pavlidou, V., & Siegal-Gaskins, J. M. 2013, MNRAS, 433, 591
- Högbom, J. A. 1974, A&AS, 15, 417
- Hutchinson, M. F. 1989, Communications in Statistics – Simulation and Computation, 18, 1059
- Iatsenko, D., Stefanovska, A., & McClintock, P. V. E. 2012, Phys. Rev. E, 85, 3101
- Jasche, J., & Kitaura, F. S. 2010, MNRAS, 407, 29
- Jasche, J., & Wandelt, B. D. 2013, ApJ, 779, 15
- Jasche, J., Kitaura, F. S., Wandelt, B. D., & Enßlin, T. A. 2010, MNRAS, 406, 60
- Jaynes, E. T. 1957a, Phys. Rev., 108, 171
- Jaynes, E. T. 1957b, Phys. Rev., 106, 620
- Jordan, M. I., Ghahramani, Z., Jaakkola, T. S., & Saul, L. K. 1999, Machine Learning, 37, 183
- Junklewitz, H., Bell, M. R., Selig, M., & Enßlin, T. A. 2014, A&A, submitted [arXiv:1311.5282]
- Kinney, J. B. 2014, Phys. Rev. E, 90, 011301
- Kirkpatrick, S., Gelatt, C. D., & Vecchi, M. P. 1983, Science, 220, 671
- Kullback, S., & Leibler, R. A. 1951, Ann. Math. Stat., 22, 79
- Malyshev, D., & Hogg, D. W. 2011, ApJ, 738, 181
- Metropolis, N., & Ulam, S. 1949, J. Amer. Statist. Assn., 44, 335

- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., & Teller, E. 1953, *J. Chem. Phys.*, 21, 1087
- Nocedal, J., & Wright, S. J. 2006, *Numerical optimization*, <http://site.ebrary.com/id/10228772>
- Oppermann, N., Selig, M., Bell, M. R., & Enßlin, T. A. 2013, *Phys. Rev. E*, 87, 032136
- Planck Collaboration VII. 2011, *A&A*, 536, A7
- Rau, U., & Cornwell, T. J. 2011, *A&A*, 532, A71
- Schmitt, J., Starck, J. L., Casandjian, J. M., Fadili, J., & Grenier, I. 2010, *A&A*, 517, A26
- Schmitt, J., Starck, J. L., Casandjian, J. M., Fadili, J., & Grenier, I. 2012, *A&A*, 546, A114
- Selig, M., Oppermann, N., & Enßlin, T. A. 2012, *Phys. Rev. E*, 85, 021134
- Selig, M., Bell, M. R., Junklewitz, H., et al. 2013, *A&A*, 554, A26
- Shewchuk, J. R. 1994, *Techn. rep.*, Carnegie Mellon University, Pittsburgh, PA
- Strong, A. W. 2003, *A&A*, 411, L127
- Transtrum, M. K., & Sethna, J. P. 2012 [[arXiv:1201.5885](https://arxiv.org/abs/1201.5885)]
- Transtrum, M. K., Machta, B. B., & Sethna, J. P. 2010, *Phys. Rev. Lett.*, 104, 0201
- Valdes, F. 1982, in *SPIE Conf. Ser.*, 331, 465
- Wandelt, B. D., Larson, D. L., & Lakshminarayanan, A. 2004, *Phys. Rev. D*, 70, 083511
- Willett, R., & Nowak, R. 2007, *IEEE Trans.*, 53, 3171
- Wingate, D., & Weber, T. 2013 [[arXiv:1301.1299](https://arxiv.org/abs/1301.1299)]