

Unified radio interferometric calibration and imaging with joint uncertainty quantification

Philipp Arras^{1,2,3}, Philipp Frank^{1,2}, Reimar Leike^{1,2}, Rüdiger Westermann³, and Torsten A. Enßlin^{1,2}

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

² Ludwig-Maximilians-Universität München (LMU), Geschwister-Scholl-Platz 1, München, Germany

³ Technische Universität München (TUM), Boltzmannstr. 3, 85748 Garching, Germany

Received 26 March 2019 / Accepted 6 June 2019

ABSTRACT

The data reduction procedure for radio interferometers can be viewed as a combined calibration and imaging problem. We present an algorithm that unifies cross-calibration, self-calibration, and imaging. Because it is a Bayesian method, this algorithm not only calculates an estimate of the sky brightness distribution, but also provides an estimate of the joint uncertainty which entails both the uncertainty of the calibration and that of the actual observation. The algorithm is formulated in the language of information field theory and uses Metric Gaussian Variational Inference (MGVI) as the underlying statistical method. So far only direction-independent antenna-based calibration is considered. This restriction may be released in future work. An implementation of the algorithm is contributed as well.

Key words. techniques: interferometric – methods: statistical – methods: data analysis – instrumentation: interferometers

1. Introduction

Radio astronomy is thriving. Super-modern telescopes such as MeerKAT, the Australian Square Kilometer Array Pathfinder (ASKAP), the Very Large Array (VLA), and the Atacama Large Millimeter Array (ALMA) are operating and the Square Kilometer Array (SKA) is in the planning stages. All these telescopes provide high-quality data on an unprecedented scale and much progress is being made instrumental-wise, which facilitates enormous improvements in sensitivity and survey speed.

Impressed by these novel facilities we would like to turn our attention to the calibration and imaging algorithms that are fed by the data from these telescopes. The amount of scientific insight that can possibly be extracted from a given telescope is limited by the capability of the employed data reduction algorithm. We suggest that there is room for improvement regarding the calibration and imaging procedure: the most widely applied algorithms view calibration and imaging as separate problems and are not able to provide uncertainty information. The latter is desperately needed to quantify the level of trust a scientist can put on any result based on radio observations. Furthermore, a statistical sound confrontation of astrophysical models to radio data requires reliable uncertainty quantification. Treating calibration and imaging as separate steps ignores their tight interdependence.

The algorithmic idea presented in this work is an advancement of the original `RESOLVE` algorithm (Radio Extended SOURces Lognormal deconvolution Estimator; [Junklewitz et al. 2016](#)) and may retain its name. The `RESOLVE` algorithm is formulated in the language of information field theory (IFT; [Enßlin et al. 2009](#); [Enßlin 2018](#)), which is a view on Bayesian statistics applicable wherever (physical) fields are supposed to be inferred. From a Bayesian point of view the question when reducing radio data is the following: given prior knowledge as well as measurement

information about the brightness distribution of a patch of the sky, what knowledge does the observer have after obtaining the data? This question is answered by the Bayes theorem in terms of a probability distribution over all possible sky brightness distributions conditional to the data.

Reconstruction algorithms may be judged based on their statistical integrity or by their performance. The first perspective ultimately leads to pure Bayesian algorithms, which are too expensive for typical problems computationally. The latter often leads to ad hoc algorithms that may perform well in applications, but these can have major shortcomings such as a missing uncertainty quantification or negative-flux pixels, which is the case, for example, for CLEAN ([Högbom 1974](#)). The `RESOLVE` algorithm attempts a compromise between these two objectives. It is based on purely statistical arguments and the necessary operations are approximated such that they can efficiently be implemented on a computer and be used for actual imaging tasks. Thus, the approximations and (prior) assumptions on which `RESOLVE` is based can be written down explicitly.

`RESOLVE` is reasonably fast but cannot compete in pure speed with algorithms like the Cotton-Schwab algorithm ([Schwab 1984](#)) as implemented in CASA. This is rooted in the fact that `RESOLVE` not only provides a single sky brightness distribution but needs to update the sky prior probability distribution according to the raw data in order to properly state how much the data has constrained the probability distribution and how much uncertainty is left in the final result. This uncertainty is defined in a fashion such that it can encode the posterior variance and also cross-correlations. Thus, the uncertainty is quantified by $\mathcal{O}(n^2)$ pieces of information where n is the number of pixels in the image. Given this massive amount of degrees of freedom it may be surprising that `RESOLVE` is able to return its results after a sensible amount of time. Having said this, there is still potential for improvement. The technical cause for the long runtime is the

complexity of the gridding/degridding operation, which needs to be called orders of magnitude more often than for conventional algorithms. This problem may be tackled from an information-theoretic perspective in the future.

Turning to the specific subject of the present publication, the data reduction pipeline of modern radio telescopes consists of numerous steps. In this paper, we would like to focus on the calibration and imaging part. Calibration is necessary because the data is corrupted by a variety of effects including antenna-based, baseline-based, and direction-dependent or direction-independent effects (Smirnov 2011). For the scope of this paper only antenna-based calibration terms are considered, a simplification which is sensible for telescopes with a small field of view such as ALMA or the VLA. The crucial idea of this paper is to view the amplitude and phase corrections for each antenna as one-dimensional fields that are defined over time. These fields are discretized and regularized by a prior which states that the calibration solution for a given antenna is smooth over time. This removes the ambiguity of an interpolation scheme in between the calibrator observations and the subsequent application of self-calibration. Because RESOLVE is an IFT algorithm, there is no notion of solution intervals, which are time bins in which traditional calibration algorithms bin the data (see, e.g., Kenyon et al. 2018). Instead IFT takes care of a consistent discretization of the principally continuous fields. Similarly, the sky brightness distribution is defined on a discretized two-dimensional space; only single-channel imaging is performed in this work.

In practice, the current approach in the IFT community is to define a generative model that turns the degrees of freedom, which are learned by the algorithm into synthetic data that can be compared to the actual data in a squared-norm fashion (in the case of additive Gaussian noise). This approach is similar to the so-called radio interferometric measurement equation (RIME; Hamaker et al. 1996; Perkins et al. 2015; Smirnov 2011). Therefore, our notation closely follows the notation defined in Smirnov (2011). Calibration effects that are part of the RIME but left out for simplicity in this publication could in principle be integrated into the RESOLVE framework.

The RESOLVE approach may be classified according to the notion of first, second, and third generation calibration established in Noordam & Smirnov (2010): it unifies cross-calibration (1GC), self-calibration (2GC), and imaging. Still it is to be strictly distinguished from existing approaches like Kenyon et al. (2018), Cai et al. (2018), and Salvini & Wijnholds (2014). This is because it focuses on a strict Bayesian treatment combined with consistent discretization (one of the main benefits of IFT) and does not use computational speed as an argument to drop Bayesian rigidity.

The actual posterior probability distribution of the joint imaging and calibration problem is highly non-Gaussian and therefore not easily storable on a computer. In order to overcome this apparent problem the posterior is approximated by a multivariate Gaussian with full covariance matrix. The algorithm prescribes to minimize the Kullback–Leibler divergence (KL divergence) between the actual posterior and the approximate one which is the information gain between the two probability distributions. We use the variant of this known as Metric Gaussian Variational Inference (MGVI) (Knollmüller & Enßlin 2019).

Together with this publication we contribute an implementation of RESOLVE that is available under the terms of GPLv3¹. It is

based on the Python library NIFTy (Arras et al. 2019), which is freely available as well.

The paper is divided into four sections. Section 2 discusses the structure of likelihood and prior for the statistical problem at hand. This defines an algorithm which is verified on synthetic data in Sect. 3 and afterward applied to real data from the VLA in Sect. 4. Section 6 finishes the paper with conclusions and an outlook for future work.

2. The algorithm

2.1. Bayes' theorem

Every reconstruction algorithm needs a prescription of how the quantity of interest s affects the data d . This prescription is called the data model. Combined with statistical information, this model defines the likelihood $\mathcal{P}(d|s)$, which is a probability distribution on data realizations conditioned on a given realization of the signal s . Bayes' theorem,

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

requires us to supplement the likelihood with a prior probability distribution $\mathcal{P}(s)$, which assigns a probability to each signal realization s . This distribution encodes the knowledge the scientist has prior to looking at the data. Since it is virtually impossible to visualize the posterior probability distribution $\mathcal{P}(s|d)$ in the high dimensional setting of Bayesian image reconstruction we may compute the posterior mean and posterior variance as

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

$$\langle |m - s|^2 \rangle_{\mathcal{P}(s|d)} := \int \mathcal{D}s \mathcal{P}(s|d) |m - s|^2. \quad (3)$$

The notation $\int \mathcal{D}s$ is borrowed from statistical physics and means integrating over all possible configurations s . For a discussion on this measure in the continuum limit see Enßlin (2018, Sect. 1.8). In practice, this integral is discretized as follows: $\int \mathcal{D}s = \int \prod_i ds_i$ where s_i refers to the pixel values of the discretized quantity s . The term $\mathcal{P}(d)$ is independent of s and serves as a normalization factor. It expresses the probability to obtain the data irrespective of what the signal is, i.e. $\mathcal{P}(d) = \int \mathcal{D}s \mathcal{P}(d, s)$.

In the following we describe the data model and implied likelihood employed by RESOLVE. This includes the assumptions RESOLVE makes about the measurement process. Afterward, RESOLVE's prior is discussed. For definiteness the notation established in Smirnov (2011) is used.

2.2. Data model and likelihood

The measurement equation of a radio interferometer can be understood as a modified Fourier transform followed by an application of data-corrupting terms, the terms which need to be solved for in the calibration procedure. Assume that the data is only corrupted by so-called antenna-based direction-independent effects. Then Smirnov (2011, Eq. (18)) is written as

$$V_{pq} = G_p \left(\int \frac{B(l, m)}{n(l, m)} e^{-2\pi i [u_{pq}l + v_{pq}m + w_{pq}(n(l, m) - 1)]} dl dm \right) G_q^\dagger, \quad (4)$$

where

¹ <https://gitlab.mpcdf.mpg.de/ift/resolve>

- l, m : Direction cosines on the sky and $n(l, m) = \sqrt{1 - l^2 - m^2}$.
- $p, q \in \{1, \dots, N_a\}$: Antenna indices where N_a is the total number of antennas of the interferometer.
- $V_{pq} \in \mathbb{C}^{2 \times 2}$: Visibility for antenna pair (pq) .
- (u_{pq}, v_{pq}, w_{pq}) : Vector that connects antenna p with antenna q . The coordinates u_{pq} and v_{pq} are aligned with l and m , respectively. The value w_{pq} is perpendicular to both and points from the interferometer toward the center of the field of view.
- $G_p \in \mathbb{C}^{2 \times 2}$: Antenna-based direction-independent calibration effect.
- $B \in \mathbb{R}^{2 \times 2}$: Intrinsic sky brightness matrix. Since only the Stokes I component is considered in this publication, this matrix is proportional to the identity matrix.

Equation (4) can be understood as a Fourier transform of the sky brightness distribution, which is distorted by the terms involving $n(l, m)$ and corrupted by the calibration terms G_p . For the purpose of this publication we make the following simplifying assumptions: First, only the total intensity I is reconstructed. Second, G_p is assumed to be diagonal, which states that there is no significant polarization leakage and especially no time-variable leakage. Finally, the temporal structure of the data is needed for the construction of the prior. Therefore, a time index is added to the above expression that is written as

$$V_{pqt} = G_p(t) \left(\int \frac{B(l, m)}{n(l, m)} e^{-2\pi i [u_{pq}l + v_{pq}m + w_{pq}(n(l, m) - 1)]} dl dm \right) G_q^\dagger(t), \quad (5)$$

where $G_p(t)$ are diagonal matrices and $B(l, m)$ is a diagonal matrix, which is proportional to unity in polarization space. We note that $G_p(t)$ needs to absorb the V -term from (4), which is possible as long as polarization leakage is not too time variable. The w -term can be taken care of by w -stacking (Offringa et al. 2014), which means that the range of possible values for w_{pq} is binned linearly such that the integral becomes an ordinary Fourier transform. Technically, the non-equidistant Fourier transform in Eq. (5) is carried out by the NFFT library (Keiner et al. 2009) in our RESOLVE implementation.

All in all, Eq. (5) prescribes how to simulate data V_{pqt} given calibration solutions $G_p(t)$ and an inherent sky brightness distribution $B(l, m)$, which is what we wanted. In order to declutter the notation in the following let us denote the quantities of interest by $s = (G_p(t), B(l, m))$ and the map R such that $V_{pqt} = R(s)$.

The commonly used data model is the following: $d = R(s) + n$. It assumes additive Gaussian noise (Thompson et al. 1986). Let N be a diagonal noise covariance matrix with the noise variances on its diagonal and $\mathcal{G}(s-m, S)$ refers to a Gaussian random field with mean m and covariance matrix S . Then, the additive noise can be marginalized over to arrive at an expression for the likelihood

$$\mathcal{P}(d|s) = \int \mathcal{D}n \mathcal{P}(d|s, n) \mathcal{P}(n) \quad (6)$$

$$= \int \mathcal{D}n \delta(n - (d - R(s))) \mathcal{G}(n, N) \quad (7)$$

$$= \mathcal{G}(d - R(s), N). \quad (8)$$

The likelihood distribution $\mathcal{P}(d|s)$ contains all information about the measurement device and the measurement process that the inference algorithm will take into account.

We conclude the discussion on data and likelihood with three remarks: First, the likelihood does not depend on the statistical method at hand. All simplifications being made are rooted

in practical reasons in the implementation process. There is no fundamental reason for not taking, for instance, a more accurate noise model or a more sophisticated calibration structure into account.

Second, the employed notation already hints at the goal of describing an algorithm that jointly calibrates and images: the generalized response function R takes at the same time the calibration parameters $G_p(t)$ and the intrinsic sky brightness distribution B as an argument.

Finally, we consider what happens if the telescope alternates between observing the science target and observing a calibration source. Then, both the data set and the intrinsic sky brightness consists of two parts and the likelihood separates into

$$\mathcal{P}(d|s) = \mathcal{P}(d_c|s) \mathcal{P}(d_s|s). \quad (9)$$

From the likelihood perspective, calibration and science source are two separate things. However, as soon as the one-dimensional calibration fields are supplemented by a prior that imposes temporal smoothness the degrees of freedom regarding the science target and calibration target interact. This solves the problem of applying interpolated calibration solutions in traditional cross-calibration in a natural way.

2.3. Prior

Turning to the prior probability distribution, we note that the technical framework in which RESOLVE is implemented allows for a variety of different priors, which may supersede that presented in this paper.

As stated before $G_p(t)$ are assumed to be diagonal,

$$G_p(t) = \begin{pmatrix} g_p^{(0)}(t) & 0 \\ 0 & g_p^{(1)}(t) \end{pmatrix}. \quad (10)$$

The elements of this matrix are functions defined over time and take the following complex nonzero values²:

$$g_p^{(i)} : [t_0, t_1] \rightarrow \mathbb{C}^*, \quad i \in \{0, 1\}, p \in \{1, \dots, N_a\}. \quad (11)$$

The natural way of parameterizing a function taking values in \mathbb{C}^* is in polar coordinates, i.e.,

$$g_p^{(i)}(t) = \exp(\lambda_p^{(i)}(t) + i\phi_p^{(i)}(t)), \quad (12)$$

where $\lambda_p^{(i)} : [t_0, t_1] \rightarrow \mathbb{R}$ and $\phi_p^{(i)} : [t_0, t_1] \rightarrow \mathbb{R}/2\pi\mathbb{Z}$. The modulus and phase of the complex gains $g_p^{(i)}$ have different physical origins. The modulus describes a varying amplification of the signal in the antenna electronics, which is rooted amongst others in fluctuating temperatures of the receiver system. Varying phases stem from fluctuations in the atmosphere. Therefore, these two ingredients of g_p have differing typical timescales a priori.

The prior knowledge on $\lambda_p^{(i)}$ and $\phi_p^{(i)}$ is the following: $\{\lambda_p^{(i)}, \phi_p^{(i)}\}$, respectively, share a typical behavior over time for all antennas p , both of which are not known a priori and need to be inferred from the data as well. This typical behavior does not change over time. Additionally, all $\lambda_p^{(i)}, \phi_p^{(i)}$ evolve smoothly over time. Mathematically, this can be captured by Gaussian random fields,

$$\mathcal{P}\left(\lambda_p^{(i)}, \phi_p^{(i)}\right)_{i,p} \Big| \Lambda, \Phi = \prod_{i,p} \mathcal{G}(\lambda_p^{(i)}, \Lambda) \mathcal{G}(\phi_p^{(i)}, \Phi), \quad (13)$$

² \mathbb{C}^* are the units of \mathbb{C} , i.e., $\mathbb{C}^* := \mathbb{C} \setminus \{0\}$.

where Λ, Φ is defined such that the Gaussian random fields obey homogeneous but still specifically unknown statistics. This means that not only the calibration solutions themselves but also their prior correlation structure is inferred. For this a prior on the covariances needs to be supplemented: $\mathcal{P}(\Lambda), \mathcal{P}(\Phi)$. In Sect. 2.4 we describe how to set up the prior on Λ and Φ such that they implement homogeneous statistics and which parameters they take.

Next, let us discuss the prior on the sky brightness distribution $B(l, m)$. We recall that the matrix $B(l, m)$ is assumed to be diagonal and proportional to unity, i.e.,

$$B(l, m) = \begin{pmatrix} b(l, m) & 0 \\ 0 & b(l, m) \end{pmatrix}, \quad (14)$$

where $b(l, m) : [l_{\min}, l_{\max}] \times [m_{\min}, m_{\max}] \rightarrow \mathbb{R}_{>0}$ map the field of view to the set of positive real numbers since sky brightness is inherently a positive quantity³. For the scope of this publication, the sky brightness contains only a diffuse component. It shall be modeled similarly to the modulus of the calibration terms: it is strictly positive a priori, smooth over its domain and may vary over large scales. Therefore, we define $b(l, m) = e^{\psi(l, m)}$ and let $\psi(l, m)$ be a two-dimensional Gaussian random field with correlation structure Ψ , which is going to be inferred as well:

$$\mathcal{P}(\psi|\Psi) = \mathcal{G}(\psi, \Psi). \quad (15)$$

All in all, the basic structure of the priors on all terms appearing in Eq. (5) has been described apart from the construction of the prior on all covariance matrices, which is the objective for the next section.

2.4. Correlated fields

To account for correlations of a Gaussian distributed field ψ the following statements are assumed to be true:

1. The autocorrelation of ψ can be characterized by a power spectrum $P_{\Psi}(|k|)$, where k is the coordinate of the Fourier transformed field.
2. The power spectrum $P_{\Psi}(|k|)$ is a positive quantity that can vary over many orders of magnitudes.
3. Physical power spectra are falling with $|k|$, typically according to a power law.
4. Given enough data, it is possible to infer any kind of differentiable power spectrum.

Note that the first assumption is equivalent to the seemingly weaker assumptions:

- In absence of data, there is no special direction in space or time, i.e., a priori the correlation of the field is invariant under rotations.
- In absence of data, there is no special point in space or time, i.e., a priori the correlation of field values is invariant under shifts in space or time.

The fact that homogeneous and isotropic correlation matrices are diagonal in Fourier space and can be fully characterized by a power spectrum is known as the Wiener–Khinchin theorem (Wiener 1949; Khintchin 1934).

It is assumed that ψ as well as its power spectrum $P_{\Psi}(|k|)$ are unknown. Therefore, both need a prior that may be formulated as generative model: an operator that generates samples for ψ and its square root power spectrum (henceforth called amplitude

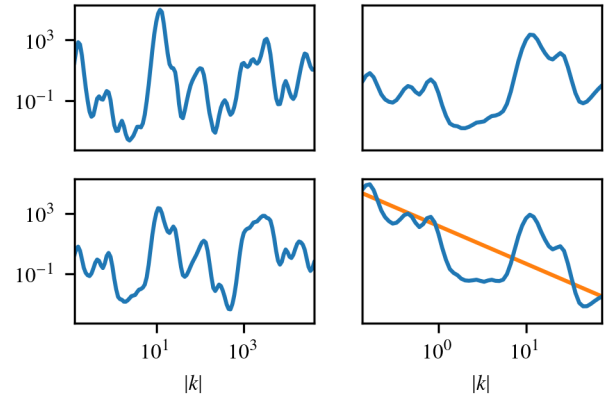


Fig. 1. Steps of the generative process defined in Eq. (16). *Top left:* smooth, periodic field defined on the interval $[t_0, 2t_1 - t_0]$. *Bottom left:* (anti-)symmetrized version of the above. *Top right:* projection of the symmetrized field to half of the original domain $[t_0, t_1]$. *Bottom right:* resulting double logarithmic amplitude spectrum after addition of the power law (orange) to the above.

spectrum) from one or multiple white Gaussian fields. Formulating a prior as a generative model has several theoretical and practical advantages (Knollmüller & EnBlin 2018).

We propose the following ansatz for an operator that converts independent normal distributed fields, ϕ and τ to the amplitude spectrum of the correlated field ψ . This operator is called amplitude operator A_C (see Fig. 1 for an illustrative example), i.e.

$$A_C(\tau, \phi) = (\text{Exp}^* \text{Exp}) \left(0.5 \cdot \left[\log(k)(\sigma_m \tau_m + \bar{m}) + \sigma_{y_0} \tau_{y_0} + \bar{y}_0 + (\text{sym} \circ \tilde{\mathcal{F}}_{\log(k)t})(\text{cp}(t) \cdot \phi(t)) \right] \right), \quad (16)$$

where $C = (a, t_0, \bar{m}, \bar{y}_0, \sigma_m, \sigma_{y_0}, \alpha, \beta)$ denotes the tuple of parameters (all real numbers), Exp^* denotes the pullback of a field by the exponential function acting on $\log(|k|)$ ⁴, Exp denotes exponentiation of the field values, $\tilde{\mathcal{F}}_{\log(k)t}$ denotes the Fourier transform of a space with coordinates t to the logarithmic coordinates $\log(k)$ of the power spectrum, \bar{m} and \bar{y}_0 are the slope and the y -intercept of the a priori mean power law, sym is an (anti-)symmetrizing operation defined to operate on a field ϕ over the interval $[t_0, t_1]$ as

$$2 \cdot \text{sym}(\phi)(x) = \phi(x) - \phi(2t_1 - t_0 - x), \quad (17)$$

for $x \in (t_0, 2t_1 - t_0)$. In words, sym mirrors the field and subtracts it from itself, then restricts the domain to half the original size. Finally, cp is the log-cepstrum,

$$\text{cp}(t) = a \cdot \left(1 + (t/t_0)^{-2} \right). \quad (18)$$

Let us show that Eq. (16) meets the requirements stated at the beginning of Sect. 2.4. Requirement 1 is trivial. Requirement 2 is met since the amplitude spectrum is constructed by applying an exponential function to a Gaussian field. Thus, all values are positive and can vary over several order of magnitudes.

To requirement 3: In absence of data, the mean of the inferred white fields ϕ and τ , to which the amplitude operator is applied, remains zero. For $\phi = 0$ and $\tau = 0$, Eq. (16) becomes

$$(\text{Exp}^* \text{Exp}) (0.5 \cdot [\bar{m} \log(k) + \bar{y}_0]), \quad (19)$$

⁴ Let $\phi : U \rightarrow V$ with $U, V \subseteq \mathbb{R}$ open and $f : V \rightarrow \mathbb{R}$ a smooth function, i.e., a field. Then $(\phi^* f)(t) := f(\phi(t))$ denotes the pullback of f by ϕ . In other words, the field f is transformed to a different coordinate system whose coordinates are related to the original one by ϕ .

³ We note the difference to Högbom's CLEAN, which has positivity not built in (Högbom 1974).

which is the equation for a power law with spectral index \bar{m} . A preference for falling spectra can be encoded by choosing the hyperparameter \bar{m} to be negative.

To requirement 4: Let us show that any differentiable function lies in the image space of the amplitude operator. This implies that any differentiable amplitude spectrum can be inferred given enough data. Let ϕ be an arbitrary smooth field over the interval $[t_0, t_1]$ and ϕ_{sym} be a smooth field that has a point symmetry at $(t_1, \phi(t_1))$ and is defined on the interval $[t_0, 2t_1 - t_0]$ as

$$\phi_{\text{sym}}(t) = \begin{cases} \phi(t) & \text{for } t \in [t_0, t_1], \\ 2\phi(t_1) - \phi(2t_1 - t) & \text{for } t \in [t_1, 2t_1 - t_0]. \end{cases} \quad (20)$$

The function ϕ_{sym} is a continuous and differentiable continuation of ϕ at t_1 . Now, we decompose ϕ_{sym} into a linear part and a residual term:

$$\phi_{\text{sym}}(t) = m \cdot (t - t_0) + y_0 + \phi_{\text{res}}(t), \quad (21)$$

where

$$y_0 = \phi_{\text{sym}}(t_0), \quad (22)$$

$$m = \frac{\phi_{\text{sym}}(2t_1 - t_0) - \phi_{\text{sym}}(t_0)}{2(t_1 - t_0)}, \quad (23)$$

$$\phi_{\text{res}}(t) = -m \cdot (t - t_0) - y_0 + \phi_{\text{sym}}(t). \quad (24)$$

The residual term is a differentiable periodic function, i.e.,

$$\begin{aligned} \phi_{\text{res}}(t_0) &= \phi_{\text{res}}(2t_1 + t_0) \\ \Leftrightarrow \phi_{\text{sym}}(t_0) - \phi_{\text{sym}}(t_0) &= -\phi_{\text{sym}}(2t_1 - t_0) + \phi_{\text{sym}}(t_0) \\ &\quad - \phi_{\text{sym}}(t_0) + \phi_{\text{sym}}(2t_1 + t_0) \end{aligned} \quad (25)$$

$$\begin{aligned} \phi'_{\text{res}}(t_0) &= \phi'_{\text{res}}(2t_1 + t_0) \\ \Leftrightarrow \phi'(t_0) - m &= \phi'_{\text{sym}}(2t_1 + t_0) - m \\ \Leftrightarrow \phi'(t_0) - m &= \phi'(t_0) - m. \end{aligned} \quad (26)$$

Thus, ϕ_{res} can be represented in Fourier space by a field that falls off at least with second order. This is exactly how ϕ_{res} is represented in Eq. (16). Assuming that the mean and the slope of the linear part are well represented by its prior distribution, it is indeed possible to represent any kind of differentiable amplitude spectrum. All in all, all four requirements are met by Eq. (16).

There remains one unconstrained degree of freedom, the value of the power spectrum at $|k| = 0$, the zero mode. As the zero mode describes the magnitude of the overall logarithmic flux, it is decoupled from the remaining spectrum and should have its own prior. This value is fixed by imposing an inverse gamma prior on the zero mode, which restricts it to be a positive quantity, while still allowing for large deviations.

To sum up, the amplitude operator depends on the following eight hyper parameters:

- a, t_0 : The amplitude parameter and cutoff scale of the log-cepstrum.
- \bar{m}, \bar{y}_0 : The prior means for the slope and the height of the power law.
- σ_m, σ_{y_0} : The corresponding standard deviations.
- α, β : The shape and scale parameter of the inverse gamma prior for the zero mode.

We note that the assumptions made at the beginning of Sect. 2.4 apply to a wide variety of processes, regardless of their dimensionality. This generic correlated field model has already been successfully used in a number of synthetic and real applications (Leike & Enßlin 2019; Knollmüller & Enßlin 2018,

2019; Knollmüller et al. 2018; Hutschenreuter & Enßlin 2019). In RESOLVE, the amplitude operator is used as a prior for the amplitude spectra of the antenna calibration fields and the image itself.

2.5. Full algorithm

In the foregoing sections, the full likelihood and prior are described. Now, we stack all the ingredients together to build the full algorithm. Let us assume that the data set consists out of two alternating observations: observations of a calibrator source and observations of the science target. This means that the likelihood splits into two parts as indicated in Eq. (9). In contrast to the sky brightness distribution of the science target that of the calibrator B_c is known: it is a point source in the middle of the field of view. The sky brightness distribution of the science target is reconstructed.

The full likelihood takes the form

$$\mathcal{P}(d_t|\xi)\mathcal{P}(d_c|\xi) = \prod_{a \in \{t,c\}} \mathcal{G}(d_a - R_a(\{G_p^{(i)}\}, B_a), N_a \otimes \mathbb{1}), \quad (27)$$

$$B_t = \exp \circ \mathcal{F} \circ (\xi_B \cdot A^B), \quad (28)$$

$$G_p^{(i)} = \begin{pmatrix} g_p^{(i)} & 0 \\ 0 & g_p^{(i)} \end{pmatrix}, \quad (29)$$

$$g_p^{(j)} = \exp(\lambda_p^{(j)} + i\phi_p^{(j)}), \quad (30)$$

$$\lambda_p^{(i)} = Z \circ \mathcal{F} \circ (\xi_{\lambda_p^{(i)}} \cdot A^\lambda), \quad (31)$$

$$\phi_p^{(i)} = Z \circ \mathcal{F} \circ (\xi_{\phi_p^{(i)}} \cdot A^\phi), \quad (32)$$

$$A^B = A_{C_B}(\xi_{A_B}), \quad (33)$$

$$A^\lambda = A_{C_\lambda}(\xi_{A_\lambda}), \quad (34)$$

$$A^\phi = A_{C_\phi}(\xi_{A_\phi}), \quad (35)$$

where C_x denote the tuple of parameters of the respective amplitude operator, Z is a padding operator. The unit matrices in Eq. (27) is a 2×2 matrix acting on the same space as the sky brightness matrix B . The tuple of all excitation fields is called ξ , where

$$\xi = \left(\xi_B, \xi_{A_B}, \xi_{A_\lambda}, \xi_{A_\phi}, \xi_{\lambda_0^{(0)}}, \dots, \xi_{\lambda_{N_a}^{(1)}}, \xi_{\phi_0^{(0)}}, \dots, \xi_{\phi_{N_a}^{(1)}} \right). \quad (36)$$

As discussed before this model is set up such that the excitation fields ξ have white Gaussian statistics a priori,

$$\mathcal{P}(\xi) = \mathcal{G}(\xi, \mathbb{1}). \quad (37)$$

The posterior probability distribution is given by

$$\mathcal{P}(\xi|d_t, d_c) \propto \mathcal{P}(d_t, d_c, \xi) = \mathcal{P}(d_t|\xi)\mathcal{P}(d_c|\xi)\mathcal{P}(\xi). \quad (38)$$

Finally, the statistical model that is employed in this publication is fully defined.

2.6. Inference algorithm

The probability distribution (38) has too many degrees of freedom to be represented on a computer. The RESOLVE algorithm solves this problem by approximating this full posterior distribution by a multivariate Gaussian distribution whose covariance is equated with the inverse Fisher information metric. The latter can be represented symbolically alleviating the need for an explicit storage and handling of otherwise prohibitively large matrices. This algorithm is called MGVI and

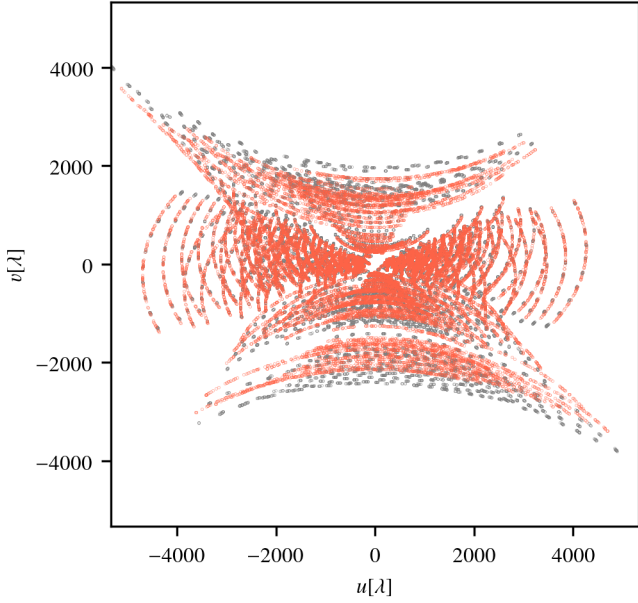


Fig. 2. Random sample (30 000 points) of uv coverage of a G327.6+14.6 (SN1006) observation with the VLA. The gray and red points indicate the uv coverage of the calibration source and science target, respectively.

is described in full length in Knollmüller & Enßlin (2019) and implemented in NIFTy⁵. The following is an inexhaustive outline of Knollmüller & Enßlin (2019).

The algorithm MGVI prescribes to minimize the KL divergence⁶ between the actual posterior and approximate posterior such that

$$\text{KL}(\mathcal{P}_1 \parallel \mathcal{P}_2) = \int \mathcal{D}s \mathcal{P}_1(s) \log \left(\frac{\mathcal{P}_1(s)}{\mathcal{P}_2(s)} \right), \quad (39)$$

where \mathcal{P}_1 is more informed compared to \mathcal{P}_2 . However, it is apparent that it is virtually impossible to perform the integration with respect to the posterior distribution as integration measure. Therefore, MGVI exchanges the order of the arguments of the KL divergence such that the integral can be approximated by samples of the approximate posterior, i.e.,

$$F[\xi] = \langle \mathcal{H}(\xi + x, d) \rangle_{x \sim \mathcal{G}(x, D(\xi))}, \quad (40)$$

where $\mathcal{H}(\xi, d) := -\log \mathcal{P}(\xi, d)$ is the information Hamiltonian and $D(\xi)$ the Fisher information. The parameter $F[\xi]$ is a cost function that can be minimized with respect to ξ . Suitable (2nd order) minimizers are provided by NIFTy.

With the help of the above approximation scheme we get a computational handle on the posterior. The drawbacks of this approach include the uncertainty estimate of MGVI sets a lower bound on the variance of the posterior and it is not suited for extremely non-Gaussian and especially multimodal probability distributions. But we note that the posterior is approximated with a Gaussian in the space on which the parameters are defined. After processing the parameters through nonlinearities as discussed in this section the actual quantities of interest such as the sky brightness distribution are not Gaussian distributed any more and may even have multiple modes. A detailed discussion on the abilities of MGVI is provided in Knollmüller & Enßlin (2019).

⁵ <https://gitlab.mpcdf.mpg.de/ift/nifty>

⁶ Also known as discrimination information.

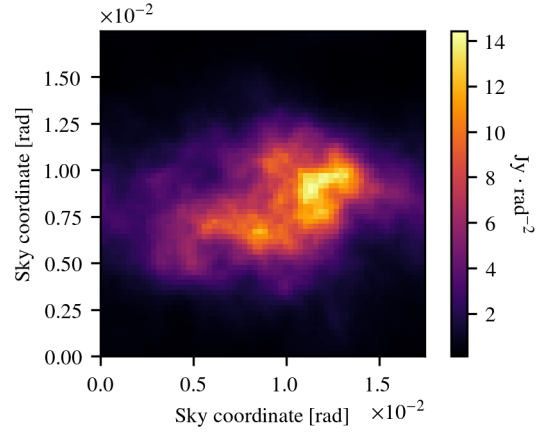


Fig. 3. Synthetic observation: ground truth sky brightness distribution $b(l, m)$ with 60' field of view.

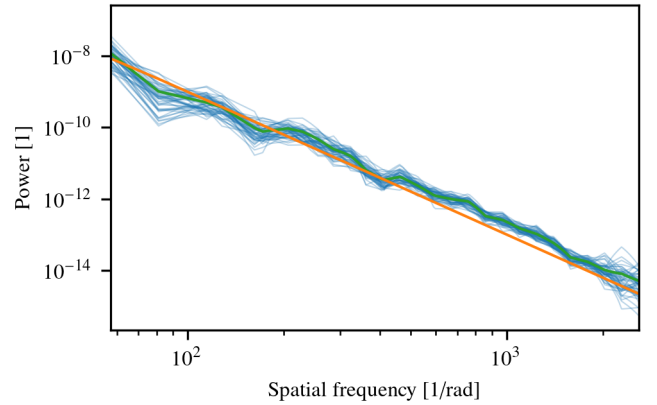


Fig. 4. Synthetic observation: power spectrum of log-sky brightness distribution. Orange: ground truth; green: posterior mean; and blue: posterior samples.

3. Verification on synthetic data

This section is devoted to the verification of the algorithm, i.e., the reconstruction of a synthetic sky brightness distribution from a simulated observation and artificial noise. The setup is described followed by a comparison of the ground truth and the reconstruction. Application to real data, where effects that are not modeled may occur and the ground truth is unknown, is presented in Sect. 4.

We employ a realistic uv coverage. It is an L -band observation of the source G327.6+14.6, also known as SN1006⁷. For the purpose of this paper we randomly select 30 000 visibilities from this data set to demonstrate that joint calibration and imaging is possible even without much data. (see Fig. 2). We use the field shown in Fig. 3 as the synthetic sky brightness distribution. It is a random sample assuming the power spectrum shown in orange in Fig. 4. The noiseless simulated visibilities are corrupted by noise whose level is visualized in Fig. 5. The resulting information source, i.e., the naturally weighted dirty image, is shown in Fig. 6.

This synthetic observation is set up in a fashion such that the calibration artifacts are stronger and the noise level is higher as compared to real data (see Sect. 4) to demonstrate the capability

⁷ VLA archive project code: AM0754, Jan 24, 2003, L -Band 1369.95 MHz, CnD configuration.

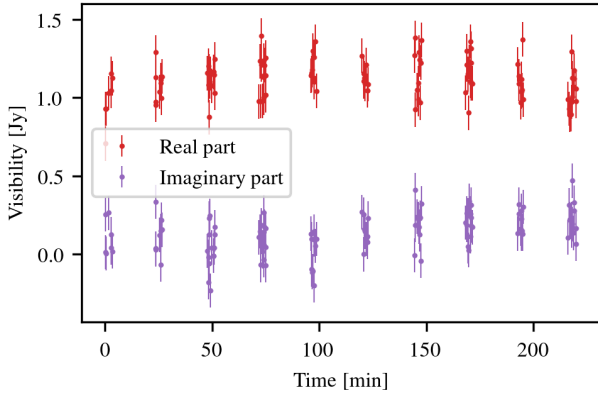


Fig. 5. Synthetic observation: visibilities of calibrator observation (polarization L, only visibilities of antennas 1 and 3). Thus, a constant value of $1 + 0j$ Jy is expected. All deviations from this are either noise or calibration errors. The error bars show the standard deviation on the data points.

of the `RESOLVE` in bad data situations. The calibration artifacts that have been applied are visualized in Fig. 7.

The `RESOLVE` algorithm is run on this synthetic data to compare its output and uncertainty estimation to the (known) ground truth. The prior parameters are listed in Table 1. Additionally, we choose a resolution of 64^2 pixels for the sky brightness distribution with a field of view of $60'$ and 256 pixels for the calibration fields that are defined on a temporal domain. As the total length of the observation was ~ 220 min one temporal pixel is ~ 50 s long. These temporal pixels should not be confused with solution intervals of traditional calibration schemes where the data is binned on a grid and then the calibration parameters are solved for. In IFT fields are by their nature continuous quantities that are discretized on an arbitrary grid. For convenience a regular grid was chosen. Then the data provides information on each pixel that is propagated to the neighboring pixels through the prior; the calibration fields are assumed to be smooth over time. Therefore, the user is free to choose the resolution of the fields in IFT algorithms as long as it is finer than the finest structure that shall be reconstructed.

As pointed out `RESOLVE` is a Bayesian algorithm whose output is not a single image of the observed patch of the sky but rather a probability distribution of all possible sky configurations. The MGVI algorithm approximates this non-Gaussian probability distribution with a Gaussian in the space of ξ , i.e., the eigenspace of the prior covariance. This again implies non-Gaussian statistics on quantities such as $b(l, m)$, $\lambda_p^{(i)}$, and $\phi_p^{(i)}$ since they depend in a nonlinear fashion on ξ . The only useful way to visualize this probability distribution is to analyze a finite number of samples from it which `RESOLVE` can generate. A given set of samples can then be analyzed with standard statistical means such as the pixel-wise mean and variance.

Figures 8–10 show the posterior mean, the absolute value of the residual, the standard deviation of the sky brightness distribution, and a histogram of the residual divided by the standard deviation computed from 100 posterior samples, respectively. The algorithm has managed to perform the calibration correctly and to reconstruct the sky brightness distribution. The total flux of the ground truth (Fig. 3) could not totally be recovered because of the noise on the synthetic measurement. Remarkably, the proposed uncertainty is a bit too small compared to the residuals which is what is to be expected from MGVI.

Since `RESOLVE` does not assume a specific power spectrum as prior for the reconstruction but rather learns it together with

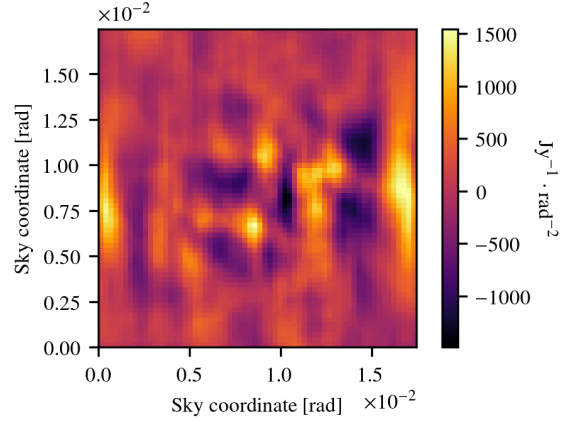


Fig. 6. Synthetic observation: information source $j = R_t^\dagger N_t^{-1} d_t$.

the sky brightness from the data, `RESOLVE` also provides the user uncertainty on the power spectrum; see Fig. 4. We note that the posterior variance on the power spectrum increases toward the boundaries of the plot. This is because interferometers do not provide information on scales larger than those that belong to the shortest baseline. On small scales an interferometer is limited by the noise level, which leads to an increased variance in the power spectrum on the right-hand side of the plot.

Next, we turn to the calibration solutions. Figure 7 shows a comparison of the ground truth and the posterior provided by `RESOLVE`. Since two polarizations are considered (LL and RR) for both the amplitude and the phase of the antenna-based calibration term, Fig. 7 has four rows. On first sight, the posterior mean and the ground truth are indistinguishable by eye and the residuals and posterior standard deviation fit together nicely. There is a significant increase of the uncertainty for, for example, antenna 2 toward the end of the observation. This is because a flagged data set was used and that simply all data points involving this antenna have been flagged from the beginning of the observation up to ~ 2 h.

To illustrate this more explicitly, Figs. 11 and 12 show the calibration solution for one antenna, respectively. The ground truth lies within the bounds of uncertainty indicated by the samples. We note that all data points have been flagged on the left-hand side of Fig. 11. Since no information about the phase is available the only constraint is the prior, which enforces temporal smoothness. Consistently, the uncertainty increases where no information is available.

Finally, we demonstrate what kind of other information posterior samples can reveal. Say, a scientist is interested in the integrated flux over a certain region. In addition to the image, this integrated flux comes with an uncertainty that can be calculated by averaging over posterior samples of the sky brightness distribution. An example is shown in Fig. 13. The scatter of the histogram is caused by the noise influence on the data, the (un)certainly of the calibration solutions, and ultimately the uv coverage. We are not aware of any other radio aperture synthesis algorithm that is able to provide this kind of probabilistic posterior information. All in all, the proposed method is able to recover the ground truth and is able to supplement it with an appropriate uncertainty estimation.

4. Application to VLA data

We continue with an application of `RESOLVE` to real data. To this end, take the VLA data set whose uv coverage and time stamps

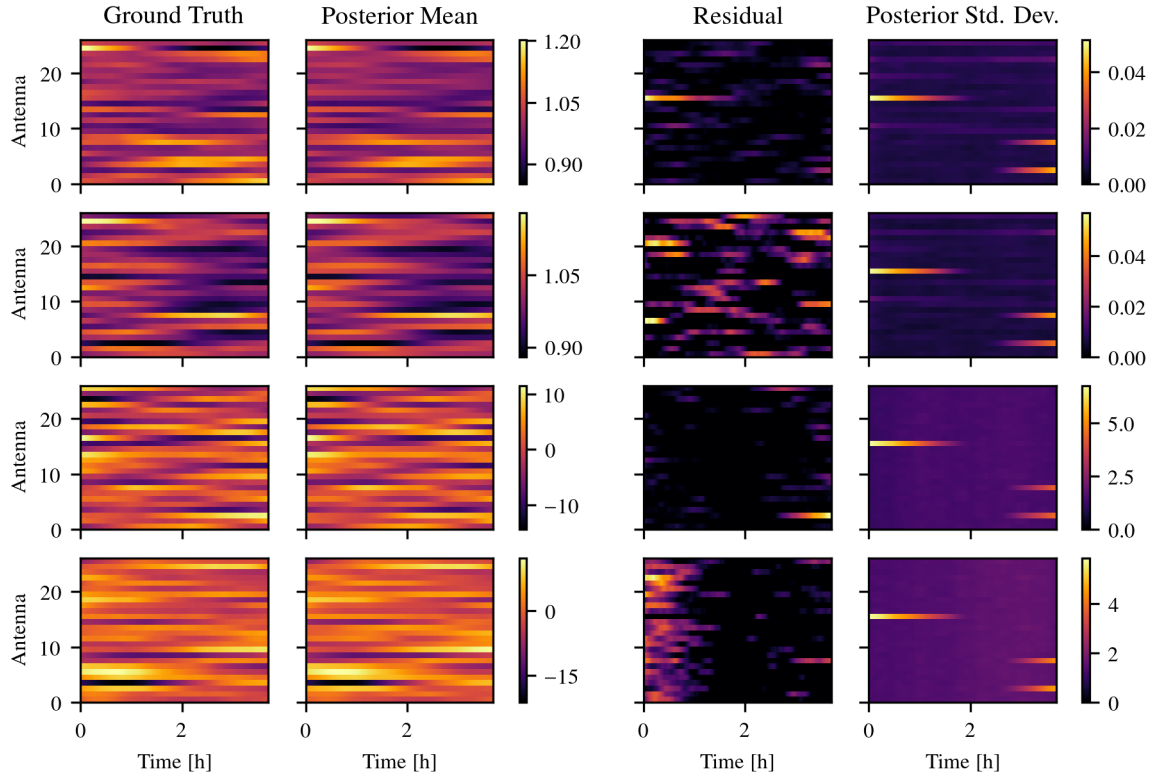


Fig. 7. Synthetic observation: calibration solutions. *First two rows:* amplitude and *bottom two rows:* phase calibration solutions. *First and third row:* LL-polarization and *second and last row:* RR-polarization. *Third column:* absolute value of the difference between posterior mean and ground truth. *Fourth column:* point-wise posterior standard deviation as provided by RESOLVE. Amplitudes do not have a unit as they are a simple factor applied to the data. Phases are shown in degrees.

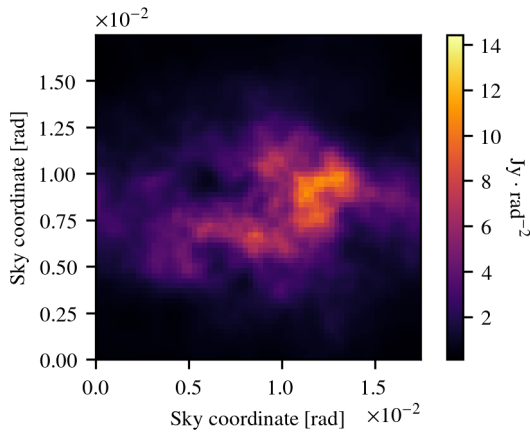


Fig. 8. Synthetic observation: posterior mean of sky brightness distribution $b(l, m)$.

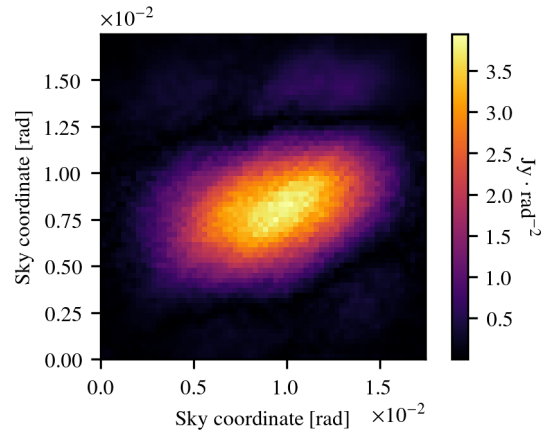


Fig. 9. Synthetic observation: absolute value of the difference between ground truth and posterior mean.

Table 1. Synthetic observation: prior parameters.

	a	t_0	\bar{m}	\bar{y}	σ_m	σ_{y_0}	α	β
A	2	2	-4	5	1	3	4	5×10^{-3}
λ	1.5	1	-4	-37	0.5	1	2	20
ϕ	1.5	1	-4	-36	0.5	1	2	20

have already been used in the preceding section. Also, the resolution of all spaces is taken to be the same.

Starting from raw data, the first thing to look at is the information source (see Fig. 14). No structure of the supernova remnant is visible whatsoever since the data is not calibrated yet. This illustrates that RESOLVE is able to operate on raw (but already flagged) visibilities that have not been processed further. Table 2 summarizes the prior parameters for the following reconstruction.

All calibration solutions are shown in Fig. 15 together with two exemplary plots in Figs. 16 and 17. The major characteristic of these solutions are hidden in the right-hand column of Fig. 15: the uncertainty on the calibration decreases whenever the calibrator source is observed as expected. Additionally,

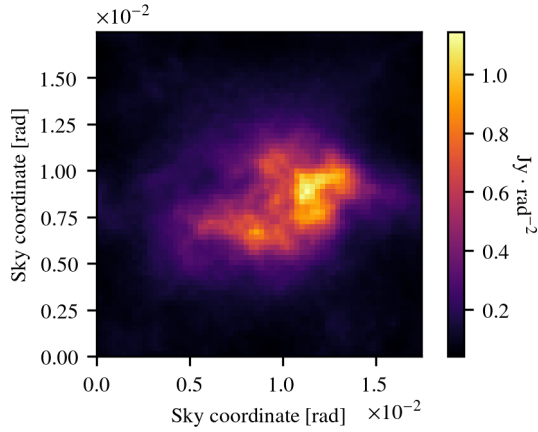


Fig. 10. Synthetic observation: pixel-wise standard deviation of posterior samples.

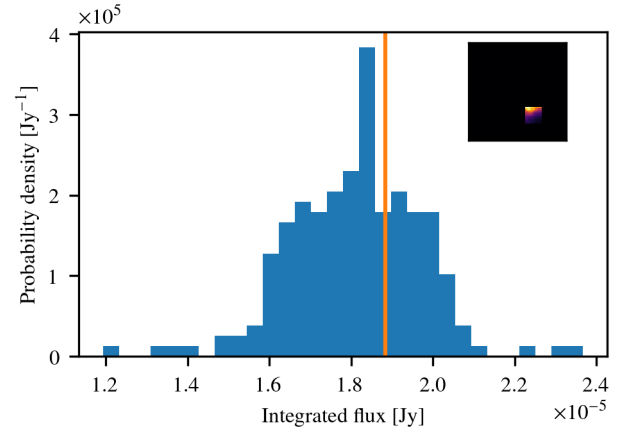


Fig. 13. Synthetic observation: histogram over samples of integrated flux in the region shown in the top right corner. Orange: ground truth.

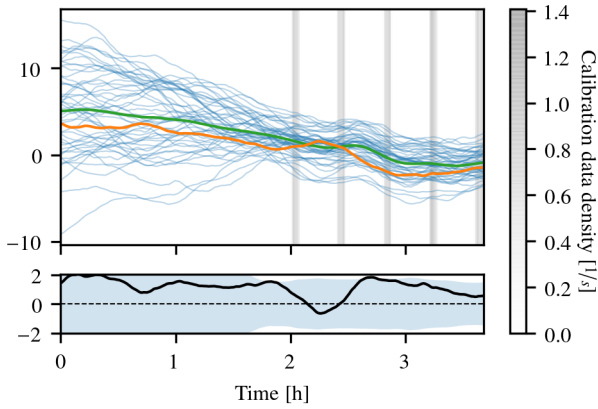


Fig. 11. Synthetic observation: phase solutions for antenna 15 and polarization R. Colors as in Fig. 12 and all phases are plotted in degrees. On the left-hand side all data points have been flagged.

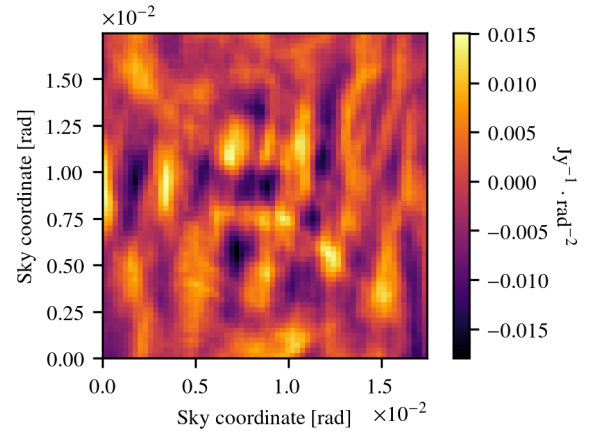


Fig. 14. SN1006: information source $j = R^\dagger N^{-1} d$.

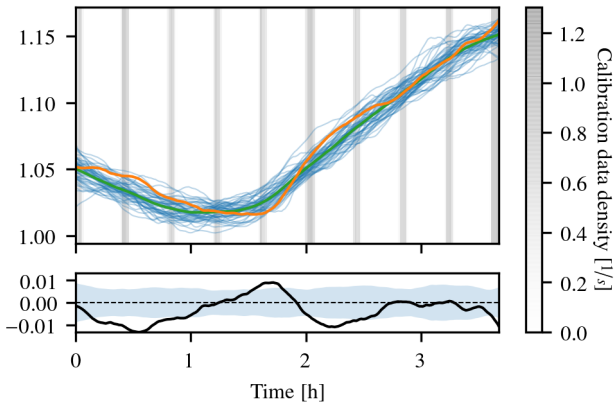


Fig. 12. Synthetic observation: amplitude solutions for antenna 0 and polarization L. Orange: ground truth; green: sampled posterior mean; and blue: posterior samples. The calibration data density shows how many data points of the calibrator observation are available. We note that a Bayesian algorithm can naturally deal with incomplete data or data from different sources. The *bottom plot* shows the residual along with the pixel-wise posterior standard deviation.

the uncertainty increases dramatically where the data has been flagged. The amplitude solutions are surprisingly stable over time although the prior would allow for more variance in the solution as can be seen from Sect. 3, where the same prior parameters have been used.

Table 2. SN1006: prior parameters.

	a	t_0	\bar{m}	\bar{y}	σ_m	σ_{y_0}	α	β
A	2	2	-4	2	1	2	4	1
λ	1.5	1	-4	-37	0.5	1	2	20
ϕ	1.5	1	-4	-36	0.5	1	2	20

There is a systematic difference between the samples for the amplitude solutions and those for the phases. The former vary around a mean solution whereas the latter exhibit a certain global offset. This is explained by the fact that the likelihood is invariant under a pixel-wise global phase shift, which is broken by the prior to a global phase shift to all temporal pixels at once. This residual symmetry is again broken by the prior on the zero-mode variance of the phase solutions. However, this prior is very weak to allow for phase solutions of arbitrary magnitude. Therefore, the phase solutions cannot have an arbitrarily large offset but still can globally vary to some degree, which is shown in Fig. 17.

Next, the posterior sky brightness is discussed. Figures 18 and 19, along with Figs. 20 and 21, show the posterior mean and pixel-wise standard deviation of $b(l, m)$. The posterior standard deviation is higher wherever more flux is detected. Therefore, Fig. 21 provides a descriptive visualization of the posterior uncertainty of the sky brightness distribution.

Last but not least the power spectrum of the logarithmic sky brightness distribution also needs to be reconstructed; this

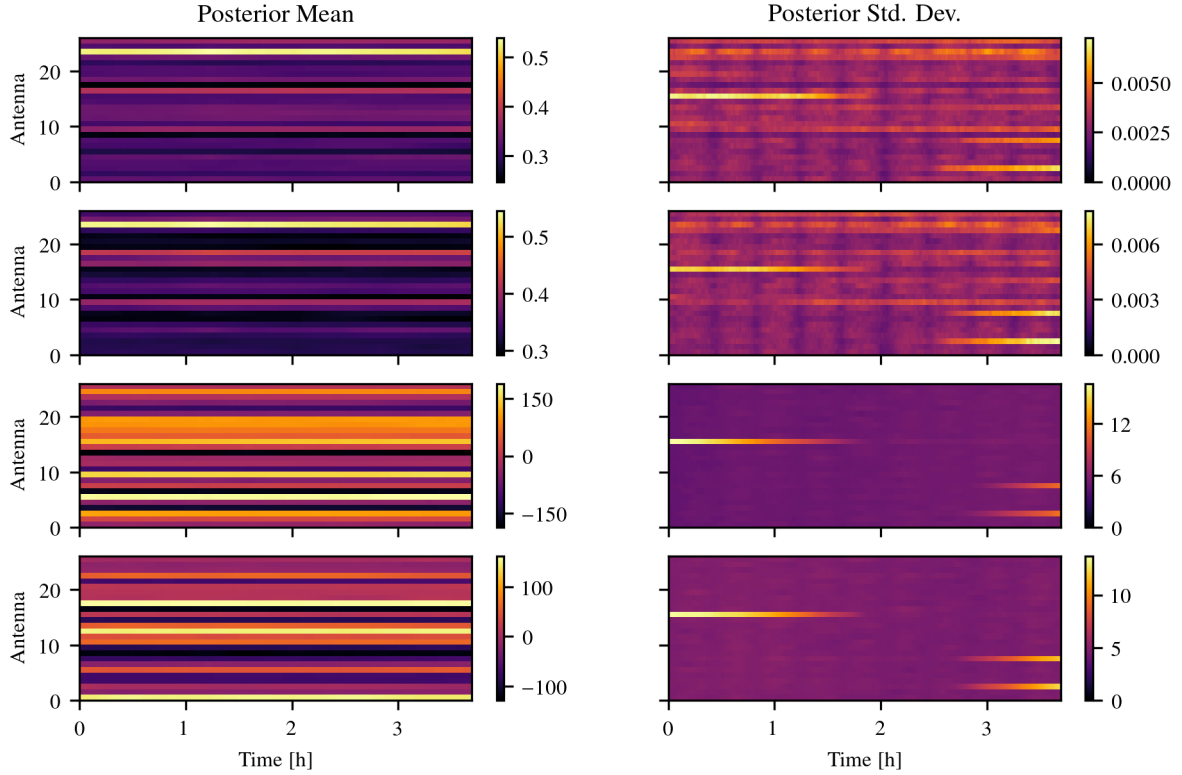


Fig. 15. SN1006: overview over calibration solutions. The four rows indicate amplitude and phase solutions for LL polarization and RR polarization as in Fig. 7.

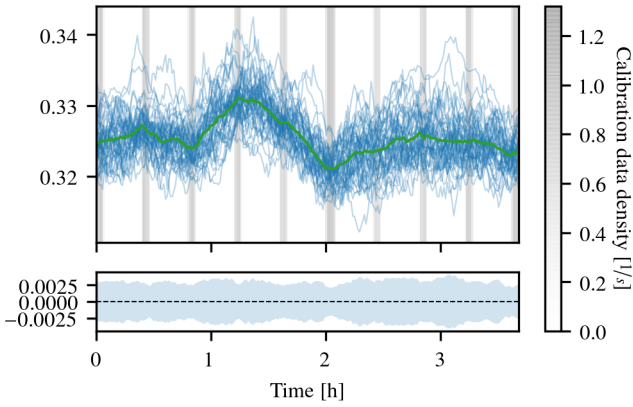


Fig. 16. SN1006: exemplary amplitude solution (see Fig. 12).

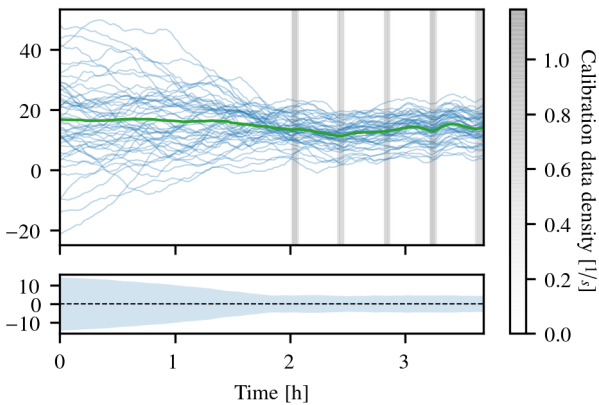


Fig. 17. SN1006: exemplary phase solution (see Fig. 11).

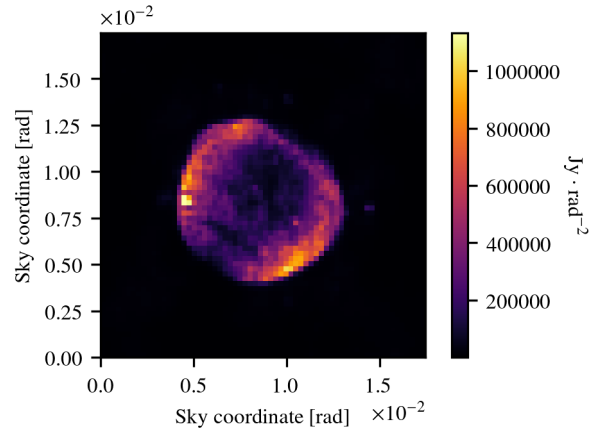


Fig. 18. SN1006: sampled posterior mean of sky brightness distribution.

is shown in Fig. 22. The power spectrum is more constrained compared to that of Sect. 3 since the noise level is much lower in this data set as compared to the synthetic data set. We might expect the posterior power spectrum to feature nodes or distinct minima because the Fourier transform of compact objects typically exhibit such. This is suppressed by the smoothness prior on the power spectrum. However, we note that this does not mean that the algorithm cannot reconstruct the object because it can still choose to not excite the respective modes in ξ_B .

All in all, this demonstrates that RESOLVE is not only able to operate on synthetic data but is actually capable of solving for the sky brightness distribution and the calibration terms at the same time for real data sets.

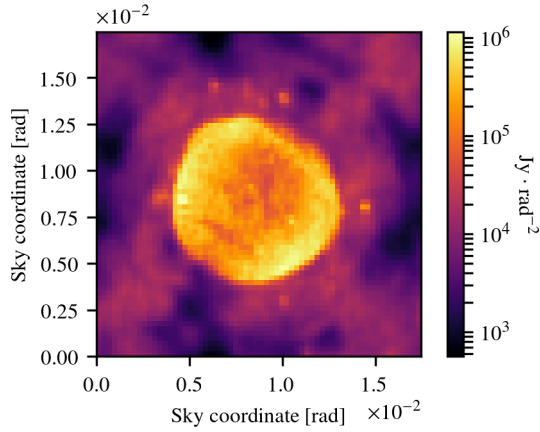


Fig. 19. SN1006: sampled posterior mean of sky brightness distribution (logarithmic color bar).

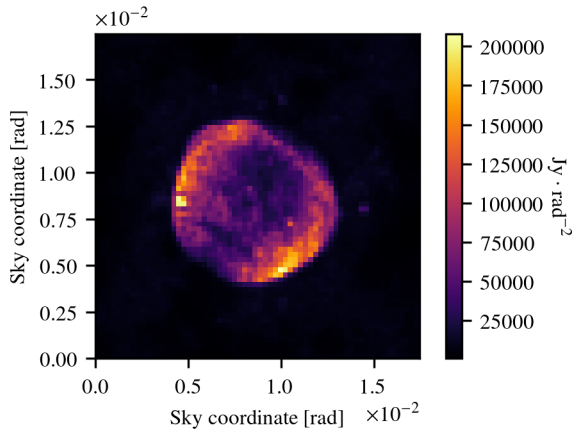


Fig. 20. SN1006: pixel-wise posterior standard deviation of sky brightness distribution.

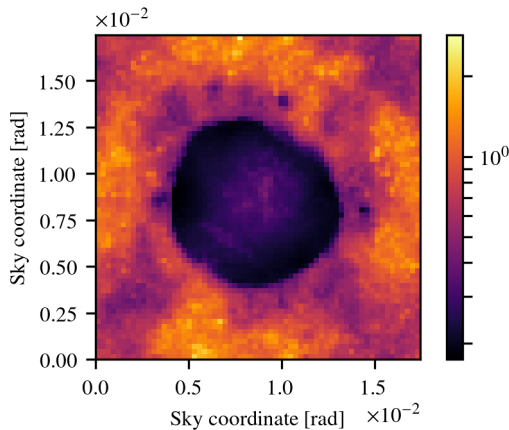


Fig. 21. SN1006: pixel-wise posterior standard deviation normalized by posterior mean of sky brightness distribution.

5. Performance and scalability

Performance and scalability are crucial aspects of the applicability of algorithms. The expensive part of the evaluation of the sky model is a fast Fourier transform (FFT), which is in $O(n \log n)$ where n is the total number of pixels of the sky model. For real-world data sets the cost for the (de)gridding exceeds the FFTs by far such that one likelihood evaluation is in $O(N)$, where N

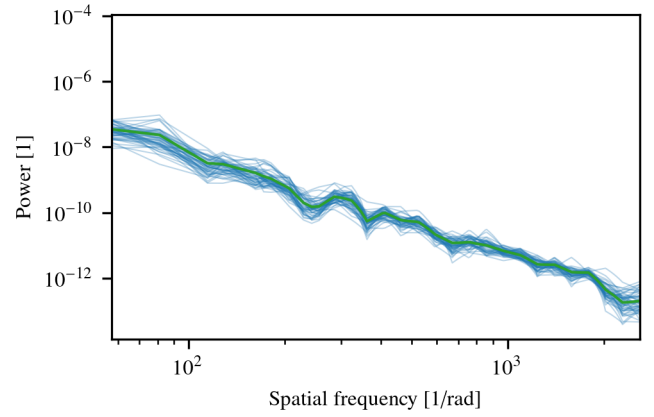


Fig. 22. SN1006: posterior power spectrum of logarithmic sky brightness distribution.

is the number of data points that need to be degridded once for each polarization. To compute the sampled KL divergence we need to compute the likelihood n_s times, where n_s is the number of samples (typically 3 – 20). The memory consumption scales linearly with the number of samples used to approximate the KL divergence, number of pixels, and number of data points. This is possible since NIFTy is designed such that no explicit matrices need to be stored.

Both reconstructions in this paper each took ≈ 60 minutes to be computed on a mobile CPU (Intel(R) Core(TM) i5-4258U CPU @ 2.40 GHz) with 4GB main memory. The response and adjoint needed to be called $\approx 30\,000$ times, respectively.

These values might improve in the future. Barnett et al. (2018) have proposed a novel gridding kernel that features speedups of several times in first experiments. This is possible since it needs relatively small support and can be computed on the fly. Also, the structure of the algorithm allows for various forms of parallelization. The gridding/degridding can be computed in parallel with OpenMP. Moreover, the data set could be split into several parts and distributed on a cluster. This is a general feature of Bayesian statistics: a likelihood can be split into the product of two likelihoods each of which contains only a subset of the data. Additionally, the evaluation of the KL divergence, which is a sum of few but expensive independent summands, can be distributed. Finally, NIFTy offers the (experimental) feature to distribute large fields on a cluster. Orthogonal to computational speedup ideas the algorithm might also benefit from compressing the likelihood itself such that fewer (de)gridding calls are necessary.

6. Conclusions

We have presented the probabilistic RESOLVE algorithm for simultaneous calibration and imaging. After a derivation from first principles of the full posterior probability distribution for the joint calibration and imaging algorithm RESOLVE, it has been shown how this distribution can be approximated by a multivariate Gaussian probability distribution to render the problem computationally solvable. This method is called MGVI and provides a prescription for how to draw samples from the approximate posterior distribution. The calibration algorithm RESOLVE has been verified on synthetic data. The results indicate that the uncertainty quantification is qualitatively sensible but should be taken with a grain of salt since MGVI systematically underestimates posterior variance. Furthermore, it has been demonstrated that the algorithm has the capability to reconstruct a sky brightness distribution of

a intricate source, the supernova remnant SN1006, together with uncertainty information from raw VLA *L*-band data.

There are many open ends to continue the investigation that we started with this paper. First, the model for the sky brightness distribution may include point source and multifrequency correlations. On top of that the response may be described more thoroughly. Direction-dependent calibration and nontrivial primary beam effects may be taken into account. Moreover, we performed the flagging by a standard CASA flagging algorithm. This can be replaced with an algorithm rooted in information theory that unifies flagging with calibration/imaging. Additionally, a major/minor cycle scheme similar to that in CLEAN may be introduced to avoid to frequent (de)gridding operations. This is necessary to apply RESOLVE to big data sets from telescopes such as MeerKAT. Finally, RESOLVE can be extended to polarization imaging. On an orthogonal track RESOLVE may be used for imaging of a variety of sources from different telescopes including ALMA and especially the Event Horizon Telescope.

Acknowledgements. We would like to thank Jamie Farnes for his workshop on calibration with CASA at the Power of Faraday Tomography 2018 conference and his script for calibrating the data set at hand with CASA, which enabled development of the IFT calibration algorithm. We thank Landman Bester, Ben Hugo, Jakob Knollmüller, Julian Rüstig, Oleg Smirnov, and Margret Westerkamp for numerous discussions, explanations, and feedback, and Martin Reinecke for his work on NIFTY, which was the crucial technical premise of the project. We acknowledge financial support by the German Federal Ministry of Education and Research (BMBF) under grant 05A17PB1 (Verbundprojekt D-MeerKAT). We thank the anonymous referee for insightful comments and the language editor for substantially improving the text quality.

References

- Arras, P., Baltac, M., Ensslin, T. A., et al. 2019, *Astrophysics Source Code Library* [record ascl:1903.008]
- Barnett, A. H., Magland, J. F., & Klinteberg, L. a. 2018, *ArXiv e-prints* [arXiv:1808.06736]
- Cai, X., Pereyra, M., & McEwen, J. D. 2018, *MNRAS*, 480, 4154
- Enßlin, T. A. 2018, *Ann. Phys.*, 1800127
- Enßlin, T. A., Frommert, M., & Kitaura, F. S. 2009, *Phys. Rev. D*, 80, 105005
- Hamaker, J., Bregman, J., & Sault, R. 1996, *A&AS*, 117, 137
- Högbom, J. 1974, *A&AS*, 15, 417
- Hutschenreuter, S., & Enßlin, T. A. 2019, *A&A*, submitted [arXiv:1903.06735]
- Junklewitz, H., Bell, M., Selig, M., & Enßlin, T. A. 2016, *A&A*, 586, A76
- Keiner, J., Kunis, S., & Potts, D. 2009, *ACM Trans. Math. Softw. (TOMS)*, 36, 19
- Kenyon, J., Smirnov, O., Grobler, T., & Perkins, S. 2018, *MNRAS*, 478, 2399
- Khintchin, A. 1934, *Math. Ann.*, 109, 604
- Knollmüller, J., & Enßlin, T. A. 2018, *ArXiv e-prints* [arXiv:1812.04403]
- Knollmüller, J., & Enßlin, T. A. 2019, *ArXiv e-prints* [arXiv:1901.11033]
- Knollmüller, J., Frank, P., & Enßlin, T. A. 2018, *ArXiv e-prints* [arXiv:1804.05591]
- Leike, R., & Enßlin, T. 2019, *A&A*, submitted [arXiv:1901.05971]
- Noordam, J. E., & Smirnov, O. M. 2010, *A&A*, 524, A61
- Offringa, A., McKinley, B., Hurley-Walker, N., et al. 2014, *MNRAS*, 444, 606
- Perkins, S., Marais, P., Zwart, J., et al. 2015, *Astron. Comput.*, 12, 73
- Salvini, S., & Wijnholds, S. J. 2014, *A&A*, 571, A97
- Schwab, F. 1984, *AJ*, 89, 1076
- Smirnov, O. M. 2011, *A&A*, 527, A106
- Thompson, A. R., Moran, J. M., Swenson, G. W., et al. 1986, *Interferometry and Synthesis in Radio Astronomy* (Springer)
- Wiener, N. 1949, *Extrapolation, Interpolation, and Smoothing of Stationary Time Series* (Cambridge: MIT Press), 2