

PolyCLEAN: Atomic optimization for super-resolution imaging and uncertainty estimation in radio interferometry

Adrian Jarret^{1,*}, Sepand Kashani¹, Joan Rué-Queral², Paul Hurley^{3,4}, Julien Fageot¹, and Matthieu Simeoni²

¹ Audiovisual Communications Laboratory, École Polytechnique Fédérale de Lausanne, Switzerland

² Center for Imaging, École Polytechnique Fédérale de Lausanne, Switzerland

³ Centre For Research In Mathematics and Data Science, Western Sydney University, Australia

⁴ International Centre for Neuromorphic Systems, Western Sydney University, Australia

Received 3 June 2024 / Accepted 10 November 2024

ABSTRACT

Context. Imaging in radio interferometry requires solving an ill-posed noisy inverse problem, for which the most adopted algorithm is the original CLEAN algorithm and its variants. Alternative explicit optimization methods have gained increasing attention, as they demonstrate excellent reconstruction quality thanks to their ability to enforce Bayesian priors. Nowadays, the main limitation to their adoption is run-time speed. Additionally, uncertainty quantification is difficult for both CLEAN and convex optimization techniques.

Aims. We address two issues for the adoption of convex optimization in radio interferometric imaging. First, we propose a method for a fine resolution setup, which scales naturally in terms of memory usage and reconstruction speed. Second, we develop a new tool to localize a region of uncertainty, paving the way for quantitative imaging in radio interferometry.

Methods. The classical ℓ_1 penalty is used to turn the inverse problem into a sparsity-promoting optimization. For efficient implementation, the so-called Frank-Wolfe algorithm is used together with a polyatomic refinement. The algorithm naturally produces sparse images at each iteration, leveraged to reduce memory and computational requirements. In that regard, PolyCLEAN reproduces the numerical behavior of CLEAN, while guaranteeing that it solves the minimization problem of interest. Additionally, we introduce the concept of the dual certificate image, which appears as a numerical byproduct of the Frank-Wolfe algorithm. This image is proposed as a tool for uncertainty quantification on the location of the recovered sources.

Results. PolyCLEAN demonstrates good scalability performance, in particular for fine-resolution grids. On simulations, the Python-based implementation is competitive with the fast numerically-optimized CLEAN solver. This acceleration does not affect image reconstruction quality: PolyCLEAN images are consistent with CLEAN-obtained ones for both point sources and diffuse emission recovery. We also highlight PolyCLEAN reconstruction capabilities on observed radio measurements.

Conclusions. PolyCLEAN can be considered as an alternative to CLEAN in the radio interferometric imaging pipeline, as it enables the use of Bayesian priors without impacting the scalability and numerical performance of the imaging method.

Key words. instrumentation: interferometers – methods: numerical – techniques: interferometric

1. Introduction

Reconstructing an image in radio interferometry amounts to solving an inverse problem, whose dimensions are determined by the number of interferometer baselines and the number of pixels used to cover the field of view. Increased observation and longer baselines enhance the quality of reconstructed images, both in terms of sensitivity and resolution, but require significantly more processing time. In the era of the Square Kilometer Array (SKA), the next generation radio interferometer aiming to contain hundreds of stations and thousands of antennas (Scaife 2020), this increased complexity causes considerable numerical challenges.

In particular, it is crucial for the imaging algorithms to be fast, reliable, and scalable. The most simple way to obtain a meaningful image consists of constructing the dirty image with an elementary inversion of the observed visibility. Although it is cheap to produce (as it only requires a single call to the adjoint of the measurement operator), the dirty image presents reconstruction artifacts due to the partial coverage of visibility space and, thus, it is rarely satisfactory. Current imaging methods are

more advanced and improve the quality of the reconstruction significantly. They can be classified into two families: adhering either to the CLEAN or the Bayesian paradigm. These methods introduce a prior image model, usually built around a notion of sparsity in the image to account for ill-posedness. However, the two families differ in terms of enforcing the model.

The CLEAN realm. CLEAN traces its origins back to 1974 with the Högbom algorithm (Högbom 1974), predating sparse reconstruction methods developed in the signal processing community by about two decades (see the historical review in (Foucart & Rauhut 2013, Notes of Chapter 1) and references therein). As one of the first sparsity-promoting methods, this pioneering algorithm has had huge impacts in radio astronomy and beyond (Cornwell 2009). CLEAN creates a model for the reconstructed image by iteratively placing elementary components called atoms. These atoms are point sources in the case of the original Högbom algorithm, or building elements with a spatial extension in the case of multi-scale reconstruction for diffuse emissions, as in Asp-CLEAN (Bhatnagar & Cornwell 2004) and MS-CLEAN (Cornwell 2008). Many developments have followed, creating a wide landscape of CLEAN-based image reconstruction techniques, improving the image model and numerical

* Corresponding author; adrian.jarret@epfl.ch

performance (Clark 1980; Schwab 1984; Offringa et al. 2014; Offringa & Smirnov 2017; Müller & Lobanov 2023b). The family of CLEAN algorithms is based on a prior model that assumes that the image to recover is sparse; that is, it can be accurately approximated with a limited number of atoms. This sparse prior model is enforced empirically along the reconstruction by controlling the number of iterations.

CLEAN algorithms are considered fast and scalable due to the low complexity of the operations performed. The long-term development efforts proposed since the first creation of CLEAN have resulted in mature software, with fast numerically optimized implementations available – see WSCLEAN (Offringa et al. 2014). This maturity is one of the reasons CLEAN-based methods remain predominant in radio astronomy applications today.

The Bayesian paradigm. The other grand family of imaging methods used in radio interferometry relies on the Bayesian statistical inference framework. There, the prior image model is explicitly described in a mathematical manner using a probability distribution. Different image estimators then lead to various reconstruction algorithms covering a wide range of scalability properties. A trade-off between the completeness of the information recovered and numerical performance can usually be readily obtained. On one end of the spectrum are methods that estimate the full posterior distribution of the data, such as Junklewitz et al. (2016); Cai et al. (2018b); Arras et al. (2021). Although these techniques recover extensive and precise information, they are numerically heavy, requiring the use of Monte-Carlo Markov chains or Langevin sampling algorithms. For this reason, they are difficult to scale when the dimensions of the problem increase. On the other end of the information-scalability trade-off are the Maximum A Posteriori estimators, proposed to the radio interferometry more than a decade ago (Wiaux et al. 2009; Li et al. 2011). These estimators turn out to be more practical as they generally amount to solving a sparsity-promoting optimization problem. Although they are, on the whole, numerically heavier to implement than CLEAN-based techniques, the runtime can still be reasonable. They are able to enforce complex priors, matching the versatility of CLEAN by recovering both point sources and diffuse emissions, and they have demonstrated excellent reconstruction quality (Schwardt 2012; Carrillo et al. 2012; Garsden et al. 2015; Chael et al. 2016; Akiyama et al. 2017; Abdulaziz et al. 2019; Müller & Lobanov 2023a). However, the optimization methods generally rely on proximal algorithms, which pose other concerns; thus limiting their overall adoption. On the one hand, proximal algorithms can be difficult to scale to large datasets. Indeed, they compute dense gradient steps at each iteration, which is not in line with the end goal of structurally sparse solutions. This mismatch leads to numerical overheads resulting in longer and less precise convergence. Some promising efforts have been made to scale out these algorithms, for instance, with the faceted distributed algorithm in Thouvenin et al. (2023a). On the other hand, most of the current imaging pipeline, including the calibration procedure, are built around atomic CLEAN-like strategies. Additional effort is thus required to integrate proximal algorithms with a different algorithmic structure into the calibration process.

It should be noted that learning approaches have also been proposed for radio interferometric imaging, including end-to-end learning (Schmidt et al. 2022), reconstruction in the image domain (Dabbech et al. 2024), and plug-and-play algorithms (Terris et al. 2023). These approaches essentially learn a prior image model from training data. Contrary to the aforementioned

families, this model is neither explicit nor fully understood. Thus, trust issues may arise when analyzing reconstructions. Deep learning models are indeed known to sometimes create artifacts, referred to as hallucinations in the machine learning community (see the discussion in Gottschling et al. 2023). As in other computational imaging applications, the learning methods offer exciting opportunities for radio interferometric imaging, together with their usual limitations. The inference task is usually performed quickly and in a scalable distributed manner, but the learning phase may be more problematic, both in terms of time and computing power. Due to the specificity of the learning methods, namely, that each application needs a bespoke trained model, learning methods have not been considered in our scalability study.

Image resolution as a quality-efficiency trade-off. Irrespective of the algorithmic framework, the reconstruction time and memory requirements are significantly impacted by the choice of pixel size. To cover the same field of view, a higher resolution uses smaller pixels and thus increases image size, degrading numerical performance. In radio interferometric imaging, the model image produced is generally convolved with the CLEAN beam, whose angular resolution corresponds to the main lobe of the point-spread function of the interferometer. For this reason, it is common to choose a coarse pixel size, at about the nominal resolution of the interferometer (defined by the diffraction limit). However, the above-mentioned sparsity-based methods, including CLEAN, actually perform nonlinear reconstruction; thus, they are able to recover information at a scale finer than this nominal resolution. Super-resolved reconstruction techniques have been proposed, both with optimization methods (Honma et al. 2014; Dabbech et al. 2018; Akiyama et al. 2017) as well as with CLEAN (Chael et al. 2016; Fried 1995). In that regard, it is relevant to develop methods that employ finer pixels than the nominal resolution, as they allow for the recovery of more accurate images from the same measurement set, at the cost of more compute time. The pixel dimension should ideally be set to the smallest recoverable scale. To take full advantage of this finer resolution, the model image should either be convolved with a representative beam sharper than the CLEAN beam or presented without convolution at all.

Contributions. PolyCLEAN, the algorithmic framework presented, takes advantage of the strengths of both CLEAN and proximal optimization families to perform scalable and fast imaging in a sub-nominal resolution setup. PolyCLEAN is based on an optimization problem and has been developed with a particular attention to numerical performance. In this regard, PolyCLEAN represents a step toward a more practical and trustful use of the optimization methods in radio interferometry imaging pipelines. The contributions are twofold. First, a Frank-Wolfe algorithm is used to solve the optimization problem of interest. The algorithm applies an atomic paradigm, in a CLEAN-like manner, which significantly differs from the proximal methods usually employed. This algorithmic structure, based on the use of sparse iterates, considerably improves the scalability of the reconstruction method. It combines a recent Polyatomic variant of the classical Frank-Wolfe algorithm (Jarret et al. 2022) with a sparsity-aware implementation of the forward operator of radio interferometry (Kashani et al. 2023). Second, the dual certificate image comes as a byproduct of the optimization problem and can be seen as a step toward quantitative imaging in radio interferometry (uncertainty and error quantification).

Section 2 presents the astronomical data model, along with a summary drawn from existing literature in radio interferometry to provide a comprehensive understanding of PolyCLEAN. Section 3 describes the CLEAN algorithm using the notation from the preceding section, with emphasis on the advantages and disadvantages of the algorithm. In a similar fashion, Sect. 4 introduces the mathematical statement of optimization methods through the prism of Bayesian inference, with a particular focus on the LASSO problem used as the recovery model for PolyCLEAN. This section also mathematically introduces the dual certificate image, which is a direct consequence of the convex analysis of the LASSO problem. This selected overview of radio interferometric imaging knowledge allows us to introduce unified notation and lays the groundwork for the detailed presentation of PolyCLEAN in Sect. 5. Section 6 reports on the numerical experiments to demonstrate the scalability of PolyCLEAN as well as its reconstruction capabilities, both on simulated and observed measurement sets.

2. Interferometric data model

Measurement equations model the data acquisition in radio interferometry. We first recall a physical measurement equation in Sect. 2.1, then derive a discretized version of this equation in Sect. 2.2, leading to the inverse problem at the heart of digital imaging in radio astronomy.

2.1. Physical model with a continuous sky

In radio interferometry, the power flux density, measured in Jansky (Jy), is the power of the received electric field. Sky images are produced by estimating the intensity or brightness which corresponds to the power flux density per unit solid angle. It is represented by a nonnegative function I defined over the celestial sphere $\mathbb{S}^2$ (Thompson et al. 2017, Sect. 1.2.1). For simplicity, equations throughout this work assume calibrated data corrected from the perturbations of both the instrumental effects and the propagation effect (Sect. 5.1 of Van der Veen et al. 2019).

Image model. The point source model assumes celestial objects emit from a single point without any spatial extension. In this model, the brightness function I is represented as a sum of Dirac impulses δ where the weight is the intensity of the source (in Jansky), so that the sky image can be modeled as

$$I(\ell, m) = \sum_i \alpha_i \delta(\ell - \ell_i, m - m_i),$$

for some weights $\alpha_i > 0$ and source locations $(\ell_i, m_i) \in [-1, 1]$. The resulting image is sparse (mostly empty with a limited number of nonzero-valued locations).

More general models include diffuse emissions, which are cosmic objects producing emissions over a spatial extent. In this case, images are generally non-sparse. Although point sources form a simpler model and are relevant for large fields of view, it is often necessary to consider extended sources as soon as the field of view is zoomed-in and narrower.

Measurement equation. The brightness distribution is sampled by means of an interferometer, which is an array of radio telescopes located on the Earth's surface. The action of the interferometer is modeled by the following linear measurement equation

$$\mathbf{V} = \Phi(I) + \boldsymbol{\varepsilon}, \quad (1)$$

where $\mathbf{V} \in \mathbb{C}^L$ is the vector of the visibility measurements and the linear functional Φ models the action of the interferometer. Furthermore, $\boldsymbol{\varepsilon} \in \mathbb{C}^L$ is the realization of a complex-valued random variable to account for instrumental noise, which will be characterized later on.

Imaging is performed by solving the inverse problem posed in Eq. (1) to find a solution, I_{sol} , that can be seen as an approximation of the actual sky. The measurement equation serves as the guiding principle of the imaging process.

Forward operator. The forward operator Φ samples the continuous visibility function $\mathcal{V}$ at the location of the baselines of the interferometer $(\mathbf{p}_k)_{k=1, \dots, L}$. These baselines are defined as the spatial difference between the receiving antennas and given in units of the observed wavelength¹, while the visibility function has the appearance of a 3D Fourier Transform with integration performed over a sphere:

$$\mathcal{V}(\mathbf{p}) = \iint_{\mathbb{S}^2} I(\mathbf{r}) e^{-j2\pi(\mathbf{r}, \mathbf{p})} d\mathbf{r}, \quad \mathbf{p} \in \mathbb{R}^3.$$

We obtain the general radio interferometry equation² (Thompson et al. 2017, Chapter 3)

$$\Phi(I)[k] = \mathcal{V}(\mathbf{p}_k) = \iint_{\mathbb{S}^2} I(\mathbf{r}) e^{-j2\pi(\mathbf{r}, \mathbf{p}_k)} d\mathbf{r} \quad (2)$$

for $k \in \{1, \dots, L\}$.

It is useful to rewrite Eq. (2) using the tangent plane coordinates around the direction of observation. Let us introduce the direction cosine coordinates $\mathbf{r} = (\ell, m, n(\ell, m)) \in \mathbb{S}^2$ where $\ell, m \in [-1, 1]$ and $n(\ell, m) = \sqrt{1 - \ell^2 - m^2}$. The measurement Equation (2) then becomes

$$\Phi(I)[k] = \iint_{[-1, 1]^2} \frac{I(\ell, m)}{n(\ell, m)} e^{-j2\pi(u_k \ell + v_k m)} \mathcal{W}(\ell, m; w_k) d\ell dm, \quad (3)$$

for $k \in \{1, \dots, L\}$, with

$$\mathcal{W}(\ell, m; w) = e^{-j2\pi w(n(\ell, m) - 1)}$$

and (u_k, v_k, w_k) as the wavelength-normalized ground coordinates of the baseline $\mathbf{p}_k$.

This forward operator corresponds to the natural weighting of the visibilities: each measurement accounts for the same weight in the reconstruction procedure. Using uniform or robust weighting schemes is not relevant with our formalism as the inverse problem is solved from the measurement equation and is independent of manual engineering of the PSF (see Rau et al. (2009); Onose et al. (2017) for more details on the weighting schemes).

Noise model. Although the observation noise can be modeled in a more accurate manner (see e.g., Simeoni & Hurley 2021), in this work, we consider a simplified noise model, as standard in radio interferometry. We assume the additive model of Eq. (1) with $\boldsymbol{\varepsilon}$ distributed according to a complex-valued Gaussian white noise.

¹ Note: this equation assumes the observation bandwidth to be small compared to the center frequency of observation.

² When I contains point sources the notation of the integrals is mathematically incorrect and is to be understood as $\iint_{\mathbb{S}^2} \delta(\mathbf{r}) d\mathbf{r} = \iint_{\mathbb{S}^2} d\delta(\mathbf{r})$.

Remark 1 (Van Cittert-Zernike). When the field of view is narrow (small values of ℓ and m) or when the baselines $\mathbf{p}_k$ are considered co-planar (coordinate $w = 0$), Eq. (3) reduces to the well known Van Cittert-Zernike equation (Thompson et al. 2017, Chapter 15). However, for larger field of views, it is crucial to consider the $\mathcal{W}$ -term in the definition of the visibility function.

Remark 2 (Well-posedness and field of view). The expressions in Eqs. (2) and (3) are valid as long as the integrals are well-defined. In practice, the radio antennas have a limited collection area, which induces an antenna primary beam of sensitivity. The beam pattern is maximum toward the zenith and rapidly decays with the angle of observation. A consequence is that antennas naturally limit the width of observation to a limited surface of the celestial sphere. To simplify notations and without loss of generality, it is assumed that the sky images I considered in the equations are supported on the observable field of view of the antennas, centered around the direction of observation. The integrals of Eqs. (2) and (3) then sum over a compact support and are well-posed.

The physical inverse problem defined by Eq. (1) deals with sky images defined on a continuous space domain. Some methods have been proposed to directly find continuous-domain solutions without resorting to explicit discretization. An example is the Bluebird algorithm (Kashani 2017), which splits radio interferometric images into different level sets of energy, or LEAP (Pan et al. 2017), which is aimed at placing point sources at arbitrary locations in a grid-free manner. Other options are available for working on the continuous-domain, for instance, by using a grid-based dictionary reconstruction, where the dictionary atoms are continuous functions but the locations of the center of these elements lie on a discrete grid (see examples of such techniques in Simeoni 2020). In this work, we classically choose to pixelize the sky images to obtain a fully discrete, finite-dimensional inverse problem. The next section describes the discretization scheme employed.

2.2. Measurement equation for digital images

In practical applications, it is common to reconstruct a 2D raster image of a given portion of the sky. The field of view coordinates $(\ell, m) \in [-1, 1]^2$ in Eq. (3) are usually discretized with a uniform fine grid $(\ell_i, m_j)_{i,j=1,\dots,n}$ of size $N = n \times n$, and the sky image is seen as a real-valued square image represented as a matrix $\mathbf{I} \in \mathbb{R}^{n \times n} = \mathbb{R}^N$. The value of a pixel sums the contributions of all the sources that lie on the area it covers, integrating both point and extended sources, divided by the angle area covered by a pixel. Each pixel is a brightness or intensity measure, whose value is in Jansky per pixel (Jy.pix^{-1}). With this model in hand, an approximate measurement equation is derived from Eq. (3) using the finite sums

$$\Phi(\mathbf{I})[k] = \sum_{i=1}^n \sum_{j=1}^n \frac{\mathbf{I}[i, j]}{n(\ell_i, m_j)} e^{-j2\pi(u_k \ell_i + v_k m_j)} \mathcal{W}(\ell_i, m_j; w_k), \quad (4)$$

for $k \in \{1, \dots, L\}$. Here, Φ is the linear analysis operator responsible for the degridging operation.

The numerical imaging methods are then based on the discrete counterpart of Eq. (1), which we can express using the above-defined linear measurement operator $\Phi : \mathbb{R}^{n \times n} \rightarrow \mathbb{C}^L$

$$\mathbf{V} = \Phi(\mathbf{I}) + \boldsymbol{\varepsilon}. \quad (5)$$

This discrete model is classical in radio interferometry. The forward operator Φ is usually defined as a sequence of linear

transformations involving the discretization of the images over a Cartesian grid and a Fourier transform (Wiaux et al. 2009; Rau et al. 2009).

Since it is useful for later stages of this study, we also state the expression of the adjoint operator $\Phi^* : \mathbb{C}^L \rightarrow \mathbb{R}^{n \times n}$. For any complex-valued measurement vector $\mathbf{h} \in \mathbb{C}^L$, we have

$$\Phi^*(\mathbf{h})[i, j] = \Re \left\{ \sum_{k=1}^L \frac{h_k}{n(\ell_i, m_j)} e^{j2\pi(u_k \ell_i + v_k m_j)} \mathcal{W}^*(\ell_i, m_j; w_k) \right\}, \quad (6)$$

for $i, j \in \{1, \dots, n\}$, where $\Re$ is the real part of a complex number. The operator Φ^* handles the gridding of the visibility: it performs the direct synthesis of an image based on the measurements. Under natural weighting the dirty image is computed as $\Phi^* \mathbf{V}$, where $\mathbf{V}$ is the vector of measured visibilities. Remarkably, Eqs. (4) and (6) can be seen as instances of a 2D non-uniform Fourier transform (3D if the $\mathcal{W}$ -term needs to be considered), which allows for a fast and scalable implementation to be developed (see more details in Sect. 5.2).

In the remainder of this article, we will stick with the discrete formulation of the sky image and the associated discrete forward operator.

Remark 3 (Cartesian grid discretization). As a mathematical object, the continuous-domain sky image $I(\ell, m)$ lives on a given area on the surface of the sphere $\mathbb{S}^2$. To represent it digitally, with a finite set of values, we need to use a tessellation of the support of the image. The classical choice of a discrete Cartesian grid – deployed here – is motivated by the use of the FFT algorithm in order to evaluate the forward operator, but other tessellations could be considered. The main limitation of a uniform Cartesian grid is that it introduces distortion effects due to the surface of the sphere represented by each pixel not being uniform. It leads to aberrations that are particularly visible toward the borders of the image for large fields of view.

Similar to its continuous counterpart, the discrete inverse problem of Eq. (5) is still ill-posed and corrupted by noise. It is usually assumed that the interesting solutions respect some properties such as positivity or sparsity, that help regularize the problem. It results in a prior model over the search space that may be explicit or not. Several methods are then available to identify a candidate solution using different ways to enforce their respective prior. In the following section, we express the classical CLEAN algorithm in the notation of this section.

3. Imaging with CLEAN

The CLEAN algorithm remains extremely popular in the radio astronomy community, with new versions and refinements being regularly proposed. From a signal processing point of view, CLEAN is an instance of a matching pursuit algorithm (Lannes et al. 1997; Mallat & Zhang 1993), which it predates by two decades. It fits a parametric signal model to the radio interferometry measurement equation while enforcing sparsity in the reconstruction coefficients.

Parametric model. In the initial algorithm, referred to as Högbom CLEAN (Högbom 1974), the unknown discrete image is assumed to be a collection of point sources. This amounts to looking for a solution of the form

$$\mathbf{I}_{\text{sol}} = \sum_k \alpha_k \delta_{(i_k, j_k)} \in \mathbb{R}^{n \times n}, \quad (\text{PS})$$

with $\delta_{(i,j)}$ being an almost empty image with only the pixel (i, j) at a value of 1. The number of point sources in the image is assumed to be small compared to the dimensions of the image, which results in few nonzero α_k coefficients and a sparse solution overall.

Later versions of CLEAN considered images expressed as combinations of elements from redundant reconstruction dictionaries, leading to more complex image models. For instance, multi-scale CLEAN (MS-CLEAN, Cornwell 2008) uses a dictionary of isotropic reconstruction atoms with different widths:

$$\mathbf{I}_{\text{sol}} = \sum_k \alpha_k \mathbf{g}_{(i_k, j_k)}^{\sigma_k}, \quad (\text{MS})$$

with $\mathbf{g}_{(i,j)}^{\sigma}$ being a discretized kernel with center on the pixel (i, j) parameterized with a spatial extension σ . For instance, $\mathbf{g}_{(i,j)}^{\sigma}$ can be a Gaussian kernel, but other functions have also been used such as the prolate spheroidal functions (which have the advantage of being compactly supported). In such a model, the sparsity is still enforced on the coefficients α_k , leading to a solution image that is usually not sparse in the image domain.

The algorithm. Matching pursuit algorithms construct images by iteratively searching for the dictionary element having the highest correlation with the residual measurements. This dictionary element is then added to the current estimate of the solution, scaled by a multiplicative gain factor. The reconstruction is stopped when the current iterate is considered good enough, for example when the residual noise is smaller than a threshold or after a given number of iterations (Offringa et al. 2014).

For the point source image model of Eq. (PS), the reference CLEAN algorithm is Cotton-Schwab CLEAN (Schwab 1984) that reads as Algorithm 1, where the reconstruction parameters are the number of iterations k_{max} and the gain α .

Algorithm 1 Cotton-Schwab CLEAN (Major cycles only)

Parameters : k_{max} (iterations), $\alpha > 0$ (gain)

Initialization : $\mathbf{I}^{(0)} = \mathbf{0}$, $\mathbf{I}_D = \Phi^* \mathbf{V}$

for $k = 1, 2, \dots, k_{\text{max}}$ **do**

1. Compute the dirty residual: $\mathbf{I}_R^{(k)} = \mathbf{I}_D - \Phi^* \Phi \mathbf{I}^{(k-1)}$
2. Find the location of the next reconstructed source:

$$s^{(k)} = \arg \max_{(i,j)} |\mathbf{I}_R^{(k)}[i, j]|$$

3. Update the iterate: $\mathbf{I}^{(k)} = \mathbf{I}^{(k-1)} + \alpha (\max \mathbf{I}_R^{(k)}) \delta_{s^{(k)}}$

Output:

Postprocess $\mathbf{I}^{(k)}$ (convolution with synthetic beam, add residual image)

Minor cycles. The most demanding step in terms of memory and computation power is the application of the forward and adjoint operators in step 1 to compute the dirty residual image. To accelerate this step, the composed operator $\Phi^* \Phi$ is approximated as a convolution with the kernel $\mathbf{B}_{\Phi}$, so-called the point spread function (PSF) of Φ and referred to as dirty beam. The cost of applying $\Phi^* \Phi$ is then reduced to simply performing a convolution, which is significantly faster than calling the exact operators but results in an approximation error. Indeed, observing a point source far from the center of the image induces a

change of intensity in addition to the spatial shift. The error is small for narrow fields of view and increases with the width of the observation window. These approximate steps are called minor cycles, which are performed for most of the iterations of CLEAN. The major cycles, with the exact computation of the operator, are performed every few hundred (or thousand) minor cycles to recompute the dirty residual with greater accuracy.

Advantages and limitations. The family of CLEAN-based algorithms has reached its popularity in radio astronomy as these methods are able to meet many of its imaging constraints. Later on, CLEAN is used as a generic term to refer to any version, whether it is Cotton-Schwab for point sources or multi-scale variants.

- ✓ Thanks to the convolution-based minor cycles, CLEAN is fast and performs few expensive calls to the measurement operator. In addition, the iterates admit a sparse representation, convenient to store and manipulate.
- ✓ CLEAN-like algorithms are flexible with respect to the image to reconstruct. The matching pursuit structure allows us to generalize the reconstruction to diffuse emissions using reconstruction dictionaries.
- ✓ CLEAN is well-suited for calibration. Indeed, the visibilities are calibrated by reconstructing a few bright and well-known point sources, which is conveniently handled with an atomic method like CLEAN.
- ✓ Last but not least, CLEAN-based algorithms have been present for a long time in the field of radio interferometric reconstruction. They thus benefit from extended development, including efficient and optimized implementations of the algorithms, and expertise from the researchers on how to use them to their full capacity.

Despite these advantages, CLEAN-based algorithms also come with a number of limitations, which we summarize below:

- ✗ CLEAN, as a matching pursuit algorithm, does not perform denoising. The effect of noise in the measurement is only handled by enforcing sparsity in the solution.
- ✗ As a consequence and because matching pursuit tends to explain all the measurements by canceling out the residuals, CLEAN needs to be stopped before complete convergence. The resulting solution is then highly sensitive to this arbitrary choice of stopping criterion.
- ✗ CLEAN often reconstructs regions with negative flux values that are physically unrealizable – as for instance reported in Arras et al. (2021). This may result from an initial weight allocation that is too greedy or tentative to explain the noise.
- ✗ Early termination of the algorithm prevents CLEAN from converging toward the final fit of matching pursuit. For this reason, the CLEAN solution can only be uncertainly interpreted as a sparse intermediate solution to a least-square problem, without any further guarantee of sparsity or recovery. See Schwarz (1978); Solo (2008); Lannes et al. (1997) for more details on the convergence of CLEAN and Locatello et al. (2018) on matching pursuit.
- ✗ CLEAN does not provide uncertainty quantification on the solution.

In summary, the CLEAN algorithm and its refinements form a set of efficient tools for radio interferometric image reconstruction, designed and powered by decades of research in the field. However, by essence, they face the same limitations as the matching pursuit algorithm for sparse recovery in its respective field of compressive sensing, that include robustness to noise and sparsity guarantees. In that regard, it is relevant to explore in radio interferometry the convex optimization methods

that are generally used in place of matching pursuit algorithms when sparse reconstruction is involved. This line of research has already proven to be fruitful, with excellent results in terms of reconstruction quality (see, e.g., [Thouvenin et al. 2023b](#)). The next section presents the Bayesian inference framework that naturally leads to the optimization problems the PolyCLEAN method is based upon.

4. Imaging with sparsity promoting Bayesian estimators

Bayesian theory provides arguably the most principled way to introduce prior information into a reconstruction problem. It relies on a user-defined prior probability distribution over the search space. Different estimators can then be used to reconstruct candidate solutions from the posterior distribution. For PolyCLEAN, we use the sparsity-promoting Laplace distribution prior and compute the maximum a posteriori (MAP) estimator, which amounts to solving a LASSO optimization problem. Here, we review how this problem appears and why it is well-suited to radio interferometric imaging.

4.1. From Bayes posterior density to LASSO: MAP

The Bayesian statistical inference framework makes the assumption that all components of our problem, namely the visibility measurements $\mathbf{V}$ and the source sky image $\mathbf{I}$, are realizations of random variables. A so-called posterior distribution of the image $\mathbf{I}$ given the observations $\mathbf{V}$ is derived, denoted by $p(\mathbf{I}|\mathbf{V})$. Due to Bayes theorem, this quantity is proportional to the product of a likelihood term $p(\mathbf{V}|\mathbf{I})$ and a prior distribution term $p(\mathbf{I})$. The likelihood is derived from the measurement Eq. (5), involving the forward operator and the measurement noise model. Regarding a prior distribution, it is common to promote sparsity in the reconstruction by means of a Laplace distribution, expressed as $p(\mathbf{I}) \propto e^{-b\|\mathbf{I}\|_1}$, where $b > 0$ is a scale parameter. We refer the reader to [Cai et al. \(2018b\)](#) for a thorough and more complete presentation of the Bayesian inference framework applied to radio interferometric imaging.

The obtained posterior distribution contains the recovered information which makes it an extremely insightful scientific tool, though possibly difficult to manipulate. In practice, this posterior distribution may be either demanding to approximate numerically or completely inaccessible. Estimators are then used to represent this distribution and produce reconstructions. In this work, we consider the maximum a posteriori (MAP) estimator, which is defined as an image for which the posterior density reaches one of its maxima. The simplicity of this estimator makes its computation conceptually simple, as it amounts to solving the following LASSO optimization problem ([Cai et al. 2018a](#)):

$$\arg \min_{\mathbf{I} \in \mathbb{R}^N} \frac{1}{2} \|\mathbf{V} - \Phi \mathbf{I}\|_2^2 + \lambda \|\mathbf{I}\|_1, \quad (7)$$

where $\lambda > 0$ is a regularization parameter. This problem is known to produce sparse solutions, namely, solutions with only a few nonzero coefficients. For this reason, and as already documented in the literature ([Wiaux et al. 2009](#); [Li et al. 2011](#); [Schwartz et al. 2012](#); [Garsden et al. 2015](#)), the LASSO problem is relevant for the reconstruction of astronomical images containing point sources. The solutions are not necessarily unique, leading to potentially many images consistent with the observed visibilities.

Remark 1 (Positivity constraint). Thanks to the flexibility of the Bayesian framework, a positivity constraint on the reconstructed

images is easily introduced by changing the prior distribution. Although the use of such a positivity constraint may be considered optional for the CLEAN reconstructions, it usually improves the quality of the images obtained using optimization methods ([Van der Veen et al. 2019](#), Sect. 4.5).

Parallel to Bayesian techniques, the LASSO problem is a classical tool in the compressed sensing community. The next section presents the grounding of LASSO in optimization theory, which motivates its use as a robust and principled model for sparse image reconstruction.

4.2. LASSO properties from optimization theory

The LASSO problem stands as a foundational method in the sparse reconstruction literature. Its name originated in the statistics community ([Tibshirani 1996](#)), also known as basis pursuit denoising in compressive sensing ([Chen et al. 2001](#)). Over the years, by means of convex optimization theory and dual analysis, many results have been derived to describe the solutions and the behavior of the problem, such as conditions for uniqueness ([Tibshirani 2013](#)) or recovery guarantees (see the discussion in the introduction of [Duval & Peyré 2015](#) for an overview). Let us focus on two notions that help to understand the relevance of the LASSO in radio interferometric imaging.

Representer theorem. The existence of sparse solutions to the LASSO problem is well-known (see for instance [Elad 2010](#)). Recent representer theorems provide a new interpretation of this property by describing the solution set of many convex optimization problems ([Boyer et al. 2019](#)). This new theory helps to build trust in the optimization-based reconstruction methods as their solutions are better understood and make physical sense. For the LASSO, the solution set is proven to be non-empty, compact and convex. The extreme points of this set are sparse elements, whose sparsity index is bounded by the number of measurements ([Unser et al. 2016](#)). In a context of high undersampling, as is the case in radio interferometry, this representer theorem ensures the existence of sparse solutions to the LASSO.

Dual certificate. Additionally, the LASSO problem benefits from the analysis of duality theory. For a given solution $\mathbf{I}_{\text{sol}}$ of Eq. (7), the so-called dual certificate³ μ_λ is defined as

$$\mu_\lambda = \frac{1}{\lambda} \Phi^*(\mathbf{V} - \Phi \mathbf{I}_{\text{sol}}). \quad (8)$$

While there are in general multiple solutions to the LASSO they all share the same measurements, that is the fit $\Phi \mathbf{I}_{\text{sol}}$ is itself unique and independent of $\mathbf{I}_{\text{sol}}$ ([Tibshirani 2013](#)). The dual certificate is thus unique and independent of the recovered solution.

This quantity can be used to extract information about all the possible solutions to the LASSO, given that at least one is known. Indeed, relying on the optimality conditions of the problem ([Duval & Peyré 2017](#), Proposition 1), the dual certificate satisfies the two following properties for any solution $\mathbf{I}_{\text{sol}}$:

$$\langle \mu_\lambda, \mathbf{I}_{\text{sol}} \rangle = \|\mathbf{I}_{\text{sol}}\|_1, \quad (9)$$

$$\|\mu_\lambda\|_\infty \leq 1. \quad (10)$$

³ The terminology is common for continuous-domain problems but the definition still holds for finite-dimensional problems ([Duval & Peyré 2017](#)).

Combined together, these properties imply that the support of any solution is contained in the so-called saturation set $\mathcal{D}_\lambda$ of the dual certificate, possibly empty, defined as the set of indices where μ_λ reaches 1 in absolute value

$$\mathcal{D}_\lambda = \{i \in \{1, \dots, N\} : |\mu_\lambda[i]| = 1\}. \quad (11)$$

This property is critical for a posteriori analysis of the reconstruction: given one solution to the LASSO, it is possible to compute the dual certificate and its saturation set and, hence, access the support of any other solution. This dual certificate takes the shape of an image, which makes it convenient to display, analyze and manipulate. We show an example in Sect. 6.3.

4.3. Versatility of the LASSO for diffuse emissions

The LASSO problem in its simple form reconstructs sparse images and thus is adapted for point source reconstruction. When diffuse emissions are involved images are not directly sparse, with lots of pixels expected to have a nonzero value. Nonetheless, the images may still be highly structured such that they can be represented by a limited number of components (and thus sparse in another domain). Such a representation can be obtained for example with wavelet transforms, which are known to produce sparse multi-scale component images. In particular, the Isotropic Undecimated Wavelet Transform is considered well-suited for representing radio astronomical images (Starck & Murtagh 2006). Using a modified LASSO problem, many methods have been developed to recover images with a sparse decomposition in a multi-scale wavelet frame (Carrillo et al. 2012; Garsden et al. 2015).

Another way to rely on sparsity for diffuse emission reconstruction is to assume that the sky image can be expressed (or approximated) with a small number of elements from a large dictionary of reconstruction atoms, which is again a form of sparsity. This is the approach used by MS-CLEAN, in which the dictionary elements are instances of prolate spheroidal functions with different scales (Cornwell 2008).

Whether a dictionary of atoms or a wavelet frame is used, the image to recover is expressed as

$$\mathbf{I} = \Psi\theta, \quad (12)$$

where Ψ is a linear synthesis operator that maps the decomposition coefficients $\theta \in \mathbb{R}^M$ to an actual image $\mathbf{I} \in \mathbb{R}^N$. The sparsity hypothesis now applies to θ . In WSCLEAN, the reconstruction coefficients are determined with an iterative greedy strategy. With PolyCLEAN, we instead solve the following LASSO problem, known as synthesis form:

$$\arg \min_{\theta \in \mathbb{R}^M} \frac{1}{2} \|\mathbf{V} - \Phi\Psi\theta\|_2^2 + \lambda \|\theta\|_1. \quad (13)$$

Remark 2 (Analysis form). It is possible to consider the problem in its so-called analysis form. The optimization variable is then $\mathbf{I}$ and the synthesis operator goes within the regularization term as the adjoint operator $\|\Psi^*\mathbf{I}\|_1$. In the general case, the problems are not equivalent and the analysis form can lead to better results (Elad et al. 2006; Tibshirani & Taylor 2011; Carrera et al. 2017). In Cai et al. (2018a), both analysis and synthesis forms demonstrated no major discernible difference. The PolyCLEAN framework relies on the synthesis form as it is more convenient to treat with a Frank-Wolfe algorithm (see Sect. 5).

4.4. Proximal solvers and their numerical limitations

LASSO, similarly to other optimization techniques in radio interferometry, has a non-differentiable objective function which cannot be minimized by simple gradient steps. These non-differentiable problems are generally solved with iterative Forward-Backward schemes, a class of algorithms that leverage the existence of a simple closed-form expression for the penalty term (Parikh & Boyd 2014). In particular, the LASSO problem is usually solved with the Accelerated Proximal Gradient Descent algorithm (APGD), also known as FISTA (Beck & Teboulle 2009). These proximal solvers are flexible and show fast convergence properties. For the LASSO, APGD is proven to have an optimal convergence speed in terms of number of iterations (Liang et al. 2022).

However, for sparse reconstruction, especially in high dimensional settings, the proximal methods may suffer from scalability limitations when compared to atomic methods such as CLEAN (Offringa & Smirnov 2017). Although the end solution is sparse, the iterates are usually dense with very little zero-valued elements. This structural mismatch between iterative method and the nature of the solution might be prejudicial for two reasons: (1) the algorithm may require a large number of iterations to sparsify the iterate, leading to a slow convergence; and (2) the solver needs to store the iterates in a dense format along the iterations, increasing memory requirements and affecting scalability.

In contrast, the PolyCLEAN method demonstrates an atomic CLEAN-like behavior, which allows sparse iterates to be maintained across iterations. In doing so, PolyCLEAN preserves the strengths of CLEAN presented in Sect. 3 while improving reconstruction thanks to the benefits of the optimization framework. The next section presents the algorithm and its scalability-oriented design.

5. The PolyCLEAN pipeline as an atomic imaging method

Let us summarize the three major bricks of the PolyCLEAN framework:

- LASSO – The imaging inverse problem of Eq. (5) is solved using the sparsity prior of the LASSO problem, namely using an ℓ_1 -norm regularizer. The dual certificate comes as a natural output of the LASSO and provides information on the localization of the reconstructed sources. PolyCLEAN can optionally enforce a positivity constraint on the image. PolyCLEAN also supports the LASSO problem formulated in synthesis form over a transformed domain, to perform sparse dictionary or sparse wavelet reconstruction (as described in Sect. 4.3).
- The Polyatomic Frank-Wolfe algorithm – The LASSO problem Eq. (7) is minimized by the Polyatomic Frank-Wolfe algorithm (PFW), recently developed by the authors (Jarret et al. 2022). This algorithm has been specifically designed to solve sparse inverse problems in large dimensional setups, which makes it relevant to radio astronomy applications.
- A sparsity-aware implementation of the interferometric forward operator – Independently of the solving method, applying the forward operator Φ is both time and memory consuming, as it is usually done via a (de-)gridding operation followed by a multidimensional fast Fourier transform. In PolyCLEAN we implement this operator using a Non-Uniform Fast Fourier Transform (NUFFT) algorithm. We rely on the HVOX package to do so, a chunking technique

for efficient computation of the NUFFT in Python (Kashani et al. 2023). The sparse structure of both the images and the frequency observations are the key ingredients in making computations faster and reducing memory requirements of the forward and adjoint operators.

With this structure, the benefits of PolyCLEAN are twofold. On the one hand, thanks to its CLEAN-like atomic behavior, PolyCLEAN provides a fast and scalable algorithm to perform sparsity-inducing Bayesian imaging, improving upon the proximal solvers in terms of computation time and memory requirements. The Polyatomic Frank-Wolfe algorithm and the HVOX-based forward operator both independently enable speed and scalability improvements. Their combination within PolyCLEAN results in an efficient and elegant way to build upon their respective computational strengths.

On the other hand, as PolyCLEAN is based on a LASSO problem, it is intrinsically connected to the LASSO dual certificate. The powerful theoretical guarantee of this dual certificate, that is the identification of the solution support, mean it can become an insightful tool for quantification of error in images. At each iteration, the Frank-Wolfe algorithm computes the empirical dual certificate that converges toward the final certificate. Although the dual certificate is a common tool in convex optimization theory, to the best of our knowledge this is the first time it is deployed for application purposes in radio interferometry. The following section presents the Polyatomic Frank-Wolfe algorithm, the NUFFT operator and the analysis of the dual certificate in the context of radio interferometry.

5.1. Solving with Polyatomic Frank-Wolfe

The backbone of PolyCLEAN consists in applying the Polyatomic Frank-Wolfe algorithm (PFW) recently proposed in Jarret et al. (2022) to solve the LASSO problem of interest.

For the sake of simplicity we limit discussion in this section to the case of point source reconstruction. When diffuse emissions are present, the interferometric problem is treated with the LASSO in synthesis form as in Eq. (13), which uses a modified forward operator.

Renewed interest in the Frank-Wolfe algorithms. The original Frank-Wolfe algorithm is a classical convex optimization algorithm (Frank & Wolfe 1956) that constructs a solution as a sum of atoms in a projection-free manner. This atomic behavior has been identified to produce a fast and scalable algorithm, particularly in the context of sparse reconstruction, leading to a revival of interest from the optimization community over the past decade (Jaggi 2013; Harchaoui et al. 2013). The Frank-Wolfe algorithm now benefits from many variants (Jaggi 2013; Lacoste-Julien & Jaggi 2015), and the connection has been drawn with other sparse reconstruction methods, including matching pursuit (Locatello et al. 2017). In particular, when Frank-Wolfe is applied to the LASSO problem (Jarret et al. 2022; Harchaoui et al. 2013), the algorithm demonstrates a striking numerical similarity with CLEAN. The reconstruction atoms are indeed point sources, identified by single pixel images. In that regard, Frank-Wolfe algorithms are good candidates to fit in a CLEAN numerical reconstruction pipeline, as they guarantee the convergence of the reconstruction toward a solution to the LASSO.

Benefits of PFW. The polyatomic variation of Frank-Wolfe is an additional refinement that allows the algorithm to place many atoms at each iteration instead of only one at a time. This

greatly reduces the solving time and is the key ingredient that brings a Frank-Wolfe algorithm up to speed with the proximal methods (Jarret et al. 2022). In a sense, this polyatomic strategy can be compared to the minor cycles in CLEAN. It is another way to populate the solution with many candidate atoms without evaluating the forward operator. PFW is proven to converge toward a minimum of the objective function of the LASSO. The name PolyCLEAN arises naturally from a portmanteau of a polyatomic update with a CLEAN-like algorithm.

Compared to APGD, PFW is more scalable and faster in the specific context of high sparsity. Indeed, as the iterates of PFW are sparse, their memory usage is limited, whereas APGD has dense iterates due to the gradient steps. Regarding the reconstruction speed, the theoretical rate of PFW is slower than APGD. However, numerical applications demonstrate significantly faster convergence for PFW (see Sect. 6.1). This is another consequence of the sparse iterates which allow the algorithm to quickly identify the few active degrees of freedom of a sparse solution.

Finally, it is noteworthy that, despite its acceleration, PFW is a reliable algorithm as it maintains convergence guarantees. As demonstrated in Jarret et al. (2022), the sequence of iterates converges toward a minimum of the LASSO objective function.

The algorithmic steps. The detailed steps are provided in Algorithm 2. The procedure starts with an empty sky. Each iteration is composed of two steps. During the first, the algorithm identifies a set of candidate atoms located at the regions of maximum intensity of the residual image. This differs from CLEAN, which only looks for a single atom placed at the maximum intensity of the residual image. These candidate atoms are stored in S_k , the set of active indices. The parameter δ controls the intensity of the polyatomic behavior.

During the second step of each iteration, PolyCLEAN simultaneously evaluates the weight to associate with each new atom and adjusts the weights already attributed to the previously selected atoms. Moreover, the algorithm prunes the set of atoms, removing those which get attributed zero weight. This second step is handled by solving another LASSO problem, defined as the restriction of the initial LASSO problem with the constraint that the support of the solution is included in S_k . This constraint significantly reduces the dimensionality of the problem so that this LASSO sub-problem can be solved efficiently both in terms of memory and runtime.

To further accelerate the second step, the solver of the sub-problem is preemptively stopped with initially low accuracy. This early stop does not affect overall reconstruction quality, as the accuracy of the reweighting is increased over the iterations, and this simply results in a speedup (see Jarret et al. 2022 for more details). An instance of an APGD algorithm is used to perform this constrained minimization.

Finally, the algorithm stops if a user-defined stopping criterion is met. As PFW is guaranteed to converge toward a minimum of the LASSO objective function, we naturally use a criterion of small relative improvement of the objective function value between the iterations.

Remark 1 (Positivity constraint). The Polyatomic Frank-Wolfe algorithm can optionally enforce positivity on the solutions, a strategy known to produce more accurate radio interferometric images as discussed in Sect. 4.1. In this case, the candidate atoms are placed only at the positive maximum values of the residual image and the reweighting subproblem is solved with a positivity constraint as well.

Algorithm 2 Polyatomic FW of quality $0 < \delta \leq 1$ **Initialize:** $\mathbf{I}_0 \leftarrow 0, \mathbf{S}_0 \leftarrow \emptyset, \Delta \leftarrow (1 - \delta) \|\Phi^* \mathbf{V}\|_\infty$ **for** $k = 0, 1, 2, \dots$ **do**Dirty residual: $\boldsymbol{\eta}_k \leftarrow \Phi^* (\mathbf{V} - \Phi(\mathbf{I}_k))$

1.a. Polyatomic exploration:

$$\mathcal{I}_{k+1} = \{1 \leq j \leq N : |\eta_{kj}| \geq \|\boldsymbol{\eta}_k\|_\infty - 2\Delta/(k+2)\}$$

1.b. Update active indices:

$$\mathcal{S}_{k+1} \leftarrow \mathcal{S}_k \cup \mathcal{I}_{k+1}$$

2. Update active weights:

$$\mathbf{I}_{k+1} \leftarrow \arg \min_{\text{Supp}(\mathbf{I}) \subset \mathcal{S}_{k+1}} \frac{1}{2} \|\mathbf{V} - \Phi \mathbf{I}\|_2^2 + \lambda \|\mathbf{I}\|_1$$

3. Prune atoms:

$$\mathcal{S}_{k+1} \leftarrow \text{Supp}(\mathbf{I}_{k+1})$$

4. Check convergence:

STOP if a stopping criterion is verified.

Output:Postprocess $\mathbf{I}^{(k)}$ (e.g. convolution with synthetic beam, add residual image)

Standard off-the-shelf implementations of radio interferometry operators perform their computations as if the images were dense. Thus, using them with PFW would require many inefficient data conversions between sparse and dense storage of the iterates. To take full advantage of the sparse iterates of PFW, we need to perform more efficient (de-)gridding operations. Our solution is to use an implementation of the forward operator that conveniently handles non-uniform data and keeps sparse inputs and outputs.

5.2. (De-)Gridding with HVOX

The degridding (computation of the visibilities) and gridding (image synthesis) operations correspond to applying respectively the forward operator Φ and its adjoint Φ^* , as given in Eqs. (4) and (6). A characteristic of the PFW algorithm is that the iterates $\mathbf{I}_k$ are sparse images, and thus stored in a sparse format to enhance scalability. Additionally, the reweighting sub-problem of the second step of PFW enforces a constraint on the support which is straightforward to apply with a sparsity-aware forward operator. For these reasons, PolyCLEAN makes use of the HVOX package (Kashani et al. 2023) to implement the forward operator. This implementation is specifically designed to efficiently handle sparse skies and sparse Fourier planes. Production gridders used in radio astronomy, such as NIFTY (Ye et al. 2022), deal only with dense images and thus are not flexible enough to be combined efficiently with PFW.

Under the hood, HVOX performs a non-uniform Fourier transform, which amounts to evaluating the following sum

$$v_k = \sum_{j=1}^M w_j e^{-j\langle \mathbf{z}_k, \mathbf{x}_j \rangle}$$

for $k \in \{1, \dots, L\}$. The arguments of the transform are $(\mathbf{x}_j, w_j) \in \mathbb{R}^3 \times \mathbb{R}$, a collection of M coordinates-intensity pairs representing the input image, and the $\mathbf{z}_k \in \mathbb{R}^3$, which are the L spatial frequencies to compute. With an appropriate choice of $\mathbf{x}_j, w_j$ and $\mathbf{z}_k$, this sum evaluates the action of Φ and the outputs v_k are the visibility measurements. The image coordinates $\mathbf{x}_j$ can be either non-uniform (support of a sparse image) or uniform (dense image), respectively leading to the so-called type 3 or type 2

NUFFTs. HVOX conveniently implements the adjoint operator Φ^* automatically.

5.3. Interpretation of the regularization parameter

We conclude this section by proposing a way to set and interpret the regularization parameter λ of the LASSO problem (7). The LASSO solution depends crucially on the value of the regularization parameter λ , and setting this value is critical to obtain a good reconstruction. Within the Bayesian interpretation, this parameter sets the ratio between the intensity of the measurement noise and the prior distribution. A high value leads to sparser solutions. However, the range of values that λ can take is arbitrary and varies significantly given the problem dimension. Some strategies exist to automatically infer the value of λ from the data, for instance including another Bayesian prior on the hyperparameters with a hierarchical model, but we do not consider these extensions in this work. Instead, we use a systematic method to scale λ with respect to the problem size, so that it only needs to be fine-tuned manually to account for the noise level. This strategy is commonly used in signal processing and is presented hereafter.

Autoscaling of the penalty parameter. There exists a critical value $\lambda_{\max}$ such that, for any larger values $\lambda \geq \lambda_{\max}$, the only solution to the LASSO problem is null (see for instance (Koulouri et al. 2021, Proposition II.1)). This value has the following closed-form expression

$$\lambda_{\max} = \|\Phi^* \mathbf{V}\|_\infty \geq 0. \quad (14)$$

For PolyCLEAN reconstructions, we propose to set the value of the regularization parameter as a fraction of this maximum value

$$\lambda = \alpha \lambda_{\max}, \quad (15)$$

with $0 \leq \alpha \leq 1$. The value of α typically belongs to the range $[0.01, 0.10]$ and needs to be specifically adapted by the user depending on the estimated noise intensity.

Interpretation of λ . Interestingly enough, we can interpret α as the ratio of measured information that is not explained by the LASSO and remains in the residual. Indeed, at convergence of the algorithm the dual certificate of the problem is computed from the solution and satisfies the conditions given in Sect. 4.2. In particular, when the solution is non-null, Eqs. (9) and (10) lead to $\|\boldsymbol{\mu}_\lambda\|_\infty = 1$. Plugging the expression of λ from Eq. (15), we obtain

$$\alpha = \frac{\|\Phi^* (\mathbf{V} - \Phi \mathbf{I}_{\text{sol}})\|_\infty}{\|\Phi^* \mathbf{V}\|_\infty},$$

where $\mathbf{I}_{\text{sol}}$ is the solution returned by the PFW algorithm. More precisely, α becomes the ratio between the absolute maximum intensity of the residual image and the dirty image. When the algorithm stops, all the pixels in the residual image have an absolute intensity smaller than the α -fraction of the initial maximum intensity of the dirty image.

6. Numerical experiments

In this section we present four experiments that investigated how PolyCLEAN compares to other imaging methods in terms of capabilities and performance.

Table 1. Datasets used in the different numerical experiments in Sect. 6.

Dataset	I	II	III
Subsection	6.1	6.2–6.3	6.4
Source	Point sources	Boötes field	M31
Data type	Simulation	Actual observations	Simulation
Antenna layout	SKA-LOW core	LOFAR core stations	SKA-LOW core
Number of stations	94–362	50	251
Longest baseline (in wavelength)	195–3241	4813	605
Frequency (MHz)	100	145.8	100
Bandwidth (MHz)	1	0.1	1
Number of integration times	7	72	11
Time resolution (s)	4800	400	2880
Total number of baselines ($\times 10^3$)	31–457	68.4	31.6
FOV (deg)	5.0	6.0	6.5
Image size (pix)	288×288–10.000×10.000	3072 × 3072	256 × 256
PSNR (dB)	20	–	20
Nominal resolution (arcsec)	See Table 2	14.3	113.6
Resolution used (arcsec)	See Table 2	7.0	91.4

- Section 6.1 reports on the scalability of PolyCLEAN. We compared the reconstruction time of different methods on simulated measurements of a sky containing point sources with different pixel resolutions.
- In Sect. 6.2 we performed comparative reconstructions of actual radio interferometry observations with CLEAN and PolyCLEAN. We demonstrated that these two methods produced sky maps that are visually similar, with the possibility to perform varying reconstruction depths.
- In Sect. 6.3, we display the dual certificate obtained at convergence of PolyCLEAN on the same observation set. To the best of our knowledge, this kind of uncertainty image had never been used in a radio interferometric imaging context.
- Finally, we evaluate PolyCLEAN for reconstruction of diffuse emissions with a sparse dictionary approach in Sect. 6.4 using simulated measurements of the Andromeda galaxy (M31).

As a common base, we used the RASCIL Python package (SKA Organization 2023) to numerically handle the visibilities, images and interferometer configurations. The code to run PolyCLEAN and to reproduce all the figures is available on a dedicated open access repository⁴. Notably, we implemented PolyCLEAN using Python and its standard scientific computing libraries. We additionally made use of the Pyxu optimization package for conveniently handling arithmetic of linear operators and the optimization sub-routines (Simeoni et al. 2024). We finally relied on the WSCLEAN software (Offringa et al. 2014) to run the different CLEAN-based algorithms, namely Cotton-Schwab for point sources reconstruction and Multi-scale for diffuse emissions, that we both refer to as CLEAN regardless of the context.

The characteristics of the datasets for the different experiments are summarized in Table 1.

In our simulations, we considered omnidirectional antennas and assumed the antenna beam to be constant on the limited field of view. This choice does not impact the scalability and speed of the algorithms.

6.1. Scalability and speed assessment

We investigated the scalability of PolyCLEAN by performing image reconstructions over simulated measurement sets of

⁴ <https://github.com/AdriaJ/polyclean>

increasing size, and compared the solving time of PolyCLEAN against two other methods: CLEAN and APGD (Liang et al. 2022). As already described, CLEAN produces a matching pursuit solution to the inverse problem, demonstrating good performance in terms of speed and result quality, thanks to years of continuous improvement. The other contender, APGD, is a proximal optimization algorithm. We used it to solve the same LASSO problem as PolyCLEAN so that the imaging inverse problem was regularized in the same manner using different numerical solvers. Notably, we recall that APGD does not use sparse iterates.

Simulation of the problem. A given area of the sky was simulated and we performed many visibility measurements of this sky using interferometers of increasing station count. The reconstructed images always covered the same area of the sky, but the resolution increased with the size of the telescope array. Longer baselines means more visibility measurements and more pixels in the reconstruction, hence the inverse problem considered gets larger (L and N increase in $\Phi : \mathbb{R}^N \rightarrow \mathbb{C}^L$). The simulation parameters are summarized in the first column of Table 1. The exact problem dimensions across the experiment are reported in Appendix A.

The visibilities were corrupted with additive Gaussian noise with PSNR nominatively set to 20 dB. The field of view had a fixed angular size of $5^\circ \times 5^\circ$. 200 point sources were randomly placed on this sky. The source intensities were drawn with a log-normal distribution. The baselines reproduced a configuration of the core SKA-LOW interferometer. Only stations closer than a given radius $r_{\max}$ from the center of the interferometer were considered. A low value of $r_{\max}$ only samples the lowest spatial frequencies while higher frequencies are observed as the radius increases.

Resolution of the images. Table 2 reports on the pixel size of the experiments. For a given field of view, the user-defined image resolution plays a significant role in solver scalability. A commonly chosen pixel size is the so-called nominal resolution $\delta\ell_0 = 1/(3B_0)$ (in radians) where $B_0 = \max_{k=\{1,\dots,L\}} \|(u_k, v_k, w_k)\|_2$ is the largest baseline. However, it has been shown that non-linear reconstruction techniques that rely on sparse priors are able to resolve objects at finer resolutions (Honma et al. 2014;

Table 2. Resolution of the images in Sect. 6.1.

$r_{\max}$ (km)	Nominal resolution	Resolution SRF 2	Resolution SRF 5	Resolution SRF 10
0.3	353.4	125.0	50.0	25.0
0.6	180.0	62.5	25.0	12.5
0.9	119.5	46.9	18.8	9.4
1.2	109.4	41.7	16.7	8.3
1.5	82.2	31.2	12.5	6.2
2.0	63.7	25.0	10.0	5.0
3.0	43.4	16.7	6.7	3.3
6.0	21.2	9.0	3.6	1.8

Notes. The resolution is defined as the size of a pixel in arcseconds. The nominal resolution is compared to the resolution of the recovered images when changing the super resolution factor.

Akiyama et al. 2017; Dabbech et al. 2018; Pan et al. 2017). We thus chose to perform the experiments at different fine resolutions which correspond to approximately 2, 5 and 10 times smaller pixel size than $\delta\ell_0$. We refer to these setups as SRF2, SRF5 and SRF10 (for Super-Resolution Factor).

Solvers parameters. For a fair comparison between the different methods, the algorithms were configured so that they output reconstructions with similar quality by adjusting the reconstruction parameters as well as the stopping criteria. The reasons for those choices are as follows.

Regarding quality assessment, model images returned by different solvers are difficult to compare to the simulated ground truth due to their pixel-sparse nature. Any localization error excessively impacts regular image metrics such as mean square error (MSE) or mean absolute deviation (MAD). Indeed, even a one pixel error on the reconstruction leads to high metrics values although the recovered intensity may be correct. For this reason, we first smoothed the output models by convolving them with the CLEAN beam – a Gaussian kernel fitted on the main lobe of the PSF – before computing MSE and MAD metrics. The resulting metrics are then affected by localization and intensity errors in a more balanced manner. These were then used to ensure that CLEAN and PolyCLEAN produced similar quality images. The average values of these metrics along repetitions of this experiment are indeed comparable and provided in Appendix B.

The choice of an appropriate stopping criterion is also critical to obtain relevant time comparisons. In contrast to CLEAN, reconstructions obtained with LASSO solvers usually benefit from finer convergence, at the cost of longer solving time. A compromise is needed between time efficiency and accuracy. The following are the three algorithms and their respective parameters.

Cotton-Schwab CLEAN was performed with minor cycles, using WSCLEAN with a `mgain` parameter of 0.7 and a `cycle gain` parameter of 0.1. The `niter` parameter was set high and never reached. WSCLEAN stopped when the residual noise was smaller than some threshold. As recommended by the authors (Offringa et al. 2014), we used the `auto-threshold` argument that self-determined the value of the threshold with respect to the standard deviation of the image. In this experiment, the `auto-threshold` was set to 1.

PolyCLEAN solves a LASSO problem, hence the quality of the reconstruction and the residual noise directly depend on the value of the regularization parameter. We used the rule proposed




Fig. 1. Schematic view of the simulation pipeline for the scalability experiment. The CLEAN procedure starts from the dirty image, while PolyCLEAN and APGD take as input the visibility measurements $\mathbf{V}$ and instantiate the initial iterate empty.

in Sect. 5.3 and set $\lambda = \alpha \lambda_{\max}$ with $\alpha = 0.01$. This value was handpicked to provide similar reconstruction quality between CLEAN and PolyCLEAN. The algorithm stopped when the relative improvement of the LASSO objective function got smaller than 10^{-4} .

APGD solved the same LASSO problem as PolyCLEAN, so the regularization parameter had the same value. The algorithm stopped as soon as the value of the objective function reached the minimum obtained with PolyCLEAN so that both solvers had equal treatment with respect to the objective function. APGD also used the HVOX implementation of the forward operator.

Implementation. CLEAN and PolyCLEAN are at different levels of maturity. The implementation of WSCLEAN is state-of-the-art C++ and has benefited from continuous development. In contrast, PolyCLEAN and APGD are based on the Python package Pyxu. Only the computation of the forward operator has been optimized using HVOX: the rest of the numerical treatment was done with `numpy`. Thus, the good performance demonstrated by PolyCLEAN in these experiment augured well for the use of sparsity-based algorithms to improve the scalability of optimization-based methods.

Simulation results. The scaling behavior of the three algorithms is shown in Fig. 2, represented by the slope of the curves. The upper x -axis represents the radius of the interferometer, while the lower x -axis displays the corresponding image size (in megapixels). Both are represented on a logarithmic scale. The reconstruction times are provided on the y -axis on a logarithmic scale as well. The markers represent the median timings over 10 repetitions with the same interferometer radius, while the shaded areas represent the inter-quartile range and illustrate the variability over repetitions. The reconstructions were run on a workstation containing a processor Intel Core i9-10900X 10-core CPU and 128 GB of DDR4 memory. Some reconstruction examples with PolyCLEAN for various problem sizes are illustrated in Appendix C.




Fig. 2. Reconstruction time for different super resolution factors SRF 2 (a), SRF 5 (b), and SRF 10. PolyCLEAN consistently outperforms the dense APGD solver. CLEAN is faster than PolyCLEAN for low SRF factor but gets outperformed when the pixel size decreases. The PolyCLEAN curve, seen as reconstruction time with respect to maximum baseline, is almost identical on the three plots irrespective of the SRF factor.

The experiment demonstrates that PolyCLEAN performed significantly better than APGD for all problem dimensions considered. This result strengthens the observations made in the original PFW article (Jarret et al. 2022).

Compared to CLEAN, PolyCLEAN also demonstrated good scalability in the high dimensional contexts considered. Indeed, we can see on the graphs Figs. 2b and 2c that PolyCLEAN matched and even outperformed WSCLEAN for images of size 3 megapixels and more. However, CLEAN was consistently faster than the LASSO-based methods for small size images. For the low super-resolution factor considered in graph Fig. 2a, CLEAN also significantly outperformed the optimization methods.

One possible explanation for the three plots is the sparsity of the simulated source images. Indeed, there were only 200 point sources spanning 200 pixels, independent of image sizes (which varied from a few thousand pixels up to 100 megapixels). Therefore, the source images got sparser as the interferometer size increased. In this context, the sparse processing and sparsity-aware operators of PolyCLEAN have enabled efficient high-resolution imaging in observational setups with large baselines, increasingly outperforming the dense operations of CLEAN.

Interestingly, we notice that the speed curves of PolyCLEAN are nearly identical on all subplots Figs. 2a, 2b and 2c irrespective of the chosen SRF. This observation shows that the reconstruction time of PolyCLEAN merely depends on the intrinsic resolution of the solution (correlated to the largest baseline) and was little affected by the actual reconstruction resolution.

6.2. LOFAR observations of the Boötes field

The second experiment compared the images reconstructed by PolyCLEAN and CLEAN when applied to the same dataset of the Boötes field observations. Consistent with experiments using other sparse optimization techniques (Wiaux et al. 2009; Li et al. 2011; Carrillo et al. 2012 to name a few), PolyCLEAN matched the image quality obtained by CLEAN. Additionally, we illustrate the ability of PolyCLEAN to perform image reconstruction at different levels of residual noise, similar to CLEAN.

The observations. Dataset II is presented in Table 1. It contains actual observations made with LOFAR (LOW-Frequency ARray, van Haarlem et al. 2013) over an observation period of 8 hours with an integration time of 8 seconds. The visibilities were subsampled in time so that only 2% of the dataset was used for image reconstruction, with one measurement every 400 seconds (one measurement kept every 50 integration times). The measurements were obtained with LOFAR using the 48 core stations and the two closest ones out of the core. The resulting uv -coverage is displayed in Fig. 3. The field of view was a square covering $6^\circ \times 6^\circ$ of the sky, and was imaged over 3072×3072 pixels, giving a pixel resolution of 7.0 arcseconds. Comparatively, the nominal resolution for the considered observations was 14.3 arcseconds (which is a super resolution factor of 2.03). The dirty image resulting from this measurement set is provided in Fig. 4. For visualization purposes, two scales were used. The background components are displayed in a symmetric linear white-to-black colormap, centered around 0, while the more intense foreground components are represented with a hot colormap (black - red - yellow) with a square root scale. The square root law conveniently compresses the higher end of the dynamic range.

Reconstruction parameters. The following configuration was used.




Fig. 3. uv -coverage of dataset II, coordinates are given in multiple of observation wavelength.




Fig. 4. Dirty image of the Boötes field obtained with dataset II. Reconstructions are provided in Figs. 5 and 6.

PolyCLEAN was used in the same setup as in the first experiment: HVOX was employed to perform the (de-)gridding operations and the regularization parameter λ was set to $\alpha\lambda_{\max}$ for some $0 < \alpha < 1$. The algorithm again stopped when the relative improvement on the objective function became smaller than 10^{-4} .

CLEAN was also applied using the standard Cotton-Schwab version implemented by WSCLEAN. The auto-threshold parameter determined the stopping criterion, as recommended by the authors.

In this experiment, the α regularization parameter for PolyCLEAN took the values $[0.05, 0.02, 0.005]$ while the autothreshold for CLEAN explored the range $[3\sigma, 2\sigma, 1\sigma]$ (with σ the standard deviation of the dirty image). The values of α were handpicked so that PolyCLEAN visually matched the reconstruction depth of CLEAN in each case.

Image reconstruction. Both CLEAN and PolyCLEAN produced a model image of point sources which can be very sparse and difficult to interpret. We used the standard radio interferometry pipeline to produce the images provided in Figs. 5 and 6.

Table 3. Error metrics compared to the catalog.

(a) MSE scores			
PolyCLEAN α	MSE	WSCLEAN MSE	autothresh
0.05	342.2	617.8	3
0.02	478.8	651.6	2
0.005	613.4	671.5	1

(b) MAD scores			
PolyCLEAN α	MAD	WSCLEAN MAD	autothresh
0.05	2.32	2.91	3
0.02	2.37	3.46	2
0.005	2.74	4.70	1

Notes. MSE and MAD between the reconstructions and a catalog image synthesized with the same CLEAN beam. The reconstructions considered are the ones provided in Fig. 5 with the same reconstruction parameters (respectively α and autothresh).

The same colormap as the dirty image is used for display. First, a CLEAN beam was computed by fitting a Gaussian kernel to the main lobe of the PSF. The model images were then convolved with this CLEAN beam, leading to Fig. 5. The dirty residual was added in Fig. 6 as it may contain some remaining information. We emphasize that the dirty residual also contained most of the noise isolated during reconstruction, hence it needs to be interpreted with care.

Analysis. At first glance the reconstructions of PolyCLEAN and CLEAN look extremely similar. Both methods are able to perform coarse and fine reconstructions by adjusting the parameters to control the level of residual noise. Indeed, the first row of Fig. 6 contains the most residual noise which leads to a grainy gray background. This background becomes smoother on the lower rows as the residual is reduced in intensity.

LASSO solutions are known to introduce a bias in recovered intensities such that some sources may undergo a shrinkage phenomenon which strengthens with the value of the regularization parameter. Additionally, CLEAN may introduce many spurious sources when trying to completely annihilate the residual. Both effects were however smoothed out when convolving with the CLEAN beam, as the latter has large spatial extent compared to the spread of the observed stars.

Note that the reconstructed images are not directly comparable side-by-side as the regularization parameters chosen for PolyCLEAN do not relate to the stopping criteria of WSCLEAN. Hence, each of the six displayed images has a different level of residual image.

We evaluated the quality of the reconstructions by comparing them to the catalog provided by Williams et al. (2016). To do so, we synthesized a reference image by convolving the sources from the catalog with the CLEAN beam from the observations of dataset II, resulting in the image provided in Fig. 7. The MSE and MAD between the reconstructions and this synthetic reference image are reported in Table 3. Both metrics indicate that the PolyCLEAN images are more faithful to the reference than the images obtained with PolyCLEAN, consistently reaching the best performance for $\alpha = 0.05$. For completeness, the difference images between the reconstructions and the reference are provided in Appendix D.




Fig. 5. Comparison without residual. *Left:* PolyCLEAN images, with decreasing α values that determine the level of regularization. *Right:* WSCLEAN images with decreasing level of residual noise at stop, determined by the `auto-threshold` parameter. As mentioned, images should not be directly compared side-by-side as they do not present the same level of regularization.




Fig. 6. Comparison with dirty residual. Same experiment as Fig. 5. PolyCLEAN images are on the left while WSCLEAN ones are on the right.




Fig. 7. Catalog image of the Boötes field corresponding to dataset II, constructed using the CLEAN beam computed with the stations of dataset II.

6.3. Dual certificate

As described in Sect. 4.2, PolyCLEAN solves a LASSO problem for which we can compute the dual certificate at convergence μ_λ with Eq. (8). μ_λ is an image with the same dimension as the sky image with values in the range $[-1, 1]^5$.

For imaging purposes, we are interested in the saturation set $\mathcal{D}_\lambda$ (defined in Eq. (11)) as this set is made of all the candidate reconstructed source locations. Remarkably, when the LASSO problem admits more than one solution, the dual certificate does not depend on the actual reconstruction returned by PolyCLEAN and delimits every possible reconstruction this LASSO could have produced.

Figure 8 shows the dual certificate associated with the LASSO reconstruction of the Boötes field from the previous section (dataset II), with a penalty parameter of $\lambda = 0.02\lambda_{\max}$. Two colorbars are given, both on a linear scale. The warm colorbar represents the highest values of the dual certificate, above a threshold at 0.8. The bright areas can be used as a proxy to identify the saturation set of the dual certificate. Indeed, the convergence is subject to numerical approximations so that the dual certificate does not reach the exact saturation set of the problem.

We observe that the bright areas indeed correspond to the center of the sources reconstructed in the previous section (see Fig. 6). An interesting phenomenon however is that the size of these bright blobs is smaller than the areas of the sources recovered after convolution with the CLEAN beam. In this sense, the dual certificate is more resolved around the recovered source locations, and the area of each blob can be interpreted as an uncertainty estimate on the position. When these uncertainty areas are small, the reconstructed sources are precisely identified by the LASSO, with more confidence than what we obtain in the CLEAN-like reconstructed image after convolution with the CLEAN beam.

Figure 9 provides a zoom on a small set of point sources. The left panel contains the dual certificate, while the right panel illustrates the PolyCLEAN reconstruction. We can see that the dual

certificate identifies more sources than the reconstruction. However, it does not inform as to the intensity of these sources, which needs to be read from the reconstructed maps. Strong sources appear brighter and more resolved on the dual certificate than fainter ones as they are easier to recover and localize.

Many use cases for the dual certificate can be considered. For instance, computing the area of each bright zone can provide a way to quantify the uncertainty on the source location. For illustration, we propose a procedure to build a representation beam, similar to the CLEAN beam, that encompasses the actual critical resolution of the LASSO reconstruction. The steps are as follows:

1. Threshold the dual certificate to a target level of accuracy, e.g. set to zero the pixels with value lower than 0.9;
2. Compute the Power Spectral Density by taking the square of the magnitude of the Discrete Fourier Transform, which corresponds to the Fourier transform of the autocorrelation of the dual certificate;
3. Perform an Inverse Discrete Fourier Transform. We obtain a sort of PSF whose main lobe has a spatial extension induced by the size of the intensity peaks of the dual certificate;
4. Fit a Gaussian kernel to the main lobe and normalize to 1 similarly to the computation of the CLEAN beam.

With this procedure, the spread of the representation beam is faithful to the resolution of the LASSO solution at convergence of the minimization algorithm. Figure 10 displays the PolyCLEAN reconstruction once convolved with the representation beam obtained with this procedure. The image is significantly more resolved than with the CLEAN beam while preserving the intensity values.

6.4. Diffuse emissions in simulations

For the final experiment, we demonstrated that PolyCLEAN can also be used to reconstruct diffuse emissions, making it as versatile and powerful as algorithms from the CLEAN family.

Simulation of the visibilities. Dataset III consists of a set of simulated measurements from an image of an Halpha region in the Andromeda galaxy (referred to as M31). The source image represents diffuse emissions and is displayed on the left panel of Fig. 11. We used a test image of size 256×256 provided by RASCIL (SKA Organization 2023) and we simulated the visibilities with the SKA-LOW core configuration. The detail of the simulation parameters are reported in Table 1. The test image used here is not meant to be realistic, as the field of view covered by the galaxy on our test image is much larger than its actual size. This choice of simulation parameters allowed us to reduce the size of the baselines so that we could illustrate the reconstruction capabilities of PolyCLEAN on a less demanding setup. The nominal resolution was 113.6 arcseconds and the image was reconstructed with a slightly finer resolution of 91.4 arcseconds.

Reconstruction methods. The two reconstruction methods considered were PolyCLEAN and WSCLEAN, with the following parameters.

PolyCLEAN solved a LASSO problem for sparse dictionary image synthesis following the description in Sect. 4.3. The dictionary was composed of four Gaussian kernels $\mathbf{g}_i$ with different scales σ_i and scale-bias factors γ_i that were discretized and cropped to a compact support. Formally, we can write $\mathbf{g}_i = \gamma_i \tilde{\mathbf{g}}_i$ with $\tilde{\mathbf{g}}_i[n, m] = (1/G_i) e^{-(n^2 + m^2)/\sigma_i^2}$ for $|n|, |m| \leq 2\sigma_i$ and G_i a normalization factor such that $\|\tilde{\mathbf{g}}_i\|_1 = 1$. We considered the scales $\sigma_i = \{0, 2, 5, 8\}$ and the associated weights $\gamma_i = (r)^i$ with

⁵ When a positivity constraint is enforced, the dual certificate may have values smaller than -1 , but still saturates at maximum 1.






Fig. 8. Dual certificate μ_λ obtained on the Boötes field reconstruction with parameter $\alpha = 0.02$. Figure 9 shows a zoom-in on the boxed region. This certificate corresponds to the reconstructions displayed on left side, second row in Figs. 5 and 6.



(a)



(b)

Fig. 9. Zoom-in: comparison of the dual certificate (a) and the PolyCLEAN reconstruction (b) with the same value of α . The reconstruction is convolved with the CLEAN beam.




Fig. 10. PolyCLEAN reconstruction convolved with a representation beam computed from the dual certificate. The area represented is the same as in Fig. 9 using the same colormap as Fig. 9b. The threshold to the dual certificate was set to 0.9.

$r = 1.2$ for $i = \{0, 1, 2, 3\}$ to balance the bias between narrow and wide components in the reconstruction⁶. Notably, this multi-scale model allows point sources to be placed with the scale $\sigma_0 = 0$. The regularization parameter was set to $\lambda = 0.02\lambda_{\max}$ (note that $\lambda_{\max}$ then needed to be computed with the synthesis operator Ψ as well). The algorithm stopped when the relative improvement of the objective function was smaller than 10^{-5} . More generally, the weights γ_i of the dictionary elements can be chosen to bias the reconstruction toward the inclusion of more point sources or more diffuse components.

WSCLEAN implemented the Multi-scale CLEAN algorithm as described in Offringa & Smirnov (2017). In particular, the reconstruction atoms were tapered quadratic functions that are also compactly supported. The autothreshold parameter was again set to one standard deviation.

⁶ The geometric growth of the scale-bias factors was directly inspired from Offringa & Smirnov (2017).




Fig. 11. Model images.




Fig. 12. Dirty image obtained with simulated measurements of the Andromeda galaxy (dataset III).

Image reconstruction. The dirty image is displayed in Fig. 12 and the reconstructions are provided in Figs. 11, 13, 14. These display respectively the model images, the result after convolution with a representation beam, and the final reconstructed images after addition of the residuals. The model images produced by both algorithms here contain diffuse sources such that these images already represent sources with a spatial extension. For this reason, we used a representation beam more resolved than the CLEAN beam (half of the spread along the two directions). This beam slightly smoothed the models while preserving the recovered multi-scale information. The resulting images are then sharper than if they had been convolved with the CLEAN beam. We also provide some quality metrics in Table 4. A more detailed view on the center of the images is provided in Appendix E. Additionally, we compare the reconstructed images with the source and between themselves in Figs. 15 and 16 using difference imaging.

Analysis. From Figs. 11 and 13, we see that the two methods reconstructed similar-looking images, both testifying to the multi-scale information of the source. The reconstructions have a similar dynamic range that is consistent with the source. This

Table 4. Metrics for reconstruction of diffuse emissions in Sect. 6.4 (rescaled, no unit).

	PolyCLEAN		WSCLEAN	
	MSE	MAD	MSE	MAD
Model images	0.1	4.3	1.7	8.2
Model convolved	12.2	45.8	5.7	38.7
+ residuals	55.5	190.0	14.1	91.4
Total weight ($\cdot/1495$ Jy)	1300		1475	
Reconstruction time (s)	57.0		5.8	

experiment was run with a pixel size close to the nominal resolution, which partly explains the slow reconstruction speed of PolyCLEAN compared to CLEAN (10 times longer).

CLEAN however tended to reconstruct many spurious point sources of small intensity. These are particularly visible on Figs. 11 and 13. This known behavior – interpreting the noise in the data as additional sources – lead to the grainy pattern seen in both figures. When the dirty residuals are added in Fig. 14, the background of the CLEAN image is more regular than that of PolyCLEAN, as it had deliberately left some information aside and identified it as noise. This behavior is also illustrated with the reconstruction metrics provided in Table 4, in which CLEAN produced better metrics than PolyCLEAN when residuals were considered. Stopping CLEAN earlier would help to alleviate this caveat but at the same time worsen the quality of the recovered emissions.

The PolyCLEAN reconstruction suffered from a loss of overall intensity in the image. We observe this phenomenon at line Total weight in Table 4, which corresponds to the sum of all the reconstructed elements in the model image. This shortcoming of PolyCLEAN is once again explained by the shrinkage effect of LASSO. Despite the fact that PolyCLEAN recovered less information in total, the end solution outperform MS-CLEAN for the model images, demonstrating the relevance of the recovered sources.

Figures 15 and 16 display difference images between various reconstructions and so take positive and negative values. Both cases have been normalized by the maximum intensity in the respective source images, hence the displayed value is interpreted as a pointwise error relative to the maximum of the source.




Fig. 13. Model images convolved with a sharp beam.




Fig. 14. Reconstructions with residuals.

Looking at the differences between model images in Fig. 15, we can once again notice the tendency of WSCLEAN to place point sources in the model image. Indeed, on the difference between PolyCLEAN and WSCLEAN (right), sharp blue pixels on the center of the image correspond to the CLEAN sources where PolyCLEAN places smoother red emissions. This creates high values of error so we used a square root scale to enhance the contrast. After convolution with the sharp representation beam (Fig. 16), the differences between the reconstructions are less intense (colorbar ranges from -15 to $+15\%$). This convolution is hence of critical importance for CLEAN-based reconstruction but at the same time has a negative impact on the resolution of the recovered image. The PolyCLEAN model image is then an appealing solution to preserve sharpness in the reconstruction.

7. Conclusions

PolyCLEAN proposes the first application of a Frank-Wolfe algorithm to image reconstruction in radio interferometry. By design, it belongs to the family of imaging optimization methods that use the LASSO problem as its guiding reconstruction principle. PolyCLEAN manages to achieve a reconstruction quality comparable to modern implementations of CLEAN, with better reconstruction of point sources and diffuse emissions, even though the model is simple in comparison to the state-of-the-art in optimization. This versatility makes PolyCLEAN already a promising and usable reconstruction framework.

The strength of PolyCLEAN lies in its competitive numerical performance in terms of scalability and reconstruction speed. Its design takes full advantage of sparsity properties of the imaging inverse problem involved in radio interferometry. In particular, when super-resolution was involved, the increased sparsity in the solutions allowed PolyCLEAN to outperform optimized CLEAN solvers such as WSCLEAN.

PolyCLEAN was developed in Python, with only the computation of the sparse forward operator having been numerically optimized. Further code development is needed to push PolyCLEAN to modern scientific computing standards. Notably, the PolyCLEAN implementation should include parallel distributed computing to tackle the numerical challenges of radio interferometry.

Additionally, we believe that the dual certificate is an insightful tool in the context of radio interferometry. Quantitative imaging techniques will play an essential role in the future of computational imaging, especially with the current rise of learning methods where trust in the reconstruction is of high concern. The dual certificate image makes a step toward an a posteriori analysis of the reconstruction.

Finally, we would like to advocate for the general use of Bayesian methods in computational imaging, as the framework enables data-fidelity and regularization terms used in modern optimization techniques to be interpreted. This results in a better understanding of the reconstruction when compared to CLEAN, for which solution characterization is limited. The Bayesian framework also opens the door for the use of some powerful




Fig. 15. Difference between model images, square root color scale. *Left:* PolyCLEAN minus source image. *Center:* WSCLEAN minus source image. *Right:* PolyCLEAN minus WSCLEAN.




Fig. 16. Difference between convolved images. *Left:* PolyCLEAN minus source image. *Center:* WSCLEAN minus source image. *Right:* PolyCLEAN minus WSCLEAN.

tools. Extensions considered to PolyCLEAN include the use of Bayesian uncertainty quantification in a similar fashion to Cai et al. (2018a) and hierarchical Bayesian models for determining the hyper-parameters in an automatic, data-driven manner.

Data availability

We used the measurement data provided in Pan et al. (2017), which can be accessed following the instructions at this <https://github.com/hanjiapan/LEAP>

Acknowledgements. The authors sincerely thank Martin Vetterli for his guidance and support throughout the writing of this article. We also thank the team at the EPFL Center for Imaging for help on computational aspects and use of their server. We finally acknowledge the anonymous referee whose review contributed to significantly improve the quality of this document. A.J. is funded by the Swiss National Science Foundation (SNSF) under grants AstroSignals: A New Window on the Universe, with the New Generation of Large Radio-Astronomy Facilities, Sinergia (CRSII5_193826) and SESAM - Sensing and Sampling: Theory and Algorithms (n° 200021_181978/1).

References

- Abdulaziz, A., Dabbech, A., & Wiaux, Y. 2019, *MNRAS*, **489**, 1230
 Akiyama, K., Kuramochi, K., Ikeda, S., et al. 2017, *ApJ*, **838**, 1
 Arras, P., Bester, H. L., Perley, R. A., et al. 2021, *A&A*, **646**, A84
 Beck, A., & Teboulle, M. 2009, *SIAM J. Imaging Sci.*, **2**, 183

- Bhatnagar, S., & Cornwell, T. J. 2004, *A&A*, **426**, 747
 Boyer, C., Chambolle, A., Castro, Y. D., et al. 2019, *SIAM J. Optimiz.*, **29**, 1260
 Cai, X., Pereyra, M., & McEwen, J. D. 2018a, *MNRAS*, **480**, 4170
 Cai, X., Pereyra, M., & McEwen, J. D. 2018b, *MNRAS*, **480**, 4154
 Carrera, D., Boracchi, G., Foi, A., & Wohlberg, B. 2017, *IEEE Signal Process. Lett.*, **24**, 1468
 Carrillo, R. E., McEwen, J. D., & Wiaux, Y. 2012, *MNRAS*, **426**, 1223
 Chael, A. A., Johnson, M. D., Narayan, R., et al. 2016, *ApJ*, **829**, 11
 Chen, S. S., Donoho, D. L., & Saunders, M. A. 2001, *SIAM Rev.*, **43**, 129
 Clark, B. G. 1980, *A&A*, **89**, 377
 Cornwell, T. J. 2008, *IEEE J. Sel. Top. Signal Process.*, **2**, 793
 Cornwell, T. J. 2009, *A&A*, **500**, 65
 Dabbech, A., Onose, A., Abdulaziz, A., et al. 2018, *MNRAS*, **476**, 2853
 Dabbech, A., Aghabiglou, A., Chu, C. S., & Wiaux, Y. 2024, *ApJ*, **966**, L34
 Duval, V., & Peyré, G. 2015, *Foundations Comput. Math.*, **15**, 1315
 Duval, V., & Peyré, G. 2017, *Inverse Prob.*, **33**, 055008
 Elad, M. 2010, *Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing* (New York, NY: Springer)
 Elad, M., Milanfar, P., & Rubinstein, R. 2006, in *2006 14th European Signal Processing Conference*, 1
 Foucart, S., & Rauhut, H. 2013, *A Mathematical Introduction to Compressive Sensing, Applied and Numerical Harmonic Analysis* (New York, NY: Springer)
 Frank, M., & Wolfe, P. 1956, *Naval Research Logistics Quarterly*, **3**, 95
 Fried, D. L. 1995, *J. Opt. Soc. Am. A*, **12**, 853
 Garsden, H., Girard, J. N., Starck, J. L., et al. 2015, *A&A*, **575**, A90
 Gottschling, N. M., Antun, V., Hansen, A. C., & Adcock, B. 2023, *SIAM Review*, submitted [arXiv:2001.01258]
 Harchaoui, Z., Juditsky, A., & Nemirovski, A. 2013, in *NIPS Workshop on Optimization for ML*
 Honma, M., Akiyama, K., Uemura, M., & Ikeda, S. 2014, *PASJ*, **66**, 95
 Högbom, J. A. 1974, *A&AS*, **15**, 417

- Jaggi, M. 2013, in [Proceedings of the 30th International Conference on Machine Learning \(PMLR\)](#), 427
- Jarret, A., Fageot, J., & Simeoni, M. 2022, [IEEE Signal Process. Lett.](#), **29**, 637
- Junklewitz, H., Bell, M. R., Selig, M., & Enßlin, T. A. 2016, [A&A](#), **586**, A76
- Kashani, S. 2017, Master's thesis, EPFL - Swiss Federal Technology Institute of Lausanne, Lausanne, Switzerland
- Kashani, S., Queralt, J. R., Jarret, A., & Simeoni, M. 2023, arXiv e-prints [arXiv:2306.06007]
- Koulouri, A., Heins, P., & Burger, M. 2021, [IEEE Trans. Signal Process.](#), **69**, 165
- Lacoste-Julien, S., & Jaggi, M. 2015, in [Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 1, NIPS'15](#) (Cambridge, MA, USA: MIT Press), 496
- Lannes, A., Anterrieu, E., & Maréchal, P. 1997, [A&AS](#), **123**, 183
- Li, F., Cornwell, T. J., & Hoog, F. d. 2011, [A&A](#), **528**, A31
- Liang, J., Luo, T., & Schönlieb, C.-B. 2022, [SIAM J. Sci. Comput.](#), **44**, A1069
- Locatello, F., Khanna, R., Tschannen, M., & Jaggi, M. 2017, in [Proceedings of the 20th International Conference on Artificial Intelligence and Statistics \(PMLR\)](#), 860
- Locatello, F., Raj, A., Karimireddy, S. P., et al. 2018, in [Proceedings of the 35th International Conference on Machine Learning \(PMLR\)](#), 3198
- Mallat, S., & Zhang, Z. 1993, [IEEE Trans. Signal Process.](#), **41**, 3397
- Müller, H., & Lobanov, A. P. 2023a, [A&A](#), **673**, A151
- Müller, H., & Lobanov, A. P. 2023b, [A&A](#), **672**, A26
- Offringa, A. R., & Smirnov, O. 2017, [MNRAS](#), **471**, 301
- Offringa, A. R., McKinley, B., Hurley-Walker, N., et al. 2014, [MNRAS](#), **444**, 606
- Onose, A., Dabbech, A., & Wiaux, Y. 2017, [MNRAS](#), **469**, 938
- Pan, H., Simeoni, M., Hurley, P., Blu, T., & Vetterli, M. 2017, [A&A](#), **608**, A136
- Parikh, N., & Boyd, S. 2014, [Foundations Trends Optimiz.](#), **1**, 127
- Rau, U., Bhatnagar, S., Voronkov, M. A., & Cornwell, T. J. 2009, [Proc. IEEE](#), **97**, 1472
- Scaife, A. M. M. 2020, [Phil. Trans. R. Soc. A.](#), **378**, 20190060
- Schmidt, K., Geyer, F., Fröse, S., et al. 2022, [A&A](#), **664**, A134
- Schwab, F. R. 1984, [AJ](#), **89**, 1076
- Schwardt, L. C. 2012, in [2012 International Conference on Electromagnetics in Advanced Applications](#), 690
- Schwarz, U. J. 1978, [A&A](#), **65**, 345
- Simeoni, M. 2020, PhD thesis, EPFL - Swiss Federal Technology Institute of Lausanne, Lausanne, Switzerland
- Simeoni, M., & Hurley, P. 2021, in [2021 IEEE International Conference on Acoustics, Speech and Signal Processing \(ICASSP\)](#), 4535
- Simeoni, M., Kashani, S., Rué-Queralt, J., & Developers, P. 2024, <https://doi.org/10.5281/zenodo.4486431>
- SKA Organization 2023, RASCIL
- Solo, V. 2008, in [2008 IEEE International Conference on Acoustics, Speech and Signal Processing \(ICASSP\)](#), 3665
- Starck, J.-L. & Murtagh, F. 2006, [Astronomical Image and Data Analysis, Astronomy and Astrophysics Library](#) (Berlin, Heidelberg: Springer)
- Terris, M., Dabbech, A., Tang, C., & Wiaux, Y. 2023, [MNRAS](#), **518**, 604
- Thompson, A. R., Moran, J. M., & Swenson, G. W. 2017, [Interferometry and Synthesis in Radio Astronomy, Astronomy and Astrophysics Library](#) (Cham: Springer International Publishing)
- Thouvenin, P.-A., Abdulaziz, A., Dabbech, A., Repetti, A., & Wiaux, Y. 2023a, [MNRAS](#), **521**, 1
- Thouvenin, P.-A., Dabbech, A., Jiang, M., et al. 2023b, [MNRAS](#), **521**, 20
- Tibshirani, R. 1996, [J. R. Stat. Soc. Ser. B](#), **58**, 267
- Tibshirani, R. J. 2013, [Electron. J. Stat.](#), **7**, 1456
- Tibshirani, R. J., & Taylor, J. 2011, [Ann. Stat.](#), **39**, 1335
- Unser, M., Fageot, J., & Gupta, H. 2016, [IEEE Trans. Information Theory](#), **62**, 5167
- Van der Veen, A.-J., Wijnholds, S. J., & Sardarabadi, A. M. 2019, [Signal Processing for Radio Astronomy](#) (Cham: Springer International Publishing), 311
- van Haarlem, M. P., Wise, M. W., Gunst, A. W., et al. 2013, [A&A](#), **556**, A2
- Wiaux, Y., Jacques, L., Puy, G., Scaife, A. M. M., & Vanderghenst, P. 2009, [MNRAS](#), **395**, 1733
- Williams, W. L., van Weeren, R. J., Röttgering, H. J. A., et al. 2016, [MNRAS](#), **460**, 2385
- Ye, H., Gull, S. F., Tan, S. M., & Nikolic, B. 2022, [MNRAS](#), **510**, 4110

Appendix A: Scaling of the simulated problems

Table A.1: Dimensions of the simulated problems in the experiment from Sect. 6.1.

$r_{\max}$ (km)	Number of visibilities (10^3)	SRF 2		SRF 5		SRF 10	
		Pix/side	Size (MPix)	Pix/side	Size (MPix)	Pix/side	Size (MPix)
0.3	30.6	144	0.02	360	0.13	720	0.5
0.6	168.6	288	0.08	720	0.52	1440	2.1
0.9	209.2	384	0.15	960	0.92	1920	3.7
1.2	235.7	432	0.19	1080	1.17	2160	4.7
1.5	269.5	576	0.33	1440	2.07	2880	8.3
2.0	293.3	720	0.52	1800	3.24	3600	13.0
3.0	344.0	1080	1.17	2700	7.29	5400	29.2
6.0	457.4	2000	4.00	5000	25.00	10000	100.0

Notes. For each value of $r_{\max}$, the number of visibilities considered increase and we have access to higher frequencies. Hence, more pixels can be used to cover the same field of view and the resolution of the image increases. This is valid for any chosen super resolution factor.

Appendix B: Error metrics




Fig. B.1: Comparative view of the reconstruction errors on the simulated images during the experiment of Sect 6.1. Median values over the 10 repetitions of the experiments are reported with the crosses and the transparent areas cover the interquartile spread. MSE appears on the first row, MAD on the second one. The metrics are computed after convolution with the CLEAN beam. Each column is a different super resolution factor: 2, 5, 10 (from left to right). We observe that the metrics have similar values between CLEAN and the LASSO-based methods (PolyCLEAN and APGD). This indicates that the different algorithms are indeed configured to return images of similar quality and thus the comparison between the solving times performed in the experiment is relevant.

Appendix C: Reconstruction examples




Fig. C.1: PolyCLEAN reconstructions of simulated skies in the experiment of Sect 6.1. Each row uses a different r_{max} and each column is with a different SRF. From top to bottom: 2000, 3000, 6000 – from left to right: SRF2, SRF5, SRF10. The field of view is kept fixed but the number of pixels increases from left to right and from top to bottom. As a consequence of including longer baselines, the CLEAN beam also gets narrower (from top to bottom).

Appendix D: Difference imaging on observations




Fig. D.1: Differences between the reconstructions and the catalog image from (Williams et al. 2016). Red and blue respectively indicate over- and under-estimation compared to the catalog. Intensities are displayed with a square root colorscale.

We observe that some strong sources are overestimated both with PolyCLEAN and WSCLEAN (bright red dots). This consistent behavior between the two methods probably originates from a mismatch between our observation data and the catalog. Location errors appear as a red dot next to a blue one but very few are present here. Additionally, we once again observe the shrinkage behavior of the LASSO problem with the PolyCLEAN reconstructions as well the tendency of CLEAN to place sources to explain the noise when the algorithm runs too long.

Appendix E: Details on reconstruction of diffuse emissions




Fig. E.1: Reconstructions of diffuse emissions (Sect. 6.4): model images (a), convolution with a sharp beam (b), and reconstructions with residuals.

We observe that WSCLEAN produced many single-pixel point sources in the reconstruction, even though it was configured to recover diffuse emissions, as we clearly see in Fig. E.1a. The model image proposed by PolyCLEAN involves more diffuse components, which make the image visually closer to the source. It resulted in better metrics, as reported in Table 4.

After convolution however, the difference between the two reconstructions is more difficult to visually analyze (Fig. E.1b). The hotspots displayed the same intensity and had the same locations, which are consistent with the source image. This representation smoothed out the singularities introduced by the tendency of WSCLEAN to place point sources. The use of a sharp representation kernel here is not conventional and common practice consists in using the CLEAN beam, spreading more evenly the information and hence loosing precision. CLEAN artifacts are still noticeable with the presence of spurious bright negative areas, even after convolution.

For completeness, we also provide the images after adding the dirty residual in Fig. E.1c (except on the source image). Interestingly, we observe that the reconstructed images (i.e., PolyCLEAN and WSCLEAN) seem to be visually closer to the convolved source in some areas. We can interpret this as indicating that the residuals still contained some relevant information.