



Heriot-Watt University
Research Gateway

Image reconstruction algorithms in radio interferometry: From handcrafted to learned regularization denoisers

Citation for published version:

Terris, M, Dabbech, A, Tang, C & Wiaux, Y 2023, 'Image reconstruction algorithms in radio interferometry: From handcrafted to learned regularization denoisers', *Monthly Notices of the Royal Astronomical Society*, vol. 518, no. 1, pp. 604-622. <https://doi.org/10.1093/mnras/stac2672>

Digital Object Identifier (DOI):

[10.1093/mnras/stac2672](https://doi.org/10.1093/mnras/stac2672)

Link:

[Link to publication record in Heriot-Watt Research Portal](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Monthly Notices of the Royal Astronomical Society

Publisher Rights Statement:

© The Author(s) 2022. Published by Oxford University Press on behalf of The Royal Astronomical Society.

General rights

Copyright for the publications made accessible via Heriot-Watt Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

Heriot-Watt University has made every reasonable effort to ensure that the content in Heriot-Watt Research Portal complies with UK legislation. If you believe that the public display of this file breaches copyright please contact open.access@hw.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Image reconstruction algorithms in radio interferometry: From handcrafted to learned regularization denoisers

Matthieu Terris,[★] Arwa Dabbech[✉], Chao Tang and Yves Wiaux[★]

Institute of Sensors, Signals and Systems, Heriot-Watt University, Edinburgh EH14 4AS, UK

Accepted 2022 September 12. Received 2022 July 26; in original form 2022 February 28

ABSTRACT

We introduce a new class of iterative image reconstruction algorithms for radio interferometry, at the interface of convex optimization and deep learning, inspired by plug-and-play methods. The approach consists in learning a prior image model by training a deep neural network (DNN) as a denoiser, and substituting it for the handcrafted proximal regularization operator of an optimization algorithm. The proposed AIRI (‘AI for Regularization in radio-interferometric Imaging’) framework, for imaging complex intensity structure with diffuse and faint emission from visibility data, inherits the robustness and interpretability of optimization, and the learning power and speed of networks. Our approach relies on three steps. First, we design a low dynamic range training data base from optical intensity images. Secondly, we train a DNN denoiser at a noise level inferred from the signal-to-noise ratio of the data. We use training losses enhanced with a non-expansiveness term ensuring algorithm convergence, and including on-the-fly data base dynamic range enhancement via exponentiation. Thirdly, we plug the learned denoiser into the forward–backward optimization algorithm, resulting in a simple iterative structure alternating a denoising step with a gradient-descent data-fidelity step. We have validated AIRI against CLEAN, optimization algorithms of the SARA family, and a DNN trained to reconstruct the image directly from visibility data. Simulation results show that AIRI is competitive in imaging quality with SARA and its unconstrained forward–backward-based version uSARA, while providing significant acceleration. CLEAN remains faster but offers lower quality. The end-to-end DNN offers further acceleration, but with far lower quality than AIRI.

Key words: techniques: image processing – techniques: interferometric.

1 INTRODUCTION

Data acquisition by interferometry in radio astronomy relies on an array of antennas measuring an incomplete coverage of the spatial Fourier domain of the image of interest, yielding an ill-posed inverse problem towards image formation. Modern and upcoming radio telescopes are designed to bring unprecedented resolution and sensitivity. In this context, the algorithms to be deployed to solve the inverse imaging problem face an ever increasing requirement to jointly deliver precision (i.e. high resolution and dynamic range), robustness (i.e. endowed with calibration and uncertainty quantification functionalities), and scalability (i.e. the capability to process sheer data volumes).

Most RI imaging pipelines utilized by astronomers are based on the CLEAN algorithm, originally proposed by Högbom (1974). Assuming a sparse sky model, this greedy algorithm iteratively removes the contribution of each point source in the dirty image (more generally known as the backprojected image, defined as the inverse Fourier transform of the data). Numerous extensions of CLEAN have been devised over the last 50 yr (e.g. Schwarz 1978; Schwab & Cotton 1983; Bhatnagar & Cornwell 2004; Cornwell, Golap & Bhatnagar 2008; Thompson, Moran & Swenson 2017). Albeit simple and computationally efficient, CLEAN-based methods often require extensive

manual intervention to ensure their stability. Furthermore, CLEAN restored images often exhibit suboptimal imaging quality. On the one hand, their resolution is limited to that of the observations due to the convolution with a smoothing beam. On the other hand, their sensitivity is limited to the noise level, through the addition of the residual image.

Since their first inception by Wiaux et al. (2009), convex and non-convex optimization algorithms underpinned by sparsity priors have emerged in radio interferometric (RI) imaging (Li, Cornwell & de Hoog 2011; Dabbech et al. 2015; Garsden et al. 2015; Onose et al. 2016; Onose, Dabbech & Wiaux 2017; Repetti et al. 2017; Dabbech et al. 2018; Pratley et al. 2018; Repetti, Pereyra & Wiaux 2018, 2019; Birdi, Repetti & Wiaux 2020; Thouvenin et al. 2022a, b) and have led to a much higher image precision than that offered by advanced versions of CLEAN. In this framework, a so-called objective function is defined, typically as the sum of a data fidelity term and a regularization term injecting a prior image model to compensate for data incompleteness. The image estimate is defined as the minimizer of this objective and is reached via provably convergent algorithms. The obtained solution can also be understood in a Bayesian framework as a maximum a posteriori (MAP) estimate with respect to a posterior distribution, the negative logarithm of which is the objective. Uncertainty quantification approaches fully powered by optimization algorithms have also been proposed (Repetti et al. 2018, 2019). The versatility of the framework has progressively

[★] E-mail: m.terris@hw.ac.uk (MT); y.wiaux@hw.ac.uk (YW)

imposed optimization algorithms as a cornerstone of current state-of-the-art techniques in RI imaging. We note that they have also proven their worth for image reconstruction in the related fields of optical interferometry (Thiébaud & Young 2017) and very long baseline interferometry (VLBI) (Akiyama et al. 2019, 2022). However, albeit robust and interpretable, optimization algorithms come at the expense of a significant increase of the computational cost over CLEAN, ultimately affecting their scalability to large data volumes.

Bayesian inference approaches have also been proposed to address the RI image formation problem (Cai, Pereyra & McEwen 2018; Arras et al. 2019), also for VLBI (Akiyama et al. 2022), naturally enabling uncertainty quantification, but remaining computationally very demanding for the data volumes expected by modern radio telescopes.

Recent works have contemplated the use of deep neural networks (DNNs) for end-to-end reconstruction or post-processing in RI imaging (Terris et al. 2019; Gheller & Vazza 2021; Connor et al. 2022). Despite promising results and tremendous scalability capabilities, these approaches have neither been thoroughly validated for high resolution high dynamic range imaging of complex structure involving diffuse and faint emission, nor compared with CLEAN or state-of-the-art optimization approaches. Moreover, end-to-end DNNs are known to be subject to robustness issues (Goodfellow, Shlens & Szegedy 2015; Nguyen, Yosinski & Clune 2015; Pang et al. 2018). Their exploitation raises questions, not only with regards to the interpretability of the obtained solution, specifically in a Bayesian context, but also their generalizability, i.e. their ability to provide accurate reconstruction quality in the case of acquisition conditions unseen during training.

In this work, we propose a new RI image reconstruction approach based on the Plug-and-Play (PnP) framework, whereby a prior image model is learned by training a DNN as a denoiser, and substituted for the so-called proximal operator enforcing regularization at the heart of proximal optimization algorithms (Venkatakrishnan, Bouman & Wohlberg 2013; Chan, Wang & Elgendy 2016; Zhang et al. 2017; Zhang, Zuo & Zhang 2019). PnP algorithms were shown to deliver outstanding image reconstruction quality in applications such as image restoration (Zhang et al. 2021) and medical imaging (Ahmad et al. 2020). Importantly, as the DNN is trained as a denoiser on an image data base to serve as a simple regularization operator, it is by construction applicable for any sensing procedure, *de facto* avoiding the generalizability issue affecting end-to-end approaches. Leveraging monotone operator theory, recent works have shown that these algorithms provably converge to well characterized solutions under non-expansiveness conditions on the denoiser, yielding similar interpretations to that of MAP estimators (Cohen, Elad & Milanfar 2021; Pesquet et al. 2021). In summary, this versatile framework simultaneously inherits the robustness and interpretability of optimization approaches, and the learning power and scalability capabilities of DNNs.

Building on Pesquet et al. (2021), we propose a PnP framework dubbed AIRI, standing for ‘AI for Regularization in radio-interferometric Imaging’. Focusing on monochromatic intensity imaging, our approach for developing a first AIRI algorithm capable of imaging complex structure with diffuse and faint emission can be summarized in three steps. The first step is to design, from publicly available optical images, a realistic but low dynamic range data base of intensity images for supervised training. The second step is to train a DNN denoiser with basic architecture ensuring positivity of the reconstructed images, at the signal-to-noise ratio of the data. The training relies on a simple ℓ_2 or ℓ_1 loss enhanced with a firm non-expansiveness term ensuring algorithm convergence, and on an on-

the-fly exponentiation procedure for dynamic range enhancement. The third step resides in plugging the DNN into the FB optimization algorithm, resulting in a simple iterative structure alternating between a denoising step and a gradient-descent data-fidelity step.

We finally validate the resulting AIRI- ℓ_2 and AIRI- ℓ_1 algorithms, implemented in Matlab, on high resolution high dynamic range simulations utilizing intensity images containing diffuse and faint emission across the field of view. Our test images are 3c353, Hercules A, Centaurus A, and Cygnus A, with size 512×512 . Our benchmark algorithms are (i) the multiscale CLEAN version implemented in the C++ WSCLEAN software package (Offringa & Smirnov 2017), (ii) two optimization algorithms from the SARA family leveraging the handcrafted ‘average sparsity’ proximal regularization operator: SARA itself (Onose et al. 2017) and its unconstrained FB-based version uSARA that we introduce here, and (iii) a UNet trained in an end-to-end fashion to reconstruct target images from dirty images. While SARA relies on an advanced primal-dual FB (PDFB) algorithm (Pesquet & Repetti 2015) enabling non-differentiable data-fidelity constraints, uSARA only differs from AIRI- ℓ_2 and AIRI- ℓ_1 by the use of the average sparsity proximal operator in lieu of a learned denoiser for regularization. It in fact corresponds to the standalone version of the imaging module of the joint calibration and imaging approach proposed in Repetti et al. (2017) and Dabbech et al. (2021). Results show that AIRI- ℓ_2 and AIRI- ℓ_1 are already competitive with SARA and uSARA in terms of imaging quality, while providing a significant acceleration. In a nutshell, AIRI does indeed inherit the robustness and interpretability of optimization approaches, and the learning power and speed of DNNs. The WSClean code remains significantly faster but offers lower reconstruction quality. The UNet offers further acceleration, but with far lower quality than AIRI.

The remainder of the paper is organized as follows. In Section 2, we review imaging approaches from proximal optimization, recall SARA, and introduce its unconstrained version uSARA. In Section 3, we review PnP algorithms through the prism of proximal algorithms and introduce AIRI. In Section 4, we study the performance of AIRI- ℓ_2 and AIRI- ℓ_1 through extensive simulation in comparison with the benchmark algorithms. In Section 5, we draw our conclusions and discuss necessary future work.

2 STATE-OF-ART OPTIMIZATION ALGORITHMS AND RI

2.1 RI imaging problem

Aperture synthesis in radio astronomy probes the sky by measuring an incomplete coverage of the spatial Fourier domain of the image of interest through an array of antennas. Focusing on monochromatic intensity imaging, assuming a narrow field of view, and in the absence of atmospheric and instrumental perturbations, each pair of antennas acquires a noisy Fourier component of the intensity image to be formed, called a visibility. The associated Fourier mode (also called uv -point) is given by the projection of the corresponding baseline, expressed in units of the observation wavelength, on to the plane perpendicular to the line of sight (Thompson et al. 2017). A discrete formulation of the resulting linear RI image formation problem, aiming to restore a target intensity image $\bar{x} \in \mathbb{R}^n$ from the measured complex visibilities $y \in \mathbb{C}^m$, reads (Onose et al. 2016)

$$y = \Phi \bar{x} + e, \quad (1)$$

where $\Phi = \mathbf{GFZ} \in \mathbb{C}^{m \times n}$ is the measurement operator, $\mathbf{G} \in \mathbb{C}^{m \times d}$ is a sparse interpolation matrix, encoding the non-uniform Fourier

transform, $\mathbf{F} \in \mathbb{C}^{d \times d}$ is the 2D Discrete Fourier Transform, $\mathbf{Z} \in \mathbb{R}^{d \times n}$ is a zero-padding operator, incorporating the correction for the convolution performed through the operator \mathbf{G} , and $\mathbf{e} \in \mathbb{C}^m$ is a realization of some i.i.d. Gaussian random noise, with zero mean and standard deviation $\tau > 0$. We refer to Appendix A for considerations regarding more general noise distributions, and the inclusion of so-called direction-dependent effects (DDEs) in \mathbf{G} .

We also note that backprojecting problem (1) into the image domain gives

$$\text{Re}\{\Phi^\dagger \mathbf{y}\} = \text{Re}\{\Phi^\dagger \Phi\} \bar{\mathbf{x}} + \text{Re}\{\Phi^\dagger \mathbf{e}\}, \quad (2)$$

where $(\cdot)^\dagger$ denotes the complex conjugate transpose and $\text{Re}\{\cdot\}$ denotes the real part. $\text{Re}\{\Phi^\dagger \mathbf{y}\}$ is known as the dirty image, $\text{Re}\{\Phi^\dagger \Phi\}$ is the operator representing the convolution of $\bar{\mathbf{x}}$ by the point spread function, also known as the dirty beam, and $\text{Re}\{\Phi^\dagger \mathbf{e}\}$ is the noise backprojected in the image domain.

We finally emphasize that in the optical interferometry and VLBI contexts, imaging is often performed from closure quantities derived from the visibilities, bringing a non-linear inverse problem affected by non-Gaussian noise (Akiyama et al. 2019, 2022). Imaging in the presence of DDEs or from closure quantities lies beyond the scope of this work.

2.2 Proximal algorithms and FB

Leaving the nature of Φ aside, problems of the likes of (1) are ubiquitous in imaging sciences, and arise for instance in image restoration (Levin et al. 2009; Yang et al. 2010; Bredies & Holler 2020), hyperspectral imaging (Wang et al. 2015; Xie et al. 2019), and medical imaging (Gupta et al. 2018; Zbontar et al. 2018; Fessler 2020), to name a few. A widespread approach to solve such a problem is to reformulate it as a convex minimization problem

$$\underset{\mathbf{x} \in \mathbb{R}^n}{\text{minimize}} f(\mathbf{x}) + \lambda r(\mathbf{x}), \quad (3)$$

where $f(\mathbf{x}) + \lambda r(\mathbf{x})$ is called the objective function, $f \in \Gamma_0(\mathbb{R}^n)$ ¹ is the term enforcing data-fidelity, and $r \in \Gamma_0(\mathbb{R}^n)$ is the regularization term introduced to address the ill-posedness of (1) by enforcing a prior image model.

The theory of optimization offers a myriad of algorithms to solve such minimization problems, ranging from the simple Forward-Backward (FB) algorithm (Bauschke & Combettes 2017), applicable if one of the two terms is differentiable, to ADMM (Boyd, Parikh & Chu 2011), or the Douglas-Rachford algorithm (Eckstein & Bertsekas 1992), applicable even if both f and r are non-differentiable, and more evolved structures, such as PDFB (Pesquet & Repetti 2015), that can handle multiterm objective functions with as many non-differentiable data-fidelity and regularization terms processed in parallel. All these algorithms are part of the same class of proximal algorithms, where differentiable and non-differentiable terms respectively translate in the algorithm into gradient and proximal operators.

Considering $r \in \Gamma_0(\mathbb{R}^n)$, its proximal operator prox_r is defined as

$$(\forall \mathbf{z} \in \mathbb{R}^n), \quad \text{prox}_r(\mathbf{z}) = \underset{\mathbf{u} \in \mathbb{R}^n}{\text{argmin}} \frac{1}{2} \|\mathbf{z} - \mathbf{u}\|^2 + r(\mathbf{u}), \quad (4)$$

and can be interpreted as a generalization of a projection operator, or interestingly, as a denoising operator. Indeed, problem (4) can be interpreted as the minimization problem to be solved for a simple

¹ $\Gamma_0(\mathbb{R}^n)$ denotes the set of convex, proper, and lower semicontinuous functions from \mathbb{R}^n to $(-\infty, \infty]$.

denoising problem of the form $\mathbf{z} = \mathbf{u} + \mathbf{w}$ where the data \mathbf{z} result from adding some i.i.d. Gaussian noise \mathbf{w} to the unknown \mathbf{u} . From this perspective, the first term of the objective function in (4) would be the standard data-fidelity term, given the Gaussian nature of the noise, and r the regularization term. A proximal operator is a denoiser!

The choice of f and r in (3) is of paramount importance as it influences the final solution to the minimization task (Mallat 1999; Selesnick, Baraniuk & Kingsbury 2005; Bredies & Holler 2020). Interpreting the reconstructed image as a MAP estimate leads naturally to choosing f as the negative log-likelihood associated with a given statistical model of the noise (typically a squared ℓ_2 norm for i.i.d. additive Gaussian noise), and r as the negative logarithm associated with a statistical prior image model (typically, ℓ_1 norm for a Laplace prior). However, optimization algorithms are not tied to this statistical interpretation. In the context of compressive sensing theory (Candès, Romberg & Tao 2006; Donoho 2006; Baraniuk 2007), where sparsity is the postulated signal model, the choice of the ℓ_1 norm as a regularization term simply emanates from it being the closest convex relaxation of the ℓ_0 norm that is the natural sparsity measure. ℓ_1 regularization was shown to yield state-of-the-art image reconstruction methods beyond the context of compressive sensing. For instance, sparsity of the gradient is enforced as the ℓ_1 norm of the magnitude of the image gradient, that is the so-called Total Variation semi-norm (TV) (Rudin, Osher & Fatemi 1992; Bredies, Kunisch & Pock 2010; Bredies & Holler 2020), whereas sparsity of the sought image in a sparsifying domain $\mathbf{A} \in \mathbb{R}^{n \times p}$ can be enforced by choosing $r(\mathbf{x}) = \|\mathbf{A}^\dagger \mathbf{x}\|_1$ where \mathbf{A} can be a wavelet dictionary (Mallat 1999), an x-let transform (Candès & Demanet 2003; Do & Vetterli 2003), or a learned dictionary (Mairal et al. 2009), to name a few. Generally, the choice of the data-fidelity term is often driven by the statistical nature of the measurement noise, while regularization terms are carefully handcrafted to meet the specificity of the applications, with sparsity models defining the state-of-the-art. In summary, the MAP interpretation has its limitation, and in particular, setting the regularization parameter λ according to the proper normalization of the statistical models of which f and r would be the negative logarithms, is known to be suboptimal.

The FB algorithm is a proximal algorithm designed to solve problems of the form (3), where $f, r \in \Gamma_0(\mathbb{R}^n)$ and f is differentiable. It is defined by the iterative sequence

$$(\forall k \in \mathbb{N}), \quad \mathbf{x}_{k+1} = \text{prox}_{\gamma \lambda r}(\mathbf{x}_k - \gamma \nabla f(\mathbf{x}_k)), \quad (5)$$

which is proven to converge towards a minimizer of the objective function (3), provided that $0 < \gamma < 2/L$ where L is the Lipschitz constant² of ∇f (Bauschke & Combettes 2017). Accelerated versions of FB, leveraging inertial terms, pre-conditioning, and stochastic approaches, also exist. We also emphasize that FB can be used to solve problems involving non-convex regularization terms r (Beck & Teboulle 2009; Attouch, Bolte & Svaiter 2013; Chouzenoux, Pesquet & Repetti 2014; Combettes & Pesquet 2015; Repetti & Wiaux 2021).

2.3 SARA and its unconstrained version

2.3.1 SARA

The SARA family represent state-of-the-art optimization algorithms for RI imaging (Carrillo, McEwen & Wiaux 2012, 2014; Abdaluziz

²The Lipschitz constant of an operator $h: \mathbb{R}^n \rightarrow \mathbb{R}^p$ is defined as $\sup_{\mathbf{x} \neq \mathbf{y}} \|h(\mathbf{x}) - h(\mathbf{y})\| / \|\mathbf{x} - \mathbf{y}\|$.

et al. 2016; Onose et al. 2016, 2017; Birdi, Repetti & Wiaux 2018; Dabbech et al. 2018; Pratley et al. 2018; Abdulaziz, Dabbech & Wiaux 2019; Thouvenin et al. 2022a). In its monochromatic intensity imaging version, SARA consists in solving the constrained minimization problem given by

$$\underset{\mathbf{x} \in \mathbb{R}^n}{\text{minimize}} \quad \iota_{\mathcal{B}(\mathbf{y}, \epsilon)}(\Phi \mathbf{x}) + \lambda r(\mathbf{x}), \quad (6)$$

where data fidelity is imposed via a constraint encoded by the indicator function³ $\iota_{\mathcal{B}(\mathbf{y}, \epsilon)}$ of the ℓ_2 -ball centred at the data \mathbf{y} and of radius $\epsilon > 0$: $\mathcal{B}(\mathbf{y}, \epsilon) = \{\mathbf{u} \in \mathbb{C}^m \mid \|\mathbf{y} - \mathbf{u}\| \leq \epsilon\}$. The ℓ_2 -bound ϵ is derived from the assumed i.i.d. Gaussian random noise statistics as $\epsilon^2 = (2m + 4\sqrt{m})\tau^2/2$ (Carrillo et al. 2012). The function r defines the SARA prior, also called the ‘average sparsity’ prior, which consists of a positivity constraint and a log-sum prior promoting average sparsity in an over-complete dictionary $\Psi \in \mathbb{R}^{n \times bn}$. More specifically, the dictionary Ψ is defined as the concatenation of $b = 9$ orthogonal bases (the first eight Daubechies wavelets and the Dirac basis), re-normalized by $\sqrt{b} = 3$ to ensure that $\Psi\Psi^\dagger = \mathbf{I}$, where \mathbf{I} denotes the identity operator. The SARA prior explicitly reads

$$r(\mathbf{x}) = \rho \sum_{j=1}^{bn} \log \left(\rho^{-1} \left| (\Psi^\dagger \mathbf{x})_j \right| + 1 \right) + \iota_{\mathbb{R}_+^n}(\mathbf{x}), \quad (7)$$

where $(\cdot)_j$ denotes the j th coefficient of its argument vector, and \mathbb{R}_+^n denotes the n -dimensional real positive orthant. The parameter $\lambda > 0$ is a regularization parameter, while $\rho > 0$ is used to avoid reaching zero values in the argument of the logarithmic terms. Thouvenin et al. (2022a) suggest to set both parameters to an estimate of the standard deviation of the measurement noise in the wavelet domain.

Minimizing (6) should typically involve the proximal operator prox_r , dubbed the ‘average sparsity’ proximal operator. However, if the log-sum prior enforces a stronger sparsity of the solution than a typical ℓ_1 prior, it is also non-convex. To address the resulting non-convex minimization task, a reweighting procedure is adopted, where a sequence of convex surrogate minimization problems is solved iteratively (Candes, Wakin & Boyd 2008; Carrillo et al. 2012), each involving a weighted- ℓ_1 prior g , given by

$$g(\mathbf{x}, \mathbf{W}) = \|\mathbf{W}\Psi^\dagger \mathbf{x}\|_1 + \iota_{\mathbb{R}_+^n}(\mathbf{x}), \quad (8)$$

where $\mathbf{W} \in \mathbb{R}^{bn \times bn}$ is a diagonal weighting matrix that needs to be updated after each resolution of the surrogate problem. The full SARA algorithm for solving (6) can be summarized as

$$\begin{aligned} & \text{for } i = 0, 1, \dots \\ & \left[\begin{aligned} \tilde{\mathbf{x}}_{i+1} &= \underset{\mathbf{x} \in \mathbb{R}^n}{\text{argmin}} \quad \iota_{\mathcal{B}(\mathbf{y}, \epsilon)}(\Phi \mathbf{x}) + g(\mathbf{x}, \mathbf{W}_i) \\ \mathbf{W}_{i+1} &= \text{Diag} \left(\rho / (\rho + \Psi^\dagger \tilde{\mathbf{x}}_{i+1}) \right), \end{aligned} \right. \end{aligned} \quad (9)$$

with $\mathbf{W}_0 = \mathbf{I}$, where $\text{Diag}(\cdot)$ denotes the diagonal matrix containing its (vector) argument on the diagonal. Minimizing each of these problems will now involve the weighted- ℓ_1 prior g via its proximal operator prox_g . One should acknowledge that the reweighting strategy also introduces significant complexity to the algorithm since a sequence of convex minimization problems needs to be solved.

Given the non-differentiability of the data-fidelity term in (6), the PDFB algorithm (Pesquet & Repetti 2015) is leveraged to solve each minimization task of the form (9). We note that because Ψ^\dagger is overcomplete and due to the presence of the positivity constraint in (8), the proximal operator prox_g does not admit a closed-form solution, thus a priori requiring a sub-iterative structure. However,

³The indicator function of a non-empty, closed, and convex set \mathcal{S} is defined as $\iota_{\mathcal{S}}(\mathbf{x}) = 0$ if $\mathbf{x} \in \mathcal{S}$, else $\iota_{\mathcal{S}}(\mathbf{x}) = +\infty$.

the full splitting functionalities of PDFB enable the decomposition of g into functions admitting simple proximal operators (a component-wise projection on the real positive orthant and a component-wise thresholding operation). These proximal operators are handled in parallel by PDFB without the need for sub-iterations (Onose et al. 2016). Interestingly, the same functionalities enable handling large data volumes via parallel processing of data blocks (Onose et al. 2016), and large image sizes via parallel processing of image facets (Thouvenin et al. 2022a), providing significant scalability to SARA.

2.3.2 Unconstrained SARA (uSARA)

Building on the works of Repetti & Wiaux (2020, 2021), we propose to focus on the unconstrained formulation of the SARA problem, that we dub unconstrained SARA (uSARA), again looked at from the prism of monochromatic intensity RI imaging only. More precisely, we consider the problem (3), with a differentiable data-fidelity term f chosen as the mean squared error loss, corresponding to the negative log-likelihood of the data under the i.i.d. Gaussian random noise assumption, given by

$$f(\mathbf{x}) = \frac{1}{2} \|\Phi \mathbf{x} - \mathbf{y}\|^2. \quad (10)$$

Adopting the prior model r described in (7), the resulting non-convex minimization task writes

$$\underset{\mathbf{x} \in \mathbb{R}^n}{\text{minimize}} \quad \frac{1}{2} \|\Phi \mathbf{x} - \mathbf{y}\|^2 + \lambda r(\mathbf{x}), \quad (11)$$

where $\lambda > 0$ is a regularization parameter. The non-convexity of the prior r from (7) is handled through a similar reweighting procedure as in (9):

$$\begin{aligned} & \text{for } i = 0, 1, \dots \\ & \left[\begin{aligned} \tilde{\mathbf{x}}_{i+1} &= \underset{\mathbf{x} \in \mathbb{R}^n}{\text{argmin}} \quad \frac{1}{2} \|\Phi \mathbf{x} - \mathbf{y}\|^2 + \lambda g(\mathbf{x}, \mathbf{W}_i) \\ \mathbf{W}_{i+1} &= \text{Diag} \left(\rho / (\rho + \Psi^\dagger \tilde{\mathbf{x}}_{i+1}) \right), \end{aligned} \right. \end{aligned} \quad (12)$$

with $\mathbf{W}_0 = \mathbf{I}$. At each reweighting iteration, problem (12) can be solved with a standard FB algorithm. The full reweighting procedure for solving the uSARA problem (11) is summarized in Algorithm 1, where $\xi_1 > 0$ is a relative variation convergence criterion, and where the gradient of f at \mathbf{x} reads as

$$\nabla f(\mathbf{x}) = \text{Re}\{\Phi^\dagger \Phi\} \mathbf{x} - \text{Re}\{\Phi^\dagger \mathbf{y}\}, \quad (13)$$

and its Lipschitz constant is the spectral norm⁴ of $\text{Re}\{\Phi^\dagger \Phi\}$, i.e.

$$L = \|\text{Re}\{\Phi^\dagger \Phi\}\|_{\text{S}}. \quad (14)$$

We recall that the stepsize γ in (5) is to be upper-bounded by $2/L$.

As appears in Algorithm 1 solving (12) (steps 5–7), the weighted- ℓ_1 prior g is involved via $\text{prox}_{\gamma \lambda g}$. The algorithmic structure does not exhibit the same full splitting functionalities as PDFB, and $\text{prox}_{\gamma \lambda g}$ must be computed iteratively. This can typically be achieved via a so-called dual FB algorithm (Combettes & Pesquet 2011), detailed in Algorithm 2, where $\xi_2 > 0$ is a relative variation convergence criterion. This algorithm alternates between a projection on the real positive orthant (Step 3), denoted by $\Pi_{\mathbb{R}_+^n}$, and the proximal operator of the dual of the weighted- ℓ_1 prior, which involves a component-wise soft-thresholding operator $\text{prox}_{\gamma \lambda \|\mathbf{W}\|_1}$ ⁵ (Step 4).

⁴The spectral norm of a linear operator is its maximum singular value, which can be computed via the power method (see e.g. Golub & Van Loan 2013).

⁵This operator, applied to \mathbf{z} boils down component-wise to $(\text{prox}_{\gamma \lambda \|\mathbf{W}\|_1}(\mathbf{z}))_j = \text{prox}_{\eta_j | \cdot|}(z_j) = \text{sign}(z_j) \max\{|z_j| - \eta_j, 0\}$ where

Formally, Algorithms 1 and 2 lead to a triply sub-iterative structure: a proximal operator loop (Step 2), inside a weighted- ℓ_1 loop (Step 5), inside a reweighting loop (Step 3). This suggests a very computationally expensive structure. None the less, first, Repetti & Wiaux (2021) have shown that only a fixed number of FB iterations K are required (Step 2), corresponding to the approximate minimization of (12), while preserving the convergence of the overall algorithm to a minimizer of the non-convex objective (11). This contrasts with SARA, where the subproblems (9) need to be solved to convergence. The value of K can in fact be optimized to provide significant acceleration of Algorithm 1, effectively removing one iteration layer. Secondly, the number of iterations in Algorithm 2 is moderate in practice when using an appropriate initialization strategy for the dual variable v_0 , leading to mild computational cost of $\text{prox}_{\gamma\lambda g}$.

We conclude this section by underlining that uSARA corresponds to the imaging module of the joint calibration imaging approach described in Repetti et al. (2017) and Dabbech et al. (2021).

3 AIRI: AI FOR REGULARIZATION IN RI IMAGING

3.1 PnP-FB

When solving an inverse imaging problem from an optimization theory viewpoint, one defines an objective function, of which the sought image would be a minimizer and which can be obtained via an iterative algorithm. The PnP approach (Venkatakrishnan et al. 2013) looks at the problem directly through the lens of the algorithm. As discussed in Section 2.2, it follows from the definition of a proximal operator that it can be interpreted as a denoiser. Borrowing a proximal optimization algorithm, the PnP approach proposes to replace the proximal regularization operator with a more general denoiser. We here note that this procedure can be applied to a wide class of algorithms, ranging from FB to PDFB. Specifically, assuming a minimization problem of the form (3) with the differentiable data fidelity term f , recall that the FB algorithm reads as in (5), with the proximal operator (4). Its PnP counterpart (PnP-FB) then simply follows as

$$(\forall k \in \mathbb{N}), \quad \mathbf{x}_{k+1} = \text{D}(\mathbf{x}_k - \gamma \nabla f(\mathbf{x}_k)), \quad (15)$$

where D is a denoising operator, i.e. an operator specifically designed to remove i.i.d. Gaussian random noise from an image.

Denoisers are ubiquitous in imaging sciences, and a large variety of denoisers have been defined and studied in the signal processing literature, including Gaussian filters, BM3D (Dabov et al. 2007), non-local means (NLM) (Buades, Coll & Morel 2011), with, more recently, deep neural networks (DNNs) leading the state-of-the-art (Zhang et al. 2017). It is therefore very tempting to plug powerful denoisers in lieu of proximal operators... and play. In this context, we aim at taking advantage of the learning capabilities of DNNs in order to learn an appropriate regularization denoiser. Just as handcrafting the regularization term r (from which the proximal regularization operator prox_r results) is crucial to the reconstruction quality for a pure optimization approach, learning a powerful denoiser is cornerstone for the reconstruction quality of PnP algorithms. Recent advances in DNNs for image denoising tasks have shown new state-of-the-art results over traditional denoisers (Zhang et al. 2017; Wang

$\eta_j = \gamma \lambda W_{j,j}$ is called the soft-thresholding parameter. All values below the threshold are set to 0 while those above are reduced by the value of the threshold (in absolute value).

et al. 2018; Zhou et al. 2020). Simultaneously, PnP algorithms with DNNs as denoisers have become the new state-of-the-art in image reconstruction (Zhang et al. 2021), significantly improving over their traditional optimization counterparts.

Yet, replacing the proximal operator with an off-the-shelf denoising operator is not inconsequential. In particular, a general non-proximal denoiser can a priori not be related with a regularization term in some overarching objective function. As a consequence, the characterization of the PnP solution (e.g. as a MAP estimator), and worse, the convergence properties of the iterative structure, are not ensured anymore, thus questioning the robustness and interpretability of PnP solutions. To overcome that issue, we follow the approach of Pesquet et al. (2021) allowing to enforce the so-called ‘firm non-expansiveness constraint’ on the denoiser D during training. This constraint ensures that D ‘contracts distances’ and yields both the convergence of Algorithm 3 and the characterization of the limit point (see Appendix B for more details). The PnP-FB algorithm with established convergence guarantees is summarized in Algorithm 3, where $\xi_3 > 0$ is a relative variation convergence criterion. In the RI imaging case of interest, the PnP counterpart to uSARA follows from defining f as in (10), with the measurement operator Φ from (1). The gradient of f at \mathbf{x} therefore reads as in (13), with Lipschitz constant given in (14).

3.2 Training noise level versus uSARA regularization parameter

3.2.1 Denoising i.i.d. Gaussian noise of specific standard deviation

Building from Pesquet et al. (2021), our AIRI denoiser D will be trained in a supervised manner to remove i.i.d. Gaussian random noise of specific standard deviation σ . In other words, D is set to tackle the denoising problem

$$\mathbf{z} = \mathbf{u} + \sigma \mathbf{w}, \quad (16)$$

where \mathbf{z} denotes the noisy data input to the DNN, \mathbf{u} is the unknown image corresponding to the target output of the DNN, and $\mathbf{w} \sim \mathcal{N}(0, \mathbf{I})$. As discussed in the next section, the input signal-to-noise ratio of the RI data provides a reliable handle on σ .

3.2.2 Equating training and target dynamic ranges

As already emphasized, AIRI will emerge from PnP-FB with f from (10) and Φ from (1). We propose to set the training noise level σ for D to an appropriate estimate of the standard deviation of some effective image-domain noise induced by the original i.i.d. Gaussian random noise of standard deviation τ on the Fourier data in (1). We propose below a procedure to estimate the standard deviation of this image-domain noise, which, for simplicity, we also approximate to be i.i.d. Gaussian random noise.

First, we resort to the image-domain formulation (2) of (1), normalized by the Lipschitz constant L in (14): $\text{Re}\{\Phi^\dagger \mathbf{y}\}/L = \text{Re}\{\Phi^\dagger \Phi\} \bar{\mathbf{x}}/L + \text{Re}\{\Phi^\dagger \mathbf{e}\}/L$. By construction, the chosen normalization ensures that the convolution operator $\text{Re}\{\Phi^\dagger \Phi\}/L$ has unit spectral norm, so that the dirty image is at the same scale as the original image. The covariance matrix of the image-domain noise $\text{Re}\{\Phi^\dagger \mathbf{e}\}/L$ thus writes $\tau^2 \text{Re}\{\Phi^\dagger \Phi\}/2L^2$. Secondly, as proposed in Thouvenin et al. (2022a), we discard the correlation structure of the image-domain noise via the approximation $\text{Re}\{\Phi^\dagger \Phi\} \simeq L\mathbf{I}$, leading to a noise covariance matrix approximation $\tau^2 \mathbf{I}/2L$. As a result the effective image-domain noise is assumed i.i.d. Gaussian, and the original standard deviation τ in the measurement domain is rescaled

to $\tau/\sqrt{2L}$ in the transfer to the image domain. Equating the standard deviations of the training noise for D and the image-domain noise leads to the following heuristic:

$$\sigma = \frac{\tau}{\sqrt{2L}}. \quad (17)$$

We note that this heuristic is algorithm agnostic, i.e. a priori independent of the algorithmic structure in which the denoiser is plugged. In this sense, while similar, it is simpler and more general than the one proposed in Pesquet et al. (2021), which was derived specifically by estimating σ around the fixed point of (15). As will be shown through simulations in Section 4, the proposed heuristic provides a very precise estimate of the optimal training noise level.

Assuming the validity of this heuristic, one must acknowledge a dependence of the training noise level σ on the measurement model, via τ and L , suggesting that a denoiser should be trained independently for each target reconstruction. Also, we consider a normalized data base of images whose maximum intensity values are lower or equal to 1 (see Section 3.3), so that the denoising performance cannot be ascertained on non-normalized images. These two generalizability questions can be solved by a simple rescaling of the original inverse problem (1) to simultaneously fit the normalization constraint (an upper bound on the peak intensity of the image being accessible from the dirty image) and rescale the noise level in the data to the training noise level of some already available denoiser. Dabbech et al. (2022) and Wilber et al. (in preparation) build from these considerations and propose both a ‘single denoiser’ approach and a more flexible ‘denoiser shelf’ approach, where denoisers are respectively trained at a single predefined noise level or on a limited number of such noise levels, to which the target inverse problems can be rescaled at will. It also appears from the above rescaling considerations that the dimension of σ is that of an inverse input image-domain peak signal-to-noise ratio rather than an absolute noise level. As the AIRI denoiser consistently categorizes σ -level values as noise to be removed, the input image-domain signal-to-noise ratio also naturally sets the target output peak signal-to-noise ratio or dynamic range of the reconstruction. In summary, our heuristic simply stipulates that the training dynamic range should be set equal to the target reconstruction dynamic range.

3.2.3 uSARA regularization parameter

We further note the very similar roles of the noise level σ of the regularization denoiser in PnP-FB and the regularization parameter λ in the FB algorithm (5) solving (3). At the algorithmic level, on the one hand, D is a denoiser and we have proposed to set σ to the estimate (17) of the effective image-domain noise level, or more generally to the inverse target reconstruction dynamic range value. On the other hand, for the traditional FB algorithm, λ is involved in (5) via $\text{prox}_{\gamma\lambda}$. Focusing on the specific uSARA instance described in Section 2.3.2, and summarized on Algorithms 1 and 2, $\gamma\lambda$ typically acts as a soft-thresholding parameter, thresholding out small values in the wavelet domain defined by the average sparsity dictionary. Setting its value to the standard deviation of the effective wavelet-domain noise, which according to Thouvenin et al. (2022a) is the same as in the image domain given the unit spectral norm of the average sparsity dictionary Ψ , yields the heuristic:

$$\gamma\lambda = \frac{\tau}{\sqrt{2L}}. \quad (18)$$

We readily note that, in contrast with the reasoning leading to (17), the above reasoning for λ is specific to the FB structure underpinning

uSARA. We also recall the parameter $\rho > 0$ in the definition of the average sparsity prior r in (7), used to avoid reaching zero values in the argument of the logarithmic terms. This parameter represents a floor level for the wavelet coefficients (Carrillo et al. 2012; Thouvenin et al. 2022a), and should naturally be set to the same effective wavelet-domain noise level estimate as λ : $\rho = \lambda$. We further note that a variation of the proposed heuristic for the wavelet-domain noise level arises from a slightly different approximation in the transfer of the noise level from the image domain to the wavelet domain. Instead of assuming noise level conservation due to the normalization of Ψ , one can analyse the noise in each of the nine bases of the dictionary separately, discarding the noise correlation structure between each pair of bases. In this case, the orthonormality of each basis and its normalization by three suggests a correction factor 1/3 in the heuristic values: $\rho = \gamma\lambda \simeq \tau/3\sqrt{2L}$.

As will be shown in Section 4, relation (18) provides a useful reference to set λ in uSARA. Even though the 1/3 correction will be shown to bring a closer-to-optimal value, it is still not as precise an estimate as (17) is for the training noise level of AIRI denoisers. We emphasize that a significant body of work has been dedicated to the question of setting the regularization parameter in problems of the form (3), with no simple and accurate solution to our knowledge (Donoho & Johnstone 1995; Luisier, Blu & Unser 2007; Vidal et al. 2020). In this context, the fact that (18) provides an appropriate reference value for λ in uSARA is an interesting result. However, the fact that the PnP-FB admits a very accurate heuristic to set the training noise level, avoiding the requirement for fine-tuning an arbitrary regularization parameter, represents a major advantage of AIRI over uSARA.

3.3 Training data base with adaptive dynamic range

3.3.1 State-of-the-art and proposed approach

Supervised training of DNNs for imaging requires the definition of a data base of images, which remains a main difficulty in translating deep learning methods to RI imaging due to the absence of groundtruths. The few works that have leveraged deep learning techniques for RI imaging so far resorted to various approaches. In Terris et al. (2019), we used a high dynamic range simulated radio map from Bonaldi & Braun (2018) of size $32\,000 \times 32\,000$ containing mainly elliptical Gaussian sources, broken down into a data base of 1100 images of size 512×512 . Gheller & Vazza (2021) relied on 1000 radio images of size 2000×2000 simulated with ENZO (Bryan et al. 2014) and augmented by random rotations, resulting in the appearance of realistic extended emission in the field of view. Connor et al. (2022) generated a data base of 900 synthetic images of size 2000×2000 containing elliptical Gaussian sources.

In this work, we are targeting high dynamic range imaging of complex structure with diffuse and faint emission across the field of view. The data base to be designed for the training of the regularization denoiser thus needs to reflect this target. We propose (i) to source publicly available optical astronomy images containing a rich combination of compact and diffuse emission, (ii) to preprocess these images, aiming to remove existing noise and artefacts via a preliminary denoising procedure, and (iii) to exponentiate their intensity values to create high dynamic range images. The procedure results in the creation of a rich and clean training data base of 2235 training images of size 512×512 . More technical details are provided below.

3.3.2 Building a rich low dynamic range data base

First, we constituted a set of 32 large grey-scale images ($\mathbf{u}_k^{\text{raw}}$) $_{1 \leq k \leq 32}$ from optical astronomical intensity images available online, containing complex structure with both compact and diffuse emission. Each image is normalized with maximum intensity value equal to 1, and when the image contained colour channels, those were averaged. The typical size of each of these images is 3500×3500 . The sample standard deviation of the residual noise in image regions with no signal, averaged over the 32 images reads $\sigma_0 \simeq 1/64$, leading to a typical dynamic range of $\sigma_0^{-1} \simeq 64$, orders of magnitude below the values of interest in modern RI observations.

Secondly, the images are preprocessed with the aim of removing the residual noise, which includes various compression artefacts, and would otherwise be learned as part of the image model during training. To that effect, for all $1 \leq k \leq 32$, a simple denoising problem is formulated for the recovery of a clean image $\bar{\mathbf{u}}_k$ from $\mathbf{u}_k^{\text{raw}}$: $\mathbf{u}_k^{\text{raw}} = \bar{\mathbf{u}}_k + \mathbf{e}_k$, assuming i.i.d. Gaussian random noise \mathbf{e}_k with standard deviation $\hat{\sigma}_k$. The denoising procedure applied consists in solving the optimization problem underpinning SARA in (6), with $\Phi = \mathbf{I}$ and $\mathbf{y} = \mathbf{u}_k^{\text{raw}}$, resulting in a denoised data base ($\mathbf{u}_k^{\text{low}}$) $_{1 \leq k \leq 32}$. We note that this procedure does not affect the (low) dynamic range of the images. In a nutshell, noise and artefacts are removed, but the floor signal level and maximum intensity remain essentially unchanged.

Finally, the cleaned images are split to create a rich and clean data base ($\mathbf{u}_s^{\text{low}}$) $_{1 \leq s \leq S}$ with $S = 2235$ images of size 512×512 containing both compact and diffuse emission. When the image dimensions are not multiples of 512, a symmetric padding is used before splitting. The resulting images exhibit a whole distribution of maximum intensity values bounded by 1, and a corresponding distribution of dynamic range values, with $\sigma_0^{-1} \simeq 64$ representing the ‘nominal’ dynamic range value of the data base.

3.3.3 Enhancing data base dynamic range to target

In order to simulate the high dynamic range of interest in RI imaging, we propose to exponentiate pixel-wise the intensity value of all groundtruth images of the low dynamic range data base. The exponentiation parameter will be set to achieve a final nominal dynamic range of the same order as the training dynamic range, which according to (17) is also equivalent to the target dynamic range of the reconstruction. Technically, the low dynamic range images $\mathbf{u}_s^{\text{low}}$ are exponentiated through pixel-wise as:

$$\mathbf{u}_s = h(\mathbf{u}_s^{\text{low}}, a) = \frac{a^{\mathbf{u}_s^{\text{low}}} - 1}{a}, \quad (19)$$

for some parameter $a \gg 1$. Recalling that the maximum span of the pixel values in any $\mathbf{u}_s^{\text{low}}$ is $[\sigma_0, 1]$, the resulting maximum span of the exponentiated images is approximately $[a^{-1}(a^{\sigma_0} - 1), 1]$. The exponentiation thus preserves the data base normalization with maximum intensity value across images bounded by 1, while providing control on the final nominal dynamic range $a(a^{\sigma_0} - 1)^{-1}$. Imposing this value to be equal to the training dynamic range leads to the following equation for a :

$$a = (1 + a\sigma)^{\sigma_0^{-1}}. \quad (20)$$

Parametrizing a as $a = b\sigma^{-1}$ leads to the following equation for the scaling factor b : $b = \sigma(1 + b)^{\sigma_0^{-1}}$. This equation can be solved numerically to estimate the exponentiation factor a as a function of σ_0 and σ . In summary, the high dynamic range groundtruth data base (\mathbf{u}_s) $_{1 \leq s \leq S}$ results from the low dynamic range groundtruth data base ($\mathbf{u}_s^{\text{low}}$) $_{1 \leq s \leq S}$ via (19), with exponentiation parameter a set in (20).

Fig. 1 summarizes the pre-processing pipeline for the creation of the low and high dynamic range data bases. Notice the diversity of images in the training samples displayed in Fig. 1(c).

3.3.4 Generating noisy and groundtruth pairs on-the-fly

The supervised training approach considered relies on image pairs: the groundtruth images (\mathbf{u}_s) $_{1 \leq s \leq S}$ corresponding to the target outputs of the DNN, and the noisy images (\mathbf{z}_s) $_{1 \leq s \leq S}$ corresponding to the DNN input. Given our choice, described in Section 3.2.1, to build a denoiser tailored to tackle i.i.d. Gaussian random noise of specific standard deviation σ , the noisy high dynamic range images read

$$\mathbf{z}_s = \mathbf{u}_s + \sigma \mathbf{w}_s, \quad (21)$$

with $\mathbf{w}_s \sim \mathcal{N}(0, \mathbf{I})$ and σ given by (17).

Importantly, as the high dynamic range data base of noisy and groundtruth pairs depends on the target dynamic range through (19) and (21), one version of the data base would need to be computed and stored for each dynamic range of interest, which is a highly impractical and unnecessary approach. Instead, \mathbf{z}_s and \mathbf{u}_s can simply be seen as functions of $\mathbf{u}_s^{\text{low}}$ and computed on the fly during training.

3.4 Training loss function

In this section, we introduce the training loss for the supervised training of the denoiser D. Building on Pesquet et al. (2021) and following technical arguments detailed in Appendix B, we regularize the training loss with an appropriate Jacobian term to enforce the firm non-expansiveness of D. By denoting $\mathbf{Q} = 2\mathbf{D} - \mathbf{I}$ where \mathbf{I} denotes the identity operator, and denoting $\theta \in \mathbb{R}^c$ the learnable parameters of D, the resulting training loss reads:

$$\underset{\theta \in \mathbb{R}^c}{\text{minimize}} \frac{1}{S} \sum_{s=1}^S \left(\mathcal{L}(\mathbf{D}_\theta(\mathbf{z}_s) - \mathbf{u}_s) + \kappa \max\{\|\nabla \mathbf{Q}_\theta(\mathbf{z}_s)\|_S, 1 - \varepsilon\} \right), \quad (22)$$

with \mathbf{z}_s and \mathbf{u}_s given in (19) and (21), respectively, as functions of $\mathbf{u}_s^{\text{low}}$ computed on the fly during training. The second term in (22) is the firm non-expansiveness regularization term, with $\kappa > 0$ the associated regularization parameter, and $\varepsilon > 0$ a safety margin parameter. The first term \mathcal{L} is the standard part of the loss function, for which practical choices involve the ℓ_2 loss i.e. $\mathcal{L}(\cdot) = 1/2 \|\cdot\|^2$, and the ℓ_1 loss, i.e. $\mathcal{L}(\cdot) = \|\cdot\|_1$; both will be studied in Section 4.

4 SIMULATIONS AND RESULTS

4.1 Overall simulation set up

4.1.1 Test images

Our test image data base consists of four radio images, namely 3c353, Hercules A, Centaurus A, and Cygnus A, all resized to a size of $n = 512 \times 512$, shown in Fig. 2. These images all have a peak value equal to 1 and dynamic ranges above 10^4 . The simulations involving these images will target reconstruction dynamic ranges of the order of 10^4 . We emphasize that this is orders of magnitude higher than typical values in the PnP literature (Zhang et al. 2017, 2021; Pesquet et al. 2021), and therefore significantly more challenging.




Figure 1. Summary of the proposed strategy for building a training data base of post-processed radio groundtruths. We first gather a set of 32 astronomical (optical) images of high quality but low dynamic range: (a) shows a subset of four images. (b) illustrates the effect of pre-processing on a small part of the image from (a). Eventually, we split the images into 2235 images of size 512×512 to create the low dynamic range training data base. Samples of our training data base, after application of the exponentiation (19) with $a = 10^3$, are shown in (c) in logarithmic scale.




Figure 2. The four test RI images, displayed in logarithmic scale.

4.1.2 Benchmark methods

We benchmark AIRI with (i) uSARA (Algorithm 1) solving the unconstrained problem (11), (ii) SARA (Onose et al. 2017; Thouvenin et al. 2022a), solving the constrained problem (6), (iii) the multiscale variant of CLEAN implemented in the WSCLEAN software (Offringa & Smirnov 2017), and (iv) a UNet (Jin et al. 2017) trained end-to-end on a large data base to recover groundtruths from a data base of dirty images.

All simulations involving uSARA rely on a value for ρ using (18) including the 1/3 correction factor, while the values of λ are fine-tuned around (18) to optimize the reconstruction quality. As for SARA, we rely on the prescriptions of Thouvenin et al. (2022a) for setting its hyperparameters. While all other methods use natural weighting of the data, CLEAN makes use of uniform weighting, which relies on the density of the Fourier sampling for enhanced resolution

(see Section 2.1 and Appendix A).⁶ We underline that the CLEAN restored image is in units of Jansky per beam area, in contrast with SARA-based methods, AIRI, and the UNet, where the associated estimated images are in the units of Jansky per pixel. Therefore, in what follows, CLEAN images are rescaled by the area of the CLEAN beam, for visualization purposes. Finally, we note that we always show the dirty image computed with a uniform weighting scheme, and rescaled in the $[0,1]$ range for visualization purposes.

Regarding the end-to-end deep learning approach, we use a Deep Residual UNet from Zhang et al. (2021) and train it to reconstruct a target image in (1) from its dirty version in (2). We generate a data base $(\mathbf{u}_s, \text{Re}\{\Phi_s^\dagger \mathbf{y}_s\})_{1 \leq s \leq S}$, where for every s , \mathbf{y}_s represents the visibility vector computed from the groundtruth \mathbf{u}_s and for a measurement operator $\Phi_s = \mathbf{G}_s \mathbf{F} \mathbf{Z}$ generated from a random pointing direction in the sky. The randomization ensures that the UNet training data base contains a rich and diverse set of sampling patterns. We borrow the procedure detailed in Section 3.3.3 to enhance the nominal dynamic range of the groundtruth data base $(\mathbf{u}_s)_{1 \leq s \leq S}$, and equate it to the target dynamic range of reconstruction. The UNet is then trained to minimize the ℓ_1 loss between the groundtruth \mathbf{u}_s and the reconstruction $\text{UNet}(\text{Re}\{\Phi_s^\dagger \mathbf{y}_s\})$. The UNet is trained for 800 epochs with the Adam algorithm.

4.1.3 Evaluation metrics

For AIRI- ℓ_2 , AIRI- ℓ_1 , uSARA, SARA, and the UNet, the image estimates $\hat{\mathbf{x}}$ are the reconstructed model images, while the image estimate considered for CLEAN is the restored image. The performance of all algorithms is assessed qualitatively via inspection of $\hat{\mathbf{x}}$, alongside the residual images, defined as $\hat{\mathbf{r}} = \beta \Phi^\dagger (\mathbf{y} - \Phi \hat{\mathbf{x}})$, where β is a normalization factor.⁷ Given the high dynamic ranges of interest, the

⁶The specific command used is: `wsclean -multiscale -niter 30000 -weight uniform -gain 0.1 -mgain 0.6 -threshold 0.001 -no-reorder -minuvw-m 0.001 -auto-mask 0.001 -mem 1.`

⁷The normalization factor is given by $\beta = 1/\max_i(\Phi^\dagger \Phi \delta)_i$, where δ is the image with value 1 at the phase centre and 0 otherwise. Hence, the dirty beam given by $\beta \Phi^\dagger \Phi \delta$, has a peak value equal to 1.




Figure 3. Architecture of the proposed denoising DNN, inspired from Zhang et al. (2017) and Pesquet et al. (2021), containing 20 convolutional layers in total. ‘Conv’ denotes convolutional layers. The number of channels is given above the arrows. Notice the ReLU layer at the output, ensuring that our reconstructed RI images are positive.

estimated images are shown in logarithmic scale, using the mapping $\text{rlog} = \mathbf{x} \mapsto \log_{10}(10^3 \mathbf{x} + 1)/3$. The residual images are displayed in linear scale.

We also consider two quantitative evaluation metrics, namely the SNR and logSNR of the estimated images $\hat{\mathbf{x}}$ with respect to the groundtruth $\bar{\mathbf{x}}$. The standard SNR metric (in dB) is defined as $\text{SNR}(\hat{\mathbf{x}}, \bar{\mathbf{x}}) = 20 \log_{10}(\|\bar{\mathbf{x}}\|/\|\bar{\mathbf{x}} - \hat{\mathbf{x}}\|)$. It simply corresponds to a logarithmic version of the norm of the error image $\bar{\mathbf{x}} - \hat{\mathbf{x}}$. The logSNR (in dB) is the SNR evaluated on a logarithmic version of the intensity images. This metric is introduced given the high dynamic ranges of interest, and assigns more relative weight to low intensities than the SNR. Our definition is $\text{logSNR}(\hat{\mathbf{x}}, \bar{\mathbf{x}}) = \text{SNR}(\text{rlog}(\hat{\mathbf{x}}), \text{rlog}(\bar{\mathbf{x}}))$.

4.1.4 DNN architecture and training particulars

First, as highlighted already, the denoising losses that we use for \mathcal{L} in our training loss (22) are chosen to be either the ℓ_2 loss or the ℓ_1 loss, giving rise to AIRI- ℓ_2 and AIRI- ℓ_1 .

Secondly, we choose a basic DNN architecture for the denoiser D, in the form of a modified DnCNN (Zhang et al. 2017) where the batch normalization layers have been removed, as in Pesquet et al. (2021). But in this work, a ReLU has been added as the last layer. This is a simple way to enforce positivity at the output of the denoiser, and therefore of the final reconstruction, much necessary as the developed AIRI algorithm targets intensity (versus polarization) imaging (see Fig. 3).

Thirdly, the DNN is trained for 10^5 epochs with the Adam algorithm (Kingma & Ba 2015) and learning rate 10^{-4} , the learning rate being divided by 2 every 2×10^4 epochs. $(\mathbf{z}_s, \mathbf{u}_s)$ in (22) are of size 46×46 , generated according to (21) and (19) from patches randomly extracted from the low dynamic range data base of 512×512 images $(\mathbf{u}_s^{\text{low}})_{1 \leq s \leq S}$, itself augmented with random rotations and zooms. The Jacobian regularization in (22) is computed with a power method with a number of iterations increased to 10 in the later stages of the training to improve the precision of the computation. As observed in Pesquet et al. (2021), starting the optimization of the network from a pretrained state with $\kappa = 0$ in (22) enhances the results significantly.

Finally, we have observed that performing, at each iteration of the PnP algorithm, random flips and 90 degrees rotations of the image before applying the DNN, and inverting the transform after the denoising step, can significantly boost the reconstruction quality. We therefore adopt this strategy, inspired by Zhang et al. (2021).

4.2 Experiment 1: validating the denoiser for denoising

4.2.1 Simulation setup

The PnP approach suggests that the best denoisers should provide the best regularizers. In this first experiment, we compare the denoising

capabilities of the AIRI- ℓ_2 and AIRI- ℓ_1 denoisers on one side, and the average sparsity proximal operator propelling uSARA and SARA on the other side. To this end, we focus on a denoising problem

$$\mathbf{y} = \bar{\mathbf{x}} + \mathbf{e}, \quad (23)$$

where $\mathbf{y} \in \mathbb{R}^n$ is the data, $\bar{\mathbf{x}} \in \mathbb{R}^n$ is the groundtruth image, and $\mathbf{e} \in \mathbb{R}^n$ is the realization of an i.i.d. Gaussian random noise with zero mean and standard deviation τ . In other words, (23) is a degenerate case of problem (1) with $\Phi = \mathbf{I}$. We consider the 3c353 image as a groundtruth $\bar{\mathbf{x}}$. As the following experiments for the full validation of AIRI- ℓ_2 and AIRI- ℓ_1 operate at a target dynamic range around 10^4 , we design this preliminary pure denoising experiment with a similarly challenging target dynamic range, and set the observation noise level to $\tau = 10^{-4}$.

First, problem (23) can be solved using AIRI via Algorithm 3. As the image peaks at 1, no rescaling of the inverse problem in (1) is needed (see Section 3.2). In this simple case, the gradient in (13) reads $\nabla f(\mathbf{x}) = \mathbf{x} - \mathbf{y}$ with Lipschitz constant in (14) $L = 1$. Taking $\gamma = 1/L = 1$, Algorithm 3 trivially boils down to a simple non-iterative application of D to the data, leading to the denoised image

$$\hat{\mathbf{x}} = \text{D}(\mathbf{y}). \quad (24)$$

The AIRI- ℓ_2 and AIRI- ℓ_1 denoisers were trained with a noise level $\sigma = \tau = 10^{-4}$ as per (17), on the proposed synthetic training data base, and following the procedure explained in Section 3.4. Following the procedure detailed in Section 3.3.3, the chosen exponentiation parameter for on-the-fly dynamic range enhancement of the data base is $a = 10^3$. Furthermore, we set the regularization parameter in the training loss (22) to $\kappa = 10^{-9}$ and $\kappa = 10^{-5}$ for AIRI- ℓ_2 and AIRI- ℓ_1 denoisers respectively. The safety margin parameter is set to $\varepsilon = 5 \times 10^{-2}$. These values are chosen here as they are the ones to ensure the stability of Algorithm 3 in our next experiments.

Secondly, problem (23) can also be solved in a proximal approach with uSARA, setting $\Phi = \mathbf{I}$. Interestingly, in this case, the uSARA objective boils down to the objective defining the average sparsity proximal operator (see definition 4) of λr , for the average sparsity prior r in (7), and $\lambda = \tau$ following the reasoning around equation (18). Therefore, applying Algorithm 1 results in computing the denoised image by simply applying the average sparsity proximal operator with appropriate λ :

$$\hat{\mathbf{x}} = \text{prox}_{\lambda r}(\mathbf{y}). \quad (25)$$

We note, given $\nabla f(\mathbf{x}) = \mathbf{x} - \mathbf{y}$ and $L = 1$, and choosing $K = 1$ and $\gamma = 1/L = 1$, that Algorithm 1 simplifies to the sequential application of $\text{prox}_{\gamma \lambda g}(\mathbf{y}, \mathbf{W}_i)$ with only the weights being updated at each iteration. The algorithm was run with convergence criteria to $\xi_1 = 6 \times 10^{-6}$ Algorithm 1, and $\xi_2 = 10^{-4}$ in Algorithm 2.

4.2.2 Experimental results

Denoising results are displayed in Fig. 4. All three denoisers appear to be highly effective, both in terms of SNR and logSNR and visual denoising quality. The SNR values are very comparable. In terms of logSNR the AIRI- ℓ_1 denoiser achieves better performance than the AIRI- ℓ_2 denoiser, itself superior to the average sparsity proximal operator. Visually, the AIRI- ℓ_1 denoiser also achieves the reconstruction with the best resolution and least amount of artefacts (see zooms).

These results suggests that the AIRI- ℓ_1 and AIRI- ℓ_2 denoisers do encapsulate a superior prior model to the average sparsity model encapsulated in the corresponding proximal operator.

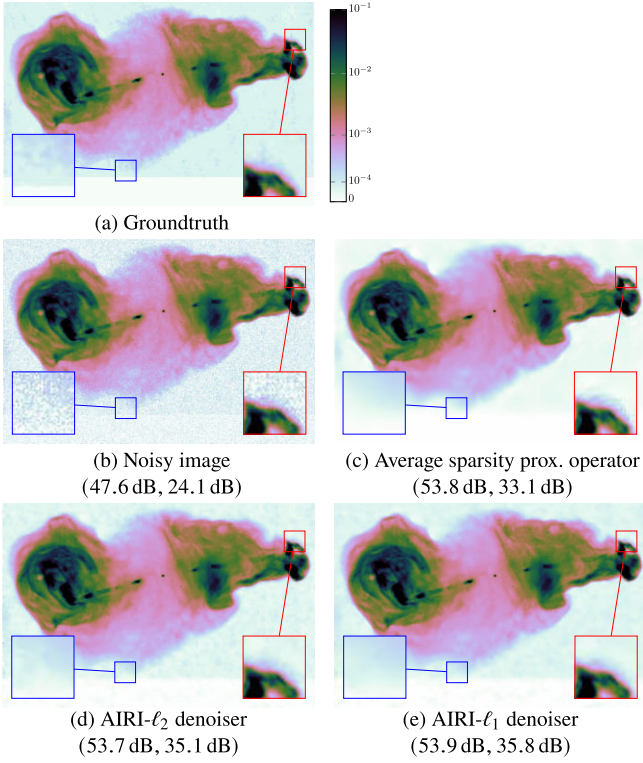


Figure 4. Experiment 1 results: Denoising results obtained with the proposed AIRI- ℓ_2 and AIRI- ℓ_1 denoisers in comparison with the averaged sparsity proximal operator. The respective groundtruth and noisy images are displayed in 4(a) and (b). Denoising results obtained with the averaged sparsity proximal operator are shown in 4(c), with AIRI- ℓ_2 in 4(d) and with AIRI- ℓ_1 in 4(e). All images are displayed in logarithmic scale. The colourbar is saturated at 10^{-1} for clarity. Values of the obtained evaluation metrics (SNR, logSNR) are indicated below their associated images.

4.3 Experiment 2: validating the training noise level heuristic

4.3.1 Simulation setup

Using our four 512×512 test images (see Fig. 2), we study the impact of the noise level σ involved in the training of the AIRI- ℓ_2 and AIRI- ℓ_1 denoisers according to the procedure set in Section 3, with the aim to validate the heuristic (17). We also study the optimal value of the regularization parameter λ of uSARA around the heuristic (18). The performance of AIRI- ℓ_2 and AIRI- ℓ_1 in terms of reconstruction quality is compared to that of uSARA.

We simulate RI observations utilizing simulated uv -coverages of the radio telescope MeerKAT. The uv -patterns correspond to five randomly selected pointing directions displayed in Figs 5 and 6(c), and two of the four total observation durations displayed in Fig. 6: $\Delta T \in \{4 \text{ h}, 8 \text{ h}\}$. The observed wavelength is 0.3 m and the simulated sampling rate of the telescope is fixed to 100 points per hour. The data size thus increases linearly with ΔT . Data are simulated following the model described in (1), with input SNR defined as $\text{iSNR} = 20 \log_{10}(\|\Phi \bar{x}\|/\tau) = 30 \text{ dB}$. We finally note that the uv -patterns extend to the edge of the considered Fourier domain as shown in Figs 5 and 6. In other words, we target mild, if any, superresolution with respect to the nominal resolution of the observation, as set by the largest baseline.

Since our groundtruth images have peak values at 1, no rescaling of the inverse problem in (1) is needed for AIRI, and a common heuristic value is reached for σ in (17) and $\gamma\lambda$ in (18). Obviously, τ




Figure 5. Experiments 1–3: MeerKAT uv -coverages simulated for four different telescope pointing directions and fixed total observation duration $\Delta T = 4 \text{ h}$. The components are measured in units of the wavelength of the observation. Each uv -pattern contains $m = 806400$ points.




Figure 6. Experiments 1–3: MeerKAT uv -coverages associated with the four observation durations and fixed telescope pointing direction (complementary to the four shown in Fig. 5). The numbers of points in the uv -patterns are (a) $\Delta T = 1 \text{ h}$: 201600, (b) $\Delta T = 2 \text{ h}$: 403200, (c) $\Delta T = 4 \text{ h}$: 806400, and (d) $\Delta T = 8 \text{ h}$: 1612800.

and L are functions of Φ , and therefore of both the pointing direction and total duration of observation. However, the heuristic values are within a 25 per cent variation range across pointing directions for each ΔT . For simplicity, a single heuristic value is considered for all five uv -patterns associated with the same ΔT , taken to be the average of the values resulting for each pointing direction. The two resulting values are as announced associated with a target dynamic range around 10^4 : $\nu_4 = 1.4 \times 10^{-4}$ for $\Delta T = 4 \text{ h}$, and

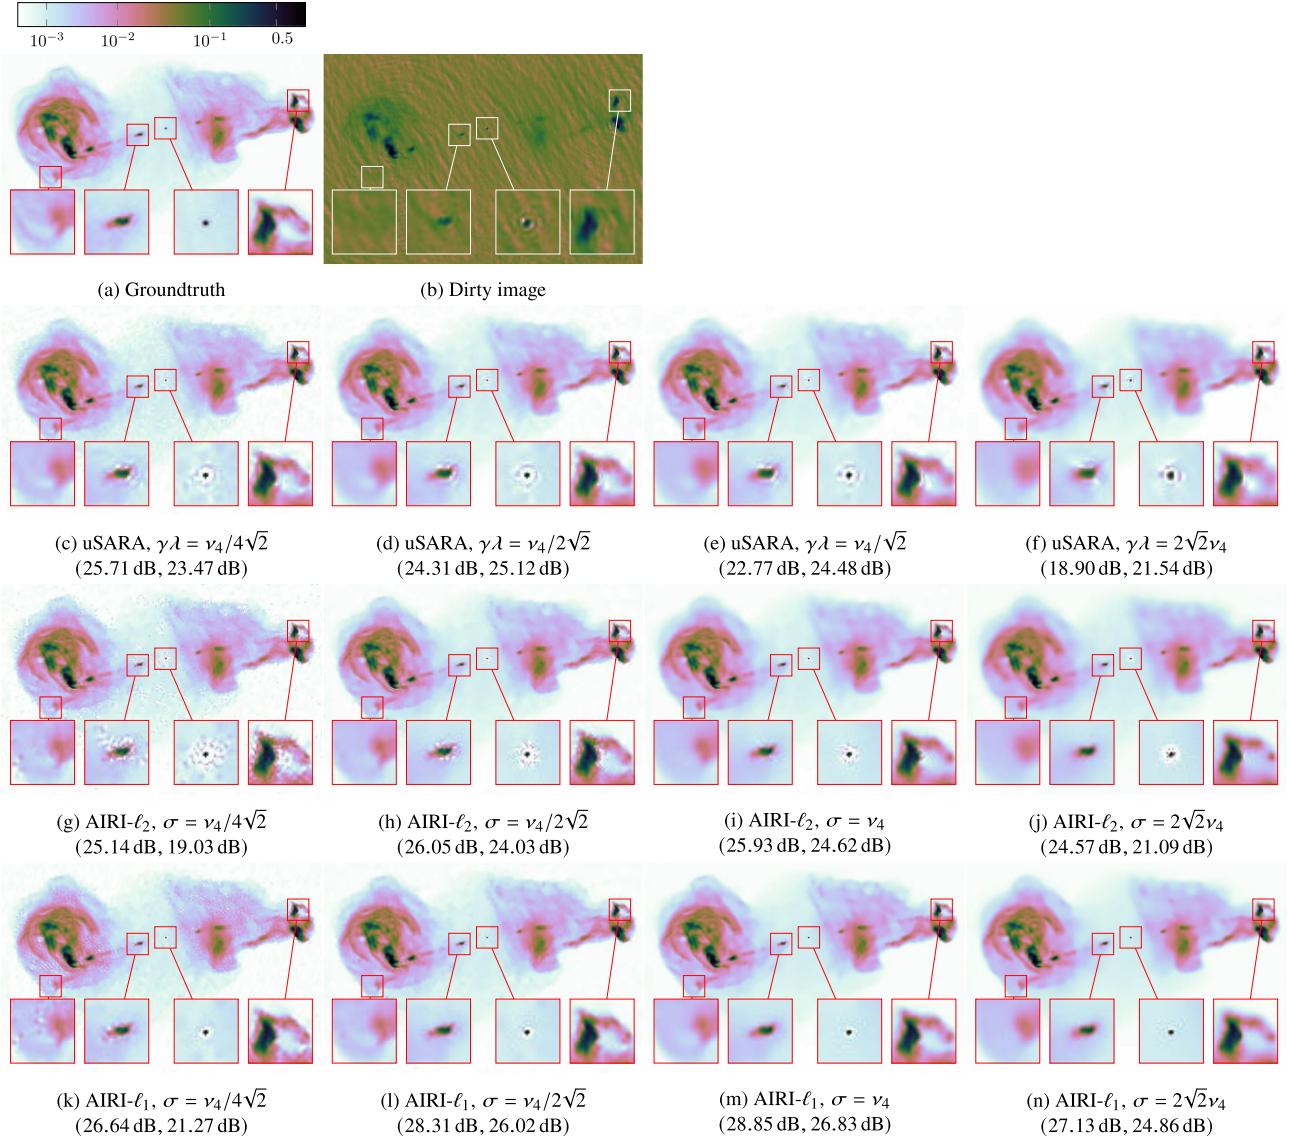


Figure 7. Experiment 2 results: Impact of the training noise level σ (resp. regularization parameter λ) in imaging with AIRI (resp. uSARA) on simulated RI data using 3c353 image as a groundtruth and the uv -pattern from Fig. 5(a) ($\Delta T = 4$ h) in comparison with the reference value ν_4 suggested by the AIRI and uSARA heuristics. Top row: the groundtruth image 7(a) and the simulated dirty image 7(b). Second row: estimated model images obtained with uSARA, with a regularization parameter λ increasing with the column index. Third (resp. fourth) row: estimated model images of AIRI- ℓ_2 (resp. AIRI- ℓ_1), with a training noise level σ increasing with the column index. Below each image we indicate the reconstruction metrics as (SNR, logSNR). Images are shown in logarithmic scale.

$\nu_8 = 9.3 \times 10^{-5}$ for $\Delta T = 8$ h. The values probed for $\gamma\lambda$ and σ are sampled a factor $\sqrt{2}$ around the heuristic and by steps of 2 thereafter.

Following the study by Pesquet et al. (2021), we choose the value of $\gamma = 1.98/L$ in Algorithms 1 and 3, which saturates the theoretical bound. The convergence criteria of Algorithms 1 and 3 are set to $\xi_1 = \xi_3 = 5 \times 10^{-6}$, with a maximum number of iterations of 6×10^3 , and we set $\xi_2 = 10^{-5}$ in Algorithm 2.

We train our AIRI- ℓ_2 and AIRI- ℓ_1 denoisers with optimized values $\kappa = 10^{-9}$ and $\kappa = 10^{-5}$, respectively, in the training loss (22), with $\varepsilon = 5 \times 10^{-2}$, and observe that this ensures the stability of Algorithm 3 despite a spectral norm of the Jacobian slightly above 1. Following the procedure detailed in Section 3.3.3, the exponentiation parameter for on-the-fly dynamic range enhancement of the data base is set to $a = 10^3$.

4.3.2 Experimental results

Considering the image of 3c353 as groundtruth and the uv -pattern from Fig. 5(a) ($\Delta T = 4$ h) for simulating the measurements, reconstructions for different values of σ in AIRI- ℓ_2 and AIRI- ℓ_1 , and λ in uSARA are displayed in Figs 7 and 8. More precisely, model images, displayed in logarithmic scale, and their associated SNR and logSNR values are provided in Fig. 7, and residual images are displayed in Fig. 8 in linear scale. We first notice that σ and λ play a similar role in the reconstruction quality: the higher their value, the smoother the recovered image; the lower their value, the stronger the artefacts. Overall, reconstructions with uSARA exhibit some wavelet artefacts, particularly noticeable around bright compact sources overlaying faint extended emission. AIRI- ℓ_2 reconstructions typically exhibit less artefacts, but appear to be generally smoother.

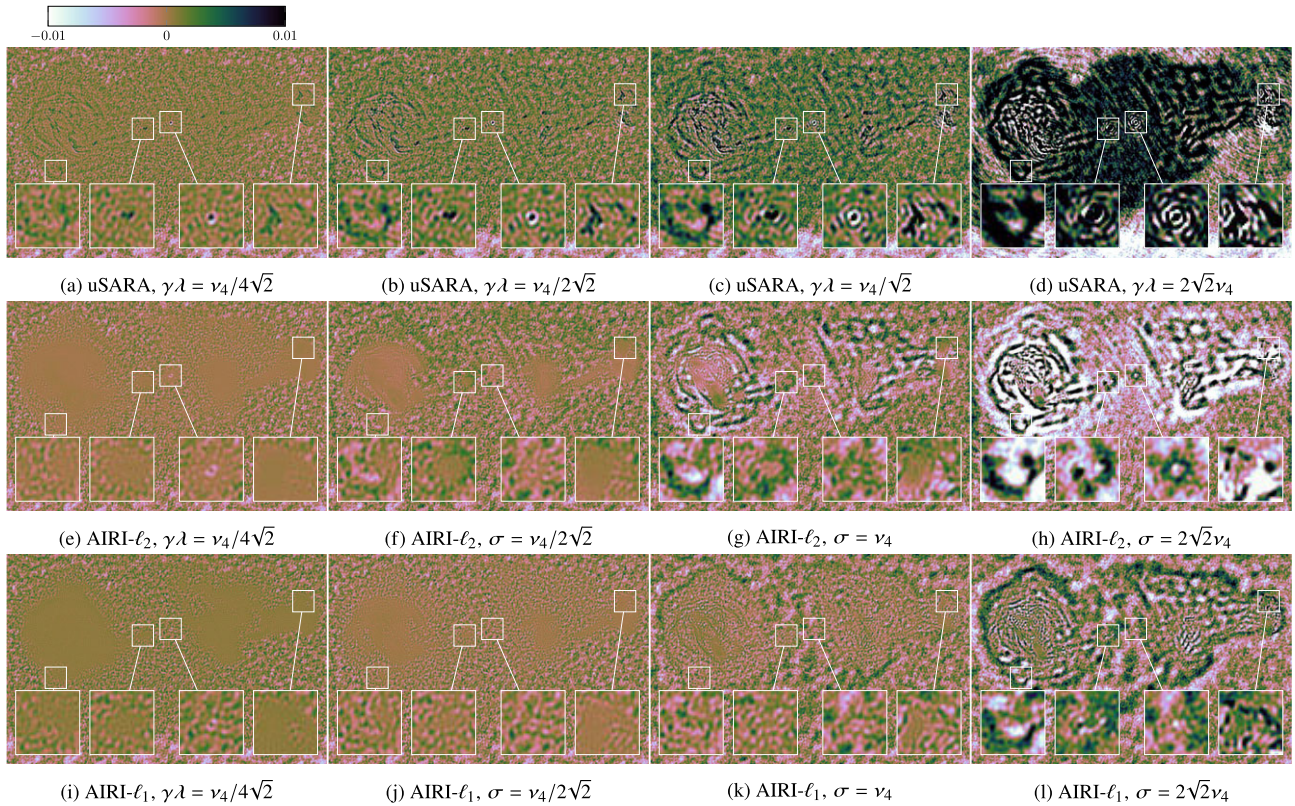


Figure 8. Experiment 2 results: Residual images associated with the reconstructions of Fig. 7. Top row: residual images for the estimated model images obtained with uSARA, with a regularization parameter λ increasing with the column index. Second (resp. third) row: respective residual images for the estimated model images obtained with AIRI- ℓ_2 (resp. AIRI- ℓ_1), with a training noise level σ increasing with the column index. All images are shown in linear scale.

AIRI- ℓ_1 reconstructions contain even less artefacts, with an achieved resolution similar to uSARA. We conclude these observations by highlighting the low-amplitude residual images obtained with low values of σ and λ . However, such results are not a token of a good quality reconstruction. They instead reflect the overfitting of the data in such settings.

Fig. 9 shows the reconstruction SNR and logSNR graphs obtained when varying σ and $\gamma\lambda$ for reconstructions with AIRI and uSARA, respectively, for both the total observation durations $\Delta T = 4$ h and $\Delta T = 8$ h. Each point on each of the four graphs gives the average value and 95 per cent confidence interval over 20 reconstructions arising when considering the four test images in Fig. 2 and five different pointing directions from Figs 5 and 6(c).

First, for both AIRI- ℓ_2 and AIRI- ℓ_1 , we observe that the SNR and logSNR metrics clearly peak at the proposed heuristic (17) (shown as the dashed vertical black line) for both $\Delta T = 4$ h and $\Delta T = 8$ h, suggesting that equating the training dynamic range to the target dynamic range is an accurate procedure. This result is also confirmed by the visual analysis in Fig. 7. For uSARA, the SNR and logSNR results and the visual analysis in Fig. 7 suggest that the relation (18) is a useful reference, but rather provides an upper-bound on the optimal value for the soft-thresholding parameter, roughly 6 to 11 times smaller when looked at from an SNR standpoint, or 3 to 6 times smaller from a logSNR perspective. In general, we find that the logSNR provides better coherence with the visual analysis, suggesting a optimal soft-thresholding value 3 to 6 times below the reference. Interestingly, the 1/3 correction factor brings a closer-to-optimal value, but does not solve the discrepancy between the SNR and logSNR metrics. In contrast to AIRI- ℓ_2

and AIRI- ℓ_1 , uSARA thus exhibits a regularization parameter of which the value cannot be set simultaneously automatically and optimally. This represents a significant advantage of AIRI over uSARA.

Secondly, considering each method at its best, SNR and logSNR values for both $\Delta T = 4$ h and $\Delta T = 8$ h confirm a superior performance of AIRI- ℓ_1 over both AIRI- ℓ_2 and uSARA, the latter two achieving similar reconstruction qualities, with virtually same logSNR, and a slight SNR advantage for AIRI- ℓ_2 .

Thirdly, these results are in line with the visual analysis when considering each method at its best (see Figs 7d, i, m, and 8b, g, k). AIRI- ℓ_2 and uSARA exhibit similar reconstruction qualities with the former, smoother but with less artefacts than the latter. AIRI- ℓ_1 exhibits less artefacts than AIRI- ℓ_2 , while preserving the resolution offered by uSARA. AIRI- ℓ_1 provides the best residuals, uSARA residuals clearly contain sources discarded by the model and relating to the artefacts, while AIRI- ℓ_2 's higher residuals correlate with the suboptimal resolution.

Fourthly, we recall that AIRI- ℓ_2 , AIRI- ℓ_1 , and uSARA are underpinned by the same FB structure, and only differ by their choice of denoiser. From this perspective, the results of the current experiment are in line with those of the previous experiment, with a better regularization denoiser naturally leading to better overall reconstruction.

Finally, we observe that the confidence intervals around the mean SNR and logSNR values in Fig. 9 are rather small for both AIRI- ℓ_2 and AIRI- ℓ_1 , and similar to those of uSARA. This is a first illustration that in the PnP approach there is no issue of generalizability to measurement conditions not considered during training as, by

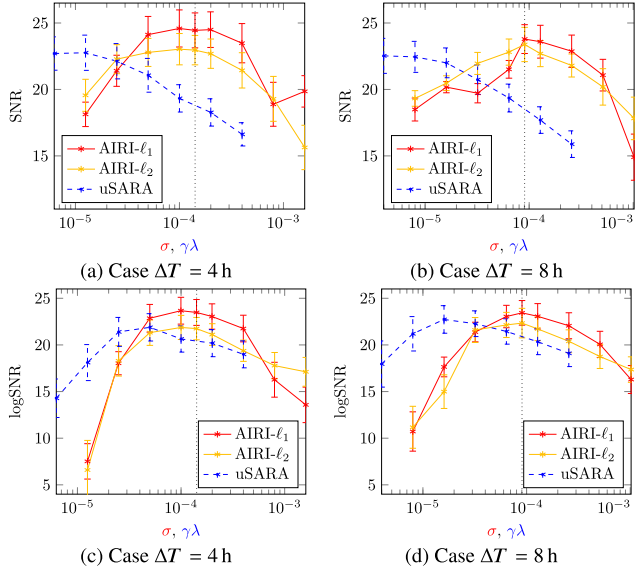


Figure 9. Experiment 2 results: Reconstruction metrics as a function of the training noise level σ in the case of AIRI- ℓ_2 and AIRI- ℓ_1 , or the thresholding parameter $\gamma\lambda$ in the case of uSARA. Top row: reconstruction SNR in the case of $\Delta T = 4$ h (a) and $\Delta T = 8$ h (b). Bottom row: reconstruction logSNR in the case of $\Delta T = 4$ h (c) and $\Delta T = 8$ h (d). Each point is an average over the 20 simulated observations built from the four different groundtruths from Fig. 2 and the five different uv -patterns shown in Figs 5 and 6(c). Error bars show the 95 per cent confidence interval. On each graph, the black vertical dotted line indicates the common heuristic value for σ and $\gamma\lambda$ (i.e. ν_4 for $\Delta T = 4$ h and ν_8 for $\Delta T = 8$ h). Values for σ and $\gamma\lambda$ are sampled a factor $\sqrt{2}$ around the heuristic and by steps of 2 thereafter.

construction, the denoisers are blind to the measurement setting (with regards to the noise level, see discussion around 17), in contrast to end-to-end networks.

4.4 Experiment 3: validating AIRI in precision and cost

4.4.1 Simulation setup

In this section, we validate AIRI- ℓ_2 and AIRI- ℓ_1 more extensively, both in reconstruction quality and computational cost, across a wider range of observation durations, and in comparison with, uSARA, SARA, CLEAN, as well as UNets trained in an end-to-end fashion. We consider the same four groundtruth images and experimental setup as for the previous experiment (see Section 4.3.1), with the same set of five pointing directions (see uv -coverages in Figs 5 and 6c). In order to further assess robustness to varying measurement conditions, the set of observation durations is enlarged to $\Delta T \in \{1 \text{ h}, 2 \text{ h}, 4 \text{ h}, 8 \text{ h}\}$ (see Fig. 6). We note that a different UNet is trained for each observation duration to alleviate the challenge of generalizability with respect to the sampling pattern. The exponentiation parameter for the UNet training data base, resulting from the procedure in Section 3.3.3, is $a = 10^3$.

We recall that our groundtruth images have peak values at 1, yielding the same heuristic value for σ and $\gamma\lambda$. These values are still within a 25 per cent variation range across pointing directions for all four ΔT values, and a single heuristic value is considered for all five uv -patterns associated with the same ΔT . The resulting four average $\nu_{\Delta T}$ are: $\nu_1 = 2.9 \times 10^{-4}$, $\nu_2 = 2.2 \times 10^{-4}$, $\nu_4 = 1.4 \times 10^{-4}$, and $\nu_8 = 9.3 \times 10^{-5}$. Given the result of the previous experiment the values of σ for AIRI- ℓ_2 and AIRI- ℓ_1 are taken exactly at the heuristic, while




Figure 10. Experiment 3 results: Reconstruction SNR (left-hand panel) and logSNR (right-hand panel) as a function of the observation duration for AIRI- ℓ_2 , AIRI- ℓ_1 , uSARA, SARA, CLEAN, and the end-to-end UNets. Each point is an average over the 20 simulated observations built from the four different groundtruths from Fig. 2 and the five different uv -patterns shown in Figs 5 and 6(c). For AIRI- ℓ_2 and AIRI- ℓ_1 , σ is chosen at the heuristic. For uSARA, $\gamma\lambda$ is fine-tuned manually to optimize the logSNR, resulting in values between 3 and 6 times below the heuristic (in line with Fig. 9). Error bars show the 95 per cent confidence interval.

the value of $\gamma\lambda$ for uSARA is fine-tuned manually to optimize the logSNR of each reconstruction, again resulting in soft-thresholding values ranging between 3 and 6 times below the heuristic, depending on ΔT . Moreover, setting $\kappa = 10^{-9}$ (resp. $\kappa = 10^{-5}$) for the AIRI- ℓ_2 (resp. AIRI- ℓ_1) denoiser in (22), with $\varepsilon = 5 \times 10^{-2}$, ensured the stability of Algorithm 3. Following the procedure detailed in Section 3.3.3, the exponentiation parameter for on-the-fly dynamic range enhancement of the data base is $a = 10^3$.

At the level of Algorithms 1, 2, and 3, the stepsize γ , convergence criteria $(\xi_i)_{1 \leq i \leq 3}$ and maximum number of iterations are set to the same values as in the previous experiment.

Finally, AIRI- ℓ_2 , AIRI- ℓ_1 , uSARA, and SARA are run in Matlab, while CLEAN is implemented in the highly optimized C++ WSCLEAN package, and the UNets are implemented in Python using PYTORCH. uSARA, SARA, and CLEAN implementations utilize 10 CPU cores (Intel Xeon E5-2695 2.1 GHz). The parallel PDFB algorithmic structure underpinning SARA (see Section 2.3.1) optimizes the distribution of the operations involved at each iteration across those 10 CPU cores: 1 for the data-fidelity term, and 1 for each of the nine bases of the average sparsity dictionary Ψ (Onose et al. 2016). For uSARA, the average sparsity dictionary is parallelized via a faceting procedure using nine facets (Pruša 2012; Thouvenin et al. 2022a). AIRI- ℓ_2 and AIRI- ℓ_1 utilize 10 CPU cores and 1 GPU (NVIDIA Tesla V100). The latter is for the application of the denoiser, as commonly done with DNNs in deep learning. The DNN can be run on CPU, but this slows down the reconstructions significantly, as will be evident from the experimental results. Implementations of uSARA, SARA, and CLEAN leveraging GPUs are not investigated here. The UNets are run utilizing 1 GPU (NVIDIA Tesla V100).

4.4.2 Reconstruction quality

Fig. 10 shows the average reconstruction SNR and logSNR as a function of the observation duration ΔT . Each point, on each of the two graphs, represents an average over 20 reconstructions arising when considering the four test images in Fig. 2 and five pointing directions, while error bars give the 95 per cent confidence intervals. As expected, the reconstruction SNR increases with ΔT , as the data size increases with ΔT . The results extend the conclusion of the

previous experiment of a superior performance of AIRI- ℓ_1 over both AIRI- ℓ_2 and uSARA across $\Delta T \in \{1 \text{ h}, 2 \text{ h}, 4 \text{ h}, 8 \text{ h}\}$. In SNR, AIRI- ℓ_1 and AIRI- ℓ_2 are respectively around 3 dB and 2 dB above uSARA. In logSNR, AIRI- ℓ_1 is between 1 dB and 2 dB above both AIRI- ℓ_2 and uSARA. The comparison with SARA reveals that AIRI- ℓ_1 is on par with SARA in SNR and slightly superior in logSNR across the range of measurement settings considered. As previously, we observe that the confidence interval around the mean SNR and logSNR values for both AIRI- ℓ_2 and AIRI- ℓ_1 is similar to that of uSARA and SARA, reflecting the robustness of the proposed method. The reconstruction quality of both CLEAN and the UNETS trained end-to-end is much inferior to those of AIRI and SARA approaches, for both SNR and logSNR and all measurement conditions considered. Interestingly, notice that the UNETS significantly improve over CLEAN in linear scale, but perform significantly worse in logarithmic scale, which confirms the difficulty of the UNETS to recover very faint emissions. This phenomenon was also noticed during training.

We recall that SARA and uSARA leverage the same prior model, encapsulated in the averaged sparsity proximal operator. However, SARA relies on a constrained data-fidelity term (see objective 6), while uSARA uses an unconstrained formulation enforcing less abruptly data fidelity, in the sense that it imposes a soft penalty rather than a hard constraint (see objective 11). A byproduct of our analysis is to confirm the superiority of the former over the latter.

Fig. 11 shows the reconstructed Hercules A images for the considered methods and the uv -coverages shown in Figs 6(a), (b), and (c). The visual analysis confirms the quantitative results from Fig. 10. We acknowledge that the AIRI- ℓ_2 and AIRI- ℓ_1 algorithms can create artefacts by overemphasizing small structures (see the red zoom box in Fig. 11 o). SARA exhibits very similar quality to AIRI- ℓ_1 , but not always with the same resolution-to-artefacts tradeoff. The images suggest that the poor reconstruction of CLEAN stems from a strong offset in intensity values (see in particular the difference between the estimated higher intensity values in the upper rectangular zoom from Figs 11(c), (k), and (q), and that of the groundtruth). Also note the important residual noise hiding low intensity features of the image, as well as remaining artefacts for $\Delta T = 1 \text{ h}$ and $\Delta T = 2 \text{ h}$. This figure also illustrates clearly the strengths and shortcomings of the UNets trained in an end-to-end manner. The relatively good SNR and low logSNR, together with the visual results, confirm that the UNets manage to recover quite accurately the high intensity emission but struggle to recover faint emission. Visually, the results of the UNets present some hallucinated point-like sources (Muckley et al. 2021), which we believe is the consequence of our synthetic training data base containing a large number of such sources.

4.4.3 Computational cost and reconstruction times

The computational costs for the various approaches considered are reported in Fig. 12. Each point on the graphs is an average over the 20 simulated observations built from the four different groundtruths from Fig. 2 and the five different uv -patterns shown in Figs 5 and 6(c). Error bars show the 95 per cent confidence intervals.

Fig. 12(a) displays the average time per iteration, independently for each of the forward and backward steps of both uSARA and AIRI, as a function of the observation duration. We recall once more that uSARA relies on the same algorithmic structure as AIRI- ℓ_2 and AIRI- ℓ_1 . In both cases, the underpinning FB algorithm alternates between a (forward) gradient-descent step enforcing data fidelity and a (backward) regularization step enforcing a prior model via the application of either the average sparsity proximal operator (uSARA,

Algorithm 1) or a learned denoiser (AIRI, Algorithm 3). We note that, in such an algorithmic structure, the computational cost of the regularization step scales linearly with the image size, while that of the gradient-descent step scales linearly with the data size.

First, we observe from the figure that the computation time of the gradient step naturally scales linearly with the observation duration, which is indeed proportional to data size. Secondly, we acknowledge that the computation of the proximal operator within uSARA is slow, which is due to its sub-iterative nature. For the image and data sizes considered in our simulations, and given our CPU implementation, it always dominates the computation time of the gradient, on average per iteration of Algorithm 1. This conclusion holds despite the fact that we have carefully optimized the external parameters of Algorithm 1 and Algorithm 2 to optimize the reconstruction time (K in Algorithm 1, and the tolerance criterion ξ_2 and dual variable \mathbf{v}_0 in Algorithm 2). In that regard, we note that an efficient initialization for the dual variable \mathbf{v}_0 consists in setting it as the dual variable from the previous use of Algorithm 2. Using such an initialization, the number of sub-iterations for Algorithm 2 progressively decreases with the iteration count in Algorithm 1. Thirdly, using a standard GPU implementation of the AIRI- ℓ_2 and AIRI- ℓ_1 denoisers, their application is two orders of magnitude faster than the computation of the proximal operator. A CPU implementation of the denoisers makes them even slower than the proximal operator in uSARA. For the image and data sizes at stake, the AIRI denoisers on GPU are also significantly faster than the computation of the gradient, on average per iteration of Algorithm 1.

Fig. 12(b) displays the total reconstruction times for uSARA and AIRI as a function of the observation duration, and in comparison with SARA and CLEAN. The AIRI- ℓ_2 and AIRI- ℓ_1 reconstruction times with denoiser on GPU scale almost linearly with the observation duration, due to the domination of the gradient step. For uSARA, and for AIRI- ℓ_2 and AIRI- ℓ_1 with denoiser on CPU, the dependency on observation duration is less severe given the domination of the regularization step at each iteration. Overall, given the speed up in the regularization step, the total reconstruction time of AIRI with denoiser on GPU is significantly (5 to 10 times) lower than the one of uSARA. AIRI- ℓ_2 and AIRI- ℓ_1 with denoiser on CPU are much slower, in fact slightly slower than uSARA itself. SARA offers reconstruction time between those of AIRI and uSARA. CLEAN remains significantly faster than AIRI (4 to 15 times), but the fastest method among all is by far the UNet trained end-to-end. We underline that this is to be expected since, not only UNets were run on GPU, but the reconstruction with the UNet only requires a single inference step, as opposed to the other methods that are iterative.

To summarize, the application of an AIRI denoiser (on GPU) is orders of magnitude faster than that of the average sparsity proximal operator of uSARA (on CPU). Substituting the former by the latter does not only provide superior imaging quality, but significantly reduces the computational cost of the regularization step. For image and data sizes where the uSARA regularization step dominates over the gradient-descent data-fidelity step, a cost reduction in the former, substituting the proximal operator for a DNN denoiser leads to significant speed up of the algorithm.

5 CONCLUSION AND FUTURE WORK

5.1 Conclusion for both AIRI and uSARA

In this work, we have proposed a new algorithmic framework for scalable precision RI imaging, dubbed AIRI. The approach consists in encapsulating a prior image model in a DNN denoiser, and

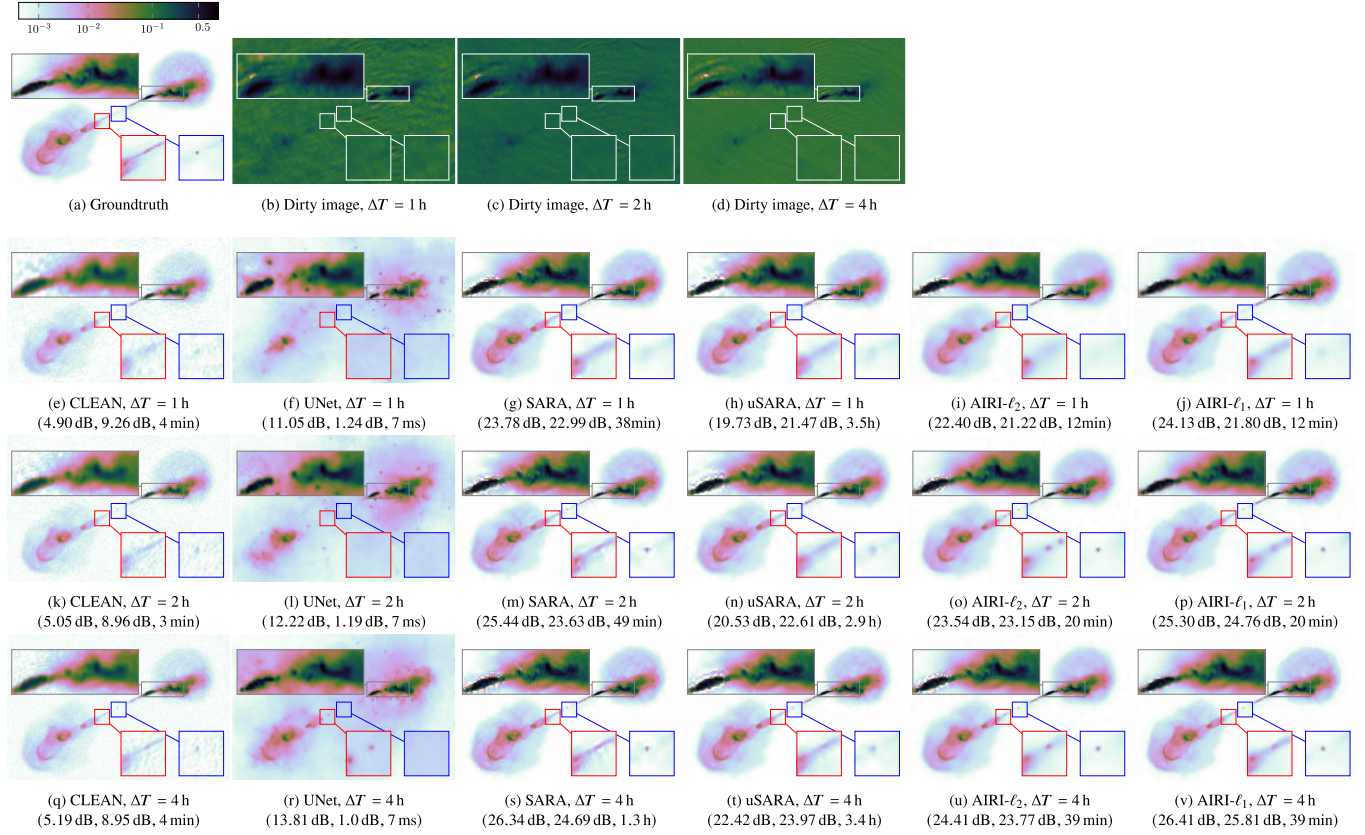


Figure 11. Experiment 3 results: Influence of the observation duration ΔT on the reconstruction quality for the different algorithms with the pointing position from Fig. 6; the associated uv -patterns are displayed in Figs 6(a), (b), and (c). Second, third, and last row show reconstruction results for $\Delta T = 1$ h, $\Delta T = 2$ h and $\Delta T = 4$ h, respectively. For these three rows, each column from left to right shows estimated model images obtained with CLEAN, UNETS, SARA, uSARA, AIRI- ℓ_2 , and AIRI- ℓ_1 , respectively. Below each reconstruction we indicate the reconstruction metrics and time (SNR, logSNR, reconstruction time). All images are shown in logarithmic scale.

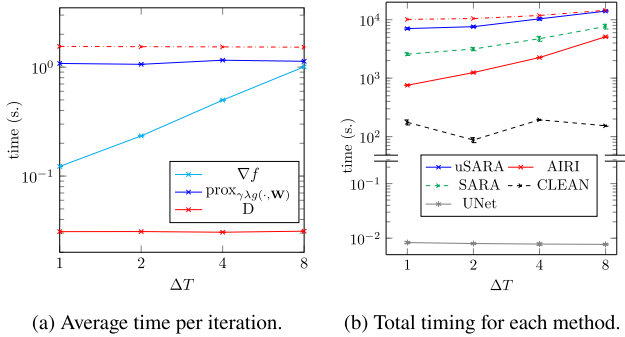


Figure 12. Experiment 3 results: reconstruction times for AIRI- ℓ_2 , AIRI- ℓ_1 , uSARA, SARA, CLEAN, and the end-to-end UNets (on GPU). AIRI- ℓ_2 and AIRI- ℓ_1 exhibit identical behaviours, with only AIRI- ℓ_2 results are reported and referred to as AIRI on the graphs. Each point is an average over the 20 simulated observations. Error bars show the (very small and virtually invisible) 95 per cent confidence interval. For AIRI- ℓ_2 and AIRI- ℓ_1 , all reconstruction times are reported for both GPU (solid red lines) and CPU implementation (dashed red lines) of D. Left-hand panel: average time per iteration in seconds (s) for the forward and backward operators involved in Algorithm 1 (∇f and $\text{prox}_{\gamma\lambda g}$) and Algorithm 3 (∇f and denoiser D), as a function of the observation duration. Right-hand panel: average total reconstruction times for AIRI- ℓ_2 , AIRI- ℓ_1 , uSARA, SARA, CLEAN, and the end-to-end UNETS, as a function of the observation duration.

plugging it in lieu of the proximal regularization operator of an optimization algorithm of the SARA family for image reconstruction. AIRI inherits the robustness and interpretability of optimization approaches and the learning power and speed of networks, while avoiding generalizability limitations of end-to-end networks, for which training involves the details of the measurement model.

More specifically, we have first designed a realistic low dynamic range training data base from optical images. We have also established a method to train firmly non-expansive denoisers at a noise level inferred from the target dynamic ranges of reconstruction, with either an ℓ_2 or ℓ_1 loss, including a procedure for on-the-fly data base dynamic range enhancement. Finally, the AIRI- ℓ_2 and AIRI- ℓ_1 algorithms resulting from plugging the denoisers into the FB algorithm were validated in simulation for reconstruction of images involving complex structure with diffuse and faint emission across the field of view (at a dynamic range of 10^4). The benchmark algorithms included CLEAN, the state-of-the-art SARA algorithm, as well as its new unconstrained version uSARA, which only differs from AIRI- ℓ_2 and AIRI- ℓ_1 by the use of the average sparsity proximal regularization operator instead of a learned denoiser in the FB algorithm.

Our results show that the first incarnations of AIRI are already competitive with, if not superior to, the advanced optimization algorithms of the SARA family in imaging quality, well beyond the capabilities of CLEAN and of the UNETS trained in an end-to-end fashion. In

particular, while the AIRI denoisers are trained on the same data base as the UNETS, the PnP solutions do not exhibit hallucinated artefacts, which illustrates the robustness of AIRI compared to the end-to-end UNETS. With regards to computation time, the AIRI approach was shown to provide a significant acceleration potential over uSARA and SARA, thanks to the inference speed of DNN denoisers over the averaged sparsity proximal operator, but remain significantly slower than CLEAN, itself orders of magnitude slower than the UNETS. The results also confirm the validity of our heuristic to set the training noise level σ , consisting in enforcing equality of the training and target dynamic range of the reconstruction, and the validity of the data base exponentiation approach (with parameter a), consisting in enforcing equality of the final nominal dynamic range of the data base and target dynamic range of reconstruction. This provides a significant advantage over uSARA, for which the regularization parameter λ requires manual fine-tuning, even though the fact that the corresponding heuristic (and its 1/3-corrected version) provides an appropriate reference upper-bound for the optimal value is an interesting result.

5.2 Current limitations and future work

In what follows, we discuss current limitations of the AIRI framework, and future work. This also includes considerations for further developments of uSARA and end-to-end DNNs.

First, the modified DnCNN denoiser architecture used, the ℓ_2 and ℓ_1 losses considered, as well as the set of optical images used to build a training data base constitute the most basic choices leading to our first AIRI- ℓ_2 and AIRI- ℓ_1 incarnations of the framework. We anticipate that future work, building a richer data base from RI observation, and considering more advanced losses, such as adversarial losses (Wang et al. 2018), and more advanced network architectures, such as UNETS (Hurault, Leclaire & Papadakis 2022), can provide a quantum jump in imaging precision with AIRI. We also acknowledge that our approach to enforce the firm non-expansiveness of the denoiser is only approximate, and relies on fine-tuning the regularization parameter κ in the training loss. Simpler and more precise approaches to ensure firm non-expansiveness should be investigated in the future (e.g. leveraging 1-Lipschitz layers). We note however that, in the context of our experiments, the proposed approach has, in practice, provided a robust way to ensure convergence of the resulting PnP algorithms.

Secondly, fully optimized and parallelized implementations of AIRI algorithms, but also SARA, or uSARA, should be investigated, leveraging GPUs not only for DNN denoisers, but also for the implementation of the linear (measurement or regularization) operators involved. This should lead to further acceleration of these approaches, and possibly participate to closing the gap with CLEAN in terms of reconstruction time. We also note that accelerated versions of uSARA are achievable, in particular leveraging preconditioning strategies (Repetti & Wiaux 2021).

Thirdly, with regards to the details of the target imaging modality, we have investigated monochromatic intensity imaging on small fields of view only, with mild, if any, superresolution factor with respect to the nominal resolution of the observation, as set by the largest baseline. Extensions of the framework to further superresolution functionality and to wideband polarization imaging should be contemplated. Also, a byproduct of the present analysis is to confirm the superiority of SARA over uSARA, due to the constrained approach to data fidelity in SARA, versus the unconstrained approach underpinning uSARA. In this context, the superior results of AIRI over uSARA are entirely related to the capability to encapsulate

a better prior model than the average sparsity model in learned DNN denoisers. On the one hand, a first possible enhancement of the AIRI approach would consist in developing a PnP version of SARA (as opposed to uSARA), simultaneously taking advantage of networks to learn prior models and of the constrained data-fidelity formulation. On the other hand, to date, the unconstrained formulation underpinning uSARA, AIRI- ℓ_2 , and AIRI- ℓ_1 is a critical building block of the only optimization-based algorithm proposed for RI imaging capable of handling jointly DDE calibration and imaging (Repetti et al. 2017; Dabbech et al. 2021). This is a strong justification for the study of unconstrained versus constrained formulations of the imaging module.

Fourthly, the proposed UNET end-to-end implementation showed interesting results, albeit significantly outperformed by the proposed AIRI algorithms. A more thorough study should be undertaken, leveraging more advanced architectures, such as unfolded networks that would enable us to incorporate knowledge about the measurement operator within the architecture.

Finally, this work is simulation-based only. Dabbech et al. (2022), Wilber et al. (in preparation) validate AIRI and uSARA on real wide-field high-resolution high-dynamic range data from MeerKAT and ASKAP, relying on a highly parallelized implementation of the measurement operator tailored to the wide-field setting.

ACKNOWLEDGEMENTS

The research of MT was funded by Heriot-Watt under the James Watt Scholarship scheme. The work of MT, AD, and YW was supported by the Engineering and Physical Sciences Research Council (EPSRC) under grants EP/T028270/1 and ST/W000970/1. The computing resources came from the Cirrus UK National Tier-2 HPC Service at EPCC (<http://www.cirrus.ac.uk>) funded by the University of Edinburgh and EPSRC (EP/P020267/1), partly through time allocation under the SUSAs project, and partly through GPU resources directly provided by EPCC (Adrian Jackson). Credits for the 32 images used in our training data base go to NOIRLab/NSF/AURA/H.Schweiker/WIYN/T.A.Rector (University of Alaska Anchorage). The radio images were taken from: for Hercules A: R. Perley and W. Cotton (NRAO/AUI/NSF), for 3c353: NRAO/VLA, for Centaurus A: NRAO/AUI/NSF/Univ.Hertfordshire/M.Hardcastle.

DATA AVAILABILITY

AIRI code will be made available as part of a later release of and joint parallel toolbox for AIRI and uSARA.

REFERENCES

- Abdulaziz A., Dabbech A., Onose A., Wiaux Y., 2016, 2016 24th European Signal Processing Conference (EUSIPCO). IEEE, Budapest, p. 388
- Abdulaziz A., Dabbech A., Wiaux Y., 2019, *MNRAS*, 489, 1230
- Ahmad R., Bouman C. A., Buzzard G. T., Chan S., Liu S., Reehorst E. T., Schniter P., 2020, *IEEE Signal Process. Magaz.*, 37, 105
- Akiyama K. et al., 2019, *ApJ*, 875, L4
- Akiyama K. et al., 2022, *ApJ*, 930, L14
- Arras P., Frank P., Leike R., Westermann R., Enßlin T. A., 2019, *A&A*, 627, A134
- Attouch H., Bolte J., Svaiter B. F., 2013, *Math. Programming*, 137, 91
- Balestrieri R., Baraniuk R., 2021, preprint ([arXiv:2104.00219](https://arxiv.org/abs/2104.00219))
- Baraniuk R. G., 2007, *IEEE Signal Process. Magaz.*, 24, 118
- Bauschke H. H., Combettes P. L., 2017, *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer, Berlin

- Beck A., Teboulle M., 2009, *SIAM J. Imaging Sci.*, 2, 183
- Bhatnagar S., Cornwell T., 2004, *A&A*, 426, 747
- Birdi J., Repetti A., Wiaux Y., 2018, *MNRAS*, 478, 4442
- Birdi J., Repetti A., Wiaux Y., 2020, *MNRAS*, 492, 3509
- Bonaldi A., Braun R., 2018, preprint ([arXiv:1811.10454](https://arxiv.org/abs/1811.10454))
- Boyd S., Parikh N., Chu E., 2011, *Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers*. Now Publishers Inc., Delft
- Bredies K., Holler M., 2020, *Inverse Probl.*, 36, 123001
- Bredies K., Kunisch K., Pock T., 2010, *SIAM J. Imaging Sci.*, 3, 492
- Briggs D. S., 1995, *American Astron. Soc. Meeting Abstracts*, 112.02
- Bryan G. L. et al., 2014, *ApJS*, 211, 19
- Buades A., Coll B., Morel J.-M., 2011, *Image Process. Line*, 1, 208
- Cai X., Pereyra M., McEwen J. D., 2018, *MNRAS*, 480, 4154
- Candès E., Demanet L., 2003, *C. R. Math.*, 336, 395
- Candès E. J., Romberg J., Tao T., 2006, *IEEE Trans. Inf. Theor.*, 52, 489
- Candès E. J., Wakin M. B., Boyd S. P., 2008, *J. Fourier Anal. Appl.*, 14, 877
- Carrillo R. E., McEwen J. D., Wiaux Y., 2012, *MNRAS*, 426, 1223
- Carrillo R. E., McEwen J. D., Wiaux Y., 2014, *MNRAS*, 439, 3591
- Chan S. H., Wang X., Elgendy O. A., 2016, *IEEE Trans. Comput. Imaging*, 3, 84
- Chouzenoux E., Pesquet J.-C., Repetti A., 2014, *J. Opt. Theory Appl.*, 162, 107
- Cohen R., Elad M., Milanfar P., 2021, *SIAM J. Imaging Sci.*, 14, 1374
- Combettes P. L., Pesquet J.-C., 2011, *Fixed-Point Algorithms for Inverse Problems in Science and Engineering*. Springer, Berlin, p. 185
- Combettes P. L., Pesquet J.-C., 2015, *SIAM J. Opt.*, 25, 1221
- Connor L., Bouman K. L., Ravi V., Hallinan G., 2022, *MNRAS*, 514, 2614
- Cornwell T. J., Golap K., Bhatnagar S., 2008, *IEEE J. Select. Topics Signal Process.*, 2, 647
- Dabbech A., Ferrari C., Mary D., Slezak E., Smirnov O., Kenyon J. S., 2015, *A&A*, 576, A7
- Dabbech A., Wolz L., Pratley L., McEwen J. D., Wiaux Y., 2017, *MNRAS*, 471, 4300
- Dabbech A., Onose A., Abdulaziz A., Perley R. A., Smirnov O. M., Wiaux Y., 2018, *MNRAS*, 476, 2853
- Dabbech A., Repetti A., Perley R. A., Smirnov O. M., Wiaux Y., 2021, *MNRAS*, 506, 4855
- Dabbech A., Terris M., Jackson A., Ramatsoku M., Smirnov O., Wiaux Y., 2022, preprint ([arXiv:2207.11336](https://arxiv.org/abs/2207.11336))
- Dabov K., Foi A., Katkovnik V., Egiazarian K., 2007, *IEEE Trans. Image Process.*, 16, 2080
- Do M. N., Vetterli M., 2003, *Studies in Comp. Math. Vol. 10*. Elsevier, Amsterdam, p. 83
- Donoho D. L., 2006, *IEEE Trans. Inf. Theor.*, 52, 1289
- Donoho D. L., Johnstone I. M., 1995, *J. Am. Stat. Assoc.*, 90, 1200
- Eckstein J., Bertsekas D. P., 1992, *Math. Programming*, 55, 293
- Fessler J. A., 2020, *IEEE Signal Process. Magaz.*, 37, 33
- Garsden H. et al., 2015, *A&A*, 575, A90
- Gheller C., Vazza F., 2021, *MNRAS*, 509, 990
- Golub G. H., Van Loan C. F., 2013, *Matrix Computations*. JHU Press, Baltimore
- Goodfellow I., Shlens J., Szegedy C., 2015, 3rd International Conference on Learning Representations. ICLR, San Diego
- Gupta H., Jin K. H., Nguyen H. Q., McCann M. T., Unser M., 2018, *IEEE Trans. Medical Imaging*, 37, 1440
- Hertrich J., Neumayer S., Steidl G., 2021, *Linear Algebra Appl.*, 631, 203
- Högbom J., 1974, *A&AS*, 15, 417
- Hurault S., Leclaire A., Papadakis N., 2022, in Chaudhuri K., Jegelka S., Song L., Szepesvari C., Niu G., Sabato S., eds, *Proceedings of the 39th International Conference on Machine Learning*, Vol. 162, Proximal Denoiser for Convergent Plug-and-Play Optimization with Nonconvex Regularization. PMLR, Baltimore, Maryland, p. 9483
- Jin K. H., McCann M. T., Froustey E., Unser M., 2017, *IEEE Trans. Image Process.*, 26, 4509
- Kingma D. P., Ba J., 2015, Adam A Method for Stochastic Optimization. Proceedings of the 3rd International Conference on Learning Representations, ICLR. Scientific Res., Wuhan
- Levin A., Weiss Y., Durand F., Freeman W. T., 2009, *IEEE Conf. on Comput. Vision and Pattern Recognition*. IEEE, New York, p. 1964
- Li F., Cornwell T. J., de Hoog F., 2011, *A&A*, 528, A31
- Luisier F., Blu T., Unser M., 2007, *IEEE Trans. Image Process.*, 16, 593
- Maddox W., Tang S., Moreno P., Wilson A. G., Damianou A., 2021, *International Conference on Artificial Intelligence and Statistics*. PMLR, Virtual, p. 2737
- Mairal J., Bach F., Ponce J., Sapiro G., 2009, 26th International Conference on Machine Learning. Association for Computing Machinery, New York, p. 689
- Mallat S., 1999, *A Wavelet Tour of Signal Processing*. Elsevier, Amsterdam
- Muckley M. J. et al., 2021, *IEEE trans. Med. Imaging*, 40, 2306
- Nguyen A., Yosinski J., Clune J., 2015, *IEEE Conf. on Comput. Vision and Pattern Recognition*. IEEE, New York, p. 427
- Offringa A. R., Smirnov O., 2017, *MNRAS*, 471, 301
- Onose A., Carrillo R. E., Repetti A., McEwen J. D., Thiran J.-P., Pesquet J.-C., Wiaux Y., 2016, *MNRAS*, 462, 4314
- Onose A., Dabbech A., Wiaux Y., 2017, *MNRAS*, 469, 938
- Pang T., Du C., Dong Y., Zhu J., 2018, *Adv. Neural Inf. Process. Syst.*, 31
- Paszke A. et al., 2017, in Wallach H., Larochelle H., Beygelzimer A., d'Alché Buc F., Fox E. B., Garnett R., eds, *Advances in Neural Inf. Processing Syst. Workshop, PyTorch: An Imperative Style, High-Performance Deep Learning Library*. Curran Associates, Inc., Red Hook, NY, p. 8024
- Pesquet J.-C., Repetti A., 2015, *J. Nonlinear Convex Anal.*, 16, 2353
- Pesquet J.-C., Repetti A., Terris M., Wiaux Y., 2021, *SIAM J. Imaging Sci.*, 14, 1206
- Pratley L., McEwen J. D., d'Avezac M., Carrillo R. E., Onose A., Wiaux Y., 2018, *MNRAS*, 473, 1038
- Pruša Z., 2012, PhD thesis, Brno University of Technology
- Repetti A., Wiaux Y., 2020, *IEEE Int. Conf. on Acoustics, Speech and Signal Processing*. IEEE, New York, p. 1434
- Repetti A., Wiaux Y., 2021, *SIAM J. Opt.*, 31, 1215
- Repetti A., Birdi J., Dabbech A., Wiaux Y., 2017, *MNRAS*, 470, 3981
- Repetti A., Pereyra M., Wiaux Y., 2018, 26th European Signal Processing Conference (EUSIPCO). IEEE, New York, p. 2668
- Repetti A., Pereyra M., Wiaux Y., 2019, *SIAM J. Imaging Sci.*, 12, 87
- Romano Y., Elad M., Milanfar P., 2017, *SIAM J. Imaging Sci.*, 10, 1804
- Rudin L. I., Osher S., Fatemi E., 1992, *Phys. D: Nonlinear Phenom.*, 60, 259
- Ryu E., Liu J., Wang S., Chen X., Wang Z., Yin W., 2019, *Int. Conf. on Machine Learning*. PMLR, Long Beach, California, USA, p. 5546
- Schwab F. R., Cotton W. D., 1983, *AJ*, 88, 688
- Schwarz U., 1978, *A&A*, 65, 345
- Selesnick I. W., Baraniuk R. G., Kingsbury N. C., 2005, *IEEE Signal Process. Magaz.*, 22, 123
- Terris M., Abdulaziz A., Dabbech A., Jiang M., Repetti A., Pesquet J.-C., Wiaux Y., 2019, *SPARS workshop: Deep PostProcessing for Sparse Image Deconvolution*. Toulouse, France
- Terris M., Repetti A., Pesquet J.-C., Wiaux Y., 2020, *IEEE Int. Conf. on Acoustics, Speech and Signal Processing*. IEEE, New York, p. 8658
- Thiébaud É., Young J., 2017, *JOSA A*, 34, 904
- Thompson A. R., Moran J. M., Swenson G. W. Jr, 2017, *Interferometry and Synthesis in Radio Astronomy*. Springer Nature, Berlin
- Thouvenin P.-A., Abdulaziz A., Dabbech A., Repetti A., Wiaux Y., 2022a, preprint ([arXiv:2003.07358](https://arxiv.org/abs/2003.07358))
- Thouvenin P.-A., Dabbech A., Jiang M., Abdulaziz A., Thiran J.-P., Jackson A., Wiaux Y., 2022b, HWU research portal ID: 52577083
- Venkatakrisnan S. V., Bouman C. A., Wohlberg B., 2013, 2013 IEEE Global Conf. on Signal and Inf. Processing. IEEE, New York, p. 945
- Vidal A. F., De Bortoli V., Pereyra M., Durmus A., 2020, *SIAM J. Imaging Sci.*, 13, 1945
- Wang L., Xiong Z., Gao D., Shi G., Wu F., 2015, *Appl. Opt.*, 54, 848
- Wang X., Yu K., Wu S., Gu J., Liu Y., Dong C., Qiao Y., Change Loy C., 2018, *European Conf. on Comput. Vision Workshops*. Munich
- Wiaux Y., Jacques L., Puy G., Scaife A. M., Vanderghenst P., 2009, *MNRAS*, 395, 1733

- Xie Q., Zhou M., Zhao Q., Meng D., Zuo W., Xu Z., 2019, IEEE/CVF Conf. on Comput. Vision and Pattern Recognition. IEEE, New York, p. 1585
- Yang J., Wright J., Huang T. S., Ma Y., 2010, *IEEE Trans. Image Process.*, 19, 2861
- Zbontar J. et al., 2018, preprint ([arXiv:1811.08839](https://arxiv.org/abs/1811.08839))
- Zhang K., Zuo W., Chen Y., Meng D., Zhang L., 2017, *IEEE Trans. Image Process.*, 26, 3142
- Zhang K., Zuo W., Zhang L., 2019, IEEE/CVF Conf. on Comput. Vision and Pattern Recognition. IEEE, New York, p. 1671
- Zhang K., Li Y., Zuo W., Zhang L., Van Gool L., Timofte R., 2021, *IEEE Trans. Pattern Analysis and Machine Intelligence*. IEEE, New York
- Zhou Y., Jiao J., Huang H., Wang Y., Wang J., Shi H., Huang T., 2020, AAAI Conf. on Artificial Intelligence. PMLR, Virtual, p. 13074

APPENDIX A: ON THE CHOSEN RI IMAGING MODEL

We discuss two points regarding the generality of the measurement model in (1). First, at the dynamic range of interest for current and future observations, direction-dependent effects (DDEs) of atmospheric and instrumental origin complicate the RI measurement equation. More precisely, at each time instant, the visibility associated with a pair of antennas is formed from the modulation of the radio sky with the product of a DDE pattern specific to each antenna. The DDEs are typically unknown and need to be calibrated, leading to a blind deconvolution problem, which requires the design of joint calibration and imaging approaches, typically alternating between imaging and calibration modules (Repetti et al. 2017; Dabbech et al. 2021). One exception to this, is the so-called w -term induced by the projection of the baseline associated with each visibility on the line of sight, which becomes non-negligible on wide fields of view, and thus acts as a known DDE. In practice, DDEs can be integrated into extended convolution kernels in \mathbf{G} (Dabbech et al. 2017). In this work, we assume a measurement operator Φ not affected by any known or unknown DDE, considering a pure imaging problem on small fields of view.

Secondly, the measurement model (1) assumes noise with constant variance across visibilities. In full generality, the noise is not white, and a diagonal whitening matrix $\Theta \in \mathbb{R}^{m \times m}$ with entries equal to the inverse noise standard deviation per visibility can be applied to the original measurement vector, ensuring that the resulting visibility vector is indeed affected by white noise. In this case problem (1) still holds, with a measurement operator $\Phi = \Theta \mathbf{G} \mathbf{F} \mathbf{Z}$, and \mathbf{y} and \mathbf{e} the visibility and noise vectors after application of the whitening matrix. This operation is known as natural weighting. This scheme is used by the algorithms of the SARA family to ensure the negative log-likelihood interpretation of its data-fidelity terms (Carrillo et al. 2012). More complex weighting schemes, such as uniform and Briggs weighting, are used by CLEAN (Briggs 1995), which involve additional multiplicative terms in Θ . Both approaches add to natural weighting by downweighting visibilities in regions of the uv -plane with high sampling density, with the aim to reduce the sidelobes of the dirty beam and improve the resolution of the CLEAN reconstruction, though at the expense of sensitivity. Interestingly, Onose et al. (2017) showed, for algorithms of the SARA family specifically, that, while keeping to natural weighting, using sampling density information in acceleration strategies such as preconditioning are very efficient at bringing jointly optimal resolution and dynamic range. In other words, the resolution-versus-sensitivity trade-off disappears asymptotically at convergence. In this work, without loss of generality, our simulations consider i.i.d. Gaussian random measurement noise so that $\Theta = \mathbf{I}$ for natural weighting, where \mathbf{I} denotes the identity

matrix. All algorithms discussed (SARA, uSARA, AIRI) use natural weighting, while CLEAN is implemented with uniform weighting (see Section 4.1.2).

APPENDIX B: ON THE PNP CONVERGENCE

Several works have recently focused on restoring the convergence of PnP algorithms using DNN denoisers. The majority of these works assume a non-expansiveness constraint on the denoiser (Romano, Elad & Milanfar 2017; Ryu et al. 2019; Terris et al. 2020; Cohen et al. 2021; Hertrich, Neumayer & Steidl 2021), but this constraint often comes with restrictive assumptions either on the algorithm, on the DNN architecture or on the operator ∇f .

In Pesquet et al. (2021), casting the problem in the framework of monotone operator theory, we have shown that the convergence of a PnP algorithm can be ensured by introducing a well-chosen regularization in the training loss of the denoiser D enforcing its firm non-expansiveness, without any form of limitation on the DNN architecture itself. Importantly, the approach also offers a generalized notion of characterization of its limit as the solution to monotone inclusion problems, more general than pure optimization problems, hence opening a path towards generalizing the usual Bayesian MAP interpretation of optimization solutions. We summarize those results in the remainder of this section.

Once again concentrating on a differentiable f such as (10), minimization problems like (3) can be reformulated (through their optimality conditions) as

$$\text{find } \hat{\mathbf{x}} \in \mathbb{R}^n \text{ s.t. } \quad 0 \in \nabla f(\hat{\mathbf{x}}) + \lambda \partial r(\hat{\mathbf{x}}), \quad (\text{B1})$$

where ∂r denotes the subdifferential of r , the subdifferential of $r \in \Gamma_0(\mathbb{R}^n)$ being the set-valued operator defined as $\partial r(\mathbf{x}) = \{\mathbf{u} \in \mathbb{R}^n | (\forall \mathbf{z} \in \mathbb{R}^n) \langle \mathbf{z} - \mathbf{x}, \mathbf{u} \rangle + r(\mathbf{x}) \leq r(\mathbf{z})\}$. When r is also differentiable, the subdifferential coincides with the gradient (i.e. $\partial r = \nabla r$). Yet, the subdifferential is also defined for non-differentiable functions such as $\iota_{\mathcal{B}(y, \epsilon)}$; we refer the reader to Bauschke & Combettes (2017) for more information. Problems as (B1) are part of the class of so-called monotone inclusion problems and encompass unconstrained minimization problems such as uSARA (11). They can however take a more general flavour and write as

$$\text{find } \hat{\mathbf{x}} \in \mathbb{R}^n \text{ s.t. } \quad 0 \in \nabla f(\hat{\mathbf{x}}) + A(\hat{\mathbf{x}}), \quad (\text{B2})$$

where A is an operator satisfying some maximal monotonicity assumptions (Bauschke & Combettes 2017). It is important to stress that this class of operators includes, but is not restricted to, subdifferentials. Problem (B2) is thus in general not equivalent to a minimization problem.

Solving such inclusion problems can be done using a generalized version of the proximal FB algorithm (5) with

$$(\forall k \in \mathbb{N}), \quad \mathbf{x}_{k+1} = J_{\gamma A}(\mathbf{x}_k - \gamma \nabla f(\mathbf{x}_k)), \quad (\text{B3})$$

where $J_{\gamma A}$ is the so-called resolvent of γA , yielding the convergence of the sequence $(\mathbf{x}_k)_{k \in \mathbb{N}}$ to $\hat{\mathbf{x}}$ satisfying (B2). As previously, the stepsize γ must satisfy $0 < \gamma < 2/L$ where L is the Lipschitz constant of ∇f .

The resolvent operator can be seen as a generalized version of the proximal operator (4) (Bauschke & Combettes 2017). In fact, (5) is a special case of (B3), since when $A = \lambda \partial r$ for $r \in \Gamma_0(\mathbb{R}^n)$, one has $J_{\gamma \lambda \partial r} = \text{prox}_{\gamma \lambda r}$. A key characterization of resolvent operators is the following: the operator J is the resolvent of a maximally monotone operator A if and only if it is firmly non-expansive. Technically, by definition, an operator J is firmly non-expansive if there exists an operator Q with Lipschitz constant smaller or equal to 1 such that

$J = (I + Q) / 2$, where I is the identity operator. Given the identical structure of PnP-FB and (B3), it follows that D in (15) should be firmly non-expansive for the convergence of the PnP algorithm to be ensured, with the PnP solutions characterized as solutions of a monotone inclusion problem (B2).

Thus, in order to impose the firm non-expansiveness of D , it is sufficient to impose that the Lipschitz constant of $Q = 2D - I$ is less than 1, or equivalently, that its Jacobian satisfies $\|\nabla Q(\mathbf{x})\|_S \leq 1$ for all \mathbf{x} . Since most deep learning frameworks rely on Jacobian-vector products (Paszke et al. 2017; Balestrieri & Baraniuk 2021; Maddox et al. 2021), one can, given some $\mathbf{x} \in \mathbb{R}^n$, compute $\|\nabla Q(\mathbf{x})\|_S$ via the power method. Thus, D would ideally be trained as a denoiser with a standard training loss, but under the constraint that $\|\nabla Q(\mathbf{x})\|_S \leq 1$ for all \mathbf{x} . Such a constraint is not trivial to impose in practice and, as proposed in Pesquet et al. (2021), we relax it and introduce a regularization term in the training loss that penalizes non firmly non-expansive networks.

Given a data base of noisy / groundtruth images $(\mathbf{z}_s, \mathbf{u}_s)_{1 \leq s \leq S}$, and denoting $\boldsymbol{\theta} \in \mathbb{R}^c$ the learnable parameters of D , the training loss ensuring the firm non-expansiveness of D reads:

$$\underset{\boldsymbol{\theta} \in \mathbb{R}^c}{\text{minimize}} \frac{1}{S} \sum_{s=1}^S \left(\mathcal{L}(D_{\boldsymbol{\theta}}(\mathbf{z}_s) - \mathbf{u}_s) + \kappa \max\{\|\nabla Q_{\boldsymbol{\theta}}(\mathbf{z}_s)\|_S, 1 - \varepsilon\} \right), \quad (\text{B4})$$

where \mathcal{L} is a loss function and $\varepsilon > 0$ is a safety margin parameter. As underlined in Pesquet et al. (2021), this penalty does not explicitly guarantee that the solution of this problem satisfies the constraint $\|\nabla Q(\mathbf{z})\|_S \leq 1$ on the data base, where it is effectively evaluated, let alone for all $\mathbf{z} \in \mathbb{R}^n$. Various countermeasures are however deployed, which have been demonstrated to be effective at ensuring PnP convergence in practice (see Pesquet et al. 2021). First, the safety margin parameter $\varepsilon > 0$ is a mean to strengthen the penalty. Secondly, the term $\|\nabla Q\|_S$ is not computed exactly at \mathbf{z}_s , but instead at a point sampled uniformly at random on the segment $[\mathbf{z}_s, \mathbf{u}_s]$, with the aim to explore a richer set of images than those of the data base only. Last but not least, the value of κ can, and should, be fine-tuned manually to optimize the balance between the regularization term and the standard loss \mathcal{L} .

We conclude this section by underlining that, contrary to other works (Ryu et al. 2019; Terris et al. 2020), the proposed approach does not require to impose layer-wise Lipschitz constraints nor architectural restrictions on the DNN.

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.