

Research Article

Image reconstruction with the JWST interferometer

Max Charles¹, Louis Desdoigts^{1,2}, Benjamin Pope^{3,4}, Peter Tuthill¹, Dori Blakely^{5,6}, Doug Johnstone^{5,6},
Shrishmoy Ray^{3,4}, K. E. Saavik Ford^{7,8,9}, Barry McKernan^{7,8,9} and Anand Sivaramakrishnan^{10,8,11}

¹Sydney Institute for Astronomy, School of Physics, University of Sydney, Camperdown, NSW 2006, Australia, ²Leiden Observatory, Niels Bohrweg 2, Leiden 2300RA, The Netherlands, ³School of Mathematical and Physical Sciences, Macquarie University, 12 Wally's Walk, Macquarie Park, NSW 2113, Australia, ⁴School of Mathematics and Physics, University of Queensland, St Lucia, QLD 4072, Australia, ⁵Department of Physics and Astronomy, University of Victoria, 3800 Finnerty Road, Elliot Building, Victoria, BC V8P 5C2, Canada, ⁶NRC Herzberg Astronomy and Astrophysics, 5071 West Saanich Road, Victoria, BC V9E 2E7, Canada, ⁷Center for Computational Astrophysics, Flatiron Institute, 162 5th Ave, New York, NY 10010, USA, ⁸Department of Astrophysics, American Museum of Natural History, New York, NY 10024, USA, ⁹Department of Science, Borough of Manhattan Community College, City University of New York, New York, NY 10007, USA, ¹⁰Space Telescope Science Institute, 3700 San Martin Drive, Baltimore, MD 21218, USA and ¹¹Department of Physics and Astronomy, Johns Hopkins University, 3701 San Martin Drive, Baltimore, MD 21218, USA

Abstract

Flying on board the *James Webb Space Telescope* (JWST) above Earth's turbulent atmosphere, the Aperture Masking Interferometer (AMI) on the NIRISS instrument is the highest-resolution infrared interferometer ever placed in space. However, its performance was found to be limited by non-linear detector systematics, particularly charge migration – or the Brighter-Fatter Effect. Conventional interferometric Fourier observables are degraded by non-linear transformations in the image plane, with the consequence that the inner working angle and contrast limits of AMI were seriously compromised. Building on the end-to-end differentiable model & calibration code *amigo*, we here present a regularised maximum-likelihood image reconstruction framework *dorito*, which can deconvolve AMI images either in the image plane or from calibrated Fourier observables, achieving high angular resolution and contrast over a wider field of view than conventional interferometric limits. This modular code by default includes regularisation by maximum entropy, and total variation defined with l_1 or l_2 metrics. We present imaging results from *dorito* for three benchmark imaging datasets: the volcanoes of Jupiter's moon Io, the colliding-wind binary dust nebula WR 137 and the archetypal Seyfert 2 active galactic nucleus NGC 1068. In all three cases, we recover images consistent with the literature at diffraction-limited resolutions. The performance, limitations, and future opportunities enabled by *amigo* for AMI imaging (and beyond) are discussed.

Keywords: Optical interferometry; astronomical detectors; James Webb Space Telescope; astrostatistics; astronomy image processing

(Received 19 October 2025; revised 26 February 2026; accepted 12 March 2026)

1. Introduction

Following its successful launch, the *James Webb Space Telescope* (JWST) immediately established itself as the world's leading observatory for infrared astronomy. Included in its suite of revolutionary instruments is the Near InfraRed Imager and Slitless Spectrograph (NIRISS; Doyon et al. 2023), hosting the Aperture Masking Interferometer (AMI; Sivaramakrishnan et al. 2023) as one of its four observing modes. The aperture mask, or Non-Redundant Mask (NRM), is a modest titanium disc small enough to fit into the NIRISS pupil wheel perforated with seven holes in a non-redundant pattern to enable interferometric observations.

The venerable technique of sparse aperture masking (or Fizeau interferometry; Fizeau 1868) blocks much of the telescope aperture, passing light only through a pattern of holes. Pairs of these create interference fringes in the point spread function, offering established calibration advantages compared to full-pupil imaging. If the array is non-redundant (Tuthill 2012), with no two holes

separated by the same vector, they form unique Fourier amplitudes (Twiss et al. 1960) and (importantly) closure phases (Jennison 1958; Baldwin et al. 1986) that are self-calibrating with respect to the effects of optical aberrations from the turbulent atmosphere and (relevant to JWST) imperfect/variable surface figure of mirrors. Closure phases can therefore provide extremely accurate constraints on faint companions and asymmetric structure in images, even in the presence of systematic optical errors.

Prior to launch, AMI was expected to debut as the most scientifically capable space interferometer to date, boasting unparalleled angular resolving power and contrast sensitivity (Sivaramakrishnan et al. 2010, 2012; Soullain et al. 2020; Tuthill, Sivaramakrishnan, & Lloyd 2022). A number of pipelines have been developed to analyse the AMI data: AMICAL (Soullain et al. 2020), SamPy (Sallum, Ray, & Hinkley 2022), and SAMPip (Sanchez-Bermudez et al. 2022) build on legacy Fizeau interferometry codes using the fast Fourier transform (FFT) to sample visibilities in the Fourier plane. In contrast, the forward-modelling code ImPlaneIA fits to the interferogram directly in the image plane (Greenbaum et al. 2015), and extracts visibility information from this fringe fit.

However, it turned out that all codes reached sensitivity limits in faint companion searches that were severely hindered

Corresponding author: Max Charles; Email: max.charles@sydney.edu.au

Cite this article: Charles M, Desdoigts L, Pope B, Tuthill P, Blakely D, Johnstone D, Ray S, Ford KES, McKernan B and Sivaramakrishnan A. (2026) Image reconstruction with the JWST interferometer. *Publications of the Astronomical Society of Australia* 43, e048, 1–15. <https://doi.org/10.1017/pasa.2026.10179>

by non-linear charge migration in NIRISS' HAWAII-2RG Complementary Metal-Oxide-Semiconductor (CMOS) detector (Sallum *et al.* 2024; Ray *et al.* 2025). Charge diffuses from brighter pixels into neighbouring pixels, resulting in bright features appearing broader or fatter; hence this charge migration has been dubbed the Brighter-Fatter Effect (BFE; Antilogus *et al.* 2014; Hirata & Choi 2020; Argyriou *et al.* 2023). AMI's image plane is barely Nyquist-sampled by the detector, which exacerbates its sensitivity to charge migration. Crucially, the non-linearity of the BFE cannot be calibrated in the Fourier plane alone. Repeated measurements sampling up-the-ramp has allowed for data selection mitigating the impact of BFE, yielding precise measurements of the proto-planetary disk and two companions in the PDS 70 system (Blakely *et al.* 2025). However, the lack of a general solution for BFE calibration has severely hindered science productivity of programs requiring precise calibration at diffraction limited scales.

In a companion paper (Desdoigts *et al.* 2025), we present a new pipeline, *amigo* (AMI Generative Optics), a Python package for end-to-end differentiable forward-modelling of AMI. This connects a differentiable optical model written in Jax (Bradbury *et al.* 2018) and ∂ Lux (Desdoigts *et al.* 2023; Desdoigts *et al.* 2024) with a neural network 'effective detector model' that models the BFE and other detector effects. These are simultaneously trained on an ensemble of calibration datasets, to arrive at a base model with fixed, fiducial optical and detector parameters that reflect the average state of the instrument. This base model can then be applied to retrieve astrophysical information from AMI data through forward-modelling and Bayesian inference with optimisation or sampling. Here we allow the surface figures of mirrors and the parameters of astrophysical scenes to vary while holding the mask metrology and electronic properties fixed. By directly modelling pixel-level data as they accumulate up-the-ramp, *amigo* is able to sidestep the challenges of Fourier-plane calibration and improve upon the performance of the JWST pipeline significantly. In the *amigo* companion paper, we apply this new pipeline to the recovery of faint companions at the diffraction *and* photon noise limit in the HD 206893 system (Hinkley *et al.* 2023) and in the binary AB Dor AC (Azulay *et al.* 2017).

In essence, *amigo* enables accurate modelling of the AMI point spread function (PSF) and astrophysical scenes composed of unresolved point sources. In this paper, we extend the capability of *amigo* to include arbitrary resolved scenes, implemented in the Python package *dorito* (Differentiable Optics for Reconstructing Images from Telescope Observations). We document and apply Regularized Maximum Likelihood (RML) image reconstruction methods fitting to either the image plane or interferometric observables. We explore several different methods for image reconstruction, finding that in most cases modelling the source distribution directly to the image plane rather than to visibilities or closure quantities yields the best performance.

We present imaging results from three public AMI datasets: the active galaxy NGC 1068, Jupiter's moon Io (previously published by Sanchez-Bermudez *et al.* 2025), and the colliding-wind binary Wolf-Rayet 137 (WR 137; previously published by Lau *et al.* 2024). These each serve to illustrate important points about the performance of imaging with AMI. The science and calibrator data are presented alongside their reconstructed images in Figure 1.

NGC 1068 is a well-known Active Galactic Nucleus (AGN; Antonucci & Miller 1985; Jaffe *et al.* 2004) whose inner regions have recently been imaged in the mid-infrared (Isbell *et al.* 2025) using the Large Binocular Telescope Interferometer (LBTI; Hinz

et al. 2016). Here, we recover a very closely matching structure in all three bands of AMI.

The innermost Galilean moon Io has a surface pocked with violent volcanism, visible from Earth in Keck adaptive optics images (Marchis *et al.* 2002; de Kleer *et al.* 2019) and close up from NASA Juno flyby observations (Mura *et al.* 2020). An independent analysis of the AMI images by Sanchez-Bermudez *et al.* (2025) using a neural network reveals bright spots at the locations of known volcanoes; we reproduce this result but at higher resolution, seeing a crisp edge to the disk of Io and well-resolved spots rotating over the hour-long observation.

Colliding-wind binaries like WR137 (Williams *et al.* 1985; Williams *et al.* 2001) produce spiral plumes of dust that were some of the first images to be resolved by aperture-masking interferometry (Tuthill, Monnier, & Danchi 1999). In AMI observations of WR137, Lau *et al.* (2024) show that its dusty environment is clearly resolved into a streak lined up with the expected elongation of the plume. An interesting and unfortunate quirk of this dataset is that, during the sequence of calibrator and science target observations with AMI, there was a 'tilt event' (Schlawin *et al.* 2023; Lajoie *et al.* 2023) in which one of the mirror segments abruptly moved, so that the wavefront encountered by the PSF reference calibrator differs from that of the science target. While in general we show in this paper that closure invariants perform poorly for anchoring image reconstructions with AMI compared to image plane modelling, in this case we find that they significantly improve the deconvolution with significant wavefront error present.

1.1. Differentiable forward-modelling of the instrument

This work employs a full end-to-end differentiable forward-model approach to image reconstruction. This section motivates this choice over alternatives such as an inverse model approach.

An instrument, or sensing system, delivers measurements $\mathbf{y}$ which encode some desired signal $\boldsymbol{\theta}$. The measurements and signal are related by some function f such that

$$\mathbf{y} = f(\boldsymbol{\theta}) + \boldsymbol{\epsilon}, \quad (1)$$

where $\boldsymbol{\epsilon}$ is some non-deterministic, independent, zero-mean noise term. Therefore an inverse problem must be solved to recover $\boldsymbol{\theta}$. The obvious approach is to construct an inverse function such that $\boldsymbol{\theta} = f^{-1}(\mathbf{y})$. However experience shows that this quickly proves arduous for complex non-linear problems and is particularly problematic when the effects of non-deterministic noise processes need to be compensated. Crucially, for the ill-posed problem of image reconstruction, an inverse function cannot be uniquely defined where f is not a one-to-one mapping.

The forward-model approach instead simply constructs the function f – an end-to-end forward-model of the instrument – to best emulate the deterministic physics at play. This is effective in a Bayesian framework as Bayesian methods can be used to infer information about the signal $\boldsymbol{\theta}$ from the data $\mathbf{y}$. The optimisation problem in eq:bayes1 can be re-expressed in terms of the forward-model as

$$\hat{\boldsymbol{\theta}} = \arg \max_{\boldsymbol{\theta}} \left[\underbrace{\mathcal{L}(\mathbf{y}|f(\boldsymbol{\theta}))}_{\text{log-likelihood}} + \underbrace{\Pi(f(\boldsymbol{\theta}))}_{\text{log-prior}} \right]. \quad (2)$$

Here $\boldsymbol{\theta}$ are the model parameters, which in the case of the *amigo* model includes the image parameters $\mathbf{c}$, other source parameters (such as detector position, flux), and optical/detector




Figure 1. Interferograms of the science targets, calibrators, and deconvolved images. *Top:* Images of the interferograms of science targets NGC 1068, Io, and WR 137. These images encode the slope of the ramp integrations, i.e. final group subtract second-to-final group. They contain noticeably less high frequency power by comparison to the *middle* row of PSF calibrators corresponding to the above sources. Bad pixels not used in the fit are set to black. *Bottom:* RML image reconstructions of all three targets. The beam size circle in the lower left of each panel indicates $\lambda/2D$. The colour tables for each image is a power law stretch ($\gamma = 0.4, 0.8, 0.3$, respectively, for each target left to right). NGC 1068 is clipped to 25% of its peak brightness in the **F480M** filter. Note WR 137 is displayed on a factor of two finer pixel scale than the other images due to its small angular size. NGC 1068 and WR 137 are both presented as RGB false-colour images weighted by the recovered flux values in each filter. However as WR 137 was only observed in the **F480M** and **F380M** filters, the green channel was assigned the average of the other two channels. The predominant white hue of these two images is a sign that the same structure is independently recovered in all three bands, a sign of successful deconvolution of a source without strongly wavelength-dependent features. The Io image is a single exposure image, ignoring the other epochs of data.

instrumental parameters (such as mirror aberrations, pixel gains, and charge migration). During optimisation, it is then possible to jointly solve for the instrumental state as well as the source parameters.

As the model complexity increases, so do the computational demands of optimisation. It is advantageous to build the forward-

model in a differentiable coding language such as Jax (Bradbury et al. 2018), PyTorch (Paszke et al. 2019), or TensorFlow (Abadi et al. 2016), as they leverage automatic differentiation (autodiff) to facilitate the use of gradient-based optimisation methods. These include straightforward optimisation engines such as gradient descent (DeepMind et al. 2020), but also encompass modern

efficient gradient-based Monte Carlo Markov Chain (MCMC) algorithms for error estimation such as Hamiltonian Monte Carlo (HMC) and No U-Turn Sampler (NUTS; Phan et al. 2019; Cabezas et al. 2024; Abril-Pla et al. 2023; Stan Development Team 2023). These languages also allow for Graphics Processing Unit (GPU)-acceleration, drastically increasing the tractability of training models with millions or billions of parameters on relatively accessible hardware.

Our RML deconvolution code *dorito* is inspired by MPoL, an RML imaging framework written in pytorch built for reconstructing radio/sub-mm interferometric data (Czekala et al. 2025). MPoL has been used to show the effectiveness of RML techniques improving image fidelity on ALMA continuum observations of protoplanetary disks (Zawadzki et al. 2023).

1.2. Data reduction and wavefront calibration

Optical interferometers must perform physical beam combination to interfere the light: for Fizeau interferometry (as performed by AMI) we refer to the fringe pattern on the detector as the interferogram. This pattern is the convolution of the PSF of the telescope with the source intensity distribution. The PSF is usually modelled as the Fourier (or Fresnel) transform of the incoming wavefront, which for high fidelity recovery requires the inclusion of phase distortions – whether from the turbulent atmosphere or, even in space, from optical path delays due to uneven or misaligned mirror surfaces.

Aperture synthesis is still possible in the presence of wavefront phase distortions by transforming the measured complex visibilities into data products that are less vulnerable to phase error. This includes closure quantities such as closure phases (Jennison 1958) or kernel phases (Martinache 2010), or extraction of squared visibilities from the power spectrum (which simply discards phase information).

In the workflow presented here, the first step of any data processing is to infer the wavefront. We load a pre-saved AMI base model trained as described in the *amigo* companion paper, and hold the parameters describing mask metrology and the effective detector model fixed. We take the PSF reference star observed along the science target, and fit for the wavefront across each aperture with the trained *amigo* base model. This is parametrised as a set of Zernike polynomials describing phase deviations across each hole of the aperture mask. We then hold these parameters fixed and analyse the science target, either directly modelling the source distribution as an array (Method 1, Section 2.1) or first extracting and appropriately calibrating complex visibilities (Method 2, Section 2.2).

1.3. Image reconstruction

While in the *amigo* companion paper we detect point source companions at high contrast, in this paper we aim to reconstruct complex resolved scenes. Image reconstruction is the process of solving the inverse problem of finding the image I which produced measured data y . Image reconstruction has a long history in both radio and optical interferometry (Baron 2016; Thiébaud & Young 2017), where a limited subset of the Fourier modes in an image are measured. In general, where information is lost to blurring or filtering (as in interferometry), the inverse problem is ill-posed; that is, there is no unique solution I which could produce y because many possible solutions fit the data equally well. Over the

last century, this has spurred the invention of many algorithmic approaches to find the best of the many possible solutions.

The simplest of these methods is direct aperture synthesis by linearly adding the interference patterns, where an array of telescopes (or telescope sub-apertures, in the case of AMI) are cophased to achieve an angular resolution equivalent to that of a single aperture the size of the longest baseline. If the sub-apertures are arranged in a non-redundant pattern such that no two are separated by the same vector, each interferometric baseline is unique and each telescope pairing will sample a different region of the u - v plane. Assuming the Fourier amplitudes are zero where the u - v plane has not been sampled, one can then inverse transform from the u - v plane back to the image plane to obtain a reconstructed image. In radio astronomy, this resultant image is referred to as the ‘dirty image’. Linear aperture synthesis alone can lead to images plagued by artefacts and sidelobes due to incomplete sampling of the Fourier plane, as is indicative by the ‘dirty image’ nomenclature. This can lead to the reconstruction of non-physical structure, however the dirty image serves as a starting point for more advanced methods.

One such older approach is the CLEAN method, which was designed to tackle interferometric data with sparse and irregular sampling of the u - v plane (Högbom 1974; Schwarz 1978; Clark 1980; Cornwell 2008). Commonly employed in radio astronomy, the CLEAN method is an iterative deconvolution algorithm that models the observed image as a collection of point sources and progressively subtracts copies of the instrument’s PSF (or dirty beam) to reconstruct a ‘cleaned’ image. However, CLEAN is sensitive to phase errors and implicitly assumes the scene is composed of point sources, so it is not optimal for reconstructing extended structure. Despite this, CLEAN is still widely and successfully used to reconstruct resolved scenes for longer wavelength observations such as ALMA (Wong et al. 2016) where u - v coverage is extensive and phase error is less destructive.

More recent Bayesian-based methods are better equipped to deal with phase errors and do not impose restrictive assumptions on the image. Within the Bayesian framework, the problem is re-framed as an optimisation or sampling problem inferring the parameters of a forward-model. One finds an image I which maximises the posterior probability, given the likelihood of that image matching the data and given prior assumptions. Mathematically, we wish to solve

$$\hat{I} = \arg \max_I \left[\underbrace{\mathcal{L}(y|I)}_{\text{log-likelihood}} + \underbrace{\Pi(I)}_{\text{log-prior}} \right], \quad (3)$$

where y is our observed data product (e.g. visibilities, closure phases) and $\hat{I}$ is the maximum a posteriori image. Different competing Bayesian approaches are generally distinguished by their implementation of the Bayesian prior, or we can step outside the Bayesian framework and add regularisers which may encode prior information but are not themselves admissible prior distributions.

1.4. Image basis

An implicit way to impose prior knowledge is by the parametrisation of the image (Thiébaud 2008). Consider a basis of n functions (also called a dictionary) used to represent an image:

$$z = \{z_j : \Omega \rightarrow \mathbb{R} \mid j = 1, 2, \dots, n\}, \quad (4)$$

with image coordinates α, δ in the field of view Ω , expressed in angular units. A continuous representation of the image I is then

written as a linear combination of these basis functions:

$$I(\alpha, \delta) = \sum_{i=1}^n c_i z_i(\alpha, \delta), \quad (\alpha, \delta) \in \Omega. \quad (5)$$

The coefficients $\mathbf{c} = (c_1, \dots, c_n)$ are the image parameters to recover in the reconstruction process. One could build a parametric model of the astrophysical scene from a set of astrophysical parameters, restricting the resulting ‘image’ to be of that form. An example of this is the AMI parametric model image of PDS 70 reconstructed by Blakely et al. (2025).

On the other hand, one could make no assumptions about the scene – the most obvious parametrisation here is the case where the image parameters are simply the pixel values. The pixel basis, sometimes referred to as the impulsion basis (Baron 2016), is expressed in the above notation as a set of Kronecker delta functions

$$z_j(\alpha, d) = \delta_{\alpha_j \alpha} \delta_{d_j d}, \quad (6)$$

where each function z_j is an impulse activated at the corresponding pixel centre (α_j, d_j) – here the latter coordinate is written as d to disambiguate from the Kronecker delta function, but is written as δ elsewhere.

Another option is a wavelet basis. This is a natural choice as most physical images, including astrophysical scenes, are sparse in a wavelet basis (Starck et al. 1994). Reconstruction in a wavelet basis with l_1 regularisation on the wavelet coefficients (see Section 1.5.1) will confine the space of possible solutions to those with more physical characteristics. The Mayonnaise image reconstruction pipeline (Pairet, Cantalloube, & Jacques 2021) – developed to reconstruct near-infrared circumstellar disk and exoplanet images in the context of high contrast angular differential imaging – successfully implements a shearlet basis, which further extends the concept of wavelets to an isotropic basis.

All images presented in this paper are reconstructed in a pixel basis, as described by Equation (6). Hence, the recovered image parameters $\mathbf{c}$ are simply the pixel values of the image, and there is no significant prior constraint on the resultant images due to the choice of image basis. We thus turn to image regularisation to apply prior constraint to the images.

1.5. Regularisation

The explicit way to incorporate prior knowledge is through the choice of log-prior function Π in Equation (3). A common implementation of Π is with a regularisation function $\mathcal{R}$, which provides a measure of the discrepancy between a proposed image and its expected characteristics (Renard, Thiébaud, & Malbet 2011). Regularisation directly addresses the ill-posed nature of the inverse problem, and is especially necessary when working with sparse data. It can help to ‘paint in’ the regions of the u - v plane which are not constrained directly by the data, as opposed to assuming they are zero as in simple aperture synthesis.

Depending on the choice of $\mathcal{R}$, regularisation penalises a candidate image for not satisfying certain desired image properties expected of the scene; e.g. smoothness, soft edges, or sparsity. This is crucial in filtering out non-physical solutions from the solution space which would satisfy the data penalty alone. Here, we optimise not just a log-likelihood, but a more general loss function. The contribution of the regularisation function to the loss function is weighted by a manually-set hyperparameter λ

$$\text{Loss} = -\mathcal{L}(\mathbf{y}|\mathbf{I}) + \Pi(\mathbf{I}) = -\mathcal{L}(\mathbf{y}|\mathbf{I}) + \lambda \mathcal{R}(\mathbf{I}), \quad (7)$$

where $\mathcal{I}$ is the image I integrated onto a discrete pixel grid, with pixel indices (i, j) . This is the backbone of RML methods and is leveraged by many existing interferometric image reconstruction packages such as *bsmem* (Buscher 1994), *Mira* (Thiébaud 2008), *Squeeze* (Baron, Monnier, & Kloppenborg 2010), *Painter* (Schutz et al. 2014), *Wisard* (Meimon, Mugnier, & Le Besnerais 2005), and *MPoL* (Czekala et al. 2025), and was used in the reconstruction of the black hole images from the Event Horizon Telescope Collaboration et al. (2019a,b).

In order for a regularisation function $\mathcal{R}$ to qualify as a true Bayesian prior distribution, it must be a valid probability distribution. Regularisers which do not meet this condition (such as maximum entropy) are not compatible with true Bayesian methods that characterise the image posterior distribution such as HMC (Willis, Jeffs, & Long 1999; Baron 2016; Sallum & Eisner 2017). However these non-Bayesian regularisers are still useful if the method only seeks a maximum a posteriori solution, such as gradient descent.

1.5.1. Common image regularisers

The choice of both $\mathcal{R}$ and λ depends on judgement. Listed below are some common choices for the regularisation function $\mathcal{R}$. Note that when reconstructing in a pixel basis, the regularisation is being applied to the pixel values themselves, which is the case we will discuss. Different image bases have significantly different responses to the same regulariser.

Maximum Entropy

$$\mathcal{R}_{\text{MaxEnt}}(\mathcal{I}) := - \sum_{i,j} \mathcal{I}_{i,j} \ln \mathcal{I}_{i,j} \quad (8)$$

A maximum entropy or ‘MaxEnt’ regulariser seeks to maximise the total Shannon entropy of the image (Shannon 1948). An image with the maximum entropy has the minimum amount of configurational information, and thus there must be evidence in the data to justify the presence of any structure (Skilling & Bryan 1984). A pioneering implementation of maximum entropy methods to optical interferometry was the code *bsmem* (Buscher 1994), with the method accumulating a proven track record.

l_p norms

$$\mathcal{R}_{l_p}(\mathcal{I}) := \left[\sum_{i,j} |\mathcal{I}_{i,j}|^p \right]^{1/p}, \quad p \geq 1 \quad (9)$$

l_p norms generalise the idea of distance to depend on different powers of the coordinates, where if $p = 2$ we recover the familiar Euclidean distance. If we introduce a regularisation term to minimise the l_2 norm of the coefficients of a model, this is often called ridge regression (Hoerl & Kennard 1970a,b) or automatic relevance determination, which can be seen as a normal prior on the model coefficients centred at zero that suppresses large values of the coefficients (Renard et al. 2011).

Regularisation with an l_1 norm, which in linear models is called Least Absolute Shrinkage and Selection Operator (LASSO; Tibshirani et al. 1996) regression, can be interpreted probabilistically as a zero-centred exponential prior on the absolute value of the coefficients. This has the effect of promoting sparsity in the parameters; i.e. if possible, many or most of the parameter values will go to zero, with a few nonzero parameter values singled out as relevant. This is an important fundamental concept in the field




Figure 2. L-curve diagram used to select the optimal regularisation parameter λ_{TV} for images of Jupiter's moon Io from AMI data. Each point on the curve is the balance between regulariser term $\mathcal{R}$ and likelihood term $\mathcal{L}$ of a converged image reconstruction for a different regularisation hyperparameter value λ_{TV} . Shown are several reconstructed images corresponding to different points along the curve. The effects of TV regularisation on the image can be seen to strengthen with increasing λ_{TV} (as regularisation increases, ringing artefacts reduce and plateaux of uniform flux with sharp edges form). The optimal value for λ_{TV} is selected from the region of the elbow in the curve (the third image). L-curves for NGC 1068 and WR 137 are included in [Appendix A](#).

of compressed sensing (Candes, Romberg, & Tao 2006; Candes & Wakin 2008). In a pixel basis, l_1 encourages a dark field populated by a few bright pixels, which is not useful for the extended structure of our three targets here; but many images are sparse in an appropriate basis, such as a wavelet basis, and l_1 regularisation on these coefficients may be useful in future work.

Quadratic Variation (QV)

$$\mathcal{R}_{QV}(\mathcal{I}) := \sum_{ij} (\mathcal{I}_{i+1,j} - \mathcal{I}_{i,j})^2 + (\mathcal{I}_{i,j+1} - \mathcal{I}_{i,j})^2 \quad (10)$$

Quadratic variation, also known as total squared variation or quadratic smoothness, is the l_2 norm on the spatial gradient of the image. It is computed over the finite difference gradient of the image as above. As QV penalises large gradient variations, it discourages neighbouring pixels to differ in value and thus smooths the image.

Total Variation (TV)

$$\mathcal{R}_{TV}(\mathcal{I}) := \sum_{ij} |\mathcal{I}_{i+1,j} - \mathcal{I}_{i,j}| + |\mathcal{I}_{i,j+1} - \mathcal{I}_{i,j}| \quad (11)$$

Total variation is the l_1 norm on the spatial gradient of the image. The expression for this is written above, however the absolute value function is not differentiable about zero. Consequently, gradient-based parameter optimisation methods such as gradient descent may exhibit unusual behaviour about this point due to numerical instabilities. It is thus preferable to instead use the following:

$$\mathcal{R}_{TV}(\mathcal{I}) := \sum_{ij} \sqrt{(\mathcal{I}_{i+1,j} - \mathcal{I}_{i,j})^2 + (\mathcal{I}_{i,j+1} - \mathcal{I}_{i,j})^2}, \quad (12)$$

as this is isotropic, differentiable at zero, and approaches the true l_1 norm with increasing parameter values.

Like QV, TV attempts to reduce variation in the image. However, it is made distinct from QV by preferring sparse spikes

in gradient between flat regions as opposed to a gentle slope. TV favours images with plateaux of flux, and consequently edges are well-preserved (Strong & Chan 2003). Over-regularisation with total variation will result in an image with cartoonish features.

1.5.2. Hyperparameter selection

To choose the hyperparameter λ is to walk the line between under-regularisation (where the solution matches the data well but may be unphysical) and over-regularisation (where the solution no longer matches the data). Visual inspection is almost always sufficient, as the negative effects of under- and over-regularisation are not subtle and there usually exists a wide interval of λ -values which strike a balance (Thiébaud & Young 2017).

A less heuristic approach is the L-curve method, which involves plotting the log-likelihood term $\mathcal{L}$ against the log-prior term Π of an optimised image for a range of λ -values (Hansen 1992). In a well-behaved loss space, varying λ should trace out a curve in the shape of an 'L', and the optimal value of λ corresponds to the elbow of the curve. This ensures the largest effect of the regulariser without incurring an excessive data penalty. An example L-curve used for regularisation hyperparameter selection in this work is shown in [Figure 2](#), for the case of Jupiter's moon Io. Additionally, L-curves for both NGC 1068 and WR 137 are included in [Appendix A](#).

2. Image reconstruction methods

In this paper, we employ two methods for image reconstruction of AMI data. Both methods use gradient descent on a pixel basis to obtain a maximum a posteriori image, with a Gaussian log-likelihood function.

To avoid covariance with the flux parameter, the source distribution is normalised such that it sums to unity on every iteration of the gradient descent loop. When fitting undithered data with a large number of gradient descent iterations, these methods tend to




Figure 3. Flow diagram depicting the Method 1 image plane-based reconstruction process. The source is modelled via a convolution with a source distribution array and is fit to the ramp data with gradient descent. The final image is obtained from the science source distribution after a specified number of gradient descent iterations are complete.

concentrate flux sparsely in a few bright pixels. To mitigate this, the source distribution is convolved with a small Gaussian kernel ($\sigma \approx 0.25$ pix) each fitting epoch. The image parameters (pixel values) are fit in log-space to enforce image positivity. After the fit, the log distribution is simply exponentiated to provide the final reconstructed image $\hat{I}$.

The primary feature that differentiates the methodologies introduced here is the data product used in the calculation of the likelihood. That is, which data product is y in Equations (2) and (3).

2.1. Method 1: Image plane method

In the image plane method, the data product y that is used for the likelihood is the ramp data. The interferogram is modelled by a convolution between the source distribution and the PSF, as opposed to complex visibilities which is the *amigo* default. Complex visibilities are not used at all in this method, thus the fit is performed entirely in the image plane. This fitting process is depicted in Figure 3. The parameters of the *amigo* model which are fit to the ramp data are the

- positions (per exposure),
- fluxes (per exposure),
- spectra (per target per filter),
- log distribution $\log_{10} \mathcal{I}$ (fit only to the science target per filter, in log space to enforce image positivity),
- mirror aberrations $\mathbf{Z}$ (fit only to the calibrator target per filter).

By avoiding the Fourier domain, Method 1 targets the scenario in which the limiting factor affecting image quality is detector non-linearities. Additionally, unlike closure phase methods, this image plane method inherently incorporates low spatial frequency information through short redundant baselines, as discussed in the following section.

2.2. Method 2: DISCO method

As opposed to Method 1, the Delay-Insensitive Subspace of Calibrated Observables (DISCO) method uses the Fourier domain to characterise the source distribution. While using Fourier quantities is a more typical approach to image reconstruction, the use of DISCOs is a novel method introduced by *amigo*.

While AMI generates closure phases (Jennison 1958) between triplets of its holes, the holes themselves have a substantial width in comparison to their separation. Despite the arrangement of the NRM holes into a non-redundant configuration, the non-zero finite hole area results in short redundant baselines within an individual hole which transmit low spatial frequency modes. We would like to access this information to accurately recover the low spatial frequencies in the image, however closure phases are only valid for non-redundant apertures. Furthermore, one of the findings of Desdoigts et al. (2025) is that the aperture mask is not quite conjugate to the telescope pupil plane and a small Fresnel defocus is required; this near-field effect means that the closure phases are no longer invariant. As a result of this redundancy and defocus, the standard closure phases are no longer robust against wavefront error, calling for a more general treatment.

We therefore adapt the kernel phase and amplitude formalism (Martinache 2010; Pope, Martinache, & Tuthill 2013; Pope 2016), which generalises closure invariants to redundant apertures:

$$\Phi^{\text{meas}} = \Phi^{\text{sky}} + \mathbf{J} \cdot \phi, \quad (13)$$

with a Jacobian matrix $\mathbf{J}$ that describes the transfer of phase errors ϕ on the pupil into u - v visibility phase measurements Φ^{meas} of the astrophysical information. That is, $\mathbf{J}$ quantifies the response of the u - v plane to perturbations of pupil plane phase error modes; here we assume these phase errors can be described with a linear combination of n Zernike modes. Thus $\mathbf{J}$ is the matrix for partial derivatives of basis vectors spanning the u - v plane with respect to the first n Zernike polynomials.

If we then generate a null operator $\mathbf{K}$ (also known as kernel operator) for the Jacobian matrix $\mathbf{J}$, we can use this to generate self-calibrating kernel phases

$$\mathbf{K} \cdot \Phi^{\text{meas}} = \mathbf{K} \cdot \Phi^{\text{sky}} + \underbrace{\mathbf{K} \cdot \mathbf{J} \cdot \phi}_{=0}, \quad (14)$$

i.e. the null operator nulls out the Zernike modes that transfer phase errors ϕ to the u - v plane. This leaves the kernel phases $\mathbf{K} \cdot \Phi^{\text{meas}}$ which contain less information than Φ^{meas} itself as all phases that could be expressed as a linear combination of the first n Zernike polynomials are mapped to zero. However this is an advantageous trade-off as the kernel phases should be invariant to wavefront phase error. Although we describe the phase as an expansion in the Zernike basis, the mathematics is the same as the classic kernel phase derivation of low amplitude piston phases on small sub-apertures. We do not treat amplitude error in the pupil; in developing *amigo* we allowed pupil transmission to be a free parameter, but did not consistently recover significantly nonzero values. Furthermore, by including a wavelength term in the Jacobian we can also null the effects of spectral miscalibration when applying the kernel operator.

We can generate the transfer matrix $\mathbf{J}$ not from a geometric model and FFT, which is prone to sampling and interpolation issues (Martinache et al. 2020), but rather by forward-modelling the system and using Jax autodiff to generate the Jacobian (following Pope et al. 2021).




Figure 4. Flow diagram depicting the Method 2 DISCO-based image reconstruction process, broken into three stages. In the first, the source distribution is forward-modelled with complex visibilities and then transformed to the DISCO basis in the second. This is identical to the fitting processes in Desdoigts et al. (2025), and is described in further detail there. Thirdly, the image reconstruction takes place as a source distribution array is fit to the reduced DISCOs.

In the DISCO method introduced in Desdoigts et al. (2025), the *amigo* forward-model with Fresnel diffraction means that phase errors in the pupil generate both phase and amplitude variations in the visibilities, and therefore kernel phase and kernel phase-to-visibility matrices. We extract complex visibilities through forward-modelling with Bayesian inference, by optimisation and then applying the Laplace approximation to estimate the covariance of these parameters (Kass, Tierney, & Kadane 1991). We keep track of covariance as we generate these kernel observables, and then re-diagonalise so that they are statistically independent following Ireland (2013). We refer to these observables as DISCOs to reflect that these are neither the previously-used kernel phase nor amplitude observables, as they include not just near-field effects but also spectral miscalibration.

To perform image reconstruction from DISCOs, there are three separate stages, with two separate fitting loops. This method is depicted in Figure 4. The first stage is a ramp fit similar to Method 1. However, rather than modelling a source distribution directly, the resolved component is modelled with complex visibilities. Thus the *amigo* model parameters that are fit to the ramp in the first stage are the

- positions (per exposure),
- fluxes (per exposure),
- spectra (per target per filter),
- log visibility amplitudes $\text{Re}(\ln V)$ (per exposure),
- visibility phases $\text{Im}(\ln V)$ (per exposure),
- mirror aberrations Z (fit only to the calibrator star).

The second stage is the transforming of the complex visibilities from the u, v basis to the DISCO basis. This involves applying the kernel operator K (which is generated from the fitted mirror aberrations Z), then diagonalising by the tracked covariance matrix. The third and final stage is where the image reconstruction is performed. Rather than propagating a source distribution through the AMIGO model as in Method 1, the source distribution is transformed to the DISCO basis and fit with the reduced DISCOs from the second stage.

DISCOs are designed to be resistant to wavefront error, and thus Method 2 is expected to be most effective when limited by residual wavefront error. Since DISCOs are invariant to Zernike tip/tilt modes by design, they are also invariant to translations

in the image plane. This introduces a heavy covariance to the pixel parameters in the third stage fit, as there is nothing to constrain the image centroid position. To convergence, a tight prior is added enforcing the image centroid position to be at the image centre.

3. Datasets

In this work, the two methods described in sec:methods are each applied to AMI datasets of three different science targets. The targets were chosen as they are publicly accessible datasets that offer a diversity of astrophysical morphologies to test our methods. NGC 1068 has an extended complex structure behind a bright central core, Io is heavily resolved with large regions of uniform flux and scattered point sources, whereas WR 137 pushes the angular resolution limits of our methods as it features a bright linear plume that is only partially resolved. The interferograms for each science target and its associated PSF reference star are displayed in Figure 1. All observations were taken on NIRISS' SUB80 sub-array with the NISRAPID readout pattern. All targets were observed through some subset of the following three near-infrared filters:

- F380M($\lambda = 3.825\mu\text{m}$, $\Delta\lambda = 0.205\mu\text{m}$),
- F430M($\lambda = 4.285\mu\text{m}$, $\Delta\lambda = 0.213\mu\text{m}$),
- F480M($\lambda = 4.815\mu\text{m}$, $\Delta\lambda = 0.303\mu\text{m}$).

The other filter available on AMI, F277W, is not Nyquist-sampled, and has not yet been used for imaging. Additionally, the neural network predicting BFE charge bleeding in the *amigo* base model was trained on calibration data which only reached $\sim 45\%$ full pixel well depth (or $\sim 30\text{ k counts}$). Hence, it is not reliable in predicting BFE above this, so here we truncate the ramp of the datasets to stay below this limit by discarding the high flux groups.

3.1. NGC 1068

NGC 1068, or M77, is the prototypical Seyfert 2 AGN located at $\sim 14\text{ Mpc}$ distance, powered by accretion onto a supermassive black hole (SMBH) of $\sim 3 \times 10^7 M_{\odot}$ (Ford et al. 2014). Seyfert 2 AGN are believed to be observed edge-on, obscuring our view of the central engine at short wavelengths but yielding an IR-bright

source in the dusty parsec-scale torus (Padovani et al. 2017). The details of how AGN such as NGC 1068 are fuelled remains unclear as does the effect of luminous feedback from an AGN on its host galaxy.

NGC 1068 is also confidently associated with the detection of TeV neutrinos, implying the interaction between hadrons in a relativistic outflow and surrounding dense material at luminosity $L_\nu \sim 3 \times 10^{42}$ erg/s (Inoue, Khangulyan, & Doi 2020; IceCube Collaboration et al. 2022), potentially from accreting stellar mass black holes embedded within the AGN disk (Tagawa, Kimura, & Haiman 2023).

NGC 1068 and its calibrator HD 17130 were observed on two separate occasions (September 2024, August 2025) by AMI as part of the GTO program 1260 (PI: K.S. Ford). Both targets were observed with all three filters (F380M, F430M, and F480M) in a five-point subpixel dither pattern. The first epoch of science data saturated the detector with 21, 25, 20 groups and 141, 112, 91 integrations for the F380M, F430M and F480M filters respectfully. The second epoch used 7 groups for all filters and 134, 139, 100 integrations; for both epochs, we only fit to the first 4 groups of the ramp to stay below ~ 30 k counts.

The region of interest to these observations has been imaged by LBTI in the mid-IR, which resolved a general N-S outflow, with clumps more evident in the N outflows (Isbell et al. 2025); we shall directly compare our image reconstruction results with that of Isbell et al. (2025).

3.2. Io

Io is the innermost of the Galilean moons of Jupiter and is hailed as the most volcanic place in the solar system. This is a result of the extreme tidal forces from Jupiter and gravitational interactions with the other Galilean moons as they exhibit orbital resonance. Erupting volcanoes are always visible on the surface of Io, shining especially bright in infrared light. The infrared surface of Io has been well characterised by the Juno spacecraft, which has been orbiting Jupiter since 2016 (Adriani et al. 2017; Davies et al. 2024a).

On August 1, 2022, Io and its calibrator HD 2236 were observed using the AMI mode. This was part of a suite of observations of the Jovian system across a broad range of observing modes to demonstrate JWST's capability for solar system science (DD-ERS program 1373, PI: I. de Pater). Observations of Io were taken only through the F430M filter in five subpixel dither positions, each consisting of 45 groups and 100 integrations. This also saturated the detector, so only the first 18 groups of the ramp were used in order to stay below ~ 30 k counts.

The JWST-Io-Sun phase angle is always small (10.3° , for these data) due to JWST's solar proximity relative to the Jovian system, so we do not expect to resolve a terminator. Each observation was taken 8.55 min apart on average, during which Io rotates 1.21° on its axis (King & Fletcher 2023). This temporal rotation is significant enough to be resolved between the dithered exposures. Thus instead of jointly fitting to all five dithered exposures, we treat each exposure as its own undithered image to create a time-evolving image. The calibrator star HD 2236 was observed with the identical observing configuration, at 4 primary dither positions each with 12 groups and 8 integrations, amounting to only 32.04 s total exposure time.

In a recent analysis of these data, Sanchez-Bermudez et al. (2025) applied several neural network and filter architectures to

image reconstruction, detecting several hot-spots and identifying these with volcanic features on Io. They do not cleanly resolve the limb of the moon against the black backdrop of space, and registration of the reconstructed image to the coordinate system of Io is made with respect to these bright features. One volcanic feature designated 'Seth Patera' was found to be extended, and the overall surface was seen to rotate over the course of the exposure.

3.3. Wolf-Rayet 137

The colliding wind binary system WR 137 is comprised of a carbon-rich (WC-type) Wolf-Rayet together with a hot O9-spectrum companion (Williams et al. 1985, 2001). Classified as an episodic dust producer, variation in the strength of the colliding wind interaction mediated by orbital phase results in a ~ 13 yr periodicity evident in the long term infrared light curve. In common with other such 'pinwheel' systems, the dust streams outwards carried on the fast winds, creating expanding structures that engrave the physics of wind and orbit in nested shells with spiral morphology that may be bright in the near- and mid-infrared.

Over three observing runs in 2022, WR 137 and its calibrator target HD 228337 were observed by AMI as part of DD-ERS program 1349 (PI: R. Lau). Primary-dithered observations were taken through the F380M and F480M filters. These observations did not significantly exceed 30 k counts so all groups up the ramp were used here. Further details about the observation configuration/schedule are given in Table 1.

This schedule is significantly more complicated due to a mirror tilt event during the first observing run, after almost all the science data had been taken but none of the calibrator data. After the mirror tilt event, this run was abandoned and rescheduled without taking any calibrator observations. Such a rapid change in optical state renders PSF reference data unhelpful: it would embody a differently aberrated wavefront to the science target. The second observing run was performed two days after the mirror tilt event, now including the calibrator and science target.

These observations have previously been analysed by Lau et al. (2024), using a variety of extraction pipelines (ImPlaneIA, AMICAL, and SAMPip: Greenbaum et al. 2015; Soulain et al. 2020; Sanchez-Bermudez et al. 2022) and image reconstruction algorithms (BSMEM, Squeeze, and IRBis: Buscher 1994; Baron et al. 2010; Hofmann, Weigelt, & Schertl 2014).

4. Results

Reconstructions of all three targets are shown in the final row of Figure 1. For all reconstructions, the source distribution was initialised as a uniform array so that the reconstruction process remains unbiased. The images of NGC 1068 and Io were reconstructed using Method 1 with TV regularisation, while WR 137 was reconstructed using Method 2 with MaxEnt regularisation.

4.1. NGC 1068 results

Figure 5 shows our reconstructed image with LBTI results from Isbell et al. (2025) overlaid in contours. The LBTI image reconstruction at $8.7 \mu\text{m}$ neatly covers our entire reconstructed image. We find excellent agreement overall, with some additional details apparent (and noting that exact agreement is not expected due to the significantly different observing wavelength). The obscuring torus at the heart of NGC 1068 lies $\sim -0.1'$ below the bright central core of this image in a SE-NW orientation

Table 1. The JWST/AMI observing schedule for WR 137 and its calibrator star (HD 228337). Each row represents an individual exposure. All dates are in 2022, when mirror tilt events were more common as the JWST mirror segments were still settling. The horizontal dashed line between July 13 and July 15 indicated the time at which the mirror tilt event occurred. The dither column refers to the dither position within the primary dither pattern, with the stare remaining in the centre of the sub-array. The number of groups and integrations of the exposures are given by n_g and n_{int} . The full integration time of an exposure is given by t . Roll is the roll angle of the telescope with respect to north.

Date	Star	Filter	Dither	n_g	n_{int}	t [min]	Roll [°]
Jul 13	WR137	F480M	1	4	400	2.0	192.0
			2				
			3				
			4				
		F380M	1	2	680	1.7	
			2				
			4				
Jul 15	CAL	F480M	Stare	7	1020	9.0	189.8
		F380M		2	2720	6.8	
	WR137	F480M		4	1600	8.0	
		F380M		2	2720	6.8	
Aug 8	WR137	F480M	1	4	400	2.0	162.7
			2				
			3				
			4				
		F380M	1	2	680	1.7	
			2				
	CAL	F480M	3	7	255	2.2	
			4				
			1				
			2				
		F380M	3	2	680	1.7	
			4				

(Gómez Rosas *et al.* 2022). The right-hand boundary of the LBTI image, coincides approximately with one edge of the $\sim 60^\circ$ opening angle $[O_{III}]$ emission cone observed by Macchetto *et al.* (1994), which presumably defines the region of low density gas exposed to the AGN radiation field. The 5 GHz radio emission from the ~ 100 pc nuclear radio jet extends from the torus location to the NNE in this image (through $0.2'$, $0.6'$). Thus, the central brightest part of our image is an off-axis hotspot around ~ 10 pc projected distance from SMBH due to outflow interaction with the dusty environment outside the torus.

Dusty gas is likely the source of much of the reconstructed emission signatures. From Mor & Netzer (2012), we might expect the dust sublimation radius R_{sub} to lie a distance

$$R_{sub} \sim 0.4 \text{ pc} \left(\frac{L_{bol}}{10^{45} \text{ erg/s}} \right)^{1/2} \left(\frac{T_{sub}}{1500 \text{ K}} \right)^{2.6} \quad (15)$$




Figure 5. A comparison of the NGC 1068 images recovered in AMI with the LBTI observations convolved by Isbell *et al.* (2025), updated with recent data by private communication. The three bands of AMI data in F480M, F430M, and F380M are represented as red, green, and blue in a false-colour image, with the LBTI 8.7 μm image flux overlaid as contours. The beam size circle in the lower left indicates $\lambda/2D$. The bright parts of the AMI colour image are close to white, indicating a consistent recovery across all three bands, and they track the bright parts of the LBTI image very closely, indicating both sets of data are independently recovering the same structure.

from the SMBH. Thus, dust at ~ 10 pc should survive being exposed to the inferred radiation field and reprocess radiation into the IR. Part of the hotspot emission may of course include both an overdensity of dusty gas as well as shock-heating and destruction of any dust overdensity. Since outflowing winds are typically $\mathcal{O}(10^3)$ km/s, a hotspot at $\mathcal{O}(10)$ pc implies that this episode of the AGN has been running for at least 10^4 yr. Likewise the overdensities within ~ 100 pc illuminated by outflow imply that this AGN episode has been running for ~ 0.1 Myr. Assuming a uniform accretion rate during that time implies a mass accretion of $\sim 10^4 M_\odot$ onto the SMBH, with a sizeable (comparable) mass of gas flung back outward in an outflow. More generally, drawing all of these observations together suggests that the feedback from this $L_{bol} \sim 10^{45}$ erg/s AGN is not powerfully impacting the surrounding galaxy. Dusty gas is still present within, or at the boundary of, the $\sim 60^\circ$ outflow region from the AGN.

4.2. Io results

We have reconstructed images of Io using Method 1 with TV regularisation, with all five epochs shown in the top row of Figure 6. Io is *heavily* resolved as seen by AMI, and visibility extraction in the forward-modelling approach has performed poorly; we have been unable to get any adequately converged image using Method 2 with a variety of initialisations and hyperparameter tunings. In the Method 1 images we detect five known Ionian volcanic surface features at their expected positions: (1) Seth Patera, (2) P197, (3) Masubi, (4) Leizi Fluctus, (5) Amirani (Davies *et al.* 2024b; Davies 2007). There is a visual agreement between the expected positions and reconstructed bright points on Io's surface, as was also found by Sanchez-Bermudez *et al.* (2025).




Figure 6. Reconstructions of each of the five exposures of Io, chronologically left to right. *Top row:* shows the reconstructed images, which were reconstructed with Method 1 using TV regularisation. The rotation of Io on its axis becomes visible when displayed as a time series animation, [hosted online](#). *Second row:* shows those same images with an ephemeris overlay, including the expected positions of five Ionian volcanic surface features. These are (1) Seth Patera, (2) P197, (3) Masubi, (4) Leizi Fluctus, (5) Amirani (Davies et al. 2024b; Davies 2007). *Bottom axis:* a time axis showing the times over which the exposures were taken.

Each volcano appears to be surrounded by a valley of flux. These are presumably not physical and are simply ringing artefacts due to minor remaining PSF miscalibration. This effect is somewhat suppressed by the TV regularisation; this can be seen in Figure 2 as the less regularised images exhibit more severe ringing artefacts.

Over time, the moon can be seen to rotate subtly on its axis – a feature for which the model had no a priori information or bias. This can be seen by tracking the Ionian meridians in the ephemeris overlay, but it is more identifiable when the exposures are played as frames in a time series. The detection of rotation is discussed in detail by Sanchez-Bermudez et al. (2025), and we do not seek to interpret this scientifically; it is unambiguously consistent with the rotating frame of Io visually in Figure 6 and in an animation of the rotation hosted online.

We do not see evidence for a more diffuse emission around the brightest feature Seth Patera, which is observed both in the Sanchez-Bermudez et al. (2025) AMI reduction and corresponding Keck observations (Marchis et al. 2002; de Kleer et al. 2019). In the Keck imaging presented therein, the emission is more extended at shorter wavelengths, and may be less apparent in F430M. It is possible that the published extended structure around this hotspot may be either miscalibrated BFE or PSF structure.

4.3. WR 137 results

We have reconstructed images of WR 137 using Method 2 with MaxEnt regularisation, shown in Figure 1. This image was jointly fit exclusively to exposures taken after the mirror tilt event on the second and third observing epochs (July 15 and August 8, see Table 1).

As shown by the contour overlay comparison in Figure 7, our results independently confirm the earlier imagery of Lau et al.




Figure 7. Comparison between a reconstructed image of WR 137 from Lau et al. (2024) (contours) and this work (image), shown on a logarithmic colour stretch. Both images are reconstructed from the same dataset using the F480M exposures after the mirror tilt event. The image from Lau et al. (2024) was reconstructed using Squeeze on observables extracted with AMICAL, which we considered the cleanest of the reconstructions given by Lau et al. It was convolved with a small Gaussian filter ($\sigma = 7.42$ mas) to declutter the contour lines for visualisation purposes. The image from this work was reconstructed with DISCOs using Method 2. The beam size circle in the lower left indicates $\lambda/2D$. The two images closely match in structure, including very faint patches of flux near $\sim 10^{-4}$ which are likely common non-physical miscalibration artefacts.

(2024). Our reconstruction shows a linear plume along an ESE-WNW orientation, present in all Lau et al. (2024) reconstructions. The length and orientation of this plume are consistent with models for a colliding-wind binary pinwheel nebula seen side-on. For




Figure 8. Images of the WR 137 dataset affected by a tilt event, deconvolved with different methods. *Left:* deconvolved with Method 1, in the pixel domain. The tilt event and consequent PSF misspecification mean that there are chromatic speckles systematically introduced into the image. *Right:* deconvolved with Method 2, from DISCOs. Even though this method has performed poorly on the other datasets, it accurately retrieves the simple plume shape of WR 137 without PSF artefacts, as the DISCOs act like closure phases and are insensitive to changes in the optical path difference between target and calibrator by design.

further astrophysical interpretation, the reader is referred to the excellent discussion found in Lau *et al.* (2024).

Additionally, to assess the robustness of DISCOs against wavefront error, we reconstructed images using both Method 1 and 2 from the science exposures taken on the first observing epoch (July 13) before the mirror tilt event. As there was no calibrator target observed that day, these science exposures were calibrated by the calibrator exposures taken *after* the mirror tilt event on July 15. Both reconstructed images are shown in Figure 8.

As expected, Method 2 proved to be superior in dealing with wavefront error. The image reconstructed from Method 1 is plagued by numerous spectrally-dependent non-physical miscalibration artefacts. By contrast, the Method 2 image is relatively free from these artefacts and is visually similar to the image created from the other two epochs in Figure 1.

5. Conclusions

This paper has presented *dorito*, an image reconstruction pipeline for recovering complex scenes from NIRISS/AMI observations. It has demonstrated recovery of diffraction-limited images, which are consistent with the highest-quality ground-based and space-based images of the same targets and of a generally improved quality compared to published results from the same data.

As the mask holes are relatively large given their separations, AMI passes a large number of Fourier modes to the interferogram and the image reconstruction problem is not nearly so under-constrained when working with forward-modelling the full u - v plane as it is when only working with the 21 baselines. In general, we have been able to deconvolve acceptable images of the moderate-contrast sources presented in this paper without regularisation at all, and have only noted modest benefits from adding these terms in. We expect regularisation to become more important at high contrast (where the signal-to-noise ratio is lower) and for very complex scenes with significant high spatial frequency

detail (where it is therefore necessary to incorporate more Fourier modes).

As noted by De Furio *et al.* (2025), the frequent wavefront sensing measurements taken by JWST make it possible – in principle – to do calibrator-free interferometry. With a drift rate of <10 nm rms per 48 h (Lajoie *et al.* 2023), aside from tilt events it is likely that our gnosis of the wavefront at any point in time is almost as good as would be achieved with a point source reference in the absence of any other information. While for the time being we recommend CAL-SCI or CAL-SCI-CAL observing sequences in line with the JWST proposal requirements, we recommend benchmarking this option as a step towards significantly reducing the time overheads and constraints required for AMI observations.

This work has revealed that it is not in general necessary to use DISCOs; rather the more straightforward approach of modelling the source in the image plane performs well in most cases. Nulling all the linear combinations of visibilities affected by wavefront miscalibration, as performed by the DISCOs algorithm, is extreme. While in the *amigo* companion paper (Desdoigts *et al.* 2025) we find this enhances our sensitivity to faint point sources, we believe that in the present paper we lose too many Fourier modes for this to be a viable approach for imaging complex scenes. Nevertheless if we appropriately propagate wavefront error from a calibrator or from routine wavefront sensing operations, we may be able to avoid fully nulling these modes while retaining benefits of self-calibration.

Our ongoing recommendations for AMI proposals agree with those set out in the companion paper: maximise the number of groups subject to keeping the maximum well depth below ~ 30 k counts. Observe calibrators of a similar brightness to the science target ensuring integration to a similar well depth. Similar recommendations were made by Sallum *et al.* (2024). We disagree with the practice in many AMI observations to date of ensuring the science target and calibrator are placed on the same subpixel position and kept as still as possible; to the contrary we strongly

recommend using as many subpixel dithers as is practical, to better separate detector effects from astrophysical information.

There is a wealth of archival observations that will benefit from *amigo*, *dorito*, or similar approaches. There are many cases where it is important to observe faint point sources at separations close to the diffraction limit, but this work shows that we can also recover complex scenes effectively. In GTO 1242, we hope to be able to refine the fits to PDS 70 published in Blakely et al. (2025), and extend this to the other transition disks HD 135344 B (Cazzoletti et al. 2018) and HD 100546 (Quanz et al. 2015) observed in the same program. Beyond archival data, there are a wealth of protoplanetary disks to observe with similarly impactful science (e.g. in the DSHARP survey: Andrews et al. 2018). The example of NGC 1068 shows that it will be valuable to observe AGN inner environments, and quasar host galaxies such as 3C 273 (following Ren et al. 2024). While it is possible to monitor the volcanism on Io with Keck adaptive optics imaging, AMI may prove to be a valuable complement.

RML imaging has long been a standard approach for image reconstruction in astronomy, and it has been shown to perform well in the present context. Nevertheless, it is far from the most modern or sophisticated approach to image deconvolution, and while a pixel basis performs well here, there is no reason to think this is optimal. Future work will include investigation of different bases, including wavelet bases which are currently being implemented in *dorito*. Machine learning also offers many paths to improvement of *dorito*. The literature here is wide-ranging, but we expect there will be great value in approaches based on score-based priors and likelihoods (e.g. Adam et al. 2025; Dia et al. 2025), deep image priors (e.g. Feng, Bouman, & Freeman 2024), information field theory (Enßlin 2019), autoencoders (notably the AMI Io work by Sanchez-Bermudez et al. 2025). This is likely to be especially impactful in resolving sources very close to the inner working angle of AMI and in generalising this approach to differentiable rendering of coronagraphic data (as in Feng et al. 2025), where the problem is much more nonlinear and there is much greater distortion of the scene close to the inner working angle.

Supplementary material. The supplementary material for this article can be found at <https://doi.org/10.1017/pasa.2026.10179>

Acknowledgements. The authors thank Imke de Pater and Joel Sánchez Bermúdez for their valuable assistance in identifying Ionian surface features and John Stansberry for his work on Io observations; Jacob Isbell for sharing the latest LBTI reconstructions of NGC 1068; and Ryan White for the consultations about colliding wind-binary morphologies.

BP, PT, and SR have been supported by the Australian Research Council grant DP230101439 and BP by DE210101639; and LD and MC have been supported by the Australian Government Research Training Program (RTP) award. We are grateful to the Australian public for enabling this science. BP and SR would like to thank the Big Questions Institute for their philanthropic support. Development of *∂Lux* has been supported by the Breakthrough Foundation through their Toliman project as a part of the Breakthrough Watch initiative. DJ is supported by NRC Canada, and DB and DJ acknowledge the support of the Natural Sciences and Engineering Research Council of Canada (NSERC). BM & KESF are supported by NSF AST-1831415, NSF AST-2206096 and Simons Foundation Grant 533845.

We acknowledge and pay respect to the traditional owners of the land on which the University of Sydney, University of Queensland and Macquarie University are situated, upon whose unceded, sovereign, ancestral lands we work. We pay respects to their Ancestors and descendants, who continue cultural and spiritual connections to Country. We acknowledge and respect the Ləkʷəŋən (Songhees and Esquimalt) Peoples on whose territory the University

of Victoria stands, and the Ləkʷəŋən and WSÁNEĆ Peoples whose historical relationships with the land continue to this day.

This work is based on observations made with the NASA/ESA/CSA *James Webb Space Telescope*. The NIRISS instrument was funded by the Canadian Space Agency and built in Canada by Honeywell. The *amigo* base model is constructed from observations associated with JWST program CAL 4481 (PI: A. Sivaramakrishnan). The data were obtained from MAST at the Space Telescope Science Institute (STScI), which is operated by the Association of Universities for Research in Astronomy, Inc., under NASA contract NAS 5-03127 for JWST.

This work was performed on the OzSTAR national facility at Swinburne University of Technology. The OzSTAR program receives funding in part from the Astronomy National Collaborative Research Infrastructure Strategy (NCRIS) allocation provided by the Australian Government, and from the Victorian Higher Education State Investment Fund (VHESIF) provided by the Victorian Government.

This research made use of *∂Lux* (Desdoigts et al. 2023, 2024); *Jax* (Bradbury et al. 2018); *equinox* (Kidger & Garcia 2021); *optax* (Hessel et al. 2020); *planetmapper* (King & Fletcher 2023); and *matplotlib* (Hunter 2007).

Code and data availability. All codes and data used to produce this work are publicly available and open source. The *amigo* model and pipeline are hosted on GitHub. The *dorito* model is hosted on GitHub, along with a series of example notebooks used in this paper's results. All data are available from the Barbara A. Mikulski Archive for Space Telescopes (MAST) under the appropriate proposal numbers; in this pipeline paper, we have relied on publicly available data with no proprietary data included, with the exception of the latter observing epoch of NGC 1068 at the time of this submission. We encourage researchers to adapt and apply *amigo* and *dorito* to their own interferometric datasets, and to contact the development team with questions and contributions.

References

- Abadi, M., et al. 2016, arXiv e-prints, [arXiv:1603.04467](https://arxiv.org/abs/1603.04467)
- Abril-Pla, O., et al. 2023, *PeerJ CS*, **9**, e1516, accessed: 2025-04-10
- Adam, A., et al. 2025, *AJ*, **169**, 254
- Adriani, A., et al. 2017, *SSRv*, **213**, 393
- Andrews, S. M., et al. 2018, *ApJ*, **869**, L41
- Antilogos, P., Astier, P., Doherty, P., Guyonnet, A., & Regnault, N. 2014, *Jl*, **9**, C03048
- Antonucci, R. R. J., & Miller, J. S. 1985, *ApJ*, **297**, 621
- Argyriou, I., et al. 2023, *A&A*, **680**, A96
- Azulay, R., et al. 2017, *A&A*, **607**, A10
- Baldwin, J. E., Haniff, C. A., Mackay, C. D., & Warner, P. J. 1986, *Natur*, **320**, 595
- Baron, F. 2016, in *Astronomy at High Angular Resolution: A Compendium of Techniques in the Visible and Near-Infrared*, ed. H. M. J. Boffin, G. Hussain, J.-P. Berger, & L. Schmidtbreick, *Astrophysics and Space Science Library* (Springer Cham), 75
- Baron, F., Monnier, J. D., & Kloppenborg, B. 2010, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Vol. 7734, *Optical and Infrared Interferometry II*, ed. W. C. Danchi, F. Delplancke, & J. K. Rajagopal, 77342I
- Blakely, D., et al. 2025, *AJ*, **169**, 137
- Bradbury, J., et al. 2018, *JAX: Composable transformations of Python+NumPy programs*. <https://github.com/google/jax>
- Buscher, D. F. 1994, in *IAU Symposium*, Vol. 158, *Very High Angular Resolution Imaging*, ed. J. G. Robertson, & W. J. Tango, 91
- Cabezas, A., Corenflos, A., Lao, J., & Louf, R. 2024, *BlackJAX: Composable Bayesian inference in JAX*, [arXiv:2402.10797](https://arxiv.org/abs/2402.10797)
- Candes, E., Romberg, J., & Tao, T. 2006, *IEEE TIT*, **52**, 489
- Candes, E. J., & Wakin, M. B. 2008, *IEEE SPM*, **25**, 21
- Cazzoletti, P., et al. 2018, *A&A*, **619**, A161
- Clark, B. G. 1980, *A&A*, **89**, 377
- Cornwell, T. J. 2008, *IEEE JSTSP*, **2**, 793

- Czekala, I., et al. 2025, arXiv e-prints, [arXiv:2502.00100](#)
- Davies, A. G. 2007, *Volcanism on Io* (Cambridge University Press)
- Davies, A. G., Perry, J., Williams, D. A., & Nelson, D. M. 2024a, in AGU Fall Meeting Abstracts, Vol. 2024, P33D–2894
- Davies, A. G., Perry, J. E., Williams, D. A., Veeder, G. J., & Nelson, D. M. 2024b, *PASJ*, **5**, 121
- De Furio, M., et al. 2025, arXiv e-prints, [arXiv:2510.05308](#)
- de Kleer, K., et al. 2019, *AJ*, **158**, 29
- DeepMind, et al. 2020, The DeepMind JAX Ecosystem
- Desdoigts, L., Pope, B., Gully-Santiago, M., & Tuthill, P. 2024, arXiv e-prints, [arXiv:2406.08704](#)
- Desdoigts, L., Pope, B. J. S., Dennis, J., & Tuthill, P. G. 2023, *JATIS*, **9**, 028007
- Desdoigts, L., et al. 2025, PASA, accepted
- Dia, N., et al. 2025, arXiv e-prints, [arXiv:2501.02473](#)
- Doyon, R., et al. 2023, *PASP*, **135**, 098001
- Enßlin, T. A. 2019, *AnPh*, **531**, 1800127
- Event Horizon Telescope Collaboration, et al. 2019a, *ApJ*, **875**, L1
- Event Horizon Telescope Collaboration, et al. 2019b, *ApJ*, **875**, L4
- Feng, B. T., Bouman, K. L., & Freeman, W. T. 2024, *ApJ*, **975**, 201
- Feng, B. Y., et al. 2025, arXiv e-prints, [arXiv:2501.01912](#)
- Fizeau, H. 1868, *CRAS*, **66**, 932
- Ford, K. E. S., et al. 2014, *ApJ*, **783**, 73
- Gómez Rosas, V., et al. 2022, *Natur*, **602**, 403
- Greenbaum, A. Z., Pueyo, L., Sivaramakrishnan, A., & Lacour, S. 2015, *ApJ*, **798**, 68
- Hansen, P. C. 1992, *SIAMRv*, **34**, 561
- Hessel, M., et al. 2020, Optax: composable gradient transformation and optimisation, in JAX!
- Hinkley, S., et al. 2023, *A&A*, **671**, L5
- Hinz, P. M., et al. 2016, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 9907, Optical and Infrared Interferometry and Imaging V, ed. F. Malbet, M. J. Creech-Eakman, & P. G. Tuthill, 990704
- Hirata, C. M., & Choi, A. 2020, *PASP*, **132**, 014501
- Hoerl, A. E., & Kennard, R. W. 1970a, *Technometrics*, **12**, 69
- Hoerl, A. E., & Kennard, R. W. 1970b, *Technometrics*, **12**, 55
- Hofmann, K. H., Weigelt, G., & Schertl, D. 2014, *A&A*, **565**, A48
- Högbom, J. A. 1974, *A&AS*, **15**, 417
- Hunter, J. D. 2007, *CSE*, **9**, 90
- IceCube Collaboration, et al. 2022, *Sci*, **378**, 538
- Inoue, Y., Khangulyan, D., & Doi, A. 2020, *ApJ*, **891**, L33
- Ireland, M. J. 2013, *MNRAS*, **433**, 1718
- Isbell, J. W., et al. 2025, *NatAs*, **9**, 417
- Jaffe, W., et al. 2004, *Natur*, **429**, 47
- Jennison, R. C. 1958, *MNRAS*, **118**, 276
- Kass, R., Tierney, L., & Kadane, J. 1991, Laplace's Method in Bayesian Analysis, doi:[10.1090/conm/115/07](#)
- Kidger, P., & Garcia, C. 2021, Differentiable Programming Workshop at Neural Information Processing Systems, arXiv e-prints, [arXiv:2111.00254](#)
- King, O. R. T., & Fletcher, L. N. 2023, *JOSS*, **8**, 5728
- Lajoie, C.-P., M., et al. 2023, OTE Science Performance Memo 2 – A Year of Wavefront Sensing with JWST in Flight: Cycle 1 Telescope Monitoring & Maintenance Summary, Technical Report JWST-STScI-008497, STScI
- Lau, R. M., et al. 2024, *ApJ*, **963**, 127
- Macchetto, F., Capetti, A., Sparks, W. B., Axon, D. J., & Boksenberg, A. 1994, *ApJ*, **435**, L15
- Marchis, F., et al. 2002, *Icar*, **160**, 124
- Martinache, F. 2010, *ApJ*, **724**, 464
- Martinache, F., et al. 2020, *A&A*, **636**, A72
- Meimon, S. C., Mugnier, L. M., & Le Besnerais, G. 2005, *OL*, **30**, 1809
- Mor, R., & Netzer, H. 2012, *MNRAS*, **420**, 526
- Mura, A., et al. 2020, *Icar*, **341**, 113607
- Padovani, P., et al. 2017, *A&ARv*, **25**, 2
- Pairat, B., Cantalloube, F., & Jacques, L. 2021, *MNRAS*, **503**, 3724
- Paszke, A., et al. 2019, PyTorch: An Imperative Style, High-Performance Deep Learning Library, [arXiv:1912.01703](#)
- Phan, D., Pradhan, N., & Jankowiak, M. 2019, arXiv preprint [arXiv:1912.11554](#)
- Pope, B., Martinache, F., & Tuthill, P. 2013, *ApJ*, **767**, 110
- Pope, B. J. S. 2016, *MNRAS*, **463**, 3573–3581
- Pope, B. J. S., Pueyo, L., Xin, Y., & Tuthill, P. G. 2021, *ApJ*, **907**, 40
- Quanz, S. P., et al. 2015, *ApJ*, **807**, 64
- Ray, S., et al. 2025, *ApJ*, **983**, L25
- Ren, B. B., et al. 2024, *A&A*, **683**, L5
- Renard, S., Thiébaud, E., & Malbet, F. 2011, *A&A*, **533**, A64
- Sallum, S., & Eisner, J. 2017, *ApJS*, **233**, 9
- Sallum, S., Ray, S., & Hinkley, S. 2022, in Optical and Infrared Interferometry and Imaging VIII, ed. A. Mérand, S. Sallum, & J. Sanchez-Bermudez, Vol. 12183 (International Society for Optics and Photonics (SPIE)), 121832M
- Sallum, S., et al. 2024, *ApJ*, **963**, L2
- Sanchez-Bermudez, J., Alberdi, A., Schödel, R., & Sivaramakrishnan, A. 2022, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 12183, Optical and Infrared Interferometry and Imaging VIII, ed. A. Mérand, S. Sallum, & J. Sanchez-Bermudez, 121831K
- Sanchez-Bermudez, J., et al. 2025, *MNRAS*, **543**, 608
- Schlawin, E., et al. 2023, *PASP*, **135**, 018001
- Schutz, A., et al. 2014, *JOSA*, **31**, 2334
- Schwarz, U. J. 1978, *A&A*, **65**, 345
- Shannon, C. E. 1948, *BSTJ*, **27**, 379
- Sivaramakrishnan, A., et al. 2010, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 7731, Space Telescopes and Instrumentation 2010: Optical, Infrared, and Millimeter Wave, ed. J. M. Oschmann, Jr., M. C. Clampin, & H. A. MacEwen, 77313W
- Sivaramakrishnan, A., et al. 2012, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 8442, Space Telescopes and Instrumentation 2012: Optical, Infrared, and Millimeter Wave, ed. M. C. Clampin, G. G. Fazio, H. A. MacEwen, & J. M. Oschmann, Jr., 84422S
- Sivaramakrishnan, A., et al. 2023, *PASP*, **135**, 015003
- Skilling, J., & Bryan, R. K. 1984, *MNRAS*, **211**, 111
- Soulain, A., et al. 2020, in SPIE Conference Series, Vol. 11446, Optical and Infrared Interferometry and Imaging VII, ed. P. G. Tuthill, A. Mérand, & S. Sallum (International Society for Optics and Photonics (SPIE)), 1144611
- Stan Development Team. 2023, Stan Reference Manual, <https://mc-stan.org>, version X.Y.Z
- Starck, J.-L., Bijaoui, A., Lopez, B., & Perrier, C. 1994, *A&A*, **283**, 349
- Strong, D., & Chan, T. 2003, *IP*, **19**, S165
- Tagawa, H., Kimura, S. S., & Haiman, Z. 2023, *ApJ*, **955**, 23
- Thiébaud, E. 2008, in Optical and Infrared Interferometry, Vol. 7013 (SPIE), 479
- Thiébaud, É., & Young, J. 2017, *JOSA*, **34**, 904
- Tibshirani, R. 1996, *JRSS*, **58**, 267
- Tuthill, P., Sivaramakrishnan, A., & Lloyd, J. 2022, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 12183, Optical and Infrared Interferometry and Imaging VIII, ed. A. Mérand, S. Sallum, & J. Sanchez-Bermudez, 121830Q
- Tuthill, P. G. 2012, in Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, Vol. 8445, Optical and Infrared Interferometry III, ed. F. Delplancke, J. K. Rajagopal, & F. Malbet, 844502
- Tuthill, P. G., Monnier, J. D., & Danchi, W. C. 1999, *Natur*, **398**, 487
- Twiss, R. Q., Carter, A. W. L., & Little, A. G. 1960, *Obs*, **80**, 153
- Williams, P. M., et al. 1985, *MNRAS*, **215**, 23P
- Williams, P. M., et al. 2001, *MNRAS*, **324**, 156
- Willis, M., Jeffs, B., & Long, D. 1999, in Conference Record of the Thirty-Third Asilomar Conference on Signals, Systems, and Computers (Cat. No.CH37020), Vol. 2, 1272
- Wong, K. T., Kamiński, T., Menten, K. M., & Wyrowski, F. 2016, *A&A*, **590**, A127
- Zawadzki, B., et al. 2023, *PASP*, **135**, 064503

Appendix A Other L-curves

L-curve plots for sources NGC 1068 and WR 137 are shown in [Figure A1](#), with the corresponding curve for Io shown in [Figure 2](#). Unregularised images of NGC 1068 largely resemble those at the elbow of the L-curve. Compared to Io, there are significantly less




Figure A1. L-curve diagram used to select the optimal regularisation parameters for NGC 1068 (top) and WR 137 (bottom). Each point on the curve is the balance between regulariser term $\mathcal{R}$ and likelihood term $\mathcal{L}$ of a converged image reconstruction for a different regularisation hyperparameter value. NGC 1068 is regularised with total variation, and WR 137 is regularised with maximum entropy. Shown are several reconstructed images corresponding to different points along the curve.

non-physical ringing artefacts needing to be suppressed with regularisation. This may imply the NGC 1068 observations were better calibrated than the Io observations, which is unsurprising given the higher data quality of the NGC 1068 calibrator exposures. Increasing TV regularisation resulted in smoother images as

expected but did not largely change the morphology of the AGN. Contrastingly, regularisation appears to be necessary for WR 137. With no regularisation WR 137 collapses into a few loosely connected points of flux which would be impossible to resolve, as a result of overfitting.