PROCEEDINGS OF SPIE

SPIDigitalLibrary.org/conference-proceedings-of-spie

Neural network based image reconstruction with astrophysical priors

Claes, R., Kluska, J., Van Winckel, H., Min, M.

R. Claes, J. Kluska, H. Van Winckel, M. Min, "Neural network based image reconstruction with astrophysical priors," Proc. SPIE 11446, Optical and Infrared Interferometry and Imaging VII, 114461U (13 December 2020); doi: 10.1117/12.2576321

SPIE.

Event: SPIE Astronomical Telescopes + Instrumentation, 2020, Online Only

Neural network based image reconstruction with astrophysical priors

R. Claes^a, J. Kluska^a, H. Van Winckel^a, and M. Min^b

^aInstituut voor Sterrenkunde (IvS), KU Leuven, Celestijnenlaan 200D, 3001, Leuven (Belgium)

^bSRON Netherlands Institute for Space Research (Netherlands)

ABSTRACT

With the advent of interferometric instruments with 4 telescopes at the VLTI and 6 telescopes at CHARA, the scientific possibility arose to routinely obtain milli-arcsecond scale images of the observed targets. Such an image reconstruction process is typically performed in a Bayesian framework where the function to minimize is made of two terms: the data likelihood and the Bayesian prior. This prior should be based on our prior knowledge of the observed source. Up to now, this prior was chosen from a set of generic and arbitrary functions, such as total variation for example. Here, we present an image reconstruction framework using generative adversarial networks where the Bayesian prior is defined using state-of-the-art radiative transfer models of the targeted objects. We validate this new image reconstruction algorithm on synthetic data with added noise. The generated images display a drastic reduction of artefacts and allow a more straightforward astrophysical interpretation. The results can be seen as a first illustration of how neural networks can provide significant improvements to the image reconstruction post processing of a variety of astrophysical sources.

Keywords: infrared interferometry, image reconstruction, deep learning, neural networks, generative adversarial networks

1. INTRODUCTION

Obtaining an angular resolution in the milli-arcsecond range is vital to study the morphology of a variety of astrophysical targets such as evolved and young stars, binaries, active galactic nuclei or solar system bodies. Current and planned single telescope facilities do not reach such an angular resolution, but thanks to the development of optical (visible and infrared) interferometry, these resolutions can now be routinely obtained.

Using this technique one does not directly observe images of the target but instead interferometric fringes that contain information about the Fourier transform of the image at given spatial frequencies u and v which depends on the projected telescope baselines on the sky. By observing with many different baselines, one can cover the measurements in the well-covered (u, v) -plane, to try to reconstruct the image of the target on the sky. With the advent of new generation of interferometric instruments which can recombine four or more telescopes, such as GRAVITY,¹ the Precision Integrated-Optics Near-infrared Imaging ExpeRiment² and the Multi AperTure mid-Infrared SpectroScopic Experiment³ at the Very Large Telescope Interferometer (VLTI) or the Michigan Infrared Beam Combiner-X⁴ at the Center for High Angular Resolution Astrophysics (CHARA), we have entered the imaging era in optical interferometry.⁵⁻⁸

Data from these instruments have revealed that the observed targets display often complex morphologies that are difficult to retrieve using geometric model fitting, as many parameters (> 10) are needed to reproduce the data. In this process, there is a high risk of including a model bias.⁹⁻¹¹ The imaging technique is, therefore, unique to unveil complex and unexpected morphology at milliarcsecond scale, provided the reconstruction does not include artifacts which hamper the astrophysical interpretation. Image reconstruction algorithms are therefore crucial to recover the best quality images.

Reconstructing an image from optical interferometric data is typically performed in a Bayesian framework where the image ($\mathbf{x}$) which maximizes the *a posteriori* probability given the observations ($\mathbf{y}$) is sought.¹²

This probability ($P(\mathbf{x} | \mathbf{y})$) can be expressed, using the Bayes rule, by:

$$P(\mathbf{x} | \mathbf{y}) = \frac{P(\mathbf{y} | \mathbf{x})P(\mathbf{x})}{P(\mathbf{y})}, \quad (1)$$

Further author information: (Send correspondence to Jacques Kluska)

Jacques Kluska: E-mail: jacques.kluska@kuleuven.be, Telephone: +32 16 32 84 51

where $P(\mathbf{y} | \mathbf{x})$ is the likelihood of the data, $P(\mathbf{x})$ is the *a priori* distribution of $\mathbf{x}$ and $P(\mathbf{y})$ is called the evidence. The evidence is not taken into account in the image optimization procedure, as it does not depend on the image $\mathbf{x}$.

In practice, instead of maximising the probability one searches the *maximum a posteriori* solution (i.e., the image with the *maximum a posteriori* probability; $\mathbf{x}_{MAP}$) that minimizes the negative logarithm of Eq. 1 that is called the cost function (f):

$$\mathbf{x}_{MAP} = \underset{\mathbf{x}}{\operatorname{argmin}} f = \underset{\mathbf{x}}{\operatorname{argmin}} (-\log(P(\mathbf{y} | \mathbf{x})) - \log(P(\mathbf{x}))) \quad (2)$$

$$= \underset{\mathbf{x}}{\operatorname{argmin}} f_{\text{data}} + \mu f_{\text{regl}}, \quad (3)$$

where f_{data} is the data likelihood cost function (e.g. χ^2), f_{regl} is the regularization and μ is the regularization weight that sets the strength of the regularization.

The regularization influences the image reconstruction by promoting images with a certain brightness distribution, independent of the data likelihood. By doing so it determines how the Fourier space is extrapolated between (u, v) -plane measurements.

Most commonly generic functions coming from the signal processing community are used such as maximum entropy, quadratic smoothness, total variation or Laplacian regularizations.^{12–16} Based on this Bayesian framework, the regularization should incorporate our prior expectations or astrophysical knowledge of the brightness distribution. This is not the case for these common generic regularizations, as such distributions are too complex to formalize in a simple equation.

We therefore present here a novel image reconstruction framework based on convolutional neural networks (CNNs).¹⁷ We employed neural networks trained in a deep convolutional generative adversarial network (GAN)¹⁸ to reconstruct images. The method is called ORGANIC: Object Reconstruction using Generative Adversarial Networks for Interferometric data. Among other properties, it allowed us to use CNNs as a regularisation, making it learn the prior distributions from images generated by models of astrophysical sources. The method can be applied to a variety of astrophysical sources if models are available. In this paper, we focus on circumstellar disks.

In section 2 the image reconstruction framework and neural network architecture is presented. The results obtained on artificial datasets are presented in section 3. The conclusions and future prospects are discussed in section 4.

2. THE ORGANIC IMAGE RECONSTRUCTION FRAMEWORK

2.1 Generative adversarial networks

A GAN is a combination of two CNNs called the generator and the discriminator (see figure 1) which compete against each other¹⁹. During the training process the generator learns to generate images that look like the model images, while the discriminator is supplied by both model images and generator images and learns to discriminate between the two by generating a number between 0 (“fake” image) and 1 (“true” image). In other words the discriminator learns to classify images produced by the generator as “fake” and those from the training dataset as “real”. Meanwhile, the generator tries to make the discriminator labeling the generator image as “real”. Therefore, during the training of a GAN, the discriminator network D and generator network G , try to achieve goals in competition of each other with the object function being defined by:

$$\underset{\theta_G \in \Theta_G}{\operatorname{argmin}} \underset{\theta_D \in \Theta_D}{\operatorname{argmax}} V(D, G), \quad (4)$$

with:

$$V(D, G) = \mathbb{E}_{\mathbf{x} \sim P_{\text{model}}} [\log(D(\mathbf{x}; \theta_D))] + \mathbb{E}_{\mathbf{z} \sim P_z} [\log(1 - D(G(\mathbf{z}; \theta_G); \theta_D))] \quad (5)$$

and θ_D and θ_G represent the trainable parameters describing the discriminator D and generator G respectively. $\mathbf{x}$ are sample images drawn from a distribution P_{model} and $\mathbf{z}$ is a noise vector sampled from distribution P_z . The global optimal value of $V(D, G)$ is achieved when the distribution of the data created by the generator $P_G(\mathbf{x})$ is equal to $P_{\text{model}}(\mathbf{x})$.¹⁹

Figure 1: Schematic representation of the used GAN

Table 1: The architecture of the Discriminator network.

layer	activation	Number of kernels	kernel size	stride	output shape	trainable parameters
2D convolution	Leaky relu $\alpha = 0.25$	32	4×4	(2,2)	(64,64, 32)	320
2D convolution	Leaky relu $\alpha = 0.25$	64	4×4	(2,2)	(32,32, 64)	18496
2D convolution	Leaky relu $\alpha = 0.25$	128	4×4	(2,2)	(16,16, 128)	73856
Fully connected with bias	Sigmoid	/	/	/	1	32769

2.2 Training the GAN

Building a GAN architecture that solves Eq. 4 requires to define both G and D , an optimization strategy and an adequate training dataset. There is, however, no clear consensus on how to best define the network architectures and training routine that will converge. Therefore, much of the information presented here has been arbitrarily chosen from common practices in the field of deep learning when possible.

2.2.1 GAN Architecture

Both of our neural networks are CNNs, as they consist of both fully connected and convolutional layers. CNN have beneficial properties when dealing with image data²⁰ as they are made with peculiar layers called the convolution layers. They allow the extraction of meaningful information from the image while limiting the amount of parameters to train. The architectures of both the discriminator and generator networks are listed in tables 1 and 2 respectively. More information about these layers can be found in appendix A

The discriminator D The discriminator takes an image 128×128 pixels as an input and gives a single value between 0 and 1 as an output. In the discriminator the leaky ReLU activation function²¹ with a leak strength of $\alpha = 0.25$ was used for all the convolutional layers. In the final, fully connected layer a sigmoid activation function is used, as this restricts the output range to be between 0 to 1. This is done as $D(x)$ represents the probability that x came from the data rather than $P_G(x)$. To avoid over-confidence, during both the training of the GAN and the image reconstruction, we used a dropout of 0.3, i.e., each of the convolutional kernels has a 30% chance of not contributing towards the output of the discriminator

Table 2: The architecture of the Generator network.

layer	activation	Number of kernels	kernel size	stride	output shape	trainable parameters
Fully connected without bias	Leaky relu $\alpha = 0.1$	/	/	/	65536	6553600
reshape	/	/	/	/	(16, 16, 256)	
2D transpose convolution	Leaky relu $\alpha = 0.1$	128	4×4	(2,2)	(32, 32, 128)	524288
Batch normalization	/	/	/	/	(32, 32, 128)	256
2D transpose convolution	Leaky relu $\alpha = 0.1$	64	4×4	(2,2)	(64, 64, 64)	131072
Batch normalization	/	/	/	/	(64, 64, 64)	128
2D transpose convolution	Leaky relu $\alpha = 0.1$	32	4×4	(2,2)	(128, 128, 32)	32768
Batch normalization	/	/	/	/	(128, 128, 32)	64
2D convolution	tanh	1	5×5	(1,1)	(128, 128, 1)	800

and the back propagation of the errors in the discriminator and generator networks. The outputs of the kernels which are not dropped are scaled up by $1/(1 - 0.3)$ such that the sum over all outputs is unchanged. The discriminator is made of 125441 trainable parameters (see Table 1).

the Generator G The goal of the generator is to produce a 128×128 image starting with a vector of 100 elements called the noise vector. To do so, the generator uses 2D-transposed convolution layers. This type of layers perform a similar operation as the traditional convolutional layer, but with the redefinition of the concepts of stride and padding in order to upscale rather than downscale an image.²² We used the leaky ReLU activation function with a leak strength of $\alpha = 0.1$ in all the transposed convolutional layers. In the final convolutional layer of the generator the tanh activation was used.¹⁸ We also included batch normalization layers^{18,23} for stable GAN training. Finally, the generator is made of 7242848 trainable parameters (see Table 2).

2.2.2 The physical models

To train the GAN we need to have model images. In this paper we focus on circumstellar disks of dust and gas as observed in the near-infrared. We produced model images of circumstellar disks using a radiative transfer code MCMMax.²⁴ MCMMax was successfully used to model several disks around post-asymptotic giant branch (post-AGB) binaries.^{25–27} It consists of an axisymmetric disk model where the disk scale-height is set by hydrostatic equilibrium. Our grid of models is based on the best fit model of the disk around IRAS08544-4431.²⁷ We therefore set the central star to have the luminosity and radius of the primary star of IRAS08544-4431.²⁷ The mass of the central star is $M = 2.4m_{\odot}$. We only varied the disk parameters that influence the image in the near-infrared: the inner disk radius R_{in} , the power law index describing the surface density as a function of radius for the inner region of the disk ρ_{in} and another such power law index describing the outer region of the disk ρ_{out} (see Table 3). This gives 455 different models.

Once each model were generated the images were produced at random inclinations (i) chosen in a flat distribution of the cosine of the inclinations. For each model we produced 12 images without the central star half of which in the continuum at $1.5\mu\text{m}$ and the other half at $2.1\mu\text{m}$. With this a total of 5460 images were generated from MCMMax. These images are sized to 128×128 pixels to meet the input size of the discriminator. Each time one of these images is sampled to be presented to the discriminator, it is given a random position angle rotation. The distribution of position angles is chosen to be flat. Upon sampling an image has a 1 in 4 chance to have a uniform ‘unresolved’ background added. The random value that is added to each pixel value when this occurs follows a uniform distribution between 0 and 0.1 relative to the maximum flux. This was done as certain disk sources are known to have an over-resolved flux which can not be accounted for by radiative transfer models (e.g.^{8–10,27}).

Before being fed to the GAN, these images are individually normalized to have pixel values in a range spanning between -1 and 1. The value -1 corresponds to a relative flux of 0 while 1 represents the maximum flux in the image. A random zoom on the images from a flat distribution ranging between -10 and 10% was also applied. Both this zoom and the random rotation of the position angle upon sampling are achieved using bi-linear interpolation.

The whole image reconstruction process is divided into two phases A and B. Phase A is the pre-training of the networks with astrophysical models. Phase B is the actual image reconstruction phase.

Table 3: The parameter range used to construct a grid of MCMa²⁴ disk models.

Parameter	Grid Range
R_{in}	[6.0; 6.5; 7.0; 7.5; 8.0; 8.5; 9.0]
ρ_{in}	[0.6; 0.3; 0.0; -0.3; -0.6; -0.9;
ρ_{out}	-1.2; -1.5; -1.8; -2.1; -2.4; -2.7; -3.0]
i	[0.5; 1.0 1.5; 2.0; 3.0.]
PA	flat distribution on $\cos i$, 12 inclinations per model, i between 0° and 70°
Zoom	flat distribution, randomly rotated each mini-batch
	flat distribution between -10 and 10, randomly chosen each mini-batch %

2.3 Phase A: pre-training

To be able to perform image reconstructions the GAN needs to go through phase A to be pre-trained on models. This GAN pre-training is performed mainly following,¹⁹ including their proposed alteration to the generators gradient. The gradients used for the discriminator were also adjusted.²⁸

We quickly recall here the main steps. First the alterations to the gradients of both the generator and discriminator are discussed.

Solving Eq. 5 may provide insufficient gradient for the generator to train well.¹⁹ To deal with this problem we maximise $\log(D(G(z)))$ for the generator instead of minimizing $\log(1 - D(G(z)))$.¹⁹ Doing so results in the same optimum as well as providing a stronger gradient during early training. In deep learning the convention is to implement such optimization problems as minimization problems.

So with the alteration to the gradient of the discriminator, our optimization problem is reformulated as trying to simultaneously optimize both:

$$\operatorname{argmin}_{\theta_D \in \Theta_D} -V(D, G) \quad (6)$$

and

$$\operatorname{argmin}_{\theta_G \in \Theta_G} -\mathbb{E}_{z \sim P_z} [\log(D(G(z)))]. \quad (7)$$

The objectives given in equation 6 can be interpreted through the binary cross entropy cost function. This cost function is given by

$$l(x) = -\frac{1}{m} \sum_{i=1}^m [y_i \log(D(x_i; \theta_D)) + (1 - y_i) \log(1 - D(x_i; \theta_D))],$$

and is commonly used for binary classification problems. The discriminator can thus be seen to be trained to achieve a target response of $y_i = 1$ for the "real" images and one of $y_i = 0$ for the "fake" images. Meanwhile, the generator is trained to achieve a target response from the discriminator of $y_i = 1$. The generator is thus trained in order to try and convince the discriminator that the images which it creates are a part of the training dataset.

Based on this formulation of the objectives,²⁸ proposed an alteration called "one sided label smoothing". When using this technique the "fake" images are labelled as 0 whereas the "true" ones are labelled as 0.9 instead of 1 when training the discriminator. The target for training the generator remains equal to be 1. Applying one sided label smoothing prevents overconfidence in the discriminator, thereby, making the GAN more likely to converge.²⁹ The procedure to optimize both 6 and 7, consists of iteratively updating the component networks and is outlined in appendix B algorithm 1 and is describe by hyper-parameters like the number of training iterations, the size of the training dataset and the number of training epochs.

During a training iteration, two mini-batches each consisting of 50 images is sampled. The first mini-batch consists of a random selection of images from the training dataset. The second mini-batch consists of "fake" images. A "fake" image is generated by giving a randomly generated input vector of 100-elements to the generator. Each element of the input vector is sampled from a Gaussian distribution centered on 0 with a standard deviation of 1. Both these mini-batches are then used to compute a learning gradient for the discriminator. After the discriminator is updated, a second mini-batch of 50 noise vectors is sampled. These noise vectors are then used to compute a learning gradient for the generator and update it.

Figure 2: A random selection of disk-images created by the the generator network after GAN-training

An epoch is a training round where almost the entire set of available images has been presented to the GAN, that is $\lfloor N_{model}/m \rfloor = \lfloor 5460/50 \rfloor = 109$ iterations. Here we choose to preform the training for 250 epochs.

We trained the GAN on a Nvidia Tesla P100 GPU, provided by the Flemish SuperComputer. GPU training of the GAN takes approximately two hours of physical time. A selection of images produced by the generator network after training can be seen on figure 2.

2.4 Phase B: image reconstruction

Once the GAN is pre-trained with MCMAX models, it can be used to reconstruct images that reproduce interferometric data. In that phase, the generator of the GAN is fine-tuned in order to reconstruct the image taking into account the data likelihood. To do so, the loss function is adapted to match Eq. 3.

2.4.1 Data likelihood

The data likelihood function is first obtained by calculating the fast Fourier transform of the reconstructed image. The complex visibilities at the relevant spatial frequencies are then obtained by applying a bi-linear interpolation to the Fourier transformed image. The squared visibilities and closure phases are computed from these complex visibilities and then compared to the measurements. The χ^2 for squared visibilities (S_{j_1, j_2}) between telescope j_1 and j_2 is:

$$\chi_{v^2}^2 = \sum_{j_1 < j_2} \frac{(S_{j_1, j_2}^{data} - S_{j_1, j_2}^{Image})^2}{Var(S_{m, j_1, j_2}^{data})} \quad (8)$$

For circumstellar disks that were observed in the near-infrared the closure phase signal is not showing any phase jumps because of the contribution of the star that is unresolved.⁸ It is, therefore, possible to neglect potential phase jumps. In this case; we can directly use this expression for the closure phases:

$$\chi_{c.p.}^2 = \sum_{j_1 < j_2 < j_3} \frac{(\beta_{j_1, j_2, j_3}^{data} - \beta_{j_1, j_2, j_3}^{Image})^2}{Var(\beta_{j_1, j_2, j_3}^{data})} \quad (9)$$

The value then used as the data likelihood is then $f_{data} = (\chi_{v^2}^2 + \chi_{c.p.}^2) / N_{observables}$ where $N_{observables}$ is the sum of the number of V^2 and CPs.

Figure 3: Example of checkerboard artifacts appearing when optimizing 10. This example occurred after 500 epochs, with a value of $\mu = 2$.

We implemented the SPARCO approach³⁰ that consists in adding the contributions of a central star or stars to the reconstructed image. The functionality to add both a uniform disk source and a point source is implemented in our code.³¹ This allows for the contributions of the central stars to be removed from the image and provide an image of the circumstellar environment with enhanced quality. It also allows to use the chromatic data from all the channels at once and reconstruct a single image that is valid at all sampled wavelengths.

2.4.2 Regularisation

The regularisation consists of the negative logarithm of the discriminators output for the reconstructed image. The objective function for which the generator is optimized is expressed as follows:

$$f = f_{data}(G(\mathbf{z}, \Theta_G)) - \mu \log(D(G(\mathbf{z}, \Theta_G)); \Theta_D). \quad (10)$$

The first term is the data-likelihood, which computes the agreement between the observations and the image produced by the generator $G(\mathbf{n}, \theta_G)$ and the observations. The idea here is that this regularization will constrain the generator to produce image which remain close to $P_{model}(\mathbf{x})$.

2.4.3 Mitigating network imperfections

When simply optimizing the objective given in Eq. 10 for a random $\mathbf{z}$ two problems occurred.

checkerboard artifacts: The first one is the creation of checkerboard artifacts in the reconstructed image (an example of such pattern is presented Fig. 3). These types of artifacts appear more often after a large number of gradient descent steps of fine-tuning the generator. For larger values of μ they appear earlier. These gradient descent steps are referred to as epochs. These types of artifacts are typical for two processes: the forward-propagation of up-sampling layers and the back propagation of convolutional layers used to compute the gradient.³² The first is present in the Generator and the second is introduced by the use of the discriminator in equation 10. These types of artifacts are clear deviations of the prior which we wish to impose. For values of μ lower than 10, the optimization of data likelihood has, however, largely obtained before these features become prominent. We have therefore chosen to fine tune the generator for 250 epochs long, thereby, avoiding these artifacts. This lower amount of epochs is also beneficial for the run time. As a consequence, Eq. 10 will never be fully minimized. A different amount of epochs, may yet improve results further. Alterations to the neural network architecture intended to avoid these types of artifacts (e.g.^{32,33}) were not explored and are beyond the scope of this contribution.

Image variability The second problem is that the images are not consistent between runs, as they have different noise vectors. This is likely a consequence of both the optimization being stopped early and the regularization not being convex, causing the training procedure to get stuck in different local minimum, depending on the start position of the generator set by the input noise vector. In order to obtain a near constant image it was chosen to restart the image reconstruction 200 times and use the median image as the final reconstructed image.

2.4.4 Actual image reconstruction process

These alterations bring us to the image reconstruction algorithm that is actually used. The procedure consists of a double iterative loop. The inner loop, iterating over a number of epochs, is made of learning iterations with a constant input vector. During this inner loop the generator network is updated to descend the gradient of Eq. 10 for a fixed input z . Once the learning have reached a certain number of iterations (called n_{epoch}) the final image is stored and we start another iteration in the general loop. When restarting a new input z is randomly chosen from the noise prior $P_z(z)$ and the generator is reset to its state from the end of the pre-training phase.

The median of these images is then adopted as our final image reconstruction. This procedure is more formally described by the pseudo-code given in appendix B algorithm 2.

3. VALIDATION ON ARTIFICIAL DATASETS

3.1 Phase 1

We created three artificial datasets to validate our method. To keep it realistic we have used a (u, v) -coverage of existing data from the PIONIER instrument of the VLTI. We added Gaussian noise to the V^2 and CPs. These real datasets are observations of IRAS08544-4431 from³¹ and HD45677 from.^{8,9} The corresponding (u, v) -coverages can be found in Appendix F. Hereafter, these datasets will be referred to as dataset A, B and C respectively. The (u, v) -coverages and noise of datasets A and C are based on the observations of IRAS08544-4431, while dataset B is based on HD45677.

The images used for datasets A and B are taken from the grid of models used to train the GAN. The images used for all datasets have sizes of 128×128 pixels images with a field of view of 70 mas. Dataset B is further discussed in the appendix C. For dataset C an image consisting of a model image from the same grid with an added Gaussian flux. This Gaussian has an amplitude of 0.5 normalized flux (relative to the maximum flux in the image), a standard deviation of 1.8 mas in both the directions on the sky. The Gaussian is centered at $\Delta\alpha = 10.8 \text{ mas}$ and $\Delta\delta = 8.4 \text{ mas}$. This element was included in order to test the capability of the framework in reconstructing unexpected features.

Since the model images do not contain central stars, the contribution of a central star was added using SPARCO with values to make the datasets realistic. Hence, the stellar-to-total flux ratio of the primary at $1.65 \mu\text{m}$ is $f_{\text{prim}} = 0.597$, the diameter of the primary is $D = 0.5 \text{ mas}$ and the spectral index for the circumstellar environment $d_{\text{env}} = 0.42$. The datasets created in this way can be found in appendix G.

To validate the quality of the reconstructed images we used two metrics for comparing the used model image and the reconstructed one. We used the mean squared error (MSE) metric¹⁵ and normalized cross correlation (NCC).³⁴ When we apply these metrics we first normalize the images such that the total flux in both images equals unity, ensuring these images occupy the same Fourier space.

The normalized cross-correlation is given by:

$$\rho(X, Y) = \frac{1}{N} \sum_i \frac{(X_i - \langle X \rangle)(Y_i - \langle Y \rangle)}{\sigma_X \sigma_Y} \quad (11)$$

Here the summation is over all the N pixels of the model image X and the reconstructed image Y . $\langle X \rangle$ and $\langle Y \rangle$ are the mean pixel values of the images. σ_X and σ_Y are the standard deviations of the pixels in the two images. The normalized cross-correlation quantifies the similarity between images. A value of -1 implies perfect anti-correlation between images, 0 implies no correlation, and 1 implies perfect correlation.

The comparison between the true images and reconstructions of dataset A, B and C can be found in appendix E. For dataset A the reconstruction with the optimal MSE and NCC with respect to the model is that with $\mu = 5$, this reconstruction can be seen on figure4 next to the model used to create the corresponding artificial dataset. Figure4 it can be seen that our image reconstruction scheme is capable of reproducing the feature seen in the model rather well. A sharp inner rim can

Figure 4: Right: The model image used for constructing the artificial dataset. The used model image has $R_{in} = 7.5au$ $\rho_{in} = -0.6$ $\rho_{out} = 1.0$ $i = 45.3$ deg and $\lambda = 1.5\mu m$. Left: The image reconstruction using $\mu = 5$. The location of the added central star is indicated by the blue star. The gray ellipse on the reconstruction displays the beamsize. The beam size is defined as twice the Gaussian FWHM fitted to the interferometric point spread function.

Figure 5: Left: the model disk image with added Gaussian used to construct artificial dataset C. The used model image has $R_{in} = 9au$ $\rho_{in} = -0.6$ $\rho_{out} = 0.5$ $i = 25$ deg and $\lambda = 1.5\mu m$. Right: The image reconstruction with $\mu = 2$ using our method. The blue star marks the position of the added central star in both images. The colorbar indicates the flux normalized to the maximum in the image. The gray ellipse on the reconstruction displays the beamsize.

be seen and the flux away from the inner most part of the disk appears flat. The results on dataset B are discussed in Appendix C. Figure 5 displays the image used to construct dataset C next to the optimal reconstruction. Both the disk and the Gaussian flux are recovered well using our method. The recovery of the Gaussian illustrates that the method is capable of recovering unexpected features. The prior thus imposed using our framework, appears to be soft enough allowing for the complexity of a source to be recovered.

4. CONCLUSION

In this work the use of a GAN to reconstruct images from optical interferometric data was explored. In this framework a GAN is first trained on a set of model images representing the prior distribution. The component networks could then be used to preform image reconstructions, by fine-tuning the generator and using the discriminator as part of the regularization.

This framework was applied to realistic, artificial datasets based on model images which were a part of the training dataset of the GAN. It was shown that for these artificial datasets, the framework provides reconstructions of high quality

which are very similar to the original images. The images reconstructed for these datasets are extremely convincing and appear as almost axis symmetric disks. Artifacts which typically appear away from the inner rim of the disk did not appear in our image reconstructions.

The method appears not only capable of doing this for data based on the models used in training, but also when an unexpected feature is introduced in the data. This is illustrated by the reconstruction on dataset C. The Gaussian addition to this data-set is reproduced. This result indicates that the reconstruction method appears capable of allowing unexpected features, which are present in the data, to appear in the image.

To the best of our knowledge, this is the first time that the use of neural networks was explored in order to implement a Bayesian prior for the image reconstruction of interferometric data. For a first exploration of such a methodology, the results obtained here, appear promising.

There are still improvements that we plan to add to the image reconstruction scheme presented here. So far the main limitation is the computational cost of the method. Currently it takes about 13 minutes on a NVIDIA Tesla P100-SXM2-16GB gpu in order to reconstruct a single image with the framework presented here. This proves to be prohibitive with regards to bootstrapping the data in order to asses the significance of the features visible in the image.

The main contribution to this large computational expense in the reconstruction is the size of the generator network which needs to be retrained. This size is necessitated by the training of the GAN itself, as a generator with insufficient capacity will be incapable of mapping the input noise vectors to images of the distribution P_{data} . Exchanging the generator, for a different network capable of producing images does not give satisfactory results, as the discriminator appears to constrain the generator, rather than provide a gradient towards P_{data} . A network capable of providing such a gradient would thus drastically improve the computational cost of such reconstructions, and could potentially even be implemented in traditional image reconstruction algorithms.

Other schemes to use neural networks to reconstruct image can also be considered. One method that may be of interest is the fitting the input of a generative model such as a GAN's generator or a variational auto-encoder.³⁵ This results in images which follow the models extremely closely. Additional freedom can be introduced by including a sparse addition to the image.³⁶

Both the speed at which new developments are made in deep learning, and the ever increasing amount of computation power available make it seems likely that deep-learning regularization's will be further developed and will very likely have an important role to play in interferometric imaging.

ACKNOWLEDGMENTS

J.K. acknowledges support from the research council of the KU Leuven under grant number C14/17/082 and from FWO under the senior postdoctoral fellowship (1281121N). VSC

APPENDIX A. USED NEURAL NETWORK LAYERS

When training a GAN two neural networks need to be defined a discriminator network and a generator network. These networks where constructed using fully connected layers, convolutional layers and batch normalization layers. Here a short description of these types of layers is given.

Fully connected layers used in our networks the pre-activation outputs are computed using either

$$o_j = \sum_i a_{ij} \cdot x_i + b_{ij} \quad \text{or} \quad o_j = \sum_i a_{ij} \cdot x_i \quad (12)$$

depending if one uses biases (b_{ij}) or not. Both $\{a_{ij}\}$ and $\{b_{ij}\}$ are trainable parameters which are optimized using a gradient based learning rule.²⁰

Convolutional layers are based on discrete convolution operations.^{20,22} a chosen number of kernels of a chosen size are used to perform these operations. It are the values in these kernels which constitute the trainable parameters of a convolutional layer. The pixel values of a pre-activation output of a convolution with such a kernel is computed as follows:

$$S(i, j, k) = \sum_m \sum_n \sum_l I(i - m, j - n, l) K(m, n, l, k), \quad (13)$$

where I represent input into the layer, this can be a grayscale image with a size equal to 1 in the direction of k , an "rgb" image with a size equal to 3 in the direction of k or the output of a previous convolutional layer which has a depth equal to the number of kernels used in the previous layer.

K is the convolutional kernel and S the output corresponding to the kernel. The indexes m and n are summed over the kernel size. In a convolutional layer a "stride" is also defined. The stride defines for which values of i and j an output is computed. An output is computed for every value of i and j which is a multiple of the chosen stride in the relevant direction. A stride larger than 1 results in a downsized output image S . When a stride equal to two is used the number of pixels in S will be half that of the input images I .

Finally, the outputs of the both types of layers are further processed by activation functions, which introduce non-linearity into the network, allowing it to tackle more complex tasks.

Batch normalization works by first normalizing the individual images produced by a convolutional layer by using:

$$\hat{x}_i = \frac{x_i - \bar{x}_B}{\sqrt{\sigma_B^2 + \epsilon}} \quad (14)$$

Here $\bar{x}_B$ and σ_B^2 are the average and variance of the output images x_i of the previous layer during training on one mini-batch of data and $\epsilon = 0.001$ is a constant added to σ_B^2 for numerical stability. These normalized outputs are then shifted and re-scaled using:

$$y_i = \gamma \hat{x}_i + \beta \equiv BN_{\gamma, \beta}(x_i) \quad (15)$$

Where γ and β are trainable parameters which are optimized by the used gradient based learning rule. In our case, batch normalization layers where applied before the activations of the previously listed convolutional layers. Two additional trainable parameters are thus added per output image of each convolutional layer as can be seen in table 2. During inference $\bar{x}_B$ and σ_B^2 are replaced by a moving average $\bar{x}_M$ and moving variance σ_M^2 . The values of $\bar{x}_M$ and σ_M^2 are updated each training iteration using

$$\bar{x}_M = \bar{x}_M \cdot \alpha + \bar{x}_B \cdot (1 - \alpha) \quad (16)$$

and

$$\sigma_M^2 = \sigma_M^2 \cdot \alpha + \sigma_B^2 \cdot (1 - \alpha). \quad (17)$$

Here $\alpha = 0.99$ was used.

APPENDIX B. THE ALGORITHMS

The image reconstruction procedure presented in this paper consists of two phases. The first preparatory phase consists of training a GAN. A pseudocode detailing this training procedure is outlined in algorithm 1. After a GAN is trained, images can be reconstructed using both the component networks of the GAN. The procedure used to reconstruct images can be found in algorithm 2.

APPENDIX C. DATASET B

As mentioned in the main text an artificial dataset was created based on the baselines of observations of HD45677 previously used in.⁸ As with the dataset presented in the main text, an L-curve was plotted for various reconstructions and these reconstructions where compared to the models on which the data was based. The MSE and NCC for the preformed reconstructions can be found in appendix E. The image reconstruction with $\mu = 1$ can be seen on figure 6. This reconstruction is not as convincing as those mentioned in the main text. The most likely cause for this is the spacer UV-coverage, as this is an important factors in determining the quality of a reconstruction.¹⁵

Algorithm 1: 1 Mini-batch stochastic gradient descent training of generative adversarial nets. The gradient-based updates can use any standard gradient-based learning rule.

Require: generator network G with parameters θ_G ;
discriminator network D with parameters θ_D ;
A chosen number of training iterations;
A chosen mini-batch size m ;
data generating distribution $p_{data}(\mathbf{x})$;
noise prior distribution $p_g(\mathbf{z})$;
A gradient-based learning rule;

for *number of training iterations* **do**

- Sample mini-batch of m noise samples $\{\mathbf{z}^{(1)}, \dots, \mathbf{z}^{(m)}\}$ from noise prior $p_g(\mathbf{z})$
- Generate a mini-batch of m examples $\{\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(m)}\}$, where $\mathbf{x}^{(i)} = G(\mathbf{z}^{(i)}; \theta_G)$ and provide corresponding output labels $y^{(i)} = 0$
- Sample minibatch of m examples $\{\mathbf{x}^{(m)}, \dots, \mathbf{x}^{(2m)}\}$ from data generating distribution $p_{data}(\mathbf{x})$ and provide corresponding output labels $y^{(i)} = 0.9$
- Update the discriminator by using the chosen learning rule to descending its stochastic gradients:

$$-\Delta_{\theta_D} \frac{1}{2m} \sum_{i=1}^{2m} [y^{(i)} \log(D(\mathbf{x}^{(i)}; \theta_D)) + (1 - y^{(i)}) \log(1 - D(\mathbf{x}^{(i)}; \theta_D))]$$

- Sample mini-batch of m noise samples $\{\mathbf{z}^{(1)}, \dots, \mathbf{z}^{(m)}\}$ from noise prior $p_g(\mathbf{z})$ and provide corresponding output labels $y_i = 1$
- Update the generator by using the chosen learning rule to descending its stochastic gradients:

$$-\Delta_{\theta_G} \frac{1}{m} \sum_{i=1}^m y^{(i)} \log(D(G(\mathbf{z}; \theta_G); \theta_D))$$

end for

APPENDIX D. ABOUT ADAM OPTIMIZATION

The chosen gradient based learning rule is the Adam optimization algorithm.³⁷ During phase A the learning rate was set to $\alpha = 0.0002$ and the first moment's exponential decay rate to $\beta_1 = 0.5$.¹⁸ During phase B the learning rate was set to $\alpha = 0.0002$ and the first moment's exponential decay rate to $\beta_1 = 0.91$ In both cases the second moment exponential decay rate and the tolerance parameter we used $\beta_2 = 0.999$ and $\mu_{ADAM} = 10^{-7}$ respectively.

APPENDIX E. COMPARISON BETWEEN TRUE IMAGES AND RECONSTRUCTION

Algorithm 2: ORGANIC image reconstruction procedure.

Require: generator network G with parameters θ_G determined by the GAN training;
discriminator network D with parameters θ_D determined by the GAN training ;
A chosen number of times to reset the generator and resample a noise vector;
A chosen number of epochs for fine tuning the generator;
The hyperparameter μ ;
noise prior distribution $p_g(z)$;
A gradient-based learning rule;

for *number of resets* **do**

- set the generator parameters to those obtained at the end of GAN training;
- Sample a noise samples z from noise prior $p_g(z)$;

for *epochs* **do**

- Update the generator by using the chosen learning rule to descend the gradient:

$$\Delta_{\theta_G} f = f_{data}(G(z, \theta_G)) - \mu \log(D(G(z, \theta_G)); \theta_D)$$

end for

- Shift the image $G(z, \theta_G)$ to have a flux range between 0 and 1. ;
 - Normalize the image $G(z, \theta_G)$ to have a total flux equal to 1. ;
 - store the normalized image of this iteration for;
- end for** • compute the median image out of the set of normalized images;
-

Figure 6: Right: The model image used for constructing dataset B. The used model image has $R_{in} = 7.0au$ $\rho_{in} = -2.4$ $\rho_{out} = 1.5$ $i = 63.7$ deg and $\lambda = 2.1\mu m$. Left: The image reconstruction using $\mu = 5$. The location of the added central star is indicated by the blue star. The gray ellipse on the reconstruction displays the beamsizes.

Table 4: The MSE and NCC between the model and reconstructions of dataset A using different values of μ . The optimal values of the MSE and NCC are indicated in bold. The data-likelihood term of the final image f_{data} , its components per type of observable (over the number of that observable type.) and the "regularization" term $-\log(D(x_f))$ are also listed.

μ	MSE	NCC	$\chi^2_{v^2}/N_{v^2}$	$\chi^2_{c.p.}/N_{c.p.}$	$f_{data}(x_f)$	$-\log(D(x_f))$
10^{-6}	$7.4979 \cdot 10^{-9}$	0.9798	1.213	1.099	1.170	0.8115067
0.0001	$7.4038 \cdot 10^{-9}$	0.9801	1.244	1.087	1.185	0.8115067
0.001	$7.9423 \cdot 10^{-9}$	0.9786	1.218	1.104	1.174	0.8115043
0.01	$7.3051 \cdot 10^{-9}$	0.9803	1.215	1.102	1.172	0.8115103
0.1	$7.0300 \cdot 10^{-9}$	0.9811	1.226	1.099	1.178	0.8115058
1	$4.1973 \cdot 10^{-9}$	0.9889	1.138	1.225	1.172	0.8114951
2	$2.9382 \cdot 10^{-9}$	0.9924	1.084	1.316	1.172	0.8114925
5	$2.7834 \cdot 10^{-9}$	0.9930	1.043	1.411	1.182	0.8114937
10	$3.3611 \cdot 10^{-9}$	0.9913	1.039	1.469	1.202	0.8114935
50	$9.1718 \cdot 10^{-9}$	0.9753	1.558	1.221	1.431	0.8114618
100	$1.503 \cdot 10^{-8}$	0.9592	1.593	2.390	1.895	0.8114126

Table 5: The MSE and between the model and reconstructions of dataset B using different values of μ . The optimal values of the MSE and NCC are indicated in bold. The data-likelihood term of the final image f_{data} , its components per type of observable (over the number of that observable type.) and the "regularization" term $-\log(D(x_f))$ are also listed.

μ	MSE	NCC	$\chi^2_{v^2}/N_{v^2}$	$\chi^2_{c.p.}/N_{c.p.}$	f_{data}	$-\log(D(x_f))$
10^{-6}	$6.2619 \cdot 10^{-9}$	0.9774	3.301	0.927	2.398	0.8113948
0.0001	$6.0988 \cdot 10^{-9}$	0.9780	3.429	0.937	2.481	0.8113982
0.001	$6.2362 \cdot 10^{-9}$	0.9775	3.428	0.925	2.476	0.8113952
0.01	$6.3580 \cdot 10^{-9}$	0.9719	3.737	0.949	2.676	0.8113961
0.1	$6.2111 \cdot 10^{-9}$	0.9776	2.446	0.839	2.566	0.8113965
1	$4.6461 \cdot 10^{-9}$	0.9861	3.060	0.941	2.254	0.8113434
2	$3.9363 \cdot 10^{-9}$	0.9856	2.389	0.928	1.833	0.8113753
5	$3.4729 \cdot 10^{-9}$	0.9873	1.970	0.895	1.578	0.8113620
10	$4.0484 \cdot 10^{-9}$	0.9856	2.569	1.002	1.972	0.8113434
50	$1.3790 \cdot 10^{-8}$	0.9601	14.672	2.154	9.908	0.8113262
100	$2.4734 \cdot 10^{-8}$	0.9174	18.587	2.450	12.446	0.8112975

Table 6: The MSE and NCC between the model and reconstructions of dataset C using different values of μ . The optimal values of the MSE and NCC are indicated in bold. The data-likelihood term of the final image f_{data} , its components per type of observable (over the number of that observable type.) and the "regularization" term $-\log(D(x_f))$ are also listed.

μ	MSE	NCC	$\chi^2_{v^2}/N_{v^2}$	$\chi^2_{c.p.}/N_{c.p.}$	f_{data}	$-\log(D(x_f))$
10^{-6}	$4.1767 \cdot 10^{-9}$	0.9746	1.016	1.006	1.012	0.8113182
0.0001	$4.2596 \cdot 10^{-9}$	0.9743	1.015	1.003	1.010	0.8113189
0.001	$3.9964 \cdot 10^{-9}$	0.9757	1.007	0.942	0.982	0.8113192
0.01	$3.5799 \cdot 10^{-9}$	0.9785	0.994	0.933	0.971	0.8113183
0.1	$3.6523 \cdot 10^{-9}$	0.9779	0.935	1.006	0.980	0.8113186
1	$3.1727 \cdot 10^{-9}$	0.9826	0.990	0.891	0.952	0.8113150
2	$2.5417 \cdot 10^{-9}$	0.9864	0.990	0.895	0.954	0.8113124
5	$2.7034 \cdot 10^{-9}$	0.9847	1.017	0.907	0.976	0.8113014
10	$4.0376 \cdot 10^{-9}$	0.9734	1.334	1.783	1.504	0.8112868
50	$4.0191 \cdot 10^{-8}$	0.6847	4.962	17.167	9.580	0.8112635
100	$5.5828 \cdot 10^{-8}$	0.5341	6.369	22.045	12.300	0.8112569

APPENDIX F. UV-COVERAGES

Figure 7: The UV-coverages of IRAS08544-4431 (left) and HD456777 (right).

APPENDIX G. COMPARISON BETWEEN DATA AND RECONSTRUCTION

Figure 8: Top: Comparison between the squared visibilities of the artificial datasets and the corresponding best image reconstruction. The contributions of a central star are present in both. Bottom: the residuals normalized by the error on the corresponding data point. The residuals are given as $(v_{data}^2 - v_{image}^2)/\sigma_{v^2}$.

Figure 9: Top: Comparison between the closure phases of the artificial datasets and the corresponding best image reconstruction. The contributions of a central star are present in both. Bottom: the residuals normalized by the error on the corresponding data point. The residuals are given as $(cp_{data} - cp_{image})/\sigma_{cp}$.

ACKNOWLEDGMENTS

We would like to thank Jeroen Audenaert, Ferréol Soulez and Giuseppe Marra for interesting discussions on the work leading to this paper. RC, JK and HVW acknowledge support from the research council of the KU Leuven under grant number C14/17/082. The resources and services used in this work were provided by the VSC (Flemish Supercomputer Center), funded by the Research Foundation - Flanders (FWO) and the Flemish Government. This research has made use of the Jean-Marie Mariotti Center OiDB service available at <http://oidb.jmmc.fr>.

REFERENCES

- [1] Gravity Collaboration, Abuter, R., Accardo, M., Amorim, A., Anugu, N., Ávila, G., Azouaoui, N., Benisty, M., Berger, J. P., Blind, N., Bonnet, H., Bourget, P., Brandner, W., Brast, R., Buron, A., Burtscher, L., Cassaing, F., Chapron, F., Choquet, É., Clénet, Y., Collin, C., Coudé Du Foresto, V., de Wit, W., de Zeeuw, P. T., Deen, C., Delplancke-Ströbele, F., Dembet, R., Derie, F., Dexter, J., Duvert, G., Ebert, M., Eckart, A., Eisenhauer, F., Esselborn, M., Fédou, P., Finger, G., Garcia, P., Garcia Dabo, C. E., Garcia Lopez, R., Gendron, E., Genzel, R., Gillessen, S., Gontte, F., Gordo, P., Grould, M., Grözing, U., Guieu, S., Haguenaue, P., Hans, O., Haubois, X., Haug, M., Haussmann, F., Henning, T., Hippler, S., Horrobin, M., Huber, A., Hubert, Z., Hubin, N., Hummel, C. A., Jakob, G., Janssen, A., Jochum, L., Jocu, L., Kaufer, A., Kellner, S., Kendrew, S., Kern, L., Kervella, P., Kiekebusch, M., Klein, R., Kok, Y., Kolb, J., Kulas, M., Lacour, S., Lapeyrière, V., Lazareff, B., Le Bouquin, J. B., Lèna, P., Lenzen, R., Lévêque, S., Lippa, M., Magnard, Y., Mehrgan, L., Mellein, M., Mérand, A., Moreno-Ventas, J., Moulin, T., Müller, E., Müller, F., Neumann, U., Oberti, S., Ott, T., Pallanca, L., Panduro, J., Pasquini, L., Paumard, T., Percheron, I., Perraut, K., Perrin, G., Pflüger, A., Pfuhl, O., Phan Duc, T., Plewa, P. M., Popovic, D., Rabien, S., Ramírez, A., Ramos, J., Rau, C., Riquelme, M., Rohloff, R. R., Rousset, G., Sanchez-Bermudez, J., Scheithauer, S., Schöller, M., Schuhler, N., Spyromilio, J., Straubmeier, C., Sturm, E., Suarez, M., Tristram, K. R. W., Ventura, N., Vincent, F., Waisberg, I., Wank, I., Weber, J., Wieprecht, E., Wiest, M., Wiezorrek, E., Wittkowski, M., Woillez, J., Wolff, B.,

Yazici, S., Ziegler, D., and Zins, G., “First light for GRAVITY: Phase referencing optical interferometry for the Very Large Telescope Interferometer,” *A&A* **602**, A94 (June 2017).

- [2] Le Bouquin, J. B., Berger, J. P., Lazareff, B., Zins, G., Haguenaue, P., Jocou, L., Kern, P., Millan-Gabet, R., Traub, W., Absil, O., Augereau, J. C., Benisty, M., Blind, N., Bonfils, X., Bourget, P., Delboulbe, A., Feautrier, P., Germain, M., Gitton, P., Gillier, D., Kiekebusch, M., Kluska, J., Knudstrup, J., Labeye, P., Lizon, J. L., Monin, J. L., Magnard, Y., Malbet, F., Maurel, D., Ménard, F., Micallef, M., Michaud, L., Montagnier, G., Morel, S., Moulin, T., Perraut, K., Popovic, D., Rabou, P., Rochat, S., Rojas, C., Roussel, F., Roux, A., Stadler, E., Stefl, S., Tatulli, E., and Ventura, N., “PIONIER: a 4-telescope visitor instrument at VLTI,” *A&A* **535**, A67 (Nov. 2011).
- [3] Lopez, B., Lagarde, S., Jaffe, W., Petrov, R., Schöller, M., Antonelli, P., Beckmann, U., Berio, P., Bettonvil, F., Glindemann, A., Gonzalez, J. C., Graser, U., Hofmann, K. H., Millour, F., Robbe-Dubois, S., Venema, L., Wolf, S., Henning, T., Lanz, T., Weigelt, G., Agocs, T., Bailet, C., Bresson, Y., Bristow, P., Dugué, M., Heininger, M., Kroes, G., Laun, W., Lehmitz, M., Neumann, U., Augereau, J. C., Avila, G., Behrend, J., van Belle, G., Berger, J. P., van Boekel, R., Bonhomme, S., Bourget, P., Brast, R., Clausse, J. M., Connot, C., Conzelmann, R., Cruzalèbes, P., Csepány, G., Danchi, W., Delbo, M., Delplancke, F., Dominik, C., van Duin, A., Elswijk, E., Fantei, Y., Finger, G., Gabasch, A., Gay, J., Girard, P., Girault, V., Gitton, P., Glazenberg, A., Gonté, F., Guitton, F., Guniat, S., De Haan, M., Haguenaue, P., Hanenburg, H., Hogerheijde, M., ter Horst, R., Hron, J., Hugues, Y., Hummel, C., Idserda, J., Ives, D., Jakob, G., Jasko, A., Jolley, P., Kiraly, S., Köhler, R., Kragt, J., Kroener, T., Kuindersma, S., Labadie, L., Leinert, C., Le Poole, R., Lizon, J. L., Lucuix, C., Marcotto, A., Martinache, F., Martinot-Lagarde, G., Mathar, R., Matter, A., Mauclet, N., Mehrgan, L., Meilland, A., Meisenheimer, K., Meisner, J., Mellein, M., Menardi, S., Menut, J. L., Merand, A., Morel, S., Mosoni, L., Navarro, R., Nussbaum, E., Ottogalli, S., Palsa, R., Panduro, J., Pantin, E., Parra, T., Percheron, I., Duc, T. P., Pott, J. U., Pozna, E., Przygodda, F., Rabbia, Y., Richichi, A., Rigal, F., Roelfsema, R., Rupprecht, G., Schertl, D., Schmidt, C., Schuhler, N., Schuil, M., Spang, A., Stegmeier, J., Thiam, L., Tromp, N., Vakili, F., Vannier, M., Wagner, K., and Woillez, J., “An Overview of the MATISSE Instrument — Science, Concept and Current Status,” *The Messenger* **157**, 5–12 (Sept. 2014).
- [4] Kraus, S., Monnier, J. D., Anugu, N., Le Bouquin, J.-B., Davies, C. L., Ennis, J., Labdon, A., Lanthermann, C., Setterholm, B., and ten Brummelaar, T., “The MIRC-X 6-telescope imager: key science drivers, instrument design and operation,” in [Optical and Infrared Interferometry and Imaging VI], *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **10701**, 1070123 (July 2018).
- [5] Kloppenborg, B., Stencel, R., Monnier, J. D., Schaefer, G., Zhao, M., Baron, F., McAlister, H., ten Brummelaar, T., Che, X., Farrington, C., Pedretti, E., Sallave-Goldfinger, P. J., Sturmman, J., Sturmman, L., Thureau, N., Turner, N., and Carroll, S. M., “Infrared images of the transiting disk in the ϵ Aurigae system,” *Nature* **464**, 870–872 (Apr. 2010).
- [6] Roettenbacher, R. M., Monnier, J. D., Korhonen, H., Aarnio, A. N., Baron, F., Che, X., Harmon, R. O., Kővári, Z., Kraus, S., Schaefer, G. H., Torres, G., Zhao, M., Ten Brummelaar, T. A., Sturmman, J., and Sturmman, L., “No Sun-like dynamo on the active star ζ Andromedae from starspot asymmetry,” *Nature* **533**, 217–220 (May 2016).
- [7] Paladini, C., Baron, F., Jorissen, A., Le Bouquin, J. B., Freytag, B., van Eck, S., Wittkowski, M., Hron, J., Chiavassa, A., Berger, J. P., Siopis, C., Mayer, A., Sadowski, G., Kravchenko, K., Shetye, S., Kerschbaum, F., Kluska, J., and Ramstedt, S., “Large granulation cells on the surface of the giant star π^1 Gruis,” *Nature* **553**, 310–312 (Jan. 2018).
- [8] Kluska, J., Berger, J. P., Malbet, F., Lazareff, B., Benisty, M., Le Bouquin, J. B., Absil, O., Baron, F., Delboulbé, A., Duvert, G., Isella, A., Jocou, L., Juhasz, A., Kraus, S., Lachaume, R., Ménard, F., Millan-Gabet, R., Monnier, J. D., Moulin, T., Perraut, K., Rochat, S., Pinte, C., Soulez, F., Tallon, M., Thi, W. F., Thiébaud, E., Traub, W., and Zins, G., “A family portrait of disk inner rims around Herbig Ae/Be stars. Hunting for warps, rings, self shadowing, and misalignments in the inner astronomical units,” *A&A* **636**, A116 (Apr. 2020).
- [9] Lazareff, B., Berger, J. P., Kluska, J., Le Bouquin, J. B., Benisty, M., Malbet, F., Koen, C., Pinte, C., Thi, W. F., Absil, O., Baron, F., Delboulbé, A., Duvert, G., Isella, A., Jocou, L., Juhasz, A., Kraus, S., Lachaume, R., Ménard, F., Millan-Gabet, R., Monnier, J. D., Moulin, T., Perraut, K., Rochat, S., Soulez, F., Tallon, M., Thiébaud, E., Traub, W., and Zins, G., “Structure of Herbig Ae/Be disks at the milliarcsecond scale. A statistical survey in the H band using PIONIER-VLTI,” *A&A* **599**, A85 (Mar. 2017).
- [10] Kluska, J., Van Winckel, H., Hillen, M., Berger, J. P., Kamath, D., Le Bouquin, J. B., and Min, M., “VLTI/PIONIER survey of disks around post-AGB binaries. Dust sublimation physics rules,” *A&A* **631**, A108 (Nov. 2019).

- [11] Gravity Collaboration, Perraut, K., Labadie, L., Lazareff, B., Klarmann, L., Segura-Cox, D., Benisty, M., Bouvier, J., Brandner, W., Caratti O Garatti, A., Caselli, P., Dougados, C., Garcia, P., Garcia-Lopez, R., Kendrew, S., Koutoulaki, M., Kervella, P., Lin, C. C., Pineda, J., Sanchez-Bermudez, J., van Dishoeck, E., Abuter, R., Amorim, A., Berger, J. P., Bonnet, H., Buron, A., Cantalloube, F., Clénet, Y., Coudé Du Foresto, V., Dexter, J., de Zeeuw, P. T., Duvert, G., Eckart, A., Eisenhauer, F., Eupen, F., Gao, F., Gendron, E., Genzel, R., Gillessen, S., Gordo, P., Grellmann, R., Haubois, X., Haussmann, F., Henning, T., Hippler, S., Horrobin, M., Hubert, Z., Jocou, L., Lacour, S., Le Bouquin, J. B., Léna, P., Mérand, A., Ott, T., Paumard, T., Perrin, G., Pfuhl, O., Rabien, S., Ray, T., Rau, C., Rousset, G., Scheithauer, S., Straub, O., Straubmeier, C., Sturm, E., Vincent, F., Waisberg, I., Wank, I., Widmann, F., Wieprecht, E., Wiest, M., Wiezorrek, E., Woillez, J., and Yazici, S., “The GRAVITY Young Stellar Object survey. I. Probing the disks of Herbig Ae/Be stars in terrestrial orbits,” *A&A* **632**, A53 (Dec. 2019).
- [12] Thiébaud, É. and Young, J., “Principles of image reconstruction in optical interferometry: tutorial,” *Journal of the Optical Society of America A* **34**, 904 (June 2017).
- [13] Thiébaud, E., “MIRA: an effective imaging algorithm for optical interferometry,” in [*Optical and Infrared Interferometry*], *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **7013**, 70131I (July 2008).
- [14] Baron, F., Monnier, J. D., and Kloppenborg, B., “A novel image reconstruction software for optical/infrared interferometry,” in [*Optical and Infrared Interferometry II*], *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **7734**, 77342I (July 2010).
- [15] Renard, S., Thiébaud, E., and Malbet, F., “Image reconstruction in optical interferometry: benchmarking the regularization,” *A&A* **533**, A64 (Sept. 2011).
- [16] Hofmann, K.-H., Heininger, M., Schertl, D., Weigelt, G., Millour, F., and Berio, P., “Image reconstruction method IRBis for optical/infrared long-baseline interferometry,” in [*Optical and Infrared Interferometry and Imaging V*], *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series* **9907**, 99073H (July 2016).
- [17] LeCun, Y., Bengio, Y., and Hinton, G., “Deep learning,” *Nature* **521**(7553), 436–444 (2015).
- [18] Radford, A., Metz, L., and Chintala, S., “Unsupervised representation learning with deep convolutional generative adversarial networks,” (2015).
- [19] Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y., “Generative adversarial networks,” (2014).
- [20] Goodfellow, I., Bengio, Y., and Courville, A., [*Deep Learning*], MIT Press (2016). <http://www.deeplearningbook.org>.
- [21] Maas, A. L., Hannun, A. Y., and Ng, A. Y., “Rectifier nonlinearities improve neural network acoustic models,” in [*Proc. icml*], **30**, 3, ICML Workshop on Deep Learning for Audio, Speech, and Language Processing (WDLASL 2013) (2013).
- [22] Dumoulin, V. and Visin, F., “A guide to convolution arithmetic for deep learning,” *arXiv e-prints*, arXiv:1603.07285 (Mar. 2016).
- [23] Ioffe, S. and Szegedy, C., “Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift,” *arXiv e-prints*, arXiv:1502.03167 (Feb. 2015).
- [24] Min, M., Dullemond, C. P., Dominik, C., de Koter, A., and Hovenier, J. W., “Radiative transfer in very optically thick circumstellar disks,” *A&A* **497**, 155–166 (Apr. 2009).
- [25] Hillen, M., Menu, J., Van Winckel, H., Min, M., Gielen, C., Wevers, T., Mulders, G. D., Regibo, S., and Verhoelst, T., “An interferometric study of the post-AGB binary 89 Herculis. II. Radiative transfer models of the circumbinary disk,” *A&A* **568**, A12 (Aug. 2014).
- [26] Hillen, M., de Vries, B. L., Menu, J., Van Winckel, H., Min, M., and Mulders, G. D., “The evolved circumbinary disk of AC Herculis: a radiative transfer, interferometric, and mineralogical study,” *A&A* **578**, A40 (June 2015).
- [27] Kluska, J., Hillen, M., Van Winckel, H., Manick, R., Min, M., Regibo, S., and Royer, P., “The perturbed sublimation rim of the dust disk around the post-AGB binary IRAS08544-4431,” *A&A* **616**, A153 (Aug. 2018).
- [28] Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., and Chen, X., “Improved techniques for training gans,” in [*Advances in neural information processing systems*], 2234–2242 (2016).
- [29] Goodfellow, I., “NIPS 2016 Tutorial: Generative Adversarial Networks,” *arXiv e-prints*, arXiv:1701.00160 (Dec. 2016).
- [30] Kluska, J., Malbet, F., Berger, J. P., Baron, F., Lazareff, B., Le Bouquin, J. B., Monnier, J. D., Soulez, F., and Thiébaud, E., “SPARCO : a semi-parametric approach for image reconstruction of chromatic objects. Application to young stellar objects,” *A&A* **564**, A80 (Apr. 2014).

- [31] Hillen, M., Kluska, J., Le Bouquin, J. B., Van Winckel, H., Berger, J. P., Kamath, D., and Bujarrabal, V., “Imaging the dust sublimation front of a circumbinary disk,” *A&A* **588**, L1 (Apr. 2016).
- [32] Sugawara, Y., Shiota, S., and Kiya, H., “Checkerboard artifacts free convolutional neural networks,” *APSIPA Transactions on Signal and Information Processing* **8**, e9 (2019).
- [33] Gao, H., Yuan, H., Wang, Z., and Ji, S., “Pixel deconvolutional networks,” *arXiv preprint arXiv:1705.06820* (2017).
- [34] Event Horizon Telescope Collaboration, Akiyama, K., Alberdi, A., Alef, W., Asada, K., Azulay, R., Bacsko, A.-K., Ball, D., Baloković, M., Barrett, J., Bintley, D., Blackburn, L., Boland, W., Bouman, K. L., Bower, G. C., Bremer, M., Brinkerink, C. D., Brissenden, R., Britzen, S., Broderick, A. E., Brogiere, D., Bronzwaer, T., Byun, D.-Y., Carlstrom, J. E., Chael, A., Chan, C.-k., Chatterjee, S., Chatterjee, K., Chen, M.-T., Chen, Y., Cho, I., Christian, P., Conway, J. E., Cordes, J. M., Crew, G. B., Cui, Y., Davelaar, J., De Laurentis, M., Deane, R., Dempsey, J., Desvignes, G., Dexter, J., Doeleman, S. S., Eatough, R. P., Falcke, H., Fish, V. L., Fomalont, E., Fraga-Encinas, R., Freeman, W. T., Friberg, P., Fromm, C. M., Gómez, J. L., Galison, P., Gammie, C. F., García, R., Gentaz, O., Georgiev, B., Goddi, C., Gold, R., Gu, M., Gurwell, M., Hada, K., Hecht, M. H., Hesper, R., Ho, L. C., Ho, P., Honma, M., Huang, C.-W. L., Huang, L., Hughes, D. H., Ikeda, S., Inoue, M., Issaoun, S., James, D. J., Jannuzi, B. T., Janssen, M., Jeter, B., Jiang, W., Johnson, M. D., Jorstad, S., Jung, T., Karami, M., Karuppusamy, R., Kawashima, T., Keating, G. K., Kettenis, M., Kim, J.-Y., Kim, J., Kim, J., Kino, M., Koay, J. Y., Koch, P. M., Koyama, S., Kramer, M., Kramer, C., Krichbaum, T. P., Kuo, C.-Y., Lauer, T. R., Lee, S.-S., Li, Y.-R., Li, Z., Lindqvist, M., Liu, K., Liuzzo, E., Lo, W.-P., Lobanov, A. P., Loinard, L., Lonsdale, C., Lu, R.-S., MacDonald, N. R., Mao, J., Markoff, S., Marrone, D. P., Marscher, A. P., Martí-Vidal, I., Matsushita, S., Matthews, L. D., Medeiros, L., Menten, K. M., Mizuno, Y., Mizuno, I., Moran, J. M., Moriyama, K., Moscibrodzka, M., Müller, C., Nagai, H., Nagar, N. M., Nakamura, M., Narayan, R., Narayanan, G., Natarajan, I., Neri, R., Ni, C., Noutsos, A., Okino, H., Olivares, H., Oyama, T., Özel, F., Palumbo, D. C. M., Patel, N., Pen, U.-L., Pesce, D. W., Piétu, V., Plambeck, R., PopStefanija, A. a., Porth, O., Prather, B., Preciado-López, J. A., Psaltis, D., Pu, H.-Y., Ramakrishnan, V., Rao, R., Rawlings, M. G., Raymond, A. W., Rezzolla, L., Ripperda, B., Roelofs, F., Rogers, A., Ros, E., Rose, M., Roshanineshat, A., Rottmann, H., Roy, A. L., Ruszczyk, C., Ryan, B. R., Rygl, K. L. J., Sánchez, S., Sánchez-Arguelles, D., Sasada, M., Savolainen, T., Schloerb, F. P., Schuster, K.-F., Shao, L., Shen, Z., Small, D., Sohn, B. W., SooHoo, J., Tazaki, F., Tiede, P., Tilanus, R. P. J., Titus, M., Toma, K., Torne, P., Trent, T., Trippe, S., Tsuda, S., van Bemmell, I., van Langevelde, H. J., van Rossum, D. R., Wagner, J., Wardle, J., Weintraub, J., Wex, N., Wharton, R., Wielgus, M., Wong, G. N., Wu, Q., Young, A., Young, K., Younsi, Z., Yuan, F., Yuan, Y.-F., Zensus, J. A., Zhao, G., Zhao, S.-S., Zhu, Z., Farah, J. R., Meyer-Zhao, Z., Michalik, D., Nadolski, A., Nishioka, H., Pradel, N., Primiani, R. A., Souccar, K., Vertatschitsch, L., and Yamaguchi, P., “First M87 Event Horizon Telescope Results. IV. Imaging the Central Supermassive Black Hole,” *ApJ* **875**, L4 (Apr. 2019).
- [35] Bora, A., Jalal, A., Price, E., and Dimakis, A. r. G., “Compressed Sensing using Generative Models,” *arXiv e-prints*, arXiv:1703.03208 (Mar. 2017).
- [36] Dhar, M., Grover, A., and Ermon, S., “Modeling Sparse Deviations for Compressed Sensing using Generative Models,” *arXiv e-prints*, arXiv:1807.01442 (July 2018).
- [37] Kingma, D. P. and Ba, J., “Adam: A method for stochastic optimization,” (2014).