Differentiable probabilistic programming for strong gravitational lensing

Marco Chianese[*][†]

Gravitation Astroparticle Physics Amsterdam (GRAPPA), Institute for Theoretical Physics Amsterdam and Delta Institute for Theoretical Physics, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands

E-mail: m.chianese@uva.nl

The difficult task of observing Dark Matter subhaloes is of paramount importance since it would constrain Dark Matter particle properties (cold or warm relic) and confirm once again the long-standing ΛCDM model. In the near future the new generation of ground and space surveys will observe thousands of strong gravitational lensing systems providing a unique probe of Dark Matter substructures. Here, we describe a new strong lensing analysis pipeline that combines deep Convolutional Neural Networks with physical models and exploits traditional sampling techniques such as Hamiltonian Monte Carlo. Using simulated strong gravitational lensing systems, we discuss first results and characterize the accuracy of the reconstruction of the main lensing parameters.

---

[*]Speaker.

1. Introduction

Although several astrophysical and cosmological observations strongly support the existence of Dark Matter (DM), its nature is yet to be understood [1, 2]. The standard ΛCDM cosmological model successfully explains the large-scale structure of the Universe and its formation through the collapse of primordial density fluctuations. The dominant contribution to the matter content of the Universe is due to the Dark Matter that is assumed to be cold (CDM) and almost collisionless. This means that at early times DM particles have non-relativistic velocities so allowing the structures to gradually form. Under these assumptions, ab-initio N-body cosmological simulations have pointed out that DM haloes are filled by DM substructures [3]. A key quantity is the subhalo mass function $\mathrm{d}N/\mathrm{d}\ln m_{\mathrm{sub}}$, that is the abundance of DM subhaloes per unit logarithmic subhalo mass m_{sub} within a host DM halo M at a given redshift.

A still viable alternative to the CDM paradigm is given by Warm Dark Matter (WDM). In this case, at the time of structure formation, WDM particles have instead non-negligible velocities that can wash out the primordial matter density fluctuations at about 1 kpc comoving scale. This results in a suppression of the formation of small-mass subhaloes. The free-streaming effects on the subhalo mass function are in general parameterized by the functional form [4]

$$\frac{\mathrm{d}N_{\mathrm{WDM}}}{\mathrm{d}\ln m_{\mathrm{sub}}} = \frac{\mathrm{d}N_{\mathrm{CDM}}}{\mathrm{d}\ln m_{\mathrm{sub}}} \left(1 + \frac{m_{\mathrm{hm}}}{m_{\mathrm{sub}}}\right)^{-1.3}, \tag{1.1}$$

where for thermally produced dark matter particles, the half-mode mass m_{hm} is equal to [5, 6]

$$m_{\mathrm{hm}} = 10^{10} \left(\frac{m_{\mathrm{DM}}}{1\,\mathrm{keV}}\right)^{-3.33} M_{\odot}\, h^{-1}, \tag{1.2}$$

with m_{DM} being the DM mass. By using the above expressions and the fitting formula reported in Ref. [7], in Fig. 1 we show the subhalo mass function with a DM host halo of $M = 10^{12}\, M_{\odot}\, h^{-1}$, for CDM and WDM paradigms with different DM masses. The smaller the mass of thermal DM particles, the fewer the DM substructures with small mass. Hence, a naive way to test the nature of dark matter is to count DM subhaloes by looking for faint and ultra-faint dwarf galaxies inside the Milky Way or other galaxies. This procedure has indeed set a lower bound on the mass of thermal DM particles: $m_{\mathrm{DM}} > 1.5$ keV [4] and $m_{\mathrm{DM}} > 2.3$ keV [8] depending on different assumptions. However, star formation in DM subhaloes with $m_{\mathrm{sub}} \lesssim 10^8\, M_{\odot}\, h^{-1}$ is expected to be suppressed by reionization, interaction with their host and self-quenching [9, 10]. An alternative way to detect small-mass subhaloes and, consequently, to discriminate between CDM and WDM is by looking for their gravitational effects in strong gravitational lensing systems [11, 12, 13, 14, 15].

Strong gravitational lensing is a gravitational effect that allows us to observe multiple images of a background galaxy (the source) through the deflection of the light due to the distribution of matter along the line-of-sight (see Ref. [16] for a review). Typically, the main contribution to the lensing comes from a foreground galaxy (the lens) that is almost aligned to the background one. In this case, we can observe the so-called Einstein ring. For these systems, the presence of DM subhaloes bounded to the lens galaxy and of small DM haloes along the line-of-sight (field haloes) affects the lensed images below the percentage level only. Hence, to detect these very small distortions, an accurate and fast analysis pipeline for strong gravitational lensing data is required.
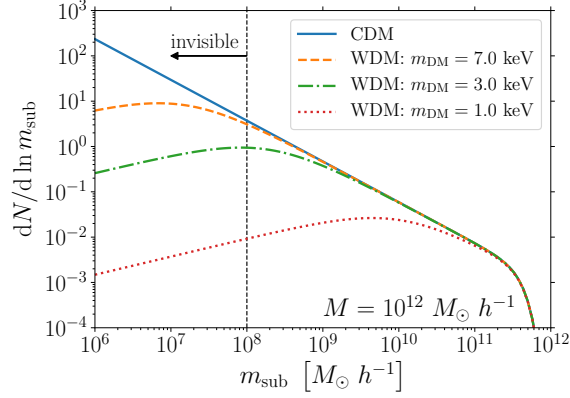
Figure 1: Subhalo mass function from N-body simulations for CDM and WDM paradigms with different DM masses. Subhalo mass function as a function of the subhalo mass within a host DM halo of $M = 10^{12} M_\odot h^{-1}$. The lines represent the expectation for CDM particles (solid blue lines) and WDM particles with different masses. Subhaloes with mass smaller than about $10^8 M_\odot h^{-1}$ (vertical line) are expected to be dark due to the suppression of star formation.

In strong gravitational lensing analyses, the zeroth order step is to reconstruct the observed lensed image by modeling at the same time the surface brightness of the source galaxy and the matter distribution of the main lens galaxy. Regarding the source, two methods have been exploited so far: the use of parametric models like the Sersic profile or the pixelated reconstruction of the source [17, 18]. While the former is not able to cover all the complex morphologies of galaxies, the latter has no *a priori* information of real galaxies and, therefore, requires a regularization according to some criteria (like smoothness among adjacent pixels). On the other hand, Machine Learning techniques can be employed to generate more realistic templates for the background source based on our knowledge and observation of real galaxies.

Deep learning methods have already been used for galaxy classification and for searches of strong lensing systems in huge observational catalogs of galaxies [19, 20]. Moreover, Convolutional Neural Networks (CNNs) have been proposed to quickly recover the main parameters describing the matter distribution of the lens (the size of the Einstein radius, the eccentricity, and so on) [21, 22]. Recently, recurrent inference machines have been shown to successfully reconstruct the background source, once they are coupled to an additional CNN that provides the lens model [23]. Both pipelines are based on supervised learning, namely the CNNs are trained by feeding them with mock lensed images for which the lens and source true models are known. This does not seem to be a viable approach if one wants to include subhaloes and field haloes since this would highly enlarge the parameter space of the system and, consequently, the required training data set.

Differently from the aforementioned methods, we develop a new analysis pipeline that combines deep CNNs with physical models in order to increase the interpretability of the results [24]. In particular, the surface brightness of the source galaxy is modeled by a Variational AutoEncoder (VAE), that is an unsupervised generative model consisting of two Neural Networks, the encoder and the decoder [25, 26, 27]. In particular, the encoder maps the input images to a latent space with a smaller dimensionality, while the decoder produces output images given the latent space vari-
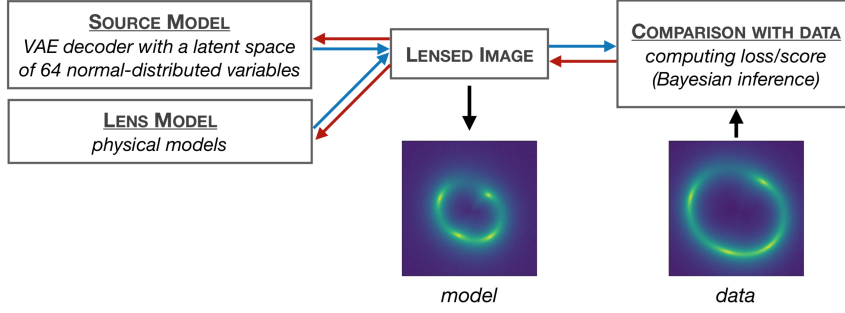
Figure 2: Our analysis pipeline implemented in a differential probabilistic programming framework. The forward model (blue lines) computes the lensed image combining the source model (the VAE decoder) and the lens model (physical models for the lens mass distribution). Then, the image is compared with the data to infer the model parameters in a Bayesian approach. The backward model (red line) computes the derivatives of the forward model with respect to all the parameters.

ables. Once the VAE is trained with a data set of observed non-lensed galaxies, the decoder can be used to generate reasonable galaxy images. The lens model is instead described by physical models that depend on parameters with explicit physical meaning. The whole analysis pipeline is implemented in a differential probabilistic programming framework using Pytorch [28] and Pyro [29] packages. This allows us to automatically and efficiently compute the derivatives of the model with respect to all the source and lens parameters. These derivatives are used to find the best-fit parameters (optimization) and to estimate the parameter posteriors with a Bayesian approach by means of traditional sampling techniques such as Hamiltonian Monte Carlo (HMC) [30, 31].

2. Lensing analysis pipeline

The lensing pipeline is depicted in Fig. 2 (see Ref. [24] for more details). The lensed image is obtained by solving the lens equation on a squared pixel grid under the thin lens approximation. The lens equation defines a relation between the two-dimensional angular coordinates in the lens plane $\boldsymbol{\theta}_{\text{lens}}$ and the ones in the source plane $\boldsymbol{\theta}_{\text{src}}$, that is

$$\boldsymbol{\theta}_{\text{src}} = \boldsymbol{\theta}_{\text{lens}} - \boldsymbol{\alpha}\left[\Sigma\right], \tag{2.1}$$

where $\boldsymbol{\alpha}$ is the displacement field that depends on the projected surface mass density of the lens, here denoted as Σ. Since the surface brightness I (photon flux density per unit angular area) is preserved by the lens equation, the lensed image is simply given by

$$I(\boldsymbol{\theta}_{\text{src}}) = I(\boldsymbol{\theta}_{\text{lens}} - \boldsymbol{\alpha}\left[\Sigma\right]), \tag{2.2}$$

The two main ingredients are therefore the surface brightness I of the source galaxy and the mass distributions Σ along the line-of-sight, which are computed in the source model and the lens model, respectively (see Fig. 2).

The lens model is completely described by semi-analytical expressions that are based on physical models. The displacement field is computed on a squared pixel grid in the lens plane as the sum of two contributions: a main halo parameterized as a Singular Power-Law Ellipsoid (SPLE)

profile [32] and an external shear. The SPLE has six parameters: the Einstein radius r_{Ein}, the slope of the profile γ, the x and y coordinates of the center, the major-to-minor axis ratio and the position angle. The external shear contribution is defined by two additional parameters. Dark matter subhaloes and field haloes will be included and discussed in a future work.

The source model is given by the decoder of a trained Variational AutoEncoder with a latent space of normal-distributed 64 variables z (the architecture will be described in detail in Ref. [24]). The VAE is trained by means of the GREAT3 catalog that contains more than 1.5×10^4 observed galaxies [33]. The training procedure is required to estimate the parameters (the weights and the biases) of the convolutional and deconvolutional neural networks (the encoder and the decoder, respectively) and the prior distributions of the latent space. Such distributions are used to sample the latent variables z in order to reproduce realistic galaxy images through the decoder. The output galaxy is defined on a squared pixel grid equal to the one defined in the lens plane. Hence, to compute the surface brightness in any position of the source plane we consider a bilinear interpolation among adjacent pixels.

The final lensed image is obtained by applying a symmetric two-dimensional Gaussian point spread function and adding uncorrelated Gaussian noise to each pixel. Then, it is compared with the data quantifying with a score of how much the two images differ each other. The loss function depends on the inference procedure: for example, it can be the pixelwise mean squared error or the so-called evidence lower bound. So far, we have described the forward model (blue arrows in Fig.2), that is the computation of the loss function given the parameters of the source and lens models. Then, the differential probabilistic programming framework allows us to compute derivatives of the loss function with respect to all the parameters. This is defined as backward model (red arrows in Fig.2).

The analysis pipeline is defined by two steps: the optimization and the inference through a Bayesian approach. During optimization, starting from a random initial configuration, the best-fit parameters are found by minimizing the loss function. The parameters are updated by means of gradient descent. For this step, we perform a maximum a posteriori probability estimation (MAP inference) by considering delta-Dirac distributions for each parameter. As soon as the optimization converges (the loss becomes almost constant), we compute the posterior distributions of the parameters starting from the best-fit points. We use the HMC sampling technique, so taking advantage of the gradient descent algorithm once again.
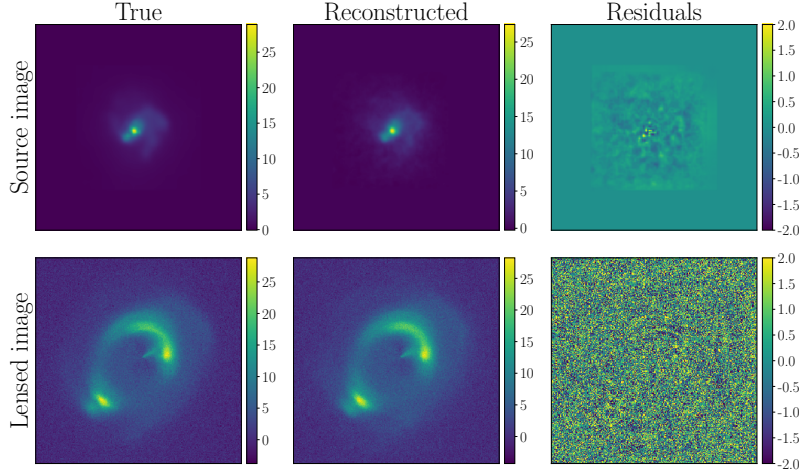
3. Source and Lensed Images Reconstruction

Here, to test our pipeline, we analyze a mock lensed image with 256×256 pixels. The mock source galaxy corresponds to a denoised galaxy of the GREAT3 data set that was not included in the VAE training. In Fig. 3, we report some parameters of the main lens halo and the loss function (last panel on the right) as a function of the number of iterations during the optimization procedure. As one can see, the convergence is achieved after much less than 2×10^4 iterations. Remarkably, the best-fit points are within 1% interval with respect to the true values, which are represented in the plots by horizontal dashed red lines.

The reconstructed source and lensed images are shown in Fig. 4 (second column) and compared with the true ones (first column). The residuals are displayed in the third column. The
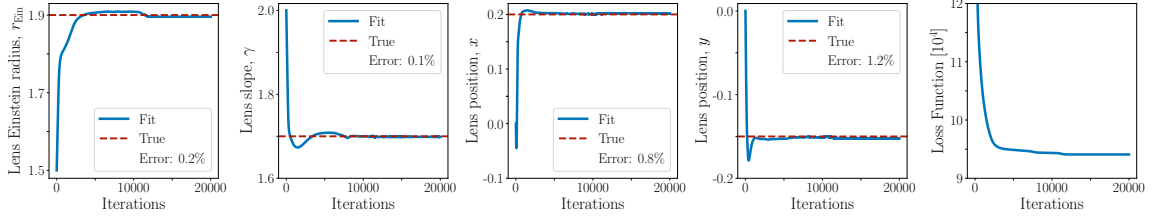
Figure 3: Optimization procedure to find the best-fit points. The main lens parameters and the loss function as a function of the number of iterations. The dashed red lines are the corresponding true parameters.

Figure 4: Reconstruction of source and lensed images after the optimization. The true source and lensed images (first column) are compared with the reconstructed ones (second column) by computing the corresponding residuals (third column).

reconstructed lensed image matches the original one at the noise level as highlighted by the residuals. On the other hand, the residuals in the source plane are higher, though the two source galaxies are qualitatively in agreement. For instance, the spiral arms are captured in the reconstruction.

Finally, in Fig. 5 we report the posteriors of the main lens parameters obtained with the HMC starting from the best-fit points. These distributions are marginalized with respect the other lens and source parameters. The ellipses show the contours containing 68% and 95% of the 500 points (grey dots) provided by the HMC. The black squares represent the true parameters.

4. Conclusions

In this paper, we have briefly described the main ingredients of a new analysis pipeline for strong gravitational lensing data. The novelties of our approach are:

- the use of a Variational AutoEncoder to generate the background source galaxy;

- the combination of deep learning techniques with physical models to increase the interpretability of the results and to make possible further generalizations of the lens model (for example, the inclusion of dark matter substructures);
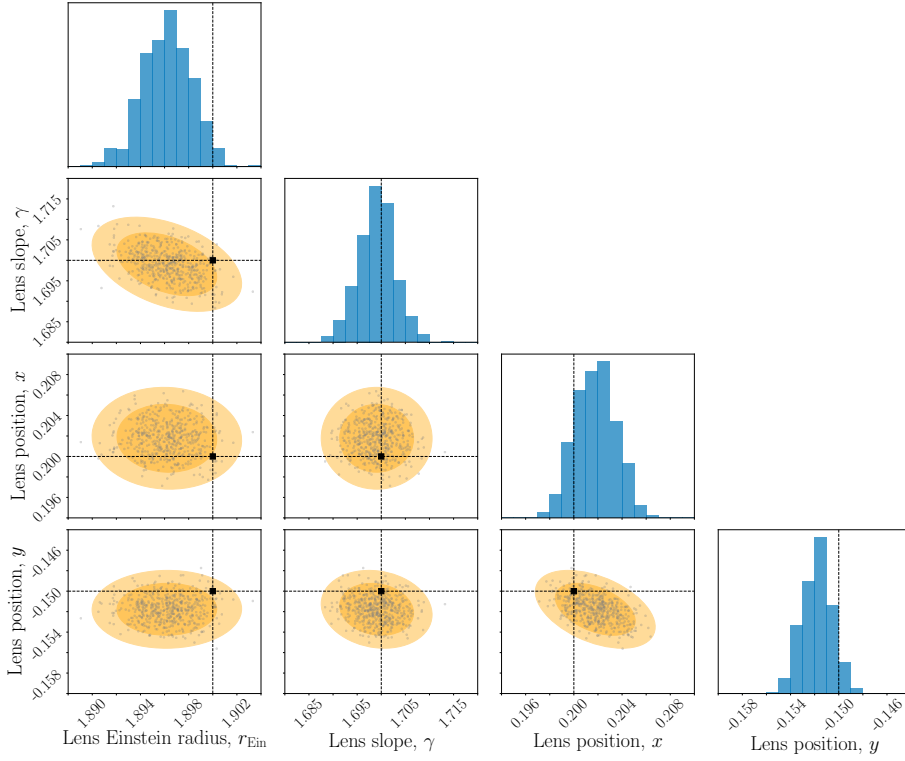
Figure 5: Bayesian inference with Hamiltonian Monte Carlo. The marginalized posterior distributions of the main lens parameters obtained by sampling 500 points with Hamiltonian Monte Carlo. The ellipses represents the contours at 68% and 95% levels. The true parameters are highlighted by the black squares.

- the differential probabilistic programming framework to automatically perform the Bayesian inference with gradient descent techniques.

The preliminary analyses of mock data are promising. In particular, we are able to infer the true parameters of the lens model at high accuracy (typically below the 1% level). Moreover, the true parameters reside within the 95% contours of the posterior distributions obtained by means of Hamiltonian Monte Carlo. However, further improvements are required on the VAE source model in order to reconstruct more complicated galaxies and reduce the residuals in the source plane. This will be discussed in detail in Ref. [24], where a more systematic study of the pipeline performance is also presented.

References

[1] G. Bertone, D. Hooper and J. Silk, Phys. Rept. **405** (2005) 279 [hep-ph/0404175].

[2] G. Bertone and T. Tait, M.P., Nature **562** (2018) no.7725, 51 [arXiv:1810.01668].

[3] M. Kuhlen, M. Vogelsberger and R. Angulo, Phys. Dark Univ. **1** (2012) 50 doi:10.1016/j.dark.2012.10.002 [arXiv:1209.5745].

[4] M. R. Lovell, C. S. Frenk, V. R. Eke, A. Jenkins, L. Gao and T. Theuns, Mon. Not. Roy. Astron. Soc. **439** (2014) 300 [arXiv:1308.1399].

[5] A. Schneider, R. E. Smith, A. V. Maccio and B. Moore, Mon. Not. Roy. Astron. Soc. **424** (2012) 684 [arXiv:1112.0330].

[6] M. Viel, G. D. Becker, J. S. Bolton and M. G. Haehnelt, Phys. Rev. D **88** (2013) 043502 [arXiv:1306.2314].

[7] C. Giocoli, G. Tormen, R. K. Sheth and F. C. van den Bosch, Mon. Not. Roy. Astron. Soc. **404** (2010) 502 [arXiv:0911.0436].

[8] E. Polisensky and M. Ricotti, Phys. Rev. D **83** (2011) 043506 [arXiv:1004.1459].

[9] A. Fitts *et al.*, Mon. Not. Roy. Astron. Soc. **471** (2017) no.3, 3547 [arXiv:1611.02281].

[10] J. I. Read, G. Iorio, O. Agertz and F. Fraternali, Mon. Not. Roy. Astron. Soc. **467** (2017) no.2, 2019 [arXiv:1607.03127].

[11] S. Vegetti, D. J. Lagattuta, J. P. McKean, M. W. Auger, C. D. Fassnacht and L. V. E. Koopmans, Nature **481** (2012) 341 [arXiv:1201.3643].

[12] S. Vegetti, L. V. E. Koopmans, M. W. Auger, T. Treu and A. S. Bolton, Mon. Not. Roy. Astron. Soc. **442** (2014) no.3, 2017 [arXiv:1405.3666].

[13] G. Despali, S. Vegetti, S. D. M. White, C. Giocoli and F. C. van den Bosch, Mon. Not. Roy. Astron. Soc. **475** (2018) no.4, 5424 [arXiv:1710.05029].

[14] S. Vegetti, G. Despali, M. R. Lovell and W. Enzi, Mon. Not. Roy. Astron. Soc. **481** (2018) no.3, 3661 [arXiv:1801.01505].

[15] E. Ritondale, S. Vegetti, G. Despali, M. W. Auger, L. V. E. Koopmans and J. P. McKean, Mon. Not. Roy. Astron. Soc. **485** (2019) no.2, 2179 [arXiv:1811.03627].

[16] T. Treu, Ann. Rev. Astron. Astrophys. **48** (2010) 87 [arXiv:1003.5567].

[17] L. V. E. Koopmans, Mon. Not. Roy. Astron. Soc. **363** (2005) 1136 [astro-ph/0501324].

[18] S. Vegetti and L. V. E. Koopmans, Mon. Not. Roy. Astron. Soc. **392** (2009) 945 [arXiv:0805.0201].

[19] Petrillo C. E. *et al.*, Mon. Not. Roy. Astron. Soc. **472** (2017) 1129.

[20] Petrillo C. E. *et al.*, Mon. Not. Roy. Astron. Soc. **482** (2019) 807.

[21] Y. D. Hezaveh, L. Perreault Levasseur and P. J. Marshall, Nature **548** (2017) 555 [arXiv:1708.08842].

[22] L. Perreault Levasseur, Y. D. Hezaveh and R. H. Wechsler, Astrophys. J. **850** (2017) no.1, L7 [arXiv:1708.08843].

[23] W. R. Morningstar *et al.*, arXiv:1901.01359.

[24] M. Chianese, A. Coogan, P. Hofma, S. Otten and C. Weniger, in preparation (2019).

[25] D. P. Kingma and M. Welling, arXiv:1312.6114.

[26] D. Jimenez Rezende , S. Mohamed and D. Wierstra, arXiv:1401.4082.

[27] A. Radford, L. Metz and S. Chintala, arXiv:1511.06434.

[28] A. Paszke *et al.*, NIPS 2017 Workshop.

[29] E. Bingham *et al.*, Journal of Machine Learning Research (2018).

[30] S. Duane, A. D. Kennedy, B. J. Pendleton and D. Roweth, Phys. Lett. B **195** (1987) 216.

[31] Betancourt, M. arXiv:1701.02434.

[32] A. Kassiola and I. Kovner, Astrophys. J. **417** (1993) 450.

[33] R. Mandelbaum *et al.*, Astrophys. J. Suppl. **212** (2014) 5 [arXiv:1308.4982].