


RESEARCH

Open Access



# CosmoGAN: creating high-fidelity weak lensing convergence maps using Generative Adversarial Networks

Mustafa Mustafa<sup>1\*</sup> , Deborah Bard<sup>1</sup>, Wahid Bhimji<sup>1</sup>, Zarija Lukić<sup>1</sup>, Rami Al-Rfou<sup>2</sup> and Jan M. Kratochvil<sup>3</sup>

## Abstract

Inferring model parameters from experimental data is a grand challenge in many sciences, including cosmology. This often relies critically on high fidelity numerical simulations, which are prohibitively computationally expensive. The application of deep learning techniques to generative modeling is renewing interest in using high dimensional density estimators as computationally inexpensive emulators of fully-fledged simulations. These generative models have the potential to make a dramatic shift in the field of scientific simulations, but for that shift to happen we need to study the performance of such generators in the precision regime needed for science applications. To this end, in this work we apply Generative Adversarial Networks to the problem of generating weak lensing convergence maps. We show that our generator network produces maps that are described by, with high statistical confidence, the same summary statistics as the fully simulated maps.

**Keywords:** Weak lensing convergence maps; Generative models; Generative Adversarial Networks; Deep learning; Machine learning

## 1 Introduction

Cosmology has progressed towards a precision science in the past two decades, moving from order of magnitude estimates to percent-level measurements of fundamental cosmological parameters. This was largely driven by successful CMB and BAO probes (Planck Collaboration et al. 2018; Alam et al. 2017; Bautista et al. 2018; du Mas des Bourboux et al. 2017), which extract information from very large scales, where structure formation is well described by linear theory or by the perturbation theory (Carlson et al. 2009). In order to resolve the next set of outstanding problems in cosmology—for example the nature of dark matter and dark energy, the total mass and number of neutrino species, and primordial fluctuations seeded by inflation—cosmologists will have to rely on measurements of cosmic structure at far smaller scales.

Modeling the growth of structure at those scales involves non-linear physics that cannot be accurately described an-

alytically, and we instead use numerical simulations to produce theoretical predictions which are to be confronted with observations. Weak gravitational lensing is considered to be one of the most powerful tools of extracting information from small scales. It probes both the geometry of the Universe and the growth of structure (Bartelmann and Schneider 2001). The shearing and magnification of background luminous sources by gravitational lensing allows us to reconstruct the matter distribution along the line of sight. Common characterizations for gravitational lensing shear involve cross-correlating the ellipticities of galaxies in a two-point function estimator, giving the lensing power spectrum. By comparing such measurements to theoretical predictions, we can distinguish whether, for example, cosmic acceleration is caused by dark energy or a modification of general relativity.

The scientific success of the next generation of photometric sky surveys (e.g. Laureijs et al. 2011, LSST Dark Energy Science Collaboration 2012, Spergel et al. 2015) therefore hinges critically on the success of underlying simulations. Currently the creation of each simulated virtual universe requires an extremely computationally expensive

\*Correspondence: [mmustafa@lbl.gov](mailto:mmustafa@lbl.gov)

<sup>1</sup>Lawrence Berkeley National Laboratory, Berkeley, USA

Full list of author information is available at the end of the article

simulation on High Performance Computing (HPC) resources. That makes direct application of Markov chain Monte Carlo (MCMC, Metropolis et al. 1953, Gelman et al. 2013) or similar Bayesian methods prohibitively expensive, as they require hundreds of thousands of forward model evaluations to determine the posterior probabilities of model parameters. In order to make this inverse problem practically solvable, constructing a computationally cheap *surrogate model* or an *emulator* (Heitmann et al. 2009; Lawrence et al. 2010) is imperative. However, traditional approaches to emulators require the use of the summary statistic which is to be emulated. An approach that makes no assumptions about such mathematical templates of the simulation outcome would be of considerable value. While in this work we focus our attention on the generation of the weak lensing convergence maps, we believe that the method presented here is relevant to many similar problems in astrophysics and cosmology where a large number of expensive simulations is necessary.

Recent developments in deep generative modeling techniques open the potential to meet this emulation need. The density estimators in these models are built out of neural networks which can serve as universal approximators (Csáji 2001), thus having the ability to learn the underlying distributions of data and emulate the observable without imposing the choice of summary statistics, as in the traditional approach to emulators. These data-driven generative models have also found applications in astrophysics and cosmology. On the observational side, these models can be used to improve images of galaxies beyond the deconvolution limit of the telescope (Schawinski et al. 2017). On the simulation side, they have been used to produce samples of cosmic web (Rodríguez et al. 2018).

In addition to generative modeling, we have recently witnessed a number of different applications of deep learning and convolutional neural networks (CNN) to the problem of parameter inference using weak gravitational lensing. Training deep neural networks to regress parameters from data simulations (Gupta et al. 2018; Ribli et al. 2019) shows that CNN trained on noiseless lensing maps can bring few times tighter constraints on cosmological parameters than the power spectrum or lensing peaks analysis methods. Similar analysis has been performed for the case of simulated lensing maps with varying level of noise due to imperfect measurement of galaxy shape distortions and finite number density of the source galaxies (Fluri et al. 2018). Although the advantage of CNNs over “traditional” methods is still present in the noisy case, the quantitative improvement is much less than in the case of noiseless simulated maps. CNNs applied to weak lensing maps have also been proposed for improved differentiating between dark energy and modified gravity cosmologies (Peel et al. 2018). Finally, an image-to-image translation method employing conditional adversarial networks was used to demonstrate

learning the mapping from an input, simulated noisy lensing map to the underlying noise field (Shirasaki et al. 2018), effectively denoising the map.

We would like to caution that all deep learning methods are very sensitive to the input noise, thus networks trained on either noiseless or even on simulated, idealized, noise are likely to perform poorly in the inference regime on the real data without a complex and domain-dependent preparation for such task (Wang and Deng 2018). This problem is exaggerated by the fact that trained CNNs are mapping input lensing map to cosmological parameters without providing full posterior probability distribution for these parameters or any other good description of inference errors.

In this work, we study the ability of a variant of generative models, Generative Adversarial Networks (GANs) (Goodfellow et al. 2014) to generate weak lensing convergence maps. In this paper, we are not concerned about inferring cosmological parameters, and we do not attempt to answer questions of optimal parameter estimators, either with deep learning or “traditional” statistical methods. Here, we study the ability of GANs to produce convergence maps that are statistically indistinguishable from maps produced by physics-based generative models, which are in this case N-body simulations. The training and validation maps are produced using N-body simulations of  $\Lambda$ CDM cosmology. We show that maps generated by the neural network exhibit, with high statistical confidence, the same power spectrum of the fully-fledged simulator maps, as well as higher order non-Gaussian statistics, thus demonstrating that such scientific data can be amenable to a GAN treatment for generation. The very high level of agreement achieved offers an important step towards building emulators out of deep neural networks.

The paper is organized as follows: in Sect. 2 we outline the data set used and describe our GAN architecture. We present our results in Sect. 3 and we outline the future investigations which we think are critical to build weak lensing emulators in Sect. 4. Finally, we present conclusions of this paper in Sect. 5.

## 2 Methods

In this section we first introduce the weak lensing convergence maps used in this work, and then proceed to describe Generative Adversarial Networks and their implementation in this work.

### 2.1 Dataset

To produce our training dataset, we use the cosmological simulations described in Kratochvil et al. (2012), Yang et al. (2011), produced using the Gadget2 (Springel 2005) N-Body simulation code and ray-traced with the Inspector Gadget weak lensing simulation pipeline (Kratochvil et al. 2012; Kratochvil et al. 2010; Yang et al. 2011) to produce weak lensing shear and convergence maps. A total

of 45 simulations were produced, each consisting of  $512^3$  particles in a box of size of  $240 h^{-1} \text{Mpc}$ . The cosmological parameters used in these simulations were  $\sigma_8 = 0.798$ ,  $w = -1.0$ ,  $\Omega_m = 0.26$ ,  $\Omega_\Lambda = 0.74$ ,  $n_s = 0.96$ ,  $H_0 = 0.72$ . These simulation boxes were rotated and translated to produce 1000 ray-traced lensing maps at the redshift plane  $z = 1.0$ . Each map covers 12 square degrees, with  $2048 \times 2048$  pixels, which we downsampled to  $1024 \times 1024$  pixels. Following the formalism introduced in Bernstein and Jarvis (2002), the gravitational lensing illustrated by these maps can be described by the Jacobian matrix

$$A(\theta) = \begin{bmatrix} 1 - \kappa - \gamma_1 & -\gamma_2 \\ -\gamma_2 & 1 - \kappa + \gamma_1 \end{bmatrix}, \quad (1)$$

where  $\kappa$  is the convergence, and  $\gamma$  is the shear.

The training data was generated by randomly cropping 200  $256 \times 256$  maps from each of the original 1000 maps. The validation and development datasets have been randomly cropped in the same manner. We have tested sampling the validation datasets from *all of* the original 1000 maps versus sampling the training from 800 maps and the validation from the remaining 200 maps. GANs are not generally amenable to memorization of the training dataset, this is in part because the generator network isn't trained directly on that data; it only learns about it by means of information from another network, the discriminator, as is described in the next section. Therefore, in our studies, it did not make any difference how we sample our validation dataset. We demonstrate that the generator network did not memorize the training dataset in Sect. 3.2. Finally, the probability for a map to have a single pixel value outside  $[-1.0, 1.0]$  range is less than 0.9% so it was safe to use the data without any normalization.

In one of the tests we report in this paper we use an auxiliary dataset, which consists of 1000 maps produced using the same simulation code and cosmological parameters, but with a different random seed, resulting in a set of convergence maps that are statistically independent from those used in our training and validation.

## 2.2 Generative Adversarial Networks

The central problem of generative models is the question: given a distribution of data  $\mathbb{P}_r$ , can one devise a generator  $G$  such that the distribution of model generated data  $\mathbb{P}_g = \mathbb{P}_r$ ? Our information about  $\mathbb{P}_r$  comes from the training dataset, typically an independent and identically distributed random sample  $x_1, x_2, \dots, x_n$  which is assumed to have the same distribution as  $\mathbb{P}_r$ . Essentially, a generative model aims to construct a density estimator of the dataset. The GAN frameworks constructs an implicit density estimator which can be efficiently sampled to generate samples of  $\mathbb{P}_g$ .

The GAN framework (Goodfellow et al. 2014) sets up a game between two players, a generator and a discriminator. The generator is trained to generate samples that aim to be indistinguishable from training data as judged by a competent discriminator. The discriminator is trained to judge whether a sample looks real or fake. Essentially, the generator tries to fool the discriminator into judging a generated map looks real.

In the neural network formulation of this framework the generator network  $G_\phi$ , parametrized by network parameters  $\phi$ , and discriminator network  $D_\theta$ , parametrized by  $\theta$ , are simultaneously optimized using gradient descent. The discriminator is trained in a supervised manner by showing it real and generated samples, it outputs a probability of the input map being real or fake. It is trained to minimize the following cross-entropy cost function:

$$J^{(D)} = -\mathbb{E}_{x \sim \mathbb{P}_r} \log D_\theta(x) - \mathbb{E}_{x \sim \mathbb{P}_g} \log (1 - D_\theta(x)). \quad (2)$$

The generator is a differentiable function (except at possibly finitely many points) that maps a sample from a noise prior,  $\mathbf{z} \sim p(\mathbf{z})$ , to the support of  $\mathbb{P}_g$ . For example, in this work, the noise vector is sampled from a 64-dimensional isotropic normal distribution and the output of the generator are maps  $x \in \mathbb{R}^{256 \times 256}$ . The dimension of the noise vector  $\mathbf{z}$  needs to be commensurate with the support of the convergence maps  $\mathbb{P}_r$  in  $\mathbb{R}^{256 \times 256}$ . In the game-theoretic formulation, the generator is trained to maximize equation (2), this is known as the minimax game. However, in that formulation, the gradients of the cost function with respect to the generator parameters vanish when the discriminator is winning the game, i.e. rejecting the fake samples confidently. Losing the gradient signal makes it difficult to train the generator using gradient descent. The original GAN paper (Goodfellow et al. 2014) proposes flipping the target for the generator instead:

$$J^{(G)} = -\mathbb{E}_{x \sim \mathbb{P}_g} \log D_\theta(x). \quad (3)$$

This “heuristically” motivated cost function (also known as the non-saturating game) provides strong gradients to train the generator, especially when the generator is losing the game (Goodfellow 2016).

Since the inception of the first GAN, there have been many proposals for other cost functions and functional constraints on the discriminator and generator networks. We have experimented with some of these but in common with a recent large scale empirical study (Lučić et al. 2018) of these different models: “did not find evidence that any of the tested algorithms consistently outperforms the non-saturating GAN”, introduced in Goodfellow et al. (2014) and outlined above. That study attributes the improvements in performance reported in recent literature to difference in computational budget. With the GAN frame-

**Table 1** Generator network architecture: layer types, activations, output shapes (channels  $\times$  height  $\times$  width) and number of trainable parameters for each layer. TransposedConv have strides = 2

	Activ.	Output shape	Params.
Latent	–	64	–
Dense	–	$512 \times 16 \times 16$	8.5M
BatchNorm	ReLU	$512 \times 16 \times 16$	1024
TransposedConv $5 \times 5$	–	$256 \times 32 \times 32$	3.3M
BatchNorm	ReLU	$256 \times 32 \times 32$	512
TransposedConv $5 \times 5$	–	$128 \times 64 \times 64$	819K
BatchNorm	ReLU	$128 \times 64 \times 64$	256
TransposedConv $5 \times 5$	–	$64 \times 128 \times 128$	205K
BatchNorm	ReLU	$64 \times 128 \times 128$	128
TransposedConv $5 \times 5$	Tanh	$1 \times 256 \times 256$	1601
Total trainable parameters			<b>12.3M</b>

work laid out we move to describing the generator and discriminator networks and their training procedure.

### 2.3 Network architecture and training

Given the intrinsic translation invariance of the convergence maps (Kilbinger 2015), it is natural to construct the generator and discriminator networks mainly from convolutional layers. To allow the network to learn the proper correlations on the components of the input noise  $\mathbf{z}$  early on, the first layer of the generator network needs to be a fully-connected layer. A class of all convolutional network architectures has been developed in Springenberg et al. (2014), which use strided convolutions to downsample instead of pooling layers, and also use strided transposed convolutions to upsample. This architecture was later adapted to build GANs in Deep Convolutional Generative Adversarial Networks (DCGAN) (Radford et al. 2015). We experimented with DCGAN architectural parameters and we found that most of the hyperparameters optimized for natural images by the original authors perform well on the convergence maps. We used DCGAN architecture with slight modifications to meet our problem dimensions, we also halved the number of filters.

The generator takes a 64-dimensional vector sampled from a normal prior  $z \sim \mathcal{N}(0, 1)$ . The first layer is a fully connected layer whose output is reshaped into a stack of feature maps. The rest of the network consists of four layers of transposed convolutions (a convolutional layer with fractional strides where zeroes are inserted between each column and row of pixels before convolving the image with the filter in order to effectively up-sample the image) that lead to a single channel  $256 \times 256$  image. The outputs of all layers, except the output one, are batch-normalized (Ioffe and Szegedy 2015) (by subtracting the mean activation and dividing by its standard deviation and learning a linear scaling) which was found to stabilize the training. A rectified linear unit (ReLU) activation (Nair and Hinton 2010) (output zero when the input less than zero and output

equal to the input otherwise) is used for all except the output layer where a hyperbolic-tangent (tanh) is used. The final generator network architecture is summarized in Table 1.

The discriminator has four convolutional layers. The number of feature maps, stride and kernel sizes are the same as in the generator. We reduced the number of filters from the DCGAN guidelines by a factor of 2, which effectively reduces the capacity of the generator/discriminator. This worked well for our problem and reduced the training time. The output of all convolutional layers are activated with LeakyReLU (Maas et al. 2013) with parameter  $\alpha = 0.2$ . The output of the last convolutional layer is flattened and fed into a fully connected layer with a 1-dimensional output that is fed into a sigmoid. Batch Normalization is applied before activations for all layers' outputs except the first layer. The final discriminator network architecture is summarized in Table 2.

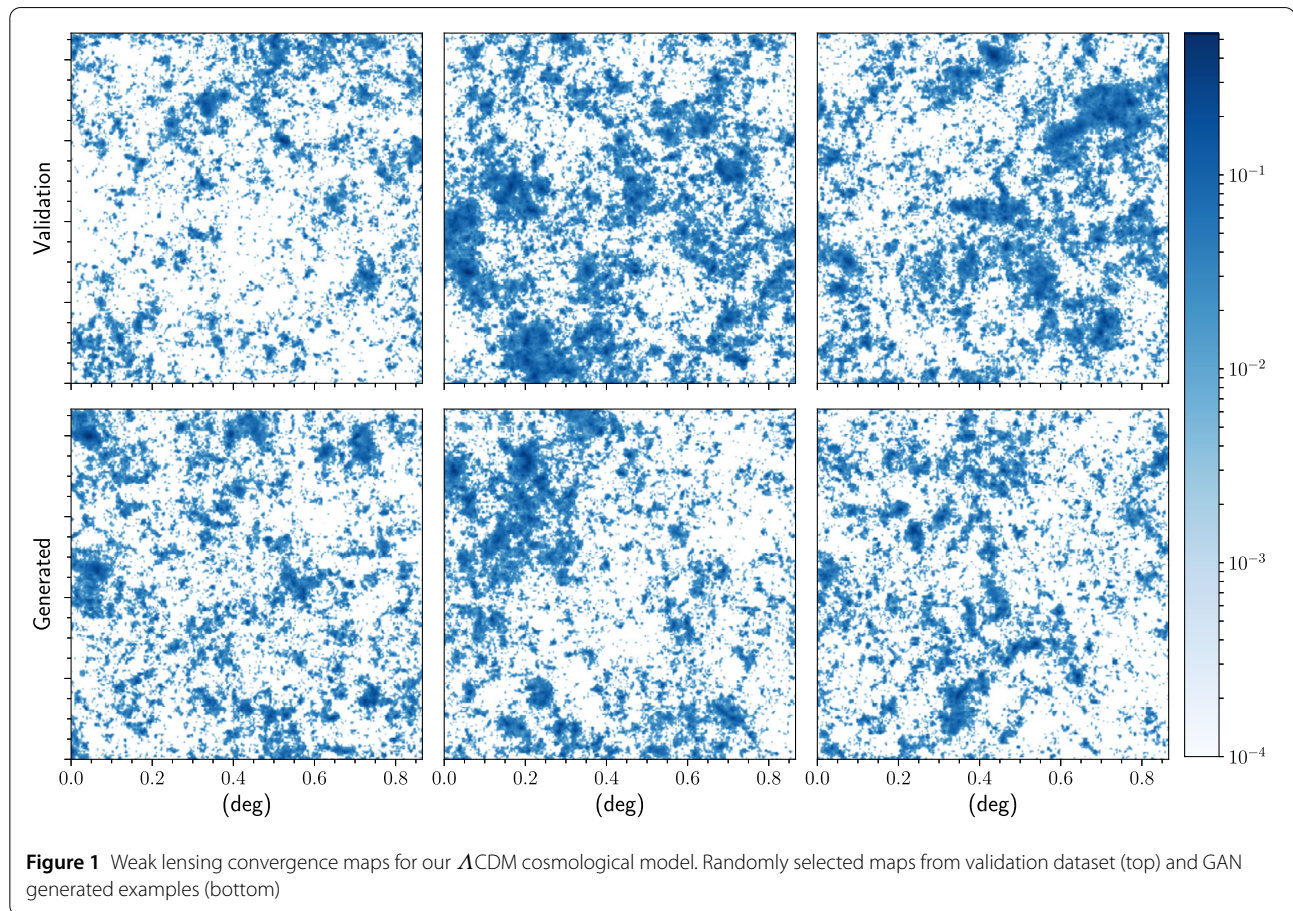
Finally, we minimize discriminator loss in Eq. (2) and generator loss in Eq. (3) using Adam optimizers (Kingma et al. 2014) with the parameters suggested in the DCGAN paper: learning rate 0.0002 and  $\beta_1 = 0.5$ . Batch-size is 64 maps. We flip the real and fake labels with 1% probability to avoid the discriminator overpowering the generator too early into the training. We implement the networks in TensorFlow (Abadi et al. 2015) and train them on a single NVIDIA Titan X Pascal GPU.

Training GANs using a heuristic loss function often suffers from unstable updates towards the end of their training. This has been analyzed theoretically in Arjovsky and Bottou (2017) and shown to happen when the discriminator is close to optimality but has not yet converged. In other words, the precision of the generator at the point we stop the training is completely arbitrary and the loss function is not useful to use for early stopping. For the results shown in this work we trained the network until the generated maps pass the visual and pixel intensity Kolmogorov–Smirnov tests, see Sect. 3. This took 45 passes (epochs) over all of the training data. Given the un-stability of the



**Table 2** Discriminator network architecture: layer types, activations, output shapes (channels  $\times$  height  $\times$  width) and number of trainable parameters for each layer. All convolutional layers have stride = 2. LeakyReLU's leakiness = 0.2

	Activ.	Output shape	Params.
Input map	–	$1 \times 256 \times 256$	–
Conv $5 \times 5$	LReLU	$64 \times 128 \times 128$	1664
Conv $5 \times 5$	–	$128 \times 64 \times 64$	205K
BatchNorm	LReLU	$128 \times 64 \times 64$	256
Conv $5 \times 5$	–	$256 \times 32 \times 32$	819K
BatchNorm	LReLU	$256 \times 32 \times 32$	512
Conv $5 \times 5$	–	$512 \times 16 \times 16$	3.3M
BatchNorm	LReLU	$512 \times 16 \times 16$	1024
Linear	Sigmoid	1	131K
Total trainable parameters			<b>4.4M</b>

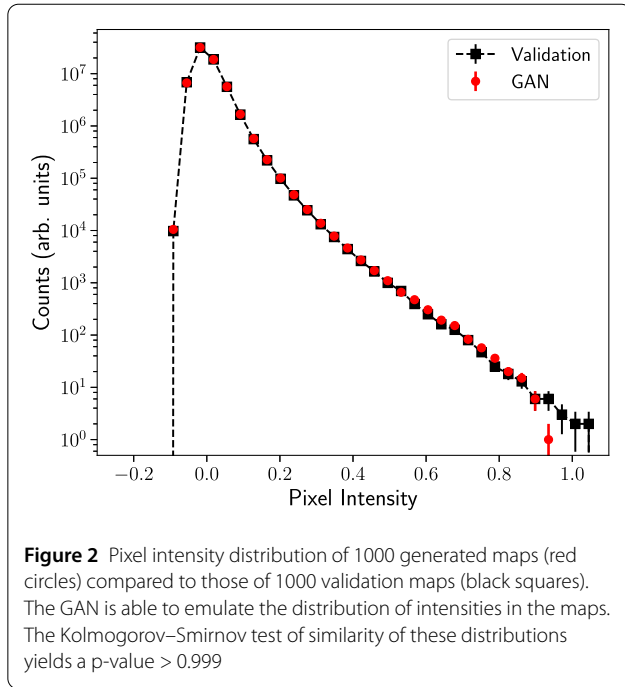


updates at this point, the performance of the generator on the summary statistics tests starts varying uncontrollably. Therefore, to choose a generator configuration, we train the networks for two extra epochs after epoch 45, and randomly generate 100 batches (6400 maps) of samples at every single training step. We evaluate the power spectrum on the generated samples and calculate the Chi-squared distance measurement to the power spectrum histograms evaluated on a development subset of the vali-

dation dataset. We use the generator with the best Chi-squared distance.

### 3 Results

Figure 1 shows examples of maps from the validation and GAN generated datasets. The validation dataset has not been used in the training or tuning of the networks. Visually, an expert cannot distinguish the generated maps from the full simulation ones.



### 3.1 Evaluation of generator fidelity

Once we have obtained a density estimator of the original data the first practical question is to determine the goodness of the fit. Basically, *how close is  $\mathbb{P}_g$  to  $\mathbb{P}_r$* ? This issue is critical to understanding and improving the formulation and training procedure for generative models, and is an active area of research (Theis et al. 2015). Significant insight into the training dynamics of GANs from a theoretical point of view has been gained in Arjovsky and Bottou (2017), Arjovsky et al. (2017) and later works. We think that when it comes to practical applications of generative models, such as in the case of emulating scientific data, the way to evaluate generative models is to study their ability to reproduce the characteristic statistic of the original dataset.

To this end, we calculate three statistics on the generated convergence maps: a first order statistic (pixel intensity), the power spectrum and a non-Gaussian statistic. The ability to reproduce such summary statistics is a reliable metric to evaluate generative models from an information encoding point of view. To test our statistical confidence of the matching of the summary statistics we perform bootstrapped two-tailed Kolmogorov–Smirnov (KS) test and Andersen–Darling (AD) test of the null-hypothesis that the summary statistics in the generated maps and the validation maps have been drawn from the same continuous distributions.<sup>a</sup>

Figure 2 shows a histogram of the distribution of pixel intensities of an ensemble of generated maps compared to that of a validation dataset. It is clear that the GAN generator has been able to learn the probabilities of pixel intensi-

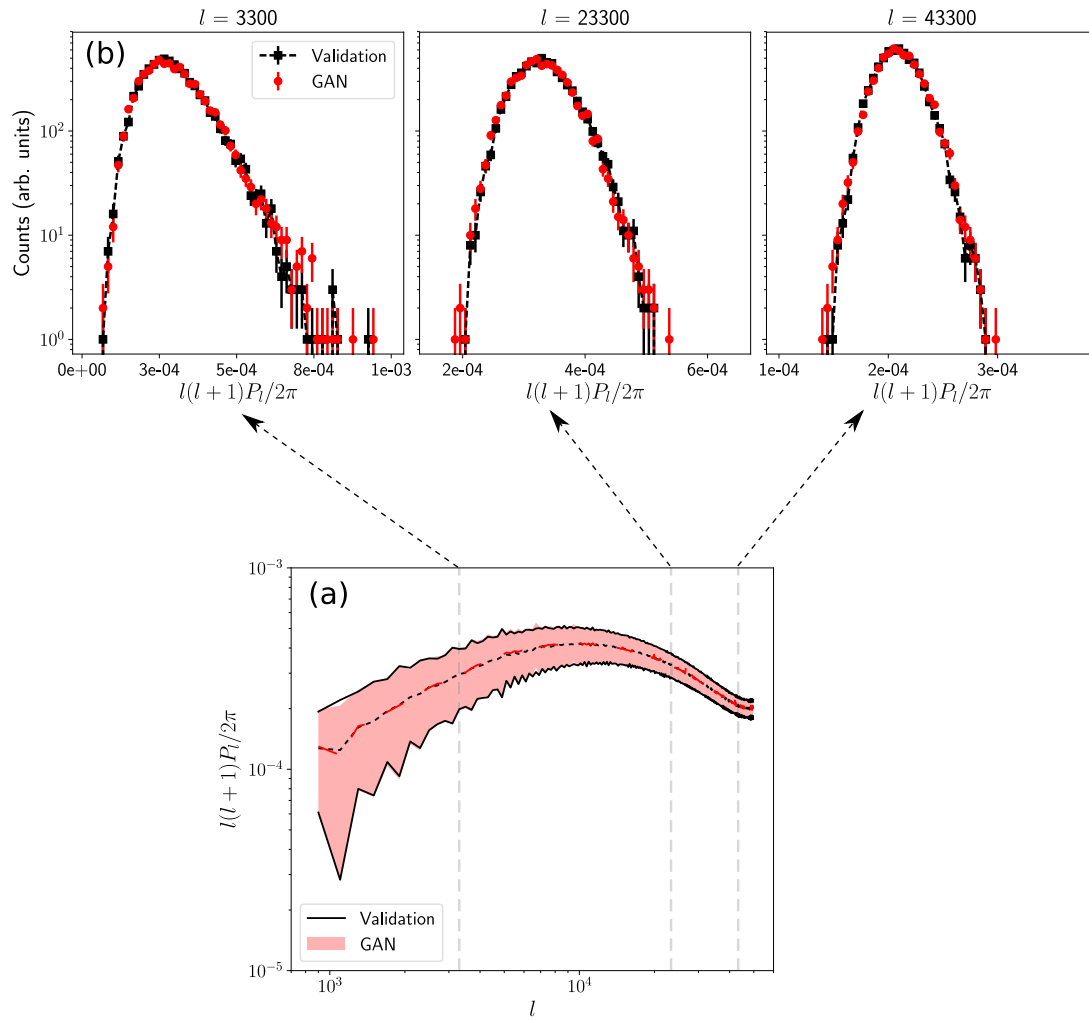
ties in the real simulation dataset, the KS p-value is  $> 0.999$ . We note that the maps generated by this GAN have the same the geometry of the simulated maps, i.e. angular size, resolution, etc.

The second-order statistical moment in gravitational lensing is the shear power spectrum. This is a measure of the correlation in gravitational lensing at different distances, and characterizes the matter density fluctuations at different length scales. Assuming we have only Gaussian fluctuations  $\delta(x)$  at comoving distance  $x$ , the matter density of the universe can be defined by its power spectrum  $P_\kappa$ :


$$\langle \tilde{\delta}(l) \tilde{\delta}^*(l') \rangle = (2\pi)^2 \delta_D(l - l') P_\kappa(l), \quad (4)$$

where  $\delta_D$  is the Dirac delta function, and  $l$  is the Fourier mode (Kilbinger 2015). The power spectrum (and its corresponding real-space measurement, the two-point correlation function) of convergence maps has long been used to constrain models of cosmology by comparing simulated maps to real data from sky surveys (Liu et al. 2015; Abbott et al. 2016; Jee et al. 2013). Numerically, the power spectrum is the amplitudes of the Fourier components of the map. We calculate the power spectrum at 248 Fourier modes of an ensemble of generated maps using LensTools (Petri 2016), and compare them to the validation dataset. Since each map is a different random realization, the comparison has to be made at the ensemble level. Figure 3(a) shows two bands representing the mean  $\mu(l) \pm$  standard deviation  $\sigma(l)$  of the ensemble of power spectra at each Fourier mode of the validation and a generated dataset. As is clear in the figure, the bands completely overlap with each other. To confirm that the range of the power spectrum at a given  $l$  is completely reproduced in the generated maps we look differentially at the underlying 1-D distributions. Samples of such distributions at equally spaced Fourier modes are shown in Fig. 3(b). The bootstrapped KS (AD) p-values for 236 (205) Fourier modes are  $> 0.99$ , of the remaining modes, 10 (26) are  $> 0.95$ , all the remaining are larger than 0.8. The power spectrum is the figure of merit for evaluating the success of an emulator for reproducing the structures of the convergence map, and we have shown with statistical confidence that GANs are able to discover and reproduce such structures. It should be noted that the power spectra of the generated maps is limited to the scale resolutions in the training dataset. For example in Fig. 3(a), one can see the statistical fluctuations at the lower modes in the validation dataset, the generator faithfully reproduces those fluctuations.

In Fig. 4 we show the correlation matrices of the power spectra shown in Fig. 3, and the difference between the correlation matrices from the validation maps and those produced by our GAN. We also show a comparison of the correlation matrices from the validation dataset and an the




**Figure 3** The power spectrum of the convergence maps, evaluated at 248 Fourier modes. We use 100 batches of generated maps (6400 in total) for this comparison. Shown in (a) are bands of the  $\mu(l) \pm \sigma(l)$ , the dashed lines represent the means  $\mu(l)$ , (b) shows the underlying distributions at 3 equidistant modes for illustration



**Figure 4** The correlation matrices of the modes of the power spectra shown in Fig. 3. The first two panels show the correlation matrices of the validation and the GAN generated maps, the third panel shows the difference between these correlation matrices. To provide a scale for comparison, in the fourth panel we also show the difference between the validation dataset correlation matrix and the correlations in an auxiliary dataset (see text for details)





**Figure 5** The Minkowski functionals of the convergence maps, which are sensitive to the non-Gaussian structures. We carried out the measurements on 100 batches of generated maps (6400 in total) and compare them to those of the validation maps. The functionals are evaluated at 19 thresholds and shown here are the bands of  $\mu \pm \sigma$  at each threshold. The dashed lines represent the mean  $\mu$

auxiliary dataset (a statistically independent dataset generated using the procedure used for the original dataset but with a different random seed, see Sect. 2.1).

There are differences between the validation and auxiliary correlations matrices of up to 7%. The differences between the validation and generated correlation matrices are of a similar level (up to 10%) but show a pattern of slightly higher correlation between high- and low- $l$  modes. We find it interesting that the GAN generator assumes a slightly stronger correlation between small and large modes, thus a smoother power spectra, than in the simulations. This is due to the large uncertainty of the power spectra at small modes in the original simulations as seen in Fig. 3(a).

The power spectrum only captures the Gaussian components of matter structures in the universe. However, gravity produces non-Gaussian structures at small scales which are not described by the two-point correlation function (Kilbinger 2015). There are many ways to access this non-Gaussian statistical information, including higher-order correlation functions, and topological features such as Minkowski functionals  $V_0$ ,  $V_1$ , and  $V_2$  (Mecke et al. 1994; Petri et al. 2015; Kratochvil et al. 2012) which characterize the area, perimeter and genus of the convergence maps. We investigate the ability of the networks to discover and reproduce these non-Gaussian structures in the maps by evaluating the three Minkowski functionals. Figure 5 compares bands of  $\mu \pm \sigma$  of the three Minkowski Functionals (calculated using (Petri 2016)) to those in the real dataset. Each functional is evaluated at 19 thresholds. As we did with the power spectrum, we show the Minkowski functionals at different thresholds in Fig. 6. The bootstrapped KS (AD) for 40 (32) threshold histograms are  $> 0.95$  p-values 7 (6) are  $> 0.9$  all the remaining are  $> 0.6$ . Again, it is clear that the GANs have successfully reproduced those non-Gaussian structures in the maps.

Finally, we measure the approximate speed of generation; it takes the generator  $\approx 15$  s to produce 1024 images on a single NVIDIA Titan X Pascal GPU. Training the network takes  $\approx 4$  days on the same GPU. Running N-Body simulations of a similar scale used in this paper requires roughly 5000 CPU hours per box, with an additional 100-500 CPU hours to produce the planes and do the ray-tracing to make the 1000 maps per box. While this demonstrates that, as expected, GANs offer substantially faster generation, it should be cautioned that these comparisons are not “apple-to-apple” as the N-Body simulations are ab-initio, while the GANs generated maps are by means of sampling a fit to the data.


### 3.2 Is the generator creating novel maps?

The success of the generator at creating maps that matches the summary statistics of the validation maps raises the question of whether the generator is simply memorizing the training dataset. Note that in the GAN framework, the generator never sees the training dataset directly. The question is if the generator learned about the training dataset through the gradients it gets from the discriminator. We conduct two tests to investigate this possibility.

In the first test we generate a set of maps and find the training maps that are closest to them. We test two distance notions, the Euclidean distance ( $L^2$ ) in the pixel space of the maps and in the power spectrum space. Some examples of the latter are shown in Fig. 7. Using both metrics, the training data points that are closest to generated maps are visually distinct. Concluding that the generator is producing maps that are not present in the training dataset.

In the second test we investigate whether the generator is a smooth function on  $\mathbf{z}$ . Essentially, does the generator randomly associate points in the prior space with maps or, learns a smooth and dense mapping?. This can





**Figure 6** The distributions of the Minkowski functionals at 3 equidistant thresholds shown for illustration

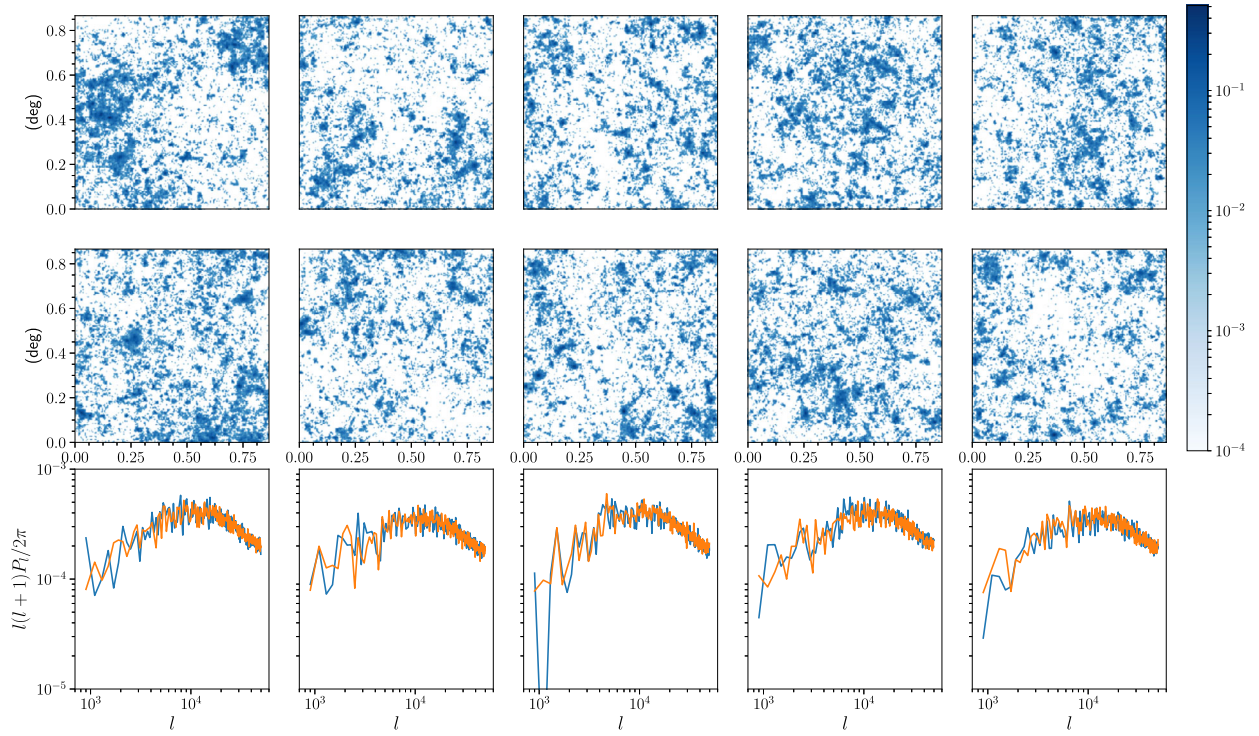
be tested by randomly choosing two noise vectors and evaluating the generator on the line connecting them. We use spherical linear interpolation to interpolate in the Gaussian prior space (White 2016). Figure 8 shows examples from this test. It is clear that the generated maps smoothly morph from one to another when the traversing the line connecting their points in the prior space.

In attempting to estimate how many distinct maps the generator can produce, we generated a few batches of 32k maps each. We then examined the pairwise distances among all the maps in each batch to try to find maps that are “close” or identical to each other. We looked at  $L_2$  distance in pixels space, in the power spectra space and in the space of the activations of the last conv-layer of the discriminator. The maps determined as “close” in this way are still noticeably distinct (visually and in the associated metric). If we are to use the “Birthday Paradox” test (Arora and Zhang 2017), this indicates that the number of distinct maps the generator can produce is  $\gg 1\text{B}$  maps. It also indicates that our generator does not suffer from mode collapse. Our conjecture here is that one or a combination


of multiple factors could explain this result: (1) the convolution kernels might be the right ansatz to describe the Gaussian fluctuations and the higher order structures encoded in convergence maps, (2) the maps dataset contain only one mode, so the normal prior, which has one mode, maps smoothly and nicely to the mode of the convergence maps (Xiao et al. 2018).

#### 4 Discussion and future work

The idea of applying deep generative models techniques to emulating scientific data and simulations has been gaining attention recently (Ravanbakhsh et al. 2016; de Oliveira et al. 2017; Paganini et al. 2017; Mosser et al. 2017). In this paper we have been able to reproduce maps of a particular  $\Lambda\text{CDM}$  model with very high-precision in the derived summary statistics, even though the network was not explicitly trained to recover these distributions. Furthermore we have explored, and reproduced the *statistical variety* present in simulations, in addition to the mean distribution. It remains an interesting question for the applicability of these techniques to science whether the increased precision achieved in this case, compared to other similar



**Figure 7** Comparison of randomly generated maps (top) with training maps (middle) that are their nearest-neighbor in terms of  $L^2$  distance in power spectrum. While the power spectra (bottom) match very well, the maps are clearly different



**Figure 8** Each row is a random example of maps produced when interpolating between two randomly chosen vectors  $\mathbf{z}$  (left to right). The generator varies smoothly between points in the prior space

work (Rodríguez et al. 2018), is due to the training procedure as outlined in Sect. 2.3 or to a difference in the underlying data.

We have also studied the ability of the generator to generate novel maps, the results shown in Sect. 3.2. We have provided evidence that our model has avoided ‘mode collapse’ and produces distinct maps from the originals that also smoothly vary when interpolating in the prior space.

Finally, the success of GANs in this work demonstrates that, unlike natural images which are the focus of the deep learning community, scientific data come with a suite of tools (statistical tests) which enable the verification and improvement of the fidelity of generative networks. The lack of such tools for natural images is a challenge to understanding the relative performance of the different generative models architectures, loss functions and training

procedures (Theis et al. 2015; Salimans et al. 2016). So, as well as the impact within these fields, scientific data with its suite of tools have the potential to provide the grounds for benchmarking and improving on the state-of-the-art generative models.

It is worth noting here that while one of the use cases of GANs is data augmentation, we do not think the generated maps augment the original dataset for statistical purposes. This is for the simple reason that the generator is a fit to the data, sampling a fit does not generate statistically independent samples. The current study is a proof-of-concept that highly over-parametrized functions built out of convolutional neural networks and optimized in by means of adversarial training in a GAN framework can be used to fit convergence maps with high statistical fidelity. The real utility of GANs would come from the open question of whether conditional generative models (Mirza and Osindero 2014; Dumoulin et al. 2018) can be used for emulating scientific data and parameter inference. The problem is outlined below and is to be addressed in future work.

Without loss of generality, the generation of one random realization of a science dataset (simulation or otherwise) can be posited as a black-box model  $S(\sigma, r)$  where  $\sigma$  is a vector of the physical model parameters and  $r$  is a set of random numbers. The physical model  $S$  can be based on first principles or effective theories. Different random seeds generate statistically independent mock data realizations of the model parameters  $\sigma$ . Such models are typically computationally expensive to evaluate at many different points in the space of parameters  $\sigma$ .

In our minds, the most important next step to build on the foundation of this paper and achieve an emulator of model  $S$ , is the ability of generative models to generalize in the space of model parameters  $\sigma$  from datasets at a finite number of points in the parameter space. More specifically, *can we use GANs to build parametric density estimators  $G(\sigma, \mathbf{z})$  of the physical model  $S(\sigma, r)$ ?* Such generalization rests on smoothness and continuity of the response function of the physical model  $S$  in the parameter space, but such an assumption is the foundation of any surrogate modeling. Advances in conditioning generative models (Mirza and Osindero 2014; Dumoulin et al. 2018) is a starting point to enable this goal. Future extensions of this work will seek to add this parameterization in order to enable the construction of robust and computationally inexpensive emulators of cosmological models.

## 5 Conclusion

We present here an application of Generative Adversarial Networks to the creation of weak-lensing convergence maps. We demonstrate that the use of neural networks for this purpose offers an extremely fast generator and therefore considerable promise for cosmology applications. We are able to obtain very high-fidelity generated images, reproducing the power spectrum and higher-order statistics

with unprecedented precision for a neural network based approach. We have probed these generated maps in terms of the correlations within the maps and the ability of the network to generalize and create novel data. The successful application of GANs to produce high-fidelity maps in this work provides an important foundation to using these approaches as fast emulators for cosmology.

## Acknowledgements

MM would like to thank Evan Racah for his invaluable deep learning related discussions throughout this project. MM would also like to thank Taehoon Kim for providing an open-sourced TensorFlow implementation of DCGAN which was used in early evaluations of this project. We would also like to thank Ross Girshick for helpful suggestions. No network was hurt during the adversarial training performed in this work.

## Funding

MM, DB and WB are supported by the U.S. Department of Energy, Office of Science, Office of High Energy Physics, under contract No. DEAC0205CH11231 through the National Energy Research Scientific Computing Center and Computational Center for Excellence programs. ZL was partially supported by the Scientific Discovery through Advanced Computing (SciDAC) program funded by U.S. Department of Energy Office of Advanced Scientific Computing Research and the Office of High Energy Physics.

## Abbreviations

HPC, High Performance Computing; GAN, Generative Adversarial Networks; DCGAN, Deep Convolutional Generative Adversarial Networks; ReLU, Rectified Linear Units; GPU, Graphical Processing Unit; KS, Kolmogorov–Smirnov; AD, Andersen–Darling.

## Availability of data and materials

Code and example datasets with pretrained weights used to generate this work are publicly available at

<https://github.com/MustafaMustafa/cosmoGAN>.<sup>b</sup> The full dataset used in this work is available upon request from the authors. We also provide the fully trained model which was used for all results in this paper, how-to at link above.

## Competing interests

Biological authors declare that they have no competing interests. Two artificial neural networks of this work: “the generator” and “the discriminator” are however directly competing against each other, effectively playing a zero-sum game.

## Authors’ contributions

MM proposed the project, designed and carried out the experiments and statistical analysis and prepared the first draft. DB guided the cosmology and statistical analysis, in addition to writing the introduction jointly with ZL and co-writing the Results section. ZL provided guidance on cosmology simulations and emulators in addition to co-writing the Discussion section. WB contributed to interpreting the Results and Discussions of future work in addition to reviewing the manuscripts. RA contributed to discussions of the deep learning aspects of the project and reviewing the manuscript. JK produced the simulated convergence maps. All authors read and approved the final manuscript.

## Author details

<sup>1</sup>Lawrence Berkeley National Laboratory, Berkeley, USA. <sup>2</sup>Google Research, Mountain View, USA. <sup>3</sup>University of KwaZulu-Natal, Westville, South Africa.

## Endnotes

<sup>a</sup> We used a ROOT (Brun and Rademakers 1997) implementation of the Andersen–Darling test and Scipy (Jones et al. 2001) for the Kolmogorov–Smirnov test.

<sup>b</sup> url date: 2019-03-15

## Publisher’s Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Received: 29 December 2018 Accepted: 16 April 2019

Published online: 06 May 2019

## References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., Zheng, X.: TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems (2015). <http://download.tensorflow.org/paper/whitepaper2015.pdf>
- Abbott, T., Abdalla, F.B., Allam, S., Amara, A., Annis, J., Armstrong, R., Bacon, D., Banerji, M., Bauer, A.H., Baxter, E., Becker, M.R., Benoit-Lévy, A., Bernstein, R.A., Bernstein, G.M., Bertin, E., Blazek, J., Bonnett, C., Bridle, S.L., Brooks, D., Bruderer, C., Buckley-Geer, E., Burke, D.L., Busha, M.T., Capozzi, D., Carnero Rosell, A., Carrasco Kind, M., Carretero, J., Castander, F.J., Chang, C., Clampitt, J., Crocce, M., Cunha, C.E., D'Andrea, C.B., da Costa, L.N., Das, R., DePoy, D.L., Desai, S., Diehl, H.T., Dietrich, J.P., Dodelson, S., Doel, P., Drlica-Wagner, A., Efstathiou, G., Eifler, T.F., Erickson, B., Estrada, J., Evrard, A.E., Fausti Neto, A., Fernandez, E., Finley, D.A., Flaugher, B., Fosalba, P., Friedrich, O., Frieman, J., Gangkofner, C., García-Bellido, J., Gaztanaga, E., Gerdes, D.W., Gruen, D., Gruendl, R.A., Gutierrez, G., Hartley, W., Hirsch, M., Honscheid, K., Huff, E.M., Jain, B., James, D.J., Jarvis, M., Kacprzak, T., Kent, S., Kirk, D., Krause, E., Kravtsov, A., Kuehn, K., Kuropatkin, N., Kwan, J., Lahav, O., Leistedt, B., Li, T.S., Lima, M., Lin, H., MacCrann, N., March, M., Marshall, J.L., Martini, P., McMahon, R.G., Melchior, P., Miller, C.J., Miquel, R., Mohr, J.J., Neilsen, E., Nichol, R.C., Nicola, A., Nord, B., Ogando, R., Palmese, A., Peiris, H.V., Plazas, A.A., Refregier, A., Roe, N., Romer, A.K., Roodman, A., Rowe, B., Rykoff, E.S., Sabiu, C., Sadeh, I., Sako, M., Samuroff, S., Sanchez, E., Sánchez, C., Seo, H., Sevilla-Noarbe, I., Sheldon, E., Smith, R.C., Soares-Santos, M., Sobreira, F., Suchyta, E., Swanson, M.E.C., Tarle, G., Thaler, J., Thomas, D., Troxel, M.A., Vikram, V., Walker, A.R., Wechsler, R.H., Weller, J., Zhang, Y., Zuntz, J.: Dark energy survey collaboration: cosmology from cosmic shear with dark energy survey science verification data. *Phys. Rev. D* **94**(2), 022001 (2016). <https://doi.org/10.1103/PhysRevD.94.022001>
- Planck Collaboration, Aghanim, N., Akrami, Y., Ashdown, M., Aumont, J., Baccigalupi, C., Ballardini, M., Banday, A.J., Barreiro, R.B., Bartolo, N., Basak, S., Battye, R., Benabed, K., Bernard, J.-P., Bersanelli, M., Bielewicz, P., Bock, J.J., Bond, J.R., Borrill, J., Bouchet, F.R., Boulanger, F., Bucher, M., Burigana, C., Butler, R.C., Calabrese, E., Cardoso, J.-F., Carron, J., Challinor, A., Chiang, H.C., Chluba, J., Colombo, L.P.L., Combet, C., Contreras, D., Crill, B.P., Cuttaia, F.D., Bernardis, P.D., Zotti, G., Delabrouille, J., Delouis, J.-M., Di Valentino, E., Diego, J.M., Doré, O., Douspis, M., Ducout, A., Dupac, X., Dusini, S., Efstathiou, G., Elsner, F., Enßlin, T.A., Eriksen, H.K., Fantaye, Y., Farhang, M., Fergusson, J., Fernandez-Cobos, R., Finelli, F., Forastieri, F., Frailis, M., Franceschi, E., Frolov, A., Galeotta, S., Galli, S., Ganga, K., Génova-Santos, R.T., Gerbino, M., Ghosh, T., González-Nuevo, J., Górski, K.M., Gratton, S., Gruppuso, A., Gudmundsson, J.E., Hamann, J., Handley, W., Herranz, D., Hivon, E., Huang, Z., Jaffe, A.H., Jones, W.C., Karaki, A., Keihänen, E., Keskitalo, R., Kiiveri, K., Kim, J., Kisner, T.S., Knox, L., Krachmalnicoff, N., Kunz, M., Kurki-Suonio, H., Lagache, G., Lamarre, J.-M., Lasenby, A., Lattanzi, M., Lawrence, C.R., Le Jeune, M., Lemos, P., Lesgourgues, J., Levrier, F., Lewis, A., Liguori, M., Lilje, P.B., Lilley, M., Lindholm, V., López-Caniego, M., Lubin, P.M., Ma, Y.-Z., Macías-Pérez, J.F., Maggia, G., Maino, D., Mandolesi, N., Mangilli, A., Marcos-Caballero, A., Maris, M., Martin, P.G., Martinelli, M., Martínez-González, E., Matarrese, S., Mauri, N., McEwen, J.D., Meinhold, P.R., Melchiorri, A., Mennella, A., Migliaccio, M., Millea, M., Mitra, S., Miville-Deschênes, M.-A., Molinari, D., Montier, L., Morgante, G., Moss, A., Natoli, P., Nørgaard-Nielsen, H.U., Pagano, L., Paoletti, D., Partridge, B., Patanchon, G., Peiris, H.V., Perrotta, F., Pettorino, V., Piacentini, F., Polastri, L., Polenta, G., Puget, J.-L., Rachen, J.P., Reinecke, M., Remazeilles, M., Renzi, A., Rocha, G., Rosset, C., Roudier, G., Rubiño-Martín, J.A., Ruiz-Granados, B., Salvati, L., Sandri, M., Savelainen, M., Scott, D., Shellard, E.P.S., Sirignano, C., Sirri, G., Spencer, L.D., Sunyaev, R., Suur-Uski, A.-S., Tauber, J.A., Tavagnacco, D., Tenti, M., Toffolatti, L., Tomasi, M., Trombetti, T., Valenziano, L., Valiviita, J., Van Tent, B., Vibert, L., Vielva, P., Villa, F., Vittorio, N., Wandelt, B.D., Wehus, I.K., White, M., White, S.D.M., Zacchei, A., Zonca, A.: Planck 2018 results. VI. Cosmological parameters. *arXiv e-prints* (2018). [arXiv:1807.06209](https://arxiv.org/abs/1807.06209)
- Alam, S., Ata, M., Bailey, S., Beutler, F., Bizyaev, D., Blazek, J.A., Bolton, A.S., Brownstein, J.R., Burden, A., Chuang, C.-H., Comparat, J., Cuesta, A.J., Dawson, K.S., Eisenstein, D.J., Escoffier, S., Gil-Marín, H., Grieb, J.N., Hand, N., Ho, S., Kinemuchi, K., Kirkby, D., Kitaura, F., Malanushenko, E., Malanushenko, V., Maraston, C., McBride, C.K., Nichol, R.C., Olmstead, M.D., Oravetz, D., Padmanabhan, N., Palanque-Delabrouille, N., Pan, K., Pellejero-Ibanez, M., Percival, W.J., Petitjean, P., Prada, F., Price-Whelan, A.M., Reid, B.A., Rodríguez-Torres, S.A., Roe, N.A., Ross, A.J., Ross, N.P., Rossi, G., Rubiño-Martín, J.A., Saito, S., Salazar-Albornoz, S., Samushia, L., Sánchez, A.G., Satpathy, S., Schlegel, D.J., Schneider, D.P., Scóccola, C.G., Seo, H.-J., Sheldon, E.S., Simmons, A., Slosar, A., Strauss, M.A., Swanson, M.E.C., Thomas, D., Tinker, J.L., Tojeiro, R., Magaña, M.V., Vazquez, J.A., Verde, L., Wake, D.A., Wang, Y., Weinberg, D.H., White, M., Wood-Vasey, W.M., Yèche, C., Zehavi, I., Zhai, Z., Zhao, G.-B.: The clustering of galaxies in the complete SDSS-III baryon oscillation spectroscopic survey: cosmological analysis of the DR12 galaxy sample. *Mon. Not. R. Astron. Soc.* **470**, 2617–2652 (2017). <https://doi.org/10.1093/mnras/stx721>
- Arjovsky, M., Bottou, L.: Towards principled methods for training generative adversarial networks. *ArXiv e-prints* (2017). [arXiv:1701.04862](https://arxiv.org/abs/1701.04862)
- Arjovsky, M., Chintala, S., Bottou, L.: Wasserstein GAN. *ArXiv e-prints* (2017). [arXiv:1701.07875](https://arxiv.org/abs/1701.07875)
- Arora, S., Zhang, Y.: Do gans actually learn the distribution? an empirical study. *CoRR* (2017). [arXiv:1706.08224](https://arxiv.org/abs/1706.08224)
- Bartelmann, M., Schneider, P.: Weak gravitational lensing. *Phys. Rep.* **340**, 291–472 (2001). [https://doi.org/10.1016/S0370-1573\(00\)00082-X](https://doi.org/10.1016/S0370-1573(00)00082-X)
- Bautista, J.E., Vargas-Magaña, M., Dawson, K.S., Percival, W.J., Brinkmann, J., Brownstein, J., Camacho, B., Comparat, J., Gil-Marín, H., Mueller, E.-M., Newman, J.A., Prakash, A., Ross, A.J., Schneider, D.P., Seo, H.-J., Tinker, J., Tojeiro, R., Zhai, Z., Zhao, G.-B.: The SDSS-IV extended baryon oscillation spectroscopic survey: baryon acoustic oscillations at redshift of 0.72 with the DR14 luminous red galaxy sample. *Astrophys. J.* **863**, 110 (2018). <https://doi.org/10.3847/1538-4357/aace5>
- Bernstein, G.M., Jarvis, M.: Shapes and shears, stars and smears: optimal measurements for weak lensing. *Astron. J.* **123**, 583–618 (2002). <https://doi.org/10.1086/338085>
- Brun, R., Rademakers, F.: Root—an object oriented data analysis framework. *Nucl. Instrum. Methods Phys. Res., Sect. A, Accel. Spectrom. Detect. Assoc. Equip.* **389**(1), 81–86 (1997). [https://doi.org/10.1016/S0168-9002\(97\)00048-X](https://doi.org/10.1016/S0168-9002(97)00048-X)
- Carlson, J., White, M., Padmanabhan, N.: Critical look at cosmological perturbation theory techniques. *Phys. Rev. D* **80**, 043531 (2009). <https://doi.org/10.1103/PhysRevD.80.043531>
- Csáji, B.C.: Approximation with artificial neural networks. MSc Thesis, Eötvös Loránd University (ELTE), Budapest, Hungary (2001)
- de Oliveira, L., Paganini, M., Nachman, B.: Learning particle physics by example: location-aware generative adversarial networks for physics synthesis. (2017). [arXiv:1701.05927](https://arxiv.org/abs/1701.05927)
- du Mas des Bourboux, H., Le Goff, J.-M., Blomqvist, M., Busca, N.G., Guy, J., Rich, J., Yèche, C., Bautista, J.E., Burtin, É., Dawson, K.S., Eisenstein, D.J., Font-Ribera, A., Kirkby, D., Miralda-Escudé, J., Noterdaeme, P., Palanque-Delabrouille, N., Pâris, I., Petitjean, P., Pérez-Ràfols, I., Pieri, M.M., Ross, N.P., Schlegel, D.J., Schneider, D.P., Slosar, A., Weinberg, D.H., Zarrouk, P.: Baryon acoustic oscillations from the complete SDSS-III Ly $\alpha$ -quasar cross-correlation function at  $z = 2.4$ . *Astron. Astrophys.* **608**, 130 (2017). <https://doi.org/10.1051/0004-6361/201731731>
- Dumoulin, V., Perez, E., Schucher, N., Strub, F., Vries, H.D., Courville, A., Bengio, Y.: Feature-wise transformations. *Distill* (2018). <https://doi.org/10.23915/distill.00011>
- Fluri, J., Kacprzak, T., Refregier, A., Amara, A., Lucchi, A., Hofmann, T.: Cosmological constraints from noisy convergence maps through deep learning. *Phys. Rev. D* **98**, 123518 (2018). <https://doi.org/10.1103/PhysRevD.98.123518>
- Gelman, A., Carlin, J.B., Stern, H.S., Rubin, D.B.: Bayesian Data Analysis, 3rd edn. Chapman and Hall/CRC, London (2013)
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N.D., Weinberger, K.Q. (eds.) *Advances in Neural Information Processing Systems*, vol. 27, pp. 2672–2680. Curran Associates, Red Hook (2014). <http://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf>
- Goodfellow, I.J.: NIPS 2016 tutorial: generative adversarial networks. *CoRR* (2017). [arXiv:1701.00160](https://arxiv.org/abs/1701.00160)
- Gupta, A., Matilla, J.M.Z., Hsu, D., Haiman, Z.: Non-Gaussian information from weak lensing data via deep learning. *Phys. Rev. D* **97**, 103515 (2018). <https://doi.org/10.1103/PhysRevD.97.103515>



- Heitmann, K., Higdon, D., White, M., Habib, S., Williams, B.J., Lawrence, E., Wagner, C.: The coyote universe. II. Cosmological models and precision emulation of the nonlinear matter power spectrum. *Astrophys. J.* **705**, 156–174 (2009). <https://doi.org/10.1088/0004-637X/705/1/156>
- Ioffe, S., Szegedy, C.: Batch normalization: accelerating deep network training by reducing internal covariate shift. In: Bach, F., Blei, D. (eds.) *Proceedings of the 32nd International Conference on Machine Learning*, vol. 37, pp. 448–456. PMLR, Lille (2015). <http://proceedings.mlr.press/v37/ioffe15.html>
- Jee, M.J., Tyson, J.A., Schneider, M.D., Wittman, D., Schmidt, S., Hilbert, S.: Cosmic shear results from the deep lens survey. I. Joint constraints on  $\Omega_M$  and  $\sigma_8$  with a two-dimensional analysis. *Astrophys. J.* **765**, 74 (2013). <https://doi.org/10.1088/0004-637X/765/1/74>
- Jones, E., Oliphant, T., Peterson, P., et al.: SciPy: open source scientific tools for Python (2001). <http://www.scipy.org/>
- Kilbinger, M.: Cosmology with cosmic shear observations: a review. *Rep. Prog. Phys.* **78**(8), 086901 (2015)
- Kingma, D.P., Ba, J., Adam: A method for stochastic optimization CoRR (2014). [arXiv:1412.6980](https://arxiv.org/abs/1412.6980)
- Kratochvil, J.M., Haiman, Z., May, M.: Probing cosmology with weak lensing peak counts. *Phys. Rev. D* **84**, 043519 (2010)
- Kratochvil, J.M., Lim, E.A., Wang, S., Haiman, Z., May, M., Huppenberger, K.: Probing cosmology with weak lensing Minkowski functionals. *Phys. Rev. D* **85**(10), 103513 (2012). <https://doi.org/10.1103/PhysRevD.85.103513>
- Laureijs, R., Amiaux, J., Arduini, S., Auguères, J., Brinchmann, J., Cole, R., Cropper, M., Dabin, C., Duvet, L., Ealet, A., et al.: Euclid definition study report. *ArXiv e-prints* (2011). [arXiv:1110.3193](https://arxiv.org/abs/1110.3193)
- Lawrence, E., Heitmann, K., White, M., Higdon, D., Wagner, C., Habib, S., Williams, B.: The coyote universe. III. Simulation suite and precision emulator for the nonlinear matter power spectrum. *Astrophys. J.* **713**, 1322–1331 (2010). <https://doi.org/10.1088/0004-637X/713/2/1322>
- Liu, J., Petri, A., Haiman, Z., Hui, L., Kratochvil, J.M., May, M.: Cosmology constraints from the weak lensing peak counts and the power spectrum in CFHTLenS data. *Phys. Rev. D* **91**(6), 063507 (2015). <https://doi.org/10.1103/PhysRevD.91.063507>
- LSST Dark Energy Science Collaboration: Large synoptic survey telescope: dark energy science collaboration. *ArXiv e-prints* (2012). [arXiv:1211.0310](https://arxiv.org/abs/1211.0310)
- Lučić, M., Kurach, K., Michalski, M., Gelly, S., Bousquet, O.: Are gans created equal? A large-scale study. In: *Advances in Neural Information Processing Systems* (NeurIPS) (2018). <https://arxiv.org/pdf/1711.10337.pdf>
- Maas, A.L., Hannun, A.Y., Ng, A.Y.: Rectifier nonlinearities improve neural network acoustic models. In: Dasgupta, S., McAllester, D. (eds.) *Proceedings of the 30th International Conference on Machine Learning*, vol. 28. PMLR, Atlanta (2013)
- Mecke, K.R., Buchert, T., Wagner, H.: Robust morphological measures for large-scale structure in the universe. *Astron. Astrophys.* **288**, 697–704 (1994). [astro-ph/9312028](https://arxiv.org/abs/astro-ph/9312028)
- Metropolis, N., Rosenbluth, A.W., Rosenbluth, M.N., Teller, A.H., Teller, E.: Equation of state calculations by fast computing machines. *J. Chem. Phys.* **21**(6), 1087–1092 (1953)
- Mirza, M., Osindero, S.: Conditional generative adversarial nets. CoRR (2014). [arXiv:1411.1784](https://arxiv.org/abs/1411.1784)
- Mosser, L., Dubrulle, O., Blunt, M.J.: Reconstruction of three-dimensional porous media using generative adversarial neural networks. (2017). [arXiv:1704.03225](https://arxiv.org/abs/1704.03225)
- Nair, V., Hinton, G.E.: Rectified linear units improve restricted Boltzmann machines. In: Fürnkranz, J., Joachims, T. (eds.) *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pp. 807–814. Omnipress, Haifa (2010). <http://www.icml2010.org/papers/432.pdf>
- Paganini, M., de Oliveira, L., Nachman, B.: CaloGAN: simulating 3D high energy particle showers in multi-layer electromagnetic calorimeters with generative adversarial networks. (2017). [arXiv:1705.02355](https://arxiv.org/abs/1705.02355)
- Peel, A., Lalande, F., Starck, J.-L., Pettorino, V., Merten, J., Giocoli, C., Meneghetti, M., Baldi, M.: Distinguishing standard and modified gravity cosmologies with machine learning. *ArXiv e-prints* (2018). [arXiv:1810.11030](https://arxiv.org/abs/1810.11030)
- Petri, A.: Mocking the weak lensing universe: the lenstools python computing package. *Astron. Comput.* **17**, 73–79 (2016). <https://doi.org/10.1016/j.jascom.2016.06.001>
- Petri, A., Liu, J., Haiman, Z., May, M., Hui, L., Kratochvil, J.M.: Emulating the CFHTLenS weak lensing data: cosmological constraints from moments and Minkowski functionals. *Phys. Rev. D* **91**(10), 103511 (2015). <https://doi.org/10.1103/PhysRevD.91.103511>
- Radford, A., Metz, L., Chintala, S.: Unsupervised representation learning with deep convolutional generative adversarial networks. CoRR (2015). [arXiv:1511.06434](https://arxiv.org/abs/1511.06434)
- Ravanbakhsh, S., Lanusse, F., Mandelbaum, R., Schneider, J., Poczos, B.: Enabling dark energy science with deep generative models of galaxy images (2016). [arXiv:1609.05796](https://arxiv.org/abs/1609.05796)
- Ribli, D., Pataki, B.A., Csabai, I.: An improved cosmological parameter inference scheme motivated by deep learning. *Nat. Astron.* **3**, 93–98 (2019). <https://doi.org/10.1038/s41550-018-0596-8>
- Rodríguez, A.C., Kacprzak, T., Lucchi, A., Amara, A., Sgier, R., Fluri, J., Hofmann, T., Réfrégier, A.: Fast cosmic web simulations with generative adversarial networks. *Comput. Astrophys. Cosmol.* **5**(1), 4 (2018). <https://doi.org/10.1186/s40668-018-0026-4>
- Salimans, T., Goodfellow, I.J., Zaremba, W., Cheung, V., Radford, A., Chen, X.: Improved techniques for training gans. CoRR (2016). [arXiv:1606.03498](https://arxiv.org/abs/1606.03498)
- Schawinski, K., Zhang, C., Zhang, H., Fowler, L., Santhanam, G.K.: Generative adversarial networks recover features in astrophysical images of galaxies beyond the deconvolution limit. *Mon. Not. R. Astron. Soc.* **467**, 110–114 (2017). <https://doi.org/10.1093/mnras/stx008>
- Shirasaki, M., Yoshida, N., Ikeda, S.: Denoising weak lensing mass maps with deep learning. *ArXiv e-prints* (2018). [arXiv:1812.05781](https://arxiv.org/abs/1812.05781)
- Spergel, D., Gehrels, N., Baltay, C., Bennett, D., Breckinridge, J., Donahue, M., Dressler, A., Gaudi, B.S., Greene, T., Guyon, O., Hirata, C., Kalirai, J., Kasdin, N.J., Macintosh, B., Moos, W., Perlmutter, S., Postman, M., Rauscher, B., Rhodes, J., Wang, Y., Weinberg, D., Benford, D., Hudson, M., Jeong, W.-S., Mellier, Y., Traub, W., Yamada, T., Capak, P., Colbert, J., Masters, D., Penny, M., Savransky, D., Stern, D., Zimmerman, N., Barry, R., Bartusek, L., Carpenter, K., Cheng, E., Content, D., Dekens, F., Demers, R., Grady, K., Jackson, C., Kuan, G., Kruk, J., Melton, M., Nemati, B., Parvin, B., Poberezhskiy, I., Peddie, C., Ruffa, J., Wallace, J.K., Whipple, A., Wollack, E., Zhao, F.: Wide-field Infrared survey telescope—astrophysics focused telescope assets WFIRST-AFTA 2015 report *ArXiv e-prints* (2015). [arXiv:1503.03757](https://arxiv.org/abs/1503.03757)
- Springel, V.: The cosmological simulation code GADGET-2. *Mon. Not. R. Astron. Soc.* **364**, 1105–1134 (2005). <https://doi.org/10.1111/j.1365-2966.2005.09655.x>
- Springenberg, J.T., Dosovitskiy, A., Brox, T., Riedmiller, M.A.: Striving for simplicity: the all convolutional net. CoRR (2014). [arXiv:1412.6806](https://arxiv.org/abs/1412.6806)
- Theis, L., van den Oord, A., Bethge, M.: A note on the evaluation of generative models. *ArXiv e-prints* (2015). [arXiv:1511.01844](https://arxiv.org/abs/1511.01844)
- Wang, M., Deng, W.: Deep visual domain adaptation: a survey. CoRR (2018). [arXiv:1802.03601](https://arxiv.org/abs/1802.03601)
- White, T.: Sampling generative networks: notes on a few effective techniques CoRR (2016). [arXiv:1609.04468](https://arxiv.org/abs/1609.04468)
- Xiao, C., Zhong, P., Zheng, C., Bourgan: Generative networks with metric embeddings. CoRR (2018). [arXiv:1805.07674](https://arxiv.org/abs/1805.07674)
- Yang, X., Kratochvil, J.M., Wang, S., Lim, E.A., Haiman, Z., May, M.: Cosmological information in weak lensing peaks. *Phys. Rev. D* **84**, 043529 (2011). <https://doi.org/10.1103/PhysRevD.84.043529>

**Submit your manuscript to a SpringerOpen<sup>®</sup> journal and benefit from:**

- Convenient online submission
- Rigorous peer review
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

Submit your next manuscript at ► [springeropen.com](https://www.springeropen.com)