Full-Sky Gravitational Lensing Simulations Using Generative Adversarial Networks

Pier Fiedorowicz [1] **William R. Coulton** [2] **Giulio Fabbian** [2] **Eduardo Rozo** [1] **Supranta S. Boruah** [3] **Shirley Ho** [2]

Abstract

We present a new method that uses a generative adversarial network to learn how to locally redistribute the mass in lognormal mass maps to achieve N-body quality full-sky weak lensing maps. Our mass maps reproduce a broad range of weak lensing summary statistics with percent level accuracy. Producing a single full-sky map requires ~ 10 seconds on an average compute node with no GPU acceleration. Relative to running a dark matter simulation, our algorithm reduces run time by more than four orders of magnitude.

1. Introduction

Weak gravitational lensing refers to the slight distortion of background galaxies due to matter inhomogeneities along the line of sight (Bartelmann & Maturi, 2016). Modern cosmology surveys measure these tiny distortions using tens of millions to billions of galaxies to reconstruct the projected matter density in the Universe (Jeffrey et al., 2021) and constrain cosmological parameters (Abbott et al., 2022).

Robust cosmological inference from weak lensing survey data requires large suites of N-body simulations. In particular, the fact that small-scale modes in the density field exhibit strong non-linear growth can only be adequately modeled using N-body simulations. However, these simulations are computationally expensive, requiring hundreds of core hours or more per simulation. This is an especially daunting prospect for map-based inference schemes (Porqueres et al., 2021; Boruah et al., 2022).

Machine learning-based methods present an attractive approach for generating N-body quality simulations in a fraction of the time. Generative adversarial networks (GANs, Goodfellow et al., 2014) have been successful in modeling complex data sets (Brock et al., 2018). A GAN is comprised of two neural networks which compete with each other through an adversarial loss. The first network is the generator, $G(\vec{z})$, which maps a random sample $\vec{z}$ from the latent space prior to the data manifold. The second network is the discriminator, $D(\vec{x})$, which learns to distinguish real samples from those created by the generator. The networks are adversarial in that the discriminator is trained to differentiate between real and fake samples, while the generator is trained to fool the discriminator.

GANs are exceptionally well suited for modeling systems in which real samples are easy to collect but cannot be analytically described. As such, they are ideally suited for modeling the outputs of N-body simulations. Indeed, GANs have already been used to generate weak lensing simulations for cosmology. However, prior work has been limited to either small flat sky patches (Mustafa et al., 2019; Perraudin et al., 2021), or curved-sky patches within a fixed footprint (Yiu et al., 2021).

This work presents a new technique for generating full-sky weak lensing simulations of the projected matter density (aka convergence) field. We first generate an approximate convergence field using the lognormal model, a well-known analytic approximation for the convergence field (Xavier et al., 2016). We then use GANs to learn how to locally redistribute mass in small patches to recover lensing maps that are statistically indistinguishable from N-body simulations. This approach has several unique advantages: 1) The lognormal model encodes a significant amount of physical information in the field, dramatically reducing the demands on the generator and allowing us to impose a physically motivated structure on the network; 2) Despite producing full-sky maps, our network is trained exclusively on small-scale patches, relying on the lognormal model to capture large scale features adequately; 3) The latent space of the model (i.e., the lognormal field) is interpretable; and 4) our model can be used to extend current lognormal-based methods of map-based cosmological inference (Boruah et al., 2022).

[1]Department of Physics, University of Arizona, 1118 E. Fourth Street, Tucson, AZ, 85721, USA [2]Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, New York, NY 10010, USA [3]Department of Astronomy and Steward Observatory, University of Arizona, 933 N Cherry Ave, Tucson, AZ 85719, USA. Correspondence to: Pier Fiedorowicz <pierfied@email.arizona.edu>.

Proceedings of the 39^{th} International Conference on Machine Learning, Baltimore, Maryland, USA, PMLR 162, 2022.

2. Method

2.1. The Lognormal Model

To generate a full-sky lensing map, we first generate a map using the log-normal model. Taruya et al. (2002) & Clerkin et al. (2016) demonstrated that convergence maps from numerical simulations can be adequately approximated as lognormal on large scales. The lognormal convergence field, $\vec{\kappa}$, can be expressed as

$$\kappa_i = e^{y_i} - \lambda, \tag{1}$$

where $\vec{y}$ is a Gaussian random field and λ is known as the "shift" parameter. The latter is fit to simulations. Given the non-linear power spectrum of the convergence field as a function of cosmology, we can readily compute the corresponding correlation function $\xi_{\kappa\kappa}$. This quantity is related to the correlation function ξ_{yy} via

$$\xi_{yy} = \ln\left(\frac{\xi_{\kappa\kappa}}{\lambda^2} + 1\right). \tag{2}$$

To ensure the mean convergence is zero, the expectation value of y is set to $\langle y \rangle = \ln\lambda - \frac{1}{2}\xi_{yy}(0)$. These equations provide a complete statistical description of the model. To generate a lognormal field, we evaluate the power spectrum C_l^{yy} of the y-field. We then randomly draw the amplitude and phases of y, and apply the non-linear transformation in equation 1 to arrive at a lognormal random field. We refer the reader to Xavier et al. (2016) for further details.

2.2. Generative Adversarial Network Implementation

We use GANs to learn how to modify lognormal maps such that they become statistically indistinguishable from those produced from numerical simulations. In a traditional GAN, the latent space is typically a comparatively low dimensional space with a Gaussian prior distribution, $z \sim N(\vec{0}, I)$. By contrast, in our method, the latent space has the same dimensionality as the output map from the generator. The latent space prior is that the input map is a random draw from the lognormal model. Because the lognormal model accurately describes the convergence field on large scales, the action of the generator G on the input map can be thought of as a local mass redistribution, implying convolutional neural networks (CNN) are particularly well suited to our problem. The end result is that the demands on the generator are drastically reduced relative to the standard approach despite the high dimensionality of our latent space ($\approx 3 \times 10^6$ for a full-sky map at our default resolution).

Based on the above discussion, we choose a CNN as our generator. Standard CNNs can only operate on flat Euclidean spaces. One option is to split the sphere into small patches, which can be approximated as flat, apply a CNN, and then recombine the patches (Han et al., 2021). However, we chose to implement CNNs directly on the surface of a sphere. Using the HEALPix pixelization of the sphere (Gorski et al., 2005), we perform rotation equivariant convolutions using the four nearest neighbors to each pixel. Rotation equivariance is desirable as the underlying physics should be rotation equivariant. Previous work has also demonstrated that rotation equivariant CNNs perform comparably or even outperform anisotropic CNNs for field-level cosmological parameter inference (Fluri et al., 2022; also see Dai & Seljak, 2022).

Because our method uses the lognormal model as a starting point, we can impose strong inductive biases with a physically motivated network architecture resulting in highly restricted networks with remarkably few parameters. The discriminator and generator networks share identical architectures. We use a ResNet (He et al., 2015) architecture with four residual blocks and 16 channels per convolutional filter. Each network contains only 3290 trainable parameters. To restrict the receptive field size of the networks, we do not include any down/upsampling layers. The results shown here are obtained using a HEALPix resolution of $N_{\rm side} = 512$, corresponding to an angular resolution of 7 arcmin. The receptive field size of each network is $1\ \rm deg^2$.

2.3. Network Training

Because the generator and discriminator only need to learn small-scale physics, it is unnecessary to train the networks on the full sky. Instead, we train the networks using small $3.5 \times 3.5\ \rm deg^2$ patches. To ensure the generator is not adversely affected by edge effects, the discriminator is restricted to the inner $2.5 \times 2.5\ \rm deg^2$ patch, ensuring all pixels contributing to the discriminator include their full receptive field within the training patch.

To avoid recomputing the pixel adjacencies at each patch on the sky, we choose to train at a patch centered on the north pole. We include patches in arbitrary positions in the sky by rotating these patches onto the north pole at a resolution of $N_{\rm side} = 2048$. The patches are then downgraded to $N_{\rm side} = 512$. This procedure avoids introducing resolution artifacts that become apparent when the rotation is done at $N_{\rm side} = 512$.

The networks are trained by optimizing the hinge loss (Lim & Ye, 2017), given by:

$$L = \max_G \min_D \mathbb{E}_{\vec{x}\sim P_{\rm data}}[\min(0, D(\vec{x}) - 1)] - \mathbb{E}_{\vec{z}\sim P_{\rm prior}}[\min(0, -D(G(\vec{z})) - 1)]. \tag{3}$$

The hinge loss is minimized relative to the discriminator parameters but maximized relative to the generator parameters, where the optimal point in parameter space is a saddle point. Because the generator's goal is to learn small perturbations to the lognormal model, we include an L1 identity loss term that limits the deviation of the generator mapping from the identity function,

$$L_{\text{identity}} = \mathbb{E}_{\kappa\sim P_{\text{prior}}}\left[|\vec{\kappa} - G(\vec{\kappa})|_1\right]. \tag{4}$$

In addition, we enforce that the generator acts as a local mass redistribution by including an L2 penalty term on the difference between the average convergence of the input and output fields.

$$L_{\text{local}} = \mathbb{E}_{\kappa\sim P_{\text{prior}}}\left[\left(\frac{1}{N}\sum_i(\kappa_i - G(\vec{\kappa})_i)\right)^2\right]. \tag{5}$$

Here, the sum is over N pixels in a test patch with the same size as the receptive field of the networks. The total loss is $L_{\text{tot}} = L - L_{\text{identity}} - L_{\text{local}}$. The "minus" appears because the loss is maximized with respect to the generator parameters.

Because GANs are challenging to train, much effort has gone into developing methods for improving the training stability (Gulrajani et al., 2017; Mescheder et al., 2018; Miyato et al., 2018). We follow Heusel et al. (2017) in training the discriminator with a higher learning rate than the generator but do not otherwise apply any other such regularization schemes. In particular, we have found that the small size of our networks combined with the L_{identity} and L_{local} penalties successfully stabilize the training.

To train our GAN, we use a set of 80 out of 108 full-sky weak lensing simulations generated for the HSC survey (Takahashi et al., 2017). The remaining 28 simulations are saved for testing. All simulations are generated using N-body simulations with a flat ΛCDM cosmology consistent with WMAP. We use the simulations to construct 80 convergence maps of $N_{\text{side}} = 512$ assuming survey properties (specifically the source redshift distribution dn/dz) consistent with those of the DES Y1 data set (Hoyle et al., 2018). Each simulation is then split into 3072 $3.5\times3.5\,\text{deg}^2$ patches. The input to the generator is a corresponding set of 80 random full-sky lognormal maps drawn from the same WMAP cosmology that went into the simulated data set, each split into the same $3.5 \times 3.5\,\text{deg}^2$ patches.

3. Results and Discussion

We define our final generator using an exponentially weighted moving average of the model parameters with a decay rate of 0.999. Once the GAN is trained, we use an independent set of 28 lognormal random mocks that are to be compared to the remaining 28 simulated maps. Figure 1 illustrates how we generate a sample. We first sample unstructured white noise. The randomly chosen spherical harmonic amplitudes are rescaled by the appropriate power spectrum, and the lognormal transform is applied. Finally, the GAN's generator locally redistributes mass to produce the output GAN sample. Figure 2 shows a comparison of a GAN map and an N-body mass map over the full sky and in two randomly selected 11.5 deg. by 11.5 deg. patches. We generated this plot several times, randomly selecting a different GAN map and a different N-body map and randomly deciding which map appeared in which hemisphere. We are unable to distinguish between the two maps by eye.

To quantitatively test the performance of our GAN, we compare the full-sky HSC mocks to the full-sky maps produced by our network across a variety of summary statistics, namely the maps': 1) power spectrum; 2) one-point function; 3) peak counts; and 4) void counts. We compute each summary statistic's average and standard deviation across the 28 test mocks. Error bars are the sample variance, not the error on the mean. The comparison between simulated and GAN map statistics is shown in Figure 3. The agreement across all the statistics we consider is remarkable. The inverse-variance weighted RMS fluctuations between the summary statistics of the simulated maps and the GAN maps is 0.65%, 0.91%, 1.29%, and 2.42% for the power-spectrum, one-point function, peak counts, and void counts, respectively.

We have demonstrated that our model allows us to generate high-quality full-sky weak lensing simulations quickly: one such simulation at a resolution of 7 arcmin requires ~ 10 sec on a 32-core machine without GPU acceleration. Thus, a single node can generate thousands of full-sky simulations in less than a day, corresponding to just shy of a five order of magnitude speedup over running N-body simulations. This makes our tool ideal for generating large samples of simulated lensing maps that can be used to characterize the statistical properties of non-Gaussian summary statistics. More importantly, a cosmology-dependent extension of our model can be used to enable map-based inference techniques. In particular, Fiedorowicz et al. (2022) & Boruah et al. (2022) model the convergence field as a lognormal random field to construct cosmology and convergence posteriors. One can add our trained generator as an extra step in the forward model to improve the quality and resolution of the resulting posteriors.

Our current GAN is limited in two important ways: 1) it does not allow for the generation of tomographic mass maps, and 2) the GAN has been trained using only one cosmology. We intend to address these deficiencies in future work to capitalize on our proposed method's remarkable promise. Our success will directly lead to fast, map-based cosmological analyses of cosmic shear data, delivering on the tremendous gains enabled by forward-modeling techniques (Porqueres et al., 2021).

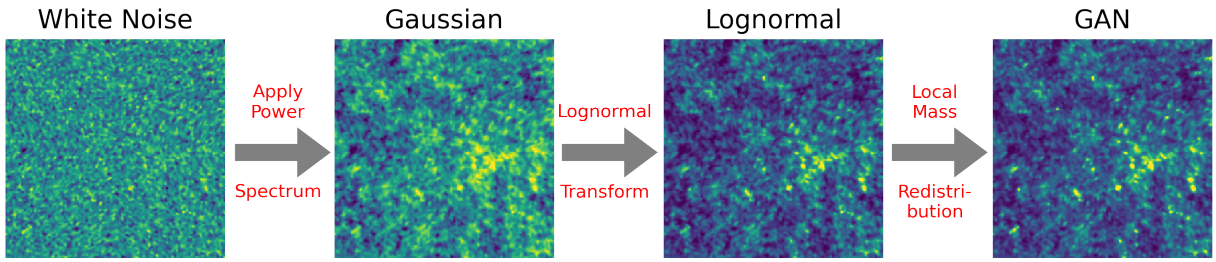

Figure 1. Outline of how we generate a GAN sample. A white noise map is turned into a Gaussian random field of the appropriate power spectrum. The lognormal transform is then applied. Finally, the GAN's generator locally redistributes mass to arrive at the final map.

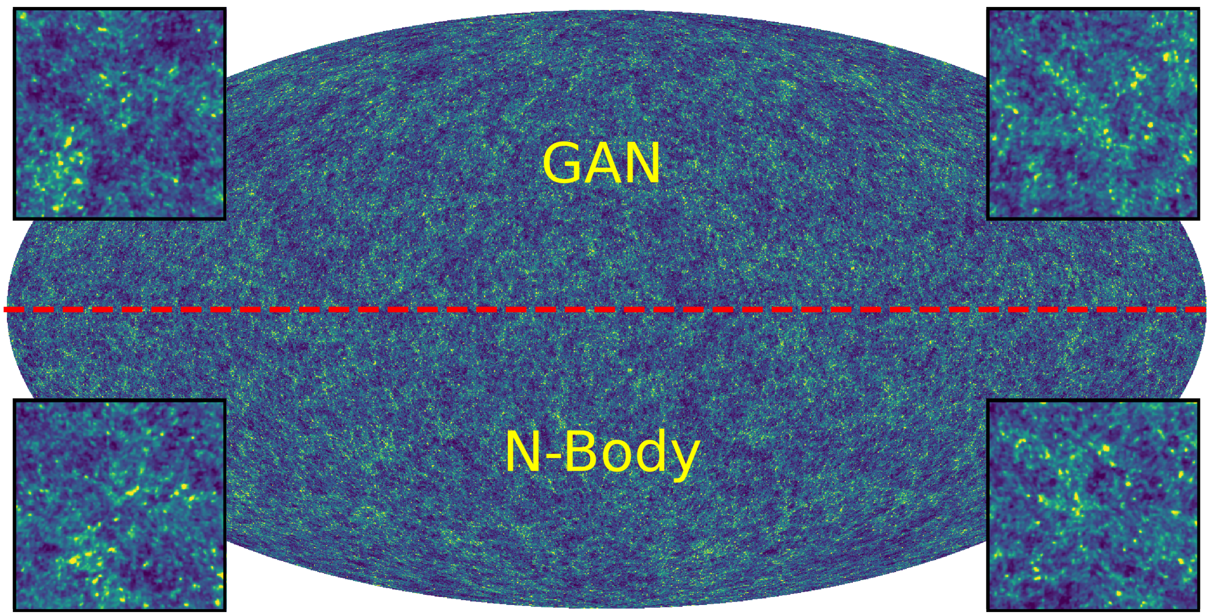

Figure 2. Comparison of a convergence map produced using our GAN (northern hemisphere) and an N-body simulation (southern hemisphere). The corner patches are 11.5 deg. by 11.5 deg. zoom-ins on random patches in each hemisphere.

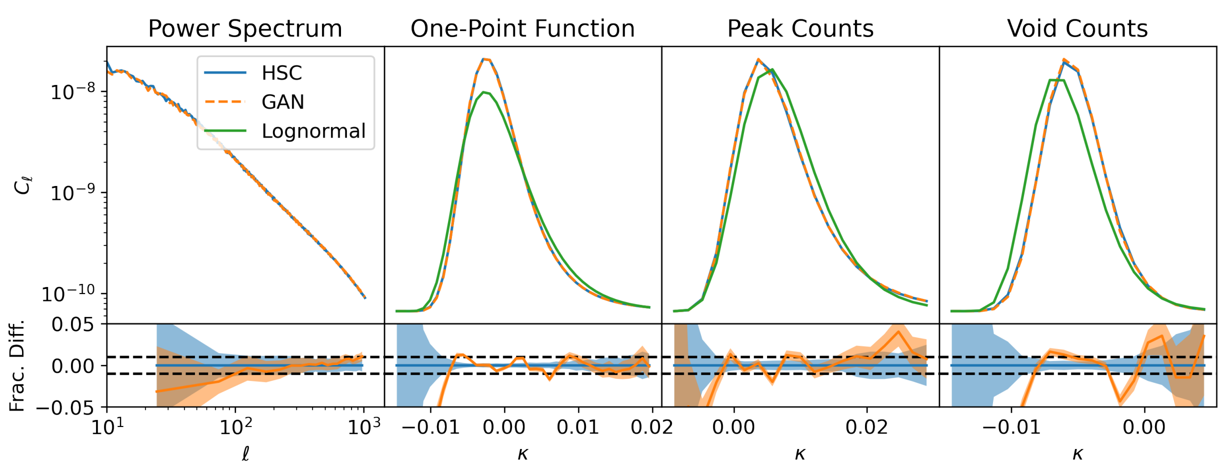

Figure 3. The mean of the multiple summary statistics across the test lognormal/GAN/N-body maps, as labeled. The bands in the lower panel show the variance in each statistic relative to the mean of the simulations. Black dashed lines corresponds to $\pm 1\%$ differences.

Acknowledgements

ER and PF are supported by NSF grant 2009401. ER also receives support from DOE grant DE-SC0009913.

References

Abbott, T., Aguena, M., Alarcon, A., Allam, S., Alves, O., Amon, A., Andrade-Oliveira, F., Annis, J., Avila, S., Bacon, D., Baxter, E., Bechtol, K., Becker, M., Bernstein, G., Bhargava, S., Birrer, S., Blazek, J., Brandao-Souza, A., Bridle, S., Brooks, D., Buckley-Geer, E., Burke, D., Camacho, H., Campos, A., Rosell, A. C., Kind, M. C., Carretero, J., Castander, F., Cawthon, R., Chang, C., Chen, A., Chen, R., Choi, A., Conselice, C., Cordero, J., Costanzi, M., Crocce, M., da Costa, L., da Silva Pereira, M., Davis, C., Davis, T., Vicente, J. D., DeRose, J., Desai, S., Valentino, E. D., Diehl, H., Dietrich, J., Dodelson, S., Doel, P., Doux, C., Drlica-Wagner, A., Eckert, K., Eifler, T., Elsner, F., Elvin-Poole, J., Everett, S., Evrard, A., Fang, X., Farahi, A., Fernandez, E., Ferrero, I., Ferté, A., Fosalba, P., Friedrich, O., Frieman, J., García-Bellido, J., Gatti, M., Gaztanaga, E., Gerdes, D., Giannantonio, T., Giannini, G., Gruen, D., Gruendl, R., Gschwend, J., Gutierrez, G., Harrison, I., Hartley, W., Herner, K., Hinton, S., Hollowood, D., Honscheid, K., Hoyle, B., Huff, E., Huterer, D., Jain, B., James, D., Jarvis, M., Jeffrey, N., Jeltema, T., Kovacs, A., Krause, E., Kron, R., Kuehn, K., Kuropatkin, N., Lahav, O., Leget, P.-F., Lemos, P., Liddle, A., Lidman, C., Lima, M., Lin, H., MacCrann, N., Maia, M., Marshall, J., Martini, P., McCullough, J., Melchior, P., Mena-Fernández, J., Menanteau, F., Miquel, R., Mohr, J., Morgan, R., Muir, J., Myles, J., Nadathur, S., Navarro-Alsina, A., Nichol, R., Ogando, R., Omori, Y., Palmese, A., Pandey, S., Park, Y., Paz-Chinchón, F., Petravick, D., Pieres, A., Malagón, A. P., Porredon, A., Prat, J., Raveri, M., Rodriguez-Monroy, M., Rollins, R., Romer, A., Roodman, A., Rosenfeld, R., Ross, A., Rykoff, E., Samuroff, S., Sánchez, C., Sanchez, E., Sanchez, J., Cid, D. S., Scarpine, V., Schubnell, M., Scolnic, D., Secco, L., Serrano, S., Sevilla-Noarbe, I., Sheldon, E., Shin, T., Smith, M., Soares-Santos, M., Suchyta, E., Swanson, M., Tabbutt, M., Tarle, G., Thomas, D., To, C., Troja, A., Troxel, M., Tucker, D., Tutusaus, I., Varga, T., Walker, A., Weaverdyck, N., Wechsler, R., Weller, J., Yanny, B., Yin, B., Zhang, Y., and and, J. Z. Dark energy survey year 3 results: Cosmological constraints from galaxy clustering and weak lensing. *Physical Review D*, 105(2), jan 2022. doi: 10.1103/physrevd.105.023520. URL https://doi.org/10.1103%2Fphysrevd.105.023520.

Bartelmann, M. and Maturi, M. Weak gravitational lensing, 2016. URL https://arxiv.org/abs/1612.06535.

Boruah, S. S., Rozo, E., and Fiedorowicz, P. Map-based cosmology inference with lognormal cosmic shear maps, 2022. URL https://arxiv.org/abs/2204.13216.

Brock, A., Donahue, J., and Simonyan, K. Large scale gan training for high fidelity natural image synthesis, 2018. URL https://arxiv.org/abs/1809.11096.

Clerkin, L., Kirk, D., Manera, M., Lahav, O., Abdalla, F., Amara, A., Bacon, D., Chang, C., Gaztañ aga, E., Hawken, A., Jain, B., Joachimi, B., Vikram, V., Abbott, T., Allam, S., Armstrong, R., Benoit-Lévy, A., Bernstein, G. M., Bernstein, R. A., Bertin, E., Brooks, D., Burke, D. L., Rosell, A. C., Kind, M. C., Crocce, M., Cunha, C. E., D'Andrea, C. B., da Costa, L. N., Desai, S., Diehl, H. T., Dietrich, J. P., Eifler, T. F., Evrard, A. E., Flaugher, B., Fosalba, P., Frieman, J., Gerdes, D. W., Gruen, D., Gruendl, R. A., Gutierrez, G., Honscheid, K., James, D. J., Kent, S., Kuehn, K., Kuropatkin, N., Lima, M., Melchior, P., Miquel, R., Nord, B., Plazas, A. A., Romer, A. K., Roodman, A., Sanchez, E., Schubnell, M., Sevilla-Noarbe, I., Smith, R. C., Soares-Santos, M., Sobreira, F., Suchyta, E., Swanson, M. E. C., Tarle, G., and Walker, A. R. Testing the lognormality of the galaxy and weak lensing convergence distributions from dark energy survey maps. *Monthly Notices of the Royal Astronomical Society*, 466(2):1444–1461, aug 2016. doi: 10.1093/mnras/stw2106. URL https://doi.org/10.1093%2Fmnras%2Fstw2106.

Dai, B. and Seljak, U. Translation and rotation equivariant normalizing flow (trenf) for optimal cosmological analysis, 2022. URL https://arxiv.org/abs/2202.05282.

Fiedorowicz, P., Rozo, E., Boruah, S. S., Chang, C., and Gatti, M. KaRMMa – kappa reconstruction for mass mapping. *Monthly Notices of the Royal Astronomical Society*, 512(1):73–85, feb 2022. doi: 10.1093/mnras/stac468. URL https://doi.org/10.1093%2Fmnras%2Fstac468.

Fluri, J., Kacprzak, T., Lucchi, A., Schneider, A., Refregier, A., and Hofmann, T. A full $w$cdm analysis of kids-1000 weak lensing maps using deep learning. *Physical Review D*, 105(8), apr 2022. doi: 10.1103/physrevd.105.083518. URL https://doi.org/10.1103%2Fphysrevd.105.083518.

Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial networks, 2014. URL https://arxiv.org/abs/1406.2661.

Gorski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., and Bartelmann, M.

HEALPix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere. *The Astrophysical Journal*, 622(2):759–771, apr 2005. doi: 10.1086/427976. URL https://doi.org/10.1086%2F427976.

Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. Improved training of wasserstein gans, 2017. URL https://arxiv.org/abs/1704.00028.

Han, D., Sehgal, N., and Villaescusa-Navarro, F. Deep learning simulations of the microwave sky. *Physical Review D*, 104(12), dec 2021. doi: 10.1103/physrevd.104.123521. URL https://doi.org/10.1103%2Fphysrevd.104.123521.

He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition, 2015. URL https://arxiv.org/abs/1512.03385.

Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. Gans trained by a two time-scale update rule converge to a local nash equilibrium. 2017. doi: 10.48550/ARXIV.1706.08500. URL https://arxiv.org/abs/1706.08500.

Hoyle, B., Gruen, D., Bernstein, G. M., Rau, M. M., Vicente, J. D., Hartley, W. G., Gaztanaga, E., DeRose, J., Troxel, M. A., Davis, C., Alarcon, A., MacCrann, N., Prat, J., Sá nchez, C., Sheldon, E., Wechsler, R. H., Asorey, J., Becker, M. R., Bonnett, C., Rosell, A. C., Carollo, D., Kind, M. C., Castander, F. J., Cawthon, R., Chang, C., Childress, M., Davis, T. M., Drlica-Wagner, A., Gatti, M., Glazebrook, K., Gschwend, J., Hinton, S. R., Hoormann, J. K., Kim, A. G., King, A., Kuehn, K., Lewis, G., Lidman, C., Lin, H., Macaulay, E., Maia, M. A. G., Martini, P., Mudd, D., Möller, A., Nichol, R. C., Ogando, R. L. C., Rollins, R. P., Roodman, A., Ross, A. J., Rozo, E., Rykoff, E. S., Samuroff, S., Sevilla-Noarbe, I., Sharp, R., Sommer, N. E., Tucker, B. E., Uddin, S. A., Varga, T. N., Vielzeuf, P., Yuan, F., Zhang, B., Abbott, T. M. C., Abdalla, F. B., Allam, S., Annis, J., Bechtol, K., Benoit-Lévy, A., Bertin, E., Brooks, D., Buckley-Geer, E., Burke, D. L., Busha, M. T., Capozzi, D., Carretero, J., Crocce, M., D'Andrea, C. B., da Costa, L. N., DePoy, D. L., Desai, S., Diehl, H. T., Doel, P., Eifler, T. F., Estrada, J., Evrard, A. E., Fernandez, E., Flaugher, B., Fosalba, P., Frieman, J., García-Bellido, J., Gerdes, D. W., Giannantonio, T., Goldstein, D. A., Gruendl, R. A., Gutierrez, G., Honscheid, K., James, D. J., Jarvis, M., Jeltema, T., Johnson, M. W. G., Johnson, M. D., Kirk, D., Krause, E., Kuhlmann, S., Kuropatkin, N., Lahav, O., Li, T. S., Lima, M., March, M., Marshall, J. L., Melchior, P., Menanteau, F., Miquel, R., Nord, B., O'Neill, C. R., Plazas, A. A., Romer, A. K., Sako, M., Sanchez, E., Santiago, B., Scarpine, V., Schindler, R., Schubnell, M., Smith, M., Smith, R. C., Soares-Santos, M., Sobreira, F., Suchyta, E., Swanson, M. E. C., Tarle, G., Thomas, D., Tucker, D. L., Vikram, V., Walker, A. R., Weller, J., Wester, W., Wolf, R. C., Yanny, B., and and, J. Z. Dark energy survey year 1 results: redshift distributions of the weak-lensing source galaxies. *Monthly Notices of the Royal Astronomical Society*, 478(1):592–610, apr 2018. doi: 10.1093/mnras/sty957. URL https://doi.org/10.1093%2Fmnras%2Fsty957.

Jeffrey, N., Gatti, M., Chang, C., Whiteway, L., Demirbozan, U., Kovacs, A., Pollina, G., Bacon, D., Hamaus, N., Kacprzak, T., Lahav, O., Lanusse, F., Mawdsley, B., Nadathur, S., Starck, J. L., Vielzeuf, P., Zeurcher, D., Alarcon, A., Amon, A., Bechtol, K., Bernstein, G. M., Campos, A., Rosell, A. C., Kind, M. C., Cawthon, R., Chen, R., Choi, A., Cordero, J., Davis, C., DeRose, J., Doux, C., Drlica-Wagner, A., Eckert, K., Elsner, F., Elvin-Poole, J., Everett, S., Ferté , A., Giannini, G., Gruen, D., Gruendl, R. A., Harrison, I., Hartley, W. G., Herner, K., Huff, E. M., Huterer, D., Kuropatkin, N., Jarvis, M., Leget, P. F., MacCrann, N., McCullough, J., Muir, J., Myles, J., Navarro-Alsina, A., Pandey, S., Prat, J., Raveri, M., Rollins, R. P., Ross, A. J., Rykoff, E. S., Sánchez, C., Secco, L. F., Sevilla-Noarbe, I., Sheldon, E., Shin, T., Troxel, M. A., Tutusaus, I., Varga, T. N., Yanny, B., Yin, B., Zhang, Y., Zuntz, J., Abbott, T. M. C., Aguena, M., Allam, S., Andrade-Oliveira, F., Becker, M. R., Bertin, E., Bhargava, S., Brooks, D., Burke, D. L., Carretero, J., Castander, F. J., Conselice, C., Costanzi, M., Crocce, M., da Costa, L. N., Pereira, M. E. S., Vicente, J. D., Desai, S., Diehl, H. T., Dietrich, J. P., Doel, P., Ferrero, I., Flaugher, B., Fosalba, P., García-Bellido, J., Gaztanaga, E., Gerdes, D. W., Giannantonio, T., Gschwend, J., Gutierrez, G., Hinton, S. R., Hollowood, D. L., Hoyle, B., Jain, B., James, D. J., Lima, M., Maia, M. A. G., March, M., Marshall, J. L., Melchior, P., Menanteau, F., Miquel, R., Mohr, J. J., Morgan, R., Ogando, R. L. C., Palmese, A., Paz-Chinchón, F., Plazas, A. A., Rodriguez-Monroy, M., Roodman, A., Sanchez, E., Scarpine, V., Serrano, S., Smith, M., Soares-Santos, M., Suchyta, E., Tarle, G., Thomas, D., To, C., and and, J. W. Dark energy survey year 3 results: Curved-sky weak lensing mass map reconstruction. *Monthly Notices of the Royal Astronomical Society*, 505(3):4626–4645, may 2021. doi: 10.1093/mnras/stab1495. URL https://doi.org/10.1093%2Fmnras%2Fstab1495.

Lim, J. H. and Ye, J. C. Geometric gan, 2017. URL https://arxiv.org/abs/1705.02894.

Mescheder, L., Geiger, A., and Nowozin, S. Which training methods for gans do actually converge? 2018. doi: 10.48550/ARXIV.1801.04406. URL https://arxiv.org/abs/1801.04406.

Miyato, T., Kataoka, T., Koyama, M., and Yoshida, Y. Spectral normalization for generative adversarial networks, 2018. URL https://arxiv.org/abs/1802.05957.

Mustafa, M., Bard, D., Bhimji, W., Lukić, Z., Al-Rfou, R., and Kratochvil, J. M. CosmoGAN: creating high-fidelity weak lensing convergence maps using generative adversarial networks. *Computational Astrophysics and Cosmology*, 6(1), may 2019. doi: 10.1186/s40668-019-0029-9. URL https://doi.org/10.1186%2Fs40668-019-0029-9.

Perraudin, N., Marcon, S., Lucchi, A., and Kacprzak, T. Emulation of cosmological mass maps with conditional generative adversarial networks. *Frontiers in Artificial Intelligence*, 4, 2021. ISSN 2624-8212. doi: 10.3389/frai.2021.673062. URL https://www.frontiersin.org/articles/10.3389/frai.2021.673062.

Porqueres, N., Heavens, A., Mortlock, D., and Lavaux, G. Lifting weak lensing degeneracies with a field-based likelihood. *Monthly Notices of the Royal Astronomical Society*, 509(3):3194–3202, nov 2021. doi: 10.1093/mnras/stab3234. URL https://doi.org/10.1093%2Fmnras%2Fstab3234.

Takahashi, R., Hamana, T., Shirasaki, M., Namikawa, T., Nishimichi, T., Osato, K., and Shiroyama, K. Full-sky gravitational lensing simulation for large-area galaxy surveys and cosmic microwave background experiments. *The Astrophysical Journal*, 850(1):24, nov 2017. doi: 10.3847/1538-4357/aa943d. URL https://doi.org/10.3847%2F1538-4357%2Faa943d.

Taruya, A., Takada, M., Hamana, T., Kayo, I., and Futamase, T. Lognormal property of weak-lensing fields. *The Astrophysical Journal*, 571(2):638–653, jun 2002. doi: 10.1086/340048. URL https://doi.org/10.1086%2F340048.

Xavier, H. S., Abdalla, F. B., and Joachimi, B. Improving lognormal models for cosmological fields. *Monthly Notices of the Royal Astronomical Society*, 459(4):3693–3710, apr 2016. doi: 10.1093/mnras/stw874. URL https://doi.org/10.1093%2Fmnras%2Fstw874.

Yiu, T. W. H., Fluri, J., and Kacprzak, T. A tomographic spherical mass map emulator of the kids-1000 survey using conditional generative adversarial networks, 2021. URL https://arxiv.org/abs/2112.12741.