

# Massive $\nu$ s through the CNN lens: interpreting the field-level neutrino mass information in weak lensing

Malika Golshan<sup>a,b</sup> and Adrian E. Bayer<sup>c,d</sup>

<sup>a</sup>Department of Physics, University of California, Berkeley, CA 94720, USA

<sup>b</sup>Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, USA

<sup>c</sup>Department of Astrophysical Sciences, Princeton University, Peyton Hall, Princeton NJ 08544, USA

<sup>d</sup>Center for Computational Astrophysics, Flatiron Institute, 162 5th Avenue, New York NY 10010, USA

E-mail: [malikagolshan@berkeley.edu](mailto:malikagolshan@berkeley.edu), [abayer@princeton.edu](mailto:abayer@princeton.edu)

**Abstract.** Modern cosmological surveys probe the Universe deep into the nonlinear regime, where massive neutrinos suppress cosmic structure. Traditional cosmological analyses, which use the 2-point correlation function to extract information, are no longer optimal in the nonlinear regime, and there is thus much interest in extracting beyond-2-point information to improve constraints on neutrino mass. Quantifying and interpreting the beyond-2-point information is thus a pressing task. We study the field-level information in weak lensing convergence maps using convolution neural networks. We find that the network performance increases as higher source redshifts and smaller scales are considered – investigating up to a source redshift of 2.5 and  $\ell_{\max} \simeq 10^4$  – verifying that massive neutrinos leave a distinct effect on weak lensing. However, the performance of the network significantly drops after scaling out the 2-point information from the maps, implying that most of the field-level information can be found in the 2-point correlation function alone. We quantify these findings in terms of the likelihood ratio and also use Integrated Gradient saliency maps to interpret which parts of the map the network is learning the most from, finding that the network extracts information from the most overdense and underdense regions.

---

## Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Method</b>	<b>2</b>
2.1	Data	2
2.2	Neural Network	4
<b>3</b>	<b>Results</b>	<b>5</b>
<b>4</b>	<b>Conclusions</b>	<b>9</b>

---

## 1 Introduction

Neutrinos have historically been among the most challenging elementary particles to detect and study, primarily because they interact only through the weak force and gravity. While their gravitational interaction is very weak due to their small mass, cosmological experiments are able to infer the neutrino mass from the distribution of radiation and matter in the Universe. Analysis of cosmic microwave background (CMB) data from Planck combined with large-scale structure (LSS) data from BOSS have presented upper limits on the neutrino mass of  $M_\nu < 0.12$  (95% CL) [1]. This is an order of magnitude lower than the upper limit on the effective electron neutrino mass measured in particle physics experiments of  $\beta$  decay which give  $m_{\nu_e}^{\text{eff}} \lesssim 1.1\text{eV}$  [2]. Given the minimum sum of the neutrino masses allowed by neutrino oscillations experiments [3–7] is  $M_\nu \gtrsim 0.06\text{eV}$  for the normal hierarchy, or  $M_\nu \gtrsim 0.1\text{eV}$  for the inverted hierarchy, there is much promise for the next generation of cosmological surveys – such as Vera Rubin Observatory LSST [8], Euclid [9], SPHEREx [10], Roman Space Telescope [11], DESI [12], PFS [13], Simons Observatory [14], CMB-S4 [15] – to measure the neutrino mass at the  $5\sigma$  level.

One particularly promising method to measure neutrino mass is by studying weak gravitational lensing. This involves analyzing the distortions of light emitted from distant galaxies caused by the gravitational influence of LSS between the galaxy and the telescope. Traditionally, information is extracted from weak-lensing convergence maps using the power spectrum, or 2-point correlation function, which is optimal in the linear regime. However, upcoming surveys will probe progressively smaller, nonlinear, scales, in which regime the 2-point statistics are no longer guaranteed to be optimal. Various simulations have been developed to understand the effects of neutrinos on nonlinear scales [16–22], and have suggested that there is much information to be found on these nonlinear scales using higher-order statistics such as voids and wavelets. For example, when considering the 3-dimensional matter field it has been shown that higher-order statistics can improve neutrino mass constraints by a factor of  $\sim 100$  compared to the power spectrum alone [23–30], while for weak lensing more modest numbers are reported [31–36]. Understanding and interpreting the source of this information is thus a pressing task.

In this work, we extend the analysis of [37], which showed that while there is much information regarding neutrino mass in the 3-dimensional matter field, almost all of the information up to  $k \sim 1 h/\text{Mpc}$  is contained in the power spectrum alone for observable fields, such as the weak lensing convergence field and galaxy field. This conclusion was

reached by considering the 3-dimensional field-level cross correlation between two simulations with matched linear physics and different neutrino masses at a single redshift, essentially assuming the distance between lensing source and observer to be small. As will be discussed, we perform an analysis using convolutional neural networks (CNNs) with ray-traced N-body simulations and reach a similar conclusion.

While higher-order statistics provide an idea of the information content in a cosmic field, (i) they are not necessarily optimal – perhaps there is much more information to be found – and (ii) it is possible that much of the information in the higher-order statistics is also available in the 2-point correlation function. In this work we address both of these points by (i) using a CNN to extract all the information at the field-level and (ii) studying how much information there is beyond the 2-point by scaling out the 2-point information from the maps before training the CNN.

CNNs have been employed in many studies to estimate cosmological parameters in weak lensing and other areas of cosmology, generally giving tighter constraints than other higher-order statistics [38–49], but have yet to be applied in the context of neutrino mass. Here, we train a classifier to distinguish between simulations with neutrino masses of  $M_\nu = 0\text{eV}$  and  $0.1\text{eV}$  (the minimum value allowed for the inverted hierarchy). We perform classification instead of full parameter inference due to the lack of currently available weak lensing simulations for neutrino mass parameter inference. We study the performance of the network as a function of maximum scale  $\ell_{\text{max}}$ , source redshift  $z$ , and noise level  $\sigma_\kappa$ .

Machine learning models are often used as black boxes, but with the wide use of neural networks in astrophysics and cosmology there has been a recent drive to interpret and explain what the networks are learning [50–56]. We study Integrated Gradient saliency maps [57] to determine which parts of weak lensing convergence maps are most informative about neutrino mass. This is of particular interest as there is much discussion about underdense regions of cosmic structure, known as voids, containing a large amount of information regarding neutrino mass in the context of 3-dimensional LSS [28–30, 37, 58, 59].

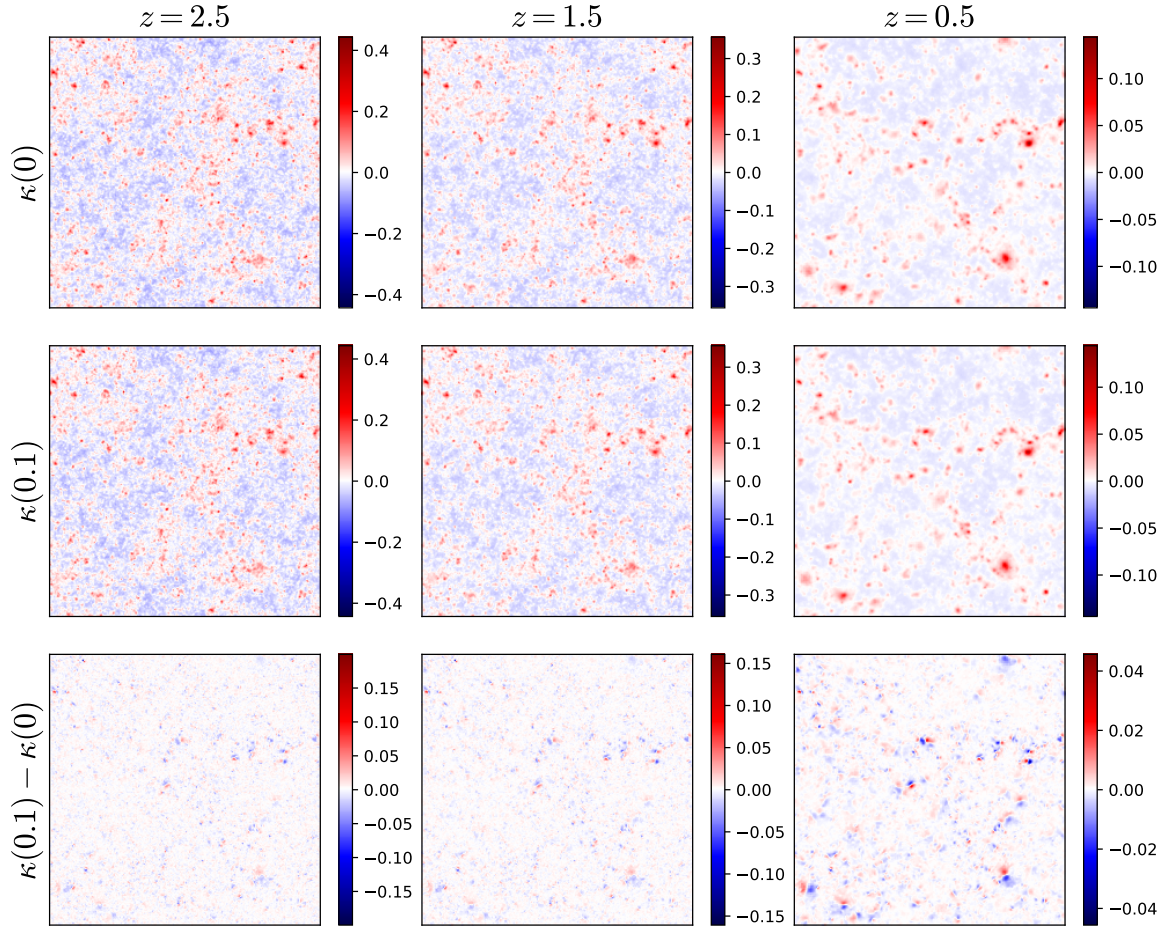
The paper is organized as follows. Section 2 describes the simulation data and neural network, Section 3 reports our results, and Section 4 provides conclusions and discussion.

## 2 Method

In this section, we describe the data, neural network architecture, and analysis methods.

### 2.1 Data


We use the **MassiveNuS** simulations [16], which are based on a modified version of the N-body code **Gadget-2** [60]. We use a set of 10,000 pairs of simulations with  $M_\nu = 0\text{eV}$  and  $0.1\text{eV}$ . The other cosmological parameters are fixed to  $A_s = 2.1 \times 10^{-9}$ ,  $\Omega_m = 0.3$ ,  $\Omega_b = 0.05$ , and  $h = 0.7$ . Each simulation contains  $1024^3$  particles in a  $(512 \text{ Mpc}/h)^3$  box, thus the Nyquist wavenumber is  $k_{\text{Nyq}} = 2\pi h/\text{Mpc}$ . Weak lensing convergence maps  $\kappa$  are produced by ray-tracing using the flat-sky approximation over a  $(3.5 \text{ deg})^2$  solid angle. We consider three source redshifts of  $z = 2.5, 1.5, 0.5$ . Figure 1 shows a particular realization of convergence maps; the first row shows a map for the  $0\text{eV}$  cosmology, the second row for a seed-matched  $0.1\text{eV}$  cosmology, and the final row shows the difference between the two. It can be seen that there is stronger, more nonlinear, structure for larger source redshifts – in particular there are sharper overdense regions and larger underdense regions, leading to sharper differences caused by the difference in neutrino mass, particularly in overdense regions.



**Figure 1.** A particular realization of convergence maps  $\kappa$  from MassiveNuS. The first row shows the maps for a 0eV cosmology, the second row for a seed-matched 0.1eV cosmology, and the final row shows the difference between the two. Different columns show different source redshifts.

Before feeding the data to the CNN we apply augmentations for various purposes:

1. **Addition of Gaussian Noise:** Gaussian noise with zero mean was added to mimic observational noise. We considered two standard deviations of  $\sigma_\kappa = 0.01$  and  $0.1$  (in the same units as  $\kappa$ ), which is around the order of magnitude of noise expected for an LSST-like survey with a galaxy number density of  $n_{\text{gal}} \simeq 5 \text{ arcmin}^{-2}$  and a shape noise of  $\sigma_\chi \simeq 0.3$  [32]. We also consider the noiseless maps.
2. **Cutting in Fourier Space:** The images were transformed to Fourier space and cut at various scales in factors of 2 of the Nyquist,  $k_{\text{max}} = 2\pi, \pi, \pi/2, \pi/4, \pi/8, \pi/16 (h/\text{Mpc})^2$ , and then transformed back into configuration space. As the maps are produced using the flat-sky approximation, we use a 2-dimensional wavenumber, denoted  $k$ , to represent the scale. To connect with curved-sky intuition, we show the relation between  $k$  and multipole  $\ell$  in Figure 2 for the various source redshifts considered.
3. **Rescaling to remove the 2-point information:** As well as considering the information in the convergence maps, we additionally analyse the maps after rescaling the



**Figure 2.** The relation between 2-dimensional wavevector  $k$  and multipole  $\ell$  for the source redshifts considered in this work. The dotted lines represent the various  $k_{\text{max}}$  cuts we perform in the analysis.

0eV maps to fake the effect of neutrinos at the level of the power spectrum. We create the maps with fake  $\nu$ s at the 2-point level by multiplying all of the 0eV convergence maps by the square-root of the ratio of the 0.1eV to 0eV power spectra as follows:  $\kappa(0) \rightarrow \sqrt{\frac{P(0.1)}{P(0)}} \kappa(0)$ , where  $P(M_\nu)$  is the 2-dimensional power spectrum averaged over all simulations with neutrino mass  $M_\nu$ . This effectively matches the 2-point information in the 0eV and 0.1eV cosmologies and enables study of the beyond 2-point information.

4. **Normalization:** The maps were normalized to have zero mean and unit standard deviation.
5. **Data Augmentation:** The images were randomly flipped and rotated to increase the variability and robustness of the training set.

## 2.2 Neural Network

We use a ResNet-18 network [61] to perform the analysis. ResNet architectures have commonly been used to achieve start-of-the-art performance on various computer-vision deep-learning tasks over the last few years. ResNets are a group of CNN models made to skip one or more layers to allow for the network to learn residual functions with reference to the layer input rather than to unreferenced functions. This is achieved via skip, or residual, connections, and enables training a network with a large depth. The architecture of the model starts with an initial convolutional layer followed by a max-pooling layer, which reduces the spatial dimensions of the input image. This is followed by a series of residual blocks grouped into four stages, each containing two residual blocks with increasing filter sizes. Each residual

block consists of two convolutional layers, batch normalization and ReLU activation, with shortcut connections that add the input of the block to its output. The final stages of the network include a global average pooling layer and a fully connected layer that outputs the class probabilities. We choose a ResNet-18 as [49] reported this to be sufficient for weak lensing cosmology applications, finding that using a deeper ResNet, or an architecture with a larger receptive field, did not improve the performance.

We modified the ResNet-18 for our applications to take only 1 channel of input, instead of the default 3. We also replaced the final layer of the network by 3 fully connected layers which output a scalar, followed by a dropout, LeakyReLU, and sigmoid layer. We then apply a binary cross entropy loss function to train a classifier between the 0eV and 0.1eV neutrino mass simulations. For optimization we used the ADAM optimizer with an exponential learning rate scheduler, using early stopping to avoid overfitting. We report results for a batch size of 32, but also found that using a batch size of 64 or 128 made no improvement to the model. We use the pre-trained ResNet-18 weights as the starting point for training, but then refit all the weights during training. The data set was divided into training, validation, and testing sets according to a 70:15:15 split.

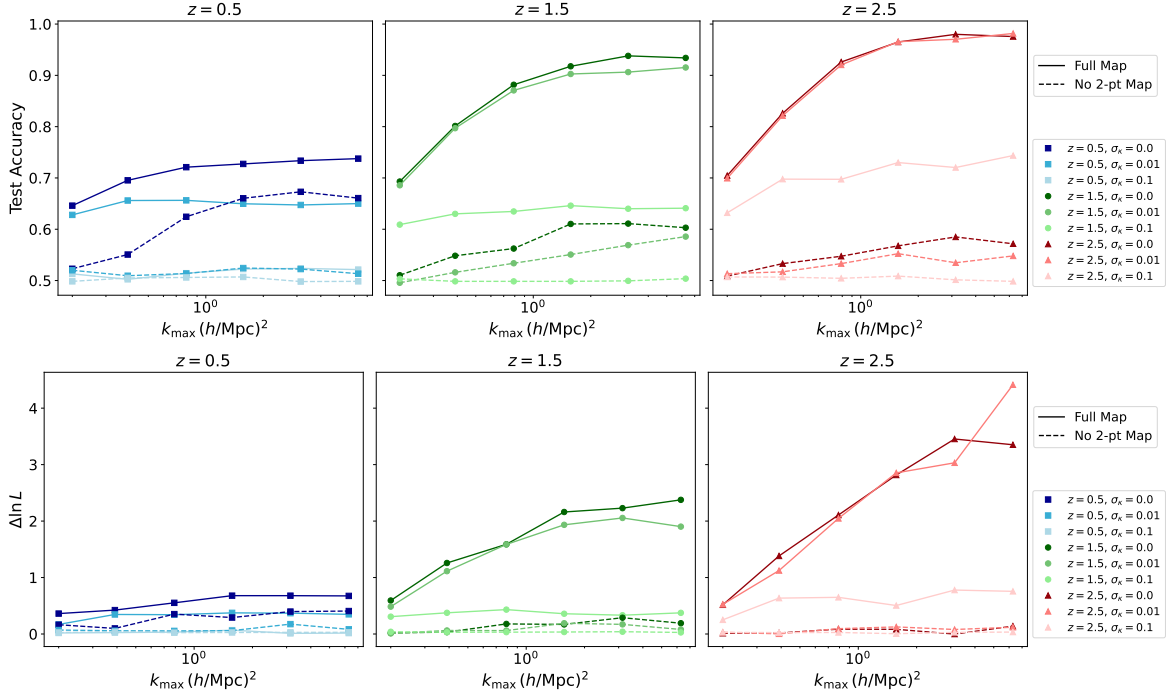
We trained the neural network for all the different source redshifts, scale cuts, noise levels, and power spectrum augmentations detailed in the previous subsection. We used `optuna` [62] to optimize the hyperparameters of the network and training, in particular optimizing the dropout rate, weight decay, learning rate, and learning rate scheduler decay rate. We performed 200 trials for every training and took the model(s) with the highest validation accuracy.

Subsequently, we applied an interpretability metric, the Integrated Gradient saliency map [57], to identify which pixels in the map the CNN is paying the closest attention to. The Integrated Gradient quantifies how the model output (the neutrino mass) changes with respect to the model inputs (the pixel values) by propagating the gradient through the network parameters. This allowed us to determine which parts of the maps contains the most information regarding neutrino mass. We implemented the saliency map using the `Captum` library [63].

### 3 Results

The upper row of Figure 3 shows the classification accuracy evaluated on the test set for the different scale cuts  $k_{\text{max}}$ , source redshifts  $z$ , and noise levels  $\sigma_\kappa$ , as well as the effect of scaling out the 2-point information. It can be seen that the accuracy of the network improves as smaller, nonlinear, scales are included; this is to be expected as the effect of massive neutrinos becomes more prominent on small scales. It can also be seen that the accuracy is higher at higher redshifts; this is to be expected as, the further away the light source is, the further the light travels before reaching the observer, experiencing more lensing and thus producing a larger signal. It can also be seen that increasing noise reduces the accuracy, which is to be expected as the massive neutrino signal is distorted by the noise. The network is able to reach almost 100% accuracy for a source redshift of 2.5 with  $\ell_{\text{max}} \simeq 5,000$ . However, most interestingly, the accuracy of the network drops to almost 50% after scaling out the 2-point information from the maps for all source redshifts and scales, except for low noise cases at small scales where it reaches  $\sim 60\%$ . This implies that the network is almost unable to tell the difference between the 0eV and 0.1eV maps, suggesting that most of the field-level information can be found using the power spectrum alone. One notable exception is in the





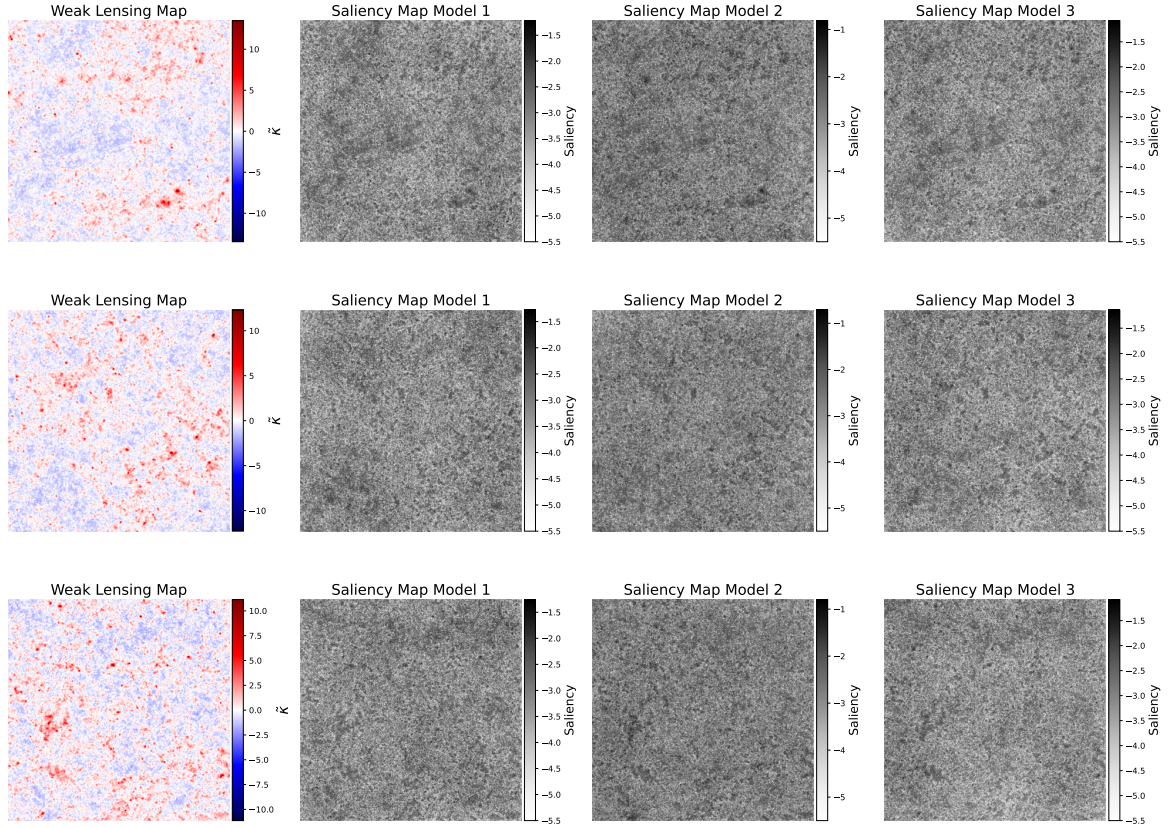
**Figure 3.** *Top:* classification accuracy evaluated on the test set as a function of  $k_{\max}$ . *Bottom:* difference in log likelihood between a  $M_\nu = 0.1$  and  $0 \text{ eV}$  for observations with  $M_\nu = 0.1 \text{ eV}$  (taken from the test set) as a function of  $k_{\max}$ . Each panel and base color corresponds to a different source redshift  $z$ . The difference in color shades corresponds to the noise  $\sigma_\kappa$ , with lighter shades corresponding to higher noise. Solid lines represent the accuracy for a model trained on the full convergence maps, while dashed lines are for maps that have had the 2-point information scaled out. It can be seen for full maps that the accuracy and difference in likelihood are typically higher at smaller scales, higher redshifts, and lower noise. However, when the 2-point information is removed the accuracy and difference in likelihood are consistently significantly worse – almost 50% and 0 respectively – indicating that much of the field-level information is contained in the 2-point alone.

case of  $z = 0.5$  where the loss of accuracy from scaling out the 2-point information is relatively small on small scales, implying that – even though the total information is relatively small (implied by the low accuracy) – a large fraction of this information exists beyond the 2-point.

Following the insight provided by the classification accuracy, we now consider a more statistically interpretable quantity, the likelihood ratio. Given the use of a binary cross entropy loss function, we can relate the output of the neural network  $f(\kappa_{\text{obs}})$  to the likelihood ratio between the  $M_\nu = 0.1 \text{ eV}$  and  $M_\nu = 0$  hypotheses as follows,

$$e^{\Delta \ln L} \equiv \frac{L(M_\nu = 0.1 \text{ eV} | \kappa_{\text{obs}})}{L(M_\nu = 0 | \kappa_{\text{obs}})} = \frac{f(\kappa_{\text{obs}})}{1 - f(\kappa_{\text{obs}})}, \quad (3.1)$$

where  $L(M_\nu | \kappa_{\text{obs}})$  is the likelihood and  $\kappa_{\text{obs}}$  is an observed map from a cosmology with true neutrino mass corresponding to the ‘1 label’ in the classification; here  $0.1 \text{ eV}$  [64]. The lower row of Figure 3 shows how the difference in log likelihood  $\Delta \ln L$  varies with  $k_{\max}$ ,  $z$ , and  $\sigma_\kappa$ , for observations with  $M_\nu = 0.1 \text{ eV}$  (taken from the test set). The overall trends are similar to the classification accuracy, with higher  $k_{\max}$ , higher  $z$ , and lower  $\sigma_\kappa$ , leading to a higher likelihood difference. In the most optimistic case of  $z = 2.5$  and  $k_{\max} = 2\pi h/\text{Mpc}$  ( $\ell_{\max} \simeq 10^4$ ) with low noise, the difference in log likelihood is  $\Delta \ln L \simeq 4.5$ . For a Gaussian




**Figure 4.** Visualization of the logarithm of the absolute value of the Integrated Gradient saliency maps for the case of  $z = 2.5$ ,  $k_{\max} = \pi/2$ , and  $\sigma_{\kappa} = 0.01$  (without scaling out the 2-point information). Each row displays a different (mean-subtracted and variance-scaled) weak lensing convergence map in the left column, followed by three saliency maps generated by the best three models from the `optuna` trials. Higher numbers in the saliency map implies that network is paying more attention to that region of the map. The saliency maps show a wide variability in terms of regions of the map they focus on, in some cases focusing on underdense regions, and in other cases focusing on overdense regions.

likelihood this would correspond to  $\Delta\chi^2 \simeq 9$  and thus a  $3\sigma$  detection.<sup>1</sup> This greatly decreases as lower source redshifts, lower scales, or higher noise is considered, dropping to  $\Delta \ln L \simeq 0$ . Most interestingly, it can also be seen that the likelihood difference is negligible when the 2-point information is removed in *all* cases, implying that there is little information beyond the 2-point.

Figure 4 shows saliency maps using the Integrated Gradient method. This enables interpretation of what the network is learning and which parts of the map the model is most sensitive to. Each row displays a different weak lensing convergence map in the left column,

<sup>1</sup>We note that, while it is instructive to consider the likelihood ratio to build intuition, this neglects any prior uncertainty on the neutrino mass as it only considers a  $M_{\nu} = 0.1\text{eV}$  alternative hypothesis relative to a  $M_{\nu} = 0\text{eV}$  null hypothesis. In a Bayesian analysis one would marginalize over a prior which spans a range of values of  $M_{\nu}$ , carefully taking into account volume effects [see e.g. 65, 66]. This in turn would reduce the significance of any detection of neutrino mass compared to the likelihood analysis here which corresponds to a delta-function prior at the observed value.





**Figure 5.** The pixel-averaged value of the logarithm of the absolute value of the Integrated Gradient saliency maps as a function of the (mean-subtracted and variance-scaled) convergence map pixel value  $\tilde{\kappa}$  for the case of  $z = 2.5$ ,  $k_{\max} = \pi/2$ , and  $\sigma_{\kappa} = 0.01$  (without scaling out the 2-point information). Higher saliency values imply that the network is paying more attention to that region of the map, thus the network is paying most attention to extreme values of  $\kappa$ , i.e. the most overdense and underdense regions.

followed by three saliency maps generated by the best three models from the `optuna` trials. Higher numbers in the saliency map implies that the network is paying more attention to that region of the map. We consider the case of  $z = 2.5$ ,  $k_{\max} = \pi/2$  and  $\sigma_{\kappa} = 0.01$  where the accuracy of the network is high, and thus the saliency maps will be informative. It can be seen that there is much variability in the saliency maps, with the networks sometimes focusing more on underdense regions and sometimes on overdense regions. This implies that there is information in various features of the map, and that using an ensemble approach could allow us to extract more information, although this is beyond the scope of this work. To study the saliency more quantitatively, Figure 5 shows the average saliency in a pixel as a function of the (mean-subtracted and variance-scaled) convergence of the pixel  $\tilde{\kappa}$ . It can be seen that the saliency is highest for extreme values of  $\tilde{\kappa}$ , i.e. in the most overdense and underdense regions. This shows the complementarity of cluster-like and void-like regions, as found in 3-dimensional LSS examples which found complementarity between halo and void statistics [28, 29]. We find similar results in the noiseless case, but after increasing the noise to  $\sigma_{\kappa} = 0.1$  the higher saliency in the clusters is washed out. This is different to the

trend observed for  $\Omega_m$  and  $\sigma_8$ , where underdense regions were found to be more salient [54], showing promise for the prospect of breaking the degeneracy between these parameters.

## 4 Conclusions

This work quantified the ability of CNNs to extract information regarding neutrino mass from weak lensing convergence maps. We studied the information content as a function of scale cut, source redshift, and noise level for an LSST-like survey. We also investigated the effect of scaling out the 2-point information from the maps to study the beyond-2-point information. We found that a ResNet-18-based classifier was able to distinguish between a 0eV and 0.1eV neutrino mass cosmology with a high accuracy that improves as higher source redshifts and smaller scales are considered, reaching almost 100% for a source redshift of 2.5 with  $\ell_{\text{max}} \simeq 5,000$ . However, the accuracy of the network dropped to almost 50% after scaling out the 2-point information from the maps, implying that most of the field-level information can be found using the power spectrum alone, in agreement with the simple cross-correlation arguments of [37]. We also went beyond using the neural network as a black box, by using Integrated Gradient saliency maps to understand which parts of the map the network is learning from. We found that the network pays most attention to the most overdense and underdense regions, with no clear preference for either. This suggests the two regions to be complementary, as found in studies that explicitly consider cluster-based and void-based summary statistics [28, 29].

In this analysis we were restricted to using a classifier due to the availability of simulations, but in the future, when more simulations are produced, it would be interesting further work to do full parameter inference, studying the beyond 2-point information in a similar manner to [67] which considered  $\Omega_m$  and  $\sigma_8$  for halo clustering. While we have performed extensive hyperparameter tuning to facilitate the optimality of the CNN, it would also be interesting future work to study the field-level neutrino mass information using different architectures and methodologies, such as differentiable forward modelling to reconstruct the initial conditions of the Universe [68–85] and normalizing flows [86–88]; although we do not expect this to change the conclusions. Furthermore, while we considered the information content in weak lensing convergence maps alone, it would be interesting to study the information when cross correlating convergence maps with other probes from LSS and the CMB using cross-survey simulations such as [89, 90] as this could provide a promising source of beyond-2-point information. Finally, it would be worthwhile to repeat this analysis while including baryonic effects to study the degeneracy between massive neutrinos and baryons at the field level [91].

## Acknowledgments

We thank Vanessa Böhm, Fraçoise Lanusse, Jia Liu, Uroš Seljak, Francisco Villaescusa-Navarro, and Benjamin Wandelt for insightful discussion. MG also thanks Vanessa Böhm for her initial mentorship. This research was conducted as part of the N3AS Undergraduate Research Program at UC Berkeley, supported by the NSF Award Number 2020275 and Heising Simons Grant 2017-228. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a Department of Energy Office of Science User Facility using NERSC award ASCR-ERCAP0029232. The computations reported in this paper were (in part) performed using resources made available by the Flatiron Institute.

The analysis of the simulations has made use of the *Pylians* library, publicly available at <https://github.com/franciscovillaescusa/Pylians3>.

## References

- [1] P. Collaboration, N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi et al., *Planck 2018 results. vi. cosmological parameters*, 2018.
- [2] M. Aker, K. Altenmüller, M. Arenz, M. Babutzka, J. Barrett, S. Bauer et al., *Improved upper limit on the neutrino mass from a direct kinematic method by katrin*, *Physical Review Letters* **123** (2019) .
- [3] Y. Fukuda, T. Hayakawa, E. Ichihara, K. Inoue, K. Ishihara, H. Ishino et al., *Evidence for oscillation of atmospheric neutrinos*, *Physical Review Letters* **81** (1998) 1562–1567.
- [4] SNO collaboration, *Direct evidence for neutrino flavor transformation from neutral current interactions in the Sudbury Neutrino Observatory*, *Phys. Rev. Lett.* **89** (2002) 011301 [[nucl-ex/0204008](#)].
- [5] KAMLAND collaboration, *Measurement of neutrino oscillation with KamLAND: Evidence of spectral distortion*, *Phys. Rev. Lett.* **94** (2005) 081801 [[hep-ex/0406035](#)].
- [6] K2K COLLABORATION collaboration, *Measurement of neutrino oscillation by the k2k experiment*, *Phys. Rev. D* **74** (2006) 072003.
- [7] F.P. An, J.Z. Bai, A.B. Balantekin, H.R. Band, D. Beavis, W. Beriguete et al., *Observation of electron-antineutrino disappearance at daya bay*, *Phys. Rev. Lett.* **108** (2012) 171803.
- [8] LSST Science Collaboration, P.A. Abell, J. Allison, S.F. Anderson, J.R. Andrew, J.R.P. Angel et al., *LSST Science Book, Version 2.0*, *arXiv e-prints* (2009) [arXiv:0912.0201](#) [[0912.0201](#)].
- [9] R. Laureijs, J. Amiaux, S. Arduini, J.L. Auguères, J. Brinchmann, R. Cole et al., *Euclid Definition Study Report*, *arXiv e-prints* (2011) [arXiv:1110.3193](#) [[1110.3193](#)].
- [10] SPHEREx collaboration, *Science Impacts of the SPHEREx All-Sky Optical to Near-Infrared Spectral Survey II: Report of a Community Workshop on the Scientific Synergies Between the SPHEREx Survey and Other Astronomy Observatories*, [1805.05489](#).
- [11] D. Spergel, N. Gehrels, C. Baltay, D. Bennett, J. Breckinridge, M. Donahue et al., *Wide-Field Infrared Survey Telescope-Astrophysics Focused Telescope Assets WFIRST-AFTA 2015 Report*, *arXiv e-prints* (2015) [arXiv:1503.03757](#) [[1503.03757](#)].
- [12] DESI Collaboration, A. Aghamousa, J. Aguilar, S. Ahlen, S. Alam, L.E. Allen et al., *The DESI Experiment Part I: Science, Targeting, and Survey Design*, *arXiv e-prints* (2016) [arXiv:1611.00036](#) [[1611.00036](#)].
- [13] N. Tamura et al., *Prime Focus Spectrograph (PFS) for the Subaru Telescope: Overview, recent progress, and future perspectives*, *Proc. SPIE Int. Soc. Opt. Eng.* **9908** (2016) 99081M [[1608.01075](#)].
- [14] SIMONS OBSERVATORY collaboration, *The Simons Observatory: Science goals and forecasts*, *JCAP* **02** (2019) 056 [[1808.07445](#)].
- [15] CMB-S4 collaboration, *CMB-S4 Science Book, First Edition*, [1610.02743](#).
- [16] J. Liu, S. Bird, J.M.Z. Matilla, J.C. Hill, Z. Haiman, M.S. Madhavacheril et al., *MassiveNuS: Cosmological massive neutrino simulations*, *Journal of Cosmology and Astroparticle Physics* **2018** (2018) .
- [17] S. Bird, Y. Ali-Haïmoud, Y. Feng and J. Liu, *An Efficient and Accurate Hybrid Method for Simulating Non-Linear Neutrino Structure*, *Monthly Notices of the Royal Astronomical Society* (2018) [[1803.09854](#)].

- [18] F. Villaescusa-Navarro, A. Banerjee, N. Dalal, E. Castorina, R. Scoccimarro, R. Angulo et al., *The imprint of neutrinos on clustering in redshift space*, *The Astrophysical Journal* **861** (2018) 53.
- [19] A. Banerjee, D. Powell, T. Abel and F. Villaescusa-Navarro, *Reducing noise in cosmological  $n$ -body simulations with neutrinos*, *Journal of Cosmology and Astroparticle Physics* **2018** (2018) 028–028.
- [20] A.E. Bayer, A. Banerjee and Y. Feng, *A fast particle-mesh simulation of non-linear cosmological structure formation with massive neutrinos*, *J. Cosmology Astropart. Phys.* **2021** (2021) 016.
- [21] J. DeRose, N. Kokron, A. Banerjee, S.-F. Chen, M. White, R. Wechsler et al., *Aemulus  $\nu$ : precise predictions for matter and biased tracer power spectra in the presence of neutrinos*, *JCAP* **07** (2023) 054 [[2303.09762](#)].
- [22] J.M. Sullivan, J.D. Emberson, S. Habib and N. Frontiere, *Improving initialization and evolution accuracy of cosmological neutrino simulations*, *JCAP* **06** (2023) 003 [[2302.09134](#)].
- [23] C. Hahn, F. Villaescusa-Navarro, E. Castorina and R. Scoccimarro, *Constraining  $m_\nu$  with the bispectrum. part i. breaking parameter degeneracies*, *Journal of Cosmology and Astroparticle Physics* **2020** (2020) 040–040.
- [24] C. Hahn and F. Villaescusa-Navarro, *Constraining  $m_\nu$  with the bispectrum ii: The total information content of the galaxy bispectrum*, 2020.
- [25] C. Uhlemann, O. Friedrich, F. Villaescusa-Navarro, A. Banerjee and S. Codis, *Fisher for complements: extracting cosmology and neutrino mass from the counts-in-cells PDF*, *Monthly Notices of the Royal Astronomical Society* **495** (2020) 4006 [[1911.11158](#)].
- [26] E. Massara, F. Villaescusa-Navarro, S. Ho, N. Dalal and D.N. Spergel, *Using the Marked Power Spectrum to Detect the Signature of Neutrinos in Large-Scale Structure*, *Phys. Rev. Lett.* **126** (2021) 011301 [[2001.11024](#)].
- [27] G. Valogiannis and C. Dvorkin, *Towards an optimal estimation of cosmological parameters with the wavelet scattering transform*, *Phys. Rev. D* **105** (2022) 103534 [[2108.07821](#)].
- [28] A.E. Bayer, F. Villaescusa-Navarro, E. Massara, J. Liu, D.N. Spergel, L. Verde et al., *Detecting Neutrino Mass by Combining Matter Clustering, Halos, and Voids*, *ApJ* **919** (2021) 24 [[2102.05049](#)].
- [29] C.D. Kreisch, A. Pisani, F. Villaescusa-Navarro, D.N. Spergel, B.D. Wandelt, N. Hamaus et al., *The GIGANTES dataset: precision cosmology from voids in the machine learning era*, *arXiv e-prints* (2021) arXiv:2107.02304 [[2107.02304](#)].
- [30] A.E. Bayer, J. Liu, C.D. Kreisch and A. Pisani, *Significance of void shape: Neutrino mass from Voronoi void halos?*, *Phys. Rev. D* **110** (2024) L061305 [[2405.12302](#)].
- [31] C.D. Kreisch, A. Pisani, C. Carbone, J. Liu, A.J. Hawken, E. Massara et al., *Massive neutrinos leave fingerprints on cosmic voids*, *Monthly Notices of the Royal Astronomical Society* **488** (2019) 4413 [[1808.07464](#)].
- [32] J. Liu and M.S. Madhavacheril, *Constraining neutrino mass with the tomographic weak lensing one-point probability distribution function and power spectrum*, *Phys. Rev. D* **99** (2019) 083508.
- [33] Z. Li, J. Liu, J.M.Z. Matilla and W.R. Coulton, *Constraining neutrino mass with tomographic weak lensing peak counts*, *Phys. Rev. D* **99** (2019) 063527 [[1810.01781](#)].
- [34] W.R. Coulton, J. Liu, M.S. Madhavacheril, V. Böhm and D.N. Spergel, *Constraining neutrino mass with the tomographic weak lensing bispectrum*, *Journal of Cosmology and Astroparticle Physics* **5** (2019) 043 [[1810.02374](#)].

- [35] G.A. Marques, J. Liu, J.M. Zorrilla Matilla, Z. Haiman, A. Bernui and C.P. Novaes, *Constraining neutrino mass with weak lensing Minkowski Functionals*, *Journal of Cosmology and Astroparticle Physics* **2019** (2019) 019 [[1812.08206](#)].
- [36] V. Ajani, A. Peel, V. Pettorino, J.-L. Starck, Z. Li and J. Liu, *Constraining neutrino masses with weak-lensing starlet peak counts*, 2020.
- [37] A.E. Bayer, A. Banerjee and U. Seljak, *Beware of fake  $\nu$ s: The effect of massive neutrinos on the nonlinear evolution of cosmic structure*, *Phys. Rev. D* **105** (2022) 123510 [[2108.04215](#)].
- [38] A. Gupta, J.M.Z. Matilla, D. Hsu and Z. Haiman, *Non-gaussian information from weak lensing data via deep learning*, *Physical Review D* **97** (2018) 103515.
- [39] D. Ribli, B.Á. Pataki, J.M. Zorrilla Matilla, D. Hsu, Z. Haiman and I. Csabai, *Weak lensing cosmology with convolutional neural networks on noisy data*, *Monthly Notices of the Royal Astronomical Society* **490** (2019) 1843.
- [40] J. Fluri, T. Kacprzak, A. Refregier, A. Amara, A. Lucchi and T. Hofmann, *Cosmological constraints from noisy convergence maps through deep learning*, *Physical Review D* **98** (2018) 123518.
- [41] J. Fluri, T. Kacprzak, A. Lucchi, A. Refregier, A. Amara, T. Hofmann et al., *Cosmological constraints with deep learning from kids-450 weak lensing maps*, *Physical Review D* **100** (2019) .
- [42] F. Villaescusa-Navarro, B.D. Wandelt, D. Anglés-Alcázar, S. Genel, J.M.Z. Matilla, S. Ho et al., *Neural networks as optimal estimators to marginalize over baryonic effects*, 2020.
- [43] N. Jeffrey, J. Alsing and F. Lanusse, *Likelihood-free inference with neural compression of  $des$   $sv$  weak lensing map statistics*, *Monthly Notices of the Royal Astronomical Society* **501** (2020) 954–969.
- [44] J.M.Z. Matilla, M. Sharma, D. Hsu and Z. Haiman, *Interpreting deep learning models for weak lensing*, *Phys. Rev. D* **102** (2020) 123506.
- [45] T.L. Makinen, T. Charnock, J. Alsing and B.D. Wandelt, *Lossless, scalable implicit likelihood inference for cosmological fields*, *Journal of Cosmology and Astroparticle Physics* **2021** (2021) 049.
- [46] T. Lu, Z. Haiman and J.M. Zorrilla Matilla, *Simultaneously constraining cosmology and baryonic physics via deep learning from weak lensing*, *Monthly Notices of the Royal Astronomical Society* **511** (2022) 1518–1528.
- [47] T. Lu, Z. Haiman and X. Li, *Cosmological constraints from hsc survey first-year data using deep learning*, *Monthly Notices of the Royal Astronomical Society* **521** (2023) 2050–2066.
- [48] SIMBIG collaboration, *Field-level simulation-based inference of galaxy clustering with convolutional neural networks*, *Phys. Rev. D* **109** (2024) 083536 [[2310.15256](#)].
- [49] D. Sharma, B. Dai and U. Seljak, *A comparative study of cosmological constraints from weak lensing using Convolutional Neural Networks*, **2403.03490**.
- [50] L. Lucie-Smith, H.V. Peiris, A. Pontzen and M. Lochner, *Machine learning cosmological structure formation*, *Mon. Not. Roy. Astron. Soc.* **479** (2018) 3405 [[1802.04271](#)].
- [51] L. Lucie-Smith, H.V. Peiris and A. Pontzen, *An interpretable machine learning framework for dark matter halo formation*, *Mon. Not. Roy. Astron. Soc.* **490** (2019) 331 [[1906.06339](#)].
- [52] L. Lucie-Smith, H.V. Peiris, A. Pontzen, B. Nord and J. Thiyaalingam, *Deep learning insights into cosmological structure formation*, *Phys. Rev. D* **109** (2024) 063524 [[2011.10577](#)].
- [53] L. Lucie-Smith, H.V. Peiris and A. Pontzen, *Explaining Dark Matter Halo Density Profiles with Neural Networks*, *Phys. Rev. Lett.* **132** (2024) 031001 [[2305.03077](#)].
- [54] J.M. Zorrilla Matilla, M. Sharma, D. Hsu and Z. Haiman, *Interpreting deep learning models for weak lensing*, *Phys. Rev. D* **102** (2020) 123506 [[2007.06529](#)].



- [55] P. Bhambra, B. Joachimi and O. Lahav, *Explaining deep learning of galaxy morphology with saliency mapping*, *Monthly Notices of the Royal Astronomical Society* **511** (2022) 5032–5041.
- [56] W. You, H. Qu, M. Gatti, B. Jain and E. Wong, *Sum-of-parts models: Faithful attributions for groups of features*, 2023.
- [57] M. Sundararajan, A. Taly and Q. Yan, *Axiomatic attribution for deep networks*, 2017.
- [58] C.D. Kreisch, F.-Y. Cyr-Racine and O. Doré, *Neutrino puzzle: Anomalies, interactions, and cosmological tensions*, *Phys. Rev. D* **101** (2020) 123505 [1902.00534].
- [59] L. Thiele, E. Massara, A. Pisani, C. Hahn, D.N. Spergel, S. Ho et al., *Neutrino mass constraint from an Implicit Likelihood Analysis of BOSS voids*, 2307.07555.
- [60] V. Springel, *The cosmological simulation code GADGET-2*, *Monthly Notices of the Royal Astronomical Society* **364** (2005) 1105 [arXiv:astro-ph/0505010].
- [61] K. He, X. Zhang, S. Ren and J. Sun, *Deep residual learning for image recognition*, 2015.
- [62] T. Akiba, S. Sano, T. Yanase, T. Ohta and M. Koyama, *Optuna: A next-generation hyperparameter optimization framework*, 2019.
- [63] N. Kokhlikyan, V. Miglani, M. Martin, E. Wang, B. Alsallakh, J. Reynolds et al., *Captum: A unified and generic model interpretability library for pytorch*, 2020.
- [64] N. Jeffrey and B.D. Wandelt, *Evidence Networks: simple losses for fast, amortized, neural Bayesian model comparison*, *Mach. Learn. Sci. Tech.* **5** (2024) 015008 [2305.11241].
- [65] A.E. Bayer and U. Seljak, *The look-elsewhere effect from a unified Bayesian and frequentist perspective*, *J. Cosmology Astropart. Phys.* **2020** (2020) 009 [2007.13821].
- [66] L.T. Hergt, W.J. Handley, M.P. Hobson and A.N. Lasenby, *Bayesian evidence for the tensor-to-scalar ratio  $r$  and neutrino masses  $m_\nu$ : Effects of uniform vs logarithmic priors*, *Phys. Rev. D* **103** (2021) 123511 [2102.11511].
- [67] BEYOND-2PT collaboration, *A Parameter-Masked Mock Data Challenge for Beyond-Two-Point Galaxy Clustering Statistics*, 2405.02252.
- [68] J. Jasche and B.D. Wandelt, *Bayesian physical reconstruction of initial conditions from large-scale structure surveys*, *MNRAS* **432** (2013) 894 [1203.3639].
- [69] U. Seljak, G. Aslanyan, Y. Feng and C. Modi, *Towards optimal extraction of cosmological information from nonlinear data*, *Journal of Cosmology and Astroparticle Physics* **2017** (2017) 009–009.
- [70] J. Jasche and G. Lavaux, *Physical Bayesian modelling of the non-linear matter distribution: New insights into the nearby universe*, *A&A* **625** (2019) A64 [1806.11117].
- [71] F. Schmidt, *An  $n$ -th order Lagrangian Forward Model for Large-Scale Structure*, *JCAP* **04** (2021) 033 [2012.09837].
- [72] V. Böhm, Y. Feng, M.E. Lee and B. Dai, *MADLens, a python package for fast and differentiable non-Gaussian lensing simulations*, *Astron. Comput.* **36** (2021) 100490 [2012.07266].
- [73] C. Modi, F. Lanusse and U. Seljak, *FlowPM: Distributed TensorFlow implementation of the FastPM cosmological  $N$ -body solver*, *Astronomy and Computing* **37** (2021) 100505 [2010.11847].
- [74] N. Porqueres, A. Heavens, D. Mortlock and G. Lavaux, *Bayesian forward modelling of cosmic shear data*, *Mon. Not. Roy. Astron. Soc.* **502** (2021) 3035 [2011.07722].
- [75] N. Porqueres, A. Heavens, D. Mortlock and G. Lavaux, *Lifting weak lensing degeneracies with a field-based likelihood*, *Mon. Not. Roy. Astron. Soc.* **509** (2021) 3194 [2108.04825].
- [76] N. Porqueres, A. Heavens, D. Mortlock, G. Lavaux and T.L. Makinen, *Field-level inference of cosmic shear with intrinsic alignments and baryons*, 2304.04785.

- [77] S.S. Boruah, E. Rozo and P. Fiedorowicz, *Map-based cosmology inference with lognormal cosmic shear maps*, *Mon. Not. Roy. Astron. Soc.* **516** (2022) 4111 [[2204.13216](#)].
- [78] S.S. Boruah and E. Rozo, *Map-based cosmology inference with weak lensing – information content and its dependence on the parameter space*, *Mon. Not. Roy. Astron. Soc.* **527** (2023) L162 [[2307.00070](#)].
- [79] S.S. Boruah, P. Fiedorowicz and E. Rozo, *Bayesian mass mapping with weak lensing data using karmma: Validation with simulations and application to Dark Energy Survey year 3 data*, *Phys. Rev. D* **110** (2024) 023524 [[2403.05484](#)].
- [80] D. Lanzieri, J. Zeghal, T.L. Makinen, A. Boucaud, J.-L. Starck and F. Lanusse, *Optimal Neural Summarisation for Full-Field Weak Lensing Cosmological Implicit Inference*, [2407.10877](#).
- [81] J. Zeghal, D. Lanzieri, F. Lanusse, A. Boucaud, G. Louppe, E. Aubourg et al., *Simulation-Based Inference Benchmark for LSST Weak Lensing Cosmology*, [2409.17975](#).
- [82] Y. Li, L. Lu, C. Modi, D. Jamieson, Y. Zhang, Y. Feng et al., *pmwd: A differentiable cosmological particle-mesh  $n$ -body library*, 2022.
- [83] A.E. Bayer, C. Modi and S. Ferraro, *Joint velocity and density reconstruction of the Universe with nonlinear differentiable forward modeling*, *J. Cosmology Astropart. Phys.* **2023** (2023) 046 [[2210.15649](#)].
- [84] A.E. Bayer, U. Seljak and C. Modi, *Field-Level Inference with Microcanonical Langevin Monte Carlo*, in *40th International Conference on Machine Learning*, 7, 2023 [[2307.09504](#)].
- [85] S.-H. Zang and H.-M. Zhu, *Constraining Neutrino Cosmologies with Nonlinear Reconstruction*, *Astrophys. J.* **961** (2024) 160 [[2311.16439](#)].
- [86] S. Hassan et al., *HIFlow: Generating Diverse  $H_i$  Maps and Inferring Cosmology while Marginalizing over Astrophysics Using Normalizing Flows*, *Astrophys. J.* **937** (2022) 83 [[2110.02983](#)].
- [87] B. Dai and U. Seljak, *Translation and rotation equivariant normalizing flow (TRENF) for optimal cosmological analysis*, *Mon. Not. Roy. Astron. Soc.* **516** (2022) 2363 [[2202.05282](#)].
- [88] B. Dai and U. Seljak, *Multiscale Flow for robust and optimal cosmological analysis*, *Proc. Nat. Acad. Sci.* **121** (2024) e2309624121 [[2306.04689](#)].
- [89] Y. Omori, *Agora: Multi-Component Simulation for Cross-Survey Science*, [2212.07420](#).
- [90] A.E. Bayer, Y. Zhong, Z. Li, J. DeRose, Y. Feng and J. Liu, *The HalfDome Multi-Survey Cosmological Simulations:  $N$ -body Simulations*, [2407.17462](#).
- [91] F. Ferlito et al., *The MillenniumTNG Project: the impact of baryons and massive neutrinos on high-resolution weak gravitational lensing convergence maps*, *Mon. Not. Roy. Astron. Soc.* **524** (2023) 5591 [[2304.12338](#)].