

# Neural network reconstruction of non-Gaussian initial conditions from dark matter halos

Jelte Bottema<sup>a,b</sup>, Thomas Flöss<sup>c,d,e,a</sup> and P. Daniel Meerburg<sup>e</sup>

<sup>a</sup>*Kapteyn Astronomical Institute, University of Groningen,  
P.O.Box 800, 9700 AV Groningen, The Netherlands*

<sup>b</sup>*Department of Physics, Stellenbosch University,  
Matieland 7602, South Africa*

<sup>c</sup>*Department of Mathematics, University of Vienna,  
Oskar-Morgenstern-Platz 1, 1090 Vienna, Austria*

<sup>d</sup>*Department of Astrophysics, University of Vienna,  
Türkenschanzstraße 17, 1180 Vienna, Austria*

<sup>e</sup>*Van Swinderen Institute, University of Groningen,  
Nijenborgh 3, 9747 AG Groningen, The Netherlands*

*E-mail: [jbott@astro.rug.nl](mailto:jbott@astro.rug.nl), [thomas.floess@univie.ac.at](mailto:thomas.floess@univie.ac.at),  
[p.d.meerburg@rug.nl](mailto:p.d.meerburg@rug.nl)*

**ABSTRACT:** We develop a machine learning approach to reconstructing the cosmological initial conditions from late-time dark matter halo number density fields in redshift space, with the goal of improving sensitivity to cosmological parameters, and in particular primordial non-Gaussianity. Using an U-Net architecture, our model achieves a cross-correlation accuracy of 44% for scales out to  $k = 0.4$  h/Mpc between reconstructed and true initial conditions of Quijote 1 Gpc<sup>3</sup> simulation boxes with an average halo number density of  $\bar{n} = 4 \times 10^{-4}$  (h/Mpc)<sup>3</sup> in the tracer field at  $z = 0$ . We demonstrate that our reconstruction is likely to be optimal for this setup and that it is highly effective at reducing redshift-space distortions. Using a Fisher analysis, we show that reconstruction improves cosmological parameter constraints derived from the power spectrum and bispectrum. By combining the power spectrum monopole, quadrupole, and bispectrum monopole up to  $k_{\max} = 0.52$  h/Mpc, our joint analysis of pre- and post-reconstructed fields from the Quijote simulation suite finds improved marginalized errors on all cosmological parameters. In particular, reconstruction improves constraints on  $f_{\text{NL}}$  by factors of 1.33, 1.88, and 1.57 for local, equilateral, and orthogonal shapes. Our findings demonstrate the effectiveness of reconstruction in decoupling modes, mitigating redshift-space distortions and maximizing information on cosmology. The

results provide important insights into the amount of cosmological information that can be extracted from small scales, and can potentially be used to complement standard analysis of observational data, upon further development.

KEYWORDS: cosmological parameters from LSS, Machine learning , non-gaussianity, physics of the early universe

ARXIV EPRINT: [2502.11846](https://arxiv.org/abs/2502.11846)

---

**Contents**

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Primordial non-Gaussianity in the dark matter halo density field</b>	<b>4</b>
2.1	Primordial bispectra	4
2.2	Scale-dependent bias	5
<b>3</b>	<b>Reconstruction with neural networks</b>	<b>6</b>
3.1	U-Net architecture	6
3.2	Training data	8
3.3	Training the model	9
3.4	Reconstruction performance	10
<b>4</b>	<b>Cosmological information content of the reconstructed field</b>	<b>12</b>
4.1	Power spectrum and bispectrum estimator	13
4.2	Covariance	14
4.3	Fisher analysis	14
4.4	Marginalization	16
4.5	Parameter constraints	18
<b>5</b>	<b>Conclusions</b>	<b>18</b>
<b>A</b>	<b>Modifying shot noise</b>	<b>21</b>
<b>B</b>	<b>Convergence of Fisher analysis</b>	<b>22</b>

---

**1 Introduction**

Primordial non-Gaussianity (pnG) refers to deviations from the Gaussian distribution of the primordial density fluctuations in the early universe. PnG is crucial for understanding the early universe, particularly during the inflationary epoch [1], providing insights into the physics in this early period and allowing us to test the validity of inflationary models [2]. Until now, pnG has been most precisely constrained using the anisotropies of the CMB [3–8]. These anisotropies are linearly linked to the small density fluctuations caused by inflation, making them an excellent observable for investigating pnG. Observations by the Planck satellite have provided the most stringent constraints on pnG, which suggest close to Gaussian initial conditions [8].

However, many models of inflation predict non-zero levels of pnG. This provides motivation to further study the subject of pnG. Although upcoming CMB experiments are expected to improve the constraints on pnG [9, 10], the CMB has inherent limitations. These limitations stem from its two-dimensional nature and the effects of Silk diffusion damping on small scales, which effectively reduce the number of observable modes [11]. Consequently, to further our understanding of the primordial universe, it is essential to explore and develop additional probes beyond the CMB. These additional methods will be crucial for placing

tighter constraints on pnG and gaining deeper insights into the early universe. The small anisotropies observed in the CMB evolve into the large-scale structure (LSS) of the universe. The three-dimensional comoving volume of the universe, spanning from the Dark Ages to the present, contains exponentially more modes than the CMB, in principle making it a powerful probe of pnG. Dark matter density fluctuations trace the primordial fluctuations from the early universe. The underlying dark matter density distribution can be studied using biased tracers such as galaxies [12, 13] and atomic and molecular spectral lines [14, 15]. Gravity evolves the primordial density fluctuations into the present-day LSS, a process that is highly non-linear due to the non-linear nature of gravity. Progress has been made in addressing the complex theoretical challenges related to studying the non-linear LSS [16–19]. Although current LSS surveys have not yet been able to place constraints on pnG that are as competitive with those provided by CMB observations [20–28], the results are promising and provide ample motivation to further explore and improve LSS as a probe of pnG.

A complication that arises due to non-linearity, is the issue of mode coupling or non-Gaussian covariance [29]. Non-linear gravitational evolution couples modes of different wavelengths, reducing the amount of unique information contained in each mode. Recently, it has been shown that this mode coupling in particular saturates constraints on pnG [30–32]. PnG is generally extracted from N-point correlation functions in Fourier space, in particular, two lowest order correlation functions; the power spectrum and bispectrum. Strong non-Gaussian covariance implies the transfer of information from the power spectrum and bispectrum into higher-order correlation functions and onto smaller scales. Measuring all higher-order correlation functions out to small scales is not feasible, because of the sheer growth in the number of possible wavelength configurations for higher-order correlation functions. However, by reconstructing the initial (linear) field, the dispersed information can be partially retrieved from higher-order correlation functions and smaller scales and reintegrated into the power spectrum and bispectrum at larger scales. Such a reconstruction is standard practice in the context of the BAO peak [33], and is shown to lead to improved parameter constraints [34, 35]. These studies primarily focused on recovering the BAO peak rather than achieving high-quality overall reconstruction of the initial conditions.

High-fidelity reconstruction at the field level can be achieved through techniques such as machine learning. Neural networks trained to emulate N-body simulations have shown promise in learning general properties of gravitational evolution [36]. Furthermore, it has been shown that neural networks are also able to reverse this evolution, reconstructing the initial conditions of a collapsed matter field [37–43]. Our recent study has demonstrated that a neural network-based approach can reconstruct the initial dark matter density distribution from the present-day dark matter density distribution [44], with an accuracy comparable to state-of-the-art iterative reconstruction algorithms [45], and once trained, operates at a fraction of the computational cost. The decoupling of different wavelength modes, due to reconstruction, reduces covariance of correlation functions. This process thus restores more unique information into the power spectrum and bispectrum of the reconstructed density field, coming from higher-order N-point correlation functions and smaller scales of the late-time density field. Using a Fisher forecast, we showed that the reconstructed power spectrum and bispectrum indeed lead to significantly improved constraints on cosmological parameters.

In reality, it is impossible to measure the dark matter distribution directly. Instead, we are limited to observing biased tracers of the underlying dark matter distribution, such as galaxies. Additionally, in previous work, the reconstruction was performed using the real-space dark matter distribution, whereas realistically, observations are obtained in redshift space. In this paper, we extend our application of neural network-based reconstruction techniques to the dark matter halo number density field in redshift space, which serves as a more realistic example of a tracer of large-scale structure. Compared to the dark matter density field, the halo field faces additional complications beyond non-Gaussian covariance. The discrete nature of dark matter halos makes them sparse tracers of the underlying density fields, introducing shot noise. Compared to the dark matter density field, where parameter constraints saturate due to non-Gaussian covariance, the saturation of parameter constraints in dark matter halo fields is primarily caused by shot noise, which dominates on smaller scales [46]. Although shot noise is the primary factor for parameter saturation, decoupling modes and incorporating more information from higher N-point statistics into the power spectrum and bispectrum can still be beneficial. This is supported by [30], who found a strong correlation among squeezed halo bispectrum configurations and between these configurations and the long-wavelength halo power spectrum. To approximate real observations, dark matter halo fields are translated from real space into redshift space. This introduces an additional source of non-linearity and, consequently, non-Gaussian covariance.

By applying neural network-based reconstruction to the redshifts-space  $z = 0$  dark matter halo fields, we will show that one can indeed reconstruct a significant part of the real-space initial conditions at the field level. Additionally, by applying our reconstruction method to the data of the Quijote simulation suite [46, 47] and measuring the power spectrum monopole, quadrupole, and bispectrum, we forecast improved constraints on pnG by factors of 1.33, 1.88, and 1.57 for local, equilateral, and orthogonal pnG, respectively, as well as improvements on other cosmological parameters. This improvement results from the mitigation of redshift-space distortions, the reduction of mode coupling and the decrease of parameter degeneracy. Such improvements are hard to achieve by including more non-linear modes due to parameter saturation caused by shot noise and non-Gaussian covariance. We argue that our reconstruction method could complement observational survey data analysis, squeezing out as much information as possible from the data, by improving constraining power without necessitating costly and observationally challenging approaches such as increasing number density, expanding the observed volume to add more modes, or increasing the maximum survey redshift.

This paper is organized as follows. In section 2, we discuss the effect of non-Gaussianity on the dark matter halo density field. We discuss the training and testing of our machine learning model and the simulations we use for training in section 3. Section 4 is dedicated to the statistical analysis, highlighting our main results and their implications for parameter constraints. We conclude in section 5.

## 2 Primordial non-Gaussianity in the dark matter halo density field

### 2.1 Primordial bispectra

For Gaussian fields, the two-point correlation function or power spectrum fully describes the field. Describing non-Gaussian fields instead requires higher-order N-point correlation functions. The lowest-order statistic revealing non-Gaussian characteristics is the three-point-correlation function or bispectrum [48]. PnG is captured by the bispectrum of the primordial potential, in Fourier space:

$$\langle \Phi(\mathbf{k}_1)\Phi(\mathbf{k}_2)\Phi(\mathbf{k}_3) \rangle = (2\pi)^3 \delta^{(3)}(\mathbf{k}_1 + \mathbf{k}_2 + \mathbf{k}_3) B_\Phi(k_1, k_2, k_3). \quad (2.1)$$

Here  $\mathbf{k}_i$  represents a 3D vector in Fourier Space,  $k_i$  is the magnitude of the vector,  $k_i = |\mathbf{k}_i|$ .  $B_\Phi(k_1, k_2, k_3)$  represents the bispectrum of the primordial potential  $\Phi(\mathbf{k})$ . Spatial homogeneity implies that the properties of space do not vary with location, which results in momentum being conserved in a system. Consequently, when considering momentum space triangles,  $\mathbf{k}_1, \mathbf{k}_2, \mathbf{k}_3$ , they must form a closed shape in Fourier space, which is enforced by the delta function in equation (2.1). The bispectrum can peak in various momentum configurations, and we refer to these unique configurations as the “shape” of the bispectrum [49]. In this paper, we consider three well-known pnG templates; the local, the equilateral, and the orthogonal shape.

The so-called local shape of primordial non-Gaussianity occurs when the primordial potential takes a form:

$$\Phi(\mathbf{x}) = \Phi_G(\mathbf{x}) + f_{\text{NL}} \left( \Phi_G(\mathbf{x})^2 - \langle \Phi_G(\mathbf{x})^2 \rangle \right) \quad (2.2)$$

where  $\Phi_G$  is purely Gaussian, resulting in a bispectrum that peaks in the “squeezed limit”, when  $k_1 \ll k_2 \approx k_3$  and is given by:

$$B_\Phi^{\text{local}}(k_1, k_2, k_3) = 2f_{\text{NL}}^{\text{local}} (P_\Phi(k_1)P_\Phi(k_2) + 2 \text{ perms.}). \quad (2.3)$$

In single-field inflationary models, such as slow-roll inflation, local pnG is necessarily suppressed [50, 51], whereas large local non-Gaussianity is expected if multiple fields play a role during inflation.

The equilateral bispectrum peaks in the equilateral triangle configuration  $k_1 \sim k_2 \sim k_3$ . Equilateral non-Gaussianities are naturally produced in the presence of non-gravitational self-interactions of the field driving inflation [52–61]. Due to the suppression of derivatives on superhorizon scales ( $k \ll aH$ ), the resulting bispectrum peaks when all modes have comparable wavenumbers at the time of their horizon exit. The equilateral bispectrum is given by the template:

$$B_\Phi^{\text{equil}}(k_1, k_2, k_3) = 6f_{\text{NL}}^{\text{equil}} (-P_\Phi(k_1)P_\Phi(k_2) + 2 \text{ perms.} - 2(P_\Phi(k_1)P_\Phi(k_2)P_\Phi(k_3))^{\frac{2}{3}} + P_\Phi(k_1)^{\frac{1}{3}}P_\Phi(k_2)^{\frac{2}{3}}P_\Phi(k_3) + 5 \text{ perms.}). \quad (2.4)$$

The self-interactions that create the equilateral bispectrum can be captured in the Effective Field Theory (EFT) [56].

Finally, EFT also predicts a third shape, which is orthogonal to the equilateral bispectrum in momentum space. A template of this type is provided in ref. [59]:

$$\begin{aligned}
 B_{\Phi}^{\text{orth}}(k_1, k_2, k_3) &= 6f_{\text{NL}}^{\text{orth}} \left( (1+p) \frac{\Delta_{123}}{k_1^3 k_2^3 k_3^3} - p \frac{\Gamma_{123}}{k_1^4 k_2^4 k_3^4} \right), \\
 \Delta_{123} &= (k_T - 2k_1)(k_T - 2k_2)(k_T - 2k_3), \\
 \Gamma_{123} &= \frac{2}{3}(k_1 k_2 + k_2 k_3 + k_3 k_1) - \frac{1}{3}(k_1^2 + k_2^2 + k_3^2), \\
 p &= \frac{27}{-21 + \frac{743}{7(20\pi^2 - 193)}}.
 \end{aligned} \tag{2.5}$$

## 2.2 Scale-dependent bias

Contrary to the dark matter field, where non-linear growth only leaves a small imprint of pnG in the power spectrum [31], for biased tracers such as halos, this imprint can be much more severe. The effect of pnG on the halo power spectrum was first predicted and validated using N-body simulations with local type pnG initial conditions in ref. [62] which showed a strong scaling of  $1/k^2$  on large scales in the power spectrum, for local-type non-Gaussian initial conditions. A clean physical interpretation and a model-independent prediction of this result can be obtained by applying the peak-background split [63–67]. The halo number density contrast,  $\delta_h(\mathbf{k}|M, z)$ , is defined as:

$$\delta_h(\mathbf{k}|M, z) = \frac{n_h(\mathbf{x}|M, z)}{\bar{n}_h(M, z)} - 1, \tag{2.6}$$

where  $n_h(x|M, z)$  is the number of halos with mass  $M$  at redshift  $z$  and  $\bar{n}_h(M, z)$  is its mean.  $\delta_h(\mathbf{k}|M, z)$  represents the halo overdensity at wavenumber  $\mathbf{k}$  with mass  $M$  at redshift  $z$ . In the absence of pnG, on the largest scales, the halo overdensity is related to the matter overdensity,  $\delta_m(\mathbf{k})$ :

$$\delta_h(\mathbf{k}|M, z) = b_1(M, z)\delta_m(\mathbf{k}, z) + \epsilon(z). \tag{2.7}$$

Here  $b_1 = \frac{1}{\bar{n}_h} \frac{\partial n_h}{\partial \delta_m}$  is the usual Gaussian Lagrangian bias that quantifies the response of the halo number count to large-scale matter fluctuations and  $\epsilon(z)$  is shot noise due to the discrete nature of the halos. In the peak-background split approach, we separate the long- and short-wavelength modes of the density field. In this context, we examine the most well-known **local** pnG example. The separation in the Gaussian primordial fluctuations reads:

$$\Phi_G(\mathbf{k}) = \Phi(\mathbf{k}_\ell) + \Phi(\mathbf{k}_s). \tag{2.8}$$

Here  $\ell$  presents long modes and  $s$  the short modes. Following equation (2.2), local pnG induces a coupling between long-wavelength modes  $k_\ell$  and short wavelength modes  $k_s$  in the primordial potential [4]:

$$\Phi(\mathbf{k}_s) \approx \Phi_G(\mathbf{k}_s) \left( 1 + 2f_{\text{NL}}^{\text{local}} \Phi_G(\mathbf{k}_\ell) \right). \tag{2.9}$$

Hence in the presence of local pnGs, the dependence of the halo overdensity in equation (2.7) does not only rely on the large-scale matter overdensity but also the primordial potential [46]:

$$\delta_h(\mathbf{k}|M, z) = b_1(M, z)\delta_m(\mathbf{k}, z) + b_{\Phi}(M, k, z)\mathcal{M}(k, z)\Phi(\mathbf{k}) + \epsilon(z). \tag{2.10}$$

$b_\Phi(M, k, z)$  represents the bias associated with the primordial potential  $\Phi(\mathbf{k})$ . Unlike  $b_1(M, z)$ ,  $b_\Phi(M, k, z)$  is a function of the wavenumber  $k$ , the source of the scale-dependent bias.  $\mathcal{M}(k, z)$  is the matter transfer function which transfers the primordial potential to the matter overdensity. We can rewrite equation (2.10):

$$\delta_h(\mathbf{k}|M, z) = [b_1(M, z) + b_\Phi(M, k, z)] \delta_m(\mathbf{k}, z) + \epsilon(z). \quad (2.11)$$

Following [64], we then obtain a relation between  $b_\Phi$  and  $b_1$  and the primordial bispectrum  $B_\Phi(k_1, k_2, k_3)$ :

$$b_\Phi(M, z) = \frac{2f_{\text{NL}}\delta_c}{\mathcal{M}(k, z)} (b_1 - 1) \frac{\sigma_W^2(k)}{\sigma_M^2}. \quad (2.12)$$

$\delta_c \approx 1.686$  is the spherical collapse threshold,  $\sigma_M^2$  represents the variance of the top-hat filtered matter field, and  $\sigma_W^2$  denotes the variance of the smoothed field in the presence of pnG

$$\sigma_M^2 = \int \frac{d^3k'}{(2\pi)^3} W^2(k') P_\delta(k'), \quad \sigma_W^2 = \int \frac{d^3k'}{(2\pi)^3} \frac{W^2(k') P_{\text{lin}}(k') B_\Phi(\mathbf{k}, \mathbf{k}', \mathbf{k} - \mathbf{k}')}{2f_{\text{NL}} (P_\Phi(k) P_\Phi(k') + 2 \text{ perms.})}. \quad (2.13)$$

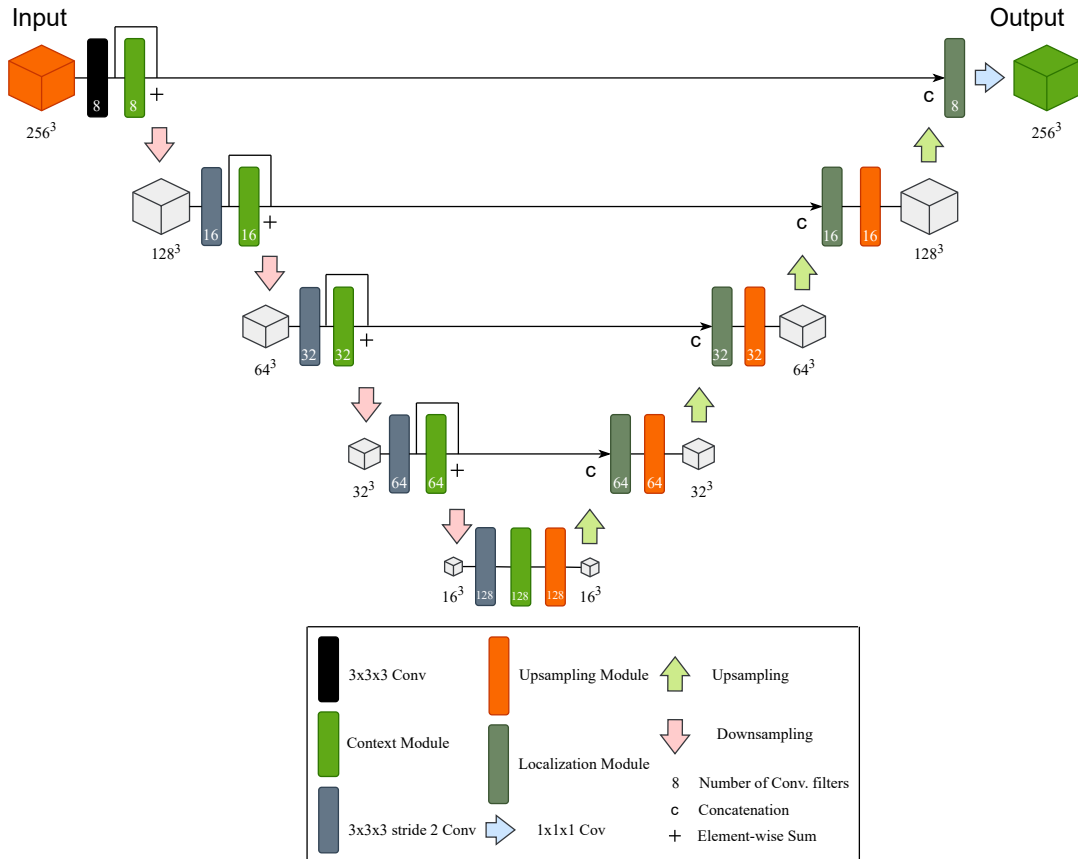
Here,  $W_M(k)$  is the Fourier transformed top-hat filter, with radius  $R^3 = \frac{3\bar{\rho}}{4\pi} M$  set by the halo mass, and  $P_{\text{lin}}(k)$  is the linear matter power spectrum. For local pnG,  $\sigma_W^2(k) \sim \text{constant}$ , and on large scales where  $\mathcal{M}(k, z) \sim k^2$ , this implies  $b_\Phi \propto 1/k^2$ . For equilateral and orthogonal pnGs,  $\sigma_W^2(k) \sim k^2$ , resulting in  $b_\Phi \propto 1$ . The correction  $f_{\text{NL}}^{\text{local}}$  of the power spectrum on observable scales primarily affects  $b_\Phi$  ( $\propto 1/k^2$ ) rather than higher order  $b_\Phi^2$  ( $\propto 1/k^4$ ).

### 3 Reconstruction with neural networks

While LSS has the potential to become the leading cosmological observable, it is inherently non-linear due to both gravitational interactions and astrophysical processes, challenging cosmological information extraction. This is especially true for pnGs, which are smaller than late-time induced non-Gaussianities. Mode coupling in non-linear gravitational evolution reduces unique information per mode [29], further complicating the detection of pnGs [30, 32, 46]. Non-linear evolution disperses the information in the initial condition to higher-order N-point statistics and smaller scales. In ref. [44] we demonstrated that neural network-based reconstruction has the potential to recover lost information, by linearizing the field, which transfers information from the non-linearly evolved modes into the recovery of the primordial modes. In this paper, we extend this method to the dark matter halo field as a realistic proxy for galaxy behaviour in LSS. We will also include redshift-space distortion effects, that are typically present in realistic tracer fields.

#### 3.1 U-Net architecture

The U-Net is a convolutional neural network originally developed for biomedical image segmentation [68]. U-Nets have recently achieved success in cosmology and astrophysics, especially in problems involving coupling of different length scales [44, 69, 70]. Unlike ordinary convolutional neural networks (CNNs), which have a smaller receptive field, U-Nets include an encoder-decoder mechanism. This feature effectively enlarges their receptive field, enabling them to efficiently process both small-scale and large-scale information. The architecture used



**Figure 1.** Diagrammatic representation of the U-Net architecture. The size of the boxes illustrates the downsampling of the density to smaller dimensions and the subsequent upsampling back to the original input dimensions.

in this work is similar to the U-Net in previous work [44], resembling the architecture originally developed in [70]. The network processes an input density field on a  $256^3$  grid, computing additional gradient fields on-the-fly according to grid-based perturbation theory [71]:

$$\mathbf{u}(\mathbf{x}) = \int d^3\mathbf{k} \frac{-i\mathbf{k}}{k^2} \delta_m(\mathbf{k}) e^{i\mathbf{k}\cdot\mathbf{x}}, \quad (3.1)$$

$$\partial_i(\mathbf{u})_j(\mathbf{x}) = \int d^3\mathbf{k} \frac{k_i k_j}{k^2} \delta_m(\mathbf{k}) e^{i\mathbf{k}\cdot\mathbf{x}}, \quad (3.2)$$

resulting in six additional input fields alongside the input density field. The U-Net architecture consists of an encoder-decoder network with skip connections. The encoder component of U-Net (left half of figure 1) starts with a convolutional layer with a  $3^3$  kernel, followed by a context block. This block includes two more convolutional layers, each of size  $3^3$ . We then combine the output of the first convolution with that of the context block, forming what is called a residual connection. Next, we introduce a stride of two in the initial convolutional layer, which downsamples the feature maps by halving their size in each spatial dimension. These downsampling operations are repeated, reducing the feature maps to a dimension of  $16^3$  cells, capturing both local and global features at different scales. The decoder component of the U-Net (right half of figure 1) reconstructs high-resolution feature maps from the

	$f_{\text{NL}}^{\text{local}}$	$f_{\text{NL}}^{\text{equil}}$	$f_{\text{NL}}^{\text{orth}}$	$h$	$n_s$	$\Omega_m$	$\Omega_b$	$\sigma_8$
$\theta_{\text{fid}}$	0	0	0	0.6711	0.9624	0.3175	0.049	0.834
$\delta\theta$	$\pm 100$	$\pm 100$	$\pm 100$	$\pm 0.02$	$\pm 0.02$	$\pm 0.01$	$\pm 0.002$	$\pm 0.015$

**Table 1.** Quijote simulations’ cosmological parameters. The first row presents the default cosmology applied for training the neural network and calculating the covariance matrix. The second row details the parameter variations in simulations employed for determining the derivatives.

encoded representations produced by the encoder. It utilizes upsampling operations to increase the dimensions of the feature maps by a factor of two and concatenate the output of the residual connections in the down-sampling part of the network, forming what are known as skip connections. Skip connections between encoder and decoder layers preserve spatial information and facilitate precise localization during reconstruction. At the end of the model, the U-Net has reconstructed a density field with the same dimensions as the input field  $256^3$ . The network is implemented using JAX and is publicly available, along with an example notebook.<sup>1</sup>

## 3.2 Training data

### 3.2.1 Simulations

For the input fields, we use the halo catalogue from the Quijote simulation suite. The initial conditions of the simulations are generated at  $z = 127$  using 2LPT. The N-body dark matter simulations are run with GADGET-III [72] as presented in [31]. The halo catalogue used in this work is constructed using the Friends of Friends (FoF) algorithm [73] on the N-body simulations at redshift  $z = 0$  [46]. The main analysis of this work focuses on the complete set of halos available in the Quijote simulation suite, characterized by a mass greater than  $M_{\text{min}} \geq 1.32 \cdot 10^{13} M_{\odot}/h$ , resulting in an average number density of  $\bar{n} = 4 \times 10^{-4} \text{ (h/Mpc)}^3$ . Additionally, the sensitivity of our reconstruction method to different number densities will be investigated. To assess the gain of doing reconstruction for cosmological parameter analysis, we will perform a Fisher forecast, using 15,000 simulations run with the fiducial cosmological model parameters, detailed in the first row of table 1. To estimate derivatives for the Fisher forecast, we employ 500 simulations for each positive and negative perturbation of every cosmological parameter, as shown in the second row of table 1, using finite differencing.

### 3.2.2 Redshift-space distortions

In galaxy dynamics, peculiar velocities are additional galaxy motions caused by gravitational influence from nearby bodies, distinct from the Hubble expansion, leading to redshift-space distortions that complicate galaxy clustering analysis. The mapping from real-space to radial position for flat models is given by [74]:

$$\mathbf{s} = \mathbf{x} + \frac{(1+z)}{H(z)} \mathbf{v}(\mathbf{x}) \hat{z}, \quad (3.3)$$

<sup>1</sup>[https://github.com/jeltebottema/dark\\_matter\\_halo\\_reconstructions](https://github.com/jeltebottema/dark_matter_halo_reconstructions).

where  $\mathbf{v}(\mathbf{x})$  is the peculiar velocity,  $H$  the Hubble parameter, and  $\hat{z}$  is the line of sight direction. The halo catalogue provides the peculiar velocities of each halo in the  $x, y, z$  axis, which is computed as a mass weighting of the halo particles [46]. For the initial training process, we distort the halo positions of the fiducial simulations along the  $z$  axis using the peculiar velocity. For the halos in the derivative simulations, we also create the catalogues with distortions applied along either the  $x$  or  $y$  direction, resulting in three catalogues per realization. Although these three catalogues obviously share the same real-space configuration, they are not perfectly correlated in redshift space, particularly on small scales, providing additional realizations for estimating derivatives [46]. This procedure is applied only to the derivative simulations, as the 15,000 simulations used for the covariance are sufficient for our analysis.

### 3.2.3 Dark matter halo density fields

We interpolate the halos from the halo catalogue, which covers a volume of  $V = 1 \text{ (Gpc/h)}^3$ , into a grid with  $256^3$  voxels using the Piecewise Cubic Spline (PCS) mass-assignment scheme [75, 76]. The overdensity of the halos is calculated by first determining the local density of halos,  $\rho_h(\mathbf{x})$ , in each voxel. This local density is first subtracted and then normalized by the mean density  $\bar{\rho}_h$ , calculated separately for each realization of the simulations, as the mean differs throughout the different catalogues:

$$\delta_h(\mathbf{x}) = \frac{\rho_h(\mathbf{x}) - \bar{\rho}_h}{\bar{\rho}_h} = \frac{\rho_h(\mathbf{x})}{\bar{\rho}_h} - 1. \quad (3.4)$$


$\delta_h$  measures the deviation of the local number density from the cosmic mean number density. In the main setup of the analysis, we omit halo mass, making  $\rho_h$  a number density,  $n_h$ .

### 3.3 Training the model

The U-Net is trained to reconstruct the initial dark matter density fields ( $z = 127$ ), used as initial conditions for the simulations, from the final dark matter halo density fields ( $z = 0$ ) by minimizing the loss function:

$$L = \frac{1}{N_{\text{sims}}} \sum_i^{N_{\text{sims}}} \left( \frac{1}{N_{\text{Cells}}} \sum_x (\delta_{\text{recon}}^{(i)}(\mathbf{x}) - \delta_{\text{init}}^{(i)}(\mathbf{x}))^2 \right). \quad (3.5)$$

The loss function is the mean-squared error of the reconstructed  $\delta_{\text{recon}}^{(i)}$  and initial  $\delta_{\text{init}}^{(i)}$  density fields of the same snapshot of simulation  $i$ , calculating the difference between the density fields pixel by pixel. The dataset consists of 380 pairs of initial dark matter ( $z = 127$ ) and final ( $z = 0$ ) dark matter halo density fields, of which we used 250 for training and 130 for validation. Our primary goal is to determine the information content of cosmological parameters from the power spectrum and bispectrum of the reconstructed density field. The FFT-based estimator used to compute the bispectrum is limited to a scale of  $k_{\text{max}} = \frac{2}{3} k_{\text{Nyq}} \approx 0.53, \text{ h/Mpc}$  [76], where  $k_{\text{Nyq}} \approx 0.82, \text{ h/Mpc}$  is the Nyquist frequency. The Nyquist frequency is given by  $k_{\text{Nyq}} = \pi \frac{N_G}{L}$ , which is the highest frequency that can be accurately resolved, determined by the number of grid points  $N_G$  and box size  $L$ . Since our focus will be on a specific subset of modes in the density fields, we only consider initial dark matter density fields up to this maximum wave number  $k_{\text{max}}$ , applying



**Figure 2.** Left: The dark matter halo field ( $z = 0$ ). Middle: the initial dark matter field ( $z = 127$ ). Right: the reconstructed density field. The fields are three-dimensional but are averaged over along a segment of the same spatial axis for visualization purposes.


an isotropic sharp cut-off in Fourier space [44]. This method ensures that the network avoids reconstructing modes that we will not use in our final analysis, thereby accelerating the training process. Note that the halo fields are not filtered, and thus contain information all the way up to  $\sqrt{3}k_{\text{Nyq}} = 1.4 \text{ h/Mpc}$ . To ensure model stability, we employ exponential moving average (EMA) of weights [77]. We train the network using an Nvidia A100 40GB GPU until the validation loss stops improving significantly, which takes approximately four hours. The training is halted when overfitting starts to occur, ensuring that the model generalizes well to unseen data.

### 3.4 Reconstruction performance

The output of the model is a reconstructed density field derived from the final dark matter halo field at  $z = 0$ . This field can be compared with the target initial dark matter field at  $z = 127$ , an example of which is provided in figure 2. As expected, the reconstructed linear matter density field is visibly more linear and less non-Gaussian, and closely resembles the true initial conditions on large scales. This figure also illustrates that the reconstruction is less accurate on smaller scales. To make these observations more quantitative, we calculate the cross-correlation between two fields  $X$  and  $Y$  as:

$$C_{X,Y}(\mathbf{k}) = \frac{P_{X,Y}(\mathbf{k})}{\sqrt{P_X(\mathbf{k})P_Y(\mathbf{k})}}. \quad (3.6)$$

Here  $P_{X,Y}(\mathbf{k})$  is the cross-power spectrum of fields  $X$  and  $Y$ , and  $P_X(\mathbf{k})$  the auto-power spectrum of field  $X$ . Here,  $X, Y$  can be the reconstructed, initial ( $z = 127$ ) matter density fields or the final ( $z = 0$ ) halo density field. A perfect reconstruction thus implies  $C_{\text{recon,init}}(\mathbf{k}) = 1$ . Figure 3 shows the cross-correlation of the final and reconstructed fields with the initial density field for a single realization, demonstrating a clear improvement in the cross-correlation after reconstruction. We find an accuracy of 44% and upward on scales below  $k = 0.4 \text{ h/Mpc}$ , a significant improvement over the 3% correlation of the final halo field with the initial dark matter field and on this scale, and the 23% correlation achieved using the standard

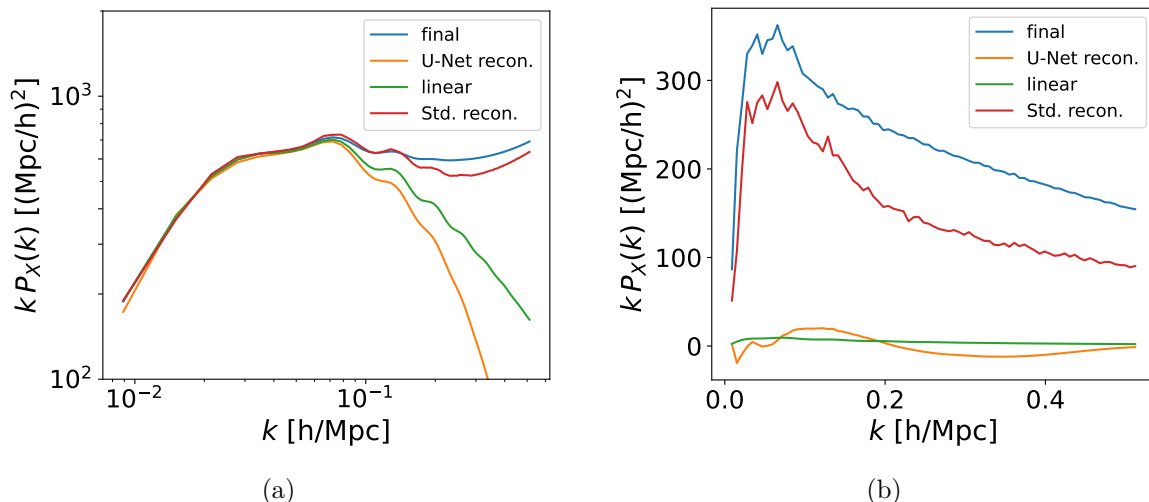


**Figure 3.** The cross-correlation of density fields at a single realization. The blue line represents the cross-correlation between the final dark matter halo fields ( $z = 0$ ) and initial dark matter fields ( $z = 127$ ), while the orange line represents the correlation between the reconstructed and initial fields. The green line represents the cross-correlation between a standard reconstructed field and the initial field.

reconstruction method [33].<sup>2</sup> We have also trained the model without the gradient velocity fields in equations (3.1) and (3.2), and is shown in the appendix figure 12. Interestingly, the model shows similar performance when provided a halo field with an isotropic sharp cut-off at the Nyquist frequency (instead of the natural cut-off of the box  $\sqrt{3}k_{\text{Nyq}}$ ), i.e. disregarding the modes along the spatial diagonal that arise due to the three-dimensional box structure of the density fields. In other words, the reconstruction up to  $k_{\text{max}} = 0.52$  Mpc/h is already saturated completely by the halo-field information up to the Nyquist frequency. This strongly suggests that our model has performed the best possible reconstruction for the given setup. At higher number densities (lower shot noise), these smaller scales could potentially provide additional information for reconstruction.

Redshift-space distortions introduce additional non-linearity into the halo fields. We can investigate the ability of our model to mitigate these distortions by looking at the power spectrum after reconstruction. The power spectra over many realizations are given in figure 4(a) and figure 4(b), for the monopole and quadrupole, respectively. In these figures, the spectra for the reconstructed fields and the linear fields at  $z = 127$  are rescaled to  $z = 0$ , and the final halo field is corrected for the linear bias term,  $b_1$ . In this way, the improvement in the spectra shown in these plots can be entirely attributed to our reconstruction method. We note that the quadrupole is largely removed by the reconstruction, demonstrating a strong reduction of redshift-space distortions. This level of reduction is not achieved with standard reconstruction, showcasing the advantage of our method in eliminating redshift-space distortions.

<sup>2</sup>We perform a single iteration of reconstruction using the ReIso scheme for removing redshift-space distortions as described e.g. in [78].



**Figure 4.** The mean power spectrum taken from 15000 fiducial cosmology simulations for the pre- and post-reconstructed density fields, compared with the initial linear power spectrum: (a) Power spectrum monopole (b) Power spectrum quadrupole. In the plot, ‘final’ denotes the pre-reconstructed halo number density field, while ‘U-Net recon’ denotes the reconstructed linear matter density fields. In these plots, the reconstructed and linear fields at  $z = 127$  are rescaled to  $z = 0$ , also the final and standard reconstruction spectra are corrected for the linear bias,  $b_1$  (for standard reconstruction the spectra are only averaged over 50 realizations).

Since the amount of information in the dark matter halo field is dependent on the shot noise, which dominates on small scales [46], it is worth exploring how varying levels of shot noise in the input fields affect the accuracy of reconstruction. To address the impact of shot noise on the reconstruction quality, we explored two modifications to the input halo density fields in appendix A. First, we test the effects of increasing the minimum halo mass threshold (decreasing the number density), which inherently raises shot noise. On large scales up to  $k = 0.1 h/\text{Mpc}$ , the reconstruction quality is approximately the same for different number densities. However, for lower number densities, the reconstruction clearly deteriorates more rapidly at smaller scales. Second, we apply binned halo mass weighting to effectively reduce shot noise. We found that incorporating additional mass information into the input density fields improves the quality of reconstruction.

#### 4 Cosmological information content of the reconstructed field

Using the reconstruction method presented in the previous section, we now turn to analysing the information content on cosmological parameters, in order to assess the effectiveness of our reconstruction method for cosmological parameter inference. First, we measure the power spectrum monopole and quadrupole, as well as the bispectrum monopole from the reconstructed, initial, and final density fields. Then we perform a Fisher analysis on the estimated spectra, utilizing a Fisher estimator first proposed in [79], which combines the standard Fisher estimator with a compressed Fisher estimator, to ensure that we obtain an unbiased forecast for the analysis.

#### 4.1 Power spectrum and bispectrum estimator

The redshift-space halo power spectrum  $P_{hh}(k, \mu)$  depends both on the magnitude of the wave vector and the angle to the line of sight,  $\mu = \mathbf{k} \cdot \mathbf{n}$ , and can be expanded into multipoles:

$$P_{hh}^{(\ell)}(k) = \frac{2\ell + 1}{2} \int d\mu \mathcal{L}_\ell(\mu) P_{hh}(k, \mu), \quad (4.1)$$

where  $\mathcal{L}_\ell(\mu)$  are the Legendre polynomials, such that the quadrupole  $P_{hh}^{(2)}(k)$  depends on the line-of-sight angle  $\mu$ , unlike the monopole  $P_{hh}^{(0)}(k)$ . We define the estimators for the power spectrum monopole and quadrupole of a density field  $\delta$  as:

$$\hat{P}^{(\ell)}(k_i) = \frac{\ell + 1}{2N_i} \sum_{\mathbf{q} \in k_i} \mathcal{L}_\ell(\mu) \delta(\mathbf{q}) \delta^*(\mathbf{q}), \quad (4.2)$$

where we sum over all momenta  $\mathbf{q}$  that are within a shell  $[k_i - \Delta k/2 + \Delta k/2]$ .  $N_i$  is a normalization factor that counts the number of modes that fall in the bin. For the power spectrum calculations, the largest scale bin starts at  $k_f$ , where  $k_f$  is the fundamental frequency given by  $k_f = \frac{2\pi}{L}$ , where  $L$  is the physical size of the simulation box (in our case 1000 Mpc). We use bins with a width of  $\Delta k = k_f$ , up to the smallest scale  $k_{\max} = 82.5k_f \approx 0.52, \text{h/Mpc}$ .

Turning to the bispectrum, similar to the halo power spectrum, redshift-space distortions cause the halo bispectrum to depend on the orientation of the modes relative to the line of sight. In this work, however, we focus exclusively on the bispectrum monopole, which is defined by averaging over all possible angles and is estimated as:

$$\hat{B}(k_1, k_2, k_3) = \frac{1}{N_{123}} \sum_{\mathbf{q}_1 \in k_1} \sum_{\mathbf{q}_2 \in k_2} \sum_{\mathbf{q}_3 \in k_3} (2\pi)^3 \delta^{(3)}(\mathbf{q}_1 + \mathbf{q}_2 + \mathbf{q}_3) \delta(\mathbf{q}_1) \delta(\mathbf{q}_2) \delta(\mathbf{q}_3). \quad (4.3)$$

Here  $N_{123}$  counts the number of triangle configurations within the bin. The bispectrum is binned with a width of  $\Delta k = 3k_f$ , where the first bin starts at  $k = 1.5k_f$ , where again the smallest scale is  $k_{\max} = 82.5k_f$ , resulting in 2276 triangle bins.

To calculate the power spectrum and bispectrum for each density field, we use the open-source code BFast,<sup>3</sup> which is GPU accelerated through JAX. Note that the code applies the usual compensation for the mass assignment PCS window function in Fourier space [80], which was used to construct the density fields.

Having measured the monopole and quadrupole power spectrum and monopole bispectrum for both pre- and post-reconstructed density fields, which includes 15,000 fiducial simulations and 1500 (including the extra 1000 rotated fields) simulations with varying cosmological parameters as listed in table 1, we have obtained our data products:

$$\mathbf{D} = \left\{ P_{\text{pre}}^{(0)}, P_{\text{pre}}^{(2)}, B_{\text{pre}}, P_{\text{post}}^{(0)}, P_{\text{post}}^{(2)}, B_{\text{post}} \right\}. \quad (4.4)$$

As the halo field is a stochastic tracer of the underlying density field, the reconstructed density field will contain a stochastic component, due to the slightly varying reconstruction quality per halo field. This stochastic component leaves an imprint in the correlation functions, similar to the shot noise contribution to the halo power spectrum, but now scale dependent.

<sup>3</sup><https://github.com/tsfloss/BFast>.

We choose not to attempt to model this stochastic contribution and thus do not perform any subtraction of the stochastic component. However, in the following Fisher forecast, the simulation-based covariance matrix does capture the additional variance coming from this implicit stochastic contribution.

## 4.2 Covariance

In order to assess the amount of cosmological information that can be extracted from a set of data, we build the covariance matrix  $C$  of the data vector  $\mathbf{D}$ :

$$C_{ij} = \frac{1}{N-1} \sum_n^N (\mathbf{D}_n - \bar{\mathbf{D}})_i (\mathbf{D}_n - \bar{\mathbf{D}})_j. \quad (4.5)$$

Here  $N$  is the number of simulations and  $\bar{\mathbf{D}}$  is the mean of the data over all  $N$  simulations. Subscripts  $i, j$  are for indexing the bins in the data vector  $\mathbf{D}$ . To illustrate the correlation between different values of the power spectrum and bispectrum, we use a correlation matrix:

$$r_{ij} = \frac{C_{ij}}{\sqrt{C_{ii}C_{jj}}}. \quad (4.6)$$

Before reconstruction, modes are coupled, leading to significant non-Gaussian covariance in the measurements. Reconstruction reduces this correlation, as evidenced by the decreased correlation among modes post-reconstruction, demonstrated by the correlation matrix in figure 5. This reduction in correlation implies the presence of more unique information in the power spectrum and bispectrum of the reconstructed density field.

## 4.3 Fisher analysis

### 4.3.1 The standard fisher estimator

The Fisher matrix formalism [81] provides a framework for estimating parameter uncertainties. Assuming Gaussian likelihood, which is accurate enough for our purposes [82, 83], the Fisher matrix for cosmological parameters  $\theta_a$  and data vector  $\bar{\mathbf{D}}$  can be expressed as:

$$F_{ab}^\theta = \frac{\partial \bar{\mathbf{D}}_i}{\partial \theta_a} (C^{-1})_{ij} \frac{\partial \bar{\mathbf{D}}_j}{\partial \theta_b}. \quad (4.7)$$

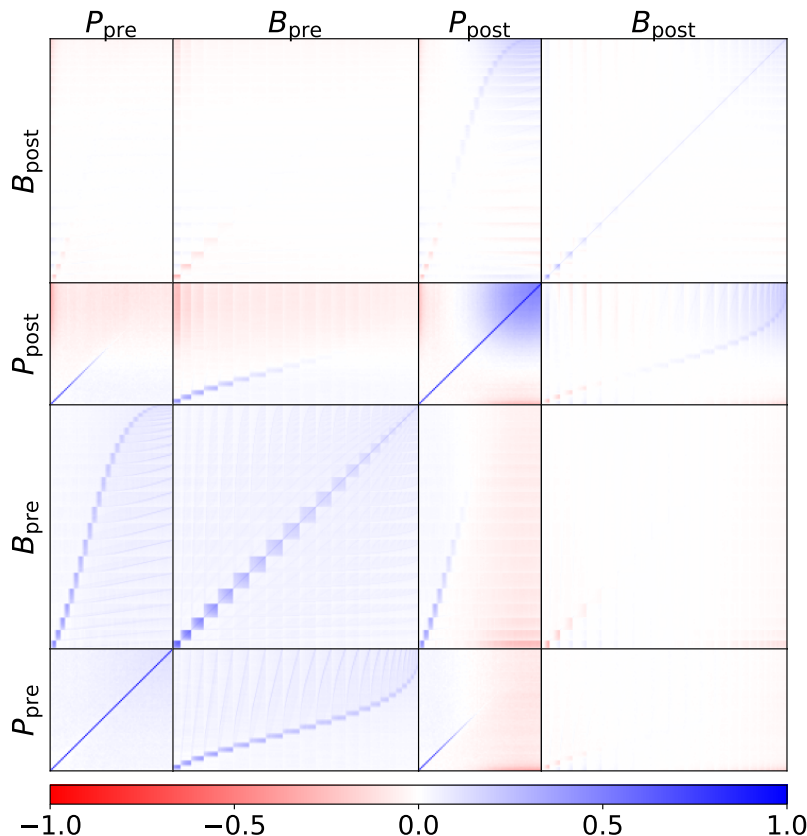
$\bar{\mathbf{D}}$  is the simulation-averaged data vector, and  $C^{-1}$  is the precision matrix derived from the inverse of the covariance matrix from equation (4.5), adjusted using the Hartlap factor [84, 85]:

$$C^{-1} \equiv \frac{N_{\text{sims}} - N_{\text{bins}} - 2}{N_{\text{sims}} - 1} C^{-1} \quad \text{for } N_{\text{bins}} < N_{\text{sims}} - 2. \quad (4.8)$$

$N_{\text{sims}}$  and  $N_{\text{bins}}$  denote the number of simulations and data points (power spectrum and bispectrum bins), respectively. We compute parameter derivatives using finite central differencing:

$$\frac{\partial \bar{\mathbf{D}}}{\partial \theta_a} = \frac{\bar{\mathbf{D}}_{\theta_a^{\text{fid}} + \delta\theta_a} - \bar{\mathbf{D}}_{\theta_a^{\text{fid}} - \delta\theta_a}}{2\delta\theta_a}, \quad (4.9)$$

where  $\bar{\mathbf{D}}_{\theta_a^{\text{fid}} \pm \delta\theta_a}$  represents the mean measured data from simulations with perturbed parameters.



**Figure 5.** Correlation matrix  $r_{ij}$  for the initial and reconstructed power spectrum and bispectrum. Bispectrum triangle configurations are sorted by increasing smallest momentum, creating visual blocks with the same shortest side in  $k$ -space. Note that we have only plotted the monopole power spectrum for visual clarity, the quadrupole power spectrum results in approximately the same correlation matrix as the monopole power spectrum.

### 4.3.2 Compressed Fisher estimator

The standard Fisher estimator can lead to overestimation due to Monte Carlo noise, particularly in derivatives. In ref. [79], Coulton & Wandelt construct an alternative Fisher estimator, based on a compressed statistic, that is oppositely biased to the standard Fisher estimator and can be combined to reduce the noise and accelerate convergence with the number of simulations. We follow the work of Coulton & Wandelt and refer the reader to that work for more details.<sup>4</sup> An optimal compression, for the parameters from data vector  $\mathbf{D}$ , would require knowledge of precisely the ingredients of the standard Fisher estimator. Instead, we use a subset  $\alpha$  number of simulations to construct the (suboptimally) compressed statistic. For a Gaussian likelihood with parameter independent covariance this is given by [79, 86, 87]:

$$t_a = \frac{\partial \bar{\mathbf{D}}_i}{\partial \theta_a} (C^\alpha)_{ij}^{-1} (\mathbf{D}_j - \bar{\mathbf{D}}_j^\alpha), \quad (4.10)$$

<sup>4</sup>We employ a slightly different notation from that of ref. [79], where we strictly reserve subscripts  $i, j$  for indexing the bins in the data vector  $\mathbf{D}$ , and subscripts  $a, b, c, d$  for indexing parameters  $\theta$ .

where  $C^\alpha$  is the covariance matrix in equation (4.5) constructed from the subset  $\alpha$  of simulations, and  $\bar{\mathbf{D}}^\alpha$  the mean over the subset  $\alpha$ . In this compressed space, the Fisher information is:

$$F_{ab}^{tt} = \frac{\partial \mu_c^t}{\partial \theta_a} (\Sigma^{tt})_{cd}^{-1} \frac{\partial \mu_d^t}{\partial \theta_b}. \quad (4.11)$$

Here  $\mu_c^t$  is the mean of the compressed statistic,  $\Sigma^{tt}$  is its covariance:

$$\Sigma_{cd}^{tt} = \frac{\partial \bar{\mathbf{D}}_i^\alpha}{\partial \theta_c} \cdot (C^\alpha)_{ij}^{-1} \cdot \frac{\partial \bar{\mathbf{D}}_j^\alpha}{\partial \theta_d}. \quad (4.12)$$

The remaining subset of simulations,  $\beta$ , is used to compute the derivatives to use in equation (4.11):

$$\frac{\partial \mu_c^t}{\partial \theta_a} = \frac{\partial \bar{\mathbf{D}}_i^\alpha}{\partial \theta_c} \cdot (C^\alpha)_{ij}^{-1} \cdot \frac{\partial \bar{\mathbf{D}}_j^\beta}{\partial \theta_a}. \quad (4.13)$$

To avoid bias in the compressed Fisher estimator due to reusing simulations, we split the simulations into two parts: 75% for estimating the compression and 25% for computing the quantities for compressed Fisher information.

### 4.3.3 Combined Fisher estimator

After calculating the standard Fisher estimator, which tends to be biased high, and the compressed Fisher estimate, which tends to be biased low, we can combine these two estimators to compute an unbiased estimator. To create a combined estimator, we use the geometric mean [88]:

$$F_{IJ}^{\text{Combined}} = G(F_{IJ}^{\text{Standard}}, F_{IJ}^{\text{comp}}), \quad (4.14)$$

with  $G(A, B)$  being:

$$G(A, B) = A^{1/2} \left( A^{-1/2} B A^{-1/2} \right)^{1/2} A^{1/2}. \quad (4.15)$$

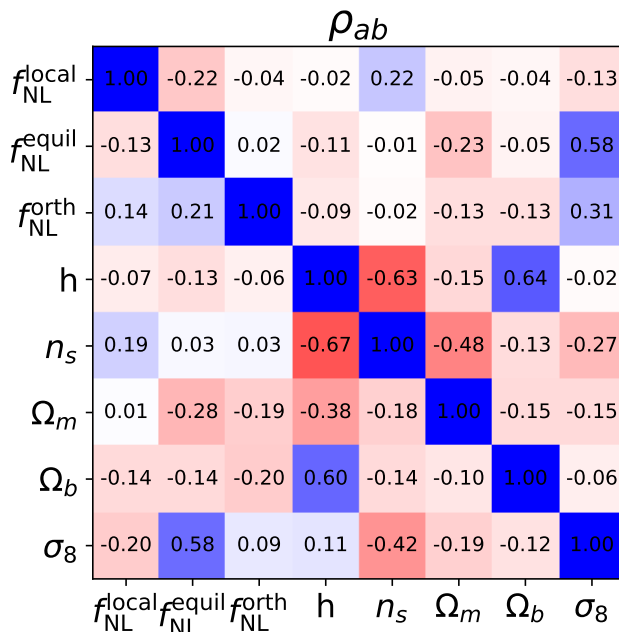
After obtaining one estimate of the combined Fisher information, we can generate a second estimate by reassigning the simulations between the compression and Fisher estimation steps. This process can be repeated multiple times, generating a set of partially correlated estimates that, when averaged, reduces the estimator variance [79]. For our purposes, we shuffled the simulations eight times to ensure that our measurements were reliable. The code for calculating the combined Fisher estimator is publicly available.<sup>5</sup> The combined Fisher estimator ensures that we avoid Monte Carlo uncertainties arising from the finite number of simulations, particularly in the derivatives.

## 4.4 Marginalization

For the Fisher matrices, the marginalized error is given by the Cramer-Rao bound [89, 90], under the assumption of a Gaussian likelihood:

$$\sigma_{\theta_a} = \sqrt{(F^{-1})_{aa}}. \quad (4.16)$$

<sup>5</sup><https://github.com/wcoulton/CompressedFisher>.



**Figure 6.** Correlation coefficients  $\rho_{ij} \in [-1, 1]$  between the different cosmological parameters when using the data products  $(P+B)_{\text{pre}}$  (lower triangular) and  $(P+B)_{\text{pre}} + (P+B)_{\text{post}}$  (upper triangular).

In the Fisher analysis, we can also quantify the degeneracy between parameters by their correlation coefficient:

$$\rho_{ab} = \frac{(F^{-1})_{ab}}{\sqrt{\sigma_a^2 \sigma_b^2}}. \quad (4.17)$$

Figure 6 shows a reduced degeneracy between the pnG parameters and between the pnG parameters and the cosmological parameters. This implies that reconstruction helps in isolating the effects of the pnG parameters.

Furthermore, the Fisher matrix can be used to construct an unbiased, minimal-variance marginalized estimator that estimates the parameter  $\theta^a$  from the observed data  $\mathbf{D}_{\text{obs}}$  given a fiducial model  $\bar{\mathbf{D}}_{\text{fid}}$ :

$$\hat{\theta}^a - \theta_{\text{fid}}^a = \sum_b (F^{-1})_{ab} \frac{\partial \bar{\mathbf{D}}}{\partial \theta^b} \cdot C^{-1} \cdot (\mathbf{D}_{\text{obs}} - \bar{\mathbf{D}}_{\text{fid}}). \quad (4.18)$$

Since the network is trained only on simulations following the fiducial cosmological model, reconstructions derived from simulations with different parameters may have inherent biases. This bias could potentially be propagated through the process when these reconstructions are used to compute derivatives, shown in equation (4.9). Despite this, the marginalization process is designed to incorporate and account for such biases. Therefore, it is expected that estimates of the parameters will remain unbiased and achieve minimal variance, even if the cosmology of the observed data deviates from the fiducial model.

	$f_{\text{NL}}^{\text{local}}$	$f_{\text{NL}}^{\text{equil}}$	$f_{\text{NL}}^{\text{orth}}$
$P_{\text{pre}}$	199.26 (74.90)	853.81 (217.63)	253.91 (75.72)
$P_{\text{post}}$	184.13 (91.38)	368.90 (86.65)	137.23 (83.03)
$P_{\text{pre}} + P_{\text{post}}$	128.41 (69.81)	227.51 (80.62)	76.31 (53.26)
$B_{\text{pre}}$	135.90 (76.31)	188.33 (217.78)	64.15 (75.49)
$B_{\text{post}}$	100.40 (69.04)	100.40 (74.35)	37.58 (31.79)
$B_{\text{pre}} + B_{\text{post}}$	80.56 (62.00)	90.29 (66.94)	34.39 (29.61)
$(P + B)_{\text{pre}}$	86.10 (55.32)	159.39 (111.19)	47.25 (40.88)
$(P + B)_{\text{post}}$	86.17 (53.78)	90.83 (64.88)	34.12 (29.26)
$(P + B)_{\text{pre}} + (P + B)_{\text{post}}$	64.95 (46.37)	84.61 (61.97)	30.13 (26.56)

**Table 2.** (Un)marginalized errors on  $f_{\text{NL}}^{\text{shape}}$  for all options of statistics estimated with the combined Fisher estimator from section 4.3.3.

$f_{\text{NL}}^{\text{local}}$	$f_{\text{NL}}^{\text{equil}}$	$f_{\text{NL}}^{\text{orth}}$	$h$	$n_s$	$\Omega_m$	$\Omega_b$	$\sigma_8$
1.33	1.88	1.57	1.75	1.44	1.55	1.76	1.51

**Table 3.** Improvement factor of parameter constraints when using the full data product  $(P + B)_{\text{pre}} + (P + B)_{\text{post}}$  as compared to just  $(P + B)_{\text{pre}}$ .


### 4.5 Parameter constraints

Table 2 displays the (un)marginalized errors on the pnG parameters, comparing the **pre**, **post**, and **(pre+post)** data points across power spectra (P), bispectra (B), and joint power spectra and bispectra (P+B). The decoupling of modes, the mitigation of redshift-space distortions, and focusing on pnG parameters the reduction of parameter degeneracy are the main factors contributing to this improvement. The improvement factors, quantified by the ratio of these marginalized errors, are presented in table 3. An exception in the improvement pattern is the local pnG parameter  $f_{\text{NL}}^{\text{local}}$ , which shows a relatively lower improvement factor compared to the equilateral  $f_{\text{NL}}^{\text{equil}}$  and orthogonal  $f_{\text{NL}}^{\text{ortho}}$  pnG parameters. This difference can be attributed to the scale-dependent bias present in the **pre**-reconstruction dark matter halo field, a phenomenon well documented in previous studies [46, 91, 92]. Figure 11 in the appendix, confirms the presence of this bias, which enhances the informative power of the power spectrum for  $f_{\text{NL}}^{\text{local}}$ .

Figure 7 displays the resulting confidence intervals and the marginalized combined Fisher errors for the eight parameters (five cosmological parameters and three pnG parameters) when fitting them simultaneously, using both power spectrum and bispectrum data.

## 5 Conclusions

We have presented a neural network-based approach to reconstructing the linear density field from late-time dark matter halo data. Our findings demonstrate that the reconstruction quality for a  $1\text{Gpc}^3$  volume with average number density of halos  $\bar{n} = 4 \times 10^{-4} (\text{h/Mpc})^3$ ,



**Figure 7.** Comparison of the constraining power for the **pre**, **post**, and **pre+post** measured power spectrum and bispectrum, jointly fitted for all 8 parameters. The coloured areas indicate the  $2\sigma$  constraints, highlighting the sensitivities and different degeneracy directions. The diagonal plots represent the marginalized probability distributions for each parameter, while the off-diagonal plots illustrate the confidence intervals for the parameter pairs.

as measured by the cross-correlation with the linear dark matter density field, achieves approximately 44% accuracy for scales  $k \leq 0.4h/\text{Mpc}$ . This is a strong improvement over the cross-correlation of the unreconstructed halo fields with the initial dark matter density field. We furthermore find strong evidence that our reconstruction is optimal for this data.

By performing the reconstruction and a subsequent Fisher analysis on a large part of the Quijote simulation suite, we show that this reconstruction leads to a significant overall improvement in cosmological parameter constraints. Focussing on primordial non-Gaussianity, we find that the marginalized errors on pnG from a joint analysis of the power spectrum (monopole and quadrupole) and bispectrum are improved by factors of 1.33, 1.88, and 1.57 for local, equilateral, and orthogonal pnG. This improvement is obtained, because the reconstruction method mitigates redshift-space distortions, decouples modes, and decreases

degeneracies among cosmological parameters. Since shot noise and non-Gaussian covariance cause saturation of parameter constraints from the dark matter halo field power spectrum and bispectrum [46], such improvements are not easily achieved by simply including more modes in the halo field analysis. The improvement is significantly lower compared to the dark matter case in [44], since reconstruction is limited by shot noise in the dark matter halo fields. Nevertheless, our reconstruction method could potentially be applied to complement existing observational data analysis methods, leading to improvements in cosmological parameter constraints. Improvements that would otherwise require an increase in number density [26], or a larger survey volume [93], both of which are costly and observationally challenging compared to our reconstruction method.

A generic limitation of the type of reconstruction performed in this work is the GPU memory constraint that restricted our grid resolution to  $256^3$  cells, limiting the amount of modes input to the network to  $k \leq 1.4$  h/Mpc. Although we find evidence that in our case (i.e. resolution and number density) the reconstruction is already optimal, for higher number densities this limited resolution could limit the accuracy of reconstruction, especially on smaller scales. However, on small scales, shell crossing of dark matter streams, and virialization of bound structures puts a fundamental limit on the reconstructability of linear modes, implying diminishing returns from including additional modes. Further research is required to determine this limit. Also, the effect of the stochastic component after reconstruction is not well understood but could lead to tighter parameter constraints when accurate noise subtraction is applied [29]. The challenge in determining the exact stochastic contribution is in part due to the low interpretability of the neural network. Moreover, while [82] justified the assumption of a Gaussian likelihood in our statistics for the Fisher analysis, [94] demonstrated that this assumption can lead to artificially tight bounds on parameter constraints when the statistics are not truly Gaussian. Also, the reconstruction is performed with a fixed bias term  $b_1$ , which means we could not marginalize over this term in the analysis. Another concern is that in simulations, the exact positions of dark matter halo tracers are known in redshift space, which may not be achievable in real observational data.

Future research could focus on applying this reconstruction method to real tracers such as galaxies or atomic spectral lines instead of dark matter halo fields. Although in this study we have used the halos as a proxy for such tracers, for practical application on observational data, this method should be further developed on even more realistic tracers. It would also be worth exploring different inputs beyond just the halo number density. For instance, incorporating velocity or void information could potentially make the input more informative for the U-Net in reconstructing the initial density distribution [40]. Studying the impact of these additions on reconstruction accuracy would be interesting and could lead to better results depending on the information available in actual observational data.

This work was motivated by a previous study on the dark matter field. We have demonstrated that reconstructing the initial conditions using a U-Net architecture neural network model on  $z = 0$  simulated dark matter halo fields can provide non-trivial improvements on cosmological parameter constraints. Our analysis shows that the U-Net can, to a certain degree, reverse the non-linear collapse of matter, and decoupling modes. The resulting field is more linear and Gaussian, and its power spectrum and bispectrum contain additional

information on cosmological parameters. We have identified several avenues for further investigation, with the ultimate goal to explore the potential of these reconstructions on real data.

## Acknowledgments

The authors like to thank Will Coulton and Florian List for comments on the manuscript. We thank the Quijote team for making the simulation suite available [47]. We thank the Center for Information Technology of the University of Groningen for providing access to the Hábrók high-performance computing clusters. We thank the developers of NumPy, JAX, Matplotlib, Optax for providing the open-source tools used in this research.


## A Modifying shot noise

Since shot noise is the main cause of saturation of cosmological parameter constraints [46], it is worth exploring how varying levels of shot noise in the input fields affect the accuracy of reconstruction, and thereby parameter constraints. We have examined two scenarios: increasing the halo mass threshold, which decreases the halo number density and thus increases shot noise by definition, and applying halo mass weighting when constructing the halo density fields, as this has been shown to effectively reduce shot noise [95]. The simulated halo catalogues provide detailed mass information, offering more insights into the matter distribution. By comparing reconstructions with and without mass information, we assess the impact of this additional information on the quality of the reconstruction.

When generating halo density fields from the halo catalogue, each halo’s mass is given, and determined by the number of dark matter particles it contains. For the main analysis in this paper, we used the lowest halo mass threshold  $M_{\min} \geq 1.32 \cdot 10^{13} M_{\odot}/h$  (20 dark matter particles) from the Quijote simulation suite, which results in an average number density of approximately  $n \geq 4 \cdot 10^{-4} h^3/\text{Mpc}^3$ . Raising the threshold decreases the number density, which inherently amplifies shot noise. This adjustment enables us to evaluate the performance of the reconstruction method under conditions of increased shot noise.

Mass information available from observational data could improve reconstruction accuracy compared to the basic setup. To investigate this, we classified halos into mass bins, with higher bin numbers indicating more massive halos. Increasing the number of bins provides more detailed mass information. Next, we can incorporate this mass information in our 3D density fields with the mass weighting using PCS, reducing the shot noise [95] in the fields, allowing us to assess the impact of varying degrees of detail in mass weighting on reconstruction accuracy.

First, we constructed density fields with an increased minimum halo mass threshold. These density fields were used to train the U-Net. To evaluate the changes in reconstruction quality by varying the number density of the initial field, we compared the cross-correlation with the initial density field. The results are shown in figure 8(a). As expected, the quality of reconstruction decreases with decreasing number density. On large scales up to  $k = 0.1 h/\text{Mpc}$ , the reconstruction quality is approximately the same for different number densities. However, for lower number densities, the reconstruction clearly deteriorated more rapidly at smaller scales.

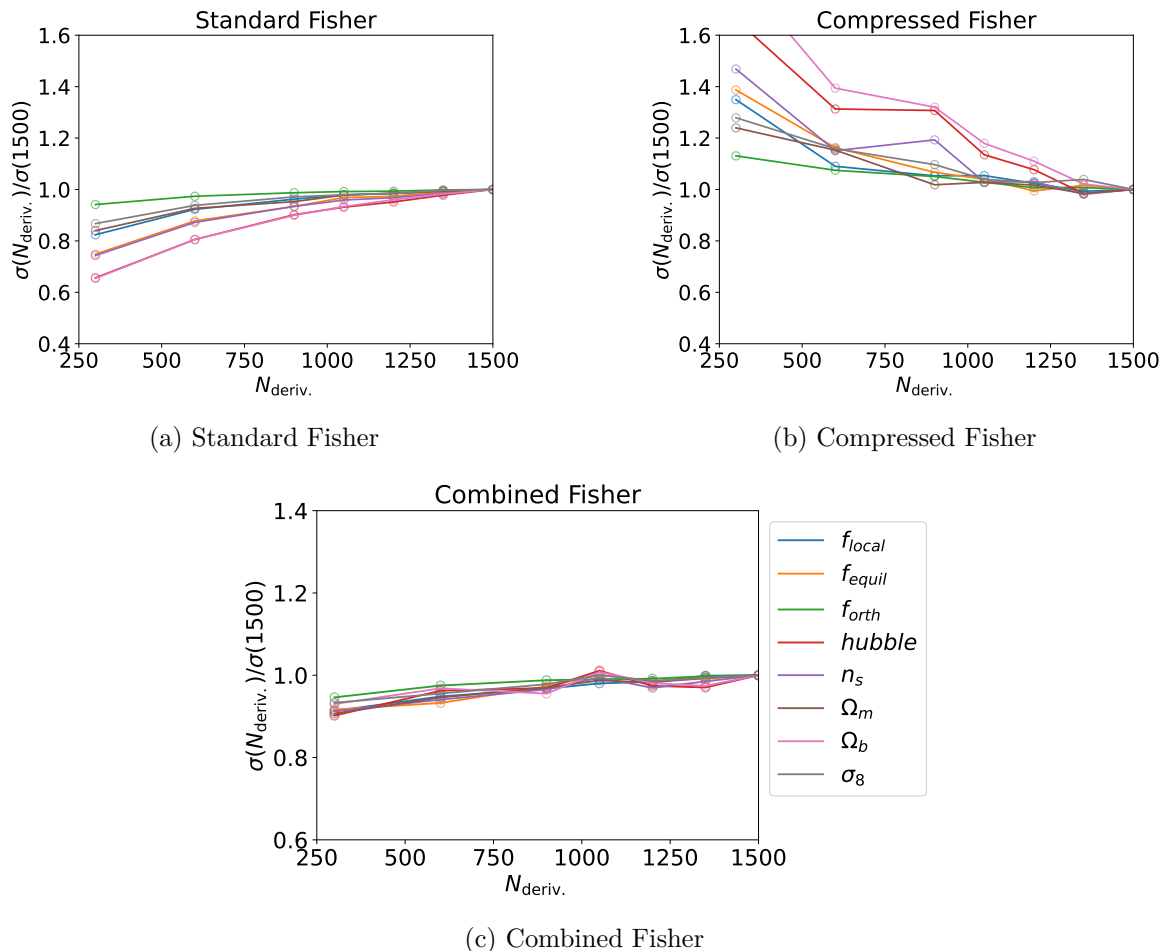


**Figure 8.** Both sub-figures display the cross-correlation of reconstructed density fields with varying shot noise. (a) shows the impact of varying halo number density on reconstruction quality, and (b) shows the effect of different mass weightings accuracy in the density fields on reconstruction quality. Here, full means that the model has the exact mass information of each dark matter halo (This differs from the underlying dark matter field because the halo threshold is set to 20 dark matter particles. Consequently, regions below this threshold are devoid of dark matter halos).

Second, we explore how our model changes with different levels of detail in mass information. Until now, we have only used the number density of dark matter halos to train the model. However, in galaxy surveys, an estimate of galaxy masses could provide additional information on the distribution of matter, effectively reducing shot noise [95]. As an analogue, the Quijote halo catalogue provides specific information about the mass of each halo. This additional mass information could potentially enhance the model’s ability to reconstruct the initial conditions. To investigate the extent of this improvement, we train different models with varying levels of mass information. By again comparing the cross-correlations for these different cases, we study how our reconstruction method benefits from having more detailed mass distribution information (reduced shot noise). Figure 8(b) illustrates that incorporating additional mass information into the input density fields improves the quality of reconstruction. This improvement arises because the model gains more insight into the underlying dark matter field compared to scenarios without mass information.

## B Convergence of Fisher analysis

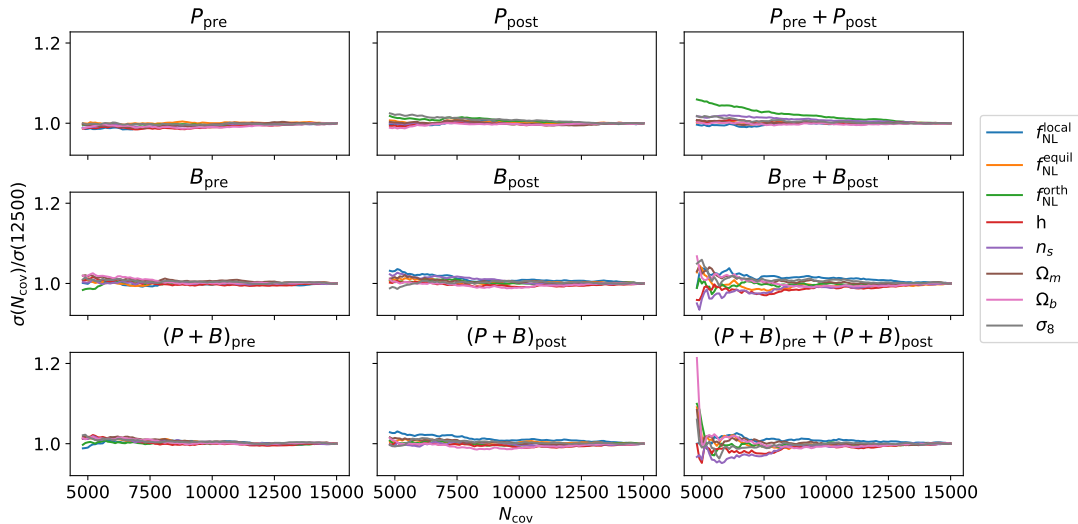
In section 4.3.3 we introduced the usage of an unbiased combined Fisher matrix that takes the geometric mean from a biased high standard Fisher matrix and a biased low compressed Fisher matrix, in order to overcome the Monte Carlo uncertainties arising from the limited number of simulations to compute the numerical derivatives [79]. Figure 9(a) demonstrates the behaviour of the standard Fisher matrix estimate as the number of derivatives,  $N_{\text{deriv}}$ , increases. As expected, the standard Fisher information tends to be biased high, indicated by the upward trend.



**Figure 9.** Convergence of Fisher matrix estimates for  $(P + B)_{\text{pre}} + (P + B)_{\text{post}}$ : (a) Standard Fisher, (b) Compressed Fisher, and (c) Combined Fisher. The combined estimator balances the biases present in the standard and compressed Fisher matrices, resulting in unbiased estimates.

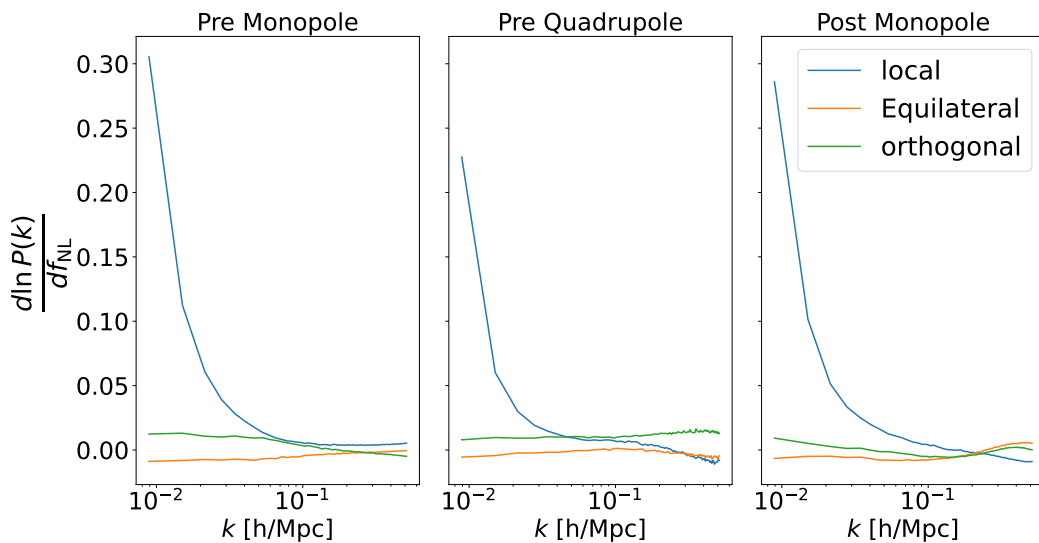
In figure 9(b), the compressed Fisher matrix estimate shows a different trend. Here, the values tend to decrease as  $N_{\text{deriv}}$  increases, indicating that the compressed Fisher information is biased low. This low bias is due to the reduction of noise in the derivatives of the compressed statistic [79].

Figure 9(c) presents the combined Fisher matrix estimate, which merges the standard and compressed estimates using the geometric mean approach. This combined estimator is converged, stabilizing around a value close to 1. The convergence observed in the combined Fisher matrix suggests that it balances the high bias of the standard Fisher matrix and the low bias of the compressed Fisher matrix. The results validate the combined estimator as an unbiased method for computing the Fisher matrix, which is also validated from the minimal-variance marginalized estimator in equation (4.18)). Furthermore, the convergence of the covariance matrix is demonstrated in figure 10.



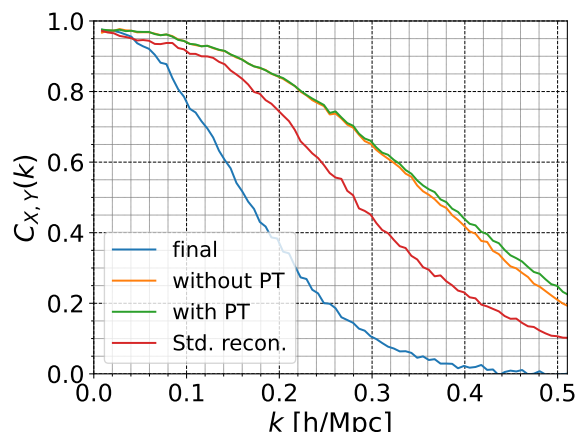
**Figure 10.** Variation of the marginalized parameter constraints as a function of the number of simulations included in the computation of the covariance matrix. The figure shows that the covariance is converged for all the scenarios analysed in this thesis.

### Scale-dependent bias



**Figure 11.** The derivative of power spectra with respect to the amplitudes of the three types of pnG, monopole, and quadrupole plotted for pre-reconstruction and monopole for post-reconstruction. It highlights the scale-dependent bias present in local pnG, which explains the low improvement factor after reconstruction. This is because information from smaller scales is already incorporated into the larger scales in the pre-reconstructed power spectrum.

## Reconstruction without gradient field contribution



**Figure 12.** The cross-correlation of density fields at a single realization. The orange line shows the performance of the model without the gradient fields in equations (3.1) and (3.2), and is compared to the model that incorporated these gradient fields.

To showcase the use of inputting the gradient fields in equations (3.1) and (3.2), we also trained the model without incorporating the gradient fields. Since these fields do not introduce additional information to the network beyond what is already contained in the density field, they were primarily introduced to improve computational efficiency and convergence. This is supported by figure 12, which shows that the model without PT achieves nearly the same reconstruction accuracy. However, more training time was required to reach this level of performance. Depending on the complexity of the network, it might still not reach the same accuracy achieved when incorporating the gradient fields.

## References

- [1] A. Achúcarro et al., *Inflation: Theory and Observations*, [arXiv:2203.08128](#) [INSPIRE].
- [2] P.D. Meerburg et al., *Primordial Non-Gaussianity*, *Bull. Am. Astron. Soc.* **51** (2019) 107 [[arXiv:1903.04409](#)] [INSPIRE].
- [3] P.G. Ferreira, J. Magueijo and K.M. Gorski, *Evidence for nonGaussianity in the DMR four year sky maps*, *Astrophys. J. Lett.* **503** (1998) L1 [[astro-ph/9803256](#)] [INSPIRE].
- [4] E. Komatsu and D.N. Spergel, *Acoustic signatures in the primary microwave background bispectrum*, *Phys. Rev. D* **63** (2001) 063002 [[astro-ph/0005036](#)] [INSPIRE].
- [5] WMAP collaboration, *First year Wilkinson Microwave Anisotropy Probe (WMAP) observations: tests of gaussianity*, *Astrophys. J. Suppl.* **148** (2003) 119 [[astro-ph/0302223](#)] [INSPIRE].
- [6] P. Creminelli et al., *Limits on non-gaussianities from WMAP data*, *JCAP* **05** (2006) 004 [[astro-ph/0509029](#)] [INSPIRE].
- [7] PLANCK collaboration, *Planck 2015 results. XVII. Constraints on primordial non-Gaussianity*, *Astron. Astrophys.* **594** (2016) A17 [[arXiv:1502.01592](#)] [INSPIRE].
- [8] PLANCK collaboration, *Planck 2018 results. IX. Constraints on primordial non-Gaussianity*, *Astron. Astrophys.* **641** (2020) A9 [[arXiv:1905.05697](#)] [INSPIRE].

- [9] SIMONS OBSERVATORY collaboration, *The Simons Observatory: Science goals and forecasts*, *JCAP* **02** (2019) 056 [[arXiv:1808.07445](#)] [[INSPIRE](#)].
- [10] CMB-S4 collaboration, *CMB-S4 Science Book, First Edition*, [arXiv:1610.02743](#) [[INSPIRE](#)].
- [11] A. Kalaja, P.D. Meerburg, G.L. Pimentel and W.R. Coulton, *Fundamental limits on constraining primordial non-Gaussianity*, *JCAP* **04** (2021) 050 [[arXiv:2011.09461](#)] [[INSPIRE](#)].
- [12] R. Catena and P. Ullio, *A novel determination of the local dark matter density*, *JCAP* **08** (2010) 004 [[arXiv:0907.0018](#)] [[INSPIRE](#)].
- [13] J.I. Read, *The Local Dark Matter Density*, *J. Phys. G* **41** (2014) 063101 [[arXiv:1404.1938](#)] [[INSPIRE](#)].
- [14] J.R. Pritchard and A. Loeb, *21-cm cosmology*, *Rept. Prog. Phys.* **75** (2012) 086901 [[arXiv:1109.6012](#)] [[INSPIRE](#)].
- [15] E.D. Kovetz et al., *Line-Intensity Mapping: 2017 Status Report*, [arXiv:1709.09066](#) [[INSPIRE](#)].
- [16] F. Bernardeau, S. Colombi, E. Gaztanaga and R. Scoccimarro, *Large scale structure of the universe and cosmological perturbation theory*, *Phys. Rept.* **367** (2002) 1 [[astro-ph/0112551](#)] [[INSPIRE](#)].
- [17] A. Taruya, F. Bernardeau, T. Nishimichi and S. Codis, *RegPT: Direct and fast calculation of regularized cosmological power spectrum at two-loop order*, *Phys. Rev. D* **86** (2012) 103528 [[arXiv:1208.1191](#)] [[INSPIRE](#)].
- [18] J.J.M. Carrasco, M.P. Hertzberg and L. Senatore, *The Effective Field Theory of Cosmological Large Scale Structures*, *JHEP* **09** (2012) 082 [[arXiv:1206.2926](#)] [[INSPIRE](#)].
- [19] M. Garny, D. Laxhuber and R. Scoccimarro, *Perturbation theory with dispersion and higher cumulants: Nonlinear regime*, *Phys. Rev. D* **107** (2023) 063540 [[arXiv:2210.08089](#)] [[INSPIRE](#)].
- [20] BOSS collaboration, *The Clustering of Galaxies in SDSS-III DR9 Baryon Oscillation Spectroscopic Survey: Constraints on Primordial Non-Gaussianity*, *Mon. Not. Roy. Astron. Soc.* **428** (2013) 1116 [[arXiv:1208.1491](#)] [[INSPIRE](#)].
- [21] B. Leistedt, H.V. Peiris and N. Roth, *Constraints on Primordial Non-Gaussianity from 800 000 Photometric Quasars*, *Phys. Rev. Lett.* **113** (2014) 221301 [[arXiv:1405.4315](#)] [[INSPIRE](#)].
- [22] S. Ho et al., *Sloan Digital Sky Survey III photometric quasar clustering: probing the initial conditions of the Universe*, *JCAP* **05** (2015) 040 [[arXiv:1311.2597](#)] [[INSPIRE](#)].
- [23] EBOSS collaboration, *Redshift-weighted constraints on primordial non-Gaussianity from the clustering of the eBOSS DR14 quasars in Fourier space*, *JCAP* **09** (2019) 010 [[arXiv:1904.08859](#)] [[INSPIRE](#)].
- [24] EBOSS collaboration, *Primordial non-Gaussianity from the completed SDSS-IV extended Baryon Oscillation Spectroscopic Survey II: measurements in Fourier space with optimal weights*, *Mon. Not. Roy. Astron. Soc.* **514** (2022) 3396 [[arXiv:2106.13725](#)] [[INSPIRE](#)].
- [25] G. Cabass et al., *Constraints on Single-Field Inflation from the BOSS Galaxy Survey*, *Phys. Rev. Lett.* **129** (2022) 021301 [[arXiv:2201.07238](#)] [[INSPIRE](#)].
- [26] G. D’Amico, M. Lewandowski, L. Senatore and P. Zhang, *Limits on primordial non-Gaussianities from BOSS galaxy-clustering data*, *Phys. Rev. D* **111** (2025) 063514 [[arXiv:2201.11518](#)] [[INSPIRE](#)].
- [27] G. Cabass et al., *Constraints on multifield inflation from the BOSS galaxy survey*, *Phys. Rev. D* **106** (2022) 043506 [[arXiv:2204.01781](#)] [[INSPIRE](#)].

- [28] T. Kurita and M. Takada, *Constraints on anisotropic primordial non-Gaussianity from intrinsic alignments of SDSS-III BOSS galaxies*, *Phys. Rev. D* **108** (2023) 083533 [[arXiv:2302.02925](#)] [[INSPIRE](#)].
- [29] K.C. Chan and L. Blot, *Assessment of the Information Content of the Power Spectrum and Bispectrum*, *Phys. Rev. D* **96** (2017) 023528 [[arXiv:1610.06585](#)] [[INSPIRE](#)].
- [30] M. Biagetti, L. Castiblanco, J. Noreña and E. Sefusatti, *The covariance of squeezed bispectrum configurations*, *JCAP* **09** (2022) 009 [[arXiv:2111.05887](#)] [[INSPIRE](#)].
- [31] W.R. Coulton et al., *Quijote-PNG: Simulations of Primordial Non-Gaussianity and the Information Content of the Matter Field Power Spectrum and Bispectrum*, *Astrophys. J.* **943** (2023) 64 [[arXiv:2206.01619](#)] [[INSPIRE](#)].
- [32] T. Flöss, M. Biagetti and P.D. Meerburg, *Primordial non-Gaussianity and non-Gaussian covariance*, *Phys. Rev. D* **107** (2023) 023528 [[arXiv:2206.10458](#)] [[INSPIRE](#)].
- [33] D.J. Eisenstein, H.-J. Seo, E. Sirko and D. Spergel, *Improving Cosmological Distance Measurements by Reconstruction of the Baryon Acoustic Peak*, *Astrophys. J.* **664** (2007) 675 [[astro-ph/0604362](#)] [[INSPIRE](#)].
- [34] Y. Wang et al., *Extracting high-order cosmological information in galaxy surveys with power spectra*, *Commun. Phys.* **7** (2024) 130 [[arXiv:2202.05248](#)] [[INSPIRE](#)].
- [35] M. Shirasaki, N.S. Sugiyama, R. Takahashi and F.-S. Kitaura, *Constraining primordial non-Gaussianity with postreconstructed galaxy bispectrum in redshift space*, *Phys. Rev. D* **103** (2021) 023506 [[arXiv:2010.04567](#)] [[INSPIRE](#)].
- [36] D. Jamieson et al., *Simple lessons from complex learning: what a neural network model learns about cosmic structure formation*, *PNAS Nexus* **2** (2023) gac250 [[arXiv:2206.04573](#)] [[INSPIRE](#)].
- [37] V. Jindal et al., *Predicting the Initial Conditions of the Universe using a Deterministic Neural Network*, [arXiv:2303.13056](#) [[INSPIRE](#)].
- [38] C.J. Shallue and D.J. Eisenstein, *Reconstructing cosmological initial conditions from late-time structure with convolutional neural networks*, *Mon. Not. Roy. Astron. Soc.* **520** (2023) 6256 [[arXiv:2207.12511](#)] [[INSPIRE](#)].
- [39] F. List, N. Anau Montel and C. Weniger, *Bayesian Simulation-based Inference for Cosmological Initial Conditions*, in the proceedings of the *37th Conference on Neural Information Processing Systems*, New Orleans, U.S.A., December 10–16 (2023) [[arXiv:2310.19910](#)] [[INSPIRE](#)].
- [40] A.E. Bayer, C. Modi and S. Ferraro, *Joint velocity and density reconstruction of the Universe with nonlinear differentiable forward modeling*, *JCAP* **06** (2023) 046 [[arXiv:2210.15649](#)] [[INSPIRE](#)].
- [41] R. Legin et al., *Posterior sampling of the initial conditions of the universe from non-linear large scale structures using score-based generative models*, *Mon. Not. Roy. Astron. Soc.* **527** (2023) L173 [[arXiv:2304.03788](#)] [[INSPIRE](#)].
- [42] X. Chen, F. Zhu, S. Gaines and N. Padmanabhan, *Effective cosmic density field reconstruction with convolutional neural network*, *Mon. Not. Roy. Astron. Soc.* **523** (2023) 6272 [[arXiv:2306.10538](#)] [[INSPIRE](#)].
- [43] X. Chen, N. Padmanabhan and D.J. Eisenstein, *Probing primordial non-Gaussianity by reconstructing the initial conditions*, [arXiv:2412.00968](#) [[INSPIRE](#)].
- [44] T. Flöss and P.D. Meerburg, *Improving constraints on primordial non-Gaussianity using neural network based reconstruction*, *JCAP* **02** (2024) 031 [[arXiv:2305.07018](#)] [[INSPIRE](#)].

- [45] M. Schmittfull, T. Baldauf and M. Zaldarriaga, *Iterative initial condition reconstruction*, *Phys. Rev. D* **96** (2017) 023505 [[arXiv:1704.06634](#)] [[INSPIRE](#)].
- [46] W.R. Coulton et al., *Quijote-PNG: The Information Content of the Halo Power Spectrum and Bispectrum*, *Astrophys. J.* **943** (2023) 178 [[arXiv:2206.15450](#)] [[INSPIRE](#)].
- [47] F. Villaescusa-Navarro et al., *The Quijote simulations*, *Astrophys. J. Suppl.* **250** (2020) 2 [[arXiv:1909.05273](#)] [[INSPIRE](#)].
- [48] P.J.E. Peebles, *The galaxy and mass n-point correlation functions: a blast from the past*, *ASP Conf. Ser.* **252** (2001) 201 [[astro-ph/0103040](#)] [[INSPIRE](#)].
- [49] D. Babich, P. Creminelli and M. Zaldarriaga, *The Shape of non-Gaussianities*, *JCAP* **08** (2004) 009 [[astro-ph/0405356](#)] [[INSPIRE](#)].
- [50] J.M. Maldacena, *Non-Gaussian features of primordial fluctuations in single field inflationary models*, *JHEP* **05** (2003) 013 [[astro-ph/0210603](#)] [[INSPIRE](#)].
- [51] P. Creminelli and M. Zaldarriaga, *Single field consistency relation for the 3-point function*, *JCAP* **10** (2004) 006 [[astro-ph/0407059](#)] [[INSPIRE](#)].
- [52] P. Creminelli, *On non-Gaussianities in single-field inflation*, *JCAP* **10** (2003) 003 [[astro-ph/0306122](#)] [[INSPIRE](#)].
- [53] M. Alishahiha, E. Silverstein and D. Tong, *DBI in the sky*, *Phys. Rev. D* **70** (2004) 123505 [[hep-th/0404084](#)] [[INSPIRE](#)].
- [54] A. Gruzinov, *Consistency relation for single scalar inflation*, *Phys. Rev. D* **71** (2005) 027301 [[astro-ph/0406129](#)] [[INSPIRE](#)].
- [55] X. Chen, M.-X. Huang, S. Kachru and G. Shiu, *Observational signatures and non-Gaussianities of general single field inflation*, *JCAP* **01** (2007) 002 [[hep-th/0605045](#)] [[INSPIRE](#)].
- [56] C. Cheung et al., *The Effective Field Theory of Inflation*, *JHEP* **03** (2008) 014 [[arXiv:0709.0293](#)] [[INSPIRE](#)].
- [57] D. Langlois, S. Renaux-Petel, D.A. Steer and T. Tanaka, *Primordial fluctuations and non-Gaussianities in multi-field DBI inflation*, *Phys. Rev. Lett.* **101** (2008) 061301 [[arXiv:0804.3139](#)] [[INSPIRE](#)].
- [58] D. Langlois, S. Renaux-Petel, D.A. Steer and T. Tanaka, *Primordial perturbations and non-Gaussianities in DBI and general multi-field inflation*, *Phys. Rev. D* **78** (2008) 063523 [[arXiv:0806.0336](#)] [[INSPIRE](#)].
- [59] L. Senatore, K.M. Smith and M. Zaldarriaga, *Non-Gaussianities in Single Field Inflation and their Optimal Limits from the WMAP 5-year Data*, *JCAP* **01** (2010) 028 [[arXiv:0905.3746](#)] [[INSPIRE](#)].
- [60] S. Renaux-Petel, *Orthogonal non-Gaussianities from Dirac-Born-Infeld Galileon inflation*, *Class. Quant. Grav.* **28** (2011) 182001 [Erratum *ibid.* **28** (2011) 249601] [[arXiv:1105.6366](#)] [[INSPIRE](#)].
- [61] E. Pajer, *Building a Boostless Bootstrap for the Bispectrum*, *JCAP* **01** (2021) 023 [[arXiv:2010.12818](#)] [[INSPIRE](#)].
- [62] N. Dalal, O. Doré, D. Huterer and A. Shirokov, *The imprints of primordial non-gaussianities on large-scale structure: scale dependent bias and abundance of virialized objects*, *Phys. Rev. D* **77** (2008) 123514 [[arXiv:0710.4560](#)] [[INSPIRE](#)].
- [63] A. Slosar et al., *Constraints on local primordial non-Gaussianity from large scale structure*, *JCAP* **08** (2008) 031 [[arXiv:0805.3580](#)] [[INSPIRE](#)].

- [64] F. Schmidt and M. Kamionkowski, *Halo Clustering with Non-Local Non-Gaussianity*, *Phys. Rev. D* **82** (2010) 103002 [[arXiv:1008.0638](#)] [[INSPIRE](#)].
- [65] J.M. Bardeen, J.R. Bond, N. Kaiser and A.S. Szalay, *The Statistics of Peaks of Gaussian Random Fields*, *Astrophys. J.* **304** (1986) 15 [[INSPIRE](#)].
- [66] N. Kaiser, *On the Spatial correlations of Abell clusters*, *Astrophys. J. Lett.* **284** (1984) L9 [[INSPIRE](#)].
- [67] M. Biagetti, *The Hunt for Primordial Interactions in the Large Scale Structures of the Universe*, *Galaxies* **7** (2019) 71 [[arXiv:1906.12244](#)] [[INSPIRE](#)].
- [68] O. Ronneberger, P. Fischer and T. Brox, *U-Net: Convolutional Networks for Biomedical Image Segmentation*, in *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III* 18, Springer (2015), pp. 234–241, [[arXiv:1505.04597](#)] [[INSPIRE](#)].
- [69] T.L. Makinen et al., *deep21: a deep learning method for 21 cm foreground removal*, *JCAP* **04** (2021) 081 [[arXiv:2010.15843](#)] [[INSPIRE](#)].
- [70] S. Gagnon-Hartman et al., *Recovering the wedge modes lost to 21-cm foregrounds*, *Mon. Not. Roy. Astron. Soc.* **504** (2021) 4716 [[arXiv:2102.08382](#)] [[INSPIRE](#)].
- [71] A. Taruya, T. Nishimichi and D. Jeong, *Grid-based calculation for perturbation theory of large-scale structure*, *Phys. Rev. D* **98** (2018) 103532 [[arXiv:1807.04215](#)] [[INSPIRE](#)].
- [72] V. Springel, *The cosmological simulation code GADGET-2*, *Mon. Not. Roy. Astron. Soc.* **364** (2005) 1105 [[astro-ph/0505010](#)] [[INSPIRE](#)].
- [73] M. Davis, G. Efstathiou, C.S. Frenk and S.D.M. White, *The Evolution of Large Scale Structure in a Universe Dominated by Cold Dark Matter*, *Astrophys. J.* **292** (1985) 371 [[INSPIRE](#)].
- [74] R. Scoccimarro, *Redshift-space distortions, pairwise velocities and nonlinearities*, *Phys. Rev. D* **70** (2004) 083007 [[astro-ph/0407214](#)] [[INSPIRE](#)].
- [75] A.K. Chaniotis and D. Poulikakos, *High order interpolation and differentiation using B-splines*, *J. Comput. Phys.* **197** (2004) 253.
- [76] E. Sefusatti, M. Crocce, R. Scoccimarro and H. Couchman, *Accurate Estimators of Correlation Functions in Fourier Space*, *Mon. Not. Roy. Astron. Soc.* **460** (2016) 3624 [[arXiv:1512.07295](#)] [[INSPIRE](#)].
- [77] A. Tarvainen and H. Valpola, *Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results*, *Adv. Neural Inf. Process. Syst.* **30** (2017) [[arXiv:1703.01780](#)].
- [78] X. Chen et al., *Extensive analysis of reconstruction algorithms for DESI 2024 baryon acoustic oscillations*, [[arXiv:2411.19738](#)] [[INSPIRE](#)].
- [79] W.R. Coulton and B.D. Wandelt, *How to estimate Fisher information matrices from simulations*, [[arXiv:2305.08994](#)] [[INSPIRE](#)].
- [80] Y.P. Jing, *Correcting for the alias effect when measuring the power spectrum using FFT*, *Astrophys. J.* **620** (2005) 559 [[astro-ph/0409240](#)] [[INSPIRE](#)].
- [81] R.A. Fisher, *On the Mathematical Foundations of Theoretical Statistics*, *Phil. Trans. Roy. Soc. Lond. A* **222** (1922) 309 [[INSPIRE](#)].
- [82] R. Scoccimarro, *The bispectrum: from theory to observations*, *Astrophys. J.* **544** (2000) 597 [[astro-ph/0004086](#)] [[INSPIRE](#)].

- [83] J. Carron, *On the assumption of Gaussianity for cosmological two-point statistics and parameter dependent covariance matrices*, *Astron. Astrophys.* **551** (2013) A88 [[arXiv:1204.4724](#)] [[INSPIRE](#)].
- [84] J. Hartlap, P. Simon and P. Schneider, *Why your model parameter confidences might be too optimistic: Unbiased estimation of the inverse covariance matrix*, *Astron. Astrophys.* **464** (2007) 399 [[astro-ph/0608064](#)] [[INSPIRE](#)].
- [85] T.W. Anderson, T.W. Anderson, T.W. Anderson and T.W. Anderson, *An introduction to multivariate statistical analysis*, 3rd edn., Wiley-Interscience (2003).
- [86] M. Tegmark, A. Taylor and A. Heavens, *Karhunen-Loeve eigenvalue problems in cosmology: How should we tackle large data sets?*, *Astrophys. J.* **480** (1997) 22 [[astro-ph/9603021](#)] [[INSPIRE](#)].
- [87] A. Heavens, R. Jimenez and O. Lahav, *Massive lossless data compression and multiple parameter estimation from galaxy spectra*, *Mon. Not. Roy. Astron. Soc.* **317** (2000) 965 [[astro-ph/9911102](#)] [[INSPIRE](#)].
- [88] R. Bhatia, *Positive definite matrices*, Princeton University Press (2009).
- [89] C.R. Rao, *Information and the accuracy attainable in the estimation of statistical parameters*, *Bull. Calcutta Math. Soc.* **37** (1945) 81.
- [90] H. Cramér, *Mathematical Methods of Statistics (PMS-9)*, vol. 9 of *Princeton Mathematical Series*, Princeton University Press (1946).
- [91] C. Wagner and L. Verde, *N-body simulations with generic non-Gaussian initial conditions II: Halo bias*, *JCAP* **03** (2012) 002 [[arXiv:1102.3229](#)] [[INSPIRE](#)].
- [92] R. Scoccimarro, L. Hui, M. Manera and K.C. Chan, *Large-scale Bias and Efficient Generation of Initial Conditions for Non-Local Primordial Non-Gaussianity*, *Phys. Rev. D* **85** (2012) 083002 [[arXiv:1108.5512](#)] [[INSPIRE](#)].
- [93] G. Cabass et al., *Constraining single-field inflation with MegaMapper*, *Phys. Lett. B* **841** (2023) 137912 [[arXiv:2211.14899](#)] [[INSPIRE](#)].
- [94] C.F. Park, E. Allys, F. Villaescusa-Navarro and D.P. Finkbeiner, *Quantification of High-dimensional Non-Gaussianities and Its Implication to Fisher Analysis in Cosmology*, *Astrophys. J.* **946** (2023) 107 [[arXiv:2204.05435](#)] [[INSPIRE](#)].
- [95] N. Hamaus et al., *Minimizing the Stochasticity of Halos in Large-Scale Structure Surveys*, *Phys. Rev. D* **82** (2010) 043515 [[arXiv:1004.5377](#)] [[INSPIRE](#)].