


PAPER

On the primordial information available to galaxy redshift surveys

To cite this article: Matthew McQuinn JCAP06(2021)024

View the [article online](#) for updates and enhancements.



IOP | ebooks™

Bringing together innovative digital publishing with leading authors from the global scientific community.

Start exploring the collection—download the first chapter of every title for free.

On the primordial information available to galaxy redshift surveys

Matthew McQuinn

Department of Astronomy, University of Washington,
Seattle, WA 98195, U.S.A.

E-mail: mcquinn@uw.edu

Received September 1, 2020

Accepted May 9, 2021

Published June 14, 2021

Abstract. We investigate the amount of primordial information that can be reconstructed from spectroscopic galaxy surveys, as well as what sets the noise in reconstruction at low wavenumbers, by studying a simplified universe in which galaxies are the Zeldovich displaced Lagrangian peaks in the linear density field. For some of this study, we further take an intuitive linearized limit in which reconstruction is a convex problem but where the solution is also a solution to the full nonlinear problem, a limit that bounds the effectiveness of reconstruction. The linearized reconstruction results in similar cross correlation coefficients with the linear input field as our full nonlinear algorithm. The linearized reconstruction also produces similar cross correlation coefficients to those of reconstruction attempts on cosmological N-body simulations, which suggests that existing reconstruction algorithms are extracting most of the accessible information. Our approach helps explain why reconstruction algorithms accurately reproduce the initial conditions up to some characteristic wavenumber, at which point there is a quick transition to almost no correlation. This transition is set by the number of constraints on reconstruction (the number of galaxies in the survey) and not by where shot noise surpasses the clustering signal, as is traditionally thought. We further show that on linear scales a mode can be reconstructed with precision well below the shot noise expectation if the galaxy Lagrangian displacements can be sufficiently constrained. We provide idealized examples of nonlinear reconstruction where shot noise can be outperformed.

Keywords: redshift surveys, cosmological parameters from LSS, galaxy surveys, baryon acoustic oscillations

ArXiv ePrint: [2008.12312](https://arxiv.org/abs/2008.12312)

Contents

1	Introduction	1
2	Characteristic wavenumbers that shape reconstruction	3
3	Reconstruction in a toy universe	5
3.1	Conditioning to select a solution and solving (the linearized toy problem)	7
3.2	A numerical realization of a simplified universe	9
3.3	Low memory optimization algorithms	10
4	Reconstruction in the linearized model	11
4.1	Comparison of linear reconstruction with the fully nonlinear reconstruction of other studies.	14
4.2	(In)sensitivity to regularization parameter σ and sub-shot errors	15
4.2.1	Motivation for singular value sizes	16
4.2.2	Small wavenumber limit of linear reconstruction	18
4.3	Altering the cosmology	19
4.4	Number density and halo mass dependences	21
5	Reconstruction in the nonlinear model	22
6	Conclusions	26

1 Introduction

As we push within sight of mining the remaining cosmological information from the cosmic microwave background, there has been an increasing theoretical emphasis on understanding the late-time growth of cosmic structure as probed by galaxy redshift surveys e.g. [1–13]. While we appear to have sufficient perturbative control to wavenumbers of $k \approx 0.15 \, h \, \text{Mpc}^{-1}$ at $z \sim 0.5$ to obtain unbiased cosmological constraints (e.g., [8, 14, 15], although see [16]), it is unsettled whether galaxy surveys can extract useful constraints to higher wavenumbers. Furthermore, at low wavenumbers it is unclear whether existing methods achieve the minimum possible error.

The traditional approach to deriving cosmological bounds from large-scale structure measurements compares lower order statistics (such as the power spectrum and, perhaps, the bispectrum) measured from observational datasets with perturbation theory e.g. [3, 14, 17], large cosmological simulations e.g. [18–20], or quicker methods that take hybrid approaches e.g. [21–24]. One disadvantage of limiting the analysis to lower order statistics is that information becomes progressively entangled in higher order correlations with increasing wavenumber. This entanglement has led to efforts that attempt to reconstruct the linear density field by in essence running the equations in reverse [1, 4, 25]. Such reconstruction methods have been applied to baryon acoustic oscillation (BAO) measurements in order to sharpen these features [26, 27], and they are anticipated to markedly improve the constraining power of future BAO surveys [28].

Inspired by this success at reconstructing the BAO, a more recent focus has been on understanding the limits for how much primordial information can be extracted (rather than

focus solely on sharpening the BAO). One class of studies use perturbative bias expansions [9, 13, 29–31] and, an even more ambitious class attempts to compare fully nonlinear theories to data [2, 5, 6, 10–12]. The simplest (albeit computationally intractable) formulation of the latter class would run simulations with all possible random initial condition fields, find the best match to the observations and, then, take the power spectrum of that simulation’s density field to constrain the cosmology. In practice, various optimizations have been devised such as assuming the displacement is a potential field and that the tracers were initially homogeneous [8, 32], improving MCMC algorithms so that the expensive posterior evaluations are more likely to be accepted [6, 32], using machine learning algorithms to extrapolate from a more-easily-simulated coarse density grid to something more akin to a halo field [11], evolving forward the density field with less expensive techniques than full simulations [6, 12, 29] and using methods that find some local posterior maximum and assuming/arguing that the biases from this being a local maximum are correctable [10, 12]. The efficacy of reconstruction tends to be similar regardless of methodology: at low redshifts, reconstruction of the initial conditions is effective to wavenumbers of $k \approx 0.5 \text{ Mpc}^{-1}$ given a dense enough galaxy survey, potentially extending the wavenumber reach over perturbative methods by a factor of a few [6, 8, 12, 32].

We will henceforth refer to these fully nonlinear efforts as “reconstruction,” even though the name was first coined in the context of spectroscopic galaxy surveys for reversing nonlinear evolution specifically in the BAO [28]. We further specialize to the class of reconstruction algorithms that start with a halo/galaxy field rather than the exploratory studies that reconstructed from a full 3D grid of the nonlinear density field. There is a significant difference between these two, as the former version of reconstruction is a highly under-constrained problem — there are many more modes that shape the galaxy field than there are observed galaxies. This aspect shapes the characteristics of the reconstructions presented here. (The best algorithms applied to the full density field are able to reconstruct wavenumbers that are $2 - 3\times$ higher to $k \sim 1 \text{ Mpc}^{-1}$ [33]. The limit of full density field algorithms is likely set instead by where information is erased by nonlinear evolution, such as shell crossings.)

Nonlinear galaxy reconstruction algorithms generically find that they are able to reconstruct the large-scale modes in a manner that appears to be roughly limited by shot noise at low wavenumbers. Above some wavenumber, their efficacy falls off a cliff, with studies finding that over a factor of ~ 2 in wavenumber the reconstructed field goes from highly correlated to uncorrelated with the input field [11, 34]. We aim to understand the principles that shape this seemingly generic behavior. Another unresolved issue is whether shot noise sets the floor for how well low wavenumber modes in the galaxy field can be reconstructed. This issue is of high import for detecting the large-scale signatures of primordial non-gaussianity [35] and neutrino mass [36]. Studies have shown that galaxy surveys, when weighting by halo mass, can have effective noises that are substantially smaller than the naive number-weighting shot noise estimate, but with a character that is still shot noise-like [14, 37, 38]. However, there is no understanding beyond that derived from brute-force numerics of the degree to which shot noise can be avoided. We present results that suggest that at low wavenumbers it may be possible to evade shot noise by a larger factor than has yet been achieved.

Density field reconstruction addresses perhaps the deepest conceptual issue in large-scale structure — the ultimate limit for reconstructing primordial information from a galaxy survey. Because the methods for reconstructing the density field are so complex and computational expensive, they have not afforded a conceptual understanding of what sets this limit, even in a simplified setting that ignores the additional complexities of redshift space

distortions and baryonic physics. We also ignore these complicating factors here. We attempt to make traction by understanding reconstruction in a toy universe in which galaxies of halo mass M reside at the peaks of a Gaussian random linear density field, with the condition that these peaks exceed the collapse threshold for a spherical system when the field is smoothed on the Lagrangian mass scale M . These peaks are then displaced with linear order Lagrangian perturbation theory (the Zeldovich approximation). This setup is motivated by the successes of (1) excursion set theory in explaining the halo mass function e.g. [39, 40] and (2) of Lagrangian perturbation theory [4, 41, 42]. We highlight the visualizations in [4], which show that the Zeldovich approximation fares excellently at describing particle displacements, erring primarily on virialized scales.

This paper is set up as follows. Section 2 discusses the scales involved in reconstruction, highlighting a curious coincidence that further motivates this work. Section 3 presents the toy problem that we aim to solve, as well as a linear simplification. Section 4 studies in detail the linear model, a setup that has bearing on understanding the limits of reconstruction. Lastly, section 5 considers our full nonlinear model.

Throughout we adopt the discrete Fourier convention in a cubic volume V that is common in cosmology with inverse transform given by $F(\mathbf{x}) = V^{-1} \sum_{\mathbf{k}} \tilde{F}(\mathbf{k}) \exp[-i\mathbf{k}\mathbf{x}]$, where the $\mathbf{k}$ cover a 3D grid with spacing $2\pi/V^{1/3}$, the sum runs over both positive and negative values of each component of the wavevector, and the tilde denotes the Fourier dual. All of our calculations consider $z = 0.5$ and, unless otherwise specified, are in the concordance Λ CDM cosmology with $\Omega_m = 0.3$, $\Omega_\Lambda = 0.7$, $n_s = 0.96$, $\sigma_8 = 0.8$, and $\Omega_b = 0.045$.

2 Characteristic wavenumbers that shape reconstruction

We are aware of several characteristic wavenumbers that may affect our ability to extract the linear modes from spectroscopic galaxy observations. Curiously, in our universe they all fall within a factor of a few of one another at low redshifts and for motivated galaxy number densities. The first is the nonlinear wavenumber k_{NL} , although it enters the least frequently in this paper's discussion. We define it as the solution to

$$\Delta_L(k)^2 \equiv \frac{k^3}{2\pi^2} D(z)^2 P_{L,0}(k) = 1. \quad (2.1)$$

Here $P_{L,0}$ is the $z = 0$ power spectrum of the linear matter overdensity, and $D(z)$ is the linear growth factor normalized to a present value of unity (such that the linear power at any redshift is $P_L = D(z)^2 P_{L,0}$). Solving for this scale in the concordance Λ CDM cosmology yields $k_{\text{NL}} = 0.2 \text{ Mpc}^{-1}$ at $z = 0.2$ and $k_{\text{NL}} = 0.5 \text{ Mpc}^{-1}$ at $z = 1.0$. The nonlinear wavenumber crudely bounds where perturbation theory is applicable.¹ Around k_{NL} , we can approximate the dimensionless linear matter overdensity power spectrum as

$$\Delta_L(k)^2 = \left(\frac{k}{k_{\text{NL}}} \right)^{3-n_{\text{eff}}}, \quad (2.2)$$

where n_{eff} is an effective power-law index. For our concordance Λ CDM cosmological model, $n_{\text{eff}} \approx 2$ at $k \sim 0.1 - 1 \text{ Mpc}^{-1}$, wavenumbers where reconstruction can be effective.

¹Formally, the density variance — which defines our nonlinear scale — is not the only parameter that shapes nonlinearity in large-scale structure in the perturbative limit. The linear displacement contributed by large and small scales relative to the wavenumber in question are additional ordering parameters e.g. [44]. The latter nonlinear parameter does enter into our discussion, setting the characteristic wavenumber k_{disp} (eq. (2.7)).

In addition to the nonlinear wavenumber, there is the wavenumber where the shot noise power is equal to the linear clustering power:

$$b_L^2 D(z)^2 P_{L,0}(k_S) = \bar{n}^{-1} \xrightarrow{\text{power law}} k_S = \left(2\pi^2 b_L^2 \bar{n} k_{\text{NL}}^{n_{\text{eff}}-3}\right)^{1/n_{\text{eff}}}, \quad (2.3)$$

$$\xrightarrow{\Lambda\text{CDM}} k_S \approx 0.4 \text{ Mpc}^{-1} \left(b_L^2 D^2 \bar{n}_{-3}\right)^{0.5}, \quad (2.4)$$

where $\bar{n}_{-3} = \bar{n}/[10^{-3}\text{Mpc}^{-3}]$, $\bar{n}$ is the galaxy number density, and b_L is the galaxy linear bias.

Additionally, there is another characteristic wavenumber that is not mentioned in discussions of reconstruction but that we argue is relevant — the “constraints wavenumber”, k_C . We define k_C as the wavenumber where the number of modes with smaller wavenumbers is equal to the number of constraints, which in the large-volume limit is solved by:

$$\mathcal{N}\bar{n}V = \frac{V k_C^3}{2\pi^2} \longrightarrow k_C = 0.4 \left(\frac{\mathcal{N}}{4}\right)^{1/3} n_{-3}^{1/3}. \quad (2.5)$$

We will motivate $\mathcal{N} \approx 3$ constraints per galaxy, owing to the three positions for each galaxy, and sometimes $\mathcal{N} \approx 4$ when the galaxies’ halo masses can be precisely estimated. In a picture where reconstruction is able to constrain large-scale modes before small-scale ones, k_C should bound the wavenumbers that can be reconstructed.

Another characteristic wavenumber is set by the Lagrangian size of halos, which approximates the maximum wavenumber that influences halo formation. We define this wavenumber to be

$$k_M = 3R_H^{-1} = 1.7 \text{ Mpc}^{-1} \left(\frac{M}{10^{12}M_\odot}\right)^{-1/3}, \quad (2.6)$$

where M is the halo mass and $R_H \equiv [3M/(4\pi\rho_m)]^{1/3}$. At k_M , the Fourier dual of a real-space tophat window function with unit support is equal to $\approx 1/3$.

Lastly, there is the scale where the displacement from wavenumbers greater than k_{disp} is larger than the halo Lagrangian radius, R_H . We estimate this wavenumber by using that the aligned pairwise variance of the Zeldovich Approximation displacement that is contributed by modes with wavenumbers greater than k ,

$$\sigma_\Psi(k)^2 = \int_k^\infty \frac{dk'}{6\pi^2} P_L(k'), \quad (2.7)$$

and then solving $\sigma_\Psi(k_{\text{disp}}) = R_H$ for k_{disp} . We will show that to the extent that this wavenumber is much smaller than the previous characteristic wavenumbers, the modes that determine the position of a halo may be reconstructed well enough to make the nonlinear reconstruction problem approximately convex (§ 5).

It is notable that in the concordance cosmology all of these characteristic wavenumbers are within a factor of several of one another. Figure 1 shows these wavenumbers as a function of the minimum halo mass surveyed, assuming 100% completeness so that all halos above the specified mass are included in the survey weighted by number density. The left panel is for $z = 0.5$, and the right for $z = 2$. Intriguingly, the constraints scale k_C and the shot noise scale k_S are nearly the same at all halo masses. This similarity is a coincidence of our cosmology, as these characteristic wavenumbers would scale differently with M in cosmologies with different




Figure 1. Visualization of the characteristic wavenumbers in the fiducial Λ CDM cosmology at $z = 0.5$ (left) and $z = 2$ (right). Four characteristic wavenumbers are shown as a function of halo mass: (1) the maximum wavenumber that contributes to the overdensity at the halo scale, $k_M = 3R_H^{-1}$, where R_H is the Lagrangian radius of a halo (black long dashed curves), (2) the wavenumber where shot noise power is equal to the linear clustering power for the specified redshifts, k_S (green dot-dashed curves), (3) the wavenumber for which the number of modes with smaller values equals $\mathcal{N} = 4$ times the number of galaxies, k_C (blue solid curves), and (4) the wavenumber above which the displacement is the Lagrangian size of a halo, k_{disp} (red dotted curves). Cumulative halo number densities, $n(>M)$, are computed using the mass function of [43] and assume 100% completeness above the specified halo mass. A wavenumber that is not shown is k_{NL} , which at $z = 0.5 - 2$ falls in the ballpark of these other wavenumbers with $k_{\text{NL}} \sim 0.2 - 0.5 \text{ Mpc}^{-1}$. Large-scale structure theory always considers k_S and never k_C , but, interestingly, k_S and k_C are very similar in the concordance cosmology at all relevant redshifts and halo masses, perhaps explaining why k_C is never considered. We argue in § 5 that nonlinear reconstruction becomes a convex problem to the extent k_{disp} is smaller than the other characteristic wavenumbers.

effective scalar spectral indices (n_{eff}). That $k_C(M) \approx k_S(M)$ in the concordance cosmology may explain why k_C never is referenced in large-scale structure literature, even though we will show it sets the maximum wavenumber that is accurately reconstructed. That k_M is a factor of ≈ 5 larger than k_C across all halo masses indicates that reconstruction is always under-constrained, with $\sim 5^3$ more modes that shape the halo field than constraints. This under-constrained property of reconstruction shapes many of the results in this study.

3 Reconstruction in a toy universe

Our aim is to reconstruct the initial conditions from 3D galaxy position measurements. This paper focuses on a toy model for the cosmos that is more tractable for understanding the limits of reconstruction, while still maintaining significant similarities to our universe. The model places galaxies with halo mass M at peaks in the smoothed-on-scale- M linear overdensity field that meet some threshold overdensity, $\delta_c(M)$.² We then displace these peaks with linear order Lagrangian perturbation theory (the Zeldovich approximation). This setup is motivated by the successes of (1) excursion set theory in explaining the halo mass function e.g. [39, 40] and (2) Lagrangian perturbation theory [4, 41, 42].

²Our approach could be straightforwardly generalized to more realistic models that consider eigenvalues of the local tidal field e.g. [45, 46].

Let us assume a survey with cubic volume V with a list of galaxies (i.e. overdensity peaks) with positions $\mathbf{x}_j$ and masses M_j for $j \in [0, N)$. The condition for peak heights reduces reconstruction to N constraint equations:

$$V^{-1} \sum_{\forall \mathbf{k}} \tilde{\delta}_{\mathbf{k}} e^{-i\mathbf{k} \cdot [\mathbf{x}_j - \boldsymbol{\psi}(\mathbf{q}_j | \tilde{\delta}_{\mathbf{k}})]} W_{M_j}(\mathbf{k}) = \delta_c(M_j) \quad \text{for } j \in [0, N), \quad (3.1)$$

where $\mathbf{q}_j$ is the Lagrangian position of the j^{th} halo and $W_M(k)$ is a window function that approximates the Lagrangian size of the halo (which hosts the ‘observed’ galaxy), here taken to be a tophat in real space, and $\boldsymbol{\psi}$ is the Zeldovich approximation displacement vector that is given by³

$$\boldsymbol{\psi}(\mathbf{q}_j | \tilde{\delta}_{\mathbf{k}}) = -iV^{-1} \sum_{\forall \mathbf{k}} \frac{\mathbf{k}}{k^2} \tilde{\delta}_{\mathbf{k}} e^{-i\mathbf{k} \cdot \mathbf{q}_j}. \quad (3.2)$$

We want to use the constraints given in eq. (3.1) (plus possibly additional constraints discussed shortly) to reconstruct the *linear theory* modes $\tilde{\delta}_{\mathbf{k}}$. As emphasized in § 2, if we take the maximum wavenumber that contributes to the formation of halos for masses typical of modern spectroscopic galaxy surveys (the wavenumber “cutoff” in W_M), this is a highly under-constrained problem with (infinitely) many solutions. We aim to select a solution that is as close as possible to the input density field.

One difficulty with the above setup is that $\boldsymbol{\psi}$ is a function of the Lagrangian coordinate of each halo $\mathbf{q}_j$, which is not an observable. One can avoid this difficulty by substituting in the Taylor expansion $\boldsymbol{\psi}(\mathbf{q}) = \boldsymbol{\psi}(\mathbf{x}) - (\boldsymbol{\psi}(\mathbf{x}) \cdot \nabla) \boldsymbol{\psi}(\mathbf{x}) + \dots$ or by some other optimization technique. However, we instead make the simplifying assumption that the displacement evaluated at the Lagrangian position of a galaxy can be calculated from $\boldsymbol{\psi}$ and the final position of the galaxy $\mathbf{x}_j$. We expect that the uncertainty in this mapping will not be what sets the efficacy of reconstruction.⁴

With this simplification, it is useful to instead formulate our equations in Lagrangian space. The master set of equations we are considering becomes

$$V^{-1} \sum_{\forall \mathbf{k}} \tilde{\delta}_{\mathbf{k}} e^{-i\mathbf{k} \cdot [\mathbf{q}_j - \Delta \boldsymbol{\psi}(\mathbf{q}_j | \tilde{\delta}_{\mathbf{k}})]} W_{M_j}(\mathbf{k}) = \delta_c(M_j) \quad \text{for } j \in [0, N), \quad (3.3)$$

where $\Delta \boldsymbol{\psi} \equiv \boldsymbol{\psi} - \boldsymbol{\psi}_{\text{TRUTH}}$. Here $\boldsymbol{\psi}_{\text{TRUTH}}$ is the *true* displacement computed from the input $\tilde{\delta}_{\mathbf{k}}$ — i.e. the field we aim to reconstruct. One might worry that now we have written the problem in terms of some unobservable quantities $(\mathbf{q}_j, \boldsymbol{\psi}_{\text{TRUTH}})$, rather than our observables $(\mathbf{x}_j, M_j)$. However, this is a sleight of hand as $\mathbf{q}_j = \mathbf{x}_j - \boldsymbol{\psi}_{\text{TRUTH}}(\mathbf{q}_j)$, noting also the simplification in the previous paragraph.

We can identify *linear* equations that, if satisfied, also solve eq. (3.3), namely for $j \in [0, N)$:

$$\delta_c(M_j) - V^{-1} \sum_{\forall \mathbf{k}} \tilde{\delta}_{\mathbf{k}} e^{-i\mathbf{k} \cdot \mathbf{q}_j} W_{M_j}(\mathbf{k}) = 0; \quad (3.4)$$

$$\boldsymbol{\psi}(\mathbf{q}_j | \tilde{\delta}_{\mathbf{k}}) - \boldsymbol{\psi}_{\text{TRUTH}}(\mathbf{q}_j) = 0. \quad (3.5)$$

³It might be more realistic to also include the halo window function $W_M(k)$ in the integral in eq. (3.2) for $\boldsymbol{\psi}$, as modes with wavelengths smaller than the halo scale do not contribute significantly to its displacement. However, we find that none of our results are appreciably changed if we include this factor. One could also add to the displacement vector an effective term $\propto k \tilde{\delta}_{\mathbf{k}}$ owing to small-scale dynamics [3, 47], but we suspect it also will have limited effect.

⁴Our results appear insensitive to ~ 2 Mpc errors in the displacement that would occur from the laziest (zeroth order) approximation $\boldsymbol{\psi}(\mathbf{q}_j) \approx \boldsymbol{\psi}(\mathbf{x}_j)$.

These are the conditions that the density smoothed on the scale W_{M_j} adds up to the collapse threshold $\delta_c(M_j)$ at the Lagrangian position $\mathbf{q}_j$ and that the displacement is the true displacement. They embody more information than an observer is able to access as they require knowledge of the $\mathbf{q}_j$ and hence $\psi_{\text{TRUTH}}(\mathbf{q}_j)$. Solutions to these conditions are not the only solutions to our nonlinear equations (eq. (3.3)), but solutions to these linear equations clearly would be preferred ones. Despite this linear system of equations also being extremely under-constrained (as these $4N$ equations should be far fewer than the total number of modes), these equations' linearity in $\tilde{\delta}_{\mathbf{k}}$ means that one can always find a solution, whereas gradient descent-like methods for solving the nonlinear equations may easily get stuck in local minima (§ 5).

“Reconstruction” from the linear eq.s 3.4 and 3.5 will be more successful than efforts starting with the nonlinear master equation (eq. (3.1)) as these equations require $\psi_{\text{TRUTH}}(\mathbf{q}_j)$, which a lowly cosmologist doing reconstruction would have no way of knowing exactly. Therefore, “reconstruction” using these linear constraint equations bounds the effectiveness of reconstruction. We later obtain solutions to our nonlinear equations that come close to saturating the bounds placed by solving the linear equations. Furthermore, we show that other reconstruction algorithms that have been applied to the density field in N-body simulations [11, 34] come close to saturating these bounds.

One might worry that our idealized setup does not capture all the information used by nonlinear reconstruction algorithms, which are trying to model the full halo distribution. For example, the setup outlined so far does not account for halos forming at density peaks in Lagrangian space:

$$-iV^{-1} \sum_{\forall \mathbf{k}} \mathbf{k} \tilde{\delta}_{\mathbf{k}} e^{-i\mathbf{k} \cdot [\mathbf{q}_j - \Delta\psi(\mathbf{q}_j | \tilde{\delta}_{\mathbf{k}'})]} W_{M_j}(k) = 0 \quad \text{for } j \in [0, N). \quad (3.6)$$

Additionally, there may be other constraints on the shape of the surrounding Lagrangian overdensity, like

$$-V^{-1} \sum_{\forall \mathbf{k}} k^2 \tilde{\delta}_{\mathbf{k}} e^{-i\mathbf{k} \cdot [\mathbf{q}_j - \Delta\psi(\mathbf{q}_j | \tilde{\delta}_{\mathbf{k}'})]} W_{M_j}(k) = -\ell_j^{-2} \delta_c(M_j) \quad \text{for } j \in [0, N), \quad (3.7)$$

such that, for example, if the characteristic size is $\ell_j \gg R_j \equiv (3M/4\pi)^{1/3}$ the halo is likely quickly growing in mass. Because the $\tilde{\delta}_{\mathbf{k}}$ in these additional constraint equations are weighted by powers of k , they are less important for constraining the low wavenumber modes that can be reconstructed accurately. Indeed, we demonstrate in § 3.2 that adding the three peak constraints per galaxy, eq. (3.6), results in a modest improvement in the reconstructed field. Lastly, if our galaxy survey is complete for halos above a given mass, then there should be constraints enforcing that there are not other peaks on this mass scale. However, in practice the solutions we find tend not to include many additional peaks, and so this condition is more or less satisfied naturally.

3.1 Conditioning to select a solution and solving (the linearized toy problem)

Our under-constrained system of equations can be cast as a least squares optimization problem. *Specializing first to the linear set of conditions*, which require godlike knowledge of $\psi_{\text{TRUTH}}(\mathbf{q}_j)$, we search for the global minimum of some or all of the terms in the following

loss function:

$$\begin{aligned}
L = \sum_{j=0}^N & \left[\underbrace{\left(\delta_c(M_j) - V^{-1} \sum_{\forall \mathbf{k}} \tilde{\delta}_{\mathbf{k}} e^{-i\mathbf{k} \cdot \mathbf{q}_j} W_{M_j}(k) \right)^2}_{\text{Lagrangian overdensity (L)}} + \underbrace{\ell_D^{-2} \left(\psi(\mathbf{q}_j | \tilde{\delta}_{\mathbf{k}}) - \psi_{\text{TRUTH}}(\mathbf{q}_j) \right)^2}_{\text{displacements (D)}} \right. \\
& \left. + \underbrace{\ell_P^2 \left(V^{-1} \sum_{\forall \mathbf{k}} \mathbf{k} \tilde{\delta}_{\mathbf{k}} e^{-i\mathbf{k} \cdot \mathbf{q}_j} W_{M_j}(k) \right)^2}_{\text{peaks (P)}} \right] + \underbrace{C(\tilde{\delta}_{\mathbf{k}})}_{\text{regularization}}. \quad (3.8)
\end{aligned}$$

We later consider the nonlinear case in a very similar setup. This problem has infinitely many minima with $L = 0$ if $C = 0$, since the number of modes that shape the density field is much greater than the number of constraint equations ($7N$ if we include all three conditions in eq. (3.8)). The regularizer C has to supplement with sufficiently many conditions to constrain the system. Furthermore, C has to be the square of an expression linear in $\tilde{\delta}_{\mathbf{k}}$ to retain the desirable property of L being positive and solvable with linear algebra (this choice is called Tikhonov regularization; [48]). Fortunately, the Gaussian mode-amplitude prior that reconstruction algorithms employ e.g. [10] fall exactly in this category, generalizing in our noiseless case to the Ridge Regression regularization condition of

$$C(\tilde{\delta}_{\mathbf{k}}) = \sigma^2 \sum_{\forall \mathbf{k}} \frac{|\delta_{\mathbf{k}}|^2}{P_L(k)V}, \quad (3.9)$$

which is a special case of Tikhonov regularization. While previous reconstruction algorithms are phrased in terms of maximizing a posterior rather than a ‘noiseless’ loss function, the loss function is analogous to the logarithm of the posterior when divided by $2\sigma^2$. The reconstruction “noise” in our case (and that we argue holds for posterior reconstruction algorithms) is not shaped significantly by the “error” parameter σ , but rather by how modes project onto the set of several times N well-constrained quantities. Modes that are not constrained are set to zero by this regularization. Indeed, we show later that our solution to linear equations depends negligibly on choice of σ once $\sigma^2 \lesssim 0.1$, which can be recast in terms of the error on e.g. the displacements for modes to be constrained at a cosmologically interesting level. Finally, the regularization does assume a cosmology to calculate the linear power spectrum, $P_L(k)$ (and we use here the $P_L(k)$ of our background cosmology), but our results are relatively insensitive to this assumption for reasons that will be discussed.

One issue with using the Gaussian regularization (eq. (3.9)) is that when a mode is not well constrained, the solution that minimizes L is $\tilde{\delta}_{\mathbf{k}} = 0$. Thus, the global minimum will be a biased solution that favors lower power in less constrained pixels. This is not a surprise and is analogous to the down-weighting of noisy pixels in optimal map making (Weiner filtering). We have also investigated a Tikhonov regularization condition that preferences a particular random phase field with power spectrum given by the input (and with the same $[VP_L(k)]^{-1}$ weighting as above). This alternative setup retains the property of linearity, but conditions to a non-Gaussian field. We find that the wavenumbers that can be reconstructed are not significantly changed in this alternative regularization scheme.

In addition to the regularization condition, we need to choose some physically motivated values for the scales that appear in L , namely ℓ_D and ℓ_P , as these will weight the equations in different manners. We take $\ell_D = \ell_P = 10$ Mpc for our fiducial values as motivated

in footnote 5.⁵ We have decrease/increased ℓ_D and ℓ_P by a factor of 5 and find that the cross correlation coefficient of our solution with the truth is negligibly altered. This lack of dependence we think owes to the under-constrained nature of problem: there are solutions that satisfy all our constraints perfectly in the limit that the regularizer normalization $\sigma \rightarrow 1$ (§ 4.2). When we turn to the nonlinear problem in § 5, ℓ_D and ℓ_P no longer are relevant.

3.2 A numerical realization of a simplified universe

We now design a realization of the universe that embodies our toy picture of a critical overdensity threshold in Lagrangian space plus its Zeldovich approximation displacement. We further desire the spectrum of halo masses to match the halo mass function found in N-body simulations. To achieve the latter, we implement a model that is motivated by extended Press-Schechter theory [39], except that it also attempts to capture halo discreteness. This algorithm is most analogous to the ‘peak patch’ scheme for generating a galaxy field of [46, 49].

To create the linear matter overdensity field, we generate a Gaussian random field with power spectrum $P_L(k)$ in a cube of volume V with N_s^3 discrete samples on a grid. We then smooth the overdensity field on different scales with a tophat in real space specified by its enclosed Lagrangian mass M , starting with the largest mass scale and moving to smaller and smaller masses. For the field filtered at each smoothing scale δ_M , we identify all the peaks that satisfy $\delta_M(z) > \delta_c$, where $\delta_c = 1.7$ as motivated by calculations in spherical collapse [50]. Starting with the highest peak, we mark off a spherical Lagrangian region of mass M ; all cells that fall within this region can no longer be used as the center of a halo. We further associate the j^{th} peak with a halo specified by its Lagrangian position q_j and mass M_j . The density field is then smoothed over a smaller scale, and the process is repeated, heeding the prior exclusions. The algorithm results in a list of ‘halo’ positions, masses, and peak overdensities. To generate the real space position of a halo $\mathbf{x}_j$, we then displace the peak with its Zeldovich approximation displacement, $\psi_{\text{TRUTH}}(\mathbf{q}_j)$, where we remind the reader that ‘TRUTH’ indicates the displacement is calculated with the linear overdensity field generated by this algorithm rather than the reconstructed linear field.

Figure 2 shows the mass function that results from this algorithm. The thick black curve is the fit to the mass function in cosmological N-body simulations of Jenkins et al. [43]. For the highest resolution calculations shown, our model does well at latching onto the Jenkins mass function, with a small overshoot at the largest masses. At the ‘fiducial’ resolution

⁵We can weakly motivate values by expanding eq. (3.10) in $\Delta\psi$ such that

$$V^{-1} \sum_{\mathbf{k}} \tilde{\delta}_{\mathbf{k}} e^{-i\mathbf{k} \cdot \mathbf{q}_j} W_{M_j}(k) \left[1 - i\mathbf{k} \cdot \Delta\psi - \frac{1}{2}(\mathbf{k} \cdot \Delta\psi)^2 + \dots \right] = \delta_c(M_j) \quad \text{for } j \in [0, N), \quad (3.10)$$

In the square brackets, the different terms bear semblance to the Lagrangian overdensity, the peak condition, and the peak shape condition. Defining $\hat{\delta}_c(\mathbf{q}, M_j) \equiv V^{-1} \sum_{\mathbf{k}} \tilde{\delta}_{\mathbf{k}} e^{-i\mathbf{k} \cdot \mathbf{q}} W_{M_j}(k)$ and ℓ_j describes the j^{th} peak’s radius of curvature (eq. (3.7)), then

$$\left(\hat{\delta}_c(\mathbf{q}_j, M_j) - \delta_c(M_j) \right) + \Delta\psi \cdot \nabla_{\mathbf{q}} \hat{\delta}_c(\mathbf{q}, M_j) \Big|_{\mathbf{q}_j} - \frac{\delta_c(M_j)}{6\ell_j^2} \Delta\psi^2 + \dots = 0 \quad \text{for } j \in [0, N), \quad (3.11)$$

assuming that the peak curvature and displacement are uncorrelated to deduce the last term. eq. (3.11) is *somewhat* analogous to our loss function (L) but where the terms related to the Lagrangian density, peaks and displacements have physically motivated coefficients rather than the parameters ℓ_D and ℓ_P . First, we anticipate that the curvature radius at peak, ℓ_j , is likely somewhat larger than the Lagrangian halo size or $R_H \sim 3$ Mpc for halos in the mass range we consider. For simplicity, and because of the factor of 6 in the denominator, we set $\ell_D = 10$ Mpc. We seldom consider the peaks constraint, nor does equation (3.11) motivate a clear choice for ℓ_P .




Figure 2. Halo mass functions in the excursion set-inspired model used to generate this study’s halo fields (§ 3.2). The colored curves show the mass function at different resolutions. The blue curve is computed using the fiducial box size and grid resolution (200 Mpc, 255^3 cells). These halo mass functions match reasonably the fit to the mass function in cosmological N-body simulations of Jenkins et al. [43] (black solid). At the fiducial resolution, the model is well converged for halos with $M > 10^{12} M_\odot$ that are our primary consideration, and it undershoots by a factor of two at the smallest halos we consider with $M = 10^{11} M_\odot$. The green curve is computed for 200 Mpc and 127^3 grid cells, which are the specifications used in § 5 for the more computationally expensive nonlinear model (but where we only consider halos with $M > 5 \times 10^{12} M_\odot$).

that we adopt for most of our calculations ($N_s = 255$, $V^{1/3} = 200$ Mpc), the mass function undershoots at $M = 10^{11} M_\odot$ by almost a factor of two, but is in reasonable agreement at higher masses.

Lastly, figure 3 provides a visualization of the algorithm’s output in a 20 Mpc deep projection through a 100 Mpc periodic box. The image colorscale saturates at linear overdensities of $[-3, 3]$. These calculations were performed for 127^3 elements, matching the fiducial resolution of our 200 Mpc box. The left panel shows the algorithm’s Lagrangian space position of halos, juxtaposed on top of the linear overdensity field. Halo sizes are encoded by the radius of the dot r_{dot} , with $r_{\text{dot}} \propto M^{1/3}$, with the left panel featuring halos with $M > 10^{11} M_\odot$. Halo exclusion around the most massive halos is evident. The middle panel shows the displaced position of halos with $M > 10^{12} M_\odot$ plus the displacement vector, again juxtaposed on top of the same linear field. Halos whose starting position is within the slice are shown. The righthand panel shows the results of one of our reconstruction algorithms (§ 4).

3.3 Low memory optimization algorithms

Our fiducial calculations attempt to constrain $255^3 - 1$ independent parameters from the positions of as many as $\sim 10^5$ galaxies. If we wrote our linear optimization problem as one of matrix algebra, we would need to store a matrix of size $\sim 255^3 \times 10^5$. Fortunately, there are fast algorithms to solve such systems that do not require holding this large matrix in memory. In particular, for the linear and nonlinear calculations presented here, we use the L-BFGS iterative algorithm, which employs a Newton’s method-like algorithm to converge towards a minimum [48], using the implementation of <http://www.chokkan.org/software/liblbfgs/>. This method requires a loss function and its gradient, the latter of which can be computed analytically for our simplified model. We terminate the iteration when the loss function has changed by less than 1% in the prior ten iteration steps, which we find is sufficient for




Figure 3. *The left panel* shows the Lagrangian space position of halos in our algorithm, juxtaposed on top of the linear overdensity field, for a 20 Mpc deep projection through a 100 Mpc box. The image of the linear overdensity saturates at values of $[-3,3]$. These calculations were performed for 127^3 elements in a 100 Mpc box, essentially matching the fiducial resolution of our 200 Mpc box. Halo sizes are encoded by the radius of the dot, which scales as $M^{1/3}$ above the minimum halo mass shown of $10^{11}M_{\odot}$. *The middle panel* shows the displaced position of halos with $M > 10^{12}M_{\odot}$ plus the displacement vector, again juxtaposed on top of the same linear field. Halos whose starting position is within the slice are shown (rather than selecting halos by their displaced position), and displacement vectors are not shown for halos that traverse a box boundary. *The right panel* shows our reconstruction from our linearized algorithm with all $7N$ constraint equations (the Lagrangian overdensity [L], displacements [D] and peak conditions [P]), as well as the same halo+displacement field as in the middle panel.

convergence. (We did find that the less conservative criteria of terminating with 10% rather than 1% sometimes did not reach convergence at low wavenumbers for the fully nonlinear reconstruction problem.) We have also checked that we obtain the same solution for the linear problem with the LSQR algorithm, which is an efficient conjugate gradient solver [51].

There are a few technical aspects of note. First, rather than treating the complex mode amplitude, we perform our optimizations treating the real and imaginary components of each mode as separate parameters. Second, we were careful to eliminate all redundancy in $\tilde{\delta}_{\mathbf{k}}$ that owes to it being the Fourier transform of a real field. The reason for the non-standard grid number of 255 in our computations rather than the more-standard $2^8 = 256$ is that odd numbered grids have less redundancy in real to complex Fast Fourier transform algorithms like FFTW.⁶ The speed that we evaluate Fast Fourier transforms is not a limiting factor for obtaining a solution. Last, we have verified that the non-conditioned loss function is zero when evaluated at the correct solution and that our analytic gradient is computed correctly by taking numerical gradients of the loss function.

4 Reconstruction in the linearized model

Linearity means that we can always find a global minimum of the loss function (L), at least once sufficiently conditioned. This section presents solutions to these linear equations for different permutations of the constraint conditions (i.e. of L, D and P in eq. (3.8) for L), different normalizations of the regularizer, and different galaxy survey specifications. These

⁶<http://www.fftw.org>




Figure 4. The results of our *linear* reconstruction (eq. (3.8)) using the Lagrangian position of all $M > 10^{12} M_{\odot}$ halos (L; left panel), this plus their displacements (L+D; middle panel), and this plus the condition that they are density peaks (L+D+P; right panel), for the same linear density field as shown in figure 3. Each slice is 20 Mpc deep projection through a 100 Mpc box, and the image saturates at overdensities of $[-3, 3]$. Halo masses and displacements are encapsulated in the same manner as in the rightmost two panels in figure 3.

equations assume that the positions, the masses of halos, and the Lagrangian displacements $\Psi_{\text{TRUTH},j}$ are known perfectly. Of course, the latter two can only be imperfectly inferred, but this exercise provides a useful bound on reconstruction. We also compare the resulting reconstructions with those of published nonlinear reconstruction algorithms in § 4.1 and show that, despite the substantial simplification to reach the linear equations, their cross correlation coefficients with the input linear overdensity field are strikingly similar to those reported in reconstruction studies. We further show that the linear solutions perform similarly to our best solutions to the nonlinear problem in § 5.

Figure 4 shows images of the reconstructions for different permutations of our linear constraints, conditioning with the Ridge Regression regularizer that is analogous to the Gaussian mode-amplitude prior in other reconstruction studies (eq. (3.9)). (Images of the linear density field that is being reconstructed are shown back in figure 3.) The normalization of the regularizer is set to our fiducial value of $\sigma = 0.01$, but we will show that our results are insensitive to σ as long as $\sigma \lesssim 0.3$. The lefthand panel in figure 5 shows the power spectra of the reconstructed input fields from these algorithms (colored curves). These should be compared with the power spectrum of the true input field (thick black solid curve). The different colored curves employ different permutations of linear constraints for the reconstruction: the Lagrangian overdensity condition (L), this condition plus the displacements condition (L+D), both of these conditions plus the peaks condition (L+D+P), and just the displacements condition (D). Generically, the algorithms reproduce the power spectrum at low k and undershoot at high k . The undershoot at high- k arises because the regularization drives modes that are not significantly constrained to zero, which are generally the high- k ones. Only in one case shown does the reconstructed power not converge to the input at low wavenumbers, when reconstruction only uses the Lagrangian overdensity condition (L). In this case, the regularizer preferences smaller mode amplitudes at higher wavenumbers, such that the power at lower wavenumbers must be increased to have sufficient variance on the halo scale to reach the threshold for collapse, $\delta_c(M)$. Including the constraints on the displace-




Figure 5. The power spectra (left panel) and cross correlation coefficients with the input linear overdensity (right panel), featuring ‘reconstructions’ that use different combinations of the linear constraint conditions. Curves labeled with ‘L’ use the condition on the Lagrangian-space overdensity threshold (eq. (3.4)), ‘D’ use the Zeldovich displacement condition (eq. (3.5)), and ‘P’ use the peaks condition (eq. (3.6)). These curves are for a reconstruction with $\bar{n} = 10^{-3} \text{Mpc}^{-3}$, which is achieved by randomly sampling halos with mass above $10^{12} M_{\odot}$, which amounts to 45% of all such halos, and they use a Ridge Regression regularization condition with $\sigma = 0.01$, although the results are essentially independent of σ for $\sigma \lesssim 0.3$. The horizontal black dashed line in the left panel is the shot noise power calculated as $\bar{n}^{-1}$, and the red vertical line in the right panel is the estimate from counting displacement constraints, $k_C(\mathcal{N} = 3)$.

ments (L+D) significantly lessens this bias and, indeed, the displacement condition shapes the solution at low wavenumbers (as displacements are the property of the halo field that is the most sensitive in the infrared). At higher wavenumbers, the Lagrangian overdensity condition shapes the reconstruction, as can be noted by the reconstructed power spectrum using only displacements (D), for which the power dives to zero at lower wavenumbers compared to when the Lagrangian overdensity condition is included.

Now turn to the righthand panel in figure 5, which shows the cross correlation coefficient,

$$r \equiv \frac{P_{XY}(k)}{\sqrt{P_{XX}(k)P_{YY}(k)}}, \quad (4.1)$$

where X and Y are the two fields being correlated. A cross correlation coefficient of unity indicates that the phase of the fields’ modes is aligned. The black thick dashed curve is the cross correlation coefficient between the true halo field and the number density-weighted nonlinear galaxy field (and r is similar in this particular case if we weight by mass). The other curves in this panel show the cross correlation coefficient between the linear density field and the various reconstructed fields. First, note that all of these reconstructions, with the exception of the Lagrangian overdensity-only (L), find a maximum wavenumber that can be reconstructed with $r \sim 0.5$ that is a factor of 2 – 3 larger than where this threshold is met when the galaxy field is simply correlated with the input linear field.⁷ This factor of 2 – 3 improvement is similar to that found in previous reconstruction studies. However, the Lagrangian overdensity constraint alone is insufficient to reconstruct even the large-scale

⁷One could argue that comparing the naive cross correlation is not a fair and one should at least compare a full perturbative bias expansion: however, while adding to the accuracy, the perturbative bias expansion does little to extend the range where $r \approx 1$ [14]. The gross profile of $r(k)$ between the input overdensity field and nonlinear halo field is set by Poisson noise.




Figure 6. Comparison of our toy reconstruction with the results of prior studies at $z \approx 0.5$. The solid curves are the cross correlation coefficients of the input field with the reconstructed presented in Modi et al. [11] and in Yu et al. [34]. The dashed curves are the results of the displacement-only linear reconstruction (D) in which we attempt to match the specifications of these studies ($M_{\min} = 5 \times 10^{12} M_{\odot}$, $\bar{n} = 3 \times 10^{-4} \text{Mpc}^{-3}$ for Modi et al. 2018 and $M_{\min} = 10^{11} M_{\odot}$, $\bar{n} = 10^{-3} \text{Mpc}^{-3}$ and $M_{\min} = 3 \times 10^{12} M_{\odot}$, $\bar{n} = 10^{-4} \text{Mpc}^{-3}$ for the two samples in Yu et al. 2017). The dotted-curve is the L+D linear reconstruction, which includes the Lagrangian-space overdensity constraint in addition to displacements.

modes well, falling below r computed from the nonlinear halo field (thick black dashed).⁸ Including the displacements immediately makes $r \rightarrow 1$ on large-scales (L+D curve). Including the peaks constraint (L+D+P) tugs r slightly upward at high wavenumbers, whereas only including the displacement condition (D) results in r transitioning more quickly from one to zero.

4.1 Comparison of linear reconstruction with the fully nonlinear reconstruction of other studies.

The cross correlation coefficients that our linear reconstruction approximation returns are similar to those found in fully nonlinear reconstruction algorithms. The solid curves in figure 6 compare the r of our linear reconstructions at $z = 0.5$ with those reported for the reconstructions in Yu et al. 2017, [34] and Modi et al. 2018, [11], which used full cosmological N-body simulations as the input for the nonlinear galaxy field. The dashed curves show our displacement-only reconstruction (D), attempting to match the halo mass specifications and number densities of the compared studies. We randomly discard halos above the halo mass threshold to match their number densities.

Specifically, we compare with the reconstructions in Yu et al. (2017) for their two highest galaxy number densities, which we approximate with the specifications $M_{\min} = 10^{11} M_{\odot}$, $\bar{n} = 10^{-3} \text{Mpc}^{-3}$ and $M_{\min} = 3 \times 10^{12} M_{\odot}$, $\bar{n} = 10^{-4} \text{Mpc}^{-3}$. They used an ‘isobaric’ reconstruction algorithm that remaps the positions of galaxies to a homogeneous initial field assuming the displacement is curl free. For both number densities, their r is closest to our

⁸This comparison is not however a fair one as the Lagrangian overdensity-only (L) reconstruction uses the Lagrangian-space halo field, for which its cross correlation coefficient with the input linear field is well below this reconstructed r on all scales.

displacement-only solution (given by the corresponding dashed curves) as expected. Yu et al. (2017) considered three redshifts and found little dependence with redshift. A lack of redshift dependence trivially results for our linear reconstruction, as only the halo mass function has redshift dependence in the linear equations and the halo mass function is of little importance for the displacement constraints used for this comparison. Shell crossing, which erases information in a redshift-dependent manner [12] and is not captured in our linear model, is one effect that could impart a redshift dependence. The lack of redshift dependence in Yu et al. suggests shell crossing may not be a principal limitation.

Next, let us consider the reconstruction in Modi et al. (2018) for their fiducial specifications of $M_{\min} = 5 \times 10^{12} M_{\odot}$ and $\bar{n} = 3 \times 10^{-4} \text{Mpc}^{-3}$. Modi et al. (2018) used a neural network to extrapolate a low resolution N-body simulation gridded in 4.5 Mpc cells onto a finer grid. We find that our displacement-only solution agrees well with their result, whereas adding the Lagrangian overdensity constraint results in our r having a tail to higher wavenumbers that outperforms the r in Modi et al. (compare the black solid curve with the black dotted curve). As their 4.5 Mpc cells are two times larger than the Lagrangian radius at their minimum halo mass, this lower resolution may make the reconstruction of Modi et al. (2018) less sensitive to the constraints owing to halo peak shapes and, hence, their results better approximated by our displacement-only reconstruction. Indeed, some calculations that follow suggest that this should be the case.

Both the Yu et al. (2017) and Modi et al. (2018) nonlinear reconstructions do not reproduce the small error found by our linear reconstruction at the lowest wavenumbers. The low- k reconstruction error of Yu et al. (Modi et al.) is more or less consistent with being limited by shot noise for number-density (mass-density) weighting. Our linear reconstructions are able to produce errors well below the shot noise floor at the lowest wavenumbers, a feature we discuss in § 4.2.

4.2 (In)sensitivity to regularization parameter σ and sub-shot errors

All of our reconstruction calculations presented so far adopt $\sigma = 0.01$ for the regularization parameter (cf. eq. (3.8)). Larger values of σ will increasingly tilt the convex loss function in our linearized problem. The more poorly a mode is constrained by the galaxy field constraint equations, the more sensitive it should be to the choice of σ . However, figure 7 shows that values of $\sigma \sim 1$ are required to have a substantial effect on the cross correlation coefficient. This independence suggests that the modes that are constrained by the galaxy constraints are typically constrained at the level $\text{var}[\tilde{\delta}_{\mathbf{k}}] \sim V P_L(k)$ or better.

This σ independence may seem puzzling, as there are wavenumbers where modes appear to be poorly constrained with $r(k) \ll 1$, yet they are insensitive to even the highest σ shown. Indeed, some modes that are reconstructed quite poorly with $r \sim 0.1$ are the least sensitive to σ ! What reconciles this apparent contradiction is that $\approx 4N$ parameters are well constrained in our reconstructions using the displacement and Lagrangian overdensity constraints (L+D), but these parameters are not necessarily the Fourier modes. Unconstrained parameters are projected to zero by the regression. The projection of the Fourier modes onto the constrained eigen-basis shapes the value of the cross correlation coefficient.

To understand this mathematically, we can recast our least squares problem in terms of matrix algebra, writing our loss function as

$$L = \sum_C \|\mathbf{A}_C \tilde{\chi} - \mathbf{b}_C\|_2^2 + \sigma^2 \|\tilde{\chi}\|_2^2, \quad (4.2)$$




Figure 7. Linear reconstruction for different regularizer normalizations: the power spectra (left panel), cross correlation coefficients between the input linear overdensity and the reconstructed field r (middle panel), and the $1-r^2$ (right panel). All reconstructed fields use the linear reconstruction approximation with N Lagrangian overdensity and $3N$ displacement constraints (L+D). These curves are computed for the case $\bar{n} = 10^{-3} \text{Mpc}^{-3}$ by randomly sampling halos with mass above $10^{12} M_{\odot}$, which requires sampling 45% of all halos in our model. The dependence of the solutions on our Ridge Regression regularization parameter σ only becomes substantial for $\sigma \gtrsim 0.3$. In the noiseless case ($\sigma \rightarrow 0^+$), the reconstruction noise is set by how the low wavenumber modes project onto the set of well constrained eigenvectors, and our estimate for this projection is given by the ‘mixing noise’ (the dot-dashed curve in the rightmost panel). The mixing noise, which approximates the performance of our linear algorithm, falls substantially below the shot noise expectation, $1-r^2 = (\hat{P}_g \bar{n})^{-1}$ where $\hat{P}_g$ is the galaxy power spectrum in the simulation, shown with the dashed curve in the rightmost panel.

where the subscript C indexes the constraint conditions (e.g. the Lagrangian overdensity [L], the j component of the displacement [Dj], etc), $\tilde{\chi} \equiv \tilde{\delta}_{\mathbf{k}} / \sqrt{V P_L(k)}$, and $\|\dots\|_2$ denotes the Euclidean norm. The second term on the right hand side of eq. (4.2) is our Ridge Regression regularization condition.

For a *single* constraint condition, which will either be $C = \text{L}$ or $C = \text{Dj}$, the solution to eq. (4.2) is $\hat{\chi}_C = (\mathbf{A}_C^\dagger \mathbf{A}_C + \sigma^2 \mathbf{I})^{-1} \mathbf{A}_C^\dagger \mathbf{b}_C$, where the dagger represents a conjugate transpose. However, if we use singular value decomposition to write $\mathbf{A}_C = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^\dagger$, where $\mathbf{\Sigma}$ is an $N \times N$ diagonal matrix and $\mathbf{U}$ and $\mathbf{V}$ are unitary matrices, then the solution to eq. (4.2) with one constraint condition for the modes is

$$\hat{\chi} = \mathbf{V} \mathbf{\Sigma}^\sigma \mathbf{U}^\dagger \mathbf{b} \quad \text{where} \quad \Sigma_{ij}^\sigma \equiv \frac{\Sigma_{ii}}{\Sigma_{ii}^2 + \sigma^2} \delta_{ij}^K, \quad (4.3)$$

dropping our C subscripts. Thus, Fourier modes that do not project onto the N well-constrained right singular vectors are set to zero. For modes that do project onto the constrained space, since we find $\hat{\chi}$ does not depend significantly on σ for $\sigma < 0.3$, eq. (4.3) indicates that the singular values, Σ_{ii} , are $\gtrsim 0.3$.

In the remainder of this section, we use this linear algebra setup to motivate why $\Sigma_{ii} \gtrsim 0.3$ and to derive a formula for the effective noise at low wavenumbers in our linear reconstruction approximation. Readers interested primarily in the gross properties of the reconstruction can skip ahead to § 4.3.

4.2.1 Motivation for singular value sizes

We motivate mathematically why the Σ_{ii} have such values by first writing out the $\mathbf{A}_C$, which we do here for our Lagrangian overdensity and the j^{th} component of the displacement

conditions:

$$[A_L]_{lm} = \sqrt{\frac{P_L(k_m)}{V}} e^{-i\mathbf{k}_m \cdot \mathbf{q}_l}, \quad (4.4)$$

$$[A_{D,j}]_{lm} = -i\ell_D^{-1} \sqrt{\frac{P_L(k_m)}{V}} \frac{k_{m,j}}{k_m^2} e^{-i\mathbf{k}_m \cdot \mathbf{q}_l}, \quad (4.5)$$

where we have ignored factors of $W_{M_l}(k)$ as this simplification aids subsequent expressions (and halo exclusion results in a similar effect), and the columns include all wavevectors. The factors of $P_L(k)$ in the above matrices arise from the definition of $\mathbf{A}_C$, chosen in order for our mode parameters to be normalized as $\tilde{\chi} \equiv \delta_k / \sqrt{V P_L(k)}$ such that our regularizer simplifies to the form in eq. (4.2). To understand the size of the singular values Σ_{ii} , we first calculate

$$\mathbf{A}_L \mathbf{A}_L^\dagger = \begin{pmatrix} \xi_L(R_{H,0}) & \xi_L(\mathbf{q}_0 - \mathbf{q}_1) & \xi_L(\mathbf{q}_0 - \mathbf{q}_2) & \cdots & \xi_L(\mathbf{q}_0 - \mathbf{q}_{N-1}) \\ \xi_L(\mathbf{q}_1 - \mathbf{q}_0) & \xi_L(R_{H,1}) & \xi_L(\mathbf{q}_1 - \mathbf{q}_2) & \cdots & \xi_L(\mathbf{q}_1 - \mathbf{q}_{N-1}) \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ \xi_L(\mathbf{q}_{N-1} - \mathbf{q}_0) & \xi_L(\mathbf{q}_{N-1} - \mathbf{q}_1) & \xi_L(\mathbf{q}_{N-1} - \mathbf{q}_2) & \cdots & \xi_L(R_{H,N-1}) \end{pmatrix}, \quad (4.6)$$

where $\xi_L(r) \equiv \langle \delta(\mathbf{x}) \delta(\mathbf{x} + \mathbf{r}) \rangle = V^{-1} \sum_{\mathbf{k}} P_L(k) e^{-i\mathbf{k} \cdot \mathbf{r}}$ is the linear correlation function, and we have approximated the effect of the window functions that we dropped previously by evaluating the diagonal at $R_{H,j}$, indicating the Lagrangian radius of the j^{th} halo. The eigenvalues of $\mathbf{A}_L \mathbf{A}_L^\dagger$ are equal to the Σ_{ii}^2 , the square of the singular values of $\mathbf{A}_L$. The form of $\mathbf{A}_L \mathbf{A}_L^\dagger$ given in eq. (4.6) suggests a well conditioned matrix with $\Sigma_{ii}^2 \sim \xi_L(R_{H,0}) \sim 1$: two rows in $\mathbf{A}_L \mathbf{A}_L^\dagger$ will be most similar when halos are separated by the minimum Lagrangian-space separation our algorithm allows, $\max[R_{H,i}, R_{H,j}]$, but we still expect ξ_L to be appreciably different at this minimum separation scale such that the eigenvalue from the eigenvector that principally arises from the subtraction of the two most-similar rows are still likely to be greater than a few tenths. For modes shaped by the Lagrangian overdensity constraint, this explains why $\sigma \sim 1$ demarcates where the normalization of the regularization condition starts to substantially affect the solution.

Now, if we repeat and compute $\mathbf{A}_{D,j} \mathbf{A}_{D,j}^\dagger$, the same expression as eq. (4.6) holds but with $\xi_L \rightarrow \ell_D^{-2} \langle \psi_j(\mathbf{x}) \psi_j(\mathbf{x} + \mathbf{r}) \rangle = \ell_D^{-2} V^{-1} \sum_{\mathbf{k}} k_j^2 k^{-4} P_L(k) e^{-i\mathbf{k} \cdot \mathbf{r}}$, i.e the correlation function of displacements along the j direction measured in units of $\ell_D = 10$ Mpc. We note that the average displacement is of the order of 10 Mpc such that again the diagonal values of $\mathbf{A}_{D,j} \mathbf{A}_{D,j}^\dagger$ are approximately unity. The matrix $\mathbf{A}_{D,j} \mathbf{A}_{D,j}^\dagger$ is not as well conditioned as $\mathbf{A}_L \mathbf{A}_L^\dagger$ because displacements are more correlated between galaxy positions. The dependence of our reconstructions on σ shown in figure 7 suggests that the smallest singular values of $\Sigma_{ii} \approx 0.1$, but with many values around $\Sigma_{ii} \sim 1$ and with the lowest wavenumbers projecting onto the singular vectors that have the largest Σ_{ii} .

Thus, reconstruction measures nearly as many numbers as constraints once $\sigma^2 \lesssim 0.1$. This statement can be related to errors on, for example, the displacements — when they are constrained to a few Mpc (as our singular values in this case are in units of 10 Mpc), $3N$ numbers are measured. Indeed, figure 8 examines the reconstructed field that results when adding an uncorrelated Gaussian error to each halo's displacement, an action which has a similar effect on r as varying σ . However, the XN well-constrained right-singular vectors are not necessarily the Fourier modes we desire to measure. Indeed, the right singular vectors for the Lagrangian overdensity problem do not appear to project well onto Fourier modes




Figure 8. Linear reconstruction when adding noise to Ψ : the same as figure 7 except, rather than varying the regularization normalization σ , we randomly add a Gaussian displacement error to the $\Psi_{\text{TRUTH}}(\mathbf{q}_j)$ in eq. (3.5) before solving the linear equations. Adding noise to the displacement has a similar effect on r as varying σ for reasons described in § 4.2.

as evidenced by the r values being substantially off unity at all wavenumbers (see the blue solid curve in the righthand panel of figure 5). When displacements are included, the well constrained subspace appears to be better represented by the lowest wavenumber Fourier modes. Interestingly, the projection is such that the lowest wavenumbers are substantially better constrained than the shot noise expectation (see the right panel in figure 7). We aim to understand this projection better in what follows.

4.2.2 Small wavenumber limit of linear reconstruction

For a better handle on this projection, rather than considering $\mathbf{A}_C \mathbf{A}_C^\dagger$ as above, we instead compute $\mathbf{A}_C^\dagger \mathbf{A}_C$. The eigenvectors of this matrix are the same as the right singular vectors of $\mathbf{A}_C$. Since the displacements shape the reconstruction at low wavenumbers, let us consider the j^{th} component of the displacements, where

$$\left[\mathbf{A}_{Dj}^\dagger \mathbf{A}_{Dj} \right]_{sm} = \ell_D^{-2} \sqrt{\frac{P_L(k_s) P_L(k_m)}{V^2}} \frac{k_{s,j} k_{m,j}}{k_s^2 k_m^2} \left[N \delta_{sm}^K + \left(\bar{n} b_{Lq} \tilde{\delta}(\mathbf{k}_s - \mathbf{k}_m) + G(\sqrt{N}) \right) (1 - \delta_{sm}^K) \right]. \quad (4.7)$$

We have used that $\sum_{j=0}^N e^{-i(\mathbf{k}_s - \mathbf{k}_m) \cdot \mathbf{q}_j} = \bar{n} b_{Lq} \tilde{\delta}(\mathbf{k}_s - \mathbf{k}_m) + G(\sqrt{N})$ for $\mathbf{k}_s \neq \mathbf{k}_m$, where b_{Lq} is the galaxy Lagrangian bias and $G(\sqrt{N})$ is a ‘random’ c-number with standard deviation $\sqrt{N}$ that owes to shot noise (i.e. you can think of the Riemann sum as a Monte-Carlo integral of a continuous overdensity field and $G(\sqrt{N})$ is the error).

We can estimate how much of a rotation of the basis is required to diagonalize $\mathbf{A}_{Dj}^\dagger \mathbf{A}_{Dj}$, where the vectors that describe these rotations are the eigenvectors. If we only consider high wavenumbers where typically $|G(\sqrt{N})| > |\bar{n} b_{Lq} \tilde{\delta}|$ such that shot noise dominates, in the plane indexed by $(\mathbf{k}_s, \mathbf{k}_m)$ a rotation by an angle of

$$\theta_{\mathbf{k}_s \mathbf{k}_m} \approx \sqrt{\frac{P_L(k_m)}{P_L(k_s)}} \frac{k_s}{k_m} \frac{G(\sqrt{N})}{N}, \quad (4.8)$$

is required to diagonalize from the Fourier basis. Let us assume that any rotation that mixes a mode $\mathbf{k}_s$ with modes with $k_m > k_C$ results in that component of the mode being lost, motivated by our numerical displacement-only reconstructions in which modes with $k \lesssim k_C$ are well constrained whereas modes with $k \gtrsim k_C$ are poorly constrained (e.g. figure 5). Under

this assumption and further assuming that shot noise dominates these modes (justified in concordance cosmology because $k_S \approx k_C$; § 2), we can use eq. (4.8) to estimate the cross correlation coefficient

$$r(k_s) = \prod_{\forall |\mathbf{k}_m| > k_C} \cos[\theta_{\mathbf{k}_s \mathbf{k}_m}] = \left(1 - \frac{1}{2} \sum_{\forall |\mathbf{k}_m| > k_C} \theta_{\mathbf{k}_s \mathbf{k}_m}^2 + \dots \right), \quad (4.9)$$

or

$$1 - r(k_s)^2 \approx \sum_{\forall |\mathbf{k}_m| > k_C} \theta_{\mathbf{k}_s \mathbf{k}_m}^2 \approx \frac{k_s^2}{\bar{n} P_L(k_s)} \int_{k_C}^{\infty} \frac{d^3 k}{(2\pi)^3} \frac{P_L(k)}{k^2}, \quad (4.10)$$

using that modes are pixelated with $(\Delta k)^3 = (2\pi)^3/V$. If we roughly expect $1 - r^2$ from the three displacement conditions to be reduced by a factor of three relative to the single component estimate given by eq. (4.10), which results in the low-wavenumber estimate of

$$1 - r(k)^2 \approx 2000 \text{Mpc}^5 \times \frac{k^2}{P_L(k)} \frac{\sigma_\Psi(k_C)^2}{2 \text{Mpc}^2} \left(\frac{10^{-3} \text{Mpc}^{-3}}{\bar{n}} \right) \quad [\text{Mixing Noise Limit}], \quad (4.11)$$

where $\sigma_\Psi(k_C)^2$ is defined by eq. (2.7), and $\sigma_\Psi(k_C)^2 = 2 \text{Mpc}^2$ for $k_C = 0.5 \text{Mpc}^{-1}$ in ΛCDM . The rightmost panel in figure 7 shows that this expression for the reconstruction noise, which we term ‘mixing noise’, captures the r we find at low k in the low σ limit in which the solution is not affected by the normalization of the regularizer. Additionally, we will show in § 4.3 that the mixing noise limit formula applies even in cosmologies where P_L is very different.

Finally, while the limiting sensitivity given by eq. (4.11) becomes independent of the regularization condition normalization, σ , once $\sigma \lesssim 0.3$, our regularization condition does shape our results as it determines the weighting via $\tilde{\chi}_j = \tilde{\delta}_{\mathbf{k}_j} / \sqrt{V P_L}$. This weighting is responsible for the factors of P_L in eq. (4.10). If we more generally use $\tilde{\chi}_j = \tilde{\delta}_{\mathbf{k}_j} / \sqrt{V \mathcal{W}(k)}$ then eq. (4.10) would instead become

$$1 - r(k_i)^2 \approx \bar{n}^{-1} \frac{k_i^2}{\mathcal{W}(k_i)} \int_{k_C}^{\infty} \frac{d^3 k}{(2\pi)^3} \frac{\mathcal{W}(k_j)}{k_j^2}, \quad (4.12)$$

with the caveat that our approximations become less good as $\mathcal{W}(k)$ is made to be more strongly increasing with k than $\sqrt{V P_L}$. While $\mathcal{W}(k) = P_L(k)$ may be the most natural weighting, eq. (4.12) suggests that one could design $\mathcal{W}(k)$ to result in even much smaller error bars by increasing its tilt towards the infrared. The reason why there is no free lunch is that the matrix $\mathbf{A}_{\text{D},j} \mathbf{A}_{\text{D},j}^\dagger$ becomes more poorly conditioned the more red-sensitive that $\mathcal{W}(k)$ becomes, amplifying modeling errors.

4.3 Altering the cosmology

Figure 9 shows how the linear reconstruction fares for different spectral tilts of the primordial matter overdensity power spectrum, changing the effective spectral index by ± 1 , but in a manner that the density variance in a tophat sphere with mass $10^{11} M_\odot$ is the same (such that the $n(> M)$ and b_L are approximately matched for galactic halos for these three cosmologies). We compare these cosmologies with $\bar{n} = 10^{-3} \text{Mpc}^{-3}$ by randomly sampling halos with mass above $10^{12} M_\odot$, which requires sampling 100% of halos in the reddest tilt, half in the fiducial, and a quarter in the bluest tilt cosmology.




Figure 9. Reconstruction varying cosmology (n_s) in the linear approximation: the power spectra (left panel), the cross correlation coefficients between the input linear overdensity and reconstructed fields r (middle panel), and $1 - r^2$ (right panel), varying the effective tilt of the input power spectrum by ± 1 from the fiducial $n_s \approx 1$ cosmology in a manner such that the overdensity variance is still the same in a tophat window with mass of $10^{11} M_\odot$ (which roughly matches $n(> M)$ and b_L for the three cosmologies). All reconstructed fields use our linear reconstruction approximation with the Lagrangian overdensity plus displacement conditions (L+D; dashed curves) or just displacements (D; dotted curves). In addition to the reconstruction results, the left panel shows the input power spectra (solid curves) and the galaxy power spectra (dot-dashed curves), the middle panel shows the constraints scale (k_C) at which $3N$ modes have smaller wavenumbers, and the right panel shows the mixing noise (eq. (4.11)) for the three cosmologies. These curves are computed for the case $\bar{n} = 10^{-3} \text{Mpc}^{-3}$ by randomly sampling halos with mass above $10^{12} M_\odot$ (which requires sampling 100% of halos in the reddest tilt, half of them in the fiducial, and a quarter of them in the bluest).

The left panel shows the power spectra of the input linear overdensity (solid curves), of the reconstruction in the linear approximation (dashed curves), and of the galaxy field (dot-dashed curve) for the three different cosmologies. The fractional contribution of shot noise is much different in these exotic ‘tilted’ cosmologies, with the one with the bluest tilt boasting a galaxy power spectrum that is significantly shaped by shot noise at all simulated scales. All cosmologies however have the same constraints scale k_C , and $r = 0.5$ is achieved at nearly the same wavenumber in all three linearized reconstruction calculations (middle panel; the dashed curves show L+D and the dotted curves show D). Even the approximate profile of the cross correlation coefficient is similar between the different calculations. Clearly shot noise is not what sets its profile. Rather, the quick transition from $r \approx 1$ to $r \approx 0$ is primarily set by mode counting, where the wavenumber that has $3N$ shorter modes, k_C , is shown as the vertical line in the middle panel.

The rightmost panel in figure 9 now shows $1 - r^2$, where the dashed and dotted curves are respectively the L+D and D linearized reconstructions. Reconstruction in the linear approximation fares far better than the shot noise ‘floor’, with a significantly steeper wavenumber scaling in the three cosmologies than if the noise were set by shot noise. This is particularly evident in the bluest tilted cosmology, where shot noise would result in an error that is nearly wavenumber independent. The blue thin dot-dashed curve show our ‘mixing noise’ estimate in this shot noise-dominated cosmology (eq. (4.11)), which remarkably reproduces the salient behavior in all three cosmologies.⁹

⁹We cap our integral to calculate σ_ψ^2 at k_M for $10^{12} M_\odot$ or otherwise the curve would shift up by almost a factor of two for the $n_s = 2$ case.




Figure 10. Reconstruction varying number density in the linear approximation: the power spectra (left panel) and cross correlation coefficients with the input (right panel) for three halo number densities in our linear reconstruction approximation. Number densities are selected by randomly sampling from the halos that have masses greater than $10^{12}M_{\odot}$. The colored solid curves are linear models that include both the Lagrangian overdensity and displacement constraints (L+D), and the dotted curves with the corresponding color only use the displacements (D). The horizontal lines in the left panel are the $\bar{n}^{-1}$, and the vertical lines in the right panel show the constraints wavenumbers, k_C .

4.4 Number density and halo mass dependences

Figure 10 shows how the reconstructed power and cross correlation coefficients vary with the galaxy number density, $\bar{n}$, using our linearized algorithm. Number densities are fixed by randomly sampling a fraction of halos that have masses greater than $10^{12}M_{\odot}$, maintaining the same spectrum of halo masses. A factor of ten increase in $\bar{n}$ results in a factor of $10^{1/3} \approx 2$ improvement in the wavenumber reach, and there is the analogous trend for a factor of ten smaller $\bar{n}$. The solid vertical lines in this figure are the constraints scale $k_C = (6\pi^2\bar{n})^{1/3}$. The low k behavior of $1 - r^2$ scales inversely with number density and is well explained by our mixing noise estimate (eq. (4.11)).

Figure 11 fixes the number density to $3 \times 10^{-4}\text{Mpc}^{-3}$ and varies the minimum halo mass, considering minimum masses of $M = 10^{11}$, 10^{12} , and $10^{13}M_{\odot}$. Also shown is the square of the square of the real-space tophat window functions that enclose a mass, $W_M(k)^2$, for these three M (black dashed curves). Generally the wavenumbers of modes that are constrained are smaller than the halo scale where $W_M(k)^2$ transitions to zero. In these cases, the displacement-only reconstruction is similar to adding the Lagrangian overdensity constraints aside from a high- k tail of low r values. When the number densities are sufficiently high such that k_C becomes more comparable to the halo scale, the Lagrangian overdensity constraints become more effective at boosting the span with $r > 0.5$ (as well as resulting in an overshoot in the reconstructed power), as can be seen for the $10^{13}M_{\odot}$ case. We suspect there is a similar behavior when the reconstruction algorithm does not resolve the Lagrangian halo scale R_H — specifically that the resulting r mimics the displacement-only reconstruction. This suspicion could explain why the displacement-only reconstruction best matches the Modi et al. (2018) [11] reconstruction, as their algorithm does not resolve R_H .

We have also investigated adding a lognormal scatter to the mass associated with halo (which enters in the Lagrangian peak condition and not the displacement condition) as would arise from uncertainty in the inference of halo masses. We find that this scatter has the effect of reducing the benefits of adding the Lagrangian overdensity condition. This scatter




Figure 11. Reconstruction varying halo masses in the linear approximation: the power spectra (left panel) and cross correlation coefficients with the input (right panel) for three minimum halo masses and a fixed number density of $\bar{n} = 3 \times 10^{-4} \text{Mpc}^{-3}$ in our linear reconstruction approximation. This number density is fixed by randomly choosing a fraction of halos above the halo mass threshold. The solid curves are linear models that include both the Lagrangian overdensity and displacement constraints (L+D), and the dotted curve use only the displacements (D), with the color indicating the halo mass threshold. The thin black dashed curves in the right panel show $W_M(k)^2$, the square of the real-space tophat window functions, for $M = 10^{13}$, 10^{12} , and $10^{11} M_\odot$ (from left to right respectively).

has little effect at low wavenumbers, as the displacements drive the reconstruction there. Reconstruction algorithms that use a gridded halo field may be more sensitive to such halo mass uncertainties (see § 5).

5 Reconstruction in the nonlinear model

So far we have concentrated on understanding reconstruction in our convex, linear limit. This section finally investigates the solution to the nonlinear problem! The main result is that, with our best nonlinear algorithms, we find a cross correlation coefficient between the input field and the reconstructed field that has a similar profile to that found with our linear approximation. Our nonlinear solutions do make larger errors at low wavenumbers that are most notable when we consider $1 - r^2$, but improvements may be possible there.

The loss function for our full nonlinear toy problem is

$$L \equiv \sum_{j=0}^N \left(V^{-1} \sum_{\forall \mathbf{k}} \tilde{\delta}_{\mathbf{k}} e^{-i\mathbf{k} \cdot [\mathbf{x}_j - \psi(\mathbf{q}_j | \tilde{\delta}_{\mathbf{k}'})]} W_{M_j}(\mathbf{k}) - \delta_c(M_j) \right)^2 + \sigma^2 \sum_{\forall \mathbf{k}} \frac{|\delta_{\mathbf{k}}|^2}{P_L(\mathbf{k})V}. \quad (5.1)$$

The first term is just repeating our model (eq. (3.1)) and the second is the same regularization as before.¹⁰ We again highlight the simplification we use for our nonlinear problem where we evaluate the estimated displacement in the exponent at the true Lagrangian position $\mathbf{q}_j = \mathbf{x}_j - \psi_{\text{TRUTH}}$ rather than estimating $\mathbf{q}_j$ from the estimated displacement field ψ and the final position of the galaxy $\mathbf{x}_j$ via some optimization method, a primitive one being Taylor expansion. We suspect that doing the optimization so that our result has no dependence on the unobservable quantity, $\mathbf{q}_j$, would not affect the solution relative to this simplification

¹⁰We have investigated including the condition that halos form at peaks at their Lagrangian position, but find no improvement (and rather somewhat poorer performance) and so do not include this condition in our nonlinear model.

(see § 3). Beyond this simplification, the inputs for solving the nonlinear model are the ‘observables’, $\mathbf{x}_j$ and M_j for $j \in [0, N)$. As in previous sections, we apply this formalism to our toy realization of the universe described in § 3.2.

This nonlinear loss function allows us to comment on what $\tilde{\delta}_{\mathbf{k}}$ are likely to fall near the convex subspace that encompasses the true solution. The wavenumbers k that contribute to $\sum_{\mathbf{k}} \tilde{\delta}_{\mathbf{k}} e^{-i\mathbf{k} \cdot \mathbf{q}_j} W_{M_j}(k)$ are generally larger than those that contribute to $\psi(\mathbf{q}_j | \tilde{\delta}_{\mathbf{k}'})$ and smaller than $\sim R_H^{-1}$, where R_H is the Lagrangian size of the halo. Once $\psi(\mathbf{q}_j | \tilde{\delta}_{\mathbf{k}})$ is estimated with accuracy R_H , the exponential factor $e^{-i\mathbf{k} \cdot \mathbf{q}_j}$ is no longer very sensitive to the mode amplitudes (allowing one to expand the part of the argument that still depends on ψ) and, to the extent $\psi(\mathbf{q}_j | \tilde{\delta}_{\mathbf{k}'})$ arises from smaller wavenumber modes than $\sum_{\mathbf{k}} \tilde{\delta}_{\mathbf{k}} e^{-i\mathbf{k} \cdot \mathbf{q}_j} W_{M_j}(k)$, the problem becomes convex. In the relevant limit where lower wavenumbers are reconstructed best, the wavenumbers that need to be accurately reconstructed to estimate $\psi(\mathbf{q}_j | \tilde{\delta}_{\mathbf{k}})$ to within R_H are $k < k_{\text{disp}}$ (§ 2). Figure 1 shows that $k_{\text{disp}} \approx 0.3 \text{ Mpc}^{-1}$ for a survey consisting of all halos with $M > 5 \times 10^{12} M_\odot$. This k_{disp} is not much larger than the wavenumbers at which the galaxy field correlates well with the input field and, thus, a reconstruction algorithm can be initialized with these modes close to their correct values. This argument helps explain the seemingly convex behavior of reconstruction in the literature in which even gradient descent-like algorithms do not seem to get stuck in minima that are far from the input field.¹¹ However, if the halo overdensities owe to modes that are entirely distinct from those that contribute to the displacements, reconstruction could ‘draw’ a halo anywhere and, therefore, would not be able to constrain the halo displacements. Thus, the intermediate modes that contribute to both ψ and $\psi(\mathbf{q}_j | \tilde{\delta}_{\mathbf{k}'})$ must play a key role.

To illustrate these points, we now present reconstructions in our simplified nonlinear problem. Because the loss function is not likely to be convex far from the true solution, the reconstruction is likely to be more successful if it starts from a point that is as good a guess as possible for the evolved field. We consider three starting points:

zero: $\tilde{\delta}(\mathbf{k}) = 0$. For this distant starting point, the L-BFGS algorithm gets trapped relatively far away from the minimum found by the linear algorithm presented in the previous section.

realistic: The strategy for this case is to guess a field that is somewhat close to the input overdensity field using quantities that are observationally accessible. Starting with an overdensity field with

$$\tilde{\delta}^{(0)}(\mathbf{k}) = \tilde{\delta}_g^{\text{mw}}(\mathbf{k})/b_g \quad \text{for } k < 0.2 \text{ Mpc}^{-1}$$

and $\tilde{\delta}^{(0)}(\mathbf{k}) = 0$ otherwise. Here $\tilde{\delta}_g^{\text{mw}}$ is the mass-weighted galaxy field.¹² Next, we take the halo field and from the position of halos displace backwards to get the Lagrangian position, with the displacements calculated from $\tilde{\delta}^{(0)}(\mathbf{k})$, and place a compensated

¹¹Of course, for modes well described by linear theory, one expects the reconstruction problem should be trivially convex in the region considered for the $\tilde{\delta}_{\mathbf{k}}$; we are referring to the more nonlinear modes.

¹²The galaxy bias b_g is estimated as $\hat{b}_g = \mathcal{N}_k^{-1} \sum_{|\mathbf{k}| < k_{\text{max}}} \tilde{\delta}_g^{\text{mw}} \tilde{\delta}_{\text{TRUTH}}^* / (\tilde{\delta}_{\text{TRUTH}} \tilde{\delta}_{\text{TRUTH}}^*)$, where $\tilde{\delta}_{\text{TRUTH}}$ is the input density field. This bias estimate uses the $\mathcal{N}_k$ lowest modes in the simulation that satisfy $k < k_{\text{max}} = 4\pi/V^{1/3}$. Using the input to calculate the bias is not something an observation would have access to, but a realistic survey would have a much larger volume than our 200 Mpc box to measure the bias.

overdensity profile at $\mathbf{q}_j^{(0)} = \mathbf{x}_j - \psi(\mathbf{q}_j | \tilde{\delta}_{\mathbf{k}}^{(0)})$ given by

$$\mathcal{W}_{M_j}^c(\mathbf{q}, \mathbf{q}_j, \delta(\mathbf{q}_j)) = 2 \left[\delta_c(M_j) - \delta(\mathbf{q}_j^{(0)}) \right] \exp \left[-\frac{(\mathbf{q} - \mathbf{q}_j^{(0)})^2}{2R_H(M_j)} \right] \left(1 - \sqrt{\frac{\pi}{2}} \frac{|\mathbf{q} - \mathbf{q}_j^{(0)}|}{2R_H(M_j)} \right), \quad (5.2)$$

where $\delta_c(M_j)$ is the overdensity that defines our halo (eq. (3.1)) and $\delta(\mathbf{q})$ is the matter overdensity field. The spatial integral over $\mathcal{W}_{M_j}^c$ is zero,¹³ and this fact also means that these profiles do not alter the power at low enough k . Thus, the matter overdensity input field for the L-BFGS algorithm is given by $\delta^{(N)}$ where

$$\delta^{(n)}(\mathbf{q}) = \delta^{(0)}(\mathbf{q}) + \sum_{j=0}^n \mathcal{W}_{M_j}^c(\mathbf{q}, \mathbf{q}_j^{(0)}, \delta^{(j)}(\mathbf{q})), \quad (5.3)$$

and we organize the placement such that the M_j that are summed are in decreasing order. For $M = 5 \times 10^{12} M_\odot$ ($1 \times 10^{12} M_\odot$), the value of L for this input $\tilde{\delta}_{\mathbf{k}}$ is reduced by a factor of 3(2) over the **zero** initialization with $\tilde{\delta}_{\mathbf{k}} = 0$.

idealized: We use the identical scheme as (2) except with $\tilde{\delta}^{(0)} = \tilde{\delta}_{\text{TRUTH}}$ for $k < 0.2 \text{ Mpc}^{-1}$ and $\tilde{\delta}^{(0)}(\mathbf{k}) = 0$ otherwise (i.e. we start with the correct solution at low wavenumbers). This idealized starting point results in a similar initial value for L to the *realistic* algorithm.

All of the results presented in this section use $\sigma = 0.01$; we find a similar dependence on σ as in our linear reconstructions.

Figure 12 shows images of the reconstruction for the *realistic* and *idealized* algorithms (middle and right panels respectively), alongside the reconstruction for the linear algorithm considered in the previous section (left panel). The reconstructed field appears similar between the linear algorithm and the *idealized* one. The *realistic* algorithm captures many of the gross features of the two other reconstructions. However, there are noticeable failures in the *realistic* algorithm where halos' true point of origination (as indicated by the orange line) is not close to any density peak, suggesting that the displacement is being significantly misestimated. Furthermore, the *realistic* algorithm's large-scale underdense regions are not as well captured as in the linear and idealized algorithms.

These similarities and differences between the reconstructions are also reflected in the cross correlation coefficients with the input linear overdensity field, which are shown in the left panel in figure 13. The *zero* algorithm does not achieve cross correlation coefficients much greater than 0.8, and it even fails to reconstruct the density modes at all in the lowest wavenumber bin. (Still, the *zero* algorithm is able to move a long way from its trivial starting point.) The *realistic* algorithm reconstructs to nearly the same maximum wavenumber as the linear one, but undershoots the precision of the linear algorithm at lower wavenumbers. A similar undershoot relative to our linear solutions was seen in the r reported in others' nonlinear reconstructions [11, 34] and is attributed to shot noise. We also show the *realistic* algorithm, where we have reduced the displacements by a factor of two and four, which makes the problem more convex. The reconstruction for these cases is somewhat improved. The *idealized* algorithm, which unrealistically starts with the input field at $k < 0.2 \text{ Mpc}^{-1}$ but where the starting point at higher wavenumbers is far from the truth, nearly saturates the bounds set by our linear calculations.

¹³In practice, our $\mathcal{W}_{M_j}^c$ only goes out $3.5R_H(M_j)$ from $\mathbf{q}_j$ as this contains 90% of the support.




Figure 12. The results of reconstruction algorithms applied to $M > 10^{12} M_{\odot}$ halos in a 100 Mpc box (see figure 3 for the input field that is being reconstructed). The left panel is our linear solution that bounds the efficacy of nonlinear reconstruction, the middle is nonlinear reconstruction in our *realistic* scenario (which initializes the reconstruction with the halo field for $k < 0.2 \text{ Mpc}^{-1}$), and the right panel in our *idealized* scenario, for the same linear field as shown in figure 3. Each slice is a 20 Mpc deep projection, and the image saturates at overdensities of $[-3, 3]$. These calculations are for 127^3 elements, matching the fiducial resolution of our 200 Mpc box. Halo positions, masses and displacements are illustrated in the same manner as in figure 3.




Figure 13. Nonlinear reconstruction results: the cross correlation coefficients r (left panel) as well as the $1 - r^2$ (right panel) between the linear input overdensity and the nonlinear reconstruction for the three different initialization methods discussed in § 5. The halo field used for these reconstructions includes all halos with $M > 5 \times 10^{12} M_{\odot}$, resulting in $\bar{n} = 5 \times 10^{-4} \text{ Mpc}^{-3}$. The nonlinear reconstructions take 127^3 elements in a 200 Mpc box, lower resolution than our linear calculations owing to their computational cost. Also shown for comparison is the linear reconstruction approximation (using the Lagrangian overdensity and displacement constraints with $\sigma = 10^{-2}$), and our *realistic* algorithm but with the displacements reduced by two or four to make the problem more convex. Finally, the black dashed curve shows the cross correlation coefficient of the mass-weighted galaxy survey with the input field: improved performance relative to this curve indicates an algorithm is outperforming the traditional shot noise floor.

One of the central unanswered questions is whether the shot noise of the galaxy field limits the low-wavenumber performance of these algorithms or whether the limit achieved by the linear algorithm is achievable. The righthand panel in figure 13 features the low wavenumber behavior, showing $1 - r^2$. Our *idealized* algorithm is able to nearly match the performance of the linear algorithm. It is interesting that *idealized* reconstruction is not pulled significantly away from the very small errors that are achieved by the linear algorithm at low wavenumbers, especially since our starting point at $k \gtrsim 0.2 \text{Mpc}^{-1}$ is still fairly distant from the true solution. However, our *realistic* algorithm's solution is not able to achieve as small errors at low k , having similar noise to the mass-weighted halo field that this algorithm's lowest k modes were initialized with. Compare with the black dashed curve, which shows $1 - r^2$ between the input linear overdensity and the mass-weighted halo field. If we artificially reduce the displacements by a factor of two or four, making the equations we are solving more convex, r does decrease significantly below the r of the mass-weighted halo field that the algorithm was initialized with. (This is despite the level of shot noise being *larger* in these reduced displacement cases since the field is less clustered.) The standard conception of shot noise has little to do with the magnitude of displacements and so the reduced noise when we reduce the displacement normalization to us suggests shot noise is not a fundamental limit (and this exercise may have some physical relevance, as at high redshifts the displacements are indeed smaller).¹⁴ We suspect that algorithmic improvements for solving the nonlinear equations may allow one to outperform shot noise in a realistic setting, as there are clear deficiencies in the reconstructed field of the *realistic* algorithm.

6 Conclusions

We have considered galaxy reconstruction in a simplified model for structure formation in which halos are the displaced Lagrangian peaks in the density field, a model that results in a deterministic nonlinear relation between the input and evolved fields. This controlled setting allowed us to investigate the ultimate limits of reconstruction, the effects of the standard Gaussian prior on mode amplitudes, and why gradient descent-like reconstruction algorithms work at all (as it was not obvious to us why they do not get stuck far from the true solution). For much of this study, we considered an intuitive linearized limit in which reconstruction is a convex problem but where the answer is also a solution to our nonlinear problem — a limit that bounds the effectiveness of reconstruction. Key findings include:

- Existing nonlinear reconstruction algorithms are close to extracting all of the accessible information. This argument rests on the tenet that our model of displaced Lagrangian peaks captures the essential information. We can then ‘linearize’ this model to make the problem convex (resulting in linear equations specifying that the modes at halo positions sum to an overdensity of 1.7 and to the Zeldovich displacements) and find the global solutions that would clearly be the best solution an algorithm applied to our full nonlinear model could hope to find. Despite the linearized model using information beyond what any observer can access, specifically the true displacements at the positions of the galaxies, we showed that the linearized solutions produce similar cross correlation coefficients to those of nonlinear reconstruction algorithms applied to cosmological N-body simulations [11, 34, particularly when we restrict to just the constraints from displacements] as well as to the best solutions we obtain for our full nonlinear problem.

¹⁴We also find that if the less constrained singular vectors are dampened by choosing $\sigma \sim 1$, the improvement in r at low wavenumbers is similar to decreasing the displacements by a factor of 2.

- Galaxy displacements generally drive the efficacy of reconstruction rather than other properties (such as the height of peaks that collapse to form halos). This result supports why the displacement reconstruction of [34] appears to be as successful as other algorithms. Additional properties beyond displacements — which are weighted to higher wavenumbers where the problem is more under-constrained — contribute an $r < 0.5$ tail to high wavenumbers in many of the mock surveys we investigated. Extracting cosmological constraints from this tail may be challenging.
- The displacements constrain $\approx 3N$ independent numbers at a cosmologically interesting level for a regularization that mimics the Gaussian prior used by reconstruction algorithms, where N is the number of galaxies (and an analogous conclusion applies to the N Lagrangian overdensity conditions). While it might seem that our setup of $\approx 3N$ equations drives this result, it did not have to be the case that the problem was sufficiently well conditioned to result in $3N$ cosmologically meaningful constraints. We showed that the conditioning was such that most of these $3N$ numbers could be reconstructed with reasonable errors on the displacements, errors that could be achieved with our *realistic* nonlinear reconstruction algorithm. The effectiveness of how a particular mode is reconstructed depends on whether it projects onto the set of well-constrained eigenvectors. For displacements, these well-constrained eigenvectors are roughly approximated by the $3N$ lowest wavenumber Fourier modes and, thus, the scale that sets the rather abrupt transition from where the cross correlation coefficient goes from one to zero is well approximated by where the number of constraints is equal to the number of modes. This constraint-counting wavenumber we showed is strikingly similar to the wavenumber where shot noise begins to dominate the power in the concordance cosmology across both redshift and halo mass threshold, possibly explaining why others had attributed this transition to shot noise.
- We found that if galaxy displacements can be sufficiently well constrained, the $1 - r^2$ of the input field with the reconstructed field falls below the shot noise expectation at low wavenumbers. For the standard Gaussian prior on mode amplitudes, the limiting standard deviation is a factor of $\sim (k/[1 \text{ Mpc}^{-1}])^2$ smaller than the naive shot noise expectation. This result was derived with a model where the displacements are Zel-dovich and phrased in terms of the error on the displacement, a displacement which a reconstruction algorithm would have to model. We could only reproduce this behavior when solving our full nonlinear problem when starting in an unphysical manner where we initialized the least squares solver with the input overdensity field at low wavenumbers. If we instead initiated the reconstruction in a more realistic way that uses the galaxy field as our starting point, a precision similar to the standard shot noise limit resulted. However, visually this reconstructed field has obvious failures, indicating to us that improvements may be possible.
- Our nonlinear setup provides intuition into why reconstruction is successful: the displacements from scales where the galaxy field does not correlate well with the linear overdensity field tend to be comparable or smaller than the Lagrangian halo size — with “smaller than” being the direction where reconstruction becomes a convex problem. As a result of this convexity, our nonlinear reconstruction was almost able to saturate the bounds set from our linear reconstruction from a starting point that is far from the true nonlinear field.

We hope to analyze our toy nonlinear problem in future work, especially to better understand the sensitivity limit at low wavenumbers. Understanding the effect of redshift space distortions and of modeling imperfections owing to astrophysics in our simplified setup are directions that may also merit study.

The jury is still out as to whether fully nonlinear reconstruction algorithms can be applied in a controlled manner to large-scale structure data. As a point of optimism, we note that the situation was similar for the perturbation theory of large-scale structure until recently, when theories were developed to the point where they could be applied to SDSS/BOSS galaxy power spectrum measurements [52, 53]. However, perturbation theory is under better control than the ungainly methods used by nonlinear reconstruction algorithms. One wonders if there is some reduction of the problem that allows more control, an example being the simpler-to-understand displacement reconstruction of [34]. It would be a shame if analyses of spectroscopic galaxy surveys were limited to low order statistics on perturbative scales — we understand the dynamics of structure formation on scales well beyond where the perturbative solutions are applicable.

Another unresolved issue regards what limits the precision of reconstruction at low wavenumbers. The finding in much the literature is that it is something shot noise-like. Shot noise sets the limit when one Fourier transforms a gridded halo field, but even there it is mitigated to the extent halos trace the comoving mass distribution [14, 37, 38]. However, in the setup of the problem discussed here (where there is a deterministic model that relates the input to a final field), it is far from apparent how shot noise should manifest in the reconstructed field. In the limit where one can sufficiently constrain the displacement of each galaxy, we showed that the error could be orders of magnitude smaller than the traditional shot estimate at low wavenumbers and in a manner that is not white noise-like. Yet, we did not resolve whether galaxy reconstruction could achieve these errors in a realistic setting. A significant reduction over shot noise would enable better constraints on primordial non-Gaussianity and neutrino masses and so we think that pursuing a more fundamental understanding is worthwhile.

Acknowledgments

We thank Yen-Chi Chen, Avery Meiksin, Uros Seljak, Blake Sherwin, Martin White and Jennifer Yeh for useful conversations. We thank Vid Iršič, Marcel Schmittfull, and especially Joanne Cohn for useful comments on an earlier version of this manuscript. We acknowledge support from NSF AAG award 2007012.

References

- [1] D.J. Eisenstein, H.-j. Seo and M.J. White, *On the Robustness of the Acoustic Scale in the Low-Redshift Clustering of Matter*, *Astrophys. J.* **664** (2007) 660 [[astro-ph/0604361](#)] [[INSPIRE](#)].
- [2] F.S. Kitaura and T.A. Ensslin, *Bayesian reconstruction of the cosmological large-scale structure: methodology, inverse algorithms and numerical optimization*, *Mon. Not. Roy. Astron. Soc.* **389** (2008) 497 [[arXiv:0705.0429](#)] [[INSPIRE](#)].
- [3] 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](#)].
- [4] S. Tassev and M. Zaldarriaga, *Towards an Optimal Reconstruction of Baryon Oscillations*, *JCAP* **10** (2012) 006 [[arXiv:1203.6066](#)] [[INSPIRE](#)].

- [5] J. Jasche and B.D. Wandelt, *Bayesian physical reconstruction of initial conditions from large scale structure surveys*, *Mon. Not. Roy. Astron. Soc.* **432** (2013) 894 [[arXiv:1203.3639](#)] [[INSPIRE](#)].
- [6] M. Ata, F.-S. Kitaura and V. Müller, *Bayesian inference of cosmic density fields from non-linear, scale-dependent, and stochastic biased tracers*, *Mon. Not. Roy. Astron. Soc.* **446** (2015) 4250 [[arXiv:1408.2566](#)] [[INSPIRE](#)].
- [7] M. McQuinn and M. White, *Cosmological perturbation theory in 1 + 1 dimensions*, *JCAP* **01** (2016) 043 [[arXiv:1502.07389](#)] [[INSPIRE](#)].
- [8] Q. Pan, U.-L. Pen, D. Inman and H.-R. Yu, *Increasing Fisher information by Potential Isobaric Reconstruction*, *Mon. Not. Roy. Astron. Soc.* **469** (2017) 1968 [[arXiv:1611.10013](#)] [[INSPIRE](#)].
- [9] M. Schmittfull, T. Baldauf and M. Zaldarriaga, *Iterative initial condition reconstruction*, *Phys. Rev. D* **96** (2017) 023505 [[arXiv:1704.06634](#)] [[INSPIRE](#)].
- [10] U. Seljak, G. Aslanyan, Y. Feng and C. Modi, *Towards optimal extraction of cosmological information from nonlinear data*, *JCAP* **12** (2017) 009 [[arXiv:1706.06645](#)] [[INSPIRE](#)].
- [11] C. Modi, Y. Feng and U. Seljak, *Cosmological Reconstruction From Galaxy Light: Neural Network Based Light-Matter Connection*, *JCAP* **10** (2018) 028 [[arXiv:1805.02247](#)] [[INSPIRE](#)].
- [12] Y. Feng, U. Seljak and M. Zaldarriaga, *Exploring the posterior surface of the large scale structure reconstruction*, *JCAP* **07** (2018) 043 [[arXiv:1804.09687](#)] [[INSPIRE](#)].
- [13] F. Schmidt, F. Elsner, J. Jasche, N.M. Nguyen and G. Lavaux, *A rigorous EFT-based forward model for large-scale structure*, *JCAP* **01** (2019) 042 [[arXiv:1808.02002](#)] [[INSPIRE](#)].
- [14] M. Schmittfull, M. Simonović, V. Assassi and M. Zaldarriaga, *Modeling Biased Tracers at the Field Level*, *Phys. Rev. D* **100** (2019) 043514 [[arXiv:1811.10640](#)] [[INSPIRE](#)].
- [15] T. Nishimichi, G. D’Amico, M.M. Ivanov, L. Senatore, M. Simonović, M. Takada et al., *Blinded challenge for precision cosmology with large-scale structure: results from effective field theory for the redshift-space galaxy power spectrum*, *Phys. Rev. D* **102** (2020) 123541 [[arXiv:2003.08277](#)] [[INSPIRE](#)].
- [16] T. Baldauf, M. Mirbabayi, M. Simonović and M. Zaldarriaga, *LSS constraints with controlled theoretical uncertainties*, [arXiv:1602.00674](#) [[INSPIRE](#)].
- [17] 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](#)].
- [18] 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](#)].
- [19] C. Modi, E. Castorina, Y. Feng and M. White, *Intensity mapping with neutral hydrogen and the Hidden Valley simulations*, *JCAP* **09** (2019) 024 [[arXiv:1904.11923](#)] [[INSPIRE](#)].
- [20] L.H. Garrison, D.J. Eisenstein and P.A. Pinto, *A High-Fidelity Realization of the Euclid Code Comparison N-body Simulation with Abacus*, *Mon. Not. Roy. Astron. Soc.* **485** (2019) 3370 [[arXiv:1810.02916](#)] [[INSPIRE](#)].
- [21] S. Tassev, M. Zaldarriaga and D. Eisenstein, *Solving Large Scale Structure in Ten Easy Steps with COLA*, *JCAP* **06** (2013) 036 [[arXiv:1301.0322](#)] [[INSPIRE](#)].
- [22] E. Lawrence, K. Heitmann, M. White, D. Higdon, C. Wagner, S. Habib et al., *The Coyote Universe III: Simulation Suite and Precision Emulator for the Nonlinear Matter Power Spectrum*, *Astrophys. J.* **713** (2010) 1322 [[arXiv:0912.4490](#)] [[INSPIRE](#)].
- [23] Y. Feng, M.-Y. Chu, U. Seljak and P. McDonald, *FastPM: a new scheme for fast simulations of dark matter and haloes*, *Mon. Not. Roy. Astron. Soc.* **463** (2016) 2273 [[arXiv:1603.00476](#)] [[INSPIRE](#)].

- [24] L.H. Garrison and D.J. Eisenstein, *Generating approximate halo catalogues for blind challenges in precision cosmology*, *Mon. Not. Roy. Astron. Soc.* **485** (2019) 2407 [[arXiv:1902.09475](#)] [[INSPIRE](#)].
- [25] N. Padmanabhan, M. White and J.D. Cohn, *Reconstructing Baryon Oscillations: A Lagrangian Theory Perspective*, *Phys. Rev. D* **79** (2009) 063523 [[arXiv:0812.2905](#)] [[INSPIRE](#)].
- [26] E.A. Kazin et al., *The WiggleZ Dark Energy Survey: improved distance measurements to $z = 1$ with reconstruction of the baryonic acoustic feature*, *Mon. Not. Roy. Astron. Soc.* **441** (2014) 3524 [[arXiv:1401.0358](#)] [[INSPIRE](#)].
- [27] BOSS collaboration, *The clustering of galaxies in the completed SDSS-III Baryon Oscillation Spectroscopic Survey: baryon acoustic oscillations in the Fourier space*, *Mon. Not. Roy. Astron. Soc.* **464** (2017) 3409 [[arXiv:1607.03149](#)] [[INSPIRE](#)].
- [28] 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](#)].
- [29] C. Modi, M. White, A. Slosar and E. Castorina, *Reconstructing large-scale structure with neutral hydrogen surveys*, *JCAP* **11** (2019) 023 [[arXiv:1907.02330](#)] [[INSPIRE](#)].
- [30] R. Hada and D.J. Eisenstein, *Application of the iterative reconstruction to simulated galaxy fields*, *Mon. Not. Roy. Astron. Soc.* **482** (2019) 5685 [[arXiv:1810.05026](#)] [[INSPIRE](#)].
- [31] F. Schmidt, F. Elsner, J. Jasche, N.M. Nguyen and G. Lavaux, *A rigorous EFT-based forward model for large-scale structure*, *JCAP* **01** (2019) 042 [[arXiv:1808.02002](#)] [[INSPIRE](#)].
- [32] H. Wang, H.J. Mo, X. Yang and F.C.v.d. Bosch, *Reconstructing the Initial Density Field of the Local Universe: Methods and Tests with Mock Catalogs*, *Astrophys. J.* **772** (2013) 63 [[arXiv:1301.1348](#)] [[INSPIRE](#)].
- [33] R. Alves de Oliveira, Y. Li, F. Villaescusa-Navarro, S. Ho and D.N. Spergel, *Fast and Accurate Non-Linear Predictions of Universes with Deep Learning*, in *34th Conference on Neural Information Processing Systems*, 11, 2020 [[arXiv:2012.00240](#)] [[INSPIRE](#)].
- [34] Y. Yu, H.-M. Zhu and U.-L. Pen, *Halo nonlinear reconstruction*, *Astrophys. J.* **847** (2017) 110 [[arXiv:1703.08301](#)] [[INSPIRE](#)].
- [35] U. Seljak, *Extracting primordial non-Gaussianity without cosmic variance*, *Phys. Rev. Lett.* **102** (2009) 021302 [[arXiv:0807.1770](#)] [[INSPIRE](#)].
- [36] M. LoVerde, *Neutrino mass without cosmic variance*, *Phys. Rev. D* **93** (2016) 103526 [[arXiv:1602.08108](#)] [[INSPIRE](#)].
- [37] U. Seljak, N. Hamaus and V. Desjacques, *How to suppress the shot noise in galaxy surveys*, *Phys. Rev. Lett.* **103** (2009) 091303 [[arXiv:0904.2963](#)] [[INSPIRE](#)].
- [38] N. Hamaus, U. Seljak, V. Desjacques, R.E. Smith and T. Baldauf, *Minimizing the Stochasticity of Halos in Large-Scale Structure Surveys*, *Phys. Rev. D* **82** (2010) 043515 [[arXiv:1004.5377](#)] [[INSPIRE](#)].
- [39] J.R. Bond, S. Cole, G. Efstathiou and N. Kaiser, *Excursion set mass functions for hierarchical Gaussian fluctuations*, *Astrophys. J.* **379** (1991) 440 [[INSPIRE](#)].
- [40] A. Cooray and R.K. Sheth, *Halo Models of Large Scale Structure*, *Phys. Rept.* **372** (2002) 1 [[astro-ph/0206508](#)] [[INSPIRE](#)].
- [41] Y.B. Zel'Dovich, *Reprint of: Gravitational instability: an approximate theory for large density perturbations*, *Astron. Astrophys.* **500** (2009) 13.
- [42] M. White, *The Zel'dovich approximation*, *Mon. Not. Roy. Astron. Soc.* **439** (2014) 3630 [[arXiv:1401.5466](#)] [[INSPIRE](#)].

- [43] A. Jenkins, C.S. Frenk, S.D.M. White, J.M. Colberg, S. Cole, A.E. Evrard et al., *The Mass function of dark matter halos*, *Mon. Not. Roy. Astron. Soc.* **321** (2001) 372 [[astro-ph/0005260](#)] [[INSPIRE](#)].
- [44] L. Senatore and M. Zaldarriaga, *The IR-resummed Effective Field Theory of Large Scale Structures*, *JCAP* **02** (2015) 013 [[arXiv:1404.5954](#)] [[INSPIRE](#)].
- [45] N. Katz, T.R. Quinn and J.M. Gelb, *Galaxy Formation and the Peaks Formalism*, *Mon. Not. Roy. Astron. Soc.* **265** (1993) 689 [[INSPIRE](#)].
- [46] J.R. Bond and S.T. Myers, *The Hierarchical peak patch picture of cosmic catalogs. 1. Algorithms*, *Astrophys. J. Suppl.* **103** (1996) 1 [[INSPIRE](#)].
- [47] D. Baumann, A. Nicolis, L. Senatore and M. Zaldarriaga, *Cosmological Non-Linearities as an Effective Fluid*, *JCAP* **07** (2012) 051 [[arXiv:1004.2488](#)] [[INSPIRE](#)].
- [48] W. Press, S. Teukolsky, W. Vetterling and B. Flannery, *Numerical Recipes: The Art of Scientific Computing*, 3rd edition, Cambridge University Press (2007).
- [49] G. Stein, M.A. Alvarez, J.R. Bond, A. van Engelen and N. Battaglia, *The Websky Extragalactic CMB Simulations*, *JCAP* **10** (2020) 012 [[arXiv:2001.08787](#)] [[INSPIRE](#)].
- [50] J.E. Gunn and J.R. Gott, III, *On the Infall of Matter into Clusters of Galaxies and Some Effects on Their Evolution*, *Astrophys. J.* **176** (1972) 1 [[INSPIRE](#)].
- [51] C.C. Paige and M.A. Saunders, *Algorithm 583: LSQR: Sparse linear equations and least squares problems*, *ACM Trans. Math. Softw.* **8** (1982) 195.
- [52] M.M. Ivanov, M. Simonović and M. Zaldarriaga, *Cosmological Parameters from the BOSS Galaxy Power Spectrum*, *JCAP* **05** (2020) 042 [[arXiv:1909.05277](#)] [[INSPIRE](#)].
- [53] G. D’Amico, J. Gleyzes, N. Kokron, K. Markovic, L. Senatore, P. Zhang et al., *The Cosmological Analysis of the SDSS/BOSS data from the Effective Field Theory of Large-Scale Structure*, *JCAP* **05** (2020) 005 [[arXiv:1909.05271](#)] [[INSPIRE](#)].