# Optimal Factorization of Cosmological Large-Scale Structure Observables

Thomas Bakx[*] and Nora Elisa Chisari[†]
*Institute for Theoretical Physics, Utrecht University, Princetonplein 5, 3584 CC, Utrecht, The Netherlands*

Zvonimir Vlah[‡]
*Division of Theoretical Physics, Ruđer Bošković Institute, 10000 Zagreb, Croatia;*
*Kavli Institute for Cosmology, University of Cambridge, Cambridge CB3 0HA, United Kingdom;*
*and Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Cambridge CB3 0WA, United Kingdom*

(Received 9 July 2024; revised 28 November 2024; accepted 4 April 2025; published 13 May 2025)

We introduce COBRA (Cosmology with Optimally factorized Bases for Rapid Approximation), a novel framework for rapid computation of large-scale structure observables. COBRA separates scale dependence from cosmological parameters in the linear matter power spectrum while also minimizing the number of necessary basis terms $N_b$, thus enabling direct and efficient computation of derived and nonlinear observables. Moreover, the dependence on cosmological parameters is efficiently approximated using radial basis function interpolation. We apply our framework to decompose the linear matter power spectrum in the standard ΛCDM scenario, as well as by adding curvature, dynamical dark energy and massive neutrinos, covering all redshifts relevant for Stage IV surveys. With only a dozen basis terms $N_b$, COBRA reproduces exact Boltzmann solver calculations to ∼0.1% precision, which improves further to ∼0.02% in the pure ΛCDM scenario. Using our decomposition, we recast the one-loop redshift space galaxy power spectrum in a separable minimal-basis form, enabling ∼4000 model evaluations per second at ∼0.02% precision on a single thread. This constitutes a considerable improvement over previously existing methods (e.g., FFTLog) opening a new window for efficient computations of higher loop and higher order correlators involving multiple powers of the linear matter power spectra. The resulting factorization can also be utilized in clustering, weak lensing, and CMB analyses.

DOI: 10.1103/PhysRevLett.134.191002

*Introduction*—Large-scale structure (LSS) surveys mapping out the three-dimensional distribution of galaxies across billions of years of cosmic history will show us a unique imprint of the laws that govern the Universe. The Stage IV era of precision cosmology aims to probe the nature of dark matter and dark energy, the geometry of the Universe, and the shape of its initial conditions of the first instance after the Big Bang [1–3]. As such, any tension with the baseline ΛCDM model (e.g., [4]) could guide us to a deeper understanding of the answers to these fundamental questions. Correspondingly, the accuracy with which the distribution of galaxies and dark matter will be charted must be matched by higher accuracy of the corresponding theoretical model computation.

Perturbation theory (PT) approaches to LSS [5–8] are a first-principle way of modeling the evolution of biased tracers of the dark matter density field such as galaxies. Correlators of biased tracers receive loop corrections that are expressable as integrals over the linear power spectrum $P_L(k)$. However, the dependence of even $P_L(k)$ on cosmological parameters is not analytically tractable, nor is the scale dependence for a given cosmology. Thus, direct implementation of these predictions (using Boltzmann solvers) is slow in any Bayesian approach, where likelihoods need to be sampled millions of times. This issue is exacerbated when considering higher-order corrections to summary statistics, which are integrals over $P_L(k)$. Conventional solutions fall into two classes. The first involves constructing an "analytical basis" of functions into which $P_L(k)$ is decomposed such that the resulting integrals can be evaluated exactly through tensor multiplications [9–12]. The second approach is to emulate the resulting integrals as functions of scale and cosmology, via, e.g., neural networks or other techniques [13–21]. These approaches each have their drawbacks: *first*, useful analytical bases are rare and typically do not approximate $P_L(k)$ well unless a large number of basis functions is used, which can lead to memory issues for higher-order statistics

[*]Contact author: t.j.m.bakx@uu.nl
[†]Contact author: n.e.chisari@uu.nl
[‡]Contact author: zvlah@irb.hr

[11,22]. Furthermore, the evaluation of the resulting tensors is still technically demanding and implementation is nontrivial, especially in redshift space [23]. Typically, techniques developed for a single observable at a specific perturbative order either lack efficient generalization to higher moments and higher perturbative orders or are rendered inapplicable altogether. In addition, such an approach still needs to be combined with a Boltzmann solver to compute $P_L(k)$ at a given cosmology. *Second*, the emulation-based approach can require substantial computational resources and suffers from a lack of efficient generalization: every next quantity requires a sufficiently dense training set across all parameters.

We pursue a different solution to this issue by finding an *optimal factorization* of the scale dependence and cosmology dependence of the linear power spectrum. That is, we decompose it as

$$P_L^{\Theta}(k) = \sum_{i=1}^{N_b} w_i(\Theta) v_i(k), \tag{1}$$

where $\Theta$ indicates a set of cosmological parameters (including redshift). The $v_i(k)$ are fixed basis functions depending only on scale, which we call *scale functions*. The *weights* $w_i(\Theta)$ encode the cosmology dependence. Optimal factorization is achieved by choosing the *smallest* number of basis functions $N_b$ (see Methodology). This decomposition allows for efficient calculation of higher-order statistics but does not rely on analytic methods for loop integrals nor a Boltzmann solver. We thus reap the benefits of both approaches while circumventing their shortcomings.

In the following sections, we obtain such a decomposition and show that it facilitates computation of perturbative corrections. We then apply it to the ΛCDM $P_L(k)$. As an illustrative example, we calculate the one-loop power spectrum of galaxies in redshift space rapidly and to high precision. Some technical aspects and extensions beyond ΛCDM are found in Supplemental Material (Secs. A–C) [24].

*Methodology*—Finding a set of scale functions that achieves a decomposition as in Eq. (1) amounts to finding a low-rank approximation of a set of *template spectra* $P_{lm} = P_L^{\Theta_l}(k_m)$ evaluated at $N_t$ fixed cosmologies $\Theta_l$ and on a fixed set of wavenumbers $k_m$. This is achieved via a truncated singular value decomposition (SVD). Prior to performing the SVD, spectra are normalized by the mean of the templates $\bar{P}(k_m)$. Writing $\hat{P}_{lm} = P_L^{\Theta}(k_m)/\bar{P}(k_m)$ and $\hat{v}_i(k_m) = v_i(k_m)/\bar{P}(k_m)$ we have

$$\hat{P} \approx \hat{U}\Sigma\hat{V}^T \tag{2}$$

where $\hat{P}$ is $N_t \times N_k$, $\Sigma$ is diagonal and small ($N_b \times N_b$ where $N_b \ll N_k$), and $\hat{V}$ is $N_k \times N_b$ containing the principal components as its orthonormal column vectors, i.e. $\hat{V}_{mi} = \hat{v}_i(k_m)$. Lastly, $\hat{U}$ is $N_t \times N_b$ and contains the weights $w_i(\Theta_l)$ (for related work, see e.g. [14,16,17,52–54]). Conducting the SVD is cheap [55] and can be done with many ($N_t > 10^7$) templates, which need not be calculated exactly—they should only mimic the shape of $P_L(k)$ to ensure that Eq. (1) is accurate. The columns of $V$ span the *optimal* $N_b$-dimensional approximation to the template set [56]. The resulting scale functions are shown in Supplemental Material (Sec. C, Fig. 6) [24].

Given scale functions, we compute weights via orthonormal projection:

$$w_i(\Theta) = \sum_{m=1}^{N_k} \hat{v}_i(k_m)\hat{P}_L^{\Theta}(k_m). \tag{3}$$

We stress that evaluating the weights $w_i(\Theta)$ is a *separate* problem, requiring either (i) exact calculation of $\hat{P}_L^{\Theta}(k_m)$ with e.g. CAMB and applying Eq. (3) or (ii) an indirect strategy using, e.g., neural networks. We opt for a different indirect strategy based on *radial basis functions* (RBFs) [57], which we describe in Supplemental Material (Sec. A) [24].

Armed with Eq. (1) it becomes simple to compute next-to-leading order corrections to observables. For example, for the power spectrum (prior to IR-resummation) in redshift space at one-loop order (see, e.g., [58–62]) one schematically has

$$\begin{aligned} P_{\text{1-loop}}^{\Theta}(k,\mu) = {} & \text{const}(k,\mu) + \mathcal{S}^l[P_L^{\Theta}](k,\mu) \\ & + \mathcal{S}^q[P_L^{\Theta}, P_L^{\Theta}](k,\mu), \end{aligned} \tag{4}$$

where $\text{const}(k,\mu)$ does not involve $P_L(k)$ while $\mathcal{S}^l$ and $\mathcal{S}^q$ are linear and quadratic operators that do not depend on cosmology. Here $l$ and $q$ superscripts refer to terms linear and quadratic in $P_L(k)$. Concretely, $\text{const}(k,\mu)$ consists of stochastic terms while $\mathcal{S}^l$ involves the linear theory part and counterterms $\propto k^2 P_L(k)\mu^{2n}$, and finally, $\mathcal{S}^q$ consists of (22) and (13)-type contributions to the loops. Plugging in Eq. (1) yields

$$\begin{aligned} P_{\text{1-loop}}^{\Theta}(k,\mu) = {} & \text{const}(k,\mu) + \mathcal{S}_i^l(k,\mu) w_i(\Theta) \\ & + \mathcal{S}_{ij}^q(k,\mu) w_i(\Theta) w_j(\Theta), \end{aligned} \tag{5}$$

where $\mathcal{S}_i^l = \mathcal{S}^l[v_i]$ and $\mathcal{S}_{ij}^q = \mathcal{S}^q[v_i, v_j]$. This reduces calculating $P_{\text{1-loop}}^{\Theta}$ to multiplications of precomputed matrices whose entries are integrals of scale functions against PT kernels. Similar arguments apply to other N-point functions and higher PT orders [10,11]. We can also extend this to the redshift-space galaxy power spectrum, including infrared (IR) resummation [see Supplemental Material (Sec. C) [24]].

*Linear power spectrum*—We decompose of $P_L(k)$ in four scenarios, varying the cosmological parameter space (ΛCDM or *generalized*) and ranges of parameters (*default*

TABLE I. Ranges and (linearly spaced) template grids for ΛCDM parameters. If a parameter is held fixed, its fiducial value is indicated.

| | Default | | Extended | |
|---|---|---|---|---|
| $\Theta$ | Range | Grid size | Range | Grid size |
| $\omega_c$ | [0.095, 0.145] | 27 | [0.08, 0.175] | 40 |
| $\omega_b$ | [0.0202, 0.0238] | 18 | [0.020, 0.025] | 20 |
| $n_s$ | [0.91, 1.01] | 12 | [0.8, 1.2] | 20 |
| $10^9 A_s$ | ⋯ | $10^9 A_s^* = 2$ | ⋯ | $10^9 A_s^* = 2$ |
| $h$ | [0.55, 0.8] | $h^* = 0.7$ | [0.5, 0.9] | $h^* = 0.7$ |
| $z$ | [0.1, 3] | $z^* = 0$ | [0.1, 3] | $z^* = 0$ |

or *extended*). We choose the range $8 \times 10^{-4} h^*/\mathrm{Mpc} < k < 4h^*/\mathrm{Mpc}$. We use $h^* = 0.7$ and compute $P_L(k)$ with CAMB [63,64] (v1.5.2). We show results for ΛCDM in the main text and defer generalized cosmologies including curvature, dynamical dark energy and neutrinos to Supplemental Material (Sec. B) [24].

In ΛCDM, we consider $\{\Theta\} = \{\omega_b, \omega_c, n_s, A_s, h, z\}$. The shape of $P_L(k)$ does not depend on the evolution parameters $\Theta_e = \{A_s, h, z\}$ when the shape parameters $\Theta_s = \{\omega_c, \omega_b, n_s\}$ are held fixed. Thus, for the SVD we only vary $\omega_c$, $\omega_b$, and $n_s$ [15,65,66]. The ranges of all parameters and choices for the SVD are indicated in Table I [67]. We thus compute $P_L(k)$ at fixed evolution parameters $\Theta_e^*$ as

$$P_L^{\Theta_s,\Theta_e^*}(k) = \sum_{i=1}^{N_b} w_i(\Theta_s) v_i(k) \tag{6}$$

and for arbitrary evolution parameters as

$$P_L^{\Theta_s,\Theta_e}(k) = \frac{A_s}{A_s^*} \frac{D_+^2(\omega_m, h, z)}{D_+^2(\omega_m, h^*, z^*)} P_L^{\Theta_s,\Theta_e^*}(k) \tag{7}$$

with $\omega_m = \omega_c + \omega_b$. The ratio of growth factors $D_+$ in Eq. (7) is also approximated using RBFs. We found it beneficial to first divide by the exact expression for a universe with $\Omega_m + \Omega_\Lambda = 1$ [68]. For the RBF approximations we use $N_n = 400$ Halton nodes [69]. We use 5000 cosmologies to test the precision of the predictions. The result is shown in Fig. 1. The default (extended) range requires $N_b = 9(13)$ basis functions for 0.01% precision for 99.7% ($3\sigma$) of the test cosmologies. Increasing $N_b$ to 12 (16) decreases the $3\sigma$ error to ∼0.02% [70]. One prediction for $P_L(k)$ takes ∼0.4 ms, while vectorized evaluation yields 250 spectra in ∼4 ms, all on one thread [73].

*One-loop galaxy power spectrum*—Using the decomposition of $P_L(k)$ from Methodology, the calculation of higher-order corrections to N-point functions is straightforward. We illustrate this using the one-loop power spectrum of galaxies in redshift space for a ΛCDM cosmology, but emphasize that this choice is irrelevant: the calculation of higher N-point functions requires only a one-time computation of a limited number of integrals, which can be done using *any method*, regardless of whether analytical techniques are available.

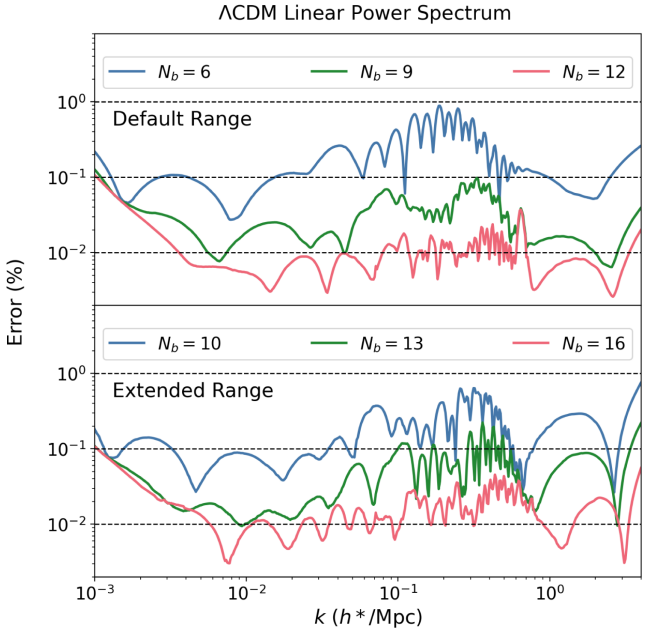


FIG. 1. The 99.7th percentile errors on the ΛCDM $P_L(k)$ for several choices of $N_b$, both for default (upper panel) and extended (lower panel) ranges. Dashed lines indicate 0.01%, 0.1% and 1% errors, respectively.

To test COBRA, we compare against the one-loop model implemented in VELOCILEPTORS [59]. All terms in this model are either constant, linear, or quadratic functions of $P_L^\Theta$ [cf. Eq. (4)], so that by Eq. (5) they reduce to matrix multiplications of size $N_b$. We employ the same parameter setup as in Methodology. The IR resummation prescription is detailed in Supplementary Material (Sec. C) [24]. We keep operator biases fixed to the values listed in [74]. We fix $A_s = 2 \times 10^{-9}$ and put all coefficients $\propto k^2\mu^{2n} P_{L,\mathrm{IR}}(k)$ equal to 40(50) in the default (extended) case and counterterms $\propto k^2\mu^{2n}$ to 3000 [75]. The $k \to 0$ contributions are subtracted to recover linear theory on large scales. We use RBF approximations for the ΛCDM growth rate and velocity dispersion [see Supplemental Material (Sec. C) [24]]; their impact on the error is small. We omit Alcock-Paczynski rescaling, but this can be included at no cost since COBRA computes the full anisotropic power.

Figure 2 displays the error for the monopole using COBRA versus using VELOCILEPTORS for $10^{-3} h^*/\mathrm{Mpc} < k < 0.5 h^*/\mathrm{Mpc}$, using the same test set as in Methodology. We obtain similar results for the quadrupole and hexadecapole. For the default (extended) range, we use $N_b = 12(16)$ for the linear part $\mathcal{S}_i^l$ from Eq. (5). We explore different choices of $N_b$ for $\mathcal{S}_{ij}^q$, which dominates computation time. For the default (extended) range, using $N_b = 9(13)$ scale functions for $\mathcal{S}_{ij}^q$ we reach ∼0.01% precision for 99.7% of all

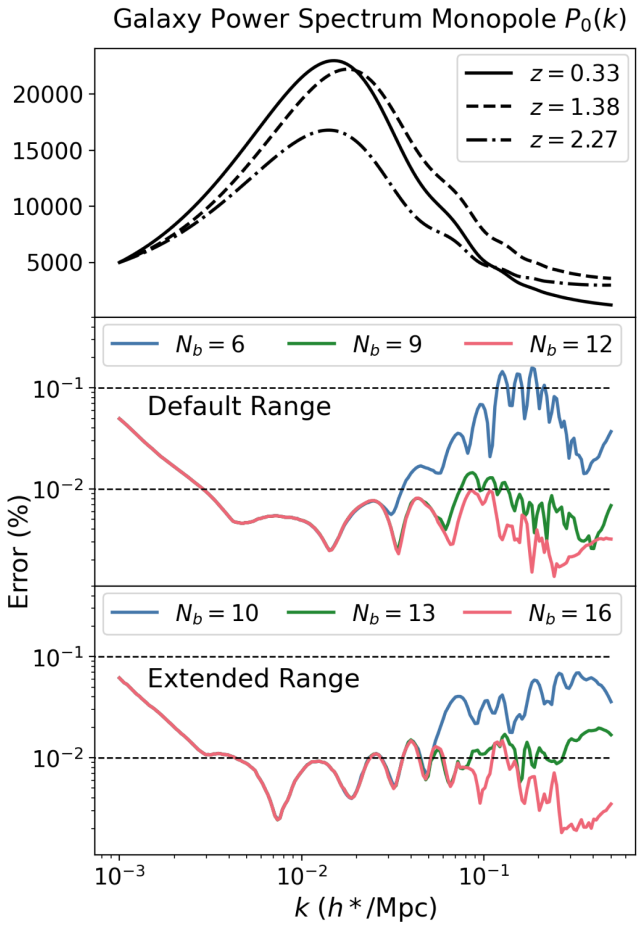


FIG. 2. Performance of COBRA for the monopole of the IR-resummed one-loop power spectrum of galaxies in redshift space. The top panel shows three random cosmologies (solid, dashed, and dot-dashed lines) and their redshifts. The second and third panels show the 99.7th percentile of the errors on the test set for the default and extended ranges. For visualization, spectra are normalized to the same overall amplitude.

test cosmologies. With $N_b = 12(16)$ for $\mathcal{S}^q_{ij}$ and 200$k$-bins, the matrices needed for all bias terms require around 25(40) MB memory. One prediction of three multipoles takes ∼2.5 ms, while 250 predictions take ∼65 ms. Thus, COBRA executes ∼4000 predictions per second. This speed is unaltered when varying bias parameters.

*Discussion*—We introduced COBRA, a method for efficient computation of large-scale structure observables, and applied it to the linear power spectrum and the one-loop power spectrum of galaxies in redshift space. Generalizing from the Introduction [and ignoring IR resummation for simplicity, see Supplemental Material (Sec. C) [24]], all polyspectra $\mathcal{P}$ take the form [10,11]

$$\mathcal{P}^{\Theta}_{\ell\text{-loop}} = \text{const} + \mathcal{S}^l[P^{\Theta}_L] + \mathcal{S}^q[P^{\Theta}_L, P^{\Theta}_L] + \mathcal{S}^c[P^{\Theta}_L, P^{\Theta}_L, P^{\Theta}_L] + \ldots, \tag{8}$$

where $\mathcal{S}^l$, $\mathcal{S}^q$, and $\mathcal{S}^c$ are linear, quadratic and cubic operators *et cetera*. Using Eq. (1),

$$\mathcal{P}^{\Theta}_{\ell\text{-loop}} = \text{const} + \mathcal{S}^l_i w_i(\Theta) + \mathcal{S}^q_{ij} w_i(\Theta) w_j(\Theta) + \mathcal{S}^c_{ijk} w_i(\Theta) w_j(\Theta) w_k(\Theta) + \ldots \tag{9}$$

with $\mathcal{S}^l_i = \mathcal{S}^l[v_i], \mathcal{S}^q_{ij} = \mathcal{S}^q[v_i, v_j]$ and $\mathcal{S}^c_{ijk} = \mathcal{S}^c[v_i, v_j, v_k]$ symmetric tensors. Most terms contain factors of $P_L(k)$ outside the loop integral. For the one-loop bispectrum, the most demanding contribution is the $B_{222}$ diagram which integrates over three $P_L(k)$. Naive counting with $N_b = 6$ and 15 bias terms yields ∼1000 integrals to be evaluated, which is much more efficient than traditional methods [10,22]. Moreover, the decay of singular values in SVD employed in COBRA [see Supplemental Material (Sec. C) [24]] implies that the number of basis functions necessary to reach a given precision can be further reduced in terms that involve products of $P_L(k)$. Integrated quantities (like mode coupling terms) might benefit from additional reduction in the number of terms without compromising precision [11]. Accuracy requirements also loosen substantially when data covariance is taken into account [20,76]. Given the small number of basis functions needed, there is ample opportunity to generalize our calculations to e.g. exact time dependence [68,77–79], scale-dependent growth [80] or higher-order N-point functions [23,81–91].

To summarize, COBRA is *fast:* the calculation for $P_L(k)$ takes ∼1 ms on one thread while the loop calculation takes ∼2 ms. Implementations are vectorized, which implies further speed-ups when computing $\mathcal{O}(100)$ predictions simultaneously. This can be exploited in conjunction with vectorized likelihood samplers [92]. COBRA is *precise*: we reach ∼0.1% precision on all observables considered. Precision can be adjusted by tuning $N_b$, which is useful if further speed-up is desired and does not require recalculating the tensors in Eq. (9). COBRA is *general*: analytical tools for calculating loop corrections are unnecessary and our method applies to any N-point function at any loop order. COBRA is *lightweight*: The tabulated scale functions and tensors for the one-loop integrals require negligible memory. Our is publicly available at [93] and requires only NumPy and SciPy.

No other method available in the literature can produce one-loop predictions with such efficiency and precision. COBRA is applicable to all orders and higher-loop computations, making it a more versatile and powerful tool as well. In particular, it would be interesting to explicitly consider the calculation of higher N-point functions such as the one-loop bispectrum and two-loop power spectrum—improvements over existing methods will be even more pronounced here. The small number of coefficients used to approximate the $P_L(k)$ facilitates direct reconstruction [94–96]. Other cosmology-independent and linear operations on observables, such as window convolution [97] or

compression schemes [52], could be incorporated in COBRA as a one-time calculation at the level of scale functions. One could explore taking analytic derivatives of RBF approximations in the context of gradient-based sampling methods [98,99]. Finally, while our focus in terms of scales and parameter ranges has been on 3D galaxy clustering on large scales, it would be valuable to extend this to sky-projected [100,101] or nonlinear quantities. These are more relevant for weak lensing and CMB probes.

*Acknowledgments*—We thank David Alonso, Marco Bonici, Giovanni Cabass, Alexander Eggemeier, and Pedro Ferreira for useful comments. We acknowledge extensive use of the open-source Python libraries NumPy, SciPy, and SCIKIT-LEARN. This publication is part of the project "A rising tide: Galaxy intrinsic alignments as a new probe of cosmology and galaxy evolution" (with Project Number VI.Vidi.203.011) of the Talent programme Vidi, which is (partly) financed by the Dutch Research Council (NWO). Z. V. acknowledges the support of the Kavli Foundation.

[1] A. Aghamousa *et al.* (DESI Collaboration), The DESI experiment part I: Science, targeting, and survey design, arXiv:1611.00036.
[2] R. Scaramella *et al.* (Euclid Collaboration), Euclid preparation. I. The Euclid wide survey, Astron. Astrophys. **662**, A112 (2022).
[3] LSST Dark Energy Science Collaboration, Large synoptic survey telescope: Dark Energy Science Collaboration, arXiv:1211.0310.
[4] A. G. Adame *et al.* (DESI Collaboration), DESI 2024 VI: Cosmological constraints from the measurements of baryon acoustic oscillations, J. Cosmol. Astropart. Phys. 02 (2025) 021.
[5] F. Bernardeau, S. Colombi, E. Gaztañaga, and R. Scoccimarro, Large-scale structure of the Universe and cosmological perturbation theory, Phys. Rep. **367**, 1 (2002).
[6] V. Desjacques, D. Jeong, and F. Schmidt, Large-scale galaxy bias, Phys. Rep. **733**, 1 (2018).
[7] J. J. M. Carrasco, M. P. Hertzberg, and L. Senatore, The effective field theory of cosmological large scale structures, J. High Energy Phys. 09 (2012) 082.
[8] D. Baumann, A. Nicolis, L. Senatore, and M. Zaldarriaga, Cosmological non-linearities as an effective fluid, J. Cosmol. Astropart. Phys. 07 (2012) 051.
[9] M. Schmittfull and Z. Vlah, Reducing the two-loop large-scale structure power spectrum to low-dimensional, radial integrals, Phys. Rev. D **94**, 103530 (2016).
[10] M. Simonović, T. Baldauf, M. Zaldarriaga, J. J. Carrasco, and J. A. Kollmeier, Cosmological perturbation theory using the FFTLog: Formalism and connection to QFT loop integrals, J. Cosmol. Astropart. Phys. 04 (2018) 030.
[11] C. Anastasiou, D. P. L. Bragança, L. Senatore, and H. Zheng, Efficiently evaluating loop integrals in the EFTofLSS using QFT integrals with massive propagators, J. High Energy Phys. 01 (2024) 002.
[12] M. Schmittfull, Z. Vlah, and P. McDonald, Fast large scale structure perturbation theory using one-dimensional fast Fourier transforms, Phys. Rev. D **93**, 103528 (2016).
[13] D. J. Bartlett, L. Kammerer, G. Kronberger, H. Desmond, P. G. Ferreira, B. D. Wandelt, B. Burlacu, D. Alonso, and M. Zennaro, A precise symbolic emulator of the linear matter power spectrum, Astron. Astrophys. **686**, A209 (2024).
[14] G. Aricò, R. E. Angulo, and M. Zennaro, Accelerating large-scale-structure data analyses by emulating Boltzmann solvers and Lagrangian perturbation Theory, arXiv:2104.14568.
[15] A. Eggemeier, B. Camacho-Quevedo, A. Pezzotta, M. Crocce, R. Scoccimarro, and A. G. Sánchez, COMET: Clustering observables modelled by emulated perturbation theory, Mon. Not. R. Astron. Soc. **519**, 2962 (2023).
[16] J. DeRose, S.-F. Chen, M. White, and N. Kokron, Neural network acceleration of large-scale structure theory calculations, J. Cosmol. Astropart. Phys. 04 (2022) 056.
[17] A. Spurio Mancini, D. Piras, J. Alsing, B. Joachimi, and M. P. Hobson, COSMOPOWER: Emulating cosmological power spectra for accelerated Bayesian inference from next-generation surveys, Mon. Not. R. Astron. Soc. **511**, 1771 (2022).
[18] M. Cataneo, S. Foreman, and L. Senatore, Efficient exploration of cosmology dependence in the EFT of LSS, J. Cosmol. Astropart. Phys. 04 (2017) 026.
[19] S.-F. Chen, Z. Vlah, E. Castorina, and M. White, Redshift-space distortions in Lagrangian perturbation theory, J. Cosmol. Astropart. Phys. 03 (2021) 100.
[20] S. Trusov, P. Zarrouk, and S. Cole, Neural network-based model of galaxy power spectrum: Fast full-shape galaxy power spectrum analysis, Mon. Not. R. Astron. Soc. **538**, 1789 (2025).
[21] S. Ramirez, M. Icaza-Lizaola, S. Fromenteau, M. Vargas-Magaña, and A. Aviles, Full shape cosmology analysis from BOSS in configuration space using neural network acceleration, J. Cosmol. Astropart. Phys. 08 (2024) 049.
[22] O. H. E. Philcox, M. M. Ivanov, G. Cabass, M. Simonović, M. Zaldarriaga, and T. Nishimichi, Cosmology with the redshift-space galaxy bispectrum monopole at one-loop order, Phys. Rev. D **106**, 043530 (2022).
[23] G. D'Amico, Y. Donath, M. Lewandowski, L. Senatore, and P. Zhang, The BOSS bispectrum analysis at one loop from the effective field theory of large-scale structure, J. Cosmol. Astropart. Phys. 05 (2024) 059.
[24] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevLett.134.191002 for technical aspects of radial basis functions (Section A), extension of LCDM cosmology (Section B) and infrared resummation (Section C), which includes Refs. [25–51].
[25] C. Rackauckas and Q. Nie, `DifferentialEquations.jl`–A performant and feature-rich ecosystem for solving differential equations in Julia, J. Open Res. Software **5** (2017).
[26] C. Tsitouras, Runge–Kutta pairs of order 5(4) satisfying only the first column simplifying assumption, Comput. Math. Appl. **62**, 770 (2011).

[27] Here and throughout, we will use *isotropic* Gaussians. It would certainly be interesting to consider generalizations with different (physically motivated) values of $\epsilon$ in different dimensions; this is mathematically straightforward [102] and could lead to improved performance with fewer nodes.

[28] B. Fornberg, E. Larsson, and N. Flyer, Stable computations with Gaussian radial basis functions, SIAM J. Sci. Comput. **33**, 869 (2011).

[29] B. Fornberg and C. Piret, A stable algorithm for flat radial basis functions on a sphere, SIAM J. Sci. Comput. **30**, 60 (2008).

[30] T. Driscoll and B. Fornberg, Interpolation in the limit of increasingly flat radial basis functions, Comput. Math. Appl. **43**, 413 (2002).

[31] G. B. Wright and B. Fornberg, Stable computations with flat radial basis functions using vector-valued rational approximations, J. Comput. Phys. **331**, 137 (2017).

[32] In fact, when $\epsilon \to 0$ the $\varphi_k^\epsilon(\theta)$ reduce to Hermite polynomials; polynomial regression techniques were also considered by [103,104] for the `EuclidEmulator`.

[33] M. Chevallier and D. Polarski, Accelerating universes with scaling dark matter, Int. J. Mod. Phys. D **10**, 213 (2001).

[34] E. V. Linder, Exploring the expansion history of the universe, Phys. Rev. Lett. **90**, 091301 (2003).

[35] We can mimic the effect of curvature on the shape on large scales by multiplying all spectra by an (empirical) "fudge factor" of the form $f(k) = 1 + a_0\Omega_K(k_0/k)^2$ for a linear grid of 12 values of $\Omega_K$. The dependence of the power spectrum on $n_s$ is trivial, meaning that templates for varying $n_s$ do not need to be recomputed either (as long as one works with a rectangular grid). An empirical correction may also be possible for the dark energy equation of state parameter $w_0$, further reducing the number of power spectrum evaluations necessary. We did not pursue this in more detail.

[36] For simplicity, we refrain from extending the range for $\Omega_K$ further since this would allow $\Omega_\Lambda < 0$ in which case the expansion from Eq. (B2) is no longer valid [105].

[37] We have attempted bringing $w_+$ closer to zero, but this resulted in significantly degraded performance, in line with findings of [104]. In any case, this region of parameter space is quite pathological and appears to be disfavored [4].

[38] A. Chudaykin and M. M. Ivanov, Measuring neutrino masses with large-scale structure: Euclid forecast with controlled theoretical error, J. Cosmol. Astropart. Phys. 11 (2019) 034.

[39] D. Bragança, Y. Donath, L. Senatore, and H. Zheng, Peeking into the next decade in large-scale structure cosmology with its effective field theory, arXiv:2307.04992.

[40] A. Chudaykin, K. Dolgikh, and M. M. Ivanov, Constraints on the curvature of the Universe and dynamical dark energy from the full-shape and BAO data, Phys. Rev. D **103**, 023507 (2021).

[41] S. Lee and K.-W. Ng, Growth index with the exact analytic solution of sub-horizon scale linear perturbation for dark energy models with constant equation of state, Phys. Lett. B **688**, 1 (2010).

[42] See, e.g., https://docs.scipy.org/doc/scipy/reference/generated/scipy.special.logit.html.

[43] See, e.g., the fast ODE solver from https://cosmologicalemulators.github.io/Effort.jl/dev/ based on odesolv1, odesolv2. We thank Marco Bonici for pointing this out to us.

[44] T. Baldauf, M. Mirbabayi, M. Simonović, and M. Zaldarriaga, Equivalence principle and the baryon acoustic peak, Phys. Rev. D **92**, 043514 (2015).

[45] Z. Vlah, U. Seljak, M. Yat Chu, and Y. Feng, Perturbation theory, effective field theory, and oscillations in the power spectrum, J. Cosmol. Astropart. Phys. 03 (2016) 057.

[46] M. M. Ivanov and S. Sibiryakov, Infrared resummation for biased tracers in redshift space, J. Cosmol. Astropart. Phys. 07 (2018) 053.

[47] D. Blas, M. Garny, M. M. Ivanov, and S. Sibiryakov, Time-sliced perturbation theory II: Baryon acoustic oscillations and infrared resummation, J. Cosmol. Astropart. Phys. 07 (2016) 028.

[48] D. Baumann, D. Green, and B. Wallisch, Searching for light relics with large-scale structure, J. Cosmol. Astropart. Phys. 08 (2018) 029.

[49] A. Chudaykin, M. M. Ivanov, O. H. E. Philcox, and M. Simonović, Nonlinear perturbation theory extension of the Boltzmann code CLASS, Phys. Rev. D **102**, 063533 (2020).

[50] D. J. Eisenstein and W. Hu, Power spectra for cold dark matter and its variants, Astrophys. J. **511**, 5 (1999).

[51] S.-F. Chen, Z. Vlah, and M. White, The bispectrum in Lagrangian perturbation theory, J. Cosmol. Astropart. Phys. 11 (2024) 012.

[52] O. H. E. Philcox, M. M. Ivanov, M. Zaldarriaga, M. Simonović, and M. Schmittfull, Fewer mocks and less noise: Reducing the dimensionality of cosmological observables with subspace projections, Phys. Rev. D **103**, 043508 (2021).

[53] L. Pathak, A. Reza, and A. S. Sengupta, Fast and faithful interpolation of numerical relativity surrogate waveforms using meshfree approximation, Phys. Rev. D **110**, 064022 (2024).

[54] L. Pathak, S. Munishwar, A. Reza, and A. S. Sengupta, Prompt sky localization of compact binary sources using a meshfree approximation, Phys. Rev. D **109**, 024053 (2024).

[55] N. Halko, P. G. Martinsson, and J. A. Tropp, Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions, SIAM Rev. **53**, 217 (2011).

[56] C. Eckart and G. Young, The approximation of one matrix by another of lower rank, Psychometrika **1**, 211 (1936).

[57] M. D. Buhmann, *Radial Basis Functions: Theory and Implementations*, Cambridge Monographs on Applied and Computational Mathematics (Cambridge University Press, Cambridge, England, 2003).

[58] L. Senatore, Bias in the effective field theory of large scale structures, J. Cosmol. Astropart. Phys. 11 (2015) 007.

[59] S.-F. Chen, Z. Vlah, and M. White, Consistent modeling of velocity statistics and redshift-space distortions in one-loop perturbation theory, J. Cosmol. Astropart. Phys. 07 (2020) 062.

[60] L. Senatore and M. Zaldarriaga, Redshift space distortions in the effective field theory of large scale structures, arXiv:1409.1225.

[61] A. Perko, L. Senatore, E. Jennings, and R. H. Wechsler, Biased tracers in redshift space in the EFT of large-scale structure, arXiv:1610.09321.

[62] R. Angulo, M. Fasiello, L. Senatore, and Z. Vlah, On the statistics of biased tracers in the effective field theory of large scale structures, J. Cosmol. Astropart. Phys. 09 (2015) 029.

[63] A. Lewis and A. Challinor, CAMB: Code for anisotropies in the microwave background (2011), https://ascl.net/1102.026.

[64] A. Lewis, A. Challinor, and A. Lasenby, Efficient computation of cosmic microwave background anisotropies in closed Friedmann-Robertson-Walker models, Astrophys. J. **538**, 473 (2000).

[65] A. G. Sánchez, A. N. Ruiz, J. G. Jara, and N. D. Padilla, Evolution mapping: A new approach to describe matter clustering in the non-linear regime, Mon. Not. R. Astron. Soc. **514**, 5673 (2022).

[66] A. G. Sánchez, Arguments against using $h^{-1}$ Mpc units in observational cosmology, Phys. Rev. D **102**, 123511 (2020).

[67] We do not extend the range for $\omega_b$ appreciably since in the context of spectroscopic clustering one typically employs a BBN prior [4].

[68] M. Fasiello, T. Fujita, and Z. Vlah, Perturbation theory of large scale structure in the ΛCDM Universe: Exact time evolution and the two-loop power spectrum, Phys. Rev. D **106**, 123504 (2022).

[69] A. B. Owen, A randomized Halton algorithm in R, arXiv:1706.02808.

[70] While the agreement between the different Boltzmann codes CAMB and CLASS [71] may not be at that level [14,72], it is a testament to the precision of our method that such small errors can be achieved.

[71] D. Blas, J. Lesgourgues, and T. Tram, The cosmic linear anisotropy solving system (CLASS). Part II: Approximation schemes, J. Cosmol. Astropart. Phys. 07 (2011) 034.

[72] J. Lesgourgues, The cosmic linear anisotropy solving system (CLASS) III: Comparision with CAMB for LambdaCDM, arXiv:1104.2934.

[73] Tests are run on an Apple M1 Pro processor (16GB RAM).

[74] https://github.com/sfschen/velocileptors/blob/master/notebooks/EPT%20Examples.ipynb

[75] We do this purely to avoid zero crossings in the monopole, which are unphysical but nevertheless could occur due to large loop contributions for some cosmologies.

[76] D. Linde, A. Moradinezhad Dizgah, C. Radermacher, S. Casas, and J. Lesgourgues, CLASS-OneLoop: Accurate and unbiased inference from spectroscopic galaxy surveys, J. Cosmol. Astropart. Phys. 07 (2024) 068.

[77] N. Choustikov, Z. Vlah, and A. Challinor, Optimizing the evolution of perturbations in the ΛCDM universe, Phys. Rev. D **108**, 023529 (2023).

[78] Y. Donath and L. Senatore, Biased tracers in redshift space in the EFTofLSS with exact time dependence, J. Cosmol. Astropart. Phys. 10 (2020) 039.

[79] M. Hartmeier and M. Garny, Minimal basis for exact time dependent kernels in cosmological perturbation theory and application to ΛCDM and $w_0 w_a$ CDM, J. Cosmol. Astropart. Phys. 12 (2023) 027.

[80] M. Levi and Z. Vlah, Massive neutrinos in nonlinear large scale structure: A consistent perturbation theory, arXiv:1605.09417.

[81] A. Eggemeier, R. Scoccimarro, and R. E. Smith, Bias loop corrections to the galaxy bispectrum, Phys. Rev. D **99**, 123514 (2019).

[82] A. Eggemeier, R. Scoccimarro, R. E. Smith, M. Crocce, A. Pezzotta, and A. G. Sánchez, Testing one-loop galaxy bias: Joint analysis of power spectrum and bispectrum, Phys. Rev. D **103**, 123550 (2021).

[83] G. D'Amico, Y. Donath, M. Lewandowski, L. Senatore, and P. Zhang, The one-loop bispectrum of galaxies in redshift space from the effective field theory of large-scale structure, J. Cosmol. Astropart. Phys. 07 (2024) 041.

[84] Y. Donath, M. Lewandowski, and L. Senatore, Direct signatures of the formation time of galaxies, Phys. Rev. D **109**, 123510 (2024).

[85] K. Osato, T. Nishimichi, F. Bernardeau, and A. Taruya, Perturbation theory challenge for cosmological parameters estimation: Matter power spectrum in real space, Phys. Rev. D **99**, 063530 (2019).

[86] K. Osato, T. Nishimichi, A. Taruya, and F. Bernardeau, Perturbation theory challenge for cosmological parameters estimation. II. Matter power spectrum in redshift space, Phys. Rev. D **108**, 123541 (2023).

[87] T. Baldauf, L. Mercolli, M. Mirbabayi, and E. Pajer, The bispectrum in the effective field theory of large scale structure, J. Cosmol. Astropart. Phys. 05 (2015) 007.

[88] T. Baldauf, M. Garny, P. Taule, and T. Steele, Two-loop bispectrum of large-scale structure, Phys. Rev. D **104**, 123551 (2021).

[89] T. Steele and T. Baldauf, Precise calibration of the one-loop trispectrum in the effective field theory of large scale structure, Phys. Rev. D **103**, 103518 (2021).

[90] J. J. M. Carrasco, S. Foreman, D. Green, and L. Senatore, The effective field theory of large scale structures at two loops, J. Cosmol. Astropart. Phys. 07 (2014) 057.

[91] D. Bertolini, K. Schutz, M. P. Solon, and K. M. Zurek, The trispectrum in the effective field theory of large scale structure, J. Cosmol. Astropart. Phys. 06 (2016) 052.

[92] J. Buchner, UltraNest—A robust, general purpose Bayesian inference engine, J. Open Source Software **6**, 3001 (2021).

[93] https://github.com/ThomasBakx/cobra

[94] L. Amendola, M. Pietroni, and M. Quartin, Fisher matrix for the one-loop galaxy power spectrum: Measuring expansion and growth rates without assuming a cosmological model, J. Cosmol. Astropart. Phys. 11 (2022) 023.

[95] L. Amendola, M. Marinucci, M. Pietroni, and M. Quartin, Improving precision and accuracy in cosmology with model-independent spectrum and bispectrum, J. Cosmol. Astropart. Phys. 01 (2024) 001.

[96] A. P. Schirra, M. Quartin, and L. Amendola, A model-independent measurement of the expansion and growth rates from BOSS using the FreePower method, arXiv:2406.15347.

[97] K. Pardede, F. Rizzo, M. Biagetti, E. Castorina, E. Sefusatti, and P. Monaco, Bispectrum-window convolution via Hankel transform, J. Cosmol. Astropart. Phys. 10 (2022) 066.
[98] J.-E. Campagne, F. Lanusse, J. Zuntz, A. Boucaud, S. Casas, M. Karamanis, D. Kirkby, D. Lanzieri, A. Peel, and Y. Li, JAX-COSMO: An end-to-end differentiable and GPU accelerated cosmology library, Open J. Astrophys. **6**, 15 (2023).
[99] J. Ruiz-Zapatero, D. Alonso, C. García-García, A. Nicola, A. Mootoovaloo, J. M. Sullivan, M. Bonici, and P. G. Ferreira, `LimberJack.jl`: Auto-differentiable methods for angular power spectra analyses, Open J. Astrophys. **7**, 11 (2024).
[100] Z. Gao, Z. Vlah, and A. Challinor, Flat-sky angular power spectra revisited, J. Cosmol. Astropart. Phys. 02 (2024) 003.
[101] A. Raccanelli and Z. Vlah, Observed power spectrum and frequency-angular power spectrum, Phys. Rev. D **108**, 043537 (2023).
[102] G. E. Fasshauer and M. J. McCourt, Stable evaluation of gaussian radial basis function interpolants, SIAM J. Sci. Comput. **34**, A737 (2012).
[103] M. Knabenhans, J. Stadel, S. Marelli, D. Potter, R. Teyssier, L. Legrand, A. Schneider, B. Sudret, L. Blot, S. Awan, C. Burigana, C. S. Carvalho, H. Kurki-Suonio, and G. Sirri (Euclid Collaboration), Euclid preparation: II. The `EuclideMulator`—A tool to compute the cosmology dependence of the nonlinear matter power spectrum, Mon. Not. R. Astron. Soc. **484**, 5509 (2019).
[104] M. Knabenhans *et al.* (Euclid Collaboration), Euclid preparation: IX. `EuclidEmulator2`—Power spectrum emulation with massive neutrinos and self-consistent dark energy perturbations, Mon. Not. R. Astron. Soc. **505**, 2840 (2021).
[105] A. J. S. Hamilton, Formulae for growth factors in expanding universes containing matter and a cosmological constant, Mon. Not. R. Astron. Soc. **322**, 419 (2001).