

Representation learning approach to probe for dynamical dark energy in matter power spectra

Davide Piras^{1,2,*} and Lucas Lombriser¹

¹*Département de Physique Théorique, Université de Genève,
24 quai Ernest Ansermet, 1211 Genève 4, Switzerland*

²*Centre Universitaire d'Informatique, Université de Genève, 7 route de Drize, 1227 Genève, Switzerland*



(Received 23 October 2023; accepted 28 May 2024; published 9 July 2024)

We present DE-VAE, a variational autoencoder (VAE) architecture to search for a compressed representation of dynamical dark energy (DE) models in observational studies of the cosmic large-scale structure. DE-VAE is trained on matter power spectra boosts generated at wave numbers $k \in (0.01-2.5) h/\text{Mpc}$ and at four redshift values $z \in (0.1, 0.48, 0.78, 1.5)$ for the most typical dynamical DE parametrization with two extra parameters describing an evolving DE equation of state. The boosts are compressed to a lower-dimensional representation, which is concatenated with standard cold dark matter (CDM) parameters and then mapped back to reconstructed boosts; both the compression and the reconstruction components are parametrized as neural networks. Remarkably, we find that a single latent parameter is sufficient to predict 95% (99%) of DE power spectra generated over a broad range of cosmological parameters within 1σ (2σ) of a Gaussian error which includes cosmic variance, shot noise, and systematic effects for a stage IV-like survey. This single parameter shows a high mutual information with the two DE parameters, and these three variables can be linked together with an explicit equation through symbolic regression. Considering a model with two latent variables only marginally improves the accuracy of the predictions, and adding a third latent variable has no significant impact on the model's performance. We discuss how the DE-VAE architecture can be extended from a proof of concept to a general framework to be employed in the search for a common lower-dimensional parametrization of a wide range of beyond- Λ CDM models and for different cosmological datasets. Such a framework could then both inform the development of cosmological surveys by targeting optimal probes, and provide theoretical insight into the common phenomenological aspects of beyond- Λ CDM models.

DOI: [10.1103/PhysRevD.110.023514](https://doi.org/10.1103/PhysRevD.110.023514)

I. INTRODUCTION

Despite its remarkable success in correctly reproducing cosmological observations, the concordance Λ CDM model still carries some unresolved questions. In particular, the nature of its two core components, namely cold dark matter (CDM) and the cosmological constant Λ , remains a mystery [1]. For example, assuming Λ as a vacuum energy density, it is currently unclear how to reconcile its value as measured by cosmological probes [2] with estimates from quantum field theory [3], invoking the concept of dark energy (DE) as the driver of the late-time accelerated expansion of the Universe. Moreover, in recent years significant tensions have emerged between different measurements of Λ CDM parameters, which could hint at modifications of the concordance model or could be attributed to experimental systematic effects. These tensions include a $\sim 5\sigma$ difference between the constraints on the expansion rate

H_0 coming from local distance ladder and CMB experiments [4], and a potential discrepancy between the values of the structure growth parameter S_8 measured with large-scale structure probes and cosmic microwave background (CMB) experiments [5,6]. Current and upcoming surveys such as the Vera Rubin Observatory¹ [7], *Euclid*² [8], the Nancy Grace Roman Space Telescope³ [9] and the Simons Observatory⁴ [10] promise to provide unprecedented insight on the nature of dark matter and dark energy, and will shed further light on these tensions.

These theoretical issues and observational discrepancies have motivated a broad variety of alternative models to Λ CDM (we refer the reader to Refs. [1,11–17] for reviews). The rich manifold of proposals ranges from dynamical dark energy and interacting dark sector models to modifications

*davide.piras@unige.ch

¹<https://www.lsst.org/>.

²<https://www.euclid-ec.org/>.

³<https://roman.gsfc.nasa.gov/>.

⁴<https://simonsobservatory.org/>.

of general relativity (GR), extra dimensions, and violations of Lorentz invariance, to name just a few. Attempts to address the Λ CDM tensions with modifications of GR also need to comply with stringent tests on small scales [16,18–22]. On the other hand, such models can exhibit complex effects that act as screening mechanisms and suppress local deviations from GR [23–27]; these are however hard to parametrize, due to their nonlinear nature (see Ref. [28] for a review). Traditionally, the impact of the manifold proposals on the matter distribution in the Universe, as measured by the galaxy correlation function or its Fourier counterpart (the power spectrum), has been of particular interest, particularly in light of the aforementioned upcoming large-scale structure surveys [29–35].

Typically, extensions to Λ CDM introduce additional parameters and functions to describe the effect of the proposed modification on cosmological observables. A variety of approaches have been introduced to develop a common parametrization for the vast space of models (see Ref. [28] for a review). In particular, in modified gravity theories the common approach is to assume a phenomenological parametrization framework where the late-time growth is modified with scale- and time-dependent functions, usually dubbed μ and Σ , which affect the Poisson equation for the gravitational and the lensing potential, respectively [32]. In order to agnostically test the plethora of models mapping onto this framework, one must then deal with large parameter spaces, and has to choose a parametric form for the additional functions; in turn, this can lead to suboptimal modeling choices, while any statistical evidence of such a hugely-extended parameter space is washed away by the large extradimensionality. For instance, ignoring scale-dependencies such as from Yukawa forces or nonlinear screening effects, only being agnostic about the time dependence of μ and Σ essentially already introduces an infinite number of parameters to constrain. This grows even further if considering dark sector interaction models, where each matter species comes with its own set of μ and Σ functions. Performing Bayesian inference in such high-dimensional scenarios gives rise to parameter degeneracies, as well as being extremely computationally expensive due to the curse of dimensionality, despite recent advances in more efficient sampling techniques like Hamiltonian Monte Carlo [36,37]. Even without taking into account systematic and observational effects that must be modeled in large-scale structure analyses, the number of theoretical parameters in Λ CDM extensions can grow infinitely due to the additional modifications being considered. Therefore, being parsimonious and finding a minimal set of parameters to successfully describe beyond- Λ CDM models is imperative.

Along these lines of research, Ref. [38] explored a data-driven parameterization of DE models based on principal component analysis (PCA). Subsequent work has developed a similar approach within a Bayesian framework [39],

and applied it in the context of weak lensing, spectroscopic and CMB surveys [40–42]. More recently, machine learning algorithms have been proposed to distinguish between different cosmological models. Reference [43] developed a deep convolutional neural network (CNN) that can discriminate between five Λ CDM models with different values of matter abundance Ω_m and matter fluctuations σ_8 given noisy convergence mass maps, with a better performance than skewness and kurtosis. Reference [44] trained a CNN on convergence maps to classify between different $f(R)$ gravities with massive neutrinos and Λ CDM, while Ref. [45] developed the Bayesian Cosmological Network (BaCoN), a Bayesian classifier that was trained to discriminate between either five different cosmological models, or between Λ CDM and non- Λ CDM, given their predicted matter power spectrum at different redshifts with 95% accuracy. More generally, machine learning algorithms have shown great capabilities at extracting and compressing information from large or high-dimensional cosmological datasets [46–51]

In this paper, we develop a variational autoencoder (VAE) to find a low-dimensional representation of a dynamical dark energy model with two additional free parameters with respect to Λ CDM; we focus on a single parametrization model of dynamical DE as a proof of concept, with the aim of extending our approach to a larger set of beyond- Λ CDM models in future work. We train the network on the theoretical matter power spectrum boost⁵ with respect to the Λ CDM prediction, and design its architecture and loss function to encourage a disentangled representation of the extra parameters. We investigate how many latent variables are needed to reconstruct the power spectra with sufficient accuracy, and investigate the links between the latents and the known additional parameters through the information-theoretic metric of mutual information (MI) and symbolic regression (SR). The latter returns a human-readable equation linking cosmological parameters and latent variables, which can then inform the development of theoretical works and new cosmological observables.

The paper is structured as follows. In Sec. II we describe the power spectrum boosts that we use to train the machine learning model, which we present in detail in Sec. III. In Sec. IV, we show the performance of the model with varying number of latent variables, and use both MI and SR to extract insight from the latent space. We discuss our results in Sec. V, and draw our conclusions in Sec. VI.

II. DATA

As highlighted in the introduction, deviations from the Λ CDM model modify the expected matter distribution in

⁵While the power spectrum is not directly observable, it serves as an observational proxy for our proof-of-concept study, given its aforementioned importance in cosmological analyses.

the Universe. Consequently, a particular quantity of interest given forthcoming surveys is the matter power spectrum (the Fourier transform of the 2-point matter correlation function), computed at different wave numbers k and redshifts z . We focus on the power spectrum boost $\mathcal{B}(k, z)$, defined as

$$\mathcal{B}(k, z) = \frac{P_{\delta\delta}^{\text{b-}\Lambda\text{CDM}}(k, z)}{P_{\delta\delta}^{\Lambda\text{CDM}}(k, z)}, \quad (1)$$

where $P_{\delta\delta}^{\Lambda\text{CDM}}(k, z)$ is the matter power spectrum predicted by ΛCDM (with constant dark energy density), while $P_{\delta\delta}^{\text{b-}\Lambda\text{CDM}}(k, z)$ represents the corresponding quantity in a beyond- ΛCDM scenario. As a proof of concept, in this work we focus on dynamical dark energy with time-varying equation of state, adopting the most common parametrization [52,53]

$$w(a) = w_0 + (1 - a)w_a, \quad (2)$$

where a is the scale factor, which introduces two extra parameters: w_0 and w_a . We refer to this model as $w_0w_a\text{CDM}$; ΛCDM is recovered for $w_0 = -1$ and $w_a = 0$, with current constraints being statistically consistent with the ΛCDM values of these parameters [54,55]. We reiterate that, while in this work we only focus on a single beyond- ΛCDM model, in future work we plan to train our proposed framework with multiple classes of nonstandard models, as well as with different cosmological probes.

We generate all matter power spectra using the Code for Anisotropies in the Microwave Background (CAMB, [56]), using HMCcode to compute the nonlinear corrections [57–59], in 400 bins in the k -range (0.01–2.5) h/Mpc and for four redshift values $z \in (0.1, 0.48, 0.78, 1.5)$, motivated by a *Euclid*-like survey. We vary five CDM cosmological parameters: the matter (Ω_m) and baryon energy density (Ω_b), the dimensionless Hubble parameter (h), the primordial amplitude (A_s) and the scalar spectral index (n_s). We consider a uniform range for all cosmological parameters, with the intervals reported in Table I. These correspond to five standard deviations around the Planck 2018 constraints [60] with each standard deviation matching *Euclid* pessimistic forecast results using a combination of weak lensing and spectroscopic galaxy clustering [61], as in Ref. [45]. We additionally clip these ranges for Ω_b , w_0 and w_a as indicated in Table I to avoid numerical instabilities of CAMB.

We assume the same Gaussian noise model on the power spectra as in Ref. [45], which is based on the expected error for a Stage IV survey like *Euclid* [62–64]:

$$\sigma(k, z) = \sqrt{\frac{4\pi^2}{k^2 \Delta k V(z)} \left(P_{\delta\delta}(k, z) + \frac{1}{\bar{n}(z)} \right)^2 + \sigma_{\text{sys}}^2}, \quad (3)$$

TABLE I. Minimum and maximum for the uniform ranges used to sample the cosmological parameters when generating the matter power spectra. These are centered around the Planck 2018 best-fit values [60] with a standard deviation from a pessimistic *Euclid* forecast [61]. We further limit the Ω_b , w_0 and w_a ranges to $[-2.4, 1.9]$, $[-3.1, 3.1]$ and $[-5, 0.9]$ standard deviations, respectively, to avoid any numerical instabilities of the Boltzmann solver we use to generate the power spectra (CAMB [56]).

	Parameter	Min	Max
CDM	Ω_m	0.27	0.36
	Ω_b	0.01	0.08
	h	0.65	0.69
	A_s	2.08×10^{-9}	2.31×10^{-9}
	n_s	0.93	1.01
Extension	w_0	-1.3	-0.7
	w_a	-1.6	0.3

where $V(z)$ is the survey volume, $\bar{n}(z)$ represents shot noise, $\Delta k = 0.055 h/\text{Mpc}$ is the bin width, and $\sigma_{\text{sys}} = 5 \text{ Mpc}/h$ is a constant term added in quadrature to account for any systematic effects. The values for $V(z)$ and $\bar{n}(z)$ are reported in Table II of Ref. [45]. The corresponding noise on the boosts $\sigma_{\mathcal{B}}(k, z)$ is obtained via linear error propagation from Eq. (1). We produce a total of 10^5 boosts; we reserve 10% of them to validate the model while training, and another 10% to test our model after training. The remainder is used to train the machine learning model. Generating the dataset is a one-time overhead that took $\mathcal{O}(1)$ hour on a single CPU core; in the future, we will extend our framework to multiple beyond- ΛCDM models by obtaining the theoretical power spectrum predictions with the efficient ReACT package [65,66], or the corresponding emulator [67].

III. MODEL

In order to obtain a compressed representation of dark energy power spectrum boosts, we employ a β -variational autoencoder (β -VAE, [68–70]). A schematic view of our DE-VAE model is reported in Fig. 1.

In particular, our model is made of two blocks, an encoder and a decoder, both parametrized as three layers of 1-D convolutional neural networks. After each convolutional layer, we apply the same trainable activation function as the one described in Ref. [71] and batch normalization [72], to make training more stable. In the encoder, the convolutional layers are followed by a dense layer with linear activation function, whose output is a set of $2d$ variables, where d is the size of the so-called latent space. These variables represent the mean and the standard deviation of d Gaussian distributions. We vary d in our analysis from $d = 1$ to $d = 3$, and report the corresponding results in Sec. IV.




FIG. 1. Schematic view of the DE-VAE model employed in this work. The input data consist of power spectrum boosts, as defined in Eq. (1), produced at four different redshift bins and for wavenumbers $k \in (0.01-2.5) h/\text{Mpc}$. These boosts are compressed by an encoder, parametrized with convolutional neural networks, to a set of disentangled latent variables (in red). These variables are then concatenated with CDM parameters (in blue) and transformed back to power spectrum boosts through the decoder. All model details are reported in Sec. III.

We sample one point from each latent distribution, and concatenate the latent variables with the five CDM parameters (described in Sec. II and Table I), to encourage the model to find a representation only of the extension to the ΛCDM model. The concatenated vector is fed through the decoder to obtain a reconstructed version $\tilde{\mathcal{D}}$ of the input data \mathcal{D} . The complete model is then trained by optimizing the loss function

$$\mathcal{L} = \left| \frac{\mathcal{D} - \tilde{\mathcal{D}}}{\sigma_B} \right|^2 + \beta D_{\text{KL}}[p(\mathbf{z}|\mathcal{D})||\mathcal{N}(\mathbf{0}, \mathbf{1})], \quad (4)$$

where $|\cdot|$ represents the ℓ^2 norm, $D_{\text{KL}}[\cdot||\cdot]$ indicates the Kullback-Leibler (KL) divergence between two distributions, $\mathcal{N}(\mathbf{0}, \mathbf{1})$ is a multivariate standard Gaussian distribution, β is a real parameter that we tune to achieve independent latent variables, and $p(\mathbf{z}|\mathcal{D})$ is the conditional distribution of the latent variables given the input data. We assume that $p(\mathbf{z}|\mathcal{D})$ is the product of d independent Gaussian distributions, namely

$$p(\mathbf{z}|\mathcal{D}) = \prod_{i=1}^d p(z_i|\mathcal{D}) = \prod_{i=1}^d \mathcal{N}(\mu_i(\mathcal{D}), \sigma_i(\mathcal{D})), \quad (5)$$

where the means μ_i and standard deviations σ_i are the output of the encoder given the input data \mathcal{D} . With this choice, the KL divergence can be computed analytically as

$$D_{\text{KL}}[p(\mathbf{z}|\mathcal{D})||\mathcal{N}(\mathbf{0}, \mathbf{1})] = \frac{1}{2} \sum_{i=1}^d (\sigma_i^2 + \mu_i^2 - 1 - 2 \log \sigma_i). \quad (6)$$

Our model thus has to balance the two components of the loss: while the first part aims to achieve a good reconstruction error weighted by the uncertainty on the power spectrum boost, the second piece encourages a disentangled latent space, namely a set of independent latent variables, each carrying information about the relevant parameters used to predict the boosts.

The training procedure is analogous to Refs. [73,74]. We set a batch size of 256, and a starting learning rate of 10^{-3} ; we train the model until the validation loss does not improve for 50 consecutive epochs, then decrease the learning rate by a factor of 10, and repeat until the learning rate reaches 10^{-6} . Our model is trained using the ADAM optimizer [75].

To quantify the dependence between the latent variables and the cosmological parameters, we estimate their mutual information (MI), an information-theoretic measure of the dependence between random variables. Given two random variables with joint probability density $p(x, y)$, their MI is defined as

$$\text{MI}(x, y) = \int_x \int_y p(x, y) \ln \frac{p(x, y)}{p(x)p(y)}, \quad (7)$$

where $p(x)$ and $p(y)$ indicate the marginal distributions, and \ln indicates the natural logarithm, so that MI is measured in natural units (nat). The MI between two variables is zero if and only if the two variables are independent; we refer the reader to Ref. [76] for a complete review of MI and its properties. We employ the GMM-MI library [77], which provides a robust and efficient estimator of MI and includes the uncertainty due to the finite sample size, to obtain MI estimates. We also use GMM-MI to ensure that the latent variables are disentangled: we tune β until we reach a MI between latent variables of $\mathcal{O}(10^{-3})$ nat

or less. Finally, in the case $d = 1$ we further search for an explicit connection between the latent variable and the (w_0, w_a) parameters using symbolic regression, which fits multiple analytic equations to the data through genetic algorithms, as implemented in the `PySR` library [78].

IV. RESULTS

We first report the accuracy of the reconstructed test power spectra $P_{\delta\delta}^{\text{b-}\Lambda\text{CDM}}(k, z)$ when varying the number of latent variables from $d = 1$ to $d = 3$. We compute the predicted spectrum $P_{\delta\delta}^{\text{b-}\Lambda\text{CDM, pred}}(k, z)$ by multiplying the output of the decoder, considered noiseless, by $P_{\delta\delta}^{\Lambda\text{CDM}}(k, z)$; we then compare the predicted and true power spectra, assuming the uncertainty on the power spectra to be as in Eq. (3). In Fig. 2, we show the 68, 95, and 99 percentiles of the statistical significance,

$$\text{Significance}(k, z) = \frac{|P_{\delta\delta}^{\text{b-}\Lambda\text{CDM, pred}}(k, z) - P_{\delta\delta}^{\text{b-}\Lambda\text{CDM}}(k, z)|}{\sigma(k, z)}, \quad (8)$$

which represents the accuracy of the prediction in units of the power spectrum error across the test set, with $\sigma(k, z)$ given by Eq. (3). We also ensure that when $d > 1$ the latent variables are disentangled, i.e. that their MI is $\mathcal{O}(10^{-3})$ nat or less. Even in the case $d = 1$, we encourage a regularized latent space by setting $\beta = 0.01$.

The trained model has a similar performance in all three cases. With one latent variable ($d = 1$, top panel), the percentile plot shows that 95% of the model predictions fall within less than 2σ of the observational error for all wave numbers and redshift bins, with a particularly good performance at $z = 0.48$. Adding a second latent variable

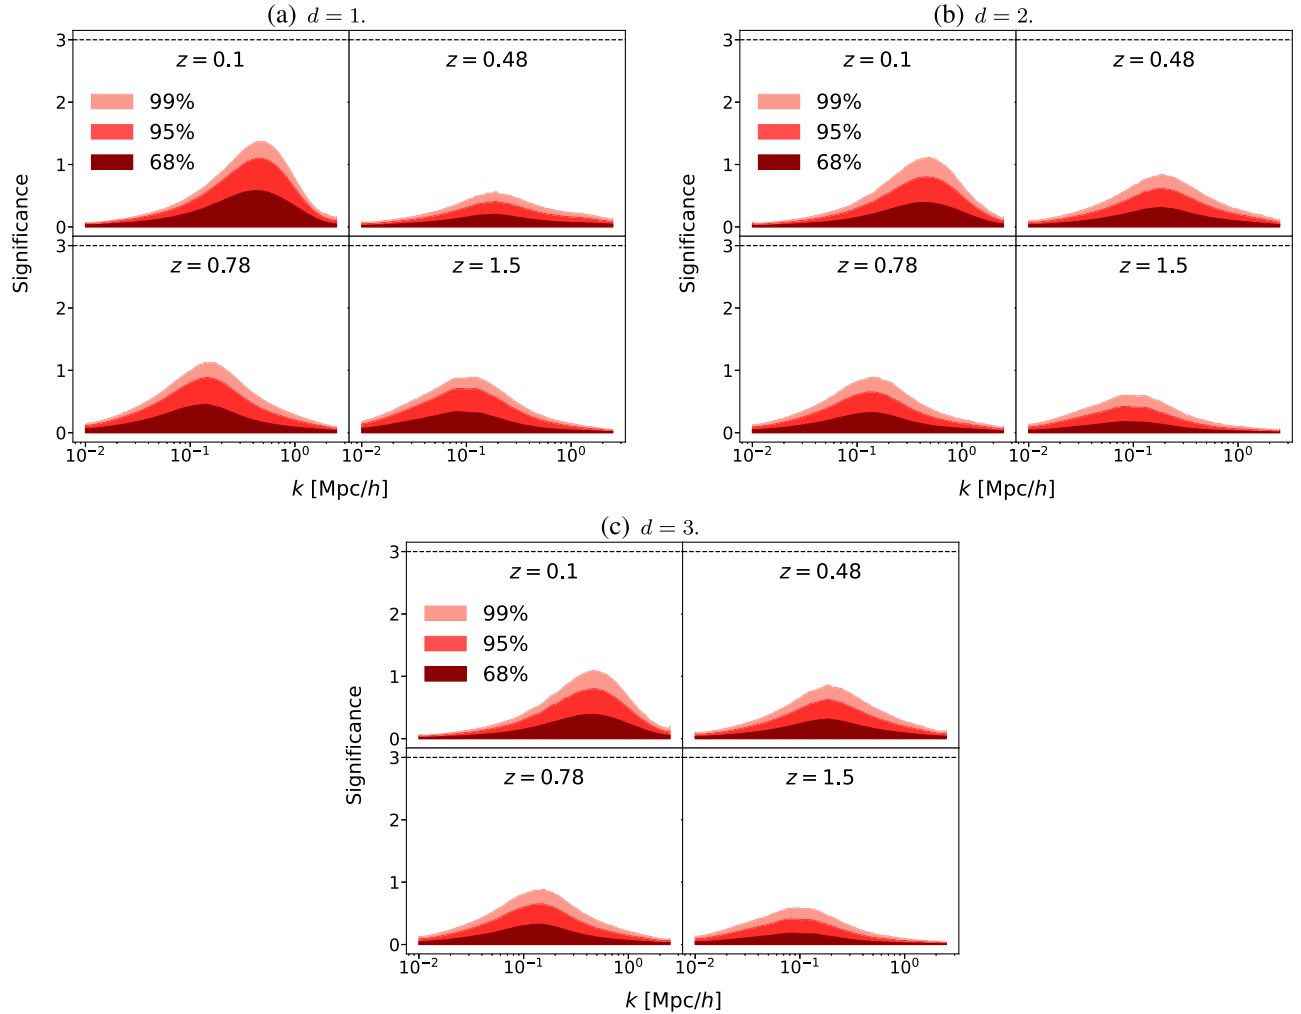


FIG. 2. Percentile accuracy for the DE-VAE model trained with different numbers of latent variables d : (a) $d = 1$, (b) $d = 2$, and (c) $d = 3$. These results are obtained for wave numbers $k \in (0.01-2.5) h/\text{Mpc}$ and redshift values $z \in (0.1, 0.48, 0.78, 1.5)$. The significance is computed using Eq. (8), which compares the test power spectra with the predictions from the model, assuming both the predicted and the ground truth power spectra have the same error as in Eq. (3). The plots show that 99% of the model predictions fall well within 1σ (2σ) of the observational error for most (all) cases. The dashed horizontal lines in each panel correspond to 3σ .

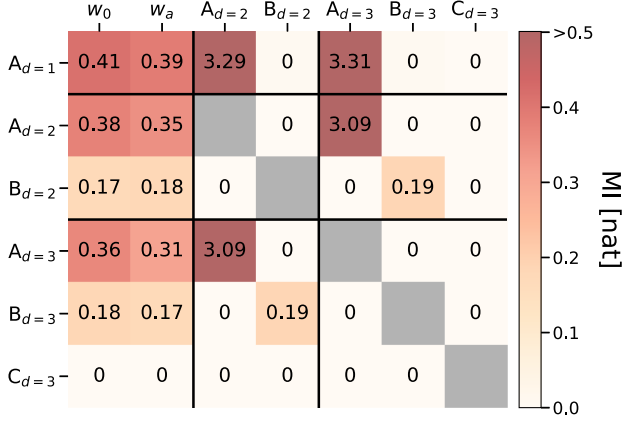


FIG. 3. Mutual information (MI) between dark energy (DE) parameters w_0 and w_a and latent variables. We consider the DE-VAE model trained with different numbers of latent variables d , and indicated with the letters A, B, or C. When $d = 1$, the only latent variable $A_{d=1}$ shows a high MI (i.e., significantly higher than 0) with the DE parameters. Adding a second disentangled variable ($A_{d=2}$ and $B_{d=2}$) slightly improves the predictions of the power spectrum (as shown in Fig. 2), and still shows a significant MI with the DE parameters. When $d = 3$, the third latent variable ($C_{d=3}$) is essentially unused, showing no MI with either the DE parameters or any other latent variables. All values lower than 0.01 nat are indicated with a zero, while we omit the MI values in the gray squares. All MI uncertainties are negligible in this instance.

($d = 2$, middle panel) only marginally improves the model accuracy, which is consistent with the case of three latent variables ($d = 3$, bottom panel). This suggests that adding a third independent latent variable has no impact on the model’s performance: this is in line with the expectations, given that only two additional parameters are used to generate the training data.

To further interpret these results, in Fig. 3 we report the MI values between the latent variables, and between the latent variables and the dynamical dark energy parameters w_0 and w_a . We refer to each latent variable as A, B, or C, with a subscript to indicate the dimensionality of the latent space. In the case with $d = 1$, the single latent variable $A_{d=1}$ shows a high MI (i.e. significantly higher than 0) with both w_0 and w_a , suggesting that $A_{d=1}$ carries information on a combination of the two dark energy parameters. This single variable is effectively capable of predicting dynamical dark energy power spectra, with 95% of the predictions falling within 1σ of the observational error, increasing to 99% for $z > 0.1$, as we showed in the upper panel of Fig. 2; we further explore the links between these variables in Sec. IV A. When training a model with two disentangled latents ($A_{d=2}$ and $B_{d=2}$), these show a significant MI with both w_0 and w_a , with $A_{d=2}$ in particular capturing most of the information content. Moreover, $A_{d=2}$ has a high MI with $A_{d=1}$: this hints at the fact that, when allowing an extra latent variable, the model only learns small corrections on

top of the $d = 1$ model, as demonstrated by the slightly improved accuracy in the middle panel of Fig. 2 as well. Training a model with a third disentangled latent variable ($C_{d=3}$) shows no advantage, since $C_{d=3}$ carries no information content about the DE parameters, and is completely independent from all other latent variables. This is confirmed in the bottom panel of Fig. 2, where we show that adding a third latent variable has minimal impact on the accuracy of the prediction of the model.

A. Symbolic regression in latent space

We further investigate the link between the dynamical DE parameters (w_0, w_a) and the single latent variable $A_{d=1}$ using symbolic regression. Without a direct link, given w_0 and w_a one would need to create the corresponding power spectrum boost and pass it through the encoder to get the corresponding latent variable. By obtaining an explicit equation, we can both bypass the encoder (which we could achieve with another neural network, too), and reuse the parametrization in other contexts, e.g. for studies looking at other observables beyond the power spectrum, as well as for the development of new theoretical models.

We fit PySR with default parameters on a subset of 10^4 training data points; we tested that the results do not change when increasing the number of iterations, or when using more data. The best equation we obtain according to the PySR score (0.92 on the test set), which balances the equation’s accuracy and its complexity, is

$$A_{d=1}(w_0, w_a) = w_0^2 + \frac{e^{w_a + \cos(w_0)}}{w_0} + e^{\cos(1)} - 1. \quad (9)$$

Note that we have manually rescaled the equation so that, for Λ CDM, $A_{d=1}(w_0 = -1, w_a = 0) = 0$. We show the 3D distribution of points for the test set, together with the graph of Eq. (9), in Fig. 4. If we use this equation to predict the value of the latent variable instead of the encoder, the 99% level of the power spectrum predictions can reach up to 5σ deviations for the test set described in Table I; this is to be expected, since symbolic regression tends to sacrifice accuracy in favor of simpler expressions [78]. However, if we restrict ourselves to spectra generated within three standard deviations of the Planck best-fit values (instead of five as the training and test data),⁶ we show that this equation actually achieves an acceptable performance, with the percentile accuracy shown in Fig. 5. We also tested that we obtain the same equation if we allow all parameters (CDM + DE) to be given as input variables to predict $A_{d=1}$, further confirming that the latent variable only depends on the DE parameters.

⁶In this case as well, we further limit the ranges of Ω_b and w_a to $[-2.4, 1.9]$ and $[-3, 0.9]$ standard deviations, respectively, to avoid numerical instabilities, as we describe in Table I.




FIG. 4. Scatter plot of w_0 , w_a and the $d = 1$ latent variable for the test set points (blue), together with the graph of the symbolic regression equation reported in Eq. (9) (orange). While only accurate within the 2σ level for 95% of the predictions on a restricted dataset (see Fig. 5), the learned analytic expression can be used to bypass the encoder and predict the value of the latent variable given w_0 and w_a . The red line shows the intersection between the symbolic regression and $A_{d=1} = 0$ planes; the latter plane describes a degeneracy with Λ CDM, which we further discuss, together with broader applications of our framework, in Sec. V.




FIG. 5. Same as Fig. 2(a), but with the latent variable predicted using Eq. (9) from w_0 and w_a , and tested on 1000 boosts generated within three standard deviations of the fiducial Planck values (instead of the five standard deviations considered in the rest of this work).

V. DISCUSSION AND OUTLOOK

Using our DE-VAE framework, we demonstrated that a single variable is sufficient to accurately predict the theoretical matter power spectrum for a particular beyond- Λ CDM model with two additional free parameters.

The architecture benefits from its representation learning formulation, which allows us to explore the lower-dimensional latent space and obtain further insight on the nature of the training data. We showcased the use of mutual information as an effective tool to interpret the latent space; moreover, we employed symbolic regression (SR) to obtain an analytic equation linking the cosmological parameters with the latent variable. We envision that our framework can be employed and extended in the following ways.

(i) Given observational constraints on the matter power spectrum, it is possible to use the DE-VAE decoder within a Markov chain Monte Carlo (MCMC) algorithm to sample the posterior distribution of the latent space (and CDM) parameters, and thus effectively classify the observed spectrum as either in agreement with Λ CDM or not. The value of the latent variable corresponding to Λ CDM can be obtained by passing constant unitary boosts at all redshifts and wave numbers through the encoder. It would then be straightforward to constrain the latent parameter $A_{d=1}$ and assess the statistical significance of any deviations from Λ CDM directly in latent space.

(ii) The particular latent variable we found with DE-VAE and the corresponding parametrization obtained using SR are optimized for a *Euclid*-like power spectrum analysis, and capture a particular degeneracy of the w_0w_a CDM model in predicting theoretical matter power spectra (indicated in red in Fig. 4). We reiterate that in our framework we are not considering systematic and observational effects from which we are not necessarily independent: these would in principle require training a different VAE for each choice of nuisance parameters, and could introduce different degeneracies based on redshift, scale and other gravitational effects. On the other hand, our goal is to capture degeneracies at the theory level, both within a single model and across multiple Λ CDM extensions. The latent variable we found for w_0w_a CDM in this work can be interpreted in the same vein as $S_8 = \sigma_8 \sqrt{\Omega_m/0.3}$,

i.e. as a parameter constructed to be particularly well-constrained by a certain dataset. Given that our neural architecture is easily adaptable to different input data, it will be therefore interesting to train our framework on other probes (e.g. on CMB spectra or weak lensing maps). While this would require to produce a new dataset (possibly including additional observational effects) to train another VAE, it could yield a latent variable capturing an orthogonal degeneracy with respect to the one we found in this work, and therefore highlight which summary statistics should be targeted to break degeneracies when constraining beyond- Λ CDM models. By analyzing different datasets and considering sky maps rather

than 2-point functions, we aim to obtain results which are complementary with respect to the matter power spectrum, as well as retain the information that is lost when compressing the full field into summary statistics [44]. In turn, this can inform the development of future cosmological missions, that could then target these particular parameters and probes to maximize their performance.

- (iii) When repeating the entire analysis using the `EuclidEmulator2` [EE2, [79]] instead of `HMCode`, we find that one latent variable is still sufficient to predict the nonlinear matter power spectra with good accuracy. While the equation found by SR when using the EE2 to generate the training data is different due to the stochasticity of the genetic algorithm, we verified that Eq. (9) still provides a decent fit to the mapping between (w_0, w_a) and the latent variable. Moreover, we note that the degeneracy direction expressed by Eq. (9) is aligned with constraints obtained with the linear matter power spectrum [80], where it is expected that $w_0 w_a$ CDM models with identical linear growth factors yield indistinguishable linear power spectra. The degeneracy is then partially preserved since quasilinear effects can be described as a function of the linear power spectrum to a good extent [81]. The degeneracy captured by our approach is thus robust to the choice of the nonlinear model, while being in agreement with results at the linear level. It can therefore be traced back to a theoretically well-understood phenomenological source, which proves the result to be both meaningful and interpretable. While for the particular model considered in this work this degeneracy is known, it suggests that similar (and perhaps more complex) degeneracies, which are theoretically interpretable, can be identified in the larger model space. They can then be utilized for new theoretically- and phenomenologically informed parametrizations.
- (iv) By showcasing the validity of our approach with a single model for which the underlying ground truth is known, we anticipate being able to train an analogous model on a broader variety of spectra obtained assuming different beyond- Λ CDM models, including Hu-Sawicki $f(R)$ [82] and dark sector interactions [83], to be used as training data. The power spectrum predictions for these models can be efficiently obtained with the `ReACT` package [65–67]. In this way, we expect to capture a common low-dimensional parametrization of a multitude of beyond- Λ CDM scenarios; in turn, this can be used to facilitate the exploration of beyond- Λ CDM models when performing MCMC analyses, and inform the theoretical development of models which try to unify these different extensions. This highlights

another advantage of our framework, since performing an MCMC analysis on each of these models separately would clearly require many more iterations than running a single MCMC on the common extension found in latent space.

VI. CONCLUSIONS

We developed DE-VAE, a framework to search for a compressed representation of beyond- Λ CDM (Λ cold dark matter) models using representation learning. As a first proof-of-concept study, we focused on dynamical dark energy (DE) as a straightforward extension to the Λ CDM cosmological model. In particular, we considered the archetypal two-variable parametrization of the DE equation of state, and produced the corresponding matter power spectrum boosts as in Eq. (1). These boosts were used as training data of a β -variational autoencoder (β -VAE) model, summarized in Fig. 1, where we encouraged a d -dimensional disentangled compressed representation independent from CDM parameters.

We found that a single latent variable, in combination with five CDM parameters, can predict the power spectra of dynamical DE at different redshifts and up to $k = 2.5 h/\text{Mpc}$ within a 1σ (2σ) error for 95% (99%) of the data; the error, reported in Eq. (3), is computed combining cosmic variance, shot noise and possible systematic effects for a Stage IV-like survey. This single latent variable shows a high mutual information (MI) with the parameters of the DE equation of state, confirming that the DE-VAE has learned a meaningful representation of the boosts. Adding a second independent latent variable only marginally improves the prediction accuracy, with both latent variables showing a significant MI with the DE parameters. Finally, we trained a model with $d = 3$ and showed that the third independent latent variable has no MI with other latent variables and with the DE parameters, confirming that it has no significant impact on the prediction accuracy. In the $d = 1$ case, we also demonstrated the use of symbolic regression to obtain an explicit equation linking the DE parameters and the latent variable, which can be used to facilitate the power spectrum prediction by skipping the encoder compression, and which sheds further light on the nature of the latent variable.

We highlighted a range of applications and future developments of our framework. To mention a few, these include applying DE-VAE to different cosmological datasets, to explore which probes should be targeted by next-generation surveys to break parameter degeneracies, as well as training the architecture on multiple extensions of Λ CDM, to shed light on their common aspects and inform the development of the underlying theories. We will explore these avenues in future work. The DE-VAE architecture and the data generated to train and test the model are available upon request.

ACKNOWLEDGMENTS

We are grateful to Benjamin Bose and Benjamin Hertzsch for their feedback on this work. D.P. and L.L. were supported by a Swiss National Science Foundation (SNSF) Professorship Grant (No. 202671). D.P. was also

supported by the SNSF Sinergia Grant No. CRSII5_193826 “AstroSignals: A New Window on the Universe, with the New Generation of Large Radio-Astronomy Facilities”. The computations underlying this work were performed on the Baobab cluster at the University of Geneva.

-
- [1] P. Bull, Y. Akrami, J. Adamek, T. Baker, E. Bellini, J. Beltrán Jiménez, E. Bentivegna, S. Camera, S. Clesse, J. H. Davis *et al.*, Beyond Λ CDM: Problems, solutions, and the road ahead, *Phys. Dark Universe* **12**, 56 (2016).
- [2] N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo *et al.* (Planck Collaboration), Planck 2018 results. VI. Cosmological parameters, *Astron. Astrophys.* **641**, A6 (2020).
- [3] J. Martin, Everything you always wanted to know about the cosmological constant problem (but were afraid to ask), *C.R. Phys.* **13**, 566 (2012).
- [4] A. G. Riess, W. Yuan, L. M. Macri, D. Scolnic, D. Brout, S. Casertano, D. O. Jones, Y. Murakami, G. S. Anand, L. Breuval *et al.*, A comprehensive measurement of the local value of the Hubble constant with $1 \text{ km s}^{-1} \text{ Mpc}^{-1}$ uncertainty from the Hubble space telescope and the SH0ES Team, *Astrophys. J. Lett.* **934**, L7 (2022).
- [5] C. Heymans, T. Tröster, M. Asgari, C. Blake, H. Hildebrandt, B. Joachimi, K. Kuijken, C.-A. Lin, A. G. Sánchez, J. L. van den Busch *et al.*, KiDS-1000 cosmology: Multi-probe weak gravitational lensing and spectroscopic galaxy clustering constraints, *Astron. Astrophys.* **646**, A140 (2021).
- [6] T. M. C. Abbott, M. Aguena, A. Alarcon, O. Alves, A. Amon, F. Andrade-Oliveira, M. Asgari, S. Avila, D. Bacon, K. Bechtol *et al.* (Dark Energy Survey and Kilo-Degree Survey Collaboration), DES Y3 + KiDS-1000: Consistent cosmology combining cosmic shear surveys, *Open J. Astrophys.* **6**, 2305.17173 (2023).
- [7] Ž. Ivezić, S. M. Kahn, J. A. Tyson, B. Abel, E. Acosta, R. Allsman, D. Alonso, Y. AlSayyad, S. F. Anderson, J. Andrew *et al.*, LSST: From science drivers to reference design and anticipated data products, *Astrophys. J.* **873**, 111 (2019).
- [8] R. Laureijs, J. Amiaux, S. Arduini, J. L. Auguères, J. Brinchmann, R. Cole, M. Cropper, C. Dabin, L. Duvet, A. Ealet *et al.*, Euclid definition study report, [arXiv:1110.3193](https://arxiv.org/abs/1110.3193).
- [9] D. Spergel, N. Gehrels, C. Baltay, D. Bennett, J. Breckinridge, M. Donahue, A. Dressler, B. S. Gaudi, T. Greene, O. Guyon *et al.*, Wide-field infrared survey telescope-astronomy focused telescope assets WFIRST-AFTA 2015 report, [arXiv:1503.03757](https://arxiv.org/abs/1503.03757).
- [10] P. Ade, J. Aguirre, Z. Ahmed, S. Aiola, A. Ali, D. Alonso, M. A. Alvarez, K. Arnold, P. Ashton, J. Austermann *et al.*, The Simons observatory: Science goals and forecasts, *J. Cosmol. Astropart. Phys.* **02** (2019) 056.
- [11] E. J. Copeland, M. Sami, and S. Tsujikawa, Dynamics of dark energy, *Int. J. Mod. Phys. D* **15**, 1753 (2006).
- [12] T. Clifton, P. G. Ferreira, A. Padilla, and C. Skordis, Modified gravity and cosmology, *Phys. Rep.* **513**, 1 (2012).
- [13] K. Koyama, Cosmological tests of modified gravity, *Rep. Prog. Phys.* **79**, 046902 (2016).
- [14] A. Joyce, B. Jain, J. Khoury, and M. Trodden, Beyond the cosmological standard model, *Phys. Rep.* **568**, 1 (2015).
- [15] A. Joyce, L. Lombriser, and F. Schmidt, Dark energy versus modified gravity, *Annu. Rev. Nucl. Part. Sci.* **66**, 95 (2016).
- [16] M. Ishak, Testing general relativity in cosmology, *Living Rev. Relativity* **22**, 1 (2019).
- [17] E. Di Valentino, O. Mena, S. Pan, L. Visinelli, W. Yang, A. Melchiorri, D. F. Mota, A. G. Riess, and J. Silk, In the realm of the Hubble tension—a review of solutions, *Classical Quantum Gravity* **38**, 153001 (2021).
- [18] J. G. Williams, S. G. Turyshev, and D. H. Boggs, Progress in lunar laser ranging tests of relativistic gravity, *Phys. Rev. Lett.* **93**, 261101 (2004).
- [19] S. Schlamminger, K.-Y. Choi, T. A. Wagner, J. H. Gundlach, and E. G. Adelberger, Test of the equivalence principle using a rotating torsion balance, *Phys. Rev. Lett.* **100**, 041101 (2008).
- [20] T. A. Wagner, S. Schlamminger, J. H. Gundlach, and E. G. Adelberger, Torsion-balance tests of the weak equivalence principle, *Classical Quantum Gravity* **29**, 184002 (2012).
- [21] C. M. Will, The confrontation between general relativity and experiment, *Living Rev. Relativity* **17**, 4 (2014).
- [22] P. J. E. Peebles, Testing general relativity on the scales of cosmology, in *General Relativity and Gravitation* (World Scientific, Singapore, 2005).
- [23] A. Vainshtein, To the problem of nonvanishing gravitation mass, *Phys. Lett.* **39B**, 393 (1972).
- [24] E. Babichev and C. Deffayet, An introduction to the Vainshtein mechanism, *Classical Quantum Gravity* **30**, 184001 (2013).
- [25] E. Babichev, C. Deffayet, and R. Ziour, k -Mouflage gravity, *Int. J. Mod. Phys. D* **18**, 2147 (2009).
- [26] J. Khoury and A. Weltman, Chameleon cosmology, *Phys. Rev. D* **69**, 044026 (2004).
- [27] K. Hinterbichler and J. Khoury, Screening long-range forces through local symmetry restoration, *Phys. Rev. Lett.* **104**, 231301 (2010).
- [28] L. Lombriser, Parametrizations for tests of gravity, *Int. J. Mod. Phys. D* **27**, 1848002 (2018).
- [29] W. Hu and I. Sawicki, Parametrized post-Friedmann framework for modified gravity, *Phys. Rev. D* **76**, 104043 (2007).

- [30] P. Zhang, M. Liguori, R. Bean, and S. Dodelson, Probing gravity at cosmological scales by measurements which test the relationship between gravitational lensing and matter overdensity, *Phys. Rev. Lett.* **99**, 141302 (2007).
- [31] L. Amendola, M. Kunz, and D. Sapone, Measuring the dark side (with weak lensing), *J. Cosmol. Astropart. Phys.* **04** (2008) 013.
- [32] E. Bertschinger and P. Zukin, Distinguishing modified gravity from dark energy, *Phys. Rev. D* **78**, 024015 (2008).
- [33] R. Bean and M. Tangmatitham, Current constraints on the cosmic growth history, *Phys. Rev. D* **81**, 083534 (2010).
- [34] L. Pogosian, A. Silvestri, K. Koyama, and G.-B. Zhao, How to optimally parametrize deviations from general relativity in the evolution of cosmological perturbations, *Phys. Rev. D* **81**, 104023 (2010).
- [35] Y.-S. Song, L. Hollenstein, G. Caldera-Cabral, and K. Koyama, Theoretical priors on modified growth parametrisations, *J. Cosmol. Astropart. Phys.* **04** (2010) 018.
- [36] S. Duane, A. Kennedy, B. J. Pendleton, and D. Roweth, Hybrid Monte Carlo, *Phys. Lett. B* **195**, 216 (1987).
- [37] M. D. Hoffman and A. Gelman, The No-U-Turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo, *J. Mach. Learn. Res.* **15**, 1593 (2014).
- [38] D. Huterer and G. Starkman, Parametrization of dark-energy properties: A principal-component approach, *Phys. Rev. Lett.* **90**, 031301 (2003).
- [39] R. G. Crittenden, L. Pogosian, and G.-B. Zhao, Investigating dark energy experiments with principal components, *J. Cosmol. Astropart. Phys.* **12** (2009) 025.
- [40] G.-B. Zhao, L. Pogosian, A. Silvestri, and J. Zylberberg, Cosmological tests of general relativity with future tomographic surveys, *Phys. Rev. Lett.* **103**, 241301 (2009).
- [41] A. Hojjati, G.-B. Zhao, L. Pogosian, A. Silvestri, R. Crittenden, and K. Koyama, Cosmological tests of general relativity: A principal component analysis, *Phys. Rev. D* **85**, 043508 (2012).
- [42] S. Asaba, C. Hikage, K. Koyama, G.-B. Zhao, A. Hojjati, and L. Pogosian, Principal component analysis of modified gravity using weak lensing and peculiar velocity measurements, *J. Cosmol. Astropart. Phys.* **08** (2013) 029.
- [43] J. Schmelzle, A. Lucchi, T. Kacprzak, A. Amara, R. Sgier, A. Réfrégier, and T. Hofmann, Cosmological model discrimination with deep learning, [arXiv:1707.05167](https://arxiv.org/abs/1707.05167).
- [44] A. Peel, F. Lalande, J.-L. Starck, V. Pettorino, J. Merten, C. Giocoli, M. Meneghetti, and M. Baldi, Distinguishing standard and modified gravity cosmologies with machine learning, *Phys. Rev. D* **100**, 023508 (2019).
- [45] M. Mancarella, J. Kennedy, B. Bose, and L. Lombriser, Seeking new physics in cosmology with Bayesian neural networks: Dark energy and modified gravity, *Phys. Rev. D* **105**, 023531 (2022).
- [46] S. K. N. Portillo, J. K. Parejko, J. R. Vergara, and A. J. Connolly, Dimensionality reduction of SDSS spectra with variational autoencoders, *Astron. J.* **160**, 45 (2020).
- [47] N. Sedaghat, M. Romaniello, J. E. Carrick, and F.-X. Pineau, Machines learn to infer stellar parameters just by looking at a large number of spectra, *Mon. Not. R. Astron. Soc.* **501**, 6026 (2021).
- [48] R. Sarmiento, M. Huertas-Company, J. H. Knapen, S. F. Sánchez, H. Domínguez Sánchez, N. Drory, and J. Falcón-Barroso, Capturing the physics of MaNGA galaxies with self-supervised machine learning, *Astrophys. J.* **921**, 177 (2021).
- [49] H. Teimoorinia, F. Archinuk, J. Woo, S. Shishehchi, and A. F. L. Bluck, Mapping the diversity of galaxy spectra with deep unsupervised machine learning, *Astron. J.* **163**, 71 (2022).
- [50] L. Lucie-Smith, H. V. Peiris, A. Pontzen, B. Nord, J. Thiyaalingam, and D. Piras, Discovering the building blocks of dark matter halo density profiles with neural networks, *Phys. Rev. D* **105**, 103533 (2022).
- [51] L. Lucie-Smith, H. V. Peiris, and A. Pontzen, Explaining dark matter halo density profiles with neural networks, *Phys. Rev. Lett.* **132**, 031001 (2024).
- [52] M. Chevallier and D. Polarski, Accelerating universes with scaling dark matter, *Int. J. Mod. Phys. D* **10**, 213 (2001).
- [53] E. V. Linder, Exploring the expansion history of the universe, *Phys. Rev. Lett.* **90**, 091301 (2003).
- [54] N. Menci, A. Grazian, M. Castellano, P. Santini, E. Giallongo, A. Lamastra, F. Fortuni, A. Fontana, E. Merlin, T. Wang *et al.*, Constraints on dynamical dark energy models from the abundance of massive galaxies at high redshifts, *Astrophys. J.* **900**, 108 (2020).
- [55] T. M. C. Abbott, M. Aguena, A. Alarcon, O. Alves, A. Amon, F. Andrade-Oliveira, J. Annis, S. Avila, D. Bacon, E. Baxter *et al.* (DES Collaboration), Dark Energy Survey year 3 results: Constraints on extensions to Λ CDM with weak lensing and galaxy clustering, *Phys. Rev. D* **107**, 083504 (2023).
- [56] A. Lewis and A. Challinor, CAMB: Code for anisotropies in the microwave background, Astrophysics Source Code Library, record ascl:1102.026 (2011), ascl:1102.026.
- [57] A. J. Mead, J. A. Peacock, C. Heymans, S. Joudaki, and A. F. Heavens, An accurate halo model for fitting non-linear cosmological power spectra and baryonic feedback models, *Mon. Not. R. Astron. Soc.* **454**, 1958 (2015).
- [58] A. J. Mead, C. Heymans, L. Lombriser, J. A. Peacock, O. I. Steele, and H. A. Winther, Accurate halo-model matter power spectra with dark energy, massive neutrinos and modified gravitational forces, *Mon. Not. R. Astron. Soc.* **459**, 1468 (2016).
- [59] A. J. Mead, S. Brieden, T. Tröster, and C. Heymans, hMCode-2020: Improved modelling of non-linear cosmological power spectra with baryonic feedback, *Mon. Not. R. Astron. Soc.* **502**, 1401 (2021).
- [60] N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Baccigalupi, M. Ballardini, A. J. Banday, R. B. Barreiro, N. Bartolo, S. Basak *et al.*, Planck 2018 results. VI. Cosmological parameters, *Astron. Astrophys.* **641**, A6 (2020).
- [61] A. Blanchard, S. Camera, C. Carbone, V. F. Cardone, S. Casas, S. Clesse, S. Ilić, M. Kilbinger, T. Kitching, M. Kunz *et al.*, Euclid preparation. VII. Forecast validation for Euclid cosmological probes, *Astron. Astrophys.* **642**, A191 (2020).
- [62] H. A. Feldman, N. Kaiser, and J. A. Peacock, Power-spectrum analysis of three-dimensional redshift surveys, *Astrophys. J.* **426**, 23 (1994).
- [63] H.-J. Seo and D. J. Eisenstein, Improved forecasts for the baryon acoustic oscillations and cosmological distance scale, *Astrophys. J.* **665**, 14 (2007).

- [64] G.-B. Zhao, Modeling the nonlinear clustering in modified gravity models. I. A fitting formula for the matter power spectrum of $f(R)$ gravity, *Astrophys. J. Suppl. Ser.* **211**, 23 (2014).
- [65] B. Bose, M. Cataneo, T. Tröster, Q. Xia, C. Heymans, and L. Lombriser, On the road to percent accuracy IV: reACT—computing the non-linear power spectrum beyond Λ CDM, *Mon. Not. R. Astron. Soc.* **498**, 4650 (2020).
- [66] B. Bose, M. Tsedrik, J. Kennedy, L. Lombriser, A. Poursidou, and A. Taylor, Fast and accurate predictions of the non-linear matter power spectrum for general models of dark energy and modified gravity, *Mon. Not. R. Astron. Soc.* **519**, 4780 (2023).
- [67] A. Spurio Mancini and B. Bose, On the degeneracies between baryons, massive neutrinos and $f(R)$ gravity in stage IV cosmic shear analyses, [arXiv:2305.06350](https://arxiv.org/abs/2305.06350).
- [68] D. P. Kingma and M. Welling, Auto-encoding variational Bayes, in *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, 2014, Conference Track Proceedings*, edited by Y. Bengio and Y. LeCun (ICLR, 2014).
- [69] D. J. Rezende, S. Mohamed, and D. Wierstra, Stochastic backpropagation and approximate inference in deep generative models, in *Proceedings of the 31st International Conference on Machine Learning*, Proceedings of Machine Learning Research, Vol. 32(2), edited by E. P. Xing and T. Jebara (PMLR, Beijing, China, 2014), pp. 1278–1286.
- [70] I. Higgins, L. Matthey, A. Pal, C. P. Burgess, X. Glorot, M. M. Botvinick, S. Mohamed, and A. Lerchner, beta-VAE: Learning basic visual concepts with a constrained variational framework, in *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, 2017, Conference Track Proceedings* (OpenReview.net, 2017).
- [71] J. Alsing, H. Peiris, J. Leja, C. Hahn, R. Tojeiro, D. Mortlock, B. Leistedt, B. D. Johnson, and C. Conroy, SPECULATOR: Emulating stellar population synthesis for fast and accurate galaxy spectra and photometry, *Astrophys. J. Suppl. Ser.* **249**, 5 (2020).
- [72] S. Ioffe and C. Szegedy, Batch normalization: Accelerating deep network training by reducing internal covariate shift, in *Proceedings of the 32nd International Conference on Machine Learning*, Proceedings of Machine Learning Research, Vol. 37, edited by F. Bach and D. Blei (PMLR, Lille, France, 2015), pp. 448–456.
- [73] 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).
- [74] D. Piras and A. Spurio Mancini, CosmoPower-JAX: High-dimensional Bayesian inference with differentiable cosmological emulators, *Open J. Astrophys.* **6**, 20 (2023).
- [75] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, in *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, 2015, Conference Track Proceedings*, edited by Y. Bengio and Y. LeCun (ICLR, 2015).
- [76] J. R. Vergara and P. A. Estévez, A review of feature selection methods based on mutual information, *Neural Comput. Appl.* **24**, 175 (2014).
- [77] D. Piras, H. V. Peiris, A. Pontzen, L. Lucie-Smith, N. Guo, and B. Nord, A robust estimator of mutual information for deep learning interpretability, *Mach. Learn.* **4**, 025006 (2023).
- [78] M. Cranmer, Interpretable machine learning for science with PySR and SymbolicRegression.jl, [arXiv:2305.01582](https://arxiv.org/abs/2305.01582).
- [79] M. Knabenhans, J. Stadel, D. Potter, J. Dakin, S. Hannestad, T. Tram, S. Marelli, A. Schneider, R. Teyssier *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).
- [80] S. G. Beauchamps, P. Baratta, J. Bel, S. Escoffier, and W. Gillard, Constraining the total neutrino mass with the power spectrum, *Ateliers Dark Energy* **1**, 22 (2021).
- [81] A. J. Mead, Spherical collapse, formation hysteresis and the deeply non-linear cosmological power spectrum, *Mon. Not. R. Astron. Soc.* **464**, 1282 (2017).
- [82] W. Hu and I. Sawicki, Models of $f(r)$ cosmic acceleration that evade solar system tests, *Phys. Rev. D* **76**, 064004 (2007).
- [83] G. Dvali, G. Gabadadze, and M. Porrati, 4D gravity on a brane in 5D Minkowski space, *Phys. Lett. B* **485**, 208 (2000).